To appear in *IEEE Potentials*

Signals and Systems II

Part III: Analytic signals and QAM data transmission

Jeffrey O. Coleman
Naval Research Laboratory, Radar Division

This six-part series is a mini-course, focused on system concepts, that is aimed at the gap between Signals and Systems and the usual first DSP course.

Part II discussed oversampling in D/A conversion and the basics of decimation, complex signals, and Nyquist signaling. This third article in the series is about analytic signals and linear data modulation. Figures are numbered in one sequence across the entire series, and gaps appear in their numbering in some individual articles. Figures are posted for instructional use on the author's website:
http://alum.mit.edu/www/jeffc

PART III

The previous article in the series concluded with Nyquist filtering for data transmission. In that discussion a single filter was Nyquist, but it is more common in practical systems for a cascade of filters to be Nyquist when considered together. But before splitting our Nyquist filter into several component filters, we need another concept.

Analytic signals and filters

Nontrivial applications involving complex signals usually hinge on the two-step process illustrated in Fig. 18. First, a signal has its conjugate added to cancel its imaginary part and double its real part. The realization effort is almost negative: the imaginary part is simply never created at all, and the real part is created with twice the amplitude that it would otherwise have been. Second, the newly created conjugate mirroring is filtered out to *restore the imaginary part that was discarded* and thereby restore the original signal. We do not expect even magicians to successfully guess the nature of discarded signal components, imaginary or otherwise, so this is quite amazing, and it is useful as well because it permits us to transmit complex signals on channels that require real ones… sometimes.

But when? It is possible precisely when, as in Fig. 18, the original complex signal and its conjugate have nonoverlapping spectra. We call such a signal *analytic*. A filter, like the one in Fig. 18, that produces an analytic output from a real input is an *analytic filter*. We will use analytic filters frequently.

There is another way to look at the requirement of input analyticity. In Fig. 18 conjugating the signal created spectral components at frequencies that were originally unoccupied, but what if the new frequency is occupied already? The curved blue line in Fig. 18 shows that the conjugate of some undesired signal, a signal otherwise of no interest, could end up overlapping the desired signal spectrally. In fact, adding the conjugate always creates an *image band* of input frequencies that must remain empty to keep the desired signal free of such corruption. The conjugate of any actual signal in the image band—such a signal is called simply an *image*—will spectrally overlap the desired signal. To say that no image is present at the system input is always equivalent to saying that the input to the "add the conjugate" operation is analytic.

For the record only, an *ideal analytic filter* has a frequency response comprising unit-gain passbands, zero-gain stopbands, and the property that its sum with its conjugate mirroring is unity except at passband-stopband boundary frequencies and their negatives, where it is zero. Such a filter has an analytic impulse response and can produce only analytic signals at its output. When discussing only analog (more properly we'd say continuous time) or only discrete-time (d.t.) signals and filters, many authors actually define analytic signals in a more restricted way as the possible outputs of one of the ideal analytic filters in Fig. 19.

Using separate voltages or currents for the real and imaginary parts of signals makes it difficult to obtain adequate stopband rejection in analog realizations of analytic filters. So filters like those in Figs. 18 and 19 are typically emulated with a combination of analog and digital filtering, as in Fig. 20, where a digital analytic bandpass filter is cascaded with a conventional analog bandpass filter. The latter is responsible for suppressing unwanted duplicate passbands implied by the digital filter's periodicity. This approach is used universally below, and the most noble identity is the key that makes it simple.

Carrierless QAM

The signaling system of Fig. 21 parallels the Nyquist signaling system of Fig. 14 from Part II, reproduced here, but with changes for practicality. The original Nyquist D/A is now split equally, in magnitude anyway, into a *root-Nyquist* transmitter D/A and a root-Nyquist analog receiver filter. Because the signal band falls completely to one side of the origin, both root-Nyquist responses are analytic, so the Fig. 18 idea has also been used to add the conjugate in the transmitter, transmit a real signal over the channel, and filter out the conjugate in the receiver. In principle the location of the Nyquist signal band is otherwise unrestricted.

For our purposes the splitting of the Nyquist filter technically need not be into halves of equal magnitude, but that extra requirement turns out to maximize system performance when Gaussian noise is added in the channel. See the matched-filter discussion in any text on detection theory or statistical communication theory.

Creating the real channel signal through filtering and conjugate addition alone makes this data-modulation system *carrierless*, because frequency shifting is not involved. (Carriers will be discussed later.) The term *carrierless QAM* (quadrature amplitude modulation) is often used for historical reasons, particularly if the constellation from which input data samples are drawn forms a square grid in the complex plane.

Let us design a realistic DSP-based architecture for such a system. Begin by splitting the root-Nyquist transmit D/A of Fig. 21 into a standard D/A and an analytic filter that has a tilted passband to compensate for the frequency-response droop of the standard D/A. Then use the Fig. 20 idea to split each of the analytic filters, one in the transmitter and one in the receiver, into an analytic digital filter followed by a real analog filter. These steps produce the more detailed system on the left in Fig. 22.

The frequency-response periods of the digital filters must be large enough for the signal-shaping root-Nyquist passband, the image band, and transition bands. In this particular design the smallest adequate oversampling ratio at the input to the digital transmit filter is five. The frequency-response periods of the transmit and receive digital filters need not be identical, although here they are. Each analog filter could be made bandpass or lowpass; one of each is shown.

The system on the left in Fig. 22, which has its steps ordered according to our design and analysis process, is reordered into the system on the right for practical realization. Wherever red arrows cross each other in going from left to right, there must be careful justification. The bottom crossing ***(circled)*** is permitted by the most noble identity. The other receiver crossing and the upper transmitter crossing are simple reorderings of filters. The other two crossings in the transmitter, however, represent something new.

Adding the conjugate commutes with the operation that precedes it if that operation is time-domain multiplication with or convolution with something real. Showing this is a bit too involved to do here but is not difficult, and the reader is encouraged to work it out. This permitted reordering will be referred to from time to time as *swapping conjugate addition with an adjacent real operation.*

In the Fig. 22 transmitter then, conjugate addition can be swapped with both the D/A and the analog filter, simplifying them in the process by providing them with real inputs. It cannot be moved past the digital filter, which has a complex impulse response, so instead it simply combines with the digital filter to simplify its realization by permitting only the real part of its output to be computed with, of course, a factor of two more gain. In the receiver, both the real and imaginary parts of the output of the complex digital filter must still be computed. But the filter input is real, and this saves substantial computation.

Carrier modulation and demodulation

Adding the conjugate makes complex signals real, and this operation on analytic signals is reversible and therefore useful. Our carrierless-QAM system used a filter to make a signal analytic, but that goal is more commonly accomplished by using a *frequency shift*: time-domain multiplication by a complex exponential. This is just frequency-domain convolution with an impulse and so slides the entire signal spectrum left or right to move its origin to the impulse frequency.

The modulator and demodulator structures in Fig. 23 use this idea. After a frequency shift turns a bandlimited input into an analytic one, adding the conjugate gives this classic modulator a real output. The classic demodulator on the left does a frequency shift and then filters out the component created by the pre-shift conjugation. In the demodulator shifting comes first so that a real post-shift filter will suffice.

Aside: To discover the more traditional but less enlightening way to draw the modulator, do the algebra: take input $i(t) + jq(t)$*, shift by frequency* $-f_c$ *by multiplying by* $e^{-j2\pi f_c t}$*, and add the conjugate by taking twice the real part. Then derive and sketch the real and imaginary parts of the (left) demodulator's output separately to reflect realization of the complex analog output as a pair of real voltages or currents.*

A frequency shift can always be interchanged with an adjacent filter, but the filter frequency response must itself be shifted as it moves "through" the shift in the direction of signal flow. But for the classic demodulator on the left in Fig. 23 to evolve into the alternate on its right, the filter must move opposite to the direction of signal flow, so its frequency response must be "unshifted" as denoted symbolically, if awkwardly, in Fig. 23 by the small horizontal red arrow. In all the Fig. 23 systems, modulation shifts the spectrum to the left and demodulation shifts it to the right. This could just as well be reversed.

Aside: Our frequency shifts for modulation and demodulation are in opposite directions here. If instead they were in the same direction, how would the output signal be different?

The next paragraph presents two sets of traditional terminology in parallel. Read it once ignoring parenthesized terms and again with parenthesized terms substituted.

Passband (baseband) signals are those in the *passband (baseband)* portions of the system, where signals are in their shifted (unshifted) form. A complex baseband signal is of the form $i(t) + jq(t)$, and the real (imaginary) part of a complex baseband signal or signal path is its *in-phase (quadrature)* or just I (Q) component. The system performs *IQ* or *quadrature modulation (demodulation)* or *IQ upconver-*

sion (downconversion), with IQ commonly written as *I-Q, I/Q,* or even *I-and-Q*.

The sinusoidal real and imaginary parts of the complex exponential are the *carriers* and are said to be at the *carrier frequency* f_c, the absolute value of the frequency shift, and because the carriers are phased 90° apart, they are said to be *quadrature carriers* or to be *in quadrature*. The modulator *modulates the I and Q signals onto quadrature carriers*.

The demodulator filter on the left in Fig. 23 is a baseband filter and is the *baseband equivalent* of the demodulator filter on the right. The latter filter is in the passband of the system on the right and is therefore the *passband equivalent* of the filter in the left system. Though baseband-passband filter equivalence here is about frequency responses related through a shift, more generally it is about having equivalent effects on the signals. If a lowpass baseband filter were to be added to the modulator input in Fig. 23, for example, it could be replaced with a passband-equivalent bandpass filter at the modulator output.

Having the root-Nyquist filters be analytic was the key to the carrierless Nyquist data-transmission concept of Fig. 21. Now the alternative is clear: use nonanalytic root-Nyquist filters and then, instead of adding the conjugate between them as shown, do IQ modulation and demodulation between them. Of the two alternatives, the IQ-modulation version is more common and generally viewed as simpler.

Each such data-modulation strategy corresponds to a path through the state-transition diagram of Fig. 24. Many systems do actually transform the signal step by step, exactly as shown, but multi-step combinations are common and for data transmission may even include the root-Nyquist filtering, the D/A operation, or the sampling step.

Part IV will continue with discussions of analog and DSP-based IQ demodulation.

Read more about it

Here are two of the many texts on communication theory. The first is more accessible, and the second is more encyclopedic.

B. P. Lathi, *Modern Digital and Analog Communication Systems*, 3rd ed. Oxford University Press: www.oup.com, 1998.

J. R. Barry, E. A. Lee, and D. G. Messerschmitt, *Digital Communication*, 3rd ed. Springer: www.springer.com, 2003.

About the author

Jeffrey O. Coleman (S'75–M'79–SM'99) joined the Radar Division of the Naval Research Laboratory (NRL) in Washington DC in 1978 then left it in 1985 for graduate studies, for a stint with The Boeing Company, and for a faculty position at Michigan Technological University from which he returned to NRL in 1997. His 1975/1979/1991 SBEE/MSEE/PhD degrees are from the Massachusetts Institute of Technology, Johns Hopkins University, and the University of Washington respectively, and his research is on theory and design methods in DSP. More: http://alum.mit.edu/www/jeffc

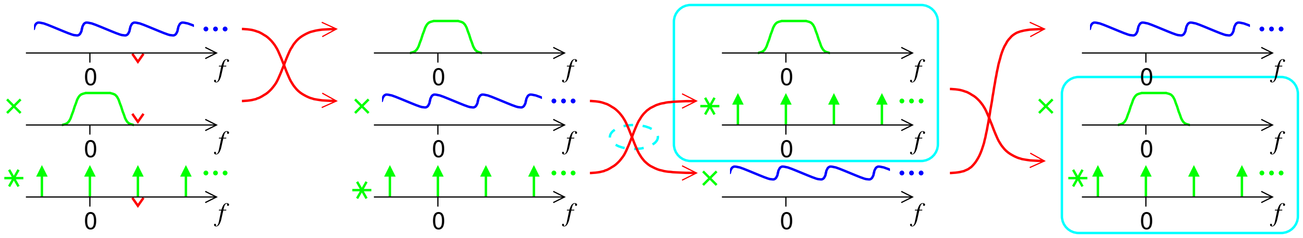

Fig. 14 (from Part II) The *most noble identity* (center) and product reordering (other red arrows) permit the top-to-bottom order of operations to be altered according to the *parenthetical groupings* and lead, on the right, to the Nyquist criterion.

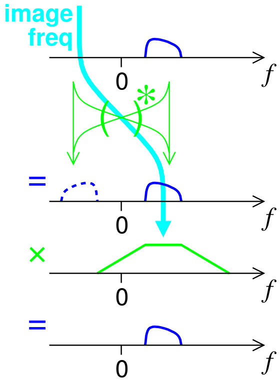

Fig. 18 Make a complex signal real by adding its conjugate. If the signal was analytic (image free) it can be restored by analytic filtering.

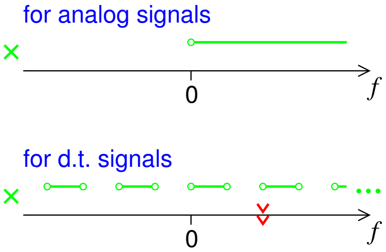

Fig. 19 The ideal analytic filters that define the most common notions of analytic signals.

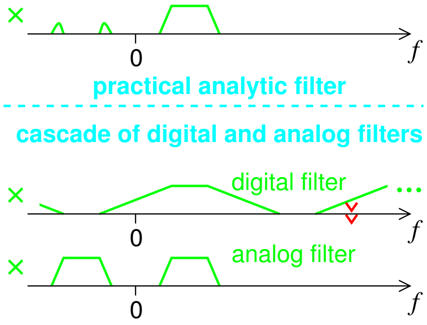

Fig. 20 A practical analytic filter system for analog inputs is nearly always equivalent to a cascade of a digital analytic filter and an analog bandpass or lowpass filter.

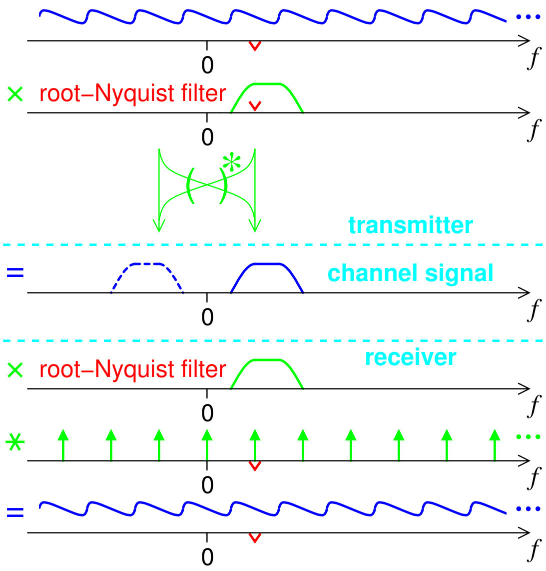

Fig. 21 A simple concept for carrierless transmission of complex data on a real analog channel results from adding the conjugate between two analytic root-Nyquist filters.

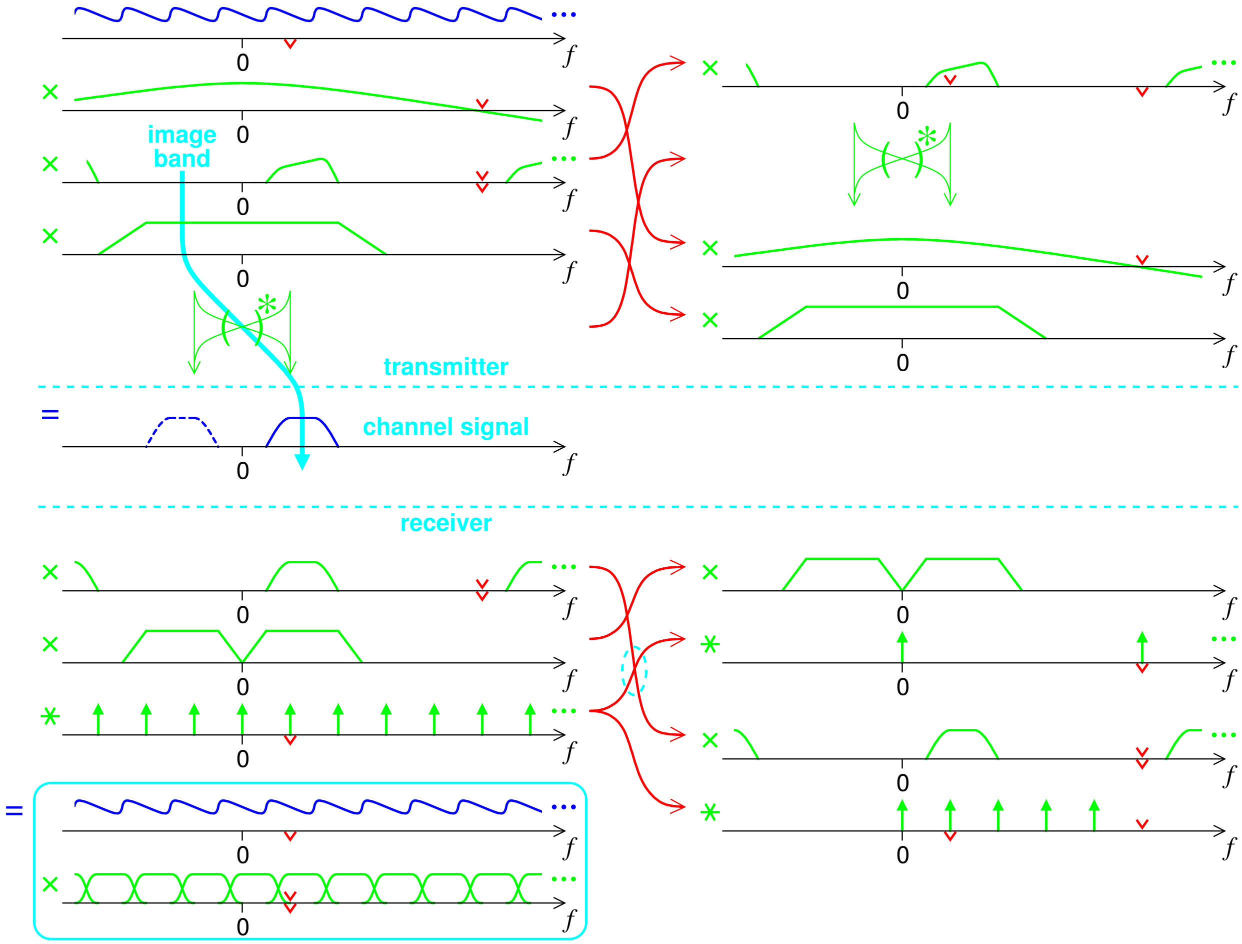

Fig. 22 Architecture of a carrierless-QAM modem. On the left: the transmit root-Nyquist filter of Fig. 21 has been replaced by an equivalent cascade of a standard D/A, an analog filter, and a digital analytic filter, with the latter two functioning together as an analog analytic filter as per Fig. 20. The receive analytic root-Nyquist filter of Fig. 21 has here been split into two filters as per Fig. 20. On the right: decimation is factored out of receiver sampling, and steps are reordered for practicality using the *most noble identity* and swapping conjugate addition with adjacent real operations. For rate compatibility input interpolation is included in the first filter.

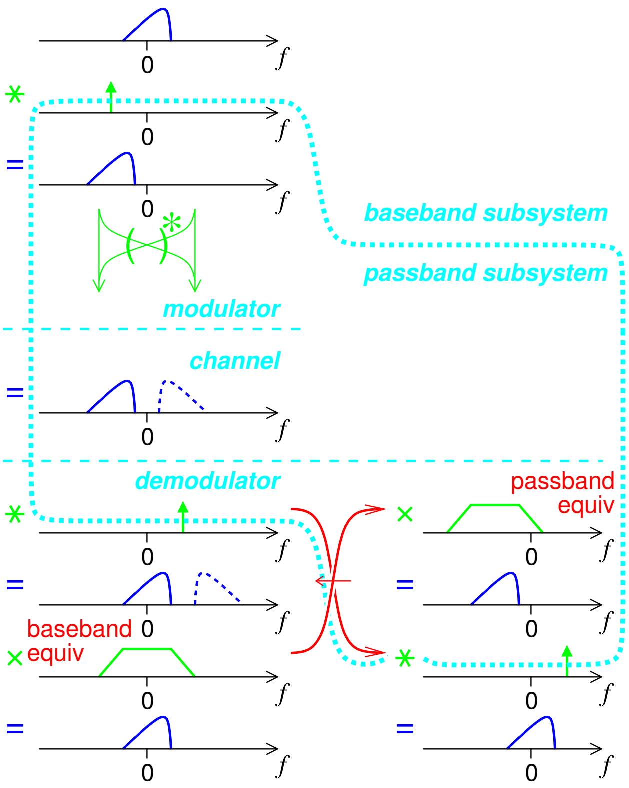

Fig. 23 A classic linear or IQ modulator and the usual shift-then-filter IQ demodulator on the left and an alternate, filter-then-shift IQ demodulator on the right.

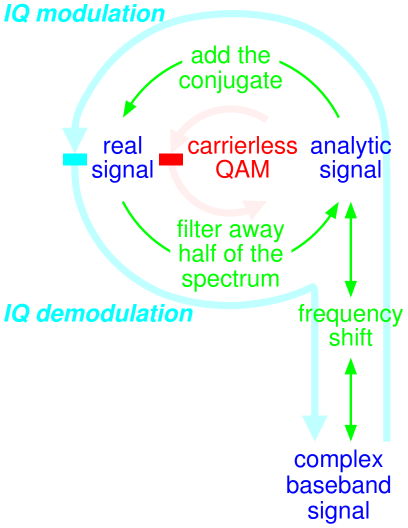

Fig. 24 IQ and carrierless-QAM modulation and demodulation as paths through a "state diagram" of signal types.