ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™
AnalogDialogue

Complex RF Mixers, Zero-IF Architecture, and Advanced Algorithms: The Black Magic in Next-Generation SDR Transceivers

By **Frank Kearney** and **Dave Frizelle**

Introduction

There is an interesting interaction between complex mixers, zero-IF architecture, and advanced algorithm development. The objective of this article is to establish the basic fundamentals of each: the principles of operation and the value they deliver in terms of system design, and then to discuss the interdependability of the three.

RF engineering is often regarded as the black art of electronics. It can be a strange mix of mathematics, mechanics, and, in some instances, just trial and error. It unsettles many a good engineer and many others settle for understanding the outcome rather than the detail. Much of the existing literature jumps straight into the theoretical and mathematical explanation without establishing the underlying concepts.

Demystifying the Complex RF Mixer

Figure 1 provides an overview of the complex mixer in an upconverter (transmitter) configuration. Two parallel paths with independent mixers are fed from a common local oscillator whose phase is offset 90° to one of the mixers. The independent outputs are then summed in a summing amplifier to produce the desired RF output.

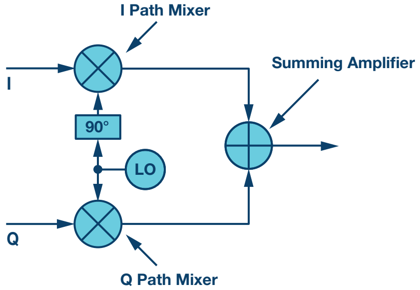

Figure 1. Basic architecture of a complex transmitter.

The configuration has a very useful application. Let's assume, as shown in Figure 2, that we feed a tone signal only on the I input, and the Q input is undriven. Given that the tone at the I input has a frequency of x MHz, the mixer in the I path produces an output at the LO frequency ±x. As there is no signal applied to the Q input, the mixer in its path produces an empty spectrum, and the output from the I mixer passes straight to the RF output.

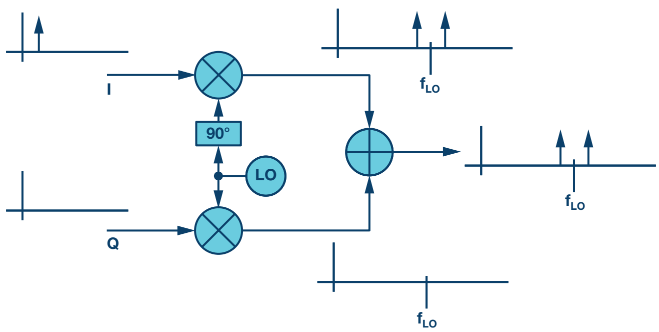

Figure 2. I path analysis.

Alternatively, let's assume that a signal tone at frequency x is applied solely to the Q input. The Q mixer in turn produces an output with tones at the LO frequency ±x. With nothing applied to the I input, its mixer output is muted and the output from the Q mixer goes straight to the RF output.

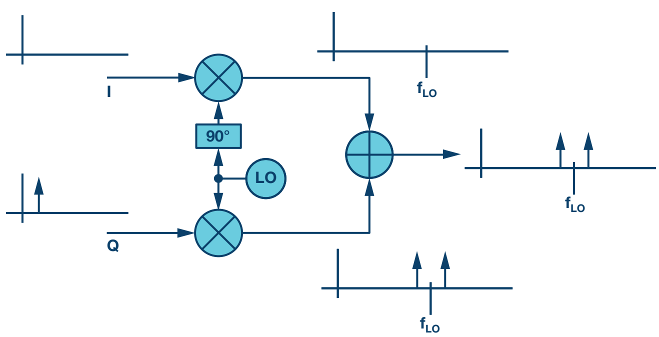

Figure 3. Q path analysis.

At first glance it may seem that the outputs from Figure 2 and Figure 3 are identical. However, there is one critical difference, namely phase. Let's assume, as shown in Figure 4, that we apply the same tone to both I and Q inputs, and that there is a 90° phase shift between the input channels.

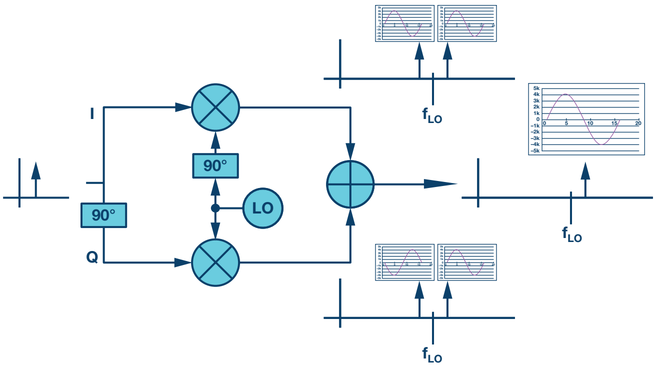

Figure 4. Simultaneous I and Q signal path analysis.

If we look closely at the output of the mixers, we observe that signals at the LO frequency plus the input frequency are in phase, whereas signals produced at the LO frequency minus the input frequency are out of phase. This results in the tones on the upper side of LO adding while the tones on the lower side cancel. Without any filtering we have removed one of the tones (or sidebands) and created an output that sits entirely on one side of the LO frequency.

The example shown in Figure 4 has the I signal leading the Q signal by 90°. If the configuration was to change such that the Q signal led the I signal by 90°, then we could expect a similar summing and cancellation, but in this instance all the signal would appear on the lower side of the LO.

Figure 5 shows the results of lab measurements of a complex transmitter. The left hand side shows the test case when I leads Q by 90°, resulting in the output tone placement on the upper side of the LO. The right hand side of Figure 5 shows the relationship swapped so that Q now leads I by 90° and the resultant output tone sits on the lower side of the LO.

In theory it should be possible to have all the energy on only one side of the LO. However, as the result from the lab experiments in Figure 5 show, in practice full cancellation may not occur, leaving some energy on the other side of LO, known as the image. Also note that energy at the LO frequency is present, known as LO leakage or LOL. Other energy is also evident in the results—these are harmonics of the wanted signal and are not discussed in this article.

For perfect image cancellation, the outputs of the I and Q mixers must be of precisely the same amplitude, and be exactly 180° out of phase with respect to each other on the image side of the LO. If the phase and amplitude requirements are not met, then the summing/cancellation process, as shown in Figure 4, becomes less than perfect and energy at the image frequency will remain.

Implications

The use of a conventional, single-mixer architecture produces LO± products. Before transmission, one of the sidebands will need to be removed, usually through the addition of a band-pass filter. The filter roll off must be such that it removes the unwanted image signal without affecting the wanted signal.

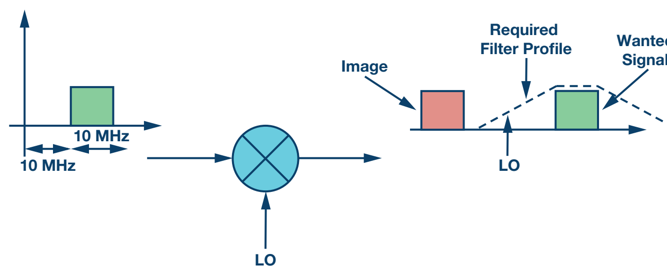

Figure 6. Single-mixer image filter requirements.

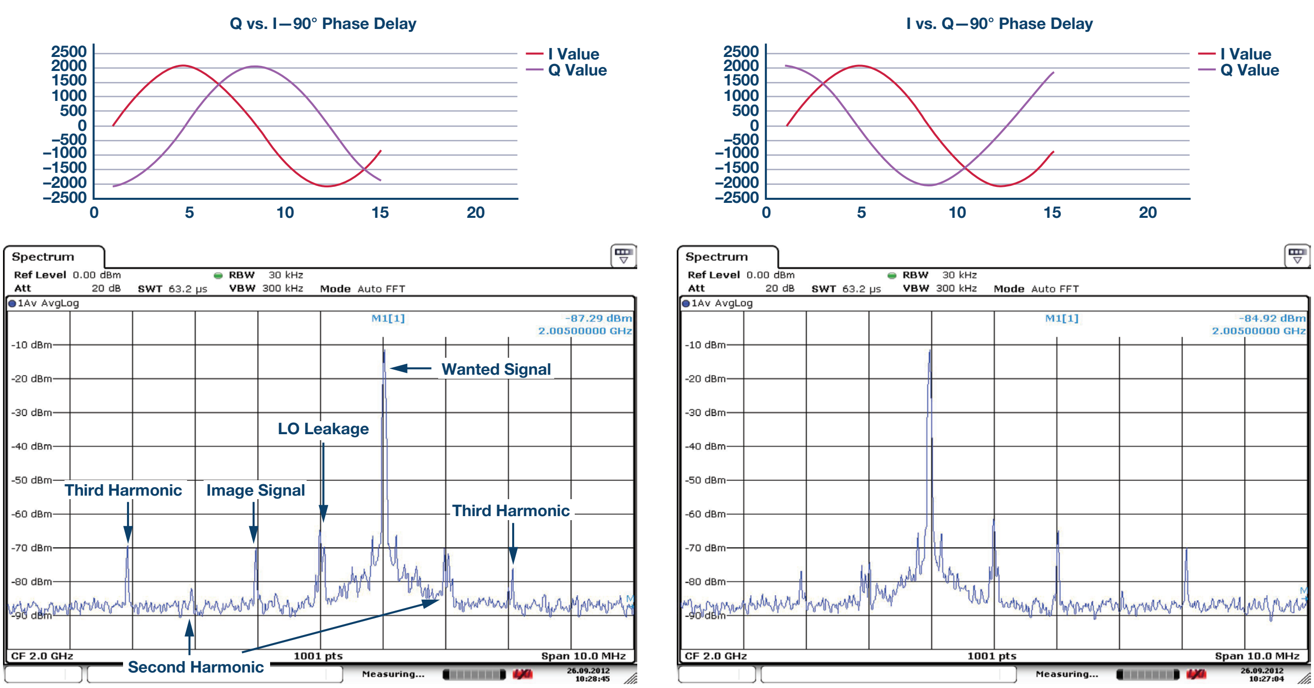

Figure 5. Tone placement dependent of the I and Q phase relationship.

The spacing between the image and the wanted signal directly affects the filter requirements. Where the spacing is large, a simple low cost filter with a gentle roll off can be used. If the spacing is narrow, then designs must implement a filter with a sharp response; typically employing multipole or SAW filters. Hence it would be correct to state that spacing must be maintained between the image and the wanted signal so that the image can be filtered without affecting the wanted signal, and that the spacing is inversely proportional to the complexity and cost of the filter. Furthermore, the filter must be tunable in frequency if the LO frequency is variable, which further increases the complexity of the filter.

The spacing between the image and the wanted signal will be determined by the signal that we apply to the mixer. The example in Figure 6 shows a 10 MHz bandwidth signal shifted 10 MHz away from dc. The resultant output from the mixer places the image 20 MHz from the wanted signal. In this configuration, to achieve a 10 MHz wanted signal spectrum at the output, we had to have a 20 MHz baseband signal path to the mixer. 10 MHz of the baseband bandwidth is unused, and the data interface rate to the mixer circuit is higher than necessary.

Returning to the complex mixer as shown in Figure 5, we know that its architecture eliminates the image without the need for external filtering. What's more, in a zero-IF architecture we can optimize efficiency so that the signal path processing bandwidth is equal to that of the wanted signal. Figure 7 shows a conceptual diagram of how this is achieved. As previously shown, if I leads Q by 90°, there will be an output on the upper side of LO only. If Q leads I by 90°, there will be an output on the lower side of LO only. Therefore, if two independent baseband signals are generated, where one is designed to produce an upper sideband output only and the other is designed to produce a lower sideband output only, they can be summed in baseband and applied to the complex transmitter. The result will be an output with different signals appearing above and below LO. In a practical application the combined baseband signal would be produced digitally. The summing nodes shown in Figure 7 are solely to illustrate the concept.

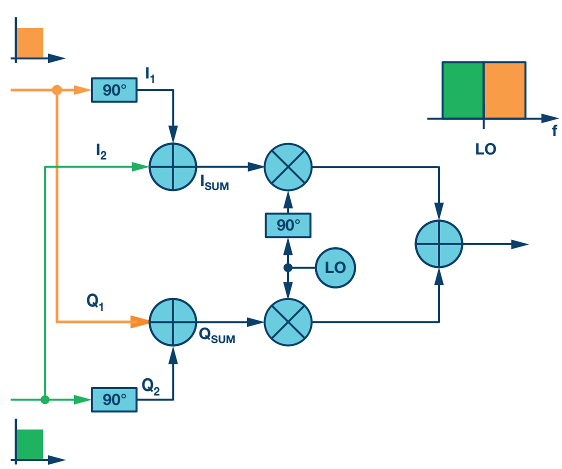

Figure 7. Zero-IF complex mixer architecture.

The Zero-IF Dividend

The use of the complex transmitter to generate a single sideband output provides substantial advantages in terms of the RF filtering required to remove the image. However, if the image cancellation performance is good enough to make the image negligible, we can exploit the architecture more by using it in zero-IF mode. Zero-IF allows us to take specially created baseband data and produce an RF output with independent signals appearing on either side of the LO. Figure 8 provides an illustration of how this might be done. We have two sets of I and Q data, where each is independent and encoded with symbol data that can be decoded at the receiver with respect to the phase of the reference carrier.

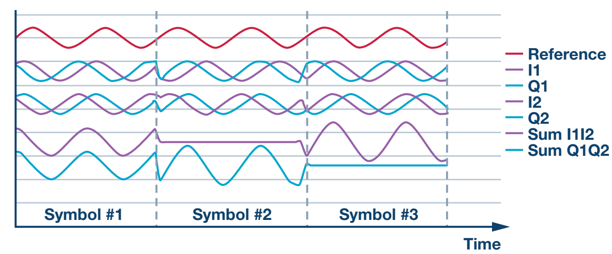

Figure 8. Taking a closer look at I/Q signaling in a zero-IF complex mixer configuration.

Initial observation shows that Q1 leads I1 by 90° and that the amplitude of both are matched. Likewise, I2 leads Q2 by 90° and their amplitudes are also matched. The independent signals are combined so that I1 + I2 = SumI1I2 and Q1 + Q2 = SumQ1Q2. The summed I and Q signals no longer exhibit phase and amplitude correlation—their amplitudes are not equal at all times and the phase relationship between them varies. The resultant output from the mixer places I1/Q1 data on one side of the carrier and I2/Q2 data on the other side of the carrier as previously explained and shown in Figure 7.

The use of zero-IF complements the advantages of the complex transmitter by positioning independent data blocks directly adjacent to each other on either side of LO. The data processing path bandwidth never exceeds that of the RF data bandwidth. So in theory, the use of a complex mixer used in a zero-IF architecture provides a solution that requires no RF filtering while also optimizing baseband power efficiency, delivering lower cost per unit of unusable signal bandwidth.

Up to this point, the focus of this article has been on the complex mixer used as a zero-IF transmitter. The same principles work in reverse and the complex mixer architecture can be used as a zero-IF receiver. The same advantages that have been described for the transmitter equally apply to the receiver. When using a single-mixer to receive a signal, the image frequency must first be filtered out using an RF filter. In the zero-IF mode of operation there is no image frequency to worry about, and signals above LO will be received independently of signals below LO.

A complex receiver is shown below. The input spectrum is applied to both I and Q mixers. One mixer is driven with LO, the other with LO + 90°. The outputs of the receiver are I and Q. In the case of a receiver, it is not as easy to prove empirically what the output will look like for a given input, but if a tone is input above LO, as shown, the I and Q outputs will be at the (tone – LO) frequency and there will be an expected phase shift between I and Q where I leads Q. Similarly, if the tone were input below LO, the I and Q outputs would again be at (LO – tone) frequency but this time Q will lead I. In this way the complex receiver can distinguish energy above LO from energy below LO.

The output of the complex receiver will be the sum of the I/Q information representing the spectrum that was received above LO and the I/Q information representing the spectrum that was received below LO. This concept was described earlier for the complex transmitter where a summed I and summed Q signal is applied to the complex transmitter. In the case of the complex receiver, the baseband processor receiving the summed I and summed Q information will easily be able to distinguish upper and lower frequencies using a complex FFT.

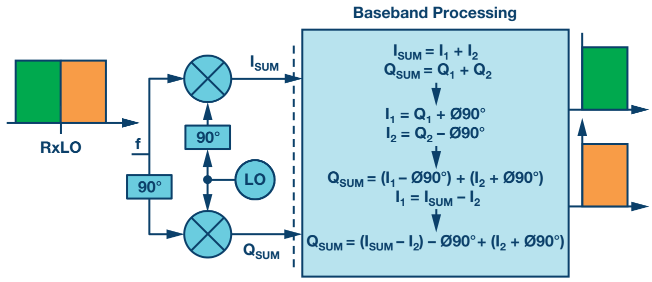

Figure 9. Zero-IF complex mixer receiver configuration.

When the summed I and summed Q signals are received, there are two knowns—the summed I signal and the summed Q signal—but there are four unknowns, namely I1, Q1, I2, and Q2. Because there are more unknowns than knowns, it would seem impossible to solve for I1, Q1, I2, and Q2. However, it is also known that I1 = Q1 + 90° and that I2 = Q2 – 90°, and with these two additional knowns it is now possible to solve for I1, Q1, I2, and Q2 using the received summed I and summed Q signals. In fact, we only need to solve for I1 and I2 because the Q signals are just copies of the I signals with a ±90° phase shift.

Limitations

In practice, the performance of the complex mixer has struggled to completely eliminate the image signal. This limitation could be considered as having two pronounced effects on radio architecture design.

Even with the performance limitation, complex IF does bring tangible benefits. Let us consider the low IF example in Figure 10. Accepting the performance limitation, we do still see an image. However, that image is heavily attenuated from that which we would expect to see from a single-mixer design (see Figure 6). Although the complex mixer continues to require a filter, the filter profile can be much more relaxed and its implementation simpler and lower cost.

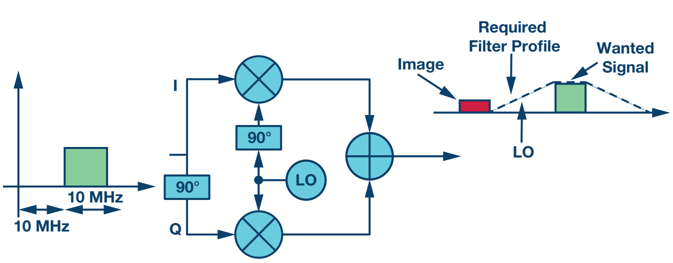

Figure 10. Practical implementation of the complex mixer. Note the attenuated image.

The filter complexity is inversely proportional to the distance between the image and the wanted signal. If we go to a zero-IF configuration, then that distance becomes zero and the image sits within the wanted signal band. The practical application of zero-IF theory has struggled, resulting in in-band image levels that degraded performance beyond an acceptable level (see Figure 11).

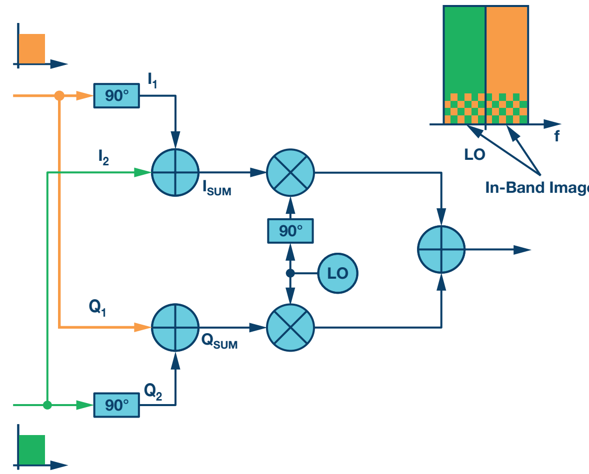

Figure 11. Zero-IF implementation restrictions.

The principles of the complex transmitter and receiver only hold true when the phase and amplitude requirements of the I and Q data paths are met. Mismatches in the signal paths will cause inaccurate cancellation of the image signals on both sides of the LO. Examples of such issues can be seen in Figure 10 and 11. In instances where zero-IF is not being used, filtering could be used to remove the image. However, if a zero-IF architecture is to be used, then the unwanted image falls directly within the spectrum of the wanted signal and a failure condition will occur if the image power is large enough. Therefore, the use of zero-IF and complex mixing can deliver an optimal system design solution but only when the design can eliminate the phase of amplitude mismatches along the signal paths.

Advanced Algorithm Enablement

The concept of the complex mixer architecture has existed for many years but the challenges of meeting the phase and amplitude requirements in a dynamic radio environment have restricted its use in a zero-IF mode. Analog Devices has overcome the challenge by a combination of smart silicon design and advanced algorithms. The design accepts that there will be signal path impairments; however, these are minimized by smart silicon design. The remaining imperfections are calibrated out by self-optimizing quadrature error correction (QEC) algorithms. Figure 12 provides a conceptual overview.

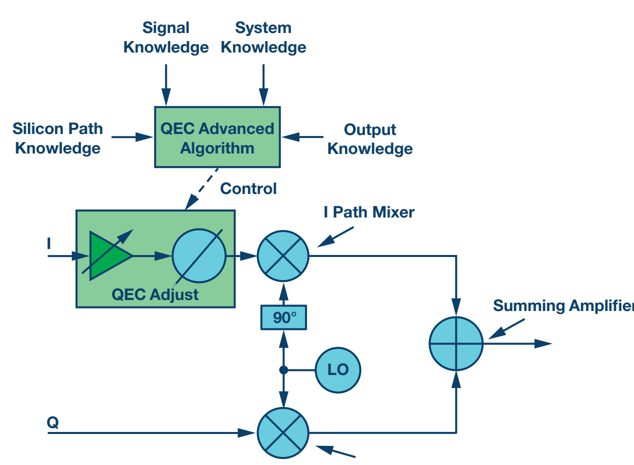

Figure 12. Advanced QEC algorithm and smart silicon design enabling zero-IF architecture.

On ADI transceiver devices such as the AD9371, the QEC algorithm sits within the on-chip ARM® processor. It has constant knowledge of the silicon signal path, the modulated RF output, and the input signal. It uses this knowledge to intelligently adapt the signal path profile in a controlled predictive fashion rather than a kneejerk reactive one. The algorithm performance is such that it can be best described as digitally assisting the performance of the analog signal path.

The dynamic QEC calibration algorithm is just one example, albeit a prominent one, of the advanced algorithms that reside and operate inside ADI transceivers. Others such as LO leakage cancellation coexist and lift the zero-IF architecture to an optimal level of performance. While these first generation of transceiver algorithms were primarily required for technology enablement, the second generation, such as digital predistortion (DPD), enhance the performance not just of the transceiver, but of the entire system.

All systems have imperfections that limit their performance. Whereas the first generation of algorithms have primarily focused on calibrating out on-chip limitations, the next generation uses the intelligence of algorithms to compensate for system performance and efficiency limitations external to the transceiver. Examples include PA distortion and efficiency (DPD and CFR), duplexer performance (TxNc), and passive intermodulation issues (PIM).

Conclusion

Complex mixers have existed for many years, but the image rejection performance that they provided did not allow them to be used in a zero-IF configuration. The combination of smart silicon design and advanced algorithms remove the performance barriers that had previously impeded the adaption of zero-IF architectures in high performance systems. With the performance limitations removed, the use of zero-IF architecture delivers saving in terms of filtering, power, system complexity, size, heat, and weight (the topic is extensively covered in an earlier article from Brad Brannon[1]).

In the case of complex mixers and zero-IF, we can consider the QEC and LOL algorithms as an enablement function. However, as the scope of the algorithmic development extends, it provides system designers with increased performance levels that allow them more flexibility in their radio designs. They may choose enhanced performance but they may also use the gains achieved from the algorithm to compensate for lower cost or size components in their radio designs.

References

[1] Brad Bannon. "Where Zero-IF Wins: 50% Smaller PCB Footprint at ⅓ the Cost." *Analog Dialogue*, Sept 2016.

Dave Frizzelle

Dave Frizelle [david.frizelle@analog.com] works as an applications manager in the Transceiver Product Group at Analog Devices Limerick supporting the integrated transceiver family of products. He has worked at ADI since graduation in 1998. His previous engineering roles include six years working in Japan and Korea supporting the development and design-in of ADI components into advanced consumer products.

Frank Kearney

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 doctoratal candidate at University College Dublin.