Chapter 4

Frequency Domain and Fourier Transforms

Frequency domain analysis and Fourier transforms are a cornerstone of signal and system analysis. These ideas are also one of the conceptual pillars within electrical engineering. Among all of the mathematical tools utilized in electrical engineering, frequency domain analysis is arguably the most far-reaching. In fact, these ideas are so important that they are widely used in many fields – not just in electrical engineering, but in practically all branches of engineering and science, and several areas of mathematics.

4.1 Frequency Content: Combining Sinusoids

The most common and familiar example of frequency content in signals is probably audio signals, and music in particular. We are all familiar with "high" musical notes and "low" musical notes. The high notes do in fact have higher frequency content than the low notes, but what exactly does this mean?

The place to start to answer this question is to consider sinusoids. Recall that the general expression for a sinusoid at frequency ω (or frequency f in Hertz) is

$$x(t) = a\sin(\omega t + \phi) = a\sin(2\pi f t + \phi)$$

When considered as an audio signal, $x(t)$ indicates the changes in air pressure on our ears as a function of time. What is important here is the time variation of the air pressure from some ambient value rather than the ambient value of the pressure itself. A negative value refers to that amount below the baseline (ambient) pressure, while a positive amount refers to a pressure higher than the baseline.

†Lecture Notes for ELE201 Introduction to Electrical Signals and Systems.
‡Thanks to Sean McCormick and Richard Radke for producing the figures.

So, $x(t)$ being a sinusoid means that the air pressure on our ears varies periodically about some ambient pressure in a manner indicated by the sinusoid. The sound we hear in this case is called a pure tone. Pure tones often sound artificial (or electronic) rather than musical. The frequency of the sinusoid determines the "pitch" of the tone, while the amplitude determines the "loudness". It turns out that the phase of the sinusoid does not affect our perception of the tone, which may not be surprising for a pure tone, but is somewhat surprising when we start combining sinusoids.

We can combine two sinusoids by adding the signals in the usual way. For example,

$$x(t) = \sin(2\pi t) + \sin(4\pi t)$$

is a combination of a sinusoid with frequency 1 Hz and a sinusoid with frequency 2 Hz. Here the amplitude of each sinusoid is 1 and the phase of each is 0. A plot of $x(t)$ is shown in Figure 4.1. The "sound" created by $x(t)$ is the combination of the two pure tones that make $x(t)$. Unfortunately, as we'll discuss in more detail in Chapter XX, humans can't hear the pure tones that comprise the signal $x(t)$ above since the frequencies are too low.

However, we can make a similar combination with signals at frequencies humans can hear. For example, consider the signal

$$d(t) = \sin(2\pi \times 350 \times t) + \sin(2\pi \times 440 \times t)$$

Each of the two sinusoids (at frequencies 350 Hz and 440 Hz) alone corresponds to a pure tone that can be heard by the normal human ear. Their combination, i.e., the signal $d(t)$, makes a very familiar sound, namely the dial tone on a standard U.S. telephone line. A plot of $d(t)$ is shown in Figure 4.2. Note that in this figure only 2 hundredths of a second are shown. Because the frequencies are high, if we showed even a whole second, the signal would oscillate so many times (350 and 440 for the constituent sinusoids) that not much useful detail would be seen.

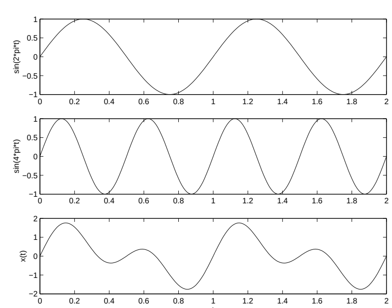

Figure 4.1: Combining sinusoids.

Although the dial tone is a simple example of a sound that still sounds artificial, by combining more sinusoids at different frequencies we can get many

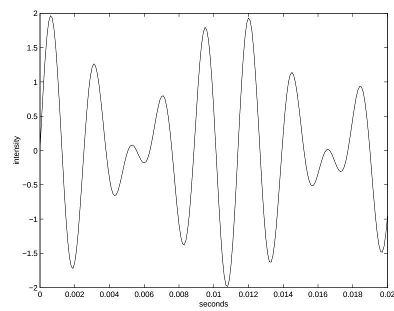

Figure 4.2: Graph of a dial tone.

other sounds. Musical notes that we find pleasing largely consist of pure tones near the pitch of the musical note, but also contain other frequencies that give each instrument its particular qualities. Voice and other natural sounds are also comprised of a number of pure tones.

Amazingly, all sounds can be built up out of pure tones, and likewise all time signals can be constructed by combining sinusoids. Similarly, starting with a general time signal, one can break this signal down into its constituent sinusoids. How to do this and the consequences of such constructions/decompositions is the subject of frequency domain analysis and Fourier transforms. First, we briefly discuss two other different motivating examples.

4.2 Some Motivating Examples

Hierarchical Image Representation

If you have spent any time on the internet, at some point you have probably experienced delays in downloading web pages. This is due to various factors including traffic on the network and the amount of data on the page requested. Images, as we will see in Chapter XX, require a substantial amount of data, and as a result, downloading images can be slow.

Suppose we are browsing through a large database/archive of images and wish to find a collection of images of a particular type. Downloading each image completely and then deciding that the bulk of the images are not what we are after can be time-consuming. Of course, we may be able to make this determination (i.e., whether or not the image suits our purposes) with much less quality than the full image possesses. Certainly we may be able to reject most of the images with only a very rough idea of their content. Wouldn't it be nice to be able to make this decision with only 1/10 of the data? That way we could download and rifle through the database much faster to find what we're after.

The problem is that if we take the standard representation of images (8 bit gray level for each pixel) and send the first 1/10 of the pixels, we will simply get a portion (namely, 1/10) of the original image, albeit at high resolution.

On the other hand if we had a way to represent the "coarse" (or low frequency) information separately from the "fine" (or high frequency) information, we could request the coarse information first and only request the additional detail if desired.

Figure 4.3 shows an example of an original image, together with xx% of the pixels in the usual format and xx% of the information at the lowest frequencies. For many purposes, the low frequency version may be adequate to make decisions, and it certainly seems more valuable than if we had only the first xx% of the pixels. This notion of a "hierarchical" representation can be formalized using ideas from frequency domain analysis.

It turns out that frequency domain ideas can help with this browsing problem in a different way as well. In addition to creating hierarchical representations that allow sending the most important information first, one might consider reducing the total amount of data in the first place. As we will see, one of the standard methods for image compression known as JPEG is based on frequency domain ideas. This is the subject of data compression, which will be discussed in Chapters XX and XX.

Radio and TV Transmission

Radio, television, and some other forms of communication (e.g., cell phones) transmit information via electromagnetic waves. The various sources in these applications can be transmitting simultaneously and in the same geographic region. But how is it that we can "tune in" to a specific radio station, television program, or individual with whom we're communicating, rather than hearing the jumble of all the various transmissions put together?

The answer is that different transmissions agree to use different frequencies. Thus, even though all the signals *are* "jumbled" together in the time domain, they are distinct in the frequency domain. With some basic frequency domain processing, it is straightforward to separate the signals and "tune in" to the frequency we're interested in.

4.3 A Trivial Frequency Decomposition

Before discussing frequency representations for general signals, we consider an example that is trivial but is still somewhat illustrative. Consider a situation in which we are interested in the values of a signal $x[0]$ and $x[1]$ at only two times $n = 0$ and $n = 1$.

Suppose we have access to the two values of the signal, and wish to convey these values to a friend. We could simply convey the values $x[0]$ and $x[1]$ themselves and we're done.

However, consider the the following alternative scheme. Define a new signal $X[0]$ and $X[1]$ by

$$\begin{aligned} X[0] &= x[0] + x[1] \\ X[1] &= x[0] - x[1] \end{aligned}$$

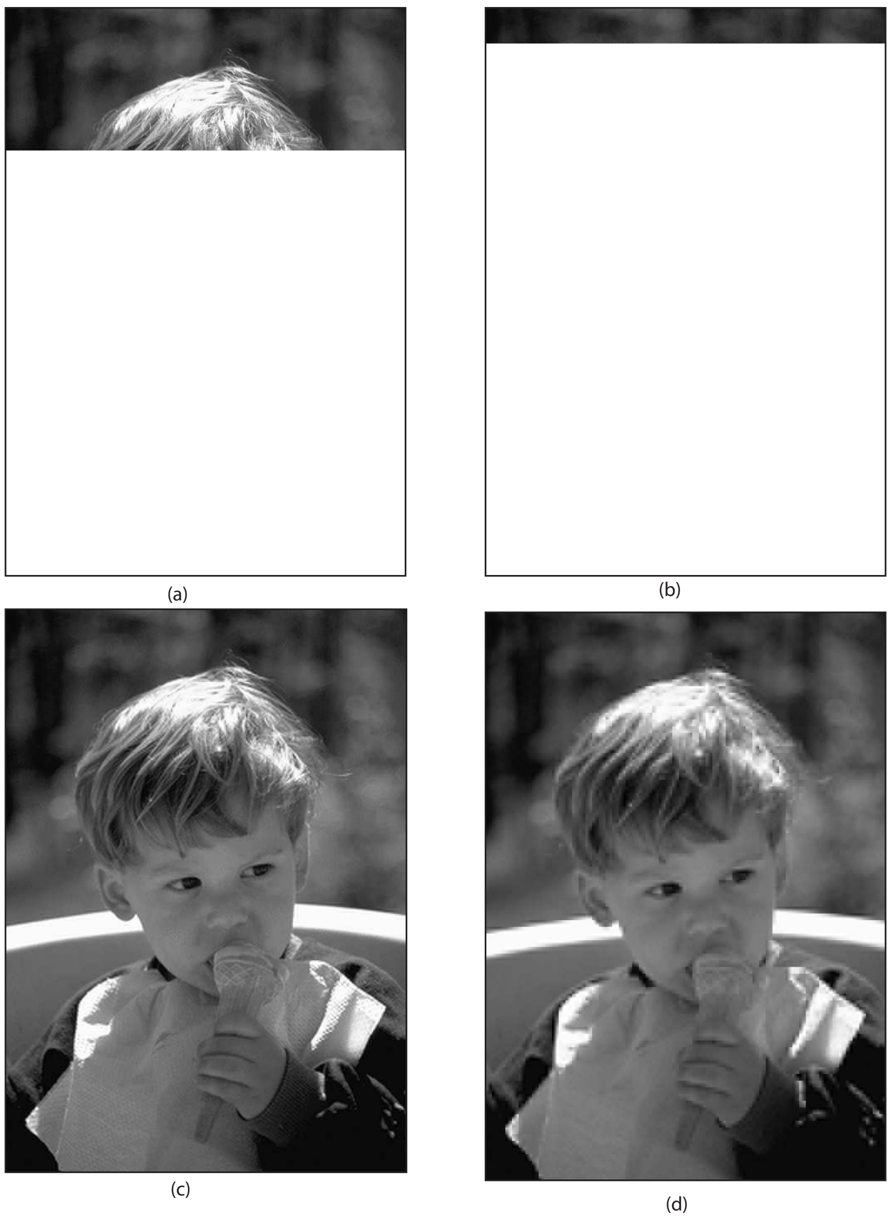

Figure 4.3: Sending image data in usual format vs. sending only low frequency data. a) 25% of data in usual format. b) 6.25% of data in usual format. c) 25% of lowest frequency data. d) 6.25% of lowest frequency data.

We can then convey the values $X[0]$ and $X[1]$ to our friend, and using these values the friend can recover the original signal $x[0]$ and $x[1]$. Namely, given $X[0]$ and $X[1]$, we get $x[0]$ and $x[1]$ by

$$\begin{aligned} x[0] &= \frac{1}{2}(X[0] + X[1]) \\ x[1] &= \frac{1}{2}(X[0] - X[1]) \end{aligned}$$

This example seems trivial, but even this simple case illustrates some interesting properties. For example, in the new scheme, information on *both* values of $x[0]$ and $x[1]$ are being sent simultaneously – partially in $X[0]$ and partially in $X[1]$. If we receive only $X[0]$ then we don't know *either* of the original signal values, but we know something about both.

Roughly, $X[0]$ is the "low frequency" content of the signal $x[0], x[1]$, This is the part of the signal that doesn't change, or that's common to both $x[0]$ and $x[1]$. It is a sort of an average between $x[0]$ and $x[1]$ (in fact, it's exactly twice the average). $X[1]$ is the "high frequency" content of the signal – the part that changes between $x[0]$ and $x[1]$. In fact, the equation shows that it is simply the difference between $x[0]$ and $x[1]$.

We can think of the signal $X[\cdot]$ as just a different representation for the signal $x[\cdot]$ since we can easily go back and forth between the two representations (using the equations above). $X[\cdot]$ will be called the frequency domain representation, while the original signal $x[\cdot]$ will be called the time domain representation. The term "time domain" refers to the fact that when describing the values of $x[\cdot]$ directly, we simply give the values of $x[n]$ where $n = 0, 1$ denotes time. On the other hand, the "frequency domain" description gives the values of $X[k]$ where k also happens to take on the values 0 and 1, but k really denotes frequency. That is, $X[0]$ is *not* the value of the original signal at time 0, but is rather the frequency content of the original signal at the low frequency (frequency 0).

As we will see, this alternate representation can be immensely useful in analyzing signals and the effects of systems on signals. Something very complicated in time domain can turn out to be very simple in frequency domain. Obviously, to be useful we need to extend this notion of frequency representations to longer signals. Fourier transforms and frequency domain analysis do just that.

4.4 Summary of the Four Cases

The details of how to formulate and extend the ideas in the sections above depend on whether the underlying signal is continuous-time or discrete-time, and whether it is periodic/finite-duration or aperiodic/infinite-duration. This results in four distinct cases, although there are many similarities and connections between these four cases.

The names of the transform in the four cases are summarized in Figure 4.4. The nomenclature is unfortunately somewhat confusing, but is rather entrenched, and so we will use the standard terminology.

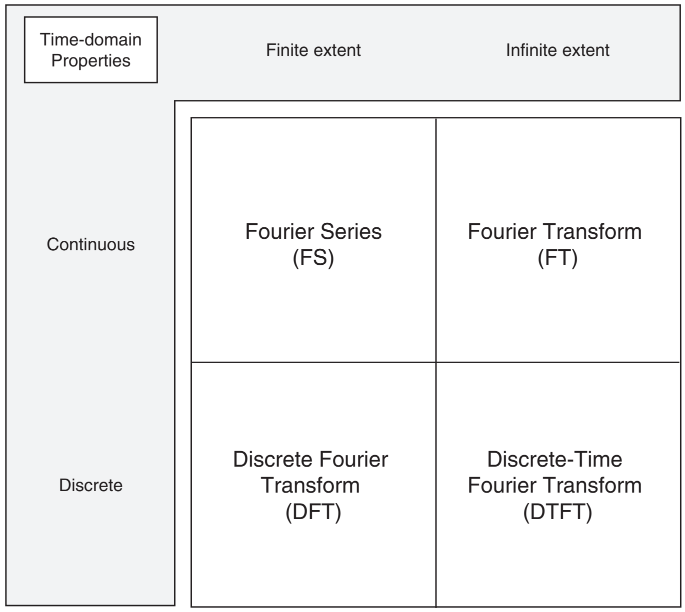

Figure 4.4: The four Fourier transforms broken down by time domain properties.

Figure 4.5 gives the definitions of the transforms in each of the four cases. First, we'll make some general remarks about all the transforms, and then we'll discuss a little more of the details for the case of the Fourier transforms (continuous-time, infinite duration) and the DFT (discrete-time, finite duration). For now it's enough to get just a rough idea of the transforms.

Notation, Frequency Variable

In each case, the original (time-domain) signal is represented by lower case x – in particular, $x(t)$ in continuous-time and $x[n]$ in discrete-time. The appropriate Fourier transform in each case is represented by upper case X.

Just as different symbols are used to denote time in the continuous-time and discrete-time cases, different symbols are used for the frequency variable depending on whether there are a continuous range of frequencies or a discrete set of frequencies. For a continuous range of frequencies, the variable ω is used. As before, ω denotes frequency in radians/sec. For the DFT and Fourier series, only a discrete set of frequencies are used. The discrete frequencies are all a multiple of a specific frequency, denoted ω_0, called the *fundamental frequency*. The variable k is used to specify the particular multiple of the fundamental frequency. Of course, the actual frequency of the complex exponential is $k\omega_0$, again in units of radians/sec.

Forward and Inverse Transform

Note that there are two equations in each case. The first equation has the transform X on the left-hand-side and either a sum or integral involving the

Fourier Series (FS)

For $x(t)$ of duration T, set $\omega_0 = \frac{2\pi}{T}$.

$$x(t): \quad 0 \leq t \leq T$$
$$X[k]: \quad k = \ldots, -2, -1, 0, 1, 2, \ldots$$

$$X[k] = \frac{1}{T}\int_{t=0}^{T} x(t)\, e^{-jk\omega_0 t} dt$$

$$x(t) = \sum_{k=-\infty}^{\infty} X[k]\, e^{jk\omega_0 t}$$

Fourier Transform (FT)

$$x(t): \quad -\infty < t < \infty$$
$$X(\omega): \quad -\infty < \omega < \infty$$

$$X(\omega) = \int_{t=-\infty}^{\infty} x(t)\, e^{-j\omega t} dt$$

$$x(t) = \frac{1}{2\pi}\int_{\omega=-\infty}^{\infty} X(\omega)\, e^{j\omega t} d\omega$$

Discrete Fourier Transform (DFT)

For $x[n]$ of length N, set $\omega_0 = \frac{2\pi}{N}$.

$$x[n]: \quad n = 0, 1, \ldots, N-1$$
$$X[k]: \quad k = 0, 1, \ldots, N-1$$

$$X[k] = \sum_{n=0}^{N-1} x[n]\, e^{-jk\omega_0 n}$$

$$x[n] = \frac{1}{N}\sum_{k=0}^{N-1} X[k]\, e^{jk\omega_0 n}$$

Discrete-Time Fourier Transform (DTFT)

$$x[n]: \quad n = \ldots, -2, -1, 0, 1, 2, \ldots$$
$$X(\omega): \quad -\pi \leq \omega \leq \pi$$

$$X(\omega) = \sum_{n=-\infty}^{\infty} x[n]\, e^{-j\omega n}$$

$$x[n] = \frac{1}{2\pi}\int_{\omega=-\pi}^{\pi} X(\omega)\, e^{j\omega n} d\omega$$

Figure 4.5: Definitions of the forward and inverse Fourier transforms in each of the four cases.

original signal x on the right-hand-side. This equation gives the expression for the transform in terms of the original signal. For each frequency, the expression tells how much of this frequency is contained in the original signal x.

The second equation in each case has the original signal x on the left-hand-side, and either a sum or integral involving the transform X on the right-hand-side. This equation tells how to recover a signal if we know the frequency content of the signal. This equation is often referred to as the *inverse transform*, since it "un-does" the transform. Sometimes to emphasize the distinction with the inverse, the term *forward transform* is used to refer to the first equation where X is obtained from knowledge of the original signal x.

Complex Values

Because of the complex exponential in the equations for X, it is possible (in fact, it is typical), that the transform X will take on complex values. This is not a problem. With a little thought this should be expected.

Recall we can write a complex number in terms of its magnitude and phase (i.e., its polar representation). These quantities determine the magnitude and phase of the underlying complex exponential. $X(\omega)$ tells how much content the original signal has at frequency ω. More precisely, $X(\omega)$ tells both the magnitude and phase to use for the complex exponential $e^{j\omega t}$ at frequency ω. If $X(\omega) = Ae^{j\phi}$ in polar form, then $X(\omega)e^{j\omega t} = Ae^{j(\omega t+\phi)}$ so that A and ϕ are the magnitude and phase, respectively, of the complex sinusoid.

If X can be complex-valued, it is natural to wonder why x isn't also complex-valued (since the forward and inverse transform equations look so similar). The simple answer is that we started with a real-valued signal x. Applying the forward transform then the inverse transform just gives us back the original signal. So, we have to end up with a real-valued signal, since that's what we started with! Of course, it also means that the transform X probably has some special structure (reflecting the special structure that the original signal is real). This is in fact the case and is listed as one of the properties in Table 4.2 later in this chapter.

This discussion also suggests that we could allow the original signal to be complex-valued. This can be done and it is sometimes useful, although in most applications the original signal, representing a physical variable of interest, is real-valued.

Interpretation as Frequency Representation

The equation for the inverse transform is natural once we accept that $X[k]$ (or $X(\omega)$) indicates the "amount" of frequency $k\omega_0$ (or frequency ω, respectively) contained in the original signal. For example, consider the case of Fourier series. In this case, the expression for $x(t)$ is just a sum of sinusoids where the amplitude of the sinusoid of frequency $k\omega_0$ is $X[k]$. The case of the DFT is similar. In the cases where there is a continuous range of frequencies (Fourier transform and DTFT), rather than a sum of terms of different frequencies, we have an integral. All of these equations show that the transform X really does give a frequency representation of x by expressing the original signal in terms of constituent sinusoids, where the amplitude and phase of the sinusoid at frequency $k\omega_0$ (or, frequency ω) is given by $X[k]$ (or, $X(\omega)$).

The fact that we have a forward and inverse transform justifies the interpretation that the transform X is really just a different representation, the frequency domain representation, of the signal x. If we know the transform then we can recover the signal using the inverse transform. Likewise, if we know the signal then we can obtain the frequency domain representation by using the forward transform. Specifying the signal in either time domain or frequency domain completely determines the signal.

Duality

The similarity in the equations for the forward and inverse transforms suggest that there may be some duality. This is indeed the case. Here, we make note only of the two attributes: whether the time/frequency variable is continuous or discrete, and whether the signal in time/frequency domain is finite-extent or infinite extent.

It turns out that being finite extent in one domain means the variable in the other domain is discrete. So, a finite extent in time (the case of Fourier Series and DFT) means the frequency variable in these cases is discrete. Likewise, infinite extent in one domain means the variable in the other domain is continuous (the case of the DTFT and FT).

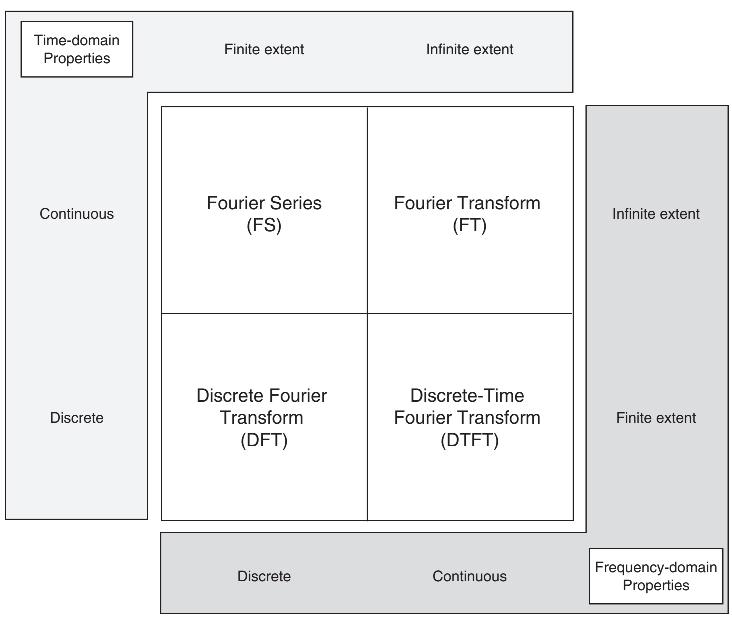

Figure 4.6: Time and frequency domain properties for the four cases.

This relationship is shown in Figure 4.6, which repeats much of Figure 4.4 but with additional labeling showing the frequency domain properties in each case.

The two cases we will discuss in more detail are either continuous and infinite extent in both domains (FT), or discrete and finite extent in both domains (DFT). For the other two cases these properties toggle as we switch domains. That is, for Fourier series, the signal is continuous and finite duration in time, but in frequency it is discrete and infinite duration. For the DTFT, the signal is discrete and infinite extent in time, but continuous and finite extent in frequency.

There are further duality properties as well, which we will discuss in the context of the FT and DFT.

4.5 FT and DFT

In this and the next two sections, we consider in more detail the continuous-time Fourier transform, or simply the Fourier transform (FT), and the discrete Fourier transform (DFT). These are the transform for continuous-time, infinite-duration signals (FT), and for discrete-time, finite-duration signals (DFT). The FT is most amenable to analytical insights and manipulations, and is used in analyzing analog systems. The DFT the workhorse that is widely used in digital computations, because of a fast algorithm called the FFT (fast Fourier transform) for computing the DFT.

For convenience, we repeat the definitions of the forward and inverse transforms here. For the FT, both the time variable t and frequency variable ω are continuous-valued and vary from $-\infty$ to ∞. The transform equations are

$$X(\omega) = \int_{-\infty}^{\infty} x(t) e^{-j\omega t}\, dt \qquad \text{forward transform} \qquad (4.1)$$

$$x(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} X(\omega) e^{j\omega t}\, d\omega \qquad \text{inverse transform} \qquad (4.2)$$

The DFT is used in the case of a finite-duration, discrete-time signal, $x[0], x[1], \ldots, x[N-1]$. The DFT $X[0], X[1], \ldots, X[N-1]$ is also finite-duration and discrete-frequency. The definitions of the forward and inverse DFT are:

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{-jk\omega_0 n} \quad k = 0, 1, \ldots, N-1 \qquad (4.3)$$

$$x[n] = \frac{1}{N} \sum_{k=0}^{N-1} X[k] e^{jk\omega_0 n} \quad n = 0, 1, \ldots, N-1 \qquad (4.4)$$

where

$$\omega_0 = \frac{2\pi}{N} \qquad (4.5)$$

Before considering some examples and properties of Fourier transforms, we introduce some notation and discuss duality a bit further. If $X(\omega)$ is the Fourier transform of a signal $x(t)$ (and so $x(t)$ is the inverse transform of $X(\omega)$), then we say that $X(\omega)$ and $x(t)$ are a Fourier transform pair. If $X(\omega)$ and $x(t)$ are a Fourier transform pair, we may write

$$x(t) \longleftrightarrow X(\omega)$$

Or, to emphasize that one is dealing with Fourier transforms (as opposed to some other transforms, of which there are many), this is sometimes written

$$x(t) \stackrel{\text{FT}}{\longleftrightarrow} X(\omega)$$

We also may use the notation $\mathcal{F}\{\cdot\}$ to denote the Fourier transform operator, and $\mathcal{F}^{-1}\{\cdot\}$ to denote the inverse transform operator. That is,

$$X(\omega) = \mathcal{F}\{x(t)\}$$

and

$$x(t) = \mathcal{F}^{-1}\{X(\omega)\}$$

Similar notation applies in the discrete case for DFT's, with the obvious changes of $x[n]$ and $X[k]$ instead of $x(t)$ and $X(\omega)$, respectively.

Regarding duality, notice again the similarity between the forward and inverse transform equations. Aside from the constants and time and frequency being interchanged, the only difference is the sign in the exponent of the complex exponential. This suggests that whenever we have a transform pair, there is a dual pair with the time and frequency variables interchanged. For example, we will see in Section XX below that the FT of a rect function in time is a sinc function in frequency. From duality, this means that the FT of a sinc function in time will be a rect function in frequency. In general, one should be a little careful making sure all the constants are worked out properly.

Duality also suggests that whenever we have a property of the FT, then there will be a dual property with the roles of time and frequency interchanged. For example, we will see in Section XX below that multiplying a signal in the time domain by a complex exponential corresponds to a shift in the frequency domain. Duality then suggests that a shift in the time domain corresponds to multiplication by a complex exponential in the frequency domain.

4.6 Some Examples of Transform Pairs

Some Fourier transform pairs can be computed quite easily directly from the definition. Given $x(t)$, we can substitute directly into the integral defining $X(\omega)$ and carry out the integration.

Example 4.1 (Fourier transform of the $\text{rect}(\cdot)$ **function)**
Consider $x(t) = \text{rect}(t)$ shown in Figure 4.7. From the definition of $X(\omega)$

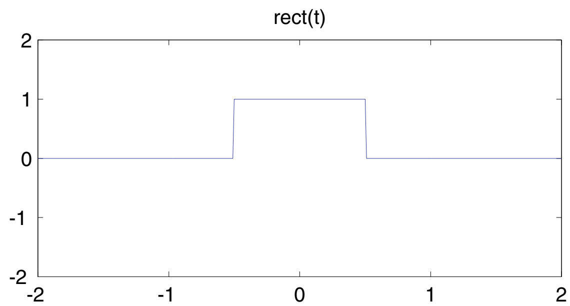

Figure 4.7: Rect function.

we have

$$\begin{aligned} X(\omega) &= \mathcal{F}\{x(t)\} = \int_{-\infty}^{\infty} x(t)e^{-j\omega t}\,dt \\ &= \int_{-\frac{1}{2}}^{\frac{1}{2}} e^{-j\omega t}\,dt \\ &= \left.\frac{e^{-j\omega t}}{-j\omega}\right|_{t=-\frac{1}{2}}^{\frac{1}{2}} \\ &= \frac{e^{j\frac{\omega}{2}} - e^{-j\frac{\omega}{2}}}{j\omega} \\ &= \frac{2\sin(\frac{\omega}{2})}{\omega} \\ &= \text{sinc}(\frac{\omega}{2\pi}) \end{aligned}$$

Therefore,

$$\mathcal{F}\{\text{rect}(t)\} = \frac{2\sin(\frac{\omega}{2})}{\omega} = \text{sinc}(\frac{\omega}{2\pi})$$

A plot of $X(\omega)$ is shown in Figure 4.8.

■

Example 4.2 (FT of the delta function)

One Fourier transform that is particularly easy to compute directly from the definition is that of the Dirac delta function introduced in Chapter XX. Using the defining property of $\delta(t)$, we get

$$\mathcal{F}\{\delta(t)\} = \int_{-\infty}^{\infty} \delta(t)e^{-j\omega t}\,dt = 1$$

■

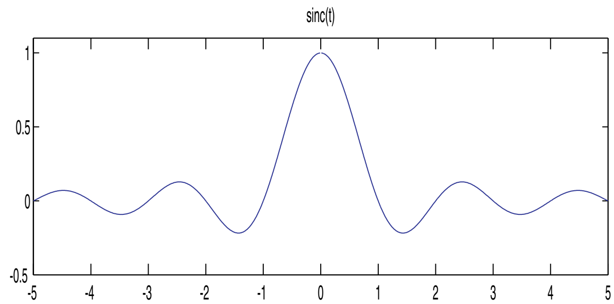

Figure 4.8:

Sometimes a Fourier transform pair can more easily be established by considering the inverse transform. Remember that since a function is uniquely determined in either time or frequency domain, considering the inverse transform is a perfectly valid way to obtain transform pairs. Also, because of the similarity between the definitions of the forward and inverse transforms, there is often a duality that can be exploited to obtain transform pairs.

For example, by considering the delta function $\delta(\omega)$ in frequency domain and using the inverse transform definition, we can see (similar to Example XX) that

$$\frac{1}{2\pi} \longleftrightarrow \delta(\omega)$$

or, equivalently

$$1 \longleftrightarrow 2\pi\delta(\omega)$$

This could also be guessed (except perhaps for the exact constants) by the duality between forward and inverse transforms.

For another example of duality, consider the $\text{rect}(\cdot)$ and $\text{sinc}(\cdot)$ functions. We saw in Example XX that the Fourier transform of the rect function in time-domain is a sinc function in frequency domain. One might guess that the Fourier transform of a sinc function in the time domain is a rect function in frequency domain. This turns out to be correct, as could be easily established by considering a rect in frequency domain and working through a calculation as in Example XX. Of course, one needs to be careful of the exact constants, but the underlying integrals are almost identical.

Table 4.1 shows several other Fourier transform pairs. Extensive tables of Fourier transforms have been developed and the pairs shown in Table 4.1 are a very small subset of known transform pairs.

Direct computation through the definition is not the only way to derive transform pairs. Once we have some transform pairs, many others can be obtained using general properties of the Fourier transform discussed in the next section.

Some Fourier Transform Pairs	
Signal	Fourier Transform
$\delta(t)$	1
$e^{j\omega_0 t}$	$2\pi\delta(\omega-\omega_0)$
$sin(\omega_0 t)$	$\frac{\pi}{j}[\delta(\omega-\omega_0)-\delta(\omega+\omega_0)]$
$cos(\omega_0 t)$	$\pi[\delta(\omega-\omega_0)+\delta(\omega+\omega_0)]$
$x(t)=1$	$2\pi\delta(\omega)$
$\frac{sin(Wt)}{\pi t}$	$X(\omega)=\begin{cases} 1, & \lvert\omega\rvert<W \\ 0, & \lvert\omega\rvert>W \end{cases}$
$step(t)$	$\frac{1}{j\omega}+\pi\delta(\omega)$
$tri(t)$	$sinc^2(\omega/(2\pi))$

Table 4.1:

4.7 Some Transform Properties

Fourier transforms satisfy a number of general and quite useful properties. These properties give some nice insights on the behavior of Fourier transforms. They are often handy when doing analysis using Fourier transforms, and can also be helpful in using known transform pairs to compute new ones. The properties tell us how the Fourier transform changes when the original signal undergoes certain changes. They are generally derived by going back to the definitions and manipulating the equations appropriately.

Example 4.3 (Linearity)

One of the simplest but most useful property is that the Fourier transform is linear. That is, if $X(\omega)$ and $Y(\omega)$ are the Fourier transforms of two signals $x(t)$ and $y(t)$, respectively, then the Fourier transform of $ax(t)+by(t)$ is $aX(\omega)+bY(\omega)$. This is often written using the notation described earlier as

$$ax(t)+by(t) \longleftrightarrow aX(\omega)+bY(\omega)$$

where it is assumed as given that $x(t) \longleftrightarrow X(\omega)$ and $y(t) \longleftrightarrow Y(\omega)$.

To show linearity, we can substitute directly in the definition and use the fact that integration is linear. Namely,

$$\mathcal{F}\{ax(t)+by(t)\} \quad = \quad \int_{-\infty}^{\infty} [ax(t)+by(t)]e^{-j\omega t}\,dt$$

$$
\begin{aligned}
&= a\int_{-\infty}^{\infty} x(t)e^{-j\omega t}\,dt + b\int_{-\infty}^{\infty} y(t)e^{-j\omega t}\,dt \\
&= a\mathcal{F}\{x(t)\} + b\mathcal{F}\{y(t)\}
\end{aligned}
$$

■

Example 4.4 (Frequency shift property)

The so-called frequency shift property states that

$$e^{j\omega_0 t}x(t) \longleftrightarrow X(\omega - \omega_0)$$

This states that if a signal is multiplied by a complex exponential at some frequency ω_0, all that happens to the Fourier transform is that it gets shifted by the same frequency ω_0.

With a little thought, this result should be expected. If the complex exponential $e^{j\omega t}$ is multiplied by $e^{j\omega_0 t}$, then we simply get $e^{j(\omega+\omega_0)t}$. That is, we still have a complex exponential, but at the new frequency $\omega + \omega_0$. So, to determine how much content the signal $e^{j\omega_0 t}x(t)$ has at a given frequency ω, we need only check what content $x(t)$ has at the frequency $\omega - \omega_0$. This result is at the heart of certain techniques in communication known as modulation, which will be described in Chapter XX.

The frequency shift property can be obtained by simply combining the complex exponentials in the integral and then noticing that the resulting expression is precisely the definition of $X(\cdot)$ evaluated at the frequency $\omega - \omega_0$. Namely,

$$
\begin{aligned}
\mathcal{F}\{e^{j\omega_0 t}x(t)\} &= \int_{-\infty}^{\infty} \left(e^{j\omega_0 t}x(t)\right) e^{-j\omega t}\,dt \\
&= \int_{-\infty}^{\infty} x(t)e^{-j(\omega-\omega_0)t}\,dt \\
&= X(\omega - \omega_0)
\end{aligned}
$$

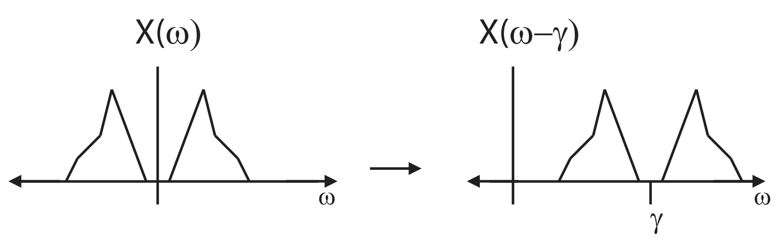

Figure 4.9: Frequency shift due to multiplication by complex exponential

■

Example 4.5 (Time shift property)
The time shift property states that

$$x(t - t_0) \longleftrightarrow e^{-j\omega t_0} X(\omega)$$

In words, shifting a signal in the time domain causes the Fourier transform to be multiplied by a complex exponential. Incidentally, this shows that if we view the Fourier transform as a mapping from input signals (the time-domain representation) to output signals (the frequency domain representation) then this mapping is *not* time-invariant, although as we saw above it *is* linear.

Because of duality, we might have guessed this result once we saw the frequency shift property above. One could prove the time shift property in a manner very similar to the way we proved the frequency shift property, but by manipulating the inverse transform equation instead of the forward transform equation.

However, a more direct method is to work with the forward transform and use a simple change of variables in the defining equation. Namely,

$$\begin{aligned} \mathcal{F}\{x(t-t_0)\} &= \int_{-\infty}^{\infty} x(t-t_0) e^{-j\omega t}\, dt \\ &= \int_{-\infty}^{\infty} x(u) e^{-j\omega (u+t_0)}\, du \qquad \text{letting } u = t - t_0 \\ &= e^{-j\omega t_0} \int_{-\infty}^{\infty} x(u) e^{-j\omega u}\, du \\ &= e^{-j\omega t_0} X(\omega) \end{aligned}$$

■

The general method used above (of writing the definition and using a change of variable and rearranging expressions) can be a very useful method for proving other properties as well. It may be reminiscent of how we proved some properties of the delta function.

A number of other properties of the Fourier transform are known, some of which are summarized in Table 4.2. We will discuss one more extremely important property in the next section.

4.8 Convolution Property and Another Look at Impulse and Frequency Response

The convolution property is so important that it deserves a separate section. Although it is not difficult to derive, it is a little more involved than the properties from the previous section.

The convolution property states that

$$x(t) * y(t) \longleftrightarrow X(\omega) Y(\omega)$$

Fourier Transform Properties
Linearity $ax(t) + by(t) \longleftrightarrow aX(\omega) + bY(\omega)$
Frequency Shift $e^{j\omega_0 t}x(t) \longleftrightarrow X(\omega - \omega_0)$
Time Shift $x(t - t_0) \longleftrightarrow e^{-j\omega t_0}X(\omega)$
Time Reversal $x(-t) \longleftrightarrow X(-\omega)$
Time Scaling $x(at) \longleftrightarrow \frac{1}{\lvert a \rvert}X(\frac{\omega}{a})$
Differentiation in Time $\frac{d}{dt}x(t) \longleftrightarrow j\omega X(\omega)$
Convolution in Time $x(t) * y(t) \longleftrightarrow X(\omega)Y(\omega)$
Multiplication in Time $x(t)y(t) \longleftrightarrow \frac{1}{2\pi}\int_{\gamma=-\infty}^{\infty} X(\gamma)\, Y(\omega - \gamma)\, d\gamma$

Table 4.2: Some Fourier Transform properties

That is, if two signals are convolved in the time domain then the Fourier transform of the convolution is just the product of the two original Fourier transforms. In other words, convolving signals in time domain is equivalent to multiplying their Fourier transforms in frequency domain. Thus, what appears to be a complicated operation in time, is simple if viewed in terms of frequency.

As with some others, this property is derived by using the definition of the Fourier transform and manipulating the expression. But, now we will have two integrals – one for the transform and one for the convolution.

$$
\begin{aligned}
\mathcal{F}\{x(t) * h(t)\} &= \mathcal{F}\left\{\int_{-\infty}^{\infty} x(\tau)h(t-\tau)\,d\tau\right\} \\
&= \int_{-\infty}^{\infty}\left(\int_{-\infty}^{\infty} x(\tau)h(t-\tau)\,d\tau\right)e^{-j\omega t}\,dt \\
&= \int_{-\infty}^{\infty} x(\tau)\left(\int_{-\infty}^{\infty} h(t-\tau)e^{-j\omega t}\,dt\right)d\tau \\
&\quad \text{by changing the order of integration} \\
&= \int_{-\infty}^{\infty} x(\tau)\left(\int_{-\infty}^{\infty} h(u)e^{-j\omega(u+\tau)}\,du\right)d\tau \\
&\quad \text{by a change of variables to } u = t - \tau
\end{aligned}
$$

$$
\begin{aligned}
&= \left(\int_{-\infty}^{\infty} x(\tau)e^{-j\omega\tau}\,d\tau\right)\left(\int_{-\infty}^{\infty} h(u)e^{-j\omega u}\,du\right) \\
&= X(\omega)H(\omega)
\end{aligned}
$$

Once we have the convolution property, duality suggests that multiplication in the time domain corresponds to convolution in the frequency domain. Specifically, this dual property is

$$x(t)y(t) \longleftrightarrow \frac{1}{2\pi}X(\omega) * Y(\omega) = \frac{1}{2\pi}\int_{-\infty}^{\infty} X(\eta)Y(\omega-\eta)\,d\eta$$

This property is called the multiplication property or the modulation property.

Recall that for an LTI system with impulse response $h(t)$, if we apply the input $x(t)$ then the ouput $y(t)$ is

$$y(t) = x(t) * h(t).$$

We can now ask what this output looks like in frequency domain. From the convolution property, we immediately see that

$$Y(\omega) = X(\omega)H(\omega).$$

Thus, the output of an LTI system can be viewed in either time domain or in frequency domain. In time domain, the output is the convolution of the input with the impulse response. In frequency domain, the LTI system does something particularly simple. It simply multiplies the FT of the input with the function $H(\omega)$, which is the FT of the impulse response.

But this is precisely how we described the frequency response of an LTI system in Chapter XX. We now see that the frequency response of an LTI system is just the Fourier transform of its impulse response. Compare Equation (XX) with the definition of the FT in Equation XX.

Although all of the properties in Table 4.2 are useful, the convolution result is *the* property to remember and is at the heart of much of signal processing and systems theory. We will further discuss and use the result throughout Chapter XX.

4.9 2-D Transforms*

So far, our discussion of Fourier transforms has focused on the case of 1-D signals – that is, signals as a function of time. Of course, nothing changes if the variable t were to represent a spatial variable instead of time. The crucial thing so far is that the signals were a function of just one independent variable.

To be able to apply the powerful tools of frequency domain analysis to images, mentioned in Section XX, we need to extend the Fourier transform to two dimensions.

In 2-D, the signal x is a function of two spatial arguments, which we will denote u and v. As discussed in Chapter XX, $x(u, v)$ denotes the value of the signal at the point (u, v).

We already saw in Section XX the notion of a sinusoid in two dimensions. For example,

$$x(u, v) = \sin(\omega_1 u) \sin(\omega_2 v)$$

is a 2-D sinusoid with frequency ω_1 in the u direction and frequency ω_2 in the v direction. This looks like an egg carton, or an array of hills and valleys. A surface plot and an image of this $x(u, v)$ are shown in Figure 4.10. In the image, the dark regions correspond to valleys while the bright regions correspond to hills.

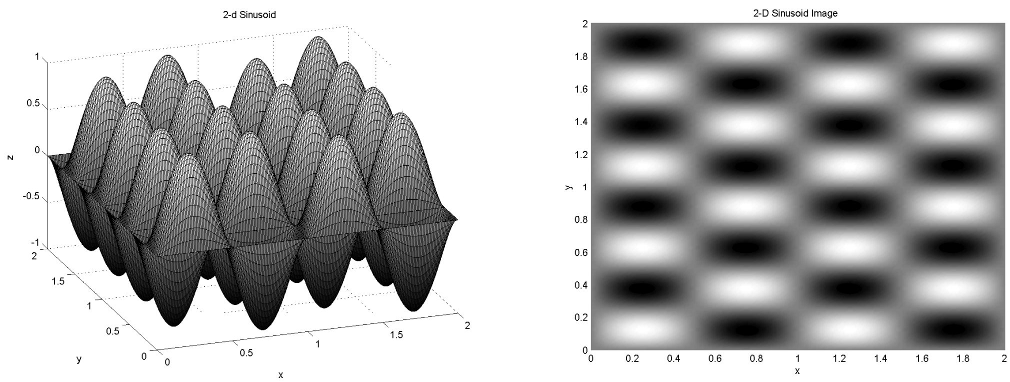

Figure 4.10: 2-d sinusoid surface plot (left) and image (right)

As in 1-D, it is more convenient to work with complex exponentials rather than the sinusoids above. Thus, the basic signals we will work with are of the form

$$e^{j\omega_1 u} e^{j\omega_2 v} = e^{j(\omega_1 u + \omega_2 v)}$$

This is a separable signal as defined in Chapter XX, in that it can be written as a product of a function of only u and a function of only v, with frequency ω_1 in the u direction and frequency ω_2 in the v direction.

One would expect that the frequency domain representation of a general signal $x(u, v)$ would be a function of two frequency variables, one for u and one for v. That is, the Fourier transform should be a function $X(\omega_1, \omega_2)$. The definitions of the forward and inverse transform in 2-D are

$$X(\omega_1, \omega_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x(u, v) e^{-j(\omega_1 u + \omega_2 v)} \, du \, dv \qquad \text{forward (4.6)}$$

$$x(u, v) = \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} X(\omega_1, \omega_2) e^{j(\omega_1 u + \omega_2 v)} \, d\omega_1 \, d\omega_2 \qquad \text{inverse (4.7)}$$

The definition of the DFT can be extended to 2-D in the natural way as well (as can Fourier series and the discrete-time Fourier transform). As in 1-D, specific transform pairs can be computed and most of the properties in 1-D have natural extensions to 2-D. When the original signal $x(u, v)$ is separable (i.e., can be written as $x(u, v) = x_1(u)x_2(v)$) then the 2-D transform is particularly simple since it reduces to two separate 1-D transforms – one in the variable u, and one in the variable v.

An example of a natural image and its FT is shown in Figure 4.12. As usual, only the magnitude of the transform is shown, where lighter areas denote higher magnitude for the corresponding frequencies, and darker areas denote smaller magnitude. Actually, since the variation in magnitude is so large, what is shown for the Fourier transform is the log of the magnitude. This allows seeing more detail in the transform as opposed to one extremely bright spot in the center and almost black everywhere else. Displaying the log of the magnitude is typical.

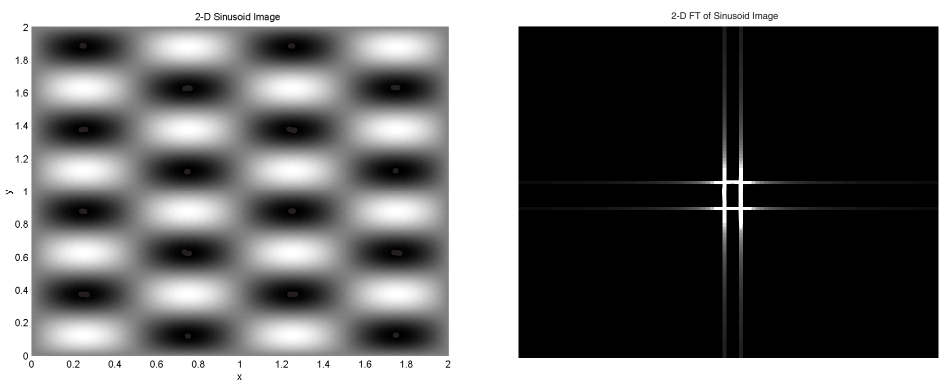

Figure 4.11: 2-D sine and its Fourier Transform

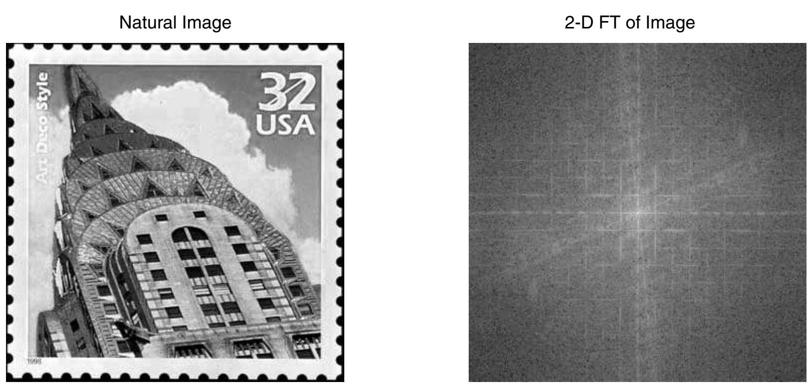

Figure 4.12: A natural image and its Fourier Transform