Theory of Operation for Digital Predistortion

Skip to end of metadata Go to start of metadata

The typical non-linear response of most power amplifiers is undesirable because it results in (in-band and out-of-band) distortion of the amplified signal. An amplifier with a linear response does not produce such distortion. Digital predistortion linearizes the non-linear response of a power amplifier over an operating region. It uses digital signal processing techniques to condition a baseband signal prior to modulation, up-conversion, and amplification by the power amplifier. As a result of this conditioning, the cascade of the digital predistortion response and the power amplifier response produces the desired linear response. Please refer to the *Digital Predistorter Bibliography* in the Linearizer DesignGuide documentation for references associated with the following explanations.

The figure below shows a high-level representation of a signal passing through a digital predistorter and then into a power amplifier.

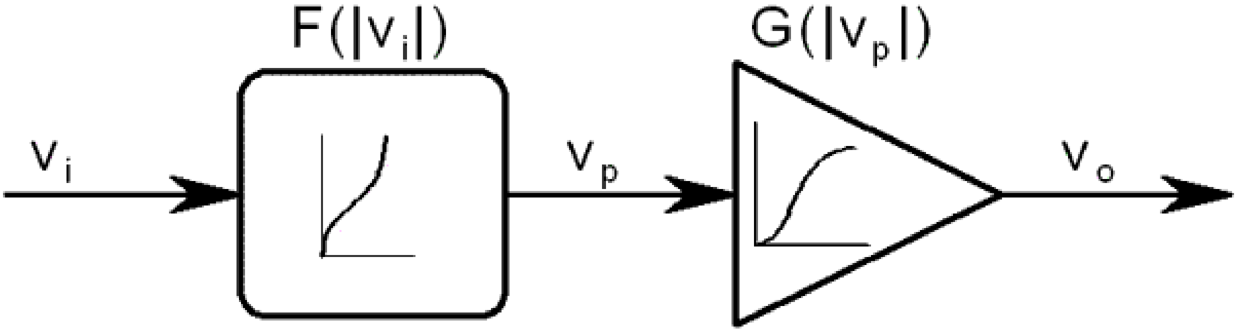

Digital Predistorter Followed by a Power Amplifier

The gain, G, of the PA is modeled as a function of the magnitude of the power amplifier input signal, v_p. The function G is, in general, memoryless and non-linear in amplitude and phase (i.e. complex, when converted to an equivalent baseband model). The use of a memoryless model that is dependent upon input signal magnitude only is a simplification of the actual response of a typical power amplifier. Other variables will affect the power amplifier response, including, most notably, frequency and operating temperature.

The digital predistorter precedes the power amplifier, conditioning the input signal, Vi, before amplification by impressing upon it a complex gain given by the digital predistortion response, F. Similar to G, F is made to be a function of the magnitude of the input signal to the digital predistortion, namely, Vi. Thus, the cascade of the digital predistortion and power amplification will result in the desired linear response when F(|Vi|) G(|Vp|) = k, where k is a constant and Vp = Vi F(|Vi|).
The following figure illustrates the typical relationship between the input power and output power of a power amplifier.

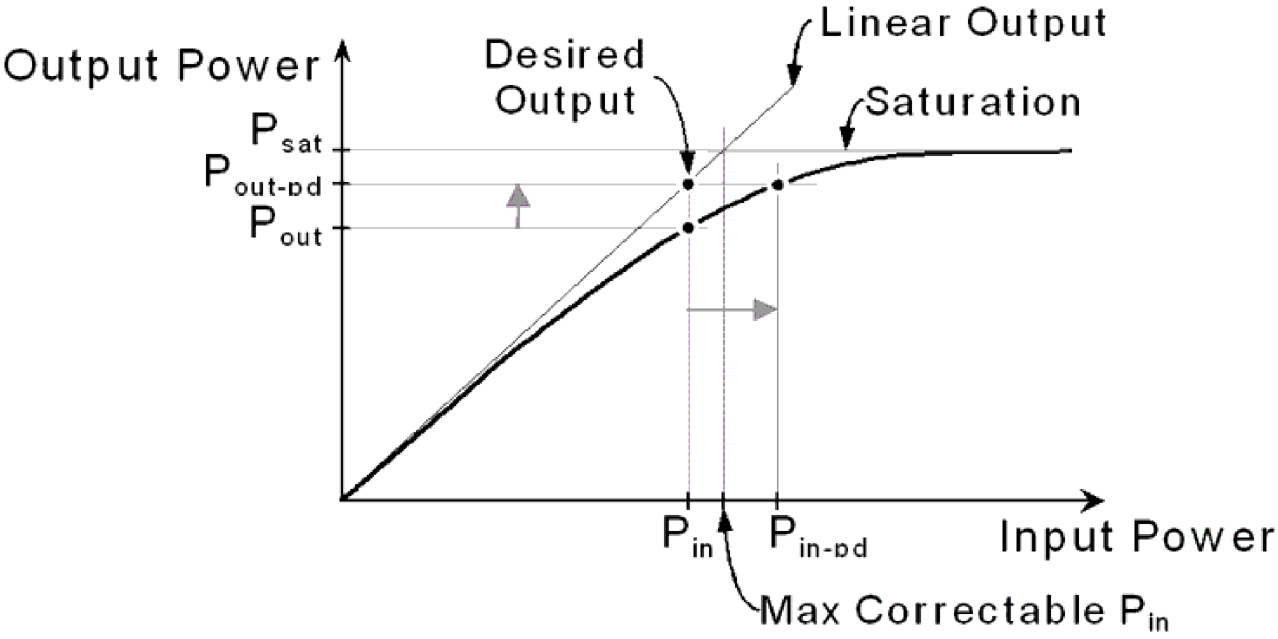

Power Amplifier Pout versus Pin and Digital Predistortion

The heavier dark curve shows that in the absence of digital predistortion the power amplifier's Pout versus Pin curve is highly non-linear. However, through the introduction of digital predistortion, the Pout versus Pin curve is made to have a linear response over a large range of input power levels. The desired linear response of the power amplifier is illustrated by the Linear Output curve shown in the preceding figure. The slope of the Linear Output curve is the desired linear gain of the power amplifier.

The manner in which digital predistortion works can be understood from the preceding figure. When the amplifier is operating in compression, the Pout versus Pin curve falls below the Linear Output curve, hence, the actual output power of the power amplifier is not sufficient for linear operation. The inclusion of digital predistortion prior to the power amplification has the effect of introducing expansion - the amplitude of the input signal is increased so that the desired output power (falling on the Linear Output curve) is achieved. The expansion effect of digital predistortion can be observed in the preceding figure where the input power, Pin (resulting in Pout before PD), is increased to Pin-pd so that the power amplifier output power is raised to Pout-pd which coincides with the Linear Output curve.

The region of the Pout versus Pin curve which can be linearized using digital predistortion is not unlimited. The maximum input amplitude that will be amplified linearly is dependent upon the desired linear gain and the saturation output power of the PA, P_{sat}. The output power cannot be expanded beyond the saturation level, thus, the intersection of the Linear Output curve with the P_{sat} level determines the maximum input amplitude that will result in linear amplification.

The amplitude of the input signal to the power amplification may be constant, as in the case of a GMSK input signal - characterized by a constant envelope. When the input signal is based upon a non-constant envelope modulation scheme, for instance, QPSK, OQPSK, or Mary-QAM, the amplitude of the input signal to the power amplification will vary continuously and rapidly. The upper and lower limits of the amplitude variation of the input signal defines an operating region of the power amplifier. Two operating regions are shown in the following figure.

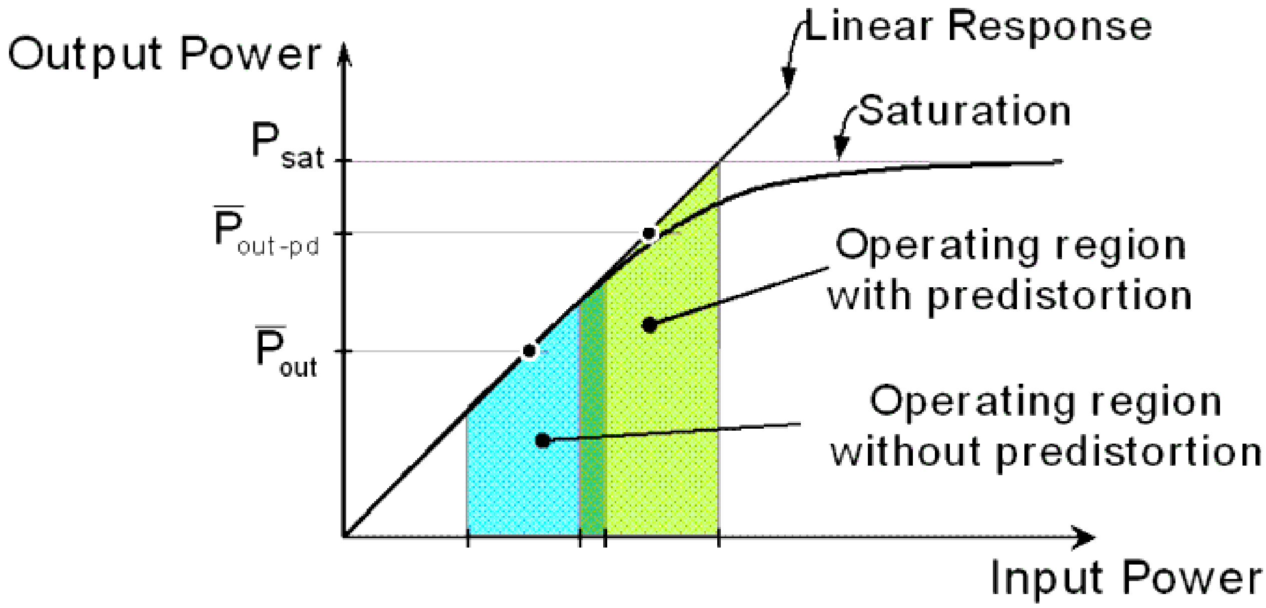

Power Amplifier Operating Region With and Without Digital Predistortion

Each operating region is characterized by an average output power level, $\overline{P_{out}}$. The magnitude of $\overline{P_{out}}$ is typically described in terms of its relation to the peak output power associated with the signal being amplified by a peak-to-average power ratio (PAR). The PAR is dependent upon the statistics of the signal being amplified.

The non-linear portion of the Pout versus Pin curve is normally located at the higher output power levels, as saturation is approached (depending on the type of amplifier, non-linear behavior may also be found within other portions of the Pout versus Pin curve). Generally, to minimize the distortion introduced by the non-linear region, the input power is curtailed or backed-off, so that the operating region falls predominantly within the linear portion of the Pout versus Pin curve.

It is desirable that the operating region be as close to saturation as possible because amplifier efficiency increases as the average output power increases. Efficiency only increases to a certain point, however, and degrades as more and more of the operating region encompasses saturation. As is shown in the preceding figure, the linearization of the Pout versus Pin curve through digital predistortion enables the average output power of the power amplifier to be elevated above the average output power of the power amplifier without digital predistortion. Raising the average output power level enables the power amplifier to be operated more efficiently.

The phase response of a power amplifier is also non-linear. Thus, Digital predistortion is also used to linearize the typical non-linear phase response (as illustrated in Digital Predistorter Followed by a Power Amplifier and Power Amplifier Pout versus Pin and Digital Predistortion) and the amplitude response. The following figure shows a typical phase response of a power amplifier and the linearizing effect of digital predistortion. The desired phase response is one that is constant at all input power levels.

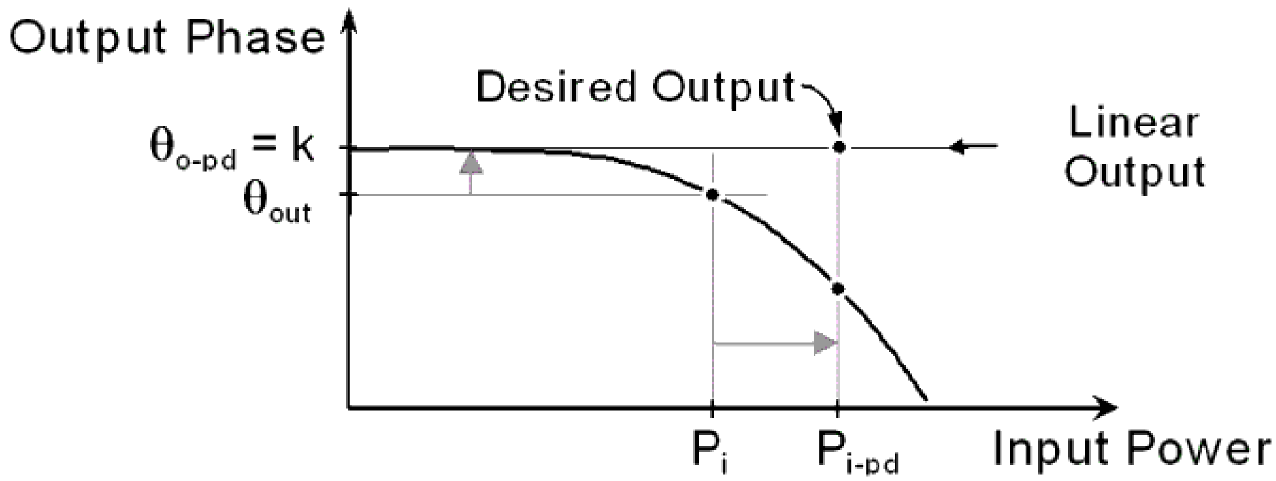

Typical Power Amplifier Phase Response

Digital Predistortion System

A system level block diagram of a digital predistortion system is shown in the following figure. The entire digital predistortion system is housed in the transmitter and requires the addition of several components to the standard transmitter components. The primary components augmenting the modulator, upconverter, and the power amplifier are:

- Complex Gain Adjust block in the baseband signal path
- Look-Up Table for storing a set of complex coefficients
- Demodulator and Downconverter whose input is drawn from the output of the power amplifier
- Adaptation Algorithm processing block that calculates the complex coefficients that are stored in the look-up table

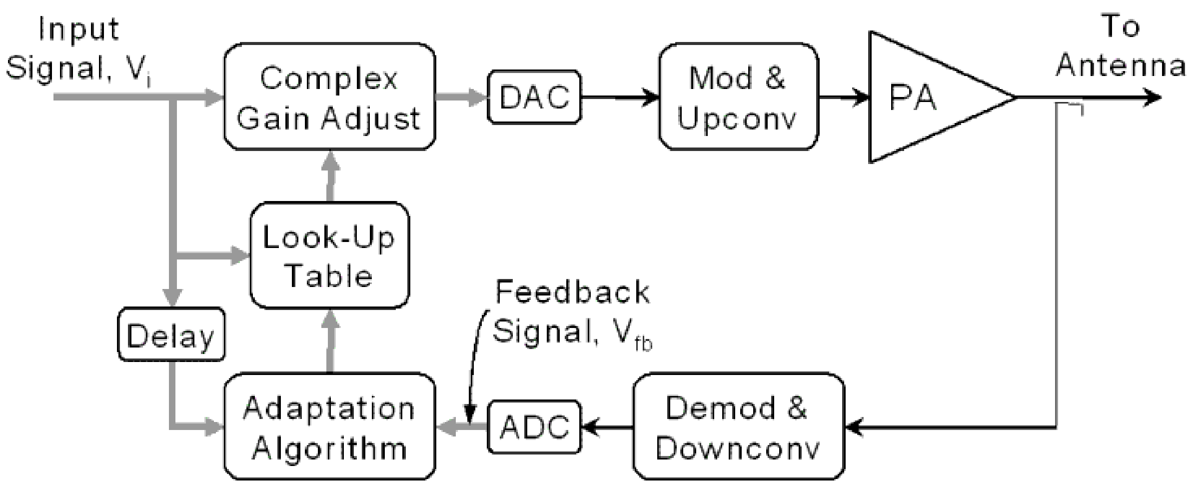

Digital Predistortion Block Diagram

With the inclusion of digital predistortion, the digital complex baseband input signal samples are multiplied prior to the DAC by complex coefficients drawn from the look-up table. The look-up table coefficients implement the predistortion function, F, shown in Digital Predistorter Followed by a Power Amplifier. The adaptation algorithm determines the values of the coefficients by comparing the feedback signal and a delayed version of the input signal.

Note

Ideally, the feedback signal should be a scaled reproduction of the input signal (i.e. linear amplification), thus, the role of the adaptation algorithm is to derive a set of coefficients that forces the error between the scaled feedback signal and the delayed input signal to zero.

A number of critical design parameters/issues can be readily identified from the digital predistortion system shown in the preceding figure:

- Sampling rate used in the DSP portions of the transmitter
- Number of bits used for quantization by DAC and ADC
- Shape of the reconstruction filter at the output of the DAC
- Bandwidth of the low pass filter at the input of the ADC
- Stability of the feedback loop/adaptation algorithm
- Precision (fixed-point) of the look-up table and DSP processing, if any
- Memory available for the look-up table and adaptation algorithm
- Complexity of the adaptation algorithm and the amount of DSP horsepower required
- Accuracy of the determination of the feedback delay

The intent of implementing the digital predistortion system in ADS is to be able to capture and analyze the impact on performance of some of these design parameters.

Implementing the Predistortion Function

The Complex Gain Adjust and Look-Up Table blocks of the preceding figure are shown in more detail in the following illustration.

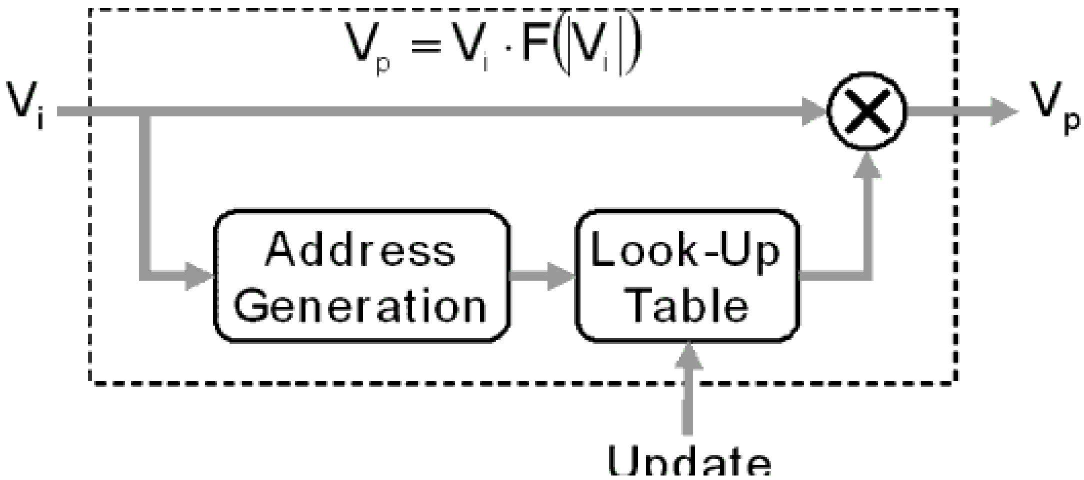

Complex Gain Adjust and Look-Up Table

The predistorting function, F, is assumed to be a function of the magnitude of the input signal. The predistorting function is implemented using a complex multiplier, a Look-Up Table and an Address Generation block that selects the appropriate coefficient from the look-up table, given the magnitude of the input signal. The coefficients stored in the look-up table are the value of the predistortion function at certain input signal magnitudes. Thus, the predistortion function is not implemented in an analytic manner, rather, it is only calculated at a specified number of points.

The size of the look-up table employed determines the number of points at which the predistortion function is calculated. In addition, the distribution of the predistortion function points need not necessarily be evenly distributed across the range of the input signal magnitude. Instead, it may be desirable to distribute the predistortion function points across the range of the input signal magnitude using a squared (power) or logarithmic relationship. The following plots illustrate how the predistortion function may be indexed in the look-up table by the magnitude of the input signal and by the power of the input signal. Both plots show the same information, however, the scale of the x-axis is in dBm in the left hand plot and in Volts in the right-hand plot.

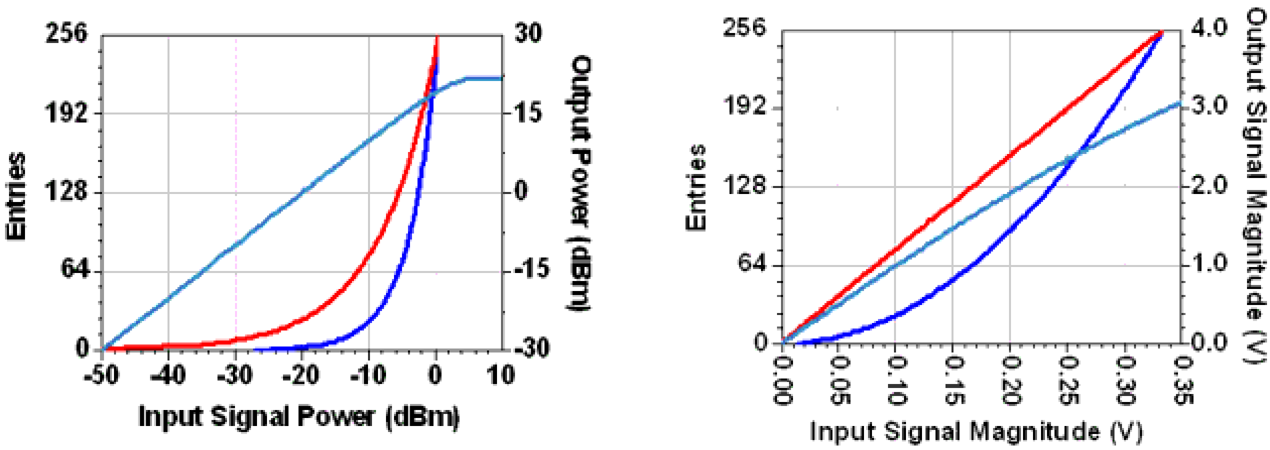

Cumulative Distribution of Look-up Table Entries for Magnitude and Power Indexing

For a look-up table with 256 entries, the plots show the cumulative distribution of the entries over the range of the input signal magnitude for the magnitude and power distributions. Overlayed upon the distributions is the amplifier response, illustrating where the entries fall along the amplifier response. The 256th entry in both distributions falls at the maximum input power level that can be linearized (approximately 0.4 dBm). The look-up table entries are equally spaced over the range of the magnitude of the input signal in the case of magnitude indexing, whereas, more look-up table entries are distributed at the higher end of the range in the case of power indexing.

Adaptation Algorithm

The function of the adaptation algorithm is to derive the predistortion function, F, i.e. the inverse characteristic of the amplifier response. The predistortion function may be derived using either a modulated signal input (random signal) or a known training signal input. The adaptation algorithm and its implementation are fundamentally different depending upon which type of input signal is utilized. The algorithms that are based upon the use of a modulated signal, employ statistical signal processing and, typically, some type of curve fitting algorithm to generate a smooth predistortion function. The complexity of the adaptation algorithm and its implementation can be significantly simplified, however, by using the alternative input signal - a known training signal.

Training Signal

The adaptation algorithm chosen for implementation is based upon the use of a training signal. The training signal is a single tone having a frequency equal to the carrier frequency and whose power is ramped up over the duration of the training period. The power of the tone is set to zero at the start of the training period and will typically peak at, or just below, the maximum correctable input power of the amplifier.

The use of a single tone whose power is ramped as a training signal greatly simplifies the adaptation algorithm and its implementation. However, the use of the training signal does require that the modulated signal being transmitted be interrupted while the training signal is transmitted. In addition, because the training signal is a single tone, the digital predistorter is only correcting for the operation of the amplifier at a single frequency, not across the entire transmission bandwidth. If the amplifier's passband is quite flat, the use of a single tone training signal in this manner will enable the predistortion function to be determined accurately. In any event, the use of a single complex coefficient to correct for distortion of the amplifier at a particular power level presupposes that any amplifier memory effects are minimal, i.e. the amplifier has a flat passband.

Implementation

The following illustration is a block diagram of the adaptation algorithm that is employed. The algorithm is based upon the determination of the open loop gain, H, of the predistorter and amplifier combination at the power level associated with each look-up table entry.

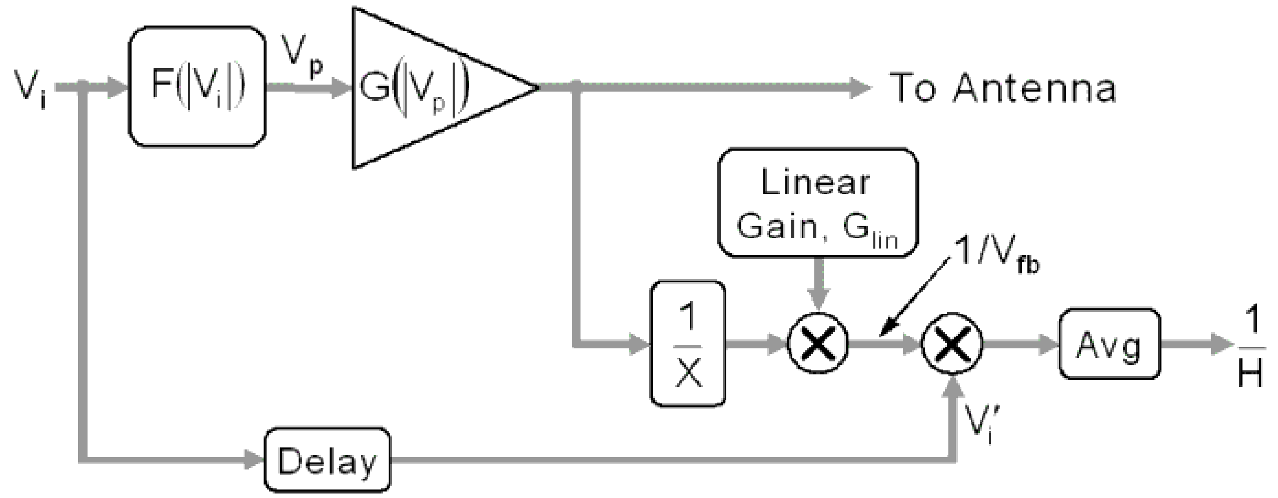

$$\frac{1}{H} = \frac{V_i'}{V_{fb}} = \frac{G_{lin}}{F(|V_i|) \cdot G(|V_p|)} \Rightarrow 1 + j \times 0$$

Implementation of the Adaptation Algorithm

Recall that the desired linear response of the predistorter and amplifier cascade requires that F(|Vi|) G(|Vp|) = k for all inputs. Hence, if Glin is set to be equal to k, the desired open loop gain of the system is unity. If the calculated open loop gain is not equal to unity, the predistortion function must be adjusted in a manner to drive the open loop towards unity. This is can be achieved as illustrated in the following manner.

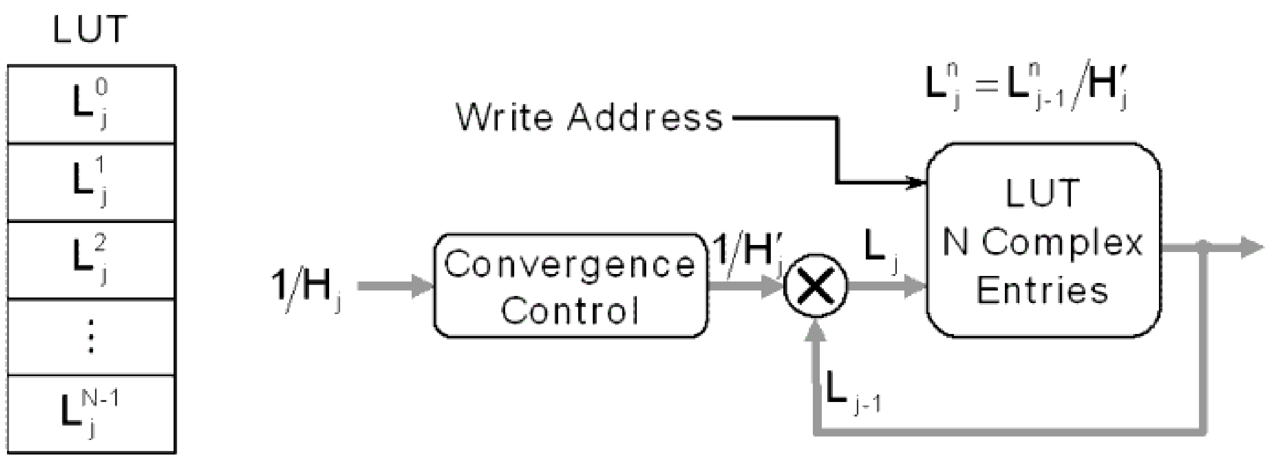

Calculation of New Predistortion Function

The predistortion function is defined by a set of coefficients stored in the look-up table, Ln, where each n corresponds to an input signal magnitude which is mapped to a look-up table address. In order to drive the open loop gain to unity, the predistortion function coefficients are updated by dividing each coefficient by the calculated open loop gain (or by the calculated open loop gain adjusted to slow the rate at which the coefficients will change).

Synchronization

The accuracy of the open loop gain calculation is dependent upon the accuracy of the estimation of the delay in the feedback path. As shown in Implementation of the Adaptation Algorithm, the input signal, Vi, must be delayed precisely by an amount equal to the delay in the feedback path.

The delay in the feedback path is estimated by calculating the correlation between the magnitude of the input signal and the magnitude of the feedback signal. The use of the magnitude of the signals has the benefit of not requiring phase synchronization in the feedback path. Because the delay in the feedback path will not necessarily be equal to an integer number of DSP sample periods, interpolation is employed to more precisely align the input and feedback signals.

The correlation between the input and feedback signal is performed on a modulated signal that precedes the training signal because the gain compression of the amplifier makes the accuracy of the correlation over the training signal suspect. In addition, because the envelope of the modulated signal will typically have a pdf such that it spends much of its time within the linear operating region of the amplifier, correlation using the modulated signal becomes more reliable. However, because the modulated signal is stochastic, the statistics of the modulated signal, as well as the size of the data block over which the correlation operation is performed will impact the accuracy of the delay estimation. In general, the accuracy of the estimation improves as the block size increases. Unfortunately, a larger block size requires more memory and takes longer to perform the estimate.

Convergence

The predistortion function cannot be exactly determined following the transmission of single training ramp and recalculation of the look-up table coefficients thereafter. A series of training ramps will have to be transmitted - although, significant improvements in the ACPR of the amplifier should be observed even after a single training ramp. Simulations have shown that the predistortion function can converge to a solution after only six training ramps.