NARRATOR: The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: Today, I want to spend a fair amount of time, to start out with, talking about this business of, how do you go from pass band to a baseband in a wireless system and some of these questions about time reference and all of these things. When we were dealing with ordinary channels that didn't have fading in them and didn't have motion or anything like that, it was a fairly straightforward thing to think of a transmitter having its time reference and the receiver locking in on its time reference, which had a certain amount of delay from what the transmitter was doing. As soon as you start having multiple paths, each with different delays in them, this starts to get confusing. If you try to write it down carefully, it gets doubly confusing. If these path lanes are changing dynamically, it gets even more confusing.

So I wanted to spend a little bit of time at the beginning of the hour trying to sort that out. It's not sorted out that well in the notes. As you'll see, it doesn't get sorted out too well here either because the problem is, whenever you write everything down, it becomes a nightmare. So this is just another effort at trying to do this.

Let's start out by just looking at a single path. This single path, we're going to assume, has a time varying attenuation, which we'll call beta of t and a time varying delay, which we'll call tau of t. Namely, this is the same thing that we did before. We started out with a fixed transmitter and a fixed receiver. Now, since we're allowing both the attenuation and the delay to vary, but only having one path. What we're essentially thinking is, that is the receiver, which is in a vehicle for example, which is moving away from or towards the transmitter.

So these are the, or in a sense an easy way of keeping track of that physical situation. So the response at passband to a real waveform, x of t, it's going to be

just beta sub t times x of t minus tau of t. In other words, here's the delay here. Here's the attenuation here and we just have a single path. Now if you start out with a baseband waveform, if you start out with a complex baseband waveform, say u of t, you're going to shift this up first to a positive frequency band, which we'll call u sub p of t, which is u of t e to the 2 pi ifct and you're going to transmit it as x of t, which is equal to two times a real part of that.

Half the time we've been dealing with going through this modulation from baseband to passband in terms of cosines and sines. Half the time we've been doing it in terms of complex exponentials. One of the problems that appears as soon as you get into wireless is that if you really like to do things in terms of cosines and sines, you really run into a lot of trouble here because with phases varying all over the place, it just becomes very, very difficult to track what's going on with all of the cosine terms and all of the sine terms.

So what we sort of have to wind up with here is this, in a sense, simpler but more abstract viewpoint where we think of modulation as being a two step process, where we first take the take the low pass waveform, move it up in frequency and then we add on the negative frequency part as an afterthought and then at the receiver, what we're going to do is have a Hilbert filter, which you would never build, of course. You would always build it in terms of cosines and sines, but conceptually have a Hilbert filter which blocks off that negative frequency part and then you just shift things down in frequency again.

So that's the way we're thinking about this here, so we're going to have u of t shifted up and then transmitted as 2 times the real part of up of t and therefore, it's going to get received -- the received waveform at the passband. We're still using the timing at the transmitter. That's the purpose of going through all this analysis, to look at what's going on at transmit time and what's going on at receive time.

So in terms of transmit time, this received waveform at passband is just the attenuation term times x of t minus tau of t, which is two times the real part of theta of t times this positive frequency part. Now when you write it out in all of its glory, it's

2 times the real part of the attenuation times the input you started with, but now this has been delayed by tau of t and e to the 2 pi ifct, but this was the carrier that was put on at the transmitter, physically at the transmitter and now we're at the receiver, so that carrier has gotten delayed by a factor of tau sub t.

The question is, how do you demodulate this? How do you come down again from transmitter to receiver? What we've thought all along and what makes a lot of sense until you start putting these delays in, which are really a pain in the neck, is that what we're going to do is just multiply this positive frequency waveform by e to the minus 2 pi i f sub ct, but that's not what we're going to do because the receiver sitting there and the receiver has to recover timing and it has to recover carrier frequency. So what's it going to do? It's going to figure out what the timing is and the timing that it wants is at time 0, it wants to be seeing what the transmitter transmitted at transmit time 0. In other words, if you look at what the transmitter was doing at time 0, it was sending, say, the first bit that was going to be sent. In receiving this, we want to have our timing so we're looking at that first bit sent. So our timing is going to be shifted from what it was then.

So we're going to take this new timing. The receiver clock time is now going to be t prime, which is t minus this delay term. If you get confused between the minuses and the pluses here, don't worry about it. I get confused about them too and the only way I can straighten them out as to put down one or the other and then spend ten minutes looking at it and try to figure out -- after I write it down -- whether it should really be a plus or a minus and I think these sines are right, but I'm not going to try to argue why in realtime. But anyway, this received waveform now, as a function of received time, is going to be the received waveform and in terms of transmit time -- and in place of t now, I'm going to have t prime so that since you t prime is equal to t minus tau of t, this is why it's t prime plus tau of t. In other words, what I'm looking at now, at time 0 at the receiver is what was being transmitted a little while ago. So that's this term and that's going to be the real part of this, where I've taken account of all of these shift terms.

Now if you look at this carefully, you see that u of t minus t of t is turned into tau of t

3

prime. This term has turned into 2 pi ifct prime, so everything is fine there. What I should have done with this term is to compensate for the fact that I'm now looking at it in a different time and at this point, what I'm going to do is to say, this quantity is changing so slowly with time that I don't care about that. If you really try to adjust this to make it right, it's a terrible mess. So we're just not going to worry about it. So what we receive then is approximately equal to this where the approximation is due to the fact that I'm not evaluating this attentuation term at exactly the right time, because it's a pain in the neck to do so.

All of this stuff with wireless, you simply have to make approximations all over the place. You have to make crazy modeling assumptions all over the place, because the physical medium is so complicated that you can't do much else and the whole question is trying to make the right approximations and get some sense of which thing's very fast and which thing's very slowly. So that's the equation that we had. What we receive now at passband, but at the received time scale is this quantity here. So what's the receiver going to do? The receiver, at its time is going to demodulate by multiplying by e to the minus 2 pi ifct prime.

First, we're going to take away this 2 times the real sine by going through the Hilbert filter, which just removes this. At that point, we have a complex positive frequency waveform and then we're going to multiply that positive frequency waveform by e to the minus 2 pi ifct prime, where t prime is the only thing the receiver knows anything about. So after you shift the baseband with a recovered carrier in receiver time, what you get in receiver time is v of t prime as equal to beta of t prime; namely, the attenuation at t prime -- or this is the approximate part of it -- times what was actually set. So now, think of this as saying, suppose this delay term is actually changing with time. Namely, it's some tau 0 minus a velocity term, vt divided by the velocity of light. In other words, the time delay that you incur, where the change in the time delay that you incur, is due to the time that it takes light to travel from where the receiver was to where the receiver is now. We can also write that as tau 0 minus the Doppler shift times the time divided by the carrier frequency.

Now here's another approximation, because what we're sending -- I mean, this

quantity here is exact. When you try to convert from the velocity, which you're going to, for what it does in frequency, which is the Doppler shift, it's a function of the frequency. Here what we're assuming is that all the frequencies we're dealing with are so close to the carrier frequency that we can just ignore everything else. So we're just writing this as the Doppler shift divided by this carrier frequency.

Now if you view the passband waveform, going into a baseband of modulation and you think of doing this in transmit time, what are you going to do? In terms of transmit time, you've got to get rid of not this term, but this thing here. So in terms of transmit time, you're going to be multiplying by e to the minus 2 pi i times carrier frequency minus the Doppler shift. So the thing that's happening is that time at the transmitter gets spread out a little bit at the receiver. In other words, as this receiving antenna is moving away from the transmitter, a period of time like this at the transmitter is a period of time like this at the receiver and therefore, a certain number of cycles of carrier at the transmitter looks like a different number of cycles of carrier at the receiver. Therefore, if you're dealing with transmit time, you're taking account of this Doppler shift. You're multiplying not by the carrier frequency, but by the carrier frequency minus the Doppler shift. If you do it in receiver time, you're just multiplying by the by this carrier frequency, because the trouble is, your clock is running a little slow. Because your clock is running a little slow, you think you're demodulating at the actual carrier frequency, whereas in terms of what the transmitter thinks, you're doing something else. If you all got confused when you studied relativity, this is exactly the same problem you got confused about there. There's nothing very mysterious about it or hard about it, except that to most of us, time is a very fundamental quantity and you don't like to monkey with that. If you can't count on time being what it's supposed to be, you're really in deep trouble. When you try to write equations that do that, it makes things very, very tough also.

So anyway, the result after we look at this in both ways, is that you get the same answer each way, but the receiver sees the carrier frequency and sees this carrier frequency in transmitter time and it sees this frequency in received time. Of course, the receiver only sees received time because the receiver sitting there trying to measure from the waveform coming in what that carrier frequency is and what that

time base is.

So now it gets to be fun. We look at multiple paths. When you look at it in transmit time, what gets received at the receiver, according to the transmit clock -- so here we are at the transmitter and we're peering at what's going on at this distant receiver and you now have the thing that we've talked about in the notes where you have all these different paths. Each path is associated with a certain attenuation and each path is associated with a time delay. This can be written in terms of what has gotten modulated up from baseband as 2 times the real part of this sum -- of beta j of t times the positive frequency part of what got transmitted and then it now has this delay in it. This can be rewritten in the same way that we did before, but now we just have all of these paths in there instead of one path.

If you look at this expression then, we have all of these propogation delay terms. We have this baseband input, which has been delayed in different ways for each path and you have have this exponential, which has been delayed also. So at this point, the receiver is in a real pickle because if the receiver is smart enough to see -- yes?

AUDIENCE: [UNINTELLIGIBLE PHRASE]

PROFESSOR: The receiver doesn't know anything. It's like you're driving in your car, you're talking on your cell phone and your cell phone -- I mean, the only way your cell phone knows that you're moving is that the cell phone is getting some waveform coming in and it has to tell from that waveform what's going on. So it's going to do things like the carrier tracking that we were talking about before and it's not going to have too much trouble because all of these changes are taking place relatively slowly, relative to this very high -- either one gigabits or two gigabits or four gigabits, or what have you, carrier frequency. So all these changes look like almost nothing, but at the same time, they're fairly important because they keep changing faces.

But anyway, in terms of the transmitter clock, this is what you actually receive. So at this point, what the receiver has to do is it has to retrieve clock time and the clock time that it retrieves is going to be -- receiver clock time is -- is t minus tau 0 sub t --

6

so it's t prime equals t minus tau 0 sub t. That's the same thing it was before, except that this tau zero of t here doesn't really amount to anything physical anymore. It's just whatever the circuitry that we have trying to recover carrier and trying to recover timing -- it's just whatever that happens to come up with, so with this best sense of the timing it should use if it's going to try to detect what the signal is. So it's some arbitrary value and again, this is varying in time. So again what we have is now this expression becomes really messy, because instead of having a receiver time variation, which cancels out what the actual variation in path length is, you really have two terms that are different and you have the same thing in terms of this phase, which is changing and this is the more important term.

But anyway, you have that. You can then demodulate. What do you do when you demodulate? You take this thing in receiver time and you multiply by e to the 2 pi i, e to the minus 2 pi i -- this carrier frequency, what you think it is, times this time reference, what you think that is -- So that term has disappeared and what we then wind up with is the carrier frequency times these two terms here. That's to the left there. These two terms turn out to have both Doppler shift terms in them and also time spread terms. So when we write that out in terms of Doppler shift, what we wind up with is the input, which is then delayed by some arbitrary amount of time, where this is our best estimate of the right time to use and this is the actual time on that path.

The thing that's happening here, which you don't see in the notes unfortunately, in the notes, it pretty much looks at things as far as transmit time is concerned. Therefore, what you see is expressions for impulse response where the impulse response is at some time much later than time 0. You have a -- I mean, you put in an impulse of time 0, what you see is stuff dribbling out over some very tiny interval, but a good deal later than what got put in. What we want to do now is to shift that so that in fact, when we look at filtering and things like that, we have a filter which starts a little bit before 0 and ends a little bit after 0 and that's the whole purpose of this. If you don't think my purposes just to confuse you, it really isn't, although I realize this is very confusing, but you sort of have to go through with it to see why these time

shifts appear in the places where they do.

But anyway, up here in the phase, you wind up with the Doppler shift terms and down here, when you're worrying about the input waveform, you're worrying about these time shift terms. We'll see we'll see how they come in a little later.

To make the expression a little bit easier, let me look at tau j prime of t as tau j minus tau 0 and let me look at the Doppler differential as the Doppler that actually occurs on path j minus what the receiver thinks the Doppler is. This is this receiver term that we've demodulated by and therefore, we've really gotten rid of this term. Then when we rewrite this received baseband waveform -- and I've now gotten rid of the t primes because everything we've done up until now, we've just automatically assumed that the receiver timing and the transmit timing, we didn't have to be careful about it and it wasn't until we got to wireless that we do have to be careful about it.

So now we're throwing it out again and what we have is these attenuations, which are very slowly varied. These terms, which are now delayed by this difference between the actual delay and what the receiver is assuming that the light ought to be. So these are really just differential delays relative to each other. These are Doppler shifts relative to each other. So that's the whole receive baseband waveform then. If you look at this, it's not apparent where the carrier frequency is coming in and the carrier frequency is coming in because these Doppler shifts have the carrier frequency in them. Namely, the Doppler shift is the velocity times the carrier frequency divided by the speed of light. One of the things you have to be very careful with in the wireless business, is that everybody thinks it's a neat idea to be able to be able to go up to higher and higher frequencies. When you go from one gigahertz up to four gigahertz, one of the things that you have to deal with is that every Doppler shift you're dealing with is now four times higher than it was before. So there's one big disadvantage in operating at extremely high frequencies. It's not necessarily a bad thing, but it's there and if you have a bunch of different paths which are all acting in different ways, you suddenly wind up with some real problems.

The received baseband waveform on path j is going to change completely over an interval -- 1 over 4 times the magnitude of this differential Doppler shift. The Doppler shift might be positive or negative, but if you look at this expression up here, that's the amount of time that it takes for this quantity to change by pi over 2. When this changes by pi over 2, this exponential goes from its maximum one down to zero or it goes from zero to minus one or it goes from minus one to zero or from zero up to plus one, and that's a pretty major change. So the interval which is required to do that is 1 over 4 times that Doppler shift.

The Doppler spread now is defined as the difference between the maximum Doppler shift and the minimum Doppler shift. If we choose this Doppler shift D sub 0 that we had, because of what the receiver was doing, seeing all these different Doppler shifts and trying to come up with something in the middle, that's pretty close to what the receiver thinks it should be using as far as receive time is concerned. So it's the maximum plus the minimum divided by 2 and then each of these differential Doppler shifts is going to be somewhere between that overall Doppler spread minus Doppler spread over 2 and plus Doppler spread over 2. It goes from -- minus D over 2 up to plus D over 2 and this difference in here, is strangely enough, D. So what we've done is by adjusting our received timing -- and slowing it down or speeding it up if necessary, is we wind up with Doppler shifts that are sort of spread around from negative Doppler shifts to plus Doppler shifts. What that means is, if you look at this expression, it takes a period of time, about -- here we had 1 over 4 times the differential Doppler, which is now 1 over 2 times this overall Doppler spread.

So what that says is, there's a coherence time on the channel of about 1 over 2 times the Doppler spread. That's important because that says the channel is going to look like the same thing for a period of time, which is about 1 over 2 times this Doppler spread term. I think this will be D instead of D 0. So you find out what the Doppler spread is and the reason you want to find out what the Doppler spread is is that that tells you how long the channel remains stable. Why do you want to know how long the channel remains stable? If you're going to do detection or anything like that, if you're going to filter the received waveform to try to find out what the input

waveform is, you would like to know what the channel is doing. You'd like to be able to measure the channel. If you're going to measure the channel, you really want to know how long it stays the same before you have to measure it again.

If you look at all these cellular systems and you look at every other wireless system in the world, all of them do one of one of two things: either periodically they measure what the channel is so they can use it and the other thing they do is every once in awhile, they send a pilot tone and they use the pilot tone to try to figure out what kind of channel they have. This is the thing which tells you how often you have to send that pilots tone. If you know about estimation theory, it's also the thing that tells you whether you have any chance or not of estimating what the channel is, because it takes you a certain amount of time in the presence of noise to estimate something and if that something you're trying to estimate it's changing faster than you can estimate it, then you don't have a prayer of a chance.

So yes, this is important. This is the prime coherence. It's one of the main parameters of a wireless channel you want to understand. It tells you how long the channel will remain essentially the same before it changes and becomes something different. Both of these quantities are very approximate. If anyone wants to tell you that it's twice what I say it is here, I certainly wouldn't quarrel with them. I mean, it's like talking about the bandwidth of a waveform them that you send. I mean, you look at the spectrum of what you send and it's going to be something which sort of tails off. If you want to call it the maximum out there where it's going to zero or the minimum, where it starts to go to zero or what have you, it doesn't make much difference. Here, it's a worse situation because you don't shape this waveform. It's just given to you and you find something which is going to be tailing off slowly. It's going to have little bits of nothingness way out there. All this is is just one very approximate number, which tells you what that Doppler spread is and therefore, which tells you how long the channel is going to remain stable and yes, in half of that time the channel starts to change and in twice that time, the channel might be still sort of like it was before. So this is only a very approximate way of looking at these things, but it gives you an idea about what you can do and what you can't do. It gives you an idea, for example, if you say, can I take this wireless system I've

been using and use it at a frequency that's 10 times as high -- and you know when you use it at a frequency 10 times as high that the stability of the channel is going to be 10 times as small. You know if you were having trouble measuring the channel before, you know you don't have a prayer of a chance when you go up to that higher frequency, so it so it tells you things like that.

There are two other sort of related quantities, which are called multipath spread and coherence frequency. So we have Doppler spread on the one hand, which is a physical phenomena and that gives rise to coherence time, which says how long the channel will stay the same and in the opposite domain -- and these are almost dual quantities of each other, you have something called multipath spread, which -- and multipath spread is pretty much what it sounds like it is. You have a bunch of different paths. You send a very short duration waveform and things start trickling in on the shortest path and then you got a big bunch of stuff in all the paths which are sort of the same lanegth and then they sort of dribble away again.

So the waveform that you see coming in is spread out from what you send because of these multiple paths and the multiple paths having different delays on them. Because you have all of these different delay paths, you're going to have an impulse response for this channel, which has a certain duration to it. When you take the Fourier transform of that, you find a certain duration for Fourier transform. As we're going to see in just a minute, that's the thing that tells you how stable this channel is in frequency.

If you measure it at one frequency, how far do you have to go in frequency before you find something which is totally separate from what you measure down here? This is something the cuts both ways. Having a large frequency coherence and having a large time coherence can be good or it can be bad. If you have a very large time coherence and a very large frequency coherence, you can measure the channel once and you can use it for a long time and you can use it over a broadband width and everything looks very nice. If the channel happens to be faded over that long period of time and over that long interval of frequency, you're really in the soup. On the other hand, if you have a very small time coherence and a very

small frequency coherence, if the channel doesn't look good there, you move to a different time interval or you moved to a different frequency interval and you can transmit again.

So both of these quantities have good aspects to them and bad aspects to them, but if you want to evaluate what a physical wireless channel is and somebody lets you ask two questions, the two questions you ought to ask are, what is the time coherence and what is the frequency coherence? Those are the first things you want to find out. Everything else is sort of secondary.

So trying to claim that the frequency coherence here almost has a matter duality, as 1 over twice the time spread. Let's take a slide to see why that is. If I take the impulse response at baseband of this channel, which I wrote down here -- here's the impulse response and you get that directly from the baseband response for giving input waveform. With all these different paths, what this impulse response is at a particular time t -- namely, this is the response at time t to an impulse tau seconds earlier. In other words, when I'm using this notation here, I'm really thinking of t as a receiver time and t minus tau as transmit time.

But anyway, this is what it is. You take the Fourier transform of this for a particular t -- in other words, we talked about this a little bit last time. You have a function of two parameters here and at any given t at the receiver, what you see is a spread of different inputs coming in from the transmitter. What we're interested in now is at that particular time t at the receiver, what's the Fourier transform, when you take the Fourier transform on tau to see what kinds of frequencies are coming in? We take the Fourier transform of a unit delta function and you just get an exponential function. So the transform of that delta is just this e to the minus 2 pi if times tau sub j prime of t. This is this differential delay that we're faced with and we still have the Doppler shift over here.

So when we look at it this way, we have both the Doppler shift sitting there and we have this delay shift sitting up here. Now the question we ask is, how much does the frequency have to change before this quantity is going to change materially? The

answer is the same as it was before. When this quantity here changes by pi over 2, you have a totally different result than you had before. So the question is, how much does a frequency have to change in order for f times tau j prime of t to the change by a factor of one-fourth. So the answer that we come up with is, this has to go from minus L over 2 to plus L over 2 and this is where we get the result that we were just talking about, that the frequency coherence is 1 over 2 times the multipath spread of the channel.

The next thing we want to do in this long process -- it's a process that we've sort of been doing all along. The whole purpose of looking at waveforms and looking at digital communication is to be able to go from sequences to waveforms -- and what do you think the next thing we're going to do is? We're going to go from these waveforms -- we've now got them down to baseband waveforms and we want to go back to looking at sequences. So what we would like to do is to take this input waveform, view it in terms of the sampling theorem -- better thing to do would be to view it in terms of input pulses, but that's too complicated for now, so we'll just take the input waveform and view it as a sequence of data that we're sending -- the u sub k's times these sinc waveforms. Same thing we've been doing all along. We take a sequence of numbers, we put little sine x over x-hats around them and that gives us a band limited waveform.

The baseband output is then -- you can look at the same way. The baseband output is going to be limited to the same frequency. Question here: You have these Doppler shifts. So when I send a frequency which is right up at the limit of w over 2, what I get back from that is going to be spread out a little bit from that. So if I send a waveform that's limited to w over 2, I'm going to get back a waveform which is a little bit bigger than that. What do I do about that? I worried about it in the notes a little bit and I'm probably going to take it out from the notes, because as a practical matter, the amount of Doppler shift that you get is such a negligible fraction of the kind of bandwidth you would be using for transmission that you just want to forget about that.

So we are going to forget about it and what's coming in is again a band limited

function with maybe a little bit of squishiness on it. We get this baseband output, which we can then view in terms of the sampling theorem. We can sample the output and what do we get when we're all done? Here's the nice thing that we can jump to: The output at this free time m is then the sum -- namely, it's the convolution, it's the discrete convolution of what went in at time m, at time m minus 1, m minus 2, m plus 1, m plus 2 -- and what we're doing is, we're designing this recovery time so the main peak is about a time 0, so k goes from minus something to plus something here and then we wind up with these channel terms, which are again, sampled at this point because we can sample them, because it's only these low frequency components of the channel response that make any difference. They're the only things that give us this output here and this output has already been filtered down to this low frequency.

When I look what these components are, these are big messes again. They're sums over all of these different paths. So what we're winding up with is a function -- if I sketch g of tau and t as a function of tau, what I'm going to find is a bunch of different terms. Now I'm going to filter this so every one of these turns into a little sinc x over x and then I'm going to sample it. So what I wind up with then is this goes into something which looks like this. It's going to look like something which goes up and comes down again, and I'm just interested in what it's sample values are and I might wind up with five or six sample values and that expresses the whole thing. I mean, this is something you have to think about a little bit because it's something that we've done about 10 times already, but now when we do it here it looks very different from what we've done before -- except this is what it is analytically. This is what it is in pictures and that's where it comes out.

The thing that's going on here then is we started to look at this physical modeling in terms of ray tracing and things like that and what we're doing at this point is we're modeling what's going on at the frequency bandwidth that we're using -- namely, not the carrier frequency anymore, but the baseband bandwidth. We're viewing this filter as having filtered these things out for us so we wind up with a baseband discrete filter that filters these impulses and aggregates them under discrete caps. So at this point, when we look at this filter for what the channel is doing, instead of

seeing a bunch of impulses, we see a rather smooth waveform. If we look at what contributes to each tap, we see a large number of different paths, which all contribute to the same tap.

If the multipath spread is not too large relative to 1 over w -- if if the multipath spread is small relative to 1 over w, then you can really represent this whole thing in one path. In other words, if the multipath spread is very, very small and the sampling time interval is big, everything just appears under one tap and that's called flat fading. When you have flat fading, the output that comes out is really just a faded version of the input and there isn't any time dispersion on it. It all just -- it's an undistorted version of what got sent, except it's very slowly fading and coming back up again and fading again and coming back up again. If you have a slightly wider bandwidth, then you need multiple taps. If you look at most of the of the systems that are being built today, the largest number of taps they ever used in this kind of implementation is three to five or something in that order, so it's not a huge number of taps that we're talking about, because the bandwidths are not huge.

But anyway, statistically what we wind up with is that if the bandwidth that we're using is relatively narrow and it's not too many times what this term 1 over L is -- namely, this multipath spread, then we wind up with a small number of paths. We can sort of assume that we have a large number of paths coming in on each tap and then what we're doing is that each tap strength is going to be a sum of a bunch of terms which are essentially random. If you add up a bunch of complex terms that are essentially random, what do you get? If you add up a very humongous number of them, you get something which is Gaussian. If you were modeling these things statistically, you would like them to be Gaussian. So what do you do? You assume that they're Gaussian, which is what everyone does. As a slight excuse for that, what we're really interested in if you're designing cellular systems, you're trying to build cellular receivers which will deal with all of the different circumstances that they come up against and as you walk around or drive around using these things, if you look at the ensemble of different situations that these phones are faced with, over that ensemble of things, each of these taps is pretty much going to look like a

Gaussian random variable. If you look at one sample path of one little cellular phone, then no, it won't look like that -- but as far as a design tool, it looks like something which you can model as Gaussian because you're going to view each of these taps in this filter, in this time varying filter, the sum of lots of unrelated components and therefore, we're going to take the density of these taps and a statistical model as being jointly Gaussian -- a proper Gaussian random variable for each of them.

What happens if different taps is jointly Gaussian? So what we wind up with is a model for wireless channels then, where the channel itself is a Gaussian random process and as a Gaussian random process both in tau and in t, it's easier to look at if we look at it in terms of these discrete samples. Namely, we're going to look at the channel now, not in terms of a wavelength type of thing. We're just going to look at it as a bunch of taps, which the receiver has to add over to get to the received waveform. We're going to model each of those taps as being Gaussian. We're going to model in time t as being stationary. We're going to model them in time tau as coming up and going back down again and existing over a multipath spread that's about L in duration.

So the phase is going to be uniform with this density. This is this two dimensional independent Gaussian density that we've looked at so many times and if you look at this, it has circular symmetry to it, which means if you look at it in an amplituding phase, the phase is random, uniform between zero and 2 pi. The energy in it is exponential and the magnitude is this, which comes from that -- which is what you call a Rayleigh density. So Rayleigh fading channels are simply channels which you have modeled as having taps, which are really random variables. Real and imaginary parts, each of them have the same distribution.

I mean, you'd be very very surprised if you looked at these taps after the way we've derived them and you didn't find the same distribution on the real part as the imaginary part. If you look at the real part and imaginary part, they're simply coming as an arbitrary phase, which comes from some arbitrary demodulating frequency that you've used with a phase in there that doesn't have any physical connotation at

all. I mean, what is real is what you choose to call real and what is imaginary is what you choose to call imaginary. If you made your time reference just a little bit, a smidgen away from where your time reference is, what is real would become imaginary and what is imaginary would become real again and you can't expect the channel to know what time reference you've chosen. I mean, this is sort of a big philosophical argument, but you certainly wouldn't expect these random variables to have a distribution which is anything other than uniform in phase.

So that's the distribution we get and what we now have is we changed this awful physical situation into a very simple analytical situation where what we're doing is modeling this complex channel as just a bunch of taps in time, at some bandwidth that we're using and each of these taps is now Gaussian and the output is going to be the discrete sum as we take a discrete sequence of inputs, pass it through this discrete filter and then we add up these outputs and at that point, we have to worry about how do we detect things from what's coming out?

So they say here -- this is a really flaky modeling assumption and it's only partly justified by looking at lots of different situations that a given cellular phone would be placed in. If we look at this whole ensemble, it makes a little bit of sense, but really what we're doing here is saying what we want is a vehicle to let us understand how to receive these waveforms. We know there is going to be fading. Since we know there is going to be fading, we have to find some kind of statistical model for it if we're going to find sensible things to do and this is the one we're going to pick.

It's very common instead of using a Rayleigh fading model to use a Rician fading model. A Rician fading model makes sense because if you have a line of sight directly to a base station, the waveform you get from the base station is going to be kind of strong and all of these waveforms you get reflected from other obstacles and so forth are going to be rather weak. So in that case, you're going to have one strong received waveform, which is one big peak and everything else is going to be very weak. The thing that that says is the kind of statistics that you want are not a sum of a lot of little tiny things, but the sum of one big thing and lots of tiny things.

That's not a good model either. What you would really like is a sum of a bunch of medium sized things most of the time, but about as far as anybody gets talking about statistical models for fading channels is talking about Rician model and talking about Rayleigh models and then worrying about what happens the different paths of these filters. The trouble with Rician random variables is they're very complicated. They have vessel functions in them. If you read the notes I passed out today, you find that when you try to detect things, when they've gone through this kind of Rician fading, everything is just a bloody mess. That's the way it is.

Since we are talking now about a statistical model for these channels, we would like to have something called a tap gain correlation function, which is going to tell us how each of these taps vary in time. That's just the expected value of Gkm and G of k prime n. What we're going to assume, and what is really a very good assumption, is if you look at the things that add up as taps under a particular value of k and under some other value of k -- in other words, all of the paths that are coming in in one range of delays and all the paths that come in on another range of delays, those things are pretty much independent of each other. So almost always, when you start calculating tap gain correlation functions, you assume that this is independent when k is unequal to k prime. So really the only thing you're concerned about here is how these quantities vary in time. We've already talked about how they vary in time. In other words, that was why we talked about all of this business about time coherence and Doppler shifts. The thing which is causing these things to change in time is these Doppler shifts. We have already decided that the amount of time the channel remain stable is about 1 over 2 times this Doppler spread.

Here we have an analytical, statistical way of doing the same thing. In other words, you can look at this correlation function here and you can see how long it takes before it drops essentially to zero. So you can define the time coherence in terms of this duration here just as well as the way we've already done. In other words, how big does n have to get before this thing gets small? What's the advantage of doing it this way instead of in terms of Doppler shifts? If you want to measure it, you don't have to go out and measure where these paths are. If you want to measure it, you just take a cellular phone and you measure where these things are coming in and

how they change in time and that gives you a direct view of what the time coherence is. Why that W was there, there's a problem in the problem set that will give you some insight into why it's there, but otherwise it's just arbitrary.

I want to spend the rest of the time talking about Rayleigh fading, because Rayleigh fading is something you can analyze easily. We're going to do something even simpler. We're going to do Rayleigh fading where you have a single tap model. In other words, where the bandwidth that you're using is small enough that all these different paths all fall in the same delay range. In other words, it's where the multipath spread, L, is small relative to 1 over W. That's what's called flat fading, so we're assuming flat fading and we're assuming that it changes in time according to a Rayleigh model.

Suppose that we do something really simple, like trying the antipodal signalling that we've been using all along, which is really neat when you have Gaussian noise. It's just about the best kind of signaling that you can use. It doesn't work at all here and it doesn't work at all because since this channel is Rayleigh fading, you don't know what the phase is of the channel, so you send a bit and what comes in is something which, sure, it has an amplitude to it, but it has a completely random phase to it, so you can't tell the difference between one and minus one. The only way you can tell is to send first one signal and then another signal, but if you send one signal and another signal, you might as well choose a single pattern which lets you do this sensibly.

So the particular thing that we'll look at is something called post position modulation. This is essentially the same as about ten other schemes and we'll talk about that a little more later, but what you're going to use is two degrees of freedom, which is two different time instants. In the first instant, you're going to send an a. In the second instant, you're going to send a zero where conversely, you're going to send a zero in the first degree of freedom and an a in the second degree of freedom. So you can do this either with two different frequencies, you can do it with two different time instants. You can do it in all sorts of different ways. Any time you have two degrees of freedom, you either use one degree of freedom or you use the other

19

degree of freedom. It's not like the situation we've talked about before, because these degrees of freedom are now complex degrees of freedom. If I send the one, it's going to come in as some complex number, which is going to have uniform phase.

So my hypotheses now are one hypothesis is a and zero and the other hypothesis is zero and a. What's going to happen with these hypotheses is that first, the channel is going to do its thing to them. In other words, if I send this, the channel is going to take that a. It's going to multiply it by a circularly symmetric Gaussian random variable, so it's going to come as some blob and in the other dimension, what I send is a zero, so what comes in there is nothing. I'm going to add noise to this, so when I send a zero, I'm going to get -- from the channel, I get a blob in this first degree of freedom. The noise adds another blob, which is also circularly symmetric. What comes in the other degree of freedom is zero on the channel and a blob of noise. So I'm going to base my detection on either a big blob or a little blob and on the other hypothesis, I got a little blob here and a big blob there. So the question is, you have to look at this pair of blobs and decide from it what you think was the most likely signals. Sounds hard, but it turns out that it's easy, because look, under the hypothesis H_0, what you're going to receive -- and these are a times a complex variable, G_0 plus Z_0. Both of these are Gaussian random variables, both zero mean, both circularly symmetric and V_1 is going to be just Z_1. So here I have a sum of two Gaussian random variables. Here I have one Gaussian random variable. V_0 is complex Gaussian and its variance is a squared plus N_0 times W, because that's what the noise is.

So I look at the probability density of these two complex Gaussian random variables, sample value v_0 and v_1, conditional on H_0. What I have is some constant here -- I don't even care about it -- times e need to the minus v_0 squared divided by a squared plus wn_0. This is the variance of V_0, because V_0 was a squared times G_0. That's the a squared there. Variance is Z_0 WN_0 and for the other random variable -- that's the little blob random variable -- probability density of that is minus v_1 squared divided by WN_0. For the alternative hypothesis, I get the

same constant out front, which I haven't even bother to write down, times e to the minus v 0 squared over WN 0. Here are the other variables: The one where we have the big blob, so that's the variance a squared plus WN 0. I take the log likelihood ratio, which is to take the logarithm of the ratio of these two hypotheses, of these two likelihoods, and when I take this quantity and divide it by this quantity, these are both up in the exponents, so I'm just taking this and subtracting this. When I subtract this quantity from this quantity and then subtract this quantity from this quantity, what I get is just this quantity here.

I mean, you can see it's simple because this difference is going to be the same as this difference, except they have a different sign. So it's just algebra to find out what happens when you when you take this minus this and then you take minus this plus this. This is what you get. Then you look at that for a bit and you say, how do I find the probability of error from this? Before we find the probability of error, the first thing we have to do is say, what's the maximum likelihood rule to use here? Maximum likelihood rule is you take the log likelihood ratio and if it's positive, you choose v0. If it's negative, you choose v1 and if it's 0 -- if 0 was zero probability, it doesn't matter what you do.

So what we're interested in now is, what is the probability that this quantity here is going to be bigger than zero? But you see, this isn't hard because when you look at the difference between v 0 squared and v 1 squared, v 0 squared is an exponential random variable. It's just the energy of a complex Gaussian. So we're taking one exponential random variable, subtracting off another exponential random variable and we're just looking at it over on the region where these things are positive. When you integrate that, this is the answer that you get. I mean, it's not hard to integrate it. I just don't -- and if you look at the notes, the notes carry it out in all of its gory detail, but it really is about two steps.

The thing we want to look at is this result here, because it looks so different from all of the error probability results that we've gotten before. I mean, before, error probabilities go down exponentially as a square of the signal amplitudes that we're using and they go down exponentially with the energy that we're using. Here, this

crazy thing is going down as 1 over 2 plus energy to noise ratio. This is 1 over 2 plus Ed over N 0.

So it really goes down slowly. So what's going on? I mean, why is it that with all of these Gaussian random variables where things go down so fast, we're really in such deep trouble here? The trouble is, every time the fading is significant and the fading is significant whenever -- under hypothesis 0, if you don't have much of a channel -- namely, if both v 0 and v 1 are both very -- if both the real part of G 0 and the imaginary part of G 0 are both very small -- which happens with reasonable probability -- then you don't have a chance of trying to decode correctly because the two signals are both going to look the same, because both of them are just Gaussian random variables.

So as a result of that, when you're dealing with Rayleigh fading, you really have to find something else to do in order to make this work. I mean, if you have Rayleigh fading and you just try to communicate in the way that we tried to communicate before -- just sending individual bits and hoping for the best, you're in very deep trouble.

So next time we'll start to talk about ways of dealing with this using diversity, using channel measurement, using all sorts of other things, but -- I don't know what happened to the person who was supposed to evaluate the class, but maybe they will come in next week. I don't Know.