SYNTHESIZER DESIGNS MINIMIZE PHASE NOISE IN CELLULAR SYSTEMS

Performance and cost tradeoffs are analyzed for a variety of single- and multiple-loop architectures.

As cellular communication systems evolve from analog-FM to fully-digital implementations, attention is being focused upon frequency-synthesizer approaches that deliver adequate spectral purity and cost effectiveness. A variety of design approaches are available that satisfy the stringent synthesizer requirements of time-domain-multiple-access (TDMA)-based systems.

While single-loop synthesizer designs have proven adequate for analog cellular applications, digital cellular systems, in particular TDMA schemes using $\pi/4$-DQPSK (differential quadrature phase-shift keying), have considerably more rigid design margins. In comparison to emerging code-division-multiple-access (CDMA)-based systems, phase noise has a more direct bearing on (uncoded) residual bit error rate (BER) in TDMA schemes. In addition, TDMA requires a fairly rapid frequency-switching speed in order to monitor alternate channels.

Stringent discrete-spurious requirements are imposed upon cellular equipment in order to prevent adjacent-channel interference. Discrete spurious performance in the receive mode is also important in the minimization of receiver desensitization in the presence of strong adjacent signals. Typical discrete spurious requirements developed by the Electronic Industries Association (EIA) are –60 and –75 dBc below and above 45-kHz offset, respectively.[1]

Given the normal 30-kHz channel spacing, it is often more difficult to meet the spurious requirements at 60-kHz offset than at 30-kHz offset because of more stringent specifications. Moreover, any additional loop filtering must normally first handle the spurious components at 30 kHz.

Phase-noise requirements for TDMA digital cellular operation must be based in principle upon a system-wide error budget which considers a wide range of design factors. All systems must also meet certain spectral-mask guidelines in order to prevent interference with other stations. Many factors contribute to the overall system BER performance, including demodulator type (e.g., coherent or noncoherent differential), quantization effects, receive-filter mismatch, post-detection processing, and phase noise. Although the emerging US cellular standard employs $\pi/4$-DQPSK, other modulation types, such as 8-PSK and 16-QAM (quadrature amplitude modulation), are also being considered for specialized services. The added throughput offered by these waveforms further complicates the phase-noise issue.

A simple (yet overly pessimistic) measure of synthesizer phase-noise

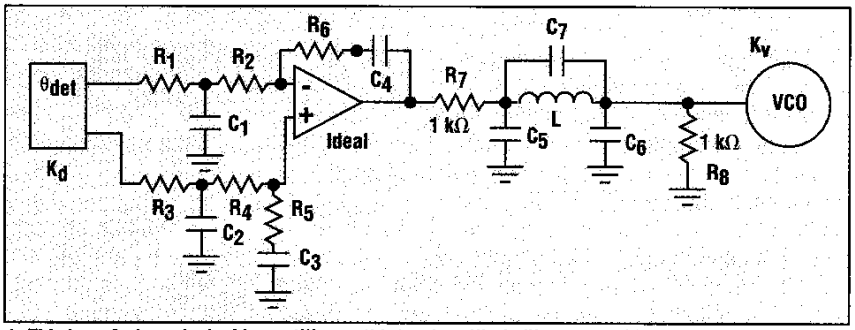

1. This type-2 phase-locked loop utilizes a third-order elliptic filter to provide excellent spurious rejection.

JAMES A. CRAWFORD, President, ComFocus Corp., 6160 Lusk Blvd., Suite C-200, San Diego, CA 92121; (619) 455-9934, FAX: (619) 455-1259.

performance can be obtained by integrating the output phase-noise spectrum over the matched-filter bandwidth to arrive at a root-mean-square (RMS) phase-jitter value for the synthesizer noise. It is not unusual to specify this value at less than 5 deg. for the π/4-DQPSK waveform. However, the actual value should be based upon the wide-ranging systems issues. If frequency-discriminator detection is used in the receiver, a more suitable measure of synthesizer phase-noise performance is the RMS phase accumulated over an individual symbol period. This may be calculated as:

$$\sigma^2 = E\{\theta(t-T)\,\theta(t)\}$$

where:

T = symbol period,

E{} denotes statistical expectation, and

θ = instantaneous carrier phase.

In general, phase-noise performance becomes increasingly demanding as the signal-constellation complexity is increased.

TYPE-2 PLL

Before proposing any new design architectures, it is helpful to examine the role of traditional approaches in the present TDMA cellular context. The type-2 phase-locked loop (PLL) has been the traditional approach to frequency-synthesizer design for many years.[2,3]

A typical type-2 PLL (Fig. 1) utilizes a three-state phase detector.[4,5] The normal post-detection RC filtering precedes the lead-lag filtering in order to limit the slew rate into the operational amplifier to function as an approximation to a zero-order hold.[6] Although this type of phase detector can exhibit excellent spurious performance, additional spurious rejection is necessarily provided by the third-order elliptic filter immediately preceding the voltage-controlled oscillator (VCO). A PLL case study with the following design parameters can be examined:

Reference frequency (F_{ref}) = 1 MHz

Final feedback divider ratio (N_{stop}) = 100

Damping factor (ζ) = 0.75

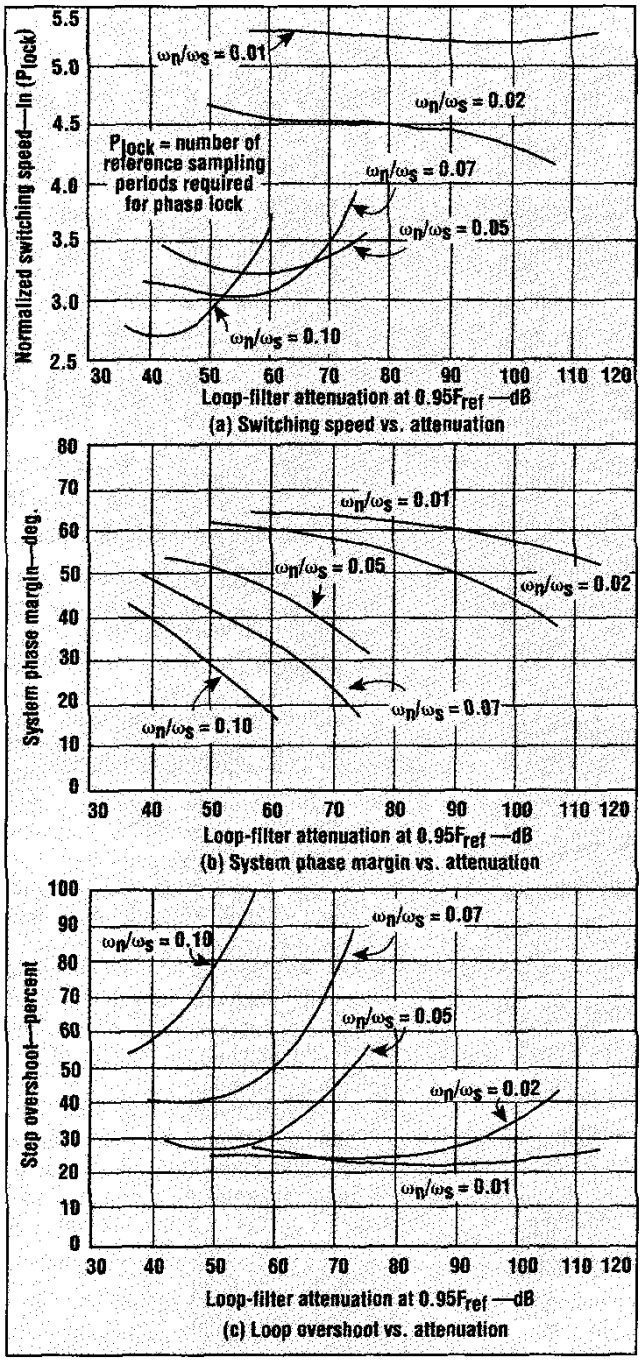

2. These plots show switching-speed (a), phase-margin (b), and VCO frequency-overshoot (c) simulation results for a type-2 PLL case study.

VCO tuning sensitivity (K_v)= 25 MHz/V

Phase-detector gain (K_d) = 0.8 V/radian

In Fig. 1, resistors R_1 and R_4 are taken to be 1 kΩ, while capacitors C_1 and C_2 are assumed to be 500 pF. Using these values, a number of transient-response simulations were run in order to estimate the switching time required for different loop bandwidths and varying degrees of reference-spur filtering. The results of this study can be normalized for any desired F_{ref} (e.g., 30 kHz for digital cellular).

The frequency step (ΔF) used to examine each loop's transient response was generated by synchronously changing the feedback di-

SYNTHESIZER DESIGN

vider ratio from a specified initial value (N_{start}) to N_{stop} = 100. As the loop bandwidth was increased, N_{start} was adjusted so that all of the switching speeds could be based upon a normalized loop bandwidth. This permitted clear observation of the elliptic filter's impact upon switching speed. The integrator time constant (τ_2) and K_v were also adjusted in order to keep ζ equal to 0.75 at N_{stop} = 100. Capacitors C_3 and C_4 were both taken to be 1000 pF.

PERFORMANCE PLOTS

Graphical results were obtained for the compiled simulation. The key design parameter for spurious performance is normally the decibel level below the carrier of the first discrete sampling spur, which for 30-kHz channel spacing would generally occur at 30-kHz offset. For this reason, all curves are plotted with respect to the loop-filter attenuation at 0.95 F_{ref} (where F_{ref} is the position of the elliptic filter zero). This characterization requires independent assessment of the phase detector's raw spurious performance in order to determine the actual spurious level at the synthesizer output.

Figure 2a shows the normalized switching speed with respect to the loop-filter attenuation at 0.95 F_{ref}. The switching speed is quantified in terms of the natural logarithm of the number of reference sampling periods required to achieve phase lock. In all cases, the magnitude of the frequency hop was assumed to be small enough so that no cycle slipping occurred at the phase detector (i.e., linear operation). The definition of phase lock was taken to be the point in time at which the VCO output phase error was less than or equal to 5 deg. with respect to the final output steady-state phase. As the plot shows, the additional discrete spurious attenuation is a strong function of the closed-loop bandwidth. As shown for all cases, a small amount of additional filtering by the elliptic filter provides a slight improvement in switching speed. Excessive filtering, however, leads to increased locking times and decreased system-stability margins.

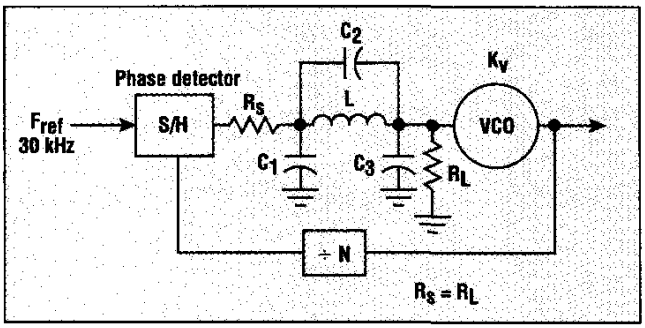

3. This loop configuration was used to study the performance of type-1-PLL-based frequency synthesizers.

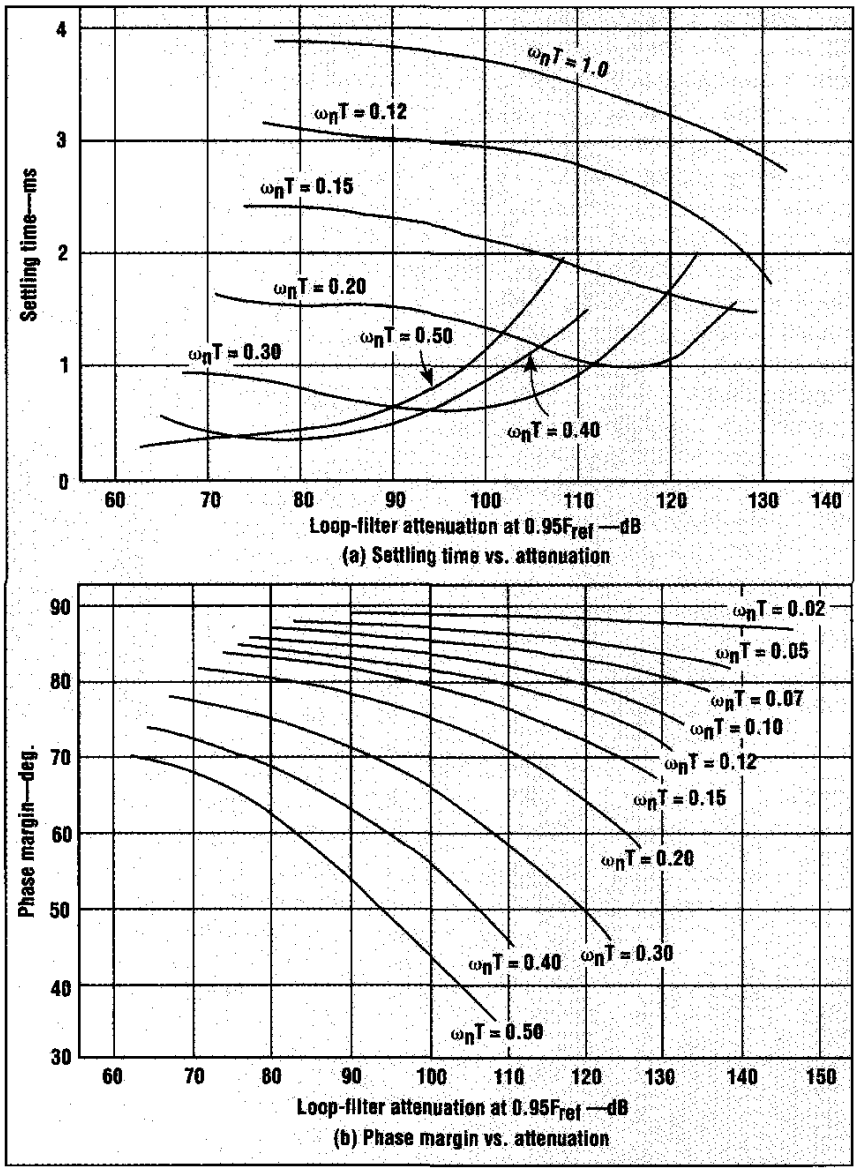

4. Settling-time (a) and phase-margin (b) performance plots were obtained for a type-1 PLL synthesizer.

In actual practice, this data cannot be viewed independently of the N_{start} value at each $\omega_n T$ because each loop case is frequency-hopped by a different amount. Additional switching time must also be included for the nonlinear portion of the switching transient present with larger frequency hops.

In Fig. 2b, the closed-loop phase margin is plotted with respect to the corresponding loop-filter attenua-

tion parameter. Although loop stability is an obvious requirement, a reasonable phase-margin value is necessary because it influences the amount of VCO frequency overshoot. In steady-state, this value directly affects how much VCO phase-noise "peaking" will be observed in the vicinity of the system's natural frequency (ω_n). A large amount of peaking leads to unacceptable phase-noise shoulders in the output spectrum.

Figure 2c shows the amount of VCO frequency overshoot for each loop case with respect to the loop-filter attenuation parameter. A large overshoot is undesirable because it indicates excessive lock-up time. In addition, if the overshoot results in saturation of the control-loop electronics (e.g., loop-filter output put on a voltage rail), substantial switching-speed decrease will be observed in comparison to the data of Fig. 2a because the linearity assumption will be violated.

Phase-noise performance comparisons for these cases require additional assumptions to be made regarding the VCO phase-noise spectrum and the noise floor of the control-loop electronics (loop filter, feedback divider, etc.). If the phase-noise contributions from the VCO self-noise and the loop-electronics noise floor are viewed as statistically-independent quantities, each contribution can be evaluated independently, with the results combined on a variance (i.e., power) basis.

In order to pursue this analysis, the closed-loop transfer functions corresponding to the PLL of Fig. 1 must be computed. The most tedious part of this task is deriving the transfer function for the third-order elliptic filter. This analysis is simplified if it is assumed that the elliptic filter is designed to accommodate an infinite load impedance rather than designed to be equally-terminated. With this assumption, a complete example computation for the closed-loop transfer functions can be performed easily using an analysis program such as MathCad from MathSoft, Inc. (Cambridge, MA).

The combined spurious, phase-noise, and switching-speed requirements are difficult to meet with off-the-shelf, highly-integrated PLL devices. As a result, a number of semiconductor suppliers are considering new approaches to deal with this design problem.

TYPE-1 PLL

Traditionally, type-2 PLLs have been used in most frequency-synthesizer designs. Prior to the availability of highly-integrated PLL devices, this approach was easily justified based on circuit complexity, let alone other factors, such as post-tuning drift.

Loop parameter values for a type-1 PLL case study

ω_n/F_{ref}	K_v (MHz/V)
0.02	2.865
0.05	7.162
0.07	10.027
0.10	14.325
0.15	21.4875
0.20	28.65
0.30	42.975
0.40	57.3
0.50	71.62

Note: K_d = 2.0 V/radian.

A number of manufacturers provide highly-integrated PLLs with a selection of phase-detector types. The phase detector of interest here is the linear sample-and-hold (S/H) phase detector. A type-1 PLL using this phase detector is the optimal choice for very-high-speed switching applications since it can theoretically achieve frequency-lock in only one reference sample period.[7,8]

Some caveats of the type-1 architecture are worth noting since most designers are not familiar with the approach. First, the S/H phase detector must generally be augmented with additional frequency-discrimination circuitry to achieve guaranteed frequency pull-in. This can be avoided if the phase-detector beatnote frequency is always well within the system's closed-loop bandwidth. Second, although normally only a problem at much higher VCO frequencies with switching speeds below 50 μs, the type-1 architecture is not nearly as immune to VCO post-tuning drift problems as the type-2 scheme. Post-tuning drift, however, should not be an issue with most cellular applications.[9]

TYPE-1 STUDY

A case study was done for a type-1 frequency synthesizer designed with a third-order elliptic lowpass filter for improved spurious performance (Fig. 3, Table).

Unlike the type-2 PLL, capture range and closed-loop bandwidth cannot be independently specified in a type-1 system. This made examination of very-small closed-loop-bandwidth systems difficult, even when employing computer simulation. To avoid VCO pretuning, the lower-bandwidth limit for type-1 systems is often set by the output-tuning bandwidth which determines the required loop-capture range.

The type-1 analysis was performed using a 30-kHz F_{ref}, with settling time plotted in milliseconds rather than as a function of sample periods. A fixed frequency hop of 10 MHz was assumed for the transient analysis. As shown in Fig. 4a, a wide range of loop bandwidths and elliptic-filter combinations can theoretically provide switching speeds below 1 ms. Loop-filter attenuations on the order of 110 dB at 0.95 F_{ref} are also attainable, which is in sharp contrast to the type-2 unnormalized performance. The type-1 system also has an inherently larger phase margin than the type-2 PLL scheme and, as Fig. 4b illustrates, is considerably more tolerant to additional filtering. In comparison to the type-2 architecture, VCO frequency overshoot during the transient response is almost nonexistent up to very high filtering levels.

Referring to Fig. 4a, a loop design utilizing an ω_n/F_{ref} of 0.30 at an attenuation level of 100 dB would be an excellent design starting point. The closed-loop bandwidth is reasonably small at 1432 Hz, which leads to a fairly small phase-noise pedestal

around the carrier. Nevertheless, the design margins for phase noise and sampling spurs are not large and therefore do not necessitate critical design consideration.

The principal drawback of the type-1 design or any basic single-loop approach is that a fairly substantial closed-loop bandwidth is required in order to obtain 1-ms switching speed. At a synthesized output frequency of 900 MHz, the large feedback divider ratio N (on the order of 30,000) leads to fairly high phase-noise levels within the loop bandwidth, thereby affecting BER performance. This noise contribution can only be reduced by decreasing N, which in effect requires an alternate synthesizer architecture.

DESIGN OPTIONS

Frequency-synthesizer design for next-generation TDMA, digital cellular systems must achieve a careful balance between cost and performance. Since market quantities are very large, special-purpose rather than general-purpose design approaches are acceptable. A number of synthesizer topologies are suitable for different levels of cellular service.

The most simple synthesizer option is the single-loop design, which has been used extensively throughout the analog-FM cellular world (Fig. 5a). For next-generation service, this architecture suffers from marginal phase-noise and switching-speed performance, particularly when using the traditional type-2 approach. Through careful attention to detail, it may be possible to achieve the needed near-term requirements in a type-1 architecture, but design margins may be unacceptably small.

The principal cost elements in any cellular synthesizer, in order of decreasing cost, are generally the VCO, the loop-filter electronics, and the phase-detector/feedback-divider electronics. With this in mind, a sensible next-generation alternative is the fractional-N approach (Fig. 5b), which requires only one VCO. The peak fractional-N spurious component can have a peak phase deviation of up to 2 radians. Therefore, in order to obtain –60-dBc fractional-N spurs, a minimum of 60-dB loop-filter attenuation would be required if no feedforward fractional-N correction is included. In general, the feedforward correction is a difficult function to implement in highly-integrated form. Without the feedforward correction, excessive loop-filtering requirements will substantially limit the achievable switching speed.

A less aggressive single-loop alternative to the full fractional-N implementation is the use of a single-sideband (SSB) mixer to offset the VCO frequency in 30-kHz incre-

ments, thereby filling in the channels at 300-kHz increments (Fig. 5c).

The mixer's unwanted sideband suppression does not have to be very good in order to avoid miscounts by the feedback frequency divider. The closest spurious component which must be combatted by this architecture is offset 30 kHz from the carrier and arises from carrier leakage through the mixer. The primary spurious term which must be eliminated is caused by the other principal mixing product, which is offset 60 kHz from the carrier.

This approach is based upon the fundamental notion that since the offset frequencies and F_{ref} are integral multiples of 30 kHz, the sampling theorem dictates that aliasing effects will not produce any spurious components closer than 30 kHz from the carrier. Loop-filtering requirements may be reduced decibel-for-decibel depending on the performance of the SSB mixer. Since the percentage-bandwidth tuning range for digital cellular is quite small, excellent sideband-mixer performance should be possible. Creation of the offset frequency from 30 to 270 kHz should be a very-low-cost procedure given modern CMOS technology.

A closely-related design alternative is an offset concept where the SSB mixer is replaced by a conventional mixer and the low-frequency offset source is replaced by an RF source in the 900-MHz range (Fig. 5d). Depending upon the overall frequency plan for the cellular equipment, it may be convenient to create the additional RF source using an injection-locked oscillator. This approach leads to exceptional phase-noise performance due to the small N. However, since the channelization is still driven by F_{ref} (30 kHz), switching speed may still be a problem. In general, some form of gain compensation would also be required within the control loop to compensate for large changes in N. Otherwise, unacceptable variations in ζ and ω_n could be experienced.

MULTIPLE LOOPS

A dual-loop option utilizes two separate PLLs, with each loop using a different F_{ref} to synthesize 30-kHz channel steps (Fig. 5e). This results in an obvious increase in complexity. However, the approach can be used to easily synthesize 5- or 10-kHz channel steps while retaining the fairly high F_{ref}. This provides fast switching speed and improved close-in phase-noise performance. It is important to note that in this architecture the phase noise which is far-removed from the carrier is solely determined by the phase-noise performance of VCO1. Since the lower-loop portion of the synthesizer is only involved with the close-in phase-noise performance of the synthesizer, VCO2 can be made completely monolithic. Injection levels from VCO2 into the mixer may be as low as –30 or –40 dBm in principle since close-in phase-noise levels are generally many decibels above the Boltzmann-level-induced phase noise. This fact can be used to minimize oscillator isolation problems.

Finally, a triple-loop configuration can be examined (Fig. 5f). In principle, PLL1 can be replaced by a direct digital synthesizer (DDS), but the phase-locked approach produces lower spurious levels, generally requires less power, and is less expensive in monolithic form. The phase-

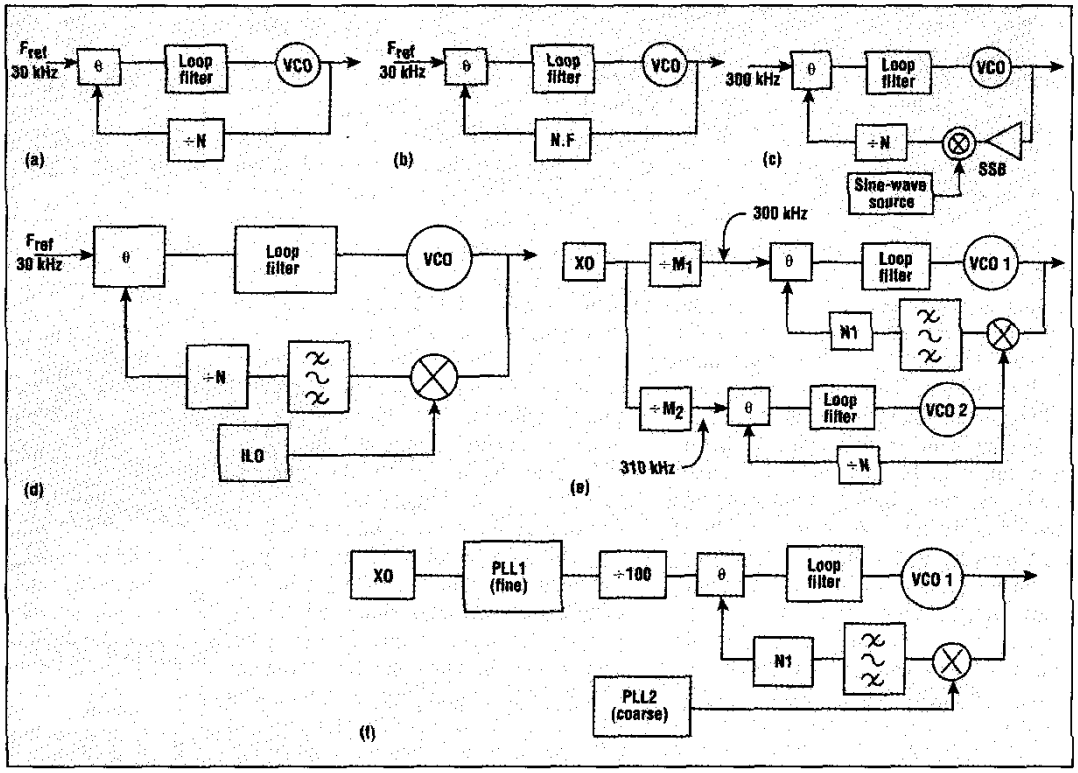

5. A number of frequency-synthesizer topologies are available for cellular applications, each providing distinctive performance advantages and drawbacks.

locked design is also attractive because no high-speed digital-to-analog converters (DACs) or anti-aliasing filters are required. Since the large-offset phase-noise performance of the synthesizer is determined by VCO1, both the fine and course PLLs can be made monolithic in order to minimize cost and size. This architecture offers exceptional phase-noise and spurious performance with fast switching speed at an attractive cost, provided that monolithic solutions are employed.

ANALYSIS TOOL

Phase-noise considerations for cellular communications require special attention if the end design is to be made as simple as possible. Although phase-noise requirements may be deduced through careful system analysis and planning, retrofitting existing equipment for cellular service can be extremely difficult when the baseline-system budget information is unavailable. Even when detailed analytical background material is available, some form of hardware validation of the frequency-synthesizer phase-noise and spurious requirements would be very useful. Other high-speed synthesis requirements, such as those needed for fiber-optic communication systems, encounter similar design problems. The NLO-100 general analysis and diagnostic tool has been developed in response to these problems.

The NLO-100 analysis tool provides a completely-programmable phase noise or discrete spurious-output spectrum covering a bandwidth of ±100 kHz around an arbitrary carrier frequency. The device uses embedded digital-signal-processor (DSP) technology and a user-friendly AT-computer interface to allow creation of arbitrary local-oscillator (LO) spectra for performance testing of cellular and other types of equipment with different phase-noise levels. This allows the enhancement or replacement of time-intensive system modeling with tangible real-time hardware for phase-noise simulation. The NLO-100 can also be used to create arbitrary jitter spectra for the evaluation of high-speed digital clock generation and recovery, such as that required in fiber-optic systems.

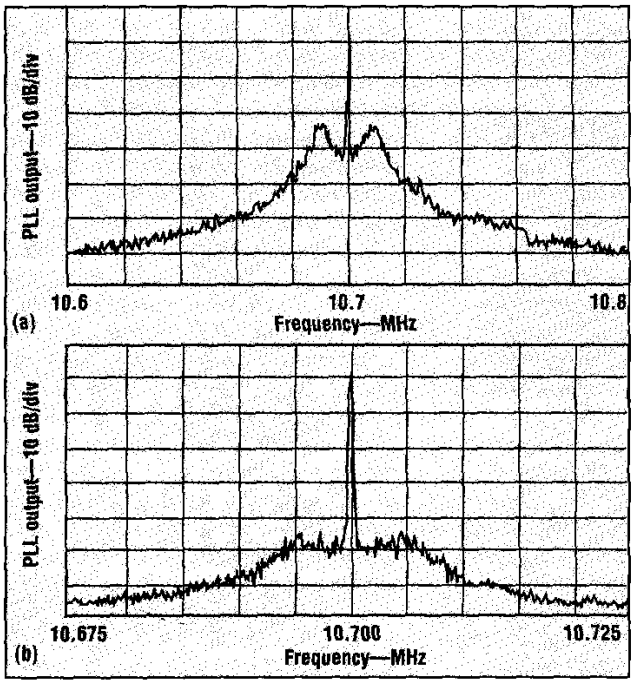

6. The versatility of the NLO-100 is illustrated by vastly-underdamped (a) and high-selectivity (b) PLL output spectra.

Two memory options provide completely-programmable phase-noise spectra in increments of 2.5 or 10 Hz. Discrete spurious elements are input in tabular form, while continuous phase-noise spectral input is greatly facilitated by allowing the user to enter the phase-noise spectrum as a template which is then mathematically generated using stationary, two-dimensional, filtered random processes. The NLO-100 can also be used to create completely-arbitrary, constant-envelope modulation waveforms over time periods up to approximately 400 ms.

The NLO-100 provides a very clean phase-noise spectrum without phase-noise impairment. Discrete spurious lines which are normally exhibited by phase-locked sources can easily be added. Almost any continuous phase-noise spectrum within the 200-kHz bandwidth can be programmed, as illustrated by the plot of a grossly-underdamped PLL output spectrum (Fig. 6a). In this plot, the close-in phase-noise floor was set to –60 dBc/Hz and a type-2 PLL was assumed with an f_n of 10 kHz and ζ of 0.20 to accentuate the noise peaking.

Low-noise, high-selectivity spectra are also handled, as shown in Fig. 6b, where the close-in phase-noise level was set to –80 dBc/Hz, while f_n and ζ were set to 5 kHz and 0.30, respectively. ••

Note

The NLO-100 is available from ComFocus Corp., with higher-bandwidth products to follow. Data sheets can be obtained by request.

References

1. Electronic Industries Association, "Cellular System Dual-Mode Mobile Base Station—Base Station Compatibility Standard," EIA/TIA/IS-54-A, March 1991.
2. V. Manassewitsch, *Frequency Synthesizer Theory and Design*, 2nd ed., John Wiley & Sons, 1980.
3. J.A. Crawford, *Computer-Aided Frequency Synthesizer Design: A Systems Approach*, Artech House, to be published.
4. A. Shipow, "Let's Clear Up the Confusion on PTD," *Microwave Systems News*, April 1977.
5. J.A. Crawford, "Understanding the Specifics of Sampling in Synthesis," *Microwaves & RF*, August 1984, p. 120.
6. J.A. Crawford, "Extending Sampling to Type-2 Phase-Locked Loops," *Microwaves & RF*, September 1984, pp. 171-174.
7. D.H. Wolaver, *Phase-Locked-Loop Circuit Design*, Prentice-Hall, 1992.
8. C.A. Sharpe, "A 3-State Phase Detector Can Improve Your Next PLL Design," *EDN*, Sept. 20, 1976.
9. J.G. Proakis, *Digital Communications*, 2nd ed., McGraw-Hill, 1983, chapter 4.

For further reading

W.F. Egan, *Frequency Synthesis by Phase Lock*, John Wiley & Sons, 1981.