A METHOD FOR COMPUTING ADJACENT-CHANNEL SPECTRAL ENERGY IN CELLULAR POWER AMPLIFIERS

Seth Pinsky

TriQuint Semiconductor, Inc.
Hillsboro, Oregon

ABSTRACT

The requirements for energy efficiency in cellular RF power amplifiers necessitate their operation in regions of significant non-linear distortion, with a resultant regrowth of undesired adjacent-channel spectral products that must be maintained within specified limits. This paper presents a method for predicting adjacent-channel spectral regrowth based on single-tone distortion curves that are easily measured in the laboratory and readily simulated in SPICE. The method affords a practical approach to predicting adjacent-channel power that has proven very useful for designing power amplifiers for cellular service under the IS-54 standard.

INTRODUCTION

The successful development of RF power amplifiers for cellular telephone use is dependent on negotiation of design trade-offs that exist among the principal performance indicators, these being output power, efficiency, and adjacent-channel power (ACP). Of these three, ACP is by far the most difficult to predict from a design point of view, the reason being that it is a function not only of the intrinsic non-linear characteristics of the amplifier, but also of the details of the encoding method and modulation scheme governing the information-containing signals supplied to the amplifier's input. The dependence of ACP on information-rich signals absolutely precludes its prediction via direct simulation in SPICE [2], for here the computing time and data storage requirements would be astronomical.

The general problem of computing the response of a non-linear amplifier to an information-rich signal with a known set of statistical properties is inordinately complex and is not readily solved. However, great computational simplification can be achieved by analyzing the response of the amplifier in the mathematical limit in which the bandwidth of the incident signal tends to zero. As we will show, this limiting case is often an acceptable approximation for predicting ACP as measured under laboratory conditions, and the tremendous simplification resulting from it can be enjoyed without accompanying great loss of accuracy.

In general, ACP is a gross result of two fundamental causes, namely non-linear amplifier distortion and information-rich signal excitation. The respective roles of each of these factors in determining ACP involves complex interrelationships that are not readily disentangled in the general case. However, under the limiting assumption of narrow-band excitation, the respective contributions of the two causes can be succinctly expressed by way of a simple mathematical relationship that we will present below. Using this approach, the problem of computing ACP is broken into four parts. First, a statistical sample of a modulation envelope is generated in a manner consistent with the communications standard. Second, the non-linear behavior of the amplifier is modeled by characterizing its response under excitation at various amplitude levels by a purely sinusoidal signal at the carrier frequency. Third, the output time-domain envelope signal is computed using an equation to be presented below. Finally, a fast Fourier transform (FFT) of the output waveform is performed, and the power residing in various spectral regions is computed by appropriate summation over the resulting frequency-domain signal.

GENERATING A STATISTICAL SAMPLE OF THE WAVEFORM

The first step in ACP computation entails generating a time-domain waveform that is statistically representative of the signals that will be supplied to the amplifier in field use. The modulated waveform incident on the amplifier is mathematically equal to the product of two components, the first being a purely sinusoidal carrier wave and the second being an envelope waveform that actually contains the signal information. Our waveform generation efforts focus on creating a waveform that can be regarded as a valid statistical sample of an IS-54 envelope signal. This sample is created by subjecting a purely random binary sequence to an encoding process consisting of two stages, the first of which is a $\pi/4$-DQPSK modulation and the second of which is a low-pass filtering.

0-7803-4471-5/98/$10.00 (c) 1998 IEEE

The π/4-DQPSK modulation is a scheme by which the binary data sequence is converted to a phase-modulation envelope. There are eight possible phases at which the (as yet unfiltered) envelope can exist at any one time, and these reside at equally-spaced intervals around the unit circle, each separated from its nearest neighbors by a 45-degree phase rotation. The envelope waveform remains in a particular phase state for a specified period of ~40 microseconds, and this period represents the transmission of a single symbol consisting of two bits of binary information. The transition of the envelope from one symbol to the next is constrained such that the phase change that takes place must be either ±45 degrees or ±135 degrees. (Changes of 0, ±90 degrees, or 180 degrees are not allowed under IS-54.) At the moment of phase change, there are four values that the envelope can assume, and thus two bits of uncorrelated information are transmitted. The selection of each symbol among its four possible values (as constrained by the previous symbol) is determined by the next bit-pair in line from the string of bits that comprise the signal being encoded.

The π/4-DQPSK time-domain envelope waveform generated above exhibits abrupt transitions from one symbol to the next, and these transitions result in spectral splatter that must be filtered as prescribed by the IS-54 standard. For this purpose, the standard prescribes that the envelope be filtered by a root-raised-cosine (RRC) low-pass filter, the frequency-domain characteristics of which are given by the following relation:

$$H(f)=\begin{cases}1 & 0\le f<\dfrac{(1-\alpha)}{2T}\\ \sqrt{\dfrac{1-\sin\left(\dfrac{\pi(2fT-1)}{2\alpha}\right)}{2}} & \dfrac{(1-\alpha)}{2T}\le f<\dfrac{(1+\alpha)}{2T}\\ 0 & f\ge\dfrac{(1+\alpha)}{2T}\end{cases} \quad (1)$$

(ref 2., p. 66.) In the above expressions, T represents the symbol period and α is a dimensionless filter constant having a specified value of 0.35 under IS-54. In our implementation, the RRC filtering is performed by taking an FFT of the unfiltered waveform, multiplying it the frequency-domain filter skirt, and then taking an inverse FFT to recover the output waveform. This FFT/filter-skirt/inverse-FFT method is an efficient way to calculate the convolution of two sequences. (For a discussion of this method, see ref. 1, pp. 110ff)

Figure 1 illustrates the time variation of the real and imaginary components of our sample of the envelope waveform as it exists after processing by the RRC filter. The RRC filter has served to dampen the transitions of the envelope (which were abrupt prior to filtering) and accordingly has reduced spectral splatter; it has also introduced an amplitude-modulation component into what was formerly a purely phase-modulation signal. It is this latter feature that will give rise to regrowth of adjacent-channel power when the signal is subjected to the non-linear characteristics of the RF amplifier.

Figure 1 -- Real v. Imaginary Component of the Input Envelope

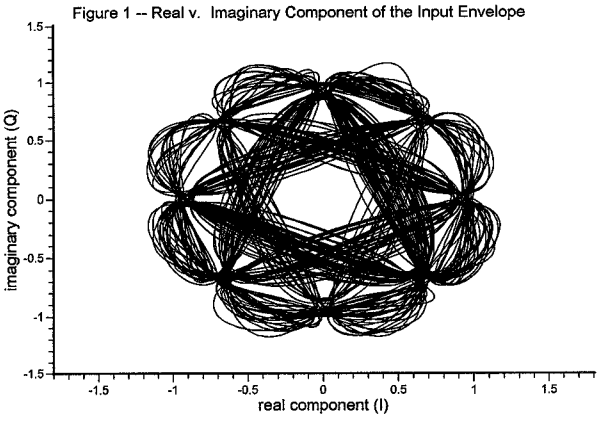

MODELING THE AMPLIFIER -- THE ENVELOPE TRANSFER CHARACTERISTIC

The envelope signal represents a time variation of the amplitude and phase arguments presented to a sinusoidal function of fixed frequency. In the most general treatment of the problem of computing ACP, it must be assumed that the time-varying amplitude and phase arguments supplied by the envelope cause electrical transients that must be analyzed. However, in our analysis it is precisely this consideration that we ignore. The deliberate neglect of transients occasioned by the time-varying nature of the envelope waveform is mathematically equivalent to the limiting case of vanishing signal bandwidth discussed above.

Our modeling of the amplifier involves a characterization of its steady-state response to sinusoidal excitation at the carrier frequency. This characterization is performed over a range of input signal amplitudes. At each level of signal excitation, the resulting output is a periodic waveform that can be decomposed by Fourier series methods into a superposition of sine waves residing at the harmonic levels. For purposes of analysis, we are interested only in the first harmonic of the output, i.e., in the fundamental Fourier component whose sinusoidal frequency is equal to that of the input excitation. At each level of sinusoidal excitation of the input, we record both the magnitude and phase of the first harmonic of the output. The resulting data set is fitted with a pair of smooth curves by way of polynomial interpolation, so that the response of the amplifier can be determined at any point residing in the range of amplitudes spanned by the data set.

In our implementation, we used a set of cubic polynomials to interpolate between adjacent data points.

Under this approach, N data points will result in N-3 polynomials that will allow prediction in the closed interval bounded by the 2nd and N-1st points of the data set. When our polynomial interpolation is complete, we have extracted the *envelope transfer characteristic* of the amplifier, this being the variation in the amplitude and phase of the first harmonic under steady-state sinusoidal excitation within the range of signal amplitudes determined by the limits of the data set.

The envelope transfer characteristic is a behavioral model of the amplifier from which we will compute ACP. As an experimental specimen for verifying our method, we selected a GaAsFET power amplifier designed for wireless service under IS-54. We used a vector network analyzer to measure the first harmonic of the output of the amplifier under sinusoidal excitation at a frequency of 836 Mhz. We collected data at 0.5 dB increments of input signal amplitude. Figure 2 shows the magnitude and phase of the measured envelope transfer characteristic of the amplifier.

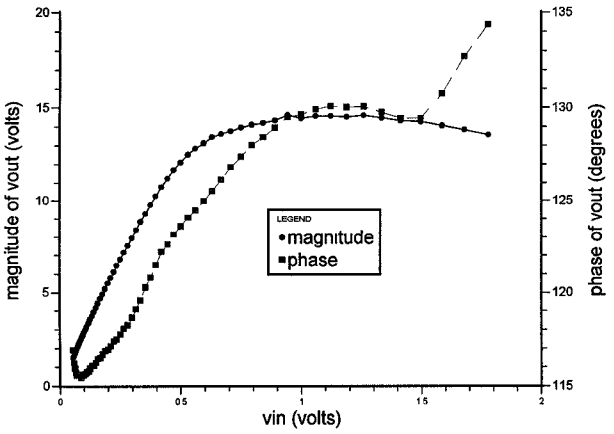

Figure 2 -- Magnitude and Phase of Envelope Transfer Characteristic

COMPUTATION OF THE RESPONSE

With the input envelope waveform generated and the envelope transfer characteristic extracted, we are now ready to calculate the output envelope waveform. The relationship is given by the following equation:

$$v_{out} = A(|v_{in}|)e^{j(\phi(|v_{in}|)+<v_{in})} \qquad (2)$$

where A represents the magnitude of the envelope transfer characteristic, ϕ represents its phase, and v_{in} and v_{out} represent the input and output envelope waveforms, respectively. The first two of these quantities are each real functions of the input signal amplitude; the last two are each complex functions of time. The simplicity of the above equation should be noted. It assumes a quasi-static, functional relationship between the instantaneous values of v_{in} and v_{out}; its conspicuous lack of integro-differential operators is a direct result of the fundamental simplifying assumption of a narrow-band signal that we postulated above.

SPECTRAL ANALYSIS OF THE OUTPUT

The v_{out} waveform calculated from the above equation must be analyzed to determine the spectral power residing in the intended channel as well as those adjacent. This analysis is performed by taking an FFT of v_{out} (figure 3) and performing a summation over the frequency-domain samples residing within a given channel. In this stage of analysis, one encounters sampling artifacts that are potentially responsible for significant distortion of the results. These artifacts must be maintained within tolerable limits by appropriate choice of sampling constants. The selection of these constants involves a trade-off between accuracy on the one hand and expenditure of computing resources (time and storage) on the other.

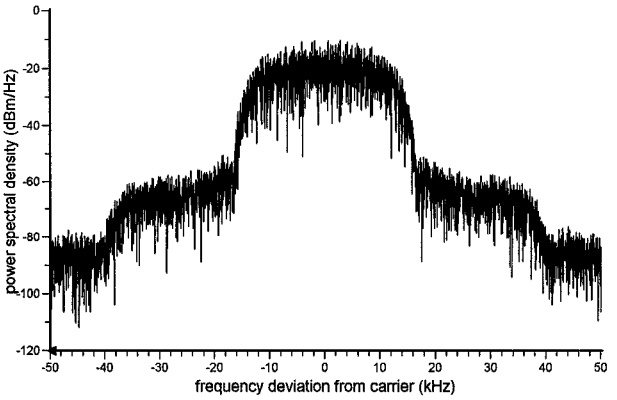

Figure 3 -- Computed Magnitude Spectrum of Output Signal

Of the sampling artifacts that exist, the most significant is that caused by random fluctuations among individual frequency-domain samples. The exact nature of these fluctuations is dependent on a detailed description of the statistical process under study, but general rules-of-thumb have been formulated for confining their impact on estimated power spectra to within desired limits. (See ref. 4, pp. 549ff). In general, the individual frequency-domain samples exhibit RMS fluctuations equal in magnitude to their expectation values, and these large fluctuations are reduced by the signal-averaging that results from the summation of the samples within a given channel. After summation, the resulting RMS error in power estimation varies inversely as the square root of the number of samples within a given channel. Application of these rules predicts that an RMS error in power estimation of 0.15 dB can be achieved with ~1000 frequency-domain samples within a given channel. Since the spacing between adjacent frequency-domain samples is equal to the reciprocal of the duration of time-domain signal under analysis, (this relationship being a general property of the FFT operation) our control of the artifact under discussion

is achieved by selecting a signal duration of appropriate length. In order to achieve the 0.15 dB accuracy desired, we chose a duration for v_{in} and v_{out} equal to 1024 symbol-durations of the IS-54 signal.

The second most significant artifact is that caused by the finite density of time-domain samples for the signals represented by v_{in} and v_{out}. If the density of time-domain samples is not sufficient, the result is frequency-domain aliasing (see ref. 1, pp. 28ff) that will cause distortion of the computed power spectra that first appears in highest-frequency regions of the spectral density distribution and gradually encroaches on lower-frequency regions as the sampling interval is relaxed. The time-domain sample interval required for virtual elimination of this artifact is equal to the reciprocal of a bandwidth figure that represents the practical extent of the power spectrum. For our computations, we settled on a sample density of 32 samples per symbol period, and this figure corresponds to a sample interval of ~1.25 microseconds under IS-54. This choice is correct under the assumption that the output power spectrum is effectively absent outside a window of width 1/1.25mS = 800 kilohertz. A visual examination of figure 3 suggests that the extent of the power spectrum resides well within this range.

Our findings concerning the management of sampling artifacts as revealed in this section must be applied to our waveform generation efforts described above. In consideration of the two artifacts discussed, our choice of a 1024-symbol waveform duration against a sample-per-symbol density of 32 yields a gross sample count of 32768. This figure represents the size of the sampled waveform and imposes a storage requirement of 512k bytes when a 64-bit floating-point data type is used. The 512k waveform size is readily handled with available computing resources and thus we anticipate that containment of the sampling artifacts can be achieved without exceeding the limits of computational feasibility.

With the sampling artifacts properly managed, the next concern is that of channel definition under the communications standard. For purposes of reckoning power in a given channel under IS-54, each frequency-domain sample is weighted by an RRC filter whose functional form is identical with that given by (1). To calculate the spectral power residing in a given channel, the RRC filter skirt is shifted horizontally so that its center coincides with that of the channel of interest. For each frequency-domain sample, the power subsequent to filtering is obtained by multiplying its square magnitude by the square magnitude of the shifted RRC filter skirt at the same frequency. The resultant filtered samples are summed over the entire width of the filter skirt. When this summation is complete, the effect of the RRC filter is that the channel boundaries under IS-54 acquire a tapered character and are gradual rather than abrupt.

EXPERIMENTAL RESULTS AND DISCUSSION

Figure 4 shows the computed variation of ACP as a function of output power for the amplifier that we used in our study. Alongside it is a plot of the measured ACP performance as a function of output power. Our computational algorithm is successful in predicting the output power that can be achieved at a given level of ACP to well within 0.5 dB. Our ability to predict ACP is not seriously degraded by the narrow-band signal approximation upon which our computational method rests.

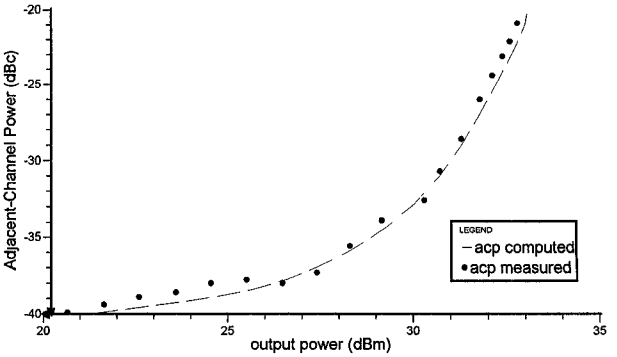

Figure 4 -- ACP Computed v. Measured

Our experiment has demonstrated the accuracy of our computational method in a case in which the behavioral model data for the amplifier were obtained by direct laboratory measurement. In day-to-day design work, the behavioral model data for amplifiers of interest are ordinarily obtained via SPICE simulations and subsequently fed to the ACP computing algorithm. Simulation of the envelope transfer characteristic is readily performed in SPICE, and computation of spectral regrowth from simulated behavioral data affords us a practical approach to bypassing the obstacles that make direct SPICE simulation of ACP impossible.

REFERENCES

1.) Alan V. Oppenheim and Ronald W. Schafer, Digital Signal Processing. Prentice-Hall, Inc., 1975

2.) Joseph Staudinger, "Specifying Power Amplifier Linearity via Intermodulation Distortion and Channel Spectral Regrowth", Applied Microwave and Wireless, July/August, 1997, pp. 62ff.

3.) TIA/EIA IS-54, "Cellular System Dual-Mode Mobile Station / Base Station Compatibility Standard," Telecommunications Industries Association, 1992.

4.) William H. Press, Saul A. Teukolsky, William T. Vettering, and Brian P. Flannery: Numerical Recipes in C: The Scientific Art of Computing. Cambridge University Press, 1988-92.