VLSI Data Conversion Circuits
Prof. Shanthi Pavan
Department of Electrical Engineering
Indian Institute of Technology, Madras

Lecture - 12
FFTs and Leakage

This is VLSI data conversion circuits lecture number12.

(Refer Slide Time: 00:13)

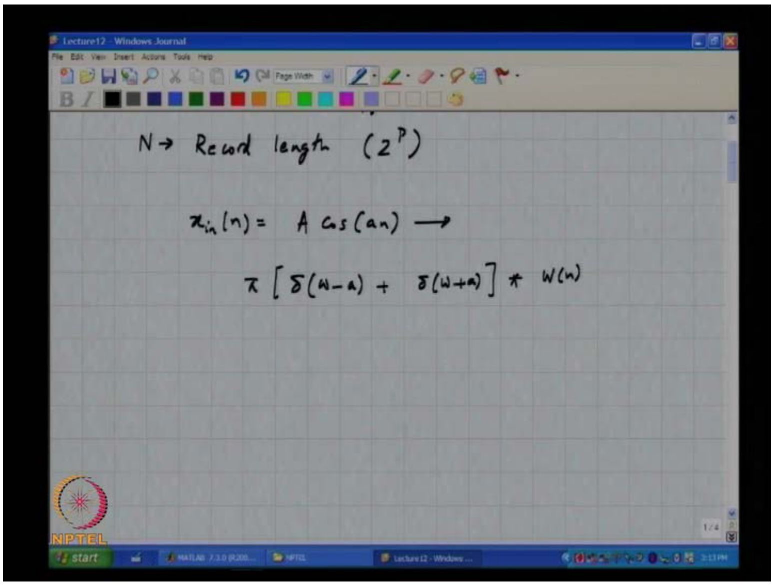

In the last class we were looking at how to estimate the spectrum of a sine wave right; given that the input frequency is known and the sampling rate is known. And, we are trying to characterize the sample and hold. And, we said that the input frequency cannot be chosen in any old way f in by f s must satisfy m by N; where N is the record length and m must be an integer ok. So, if this condition is satisfied then if we find the discrete Fourier series corresponding to the length N; you will get N coefficients of the discrete Fourier series.

And, if this relationship is valid then you will see only 1 coefficient being non 0, all the other coefficients will be 0. In practice the other coefficients will be very small corresponding to numerical precision of the F F T computation right. And, as I was mentioning to you last time N is called the record length. And, often chosen to be 2 to the power of p where p is an integer. The advantage being that it turns out that their very fast algorithms to compute the Fourier series coefficients; if the record length happens to be a power of 2, right.

And, as we saw each one of these Fourier coefficients is what is called a frequency bin right. And, therefore since there are N coefficients and this spans a complete a frequency range of 0 to f s by 2, right. In f s by 2 there will be N by 2 bins; which therefore means that the width of a bin right in real frequency terms is f s by N. So, we were wondering what would happen if the input frequency or the input tone did not lie on a bin in other words if f in by f s is not m by N; but there is a fractional part to N.

And, we saw the last time around that when the input tone does not lie exactly on a bin there is something called leakage right. And, that is happening because one can think of it as follows the discrete Fourier series is nothing but a sampled version of the discrete time Fourier transform. If the sequence is a finite length 1; the discrete Fourier transform of the sequence is simply the infinite extension multiplied by a rectangular window of length N which in the frequency domain corresponds to convolution of the Original spectrum with non periodic spectrum Periodic term with that of the window correct.

So, in other words if we had a sinusoid which is not on a bin; let us say. So, the input is of the form A cos of (a n) where small a denotes the frequency. And, we are only taking a finite number of samples say capital N samples. Then, the Fourier transform the discrete time Fourier transform of this sequence is nothing but the delta function convolved with the spectrum or Fourier transform of the window function. Since, you are taking only a finite window of samples it is equivalent to having a window which is the so called rectangular window.

And, as we saw the last time around the spectrum or that or the Fourier transform of the rectangular window is a sync pulse. So, what we are in effect having for the finite length sequence is one which has a Fourier transform, which is this impulse convolved with Sync pulse ok.

(Refer Slide Time: 05:34)

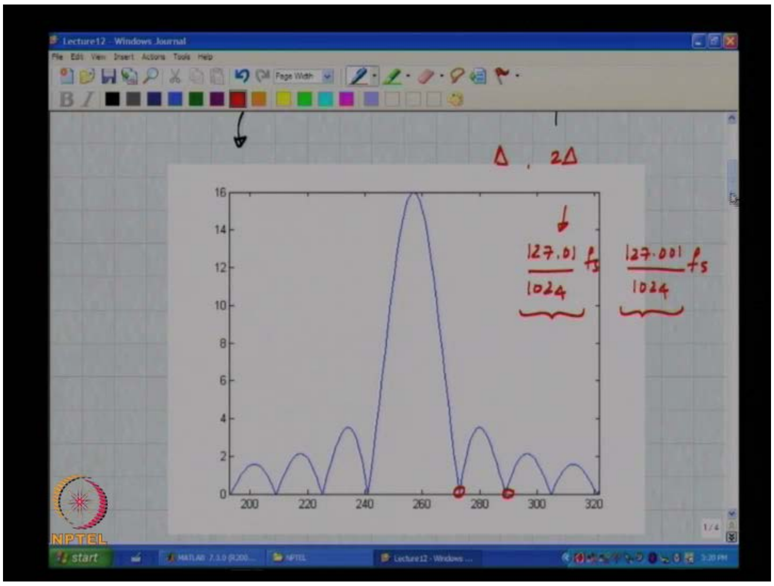

So, in other words it will look like this; let me show you a picture right. And, if the input lies exactly on a bin what happens? This is the discrete time Fourier transform of the windowed sequence; what we are doing is computing the discrete Fourier series coefficients of this sequence right. And, as we saw the last time around that is nothing but a sampled version of this. Now, if the input lies exactly on a frequency bin; then, we will be sampling these points right. And, which is why it makes sense that when you have a rectangular window. And, the input lies on a frequency bin which is equivalent to saying that f in by f s these are the form m by N; you will see only 1 coefficient which is non 0 when you compute the discrete Fourier series of that sequence.

Now, if the input is not on a frequency bin in other words f in by f s is not of the form m by N where m is an integer; then, what is happening the this small number A is no longer of the form 2 pi by n. So, what we are doing is essentially sampling the spectrum slightly of the centre if you want to put it that way; so, for example here, here and so on. And, when you look at the discrete Fourier series coefficients, please note that this is a linear scale; what you will see is something of you join these points, you will see something like this, this is clear.

Now, as a side point if the deviation from the bin was originally delta; that corresponds to this deviation right. And, if the deviation increased to 2 deltas what can you say about the level of the leakage.

Student: Increases.

It increases right by how much? Earlier, let us say the input tone was 127 by 1024 plus say point over 1 times f s and naturally we expect leakage; because the input does not lie on the exactly on the bin correct. Now, if I make this 127.001 by 1024; now, also we expect leakage for precisely the same reason. Now where do you think the leakage will be more? We had 2 frequencies let us say 127.01 by 1024 times f s and 127.00 1 by 1024 times f s; where do you think the leakage will be more?

Student: ((Refer Time: 09:50))

Obviously, this will leak more right; the question is how much more? This is just a side point. But I thought it is quite interesting too; can you look at the diagram and tell me this is a linear scale.

Student: This is the ratio of the sync pulse sine square sir.

Yeah, of course it is a ratio of the sync values. But can you get more specific then that; if you are exactly on a bin what will the leakage be?

Student: 0.

0 right; so, that is you are here; if you deviate slightly from that position what happens to the Fourier coefficient as you can see from the picture.

Student: It increases.

It will increase; for small deviations?

Student: Linear.

It will be linear. So, whether I mean; so, when you compare the leakage of 127.01 versus 127.001 what should we expect?

Student: ((Refer Time: 11:23))

You should see.

Student: ((Refer Time: 11:25))

Expect to see a leakage which is.

Student: 10 times more than.

10 times more than 1 with respect to the other right; that day I thought we were going through this example ok.

(Refer Slide Time: 11:50)

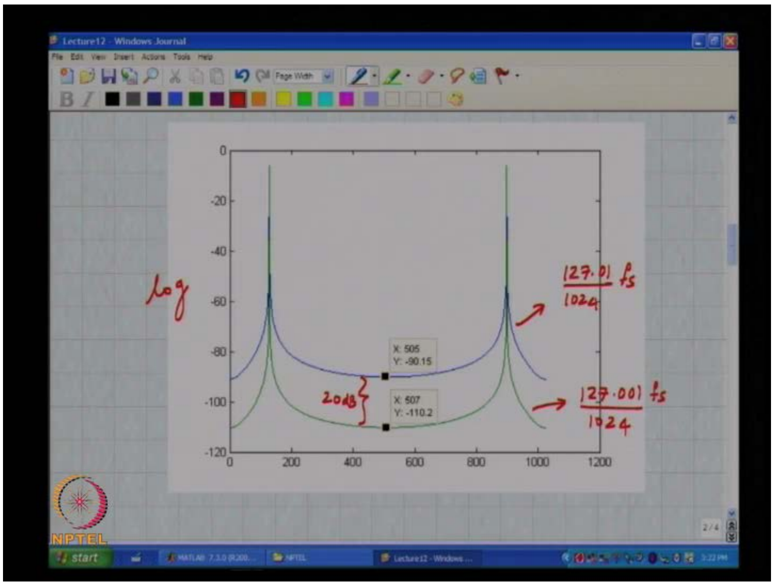

We discussed this other day right. I mean one of them was a tone of frequency 127.01 by 1024 times of f s and the other one is 1 27.001 by 1024 times f s. And, as we expect the leakage is smaller; the other day we made wild guesses and said one is 127.001 right. But it indeed makes sense that what we should expect is the leakage should go down by a factor of 10 which is.

Student: 20 d B.

20 d B this is a log scale remind you and definitely the difference is 20 d B is this clear all right. So, the moral of the story so far is that if the input and the sampling rate are synced. In other words it is possible to somehow make sure that f in by f s equal to m by N where N is the record length and m is an integer. If it is possible to do that then the input frequency will lie on a bin and there is no problem right. And, please mind you this is just simply trying to capture an input sinusoid; we did not talked about more things which happen in practice which is noise and all this other stuff. But simply to capture a

sinusoid and make sure that from the samples we are able to figure out what sinusoid it is all this has to happen right.

And, since we are on the topic; you might as well discuss what happens in those practical cases where for some strange reason or the other it is not possible to make the input and clock in sync. I mean as we discussed the other day to make the input and clock in sync it is trying to set your watch and your friends watch to be ((Refer Time: 14:16)) exactly, you know exactly in sync which means that periodically one would have to sync up both their watches either to a master reference or to each other; you understand.

And, that is a negative feedback system if you are going too slow your speed yourself up, if you are going too fast you slow yourself down; such a negative feedback system is what is called a phase lock loop and we will not go into the details in this class right. But sometimes it is simply not possible in order I mean that the clock source and the input sinusoid are able to talk to each other and sync themselves up. Then, we need to wonder you know what to do right. So, if I had chance that the input tone and the sampling rate will be exactly related by m by N; even if there were related according to an integer ratio at some point in time soon enough they will drift away and the relationship will no longer be valid.

So, we need to figure out what we need to do to fix this problem right; in other words if this ratio is not satisfied how do we, I mean what do we do with the sequence. So that we at least try to mitigate this so; called leakage to some degree right. So, before we get to the solution; let us kind of look at the problem in another light which is in the let us try and interpret this in the time domain right.

(Refer Slide Time: 16:05)

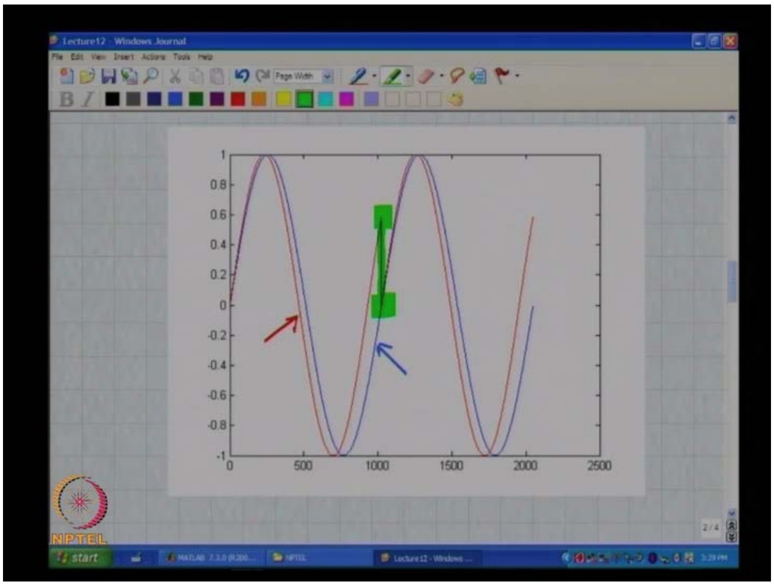

So, if the input sinusoid lies exactly on a bin correct; please note that the discrete Fourier series all it is doing is making a periodic extension of the sequence that you have; and computing a Fourier series of this periodic sequence correct. Now, in other words if the input was on a bin like the sinusoid in blue here right. Let me show that to you when you make a periodic extension what do you think will happen at the end points.

Student: ((Refer Time: 17:08))

I mean in other words when you put make a periodic extension they will be no kinks at the right; whereas, if the input sequence does not or the input sinusoid, does not lie on a bin what happens at the end points? What you will see is some curve in an example of that is shown in red right where at the end point I mean it is not periodic obviously. So, at the end points there is a discontinuity; can you comment on the magnitude of this discontinuity as you as time progresses? Let us say I took another set of 1024 samples; do you think the discontinuity would be the same?

Student: ((Refer Time: 17:55))

It will be different why simply because the sequence is not periodic correct. So, when the input tone does not lie on a frequency bin what you must remember is you are finding the discrete Fourier series corresponding to a sequence like the one shown in red; where at the ends of the sequence there is a discontinuity right. A discontinuity is nothing but an impulse correct. So, there is a jump here and how frequently does this jump occur?

Student: ((Refer Time: 18:41))

It occurs every N samples this corresponds to a fundamental frequency of.

Student: ((Refer Time: 18:50))

I mean 1 by N samples; the periodicity of this jump is 1 by N samples. And, whenever you have a jump in a signal what can you say about it is frequency components?

Student: Which have all the frequency?

There are lot of high frequency components; in other words even though the fundamental frequency of this jump is has I mean is 1 by N samples it will have a whole bunch of harmonics correct. So, this jump is what is responsible for causing leakage when you take a rectangular window does it make sense right. So, every N samples you have a jump correct. And, the jump consists of impulses which is got a lot of high frequency content and that will simply the jump and it is harmonics will now add to the spectrum of the sinusoid is this clear yes ok. So, now that we understand this let us try and figure out; what we do about this problem.

(Refer Slide Time: 20:24)

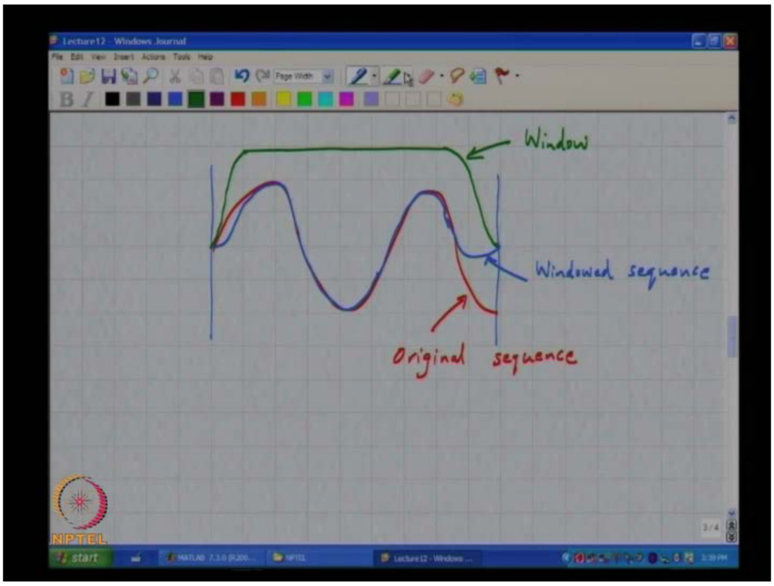

So, we have a record length; please note that even though these are all sequences I will just draw them as a continuous wave otherwise plotting the sequences is a takes a lot of time you understand. So, even though I show a continuous wave form; you must understand that these are actually samples you may perhaps think that they are. So, many samples that when you plot all those things together it looks like a like a curve. So, now we have a situation where we have a fixed window length. And, the sinusoid is not quite periodic; in the sense that there are discontinuities at the ends when I make a periodic extension of this sequence there are kinks occurring at every N samples right. So, the question is what do we do about this how do we fix the problem any suggestions?

Student: We try to make that discontinuity go to 0.

So, we know that this is all the data in the middle is all seems quite right; it is only the data at the ends which is problematic right when we were taking, when we were using a rectangular window what was happening was that all samples of the data are equally important correct. But we know that in reality the leakage is occurring; because of this discontinuity all right. It is not fundamentally a problem with most of my data there are only a few guys in the end who are contributing a lot to this leakage right. Any time, you have a kink in a wave form it will show up in the Fourier spectrum as high frequency components or as leakage ok.

Now, since we want this part and we do not want this part; because that is the 1 that is causing leakage what do you think we can do? In other words the samples in the middle are important, the samples in the end are not so important.

Student: ((Refer Time: 23:11))

I mean if you remove them I mean if you chop of the samples in the end; then, what you chopped off will become the end is not it that is clearly not a solution; you understand it is like saying some the some guy who is getting off at a you know travelling by train. And, then who got off at the platform to fill water right and there was a long line and by the time he filled water the train started moving. And, then you know he missed the train, he tried to catch the last compartment and he could not make it. And, then he missed the train and then he lost a complaint saying or they should be the problem is with the last compartment.

So, they should not be a last compartment in the train you know; you understand that is a I mean chopping off the one is that you do not want right does not help you at all ok. So, what do you think we should do?

Student: Multiply by.

So, you say that these are I mean all are equal; but some are more equal than the others right. So, you say I will weight all these since all these seem equally important by the same factor; but I know my end there is a not so, important. So, I will taper off my weighting function like this and like this does it make sense. So, in other words I will take this sequence, I will multiply it by another sequence; which remains fairly constant for all the middle samples right and tapers off smoothly to 0 at the ends you understand.

Now, if you do this what has happened to the resulting sequence; the resulting sequence looks like this does it make sense? It is simply this is the so called windowed sequence; the one in red is the original sequence and this is the window function. Does it make sense? I mean you can think of it as a instead of using a rectangular window where all samples have the same weight; we realize that the leakage is a problem arising due to discontinuities at the ends.

Therefore, what you want to do is to not treat all samples with the same levels of importance; the one is in the middle are more important than the one is in the edges. So, in other words you multiply this by another sequence which tapers off smoothly to 0 at the ends; it is very important for this window function to taper off smoothly why?

Student: The idealities.

I mean if there are discontinuities in the window function then.

Student: ((Refer Time: 27:27))

I mean they will also be present in the in the windowed sequence which will again lead to leakage you understand. So, this is the basic idea behind what is called windowing the sequence before taking it is computing it is discrete Fourier series coefficients. And, clearly the spectrum of the windowed sequence is not the same as the spectrum of the.

Student: Original.

Original sequence; it has to be different because I have taken the original sequence and modified it by multiplying it with some within codes known function right. But what one must intuitively expect is that the leakage will be less compared to using a rectangular window. Because the discontinuities are now gone; I mean it might also do something to our desired signal, it may make the sinusoid look different right. But at least we know the window we are using. So, hopefully we will be able to make sense out of what we get at the output; is the motivation clear for using windowing ok.

So, if I just draw a window function like this which says it must be you know the samples in the middle must roughly be given the equal weight and the samples at the end must be given 0 weights. And, it must taper off smoothly to 0; what do you think I mean how many ways are there do you think of coming up with a window sequence which can do this.

Student: for any number sir.

For any number of ways of coming up with a window sequence which gives approximately the same weight to samples in the middle and smoothly tapers off to 0 at the ends right. So, corresponding to each one of these things; there are any number of windows which are called spectral windows, which are around in the literature. People keep you know have discovered and keep discovering new window functions where there is a minor variation here and there right.

But the basic idea is something where which is flat on the top and slowly tapers off to 0 at the ends.

(Refer Slide Time: 30:30)

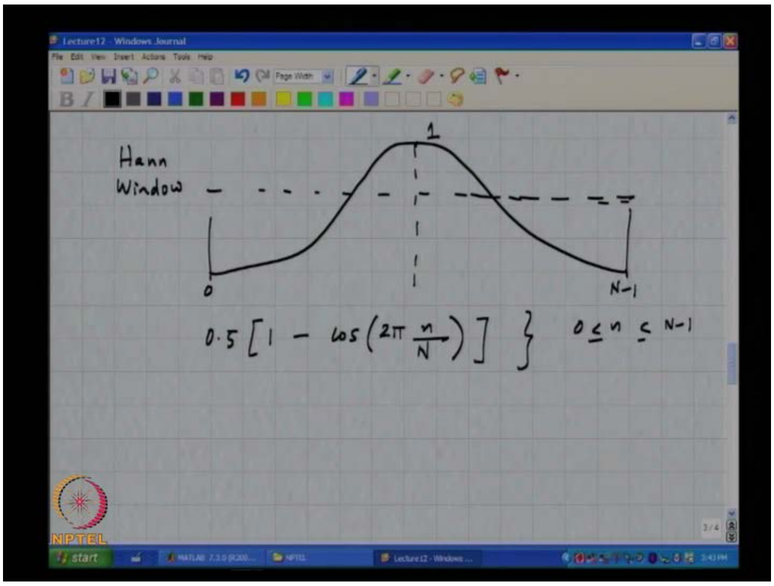

So, let us take a look at some common window functions and see their properties; as I said a window function is job is to start at 0, end at 0 and kind of emphasize the samples in the middle, more than the samples in the end. So, one function which kind of looks does this you know I have drawn a very smooth looking curve; can somebody tell me what kind of, can somebody give me a functional relationship for.

Student: Sir this is a Gaussian.

Well, it is not a Gaussian can never go to 0 right.

Student: ((Refer Time: 31:07))

So, this looks like a.

Student: ((Refer Time: 31:11))

I mean cosine wave you know pushed up; instead of going from minus 1 to 1 it goes from 0 to 1. So, see you have discovered a window yourself. So, it is. So, this is what is called a race cosine or a hann window all right. And, is given by half times 1 plus cos or 1 minus cos n by N times this is what 2 pi. So, it starts off with 0 and ends at 0 and in the middle it reaches. So, the average is half and then this peak is 1 then goes from 0 to 0 all right.

Now, let us see what the raised cosine window does to a sinusoid which lies on a bin right. And, what do you expect; earlier when we had a rectangular window what do we the discrete Fourier series was sampling the discrete time Fourier transform or the window multiplied by the sinusoid correct. It is the same thing now it is except that earlier the discrete time Fourier transform of the rectangular window was a sync function. Now, it is not a sync it is it is something else. So, this is I mean please note that this is valid only within 0 to N minus 1; does it make sense?

(Refer Slide Time: 34:14)

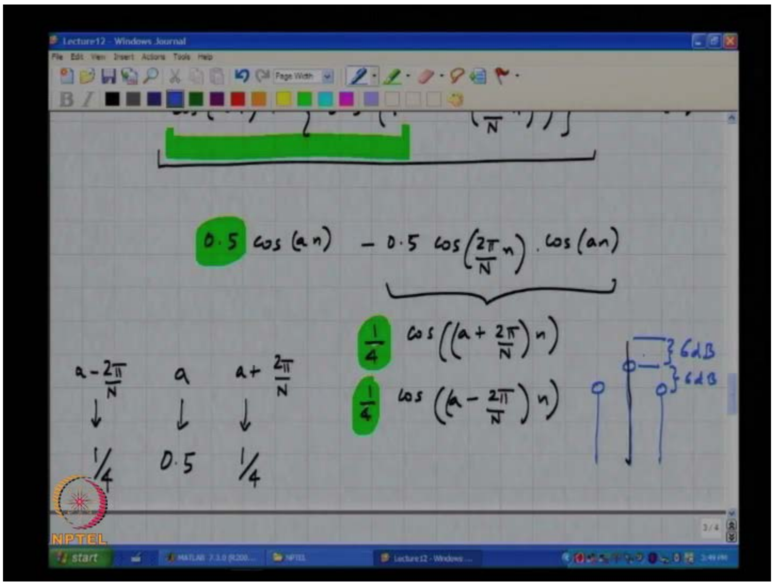

Now, So if you had A cos (a n) and then you multiplied this by this window function which is half times 1 minus cos 2 pi by N times N; this is the sequence who is discrete time fourier transform is being sampled when you compute the F F T is this clear ok. So, now can you tell me if the input lies on a bin what do you think you must get?

Student: ((Refer Time: 35:04))

So, I mean this you can also think of it as being multiplied by since this only is valid between 0 and N minus 1; you can think of this as being multiplied by a rectangular window of size n. So, what do you think you will?

Student: ((Refer Time: 35:28))

And, those bins corresponds to.

Student: Correspond to the fundamental and 1 to the left and 1 to the right.

1 to the right; so, as you can see here cos (a n) is getting multiplied by half which basically means that it is simply half cos (a n). And, it is also getting multiplied by half cos 2 pi by N; times N times cos (a n) of correct and how does this look like; cos a, cos b.

Student: a plus b.

You will have only 2 components; quarter times cos of a plus 2 pi by N and quarter times cos of a minus 2 pi by N right. So, if a lies on a bin a plus 2 pi by N will also lie on a bin and a minus 2 pi by N will also lie on a bin correct and all the other bins are 0 is that clear. So, in other words if the input lies on a bin; then, for a rectangular window you would see only how many of those coefficients should be non 0.

Student: ((Refer Time: 37:18))

1 if you take only 1 half they will only be 1 coefficient which is non 0; when you use a hann window how many do you expect to see? Expect to see 3 and where are the 3? 1 corresponds to the input itself, 1 bin to the left 1 bin to the right. And, please note also that there is a very definite amplitude relationship between the main bin and 2 sidekicks all right. So, the main bin is has a coefficient or an amplitude which is half the input amplitude correct; in other words if I had used a rectangular window I would have got some strength for the bin a, assuming that a was on the on 1 of the bins.

Now, I will only get half a. And, then the so, this is a; a plus 2 pi by N what will I get?

Student: ((Refer Time: 38:34))

I will get an amplitude which is one-forth and a minus 2 pi by N; I will get an another amplitude which is one-fourth, does it make sense correct. So, in other words if the rectangular window gave me something like this for that bin a hann window will give me non 0 values for 3 bins. This bin will be 6 d B below what I would get if I used a rectangular window. And, similarly this would also be 6 d B below is this clear half is 6 d B.

Student: Sir what is that the half we do not need a half?

Yeah, I mean we do not need a half; I mean the question is why are we getting these half is right. I mean thankfully and not surprisingly; the these multiplication factors you get are all dependent on the input I mean these are the input amplitude multiplied by these factors. So, in that sense this is not a disaster; because if I know the window I am using. And, if I find you know a by 2 a and a by 2; I know that the true amplitude is 2 a is that clear.

(Refer Slide Time: 40:50)

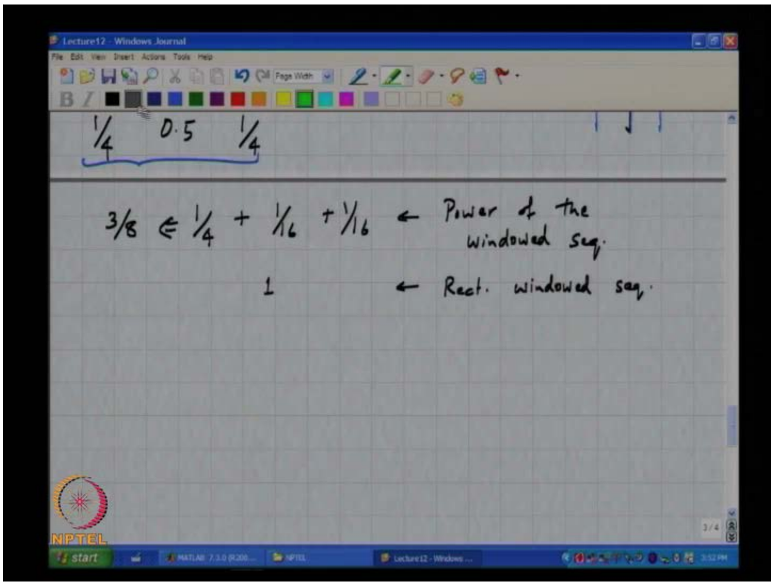

And, what can you say about the power of the sequence? The power of the windowed sequence is one-forth plus one-sixteenth plus one-sixteenth is the power of the windowed sequence. And, when you compare it with the input power, I am neglecting the half factors by which have to be there but the ratio of the powers is what?

Student: ((Refer Time: 41:24))

It is.

Student: ((Refer Time: 41:24))

So, rectangular windowed; so, is ratio is 3 by 8 in other words the power of the windowed sequence is 3 h; the power of the same sequence had it been windowed by a rectangular window why does this make sense?

Student: ((Refer Time: 41:56))

I mean is it something which is surprising or is it something to be expected.

Student: It has to be expected.

It has to be expected why?

Student: ((Refer Time: 42:07))

Great; because we are basically waiting only some samples in the middle were something close to one right. All the samples other than those are getting waited by factors which are.

Student: Less than.

Less than 1 which means that they contribute lesser amounts to the power. Therefore, it is it definitely make sense that the power of the windowed sequence is smaller than or rather to be more correct I must say; the power of the hann windowed sequence is much smaller than the power of the rectangular windowed sequence.

(Refer Slide Time: 43:19)

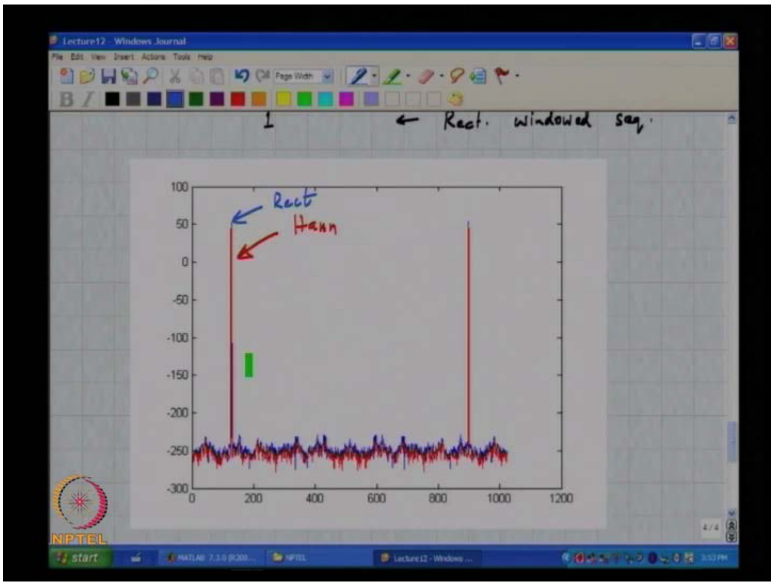

Now, let me show you a picture ok. So, the red one here is hann whereas the blue is the rectangular windowed version; this is intended to show you that clearly the blue one is sticking slightly above the red one right they zoom in there and take a look. But I wanted to show to you that all the other coefficients are.

Student: 0.

0; that is only possible if the input lies on a frequency bin correct. So, when I zoom this is what you get.

(Refer Slide Time: 44:06)

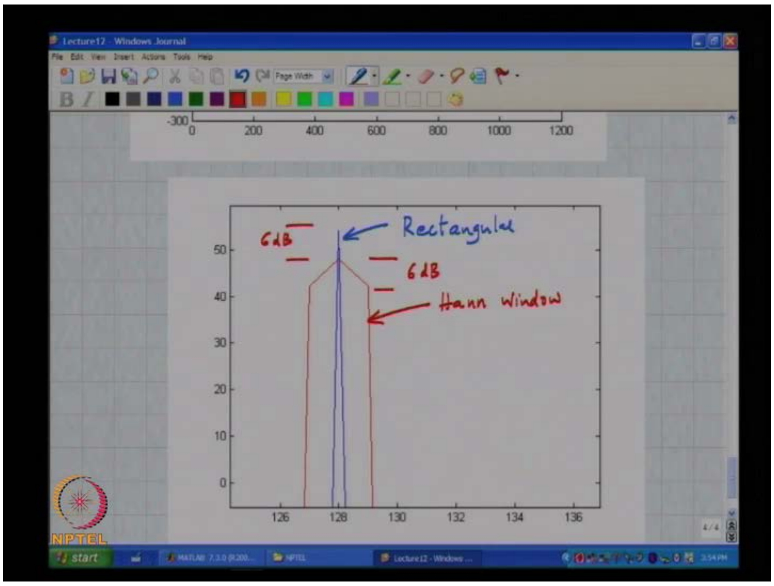

So, this is the rectangular window and this is the hann window right. And, as you can see this is 6 d B and so, are these is this clear all right. Let us see what happens when I mean finally; please note that the whole idea behind coming out of this window stuff is to see what happens when the input is not lying on a bin you understand ok. So, and before I go there let me retreat the fact that; if the input is on a bin the input power is now spread across 3 bins. So, if you want to find the true input power from the hann windowed FFT what you need to do is find the power in all the 3 bins and then multiply that by what?

Student: 8 by 3.

8 by 3; because 8 by 3 corresponding corresponds to the power loss of the hann window because of this waiting is this clear.

(Refer Slide Time: 46:46)

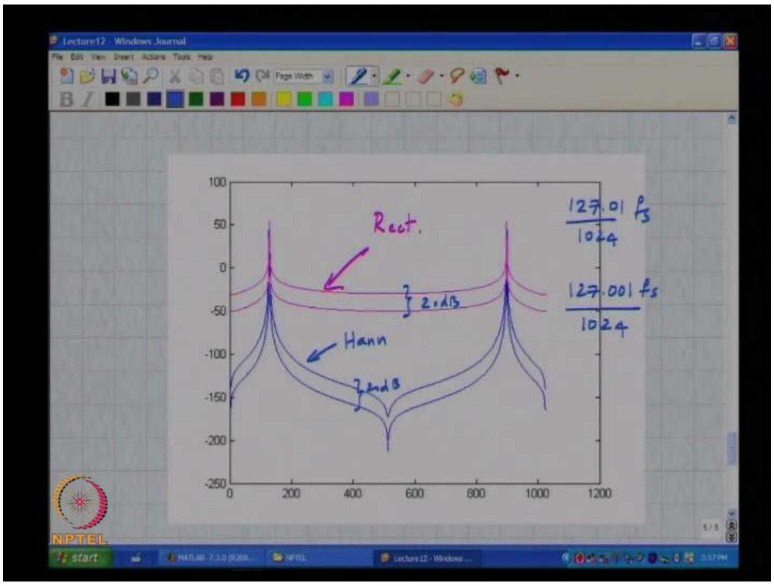

Now, let us see what happens when a hann window is used and the input is slightly off bin ok. So, the one is in magenta correspond to.

Student: Rectangular window.

Rectangular window and the one is in blue correspond to hann windows. And, in both cases the 2 tones used are 1 27.0 1 by 1024 f s and 127.001 f s by 1024. And, as you can see still there is the I mean there is leakage in both cases right. But the leakage when you use a hann window is?

Student: Less.

Is much smaller than the leakage you get when you use a rectangular window. And, this makes sense I mean how much improvement you must get is not immediately obvious right. But the fact that you do get an improvement is and makes sense all right. And, again as you can see as the frequency deviation goes from 0.001 to 0.0001; the leakage in both cases reduces by

Student: ((Refer Time: 48:23))

20 degree and you know why all right. Now, that we have considered, you have seen the advantages of windowing one might ask you know what is so, scared about the hann window. And, the answer is there is nothing really scared about it; except the fact that if

the input lies on bins then, the I mean the hann windowed spectrum will have just 3 components. And, not surprisingly there are any number of windows that are available right. In data conversion work I mean this is what is often done if you are able to sync the input with the I mean with the clock; and you are building what is called a Nyquist rate converter; then, there is no problem at all leakage does not happen because they are synced properly and you use a rectangular window right.

Now, there are cases when you are not able to sink; in which case you use one of these window functions we will discuss more in detail what that is right. There are also cases where the input is lying on a bin right. But in addition to the signal there is noise; in which case we had to worry about? What to do? I mean why it might be for as far as the signal is concerned; the signal may not leak because it is sitting on a bin. But there is noise in addition and the noise being what it is right there is discontinuities at the beginning and the end of the record and that will cause leakage right. So, we will see all this in the next class.