

# Ultrawideband Digital Predistortion (DPD): The Rewards (Power and Performance) and Challenges of Implementation in Cable Distribution Systems

By **Patrick Pratt** and **Frank Kearney**

Share on   

## Introduction

The first cable systems in the U.S. started to appear in the early '50s. Even with the rapid changes in technology and distribution methods, cable has maintained a prominent position as a conduit for the distribution of data. New technologies have layered themselves on the existing cable network. This article focuses on one aspect of that evolution—power amplifier (PA) digital predistortion (DPD). It's a term that many involved in cellular system networks will be familiar with. Transitioning the technology to cable brings substantial benefits in terms of power efficiency and performance. With these benefits come substantial challenges; this article dives deep into some of these challenges and provides an overview as to how they may be solved.

## Understanding the Requirements

When power amplifiers are operated in their nonlinear region, their output becomes distorted. The distortion can affect the in-band performance and may also result in unwanted signal spilling over into adjacent channels. The spill-over effect is particularly important in wireless cellular applications, and adjacent channel leakage ratio—or ACLR, as it is termed—is tightly specified and controlled. One of the prominent control techniques is digitally shaping or predistorting the signal before it gets to the power amplifier so that the nonlinearities in the PA are cancelled.

The cable environment is very different. Firstly, it can be regarded as a closed environment; what happens in the cable stays in the cable! The operator owns and controls the complete spectrum. Out-of-band (OOB) distortions are not a major concern. In-band distortions are, however, of critical importance. The service providers have to ensure the highest quality in-band transmission conduit so that they can leverage the maximum data throughputs. One of the ways that they ensure this is by running the cable power amplifier strictly within its linear region. The trade-off for this mode of operation is very poor power efficiencies.

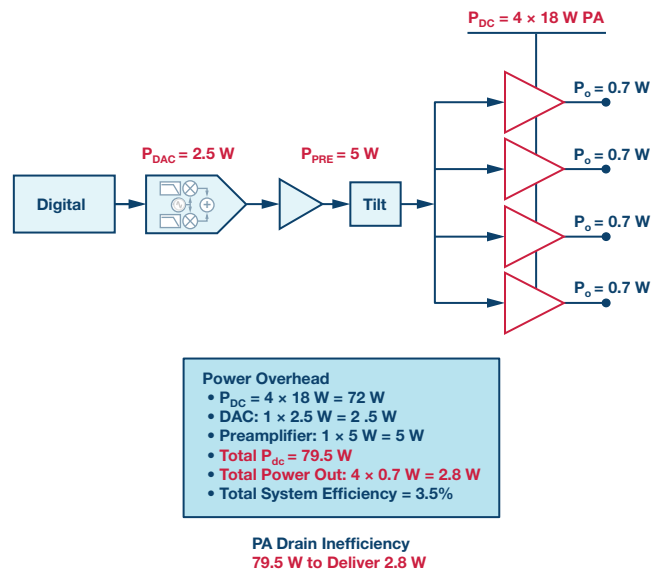


Figure 1. Power efficiencies in cable power amp drivers.

Figure 1 provides an overview of a typical cable application. Although the system consumes nearly 80 W of power, just 2.8 W of signal power is delivered. The power amplifiers are very low efficiency Class A architectures. The maximum instantaneous peak efficiency can be calculated to be 50% (when the signal envelope is at maximum, assuming inductive loading). If the PA is to operate entirely in its linear region, then taking into account the very high peak to average ratio of the cable signals (typically 14 dB) means that the amplifiers need to operate on average 14 dB below the start of compression, hence ensuring that no signal compression occurs even at the peaks of the signal. There is a direct correlation between the back-off and the amplifier operating efficiency. As the amplifier is backed off 14 dB to accommodate the full range of cable signals the operating efficiency will reduce by  $10^{-14/10}$ . Hence, the operating efficiency drops from its theoretical max of 50% to  $10^{-14/10} \times 50\% = 2\%$ . Figure 2 provides an overview.

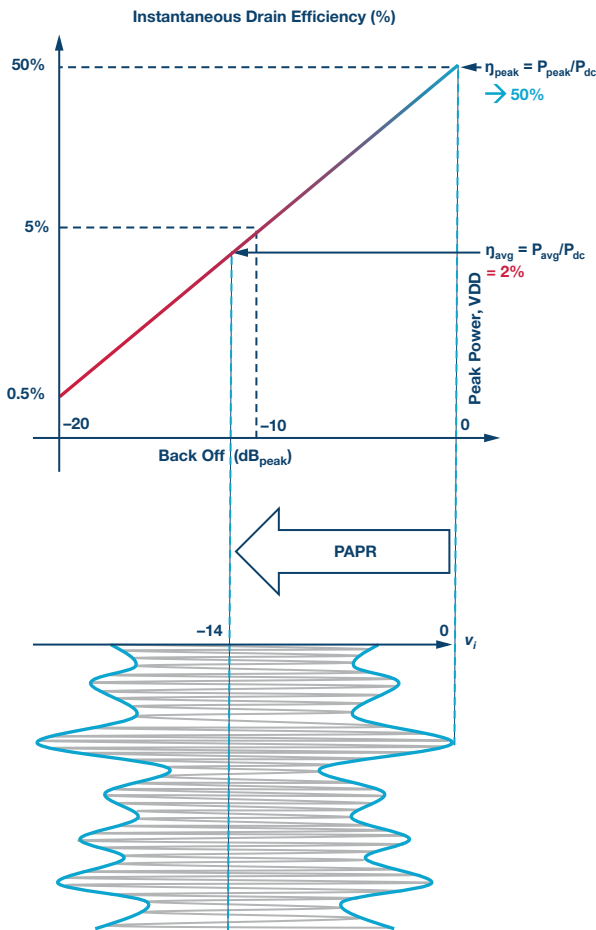


Figure 2. High peak to average ratios push a back-off operational mode and a dramatic decrease in efficiency.

In summary, power efficiency is the major issue. The lost power has a cost implication, but, just as importantly, it also uses up a scarce resource within the cable distribution system. As cable operators add more features and services, they require more processing, and the power for that processing may be constrained within existing power budgets. If wasted power can be retrieved from the PA inefficiencies, then it can be reallocated to those new functions.

The proposed solution to the PA inefficiency is digital predistortion. It's a method universally adopted and employed right across the wireless cellular industry. DPD allows the user to operate the PA in a more efficient, but more nonlinear region, and then pre-emptively correct for the distortions in the digital domain before the data is sent to the PA. DPD is essentially shaping the data before it gets to the PA to counteract the distortions the PA will produce, and hence extend the linear range of the PA, as shown in Figure 3. That extended linear range can be used to support higher quality processing, deliver lower modulation error rates (MER),<sup>1</sup> or allow the PA to run at a reduced bias setting—thus saving power. Although DPD has been widely used in wireless cellular infrastructures, implementing DPD in a cable environment has unique and challenging requirements.

As shown in Figure 4, actual operating efficiencies for the cable application sit at approximately 3.5%! Implementing DPD results in the power requirements of the system dropping from 80 W to 61 W—a power saving of 19 W, which is a 24% reduction. Previously each PA required 17.5 W of power; now that drops to 12.8 W.

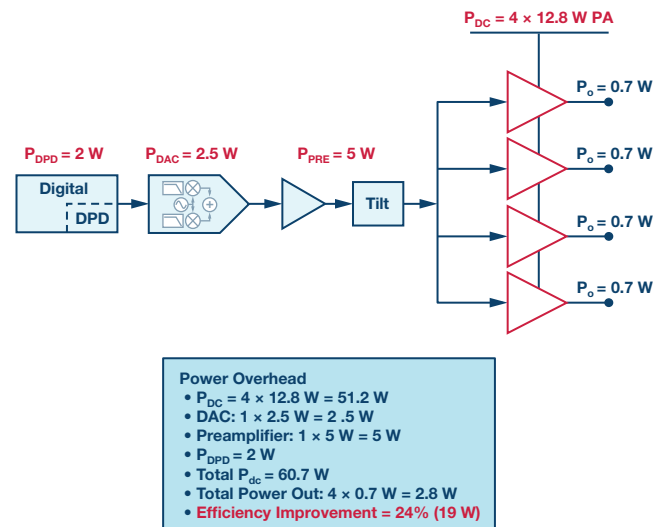


Figure 4. Overview of power saving through DPD implementation.

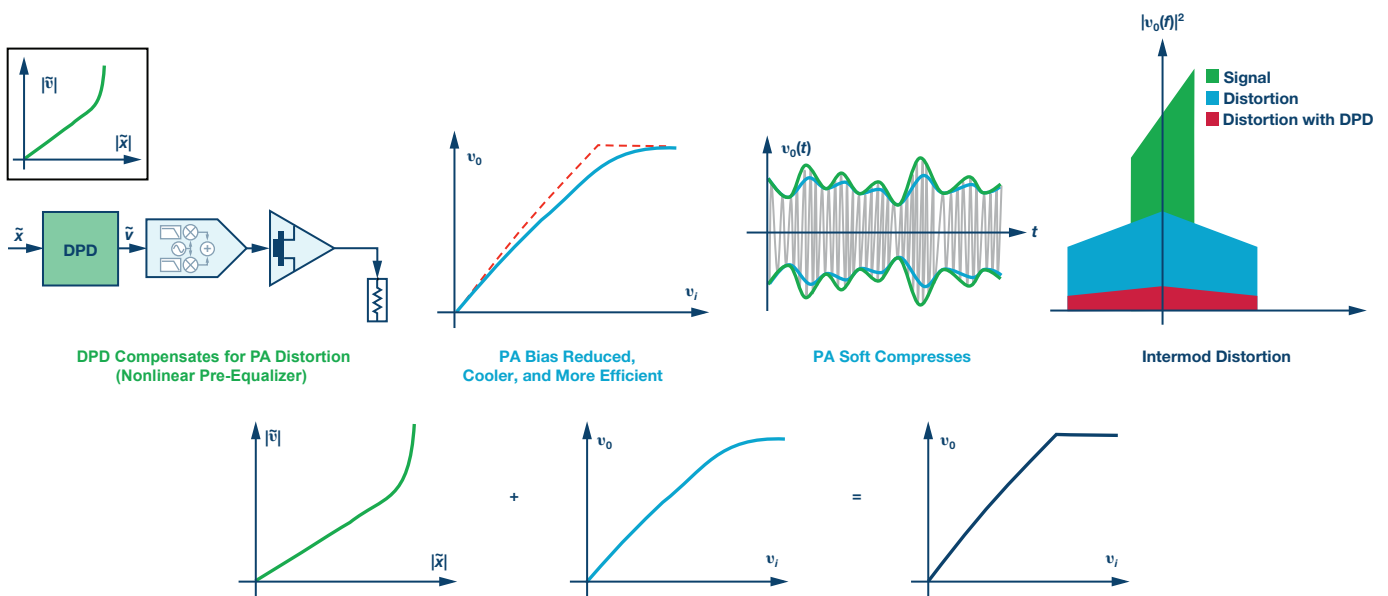


Figure 3. Digital predistortion overview.

## The Challenges of Implementation

The value of DPD is clear but the cable application poses many unique challenges to its implementation. Those technical challenges have to be met within the confines of the available resources. For example, the solution itself must be power efficient, as there is little value in optimizing the PA efficiency if the power saved gets shifted into powering the solution. Likewise, the digital processing resources need to be appropriate so that they can reside efficiently within the current FPGA architectures. A very large/complex algorithm with nonstandard hardware requirements and extensive architectural change is unlikely to get adapted.

## Ultrawide Bandwidth

Perhaps the most prominent difference between the cable application and that of the wireless cellular environment is the bandwidth of operation. In cable, there is approximately 1.2 GHz of bandwidth to linearize. The wide bandwidth challenge is compounded by the fact that the spectrum starts just 54 MHz from dc and that the signal bandwidth is larger than the channel center frequency. We must remember that the power saving is going to come by driving the PA into its nonlinear region of operation; that delivers better efficiency, but at the cost of generating nonlinear products. The DPD has to cancel the nonlinearities created by the PA, with a particular focus on those that fall back in-band of the wanted signal. That poses a unique challenge in the cable application.

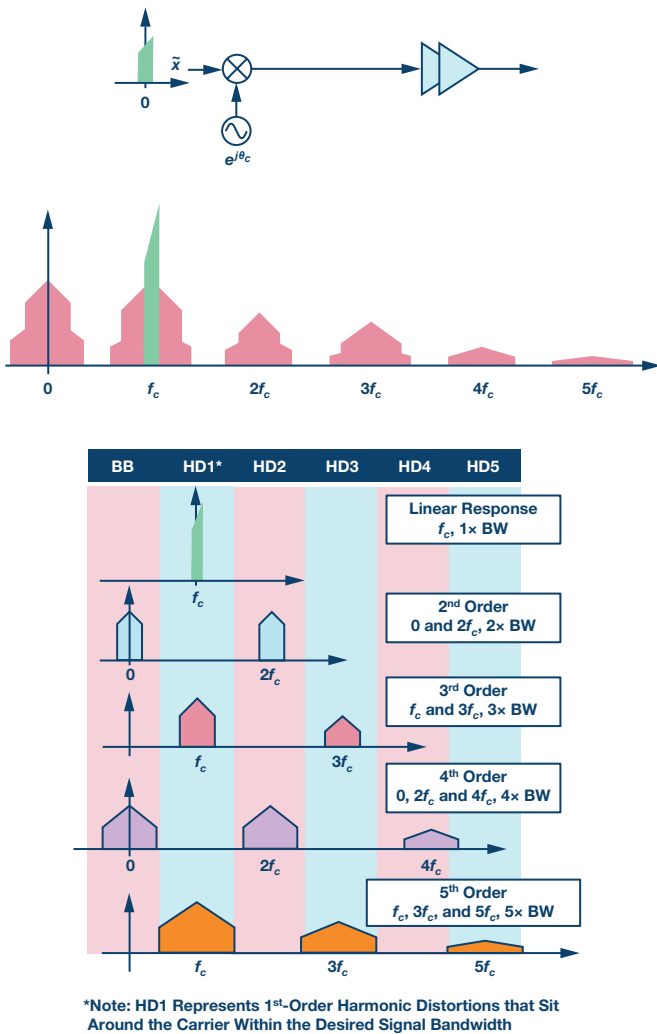


Figure 5. Harmonic distortion terms in conventional narrow-band interpretation.

Figure 5 provides an overview of the wideband harmonic distortion terms we might expect for a conventional, narrow-band (narrow-band is defined later in this section) upconverted baseband signal that goes through a nonlinear amplification stage. The nonlinear PA output is typically described by a power series expression, such as the Volterra series having the form of

$$y(n) = \sum_{m_1=0}^M \dots \sum_{m_k=0}^M h_{k,m_1,\dots,m_k} \prod_{q=1}^k x(n-m_q)$$

which can be understood as a generalization of the Taylor power series to include memory effects. Of fundamental importance to note is that each nonlinear term ( $k = 1, 2, \dots, K$ ) generates multiple harmonic distortion (HD) products. For example, the 5<sup>th</sup> order has 3 terms: 5<sup>th</sup> order at 1<sup>st</sup> harmonic, 5<sup>th</sup> order at 3<sup>rd</sup> harmonic, and 5<sup>th</sup> order at 5<sup>th</sup> harmonic. Note also that harmonic bandwidth is a multiple of its order; for example, 3<sup>rd</sup>-order terms are three times wider than the stimulus bandwidth.

In cable, it is not so much the large signal bandwidth, but its positioning on the spectrum (just 54 MHz from dc), that poses a particular challenge to DPD. Harmonic distortion occurs in all nonlinear systems; the focus of cable DPD is the harmonic distortion that fall in band. Looking at Figure 5, it can be seen that in conventional narrow-band applications, the focus will be the 3<sup>rd</sup>- and 5<sup>th</sup>-order harmonics. Although others are created, they fall out of the band of interest and can be removed through conventional filtering. We can define wideband and narrow-band applications by their fractional bandwidth where the fractional bandwidth is defined as

$$\frac{(f_n - f_l)}{f_c}$$

( $f_n$  = highest frequency,  $f_l$  = lowest frequency, and  $f_c$  = center frequency). When the fractional bandwidth goes above 1, the application can be deemed wideband. Most cellular applications have fractional bandwidth of 0.5 or lower. Therefore, their HD behavior adheres to the characteristics shown in Figure 6.




Figure 6. Narrow-band simplification; only products around the 1<sup>st</sup> harmonic need be considered.

For such narrow-band systems, only in-band distortion around the 1<sup>st</sup> harmonic needs be cancelled by the DPD since a band-pass filter can be employed to remove all other products. Note also, as no even order products fall in band, the DPD needs only treat odd order terms.

In the cable application we can approximate  $f_n \sim 1200$  MHz,  $f_l \sim 50$  MHz, and  $f_c \sim 575$  MHz, thus giving us a fractional bandwidth of 2. To determine what minimum HD order needs to be corrected, the equation

$$f_l K_{min} < f_n$$

( $K_{min}$  is the lowest nonlinear order to be considered) can be used, or numerically  $50 \text{ MHz} \times 2 = 100 \text{ MHz}$ , which is less than 1200 MHz—hence the 2<sup>nd</sup>-order HD falls well within the band of operation and must be corrected. Therefore, if the decision is made to operate the cable PA outside of its very safe and linear operation, the resultant harmonic distortion will be as indicated in Figure 7.

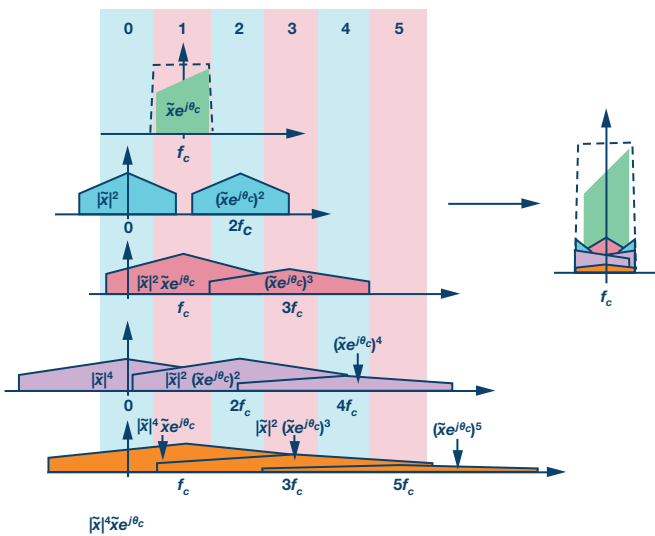


Figure 7. The effects of broadband harmonic distortion in wideband cable applications.

In contrast to wireless cellular where only the odd order harmonics were of concern, within the cable application both even and odd terms fall in band, creating multiple, overlapping zones of distortion. This has some serious implications for the complexity and sophistication of any DPD solution, as the algorithm has to move past simple narrow-band assumptions. The DPD solution must accommodate terms for each order of harmonic distortion.

In a narrow-band system, the even order terms can be ignored and the odd order produce 1 term each within the band of interest. DPD in the cable application has to concern itself with both odd and even order harmonic distortions and it also must consider that each order can have multiple overlapping in-band elements.


### Positioning the Harmonic Distortion Corrections

Considering a conventional narrow-band DPD solution where processing is done at complex baseband, we are principally concerned with harmonic distortions that sit symmetrically around the carrier. In wideband cable systems, although that symmetry is maintained for those terms located around the 1<sup>st</sup> harmonic, that symmetry no longer holds for the higher harmonic products.


As shown in Figure 8, conventional narrow-band DPD is done at complex baseband. In these instances, only the 1<sup>st</sup> harmonic products fall in band such that their baseband representation translate directly to RF. When we consider wideband cable DPD, the higher harmonic distortions must be frequency offset so that baseband representation after upconversion is positioned correctly in the real RF spectrum.

### Loop Bandwidth Limitations:

Closed-loop DPD systems employ a transmission and an observation path. In an idealized model, neither of the paths would be bandwidth restricted and both would be wide enough to pass all DPD terms; that is, both in-band and out-of-band terms are passed.



(a) Conventional Narrow-Band DPD Processing Done at Complex Baseband.



(b) Wideband Cable DPD, OOB HD Must be Frequency Offset to Allow for RF Upconversion.

Figure 8. Note the frequency offset requirements in the complex baseband processing for wideband DPD.

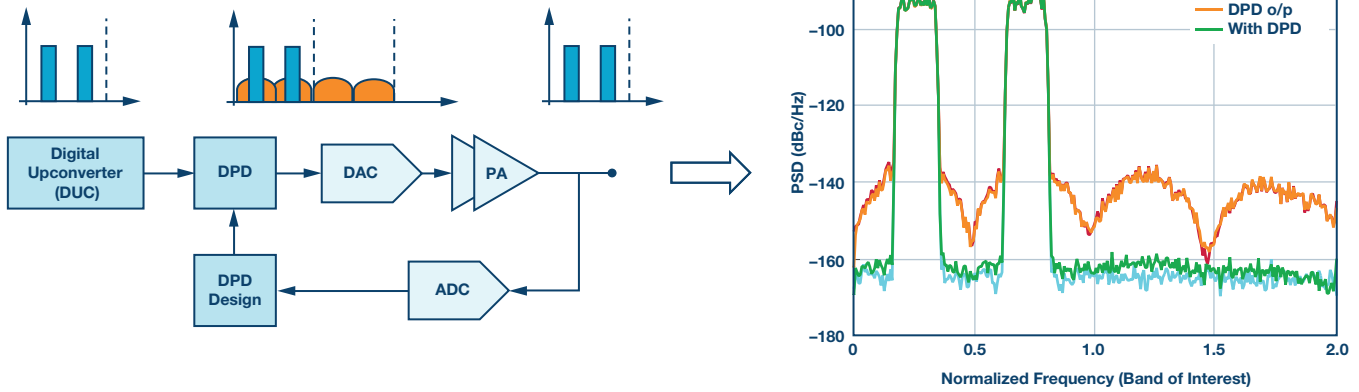


Figure 9. Idealized DPD implementation with no bandwidth restrictions.

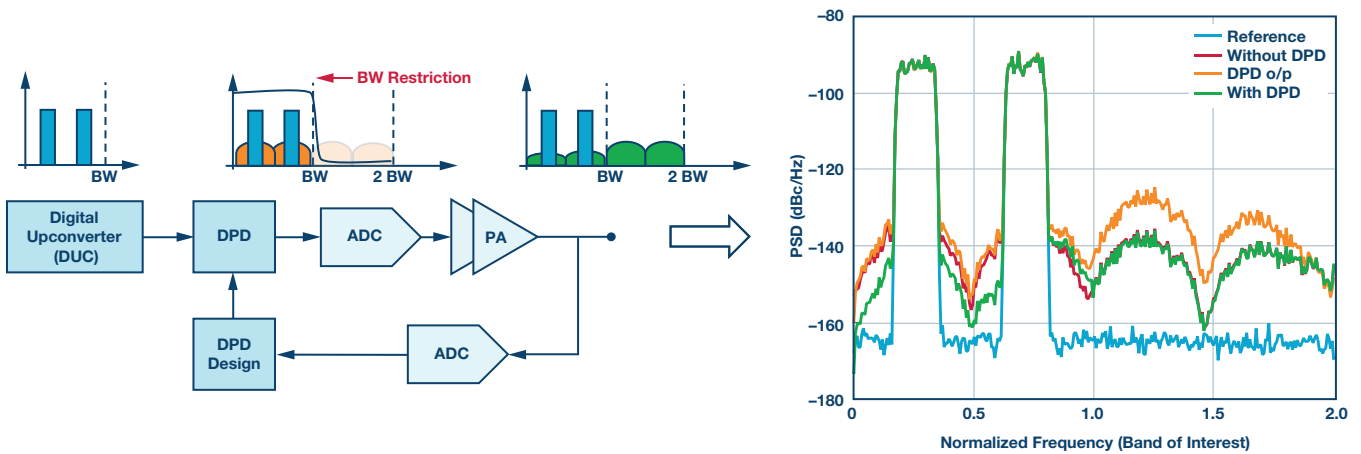


Figure 10. Decreased performance of DPD as bandwidth limitations in the signal path limit the OOB terms.

Figure 9 provides an overview of a DPD implementation. In an ideal situation, the path from the digital upconverter (DUC) via the DPD to the DAC and through the PA will have no bandwidth restrictions. Likewise, the ADC on the observation path will digitize the full bandwidth (note that for illustration purposes we are showing a 2× bandwidth signal path; in some wireless cellular applications that may extend to 3× to 5×). The ideal implementation has the DPD produce terms both in-band and out of band that totally cancel the distortion introduced by the PA. It is important to note that for accurate cancellation, terms are created well outside the bandwidth of the signal of interest.

In a practical implementation, the signal path has bandwidth limitations that modify the DPD performance from the ideal implementation.

In the cable application, the bandwidth limitations can come from a variety of sources: The JESD link between the FPGA and the DAC, the DAC anti-imaging filter, and the PA input matching. The most notable effect of these limitations is the OOB performance. As can be observed on the simulation shown in Figure 10, the DPD fails to correct the OOB distortions. In cable, where OOB distortions contribute to in-band performance degradation, this can be of particular significance; bandwidth restrictions in the signal path can, and do, affect the in-band performance.

The cable environment is unique in that the operator owns the entire spectrum. Emissions that fall out of the band of interest (54 MHz to 1218 MHz) are in a portion of the spectrum not used by anyone else and are also subject to decay due to the inherent cable losses at high frequency. The observation path need only be concerned with monitoring what is happening within the band of operation.

An important distinction needs to be made here; emissions that fall out of band are not of concern, but those that are generated out of band and extend back down into band are. Hence, although OOB emissions are not of concern, the terms that create them are. The implementation is very different to the wireless cellular application where the observation bandwidth requirement is typically 3× to 5× that of the band of operation. In cable, the focus is in-band performance, and OOB terms need only be considered regarding their effect on the in-band performance.

Cable DPD only needs to correct for in-band products: 54 MHz to 1218 MHz for DOCSIS 3. The DPD generates 2<sup>nd</sup>, 3<sup>rd</sup>, ... cancellation terms. Although we only need to correct over the cable BW, within the DPD actuator these terms extend over a wider BW (for example, the 3<sup>rd</sup> order extends out to 3× 1218 MHz). To preserve stability of a conventional DPD adaptation algorithm, these OOB terms should be preserved around the loop. Any filtering of the DPD terms tends to destabilize the adaptation algorithm. In a cable system there is band limiting and thus a conventional algorithm can fail.

## DPD and Cable Tilt Compensation

As with all other mediums of transmission, cable introduces attenuation. In general, this attenuation can be considered as functions of the cable quality, the cable run distance, and the frequency of transmission. If a relatively uniform received signal strength is to be achieved at the receiving end of the cable, right across the spectrum of operation, then pre-emphasis (tilt) must be added at the transmit side. The tilt can be regarded as an inverse transfer function of the cable. It applies pre-emphasis, or shaping, that is proportional to the frequency of transmission.

The shaping is achieved by a low power passive analog equalizer known as a tilt compensator, located just prior to the power amplifier. Little or no attenuation is applied at the high frequencies, whereas the maximum attenuation is applied to the lower frequencies. Signals at the output of the tilt compensator can have level variations of up to 22 dB across the spectrum of operation.




Figure 11. Tilt compensator implementation.

The tilt compensator shapes the signal, and that shaping profile is maintained as the signal is processed through the PA. Conventional DPD implementation would see the shaping as an impairment and try to correct for it as DPD is an (nonlinear) equalizer. It might seem reasonable to suggest that if the inverse to the tilt was added to the observation path that it would mitigate the effect. However, this is not true. Because the PA is nonlinear, commutativity doesn't hold—or, in other words,

$$PA\{T\} \neq T\{PA\}$$

(PA is the model of the power amplifier and T is the model of the tilt compensator).

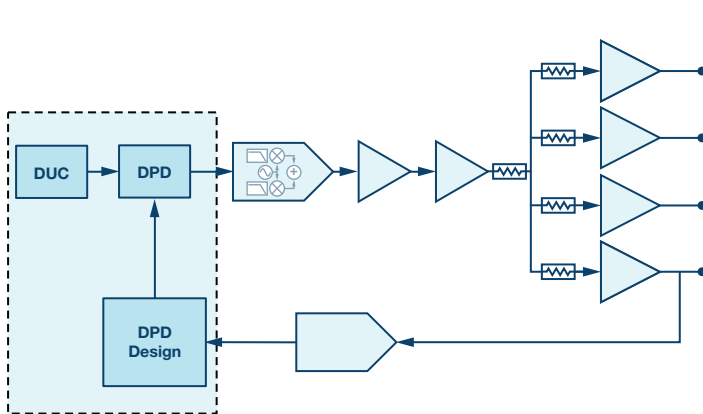


Figure 12. Single DPD with multiple PAs (simulated results).

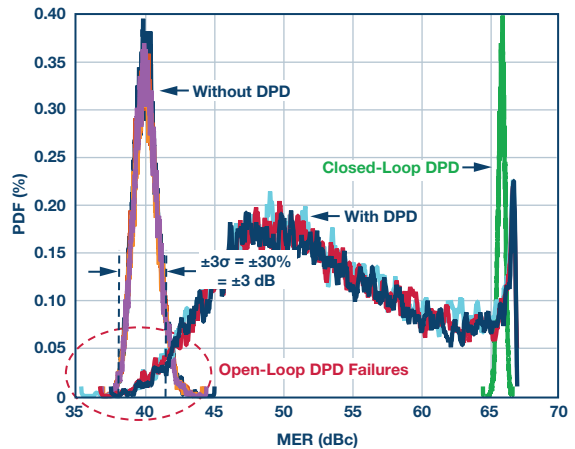
For optimal operation, the DPD processing block needs an explicit knowledge of the signal that will present itself at the input to the PA. In a cable DPD application, the tilt compensation must be maintained while simultaneously having the DPD algorithm model the PA. This presents some very unique and difficult challenges. We need a low cost, stable solution that doesn't equalize out the tilt. While the nature of the solution can't be disclosed in this article, ADI has found an innovative solution to this problem, which may be described in detail in future publications.

## DPD and Cable PA Architecture

As shown in Figure 4, the typical cable application will have the output from one DAC split and supplied to four separate PAs. To get maximum power savings, the DPD needs to be implemented on all of those PAs. One possible solution could be the implementation of four independent DPD and DAC blocks. The solution works but the efficiencies are reduced and the system implementation costs increased. The additional hardware has a dollar and a power cost.

Not all PAs are created equal and, although process matching (during manufacturing) may deliver units that have similar personalities, differences will persist and may grow larger with aging, temperature, and supply variations. Having said that, using one PA as the master and developing an optimized DPD for it, which is then applied to other PAs, does deliver a system performance benefit as shown from the simulated results in Figure 12.

The plots to the left-hand side indicate the PA performance with no DPD applied. The nonlinear operational mode causes distortions and this is reflected in the MER<sup>1</sup> performance, having a range of 37 dBc to 42 dBc. Closed-loop DPD is applied by observing the output from the master PA; the green plot to the right-hand side of the diagram shows the enhanced performance. The DPD has corrected for the PA distortions and the result is that the overall performance has been shifted to deliver a MER ranging from 65 dBc to 67 dBc. The remaining plots in the middle show the performance of the slave PAs—that is, the PAs that have been corrected based on the master PA. As can be seen, implementing closed-loop DPD by just observing one PA has benefited the performance of all the PAs. However, the performance of the slave PAs continues to have points of operation that will fail. The span of performance on the slave PAs ranges from 38 dBc to 67 dBc. The wide range itself is not the concern, but a portion of that range falls below the acceptable operational threshold (typically 45 dBc for cable).



The unique system architecture in cable provides an extra challenge for DPD. Optimized performance requires closed-loop DPD implementation. However, conventional thinking would suggest that to do so in cable would require additional hardware in each of the PA paths. An optimal solution needs to deliver the enhancements of closed-loop DPD to each PA but without the additional hardware costs.

### Solving the Challenges with SMART Algorithms

As described earlier in this article, cable DPD presents the designer with particularly unique and difficult challenges. The challenges have to be resolved, but within the constraints of power and hardware so that the advantages are not eroded; there is little value in saving PA power if that power is used in additional DACs or the FPGA. Likewise, the power saving must be balanced with a hardware cost. ADI solves the challenge using a combination of high performance analog signal processing, coupled with advanced algorithm implementation.

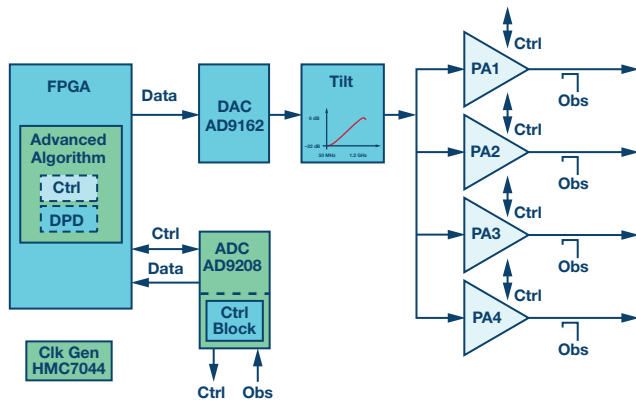


Figure 13. Cable DPD implementation using advanced converters and SMART algorithms.

A high level overview of the ADI implementation is shown in Figure 13. The solution can be regarded as having three key elements: The use of advanced converter and clocking products, an architecture that supports a comprehensive signal chain monitoring/control, and, lastly, an advanced DPD algorithm that can utilize the former knowledge to deliver optimal performance.

The algorithm sits at the heart of the solution. It uses its extensive knowledge of the signal being processed and the transfer function of the signal path to shape the output while simultaneously adjusting the dynamic control of some aspects of the signal path. The dynamic system solution means not only does the system designer have the capability to obtain

substantial power saving, but that those power savings can be directly traded off against performance. The algorithm is such that once the user defines the MER<sup>1</sup> performance level to which the system must operate, the system tuning is implemented so that performance is achieved across all the outputs. It's important to note that the algorithm also ensures that the performance threshold is achieved while maintaining optimal power usage at each PA; no PA takes more power than required to achieve the target performance.

The previous paragraph provides an overview of the solution implementation. The specifics of the algorithm itself are ADI proprietary IP and beyond the scope of this article. The SMART algorithm has the ability of learning the system path and then changing both the nature of the data being transmitted through the path and the characteristics of the path itself to deliver optimal results. We define optimal results as maintaining the quality of the MER while simultaneously lowering the power requirement.

The path characteristics, along with the nature of the transmitted signal, are in constant flux. The algorithm has the self-learning capability to deal with that dynamic adaptability. What's more, the adaption happens while the system is live without interrupting or distorting the transmitted streams.

### Conclusion

The cable environment continues to be an important infrastructure for the delivery of data services. As the technology evolves, so too do the pressures on spectrum and power efficiencies. The next generation of developments is calling for ever increasing demands and pushing higher order modulation schemes and better power efficiencies. These enhancements must come without affecting system performance (MER), and while DPD provides one possible avenue of enablement, its implementation in the cable application poses unique and difficult challenges. ADI has developed a total system solution to address these challenges. That solution encompasses silicon (DACs, ADCs, and clocks), PA control, and advanced algorithms. The combination of all three technologies provides users an adaptable solution where they can easily trade off power and performance requirements with minimal compromise. This software-defined solution also supports a painless transition to the next generation of cable technologies that are expected to incorporate full duplex (FD) and envelope tracking (ET).

*Note 1: Modulation error rate is a measure of the quality of modulation. It represents the difference between the target symbol vector and the transmitted symbol vector.  $MER = 10\text{Log}(\text{average signal power}/\text{average error power})$ . It can be regarded as a measure of accuracy of the symbol placement within the constellation.*

Patrick Pratt [patrick.pratt@analog.com] is a senior research scientist heading up Analog Devices cable DPD algorithm development. His career spans over 30 years and encompasses algorithm research and development activity, both in private organizations and academic institutions. Patrick holds a Ph.D. in electronic engineering from Cork Institute of Technology.



Patrick Pratt

Frank Kearney [frank.kearney@analog.com] works as algorithm development manager in the Communication Systems Engineering Team at Analog Devices Limerick. He has worked at ADI since graduation in 1988. He recently returned from China where he held the position of senior applications manager for the Systems Engineering Team in Asia Pacific. He is currently a doctoral candidate at University College Dublin.



Frank Kearney

Also by this Author:

[Passive Intermodulation \(PIM\) Effects in Base Stations: Understanding the Challenges and Solutions](#)

Volume 51, Number 1