

DesignCon 2012

The Relationship Between Discrete-Frequency S-parameters and Continuous-Frequency Responses

Peter J. Pupalaikis, LeCroy Corporation

PeterP@LeCroy.com

Abstract

We explore in detail the relationship between discrete-frequency responses connected with s-parameters and the implied continuous time response. This is done in both the frequency- and time-domain to develop the proper insight and explore issues with real time-domain responses, time-aliasing, causality, interpolation and re-sampling of discrete-frequency data. Using the insight gained, we identify the conditions for sufficiency of sampling and the side-effects of the invariable practical conditions when these side-effects cannot be completely dispelled.

This paper is especially useful for understanding issues involved in direct application of s-parameters in linear simulations, like those used in virtual probing applications in scopes.

Author(s) Biography

Pete Pupalaikis was born in Boston, Massachusetts in 1964 and received the B.S. degree in electrical engineering from Rutgers University, New Brunswick, New Jersey in 1988.

He has worked at LeCroy Corporation for 16 years and currently manages integrated-circuit development and intellectual property as Vice President of Technology Development. Prior to LeCroy, he served in the United States Army and as a consultant in embedded systems design.

Mr. Pupalaikis has numerous patents in the area of measurement instrument design and is a member of Tau Beta Pi, Eta Kappa Nu and the IEEE signal processing, instrumentation, and microwave societies.




Figure 1 - Virtual Probe Simulation Environment




Figure 2 - RC Network in Buffered Arrangement

Introduction

A common and exemplary situation in a simulation is to determine the effect at a load of a circuit driven at its source by a known waveform. One way to do this is to define a source waveform under a normal set of conditions and to determine a transfer function that is a ratio of the voltage at the load under conditions containing the circuit to the voltage at the source under more normal and deterministic conditions.

A virtual probing¹ simulation environment is shown in Figure 1 where a two port network is driven with a source impedance Z_s and a load impedance Z_l . The relationship between the voltage at the load with the two-port network in the system and the the voltage at the source and load without the two-port network (i.e. source and load directly connected) can be shown to be (1):

$$H = \frac{V_{lo}}{V_{sm}} = \frac{S_{21}(1 - \Gamma_l \Gamma_s)}{1 - \Gamma_s S_{11} - \Gamma_l S_{22} + \Gamma_s \Gamma_l |S|} \quad (1)$$

In (1), the reflection coefficients are related to the impedance Z by an arbitrary reference impedance Z_0 , usually 50Ω as (2):²

$$\Gamma = \frac{Z - Z_0}{Z + Z_0} \quad (2)$$

¹See P. J. Pupalaikis et al., "Virtual Probing", US Patent 7,660,685, filed Aug 2007

²Note that if the source and load impedance in (1) is matched to the reference impedance Z_0 , then (1) becomes $H = S_{21}$. This is the source of the simplification that the transfer function equals S_{21} .

Setting $\Gamma_l=1$ and $\Gamma_s=-1$ in (1) to correspond to a low-impedance source and a high-impedance load, we obtain (3):

$$H = \frac{2S_{21}}{1 + S_{11} - S_{22} - |\mathbf{S}|} \quad (3)$$

Considering (3), here we have a transfer function that provides for a two-port circuit or network situated between two buffers.³ The transfer function is providing the frequency response of the system that supplies the relationship to the waveform at a receiver buffer to a waveform that would be present when the source buffer is terminated in high-impedance. Since the network is described by s-parameters, we know that we are provided a finite number of two-by-two matrices containing complex numbers corresponding to a number of chosen frequency points. The transfer function as shown is therefore valid for the same discrete frequencies as the s-parameters and in this case is a single complex response per frequency point. Since the system is linear, it is reasonable to convert this complex frequency response into an impulse response and to convolve this impulse response with the time-domain waveform representing the output from the source to obtain the waveform at the receive side. One reasonable way to do this is to use inverse discrete Fourier transform (IDFT) methods.

Consider, for example purposes, that the two-port network is actually a series resistance followed by a shunt capacitance as shown in Figure 2. The s-parameters of a two port network containing a series impedance Z_r connected to a shunt impedance to ground Z_h is shown in (4):

$$\mathbf{S} = \frac{1}{Z_r Z_h + 2 Z_h Z_0 + Z_r Z_0 + Z_0^2} \begin{pmatrix} Z_r Z_h + Z_s Z_0 - Z_0^2 & 2 Z_h Z_0 \\ 2 Z_h Z_0 & Z_r Z_h - Z_r Z_0 - Z_0^2 \end{pmatrix} \quad (4)$$

Substituting (4) into (3) we obtain (5):

$$H = \frac{Z_h}{Z_r + Z_h} \quad (5)$$

For a series resistance $Z_r = R$ and a shunt capacitance $Z_h = \frac{1}{Cs}$ and a time-constant defined as $\tau = RC$, we have (6), as expected.

$$H = \frac{\frac{1}{\tau}}{s + \frac{1}{\tau}} \quad (6)$$

Here, $s = j 2 \pi f$ is the Laplace variable. The concept here is that (1) and (3) present discrete-frequency transfer functions of s-parameters and (6) and (5) are related continuous-frequency

³A buffer is considered as a device with unity forward gain, infinite input impedance and zero output impedance and perfect reverse isolation.

functions. Since (3) and (1) involve only s-parameters, the concept of an RC network is lost. This is one of the features and drawbacks of the direct use of s-parameters.

To summarize:

1. s-parameter based transfer functions involve discrete frequencies.
2. In s-parameter based transfer functions, the underlying topology of a network is lost.

These are both features and drawbacks of the use of s-parameters because while the circuit topology is lost, it allows for the development of transfer functions that are independent of topology. But it is assumed that the discrete-frequency nature of the system still allows for proper simulation of systems *as if* the transfer function was continuous. Engineers rely on this being true otherwise there would be no point in simulations in this manner. It will be shown that discrete frequency s-parameters can provide valid simulations as long as certain criteria are met to a large extent. When these criteria are not met, invalid and often nonsensical results come from simulation. This paper is designed to help the user of s-parameters in simulation avoid certain pitfalls.

It should be pointed out that all of the problems endemic to the direct usage of s-parameters can be avoided by fitting models to a system. This means that given a knowledge of the circuit topology, one can fit the values of the circuit elements such that the resulting model matches the s-parameters provided. In this example, if we knew or assumed that the circuit consisted of an RC network, we could utilize (4) and calculate the partial derivatives with respect to each presumed circuit element, namely the series and shunt impedance. Better yet, and simpler, would be to choose a value of τ that makes the transfer function in (6) match the transfer function of (3) based on s-parameters provided.

There are many problems and benefits in fitting of models to s-parameters. This paper does not deal with that topic and deals only with the direct usage of s-parameters in simulation.

While an RC network is not necessarily the most practical s-parameter usage example, we will use an RC network for all of the examples in this paper because the expected characteristics of the network are well known and because this type of network has certain characteristics that highlight the type of problems that can crop up.

The Relationship Between Discrete-Frequency and Continuous-Frequency Responses

The relationship between a continuous-frequency response $H(f)$ and it's corresponding discrete-frequency response $\hat{H}(f)$ can be shown to be (7):

$$\hat{H}(f) = \left(H(f) \cdot \Pi(fT) \cdot \Delta_{\frac{1}{KT}}(f) \right) * \Delta_{\frac{1}{T}}(f) \quad (7)$$

Equation (7) shows a very complicated relationship between the continuous- and discrete-frequency responses and introduces certain notation. While stepping through this equation, we will define various intermediate responses that will help our understanding.

The first step on the left inside the parentheses is to multiply the continuous-frequency response $H(f)$ by the rectangle, or boxcar function $\Pi(fT)$. The boxcar function is defined with its corresponding inverse Fourier transform as (8):

$$\Pi(f) = \text{rect}(f) = \begin{cases} 0 & \text{if } |f| > \frac{1}{2} \\ \frac{1}{2} & \text{if } |f| = \frac{1}{2} \\ 1 & \text{if } |f| < \frac{1}{2} \end{cases} \quad F^{-1}\{\Pi(f)\} = \text{sinc}(\pi t) = \frac{\sin(\pi t)}{\pi t} \quad (8)$$

The boxcar function limits the bandwidth of $H(f)$ to $1/2T$ where T is defined as the sample period⁴. Therefore we can define an intermediate band-limited response as (9):

$$H_{bwl}(f) = H(f) \cdot \Pi(fT) \quad (9)$$

Substituting (9) in (7) allows us to write (10):

$$\hat{H}(f) = \left(H_{bwl}(f) \cdot \Delta_{\frac{1}{KT}}(f) \right) * \Delta_{\frac{1}{T}}(f) \quad (10)$$

The next step is to multiply the band-limited response $H_{bwl}(f)$ by the sampling function $\Delta_{\frac{1}{KT}}(f)$ where the sampling function is defined for both the time- and frequency-domain as (11) and (12) and K is defined as the integer impulse-response sample length:

$$\Delta_T(t) = \sum_{k=-\infty}^{+\infty} \delta(t - kT) = F^{-1}\left\{ \Delta_{\frac{1}{T}}(f) \right\} \quad (11)$$

$$\Delta_{\frac{1}{T}}(f) = \frac{1}{T} \sum_{n=-\infty}^{+\infty} \delta\left(f - \frac{n}{T}\right) = F\left\{ \Delta_T(t) \right\} \quad (12)$$

The application of the sampling function effectively sampled the band-limited response such that we are now left with discrete frequencies allowing the definition in (13):

$$H_{disc}(f) = H_{bwl}(f) \cdot \Delta_{\frac{1}{KT}}(f) \quad (13)$$

Substituting (13) in (10) allows us to write (14):

$$\hat{H}(f) = H_{disc}(f) * \Delta_{\frac{1}{T}}(f) \quad (14)$$

⁴Note that the boxcar function limits both the positive and negative frequencies. The continuous-frequency response is assumed to have both positive and negative frequency content.




Figure 3 - Continuous- and Discrete-Frequency Response Relationship

The final step, therefore is to convolve the discrete-frequency response $H_{disc}(f)$ with the sampling function $\Delta_{\frac{1}{T}}(f)$. Convolution with the sampling function causes the response to repeat over and over again with a spacing of $1/T$.

The frequency-domain implications of (7) are shown graphically in Figure 3. In Figure 3, we see the conversion process from a continuous-frequency function to a repeating discrete-frequency function.

Now let's examine each of these steps in detail from the time-domain perspective.

In the first step, we band-limited the response using a boxcar function. This boxcar function limited the response to $\pm 1/2T$ in frequency content. When the discrete-frequency response that arises from s-parameter models is used, there is an implicit sample period T and corresponding

sample rate $1/T$ associated with the response which is roughly⁵ twice the final frequency point. As shown in (8), the boxcar function has the sinc function as its inverse Fourier transform. Since we multiplied the response by the boxcar function in the frequency domain, this is the same as convolving the impulse response with the sinc function in the time-domain. This band-limiting step is the cause of most causality issues associated with s-parameter based simulations although it will not cause tremendous issues unless you choose to simulate a sample rate that is different than that specified by your s-parameters⁶. Obviously, the band-limiting step will not cause any problems if the response itself is actually band-limited already.

Let's now define some time-domain responses:

$$h(t) = F^{-1}\{H(f)\} \quad (15)$$

and:

$$h_{bwl}(t) = F^{-1}\{H_{bwl}(f)\} = h(t) * \text{sinc}\left(\pi \frac{t}{T}\right) \quad (16)$$

In the first step, we saw that the effect of limiting the bandwidth of the continuous-frequency response which is the same as choosing the highest frequency point in the s-parameter data is a first point of divergence from good s-parameter based models. The second step also has a large and often ill-understood effect. We saw that the second step is to frequency sample the band-limited response with a frequency spacing of $1/KT$ by multiplying $H_{bwl}(f)$ by the sampling function $\Delta_{\frac{1}{KT}}(f)$ to produce $H_{disc}(f)$. K is the total number of frequency points that will be present after sampling – but it contains both positive and negative frequencies of the continuous-frequency response. Multiplication in the frequency domain is the same as convolution in the time domain so we are effectively convolving the band-limited response with a train of pulses:

$$h_{disc}(t) = F^{-1}\{H_{disc}(f)\} = h_{bwl}(t) * \Delta_{KT}(t) = h_{bwl}(t) * \sum_{k=-\infty}^{+\infty} \delta(t - kKT) \quad (17)$$

Examine how the last part of (17) is written and realize that $h_{disc}(t)$ is periodic and is an infinite sum of delayed and advanced values of $h_{bwl}(t)$ that are separated by a time KT as shown in Figure 4. We call $h_{disc}(t)$ the *impulse train response*. It is important to understand that periodicity in one domain is always associated with discreteness in another domain. It is the frequency sampling (i.e. the discreteness in frequency) that caused the time-domain periodicity.

The overlapping responses are another source of problems in s-parameter based simulation models. We see that the frequency spacing of $1/KT$ as shown in Figure 3 and provided in (44) determines the time spacing KT in the impulse train that produces the impulse train response. If

⁵Strictly speaking, the sample rate might not be exactly twice the final frequency point depending on whether K is chosen as even or odd.

⁶While there are certainly other possibilities for causality issues, the effect outlined here is specifically caused by using s-parameters because of their band-limited nature.




Figure 4 - The Impulse Train Response

KT is chosen to be too short a time duration (or similarly if $1/KT$ is chosen as too long a frequency duration), the impulse responses interfere with each other significantly. This effect is called *time-aliasing*. Time-aliasing is caused by sampling the frequency response, just as the more familiar effect of frequency-aliasing is caused by sampling time-domain waveforms.

The final step was to convolve the discrete-frequency response $H_{disc}(f)$ with the sampling function $\Delta_{\frac{1}{T}}(f)$ which caused the response to repeat over and over again in frequency with a spacing of $1/T$. Again looking at the time-domain, this is the same as multiplying the time-domain response $h_{disc}(t)$ with the sampling function $\Delta_T(t)$:

$$\hat{h}(t) = F^{-1}[\hat{H}(f)] = h_{disc}(t) \cdot \Delta_T(t) \quad (18)$$

From the time-domain perspective, this final step serves only to sample the continuous-time response $h_{disc}(t)$ to make it into discrete elements.⁷ What is interesting about this effect is the effect already mentioned in the frequency-domain – that of causing the frequency response to repeat over and over with a period $1/T$. If the response were not already band-limited to $1/2T$, frequency aliasing would occur because the overlapping responses would be summed together in

⁷It's interesting that while $H_{disc}(f)$ was named for being discrete, it is discrete in the frequency-domain only. $h_{disc}(t)$ is not discrete.

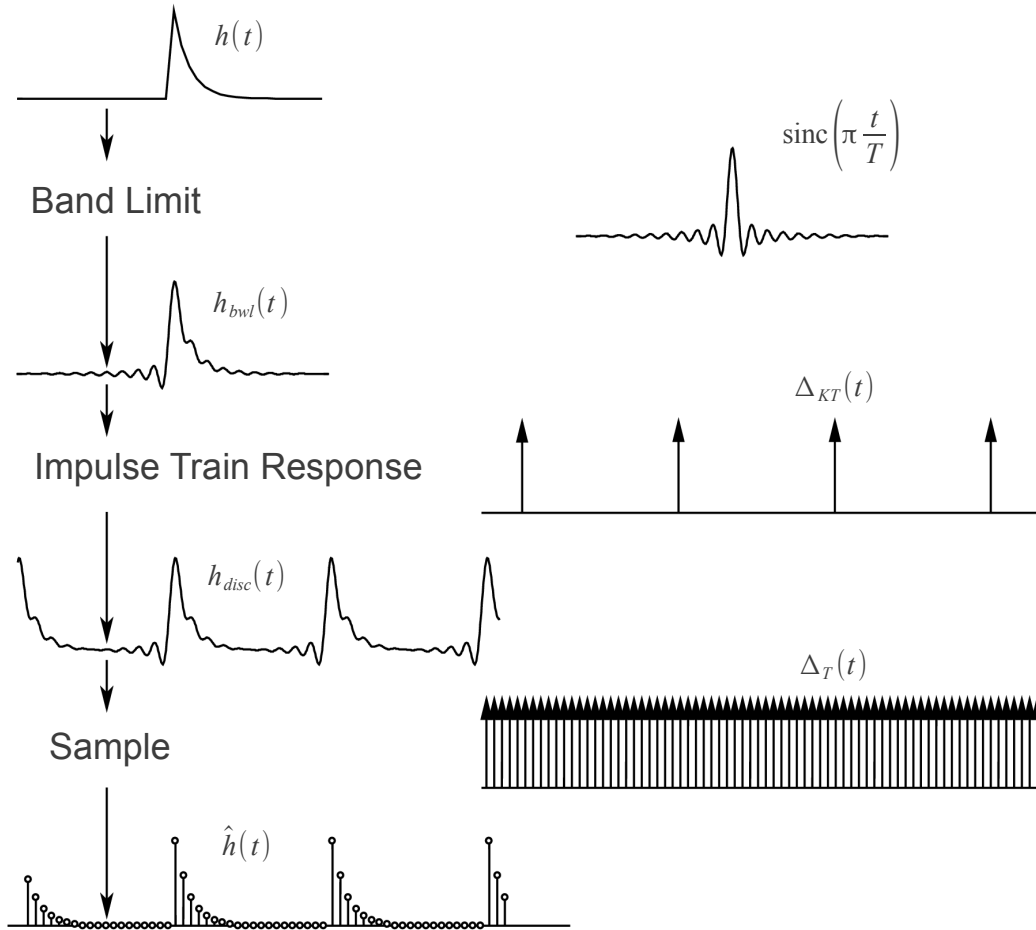


Figure 5 - Continuous- and Discrete-Frequency Response Time-Domain Relationships

the frequency-domain this time. But the band-limited nature causes only repetition. As it turns out, although the response $\hat{h}(t)$ is a sampled response, there is no possibility for frequency aliasing as the band-limit frequency is always the same as pulse-train period in the frequency domain. This makes sense because intuitively it is impossible to cause frequency aliasing by taking frequency samples, just as it is impossible to cause time aliasing by sampling in the time-domain. Sampling in one domain only has the possibility of aliasing in the other domain.

The time-domain implications of (7) are shown graphically in Figure 5.

In summary, in this section we have provided an equation (7) that provides both the frequency and time-domain effects in the process of equating a continuous-frequency with its discrete-frequency counterpart. During the examination of the steps in this process, we've identified two specific areas of concern. The first area of concern is the band-limiting of the response. The second area of concern is the possibility for time-aliasing. One area not of concern but

highlighted here is the implicit repetitive nature of both the time- and frequency-domain response associated with the discrete-time and discrete-frequency responses.

The band-limiting of the response, although seemingly obvious, requires some discussion because it creates some nonintuitive results. We see by examining Figure 5 that the band-limited response of the RC network looks quite different from the continuous response as it exhibits the familiar Gibbs phenomena with its pre-shoot and ringing. But note that this pre-shoot and ringing will not be present in the final response $\hat{h}(t)$ to a large degree. This seems strange until you realize that the band-limit frequency is exactly one-half the reciprocal of the sample period. In other words, if we sample the sinc function used to band-limit response which is $\text{sinc}(\pi t/T)$ with the sampling function used to sample the response which is $\Delta_T(t)$ which is a train of sample pulses separated by time T , we are sampling with a sample time $t = kT$ and therefore we have $\text{sinc}(\pi k)$ which is unity at time zero and zero elsewhere. There will still be slight errors at the edges due to the phase, which we will talk about later.




Figure 6 - Sinc Function and Sampled Points

We have chosen here to associate the sample period of the impulse response with the bandwidth limit associated with the frequency response but this association does not hold when the discrete-frequency response is re-sampled onto a different frequency grid for simulation at sample rates other than that associated with the discrete-frequency response. When the discrete-frequency is re-sampled, it's analogous to sampling the sinc function at a different rate which exposes the Gibbs phenomena.

The second area of concern highlighted is the possibility for time-aliasing. We have shown that this effect causes the continuous-time impulse responses to overlap.

Both of these areas of concern cannot usually be totally mitigated. The band-limiting effect is mitigated by choosing the highest frequency for the discrete-frequency response as high as possible. Generally, this means choosing it such that the response is significantly attenuated prior to the band-limit response. In some situations this is easy, but when working with s-parameter data, things become more difficult. This is because in high-frequency systems, while the thru-response (e.g. S_{21} of a two-port) drops with frequency, the return-loss (e.g. S_{11}) generally increases. In (1) and (3), we see a situation that depends on S_{11} and S_{22} . Often, in microwave systems, it is impossible to choose a frequency that completely mitigates the band-limit effects. The best that can be done is to ensure that the highest frequency point is equal or above one-half the maximum sample rate used in simulation, to ensure that the frequency content




Figure 7 - Time Sense of Impulse Response

of any stimulus applied to the system is limited to this highest frequency and to examine transfer functions to ensure that they are not overly sensitive to impedance match at high frequencies.

Mitigating the time-aliasing effect is achieved by taking enough frequency points with a small enough frequency spacing such that the impulse responses that form the impulse train response do not overlap. We can see that this is governed by the choice of K once the sample period T has been chosen. This choice can often be inferred by knowledge of the electrical length of the device. While KT is the impulse response length, it is not the electrical length. The impulse-response length must usually be chosen such that it is significantly longer than the electrical length. To understand this, consider a long cable that is significantly mismatched in impedance with the source and load impedance. Stimulating this cable causes waves to reflect back and forth from the interfaces. Fortunately, because of loss, these reflections eventually die down sufficiently, but in cases with low loss, many multiples of the electrical length must be chosen for KT . In highly lossy systems, especially systems that are well matched at frequencies where loss is low, sometimes four times the electrical length is sufficient. One thing that should be pointed out is that while it appears that the requirement is for the impulse response to die out before KT time expires, the requirement is actually more stringent. This is because the impulse train response causes us to lose absolute sense of time and we must make a choice. We must decide whether the end of the impulse response near the time KT is the end of the impulse-response or the beginning of the next in the impulse train response. This becomes important when the band-limiting effect outlined previously causes causality violations, and the end of the impulse response contains the causality violation in the next impulse-response in the train. To deal with this, we enforce that the second half of the discrete-time impulse-response contains response that occurs before time zero and that the first half of the discrete-time impulse-response contains response that occurs before time zero. This is shown for even K in Figure 7.

If a continuous-frequency system has zero frequency response after $1/T$ and has an impulse response strictly limited to $KT/2$, then the discrete-frequency response, while in discrete form contains all of the information necessary to completely replicate the continuous-frequency response. This statement is analogous to the sufficiency of sampling rules most are familiar with with regard to sampling of time-domain waveforms for no aliasing and perfect reconstruction.

It should be noted here that the relationship provided in (7) holds strictly for evenly spaced frequency and time points. This is because the sampling function applied in the frequency-domain is $\Delta_{\frac{1}{KT}}(f)$. This situation has been chosen because, while not required, causes the discrete-frequency response and the associated discrete-time impulse-response to be related through the discrete Fourier transform. We will show ways later for dealing with unevenly spaced frequency points.

S-Parameters, Responses and the DFT

S-parameters represent a special case of discrete-frequency response. They are represented by a matrix per frequency point whereby each matrix element represents a response at that frequency point at reflected and incident ports corresponding to the rows and columns of each matrix. For a P port device and $N+1$ frequency points, $n \in 0 \dots N$, each matrix element $(S_n)_{r,c}$, $r, c \in 1 \dots P$ represents the reflected wave from port r when an incident wave is driven at port c at frequency point $f_n = n \Delta f$ while the driven port is driven through and all other ports are terminated in the reference impedance Z_0 . The incident waveform is assumed to be a cosine wave of unity amplitude and zero phase and each matrix element is a complex number where the magnitude represents the amplitude and the argument represents the phase of the reflected cosine wave. Each matrix element is the response to the incident wave. Therefore, each s-parameter element represents a response. For subsequent discussion, we will refer to a response element as H_n corresponding to the frequency f_n . When dealing with s-parameters, all of the rules associated with a given matrix element hold over frequency. For example, the response at port r for an incident wave at port c under the conditions previously outlined is defined as $H_n = (S_n)_{r,c}$ allowing us to no longer talk about s-parameters specifically but instead discrete-frequency responses where it is understood that there are P^2 responses associated with the s-parameters of a P port device.

S-parameters and all responses for positive frequencies have an implied negative frequency relationship for real time-domain responses. While the details are shown in Appendix A, basically for each positive frequency, there is an implied negative frequency which is the complex conjugate of the positive frequency response. In other words, for each response element H_n corresponding to a positive frequency f_n , there is an implied element equal to H_n^* at frequency $-f_n$. We will not prove this here, but this is a condition required for real corresponding discrete-time impulse responses h_k .

Also, since we will be using discrete Fourier transform (DFT) and IDFT methods to go back and forth between the frequency- and time-domain, there are some additional requirements on the mapping of the discrete frequency-response to a DFT response. The notation \mathbf{H} is used for a

vector of responses at discrete-frequencies and the notation \mathbf{X} is used to refer to DFT elements, which are slightly different. We will refer to \mathbf{h} as the vector of impulse-response points and the notation \mathbf{x} as the IDFT of \mathbf{X} . We will use either notation depending on whether we're talking about impulse responses or the IDFT.

With this said, before we talk about the two main choices of last frequency point and frequency spacing, we should recognize that the specifics of both of these choices insofar as DFT methods are concerned are intricately linked with the concept of sample rate and impulse-response length from a digital simulation standpoint. Instead of providing pages of explanation on this relationship, a concise set of equations are provided in *Appendix A - Relationship Between Response Elements and DFT Elements*. Here all of the details are provided which are quite wrapped up in choices of whether the final impulse-response length is chosen even or odd. It is an arbitrary choice because for $N+1$ frequency points, we can always choose $K=2N$ for an even K or $K=2N+1$ for an odd K . Since it is undesirable to carry all of the details associated with even or odd K forward through this paper, even K is chosen to illustrate all subsequent points.

The first choice needed in dealing with a discrete-frequency response is to first choose the upper frequency point. The nature of the choice is to choose the upper frequency high enough such that the response is attenuated enough at and beyond the upper frequency for the band-limiting effect to be as insignificant as possible. Unfortunately, the choice of upper frequency insofar as s-parameter measurements are concerned are often dictated by the limitation of a measurement instrument when s-parameter measurements are concerned. Another inconvenient fact regarding high-frequency systems is that while the off-diagonal responses (the response from one port to another) generally drop in magnitude as frequencies increase due to loss in the device, the responses on the diagonal which determine the match of the port to the reference impedance generally increase with frequency. In many cases, there is no real limit to how high of a response you need and as you can see in systems like described in (1) and (3), the response depends on these diagonal elements (S_{11} and S_{22} in the equations referenced). This can be a real problem. The RC network example is really problematic (which is why it was chosen) because it can be shown that the upper frequency required for the response in (6) to be attenuated by a reasonable amount like -40 dB is $16/\tau$ meaning that for a response with a time-constant of 100 ps, a 160 GHz upper frequency is theoretically required!

For even K , this upper frequency point f_N is exactly half the sample rate used for digital simulations, so often the exact choice of sample rate factors into the determination of f_N as well. In other words, we choose a sample rate for the simulation such that half the sample rate meets the attenuation criteria already stated.

Next, K is actually chosen in conjunction with the sample period T such that the impulse response \mathbf{h} is settled to near zero at and after time $KT/2$. Note that in the case of the RC network, five time constants is adequate to settle to less than 1/100th the initial voltage. If you are provided with s-parameter data that has not been sufficiently sampled frequency-wise, there is no good method of recovery. A good way to check the data is to verify whether more s-parameter frequencies are needed is to verify that as the frequency spacing is changed that the impulse response features all move in the right direction. For example, if one were to increase

KT and find that features in the impulse response move to the left, this is a sure sign that time-aliasing is occurring. The goal is to produce data whose impulse train response does not exhibit time-aliasing and whose response is settled to zero or near zero in the latter half of the response. The reason why we want this is because of the time relationships offered in (37) and we will want to be able to distinguish clearly between time before and after the impulse response to deal with possible causality issues. In general, we say that the the latter half of the time domain response $h_{disc}(t)$ contains the response before time zero and that the first half of the response contains the response after time zero.

One note here is that engineers are often asking how many frequency points should they take for the s-parameters. Since the end frequency is usually set, they are really asking about the frequency spacing. Often the choice is made surrounding a notion of frequency resolution needed to resolve various nooks and crannies of the frequency response but here we see that a more compelling method is to base the choice on the time-domain! Note also that our example here involved a simple RC network – not something that would seem to require large amounts of frequency resolution. Clearly we can see that the value of KT should be chosen large enough to mostly isolate the impulse responses.

Real Time-domain Responses and Causality Issues

If we examine carefully the information provided in Appendix A used to map DFT based impulse-responses and frequency-responses, we notice a troubling aspect. This aspect is the enforcement of a real final frequency response element X_N in (38). While this is enforced only for even K , it does not change the underlying problem which will be described in this section.

It can be shown that the enforcement of $X_N \in \mathbb{R}$ is a requirement for the IDFT to produce a real time-domain response. It arises from pairing the positive and negative frequencies to produce cosine waves. As the rules show in Appendix A, for a positive frequency with response $X_{N-\sigma} = A_n e^{j\theta_n}$ is paired with a negative frequency whose response is the complex conjugate $X_{N+\sigma} = X_n^* = A_n e^{-j\theta_n}$ such that the response at a given frequency is (19):

$$H_{N-\sigma} = A_n \left(\frac{e^{j\theta_n} + e^{-j\theta_n}}{2} \right) \quad (19)$$

Leads to a time-domain response as in (20):

$$h_k = A_n \cos(2\pi f_n t_k + \theta_n) \quad (20)$$

Of course, the cosine waves are all summed together for each n for form the total time-domain impulse-response. For K even, the last point is not paired, however, and the response $H_N = A_N e^{j\theta_N}$ must be real for this summation to lead to real values of h .

This seems like a problem only for even K , but in reality, what this requirement really wants is for the phase to be a continuous function, but strictly speaking, the continuity of phase is only important when the magnitude response is non-zero near the discontinuity. In other words, the smaller the magnitude response, the less important the phase discontinuity.




Figure 8 - Magnitude and Phase of an RC Network

To illustrate this, Figure 8 shows the magnitude and phase for an RC network. The negative frequencies have been plotted to the right of half the sample rate to illustrate the discontinuity. The magnitude is continuous at half the sample rate, but derivative discontinuous. The phase is not continuous. As an exercise for the reader, try taking the DFT of samples from the impulse-response of an RC network. You will find that the phase response starts it's downward trend as in Figure 8, but then bends upward towards 0 at the end!

Given the discrete-frequency response H , we plot the discrete-time impulse response by computing $h_k = |H_n| \cos(2\pi f_n t_k + \arg(H_n))$. Note we have side-stepped for the moment the fact that H_N is not real by not trying to compute the IDFT, but rather by simply summing cosinusoids. The result is plotted in Figure 9 where we see a causality issue.

To show where this is coming from, we determine a value Td such that:

$$e^{-j2\pi f_N Td} H_N \in \mathbb{R} \quad Td = \frac{\arg(H_N)}{2\pi f_N} \quad H'_n = e^{-j2\pi f_n Td} H_n \quad (21)$$

The phase response of H' is plotted in Figure 10 where we can see that H_N is indeed real.

Now we compute $h'_k = |H'_n| \cos(2\pi f_n t_k + \arg(H'_n))$ which is plotted in Figure 10 where we see that the causality issue has been largely resolved. There is an implication here that everything




Figure 9 - Causality Issues from Discontinuous Phase




Figure 10 - Phase and Impulse Response of Time Delayed RC Network

can be fixed by correcting the phase so that it reaches zero at the end frequency but in fact it cannot. The phase is very smooth for the example used (i.e. nearly derivative continuous) which will not generally be the case. When the phase is not smooth often no amount of fiddling with phase through introduction of a time-delay can completely resolve the non-causality. Often, this effect can only be minimized.

In this example provided, we advanced the impulse response⁸ to improve causality. This is because the causality issue here was not caused by true causality violation, but rather because the previously non-continuous group delay caused the main impulse point to appear inside of a fractional sample. We could have just as well delayed the waveform to cause this same improvement.

Dealing with causality this way requires us to remember the amount of advance or delay applied so that it can be applied later in simulation. It turns out that this is only pushing the problem upstream. This can be understood by realizing that to put the delay back in, in this case, would require a fractional delay filter and a fractional delay filter is simply a sinc function again. In this case, the frequency- and impulse-response of the fractional delay filter would be (22):

$$H_{frac_n} = e^{-j2\pi f_n Td} \quad h_{frac_k} = \frac{1}{2} \sum_{n=0}^N \cos[2\pi f_n (t_k - Td)] = \text{sinc}[\pi(t_k - Td)] \quad (22)$$

The impulse-response of the filter in (22) is shown for the example in Figure 11 where it is clear that the fractional delay filter simply puts back the causality issues.

Although it seems that we're making a big deal of this, these issues do not generally cause large problems in simulations geared toward signal-integrity. That's because generally, the frequency content of the source waveform is attenuated enough such that the effects shown are not seen. It is easy to get fixated on the impulse-responses and the purpose of the section was to illustrate the reasons for the effects.

⁸While the equations look like a delay was applied, the value in this case was a negative delay. It's clear from the phase that the response was in fact advanced.




Figure 11 - Fractional Delay Filter

Before leaving this section, it is important to point out a few items. First, we side-stepped the issue of dealing with $H_N \notin \mathcal{R}$ by directly sampling the continuous-time waveform resulting from the sum of co-sinusoids provided. This is always an option and is also an option for dealing with unevenly spaced points provided as a response. In fact, we can always then compute the DFT of the resulting discrete-time impulse response to compute a new response that has been sampled onto an evenly spaced frequency grid.

Second, while we saw that a reason for ringing and causality issues in the time-domain response were due to fractional sample group delay in the impulse response that was correlated with the potential non-real value of H_N and that the determination of a time-delay that makes H_N real did more to explain things than fix things, you will find that this technique is useful when going in and out of the time-domain for the purpose of re-sampling frequency-domain data.

At this point, before addressing the topic of interpolation which is really at the heart of connecting discrete-frequency and continuous-frequency responses, we diverge slightly to address a related topic – that of the relationship between periodicity and discreteness.

Periodicity and Discreteness or (Repetition and Zero Insertion)

While going through the process of converting a continuous-frequency response given by (7) we observed two phenomena that were not fully commented on. In (13) we saw the production of $H_{disc}(f)$ in the frequency domain created by sampling $H_{bwl}(f)$ and we saw that in (18), the time-domain view of this same process caused $h_{disc}(t)$ to be periodic. This makes sense because we know that the Fourier transform of a periodic signal is computed using the Fourier series and the Fourier series produces results for discrete frequencies. We saw the similar process in the production of $\hat{h}(t)$ by sampling $h_{disc}(t)$ in (18) which caused the less obvious effect of making $\hat{H}(f)$ become periodic in frequency by the analogous frequency-domain process shown in (14).

In fact, these are examples of a relationship between each domain in that periodic waveforms in one domain are always discrete waveforms in the other. This is an important and often misunderstood feature of the very much relied upon DFT and IDFT. Since x is discrete-time,

we know that X must repeat in frequency and since X is discrete-frequency, then x must be periodic in time. In other words, if they are discrete in both domains, they must be periodic in both domains. The reason this is important is because wrapped up in this feature is the concept of aliasing, upsampling and finally the important point of this paper – interpolation.

Note that when dealing with the DFT and the IDFT, both x and X are a finite number of points giving the user the indication that the impulse response is limited and isolated in time and the frequency response is limited in frequency. This is not the case and both finite sets of data are finite only because the additional repeating data is redundant.

To see how repetition and discreteness are interrelated mathematically, consider a K element time sequence x whose K element DFT with $n \in 0 \dots K-1$ is given by (23) as:

$$X_n = \sum_{k=0}^{K-1} x_k e^{-j2\pi \frac{nk}{K}} \quad (23)$$

The DFT appears to produce a spectrum restricted to K points. Let's examine the effect of repeating this spectrum on a new time sequence. Let M be the number of repetitions of the frequency spectrum such that we now have $K' = MK$ frequency points such that for $k', n' \in 0 \dots K'$ we have:

$$X'_{n'} = X_{n' \bmod K} \quad (24)$$

The mod function is defined as:

$$a \bmod b = a - \text{floor}\left(\frac{a}{b}\right)b \quad (25)$$

In other words, X' is X concatenated or repeated M times. The inverse DFT of this new sequence is given by:

$$\begin{aligned} x'_{k'} &= \frac{1}{K'} \sum_{n'=0}^{K'-1} X'_{n'} e^{j2\pi \frac{n'k'}{K'}} = \frac{1}{K'} \sum_{n'=0}^{K'-1} X_{n' \bmod K} e^{j2\pi \frac{n'k'}{K'}} = \frac{1}{K'} \sum_{n=0}^{K-1} X_n \sum_{m=0}^{M-1} e^{j2\pi \frac{(Km+n)k'}{KM}} \\ &\dots = \frac{1}{KM} \sum_{n=0}^{K-1} X_n e^{j2\pi \frac{nk'}{KM}} \sum_{m=0}^{M-1} e^{j2\pi \frac{mk'}{M}} \end{aligned} \quad (26)$$

In order to make some important observations about (26), define $r = e^{j2\pi \frac{k'}{M}}$ causing the summation at the end of (26) to become:

$$S = \sum_{m=0}^{M-1} r^m \quad (27)$$

Our first observation is that when k' is a multiple of M , then $r=1$ making $S=M$ in (27) and (26) becomes:




Figure 12 - Frequency-Domain Repetition Effect on Time-Domain

$$x'_{k'} = \frac{1}{K} \sum_{n=0}^{K-1} X_n e^{j2\pi \frac{nk'}{KM}} = x_{\frac{k'}{M}} \quad (28)$$

In other words, (28) says that whenever k' is a multiple of M , $x'_{k'}$ is one of the original points in x .

Our second observation is that when k' is not a multiple of M , then we rewrite the summation in (27) in its analytic form as:

$$S = \sum_{m=0}^{M-1} r^m = \frac{r^M - 1}{r - 1} = \frac{e^{j2\pi k'} - 1}{e^{j2\pi \frac{k'}{M}} - 1} \quad (29)$$

In (29) we see that the numerator is always zero and the denominator cannot be zero, meaning that $S=0$ in this case which means that when k' is not a multiple of M , we have $x'_{k'}=0$.

We see that repeating the spectrum in one domain M times is the same as inserting a zero in between every M^{th} point in the other domain. In fact the sequence in each domain is assumed to repeat infinitely meaning that nothing exists between the points in the other domain. This is what is meant by periodicity in one domain ensuring discreteness in the other.

Zero Padding and Interpolation

So far, we've shown that a continuous-frequency response undergoes a complex procedure to become a corresponding discrete-frequency response and that although the sample points of the discrete-frequency response are exact samples of the continuous-frequency response, the corresponding discrete-time response is not necessarily exact samples of the associated continuous-time response (due to time-aliasing). But a more fundamental problem philosophically is not the error in the discrete-time response, but the fact that it is discrete in the first place. One must ask the question, "how can one make inferences on continuous-time and

continuous-frequency systems with only discrete-time and discrete-frequency data?” The answer to this question is a question of interpolation. Most have heard of the famous Nyquist criteria which says that if a time-domain waveform is sampled with sufficient conditions met, then the continuous-time waveform can always be recovered.

We said earlier that if a system is band-limited in frequency response and time-limited in impulse-response, we can achieve no error between the discrete-time response and the actual underlying continuous-time response. In other words, the discrete-time response will be exact samples of the continuous-time response. Not only that, If we can meet the band- and time-limit criteria, we can reconstruct both the continuous-time and continuous-frequency responses from the discrete information provided.

To see how this works, we'll draw on all of the information provided previously in this paper.

In the last section we saw the effect of repetition. To summarize the effect, if one is provided a discrete-frequency response (which we know is representative of a response that repeats over and over with frequency) and we repeat the response, then the discrete-time response is the same as the original discrete-time response except that zeros are placed in between every sample. If we knew or assumed, however that the original discrete-frequency response were truly limited in frequency, then we can brick-wall filter the repeated discrete-frequency response to the original band-limit prior to computing the inverse transform that produces the discrete-time response. When we do this, we find that instead of a zero inserted between each point, we have inserted an *interpolated* point. The interpolated point is interpolated using the commonly known method of sinc interpolation and the interpolated point lies directly on the continuous time impulse-response. This process of repetition followed by brick-wall filtering is the same as zero padding the discrete-frequency response.

In the mathematics vocabulary we've used so far, this process creates a version of $\hat{h}(t)$ which is the same time length, but with double the number of points, and these extra points are samples of the discrete-time response $h_{disc}(t)$. But since we also said that the response was time-limited also, a period of $h_{disc}(t)$ is simply the isolated impulse response $h_{bwl}(t)$ and because of the band limiting is also exact samples of $h(t)$.

Finally, because we are allowed to make as many repetitions as we want of the discrete-frequency response, we can obtain an arbitrarily continuous version of $h(t)$.

To summarize, if we have the discrete-frequency response which is band- and time-limited, then we can perfectly reconstruct the continuous function $h(t)$ by zero padding the discrete-frequency response.

A similar situation exists with regard to the interpolation of the discrete-frequency response. If we were to insert a zero between every discrete-frequency point and convert back to the time-domain, this is the same as taking two periods of the response $\hat{h}(t)$. But we said that the response was actually time-limited, so we can set the repeated points to zero with no bad effect except to double the period of $\Delta_{KT}(t)$ used to produce the impulse-train response. Converting this zero padded impulse-response back to the time-domain has the effect of converting the zeros that we inserted into the discrete-frequency response into interpolated points. By virtue of the

time- and band-limiting, these points are actual points on $H(f)$ just as the original points were. Since we can zero-pad the response $\hat{h}(t)$ as much as we'd like, we can determine to an arbitrary degree the continuous response $H(f)$.

In summary, if a discrete-frequency response is both time- and frequency-limited and we meet the sufficiency of sampling criteria in both frequency and frequency spacing, then the continuous response associated with the discrete-frequency response can be completely reconstructed. Of course real systems will not meet this criteria to some degree, but to the degree it does, this capability for perfect reconstruction holds.

Two implementation notes regarding the DFT and IDFT are worthwhile to discuss.

Here, we have basically provided a method for interpolation of both the frequency-response and the impulse-response. Regarding the frequency response, a method of interpolation would be to compute the IDFT of X corresponding to H , pad the resulting x , and compute the new DFT X' which corresponds to the new interpolated response H' . If you do this, you need to be careful about where the padding is applied insofar as x is concerned. Remember that in x the first half of the samples correspond to positive time and the second half of the samples correspond to negative time, according to the time rules outlined in (37). The padding must go in between the boundary between the positive time and negative time samples to have the proper effect. If this is not done, you will get wrong results.

The second detail involves interpolation of the impulse response through zero padding of the discrete-frequency response. We have similar problems as in the previous paragraph in that the second half of X represents negative frequencies. The padding must go between the boundary between the positive and negative frequencies in X . Sometimes it is easier to convert X back to H , apply the zero padding to the end to produce an H' and convert back to an X' using the rules in Appendix A.

One final note regarding interpolation is that you are not obliged to interpolate through brick-wall filtering of the repeated response in the frequency-domain, which is equivalent to zero padding in the frequency-domain which is equivalent to convolution with a sinc function in the time-domain. There are other possibilities that don't introduce the Gibbs phenomena when you cannot utilize a high enough end frequency. This is why we explained zero padding as brick-wall filtering of the repeated response in the first place. We can choose to filter the repeated response in alternate fashions. For example, Figure 13 shows a comparison between sinc interpolation with linear interpolation for a doubling of the end frequency. In Figure 13, the top row shows first the sinc interpolating function which can be seen to have no effect on the original time-domain response and to interpolate the points in between. Next, the frequency-domain view is shown with the original RC network response repeated shown in the dotted line, the brick-wall filter equivalent to the sinc interpolating function and the cross-hatched area which shows the resulting frequency response. Finally, the interpolated impulse-response is shown. Note that the original points fall on the faint gray response which is the actual continuous-time response for a system which is not actually band-limited. The interpolated points exhibit the Gibbs phenomena. In the next line we repeat the exercise for linear interpolation. First we see that the linear interpolating filter also has the effect of not changing

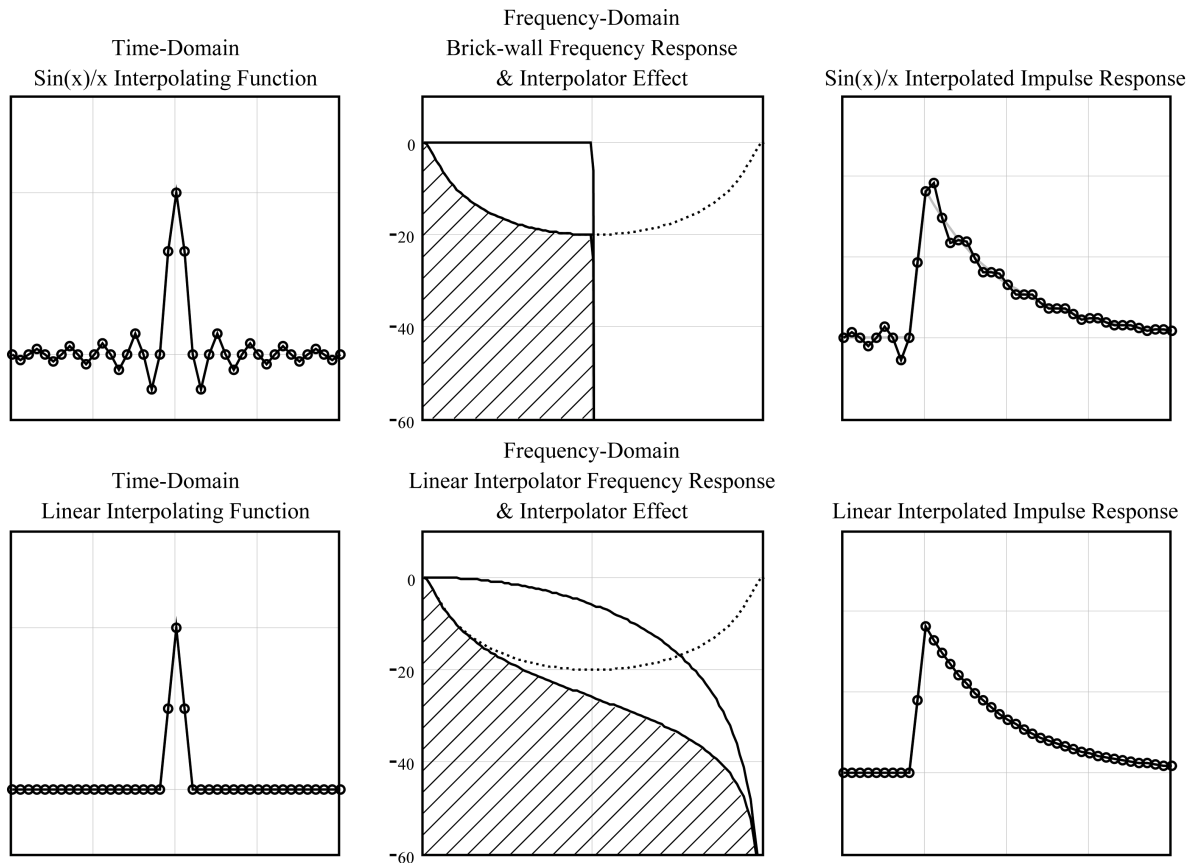


Figure 15 - Comparison of Sinc Interpolation with Linear Interpolation

the original frequency points and taking the average of the two adjacent points for interpolation. We can see also that to some extent it is a short, crude approximation of the sinc interpolating function. Next we see the filter response and the effect on the dotted line repeated response in the cross-hatched area. Here we see that we did not strictly limit the response, but instead altered the original response and extrapolated to some extent the response out in frequency. Finally, the time-domain effect is shown where we see the linearly interpolated impulse-response. There are many issues to deal with in considering this type of interpolation method as a general practice, most importantly that it causes non-linear effects. These non-linear effects stem from the fact that repeating a spectrum is not a linear operation. Understand that linear operations cannot cause a single frequency input to a system to produce other frequencies and clearly repeating a spectrum introduces this effect. When the repeated spectrum is brick-wall filtered, this non-linearity possibility is not present, but in other interpolation schemes, portions of the repeated spectrum are retained.




Figure 14 - Chirp Z Transform

Re-sampling and the Chirp Z-Transform

Given a set of s-parameters, one goal is to perform simulations directly using them, employing transfer functions calculated based on them. Two problems immediately arise. The first is that not all s-parameter files have been generated using the same frequency spacing or end frequency. The second is that even with a common end frequency, simulation at a different sample rate than that dictated by the end frequency is also often desired. These problems require re-sampling the s-parameters onto a frequency grid different than that already provided. Fortunately, there is a rigorous way to perform this re-sampling and it is based on the chirp z-transform. The method presented here is to first place the s-parameter based responses into DFT form as provided in Appendix A. Then, the IDFT is computed which will provide a time-domain impulse-response that has a sample rate dictated by the s-parameters provided. Finally, the chirp z-transform is computed to produce a DFT that matches the desired frequency spacing and end frequency. All we need now is to understand the chirp z-transform.

One definition of the DFT is that it is the z-transform of a finite duration time sequence evaluated at equally spaced points on the unit circle. To see what is meant by this definition, consider that the definition of the z-transform is:

$$H(z) = \sum_{k=-\infty}^{+\infty} x_k z^{-k} \quad (30)$$

For a finite duration sequence x consisting of K points, (30) reduces to:

$$H(z) = \sum_{k=0}^{K-1} x_k z^{-k} \quad (31)$$

If we define, for $n \in 0 \dots K-1$, $z_n = e^{j2\pi \frac{n}{K}}$, we have z defined as points on the unit circle that start at 1 and are equi-angularly spaced at angular increments of $2\pi/K$ and we have:

$$X_n = \sum_{k=0}^{K-1} x_k e^{-j2\pi \frac{nk}{K}} \quad (32)$$

The equation in (32) is, of course, the definition of the DFT.

The chirp z transform (CZT) is a similar computation as the DFT except that for $M+1$ elements, $m \in 0 \dots M$, it allows for a set of points ζ to consist of points on an equi-angularly spaced spiral or arc. The location of these points are described such that:

$$\zeta_m = A W^m \quad A = A_0 e^{j2\pi\theta_0} \quad W = W_0 e^{j2\pi\phi_0} \quad (33)$$

The chirp z transform is defined as:

$$X_m = H_{czt}(\zeta_m) = \sum_{k=0}^{K-1} x_k \zeta_m^{-k} = \sum_{k=0}^{K-1} x_k [A W^m]^{-k} = \sum_{k=0}^{K-1} x_k A^{-k} W^{-mk} \quad (34)$$

In (34), A_0 , θ_0 , W_0 , and ϕ_0 are arbitrary constants that define the arc. To see how the arc is defined, consider that when $m=0$, $\zeta_0 = A$ that the beginning of the arc is at radius A_0 and at a beginning angle of $2\pi\theta_0$ and that when $m=M$, $\zeta_M = A_0 W_0^M e^{j2\pi[\theta_0 + M\phi_0]}$ meaning that the end of the arc is at radius $A_0 W_0^M$ and at an angle $2\pi[\theta_0 + M\phi_0]$ and therefore the arc spans the angle $2\pi M\phi_0$. Thus for each $m \in 0 \dots M$:

$$\zeta_m = A_0 W_0^m e^{j2\pi[\theta_0 + m\phi_0]} \quad (35)$$

Here we see that the angle between each point on the arc is $2\pi\phi_0$. This is all shown graphically in Figure 14.

The CZT has many uses and has a fast implementation⁹. The use we are concerned with regarding s -parameters is only the ability to re-sample onto a desired equally spaced frequency scale. For our purposes, we define $A_0=1$, $W_0=1$, $\theta_0=0$, $\phi_0=\Delta f$, so for $M+1$ frequency points, we have:

$$X_m = \sum_{k=0}^{K-1} x_k \zeta_m^{-k} = \sum_{k=0}^{K-1} x_k e^{-j2\pi km \Delta f} \quad (36)$$

⁹ See M. T. Jong, *Methods of Discrete Signal and Systems Analysis*. 1982. 198

Be aware that if you use this method of interpolation, you will need to deal with any of the points in the CZT that involve frequencies beyond those provided in the original data, because the CZT will wrap around the unit circle. These frequencies points can either be set to zero or filtered with a different interpolation function as outlined in the previous section.

When the s-parameters or response provided are unevenly spaced, we can always express the IDFT of the original response in the sum of cosines form, sample the continuous-time sum of these cosines and apply the CZT.

Summary

We have explored the relationship between continuous-frequency and discrete-frequency responses in both the frequency- and time-domains and have found that thinking about this relationship in both domains provides the most insight. We found that to the extent that the continuous-frequency response is band-limited and its impulse response is time-limited, then if the discrete-frequency response is properly sampled, then not only are the sampled frequency response exact points on the continuous-frequency response, but the implied impulse-response are exact points on the continuous-time impulse-response. We examined the mechanics of the relationship and have seen the side-effects when the sufficiency of sampling conditions are not met. We explored issues in handling discrete-frequency responses in providing real impulse-responses. We examined the relationship between periodicity and discreteness in both domains and expanded upon this to develop techniques for interpolation in one domain through zero-padding in the other. Finally we provided some alternate interpolation techniques and provided some re-sampling techniques that are necessary in utilizing s-parameters directly in simulation.

Acknowledgements

This year I experimented with writing my DesignCon paper in a free software environment and was astounded by how good the many free tools are. I used:

- Ubuntu Linux 11.10 Oneiric Ocelot maintained by Canonical.
- LibreOffice Writer and Impress provided by the OpenOffice.org.
- WINE sponsored by CodeWeavers for emulating certain Windows programs under Linux like MathCad 2000 by MathSoft.
- VirtualBox from Oracle for running a Windows 2000 environment under Linux.
- TpX written by Alex Tsyplakov for all of the drawings.

Thanks for all of this excellent free software!

Appendix A - Relationship Between Response Elements and DFT Elements

Given a K element real time-domain sequence \mathbf{x} , where $k \in K-1$, $x_k \in \mathbb{R}$ which is related to a K element complex frequency-domain sequence \mathbf{X} through the DFT (i.e. $\mathbf{X} = \text{DFT}(\mathbf{x})$), then we define a sample period T and a corresponding sample rate F such that the sample times are defined as (37):

$$t_k = \begin{cases} kT = \frac{k}{F} & \text{if } k \leq \frac{K}{2} \\ (k-K)T = (k-K)\frac{1}{F} & \text{if } k > \frac{K}{2} \end{cases} \quad (37)$$

The following relationships hold between the DFT \mathbf{X} and the response \mathbf{H} :

For K even, define $N = K/2$ such that $K = 2N$, $n \in 0 \dots N$:

$$H_n = X_n \quad X_n = \begin{cases} \Re(H_0) & \text{if } n=0 \\ \Re(H_N) & \text{if } n=N \\ H_n & \text{otherwise} \end{cases} \quad (38)$$

$$f_n = \frac{n}{K}F = \frac{n}{N}\frac{F}{2} \quad F = \frac{1}{T} = 2f_N \quad (39)$$

For enforcement of realness when K even, define $\sigma \in 1 \dots N-1$:

$$X_0 \in \mathbb{R} \quad X_N \in \mathbb{R} \quad X_{N+\sigma} = X_{N-\sigma}^* \quad f_{N+\sigma} = -f_{N-\sigma} \quad (40)$$

For K odd, define $N = (K-1)/2$ such that $K = 2N+1$, $n \in 0 \dots N$:

$$H_n = X_n \quad X_n = \begin{cases} \Re(H_0) & \text{if } n=0 \\ H_n & \text{otherwise} \end{cases} \quad (41)$$

$$f_n = \frac{n}{K}F = \frac{n}{2N+1}F \quad F = \frac{1}{T} = \frac{K}{N}f_N = \frac{2N+1}{N}f_N \quad (42)$$

For enforcement of realness when K even, define $\sigma \in 1 \dots N$:

$$X_0 \in \mathbb{R} \quad X_{N+\sigma} = X_{N-\sigma+1}^* \quad f_{N+\sigma} = -f_{N-\sigma+1} \quad (43)$$

$$\text{The frequency spacing is always defined as: } \Delta f = \frac{F}{K} = \frac{1}{KT} \quad (44)$$