TheFourierTransformAndItsApplications-Lecture20

Instructor (Brad Osgood):There we go. All right, the exams – 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 website, so I will do that after I get back from class today. And you can pick up your exams from Denise Murphy, my admin in Packard 267. She has all of them. I'll also send that announcement out to the class. Any questions or comments on that? Good. Today we're gonna continue with our discussion of the DFT. This is "Getting to know your DFT, your discrete Fourier transformation." Now the subtitle I would say actually 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 you built up for the continuous Fourier transform, the formulas that you've learned to work with and so on, really all have analogs 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 gonna take the point of view that – and I'm gonna try to make the – 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 again it allows us to leverage off of what we have done in the continuous case. Now before doing – so let me recall the definition. Here's the definition, and where we ended up last time was sort of a definition that where the continuous case, the idea of sampling a continuous signal to get a discrete signal, sampling its Fourier transform to get a discrete Fourier transform, all that had pretty much vanished by the final definition. So the definition that we wound up with – ultimately wound up with – ultimately had makes the continuous case almost invisible, vanish. It looked like this. You have a discrete signal. So I'm using sort of 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 – you can think of this an n-tuple of numbers, or a discrete signal whose value at the nth point is just the nth component here. Okay? Oops, that doesn't look very – that's not much of a statement. And since you can either consider it as a discrete signal that's defined on the integers, where the integers from zero to N minus one, or you can think of it as an n-tuple or a vector. If you have a discrete signal and its Fourier transform is another discrete signal, its DFT is the discrete signal – I'll call it capital F, but I'll also use the notation to make the connection with the 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.

And it's defined by its nth component, so the nth component of the Fourier transform is the sum from say N equals zero to N minus one of the Nth component of F times E to the minus two pi I N M over N. All right? 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 points, F of zero, F of one, F of two and so on. That's the definition. They say 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. And that’s pretty much how we’re gonna deal with it. But before embarking on that path forever, I wanna have one nod back to the continuous case, and talk a little bit about how the DFT is employed and the kind of things you have to know when you’re gonna use it in practice because you already have to some extent, and you certainly will more in the future. So one look back at the continuous case to talk about one additional – one sort of phenomena 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, continuous roots of the DFT, and reciprocity of the two domains, time and frequency.

So how did it work? We had a signal in the time domain that we discretized, and we had a signal in the frequency domain that we discretized. So imagine that you have two grids. You have a grid in the time domain. You have N sample points say that are delta T apart, N points. I’m running out of room to write things, so let me just write it here and then write some of the notation up above. And this is the frequency domain, and we have spacing in the frequency domain say delta S apart. So we have N sample points T zero up to T N minus one spaced delta T apart, and N sample points in time that is, and we have N sample points in frequency, S not up to S N 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, and there’s a relationship between them, so you have three quantities of interest, delta T, delta S, and N. And they’re not independent. That’s the important point. Let me remind you of again – I’m not gonna go through the derivation again, but let me remind you what the relationship is. We had at N delta T – that’s the spacing in the time domain that has to be [inaudible] – just let me sure I get this right. I want to say this right – is L, that’s the spacing in the time domain. I want to use the notation I used last night. This is a sort of a time limitedness. And we had N times delta S is the band limitedness that we called that two B. That was the band limitedness, 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 B L over N squared, but 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 B L 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 gonna take the same number of sample points in both domains, but that’s what we found. There’s a relationship there. So here let me highlight it. That is delta S delta T or delta T delta S is one over N. 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 the reciprocity relationship. This again has practical significance and practical consequences when you’re applying the DFT. You can imagine you have a continuous signal. You wanna 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.

And if you do that, once you do that then delta S is fixed, and delta S is determined. The spacing in the frequency domain is determined. That is to say maybe one way of putting this is the accuracy of the resolution or how fine the resolution is in frequency is fixed by the choices you make in the time domain. So let's say – I don't know how you would say the resolution in frequency is determined – is fixed by the choices you make in time. Conversely, you could say I want a certain accuracy in – or I want a certain resolution in frequency. I want delta S to have a certain fineness. That's gonna determine then how many sample points I want to take and how I space them in the time domain. Or at least you have to choose two of those things and the third one is determined. And 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? Ain't that the way things go? The freedom you have typically is how many measurements you make and how frequently you make them. And once you do that, then that determines what the resolution is like in the frequency domain. There's ways of getting around this. There's zero padding and special – not getting around it, but ways of understanding it or massaging it. So-called zero padding, you have a problem on that. You'll have other chances to experiment with that and other sorts of things, but 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, the sort of stretched in one domain means shrunk in the other domain, or reciprocity between time and frequency. I say we've seen many different instances of that in the continuous case, and here's an example of how that carries over to the discrete case. So 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 often in most common contexts where it's associated somehow with some sort of continuous process that you're sampling. You wanna take the Fourier transform. You don't have the formula for it. You have to use numerical algorithms to actually compute it, and that means your sampling and these – the reciprocity relationship puts certain restrictions or limitations on what you should expect to get. So that is my final nod at least for now back to the continuous case. So let's go back now to the discrete world and pretty much stay there. So back to this formula and its consequences. Let me erase it and write it down again. We're back to the discrete world, discrete setting.

Again, my goal here is to make the discrete Fourier transform look as much possible like the continuous Fourier transform. So here's how to do that. I wanna make the DFT and associated formulas look like the continuous case. Maybe I should say we're not really abandoning the continuous case, but we're doing it – we're situating ourselves firmly in the discrete side of things. Now to help out in that, partly it's a matter of notation. 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 wanna do is introduce a symbol, introduce a way of thinking about the exponentials that occur in the definition of the DFT. So let me write down the formula one more time. So again F is a discrete signal, F zero up to F N minus one. And its Fourier transform is another discrete signal whose Nth component is given by this formula: sum from N equals zero N minus one, the Nth component of F, little F times E to the minus two pi I N M over N. Okay, that's a two pi I. The first thing I wanna do is introduce the notation. I wanna 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 discrete Fourier transform that would be very difficult to do otherwise. So I want you 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 the discrete complex exponential or the vector complex exponential. You pick what you wanna call it.

So let me give you the definition. I'm gonna write – I'll leave that board up there, so let me write it over here. I'm gonna 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, so let me define it in terms of a plus sign. Then I'll start taking powers to get a minus sign. So E to the two pi I over N. We'll see where this is coming from in just a second. E to the two pi I times two over N, and so on and so on, all the way to final term. The Nth term is going to be E to the two pi I N minus one over N. That's the basic vector or discrete complex exponentials. So its Nth component – so omega, the Nth component of this is E to the two pi I M over N. Okay? It doesn't take a great leap in the imagination to do this. Now I wanna 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 after the minus sign, but it views it slightly differently. Its self is a discrete signal, so you can either view it as an N vector, or you can view it as a discrete signal defined on the integers from zero to N minus one. I wanna take powers. 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 MATLAB it does, I suppose. And it certainly makes sense to take powers of a discrete function, so omega to the – I wanna 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 entries of 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 N minus one over N. And of course, if I can take positive powers, I can also take negative powers. N is not meant to be either – we're sticking to the 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 to 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. The way this notation – we're defining the vector or 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 of little F, the function you're operating on, times omega to the minus N, the Nth component of that. I haven't done anything different here. I've just rewritten 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 M N over N. Okay? Or – let me put 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 at 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 F N, the values of the signal you're inputting at N times omega to the minus N. This expression is nothing but this expression evaluated on the index M. Now I actually find this actually a pretty convenient way of writing the DFT. It's a question about when you write your – it's always a question in this subject, whether it's the continuous case of the discrete case, when you write your variables and when you wanna avoid writing your variables. And 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 Mth index. Again, this expression is nothing but this expression evaluated at the index M, or evaluated at the component – finding the Mth component.

So we're gonna work with this expression a fair amount or expression evaluate it, and try to make the – 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 E to the two pi I – minus two pi I S T in the continuous case is replaced by the discrete exponential. That's about as close as you can get, and it's pretty close actually. For a lot of practical purposes, for a lot of formulas, for a lot of computations, it's pretty close. We're gonna make a lot of use of it. Now I told you last time that there are just lots of little things that you have to kind of digest, most of which are analogous to properties in the continuous case. So I'm afraid that really what I have to do now is go through almost a list of different points, little points that come up, and I can't do all of them. So again as I pleaded last time that 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, and I decided actually to give a slightly – to reorder things slightly from the notes just for a little variety here. So I wanna derive many of the – I'll derive all the same formulas. Nothing's gonna be different, of course, but I wanna do it from a slightly different tack. I mean I like the way I actually wrote it all out in the notes, but it does take a little more time, and I wanted to here in class try to sort of hit the high points a little bit more quickly. So again we have – 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 sort of use the DFT day to day. So the very first one – so list, many properties, little properties.

The first one actually is something – I'll mention it now, but I'll talk about it next time – is something that's different between the discrete case and the periodic and the continuous case, and 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 sort of the first thing you should establish about the DFT. And 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 again 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 from zero – 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 period of period capital N, that is to say to be defined on all the integers. So we'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, its discrete Fourier transform as not just defined on the integers from zero to N minus one but as periodic discrete signals of period N. That is so because the vector exponential itself turns out to be naturally a periodic signal of period N, discrete periodic signal of period N. But again, I'll do the details next time, but I wanted to highlight this because this is how – this is sort of the approach taken in the notes, and it really [inaudible] high points [inaudible] this is so [inaudible] omega itself, its discrete complex exponential should be – is naturally a periodic discrete signal of period N. And more on this next time.

This actually turns out to have significant consequences. The fact that you really have to – you know, it's all – whether or not you use it in a particular problem or a particular setting, it's lurking in the background. If the input signals are periodic and the output signals are periodic. And 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, and it usually has to do with not taking into proper account the periodicity. Sometimes you have to shape a signal a little bit differently to take that into account and so on, so more on this next time, but I did wanna mention it. The main thing that I wanna talk about today is the orthogonality of the complex – of the vector – of the discrete complex exponentials and its consequences. So two, the other sort of little fact – not so little fact, actually – is the orthogonality of the discrete complex exponentials. Most sort of 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. Properties of the DFT involve this property. Now what is the property? So let me give you a set up, and I'll tell you what I mean. This would be harder to – or not as – couldn't be formulated in as easy a way if I didn't introduce this vector complex exponential. If I stayed writing just 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 way, 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 over N – N minus one, sorry, over N. And 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'll write one more term out here – 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. Now I wanna show you why that works, and also what happens when K is equal to L because that's where much heartache comes.

So let's compute omega to the K – to say they're orthogonal means I have to compute their inner product. All right, we haven't looked at inner products much since the very beginning of the course when we talk about Fourier series, but they're coming back now, and 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 of these two vectors, so what is that? That's the sum from say N equals zero to N minus one. Let me just write it, 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 the two vectors is the sum of the products of the components, but in the 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 through with the complex exponentials now, the usual complex exponentials. E to the two pi I K N over N times E to the two pi I L N over N conjugate which just puts a minus sign in the 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 K N over N times E to the minus two pi I L N over N. Right? And you have to realize that what you have here is a geometric series. That is this is – let me write this a little bit differently. This is 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 either the two pi I K N over N, either the two pi I – minus E to the two pi I L N 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, and we know what its sum is. Okay? 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 – let me write it like this – divided by one minus E to the two pi I K minus L divided by N. Nothing up my sleeve. Make sure I did this right. 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 minus one where R is this exponential. So it's one minus the thing to the Nth power at N terms – zero to N minus one is N terms – raised 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. There's no problem with the denominator here because K is different form 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. [Inaudible] 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, is 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, and 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. It doesn't really – that follows out from the general discussion. We'll get to that next time, but this is the sort of encounter with it. This is probably the easiest way of thinking about it, and it's an extremely important property, the orthogonality of the vector exponentials. Now it's a fact or it has been a fact – traditionally in the electrical engineering department – I say this. Listen very carefully. It's a Quals tip. Bob Gray, 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. All right. So now you know. I don't know if he still does it since I've been making a big deal out of this the last couple 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. Now the thing that makes – now 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, all right? The length of the vectors is not N, it's the – 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 product of the vector with 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 – and the length is equal to the square root of N.

This fact causes an extra factor of N or one over 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, and it always traces back to this. Causes a factor of N or one over N to appear in many formulas. Sorry. The way we define the DFT it does. There are ways of sort of getting around it, but it's sort of awkward, and once again like in most things, there's no particular consensus as to what is best here, but the way we're defining the DFT which is pretty standard, 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 a little bit differently, and I sort of liked it again because it was a way of discovering what the formula should be independent actually of the discrete – although that comes into it – of the orthogonality of discrete complex exponentials and so on. But again just to do things a little bit 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 the sort of deus ex machina aspect of it where you write down this formula and say, "Son of a bitch, it works." I sort of instinctively don't like that, but I'm gonna do it anyway. Son of a bitch, I'm gonna do it. And that is – I don't wanna – because I wanna show you how it's a consequence of this orthogonality relationship. It's quite nice. Here's what we find. The inverse DFT is given by the inverse Fourier – I'll just write down the formula. I'll write it over here in its full glory. Again, I'll put an underline under it, so again apply to some signal F is the Mth component of that is one over N times the sum from N equals zero to capital N minus one of little F and N, and this time you take positive powers, E to the plus two pi I N M 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 is one over N times the sum from N is equal to zero to N minus one of F N – omega to the N, its Mth component, same thing, omega to the positive power 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 F N times omega to the Nth. Oh, 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, and it comes in there exactly because the vector complex exponentials are orthogonal but not quite orthonormal. Okay? But now I'll show you why it works. I'll show you why this is the inverse. Again, if you look at the discussion in the notes, this formula emerges from an analog of the continuous case. We talk about reverse signals, and duality, and all those things we're gonna talk about – well, 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 sort of digest. And if you do that, if you pursue the analogy actually with the continuous case via duality and all the rest of that stuff, actually this formula emerges quite nicely from that, independent of first talking about the orthogonality of the vector complex exponential. But since we've done it this way, let me show you how they come in. All right, so I have to show you if I have time here that the inverse Fourier transform of the Fourier transform of F is equal to F. Or I have to show you the inverse Fourier transform of the Fourier transform – let me put a variable in here – at the M, the Mth component is the Mth component. That's what it means for the thing to be the inverse Fourier transform. Okay. So let's do that. We have no recourse here other than to appeal to the formula, and I want you to see how this orthogonality comes in. I have to be careful with my indices here. So the inverse Fourier transform of the Fourier transform of F is one over N times [inaudible] at M – is one over N times the sum from N is equal to zero up to N minus one of the coefficients of the thing I'm putting into it which is the Fourier transform times omega two – or maybe I'll write in terms of complex exponentials here – times E to the two pi I N M over N. Right? Okay. Right? Is that correct? That's what the Fourier – the inverse Fourier transform looks like that?

Student:[Inaudible].

Instructor (Brad Osgood):Right. Okay. Now I can begin the formula for the Fourier transform. I'm just sort of debating in my own mind here, 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 omegas. I started with this, so I'll keep this us. It may be a little cleaner if I use the other notation for it, but never mind. Let's go forward. So now I have to write this sum from N equals zero up to capital N minus one, so I need to bring in the formula for the Fourier transform. That's the sum say 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 to N minus one F of K E to the minus two pi I K N over M, capital N, and then times E to the two pi I N M over capital N. Check me. Make sure I haven't slipped anything past you here. Did I remember my factor one over N in front? I did. 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. The power here N means I'm taking the Nth component of the Fourier transform of the left. Good. And now you can combine everything here. What I'm gonna do – guess. I'm gonna swap the order of summation. Well, I'm gonna put everything together and then swap the order of summation. So that's the sum from – it's one over N times from the N equals zero – K equals zero one over N – sum from N equals zero to N minus one, sum from K equals zero to N minus one, and F of K times E to the two pi I N, K over N times E to the two pi I N M over N, and I see I got this wrong. This is a minus here, right? Minus, okay.

Okay, now swap the order of summations, like swapping the order of integration, a technique we used many times. So this is one over N sum from K equals zero of the sum from L and the sum from N equals zero. Now this depends only on F of K. It does not depend on N, so I'm gonna bring that out of the sum, F of K times the sum from N equals zero up to N minus one, and what remains are these products of complex exponentials, E to the minus two pi I N K over N times E to the two pi I N M over N. 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 N M over N, E to the minus two pi I N K over N, it is the Mth vector complex exponential – the Mth power inner product with the Kth power. That's what this sum is summed over N, the Kth power of omega inner product with the Mth power of omega. And that is either equal to zero if M is different from N – 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 N K over N times E to the plus two pi I N M over N is either equal to zero if K is equal to M, and it's equal to N if – if K is different from M, and equal to N if K is equal to M. 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 the one over N, and if the only term that survives is when K is equal to M, you get F of M. And we're done about. So the only 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, and 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 exponential, crucially on that, like many of the properties. And the only thing that makes it a little bit painful is this extra factor of N that comes in but that's just life. That's just the way it goes. That's it for today. Next time, we'll do more little interesting facts about the DFT.

[End of Audio]

Duration: 53 minutes