Understanding digital signal processing's frequency domain

DSP has a reputation for being difficult to learn. A few pointers about the math and notation of discrete spectra can help.

By Richard Lyons

RF engineers face many obstacles when learning digital signal processing (DSP). One major challenge is understanding how spectral components are defined and illustrated in the discrete frequency domain. When RF folk begin to read DSP literature, they encounter strange new terminology describing discrete spectra such as folding frequency, aliasing, orthogonal images, and Nyquist frequency.

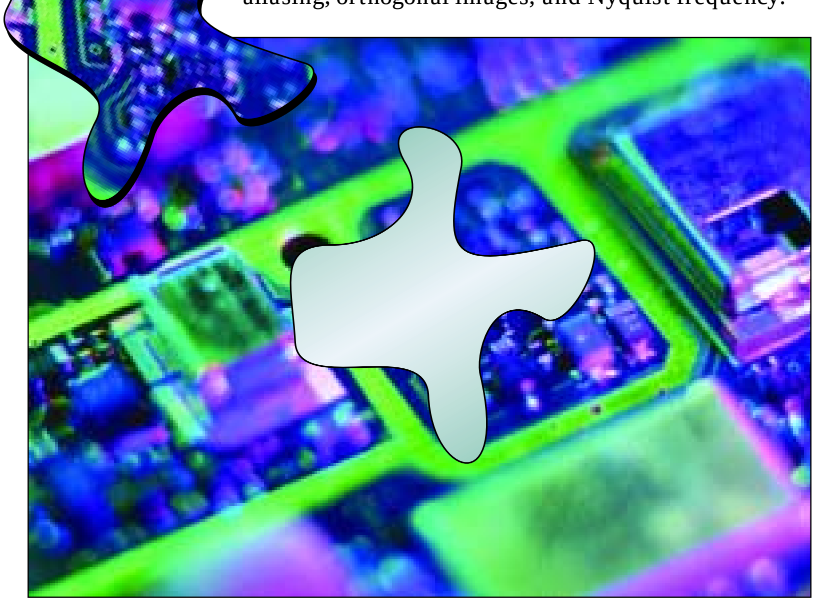

Understanding digital signal processing can be very confusing. Figuring out the DSP puzzle begins with an understanding of the math and notation of discrete spectra.

It's confusing

Typical DSP spectral diagrams initially seem peculiar because they often show negative-frequency spectral components and what appear to be replicated spectral components.

Making matters worse for the inquisitive RF engineer, various DSP authors use different, and sometimes puzzling, notation in labeling frequency axes in their spectral plots—the frequency dimension of hertz is often not used at all in discrete spectral diagrams.

For example, many university DSP textbooks actually label the discrete frequency axis from $-\pi$ to $+\pi$. This perplexing frequency-domain terminology and notation originate from a kind of frequency ambiguity inherent in discrete (sampled) systems and the fact that in DSP, all signals are generally described as if they were complex (with real and imaginary parts).

Understanding the differences between analog and discrete spectra is one of the reasons DSP has the reputation for being difficult to learn. Fortunately, several books have been published that ease the RF engineer's burden of learning DSP[1-3].

The first step on the road to understanding

The short journey to understanding the math and notation of discrete spectra begins by discussing the frequency-domain ambiguity associated with discrete signals. It ends at the point of understanding all the subtle aspects, notations, and language of the DSP's discrete frequency domain. Briefs stops will be made along the way to review complex signals, negative frequency, and discrete spectrum analysis using Fast Fourier transforms (FFT).

Frequency-domain ambiguity

Begin by reviewing a source of one difficulty: the frequency-domain ambiguity that exists when digitizing a continuous (analog) signal $x(t)$ with an analog-to-digital (A/D) converter (see Figure 1).

This process samples the continuous $x(t)$ signal to produce the $x(n)$ sequence of binary words stored in the computer for follow-on processing. (Variable 'n' is a dimensionless integer used as the independent time-domain index in DSP, just as the letter 't' is used in continuous-time equations.) The $x(n)$ sequence represents the voltage of $x(t)$ at periodically spaced instants in time. This process is called periodic sampling. The time period between samples is designated as t_s. It will be measured in seconds and defined as the reciprocal of the sampling frequency f_s, i.e., $t_s = 1/f_s$. In the literature of DSP, the f_s sampling frequency is often given the dimensions of 'samples/second,' but sometimes its dimension is indicated as Hz because f_s shows up on the axis of the discrete spectral diagrams.

Looking at an example (see Figure 2), consider the effect of sampling a 400 Hz sinusoidal $x(t)$ waveform at a sampling frequency f_s = 1 kHz shown in Figure 2(a). The $x(n)$ discrete-time samples from the A/D converter are plotted as the dots, and they are separated in time by one millisecond ($1/f_s$). The first three samples of the $x(n)$ sequence are $x(0) = 0$, $x(1) = 0.59$, and $x(2) = -0.95$.

The frequency-domain ambiguity under review

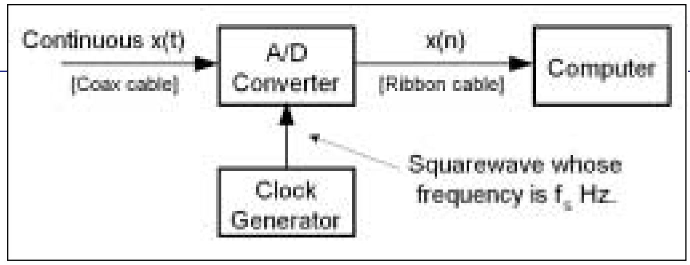

Figure 1. Periodic sampling, or digitizing, of a continuous signal.

here is illustrated in Figure 2(b), where the $x(n)$ samples would be unchanged if the A/D's $x(t)$ input was a 1400 Hz continuous sinusoid. Another example is shown in Figure 2(c), where the continuous $x(t)$ is a –600 Hz sinusoid, again resulting in identical $x(n)$ samples. This means that, given the $x(n)$ samples alone, it cannot be determined whether the continuous $x(t)$ sinewave's frequency was 400 Hz, 1400 Hz, or −600 Hz. That uncertainty is the frequency-domain ambiguity. (The concept of negative frequency will be justified later. For now, the definition of the –600 Hz sinewave is one whose phase is shifted by 180° relative to a +600 Hz sinewave.)

Based on the previous discussion, it is reasonable to assume that there are an infinite number of other frequencies that a sinusoidal $x(t)$ could have, and still result in the same $x(n)$ samples of Figure 2. Those other frequencies, defined as images with frequencies $f_i(k)$, can be identified by:

$$f_i(k) = 400 + k f_s \text{ Hz} \quad (1)$$

where k is an integer, and the 'i' subscript means image. Equation (1) indicates that any continuous sinewave whose frequency differs from 400 Hz by an integer multiple of f_s is indistinguishable from a 400 Hz sinewave in the world of discrete samples. A few of the images of 400 Hz, when f_s = 1 kHz, are listed in Figure 3(a).

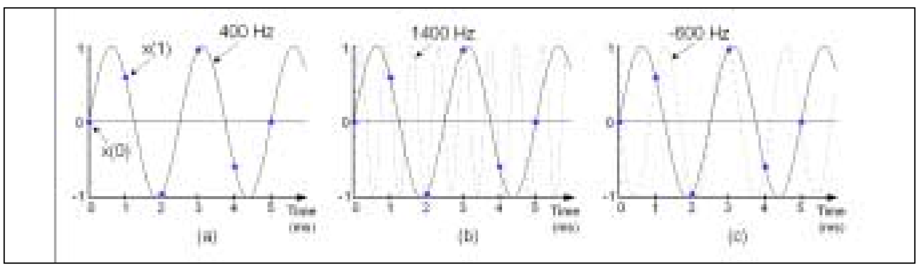

Figure 2. Frequency-domain ambiguity shown while digitizing sinusoids whose frequencies are (a) 400 Hz; (b) 1400 Hz, dashed curve; and (c) -600 Hz dashed curve.

Terminology and presentation

When the pioneers of DSP discovered and understood this frequency-domain ambiguity they were faced with the questions of what terminology to use in its description and how discrete spectral diagrams should be drawn. One reasonable frequency-domain depiction of this situation is shown in Figure 3(b), where it can be said that the spectrum of our discrete $x(n)$ sequence is an infinite set of spectral impulses periodically spaced in the frequency domain. (Mathematical proofs justifying Figure 3(b) are available[4,5].

Keep three thoughts in mind:

First, the notion that the spectrum of the discrete $x(n)$ sequence is an infinite set of spectral impulses does not imply that $x(n)$ has infinite energy. Those multiple impulses in Figure 3(b) merely indicate that $x(n)$ could be a sampled (discrete-time) version of any one of infinitely many continuous sinewaves.

Second, resist the temptation to call those spectral impulses harmonics. The word harmonic has a specific meaning in the analog world—related to spurious tones generated by nonlinear hardware components—that is not implied by Figure 3(b). For now, define those frequency-domain impulses as spectral replications. Third, notice that the spacing between the spectral replications is the sampling rate f_s.

The next step is to make the discrete spectral diagram more accurate and consistent with real-world signals by tackling the concepts of complex (quadrature) signals and negative frequency. This issue is critical because so much of DSP deals with complex numbers such as the complex-valued (magnitude and phase) spectra of discrete time-domain sequences, the complex-valued frequency responses a digital filters, and the complex signals needed to build modern digital communications systems.

k	$f_i(k)$
...	...
-2	-1600 Hz
-1	-600 Hz
0	400 Hz
1	1400 Hz
2	2400 Hz
...	...

(a)

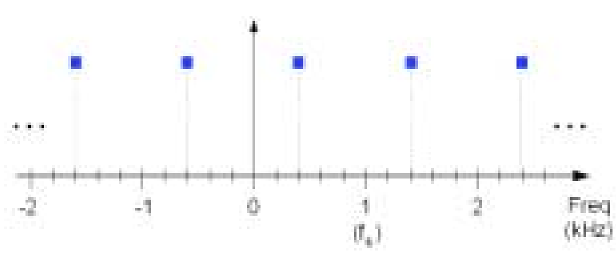

(b)

Figure 3. Examples of image frequencies of 400 Hz when fs is 1 kHz, (a) a few examples, (b) one possible frequency-domain depiction.

Complex signals and road maps

For this discussion, the focus is on a complex signal having a real and an imaginary part that is a function time.

Recall, as Karl Gauss first recommended, a single complex number can be represented by a point on the two-dimensional complex plane with its two axes (real and imaginary) orthogonal to one another. This means there is a 90° difference in the axes' orientations.

Consider a complex number whose magnitude is one, and whose phase angle increases with time. That complex number is the $e^{j2\pi f_o t}$ point shown on the complex plane in Figure 4. (Here the $2\pi f_o$ term is frequency in radians/second, and it corresponds to a frequency of f_o cycles/second where f_o is measured in Hz.) As time t gets larger, the complex number's phase angle $\phi = 2\pi f_o t$ increases and the number orbits the origin of the complex plane in a counterclockwise direction. Figure 4 shows that number, represented by the solid dot, frozen at some arbitrary instant in time. (That rotating $e^{j2\pi f_o t}$ complex number goes by two names in DSP literature; it's often called a complex exponential, and it's also referred

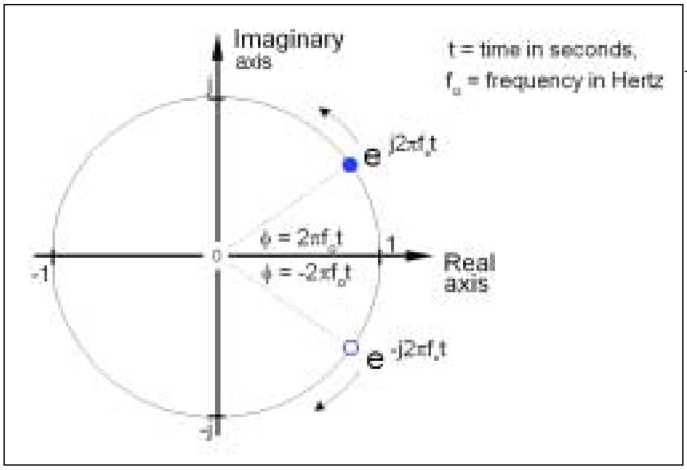

Figure 4. A snapshot, in time, of two complex numbers whose exponents, and thus their phase angles, change with time.

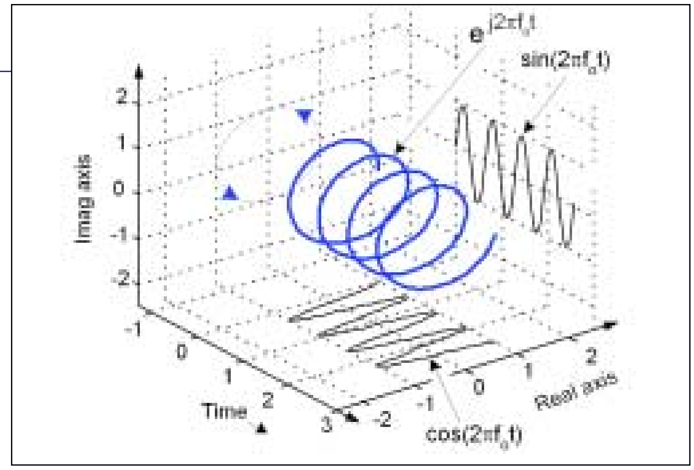

Figure 5. The value of the $e^{j2\pi f_o t}$ complex exponential signal.

to as a quadrature signal.) If the frequency f_o = 2 Hz, then the solid dot would rotate around the circle two times, or two cycles, per second.

Because complex numbers can be represented in both polar and rectangular notation, the polar $e^{j2\pi f_o t}$ quadrature signal (using one of Leonhard Euler's identities) in rectangular form is:

$$e^{j2\pi f_o t} = \cos(2\pi f_o t) + j\sin(2\pi f_o t) \qquad (2)$$

Equation (2) indicates that as $e^{j2\pi f_o t}$ rotates around the origin, its real part, the east-west distance from the origin, varies as a cosine wave. Its imaginary part, the north-south distance from the origin, varies as a sinewave. (Now the association between road maps and quadrature sinusoids becomes more clear.)

The attributes of our two-dimensional $e^{j2\pi f_o t}$ complex exponential are best illustrated with a three dimensional time-domain (see Figure 5). Notice how the $e^{j2\pi t}$ signal spirals beautifully along the time axis with its real part being a cosine wave and its imaginary part being a sinewave. At time t = 0 the signal has a value of 1 + j0 as we would expect. (Equation (2) allows the representation of a single complex exponential as the orthogonal sum of real cosine and real sine functions.)

There really *is* a signal

That $e^{j2pf_o t}$ signal is not just mathematical mumbo jumbo. It can be physically generated and transmitted to a laboratory down the hall. All that is required are two sinusoidal signal generators, set to the same frequency, f_o. (However, there is the requirement that, somehow, the generators must be synchronized so their relative phase shift is fixed at 90°. Their outputs need to be orthogonal.)

> ***Pop quiz:***
>
> *Q. What would be seen on the scope's display if the cables were mislabeled and the two real signals were inadvertently swapped?*
>
> *A. Another circle orbiting in a clockwise direction.*

Next, coax cables labeled 'cosine' for the cosine signal and 'sine' for the sinewave signal are connected to the generators' output connectors and ran down the hall to their destination. At the receiving end in the other lab, connect the real signals to the horizontal and vertical input channels of an oscilloscope (see Figure 6). The signals would show as rotating counterclockwise in a circle on the scope's display. (Remembering, of course, to set the scope's horizontal sweep control to the 'external' position.)

This oscilloscope example helps us answer the important question, "When working with quadrature signals, how is the j-operator implemented in hardware?" The answer is that the j-operator is implemented by how the two real signals are treated in relation to one another. They must be treated orthogonally such that the cosine signal represents the real (east-west) value, and the sinewave signal represents the imaginary (north-south) value. So in the oscilloscope example, the j-operator is implemented merely by how the connections are made to the scope. The result is a two-dimensional quadrature signal represented by the instantaneous position of the dot on the scope's display.

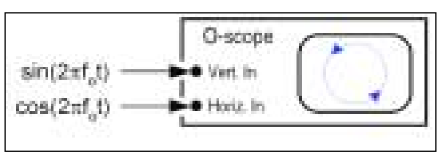

Figure 6. Displaying a quadrature e^{j2pfot} signal using an oscilloscope.

(If the instantaneous phase of the $e^{j2\pi f_o t}$ signal is controlled based on some bipolar binary data (+1 and –1), the other lab could measure that phase at certain instants in time and extract that binary data. Many digital communications systems operate on this principle).

OK, back to business—at this point the question may arise, "Where does the idea of negative frequency come in here?" Well, there's a "negative frequency" signpost up ahead, and the answer lies just up the road.

Don't negate negative frequencies

The notion of negative frequency is often troubling to RF engineers who have spent so much time examining the spectra displayed on analog spectrum analyzers. Some RF engineers think of frequency, by its very nature, as something that cannot be negative—such as getting into a car and driving minus ten miles. Negative frequency can be given a solid physical meaning by defining it properly in the context of complex, or quadrature, signals.

Referring again to Figure 4, notice another complex exponential, e–j2πfot, the white dot, orbiting in a clockwise direction. This occurs because its phase angle, f = –2πfot, gets more negative as time increases. Again, if the frequency fo = 2 Hz, then the white dot would rotate around the circle two times, or two cycles, per second in the clockwise direction. By definition, it is the rotational frequency minus two cycles per second. The two complex exponentials in Figure 4 are of great interest because of what is obtained when they are combined algebraically.

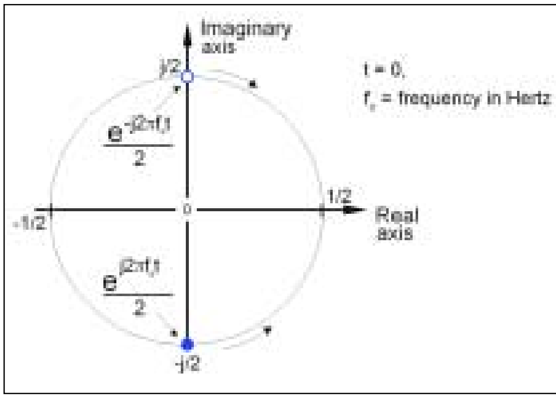

Figure 7. The two complex exponentials, at time t = o, that comprise a sinewave.

For example, what is the sum of the positive-frequency counterclockwise rotating $e^{j2\pi f_o t}$ and the negative-frequency clockwise rotating $e^{-j2\pi f_o t}$ when their real and imaginary parts are added separately? The resulting sum is an oscillating function whose imaginary part is always zero. That real-only sum is a cosine wave whose peak amplitude is 2. If the magnitudes of the complex exponentials in Figure 4 had been 0.5 instead of 1, they would graphically depict another important Euler identity:

$$\cos(2\pi f_o t) = \frac{e^{j2\pi f_o t}}{2} + \frac{e^{-j2\pi f_o t}}{2} \qquad (3)$$

Equation (3) provides the ability to represent a real cosine wave as the sum of positive-frequency and negative-frequency complex exponentials. By previous definitions, a positive-frequency complex exponential's exponent is positive, and a negative-frequency complex exponential is one whose exponent is negative.

Another Euler identity, Equation. (4), gives the relationship of a real sinewave as the sum of positive-frequency and negative-frequency complex exponentials.

$$\sin(2\pi f_o t) = j\frac{e^{-j2\pi f_o t}}{2} - j\frac{e^{j2\pi f_o t}}{2} \qquad (4)$$

The j-operators in Equation (4) merely explain the relative phase of the complex exponentials at time $t = 0$ as illustrated in Figure 7. At time $t = 0$, Equation (4) becomes:

$$\sin(2\pi f_o t)\big|_{t=0} = j\frac{e^{0}}{2} - j\frac{e^{0}}{2} = \frac{j}{2} - \frac{j}{2} = 0 \qquad (5)$$

complying with our knowledge that a sinewave's amplitude is zero at time $t = 0$.

Don't worry if these concepts of the j-operator and complex exponentials seem a little perplexing at this point. They take some time to get used to. (Even the great Karl Gauss struggled with these ideas at first. He called the j-operator "the shadow of shadows.")

Back on track

Remember that the ultimate goal is to understand the nature of the spectral diagrams used in DSP. In doing so, it was necessary to define the notion of negative frequency. That definition is inherent in the complex-valued (real and imaginary) representation used for discrete spectra in DSP. Unlike the amplitude-only results seen when using an analog spectrum analyzer, in the DSP world spectrum analysis provides complex-valued results. That is, discrete spectra show the relative phase shift between spectral components.

Yet more complexities

For the following discussion, refer to the complex spectra of a few simple sinusoids, from the viewpoint of Euler's identities, as shown in Figure 8.

The time-domain waveform and the complex spectra of a sinewave defined by $\sin(2\pi f_o t)$ is shown in Figure 8(a). Shifting that sinewave in time by 90° develops the cosine wave shown in Figure 8(b). Another shift in time by $\phi°$ results in the arbitrary-phase cosine wave of Figure 8(c).

Remember now, the positive and negative-frequency spectral components of the sinewave rotated counterclockwise and clockwise, respectively, by 90° in going from Figure 8(a) to Figure 8(b). If those cosine waves' spectral components continued their rotation by $\phi°$, the result will be the situation shown in Figure 8(c). These three-dimensional frequency-domain spectra are shown replete with phase information because, in the world of DSP, interest in spectral phase relationships should be observed. The FFT algorithm is used to measure spectral magnitude and phase the way an RF engineer uses a network (vector) analyzer. (Note that Figure 8 illustrates an important signal processing principle. A time-domain shift of a periodic signal results only in phase shifts in the frequency domain; spectral magnitudes do not change.)

The top portion of Figure 8 illustrates Equation (3) and the center portion is a graphical description of Equation (4). The goal is now looming on the horizon. Figure 8 is a reminder that one legitimate, and useful, way to show the spectrum of a real cosine wave—one that could be transmitted using a 50Ω coax cable—is to include both positive and negative-frequency spectral components.

With this thought in mind, refer to the drawing of the spectral magnitude (ignoring any phase information) of a continuous 400 Hz sinusoid as shown in Figure 9(a). This figure shows the inherent spectral symmetry at about 0 Hz when representing real signal spectra with complex exponentials. In this case, "real signal" is defined by an $x(t)$ signal having a non-zero real part

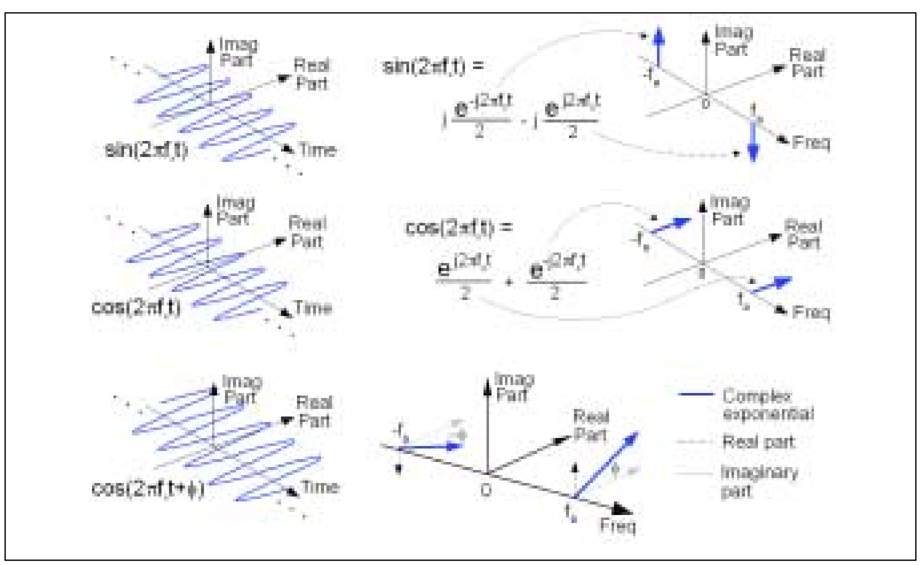

Figure 8. Complex frequency domain representation of three sinusoids.

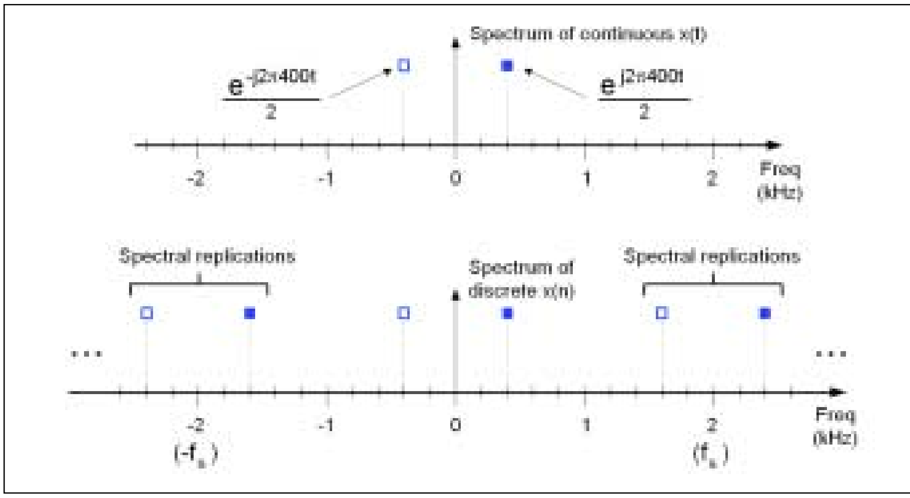

Figure 9. The spectral magnitude plot of (a) a 400 Hz continuous sinusoid, and (b) a discrete sequence of a 400 Hz sinusoid sampled at a 2 kHz sampling rate.

but whose imaginary part is always zero. (Convention developed in this article treats all signals as complex and regards real signals as a special case of complex signals.) Figure 9(a) is another graphical representation of Euler's identity in Equation (3).

Applying this convention of "spectral replications due to periodic sampling," to this analysis illustrates the spectral magnitude of discrete samples of a 400 Hz sinusoid, sampled at an f_s = 2 kHz sampling rate, as that of Figure 9(b).

Finally, closing the loop, Figure 9(b) is typical of the spectral magnitude representations used in the DSP literature. It combines the spectral replications (centered about integer multiples of f_s) due to periodic sampling, as well as the use of negative frequency components resulting from representing real signals in complex notation.

One more time

To further cement the concept, a review of the spectrum of another discrete sequence is presented. Figure 10(a) shows the spectral magnitude of a continuous $x(t)$ signal having four components in the range of 100 Hz to 700 Hz . Dark and light squares distinguish the positive- and negative-frequency spectral components. Figure 10(b) shows the spectral replication for a discrete $x(n)$ sequence that's $x(t)$ sampled at 2 kHz.

The sole purpose of this article is to provide the meaning, relevance, and validity of Figure 10(b) in representing the spectrum of discrete samples of a real sinusoid in the complex-valued world of DSP. This figure is a reminder of the following important properties: Continuous real signals have spectral symmetry of about 0 Hz, and discrete real signals have spectral symmetry of about 0 Hz and $f_s/2$ Hz.

Figure 10 illustrates the reason that Nyquist Criterion for lowpass signals—signals whose spectral components are centered about 0 Hz—states that the f_s sampling rate must be equal to or greater than twice the highest spectral component of $x(t)$. Because $x(t)$'s highest spectral component is 700 Hz, the f_s sample rate must be no less than 1.4 kHz. If f_s were 1.3 kHz, as in Figure 10(c), the spectral replications would be too close together and spectral overlap would occur. It can be seen that the spectrum in the range of –1 kHz to +1 kHz in Figure 10(c) does not correctly represent the original spectrum in Figure 10(a).

This situation is typically called aliasing, and it results in $x(n)$ sample values that contain amplitude errors. Unfortunately, for meaningful, information-carrying signals, there is no way to correct for those errors.

It can be seen that the spectral overlap is centered about $f_s/2$ and that this particular frequency is important enough to have its own nomenclature. It is sometimes called the folding frequency, but more often it's called the Nyquist frequency. Because of this anomaly, the following important statement relating continuous and discrete signals can be made. "Only continuous frequency components as high as the Nyquist frequency ($f_s/2$) can be unambiguously represented by a discrete sequence obtained at an f_s sampling rate." Figure 10(c) also presents another fundamental connection between the worlds of continuous and discrete signals. All of the continuous $x(t)$ spectral energy shows up in the discrete $x(n)$ sequence's spectral frequency range of $-f_s/2$ to $+f_s/2$.

The purpose for showing replicated spectra, as in Figure 10, is not to cause complication or confusion, but to provide a straightforward explanation for the effects of overlapped spectra due to aliasing. Drawing replicated spectra is useful in illustrating the spectral translation that takes place in bandpass sampling, and describing the result of frequency translation operations such digital downconversion. With that said, this article concludes with an explana-

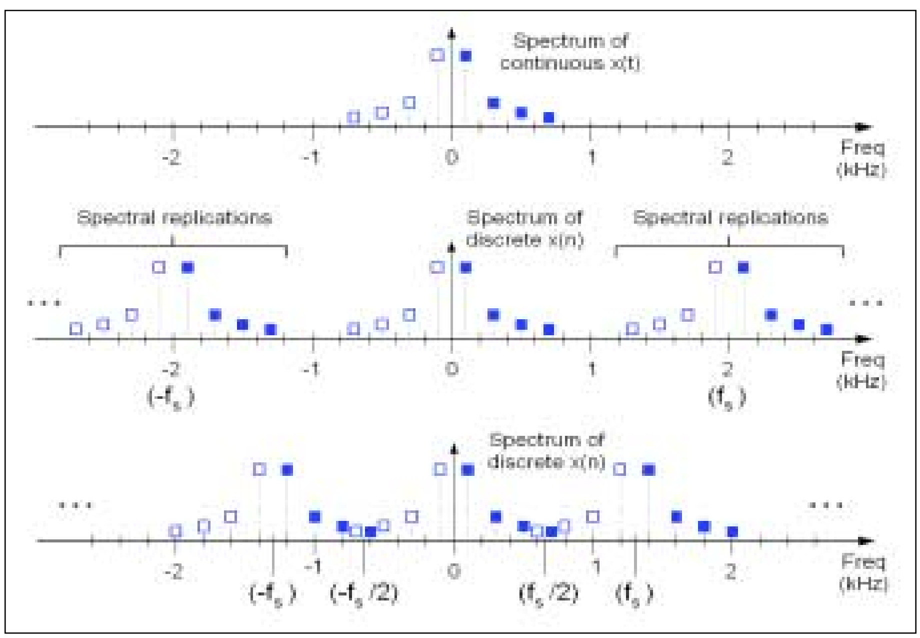

Figure 10. Spectrum of a signal with four components in the range of 100 Hz to 700 Hz. (a) Spectral magnitude of the continuous signal. (b) Spectrum of a sampled sequence when fs = 2 kHz, and (c) when fs = 1.3 kHz.

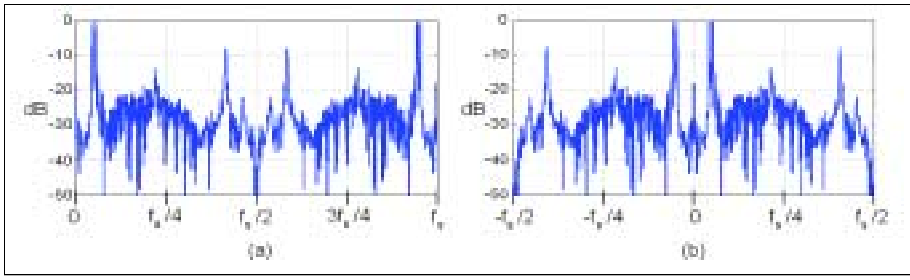

Figure 11. Example spectral magnitude plots; (a) 0 Hz on the left, or (b) 0 Hz in the center.

tion of the various, and sometimes puzzling, notation used in frequency-axis labeling encountered.

Discrete frequency-axis notation

For convenience, frequency-domain drawings in the DSP world are often labeled using the f_s sampling rate. This value of one, which leads to the notation that $\omega_s = 2\pi$. Thus, in their DSP books, frequency-domain plots are often shown as in Figure 12. These show the frequency axis is a normalized angle with $-f_s/2$ replaced with $-\pi$, and $f_s/2$ replaced with π.

A common scheme for labeling the

Frequency axis notation:	dimensions	frequency axis range	period of repetition
Cyclic freq.	Hz	$-f_s/2$ to $+f_s/2$	f_s
Radial freq.	Radians/sec	$-\omega_s/2$ to $+\omega_s/2$	ω_s
Normalized angle	Radians	$-\pi$ to $+\pi$	2π
Normalized discrete freq.	Radians/sample	–0.5 to +0.5	1

Table 1. Characteristics of various frequency-axis notations.

convention is best explained with a couple of examples; the first of which is performing spectrum analysis (using the FFT) of, for example, a real time-domain audio sequence obtained at an f_s = 11.025 kHz rate. Using either convention, a plot of the spectral magnitude results is shown in Figure 11. If it is later discovered that the sample rate was actually f_s = 22.05 kHz, it would not be necessary to repeat the analysis nor redraw spectral plots because the frequency axis is referenced to f_s.

Another example of referencing discrete frequency-domain plots to the f_s sample rate is in describing digital filters. A five-point moving average digital filter has the frequency magnitude response shown in Figure 12. That frequency response is the same whether the filter is used in an f_s = 8 Msample/s digital communications system or in an f_s = 1 sample/day stock market price analysis.

DSP authors have several choices in labeling the frequency axis of their frequency-domain plots. Some reference their frequency-axis to the f_s sample rate, as in Figures 11 and 12. Others reference their frequency-axis to a sample rate measured, not in cycles/second, (Hz) but in radians/second so that the sample rate is called ω_s, where $\omega_s = 2\pi f_s$[4,5]. DSP purists often, to make the notation more concise, assign f_s a discrete frequency axis normalizes all frequencies to the f_S sampling rate. The justification for doing so is as follows: If a representation of a sinewave whose frequency is f_o Hz by $x(t) = \sin(2\pi f_o t)$. Discrete samples of $x(t)$ is:

$$x(n) = x(t)\big|_{t=nt_s} = \sin(2\pi f_o t)\big|_{t=nt_s} = \sin(2\pi f_o n t_s) \quad (6)$$

then, with the factors $2\pi f_o$ having the dimension of radians/second, and t_s having the dimension seconds/sample, the resultant angle in Equation (6) has the dimension of radians/sample. If Equation (6)'s t_s is replaced with $1/f_s$, the discrete sinusoidal samples can be represented by:

$$x(n) = \sin\left(2\pi \frac{f_o}{f_s} n\right) = \sin(\theta_o n) \quad (7)$$

where θ_o becomes a normalized discrete frequency in the range of –0.5 to +0.5 measured in radians/sample. This definition is why some authors like to say, "For continuous signals, frequency is measured in radians/second. For discrete signals, frequency is measured in radians/sample." Redrawing the spectrum from Figure 11(b) illustrates the normalized angle and normalized frequency axis representations in Figure 13. Those four popular frequency-axis notations are shown in Table 1.

While it takes DSP beginners some time to become comfortable with these various frequency-axis notations, commercial signal processing modeling software packages, like SystemView, Mathcad, and MATLAB, allow for convenient labeling of the frequency-domain plots in hertz.[6-8].

Conclusion

This discussion has reviewed the graphical depictions and terminology of DSP to explain the differences between continuous (analog) and discrete spectrum analysis with regard to spectral replications and the idea of negative frequency.

The use of spectral replications in DSP diagrams is a way of accounting for the inherent frequency-domain ambiguity when performing periodic sampling of a continuous signal. These replications provided a consistent explanation for errors due to aliasing, and illustrated that the spectral replications should not be interpreted as harmonics.

The use of Euler's identities relating real-only signals and complex exponentials lead to a definition of negative frequency that allows the representation of all signals in complex notation. (Thankfully, the developed definition for negative frequency did not violate past experience in analog RF signal processing.) However, three-dimensional time and frequency domain drawings were required to give a solid physical meaning to those complex signals in the time domain and to show the relative phase of their spectral components.

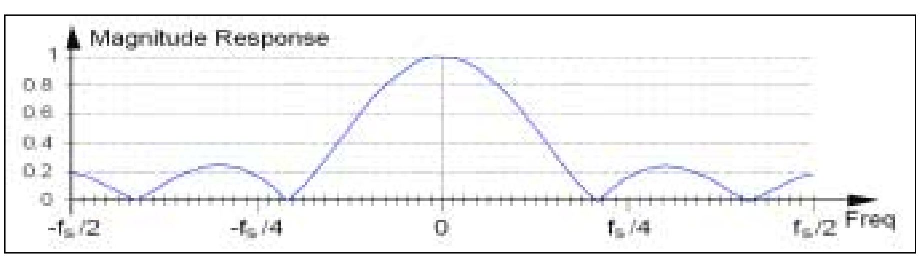

Figure 12. Frequency magnitude response of a five-point moving average ("boxcar") digital filter.

Finally, the relationship between four popular methods for labeling the frequency axis in the world of DSP were defined and reviewed.

RF

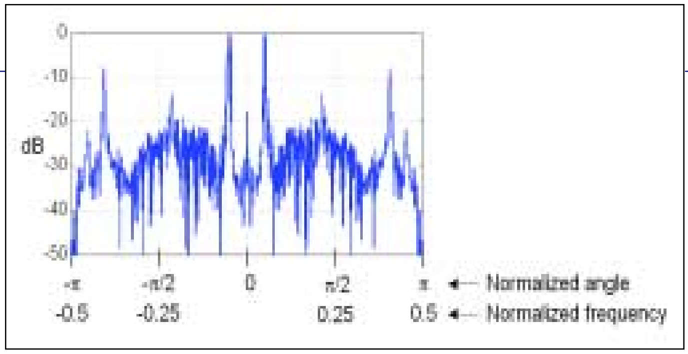

Figure 13. A discrete spectral magnitude plot using the normalized angle and normalized frequency axis representations.

References

[1] R. Lyons, "Understanding Digital Signal Processing," Addison Wesley Longman, Reading, Massachusetts, 1997.

[2] D. Glover and J. Deller, "Digital Signal Processing and the Microcontroller," Prentice Hall, Upper Saddle River, New Jersey, 1998.

[3] S. Smith, "The Scientist and Engineer's Guide to Digital Signal Processing," California Technical Publishing, San Diego, California, 1997.

[4] A. Oppenheim, et al, "Discrete-Time Signal Processing," Prentice Hall, Upper Saddle River, New Jersey, 1999, Section 4.2.

[5] J. Proakis and D. Manolakis, "Digital Signal Processing-Principles, Algorithms, and Application," Prentice Hall, Upper Saddle River, New Jersey, 1996, Section 4.2.9.

[6] SystemView - Elanix, Westlake Village, CA, **http://www.elanix.com.**

[7] Mathcad - MathSoft Eng. & Education, Cambridge, MA, http://www.mathsoft.com.

[8] MATLAB - The Mathworks, Natick, MA, **http://www.mathworks.com.**

About the author

Richard Lyons is a consulting systems engineer and lecturer with Besser Associates in Mt. View, CA. He has been the lead hardware engineer for numerous multimillion dollar signal processing systems for both the National Security Agency (NSA) and TRW. An experienced lecturer and instructor at the University of California Santa Cruz Extension, Lyons has delivered signal processing seminars and training courses throughout the United States and Europe. He has written numerous articles on DSP topics, and wrote the best-selling DSP book "Understanding Digital Signal Processing." Lyons is a member of the IEEE, Eta Kappa Nu, the electrical engineering honor society, and rides a 1981 Harley Davidson. He can be contacted at: *ricklyon@onemain.com*