TheFourierTransformAndItsApplications-Lecture21

Instructor (Brad Osgood): Here we go. All right. The exams, you remember the exams, I think they've all been graded, and the scores have all been entered. Although, I don't think we've made the scores yet visible on the web site. I will do that after I get back from class today. You can pick up your exams from Denise Murphy, my admin in Packard 267. She has them all, so I'll send that announcement out to the class. Any questions or comments on that?

Good. Okay. Today, we're going to continue with our discussion of the DFT. This is getting to know your DFT. Your Discrete Fourier Transform. Now, the subtitle, I would say for this, should be – you already know it. The point of the way we're talking about the Discrete Fourier Transform is that it can be made to resemble the Continuous Fourier Transform in a great many ways. So the intuition that you built up for the Continuous Fourier Transform, the formulas that you've learned to work with and so on really all have analogues in the Discrete case. Now, not in all cases. There are some things that don't quite match up, and that's interesting, too, but to realize what doesn't match up so carefully and to realize why it's interesting, I think, goes along with seeing in many ways how much things are the same. So we're going to take the point of view that – and I'm going to try to take the route to make the Discrete Fourier Transform look as much as possible like the Continuous Fourier Transform. That's our point of view. You don't have to do it this way, but I think it's most satisfactory, and it allows us to leverage off what we have done in the continuous case. Now, before doing this, let me recall the definition. Here's the definition. Where we ended up last time was sort of a definition that, for the Continuous case, the idea of sampling a continuous signal to get a discrete signal, sampling a Fourier Transform to get a Discrete Fourier Transform, all that had pretty much vanished by the final definition. So the definition that we ultimately wound up with and ultimately had makes the continuous case almost invisible, vanish.

It looked like this. You have a Discrete signal. So I'm using both the signal notation and vector notation here, and I'll continue to do that, sort of mix the two up, because I think they're both useful. So the idea is you have either an N-tuple of numbers or a Discrete signal whose value at the nth point is just the value here, FM. Okay? Oops, that doesn't look right. That's not much of a statement. So you can either consider it as a Discrete signal who's defined on the integers or the integers from zero to N minus one, or you can think of it as an N-tuple or as a vector. So you have a Discrete signal, and the Fourier Transform is another Discrete signal. It's DFT. It's the discrete signal – I'll call if capital F, but I'll also use the notation to make the connection with a continuous case, a script F with a little underline under it indicating it's supposed to be a vector quantity or something a little bit different from the continuous case. It's defined by its nth component. So the nth component of the Fourier Transform is the sum from N equals zero to N minus one. The nth component of F times E to the minus two pi I, NM over N. So everything is defined here in terms of the indices in the exponential, and these are the values of the Discrete Function at the index point, F of zero, F of one, F of two and so on.

That's the definition. You don't see at all the fact that in our derivation, this came from, starting with a Continuous signal, sampling it, sampling the Fourier Transform and then somehow, ultimately, leading to this definition here. It's just as an operation on one Discrete signal producing another Discrete signal. That's pretty much how we're going to deal with it. But before embarking on that path forever, I want to have one nod back to the Continuous case. I'm going to talk a little bit about how the DFT is employed and the kind of things you have to know when you use it in practice. You already have to some extent, and you certainly will more in the future. So we want to look back at the Continuous case to talk about one additional phenomenon of reciprocity that comes – the reciprocal relationship between the time domain and the frequency domain that comes up also in the Discrete case. So one more look back at the Continuous case, the Continuous roots of the DFT and reciprocity of the two domains, time and frequency. So how did it work? We have a signal on the time domain that we discretize, and we have a signal in the frequency domain that we discretize. So imagine that you have two grids. You have a grid in the time domain. You have N sample points that are delta T apart. I'm running out of room to write things. Let me just write it here, and then I'll write some of the notation up above. This is the frequency domain, and we have spacing in the frequency domain of delta S apart. So N sample points, T zero up to TN minus one, space, delta T apart, and N sample points in time that is. And we have N sample points in frequency. S one [inaudible] say up to SN minus one, spaced, delta S apart. So there are three quantities here. There's the spacing in the time domain, the spacing in the frequency domain and the number of sample points, N. There's a relationship between them. You have three quantities. Quantities of interest. Delta T, delta S and N. They're not independent. They're not independent. That's the important point. Let me remind you again – I'm not going to go through the derivation again, but let me remind you what the relationship is. We had an N delta T. That's the spacing in the time domain that has to be – I want to say this right – is L. That's the spacing of the time domain. I want to use the notation I used last night.

This is sort of a time limitedness. We had N times delta S is the band limitedness that we called that two B. That was the bandwidth. So delta T times delta S is equal to – delta T is L over N times two B over N. That's L times two. That's two BL over N squared. You will recall that when we set up the sampling, two times B times L was equal to N, the number of sample points. This is N over N squared. That is one over N. Okay? The way we did the sampling, making this not completely justifiable, appeal to the sampling theorem and so on, gave us a relationship between how you sample in the time domain, how you sample in the frequency domain, the number of sampling points you took, and it was exactly that two BL was equal to N. It was sort of surprising or not clear, a priori, that if you carry out this procedure of sampling in the time domain and sampling in the frequency domain, you're going to take the same number of sample points in both domains.

That's what we found as a relationship there. Let me highlight it. That is delta S delta T, or delta T delta S is one over M. The spacing in the frequency domain and the spacing in the time domain are reciprocally related to the number of sample points. This is called a reciprocity relationship. All right. Now this again has practical significant and practical consequences when you're applying the DFT. You can imagine, when you have a continuous signal, you want to sample it. What can you choose and what is forced upon you? Well, you can imagine choosing how frequently you sample, right? That's the delta T, and you can imagine how many measurements you make, how often you sample, how many samples you take, which is N. So you can imagine choosing delta T. That's how frequently you sample, or the sampling rate, and N, the number of samples. If you do that, once you do that, then delta S is fixed. Delta S is determined. The spacing in the frequency domain is determined. One way of putting this is the accuracy of the resolution, or how find the resolution is in frequency, is fixed by the choices you make in the time domain.

So let's say the resolution in frequency is detailed, is fixed, by the choices you make in time. Conversely, you could say, I want a certain accuracy in – I want a certain resolution in frequency. I want delta S to have a certain fineness. That's going to detail, then, how many sample points I want to take and how I space them in the time domain. At least, you have to choose two of those things, and the third one is detailed. You can imagine, if you're doing this with real data, you have a certain freedom here, but the freedom also carries with it certain restrictions. Ain't that life? That's the way things go. The freedom you have is typically how many measurements you make and how frequently you make them. Once you do that, that determines what the resolution is like in the frequency domain. There are ways of getting around this. There's zero padding and special techniques – not getting around it, but there's ways of understanding it or massaging it. So-called zero padding, you'll have a problem on that, and you'll have other chances to experiment with that. It's built into the system. This sort of reciprocity relationship is another example of the same sort of thing we've seen so many times in the Continuous case, now carried over to the Discrete case.

It's stretched in one domain means shrunk in the other domain. Or it's reciprocity between time and frequency. As I said, we've seen many different instances of that in the Continuous case, and here's an example of how it caries over to the Discrete case. All right. I wanted to say that because it is the sort of thing that you meet, and you have to understand, when you actually apply the DFT in most common context where it's associated somehow with some sort of Continuous process that you're sampling. You want to take the Fourier Transform and you don't have a formula for it, you have to use numerical algorithms to actually compute it. That means you're sampling. The reciprocity relationship puts certain limitations on what you should expect to get. Okay. So that is my final nod, at least for now, back to the Continuous case. Let's go back now into the Discrete world and pretty much stay there. So back to this formula and it's consequences. Let me erase it and write it down again. Back to the Discrete world, the Discrete setting.

Again, my goal here is to make the Discrete Fourier Transform look as much as possible like the Continuous Fourier Transform. So here's how to do that. I want to make the DFT and associated formulas look like the Continuous case. So maybe I should say we're not really abandoning the Continuous case, but we're situating ourselves firmly on the Discrete side of things. Now, to help out in that, partly it's a matter of notation, but even more so, it's what you do with the notation and the consequences of thinking about things a certain way. So the first thing I want to do is introduce a symbol, introduce a way of thinking about exponentials that occur in the definition of a DFT. So let me write down the formula one more time. Again, F is a Discrete signal. F zero up to FN minus one. Its Fourier Transform is another Discrete signal whose nth component is given by this formula. N equals zero, N minus one, the nth component of little F times E to the minus two pi I, NM over N. Okay. That's a two pi I. First thing I want to do is I want to introduce the notation. I want to view these complex exponentials as themselves, coming from a Discrete signal, as values of a Discrete signal.

This turns out to be a very helpful thing to do for a number of reasons. It gives you a compact way of writing things. It also gives you a way of enunciating certain properties of the Fourier Transform that would be very – the Discrete Fourier Transform that would be very difficult to do otherwise. So I want to realize the complex exponentials that come in the definition here as arising from also a Discrete signal, the Discrete or vector complex exponential. People either refer to it as a Discrete complex exponential or the vector complex exponential. You pick what you want to call it. So let me give you the definition. I'll leave that word up there. Let me write it over here. I'm going to write omega as a vector or as a discrete signal to be the N-tuple of powers of the complex exponential that appear in the definition. So the zeroth coefficient, the first entry is one. That's E to the zero. Then it's E to the two pi I – there's a minus sign up here. Let me define it in terms of a plus sign, and then I'll start taking powers to get a minus sign. So either the two pi I over N. You'll see where this is going from in just a second. E to the two pi I times two over N, and so on, so on, all the way up to the final term. The nth term is going to be E to the to pi I, N minus one over N. All right?

That's the basic vector or Discrete complex exponential. So its nth component – so omega, the nth component of this, is either the two pi I, M over N. Okay? Doesn't take a great leap of the imagination to do this. Now, I'm going to define powers of this. Powers of omega. This just collects in one place the powers of the complex exponential that appear in the definition of the DFT up to the minus sign, but it views it slightly differently. Itself is a Discrete signal. So again, you can either view it as an N-vector, or you can view it as a discrete signal to find out the integers from zero to N minus one. I want to take power. Now you say to yourself, what does it mean to take powers of a vector? Well, it doesn't make sense to take powers of a vector, but if you believe in mat lab it does, I suppose. It certainly makes sense to take powers of Discrete function. So omega to the – I want to use the same notation I have in the notes here. Omega to the N is just the Discrete signal whose entries are the nth powers of the entry-level omega. So it's one E to the two pi I, N over N, E to the two pi I two, N over N, up to E to the two pi I, N-1 over N.

Of course, if I can take positive powers, I can also take negative powers. N is not restricted to be positive here, but just to write it down, omega to the minus N, same thing, is one E to the minus two pi I, N over N, E to the minus two pi I, two N over N. Then all the way up to E to the minus two pi I, N times N minus one over N. Just replacing N by minus N. Okay. So with this notation, the Discrete Fourier Transform looks a little bit more compact. It's not the only reason for doing it, but it's not a bad reason for doing it. So with this notation, we're defining the Discrete complex exponential this way, we can rewrite the definition of the Discrete Fourier Transform. So with this, you can write the DFT as the Fourier Transform of F. The nth component is the sum from N equals zero up to N minus one of the nth component, a little F, the function you're arbitrating on, time omega to the minus N, the nth component of that. I haven't done anything different here. I've just rewritten it in terms of this Discrete vector exponential. Omega to the minus N of M is the nth entry of that quantity. Omega to the minus N, the nth component of this is E to the minus two pi I, MN over N. Okay?

My underlines there so we realize we're taking Discrete case here. Everything's Discrete. Or even more compactly, this is the Fourier Transform evaluated in M. The Fourier Transform of F sort of written just as an operator, is the sum from N equals zero up to N minus one of FN. The values of the signal you're inputting at N times omega to the minus N. All right? This expression is nothing but this expression evaluated on the index, M. Now, I find this, actually, a pretty convenient way of writing the DFT. It's always a question in this subject whether it's the Continuous case or the Discrete case. When you write your variables and when you want to avoid writing your variables. The same sort of thing applies in the Discrete case as well as the Continuous case. So this is sort of as far as you can go in writing the DFT without writing variables. In this case, the variable that we're not writing is the nth index. Again, this expression is nothing but this expression evaluated at the index M, or evaluated at the component, finding the nth component.

Okay. Now, we're going to work with this expression a fair amount, or the expression evaluated and try to make the decision – so this is probably about as close as you can make the Discrete Fourier Transform look like the Continuous Fourier Transform. The integral's replaced by the sum. The Continuous exponential, either the minus two pi I ST in the Continuous case is replaced by the Discrete exponential. That's about as close as you can get. It's pretty close, actually. For a lot of practical purposes, for a lot of computations, for a lot of formulas, it's pretty close. We're going to make a lot of use of it. Now, I told you last time that there are lots of little things that you have to digest, most of which are analogous to properties in the Continuous case. So I'm afraid that what I have to do now is just really go through almost a list of different, little points that come up, and I can't do all of them. So again, as I pleaded last time, I want you to read through the notes and sort of hit those points, some of which we'll talk about, some of which we won't talk about. I decided, actually, to reorder things slightly from the notes, just for a little variety here. So I'm going to derive many of the – I'm going to derive all the same formulas. Nothing's going to be different, of course, but I want to do it from a slightly different tack. I mean, I like the way I wrote it all out in the notes, but it does take a little more time, and I wanted to, here in class, try and hit the high points a little bit more quickly.

So again, there's no way of getting around this. We just have to have a certain list of properties that come out as consequences of that formula that we need in order to be able to use the DFT day to day. So the very first one, mini-properties, little properties. So the first one, actually, is something I'll mention now, but I'll talk about it next time. It's something that's different between the Discrete case and the Continuous case. That is the periodicity of the inputs and outputs. So I'll do this next time. I'll talk about this next time, but I wanted to mention it here because it really is probably the first thing you should establish about the DFT. Here, despite all my big build-up about how similar they are, this is a difference, actually, between Continuous and Discrete cases. Now, what I mean by this, and I'll say this in more detail next time, is that initially, you have a Discrete signal you feed in, and you get a Discrete signal out. You have a signal defined or indexed from zero to N minus one. The signal that comes out of that is also indexed from zero to N minus one. But as it turns out, you are really compelled by the definition of the DFT to extend those signals to be periodic of period capital N. That is to say to be defined on all the integers.

So I'll leave that up there because it really depends on the formula. The definition of the DFT compels you to regard the input, little F, and the output, capital F, as not just defined on the integers from zero to N minus one but as periodic Discrete signals of period N. All right? That is so because the vector exponential itself turns out to be a Discrete periodic signal of period N. Again, I'll do the details next time, but I just wanted to highlight this because this is the approach taken in the notes, and it [audio cuts in and out]. This is so because omega itself is a Discrete complex exponential is naturally a periodic Discrete signal of period N. Okay? More on this next time. This actually turns out to have significant consequences, the fact that you have to – I mean, whether or not you use it in a particular problem or a particular setting, it's lurking in the background. The input signals are periodic, and the output signals are periodic. In some cases, actually, it has consequences for computation. Some computations, if it's not taken into account, you can get results that are not in accord with what you might expect. It usually has to do with not taking into proper account the periodicity. Sometimes you have to shape this thing a little bit differently to take that into account and so on. So more on this next time, but I did want to mention it.

The main thing that I want to talk about today is the orthogonality of the Discrete complex exponentials and its consequences. The other sort of little fact – not so little fact, actually, is the orthogonality of the Discrete complex exponentials. Most non-trivial or slightly less than trivial – most interesting properties of the DFT can be traced back to what I'm about to talk about now. I'll put that in quotes. The properties of the DFT involve this property. What is the property? Okay. Let me get you set up, and I'll tell you what I mean. This couldn't be formulated in as easy a way if I didn't introduce this vector complex exponential. If I stayed writing the exponential terms themselves, either the two pi I, blah, blah, blah, then it would be a little more difficult, a little more awkward to formulate the property. But if you introduce the Discrete complex exponential, there's a very nice property that they have, that turns out to be very important for a number of reasons. So again, omega is one E to the two pi I over N, up to E to the two pi I, N minus one over N. The powers that we look at are omega to the K, say, is one E to the two pi I, K over N, E to the two pi I, two K over N, up to E to the two pi I, K times N minus one over N.

The orthogonality of the Discrete complex exponentials is really the orthogonality of the powers of these complex exponentials. That is to say if K is different from L, then omega to the K and omega to the L are orthogonal. Now here, when I say they're orthogonal, I mean not so much thinking of them as Discrete signals, but thinking of them as N vectors. All right. Now, I want to show you why that works and what happens when K is equal to L because that's where much heartache comes. All right. So let's compute omega to the K. So to say the [inaudible], they have to compute their inner product. We haven't looked at inner products much since the beginning of the course, when we talked about Fourier series, but they're coming back now. I'll remind you of the definition in this particular case. So I'm taking omega to the K dot omega to the L. It's a complex inner product. It's the inner product of these two vectors. So what is that? That's the sum from N equals zero to N minus one. Let me write it out in terms of – well, I'll just write it, and then I'll say more about it in a second. It's the nth component of omega to the K times the nth component of omega to the L conjugate. If I use this notation to indicate the nth component, then the inner product of two vectors is the sum of the products of the component. But with a complex inner product, you take the complex conjugate of the second term. So what is this? This is equal to the sum from N is equal to zero to N minus one.

Let me write it in terms of the complex exponentials now. E to the two pi I KN over N times E to the two pi I LN over N, conjugate, which is – it just puts a minus sign in the second complex exponential. Keep me honest on my algebra here so I don't make any slips. So that's the sum from N is equal to zero to N minus one of E to the two pi I KN over N times E to the minus two pi I of LN over N. You have to realize that what you have here is a geometric series. That is – I'm going to write this a little bit differently. The sum from N is equal to zero to N minus one of E to the two pi I, K minus L raised to the N – oh, divided by N, sorry. K minus L divided by N, all raised to the N. Exponentials being what they are, this is E to the two pi I, KN over N, minus E to the two pi I, LN over N. I group those terms, two pi I, K minus L divided by N, raised to the nth power. That's a geometric series, finite geometric series.

We know what its sum is. We know how to sum that. So if K is different from L, if K is not equal to L, then this sum is equal to one minus E to the two pi I, K minus L, divided by N, raised to the nth power. Divided by one minus E to the two pi I, K minus L divided by N. Okay? Nothing up my sleeve, make sure I did this write. I believe so. Make sure I wrote it neatly enough. So it's one plus R plus R-squared plus R-cubed up to the R to the N minus one, where R is this exponential. So it's one minus the thing to the nth power. I've [inaudible]. Zero to N minus one is N terms. Raise it to the nth power, divided by the thing that's getting raised to the power, one minus the thing that's getting raised to the power. It's no problem with the denominator here because K is different from L, and K minus L is always less than N, so these are distinct. This is never equal to one down below. But on top, if I raise it to the nth power, it is one minus E to the two pi I, K minus L, divided by one minus E to the two pi I, K minus L, divided by N. K minus L is an integer. E to the two pi I, K minus L, is one. So this baby is zero.

So if K is different from L, omega to the K dot omega to the L, the inner product of the Discrete complex exponentials is equal to zero. What happens if K is equal to L? If K is equal to L, go back to the sum here. If K is equal to L, then I'm getting E to the zero, so I'm just getting a sum of a bunch of ones. How many ones do I have? N of them. So if K is equal to L, omega to the K dot omega to the L is equal to N. To summarize, omega to the K dot omega to the L is equal to zero, if K is different from L. It's equal to N if K is equal to L. In the notes, this is done a little bit more generally because I allow for periodicity. Instead of saying K is equal to L or K is different from L, you say K is different from L, modulo, capital N. Here you say K is equal to L, modulo, capital N. That falls out from the general discussion, but we'll get to that next time. This is the first sort of encounter with it, and this is probably the easiest way of thinking about it. It's an extremely important property. The orthogonality of the vector exponentials. Now, it has been a factor traditionally in the electrical engineering department – I say this. Listen very carefully. It's a quals tip. My dear colleague, Bob Gray, used to always ask quals questions that somehow reduced to the orthogonality of the Discrete complex exponentials. Somehow, that was always involved in his quals question.

So now you know. I don't know if he still does this since I've been making a big deal out of this the last couple of years, but somehow, he would always manage to ask a question that reduced to this or somehow involved this in a crucial way. It's very important. The thing that makes – here's another difference between the Continuous case and the Discrete case that shows up in a lot of formulas, and it leads to much heartache and much grief. The fact that they're orthogonal but not orthonormal. The length of the vectors is not one. It's the square root of N. What this says is, to put this another way, the length, the norm of this vector exponential, Discrete complex exponential, is the inner prerogative of the vector itself. So omega K squared is omega to the K dot omega to the K, and that's equal to N. So the length squared is N. The length is equal to the square root of N.

This fact causes an extra factor of N or one of N to appear in many formulas involving the DFT. It can all be traced back to this. It is a royal pain in the ass. I'm sorry to have to report that to you, but it is. It always traces back to this. It causes a factor of one over N to appear in many formulas. The way we define the DFT, it does. There are ways of sort of getting around it, but it's awkward. Once again, like in most things, there's no particular consensus as to what is best. But the way we're defining the DFT, you have this sad fact that the vector complex exponentials are orthogonal but not quite orthonormal. They cause that extra factor to come in. Now, let me give you the first important consequence of this. The first important consequence of the orthogonality of these Discrete complex exponentials is the inverse DFT. So a consequence of the orthogonality is a simple formula for the inverse DFT.

Now again, in the notes, I did this a little bit differently. I mean, I wound up with the same formula, but I did it differently. I sort of liked it because it was a way of discovering what the formula should be independent of the orthogonality of Discrete complex exponentials and so on. But again, just to do things differently so you have two different points of view here, let me just give you the formula and show why it works. I don't like doing that because I don't like this [inaudible] aspect of it where you write down this formula and say, son of a bitch, it works. I instinctively don't like that, but I'm going to do it anyway. I'm going to do it. I want to show you the consequence of this orthogonality relationship. It's nice. Here's what we find. The inverse DFT is given by the inverse – I'll just write down the formula in its full glory. I'll put an underline under it. So apply to some signal F is the nth component of that. It's one over N times the sum from N equals zero to NM minus one of little F and N. This time, you take positive powers. E to the plus two pi I NM over N.

Or, written a little bit more compactly, in terms of vector exponentials, the inverse Fourier Transform of F is one over N at M. It's one over N times the sum from N is equal to zero to N minus one of FN, omega to the N, its nth component. Same thing. Omega to the positive power of N. Or written more compactly still, if I drop the variable, inverse Fourier Transform of F is the sum from N is equal to zero to capital N minus one of FN times omega to the N. I forgot the painful factor, one over N. So this is as close as you can get for the inverse Discrete Fourier Transform to look like the inverse Continuous Fourier Transform and the difference, among other things, is this irritating factor of one over N. It comes in there exactly because the vector complex exponentials are orthogonal but not quite orthonormal. But now, I'll show you why this is the inverse. Again, if you look at the discussion in the notes, this formula emerges from an analogue of the Continuous case. You talk about reverse signals and duality. All those things, we're going to talk about. Maybe not talk about all of them, but I want you to read all of them because it's another series of little points that you have to digest. If you pursue the analogy with the Continuous case, via duality and all the rest of that stuff, this formula emerges quite nicely from that, independent of first talking about the orthogonality of the vector complex exponentials.

But since we're done it this way, let me show you how they come in. So I have to show you, if I have time here, that the inverse Fourier Transform or the Fourier Transform of F is equal to F. Or I have to show to the inverse Fourier Transform or the Fourier Transform – let me put a variable in here – at the nth component is the nth component. That's what it means for the [inaudible] inverse Fourier Transform. Let's do that. We have no recourse here other than to appeal to the formula. I want you to see how this orthogonality comes in. I have to be careful of my indices here. So the inverse Fourier Transform of a Fourier Transform of F at M is one over N times the sum from N is equal to zero up to N minus one, up to the coefficients of the things I'm putting into it, which is the Fourier Transform, times omega to – maybe I'll write it in terms of complex exponentials here. Times E to the two pi I NM over N, right? Okay. Is it correct? The inverse Fourier Transform looks like that. Okay. Now, I can begin the formula for the Fourier Transform. I'm just sort of debating in my own mind, which you can't hear, whether or not I should write it in terms of the complex exponential or write it in terms of the omega. I'll keep this up. It might be a little cleaner if I use the other notation form, but never mind. Let's go forward.

Now I have to write this sum from N equals zero up to capital N minus one. So I have to bring in the formula for the Fourier Transform. That's the sum of K equals zero up to N minus one of F of K. The nth component of the Fourier Transform is the sum from K equals zero minus one, F of KE to the minus two pi I, KN over M. Then times E to the two pi I, NM over M. Check me. Make sure I haven't slipped anything past you here. Did I remember my factor one over N in front? I did. All right. So once again, all I did here was substitute the formula for the Fourier Transform. The nth component of the Discrete Fourier Transform of F is this sum, right? The power here, N, means I'm taking the nth component of the Fourier Transform of little F. Good.

Now, you can combine everything here. What I'm going to do, guess. I'm going to swap the order of summation. I'm going to put everything together, and then I'm going to swap the order of summation. That's the sum from – it's one over N times the sum from N equals zero to K equals zero. One over N, sum from N equals zero to N minus one. Sum from K equals zero to N minus one, F of K, times E to the two pi I, N – K over N times E to the two pi I. NM over N. And I see I got this wrong. This is a minus here. Now, swap the order of summation. It's like swapping the order of integration, a technique we've used many times. So this is one over N, sum from K equals zero of the sum from LN – the sum from N equals zero. Now this depends only on – F of K does not depend on N, so I'm going to bring that out of the sum. F of K times the sun from N equals zero up to N minus one. What remains are these products of complex exponentials. E to the minus two pi I, NK over N, times E to the two pi I, NM over N. Okay? Wonderful. You have to have a certain taste for this. Now, what you should recognize in that sum is the inner product of – this is the powers of the vector complex exponential times its conjugate. E to the plus two pi I, NM over N. E to the minus two pi I, NK over N. It is the nth power inner product with the Kth power. That's what this sum is, summed over N. All right? The Kth power of omega inner product with the nth power of omega. That is either equal to zero if M is different from K, or N if M is equal to K. So that sum, N equals zero up to N minus one of E to the minus two pi I, NK over N, times E to the plus two pi I, NM over N, is either equal to zero, if K is different to M. And it's equal to N if K is equal to M. Okay?

All right. So only one term survives here, and then it's the outside sum here. So the only term that survives is when K is equal to M, in which case you get N times one over N, times – and if the only term that survives is when K is equal to M, you get F of M. And we're done. The [inaudible] surviving term is K equals M, and you get the inverse Fourier Transform of the Fourier Transform of F at M is equal to F of M. That is to say the inverse Fourier Transform given by this formula really is the inverse of the Fourier Transform. Works like a charm, and it depends crucially on the orthogonality of the complex exponentials. Crucially on that, like many other properties. The only thing that makes it a little bit painful is this extra factor of N that comes in, but that's just life. It's just the way it goes. Okay. That's it for today. Next time, we'll do more interesting facts about the DFT.

[End of Audio]

Duration: 53 minutes