Time and Frequency Domains

In this chapter, we explore the basic properties of signals in preparation for looking at how they interact with interconnects. We will find that there are multiple ways of looking at a signal, each providing a different perspective. The quickest path to the answer may not be the most obvious path. The different perspectives we will use to look at signals are called domains. In particular we'll use the time domain and the frequency domain.

We will find that while we may generally be more familiar with the time domain, the frequency domain can provide valuable insight to understand and master many signal-integrity effects such as impedance, lossy lines, the power-distribution network, measurements, and models.

After introducing the two domains, we will look at how to translate between the two for some special cases. We will apply what we learn to relate two important quantities: rise time and bandwidth. The first is a time-domain term and the second a frequency-domain term. However, as we will see, they are intimately related.

Finally, we'll apply this concept of bandwidth to interconnects, models, and measurements.

2.1 The Time Domain

We use the term a lot—*the time domain*. But what do we really mean? What is the time domain? What are the features that are special about the time domain that make it useful? These are surprisingly difficult questions to answer because they seem so obvious and we rarely think about what we really mean by the time domain.

> **TIP** The time domain is the real world. It is the only domain that actually exists.

We take it for granted because from the moment we are born, our experiences are developed and calibrated in the time domain. We are used to seeing events happen with a time stamp and ordered sequentially.

The time domain is the world of our experiences and is the domain in which high-speed digital products perform. When evaluating the behavior of a digital product, we typically do the analysis in the time domain because that's where performance is ultimately measured.

For example, two important properties of a clock waveform are clock period and rise time. Figure 2-1 illustrates these features.

The clock period is the time interval to repeat one clock cycle, usually measured in nanoseconds (nsec). The clock frequency, F_{clock}, or how many cycles per second the clock goes through, is the inverse of the clock period, T_{clock}.

$$F_{clock} = \frac{1}{T_{clock}} \tag{2-1}$$

where:

F_{clock} = the clock frequency, in GHz

T_{clock} = the clock period, in nsec

For example, a clock with a period of 10 nsec will have a clock frequency of 1/10 nsec = 0.1 GHz or 100 MHz.

The rise time is related to how long it takes for the signal to transition from a low value to a high value. There are two popular definitions of rise time. The 10–90 rise time is how long it takes for the signal to transition from 10% of its final value to 90% of its final value. This is usually the default meaning of rise time. It can be read directly off the time domain plot of a waveform.

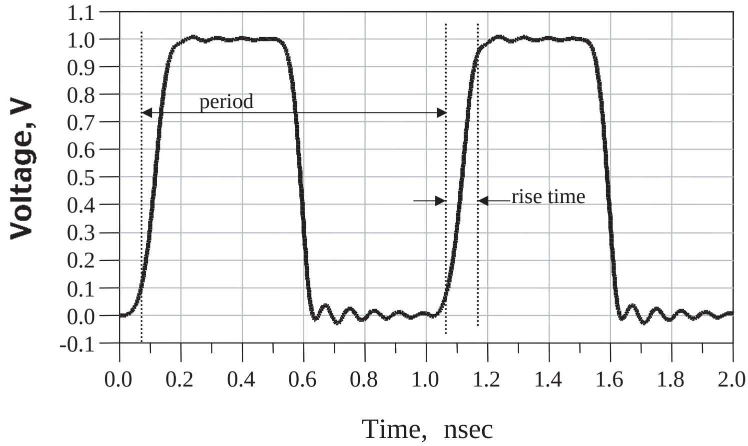

Figure 2-1 Typical clock waveform showing the clock period and the 10–90 rise time for a 1-GHz clock. The fall time is typically slightly shorter than the rise time and sometimes creates more noise.

The second definition is the 20–80 rise time. This is the time it takes for the signal to transition from 20% of its final value to 80% of its final value. Of course, for the same waveform the 20–80 rise time is shorter than the 10–90 rise time. Some IBIS models of real devices use the 20–80 definition of rise time. This makes it confusing. To remove ambiguity, it's often good practice referring explicitly to the 10–90 rise time or the 20–80 rise time.

There is a corresponding value for the fall time of a time-domain waveform. Depending on the logic family, the fall time is usually slightly shorter than the rise time. This is due to the design of typical CMOS output drivers. In a typical output driver, a p and an n transistor are in series between the V_{CC} (+) and the V_{SS} (–) power rails. The output is connected to the center, between them. Only one transistor is on at any one time, depending on whether the output is a low or a high.

When the driver switches from a low to a high (i.e., rising edge), the n transistor turns off and the p transistor turns on. The rise time is related to how fast the p transistor can turn on. When switching from the high to the low state (i.e., a falling edge), the p transistor turns off and the n transistor turns on. In general, for the same feature-size transistor, an n transistor can turn on faster than a p transistor. This means switching from high to low, the falling edge will be shorter than the rising edge. In general, signal-integrity problems are more likely to occur when switching from a high to low transition than from a low to high transition. By making the n channel transistor larger than the p channel, the rising and falling edges can be closely matched.

Having established an awareness of the time domain as a distinct way of looking at events, we can turn our attention to one of a number of alternative ways of analyzing the world—the frequency domain.

2.2 Sine Waves in the Frequency Domain

We hear the term *frequency domain* quite a bit, especially when it involves radio frequency (rf) or communications systems. We will also encounter the frequency domain in high-speed digital applications. There are few engineers who have not heard of and used the term multiple times. Yet, what do we really mean by the frequency domain? What is the frequency domain and what makes it special and useful?

TIP The most important quality of the frequency domain is that it is not real. It is a mathematical construct. The only reality is the time domain. The frequency domain is a mathematical world where very specific rules are followed.

The most important rule in the frequency domain is that the only kind of waveforms that exist are sine waves. Sine waves are the language of the frequency domain.

There are other domains that use other special functions. For example, the JPEG picture-compression algorithm takes advantage of special waveforms that are called wavelets. The wavelet transform takes the space domain, with a lot of x-y amplitude information content, and translates it into a different mathematical description that is able to use less than 10% of the memory to describe the same information. It is an approximation, but a very good one.

It's common for engineers to think that we use sine waves in the frequency domain because we can build any time-domain waveform from combinations of sine waves. This is a very important property of sine waves. However, there are many other waveforms with this property. It is not a property that is unique to sine waves.

In fact, there are four properties that make sine waves very useful for describing any other waveform. These properties are as follows:

1. Any waveform in the time domain can be completely and uniquely described by combinations of sine wave.
2. Any two sine waves with different frequencies are orthogonal to each other. If you multiply them together and integrate over all time, they integrate to zero. This means you can separate each component from every other.
3. They are well defined mathematically.
4. They have a value everywhere with no infinities and they have derivatives that have no infinities anywhere. This means they can be used to describe real world waveforms, since there are no infinities in the real world.

All of these properties are vitally important, but are not unique to sine waves. There is a whole class of functions called *orthonormal functions*, or sometimes called *eigenfunctions* or *basis functions*, which could be used to describe any time-domain waveform. Other orthonormal functions are Hermite Polynomials, Legendre Polynomials, Laguerre Polynomials, and Bessel Functions.

Why did we choose sine waves as our functions in the frequency domain? What's so special about sine waves? The real answer is that by using sine waves, some problems related to the electrical effects of interconnects will be easier to understand and solve using sine waves. If we switch to the frequency domain and use sine-wave descriptions, we can sometimes get to an answer faster than staying in the time domain.

TIP After all, if the time domain is the real world, we would never leave it unless the frequency domain provides a faster route to an acceptable answer.

Sine waves can sometimes provide a faster path to an acceptable answer because of the types of electrical problems we often encounter in signal integrity. If we look at the circuits that describe interconnects, we find that they will often include combinations of resistors (R), inductors (L), and capacitors (C). These elements in a circuit can be described by a second-order linear differential equation. The solution to this type of differential equation is a sine wave. In these circuits, the naturally occurring waveforms will be combinations of the waveforms that are solutions to the differential equation.

We find that in the real world, if we build circuits that contain Rs, Ls, and Cs and send any arbitrary waveform in, more often than not, we get waveforms out that look like sine waves and can more simply be described by a combination of a few sine waves. An example of this is shown in Figure 2-2.

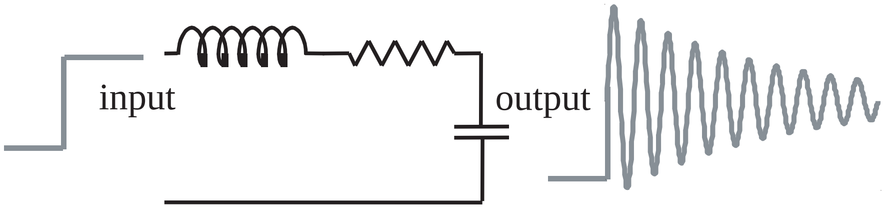

Figure 2-2 Time-domain behavior of a fast edge interacting with an ideal RLC circuit. Sine waves are naturally occurring when digital signals interact with interconnects, which can often be described as combinations of ideal RLC circuit elements.

2.3 Shorter Time to a Solution in the Frequency Domain

TIP The only reason we would ever want to move to another domain is to get to an acceptable answer faster.

In some situations, if we use the naturally occurring sine waves in the frequency domain rather than in the time domain, we may arrive at a simpler description to a problem and get to a solution faster.

It is important to keep in mind that there is fundamentally no new information in the frequency domain. The time- and the frequency-domain descriptions of the same waveforms will each have exactly the same information content.

However, some problems are easier to understand and describe in the frequency domain than in the time domain. For example, the concept of bandwidth is intrinsically a frequency-domain idea. We use this term to describe the most significant sine-wave frequency components associated with a signal, a measurement, a model, or an interconnect.

Impedance is defined in both the time and the frequency domain. However, it is far easier to understand, to use, and to apply the concepts of impedance in the frequency domain. We need to understand impedance in both domains, but we will often get to an answer faster by solving an impedance problem in the frequency domain first.

Looking at the impedance of the power and ground distribution in the frequency domain will allow a simpler explanation and solution to rail-collapse problems. As we shall see, the design goal for the power-distribution system is to keep its impedance below a target value from direct current (DC) up to the bandwidth of the typical signals.

When dealing with EMI issues, both the FCC specifications and the methods of measuring the electromagnetic compliance of a product are more easily performed in the frequency domain.

With today's current capabilities of hardware and software tools, the quality of the measurements and the computation speed of the numerical-simulation tools can sometimes be better in the frequency domain.

A high signal-to-noise ratio (SNR) means higher quality measurements. The SNR of a vector-network analyzer (VNA), which operates in the frequency domain, is constant over its entire frequency range, which can be –130 dB from 10 MHz up to 50 GHz and more. For a time-domain reflectometer (TDR), the effective bandwidth may be as high as 20 GHz, but the SNR starts at –70 dB at low frequency and drops to as low as –30 dB at 20 GHz.

Many of the effects related to lossy transmission lines are more easily analyzed, measured, and simulated by using the frequency domain. The series resistance of a transmission line increases with the square root of frequency, and the shunt AC leakage current in the dielectric increases linearly with frequency. The transient (time-domain) performance of lossy transmission lines is often more easily obtained by first transforming the signal into the frequency domain, looking at how the transmission line affects each frequency component separately, and then transforming the sine-wave components back to the time domain.

2.4 Sine Wave Features

As we now know, by definition, the only waveforms that exist in the frequency domain are sine waves. We should also be familiar with the description of a sine wave in the time domain. It is a well-defined mathematical curve that has three terms that fully characterize absolutely everything you could ever ask about it. An example is shown in Figure 2-3.

The following three terms fully describe a sine wave:

- Frequency
- Amplitude
- Phase

The frequency, usually identified using a small f, is the number of complete cycles per second made by the sine wave, in Hertz. Angular frequency is measured in radians per second. A radian is like degrees, describing a fraction of a cycle. There are $2 \times \pi$ radians in one complete cycle. The Greek letter ω is often

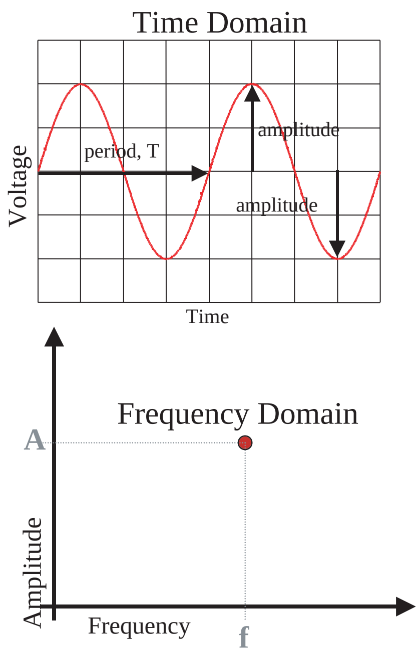

Figure 2-3 Top: Description of a sine wave in the time domain. It is composed of over one thousand voltage-versus-time data points. Bottom: Description of a sine wave in the frequency domain. Only three terms define a sine wave, which is a single point in the frequency domain.

used to refer to the angular frequency, measured in radians per second. The sine-wave frequency and the angular frequency are related by:

$$\omega = 2\pi \times f \tag{2-2}$$

where:

ω = angular frequency, in radians/sec

π = constant, 3.14159...

f = sine-wave frequency, in Hz

For example, if the frequency of a sine wave is 100 MHz, the angular frequency is $2 \times 3.14159 \times 100$ MHz ~ 6.3×10^8 radians/sec.

The amplitude is the maximum value of the peak height above the center value. The wave peak goes below the horizontal just as much as it goes above.

The phase is more complicated and identifies where the wave is in its cycle at the beginning of the time axis. The units of phase are in cycles, radians, or degrees, with 360 degrees in one cycle. While phase is important in mathematical analysis, we will minimize the use of phase in most of our discussion to concentrate on the more important aspects of sine waves.

In the time domain, describing a sine wave requires plotting a lot of voltage-versus-time data points to draw the complete sine-wave curve. However, in the frequency domain, describing a sine wave is much simpler.

In the frequency domain, we already know that the only waveforms we can talk about are sine waves, so all we have to identify are the amplitude, frequency, and phase. If there is only one sine wave we are describing, all we need are these three values and we have identified a complete description of the sine wave.

Since we are going to ignore phase for right now, we really only need two terms to completely describe a sine wave: its amplitude and its frequency. These two values are plotted with the frequency as one axis and the amplitude as the other axis, as shown in Figure 2-3. Of course, if we were including phase, we'd have a third axis. A sine wave, plotted in the frequency domain, is just one single data point. This is the key reason why we will go into the frequency domain. What might have been a thousand voltage-versus-time data points in the time domain is converted to a single amplitude-versus-frequency data point in the frequency domain.

When we have multiple frequency values, the collection of amplitudes is called the spectrum. As we will see, every time-domain waveform has a particular pattern to its spectrum. The only way to calculate the spectrum of a waveform in the time domain is with the Fourier Transform.

2.5 The Fourier Transform

The starting place for using the frequency domain is being able to convert a waveform from the time domain into a waveform in the frequency domain. We do this with the Fourier Transform. There are three types of Fourier Transforms:

- Fourier Integral (FI)
- Discrete Fourier Transform (DFT)
- Fast Fourier Transform (FFT)

The Fourier Integral (FI) is a mathematical technique of transforming an ideal mathematical expression in the time domain into a description in the frequency domain. For example, if the entire waveform in the time domain were just a short pulse, and nothing else, the Fourier Integral would be used to transform to the frequency domain.

This is done with an integral over all time from – infinity to + infinity. The result is a frequency-domain function that is also continuous from 0 to + infinity frequencies. There is a value for the amplitude at every continuous frequency value in this range.

For real-world waveforms, the time-domain waveform is actually composed of a series of discrete points, measured over a finite time, T. For example, a clock waveform may be a signal from 0 v to 1 v and have a period of 1 nsec and a repeat frequency of 1 GHz. To represent one cycle of the clock, there might be as many as 1000 discrete data points, taken at 1-psec intervals. An example of a 1-GHz clock wave in the time domain is shown in Figure 2-4.

To transform this waveform into the frequency domain, the Discrete Fourier Transform (DFT) would be used. The basic assumption is that the original time-domain waveform is periodic and repeats every T seconds. Rather than integrals, just summations are used so any arbitrary set of data can be converted to the frequency domain using simple numerical techniques.

Finally, there is the Fast Fourier Transform (FFT). It is exactly the same as a DFT, except that the actual algorithm used to calculate the amplitude values at each frequency point uses a trick of very fast matrix algebra. This trick works only if the number of time-domain data points is a power of two, for example 256 points, or 512 points, or 1024 points. The result is a DFT, only calculated 100–10,000 times faster than the general DFT algorithm, depending on the number of voltage points.

In general, it is common in the industry to use all three terms, FI, DFT, and FFT, synonymously. We now know there is a difference between them, but they have the same purpose—to translate a time-domain waveform into its frequency-domain spectrum.

TIP Once in the frequency domain, the description of a waveform is a collection of sine-wave frequency values. Each frequency component has an amplitude and phase associated with it. We call the entire set of frequency values and their amplitudes the spectrum of the waveform.

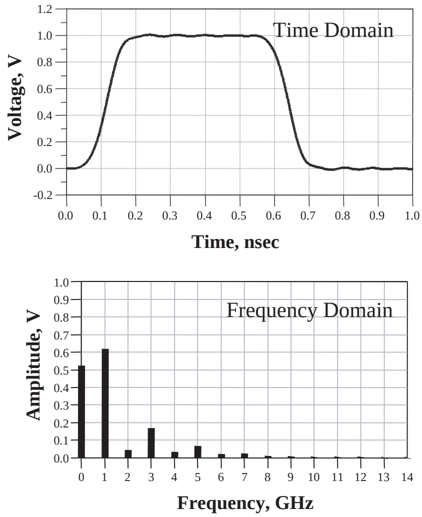

Figure 2-4 One cycle of a 1-GHz clock signal in the time domain (top) and frequency domain (bottom).

An example of a simple time-domain waveform and its associated spectrum, calculated by using a DFT, is shown in Figure 2-4.

At least once in his or her life, every serious engineer should calculate a Fourier Integral by hand, just to see the details. After this, we never again need to do the calculation manually. We can always get to an answer faster by using one of the many commercially available software tools that calculate Fourier Transforms for us.

There are a number of relatively easy-to-use, commercially available software tools that calculate the DFT or FFT of any waveform entered. Every version of SPICE has a function called the .FOUR command that will generate the amplitude of the first nine frequency components for any waveform. Most versions of the more advanced SPICE tools will also compute the complete set of amplitude and frequency values using a DFT. Microsoft Excel has an FFT function, usually found in the "engineering add-ins."

2.6 The Spectrum of a Repetitive Signal

In practice, the DFT or FFT is used to translate a real waveform from the time domain to the frequency domain. It is possible to take a DFT of any arbitrary, measured waveform. A key requirement of the waveform is that it be repetitive. We usually designate the repeat frequency of the time-domain waveform with the capital letter F.

For example, an ideal square wave might go from 0 v to 1 v, with a repeat time of 1 nsec and a 50% duty cycle. As an ideal square wave, the rise time to transition from 0 v to 1 v is precisely 0 sec. The repeat frequency would be 1/1 nsec = 1 GHz.

If a signal in the time domain is some arbitrary waveform over a time interval from t = 0 to t = T, it may not look repetitive. However, it can be turned into a repetitive signal by just repeating the interval every T seconds. In this case, the repeat frequency would be F = 1/T. Any arbitrary waveform can be made repetitive and the DFT used to convert it to the frequency domain. This is illustrated in Figure 2-5.

For a DFT, only certain frequency values exist in the spectrum. These values are determined by the choice of the time interval or the repeat frequency. When using an automated DFT tool, such as in SPICE, it is recommended to choose a value for the period equal to the clock period. This will simplify the interpretation of the results.

The only sine-wave frequency values that will exist in the spectrum will be multiples of the repeat frequency. If the clock frequency is 1 GHz, for example, the DFT will only have sine wave components at 1 GHz, 2 GHz, 3 GHz, etc.

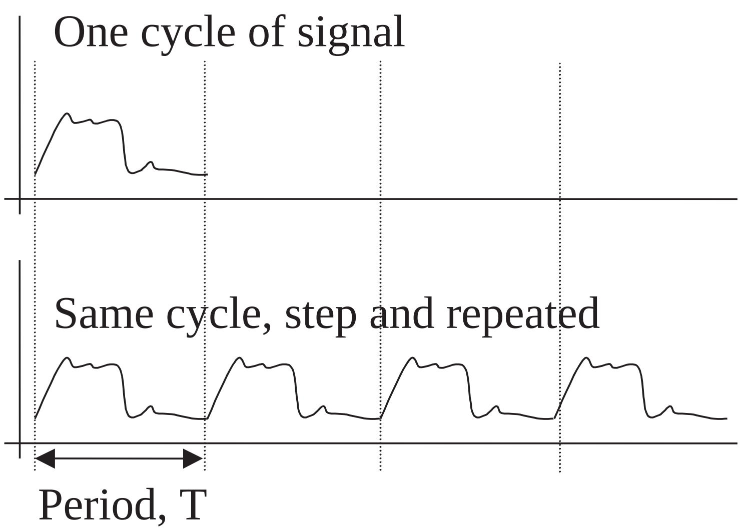

Figure 2-5 Any arbitrary waveform can be made to look repetitive. A DFT can be performed only on a repetitive waveform.

The first sine-wave frequency is called the first harmonic. The second sine-wave frequency is called the second harmonic, and so on. Each harmonic will have a different amplitude and phase associated with it. The collection of all the harmonics and their amplitudes is called the spectrum.

The actual amplitudes of each harmonic will be determined by the values calculated by the DFT. Every specific waveform will have its own spectrum.

2.7 The Spectrum of an Ideal Square Wave

An ideal square wave has a zero rise time, by definition. It is not a real waveform; it is an approximation to the real world. However, useful insight can be gained by looking at the spectrum of an ideal square wave and using this to evaluate real waveforms later. An ideal square wave has a 50% duty cycle, is symmetrical, and has a peak voltage of 1 v. This is illustrated in Figure 2-6.

If the ideal square-wave repeat frequency is 1 GHz, the sine-wave frequency values in its spectrum will be multiples of 1 GHz. We expect to see components at f = 1 GHz, 2 GHz, 3 GHz, and so on. But what are the amplitudes of each sine wave? The only way to determine this is to perform a DFT on the ideal square wave. Luckily, it is possible to calculate the DFT exactly for this special case of an ideal square wave. The result is relatively simple.

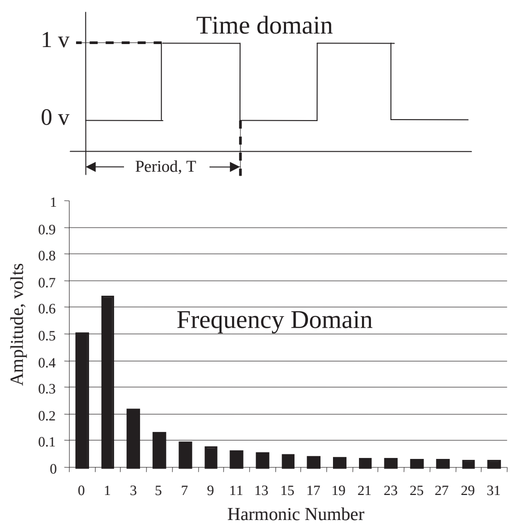

Figure 2-6 Time and frequency domain views of an ideal square wave.

The amplitudes of all the even harmonics (e.g., 2 GHz, 4 GHz, 6 GHz) are all zero. It is only odd harmonics that have values. The amplitudes, A_n, of the odd harmonics are given by:

$$A_n = \frac{2}{\pi \times n} \tag{2-3}$$

where:

A_n = the amplitude of the n^{th} harmonic

π = the constant, 3.14159...

n = the harmonic number, only odd allowed

For example, an ideal square wave with 50% duty-cycle and 0 v to 1 v transition has a first harmonic amplitude of 0.63 v. The amplitude of the third harmonic is 0.21 v. We can even calculate the amplitude of the 1001^{st} harmonic. It is 0.00063 v. It is important to note that the amplitudes of higher sine-wave-frequency components decrease with 1/f.

If the transition-voltage range of the ideal square wave were to double to 0 v to 2 v, the amplitudes of each harmonic would double as well.

There is one other special frequency value, 0 Hz. Since sine waves are all centered about zero, any combination of sine waves can only describe waveforms in the time domain that are centered about zero. To allow a DC offset, or a nonzero average value, the DC component is stored in the zero-frequency value. This is sometimes called the zeroth harmonic. Its amplitude is equal to the average value of the signal. In the case of the 50% duty-cycle square wave, the zeroth harmonic is 0.5 v.

To summarize:

- The collection of sine-wave-frequency components and their amplitudes is called the spectrum. Each component is called a harmonic.
- The zeroth harmonic is the DC value.
- For the special case of a 50% duty-cycle ideal square wave, the even harmonics have an amplitude of zero.
- The amplitude of any harmonic can be calculated as $2/(\pi \times n)$.

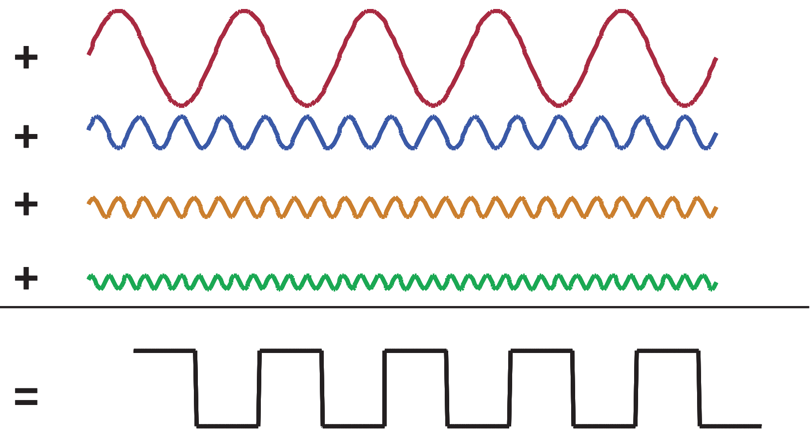

Figure 2-7 Convert the frequency-domain spectrum into the time-domain waveform by adding up each sine-wave component.

2.8 From the Frequency Domain to the Time Domain

The spectrum, in the frequency domain, represents all the sine-wave-frequency amplitudes of the time-domain waveform. If we have a spectrum and want to look at the time-domain waveform, we simply take each frequency component, convert it into its time-domain sine wave, then add it to all the rest. This process is called the Inverse Fourier Transform. It is illustrated in Figure 2-7.

Each component in the frequency domain is a sine wave in the time domain, defined from t = – infinity to t = + infinity. To re-create the time-domain waveform, we take each of the sine waves described in the spectrum and add them up in the time domain at each time-interval point. We start at the low-frequency end and add each harmonic based on the spectrum.

For a 1-GHz ideal-square-wave spectrum, the first term in the frequency domain is the zeroth harmonic, with amplitude of 0.5 v. This component describes a constant DC value in the time domain.

The next component is the first harmonic, which is a sine wave in the time domain with a frequency of 1 GHz and an amplitude of 0.63 v. When this is added to the previous term, the result in the time domain is a sine wave, offset to 0.5 v. It is not a very good approximation to the ideal square wave. This is shown in Figure 2-8.

The next term is the third harmonic. The amplitude of the 3-GHz sine-wave-frequency component is 0.21 v. When we add this to the existing time-domain

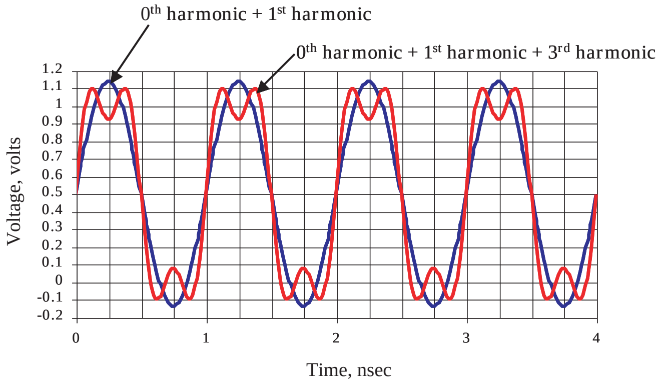

Figure 2-8 The time-domain waveform is created by adding together the zeroth harmonic and first harmonic and then the third harmonic, for a 1-GHz ideal square wave.

waveform, we see that it changes the shape of the new waveform slightly. The top is a bit more flat, better approximating a square wave, and the rise time is a little sharper. As we go through this process, adding each successive higher harmonic to re-create the ideal square wave, the resulting waveform begins to look more and more like a square wave. In particular, the rise time of the resulting time-domain waveform changes as we add higher harmonics.

To illustrate this in more detail, we can zoom in on the rise time of the waveform, centered about the beginning of a cycle. As we add all the harmonics up to the seventh harmonic, and then all the way up to the nineteenth, and finally, all the way up to the thirty-first harmonic, we see that the rise time of the resulting waveform in the time domain continually gets shorter. This is shown in Figure 2-9.

Depending on how the DFT was set up, there could be over 100 different harmonics listed in the spectrum. The logical question to ask is, do we have to include all of them, or can we still re-create a "good enough" representation of the original time-domain waveform with just a limited number of harmonics? What really is the impact of limiting the highest harmonic included in the re-created time-domain waveform? Is there a highest sine-wave-frequency component at which we can stop?

2.9 Effect of Bandwidth on Rise Time

The term *bandwidth* is used for the highest sine-wave-frequency component that is significant in the spectrum. This is the highest sine-wave frequency we need to include to adequately approximate the important features of the time-domain

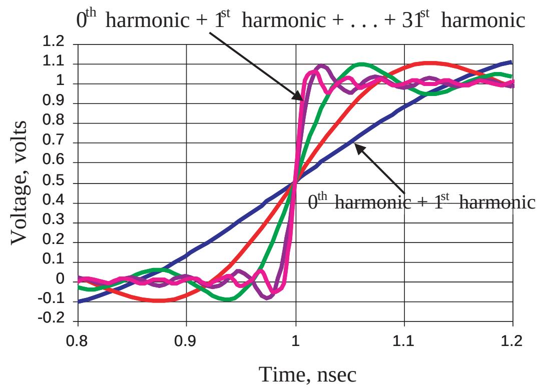

Figure 2-9 The time domain waveform created by adding together the zeroth harmonic and first harmonic, then the third harmonic and then up to the seventh harmonic, then up to the nineteenth harmonic, and then all harmonics up to the thirty-first harmonic, for a 1-GHz ideal square wave.

waveform. All frequency components of higher frequency than the bandwidth can be ignored. In particular, as we will see, the bandwidth we choose will have a direct effect on the shortest rise time of the signal we are able to describe in the time domain.

The term *bandwidth* historically is used in the rf world to refer to the range of frequencies in a signal. In rf applications, a carrier frequency is typically modulated with some amplitude or phase pattern. The spectrum of frequency components in the signal falls within a band. The range of frequencies in the rf signal is called the bandwidth. Typical rf signals might have a carrier frequency of 1.8 GHz with a bandwidth about this frequency of 100 MHz. The bandwidth of an rf signal defines how dense different communications channels can fit.

With digital signals, bandwidth also refers to the range of frequencies in the signal's spectrum. It's just that for digital signals, the low frequency range starts at DC and extends to the highest frequency component. In the world of digital signals, since the lowest frequency will always be DC, bandwidth will always be a measure of the highest sine wave frequency component that is significant.

When we created a time-domain waveform from just the zeroth, the first, and the third harmonics included, as in Figure 2-8, the bandwidth of the resulting waveform was just up to the third harmonic, or 3 GHz in this case. By design, the highest sine-wave-frequency component in this waveform is 3 GHz. The amplitude of all other sine-wave components in this time-domain waveform is exactly 0.

When we added higher harmonics to create the waveforms in Figure 2-9, we designed their bandwidths to be 7 GHz, 19 GHz, and 31 GHz. If we were to take the shortest rise-time waveform in Figure 2-9 and transform it back into the frequency domain, its spectrum would look exactly like that shown in Figure 2-6. It would have components from the zeroth to the thirty-first harmonics. Beyond the thirty-first harmonic, all the components would be zero. The highest sine-wave-frequency component that is significant in this waveform is the thirty-first harmonic, or the waveform has a bandwidth of 31 GHz.

In each case, we created a waveform with a higher bandwidth, using the ideal-square-wave's spectrum as the starting place. And, in each case, the higher-bandwidth waveform also had a shorter 10–90 rise time. The higher the bandwidth, the shorter the rise time and the more closely the waveform approximates an ideal square wave. Likewise, if we do something to a short rise-time signal to decrease its bandwidth (i.e., eliminate high-frequency components), its rise time will increase.

For example, it is initially difficult to evaluate the time-domain response of a signal propagating down a lossy transmission line in FR4. As we will see, there are two loss mechanisms: conductor loss and dielectric loss. If each of these processes were to attenuate low-frequency components the same as they do high-frequency components, there would simply be less signal at the far end, but the pattern of the spectrum would look the same coming out as it does going in. There would be no impact on the rise time of the waveform.

However, both conductor loss and dielectric loss will attenuate the higher-frequency components more than the low-frequency components. By the time the signal has traveled through even four inches of trace, the high-frequency components, above about 8 GHz, can have lost more than 50% of their power, leaving the low-frequency terms less affected. In Figure 2-10 (top), we show the measured attenuation of sine-wave-frequency components through a four-inch length of transmission line in FR4. This transmission line happens to have a 50-Ohm characteristic impedance and was measured with a network analyzer. Frequency components below 2 GHz are not attenuated more than –1 dB, while components at 10 GHz are attenuated by –4 dB.

This preferential attenuation of higher frequencies has the impact of decreasing the bandwidth of a signal that would propagate through the interconnect. Figure 2-10 (bottom) is an example of the measured rise time of a 50-psec signal entering a 36-inch-long trace in FR4 and this same waveform when it exits the trace. The rise time has been increased from 50 psec to nearly 1.5 nsec, due to the higher attenuation of the high-frequency components. Thirty-six inches is a typi-

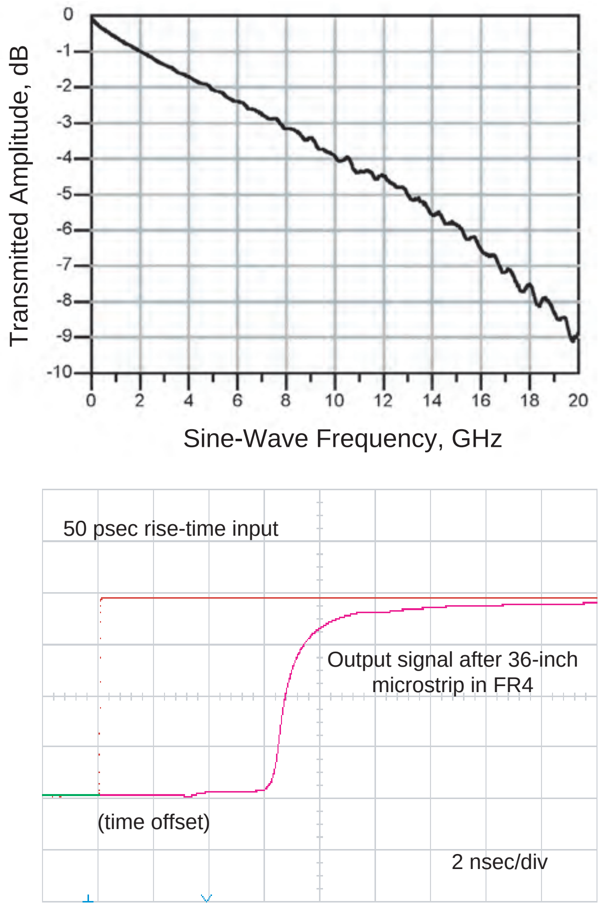

Figure 2-10 Top: The measured attenuation through a 4-inch length of 50-Ohm transmission line in FR4 showing the higher attenuation at higher frequencies. Bottom: The measured input and transmitted signal through a 36-inch 50-Ohm transmission line in FR4, showing the rise time to have degraded from 50 psec to more than 1.5 nsec.

cal length for a trace that travels over two 6-inch-long daughter cards and 24 inches of backplane. This rise-time degradation is the chief limitation to the use of FR4 laminate in high-speed serial links above 1 GHz.

TIP In general, a shorter rise-time waveform in the time domain will have a higher bandwidth in the frequency domain. If something is done to the spectrum to decrease the bandwidth of a waveform, the rise time of the waveform will be increased.

The connection between the highest sine-wave-frequency component that is significant in a spectrum and the corresponding rise time of the waveform in the time domain is a very important property.

2.10 Bandwidth and Rise Time

The relationship between rise time and bandwidth for a re-created ideal square wave can be quantified. In each synthesized waveform in the previous example re-creating an ideal square wave, the bandwidth is explicitly known because each waveform was artificially created by including sine-wave-frequency components only up to a specified value. The rise time, defined as the time from the 10% point to the 90% point, can be measured from time-domain plots.

When we plot the measured 10–90 rise time and the known bandwidth for each waveform, we see that empirically there is a simple relationship. This is a fundamental relationship for all signals and is shown in Figure 2-11.

For the special case of a re-created square wave with only some of the higher harmonics included, the bandwidth is inversely related to the rise time. We can fit a straight-line approximation through the points and find the relationship between bandwidth and rise time as:

$$\mathrm{BW} = \frac{0.35}{\mathrm{RT}} \tag{2-4}$$

where:

BW = the bandwidth, in GHz

RT = the 10–90 rise time, in nsec

For example, if the rise time of a signal is 1 nsec, the bandwidth is about 0.35 GHz or 350 MHz. Likewise, if the bandwidth of a signal is 3 GHz, the rise time of the signal will be about 0.1 nsec. A signal with a rise time of 0.25 nsec, such as might be seen in a DDR3-based system, has a bandwidth of 0.35/0.25 nsec = 1.4 GHz.

There are other ways of deriving this relationship for other waveforms, such as with Gaussian or exponential edges. The approach we took here for square waves is purely empirical and makes no assumptions. It is one of the most useful rules of thumb in our toolbox.

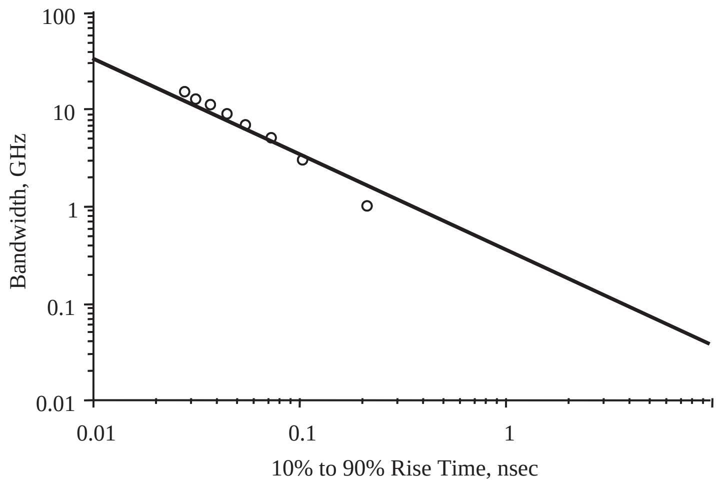

Figure 2-11 Empirical relationship between the bandwidth of a signal and its 10–90 rise time, as measured from a re-created ideal square wave with each harmonic added one at a time. Circles are the values extracted from the data; line is the approximation of BW = 0.35/rise time.

It is important to keep the units straight. When rise time is in microseconds, the bandwidth is in MHz. For example, a very long rise time of 10 microseconds has a bandwidth of about 0.35/10 microsec = 0.035 MHz. This is equivalent to 35 kHz.

When the rise time is in nanoseconds, the bandwidth is in GHz. A 10-nsec rise time, typical of a 10-MHz clock frequency, has a bandwidth of about 0.35/10 nsec = 0.035 GHz or 35 MHz.

2.11 What Does *Significant* Mean?

We defined the bandwidth of a signal as the highest sine-wave-frequency component that is significant. In the example, where we started with an ideal square wave and limited the high-frequency components, there was absolutely no ambiguity about what *significant* meant. We explicitly cut off all higher frequency sine-wave components in the frequency domain so that the highest significant component was the last harmonic in the spectrum.

We simply showed that if we include 100% of all the frequency components of an ideal square wave, up to the bandwidth, we would be able to re-create a square wave with a limited rise time, where the relationship of rise time = 0.35/BW. But what is the impact from adding only a fraction of the next component?

For example, if we take an ideal-square-wave clock signal with clock frequency of 1 GHz, its first harmonic will be a 1-GHz sine-wave frequency. If we were to include 100% of every component up to the twenty-first harmonic, the bandwidth would be 21 GHz and the resulting rise time of the re-created signal would be 0.35/21 GHz = 0.0167 nsec or 16.7 psec.

How would the rise time change if we added the twenty-third harmonic? The rise time would be 0.35/23 GHz = 0.0152 nsec or 15.2 psec. The rise time dropped by 1.5 psec. This is about 10% of the rise time, which is consistent, because we increased the bandwidth by 10%. The magnitude of the component we added was just 0.028 v, compared with the first harmonic of 0.63 v. Even though this amplitude is a small amount, less than 5% of the first harmonic amplitude and less than 3% of the peak value of the original square wave, it had the impact of dropping the rise time by 10%.

The spectrum of an ideal square wave has components that extend to infinite frequency. In order to achieve the zero rise time of an ideal square wave, each of these components is needed and is significant.

For a real time-domain waveform, the spectral components will almost always drop off in frequency faster than those of an ideal square wave of the same repeat frequency. The question of significance is really about the frequency at which amplitudes of the higher harmonics become small compared to the corresponding amplitudes of an ideal square wave.

By "small," we usually mean when the power in the component is less than 50% of the power in an ideal square wave's amplitude. A drop of 50% in power is the same as a drop to 70% in amplitude. This is really the definition of significant. Significant is when the amplitude is still above 70% of an ideal square wave's amplitude of the same harmonic.

TIP For any real waveform that has a finite rise time, *significant* refers to the point at which its harmonics are still more than 70% of the amplitude of an equivalent repeat-frequency ideal square wave's.

In a slightly different view, we can define *significant* as the frequency at which the harmonic components of the real waveform begin to drop off faster than 1/f. The frequency at which this happens is sometimes referred to as the *knee frequency*. The harmonic amplitudes of an ideal square wave will initially drop off similarly as 1/f. The frequency at which the harmonic amplitudes of a real waveform begin to significantly deviate from an ideal square wave is the knee frequency.

To evaluate the bandwidth of a time-domain waveform, we are really asking what is the highest frequency component that is just barely above 70% of the same harmonic of an equivalent ideal square wave. When the harmonic amplitudes of the real waveform are significantly less than an ideal square wave's, these lower amplitude harmonics will not contribute significantly to decreasing the rise time and we can ignore them.

For example, we can compare, in the time-domain waveform, two clock waves with a repeat frequency of 1 GHz: an ideal square wave and an ideal trapezoidal waveform, which is a non-ideal square wave with a long rise time. In this example, the 10–90 rise time is about 0.08 nsec, which is a rise time of about 8% of the period, typical of many clock waveforms. These two waveforms are shown in Figure 2-12.

If we compare the frequency components of these waveforms, at what frequency will the trapezoid's spectrum start to differ significantly from the ideal square wave's? We would expect the trapezoid's higher frequency components to begin to become insignificant at about 0.35/0.08 nsec = about 5 GHz. This is the fifth harmonic. After all, we could create a non-ideal square wave with this rise time if we were to take the ideal-square-wave spectrum and drop all components above the fifth harmonic, as we saw earlier.

When we look at the actual spectrum of the trapezoid compared to the square wave, we see that the first and third harmonics are about the same for each. The trapezoid's fifth harmonic is about 70% of the square wave's, which is still a large fraction. However, the trapezoid's seventh harmonic is only about 30% of the ideal square wave's. This is illustrated in Figure 2-12.

We would conclude, by simply looking at the spectra of the trapezoid, that harmonics above the fifth harmonic (i.e., the seventh and beyond) are contributing only a very small fraction of the amount of voltage as in the ideal square wave. Thus, their ability to further affect the rise time is going to be minimal. From the spectrum, we would say that the highest sine-wave-frequency component that is significant in the trapezoid, *compared to that in the ideal square wave*, is the fifth harmonic, which is what our approximation gave us.

There are higher harmonics in the trapezoid's spectrum than the fifth harmonic. However, the largest amplitude is 30% of the square wave's and then only a few percent after this. Their magnitude is such a small fraction of the amplitude of the ideal square wave's that they will contribute very little to the decrease of the rise time and can be ignored.

The bandwidth of any waveform is always the highest sine-wave-frequency component in its spectrum that is comparable in magnitude to a corresponding ideal

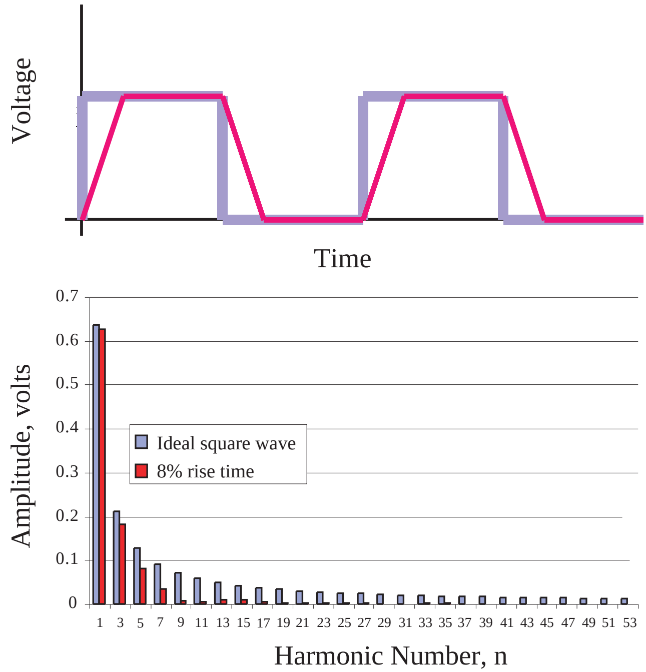

Figure 2-12 Top: Time domain waveforms of 1-GHz repeat frequency: an ideal square wave and an ideal trapezoidal wave with 0.08-nsec rise time. Bottom: Frequency-domain spectra of these waveforms showing the drop-off of the trapezoidal wave's higher harmonics, compared to the square wave's.

square wave. We can find out the bandwidth of any waveform by using a DFT to calculate its spectrum and compare it to an ideal square wave. We identify the frequency component of the waveform that is less than 70% of the ideal square wave, or we can use the rule of thumb developed earlier, that the BW is 0.35/rise time.

TIP It is important to note that this concept of bandwidth is inherently an approximation. It is really a rule of thumb, identifying roughly where the amplitude of frequency components in a real waveform begin to drop off faster than in an ideal square wave.

If you have a problem where it is important to know whether the bandwidth of a waveform is 900 MHz or 950 MHz, you should not use this term *bandwidth*. Rather, you should use the whole spectrum. The entire spectrum is always an accurate representation of the time-domain waveform.

2.12 Bandwidth of Real Signals

Other than the approximation for the bandwidth of a waveform based on its rise time, there is little calculation we can do by hand. Fourier Transforms of arbitrary waveforms can only be done using numerical simulation.

For example, the spectrum of a good-quality, nearly square wave signal has a simple behavior. If a transmission line circuit is poorly terminated, the signal may develop ringing. The resulting spectrum will have peaks at the ringing frequency. The amplitudes of the ringing frequency can be more than a factor of 10 greater than the amplitudes of the signal without ringing. This is shown in Figure 2-13.

The bandwidth of a waveform with ringing is clearly higher than one without. When ringing is present in a waveform, the bandwidth is better approximated

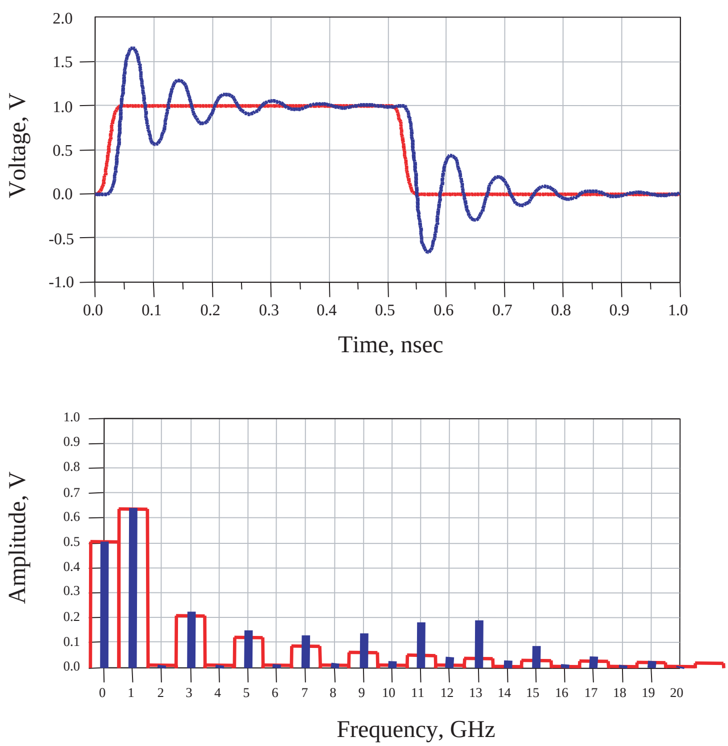

Figure 2-13 Top: The time-domain waveform of a near-square wave and one that has significant ringing due to poor termination. Bottom: the resulting DFT spectrum of these two waves, showing the effect of the ringing on the spectrum. The wide bars are for the ideal waveform while the narrow bars are for the ringing waveform.

by the ringing frequency. Just using the bandwidth to characterize a ringing signal, though, may be misleading. Rather, the whole spectrum needs to be considered.

EMI arises from each frequency component of the currents radiating. For the worst offender, the common currents, the amount of radiated emissions will increase linearly with the frequency. This means that if the current had an ideal-square-wave behavior, though the amplitude of each harmonic drops off at a rate of 1/f, the ability to radiate would increase at the rate of f, so all harmonics contribute equally to EMI. To minimize EMI, the design goal is to use the absolute lowest bandwidth possible in all signals. Above the bandwidth, the harmonic amplitudes drop off faster than 1/f, and would contribute to less radiated emissions. By keeping the bandwidth low, the radiated emissions will be kept to a minimum.

Any ringing in the circuits may increase the amplitudes of higher-frequency components and increase the magnitude of radiated emissions by a factor of 10. This is one reason why solving all signal-integrity problems is usually a starting place to minimize EMI problems.

2.13 Bandwidth and Clock Frequency

As we have seen, bandwidth relates to the rise time of a signal. It is possible to have two different waveforms, with exactly the same clock frequency but different rise times and different bandwidths. Just knowing the clock frequency cannot tell us what the bandwidth is. Figure 2-14 shows four different waveforms, each with exactly the same clock frequency of 1 GHz. However, they have different rise times and hence different bandwidths.

Sometimes, we don't always know the rise time of a signal but need an idea of its bandwidth anyway. Using a simplifying assumption, we can estimate the bandwidth of a clock wave from just its clock frequency. Still, it is important to keep in mind that it is not the clock frequency that determines the bandwidth, it is the rise time. If all we know about the waveform is the clock frequency, we can't know the bandwidth for sure; we can only guess.

To evaluate the bandwidth of a signal from just its clock frequency, we have to make a very important assumption. We need to estimate what a typical rise time might be for a clock wave.

How is the rise time related to the clock period in a real clock waveform? In principle, the only relationship is that the rise time must be less than 50% of the period. Other than this, there is no restriction, and the rise time can be any arbitrary fraction of the period. It could be 25% of the period, as in cases where the

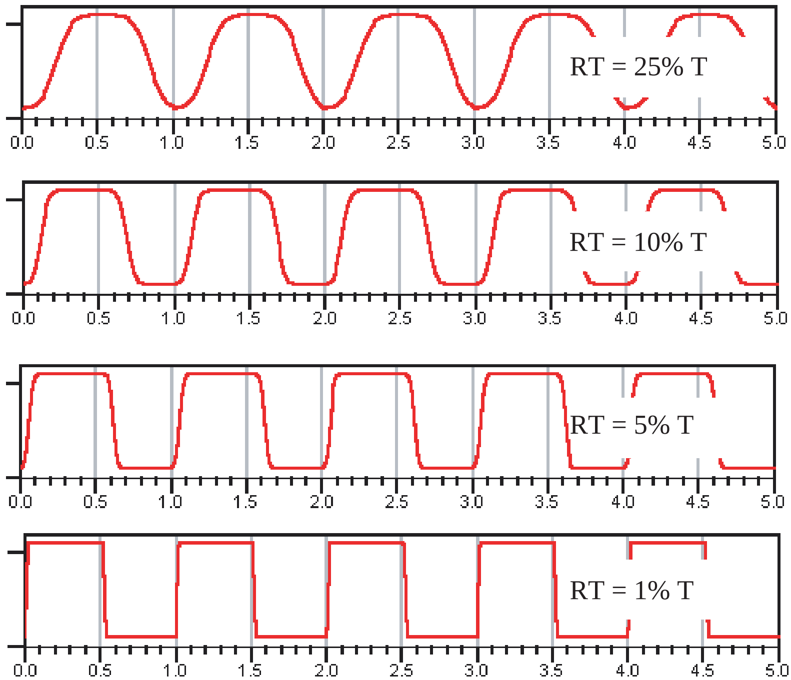

Figure 2-14 Four different waveforms, each with exactly the same 1-GHz clock frequency. Each of them has a different rise time, as a fraction of the period, and hence different bandwidths.

clock frequency is pushing the limits of the device technology, such as in 1-GHz clocks. It could be 10% of the period, which is typical of many microprocessor-based products. It could be 5% of the period, which is found in high-end FPGAs driving external low-clock-frequency memory buses. It could even be 1% if the board-level bus is a legacy system.

If we don't know what fraction of the period the rise time is, a reasonable generalization is that the rise time is 7% of the clock period. This approximates many typical microprocessor-based boards and ASICs driving board-level buses. From this, we can estimate the bandwidth of the clock waveform.

It should be kept in mind that this assumption of the rise time being 7% of the period is a bit aggressive. Most systems are probably closer to 10%, so we are assuming a rise time slightly shorter than might typically be found. Likewise, if we are underestimating the rise time, we will be overestimating the bandwidth, which is safer than underestimating it.

If the rise time is 7% of the period, then the period is 1/0.07 or 15 times the rise time. We have an approximation for the bandwidth as 0.35/rise time. We can relate the clock frequency to the clock period, because they are each the inverse of the other. Replacing the clock period for the clock frequency results in the final relationship; the bandwidth is five times the clock frequency:

$$BW_{clock} = 5 \times F_{clock} \tag{2-5}$$

where:

BW_{clock} = the approximate bandwidth of the clock, in GHz

F_{clock} = the clock repeat frequency, in GHz

For example, if the clock frequency is 100 MHz, the bandwidth of the signal is about 500 MHz. If the clock frequency is 1 GHz, the bandwidth of the signal is about 5 GHz.

This is a generalization and an approximation, based on the assumption that the rise time is 7% of the clock period. Given this assumption, it is a very powerful rule of thumb, which can give an estimate of bandwidth with very little effort. It says that the highest sine-wave-frequency component in a clock wave is typically the fifth harmonic!

It's obvious, but bears repeating, that we always want to use the rise time to evaluate the bandwidth. Unfortunately, we do not always have the luxury of knowing the rise time for a waveform. And yet, we need an answer *now*!

TIP Sometimes getting an OK answer is often more important than getting a BETTER answer LATE.

2.14 Bandwidth of a Measurement

So far, we have been using the term *bandwidth* to refer to signals, or clock waveforms. We have said that the bandwidth is the highest significant sine-wave-frequency component in the waveform's spectrum. And, for signals, we said *significant* was based on comparing the amplitude of the signal's harmonic to the amplitude of an equivalent repeat frequency ideal square wave's.

We also use this term *bandwidth* to refer to other quantities. In particular, it can relate to the bandwidth of a measurement, the bandwidth of a model, and the bandwidth of an interconnect. In each case, it refers to the highest sine-wave-frequency component that is significant, but the definition of significant varies per application.

The bandwidth of a measurement is the highest sine-wave-frequency component that has significant accuracy. When the measurement is done in the frequency domain, using an impedance analyzer or a network analyzer, the bandwidth of the measurement is very easy to determine. It is simply the highest sine-wave frequency in the measurement.

The measured impedance of a decoupling capacitor, from 1 MHz up to 1 GHz, shows that below about 10 MHz, the impedance behaves like an ideal capacitor, but above 10 MHz, it looks like an ideal inductor. Such a measurement is shown in Figure 2-15. There is good, accurate data up to the full range of the network analyzer, in this case, up to 1 GHz. The bandwidth of the measurement is 1 GHz in this example. The measurement bandwidth is not the same as the useful application bandwidth of the device.

When the measuring instrument works in the time domain, such as a time-domain reflectometer (TDR), the bandwidth of the measurement can be found by the rise time of the fastest signal that can be launched into the DUT. After all, this is a rough measure of when the higher-frequency components are small.

In a typical TDR, a fast step edge is created and its change due to interaction with the DUT is measured. A typical rise time entering the DUT is 35 psec to 70

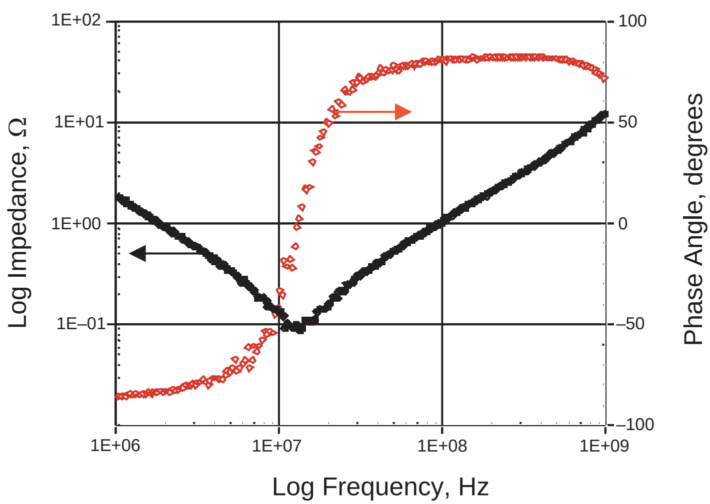

Figure 2-15 Measured impedance of a small 1206 ceramic decoupling capacitor. The measurement bandwidth for this data is 1 GHz.

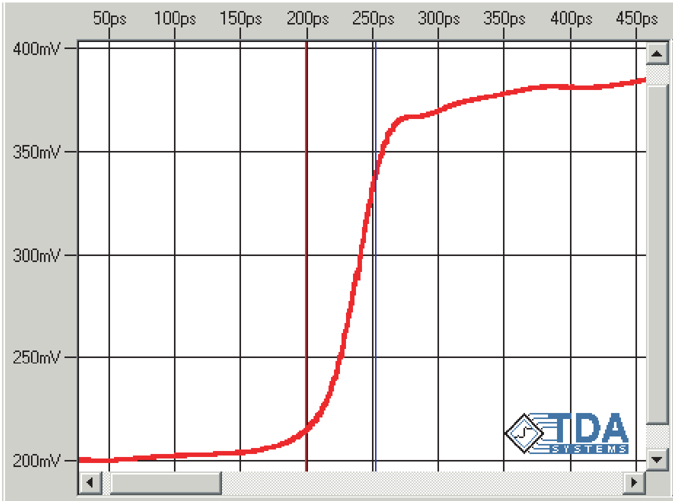

Figure 2-16 Measured TDR profile from the output of a 1-meter cable and microprobe tip, open at the end. The TDR rise time after the cable and probe is about 52 psec. The bandwidth of the measurement is about 0.35/ 52 psec = 7 GHz. The measurement was recorded with TDA Systems IConnect software, using a GigaTest Labs Probe Station.

psec, depending on the probes and cables used. Figure 2-16 shows the measured rise time of a TDR as about 52 psec. The bandwidth of the edge is 0.35/52 psec = 0.007 THz or 7 GHz. This is the bandwidth of the signal coming out of the TDR and is a good first order measure of the bandwidth of the measurement.

In state of the art TDRs, calibration techniques allow the bandwidth of the measurement to exceed the bandwidth of the signal. The bandwidth of the measurement is set by when the signal-to-noise ratio of a frequency component is below a reasonable value, like 10. The bandwidth of the measurement of some TDRs can exceed the signal's bandwidth by a factor of 3-5, making the bandwidth of a TDR's measurement as high as 30 GHz.

2.15 Bandwidth of a Model

TIP When we refer to the bandwidth of a model, we are referring to the highest sine-wave-frequency component where the model will accurately predict the actual behavior of the structure it is representing. There are a few tricks that can be used to determine this, but in general, only a comparison to a measurement will give a confident measure of a model's bandwidth.

The simplest starting equivalent circuit model to represent a wire bond is an inductor. Up to what bandwidth might this be a good model? The only way to really tell is to compare a measurement with the prediction of this model. Of course, it will be different for different wire bonds.

As an example, we take the case of a very long wire bond, 300 mils long, connecting two pads over a return-path plane 10 mils below. This is diagrammed in Figure 2-17. A simple starting circuit model is a single ideal inductor and ideal resistor in series, such as shown in Figure 2-18. The best values for the L and R give a prediction for the impedance that closely matches the measured impedance up to 2 GHz. The bandwidth of this simple model is 2 GHz. This is shown in Figure 2-18.

We could confidently use this simple model to predict performance of this physical structure in applications that had signal bandwidths of 2 GHz. It is surprising that for a wire bond this long, the simplest model, that of a constant ideal inductor and resistor, works so well up to 2 GHz. This is probably higher than the useful bandwidth of the wire bond, but the model is still accurate up to this high a frequency.

Suppose we wanted a model with an even higher bandwidth that would predict the actual impedance of this real wire bond to higher frequency. We might add the effect of the pad capacitance. Building a new model, a second-order model, and finding the best values for the ideal R, L, and C elements result in a simulated impedance that matches the actual impedance to almost 4 GHz. This is shown in Figure 2-18.

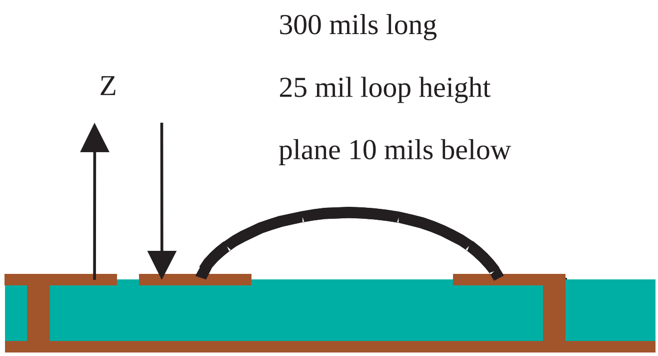

Figure 2-17 Diagram of a wire-bond loop between two pads, with a return path about 10 mils beneath the wire bond.

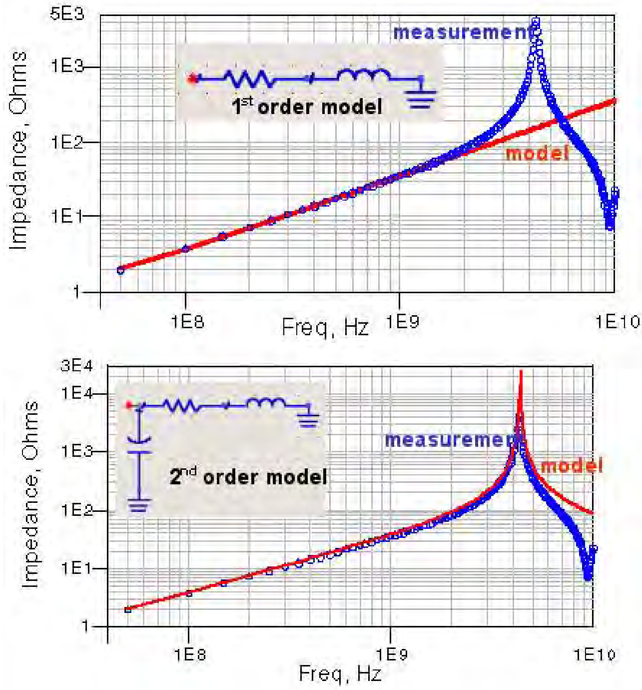

Figure 2-18 Top: Comparison of the measured impedance and the simulation based on the first-order model. The agreement is good up to a bandwidth of about 2 GHz. Bottom: Comparison of the measured impedance and the simulation based on the second-order model. The agreement is good up to a bandwidth of about 4 GHz. The bandwidth of the measurement is 10 GHz, measured with a GigaTest Labs Probe Station.

2.16 Bandwidth of an Interconnect

The bandwidth of an interconnect refers to the highest sine-wave-frequency component that can be transmitted by the interconnect without significant loss. What does *significant* mean? In some applications, a transmitted signal that is within 95% of the incident signal is considered too small to be useful. In other cases, a transmitted signal that is less than 10% of the incident signal is considered usable. In long-distance cable-TV systems, the receivers can use signals that have only 1% of the original power. Obviously, the notion of how much transmitted signal is significant is very dependent on the application and the particular specification. In reality, the bandwidth of an interconnect is the highest sine-wave frequency at which the interconnect still meets the performance specification for the application.

TIP In practice, *significant* means when the transmitted frequency-component amplitude is reduced by –3 dB, which means that its amplitude is reduced to 70% of the incident value. This is often referred to as the 3-dB bandwidth of an interconnect.

The bandwidth of an interconnect can be measured in either the time domain or the frequency domain. In general, we have to be careful interpreting the results if the source impedance is different than the characteristic impedance of the line, due to the complication of multiple reflections.

Measuring the bandwidth of an interconnect in the frequency domain is very straightforward. A network analyzer is used to generate sine waves of various frequencies. It injects the sine waves in the front of the interconnect and measures how much of each sine wave comes out at the far end. It is basically measuring the transfer function of the interconnect, and the interconnect is acting like a filter. This is also sometimes referred to as the *insertion loss* of the interconnect. The interpretation is simple when the interconnect is 50 Ohms, matched to the network analyzer's impedance.

For example, Figure 2-19 shows the measured transmitted amplitude of sine waves through a 4-inch length of a 50-Ohm transmission line in FR4. The

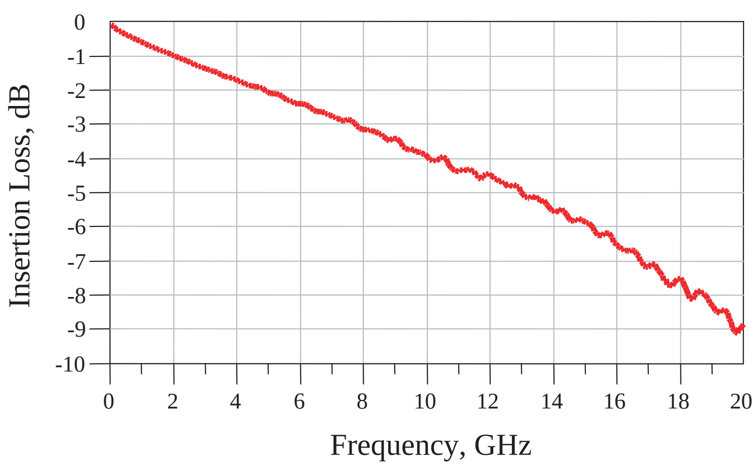

Figure 2-19 Measured transmitted amplitude of different sine-wave signals through a 4-inch-long transmission line made in FR4. The 3 dB bandwidth is seen to be about 8 GHz for this cross section and material properties. Measured with a GigaTest Labs Probe Station.

measurement bandwidth is 20 GHz in this case. The 3-dB bandwidth of the interconnect is seen to be about 8 GHz. This means that if we send in a sine wave at 8 GHz, at least 70% of the amplitude of the 8-GHz sine wave would appear at the far end. More than likely, if the interconnect bandwidth were 8 GHz, nearly 100% of a 1-GHz sine wave would be transmitted to the far end of the same interconnect.

The interpretation of the bandwidth of an interconnect is the approximation that if an ideal square wave were transmitted through this interconnect, each sine-wave component would be transmitted, with those components lower than 8 GHz having roughly the same amplitude coming out as they did going in. But the amplitude of those components above 8 GHz would be reduced to insignificance.

A signal that might have a rise time of 1 psec going into the interconnect would have a rise time of 0.35/8 GHz = 0.043 nsec or 43 psec when it came out. The interconnect will degrade the rise time.

TIP The bandwidth of the interconnect is a direct measure of the minimum rise-time signal an interconnect can transmit.

If the bandwidth of an interconnect is 1 GHz, the fastest edge it can transmit is 350 psec. This is sometimes referred to as its intrinsic rise time. If a signal with a 350-psec edge enters the interconnect, what will be the rise time coming out? This is a subtle question. The rise time exiting the interconnect can be approximated by:

$$RT^2_{out} = RT^2_{in} + RT^2_{interconnect} \tag{2-6}$$

where:

RT_{out} = the 10–90 rise time of the output signal

RT_{in} = the 10–90 rise time of the input signal

$RT_{interconnect}$ = the intrinsic 10–90 rise time of the interconnect

This assumes that both the incident spectra and the response of the interconnect correspond to a Gaussian-shaped rise time.

For example, in the case of this 4-inch-long interconnect, if a signal with a rise time of 50 psec were input, the rise time of the transmitted signal would be:

$$\mathrm{sqrt}(50\ \mathrm{psec}^2 + 43\ \mathrm{psec}^2) = 67\ \mathrm{psec}. \qquad (2\text{-}7)$$

This is an increase of about 17 psec in the rise time of the transmitted waveform compared to the incident rise time.

In Figure 2-20, we show the measured time-domain response of the same 4-inch-long, 50-Ohm interconnect that was measured in the frequency domain above. The input waveform has been time shifted to lie directly at the start of the measured output waveform.

The rise time of the waveform going into the PCB trace is 50 psec. The measured 10–90 rise time of the output waveform is about 80 psec. However, this is somewhat distorted by the long roll to stabilize at the top, characteristic of the behavior of lossy lines. The extra delay at the 70% point is about 15 psec, which is very close to what our approximation above predicted.

If a 1-nsec rise-time signal enters an interconnect with an intrinsic rise time of 0.1 nsec, the rise time of the signal transmitted would be about sqrt(1 nsec^2 + 0.1 nsec^2), or 1.005 nsec, which is still basically 1 nsec. The interconnect would not affect the rise time. However, if the interconnect intrinsic rise time were 0.5 nsec, the output rise time would be 1.1 nsec, and would start to have a significant impact.

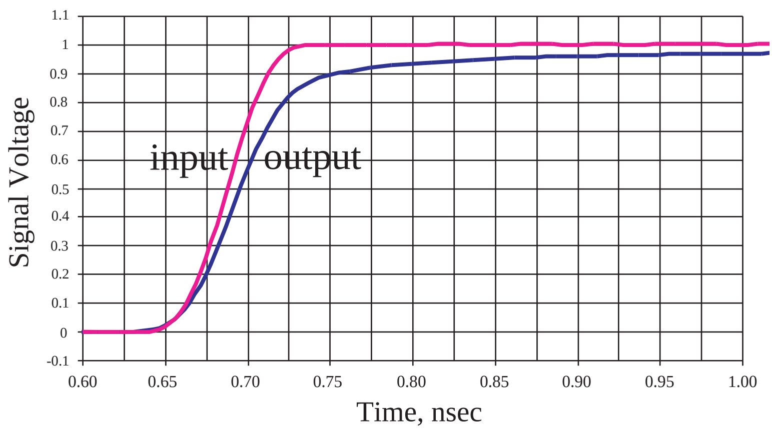

Figure 2-20 Measured input and transmitted signal through a 4-inch long, 50-Ohm transmission line in FR4 showing the rise-time degradation. The input rise time is 50 psec. The predicted output rise time is 67 psec based on the measured bandwidth of the interconnect. Measured with a GigaTest Labs Probe Station.

TIP As a simple rule of thumb, in order for the rise time of the signal to be increased by the interconnect less than 10%, the intrinsic rise time of the interconnect should be shorter than 50% of the rise time of the signal.

TIP In the frequency-domain perspective, to support the transmission of a 1-GHz bandwidth signal, we want the bandwidth of the interconnect to be at least twice as high, or 2 GHz.

It is important to keep in mind that this is a rule of thumb and it should not be used for design sign-off. It should be used only for a rough estimate or to identify a goal. If the bandwidth of an interconnect is within a factor of two of the bandwidth of the signal, it would probably be important to perform an analysis of how the interconnect affected the entire signal's spectrum.

2.17 The Bottom Line

1. The time domain is the real world and is typically where high-speed digital performance is measured.
2. The frequency domain is a mathematical construct where very specific, specialized rules apply.
3. The only reason to ever leave the time domain and use the frequency domain is to get to an answer faster.
4. The rise time of a digital signal is commonly measured from 10% of the final value to 90% of the final value.
5. Sine waves are the only waveform that can exist in the frequency domain.
6. The Fourier Transform converts a time-domain waveform into its spectrum of sine-wave-frequency components.
7. The spectrum of an ideal square wave has amplitudes that drop off at a rate of 1/f.
8. If the higher-frequency components are removed in the square wave, the rise time will increase.
9. The bandwidth of a signal is the highest sine-wave-frequency component that is significant, compared to the same harmonics in an ideal square wave with the same repeat frequency.

10. A good rule of thumb is that the bandwidth of a signal is 0.35/rise time of the signal.
11. Anything that decreases the bandwidth of a signal will increase its rise time.
12. The bandwidth of a measurement is the highest sine-wave frequency where the measurement has good accuracy.
13. The bandwidth of a model is the highest sine-wave frequency where the predictions of the model give good agreement with the actual performance of the interconnect.
14. The bandwidth of an interconnect is the highest sine-wave frequency where the performance of the interconnect still meets specifications.
15. The 3-dB bandwidth of an interconnect is the highest sine-wave frequency where the attenuation of a signal is less than –3 dB.