Signal and Information Processing

Alejandro Ribeiro

February 15, 2016

Contents

Chapter 1

Fourier Transforms

1.1 Continuous Time Signals - Inner product, norm, energy

Until now we have only dealt with **discrete** time signals. This is at odds with the reality of the universe. Time is **continuous** and never ceases. Time is not finite. Thus, it is necessary to have a mathematical framework to work with continuous time signals. In this chapter, we will explore the continuous extension of the Discrete Fourier Transform known as the Fourier Transform (FT).

1.1.1 Definition of continuous signal

Discrete time signals are defined for values at indices. $x : [0, N - 1] \to \mathbb{C}$. Notice that the domain of x(n) is integers ($n \in \mathbb{Z}$). The signal $x(n)$ is only defined for integer values of n. For continuous time analysis we fill in the gaps between samples. $x(t)$ is defined for all real values of t ($t \in \mathbb{R}$).

Put formally, a continuous time signal $x(t)$ maps values of $t \in \mathbb{R}$ to values of $x(t) \in \mathbb{C}$. Any real valued input has a corresponding output.

1.1.2 Inner Product

The inner product of two continuous signals is analogous to the discrete case - the only difference is the integral. This will be a common trend. In the discrete case we **sum**; in the continuous case we **integrate**.

$$\langle x, y \rangle := \int_{-\infty}^{\infty} x(t) y^*(t) dt \tag{1.1}$$

The concept of inner product as *relatedness* still holds. If $\langle x, y \rangle = 0$ the signals are **orthogonal** (completely unrelated to one another).

1.1.3 Norm and Energy

The norm of a signal x(t) is the square root of its energy. The energy of a signal is its norm squared. The energy of a signal is defined as:

$$\|x\|^2 := \int_{-\infty}^{\infty} |x(t)|^2 dt \tag{1.2}$$

The norm of a signal $x(t)$ is therefore:

$$\|x\| := \left[\int_{-\infty}^{\infty} |x(t)|^2 dt \right]^{1/2} \tag{1.3}$$

Using the property of complex numbers that $x(t)x^*(t) = \|x\|^2$ we can calculate the energy of a signal x(t) as **the inner product of the signal with itself**.

$$Energy\left[x(t) \right] = \langle x, x \rangle \tag{1.4}$$

1.1.4 Cauchy Schwarz Inequality

The largest an inner product can be is when the signals (vectors) are collinear

$$-\|x\| \, \|y\| \leq \langle x, y \rangle \leq \|x\| \, \|y\| \tag{1.5}$$

Written in terms of energy we get:

$$\langle x, y \rangle^2 \leq \|x\|^2 \, \|y\|^2 \tag{1.6}$$

1.1.5 Other notes

Let's define a signal x(t) seen in Figure 1.1. What does it mean to apply a constant τ to the input? Experimentation shows that the signal $x(t-\tau)$ is the original signal shifted to the right (positive) by τ (Figure 1.2). How can we arrive at this conclusion intuitively? Plugging in τ to $x(t-\tau)$ gives $x(0)$. Importantly, **We HAVE NOT changed the function** $x(t)$. Instead, we are shifting the **input** to the function. The function $x(t-\tau)$ is the same as the original but centered at τ. Likewise, the function $x(t+\tau)$ is the same as the original but centered at $-\tau$.

The original signal:

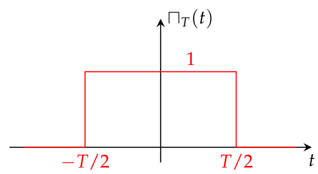

Figure 1.1: Energy Square Pulse of Duration T centered at $t = 0$

The signal with input shifted by $-\tau$:

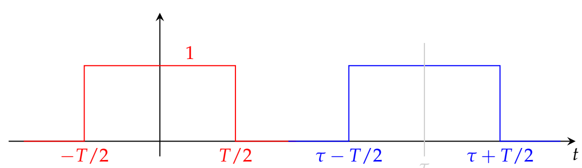

Figure 1.2: Energy Square Pulse of Duration T centered at $t = \tau$

1.2 Fourier Transform (FT)

From here on, the continuous Fourier Transform will be referred to as **FT** for brevity.

1.2.1 Definition of the Fourier Transform (FT)

Definition of FT:

$$X(f) := \int_{-\infty}^{\infty} x(t)e^{-j2\pi ft}dt \tag{1.7}$$

Let's try to gain some intuition about what the FT is doing. We are determining how much the input signal corresponds with complex exponentials of every possible frequency. When you hear "*how much signal x relates to signal y*" you should immediately think of the inner product. It turns out that we can express the FT as an inner product. If we define $e_f(t) := e^{2\pi ft}$ we can write the FT below. All expressions below are the **same**.

$$\mathcal{F}(x) = X(f) := \langle x, e_f \rangle = \int_{-\infty}^{\infty} x(t)e_f^*(t) = \int_{-\infty}^{\infty} x(t)e^{-j2\pi ft}dt \tag{1.8}$$

One important observation, we can't compute the FT using a computer. Integrals can't be computed as sums because we could never capture the entire real number line - the set of real numbers is "infinitely dense". As the slides note, the FT is an analytical tool, not a computational tool.

1.2.2 The sinc function

Refer to the slides for the calculation - the FT of a square pulse is the sinc function. The Inverse Fourier Transform (iFT - discussed in Section 3) of the sinc function is a square pulse. The FT of the sinc function is a square pulse (in frequency). The iFT of a square pulse is the sinc function (in time). The sinc function and the square pulse are **DUALS**. They are an indivisible pair - if you have one in the time domain, you have the other in frequency.

The sinc function is defined as:

$$sinc(u) = \frac{\sin(u)}{u} \tag{1.9}$$

Let's get a visual of the square pulse - sinc function duality

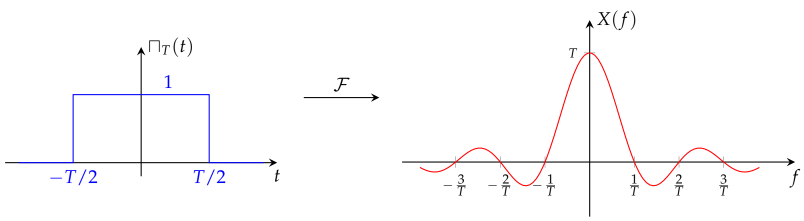

Figure 1.3: FT of square pulse is sinc function

Let's apply some intuition. **Wider square pulses give narrower sinc functions**. How do we know this? Say I have a square pulse of infinite duration. The pulse has a constant value over the domain $[-\infty, \infty]$. I want to use the FT to put this pulse in the **frequency** domain. The value of the function is constant in time for all time. Therefore, I can describe the entire signal using one frequency coefficient corresponding to $f = 0$ (the constant exponential). The FT of an infinite square pulse is just a single point (called a "delta" function) at $f = 0$. That is, I only need the frequency zero to describe the infinite square pulse because it never changes.

Let's say we have a finite square pulse that goes from high to low very quickly. Our intuition tells us: "The pulse changes quickly" $\rightarrow$ HIGH FREQUENCY. So the square pulse is high frequency. Therefore, we should have coefficients at high frequency values. When we plot the FT, high frequency values are farther away from zero, so the sinc pulse **must be wider** in this case.

1.2.3 The DFT and the FT

The FT is a great analytical tool. Given nice inputs we can do the calculus by hand and compute the FT. This hardly ever happens in reality - we need techniques that translate easily to computations. Let's explore how we can approximately compute the FT.

dt represents infinitesimal time increments that approach zero. We can't do real computations with infinitesimal increments. The first workaround is to use the fixed time increment T_s to separate our samples.

$$X(f) \;=\; \int_{-\infty}^{\infty} x(t)\, e^{-j2\pi f t} dt \;\approx\; T_s \sum_{-\infty}^{\infty} x(nT_s)\, e^{-j2\pi f n T_s} \tag{1.10}$$

This is better but still not computable. We need to handle the infinite summation. Let's restrict the sum to only consider N samples from 0 to $N-1$. We assign each of these samples occurring at nT_s to a value of k.

$$X(f) \;\approx\; T_s \sum_{k=0}^{N-1} x(nT_s)\, e^{-j2\pi f n T_s} \tag{1.11}$$

This is computable on the right hand side of the equation. $X(f)$, however, is still infinitely dense - we haven't specified which values of f we want to compute for, so we have to assume all of them. Let's restrict the values of $X(f)$ we need to compute to $f = (k/N)f_s$.

$$X\left(\frac{k}{N} f_s\right) \;\approx\; T_s \sum_{k=0}^{N-1} x(nT_s)\, e^{-j2\pi (k/N) f_s n T_s} \;=\; T_s \sum_{k=0}^{N-1} x(nT_s)\, e^{-j2\pi k n/N} \tag{1.12}$$

This is the definition of the DFT! We've shown how we can somewhat approximate the FT using the DFT. This is a very important result.

The graphic below demonstrates how the DFT can be used to approximate the FT. Notice that as the sampling time (T_s) approaches zero, and the number of samples (N) approaches ∞, the DFT approaches the FT.

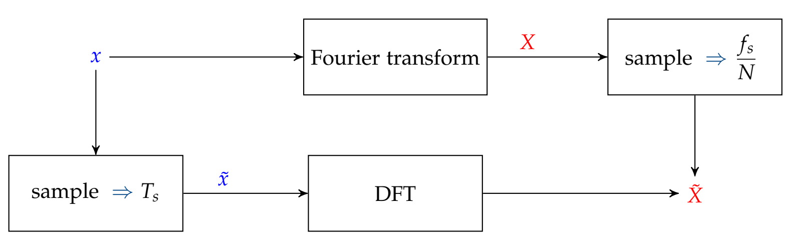

Figure 1.4: The DFT approximates FT as we approach limits of T_s and N

1.2.4 Notes on FT of complex exponential

The FT of a complex exponential does not exist because the complex exponential does not have finite energy. See the lecture slides for proof. However, if we truncate the complex exponential using a square pulse, the FT gives a shifted sinc function. This consideration is actually backwards. A standard property of the FT is that multiplying the original signal by a complex exponential in the time domain results in a shift in the frequency domain. In the original case, we multiply a square pulse in time by a complex exponential, which shifts the resulting sinc function in the frequency domain. Multiplying the input by a complex exponential of frequency f_0 shifts the output FT frequency by f_0. Basically, we replace f by $f - f_0$ in the output.

$$\tilde{X}(f) = T \cdot sinc(\pi(f - f_0)T) \tag{1.13}$$

1.3 Inverse Fourier Transform (IFT)

From here on, the continuous Inverse Fourier Transform will be referred to as **IFT** for brevity.

1.3.1 Definition of the Inverse Fourier Transform (IFT)

Definitions of the IFT:

$$x(t) := \int_{-\infty}^{\infty} X(f)\, e^{j2\pi f t} df = \mathcal{F}^{-1}(X) \tag{1.14}$$

We can also view this as an inner product with a complex exponential of frequency $-t$. This interpretation will not be that useful to us, but serves as a good refresher on inner products.

$$x(t) := x(t) = \langle X, e_{-t} \rangle \tag{1.15}$$

We can see that the IFT is, indeed, the inverse of the Fourier Transform. See lecture slides for a detailed proof.

1.3.2 Frequency Decomposition of a Signal

The IFT allows us to conceptualize a signal as a sum of oscillations, as did the DFT. This interpretation allows us to see the frequency composition of a signal: that is, does the signal change slowly (concentrated around $f = 0$) or quickly (concentrated away from $f = 0$)?

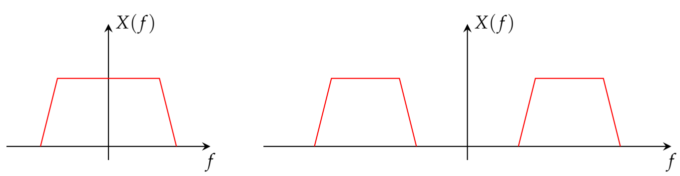

Figure 1.5: The left signal changes slowly, and the right changes quickly.

1.3.3 Fourier Transform Pairs

We have already seen that the FT of a square pulse is a sinc pulse. The inverse is also true. That is, the square pulse and the sinc pulse are a Fourier Transform pair. This is due to the symmetry between the FT and IFT formulas (only a minus sign changes). In general, these pairs will be conjugate symmetric. However, with real signals, you don't have to worry about this.

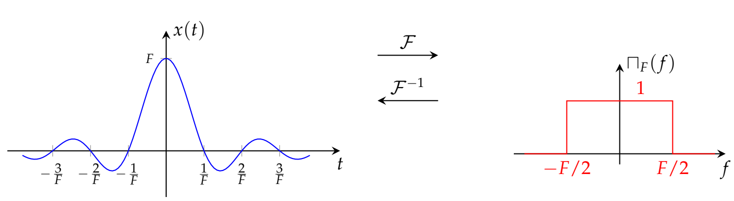

Figure 1.6: We can see that this is the dual of Figure 1.3.

1.4 Delta Function

We have already seen the discrete delta function, which took a value of 1 at $t = 0$ and 0 everywhere else. We will investigate the analogous function in continuous time.

1.4.1 Sequence of Progressively Taller Sinc Pulses

We define the continuous delta function as the limit of a sinc pulse.

$$\delta(t) = \lim_{F \to \infty} Fsinc(\pi F t) \tag{1.16}$$

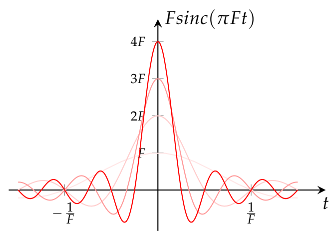

Figure 1.7: We can see that, at the limit, this becomes an infinitely tall and narrow spike.

Intuitively, the delta function, then, is infinity when F is 0 and null everywhere else. However, we want to be a little more mathematically rigorous about it. What would happen if we integrate? We know that essentially all the "mass" of this signal is concentrated at $F = 0$. So if we integrate the product of a signal and this delta function, we get:

$$\int_{-\infty}^{\infty} x(t) Fsinc(\pi F t) dt \approx x(0) \tag{1.17}$$

Taking the formal limit, we see that:

$$\lim_{F \to \infty} \int_{-\infty}^{\infty} x(t) Fsinc(\pi F t) dt = x(0) \tag{1.18}$$

We recognize the integrand as an inner product of x with the delta function as we have just defined it. We can then redefine the delta function using this property:

$$\langle x, \delta \rangle = x(0) \tag{1.19}$$

In terms of integrals, this is:

$$\int_{-\infty}^{\infty} x(t) \delta(t) dt = x(0) \tag{1.20}$$

So this integral definition of the delta function agrees with our intuition, except for one key difference. δ takes a value of infinity at 0 and is null everywhere else. Elementary calculus would suggest that the area under this curve would be 0. But it's not! The area under this curve must be 1 for the integral definition to work.

So let's summarize. The delta function, when integrated against a signal, outputs only the 0th element of that signal. This weird function has a value of infinity at 0, and is null everywhere else, yet its area is 1. Why? The short answer is because math is weird and that's just how it works out.

Of course, we can't really call the delta function a "function" if it breaks the fundamental properties of functions that we've learned since calculus. It's actually an abstract, generalized distribution.

It does not exist in real life, but once we pass it through an integral we can observe its effect on other signals.

1.5 Generalized Orthogonality

With our new knowledge of delta functions, we will revisit one of the most important concepts of this course: orthogonality.

We have seen how continuous time and discrete time are analogues of one another. Let's think about what would happen if we wanted to take the inner product of two continuous time complex exponentials with frequencies f and g.

We know that the inner product tells us how much two signals are like each other (love, hate, and indifference). Taking the discrete analog and moving into continuous time, we *define* the inner product of these complex exponentials to be a delta function.

$$\langle e_f, e_g \rangle = \delta(f - g) \tag{1.21}$$

Note that this is just a definition at this point. But if we think about it, we see why it makes sense. If $f = g$, the complex exponentials are "infinitely" related (i.e. they are the same), and if $f \neq g$, then they are "indifferent" - that is, orthogonal. As we have already seen, the orthogonality of complex exponentials is very important and makes our lives substantially easier.

Again, the delta function has some tricky mathematical technicalities. We can't technically take the inner product of two complex exponentials, because it doesn't exist. To be formal, we can truncate the exponentials and take the limit. See the lecture slides for a thorough treatment of this. However, for an intuitive sense of how the delta function behaves, this will suffice.

1.6 Generalized Fourier Transforms

1.6.1 Fourier Transform of a Complex Exponential

We're going to again start with a definition, and then think about why it makes sense. What is the FT of a complex exponential? If we think about the analogous structure between the FT and the DFT, we should already know the answer.

$$X_g(f) = \delta(f - g) \tag{1.22}$$

Let's think about what this means. A FT is an inner product, and an inner product tells us how related two things are. So the FT should be all zeros (because *complex exponentials are orthogonal*) except for the one frequency that would match that of the complex exponential. Indeed, this is the case.

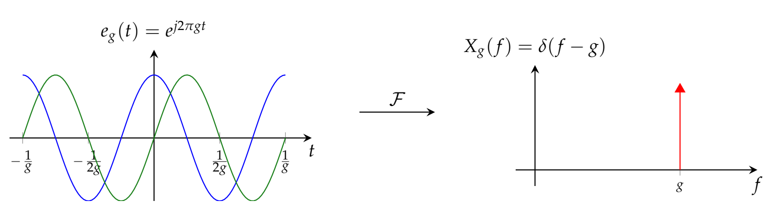

Figure 1.8: What we've just described is the exact definition of a delta function that we've constructed.

Furthermore, this definition is consistent with the IFT.

$$e_g(t) = \int_{-\infty}^{\infty} X_g(f)e^{j2\pi ft}df = \int_{-\infty}^{\infty} \delta(f-g)e^{j2\pi ft}df = e^{j2\pi gt} \tag{1.23}$$

So, even if we don't want to get too deep into the math, we can intuitively reason why the delta function allows us to define the FT of a complex exponential.

1.6.2 Fourier Transform of a Delta Function

Let's think back to what we already know about the duality of Fourier transform pairs. We just saw that the FT of a complex exponential is a shifted delta function. What, then, would be the FT of a shifted delta function?

$$X_u(f) = \int_{-\infty}^{\infty} \delta(t-u)e^{-j2\pi ft}dt = e^{-j2\pi fu} = e_{-u}(f) \tag{1.24}$$

As our intuition would suggest, its a complex exponential with frequency u.

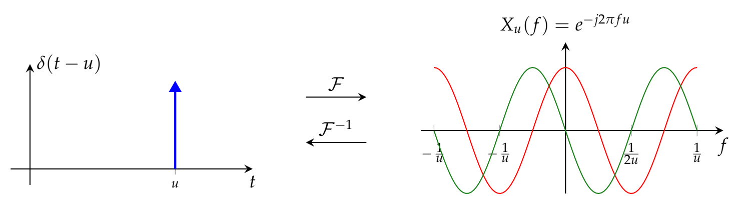

Figure 1.9: Why does it make sense that the FT of a delta function is a complex exponential with frequency u?

To make this idea clearer, let's consider a particular case: an unshifted delta function. We know that the FT of an unshifted delta function should be a complex exponential with frequency $f = 0$. What does this look like visually?

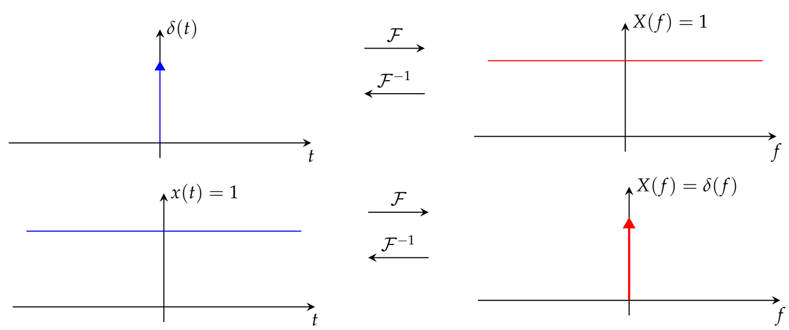

Figure 1.10: A complex exponential with frequency 0 is a constant, which is what we would expect to see.

1.6.3 Fourier Transform of a Cosine

We recall that the DFT of a discrete cosine was symmetric pair of discrete deltas. Following the development of information we've seen thus far, what do we think the FT of a cosine will be?

Let's use the fact that we can write a cosine as a pair of complex exponentials and that the FT is a linear operator.

$$\cos(2\pi gt) = \frac{1}{2}\left[e^{j2\pi gt} + e^{-j2\pi gt}\right] \tag{1.25}$$

$$X(f) = \frac{1}{2}\left[\delta(f-g) + \delta(f+g)\right]$$

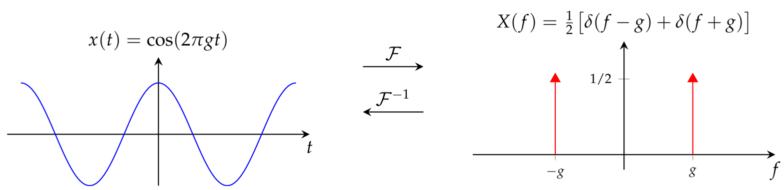

Figure 1.11: Again, we get what we expected.

1.7 Properties of the Fourier Transform

Many of the properties of the FT are identical to those of the DFT. This should make sense, since the DFT is an approximation to the FT. We will briefly reiterate some of the properties we've already covered, but will not go into too much detail. To see proofs, refer to the lecture slides.

Conjugate Symmetry:

$$X(-f) = X^*(f) \tag{1.27}$$

Linearity:

$$\mathcal{F}(ax + by) = a\mathcal{F}(x) + b\mathcal{F}(y) \tag{1.28}$$

Parseval's Theorem (Energy Conservation):

$$\int_{-\infty}^{\infty} |x(t)|^2 dt = \|x\|^2 = \|X\|^2 = \int_{-\infty}^{\infty} |X(f)|^2 df \tag{1.29}$$

1.7.1 Shift and Modulation

Shift and modulation also exist for the DFT, but are much more difficult to see. They are dual operations of each other, much like the FT and IFT. We will see that multiplying a signal by a complex exponential in time is equivalent to shifting it in the frequency domain. The inverse is also true, as we would expect. Shifting a signal in time is equivalent to multiplying it by a complex exponential in the frequency domain.

$$x_\tau = x(t-\tau) \qquad \Longleftrightarrow \qquad X_\tau(f) = e^{-j2\pi f\tau}X(f) \tag{1.30}$$

We can see that while the *phase* of the signal in time changes, its *modulus* doesn't. Why? Because a complex exponential has modulus 1!

$$|X_\tau(f)| = |e^{-j2\pi f\tau}X(f)| = |e^{-j2\pi f\tau}| \times |X(f)| = |X(f)| \tag{1.31}$$

So why is this important? Let's think about real-world applications. If we want to input, say, a voice recording (e.g. Siri), do we want the user to have to start speaking exactly at time $t = 0$? No, we want the user to be able to start talking at any point. This property tells us that we can treat a time-shifted signal (e.g. start talking at time $t = 2$ seconds) the same as if it were not shifted at all, only by multiplying by a complex exponential in frequency!

When we refer to modulation, we are really just referring to the dual of this property. That is, you can shift a signal's spectrum (frequency composition) simply by multiplying by a complex exponential in time.

$$x_g = e^{j2\pi gt}x(t) \qquad \Longleftrightarrow \qquad X_g(f) = X(f-g) \tag{1.32}$$

1.7.2 Modulation of Bandlimited Signals

A *bandlimited signal* is one that has its entire spectrum within a certain range, called the bandwidth, W. That is, $X(f) = 0$ for $f \notin [-W/2, W/2]$. Since all of the frequency information is contained within a certain range, if we shift (modulate) this range, we can have two signals on the same frequency axis with *no overlap*.

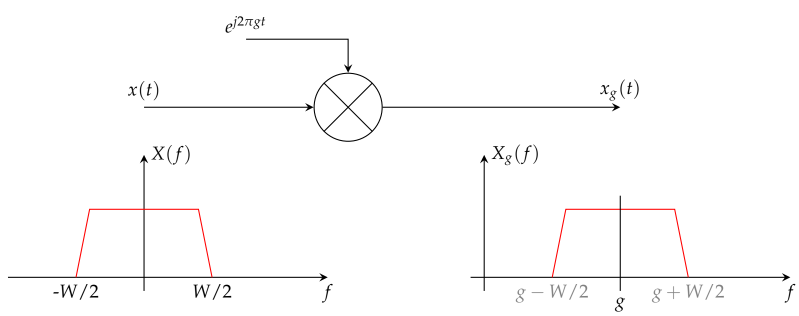

Figure 1.12: A bandlimited signal (left) can be modulated (put) anywhere on the frequency axis by multiplying by a complex exponential in the time domain.

It makes sense, then, that to recover our original signal from the modulated version, we can just multiply by the conjugate complex exponential in time. That is, multiply by $e^{-j2\pi gt}$ in time.

Let's think about an application where this would be useful. What would happen if we wanted to send two phone calls through the same telephone line? We couldn't just smash them together, because then we wouldn't be able to tell which is which. So, assuming the phone calls are *bandlimited* (we will test this assumption later), we can modulate them such that their spectra do not overlap. Then, there will be no mixup of signal information, and we can simply demodulate them at the end!

1.8 Convolution

We will start this section with a soundbyte that is fundamental to this course: *convolution in time is equivalent to multiplication in frequency* (and vice versa).

The convolution of two signals $x(t)$ and $h(t)$ is defined as follows:

$$[x * h](t) = y(t) = \int_{-\infty}^{\infty} x(u)h(t-u)\, du \tag{1.33}$$

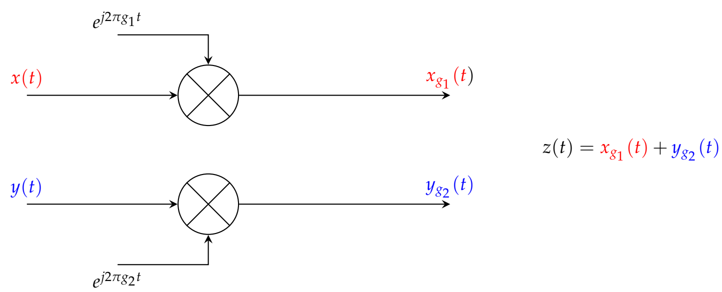

Figure 1.13: Multiply by complex exponentials in time to modulate frequency spectra.

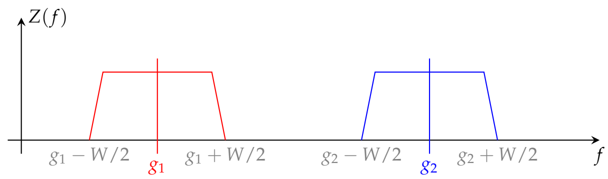

Figure 1.14: Now there will be no mixing of information between the signals, and we can simply multiply by the respective conjugate complex exponentials to recover the original signals.

Although this is a symmetric operation, it is easier for us to think of it asymmetrically. That is, we say that x *hits* h.

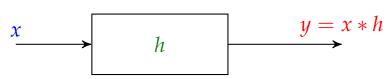

1.8.1 Convolution with a Delta Function

Based on our definition of the delta function above, we should already have a good idea about how this convolution behaves.

$$y(t) = \int_{-\infty}^{\infty} \delta(u)h(t-u)\,du = h(t) \tag{1.34}$$

As we would expect, it just gives us back our original signal $h(t)$. What would happen if we convoluted with a scaled delta function? Since the convolution is a linear operator, we would expect to just get a scaled version of $h(t)$. This is indeed the case.

$$y(t) = \int_{-\infty}^{\infty} \alpha\delta(u)h(t-u)\,du = \alpha \int_{-\infty}^{\infty} \delta(u)h(t-u)\,du = \alpha h(t) \tag{1.35}$$

What about a convolution with a shifted delta function? Again, as we would expect, the result is just a shifted version of $h(t)$.

$$y(t) = \int_{-\infty}^{\infty} \alpha\delta(u-s)h(t-u)\,du = \alpha \int_{-\infty}^{\infty} \delta(u-s)h(t-u)\,du = \alpha h(t-s) \tag{1.36}$$

This leads us to a more intuitive interpretation of this abstract convolution operation. We can view the convolution as a conglomeration of the scale, sum, and shift operations across the time axis.

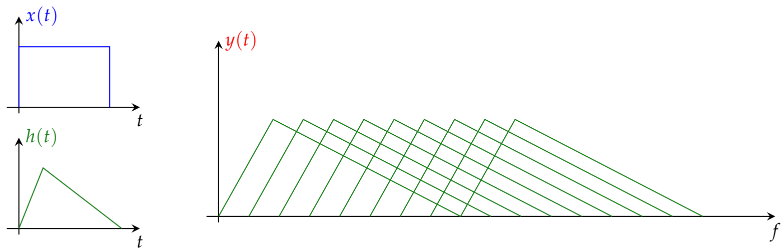

Figure 1.15: The convolution can be viewed as a scale, shift, and sum of a signal across the input domain.

We'll reiterate the convolution theorem because of its importance: *convolution in time = multiplication in frequency*.

$$z = x * y \qquad \Longleftrightarrow \qquad Z = XY \tag{1.37}$$

For a detailed proof of the convolution theorem, see the lecture slides. Otherwise, remembering and understanding the conclusion is sufficient.

We've heavily emphasized the importance of this theorem. Why? Well, this theorem essentially says that we can take a hard operation (convolution) and turn it into an easy operation (multiplication) solely by changing domains. So if we want to design a system, we can do so in the frequency domain, where things are nicer. Then, we can perform the actual implementation in the time domain (a computer doesn't care how hard the convolution is).

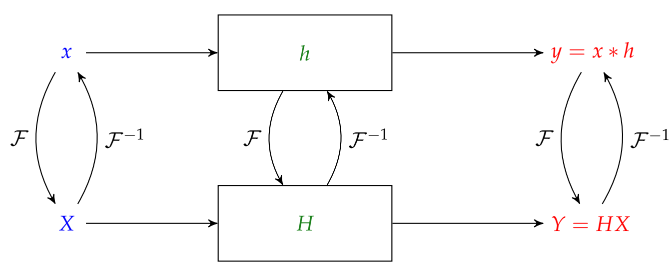

Figure 1.16: We take the bottom path for *design* and the upper path for *implementation.*

.

1.9 The Signal and the Noise

Now, let's take a look at one final example that ties together everything we've learned so far. Let's return to the noisy signal we've seen before.

We saw that to "clean up" this noisy signal, we took its Fourier transform to move into the frequency domain, implemented a low-pass filter $H(f) = \sqcap_W(f)$ with $W = 200$Hz in the frequency domain, and then took the inverse Fourier transform to return to the time domain. This worked, but we can now make this approach more efficient.

We know that we want to implement this same low pass filter. But instead of doing so in the frequency domain and then taking the IFT, let's just implement it directly in the time domain. How do we do this? We convolute our time signal with the IFT of our filter.

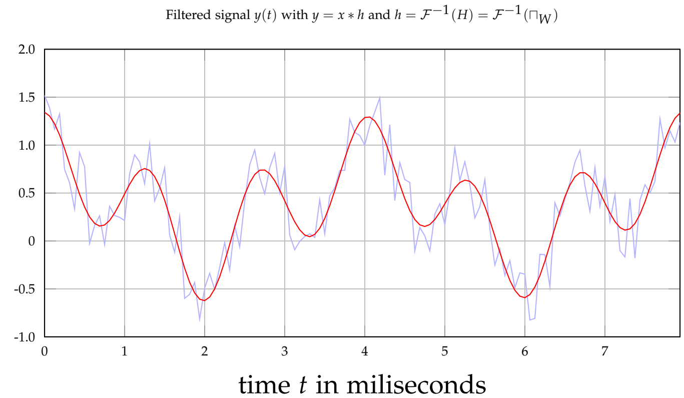

We've already seen that the IFT of a square pulse is a sinc pulse. So let's convolute our signal with a sinc pulse to achieve this same result.

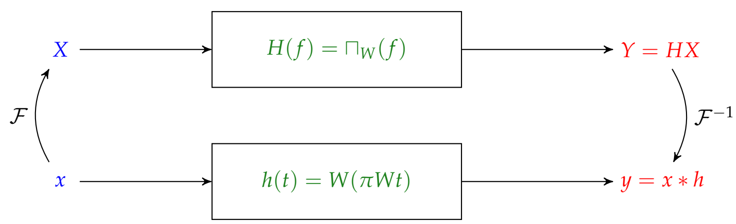

Figure 1.17: We design in the frequency domain, then implement in the time domain.

We will see the importance of this if we want to have a system that can detect signals inputted in real time, as we will do in lab.