Solving the Interference Curse

By Barry Manz

Interference is the bane of every system, from EW to radar, communications and GPS, and it is one of the most difficult problems to defeat because it can result from an enormous number of sources. Interference comes not just from adversaries (in the form of jamming) but also from friendly forces, from equipment interfering with itself (i.e., self-interference), and from passive intermodulation distortion (PIM) that is often not only exceedingly difficult to locate but appears seemingly out of nowhere. Facing challenges like this, it is no wonder that developers of military technology are addressing interference mitigation more aggressively than ever.

AN EARLY PROBLEM

From an historical perspective, interference has been around since 1844, when Samuel F. B. Morse tapped "What hath God wrought" and sent it over wires from Washington to Baltimore. Once the table radio became widely available, beginning in the 1920s, it was a must-have household appliance, and listeners were happy enough just to receive broadcasts despite a level of interference that would intolerable today.

Early radio listeners had to contend with atmospheric noise, as well as electrical interference created by household electrical appliances with no shielding, along with the buzz from power lines. In addition, AM broadcast transmitters, which generated increasing amounts of RF power, created very strong harmonics and spurious emissions that became worse after sundown, when stations could be heard for thousands of miles.

The invention of the superheterodyne receiver by French officer Lucien Lévy and American officer Edwin H. Armstrong (who got there first is still disputed) allowed radios to receive signals from much further away, paving the path for more advanced commercial and military RF systems. Unfortunately, with longer reception came even more interference.

It became such an important issue that in 1933 the International Special Committee on Radio Interference (CISPR) was created within the International Electrotechnical Commission to harness it. In 1934, CISPR began to release the first of what would ultimately become a massive number of requirements for allowable emissions levels and immunity limits for electronic devices that today govern every type of electrical product and covers frequencies from DC through 400 GHz.

In 1967, the Department of Defense (DOD) created Mil-STD-461A, which established testing and verification requirements for electronic devices being used in defense systems, as well as emissions and susceptibility limits for new military electronic equipment. The FCC in 1979 added legal limits on EMI for digital equipment. Avionics and aerospace engineers investigated EMI-related issues to better understand interference sources, first using Faraday cages and eventually implementing better electronic designs and layouts, as well as shielding and filtering. These regulations have continued to evolve over the years as systems have become faster, smaller and more powerful with a greater propensity to interfere with each other.

Nevertheless, even with all these requirements, the world is still plagued by interference – and it is getting worse as the densely populated spectrum below 6 GHz becomes saturated with more signals from a variety of military and commercial radars, telecommunications devices and towers, as well as other systems and devices. As these technologies expand their use to higher frequencies, greater interference will come with them.

THE BASIC PROBLEM

Interference is most widely discussed in terms of its effect on commercial wireless systems whose problems are inherently well defined. In contrast, those of defense systems are as unpredictable as the adversaries they face, and they are deluged by interference from licensed, unlicensed, unintentional, self-induced, and sometimes outright bizarre sources, such as PIM. The transition from analog to digital modulation techniques, as well as direct RF sampling receivers, have created new hurdles to overcome, too.

At a high level, the problem appears simple: keeping interfering signals from degrading receiver performance. But achieving this is extremely difficult. For example, transmitters cause intermodulation distortion, harmonics and spurious emissions that can overwhelm the receiver's front end. When this occurs, its signal-noise ratio is degraded and it can, in the worst cases, make it impossible for the receiver to detect the desired signals.

Well-known techniques for reducing these emissions include operating amplifiers below their saturation point (back-off), as well as analog and digital pre-distortion and reducing phase noise. Increasing the linearity of the receiver also aids in preventing intermodulation products from interfering with the desired signals.

In the receive chain, the most effective solution for reducing interference has long been the venerable analog bandpass filter, whose ability to reject out-of-band signals is exceptional. The bandpass filter **(Figure 1)** passes frequencies within a narrow range while rejecting (attenuating) frequencies outside that range. The level of rejection can be very high – 50 dB or more – effectively reducing the signal strength of interferers to manageable levels. But a bandpass filter rejects signals only over a specific range of frequencies. Switched filter banks are another solution often used to cover multiple frequencies, but they can become large. Digitally tunable filters, which can provide greater protection over wider bandwidths, offer yet another approach.

Even though they are sometimes considered archaic, filter technology offers an effective solution to interference. Filters, addition to interference cancel-

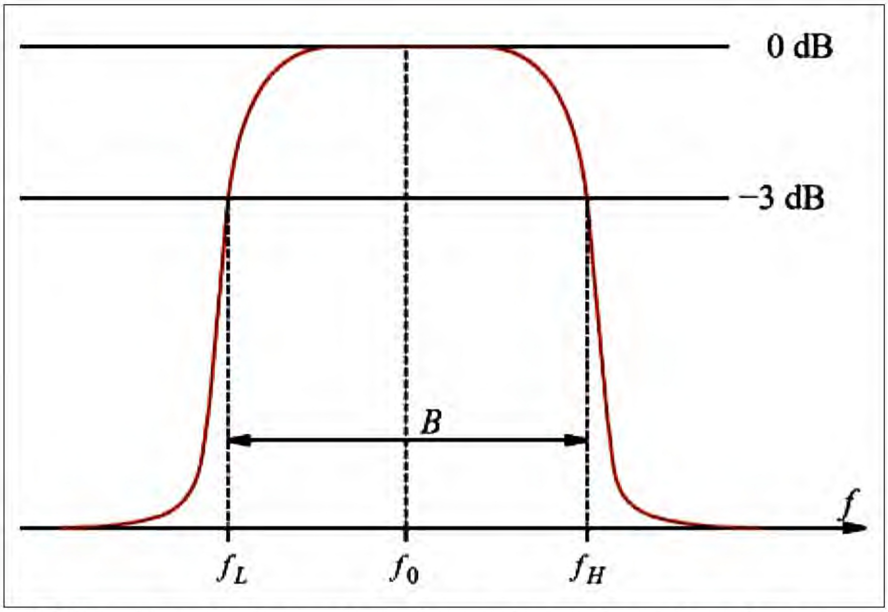

Fig. 1: The bandpass filter has served every RF application well for decades, as it reduces interference to a significant degree outside the frequency of interest. WIKIPEDIA

lation, are the primary mechanisms by which the most advanced interference mitigation systems are constructed.

Remedies such as these can obviously be applied only to systems over which the owner has control, and obviously are not applicable to those of adversaries or even operators of licensed and unlicensed systems. As a result, they represent a solution to just one set of interference problems, but this is more than enough to keep designers up at night.

THE COMPLEX PROBLEM

There is also the onerous problem of interference caused by PIM that is different from other forms of intermodulation distortion. PIM is not generated by active (non-linear) components, such as amplifiers, but rather by components normally considered linear, such as antennas, attenuators, cables, duplexers, diplexers, filters, and connectors, in which oxidation or other effects may cause them to become non-linear.

PIM appears when two or more signals are present in a passive non-linear device and they mix to produce other signals related to them. Unlike mixers, in which such a result is desired, in the case of PIM the effect is undesirable. If these unintentionally mixed signals are significantly weaker than the desired signal, they may have no effect on a receiver. But if they occur after the combining stage, for example, they can reach levels high enough to form distortion products at amplitudes high enough to degrade receiver performance. While active components like amplifiers can produce extremely high levels of distortion, this type of interference can be removed by filtering. PIM typically cannot be removed in this way, because it can be generated late in the transmit signal path or outside the system entirely.

PIM has typically been considered a problem for systems generating high RF power levels, a misnomer that has remained in its definition primarily because high-power systems, such as the radars and satellite communications systems on surface combatants and AM, FM, and TV broadcast installations, were

Fig. 2: The rusty bolt effect has been a major source of problems on navy ships for decades. WIKIPEDIA

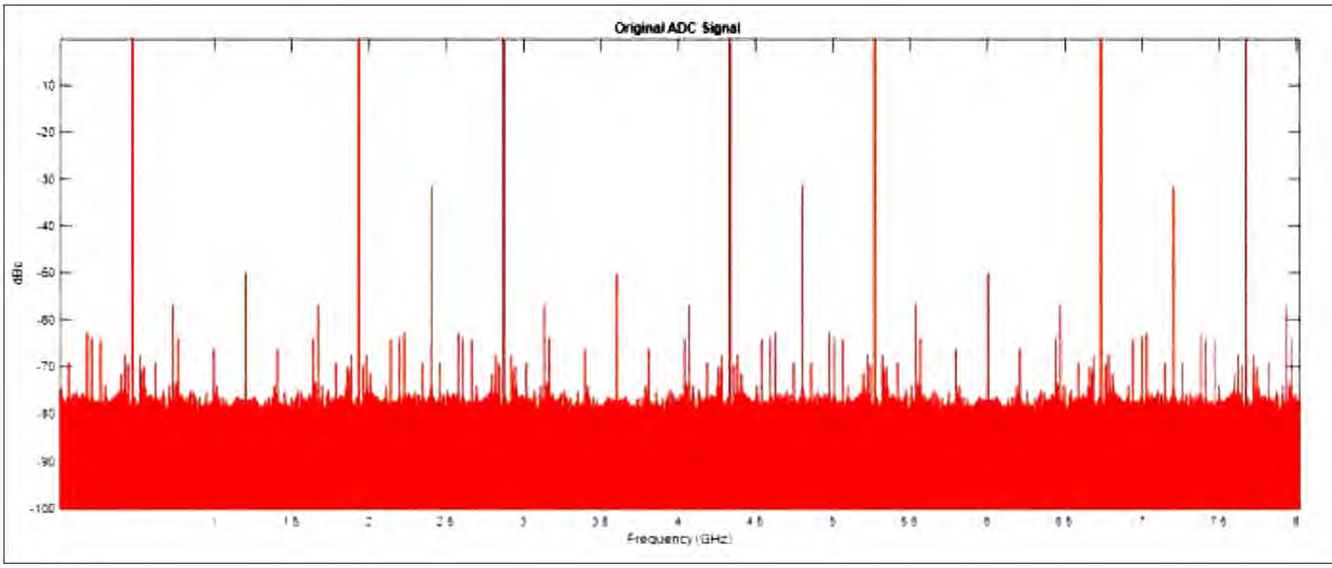

Fig. 3a

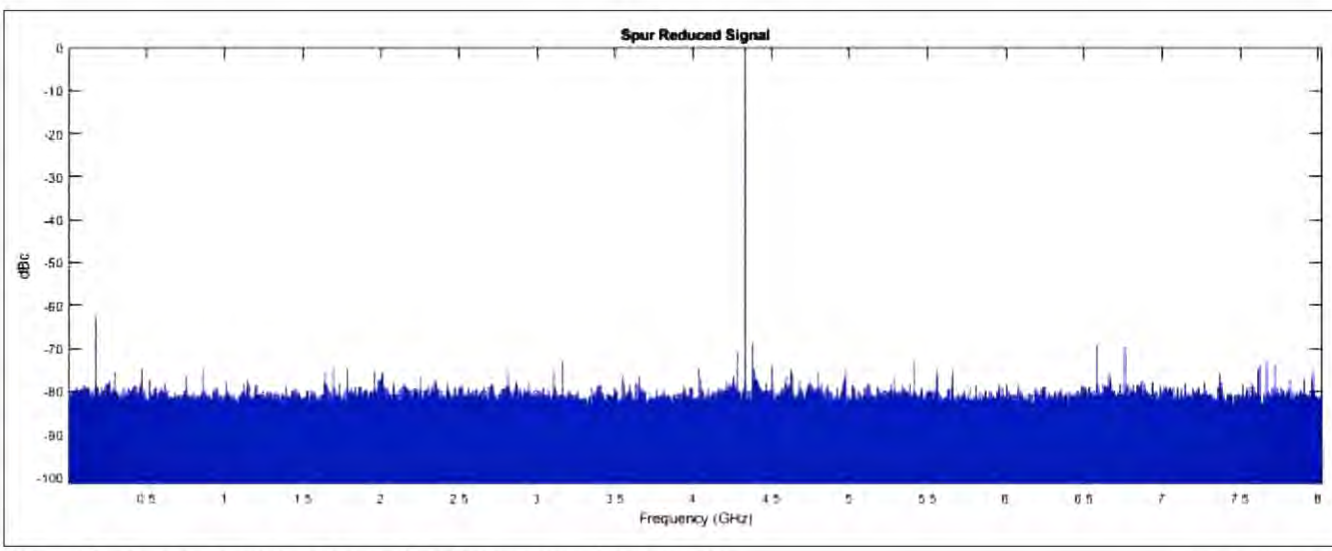

Fig. 3b: *Spurious signals without the use of HDRR (a) and with it employed (b).*

PRECISION RECEIVERS, INC.

the first to experience it. In practice, studies have shown that components producing PIM can cause damaging effects several miles away, even when their signal strength is low.

The problem on surface combatants has been so severe for so long it received its own moniker – the rusty bolt effect **(see Figure 2)**, in reference to the corrosion and rust present almost anywhere on a ship where interactions of electromagnetic energy with dirty connections or corroded parts cause the same effect as a diode. As more and more discrete transmit and receive systems have been added to surface combatants over the years, in combination with the ship's structure they are a breeding ground PIM generation.

The wide variety of ways PIM can be generated makes it extremely difficult to detect. However, over the last 15 years or so, the availability of portable PIM analyzers have made identifying this problem much easier. Resolution of the problem typically requires inspection and repair and replacement of the offending component. When PIM is generated by a ship's infrastructure, however, it can require extensive measures, such as repainting large surfaces of the superstructure.

TACKLING INTERFERERS HEAD ON

Generally speaking, there are two primary means of achieving interference mitigation: filtering and interference cancellation. These approaches can be used alone and together, and the solutions described later in this article use either one or both of them. However, as noted earlier, the design of receivers is changing rapidly from their traditional architecture, which relies on analog components that include a mixer and local oscillator in the front end, to direct RF sampling that eliminates one of both of them in favor of digitizing the incoming signal directly when it is captured over the air. While this approach has enormous benefits, it also introduces a problem related to the use of very-high-performance analog-to-digital converters.

Direct RF sampling is becoming the architecture of choice for applications ranging from electronic warfare (EW) and signals intelligence (SIGINT) systems to radar and communications systems. As this approach eliminates nearly all RF components in the signal chain before the input signal is digitized, a system can be smaller and less complex, because once in the digital domain, the functions traditionally performed by several analog components can instead be performed digitally in a digital signal processor or an FPGA.

In these receivers, high performance can only be obtained when its signal-to-noise ratio (SNR) and spurious-free dynamic range (SFDR) are extremely high. The component most important for achieving this is the analog-to-digital converter (ADC) because, as the first signal processing component after signal capture, it defines the performance that the entire receiver can achieve.

One of the most common and useful measurements of ADC performance is its effective number of bits (ENOB). Ideally, in a 12-bit ADC for example, the ENOB figure is close to 12. But there are a number of factors that can diminish the ENOB in an ADC. For instance, an ADC encounters errors in quantization, offset, gain, linearity and timing that create spurious signals in its output. Even the most impressive examples of today's ADC boards pose significant design challenges, especially as sampling rates and instantaneous bandwidth increase and when time-interleaving of ADCs is employed.

If the strength of these spurious signals is high enough, it becomes difficult and sometimes impossible to separate the signals of interest from the noise. Techniques, such as clock dithering, calibration and commutating the ADC at lower rates, have been used to mitigate the issues in the ADC, but each one of these techniques has significant drawbacks that can often cause as many problems as they attempt to solve and require considerable computing power, as well.

One of the most recent approaches to solving this problem is the High Dynamic Range Receiver (HDRR) solution developed by Precision Receivers Inc. (Marshall, VA). The HDRR does not have the shortcomings of other methods, according to the company, and it is most effective when acting on signals with high dynamic range and wide bandwidth. It can achieve an order-of-magnitude improvement in reducing unwanted spurious signals, improving spurious-free dynamic range (SFDR) by up to 15 dB and optimizes the ADC's ENOB as well.

HDRR can be used in any direct-sampling system regardless of its ADC and at any frequency of interest. It also does not require self-calibration, reduces anti-aliasing filter complexity, and minimizes the required amount of post-processing and signal analysis. It is based on the fact that an RF input signal from an antenna consists of desired signals and noise, and that digitization in the ADC introduces another noise-distortion component. HDRR modifies this additional noise component in part by manipulating the ADC's control signals using an approach developed by the company over several years. The process effectively removes the noise contributed by the ADC, producing an output signal with much less distortion, which is then passed to an FPGA for processing or to a mass storage device, depending on the application (see **Figure 3**).

YET ANOTHER PROBLEM TO SOLVE

Defense and commercial wireless systems both face the challenge of increasing spectral efficiency, as there is precious little of it to waste within the communication sweet spot between about 600 MHz to 7 GHz. The "ideal" of achieving greater spectral efficiency is through full duplex rather than half-duplex communication. By enabling transmission and reception to be performed in the same channel at the same time – known as same-frequency simultaneous transmit and receive or SF-STAR – a system can effectively double its spectral efficiency.

However, achieving full-duplex operation is extremely challenging and, owing to the emergence of 5G and the requirements of next-generation military communications systems, it has become an extremely important problem to solve. The result has been a flurry of journal articles extolling the virtues of various approaches taken by the wireless industry and academia.

The most difficult challenge is the self-interference from the transmitted signal power to the simultaneously received signal. Unfortunately, this creates an overlapping of the strong signal with the much weaker received signal of interest, producing considerable self-interference. This signal can theoretically be removed by subtracting the transmitted signal from the received waveform, but the signal will be linearly and nonlinearly distorted while propagating to the receiver. The problem is the result of RF power amplifier non-linearities, transmitter and receiver in-phase/quadrature (I/Q) imbalance, the phase noise of the local oscillator, and ADC quantization noise. To be effective, an approach must ensure that the self-interference power level is below the noise floor, and if not, then it must be self-canceled by at least 70 dB.

If the approach works, the system can cancel its own transmitted signal in its receiver, and what it transmits does not impact what it simultaneously receives, which requires that changes in the time-varying self-interference channel must be tracked in real-time, so the system must be self-adaptive. Some of the most impressive STAR work has been conducted at Stanford University, whose researchers have achieved self-interference cancelation greater than 100 dB, and its architecture is widely used as the foundation of other approaches.

Full-duplex capability will have a profound effect on spectral efficiency, and it could not come at a better time for both the wireless industry and DOD, which are inextricably intertwined. The DOD "owns" an enormous amount of spectrum, some of which it has reluctantly offered to the FCC for auction, which will help the wireless industry in the US still where there is less commercially unused "mid-band" spectrum available than almost every other first-world country. Every bit of it will be required for 5G to be realized.

OTHER APPROACHES

These efforts toward interference cancellation are far from the only ones either in progress or in service. One recent initiative is the Defense Research Projects Agency's (DARPA's) Wideband Adaptive RF Protection (WARP) program. The objective of WARP is to "harden" wideband receivers operating in congested and contested EM environments using adaptive filters and signal cancellers that selectively attenuate or cancel signals. In addition to countering intentional jamming, WARP addresses interference caused by a transceiver's transmitter, sniffing the EM environment, and using wideband tunable filters to maintain the receiver's dynamic range without decreasing its sensitivity. That is, a WARP solution would reduce the effect of large signals on the desired signal without reducing the receiver's performance.

In addition to filtering, the program also addresses adaptive signal cancellation technology. The traditional means of doing this is by using a duplexer to separate signals on different bands. But this approach is not effective when radios transmit and receive on the same frequency (i.e. SF-STAR) in order to increase spectral efficiency. Unfortunately, some SF-STAR systems can create the risk of self-interference caused by the transmitted signal unintentionally disrupting the receiver input.

Conventional techniques for self-interference mitigation in these situations requires minimizing the coupling between the transmitter and receiver or employing a controllable auxiliary path between the two ports to cancel the signal coupling. But these techniques sometimes fail to achieve the required amount of cancellation and are not reconfigurable or scalable across a wide range of frequencies.

Researchers have already proposed using a full-duplex antenna that achieves self-interference cancellation through polarization duplexing, and the WARP program is attempting to improve on this by sampling and canceling it in the digital domain. This technique is already being used by several companies, such as L3 Harris, which has been working on interference mitigation for many years. Its solution employs various tools depending on the type of threat, such as multiplexing and RF filters, as well as interference cancelation.

The latter samples interfering signal and creates the opposite of it in real-time and, when combined, they cancel each other out, leaving only the desired signal. It also cancels out noise and spurious signals even if they are on the same frequency. The result can be 100 dB or of interference cancellation. Its latest versions use the company's Advanced Interference Mitigation System (AIMS) that makes it possible to quickly enable solutions for a wide range of interference profiles.

Cobham's mINCAN interference cancellation system injects an anti-phase version of the interfering signal, similar to the approach used for acoustic noise cancellation in smartphones, headsets and hearing aids. However, this system can handle very high levels of received interference and accommodates the high speeds of frequency-agile transmitters.

A sample of the signal from the transmitter in a co-located situation is taken using a directional coupler in the transmit antenna path and is scaled in phase and amplitude before being mathematically being added to the receive path to cancel out the interference. Multiple interferers can be canceled by allocating a module to each one, and cancellation takes place before the interfering signal reaches its full amplitude to avoid receiver blocking.

Another technique developed at MagiQ called Agile Interference Mitigation System (AIMS – not to be confused with L3Harris's solution) identifies high-level signals that would otherwise lead to distortion in the receiver and sets up frequency-agile, very-high-Q filters to quickly suppress them. The AIMS device is located between the antenna and receiver and discriminates signals of interest from interferers and suppresses the latter signals before reaching the receiver without effect on the desired signal.

AIMS monitors the spectrum in real-time, identifies unwanted interference signals, invokes the high-Q filter technology to allow removal of about 60 dB of interference while leaving nearby signals of interest uncompromised. The tunable filters can be continually updated to adapt to rapidly changing EM environments.

Septentrio's AIM+ system is focused on satellite-based navigation systems and mitigates the effects of narrow-band interference using three notch filters that remove a narrow part of the RF spectrum around the interfering signal. The power spectrum plot AIM+ technology detects and neutralizes interference resulting in faster set-up, reduced downtime and secure operation. AIM+ protects against simple narrowband interference, as well as more complex wideband interference, including jamming and spoofing.

Kumu Network's self-Interference cancellation technology is also designed to allow radios to transmit and receive on the same or adjacent channels. Its canceller suppresses the interference a transmitter presents to a co-located receiver even if the two radios operate with no guard band between them. The solution adapts in real-time to the changing environment to ensure consistently high isolation between the two radios.

The technology is also available in chip form with frequency-agnostic taps for self-interference cancellation or FIR filtering. Digital cancellation taps can be used where analog cancellation alone does not provide sufficient cancellation or where longer delay reflections impact the receiver and need to be canceled. It has four FIR filter chains organized as 12 programmable taps per chain with a maximum of 350-nsec aggregate delay through each chain.

Two chains can be cascaded for a maximum of 24 taps for a total delay of up to 700 nsec, or the chains can be configured to support 2x2 MIMO operation. The IC supports a range of signal processing applications where analog signal manipulation is required to avoid the delay and resolution problems that digital conversion introduces. For example, the chip can be used to implement full-duplex systems or to improve co-existence between co-located radios.

CUTTING THROUGH THE NOISE

At some time in the very (very) distant future, all types of interference may somehow be resolved, but until then, DOD, industry and academia will continue to throw every possible resource into reducing it. This will not be easy, as there are so many types available to work on at every level. The interference cancelation techniques described in this article may be the most widely used or potentially promising, but there are others as well, and with so much work becoming conducted to eliminate it, new ones will continue to be continuously revealed.