

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

[MUSIC PLAYING]

PROFESSOR: In the last lecture, we discussed a number of general properties for systems, which, as you recall, applied both to continuous-time and to discrete-time systems. These properties were the properties associated with a system having memory. The issue of whether a system is or isn't invertible, we talked about causality and stability, and finally we talked about when linearity and time invariance. In today's lecture, what I'd like to do is focus specifically on linearity and time invariance, and show how for systems that have those properties, we can exploit them to generate a general representation.

Let me begin by just reviewing the two properties again quickly. Time invariance, as you recall, is a property that applied both to continuous-time and discrete-time systems, and in essence stated that for any given input and output relationship if we simply shift the input, then the output shifts by the same amount. And of course, exactly the same kind of statement applied in discrete time. So time invariance was a property that said that the system didn't care about what the time origin of the signal is.

Linearity was a property related to the fact that if we have a set of outputs associated with a given set of inputs-- as I've indicated here with the inputs as ϕ_k and the outputs ψ_k , then the property of linearity states that if we have an input which is a linear combination of those inputs, then the output is a linear combination of the associated outputs. So that linear combination of inputs generates, for a linear system, an output which is a linear combination of the associated outputs.

Now the question is, how can we exploit the properties of linearity and time

invariance? There's a basic strategy which will flow more or less through most of this course. The strategy is to attempt to decompose a signal, either continuous-time or discrete-time, into a set of basic signals. I've indicated that here.

And the question then is, what basic signal should we pick? Well, the answer, kind of, is we should pick a set of basic signals that provide a certain degree of analytical convenience. So we choose a set of inputs for the decomposition that provide outputs that we can easily generate.

Now, as we'll see, when we do this, there are two classes of inputs that are particularly suited to that strategy. One class is the set of delayed impulses, namely decomposing a signal into a linear combination of these. And as we'll see, that leads to a representation for linear time-invariant systems, which is referred to as convolution.

The second is a decomposition of inputs into complex exponentials-- a linear combination of complex exponentials-- and that leads to a representation of signals and systems through what we'll refer to as Fourier analysis. Now, Fourier analysis will be a topic for a set of later lectures. What I'd like to begin with is the representation in terms of impulses and the associated description of linear time-invariant systems using convolution.

So let's begin with a discussion of discrete-time signals, and in particular the issue of how discrete-time signals can be decomposed as a linear combination of delayed impulses. Well, in fact, it's relatively straightforward. What I've shown here is a general sequence with values which I've indicated at the top.

And more or less as we did when we talked about representing a unit step in terms of impulses, we can think of this general sequence as a sequence of impulses-- delayed, namely, occurring at the appropriate time instant, and with the appropriate amplitude. So we can think of this general sequence as an impulse occurring at $n = 0$ and with a height of $x[0]$, plus an impulse of height $x[1]$ occurring at time $n = 1$, and so that's $x[1] \delta[n-1]$, an impulse at -1 with an amplitude of $x[-1]$, etc.

So if we continued to generate a set of weighted, delayed unit samples like that, and if we added all these together, then that will generate the total sequence.

Algebraically, then, what that corresponds to is representing the sequence as a sum of individual terms as I've indicated here or in terms of a general sum, the sum of $x[k] \delta[n-k]$.

So that's our strategy-- the strategy is to decompose an arbitrary sequence into a linear combination of weighted, delayed impulses. And here again is the representation, which we just finished generating.

Now, why is this representation useful? It's useful because we now have a decomposition of the sequence as a linear combination of basic sequences, namely the delayed impulses. And if we are talking about a linear system, the response to that linear combination is a linear combination of the responses. So if we denote the response to a delayed impulse as $h_k[n]$, then the response to this general input is what I've indicated here, where $y[n]$, of course, is the output due to the general input $x[n]$. $h_k[n]$ is the output due to the delayed impulse, and these are simply the coefficients in the weighting.

So for a linear system, we have this representation. And if now, in addition, the system is time-invariant, we can, in fact, relate the outputs due to these individual delayed impulses. Specifically, if the system is time-invariant, then the response to an impulse at time k is exactly the same as the response to an impulse at time 0 , shifted over to time k . Said another way, $h_k[n]$ is simply $h_0[n-k]$, where h_0 is the response of the system to an impulse at $n = 0$. And it's generally useful to, rather than carrying around $h_0[n]$, just simply define $h_0[n]$ as $h[n]$, which is the unit sample or unit impulse response of the system.

And so the consequence, then, is for a linear time-invariant system, the output can be expressed as this sum where $h[n-k]$ is the response to an impulse occurring at time $n = k$. And this is referred to as the convolution sum.

Now, we can-- just to emphasize how we've gone about this, let me show it from another perspective. We of course have taken the sequence $x[n]$, we have

decomposed it as a linear combination of these weighted, delayed impulses. When these are added together, those correspond to the original sequence $x[n]$. If this impulse, for example, generates a response which is $x[0] h[n]$, where $h[n]$ is the response to a unit impulse at $n = 0$, and the second one generates a delayed weighted response, and the third one similarly, and we generate these individual responses, these are all added together, and it's that linear combination that forms the final output.

So that's really kind of the way we're thinking about it. We have a general sequence, we're thinking of each individual sample individually, each one of those pops the system, and because of linearity, the response is the sum of those individual responses.

That's what happens in discrete time, and pretty much the same strategy works in continuous time. In particular, we can begin in continuous time with the notion of decomposing a continuous-time signal into a succession of arbitrarily narrow rectangles. And as the width of the rectangles goes to 0, the approximation gets better. Essentially what's going to happen is that each of those individual rectangles, as they get narrower and narrower, correspond more and more to an impulse.

Let me show you what I mean. Here we have a continuous-time signal, and I've approximated it by a staircase. So in essence I can think of this as individual rectangles of heights associated with the height of the continuous curve, and so I've indicated that down below. Here, for example, is the impulse corresponding to the rectangle between $t = -2\Delta$ and $t = -\Delta$. Here's the one from $-\Delta$ to 0, and as we continue on down, we get impulses, or rather rectangles, from successive parts of the wave form.

Now let's look specifically at the rectangle, for example, starting at 0 and ending at Δ , and the amplitude of it is $x(\Delta)$. So what we have-- actually, this should be $x(0)$, and so let me just correct that here. That's $x(0)$. And so we have a rectangle height $x(0)$, and recall the function that I defined last time as $\delta_{\Delta}(t)$, which had height $1/\Delta$ and width Δ . So multiplying finally by this last little Δ ,

then, this is a representation for the rectangle that I've shown there.

Now there's a little bit of algebra there to kind of track through, but what we're really doing is just simply representing this in terms of rectangles. What I want to do is describe each rectangle as in terms of that function $\delta_{\Delta}(t)$, which in the limit, then, becomes an impulse.

So let's track that through a little further. When we have that linear combination, then, we're saying that $x(t)$ can be represented by a sum as I indicate here, which I can then write more generally in this form, just indicating that this is an infinite sum. We now want to take the limit as Δ goes to 0, and as Δ goes to 0, notice that this term becomes arbitrarily narrow. This goes to our impulse function, and this, of course, goes to x of τ . And in fact, in the limit, a sum of this form is exactly the way an integral is defined. So we have an expression for $y(t)$ in terms of an impulse function. There, I have to admit, is a little bit of detail to kind of focus on at your leisure, but this is the general flow of the strategy.

So what we have now is an integral that tells us that tells us how $x(t)$ can be described as a sum or linear combination involving impulses. This bottom equation, by the way, is often referred to as the sifting integral. In essence, what it says is that if I take a time function $x(t)$ and put it through that integral, the impulse as it zips by generates $x(t)$ all over again.

Now, at first glance, what it could look like is that we've taken a time function $x(t)$ and proceeded to represent it in a very complicated way, in terms of itself, and one could ask, why bother doing that? And the reason, going back to what our strategy was, is that what we want to do is exploit the property of linearity. So by describing a time function as a linear combination of weighted, delayed impulses, as in effect we've done through this summation that corresponds to a decomposition in terms of impulses, we can now exploit linearity, specifically recognizing that the output of a linear system is the sum of the responses to these individual inputs.

So with $h_{\Delta}(t)$ corresponding to the response to $\delta_{\Delta}(t-k\Delta)$ and the rest of this stuff, the $x(k\Delta)$ and this little Δ are basically scale factors-- for a

linear system, then, if the input is expressed in this form, the output is expressed in this form, and again taking the limit as Δ goes to 0, by definition, this corresponds to an integral. It's the integral that I indicate here with $h_{\tau}(t)$ corresponding to the impulse response due to an impulse occurring at time τ .

Now, again, we can do the same thing. In particular, if the system is time-invariant, then the response to each of these delayed impulses is simply a delayed version of the impulse response, and so we can relate these individual terms. And in particular, then, the response to an impulse occurring at time $t = \tau$ is simply the response to an impulse occurring at time 0 shifted over to the time origin τ . Again, as we did before, we'll drop this subscript h_0 , so $h_0(t)$ we'll simply define as $h(t)$. What we're left with when we do that is the description of a linear time-invariant system through this integral, which tells us how the output is related to the input and to the impulse response.

Again, let's just quickly look at this from another perspective as we did in discrete time. Recall that what we've done is to take the continuous function, decompose it in terms of rectangles, and then each of these rectangles generates its individual response, and then these individual responses are added together. And as we go through that process, of course, there's a process whereby we let the approximation go to the representation of a smooth curve.

Now, again, I stress if there is certainly a fair amount to kind of examine carefully there, but it's important to also reflect on what we've done, which is really pretty significant. What we've managed to accomplish is to exploit the properties of linearity and time invariance, so that the system could be represented in terms only of its response to an impulse at time 0. So for a linear time-invariant system-- quite amazingly, actually-- if you know its response to an impulse at $t = 0$ or $n = 0$, depending on discrete or continuous time, then in fact, through the convolution sum in discrete time or the convolution integral in continuous time, you can generate the response to an arbitrary input.

Let me just introduce a small amount of notation. Again, reminding you of the

convolution sum in the discrete-time case, which looks as I've indicated here-- the sum of $x[k] h[n-k]$ will have the requirement of making such frequent reference to convolution that it's convenient to notationally represent it as I have here with an asterisk. So $x[n] * h[n]$ means or denotes the convolution of $x[n]$ with $h[n]$.

And correspondingly in the continuous-time case, we have the convolution integral, which here is the sifting integral

as we talked about, representing $x(t)$ in terms of itself as a linear combination of delayed impulses. Here we have the convolution integral, and again we'll use the asterisk to denote convolution.

Now, there's a lot about convolution that we'll want to talk about. There are properties of convolution which tell us about properties of linear time-invariant systems. Also, it's important to focus on the mechanics of implementing a convolution-- in other words, understanding and generating some fluency and insight into what these particular sum and integral expressions mean.

So let's first look at discrete-time convolution and examine more specifically what in essence the sum tells us to do in terms of manipulating the sequences. So returning to the expression for the convolution sum, as I show here, the sum of $x[k] h[n-k]$ -- let's focus in on an example where we choose $x[n]$ as a unit step and $h[n]$ as a real exponential times the a unit step. So the sequence $x[n]$ is as I indicate here, and the sequence $h[n]$ is an exponential for positive time and 0 for negative time. So we have $x[n]$ and $h[n]$, but now let's look back at the equation and let me stress that what we want is not $x[n]$ and $h[n]$ -- we want $x[k]$, because we're going to sum over k , and not $h[k]$ but $h[n-k]$.

So we have then from $x[n]$, it's straightforward to generate $x[k]$. It's simply changing the index of summation. And what's $h[n-k]$? Well what's $h[-k]$? $h[-k]$ is $h[k]$ flipped over.

So if this is what $h[k]$ looks like, then this is what $h[n-k]$ looks like. In essence, what the operation of convolution or the mechanics of convolution tells us to do is to take

the sequence $h[n-k]$, which is $h[-k]$ positioned with its origin at $k = n$, and multiply this sequence by this sequence and sum the product from $k = -\infty$ to $+\infty$. So if we were to compute, for example, the output at $n = 0$ -- as I positioned this sequence here, this is at $n = 0$ -- I would take this and multiply it by this and sum from $-\infty$ to $+\infty$. Or, for $n = 1$, I would position it here, for $n = 2$, I would position it here.

Well, you can kind of see what the idea is. Let's look at this a little more dynamically and see, in fact, how one sequence slides past the other, and how the output $y[n]$ builds up to the correct answer.

So the input that we're considering is a step input, which I show here, and the impulse response that we will convolve this with is a decaying exponential. Now, to form the convolution, we want the product of $x[k]$ -- not with $h[k]$, but with $h[n-k]$, corresponding to taking $h[k]$ and reflecting it about the origin and then shifting it appropriately. So here we see $h[n-k]$ for $n = 0$, namely $h[-k]$, and now $h[1-k]$, which we'll show next, is this shifted to the right by one point. Here we have $h[1-k]$ -- shifting to the right by one more point is $h[2-k]$, and shifting again to the right we'll have $h[3-k]$.

Now let's shift back to the left until n is negative, and then we'll begin the convolution. So here's $n = 0$, $n = -1$, $n = -2$, and $n = -3$.

Now, to form the convolution, we want the product of $x[k]$ with $h[n-k]$ summed from $-\infty$ to $+\infty$. For n negative, that product is 0, and therefore the result of the convolution is 0. As we shift to the right, we'll build up the convolution, and the result of the convolution will be shown on the bottom trace.

So we begin the process with n negative, and here we have $n = -1$. At $n = 0$, we get our first non-zero contribution. Now as we shift further to the right corresponding to increasing n , we will accumulate more and more terms in the sum, and the convolution will build up. In particular for this example, the result of the convolution increases monotonically, asymptotically approaching a constant, and that constant, in fact, is just simply the accumulation of the values under the exponential.

Now let's carry out the convolution this time with an input which is a rectangular pulse instead of a step input. Again, the same impulse response, namely a decaying exponential, and so we want to begin with $h[n-k]$ and again with n negative shown here. Again, with n negative, there are no non-zero terms in the product, and so the convolution for n negative will be 0 as it was in the previous case. Again, on the bottom trace we'll show the result of the convolution as the impulse response slides along.

At $n = 0$, we get our first non-zero term. As n increases past 0, we will begin to generate an output, basically the same as the output that we generated with a step input, until the impulse response reaches a point where as we slide further, we slide outside the interval where the rectangle is non-zero. So when we slide one point further from what's shown here, the output will now decay, corresponding to the fact that the impulse response is sliding outside the interval in which the input is non-zero. So, on the bottom trace we now see the result of the convolution.

OK. So what you've seen, then, is an example of discrete-time convolution. Let's now look at an example of continuous-time convolution. As you might expect, continuous-time convolution operates in exactly the same way. Continuous-time convolution-- we have the expression again $y(t)$ is an integral with now $x(\tau)$ and $h(t-\tau)$. It has exactly the same kind of form as we had previously for discrete-time convolution-- and in fact, the mechanics of the continuous-time convolution are identical.

So here is our example with $x(t)$ equal to a unit step and $h(t)$ now a real exponential times a unit step. I show here $x(t)$, which is the unit step function. Here we have $h(t)$, which is an exponential for positive time and 0 for negative time.

Again, looking back at the expression for convolution, it's not $x(t)$ that we want, it's $x(\tau)$ that we want. And it's not $h(t)$ or $h(\tau)$ that we want, it's $h(t-\tau)$. We plan to multiply these together and integrate over the variable τ , and that gives us the output at any given time. If we want it at another time, we change the value of t as an argument inside this integral.

So here we have $x(t)$, and here we have $h(t)$, which isn't quite what we wanted. Here we have $x(\tau)$, and that's fine-- it's just $x(t)$ with t relabeled as τ . Now, what is $h(t-\tau)$? Well, here's $h(\tau)$, and if we simply turn that over, here is $h(t-\tau)$. And $h(t-\tau)$ is positioned, then, at τ equal to t .

As we change the value of t that change the position of this signal, now we multiply these two together and integrate from $-\infty$ to $+\infty$ with $h(t-\tau)$ positioned at the appropriate value of t . Again, it's best really to see this example and get the notion of the signal being flipped and the two signals sliding past each other, multiplying and integrating by looking at it dynamically and observing how the answer builds up.

Again, the input that we consider is a step input. And again, we use an impulse response which is a decaying exponential. To form the convolution, we want the product of $x(\tau)$ -- not with $h(\tau)$, but with $h(t-\tau)$. So we want $h(t)$ time-reversed, and then shifted appropriately depending on the value of t .

Let's first just look at $h(t-\tau)$ for t positive corresponding to shifting $h(-\tau)$ out to the right, and here we have t increasing. Here is t decreasing, and we'll want to begin the convolution with t negative, corresponding to shifting $h(-\tau)$ to the left.

Now to form the convolution, we want the product of these two. For t negative, there are no non-zero contributions to the integral, and so the convolution will be 0 for t less than 0. On the bottom trace, we show the result of the convolution, here for t negative, and for t less than 0, we will continue to have 0 in the convolution.

Now as t increases past 0, we begin to get some non-zero contribution in the product, indicated by the fact that the convolution-- the result of the convolution-- starts to build up. As t increases further, we will get more and more non-zero contribution in the integrand. So, the result of the convolution will be a monotonically increasing function for this particular example, which asymptotically approaches a constant. That constant will be proportional to the area under the impulse response, because of the fact that we're convolving with a step input.

Now let's carry out the convolution with an input which is a rectangular pulse-- again, an impulse response which is an exponential. So to form the convolution, we want $x(\tau)$ with $h(t-\tau)$ -- $h(t-\tau)$ shown here for t negative. To form the convolution, we take the integral of the product of these two, which again will be 0 for t less than 0. The bottom trace shows the result of the convolution here, shown as 0, and it will continue to be 0 until t becomes positive, at which point we build up some non-zero term in the integrand.

Now as we slide further, until the impulse response shifts outside the interval in which the pulse is non-zero, the output will build up. But here we've now begun to leave that interval, and so the output will start to decay exponentially. As the impulse response slides further and further corresponding to increasing t , then the output will decay exponentially, representing the fact that there is less and less area in the product of $x(\tau)$ and $h(t-\tau)$. Asymptotically, this output will then approach 0.

OK. So you've seen convolution, you've seen the derivation of convolution, and kind of the graphical representation of convolution. Finally, let's work again with these two examples, and let's go through those two examples analytically so that we finally see how, analytically, the result develops for those same examples.

Well, we have first the discrete-time case, and let's take our discrete-time example. In general, the convolution sum is as I've indicated here. This is just our expression from before, which is the convolution sum. If we take our two examples-- the example of an input which is a unit step, and an impulse response, which is a real exponential multiplied by a unit step-- we have then replacing $x[k]$ by what we know the input to be, and $h[n-k]$ by what we know the impulse response to be, the output is the expression that we have here.

Now, in this expression-- and you'll see this very generally and with some more complicated examples when you look at the text-- as you go to evaluate these expressions, generally what happens is that the signals have different analytical forms in different regions. That's, in fact, what we have here.

In particular, let's look at the sum, and what we observe first of all is that the limits

on this sum are going to be modified, depending on where this unit step is 0 and non-zero. In particular, if we first consider what will turn out to be the simple case-- namely, $n < 0$ -- for $n < 0$, this unit step is 0 for $k > n$. With $n < 0$, that means that this unit step never is non-zero for k positive.

On the other hand, this unit step is never non-zero or always 0 for k negative. Let me just stress that by looking at the particular graphs, here is the unit step $u[k]$. Here is the unit step $u[n-k]$, and for $n < 0$, so that this point comes before this point, the product of these two is equal to 0. That means there is no overlap between these two terms, and so it says that $y[n]$, the output, is 0 for $n < 0$.

Well, that was an easy one. For $n > 0$, it's not quite as straightforward as coming out with the answer 0. So now let's look at what happens when the two unit steps overlap, and this would correspond to what I've labeled here as interval 2, namely for $n > 0$. If we just look back at the summation that we had, the summation now corresponds to this unit step and this unit step, having some overlap.

So for interval 2, corresponding to $n > 0$, we have $u[k]$, the unit step. We have $u[n-k]$, which is a unit step going backward in time, but which extends for positive values of n . If we think about multiplying these two together, we will get in the product unity for what values of k ? Well, for k starting at 0 corresponding to one of the unit steps and ending at n corresponding to the other unit step. So we have an overlap between these for k equal to 0, et cetera, up through the value n .

Now, that means that in terms of the original sum, we can get rid of the unit steps involved by simply changing the limits on the sum. The limits now are from 0 to n , of the term $\alpha^{(n-k)}$. We had before a $u[k]$ and $u[n-k]$, and that disappeared because we dealt with that simply by modifying the limits. We now pull out the term α^n , because the summation is on k , not on n , so we can simply pull that term of the sum. We now have $\alpha^{(-k)}$, which we can rewrite as $(\alpha^{(-1)})^k$.

The upshot of all of this is that $y[n]$ now we can reexpress as α^n , the sum from 0 to n of $(\alpha^{(-1)})^k$. The question is, how do we evaluate that? It essentially

corresponds to a finite number of terms in a geometric series. That, by the way, is a summation that will recur over and over and over and over again, and it's one that you should write down, write on your back pocket, write on the palm of your hand, or whatever it takes to remember it. What you'll see is it that will recur more or less throughout the course, and so it's one worth remembering.

In particular, what the sum of a geometric series is, is what I've indicated here. We have the sum from 0 to r , of β^k . It's $1 - \beta^{(r+1)}$ -- this is one more than the upper limit on the summation-- and in the denominator is $1 - \beta$. So, this equation is important. There's no point in attempting to derive it. However you get to it, it's important to retain it.

We can now use that summation in the expression that we just developed. So let's proceed to evaluate that sum in closed form. We now go back to the expression that we just worked out-- $y[n]$ is α^n , the sum from 0 to n , $(\alpha^k)^k$. This plays the role of β in the term that I just-- in the expression then I just presented. So, using that result, we can rewrite this summation as I indicate here.

The final result that we end up with after a certain amount of algebra is $y[n]$ equal to $(1 - \alpha^{(n+1)}) / (1 - \alpha)$. Let me just kind of indicate with a few dots here that there is a certain amount of algebra required in going from this step to this step, and I'd like to leave you with the fun and opportunity of doing that at your leisure.

The expression we have now for $y[n]$, is $y[n] = (1 - \alpha^{(n+1)}) / (1 - \alpha)$. That's for n greater than 0. We had found out previously there it was 0 for n less than 0. Finally, if we were to plot this, what we would get is the graph that I indicate here. The first non-zero value occurs at $n = 0$, and it has a height of 1, and then the next non-zero value at 1, and this has a height of $1 + \alpha$, and this is $1 + \alpha + \alpha^2$. The sequence continues on like that and asymptotically approaches, as n goes to infinity, asymptotically approaches $1 / (1 - \alpha)$, which is consistent with the algebraic expression that we have, that we developed, and obviously of course is also consistent with the movie.

That's our discrete-time example, which we kind of went through graphically with the

transparencies, and we went through graphically with the movie, and now we've gone through analytically. Now let's look analytically at the continuous-time example, which pretty much flows in the same way as we've just gone through.

Again, we have the convolution integral, which is the integral indicated at the top. Our example, you recall, was with $x(t)$ as a unit step, and $h(t)$ as an exponential times a unit step. So when we substitute those in, this then corresponds to $x(t)$ and this corresponds to $h(t-\tau)$. Again, we have the same issue more or less, which is that inside that integral, there are two steps, one of them going forward in time and one of them going backward in time, and we need to examine when they overlap and when they don't. When they don't overlap, the product, of course, is 0, and there's no point doing any integration because the integrand is 0.

So if we track it through, we have again Interval 1, which is $t < 0$. For $t < 0$, this unit step, which only begins at $\tau = 0$, and this unit step which is 0, by the time τ gets up to t and beyond. For $t < 0$, there's no overlap between the unit step going forward in time and the unit step going backward in time. Consequently, the integrand is equal to 0, and consequently, the output is equal to 0.

We can likewise look at the interval where the two unit steps do overlap. In that case what happens again is that the overlap, in essence of the unit step, tells us, gives us a range on the integration-- in particular, the two steps overlap when t is greater than 0 from $\tau = 0$ to $\tau = t$. For Interval 2, for $t > 0$ -- again, of course, we have this expression. This product of this unit step and this unit step is equal to 1 in this range, and so that allows us, then, to change the limits on the integral-- instead of from $-\infty$ to $+\infty$, we know that the integrand is non-zero only over this range.

Looking at this integral, we can now pull out the term which corresponds to e^{-at} , just as we pulled out a term in the discrete-time case. We notice in the integral that we have $e^{-a\tau}$, so that gives us the integral from 0 to t of $e^{-a\tau}$. If we perform that integration, we end up with this expression. Finally, multiplying this by

e^{-at} , we have for $y(t)$, for t greater than 0, the algebraic expression that I've indicated on the bottom.

So we had worked out t less than 0, and we come out with 0. We work out t greater than 0, and we come out with this algebraic expression. If we plot this algebraic expression as a function of time, we find that what it corresponds to is an exponential behavior starting at zero and exponentially heading asymptotically toward the value $1/a$.

We've gone through these examples several ways, and one is analytically. In order to develop a feel and fluency for convolution, it's absolutely essential to work through a variety of examples, both understanding them graphically and understanding them as we did here analytically. You'll have an opportunity to do that through the problems that I've suggested in the video course manual.

In the next lecture, what we'll turn to are some general properties of convolution, and show how this rather amazing representation of linear time-invariant systems in fact leads to a variety of properties of linear time-invariant systems. We'll find that convolution is fairly rich in its properties, and what this leads to are some very nice and desirable and exploitable properties of linear time-invariant systems. Thank you.