Digital Communication using GNU Radio

Prof. Kumar Appaiah

Department of Electrical Engineering

Indian Institute of Technology Bombay

Week-04

Lecture-20

Detection and Optimal Decision for On-Off Signaling in AWGN Channel

Hello, and welcome back to our lecture series on Digital Communication using GNU Radio. I am Kumar Appiah from the Department of Electrical Engineering at IIT Bombay. In this session, we will continue our discussion from the previous lecture, where we explored the topic of optimal reception in an additive white Gaussian noise (AWGN) channel.

(Refer Slide Time: 01:14)

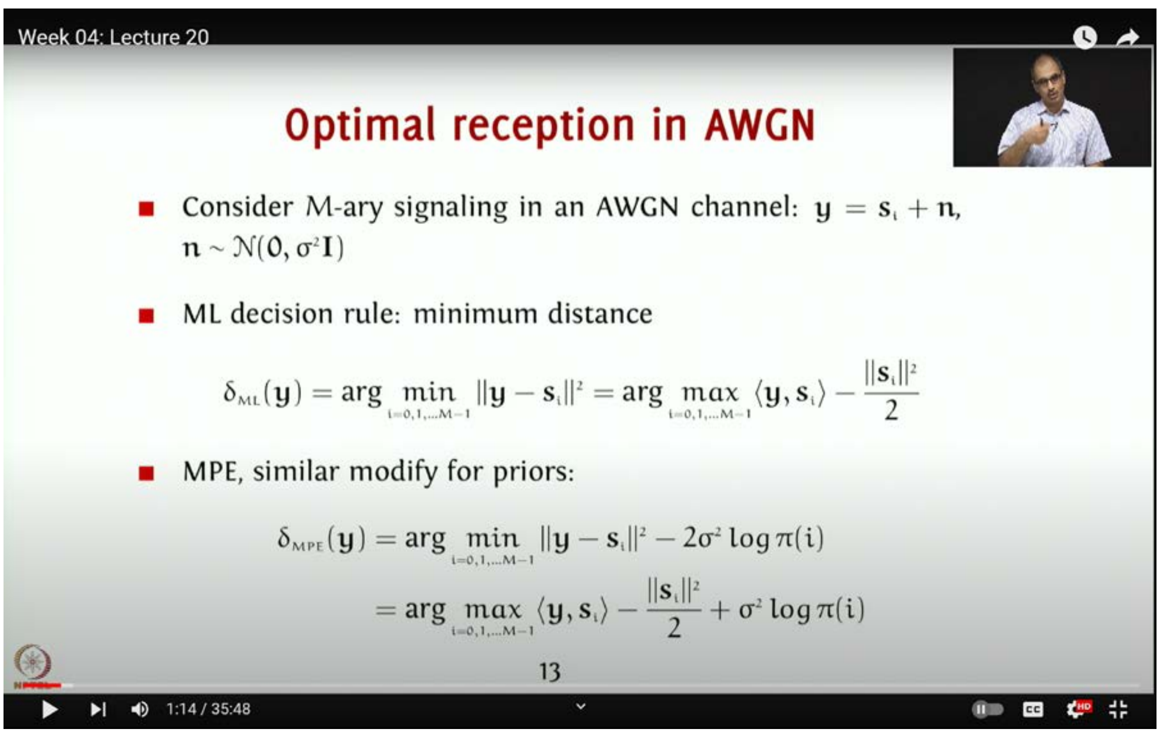

To recap briefly, we have been examining binary signaling, and before we dive deeper into that, it is essential to revisit some of the foundational tools we developed earlier.

Specifically, let us consider M-ary signaling within an AWGN channel. Here, the received vector y is modeled as y = s_i + n, where s_i represents the transmitted signal, and n is the noise.

Now, how do we arrive at these vectors? We derive them by projecting the signals onto the basis functions, effectively allowing us to treat the signals as vectors in a multidimensional space. Importantly, the noise that affects our signaling is also projected onto these basis signals. The resulting noise vector is what directly influences detection. Thanks to the theorem of irrelevance, the component of noise orthogonal to the signal space, denoted as $n_{\perp}$, does not interfere with our detection process. Thus, the relationship simplifies to y = s_i + n, where n is a zero-mean Gaussian noise vector with a covariance matrix σ^2 I. This means that each component of the noise has variance σ^2, and these components are independent and uncorrelated.

From this setup, we derived the maximum likelihood (ML) detection rule. The essence of ML detection lies in maximizing the likelihood function, which in this case is Gaussian. The key observation was that the relationship between y and s_i appears only in the exponent, specifically in the term $e^{-|y-s_i|^2}$. Therefore, maximizing the likelihood function is equivalent to minimizing the squared Euclidean distance between y and s_i. Consequently, the ML detection rule reduces to a minimum distance detection rule.

Additionally, we explored a simplification where the detection criterion can be expressed in terms of inner products. By expanding $|y - s_i|^2$, we can eliminate the common term $|y|^2$ and focus on maximizing the inner product $\langle y, s_i \rangle$ minus $|s_i|^2$. This alternative form is often convenient and interchangeable with the minimum distance rule, depending on the specific problem.

Now, when considering the minimum probability of error detection, an extra term related to the prior probability π_i of the transmitted symbols comes into play in the likelihood function. This is because the prior probabilities affect the detection rule, contributing an additional logarithmic term. If the prior probabilities $\log \pi_i$ are the same for all symbols i, the decision rule simplifies back to the ML detection. Therefore, under the assumption of equally likely transmitted symbols, the ML and minimum probability of error detection rules are equivalent.

(Refer Slide Time: 05:27)

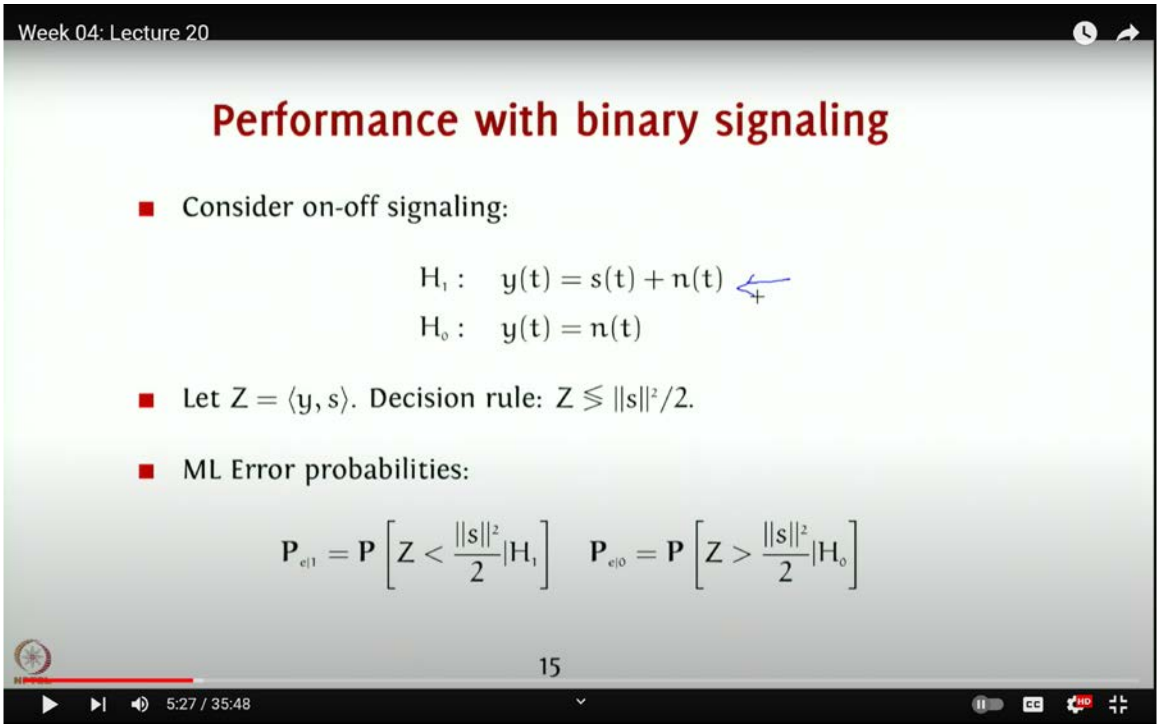

To illustrate the significance of considering prior probabilities, imagine a scenario where the probability of transmitting a '1' is much higher than that of transmitting a '0'. For instance, if the probability of sending a '1' is 0.9 and the probability of sending a '0' is only 0.1, the detection strategy must account for this imbalance.

Let me give you a simple exercise to think about. Suppose the probability of transmitting a '1' is 0.9. This implies that if you were to blindly assume that a '1' is always sent, regardless of any information from the receiver, you'd still be correct 90% of the time! In this case, it becomes clear that performing maximum likelihood (ML) detection is not necessarily optimal. If you work through the math, you’ll find that the minimum probability of error (MPE) would actually result in a success rate slightly higher than 90%.

This is something worth contemplating. The proof of this was covered in our previous class, so we won't rehash that now. Instead, let's shift our focus to binary signaling. Binary signaling forms the foundation for more complex communication systems, essentially, it is the process of transmitting bits. Although it might seem simple, binary signaling serves as a critical building block for signaling schemes that involve a higher number of signals.

To better understand binary signaling and evaluate its performance, let's consider a straightforward case: on-off signaling. In this method, either you transmit a signal by switching on a voltage, or you transmit nothing at all. These two modes are analogous to transmitting a '0' or a '1'. The idea here is that, although we're focusing on a single symbol in this instance, the transmission pattern repeats itself at every symbol interval T, as we discussed in the context of pulse shaping.

(Refer Slide Time: 08:53)

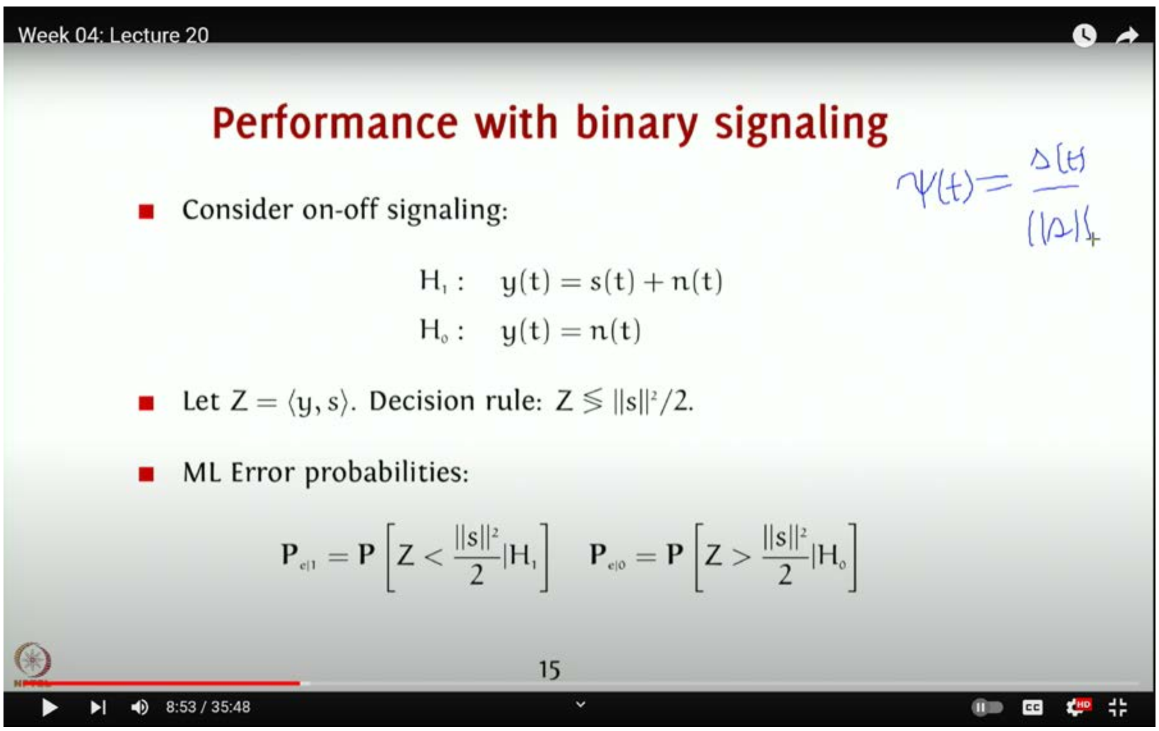

In this scenario, let’s define two hypotheses. Under hypothesis H_1, we have the relationship Y(T) = S(T) + N(T), where S(T) is the signal sent, and Y(T) is the received signal, which is the sum of the transmitted signal S(T) and the noise N(T). This corresponds to the case when a '1' is transmitted. On the other hand, under hypothesis H_0, which represents sending a '0', no signal is transmitted, and thus the received signal is purely noise: Y(T) = N(T).

The challenge now is to use the knowledge we've gained so far to determine the optimal detection strategy. We will assume that '0' and '1' are transmitted with equal probability, meaning that minimizing the probability of error is equivalent to performing maximum likelihood detection. In our last lecture, we made an educated guess that the decision rule could be based on the inner product $\langle Y, S \rangle$. The idea is that if this inner product, denoted as Z, exceeds a certain threshold, we decide that a '1' was transmitted; otherwise, we decide that a '0' was transmitted. The threshold in this case is $|S|^2 / 2$.

This guess is grounded in reason: if there were no noise, the inner product $\langle Y, S \rangle$ would simply be $|S|^2$ when a '1' is transmitted, and it would be zero when a '0' is transmitted. So, intuitively, the decision threshold lies at the midpoint between these two values. Once we establish this decision rule, we can then compute the associated ML error probabilities to assess its performance.

Before delving deeper into error probability calculations, it's always a good idea to conceptualize the problem in terms of vectors. Visualizing the signaling in terms of vector space helps solidify your understanding of how signals behave under different conditions. In this particular case, we only have one signal of interest, S(T), making it an ideal example to solidify these concepts further.

Let's dive into the discussion. Suppose we have a signal S(T), and let's say it has a value of 1 between 0 and 1 for simplicity. Now, you could choose a different signal, like a sinc function, but for now, we'll keep it simple and stick with this. The way we've chosen S(T) allows for an easy application of Gram-Schmidt orthonormalization because there's only one signal involved. In this case, S(T) is equivalent to our basis function ψ(T).

If you wanted to experiment a bit, you could scale S(T) by some factor, say 10. In that case, ψ(T), which is our normalized basis function, would be S(T)/10, and you can scale accordingly. Essentially, imagine that in the background, there is a basis function, ψ(T), which is equal to S(T) divided by its norm. To clarify, the norm of S(T) is the square root of the integral of $S^2(T)\,dt$. Thus, ψ(T) is simply S(T) scaled to have unit energy. This gives us a clear picture: we are dealing with one-dimensional signaling since we only have one signal S(T). If we had additional signals, say S_1(T) and S_2(T), spanning different dimensions, we would then be in a multi-dimensional signaling scenario, but here, it's just one-dimensional.

(Refer Slide Time: 12:57)

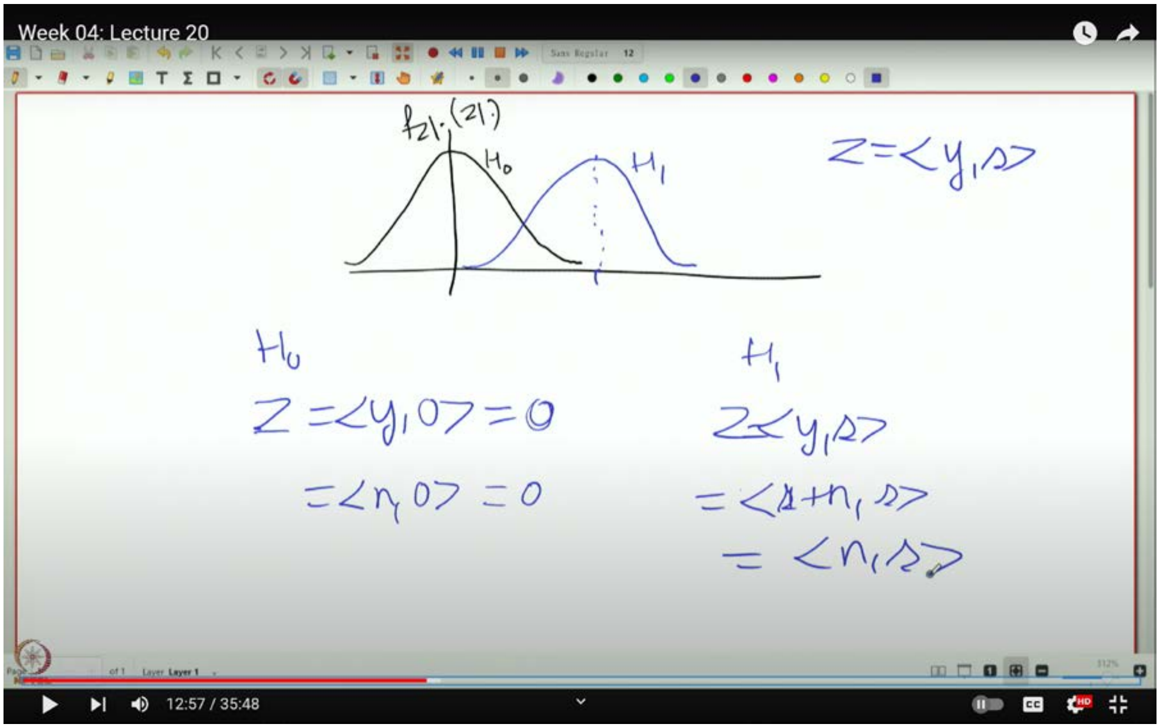

Now, let's work through this in a clear and systematic manner, starting with the decision rule and verifying whether it holds true. Let's visualize this by sketching the probability density functions (pdfs) of the received signal Y under two hypotheses: H_0 and H_1.

Hypothesis Testing

Under hypothesis H_0, nothing is sent, so the receiver only picks up noise. Under hypothesis H_1, S(T) is transmitted, and the received signal is the transmitted signal plus noise, S(T) + N(T). In both cases, we project the received signal onto the basis function ψ(T) by calculating the inner product < Y, S >, which gives us a scalar Z.

Let's now visualize the pdfs of Z for each hypothesis. Under H_0, where no signal is sent, Z corresponds purely to noise, so the pdf of Z, denoted $f_Z(Z \mid H_0)$, will be centered around zero with a certain variance. Under H_1, where the signal plus noise is received, the pdf of

Z, denoted $f_Z(Z \mid H_1)$, will be centered around the mean of the inner product $\langle S+N, S\rangle$, which corresponds to $\langle S, S\rangle = |S|^2$, again with some variance due to the noise.

Mean and Variance of the Gaussians

Now, let's explore the means and variances of these two Gaussian distributions. The key question is why these distributions are Gaussian in the first place. The reasoning is straightforward: S(T) is a known deterministic signal, and N(T), the noise, is Gaussian. When we calculate the inner product Z = < Y, S >, the result is a linear combination of the Gaussian noise, and any linear combination of Gaussian variables is also Gaussian.

Under H_0, the signal isn't transmitted, so Z reduces to the inner product of noise with the basis function, Z = < N, S >. Since the noise is zero-mean, the expected value of Z is zero. However, under H_1, the transmitted signal contributes to Z, so we have $Z = \langle S+N, S\rangle$, which simplifies to $|S|^2 + \langle N, S\rangle$. Thus, the expected value of Z is $|S|^2$ under H_1.

(Refer Slide Time: 17:45)

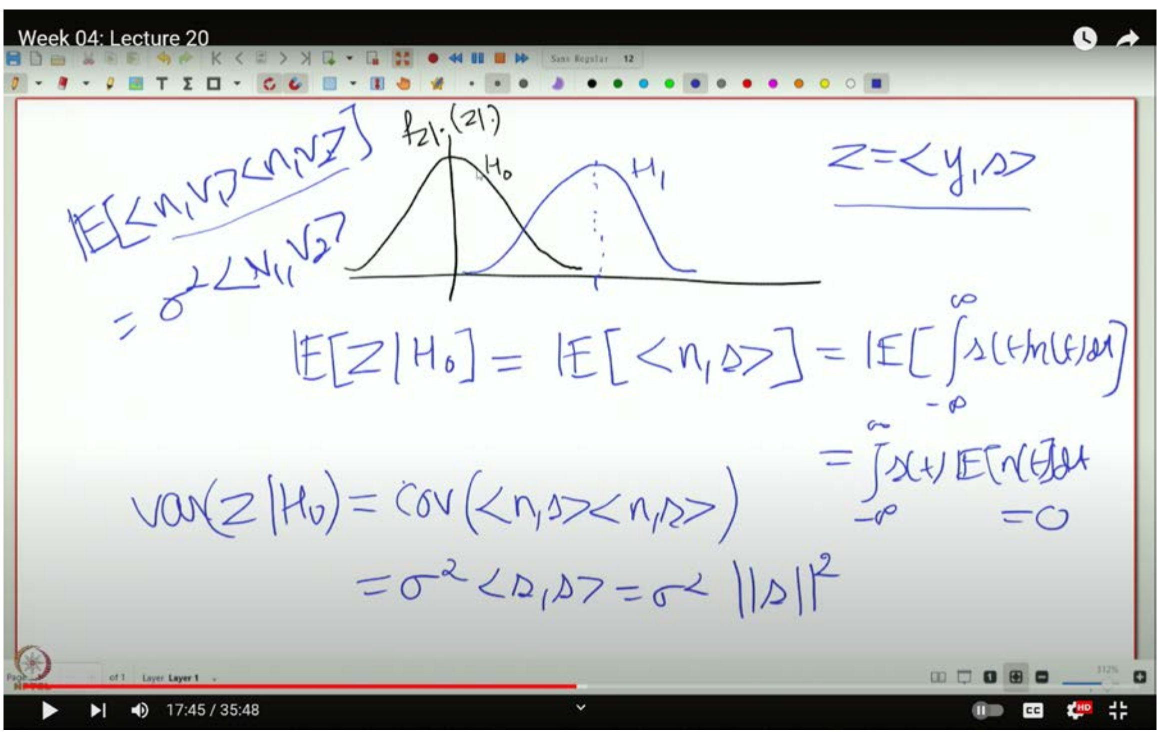

Thus, in both cases, Z is conditionally Gaussian: under H_0, it has a mean of zero, and under H_1, it has a mean of $|S|^2$. The variances in both cases depend on the power of the noise projected onto the signal space. With this understanding, we can move forward to calculate the ML error probabilities based on the value of Z.

Before we proceed with those calculations, it's essential to ensure that this setup is clear. We've established that Z follows a Gaussian distribution under both hypotheses, with different means and variances. From here, we can work through the decision rule, verifying that it minimizes the probability of error, and ultimately derive the corresponding error probabilities.

We are adding a fixed number, and when we take the inner product ⟨Y, S⟩, Y contains N(t), which is the Gaussian noise, while all other terms are fixed numbers. Hence, the expectation of Z given H_0 is the expectation of N over S, i.e., $E[\langle N, S\rangle]$, which is zero. The variance of Z given H_0 is the covariance of $\langle N, S\rangle$, which we've previously established as $\sigma^2|S|^2$. This implies that the expectation of Z given H_1 is the expectation of $< S + N, S >$, which is $|S|^2$. Similarly, the variance of Z given H_1 is the covariance of $< S + N, S >$, and we will prove this using known results. Let's first revisit the key concepts before moving forward.

Our definition states that Z is the inner product ⟨Y, S⟩. Since Z depends on N, it is a random variable. Now, what does the expectation of Z given H_0 represent? Under H_0, the value of Y is just N(t), so the expectation of Z given H_0 is simply $E[<N, S>]$. This is unsurprising because in this case, we are only measuring the noise, so Z contains only the noise component overlapping with S. When we express this as an integral, it becomes:

$$E\left[\int_{-\infty}^{\infty} S(t)N(t)\ dt\right] = \int_{-\infty}^{\infty} S(t)E[N(t)]\ dt$$

Since the expectation of N(t) is zero (because the noise has a mean of zero), the entire expression evaluates to zero.

I'll leave the integral form as is, but I won't repeat it for the subsequent cases. Now, because we need to find the error probability, we must carefully characterize the Gaussian distribution under H_0. We know the mean is centered around zero. So, that guess was correct. Now, let's determine the variance, which gives us the "spread" of the Gaussian. To do this, recall a previous result: the expectation of the inner product $< N, V_1 >< N, V_2 >$, where V_1 and V_2 are any two fixed signals, is given by $\sigma^2 < V_1, V_2 >$. We derived this result in a previous lecture, and it will be useful in the calculations ahead.

Thus, the variance of Z given H_0 can be determined. Under hypothesis H_0, Y is simply N. Therefore, the variance is the covariance of $< N, S >< N, S >$. Instead of writing this as an expectation, we can directly express it as a covariance, which simplifies things. We know that the covariance of X with itself, $Cov(X, X)$, is the variance. So, we write this as Cov(< N, S >, < N, S >), which simplifies to $\sigma^2|S|^2$. This shows that the variance of the Gaussian under hypothesis H_0 is characterized by $\sigma^2|S|^2$. In other words, when zero is transmitted, the distribution of Z is a Gaussian with mean 0 and variance $\sigma^2|S|^2$.

(Refer Slide Time: 20:02)

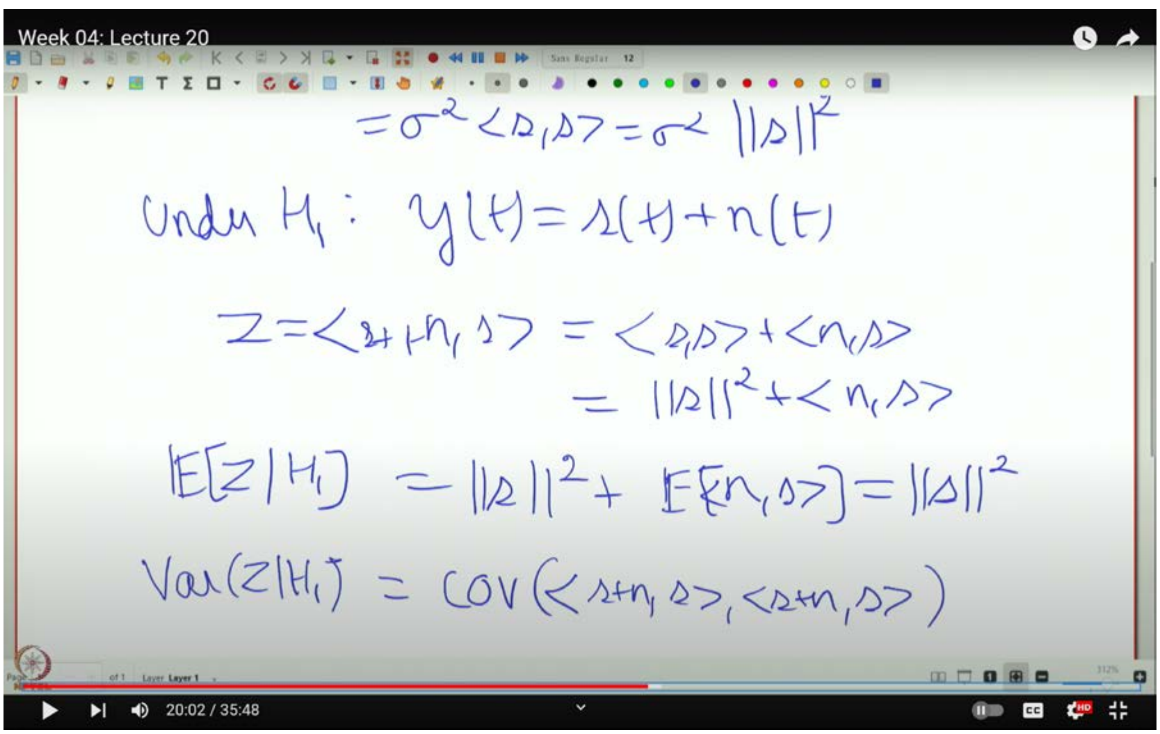

Our next task is to analyze the case under H_1. In this scenario, Y(t) = S(t) + N(t). Therefore, $Z = \langle Y, S\rangle = \langle S + N, S\rangle$, which expands to $\langle S, S\rangle + \langle N, S\rangle, or\ |S|^2 + \langle N, S\rangle$. What is the expectation of Z under H_1? Since $|S|^2$ is a fixed number (because S is a deterministic signal), this becomes $|S|^2 + E[\langle N, S\rangle]$. We've already shown that $E[\langle N, S\rangle] = 0$. Thus, the expectation of Z under hypothesis H_1 is $|S|^2$.

To compute the variance of Z given H_1, we again use the covariance approach. We need to compute the covariance of Z with itself, which is Cov(< S + N, S >, < S + N, S >). Using the same covariance formula we derived earlier, this can be simplified. However, since the noise is zero-mean, we can directly apply the covariance formula, leading to the result.

So, it doesn't really matter – this result is the same as the covariance of N < V_1, V_2 >, which is $\sigma^2\langle V_1, V_2\rangle$. Now, although the formula here isn't in the same form due to the extra S term, we can still utilize the linearity of the inner product. In fact, we can express this covariance as:

$$Cov(\langle S, S\rangle + \langle N, S\rangle, \langle S, S\rangle + \langle N, S\rangle)$$

Given that $|S|^2$ is fixed, we can simplify this by subtracting the fixed terms out when considering the covariance. So, what we're left with is:

$$Cov(\langle N, S\rangle, \langle N, S\rangle) = \sigma^2|S|^2$$

This is, unsurprisingly, the same as our previous result. Therefore, we conclude that the Gaussian distribution under hypothesis H_1 has a mean of $|S|^2$ and the same variance as the Gaussian distribution under H_0.

By symmetry, we can safely deduce that the midpoint between 0 and $|S|^2$ is $|S|^2/2$. Now, when it comes to making a decision based on the observation y, how do you determine which hypothesis is more likely? It turns out the decision process is quite simple: You just need to check which side of $|S|^2/2$ you are on. If you are on the right side of this midpoint, it is more likely that H_1 occurred. Conversely, if you are on the left, H_0 is more probable. In other words, if Z is closer to $|S|^2$, it is more likely that hypothesis H_1 is true, meaning '1' was sent. On the other hand, if Z is closer to 0 – that is, to the left of $|S|^2/2$, you should conclude that '0' was sent. This is the optimal decision rule for binary signaling in this scenario.

(Refer Slide Time: 26:47)

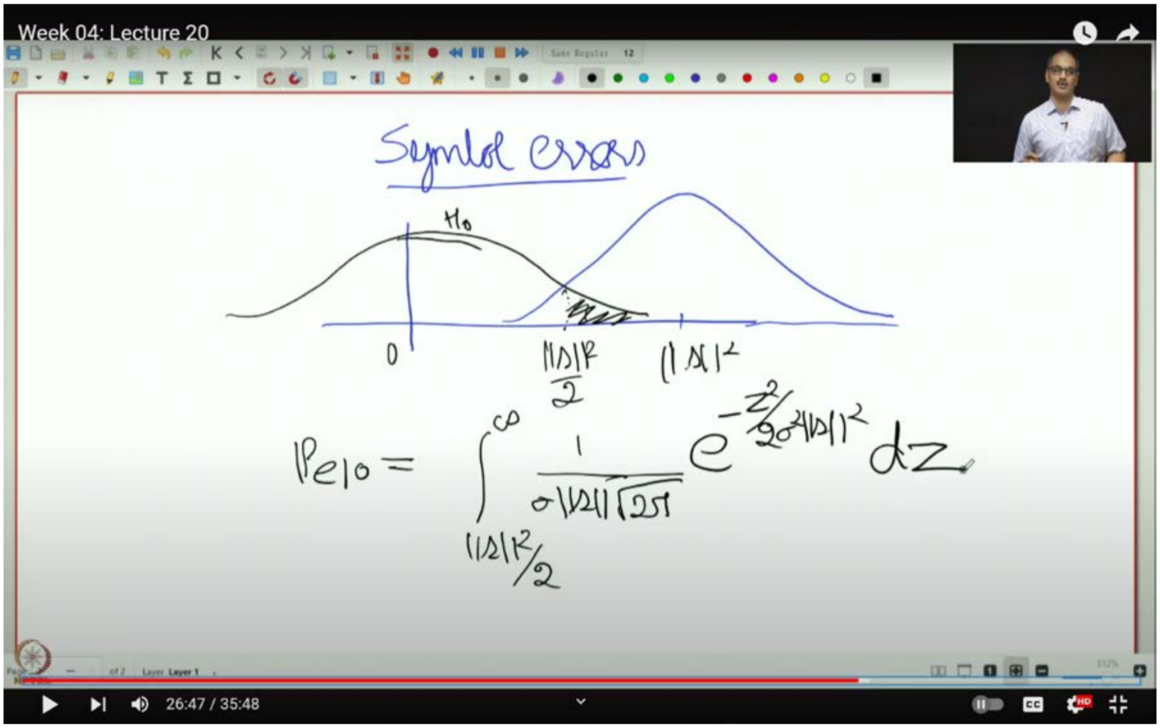

However, even with the optimal decision rule, errors can still occur. Why is that? The noise might cause errors in the decision-making process. For example, under hypothesis H_0, when '0' is sent, the noise could push the decision to the right side of $|S|^2/2$, leading you to incorrectly decide that '1' was sent. Similarly, under hypothesis H_1, if '1' is sent, a large negative realization of the noise could cause the decision to fall to the left of $|S|^2/2$, making you incorrectly decide that '0' was sent. Thus, in both cases, there is a possibility of making an incorrect decision due to noise interference.

These incorrect decisions are referred to as symbol errors. So, what we want to calculate now is the probability of making symbol errors under these conditions. Let's go ahead and work this out.

We'll start with defining the symbol error probability. To make this clearer, let me sketch it out. Here, I'll draw the decision regions. First, under hypothesis H_0, and then in another color, I'll depict the region for hypothesis H_1.

Now, let's label these key points: This point here represents $|S|^2$, which is the same as $\int|S(t)|^2dt$. The midpoint is $|S|^2/2$, and here is the origin at 0.

What is the probability of an error under hypothesis H_0? An error occurs when '0' is sent, but the observation lands in the region beyond $|S|^2/2$. To determine this symbol error probability, we need to calculate the probability that, under hypothesis H_0, the observation crosses $|S|^2/2$.

Now, let's calculate the probability of error. To do that, recall the characteristics of Z under hypothesis H_0: it has a mean of 0 and a variance of $\sigma^2|S|^2$. Given this, we aim to compute the probability of error under hypothesis H_0, denoted as P_{E0}. This probability corresponds to the area under the tail of the Gaussian distribution from $|S|^2/2$ to infinity. The formula for the error probability is:

$$P_{E0} = \int_{|S|^2/2}^{\infty} \frac{1}{\sigma|S|\sqrt{2\pi}} e^{-\frac{z^2}{2\sigma^2|S|^2}} dz$$

Here, we're applying the Gaussian distribution formula $\frac{1}{\sigma\sqrt{2\pi}} e^{-x^2/2}$, adapted for our situation where the variance is $\sigma^2|S|^2$. The goal is to determine the probability that the random variable Z exceeds $|S|^2/2$.

Now, to simplify the evaluation of this integral, we transform it into the standard normal form. Let's make a substitution:

$$u = \frac{z}{\sigma|S|}$$

This substitution standardizes the variable Z by dividing it by its standard deviation. When $z = |S|^2/2$, the corresponding value of u becomes:

$$u = \frac{|S|}{2\sigma}$$

Thus, the error probability P_{E0} becomes:

$$P_{E0} = \int_{|S|/2\sigma}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-u^2/2} du$$

(Refer Slide Time: 29:00)

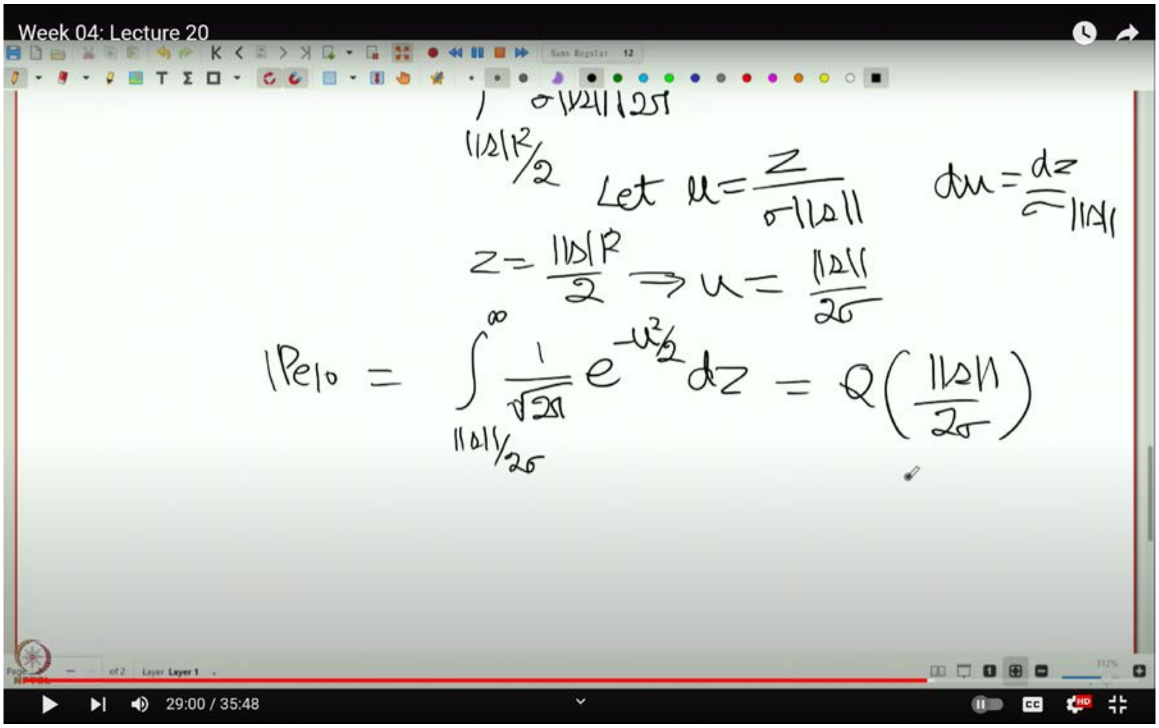

At this point, we recognize this as the Q-function, which represents the tail probability of the standard normal distribution. Therefore, we can express the error probability as:

$$P_{E0} = Q\left(\frac{|S|}{2\sigma}\right)$$

The Q-function gives us the probability of a symbol error occurring when '0' is sent. You might also refer to this as the bit error probability, since we're dealing with binary signals ('0' and '1').

Now, let's shift focus to the probability of error under hypothesis H_1. Recall that, under H_1, the distribution of Z has a mean of $|S|^2$ and the same variance, $\sigma^2|S|^2$. In this case, an error occurs if the observation falls to the left of $|S|^2/2$. Due to the symmetry of the Gaussian distributions, the area representing the error under H_1 is the same as that under H_0.

Nevertheless, it is instructive to explicitly calculate this probability to verify the symmetry and ensure accuracy.

Thus, by symmetry, the probability of error under hypothesis H_1 is:

$$P_{E1} = Q\left(\frac{|S|}{2\sigma}\right)$$

In both cases, whether hypothesis H_0 or H_1 holds, the probability of making a symbol error is captured by the same Q-function. This confirms that our intuition about the symmetry of the decision regions is correct.

Let me write this down quickly so we can confirm that the error probability remains consistent. Under hypothesis H_1, we know that the mean is $|S|^2$ and the variance is $\sigma^2|S|^2$. Therefore, the probability of error given hypothesis 1, denoted as P_{E1}, is found by evaluating the probability that the observed value Z falls to the left of $|S|^2/2$.

(Refer Slide Time: 33:16)

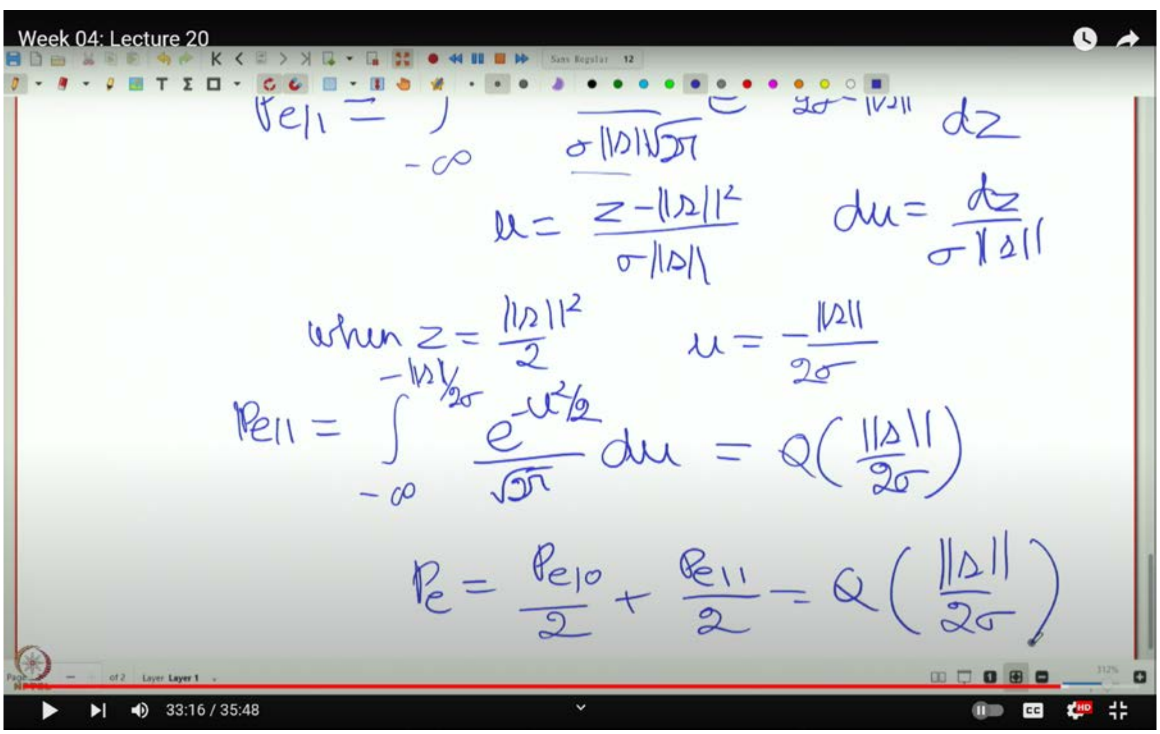

Mathematically, this can be expressed as the integral:

$$P_{E1} = \int_{-\infty}^{|S|^2/2} \frac{1}{\sigma|S|\sqrt{2\pi}} e^{-\frac{(z-|S|^2)^2}{2\sigma^2|S|^2}} dz$$

Again, we apply a substitution to simplify this expression. Let:

$$u = \frac{z - |S|^2}{\sigma|S|}$$

Then, $du = \frac{dz}{\sigma|S|}$, which takes care of the differential part of the integral. Now, we adjust the limits of integration. When $z = |S|^2/2$, the corresponding value of u becomes:

$$u = \frac{|S|^2/2 - |S|^2}{\sigma|S|} = -\frac{|S|}{2\sigma}$$

Thus, the probability of error P_{E1} becomes:

$$P_{E1} = \int_{-\infty}^{-|S|/2\sigma} \frac{1}{\sqrt{2\pi}} e^{-u^2/2} du$$

However, due to the symmetry of the Gaussian distribution, integrating from -∞ to -$|S|/2\sigma$ is equivalent to integrating from $|S|/2\sigma$ to ∞. This results in:

$$P_{E1} = Q\left(\frac{|S|}{2\sigma}\right)$$

So, the probability of error under hypothesis H_1 is also $Q\left(\frac{|S|}{2\sigma}\right)$.

Given that both hypotheses are equally probable, the overall probability of a symbol error (or bit error) is the average of P_{E0} and P_{E1}, which are identical. Thus, the total probability of error becomes:

$$P_E = \frac{1}{2}(P_{E0} + P_{E1}) = Q\left(\frac{|S|}{2\sigma}\right)$$

In conclusion, this lecture highlights a critical point: whenever you are dealing with an additive white Gaussian noise (AWGN) channel and binary signaling, the relationship between the signal and noise allows you to compute error probabilities quite easily. Symmetry is a powerful tool, wherever possible, leverage it to simplify your calculations. You won’t always need to work out the full integrals; sometimes, by inspection, you can directly identify the error probability. However, it's essential to proceed carefully.

(Refer Slide Time: 34:21)

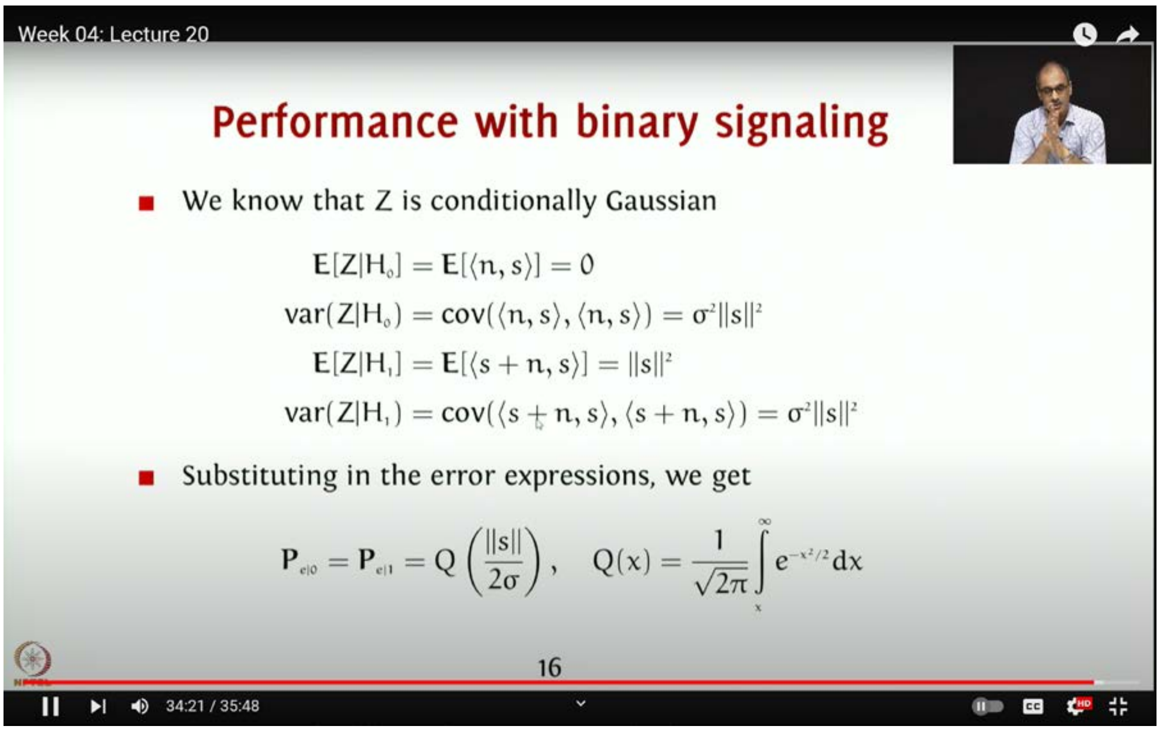

Moreover, when evaluating metrics like Z (the inner product $<y, S>$), finding the decision threshold is key to determining the error probability. In our case, the decision point is at $|S|^2/2$. If the noise pushes the observation Z exactly to this point, the decision is ambiguous, although the probability of the noise doing so is effectively zero. If the observation falls to the left, you conclude that '0' was sent; if it falls to the right, you conclude that '1' was sent.

When the noise is large enough, symbol errors become inevitable, and the probability of making such errors is given by $Q\left(\frac{|S|}{2\sigma}\right)$, as we've calculated under the binary signaling framework, where '0' corresponds to zero being sent and '1' corresponds to $|S|$ being sent.

In the next lecture, we'll extend this concept to binary signaling with two distinct signals, allowing us to compare and contrast the error probabilities. We'll also introduce discussions about energy per symbol, energy per bit, and how these factors influence the error rates in signaling systems. Thank you.