

# How to Make a Digital Predistortion Solution Practical and Relevant

For digital predistortion implementations, static quantitative data fails to capture many of the challenges, risks, and performance tradeoffs of real-world scenarios. Here's how to get beyond fundamentals and into considerations for complex 5G environments.

In an ideal world, the output of a power amplifier (PA) would be an identical scaled version of the input and most of the power consumed by the amplifier would reside in the output signal. Hence, we would have maximum efficiency and no distortion. In the real world, though, we fall short—real linear amplifiers tend to have very poor efficiencies.

Amplifiers used in cable distribution systems, for example, have excellent linearity, but this comes at the cost of efficiency. In most cases, the efficiency struggles to achieve greater than 6% with the balance of the power (94%) being wasted, which imposes economic, environmental, and application costs. In cellular base stations, electricity accounts for over 50% of the operating-expense (OPEX) costs.

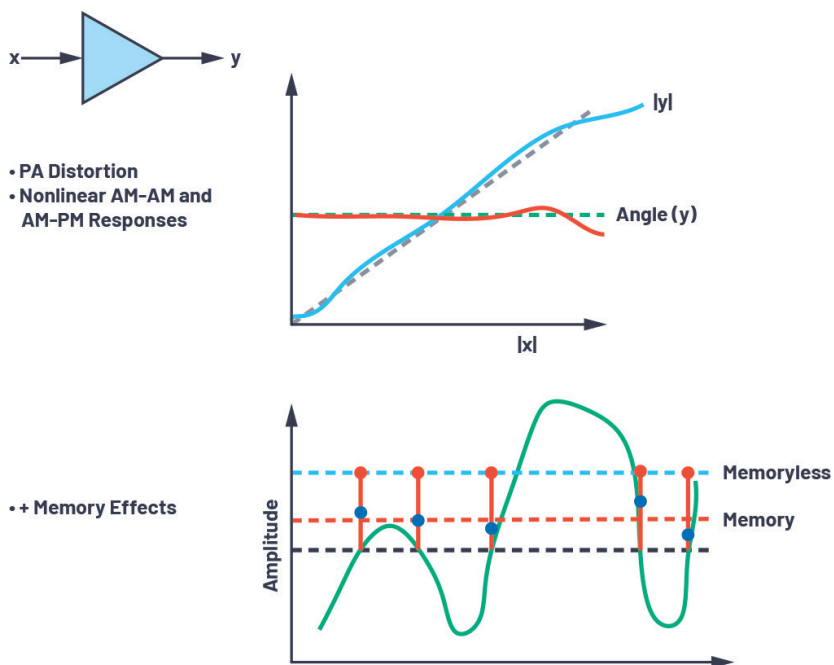
Wasted power increases electricity usage and produces greenhouse gases, while much of the power that isn't emitted as radio waves must be dissipated as heat. Consequently, active and passive thermal management is needed.

Over the last several decades, the cellular industry has pushed the efficiency of the PA to a performance level more than 50%. This has been achieved by adopting smart architectures such as the Doherty architecture and advanced process technologies like GaN.

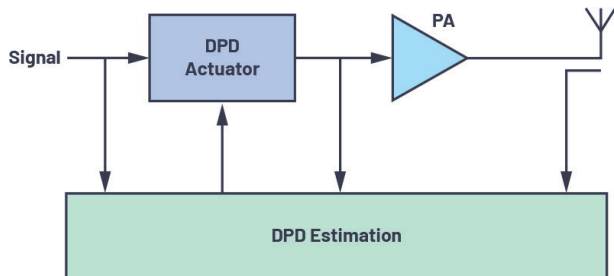
However, higher efficiency comes at a cost—linearity. Poor linearity in cellular systems has two principal consequences: in-band distortions and out-of-band

emissions. In-band distortions disrupt the fidelity of the transmitted signal and can be represented by a degradation in error-vector-modulation (EVM) performance. Out-of-band emissions break the 3GPP emissions mask and may cause unwanted interference to operators occupying adjacent channel frequency allocations. We typically measure this aspect of performance in terms of adjacent-channel leakage ratio (ACLR).

GaN PAs offer an additional challenge in that in-band



1. These plots depict a power amplifier's dynamic transfer function with memory effects.



2. Shown is a conceptual representation of a digital predistortion system.

distortions also are produced by the charge-trapping effect. They're dynamic in nature and unrelated to any signal-to-noise ratio (SNR) implied from the ACLR.

Correcting the PA's nonlinearity is essential. It's a reasonable assumption that if one knew the transfer function of the PA, employing its inverse on the data would nullify the nonlinearities. However, the PA has what may be considered a dynamic transfer function; its output-to-input characteristics can be thought of as continuously in flux.

Furthermore, the dynamic transfer function (*Fig. 1*) depends on a combination of the PA characteristics (including power, voltage, and temperature), the input signal presented to the PA, and prior signals that the PA has processed (memory effects). The dynamic nonlinear behavior of the PA must be modeled before it can be corrected, hence the requirement for digital predistortion (DPD). Moreover, the DPD needs to be adaptive to the dynamics of the environment.

*Figure 2* depicts the core elements for many DPD systems: observation, estimation, and actuation. The concept in *Fig-*

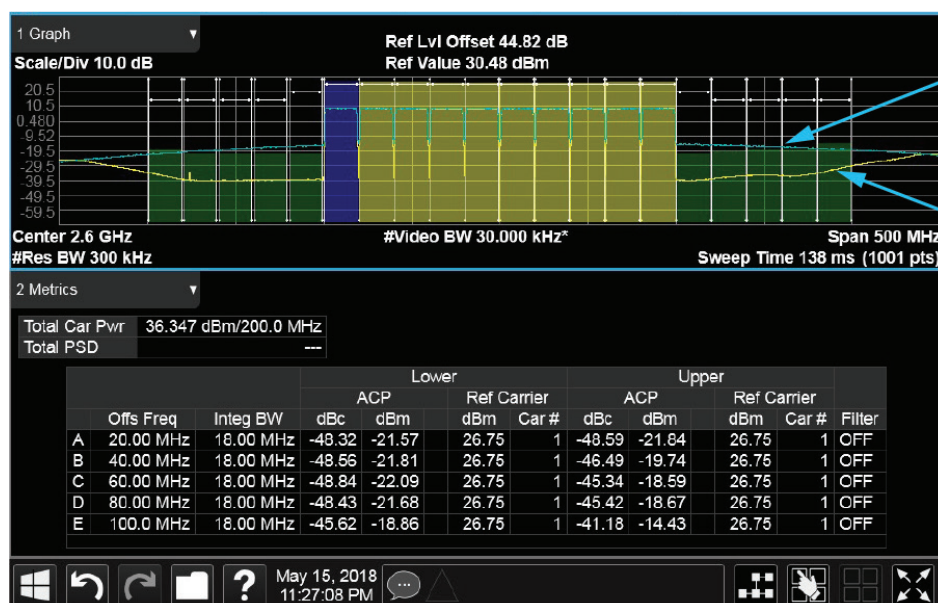
*ure 2* generates a model that tracks the expected response of the PA so that an appropriate cancellation signal can be generated to nullify the predicted nonlinear behavior of the PA. There are many models, such as the ubiquitous generalized memory polynomial (GMP).

A PA operating in its linear region generates less out-of-band distortions and as shown in *Figure 3*, has a notable reduction in the level of noise that leaks into the adjacent channels. The screenshot is from a spectrum analyzer on a typical DPD test bench, which is used to demonstrate static DPD performance that meets the standards required by many ACLR compliance tests.

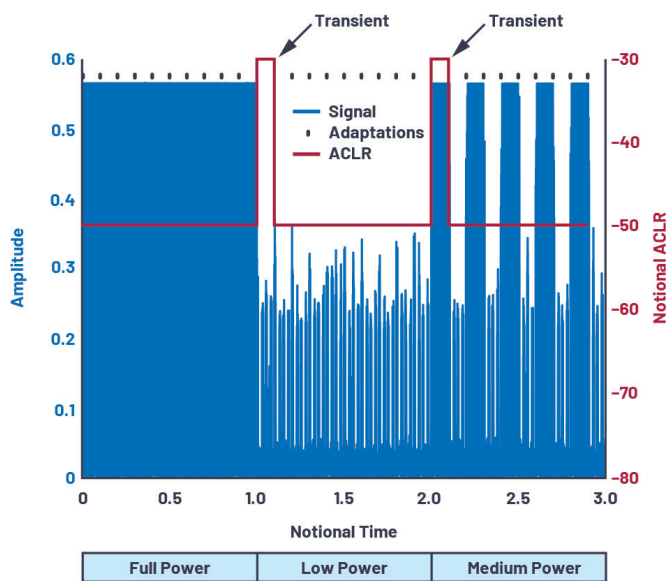
### Market Evolution, Performance Enhancement, and a Moving Target

DPD has been utilized commercially in some 8 million cellular base stations since the 1990s. As the technology and generational requirements of the cellular market have changed (2G, 3G, 4G, and now 5G), so too have the requirements placed on DPD. Those challenges include, but are not limited to, wider bandwidths, higher powers, carrier placements, higher peak-to-average signal ratios, and densification in the number and proximity of base stations.

Equipment vendors are anxious to differentiate their product offerings and continue to push for performance enhancement in terms of efficiency relative to the relevant 3GPP specification. PA efficiency continues to present a challenge. Whereas traditional drivers of change would have been OPEX costs and thermal management (including the hardware/weight costs associated with it), environmental considerations are now accelerating that change.



3. These plots illustrate adjacent channel leakage before and after application of digital predistortion.



4. Shown in this plot is an example of abrupt signal changes and how the ACLR might respond to a continuously adapting DPD.

PAs and DPD share a partially symbiotic relationship. In some instances, that relationship can be harmonious and in others more difficult. A PA that is DPD-friendly with DPD from one supplier may struggle with that from another. Often, optimal performance is achieved when both DPD and the PA are configured and tuned to match the specific application. However, because PA design continuously evolves to meet the aggressive requirements of 5G and beyond, DPD must evolve accordingly.

As wideband and dual-band applications become the norm, PA developers are challenged on how to achieve wider bandwidths at higher frequencies while maintaining performance expectations. Developing a PA with a bandwidth capability of 200 MHz and beyond is a challenge. Ensuring that it also can meet 3GPP specifications and efficiency creates further challenges. These challenges, in turn, fall back on the DPD developers.

### Understanding the Challenge

Quantifying DPD performance isn't a straightforward task. A matrix of conditions and scenarios should be considered—in addition to the PA, there's also a slew of other mitigating dependencies. When we consider performance, the specifics of the test conditions must be clearly defined: Achieving >50% ef-

iciency at a bandwidth of 200 MHz is a much greater challenge than the same level of efficiency at an operating bandwidth of 20 MHz.

The situation becomes more complex when we consider carrier placement within the allocated spectrum. It may be a contiguous signal, but it also may be a segmented carrier allocation in which portions of the spectrum are occupied.

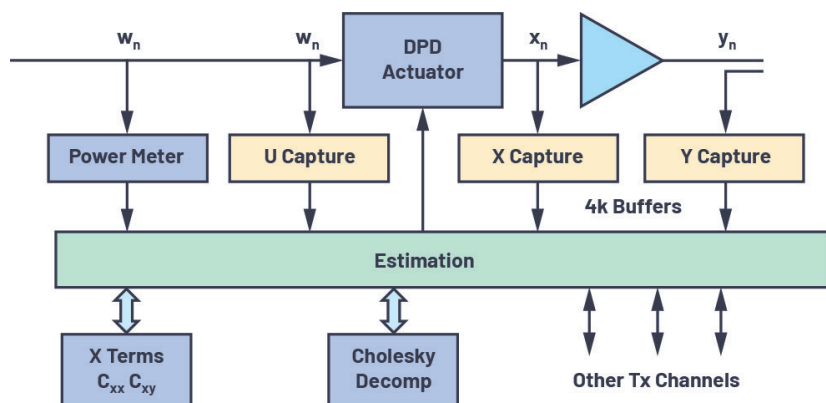
At a high level, there are quantitative indicators of DPD performance—the data points primarily defined by the 3GPP specification or operator requirements: ACLR, EVM, and efficiency. Meeting these are just the tip of the DPD performance iceberg. If we add stability and robustness to the mix, the enormity of the challenge starts to surface. Two critical aspects define DPD performance: the static bench-level performance and the real-world operational dynamic performance.

To characterize the challenge of dynamics, *Figure 4* illustrates signal evolution in a dynamic environment and shows how the ACLR might respond to a continuously adapting DPD. The numbers are notional. The plot provides an example of the effect of abrupt signal changes, which are extreme but legitimate. As the signal changes, the DPD model adapts to it. Adaptation events are indicated as dots.


In the transition time between a signal change and the next adaptation, a mismatch occurs between the model and the signal. Therefore, the ACLR value can rise, increasing the risk of exceeding the emissions specification for the duration of the transient.

Adaptation takes a finite time, so there will always be a transient. The challenge for high-performance DPD is to reduce that model mismatch time to a minimum while also ensuring a smooth transition between both states. The process must be managed so that speed of adaptation and disruption to ACLR are both considered.

It's important to understand how the model mismatch depends on the nature of the signal transitions. When the mis-



5. This DPD implementation features more extensive data capturing/observation.



6. Long-term gain errors were introduced by GaN PA charge trapping.

match is high, DPD risks degrading performance or, even worse, the stability of the radio. Instability, should it occur, can see the DPD algorithm spiral out of control, blasting emissions masks and, in worst-case scenarios, damaging the radio hardware. On the seesaw of performance vs. stability, stability will always be the prominent design consideration. A DPD design must be robust to ensure stability and error recovery under normal and abnormal operating conditions.

The challenge for a high-performance, practical DPD solution can be summarized in these requirements:

- Static performance (compliance testing or where the BTS traffic load is approximately constant): ACLR and error-vector magnitude (including GaN as a special case)
- Dynamics
- Robustness

In addition, because Analog Devices is a third-party vendor of DPD, the following also must be considered:

- *Maintenance*: The resolution of performance issues that occur when our customer (the OEM) deploys to its cus-

tomers (the operator).

- *Evolution*: During its lifetime in the field, the PA technology and signal-space application can change.
- *Generality*: An OEM can fine-tune its DPD to each product. We don't have that luxury. We must meet the needs of many applications while minimizing configurability and redundancy.

### Progressing DPD Performance to Meet the Challenges

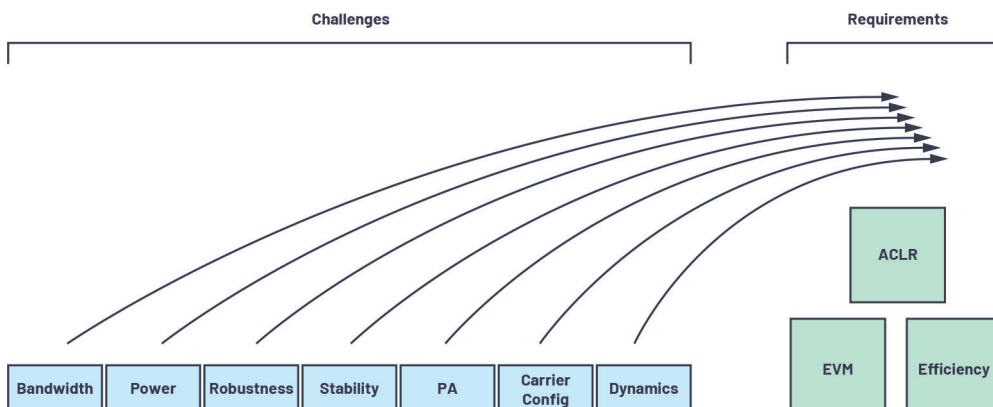
If we consider static performance alone, there's an element of linear progression to DPD development. Namely, if we provide more resources, then we enhance performance. For example, more GMP coefficients help to model the PA behaviors more accurately. Thus, as bandwidths widen, this becomes one element of a strategy to maintain, if not improve, performance.

However, such an approach has its limitations. A point of diminishing returns will be reached at which additional resources provide little or no benefit. DPD algorithm developers need to take more creative approaches to eke out further enhancements. ADI's approach is to augment the base algorithm GMP with more general basis functions and higher-order Volterra products.

As developers attempt to create a model that will accurately predict the PA behavior, data accumulation and data manipulation are core essential elements. Capturing data at successive time and power levels gives developers a more complete reservoir or armory from which to make their assessments and shape model behavior.

Figure 5 provides a conceptual overview of a system adopting such an approach. Note the more extensive data-capturing/observation nodes coupled with the digital power monitoring. Power monitoring helps with dynamics. Previously stored models can be brought into play in various ways to mitigate the dynamic transients discussed above.

In recent years, GaN PA technology has brought about an additional challenge for DPD developers: long-term memory effects. GaN process technology brings with it many



7. A holistic approach to DPD assessment balances all of the elements of DPD performance with the challenges.

distinct advantages in terms of efficiency, bandwidth, and operating frequency. It does exhibit what's known as the charge-trapping effect, though.

Charge trapping in GaN is a long-term memory effect, where there's a trap and then a thermal de-trap. GMP-based DPD corrects some of the error. However, residual error continues to impact signal quality. This distortion induces a corresponding rise in EVM. *Figure 6* provides a graphical representation of the phenomenon. Note the PA gain fluctuations and the temporal nature of those fluctuations. Also note the trap and de-trap states and that de-trapping occurs on the lower power symbols.

As the temporal effect is long-term, traditional approaches would suggest the acquisition of a very large number of sample points and, hence, a large amount of data to be stored and processed. Memory costs, silicon area, and processing costs mean that this approach isn't a feasible option for commercial DPD deployments. DPD developers must negate the effects of charge trapping but do so in a way that lends itself to efficient implementation and operation.

Charge-trap correction (CTC) is a feature supported at low cost in terms of power and compute time in ADI's ADRV9029 transceiver. It's been shown to recover the EVM to a level that's within the EVM 3GPP specifications. A next-generation transceiver, the forthcoming ADRV9040, boasts a more elaborate solution that's planned to deliver enhanced performance in dynamic scenarios and better coverage against what are an increasing number of GaN PAs with unique charge-trap personalities.

As stated, the stability of a DPD implementation is of utmost importance. Robustness is addressed by continuously monitoring the internal state and providing rapid responses to unusual conditions.

The generality of ADI's solutions is addressed by testing on a wide sample of PAs from many vendors—a large percentage of whom a symbiotic technical relationship is established.

## Conclusion

All too often when DPD performance is being presented, the focus is on the static elements of performance. While the yardstick of measurement in terms of EVM and ACLR remain valid, more attention must be paid to the matrix of operating conditions and requirements that frame those measurements. The demands of 5G NR continue to push application requirements. This, coupled with the desire for higher PA efficiencies, compounds the challenge of DPD algorithm development.

As we start to qualify DPD performance (*Fig. 7*), we need a holistic approach that handles:

- Static performance
- Dynamic performance

- Robustness
- Stability

DPD that has narrow margin to the specification may not be welcomed, while DPD that causes temporary specification extrusions may unsettle operators. DPD that goes unstable and results in illegal emissions and possible PA failure is disastrous.

A DPD algorithm should not be considered an off-the-shelf item. Optimal performance is achieved when the DPD is pruned to the specifics of the PA and the application. Hence, algorithm agility and development/field support also are important considerations. An effective DPD algorithm can deliver substantial system benefits. The complexity of the requirements and the performance assessment should not be underestimated.

*Steve Summerfield is a director at Analog Devices and leads algorithm and architecture development for wireless infrastructure applications. He joined ADI in 2017, having previously held senior positions in various telecommunications and semiconductor companies. Steve also has a long prior career in the university sector. He is widely published, holds many patents, and has a Ph.D. in theoretical physics. He can be reached at [steve.summerfield@analog.com](mailto:steve.summerfield@analog.com).*

*Frank Kearney joined Analog Devices after graduation in 1988. During his time with the company, he has held diverse engineering and management roles. He currently manages a team of architects and algorithm developers in the Wireless Systems Group. The group's focus is on transmit path efficiency and system-level enhancements for O-RAN radio architectures. Frank holds a doctorate from University College Dublin. He can be reached at [frank. Kearney@analog.com](mailto:frank. Kearney@analog.com).*