AN-1371
APPLICATION NOTE

One Technology Way • P.O. Box 9106 • Norwood, MA 02062-9106, U.S.A. • Tel: 781.329.4700 • Fax: 781.461.3113 • www.analog.com

Variable Dynamic Range

by Brad Brannon and Jonathan Harris

INTRODUCTION

Variable dynamic range (VDR) provides an efficient, non-obstructive technique to provide wideband, high resolution, and fast sample rate to digital predistortion (DPD) observation receivers. VDR achieves the following:

1. Full bandwidth and dynamic range of the analog-to-digital converter (ADC) is almost always available for the DPD feedback receiver because all converter bits pass through when the signal behaves like a power amplifier (PA) output, converged or nonconverged.
2. DPD convergence is improved over other export methods because full resolution of the converter is available under all DPD operating conditions.
3. DPD loop can respond to peak conditions immediately because the full range and bandwidth of the converter are available for DPD operating conditions.
4. Synchronization with external hardware is not required because full resolution is automatically available except when non-DPD signals are present, reducing signal routing and system complexity.

Export controls limit the availability of data converters to certain countries. However, the growing commercial telecommunications market demands access to faster and higher resolution converter technology to build these platforms, particularly for digital predistortion of power amplifiers. In this application, nonlinearities are measured and digitally corrected such that overall intermodulation performance is significantly better than that of the core signal chain while at the same time boosting amplifier efficiency. Depending on the implementation and system requirements, end adjacent channel power ratio (ACPR) can reach 80 dB to 85 dB and efficiencies approaching 50% in some new amplifier topologies.

However, to achieve these levels of performance, high performance observation receivers are required that exceed the overall output requirements from the PA. Compounding issues include the wide bandwidths required from these systems. Many core systems today require operational bandwidths in excess of 100 MHz and correction bandwidths for fifth- and seventh-order intermodulation products that drive the bandwidth to 500 MHz and more. Previously available sampling receivers met export restrictions by limiting the bandwidth by analog or digital means and are unsuitable for DPD applications. Existing digital options available include digital noise shaping, FIR filtering, and triggered resolution adjustment. Each of these options introduces unnecessary impairments that must be overcome. However, to accurately address these impairments, it is best to start with a better understanding of the needs of the application.

TABLE OF CONTENTS

REVISION HISTORY

9/15—Revision 0: Initial Version

TYPICAL DPD APLLICATION

Figure 1 shows a typical DPD application. Although not shown, zero IF (ZIF) and complex IF (CIF) sampling architectures are also common. A typical transmitter output consists of between 20 MHz and 100 MHz of useful signal information; however, new designs are becoming available that support greater than 100 MHz of bandwidth.

In addition to the useful signal information, these transmitters typically generate significant harmonics and intermodulation products. To compound these linearity issues, the PA designs are typically optimized for efficiency, which means that linearity is sacrificed. However, it is common for overall dc efficiency to be 33%, meaning that dc power is three times the output RF power. This means that for 50 watts out, 150 watts are dissipated. Newer designs are moving toward 50% efficiency, which indicates that intermodulation of the core amplifier is even worse.

To keep the overall linearity at an acceptable level and to keep spurious out of neighboring bands, the baseband signals are typically predistorted by digital techniques (Morgan, Ma, Kim, Zierdt, and Pastalan, "A General Memory Polynomial Model for Digital Predistortion of RF Power Amplifiers," *IEEE Transactions on Signal Processing*, Vol. 54, No. 10, October 2006) that seek to model and generate antidistortion of the transmitted data such that the resulting spectrum is significantly better than the core amplifier itself.

As shown in Figure 1, these techniques are closed-loop and require a receiver of higher linearity than the overall requirements. Otherwise, the distortion of the receiver is interpreted as distortion in the output and limits the overall convergence of the system. Therefore, the downconversion and digitization of the transmitter output require very wide bandwidth, high linearity, as well as high resolution, which requires a core ADC that offers good performance in all three parameters.

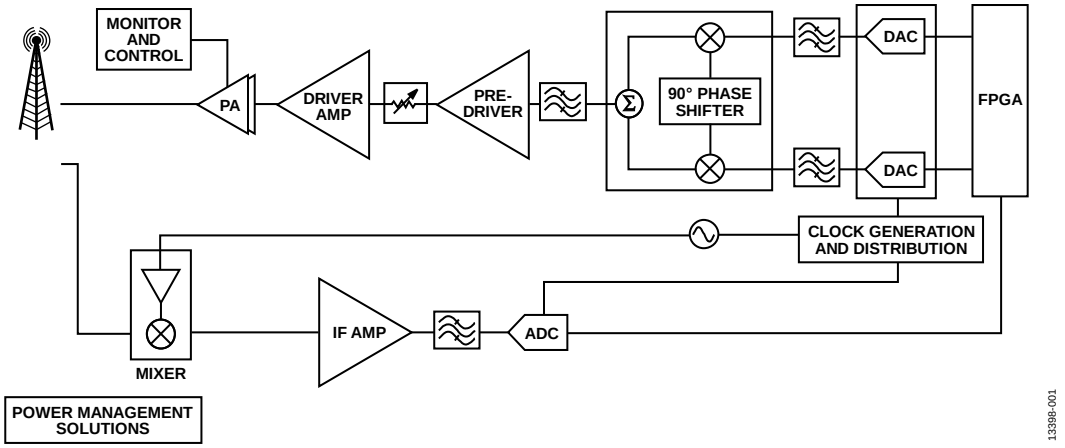

Figure 1. Typical DPD System Diagram

TYPICAL PA OUTPUTS

A typical uncorrected power amplifier exhibits poor intermodulation performance. An example is shown in Figure 2. The intermodulation products are in black, approximately 30 dB below the main carriers and extend approximately 100 MHz on either side of the useful signal information.

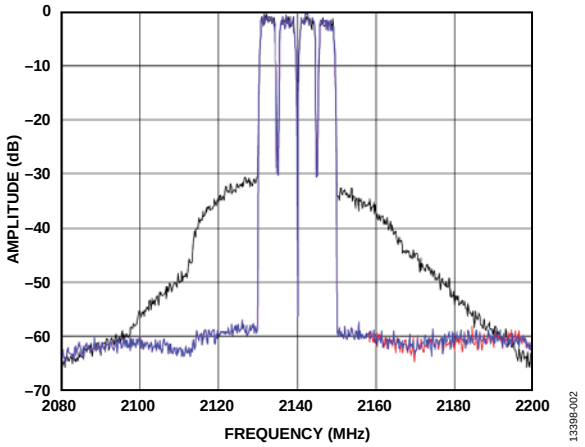

Figure 2. Uncorrected Amplifier ACPR

These intermodulation products not only detract from the useful signals, but they spill energy into adjacent wireless services, potentially causing disruption to those services. These signals are often not easily filtered because the transition bands of transmit filters often extend over 20 MHz or more, as illustrated by Figure 3, and these bands are often protected. Therefore, DPD is required to ensure compliance with the operating license.

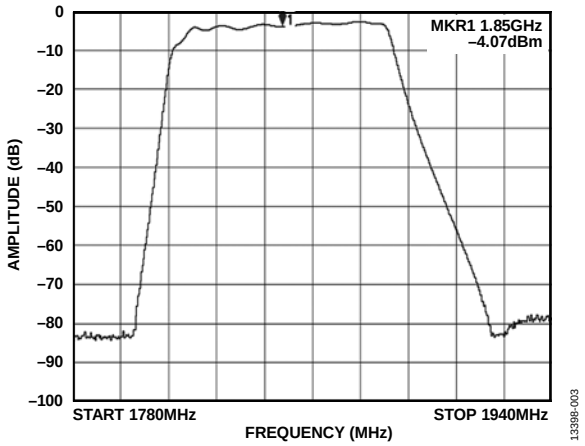

Figure 3. Typical Transmit Filter Response

REQUIREMENTS OF CORE ADC

There are three basic requirements of the ADC for a DPD application: sample rate, spurious-free dynamic range (SFDR), and signal-to-noise ratio (SNR). SNR is also expressed as the effective number of bits (ENOB).

The sample rate of the converter determines the bandwidth that can be corrected. The Nyquist criterion is often cited as dictating a requirement for determining the sample rate given the bandwidth. If the bandwidth needs to be x, the Nyquist criterion dictates that 2x is required for the sample rate. There are other practical concerns as well. Where analog filtering is concerned, the ratio is extended to 3x to ensure that filter roll-off is not overly challenged. However, many DPD applications use very little filtering aside from the duplexer to ensure low impact on group delay. Additionally, there is a very wide range of DPD intellectual property (IP) available. Much of that IP works to relax the sample rate requirements, oftentimes cheating on the Nyquist requirements. Regardless, very wide bandwidth analog signals require very high sample rates.

For SFDR, the typical performance of the observation path must be much better than the target system performance. Typically, the converter is expected to be 10 dB to 15 dB better than the transmitter output. For MC-GSM, performance goals are typically 60 dB for third-order terms and 70 dB for fifth-order terms. To not bias the results, the observation path must therefore provide performance of 75 dB to 85 dB, respectively, for these terms. Whereas 3G and 4G platforms may have lower targets, similar levels of performance are expected to ensure robust correction.

In the typical DPD loop, there is a direct relationship between noise and loop convergence. The higher the noise floor, the longer it takes the loop to converge. Conversely, the lower the noise floor, the quicker convergence is reached. Because regulations must be met and out of band energy minimized, system designers must design a system that responds as quickly as possible to out of band signals and, therefore, trade off loop response for converter resolution.

Figure 4, Figure 5, Figure 6, and Figure 7 demonstrate this concept. Figure 4 and Figure 5 represent loops where the convergence is made by slowly integrating the information in the feedback path; 9-bit and 14-bit converters are represented, respectively. As shown, there is virtually no difference between 9-bit and 14-bit convergence rates for a slow loop.

Figure 6 and Figure 7 represent loops that are fast enough to dynamically respond to signal conditions, as is expected in a realistic scenario. In this case, there is a marked difference in the performance between the 9-bit and 14-bit receiver paths.

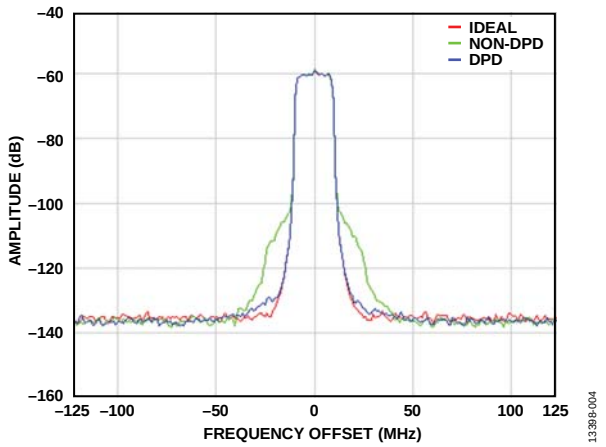

Figure 4. 9-Bit Converter with Slow Loop

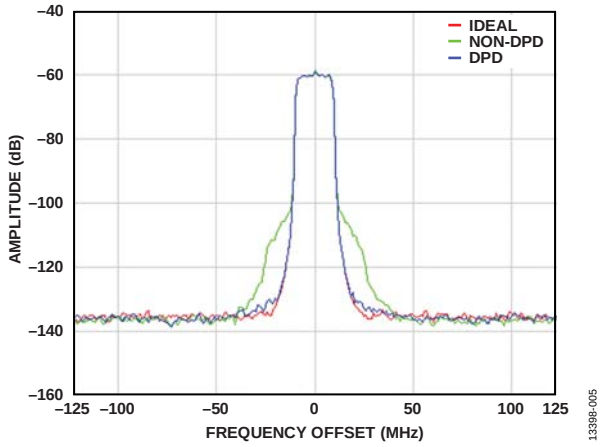

Figure 5. 14-Bit Converter with Slow Loop

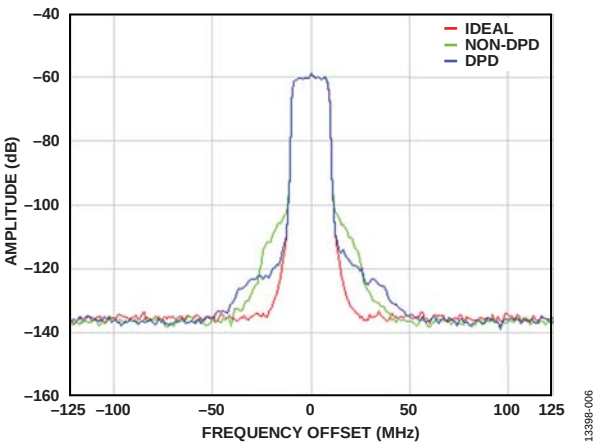

Figure 6. 9-Bit Converter with Fast Loop

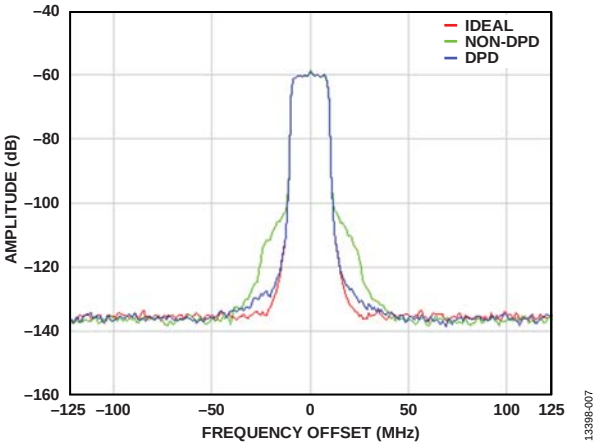

Figure 7. 14-Bit Converter with Fast Loop

TRADITIONAL TECHNIQUES

From the previous section, it is clear that sample rate, resolution, and linearity are all important factors in most DPD applications. Many of the techniques used in export devices throttle these aspects in one way or another.

One of the most common techniques is noise shaping, where the digital data is noise shaped such that in band noise is shifted to regions outside those of primary interest, as shown in Figure 8. Noise shaping is typically done for main receiver functions where information outside the main band primarily consists of out of band blocker information.

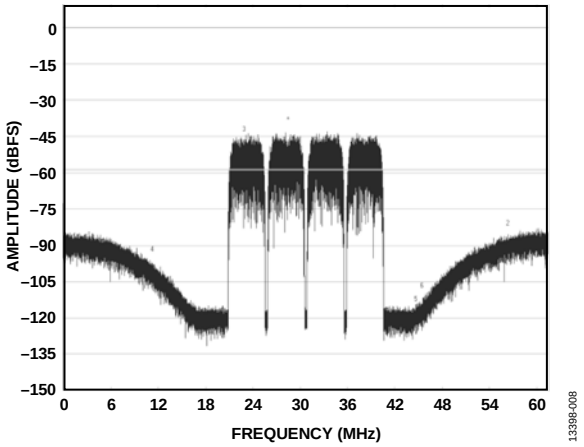

Figure 8. Noise Shaping in an FFT

However, for DPD applications, high-order intermodulation products are contained in this out of band region, as shown in Figure 9, and it is important to adequately suppress them in the output spectrum.

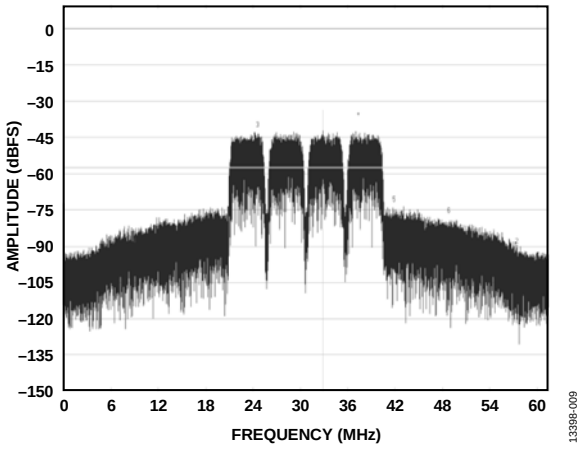

Figure 9. Intermodulation Products in the Output Spectrum

If noise shaping is used in DPD applications, these intermodulation products become obscured and inaccessible. Therefore, noise shaping is not suitable for DPD applications.

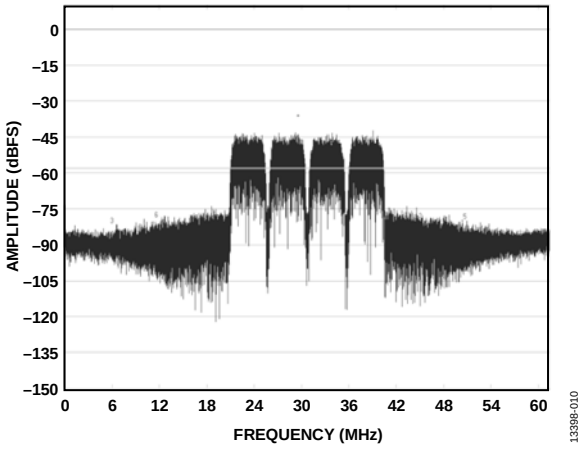

Figure 10. Intermodulation Products Obscured by Noise Shaping

Decimation filtering is also often used in exportable receivers; however, decimation filtering is fundamentally counter to the need for wider bandwidths and is not a solution for DPD applications.

APPLICATION SPECIFIC SOLUTIONS

A careful study of the DPD application provides clues to potential solutions. Figure 11 shows what typically happens in a DPD loop. At onset, the PA output is uncorrected and intermodulation products can be quite large. As the loop converges, the intermodulation products gradually reduce and approach a lower limit determined by the initial error and the sophistication of the correction algorithm. A typical DPD loop convergence is shown in the series of curves in Figure 11. By analyzing both the initial conditions and the convergence process, a number of key points emerge.

First, the total power is shared between several carriers. The ACPR is typically measured relative to one of these signals and not the total power (although they track one another). Therefore, relative to full scale, each time the number of carriers is doubled, the ACPR relative to full scale is reduced by 3 dB.

Second, the composite signal exhibits some peak to rms value. For raw data, this can be up to 13 dB or more depending on the signal statistics. However, to gain the most efficiency from the power amplifier, most systems run some form of crest factor reduction or amplitude compression, which allows higher rms power output without running the PA too far into compression. The dynamic range requirements for the ADC core are similarly relaxed. A round number for post compression peak to rms is about 6 dB. Running too much compression has a detrimental effect on signal quality, potentially disrupting the constellation and increasing the error vector magnitude.

Third, most current generation power amplifiers have raw ACPR in the range of 25 dB to 30 dB. Typical performance is shown in Figure 2. Typical performance varies somewhat by amplifier topology and transistor type; however, Figure 2 is a good reference point. As amplifier efficiency improves, the uncorrected performance typically gets worse and exhibits even worse wideband behavior.

Combining the peak to rms and the ACPR gives typical values for intermodulation of −31 dBFS. Most systems add additional margin and backoff to prevent clipping and to allow other variations. This information can be used to construct a signal model of an uncorrected PA output as shown in Figure 11. This figure shows two types of regions, one consisting of the desired signals shown as the prominent inner most signals, and the other is the intermodulation products shown in the outer regions.

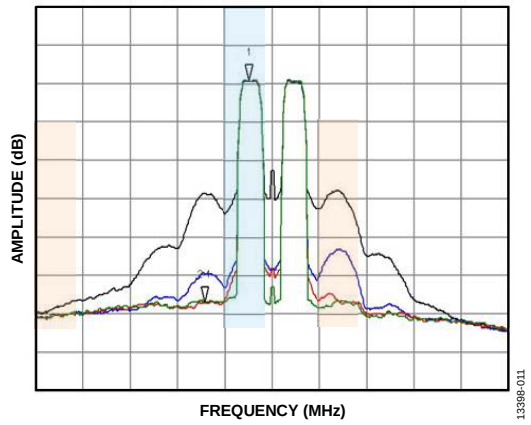

Figure 11. Typical DPD Loop Convergence Example

A receiver offering maximum resolution, even when uncorrected, provides large benefits to improved convergence and performance. If intermodulation products extend above the mask defined by the general DPD model typified in Figure 12, resolution of the receiver must be reduced to comply with export guidelines. Otherwise, full resolution can be provided to the application.

Other solutions exist where full resolution is provided only for short periods of time; however, this action must be synchronized with an external event such as amplifier peaking. Whereas this is straightforward to accomplish because the baseband data is known, it requires signal pins on both the DPD IC and the receiver. Additionally, when the peak event expires, the receiver must revert to low resolution again. Even if the DPD loop was successful at converging during the high resolution mode, updates to the DPD coefficients during low resolution mode run the risk of reducing ACPR and destabilizing the loop when the low resolution mode returns. Alternatively, providing full range and resolution under closed-loop conditions not only improves convergence, but improves loop stability and allows convergence to high levels of performance for the closed-loop system. This is the key advantage of using a technique such as variable dynamic range.

VARIABLE DYNAMIC RANGE

A receiver equipped with variable dynamic range (VDR) provides full converter resolution, sample rate, and dynamic range to the application described. Core converter performance is not impacted in any way other than potential changes to the resolution under conditions that are extraneous to the DPD application. Because the core converter data does not pass directly through any class of filtering, there is no impact on bandwidth, group delay, and latency. Original signal integrity is completely maintained.

VDR works for both real and complex signal chains. When implemented in a real receiver architecture, the main signal region is defined as 25% of the sample rate. In complex receiver architectures, the region can be defined as either 25% of the sample rate or 43% of the sample rate. This region is only the region of the active carriers. The intermodulation products can fall outside of this percentage without impacting performance of the ADC.

A simplified explanation is shown in Figure 12. The central region shown in the white background is where the main PA output signals are placed. The intermodulation products fall into the outer regions. As long as the intermodulation products stay below the red mask, full resolution is provided in the receiver.

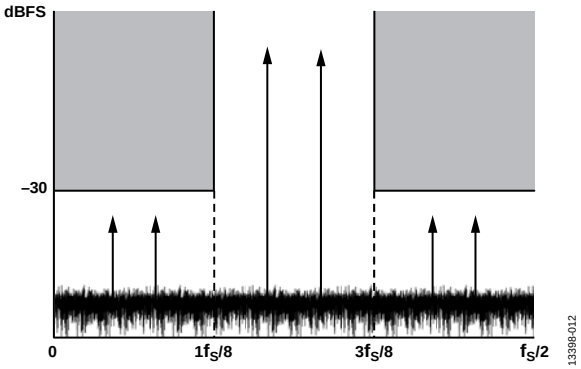

Figure 12. Simplified Diagram of VDR Frequency Bands

Even when the mask is crossed, the resolution is not dropped to the lowest resolution. In fact, the resolution is gradually removed, which further softens the impact to the signal processing. Table 1 shows the defined VDR mask. As the signal in the out of band region increases above −30 dBFS, the resolution is reduced by 1 bit. For the AD6674 and AD6679, this transitions from 14 bits to 13 bits. For each 6 dB increase in signal level above the mask, one additional bit is removed. To further soften the impact, dithering is used to create a continuous transition between levels, with the goal of preventing instability of the loop during this process. Note that within the protected zone of 25% or 43%, full-scale signals do not cause reduction of resolution. This only applies to signals outside of the protected zone, as defined in Figure 12.

Table 1. VDR Reduced Output Resolution Values

Signal Amplitude Violating Defined VDR Mask	Output Resolution
Amplitude ≤ −30 dBFS	14 bits
−30 < Amplitude ≤ −24 dBFS	13 bits
−24 < Amplitude ≤ −18 dBFS	12 bits
−18 < Amplitude ≤ −12 dBFS	11 bits
−12 < Amplitude ≤ −6 dBFS	10 bits
−6 < Amplitude ≤ 0 dBFS	9 bits

Figure 13 to Figure 18 demonstrate what might happen as a DPD loop converges. Figure 13 shows the noise floor when intermodulation extends as high as 6 dBFS. Figure 14 through Figure 18 show what happens as the loop converges and the intermodulation products drop. Although subtle in this representation, the noise floor starts at −149 dBFS/Hz and falls to −153 dBFS/Hz through the series.

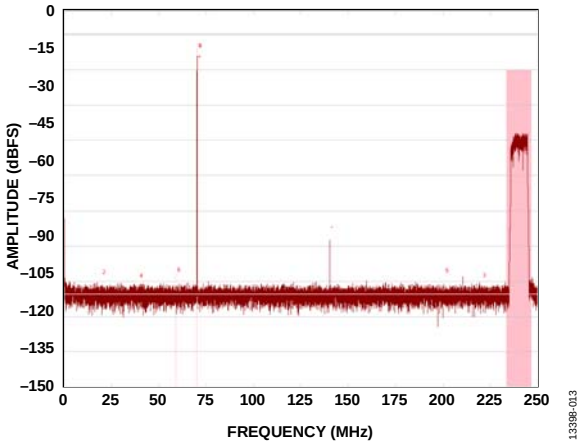

Figure 13. −6 dBFS Out of Band

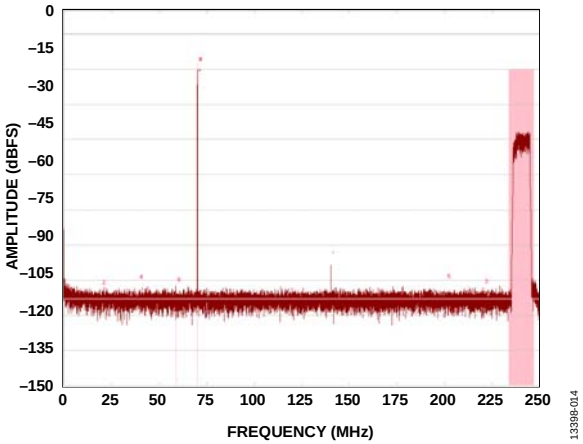

Figure 14. −12 dBFS Out of Band

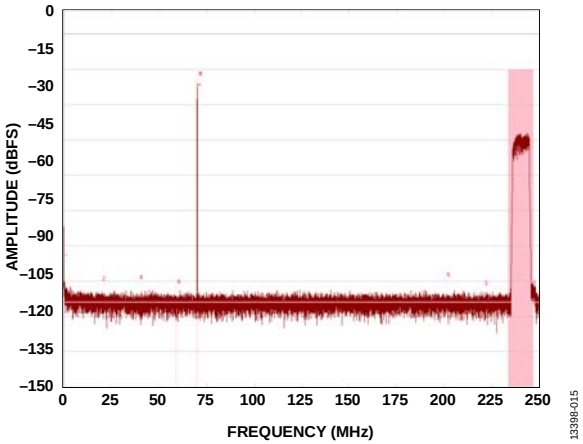

Figure 15. −18 dBFS Out of Band

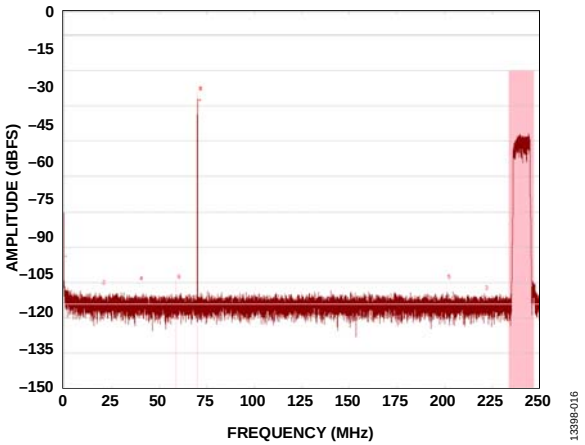

Figure 16. −24 dBFS Out of Band

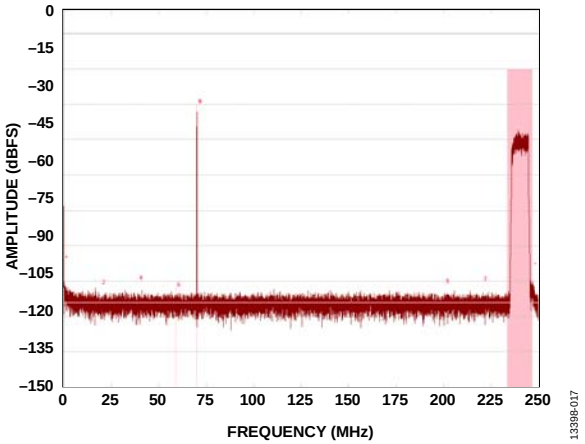

Figure 17. −30 dBFS Out of Band

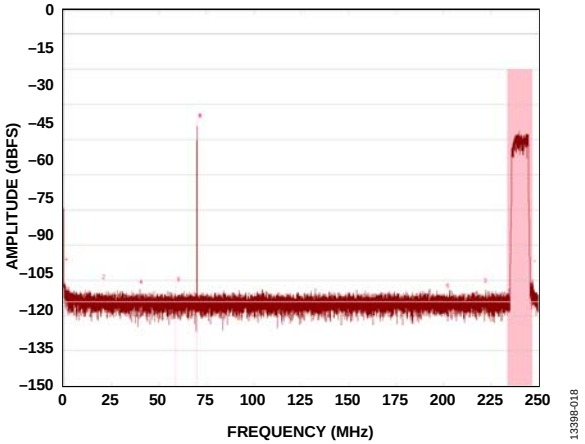

Figure 18. −36 dBFS Out of Band

The protected zone need not be in the middle of the Nyquist zone. This zone can be fully tuned within the Nyquist band. On the AD6674 and AD6679, it can be tuned in increments of $f_S/16$ and can reside on either side of dc or at dc when implemented as a complex receiver, providing significant flexibility in the receiver design.

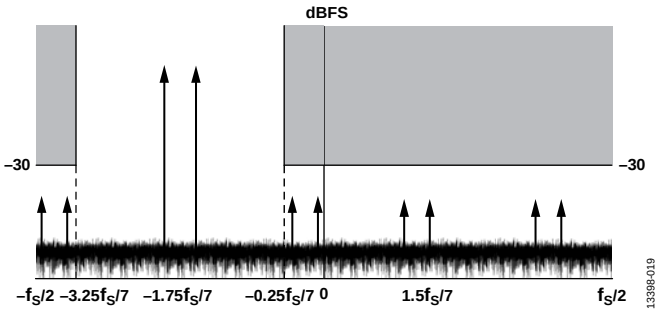

Figure 19. Complex Mode VDR Tuned for Negative Frequencies

Figure 19 shows VDR tuned to the negative side of dc, which is useful for applications where complex IF is implemented in the observation path and results in a negative IF.

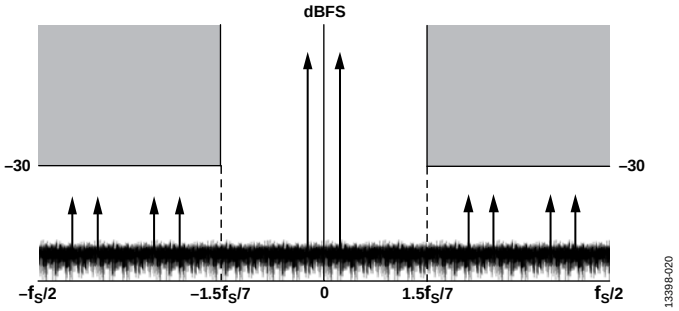

Figure 20. Complex VDR Tuned for ZIF Operation

Figure 20 shows VDR tuned to dc, which is useful for applications implementing a ZIF observation path.

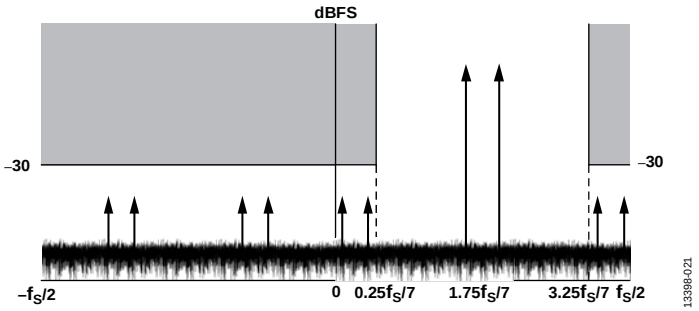

Figure 21. Complex Mode VDR Tuned for Positive Frequencies

Figure 21 shows VDR tuned to the positive side of dc, which is useful for applications where complex IF is implemented in the observation path and results in a positive IF.

VDR IN ACTION

VDR is designed to create little or no impact on the digital predistortion loop. The window created for VDR is designed to overlay with current DPD applications in such a way that full resolution is presented to the processor block under nearly all conditions, including most preconverged states.

In the rare case where the initial DPD samples are so poor that the intermodulation products are outside of the VDR window, resolution of the converter may be limited. However, after seed or first iteration values are applied to the polynomials, performance of the loop improves enough that performance is below the mask. Even when signal amplitudes are above the mask, significant resolution remains, as shown in Table 1. Under the worst conditions for the AD6674 and AD6679, resolution is reduced to nine bits. In this case, intermodulation products must be greater than −6 dBFS, which is significantly above the noise floor created by a 9-bit converter and, therefore, the converter does not dominate performance.

As the system converges, additional resolution is added to the converter. As the intermodulation products fall below −6 dBFS, 10 bits are provided; below −12 dBFS, 11 bits are provided; below −18 dBFS, 12 bits are provided; below −24 dBFS, 13 bits are provided; and, finally, below −30 dBFS, full resolution is available.

In each case, there is significant headroom between the intermodulation product and the converter noise floor so that convergence is not limited. This is loosely demonstrated in Figure 13 through Figure 18.

As previously outlined, resolution largely determines how fast the loop can converge. For a generic algorithm, this is shown in Figure 22. There are several key points to notice. First, the lower the resolution, the slower the convergence. The second point is that at the start of convergence, there is very little difference in the slope of the curve, meaning that for the first few samples, it is of little importance what the resolution actually is.

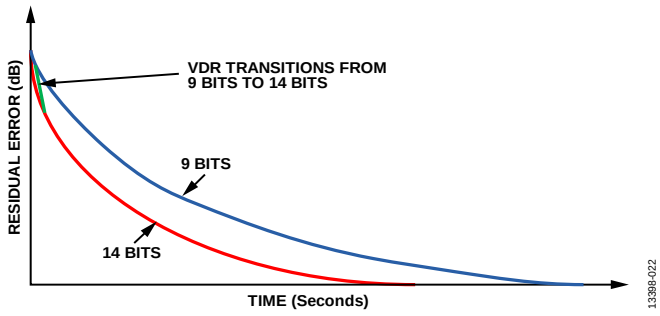

Figure 22. Residual Error vs. Time for 9-Bit and 14-Bit Resolutions

In a worst case scenario, the 9-bit curve may initially be followed. However, the slope is quite high and there is little difference in this region of the curve, and very quickly higher resolution is provided.

For applications that utilize every data sample, such as those from OP6180, VDR provides full resolution in times as short as four converter samples. In applications that provide block processing, full resolution occurs after the first iteration. For both types of applications, the slope transitions from 9 bits to 14 bits, providing acceleration of the DPD loop such that convergence time improves. In the typical case, VDR provides full resolution from the beginning because the VDR mask is defined to allow full resolution even from most unconverged PA outputs.

Perhaps more important than start-up conditions is VDR response to a transient events. VDR continues to provide full resolution during peaking events, whereas other applications remain in low resolution mode or must be synchronized to the DPD application to work properly. VDR has the benefit that there is no need for synchronization between the DPD application and the receiver.

Additionally, VDR is tolerant to sparse events. Because the mask is based on a FIR filter, infrequent events that drop the resolution of the receiver are removed by filtering. Therefore, VDR provides near 100% operation in high resolution mode, improving convergence as shown in Figure 22.

Because VDR provides a very high percentage of operation in high resolution mode, overall DPD performance is much improved over low resolution systems. This performance is shown in Figure 23, which represents residual error. The blue trace (bottommost plot line) shows 9-bit residual error and the red trace (uppermost plot line) shows 14-bit residual error. Because the residual error is lower for high resolution operation, overall convergence is better with the difference clearly shown in Figure 6 and Figure 7.

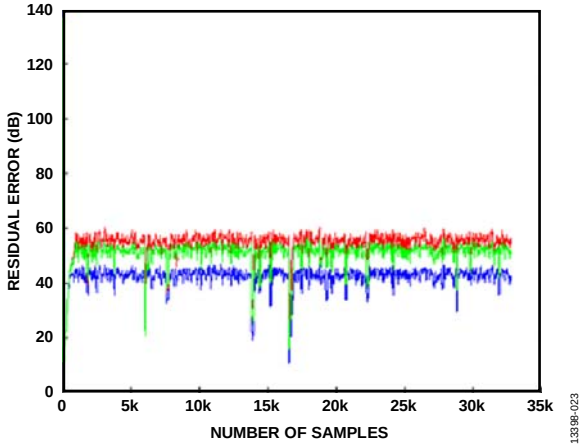

Figure 23. Residual Error for 9-Bit to 14- Bit Resolutions

Figure 24 shows an example where the resolution of VDR is moved from high resolution to low resolution. This extreme case shows how the residual error changes as a function of resolution. One key point to note is that the typical loop must not lose stability as the resolution is changed. In fact, most intermodulation events must not cause disruption of high resolution performance. From a mask point of view, VDR selects between one of two filters: one is 7 taps long and the other is 23 taps. The number of required taps is moderated by the external analog filter step response.

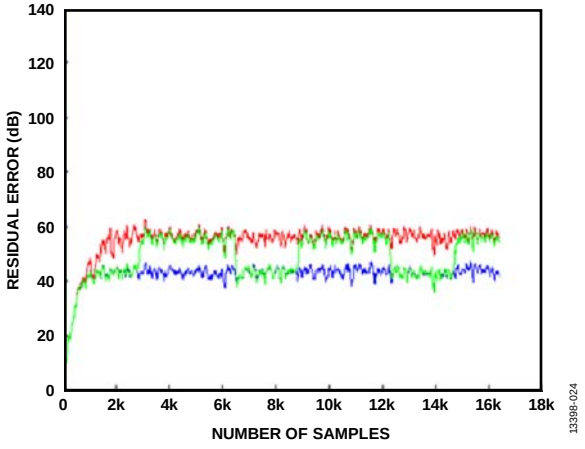

Figure 24. Residual Error with High and Low Resolution with VDR

As shown in Figure 25, the spectral mask is outside of the main signal path. As noted previously, the output data is not directly encumbered by the digital filter. The digital filter is only used to perform spectrum analysis on the data. The response is then compared to the thresholds outlined in Table 1 to determine what output resolution to use.

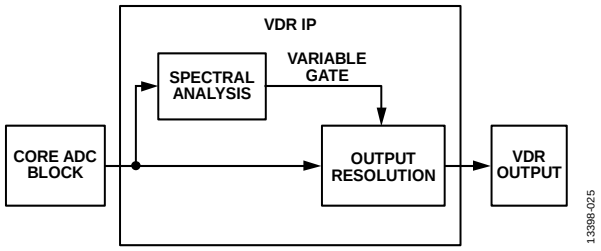

Figure 25. VDR Signal Flow Diagram

The result is that the full dynamic range of the converter is provided when spectral components fall within one of the two defined masks, as shown in Figure 26. As shown, the wideband mask covers most (86%) of the Nyquist band and is completely tunable because internally, data is treated as complex. Whereas all signals, including the intermodulation products, fall within this mask, full dynamic range of the converter is provided. As shown in Figure 26, for nearly all conditions, converter performance is maximum. Even when signals fall in the 14% of spectrum in the transition band of the mask, it must be above the mask to reduce the resolution of the converter.

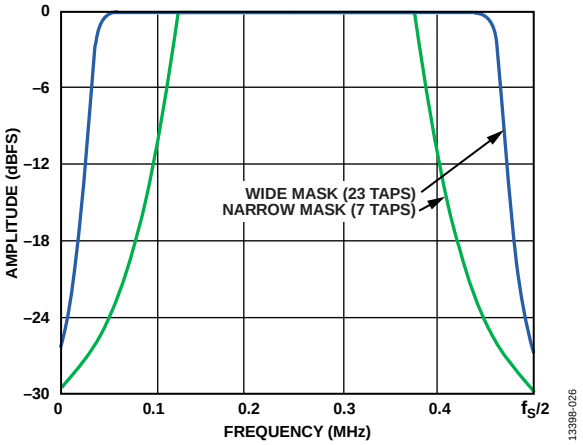

Figure 26. Wide and Narrow VDR Masks

To help explain this, if a signal occurs at $f_S/2$, as long as it is below the mask, full resolution of the receiver is available. This equates to energy below −30 dBFS at this frequency passing without notice. If the energy is 6 dB above the mask, the resolution is reduced by one bit. If the energy in this frequency is 12 dB above the mask, the resolution is reduced by two bits, and so forth. Because the restricted band is relatively narrow, it is not expected to impact DPD applications or performance.

To help visualize this, the devices were characterized as shown in Figure 27 to Figure 29 by sweeping the analog input across both frequencies and amplitude. The resolution curves represent the resulting resolution for the stimulus. Figure 27

shows an example for real mode operation, whereas Figure 28 and Figure 29 show examples for complex mode operation. In these examples, the sample rate is set to a commonly used value of 368.64 MSPS. In Figure 27 and Figure 28, the VDR is tuned to midband. A third sweep is shown in complex mode using the wideband filter mode. In this case, the VDR is tuned to a center frequency offset from midband to show how IF may be shifted if desired.

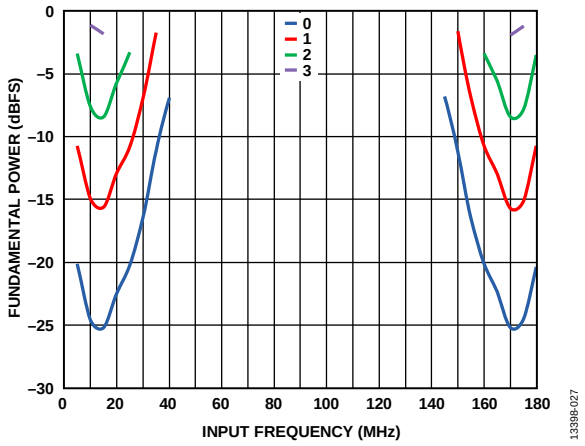

Figure 27. Frequency Sweep in VDR Real Mode

In Figure 27, the primary transmit signal resides between approximately 45 MHz and 145 MHz, yielding 100 MHz of open spectrum. In this range, full resolution is always provided to the application. As long as intermodulation products remain below the blue curve, full resolution continues to be provided to the application. If spectral components increase above the blue curve (bottommost plot line) but below the red curve (third plot line from the top of the graph), the LSB (Bit 0) is set to 0, and the remaining bits are rounded. If spectral components increase above the red curve but below the green curve, the lower two bits (Bit 0 and Bit 1) are set to 0, and the remaining bits are rounded. If the spectral component increases above the green curve (second plot line from the top of the graph) but below the purple curve (uppermost plot line), three bits are removed (Bit 0, Bit 1, and Bit 2). Finally, if the signal increases above the purpose, four bits are removed.

In Figure 28 and Figure 29, for complex operation, the exact same process is applied. In narrow-band mode, 100 MHz of clear spectrum are also provided. In the wideband complex mode shown in Figure 30, approximately 160 MHz of protected spectrum are provided, which is sutiable for time division duplex (TDD) applications.

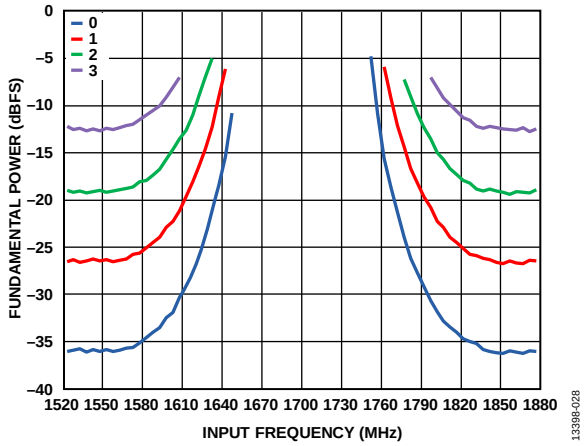

Figure 28. Frequency Sweep in VDR Complex Mode with Narrow-Band Mask

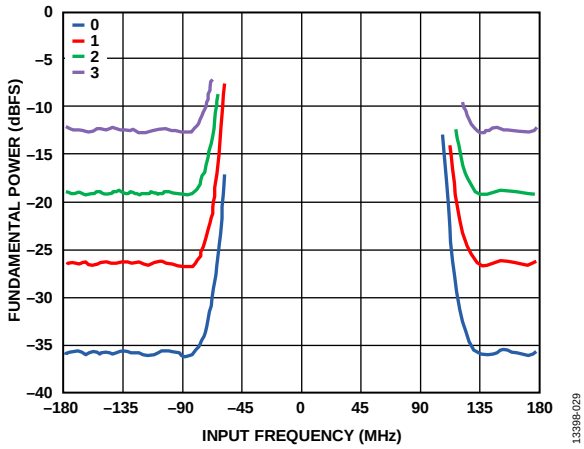

Figure 29. Frequency Sweep in VDR Complex Mode with Wideband Mask

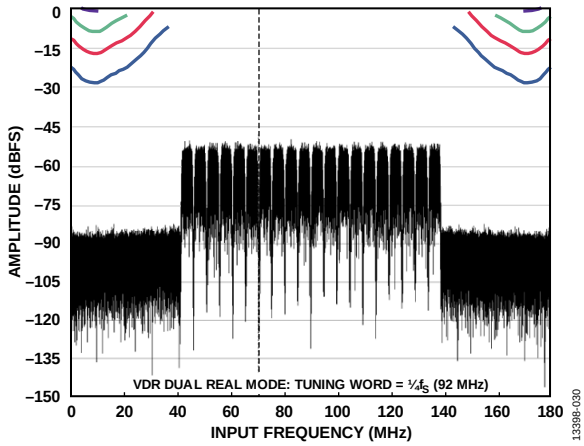

Figure 30. VDR Real Mode (1/4 f_S Tuning) with Multicarrier LTE Signal Overlay

AN13398-0-9/15(0)