

DIRECT DIGITAL SYNTHESIS OF PRECISE, STABLE ANGLE MODULATED RF SIGNAL

The present invention relates to radio frequency modulators, particularly digital radio frequency modulators.

Modulation can be defined as the alteration of some characteristic of a known signal or waveform, i.e., a carrier, as a function of some unknown signal or waveform that conveys information. In radio-frequency (RF) communication systems, the carrier is typically a sinusoid, and there are several methods of modulating the carrier. These include linear modulation, angle modulation, and various types of pulse modulation. Given a sinusoidal carrier described by the equation $A(t)\cos(\omega_c t + \phi(t))$, there are two parameters, the amplitude and the phase angle, that can be varied in accordance with an information signal. Linear modulation results when the amplitude is varied as a linear function of the information signal. Angle modulation includes phase modulation and frequency modulation. If a term is included in the argument of the sinusoidal function that varies in proportion to the information signal, the result is phase modulation. If the argument is such that the difference in the instantaneous frequency and the carrier frequency is proportional to the information signal, the result is frequency modulation.

Demodulation of RF signals has typically involved a quadrature detector having two branches, an I ("in-phase") branch and a Q ("quadrature" or 90° phase-shifted) branch. In the I branch, a received signal is multiplied by the cosine form of the carrier signal and then passed through a low-pass filter. In the Q branch, the received signal is multiplied by the sine form of the carrier signal and pass through a low-pass filter. Quadrature detectors of this type are linear, well-understood, and almost universally used. To obtain the information signal from the I and Q components produced by the respective I and Q branches of the quadrature detector, signal processing is performed. In particular, the phase of the signal may be obtained by taking the inverse tangent of the ratio of Q to I. The amplitude of the signal may

be obtained according to the Pythagorean theorem by taking the square root of the sum of the squares of I and Q. These mathematical operations are non-linear.

Two salient observations may therefore be made concerning quadrature detection. First, detection proceeds in two steps, a first mixing step (to obtain I and Q) that is linear and a second signal processing step to which non-linearities are relegated. Second, a coordinate system conversion is first performed and then reversed. That is, the received signal, which may be readily described in polar coordinates in terms of the desired quantities of amplitude and phase, is first converted to rectangular coordinates by projecting the instantaneous signal vector in polar coordinates onto the X (I) and Y (Q) axes, and is then converted back to polar coordinates to obtain amplitude and phase. Such conversions require circuitry that occupies space and consumes power—both of which may be precious commodities, especially in mobile applications such as cellular telephones, pagers, etc. Such conversions may also entail substantial inaccuracies.

A similar situation exists on the transmit side with respect to RF modulators. That is, amplitude and phase information is originally represented in polar form. I and Q processing is then performed in which amplitude and phase information is encoded as I and Q signals which are then converted back to polar form and summed to form the final output signal. This process is illustrated in Figure 1. One widely used type of modulation, Phase Shift Keying (PSK), requires the transmit signal to undergo accurate phase shifts and to exhibit good phase stability. PSK is used in GSM cellular telephones, for example. A conventional modulator of the quadrature type, while it exhibits good phase stability, achieves only marginal accuracy. Sophisticated coding schemes may be required to compensate for inaccuracies of the modulator, and performance under noisy conditions may noticeably deteriorate.

Beside conventional quadrature techniques, various other modulation techniques are known. In one such technique, a phase lock loop (PLL) is used to multi-

ply a modulated signal in frequency and phase to obtain a high frequency signal (e.g., 900 MHz). Referring to Figure 2, a PLL includes a phase detector 201, a loop filter 203, a Voltage Controlled Oscillator (VCO) 205, and a divide by N counter 207. In the case of a 900 MHz output signal, a typical divisor might be 64, for example. If a modulated signal $\cos(\omega_c t + \phi(t))$ is applied to the input of the circuit, then ideally, a modulated signal $\cos(N\omega_c t + N\phi(t))$ is produced at the output of the circuit. The stability of the circuit, however, is poor unless the loop filter is made to have a narrow bandwidth. With a narrow bandwidth loop filter, the circuit is no longer able to track rapid changes in the modulated input signal, producing inaccuracies in the modulated output signal.

To overcome the foregoing difficulty, an arrangement shown in Figure 3 has been devised. A baseband modulation signal is applied to a Voltage Controlled Crystal Oscillator (VCXO) 309 to produce a modulated signal which is applied to the input of a PLL, as before. An additional feedforward path is used to inject the baseband modulation signal into the PLL at a point past the loop filter. More particularly, the baseband modulation signal is applied to an adjustable gain amplifier 311. The output of the adjustable gain amplifier is applied to an adder 313 situated between the loop filter and the VCO. The gain of the amplifier is adjusted so that the effects of the loop filter in removing some of the modulation is precisely offset by reinjecting the modulation signal. Unfortunately, achieving precisely the correct adjustment is a painstaking manual process. Furthermore, although sufficient precision may be obtained for FM radio communications, the precision required for PSK radio communications is lacking.

The present invention, generally speaking, provides an RF modulator that allows precise, stable phase shifts to be obtained. The modulator uses a PLL structure including an auxiliary feedforward path used to inject a baseband modulation signal into the PLL at a point past a loop filter of the PLL. A phase demodulator recovers phase information from the output signal of the PLL. The recovered phase

information is compared to the phase information of the baseband modulation signal. A resulting error signal is used to control injection of the baseband modulation signal into the PLL, to automatically achieve the correct "dosage." A precise, adaptive, phase-stable modulator results. The adaptation of the modulator compensates for variability of the VCO and other components of the PLL.

Figure 1 is a block diagram of a conventional quadrature RF modulator;

Figure 2 is a block diagram of a conventional PLL-based RF modulator;

Figure 3 is a block diagram of a conventional PLL-based RF modulator having a baseband modulation signal injection feature;

Figure 4 is a block diagram of an RF modulator in accordance with the present invention;

Figure 5 is a first timing diagram illustrating operation of the frequency sampling circuit of Figure 21;

Figure 6 is a second timing diagram illustrating operation of the frequency sampling circuit of Figure 21;

Figure 7 is a graph of two alternative weighting functions that may be used to perform digital filtering of a digital bit stream produced by a circuit such as that of Figure 21;

Figure 8 is a graph illustrating the accuracy obtained from a digital frequency discriminator using a constant weighting function;

Figure 9 is a graph illustrating the accuracy obtained from a digital frequency discriminator using a triangular weighting function;

Figure 10 is a block diagram of one example of a digital filter that may be used in conjunction with a frequency sampling circuit such as that of Figure 21;

Figure 11A is a tabulation illustrating one method of digital phase discrimination;

Figure 11B is a plot showing results of the method of Figure 11A;

Figure 11C is a plot of a weighting function used in connection with Figures 11A and 11B;

Figure 12A is a tabulation illustrating another method of digital phase discrimination;

Figure 12B is a plot showing results of the method of Figure 12A;

Figure 12C is a plot of a weighting function used in connection with Figures 12A and 12B;

Figure 13 is a block diagram of digital phase discrimination hardware in accordance with the technique of Figure 12;

Figure 14A is a tabulation illustrating yet another method of digital phase discrimination;

Figure 14B is a plot showing results of the method of Figure 14A;

Figure 14C is a plot of a weighting function used in connection with Figures 14A and 14B;

Figure 15 is a block diagram of digital phase discrimination hardware in accordance with the technique of Figure 14;

Figure 16A is a tabulation illustrating still another method of digital phase discrimination;

Figure 16B is a plot showing results of the method of Figure 16A;

Figure 16C is a plot of a weighting function used in connection with Figures 16A and 16B;

Figure 17 is a block diagram of digital phase discrimination hardware in accordance with the technique of Figure 16.

Figure 18 is a block diagram illustrating a sampled-data model of a Sigma-Delta modulator and of a sampling circuit applied to frequency sampling in accordance with one embodiment of the present invention;

Figure 19 is a table helpful in explaining operation of the circuit model of Figure 18 in the instance of an input frequency that is 0.6875 times a reference frequency;

Figure 20 is a timing diagram illustrating the principle of operation of the circuit model of Figure 18 as applied to frequency sampling;

Figure 21 is a schematic diagram of one example of a frequency sampling circuit described by the circuit model of Figure 18.

Referring now to Figure 4, a block diagram of an RF modulator in accordance with the present invention is shown. A portion of the modulator indicated within dashed lines is substantially the same as a corresponding portion of the conventional modulator of Figure 3. Optionally, the modulator of Figure 4 may

include a downconverter 420 consisting of a frequency synthesizer, a mixer and a low pass filter. In instances where the output frequency of the modulator is very high, the downconverter reduces the frequency to one that may be more readily handled by the divide by N counter. Note, however, that in other instances N may be equal to one. Also, the modulator of Figure 4 may optionally include an amplitude modulation stage 421. In some modulation schemes, both the amplitude and phase characteristics of the signal may be modulated, in which instance the amplitude modulation stage is used to impart the desired amplitude modulation.

Preferably, instead of the VCXO of Figure 3, the modulator of Figure 4 uses a Direct Digital Synthesis (DDS) circuit 415 to produce a modulated input signal that is applied as a reference signal to the PLL. Such a DDS circuit may be based upon a known Number Controlled Modulated Oscillator (NCMO) described in U.S. Patent No. 4,746,880 issued May 24, 1988, entitled Number Controlled Modulated Oscillator, incorporated herein by reference. A baseband modulation signal is applied to the DDS circuit to produce the modulated input signal, which is then applied as a reference signal to the PLL.

The stability and precision of the present modulator derive from an additional feedback loop including a phase demodulator 419 and a comparison circuit 417. The phase demodulator may be of a type described in U.S. Patent Application No. 09/006,938 (Atty. Dkt. No. 32219-003), entitled DIGITAL PHASE DISCRIMINATION BASED ON FREQUENCY SAMPLING, filed January 14, 1998, incorporated hereinafter.

The phase demodulator 419 recovers phase information from the output signal of the PLL. Using a comparison circuit 417, the recovered phase information is compared to the phase information of the baseband modulation signal. A resulting error signal is used to control injection of the baseband modulation signal into the PLL, to automatically achieve the correct "dosage." That is, an output signal of the comparison circuit 417 sets the gain of the adaptive gain amplifier 411 to

the value required in order to match the phase information carried by the output signal to the desired phase information. The painstaking adjustment process characteristic of the prior art is eliminated. Furthermore, the adaptation process performed by the present modulator is continuous, operating whenever the modulator is turned on. The effects of environmental conditions, component variations, component aging, etc., can all be compensated for.

The phase demodulator of U.S. Patent Application No. 09/006,938 (Atty. Dkt. No. 32219-003), entitled DIGITAL PHASE DISCRIMINATION BASED ON FREQUENCY SAMPLING, filed January 14, 1998, will now be described.

The approach followed by the digital frequency discriminator of the present invention may be appreciated by analogy to Sigma-Delta A/D conversion, well-documented in the prior art by such references as "Oversampling Delta-Sigma Data Converters", Candy, et al., *IEEE Press*, pages 1-6, Piscataway, NJ (1992). A Sigma-Delta converter modulates a varying-amplitude analog input signal into a simple digital code at a frequency much higher than the Nyquist rate. The design of the modulator allows resolution in time to be traded for resolution in amplitude. A sampled-data circuit model of a Sigma-Delta modulator, shown in Figure 1, may be directly applied to frequency sampling as described herein.

Referring to Figure 18, an input signal x_i occurring at sample time i has subtracted from it the output signal y_i at sample time i . The result is applied to an accumulator having an output signal w_i . A "new" input signal of the accumulator at sample time i is combined with the "old" output signal of the accumulator to form a new output signal of the accumulator. The output signal of the accumulator is quantized, the quantization being represented as the addition of an error e_i . The output signal of the quantizer is the final output signal y_i .

Assume now that x_i is the ratio of two frequencies and that the quantizer is a two-level quantizer. Further assume that the ratio of the two frequencies for the time period in question is, say, 0.6875. As shown in Figure 19, the latter value is accumulated a first time, giving an accumulated value of 0.6875. This valuing being less than 1, the value 0.6875 is again added to the accumulated value, giving a new accumulated value of 1.375. Since this value is now greater than 1, 1 is subtracted from 0.6875 and the result ($0.6875 - 1 = -0.3125$) added to the accumulator

to give a value of 1.0625. Operation proceeds in this fashion. During the foregoing sequence of operations, a data stream is produced by taking the integer portion, 1 or 0, of each accumulated value.

Referring to Figure 20, the interpretation of the sequence of numbers shown in Figure 19 may be appreciated. Two clock signals are shown. Again, it is assumed that the ratio of the frequency of the upper clock signal to that of the lower clock signal during the period of interest is 0.6875. At time $t = 0$, rising edges of both clock signals coincide. At the first subsequent rising edge of the lower clock signal, 0.6875 periods of the upper clock signal have elapsed. At the next rising edge of the lower clock signal, 1.375 periods of the upper clock signal have elapsed. At the next rising edge of the lower clock signal, 1.0625 periods of the upper clock signal have elapsed since the elapse of the first period of the upper clock signal, and so on.

A schematic diagram of a capture circuit, or frequency sampling circuit, that may be used to data samples corresponding to the data stream described in the foregoing example is shown in Figure 21. In the illustrated embodiment, it is assumed that the ratio of the clock signals is such that no more than one rising edge of the faster clock will occur during a single period of the slower clock. In other embodiments, this assumption need not apply.

The capture circuit includes an input portion 2101 and an output portion 2103. The input portion includes two sections Ch1 and Ch2 that must be carefully matched to minimize errors. Each section comprises a chain of two or more D flip-flops coupled in series. In the following description, the same reference numerals will be used to reference the respective flip-flops themselves and their respective output signals.

Within each section, the first flip-flop in the chain is clocked by a sampled clock signal F_x . The succeeding flip-flops in the chain are clocked by a sampling clock signal F_s . The D input of the first flip-flop Q1 in the upper section is coupled to the $\overline{Q}$ output of the same. The D input of the first flip-flop in the lower section is coupled to the Q output of the first flip-flop in the upper section. The remaining flip-flops in both sections are coupled in series--i.e., Q to D, Q to D.

The function of the input portion is to 1) produce two signals, logical inverses of one another, that transition on rising edges of the clock signal F_x ; 2) to latch the values of the two signals on the rising edge of the clock signal F_s ; and 3) to detect transitions from one clock to the next. Additional intermediate stages in

series with Q3 and Q4 may be required to minimize metastability resulting from the asynchrony of the two clock signals, and in fact multiple such stages may be desirable in a particular design.

The output portions include, in an exemplary embodiment, three two-input NAND gates. Respective NAND gates N1 and N2 are coupled to the D and $\bar{Q}$ signal of the final flip-flop stages of the input sections. Output signals of the NAND gates N1 and N2 are combined in the further NAND gate N3 to form the final output of the capture circuit.

The function of the output portion is to detect a change in the input clock signal level from one sample clock to the next in either of two channels formed by the two input sections. The two input sections function in a ping-pong fashion, alternately detecting changes in the input clock signal level.

Operation of the capture circuit of Figure 21 may be more fully appreciated with reference to the timing diagram of Figure 5. The first stages of the two channels form inverse signals Q1 and Q2 approximately coincident with (but slightly delayed from) rising edges of the input clock signal. The signals Q3 and Q4 are formed by sampling the signals Q1 and Q2, respectively, in accordance with the sample clock. The signals Q5 and Q6, respectively, are delayed replicas of the signals Q3 and Q4. The NAND gates together realize the logic function $X = Q3 \cdot \bar{Q5} \vee Q4 \cdot \bar{Q6}$.

In the example of Figure 5, the illustrated signals are all idealized square-wave signals. In actuality, the signals will have finite rise and fall times. The possible effect of the finite rise and fall times of the signals Q1 and Q2 and the asynchrony of the circuit is metastability, as illustrated in Figure 6. Here, the signals Q3 and Q5 and the signals Q4 and Q6 are each in an indeterminate state for one cycle. The resulting output of the circuit may or may not be correct. However, because the decision was a "close call" to begin with, the effect of an occasional erroneous decision on the overall operation of the circuit is negligible. The time window of instability is reduced by increasing the overall gain in the path. If the gain in Q3 and Q9 is sufficient to reduce the probability of an error to an acceptable level, then no additional circuitry is required. If not, then additional circuitry will be required to increase the gain.

In order to recover the ratio of the frequencies of the two clock signals from the data stream produced by a capture circuit such as the one of Figure 21, digital

filtering is applied. Advantageously, an extensive body of digital filtering techniques applicable to Sigma-Delta (or Delta-Sigma) A/D converters may be applied directly to the digital stream. Furthermore, by using an appropriately-chosen weighting function, high accuracy may be obtained.

The weighted sum of products is an example of an FIR filter. The weighting function described heretofore is therefore that of an FIR filter in digital filtering theory. It should be recognized, however, that IIR filters can also be used. In the process of FIR digital filtering, the weighting function is applied to a "window" of data samples to obtain an estimate of the ratio of frequencies in the center of the window. The window is then "picked up and moved" to the next sequence of samples. Windowing will typically overlap. A window may include 256 samples, for example.

Referring to Figure 7, two alternative weighting functions are shown for a window of 256 samples. The weighting functions are normalized, meaning that the area under the weighting function is unity. One weighting function, indicated in dashed lines, is a straight-line, constant weighting function. Another weighting function, indicated in solid line, is a triangular weighting function. The weighting function is the impulse response function in digital filters.

Results of digital filtering using the straight-line weighting function and the triangular weighting function respectively, are shown in Figure 8 and Figure 9. In the case of both Figure 8 and Figure 9, the frequency ratio was increased from just under 0.687 to just over 0.693. As seen in Figure 8, using a straight-line weighting function, the quantized signal oscillates between two levels that are adjacent to the input in such a manner that its local average equals the average input. The average error was calculated to be 1772ppm. As seen in Figure 9, using a triangular weighting function, the quantized signal tracks the input with an average error of 83ppm.

A schematic diagram of an exemplary frequency accumulator that applies a triangular weighting function and that may be used to accomplish the desired digital filtering is shown in Figure 10. In the example shown, the frequency accumulator uses a 7-bit counter 101, a 14-bit adder 103 and a 14-bit register 105. The 7-bit counter is clocked by the sample frequency F_s . The output of the 7-bit counter is provided to one input of the adder. The function of the 7-bit counter is to count up from 0 to 127 and then down from 127 to 0. The count of 127 occurs twice in succession. This behavior is achieved using a flip-flop 107. The flip-flop is clocked

by the sample frequency F_s . A Terminal Count signal of the 7-bit adder is input to the flip-flop. The output of the flip-flop is coupled to a Count Down input of the 7-bit counter.

The "oversampled" data stream is coupled to a control input of the adder. When the current bit of the data stream is a 1, an addition is performed. When the current bit is a 0, no addition is performed. A Carry In input of the adder is tied high, effectively causing the range of weights to be 1 to 128.

The 14-bit register is clocked by the sample frequency F_s . Its output is applied to the other input of the adder. Its input receives the output word produced by the adder. The function of the 14-bit adder is to perform an accumulation operation for 256 clocks. At the conclusion of the 256 clocks, the output of the 14-bit adder is used as an estimator for the frequency ratio. More particularly, in the example shown, the output of the accumulator is equal to $R \times 128 \times 129$, where R is the frequency ratio estimator.

The foregoing technique may be readily extended to phase discrimination. Various different methods and apparatus for digital phase discrimination will be described entailing different design tradeoffs.

The first method is conceptually straightforward but computationally expensive. Referring to Figure 11A, the same observed frequency data stream and the same set of weights corresponding to a triangular weighting function (Fig. 11B) are used. The ratio of the reference frequency to the sampled frequency over a relatively long period of time is first determined using the technique described previously. Having obtained this frequency ratio estimator, short-term frequency deviations are estimated by calculating the same frequency estimate as before but at a relatively high rate, as often as once per sample period. That is, successive samples are all taken using the circuit of Figure 10, as often as each sample period. The difference (ΔF) of each frequency estimate (F) from the previously-determined frequency ratio (F_r) is calculated, multiplied by an appropriate scale factor k and accumulated to obtain a corresponding phase estimate P_f . (The first value of P_f is an arbitrarily chosen initial condition, chosen for comparison to an ideal estimate. In practice, the phase may be initialized to a value based on a priori knowledge of signal characteristics, or, absent such a priori knowledge, may be set to zero upon detection of a phase inflection point.)

A phase-plot simulation comparing actual phase of a specified waveform (solid-line) with estimated phase using the foregoing phase estimation method

(dashed-line) is shown in Fig. 11C.

The foregoing "frequency difference" phase estimation method is computationally expensive because of the need to calculate frequency estimates at a relatively high rate. A "pre-summation difference" phase estimation method obviates this requirement. Referring to Figure 12A, instead of subtracting the frequency ratio from a frequency estimate, the frequency ratio Fr is subtracted from the sampled data stream itself. Assuming that the data stream is a bit stream of ones and zeros only, and assuming a frequency ratio $Fr = 0.6875$, then the pre-summation difference Y will have one of only two values, $Y = 1 - 0.6875 = .3125$ or $Y = 0 - 0.6875 = -0.6875$. The Y values are accumulated to obtain corresponding values PX . Phase estimates PP_n are obtained by filtering the PX values in substantially the same manner as described previously in relation to forming frequency estimates (using the identical weighting function, Fig. 12B, for example) with the exception that the filtered values are scaled by the scale-factor k .

The pre-summation difference phase-calculation may be shown to be mathematically equivalent to the frequency difference phase calculation. Simulation results, shown in Fig. 12C, are therefore the same as in Fig. 11C. The hardware realization, however, may be considerably simpler using the pre-summation difference phase calculation, since only one computation is required per phase point. Such a hardware realization is shown in Figure 13.

The pre-summation difference phase estimator of Figure 13 includes generally a first accumulator ACC1, a weight generator WG similar or identical to the weight generator previously described in relation to Figure 10, and a second accumulator ACC2.

The accumulator ACC1 functions to produce phase numbers PX_i in correspondence to bits (or in other embodiments, symbols) of the observed frequency data stream and includes a multiplexer 1301, an adder 1303 and a register (e.g., a 16-bit register) 1305. The multiplexer 1301 selects one of the two possible values of Y_i in accordance with the value of X and applies Y_i to the adder 1303. The register value is added to Y_i to form PX_i , which is then strobed into the register. The adder 1303 and register 1305 therefore accumulate the PX_i values.

The PX_i values are then filtered in the accumulator ACC2, which includes a multiplier 1307, an adder 1309 and a register 1311. The multiplier receives weights from the weight generator WG and PX_i values from the accumulator

ACC1. Respective weights and PX_i values are multiplied and the products accumulated, e.g., for 128 clock cycles, to produce a phase estimator PP. The multiplier may be constructed so as to apply the scale factor k to each product during the accumulation process.

An even simpler realization may be achieved using an integer difference phase calculation. The integer difference phase calculation is not mathematically equivalent to the foregoing methods, but is very close. Referring to Fig. 14A, this method uses, in addition to the observed frequency data stream, a reference frequency data stream that would result if the reference frequency were applied to the capture circuit of Figure 4 (with the same clock). A running sum D_i is then formed of the integer difference $X_i - R_i$. In many practical applications, such as the one illustrated in Figures 14 and 15, D_i will have the values 1, 0 and -1 exclusively. The general case in which D_i takes on other values may be appreciated and understood, however, from the present example, and is embraced by the present description.

Phase estimates are formed by filtering the D_i values in the same or similar manner as previously described. The same triangular weighting function may be used Fig. 14B. The integer difference phase calculation method produces identical simulation results, Fig. 14C, as the preceding methods.

Referring to Figure 15, in the instance where D takes on the values 1, 0 and -1 exclusively, the corresponding hardware realization may be substantially simplified (as compared to that of Figure 13, for example).

The integer difference phase estimator of Fig. 15, like that of Fig. 13, includes generally a first accumulator ACC1, a weight generator WG, and a second accumulator ACC2. The accumulator ACC1 is of considerably different construction than the corresponding structure of Figure 13. The accumulator ACC1 of Figure 15 includes a reference pattern generator 1501, a 1-bit subtractor 1503, a 2-bit adder 1505 and a 2-bit register 1507. The 1-bit subtractor subtracts respective R values from respective X values. The 2-bit adder and the register accumulate the resulting D_i values which, as explained previously, may be constrained to 1, 0, -1 only.

The weight generator WG and the accumulator ACC2 are substantially the same as in Figure 13, described previously. However, because D_i takes on the val-

ues 1, 0 and -1 exclusively, no multiplier is required. Instead, if $D_i = 1$, the weight value is added to the accumulated value, and if $D_i = -1$, the weight value is subtracted. (If $D_i = 0$, the accumulated value remains unchanged.) The savings of a hardware multiplier is a particular advantage of the implementation of Fig. 15.

A further method of phase estimation is referred to as the clock measure phase calculation method. Referring to Figure 16A, this method is similar to the previous integer difference phase calculation method insofar as R, X and D are concerned. This method, however, uses in addition to the reference frequency data stream R, "clock measure" numbers RG, which are the same as the numbers appearing in Fig. 2. Moreover, the weight function used is distinctly different, as shown in Figure 16B. Clock measure phase estimate values PC are obtained using the following formula:

$$PC_n = k \cdot \left(D_n - \text{frac}RG_n + 0.5 + \sum_i (W_i \cdot X_{i+n-64}) \right)$$

Simulation results using the clock measure phase calculation method are shown in Figure 16C.

Referring to Figure 17, the clock measure phase estimator includes generally a first accumulator ACC1, a weight generator WG, and a second accumulator ACC2. The estimator additionally includes a summation block 1701.

The accumulator block ACC1 is substantially the same as the accumulator block ACC1 of Figure 15. Note, however, that the reference pattern generator generates both the reference frequency data stream R, used within the accumulator ACC1, and the clock measure data stream RG which is input to the summation block 1701.

The weight generator includes a counter 1703 and weight generator logic 1705.

The accumulator ACC2 includes an adder 1707 and a register 1709. When $X = 1$, the weight value from the weight generator is added to the contents of the register 1709. The output of the adder becomes the new input of the register, which performs an accumulator operation for, e.g., 128 clock cycles.

At the conclusion of the accumulation operation of ACC2, the outputs of ACC1 and ACC2, together with the corresponding RG value, are summed in the summation block 1701.

What is claimed is:

1. An RF modulator comprising:
a phase lock loop including a loop filter and receiving as an input signal a baseband modulation signal and producing as an output signal a modulated RF signal;
circuitry for producing an injection modulation signal;
circuitry for injecting the injection modulation signal into the phase lock loop at a point past the loop filter; and
control circuitry, coupled to the circuitry for injection the injection modulation signal, for controlling the amplitude of the injection modulation signal.
2. The apparatus of Claim 1, wherein said circuitry for injecting comprises an adaptive gain amplifier.
3. The apparatus of Claim 2, wherein said control circuitry includes a phase demodulator.
4. The apparatus of Claim 3, wherein said phase demodulator is digital.
5. The apparatus of Claim 4, wherein said control circuitry further comprises a comparison circuit.
6. The apparatus of Claim 5, wherein the comparison circuit compares the baseband modulation signal to a phase information signal recovered by the phase demodulator and produces an error signal.
7. The apparatus of Claim 6, wherein said error signal is used to set the gain of the adaptive gain amplifier.



8. The apparatus of Claim 1, further comprising a direct digital synthesis circuit for receiving the baseband modulation signal, producing a corresponding modulated signal, and applying the modulated signal to the phase lock loop.

9. The apparatus of Claim 8, wherein the direct digital synthesis circuit include a number controlled modulated oscillator.

10. The apparatus of Claim 1, wherein said phase lock loop includes a divide by N counter.

11. The apparatus of Claim 1, wherein said phase lock loop includes a downconverter.

12. The apparatus of Claim 1, further comprising a amplitude modulation stage coupled to the output signal of the phase lock loop.

13. A method of producing precise, stable phase shifts in a modulated RF signal, comprising the steps of:

using a baseband modulation signal, producing a modulated signal;
applying the modulated signal as a reference signal to a phase lock loop; and

automatically deriving an adaptive gain factor and applying the adaptive gain factor to the baseband modulation signal to produce an injection signal and injecting the injection signal into the phase lock loop;

wherein the injection signal compensates for the phase lock loop partially removing modulation introduced by the baseband modulation signal.



14. The apparatus of Claim 13, wherein automatically deriving an adaptive gain factor comprises:

recovering phase information from an output signal of the phase lock loop;

comparing the phase information to the baseband modulation signal and producing an error signal, the adaptive gain factor being derived from the error signal.

