Digital Communication using GNU Radio

Prof. Kumar Appaiah

Department of Electrical Engineering

Indian Institute of Technology Bombay

Week-10

Lecture-48

Parallelising Frequency Selective Channels

Welcome to this lecture on Digital Communication Using GNU Radio. I am Kumar Appiah from the Department of Electrical Engineering at IIT Bombay. In today's lecture, we will introduce the concept of wireless communication and delve into OFDM, which stands for Orthogonal Frequency Division Multiplexing.

(Refer Slide Time: 02:07)

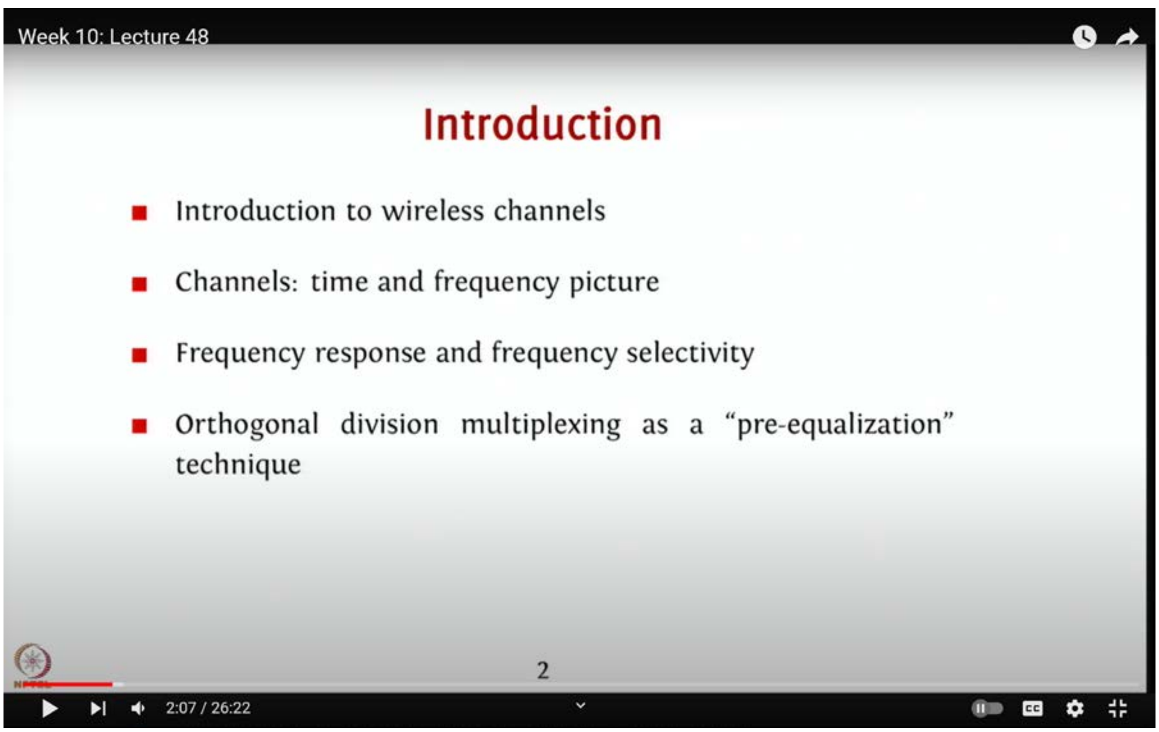

To provide context, in the previous lecture, we discussed the impact of a channel's impulse response on system performance. A significant challenge arises when the channel introduces distortions that require equalization at the receiver. Now, what if we could simplify this process? Imagine applying some degree of equalization at the transmitter itself, before the signal is even sent, and doing so without prior knowledge of the channel characteristics. This is where OFDM shines, it enables us to parallelize the channels, leading to a much simpler equalization process at the receiver.

(Refer Slide Time: 03:31)

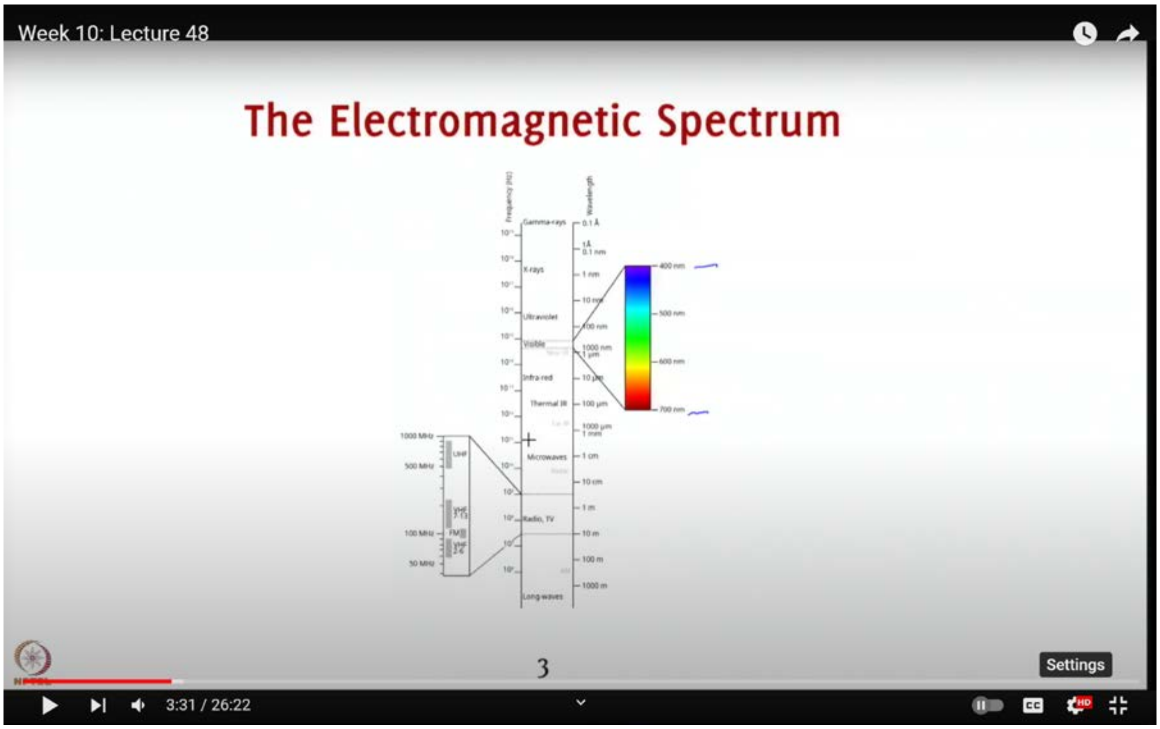

We will explore this in greater detail, but first, let me motivate the discussion with a basic channel model for wireless communication systems. It’s worth noting that this model isn’t limited to wireless systems; the same approach can be applied to a variety of communication channels. These include optical channels, underwater communication systems, fiber-optic cables, free-space communications, and more. Wireless communication is simply our point of focus today, but the principles extend to other systems as well.

The core idea is that we will examine the channel behavior from both time and frequency domain perspectives. We will then discuss the concepts of frequency response and frequency selectivity, which describe how the channel behaves differently across various frequency ranges. Finally, we will introduce Orthogonal Frequency Division Multiplexing (OFDM) as a pre-equalization technique, effectively simplifying the equalization process at the receiver.

(Refer Slide Time: 04:11)

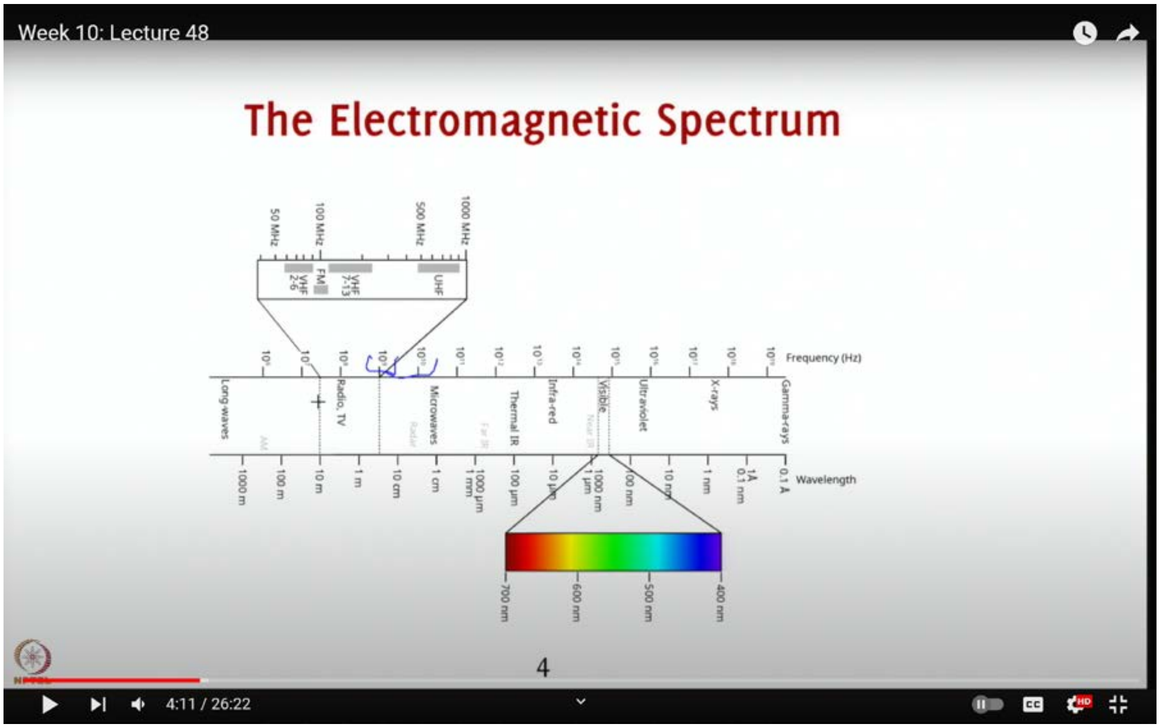

This lecture will allow us to gain a closer look at how OFDM aids in managing frequency-selective channels, enabling a more efficient communication process.

If you recall from your knowledge of the electromagnetic spectrum, this particular chart is sourced from Wikipedia. You can verify that the visible light spectrum lies within the range of 400 to 700 nm. So, between 400 and 700 nm, we have what we perceive as visible light. Now, when you look at the spectrum either in terms of frequency or wavelength, it helps you understand the kinds of signals you might expect across different parts of the spectrum. Below 400 nm, the signals have extremely high frequencies and correspondingly short wavelengths. These regions of the spectrum are where ultraviolet (UV) and X-rays dominate.

However, most of our communication systems operate in the infrared region. For instance, in optical communication, you typically work with wavelengths around 1.5 microns or even 10 to 100 microns. As you shift further down the spectrum, moving into the domain of tens of megahertz (MHz) and kilohertz (kHz), you enter the frequency ranges where most wireless communication takes place.

(Refer Slide Time: 05:31)

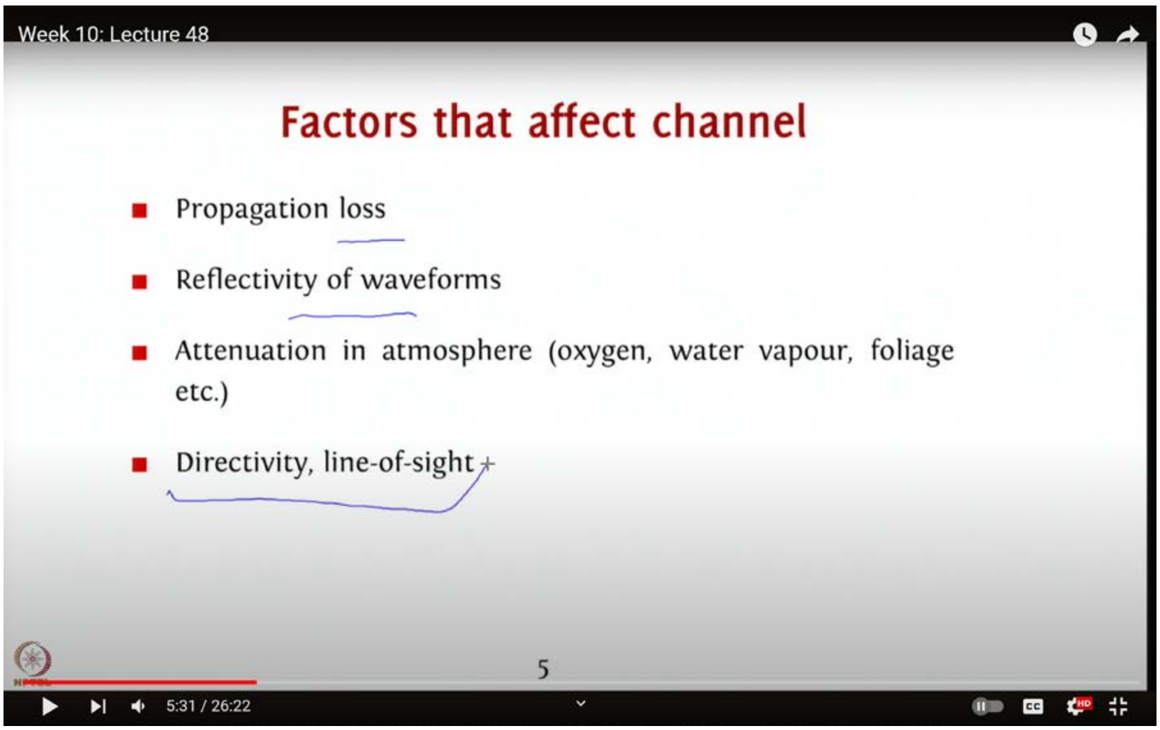

Let me give you a broader perspective. If you examine the frequency spectrum more closely, particularly in the gigahertz (GHz) range or even down to tens of kilohertz, you'll find this is where systems like radar, radio, television, and FM radio operate. In this range, propagation through the air is very effective, meaning that you can transmit signals at these frequencies over long distances with relatively minimal energy loss.

Now, why do we choose specific frequencies for communication, or why might one frequency be preferred over another? There are several factors at play.

One consideration is regulatory policy. Governments may dictate which channels or frequencies are available for use, and in many cases, you might even have to pay for access to a particular frequency band. Beyond policy, you must also consider the propagation characteristics of the signal, such as how much energy is lost as the signal travels through the air at a given frequency or wavelength.

(Refer Slide Time: 08:21)

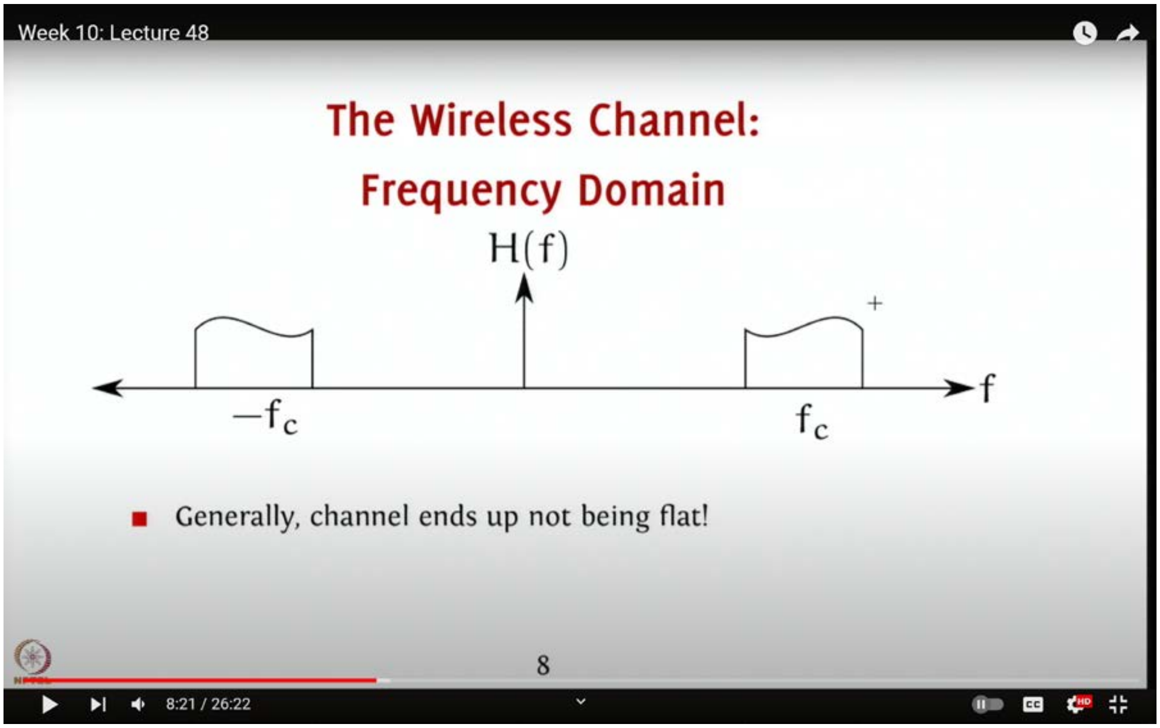

Another key factor is the reflectivity of the waveforms. Depending on the medium through which the signal travels, reflections may either be beneficial or detrimental. Attenuation in the atmosphere is another concern, while propagation loss gives you a general sense of signal degradation, specific atmospheric conditions can affect certain wavelengths more than others.

For example, some wavelengths are more heavily impacted by the presence of atmospheric oxygen, while others are more susceptible to absorption by water vapor or even the presence of foliage. The environment in which you plan to deploy your communication system will influence the choice of wavelength, as some wavelengths are far more optimal depending on the conditions of the surrounding atmosphere.

Finally, another critical factor you might be interested in is the directivity or line of sight between the transmitter and receiver. For some systems, line of sight isn't necessary. A good example of this is traditional wireless communication systems like GSM or 4G-based mobile devices. These devices work perfectly fine indoors, even if there is no direct visible path between the handset and the base station. The signal is robust enough to penetrate walls and other obstructions without requiring a clear, unobstructed line of sight.

(Refer Slide Time: 09:32)

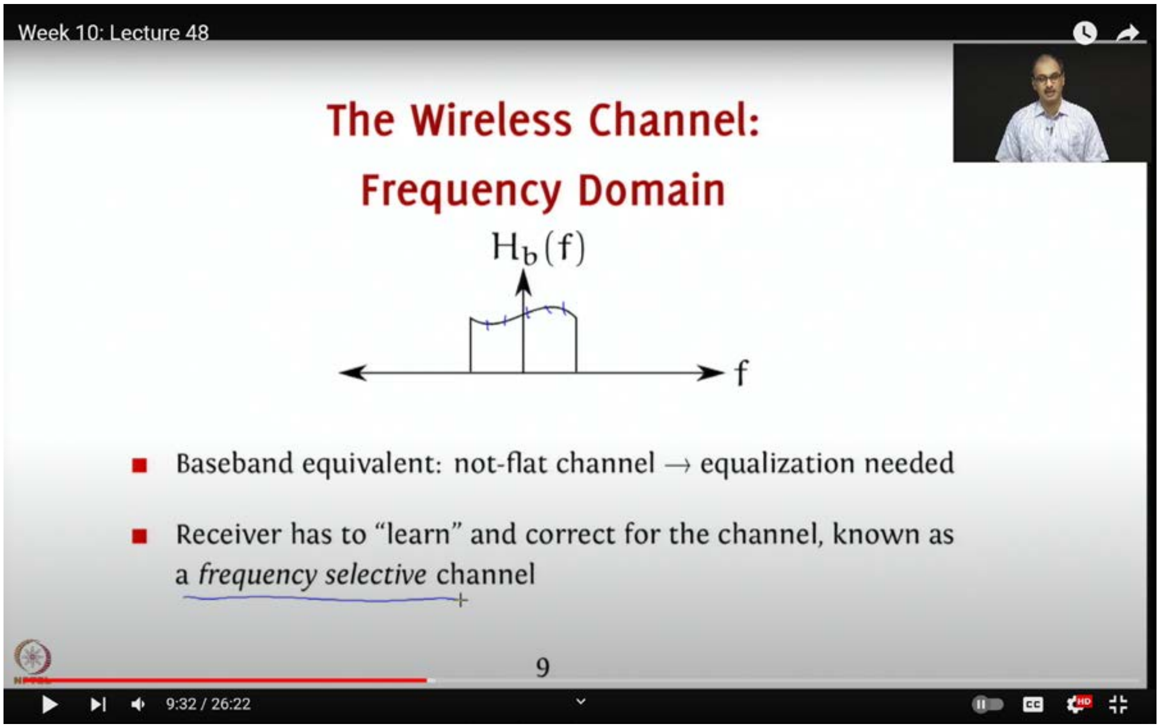

However, there are modern systems, particularly those operating in the millimeter-wave range, where line of sight is essential for effective communication. In such systems, both the transmitter and receiver need to "see" each other directly in order to maintain a reliable connection. This is a design consideration that must be factored in when working with such high-frequency systems.

Now, let's move to the frequency domain picture. So far, we've been focusing on the time domain, where issues like inter-symbol interference (ISI) have been prominent. In the time domain, if there's significant ISI, the channel no longer behaves like an ideal impulse channel. Its frequency response becomes distorted, no longer remaining flat across the spectrum.

This time, however, we'll shift our focus to the frequency response picture directly, as it will help us better understand the pre-equalization technique we're discussing. Choosing the appropriate frequency range becomes crucial here. You have to consider factors such as the frequency range you are licensed or permitted to transmit on, the acceptable levels of signal loss, and the compatibility of the transmitter and receiver.

(Refer Slide Time: 12:41)

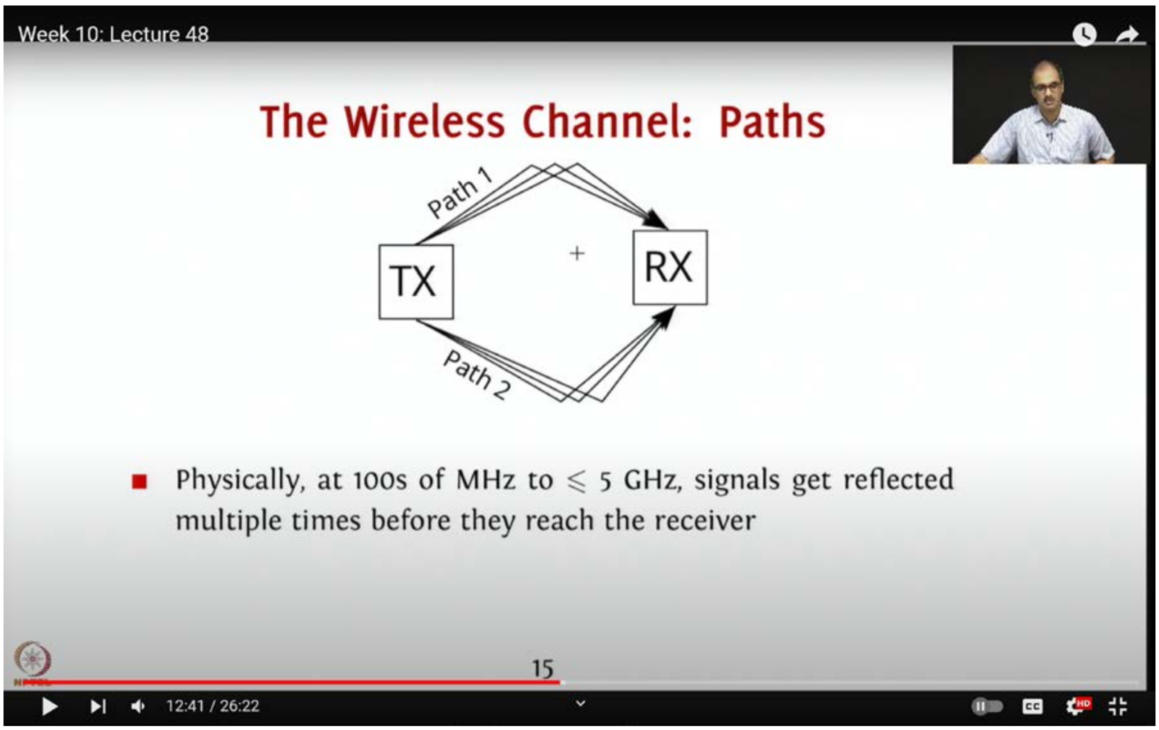

Once you've settled on a frequency range, say around a central frequency f_c, and a bandwidth $\pm\frac{W}{2}$, this becomes the band at which you will transmit your signal. Now, when you start transmitting, let's assume we are working with a baseband signal. A flat spectrum would imply no data is being transmitted. However, as soon as you introduce data, whether using a sinc pulse, root-raised cosine pulse, or any other waveform, the spectrum will no longer be flat. Ideally, you want the channel to have a flat frequency response to ensure smooth data transmission, but in practice, this is rarely the case without proper equalization techniques in place.

You might feel confident that your data will pass through without any issues, but the reality is that the channel often doesn't behave as a flat one. Why? Because certain frequencies experience reflections and losses more than others. Some frequencies get reflected back, causing distortions, so you rarely encounter a truly flat channel in practice. A non-flat channel is essentially equivalent to dealing with inter-symbol interference or similar effects.

(Refer Slide Time: 16:28)

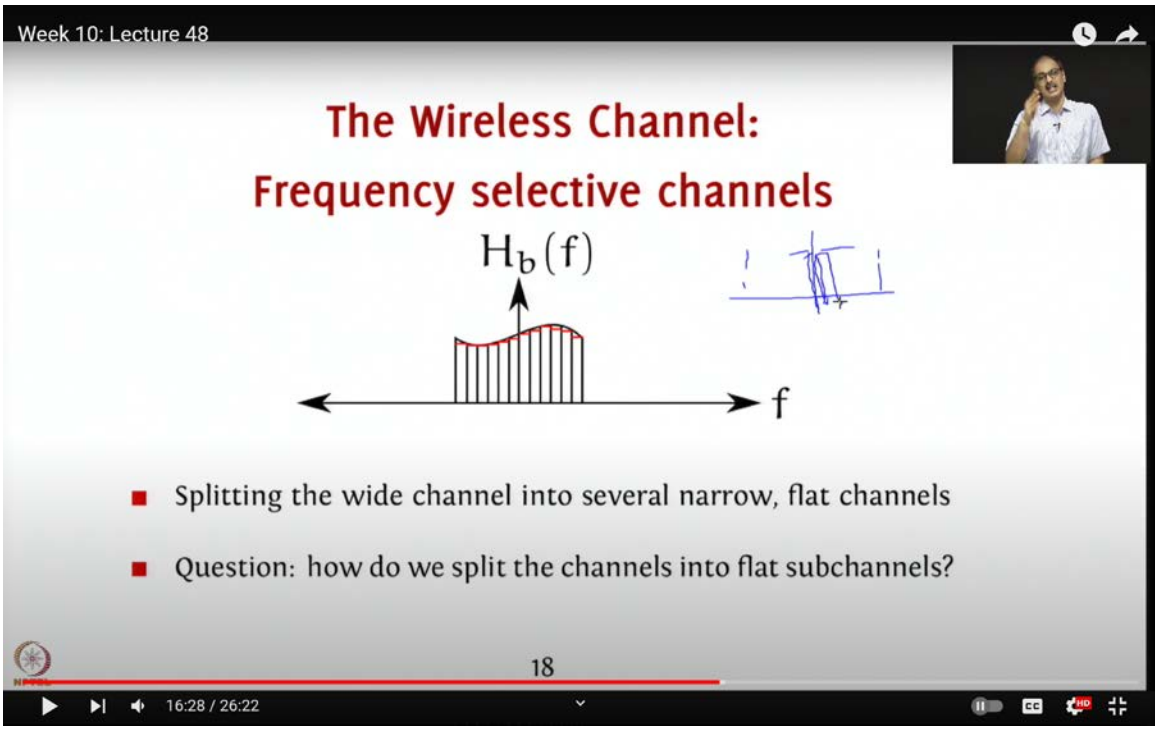

Ideally, we hope the channel exhibits a linear transformation, but as we've discussed in previous lectures, when the channel transforms your waveform, it requires equalization, essentially undoing the effects of the channel by applying the appropriate filtering. In this context, we need to recognize that a non-flat channel response introduces challenges that need to be mitigated through equalization.

Let's take a look at this in the baseband picture. After passing your signal through the channel at frequency f_c, you downconvert it back to baseband. Now, instead of a flat frequency response, the channel is distorted. One solution would be to use linear equalizers like MMSE (Minimum Mean Square Error), zero forcing, or one of the adaptive equalization techniques to compensate for the channel effects. The receiver needs to learn the channel's behavior and adjust accordingly to restore performance.

In cases where the channel affects different frequencies in different ways, we refer to this as a frequency-selective channel. This term is commonly used when describing channels that perform some non-trivial filtering. If the channel merely scales the input uniformly across frequencies, it is not frequency selective. But if it behaves differently across the spectrum, it is considered frequency selective.

The behavior of frequency-selective channels must be carefully analyzed and learned. Let's consider a typical wireless channel operating in the range of several hundred MHz to a few GHz, say, from 800 MHz to around 2.5 or 3 GHz. In fact, we could even extend this range up to 5 GHz. Within this range, signals often undergo multiple reflections before finally reaching the receiver, introducing significant complexity in how the channel shapes the transmitted signal.

The same model essentially applies to other frequencies as well, though the specific impacts may vary, either becoming more pronounced or less significant depending on the frequency range. Let's consider a practical example. Imagine you're using your cell phone inside a building, and it's not directly facing the base station. Here's your base station, and here's your cell phone.

Typically, the signal doesn't travel directly to the phone; instead, it gets reflected off various surfaces. These reflections can come from a nearby building, a wall, a tree, or even a passing car. Some of these objects, like buildings, are static, while others, like vehicles, are in motion. Now, an important thing to keep in mind is that even static objects aren't entirely stationary. For instance, a short building may seem immovable, but a very tall building can sway slightly due to wind or other factors. This movement may only be a few millimeters or centimeters, and while that may seem insignificant at first glance, in the context of communication signals, even small movements can be meaningful.

Consider this: if you're dealing with signals that have wavelengths on the order of tens of centimeters, which is common for many wireless communication frequencies, these shifts matter. For example, at 1 GHz, the wavelength is approximately 30 centimeters. You can verify this, but what it means is that even millimeter-scale movements in the environment can alter the signal's propagation.

Moreover, you typically don’t just get a single reflection; you get multiple reflections from different sources or even from the same source but at different points in time. This results in what we call multipath propagation. So, the signal reaching your device may include not only a direct line-of-sight transmission (if one exists) but also several reflected copies from various surfaces in the environment. These reflected signals can interfere with each other, sometimes reinforcing the signal and sometimes causing fading or distortion.

The challenge with these reflected copies is that they take varying amounts of time to arrive at the receiver. For instance, the first copy may reach you quickly, while the second arrives with a slightly different gain and delay, and the third arrives with yet another variation in gain and delay. Each of these reflected signals adds up, contributing to an effective impulse response that the channel produces for your system. Our task is to handle this impulse response, learn its characteristics, and correct for it to ensure optimal system performance. This has been a recurring theme in our discussions.

In most cases, you won’t just have two paths; there might be 10, 12, or even 20 paths depending on the system and environment. Alternatively, there could be a single dominant path, and you may be able to ignore the rest. Different scenarios exist, and each has to be managed appropriately.

This brings us to the concept of frequency selectivity. The non-uniform gain across different frequencies arises due to the presence of these multiple paths. If you only had a single path, you would be dealing with a flat channel, one where the gain is consistent across the entire frequency range. However, when you have multiple delayed paths, each with different gains, you are dealing with a frequency-selective channel. So how do we handle this?

Yes, we can use techniques like zero forcing, MMSE (Minimum Mean Square Error), or adaptive equalization to correct the channel’s behavior. But what if we approached it differently? What if we could break down the channel into smaller segments? Imagine dividing the channel's bandwidth, let's say from -W/2 to W/2 in the baseband, into several smaller, narrowband channels.

This division offers a significant advantage. When you split the channel into several narrow segments, you can make an approximation: if each segment is narrow enough, you can assume the gain is approximately constant within that range. Now, of course, one could argue that the gain might not be perfectly constant, but if you make the segments sufficiently narrow, this approximation holds reasonably well. This is because the channel's characteristics are influenced by natural processes, which typically vary smoothly across the frequency domain rather than exhibiting abrupt changes.

By breaking the channel into smaller parts, we simplify the problem, making it easier to model and correct for the variations introduced by multipath propagation.

(Refer Slide Time: 18:11)

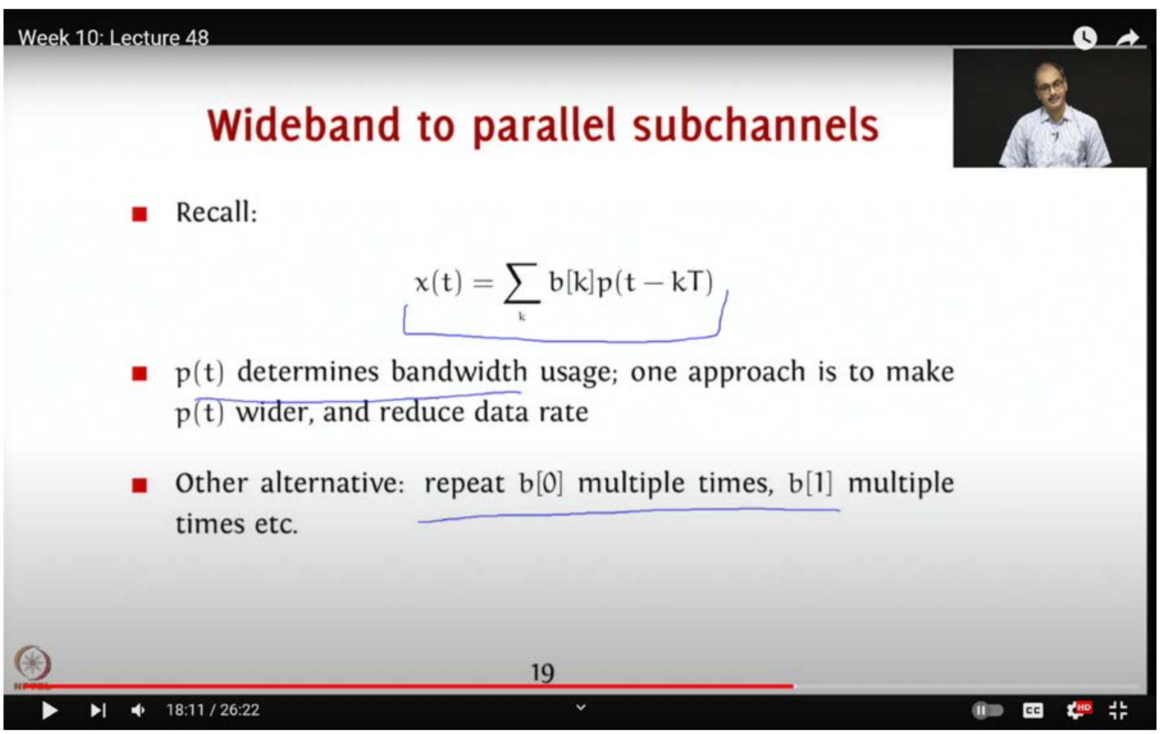

If you manage to divide the channel into smaller segments, within each of these narrow bands, you can safely assume that the gain remains relatively constant. This stability in gain eliminates the need for the kind of complex equalization that compensates for inter-symbol interference and other distortions. Why is that? Well, when you have a wide bandwidth, equalization becomes necessary to correct the varying gain across that range. However, by narrowing the bandwidth significantly, the channel's variation in the frequency domain becomes minimal, it essentially flattens out. As a result, the assumption that the channel is flat within these narrow bands becomes a reasonable and effective approximation. In the time domain, this translates to an impulse response that closely resembles a delta function, no longer plagued by multiple paths or significant distortions.

By splitting a frequency-selective channel into several frequency-flat channels, you achieve a simpler and more manageable situation. This seems like a clever technique, but how do you implement it? The intuitive way to approach this is by using much wider pulses. For instance, if you were to use a sinc pulse, a raised cosine pulse, or a root-raised cosine pulse, you would choose them based on the range -w/2 to w/2. However, doing so means that you must then employ another sinc or root-raised cosine pulse of similar width but operating at a slightly different neighboring frequency.

This method is essentially what we call frequency division multiplexing (FDM). Instead of occupying a broad channel, you would focus on narrowband channels. You would select one narrowband channel, and then another, and another, each centered around a slightly different frequency f_c. You would continue this process, shifting them slightly in the frequency domain. This technique effectively allows you to achieve the desired outcome, breaking down the wideband channel into manageable frequency slices.

Now, think back to how we typically signal: x(t) was expressed as a summation, $x(t) = \sum b_k p(t - kT)$, where p(t) represented the effective pulse that encapsulates both the transmitted pulse and the channel effects. This pulse p(t) determines the bandwidth usage. Remember, the bandwidth is dictated by g_{tx} (t), the transmit filter, because g_{tx}(t), when convolved with the channel, cannot expand the frequency footprint. This is a fundamental property of linear time-invariant (LTI) systems, convolution does not increase the bandwidth.

So, one approach to managing this is to make the pulse p(t) wider, as I mentioned earlier.

This adjustment allows you to work within narrower bands and avoid the complexities associated with a broader channel, while still maintaining effective communication.

When you make p(t) wider and wider, in the frequency domain it becomes narrower and narrower. However, this approach can be quite cumbersome because it requires creating multiple parallel streams of data that must be modulated precisely to fit into specific frequency slots. Essentially, you would need multiple radio frequency translations to position each segment correctly, which can be complex and unwieldy. Is there a simpler solution?

(Refer Slide Time: 20:00)

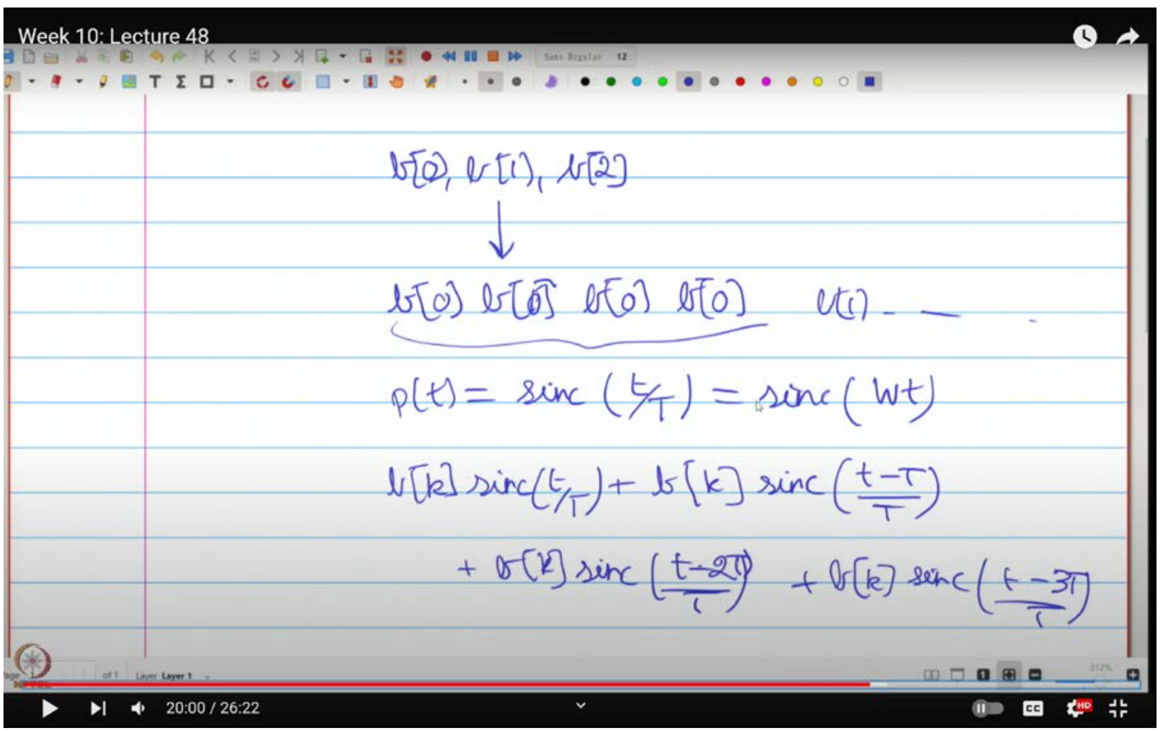

Yes, there is a more straightforward approach. Instead of creating multiple parallel streams, why not simply repeat each symbol multiple times? This might sound a bit unconventional, but the idea is to use the same p(t), which spans the bandwidth from -w/2 to w/2, and repeat each symbol b_0, b_1, and so on multiple times.

How does this work? Let's walk through a hypothetical example. Suppose you have a sequence of symbols, such as b_0, b_1, b_2, etc. Now, if you repeat b_0 four times, then b_1 four times, and so forth, you effectively create a new sequence where each symbol is repeated multiple times in a periodic pattern.

You might argue that this reduces the data rate to one-fourth of the original, and that's a valid point. However, this approach simplifies the frequency usage pattern and can be more manageable. We will address the reduction in data rate and explore how to handle it effectively in due course. For now, this method provides a simpler way to manage bandwidth and channel utilization.

Let's examine what happens to the spectrum when we repeat b_k four times. Suppose p(t) is a sinc function given by $\mathrm{sinc}(t/T)$, where T is the symbol duration and w = 1/T.

For our example, if we have a sequence where each b_k is sent four times, this translates to a signal composed of $b_k \mathrm{sinc}(t) + b_k \mathrm{sinc}(t-T) + b_k \mathrm{sinc}(t-2T) + b_k \mathrm{sinc}(t-3T)$. This results in an equivalent pulse p'(t) defined as:

$$p'(t) = \mathrm{sinc}(t/T) + \mathrm{sinc}\big((t-T)/T\big) + \mathrm{sinc}\big((t-2T)/T\big) + \mathrm{sinc}\big((t-3T)/T\big)$$

(Refer Slide Time: 21:32)

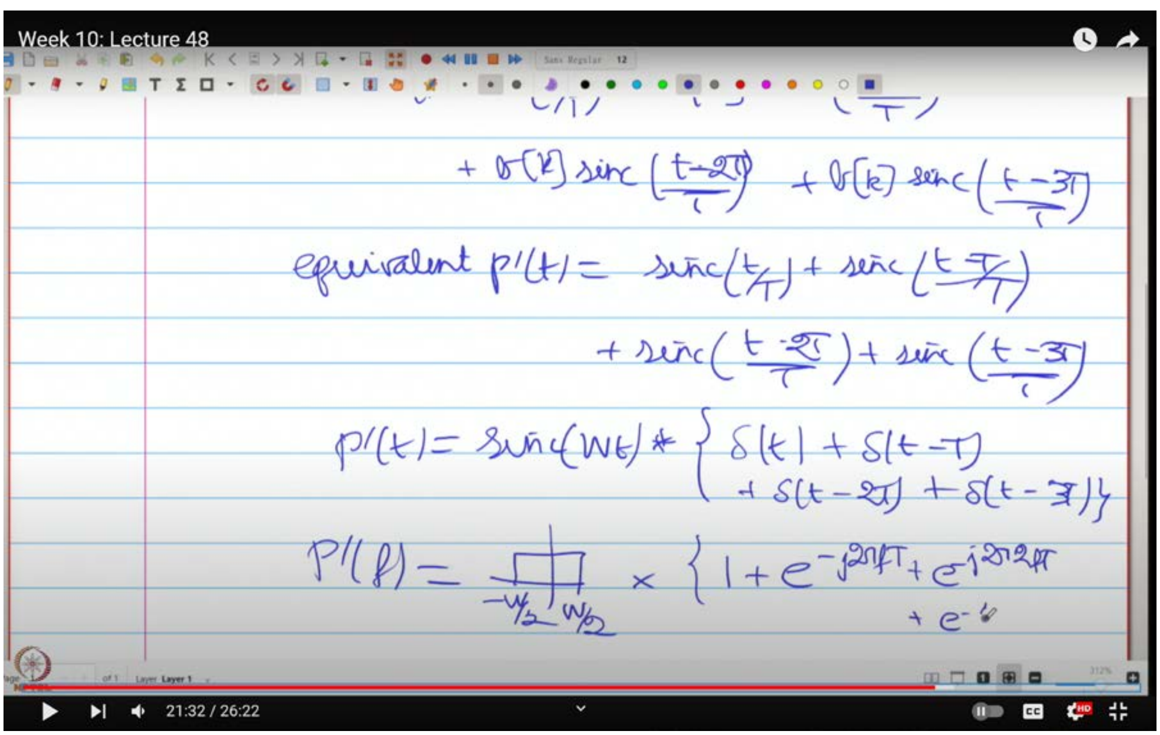

To analyze the frequency domain effect, we need to compute the Fourier transform of p'(t). The Fourier transform of a sinc function is a rectangular function, so:

$$p'(f) = \mathrm{rect}(f/w) * [\delta(f) + \delta(f - 1/T) + \delta(f - 2/T) + \delta(f - 3/T)]$$

Where $\mathrm{rect}(f/w)$ is a rectangular function spanning from -w/2 to w/2. Taking the convolution with the delta functions results in:

$$P'(f) = \mathrm{rect}(f/w)\left[1 + e^{-j2\pi fT} + e^{-j4\pi fT} + e^{-j6\pi fT}\right]$$

(Refer Slide Time: 23:51)

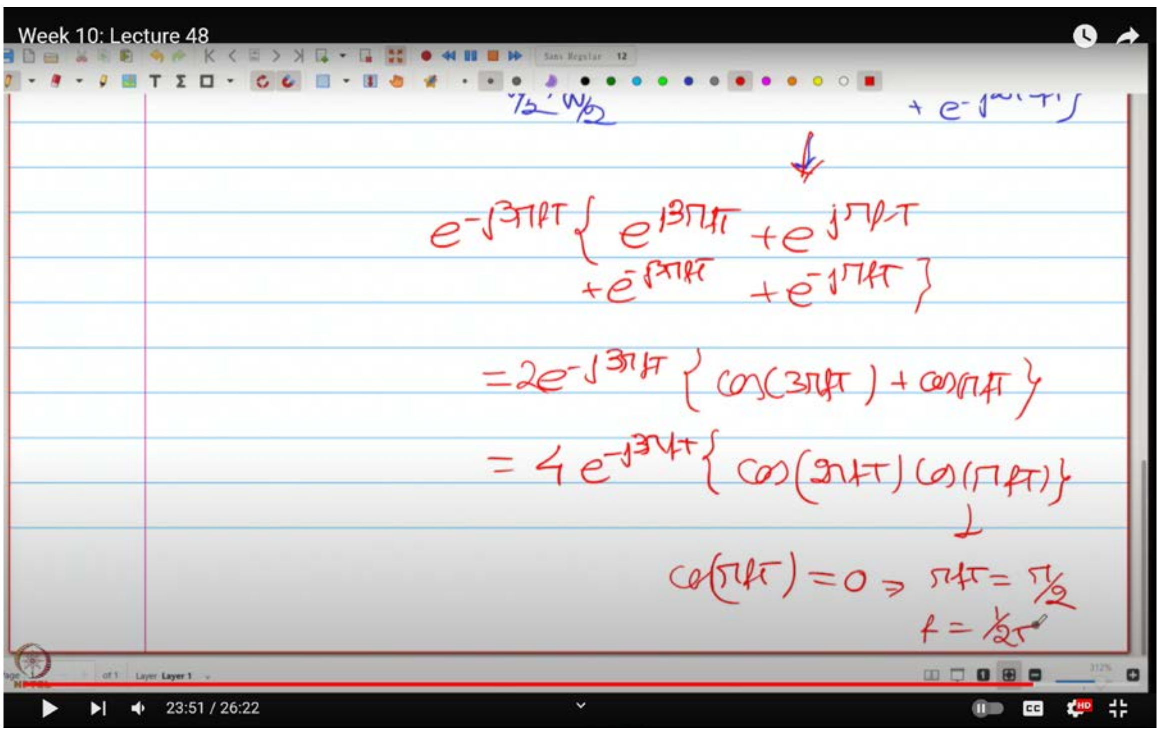

To simplify, factor out $e^{-j6\pi fT}$:

$$P'(f) = \mathrm{rect}(f/w)\left[e^{-j6\pi fT}\left(e^{j6\pi fT} + e^{j4\pi fT} + e^{j2\pi fT} + 1\right)\right]$$

$$P'(f) = \mathrm{rect}(f/w)\left[e^{-j6\pi fT}(2\cos(3\pi fT) + 2\cos(\pi fT))\right]$$

Using the trigonometric identities, this simplifies further to:

$$P'(f) = \mathrm{rect}(f/w)[2(\cos(3\pi fT) + \cos(\pi fT))]$$

The spectrum of p'(f) indicates that the spectrum is now periodic with zeros at certain frequencies. Specifically, $cos(\pi f T)$ has zeros where $\pi f T = \frac{\pi}{2}$, which simplifies to:

$$f = \frac{1}{2T}$$

(Refer Slide Time: 25:10)

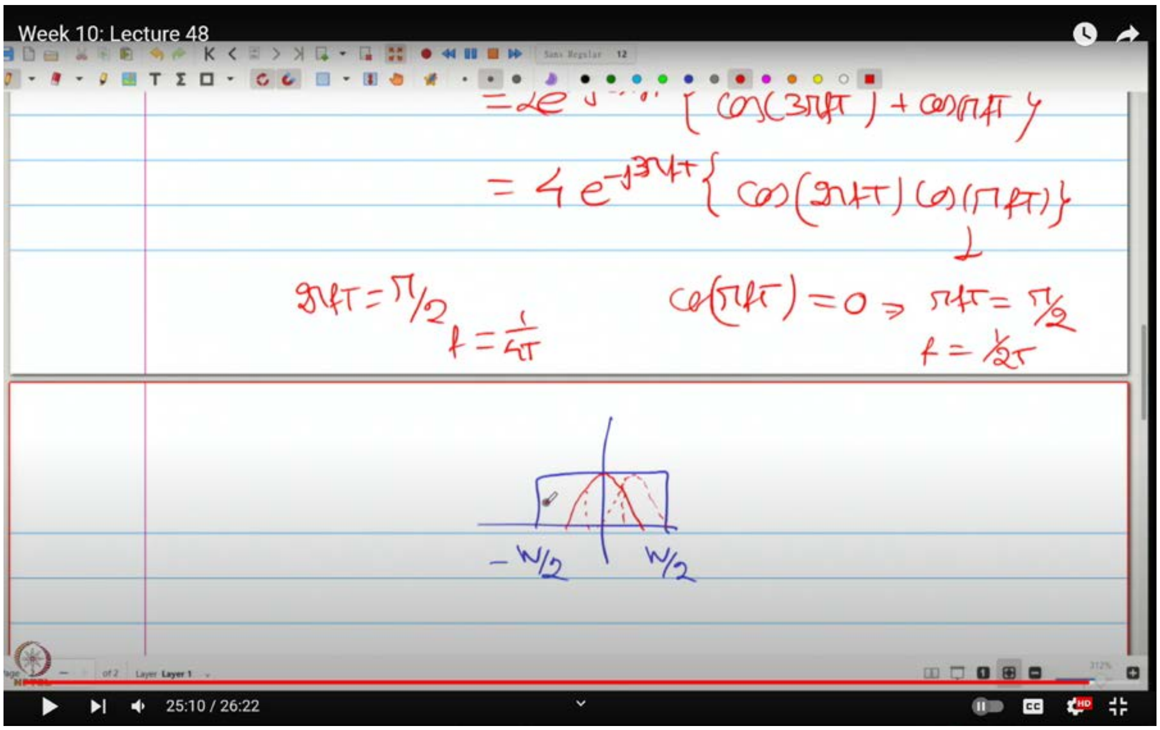

Similarly, for $cos(2\pi f T)$, zeros occur where:

$$f = \frac{1}{4T}$$

So, the spectrum of p'(t) will have its main lobe between -w/2 and w/2, but with additional nulls or zeros at specific frequencies based on the period of the cosine functions. This results in a modified spectrum that is concentrated within a specific range, with the rest of the spectrum approaching zero. This technique effectively divides the bandwidth into smaller, narrower bands, simplifying the problem of frequency-selective channels.

It's as if you're utilizing only a quarter of the available spectrum. The key idea here is that by using additional spectral components to occupy adjacent frequency bands, you can effectively cover four separate frequency ranges. This concept leads us to Orthogonal Frequency Division Multiplexing (OFDM).

The central principle is that if you repeat your symbol four times, you essentially reduce the spectral footprint to one-fourth of the original width. This aligns with concepts from digital signal processing (DSP), where repeating signals decreases their spectral usage. Although it’s not precisely one-fourth in practice, it’s a close approximation that proves effective. Notably, interference-related issues are minimized.

In the next lecture, we will formalize this concept and explore how it evolves into Orthogonal Frequency Division Multiplexing. This approach offers significant advantages for the equalization process at the receiver. Thank you.