

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.

PROFESSOR: Ladies and gentlemen, welcome to this lecture on non-linear finite element analysis of solids and structures. In this lecture, I'd like to continue with our discussion of the methods of solution for the finite element equations that we are encountering in non-linear dynamic analysis. In the previous lecture, I pointed out that the techniques used for the solution of these finite element equations in non-linear dynamic analysis can be classified as data integration techniques, methods of multiple position, and substructuring. We talked in the previous lecture about these types of methods, and in particular, we discussed in detail an explicit integration method, namely the central difference method, and an implicit integration method, namely the trapezoidal rule. We talked about how these methods are applied, how time steps might be selected, what the restrictions of these methods are, and so on. I'd like to now in this lecture discuss with you the method of multiple position, and also substructuring technique.

After the discussion off these two techniques, I'd like to talk with you about some examples. As a first example, we want to consider the wave propagation in a rod. As a second example, I'd like to share with you experiences regarding the response of a 3 degree of freedom system. As a third example, then, we want to consider the analysis of a 10-story tapered tower. Example 3 and example 1 are really solutions of a linear system, but we learn quite a bit, even regarding non-linear analysis, from these examples.

As the fourth example, then, I'd like to share with you experiences obtained in the analysis of a pendulum that goes through very large displacement. As a fifth example, then, we consider the pipe whip response solution. This is of much practical interest. And after these examples, which I've prepared view graphs for, I'd like to consider with you some slides. On these slides, we will have the analysis of a

control rod drive housing with the response of a spherical cap, a highly non-linear problem, large displacement, elastoplastic response, and then the analysis of a fluid structure interaction problem, a pipe test in which we are comparing our computer's results with experimental results.

I should point out right now that the details regarding these example solutions are given in some papers that I had also earlier. And please look at the study guide, where the references to these papers are given.

Let me now walk over here to the view graph that we have prepared regarding this lecture. And the first method of solution that I wanted to discuss with you is the method of multiple position. Quite a number of you probably have used this technique already in linear analysis. In fact, the word "superposition" somehow implies that you are dealing with a linear analysis. However, we can extend that method directly to non-linear analysis if we recognize that the mode shape vectors of the system at a particular time can be used as generalized displacements, or basis vectors, to express the response to any other time, as well. The method is effective when, in non-linear analysis, the response lies in only a few vibration modes, just the same way as we are applying multiple position also in linear analysis, namely, it's only effective really when a few vibration modes capture the total dynamic response. The same holds here, of course, in non-linear analysis. And when the system has possibly only local non-linearities.

Let's look at is equations that we're using in this multiple position technique. The basic starting point are the equations that we have seen before regarding dynamic analysis. But notice in this particular equation, I have assumed now that the damping matrix is not present, so we have no damping in the system coming from the damping matrix. There could, of course, be damping due to the material non-linearities, which are captured in the K and the F matrix.

Well, this is the set of equations that we want to solve. And usually, we just iterate on these equations until the right-hand side is 0, where of course we also would bring this one here onto the right-hand side. So think of this one here as a minus on

the right-hand side, and if, then, this right-hand side, with M on that right-hand side is 0, meaning the data U_K is 0, no more updating to the incremental displacements, then we have established equilibrium configuration corresponding to time t , of course, including inertia effects.

In multiple position, we don't want to directly to deal with these equations in this form. Instead, we want to transform them into a new form, and into a form such that we obtain less unknowns to deal with. Let's assume for the moment that τ is equal to 0 so that we are dealing with the solution of the equations using the initial stress method. We talked about initial stress method in an earlier lecture. Then using this transformation here, where ϕ_r are the eigenvectors of this problem down here, And the x_i are the generalized displacements corresponding to these eigenvectors.

Notice this equation is exactly the same that we're used to seeing in linear analysis. Of course, in linear analysis, we are usually not putting this $t + \Delta t$ up there. It's understood that the time dependency is in x_i , but the space dependency is in ϕ_i . And notice we are summing here from r to s . r would have to be selected by the analyst, and s the same way.

Having introduced this transformation with this eigenvalue problem here, and substituting for the U displacements in the general equilibrium equation, and using the same approach as we are used to in linear analysis, we directly obtain this set of equations. Now, notice here that in this vector \ddot{x} , and in this vector x , Δ meaning increment, we have only the components corresponding to the r mode up to the s mode. These are the generalized displacements corresponding to ϕ_r up to ϕ_s . Here we have ϕ_r to ϕ_s listed in this matrix, capital ϕ . Lower-case ϕ for these here. These are the columns in this matrix that is denoted as a capital ϕ .

Notice here we define an ω^2 matrix, 0 elements in the off-diagonals. And on the diagonal, we have the frequencies of the system, of the actually linear system, because we use U_K in the eigenvalue solution. These are the entries for this equation up here. And notice that we have now here much less equations than

the total number of finite element degrees of freedom. Typically, r might be 1, s might be 20, and then we would have 20 equations here, 20 equations only. Of course, r could also be 25, and s 30, in which case we have here a multiple position corresponding to the 25th up to the 30th mode. In any case, in general, we will have much less equations than finite element degrees of freedom that are listed, for example, in this vector here, R , I should say. The externally applied load corresponding to these degrees of freedom are listed in here.

Of course, the F vector is the vector of nodal point forces corresponding to the element stresses at time t plus Δt , and at the end of iteration, k minus 1. This matrix times that vector gives us a vector which has the same length as this vector has and that vector has. So we are not dealing anymore with all the degrees of freedom that are listed for which we have the forces here in the solution of the equation, but with much less degrees of freedom, the degrees of freedom being equal to the number of mode shapes that are being employed.

An interesting and most important point is that in non-linear analysis, we have a coupling, due to this vector appearing in here, between the different mode shape equations, or modal equations. We cannot solve, as we usually do in linear analysis, one equation of the modes of a position after the other. In linear analysis, you would typically solve the dynamic response equation in mode 1, then the one in mode 2, then the one in mode 3, and like that, you'd sweep over all the generalized degrees of freedom, one after the other, for the total response from time 0 to the end of the response time that you're interested in. In non-linear analysis, we have to go step by step still forward in time, solving all equations, all modal equations at time t , then all modal equations at time t plus Δt , then at time t plus $2\Delta t$, and so on, because there is this coupling, and because we don't know f unless we have all the information coming from all the mode shapes.

Well, a typical problem that you might want to solve with this approach is the pipe whip problem. Here we have a very long, slender pipe. There is here a restraint that has a gap. And suddenly, this end of the pipe is subjected to a very large force. The pipe undergoes dynamic motion, a wave travels along here, goes plastic, hits this

stop, et cetera. It's a very complicated problem, and this approach of multiple position is it cheap way of obtaining a solution to this problem.

Of course, the accuracy of the solution depends on how many mode shapes you include. If you only include a few mode shapes, the solution is quite inexpensive, but the accuracy is not very high. If you include as many mode shapes as there are finite element degrees of freedom, then of course you get the exact response. All you have done is transformed from one basis, the finite element degrees of freedom basis, to the mode shape basis. But then, of course, the solution is very expensive, and you don't want to ever do is that. The point is that you want to perform such analysis if you only need to include a few modes in the response solution. We would actually look at this problem, or such a problem, a little later, when we discuss example analyses.

As a second technique that can also be very useful in non-linear dynamic analysis, I'd like to discuss with you the substructuring technique. This procedure is used is implicit time integration, and the important aspect is that prior to the analysis, we are setting u this effective stiffness matrix corresponding to the linear degrees of freedom only, and we are statically condensing out all the degrees of freedom that will remain linear throughout the response solution.

It is quite effective for local non-linearities. Here, for example, we have a tower structure, and the only non-linearity that this structure sees, say, is the slip down here. This typically would mean that we have a large number of linear degrees of freedom that we can deal with very effectively, using the substructuring approach. The non-linear degrees of freedom here, of course, require iteration and updating of matrices, and those we deal with in the usual way.

Let me show you in more detail what I mean here. Let's look schematically at an example. Here we have a 10-story building, and that 10-story building is discretized, idealized, as shown here. We also want to include the elasticity of the foundation, and that is modeled very simply here as an assemblage of springs.

We deal here with two substructures, shown in red for the substructure degrees of

freedom and in black for the master degrees of freedom. This is also part of the master degrees of freedom. Notice that here we have one substructure-- in other words, this part here is nothing else than that part right here. That's what is one substructure. And of course, you see here another substructure. But we have a repetition here-- in other words, this substructure here is identically equal to that substructure. This, of course, is the fact because of the particular design of this building. The upper stories look the same way as the lower stories. That's the reason why we can say, this is one substructure model which can model this part of the structure and that part of the structure, as well.

If we look at the stiffness matrix, or the effective stiffness matrix of this structure that he would encounter in a direct integration solution using implicit time integration, schematically, we see the following. Notice here we have the entries coming from the substructures, and here we have entries from the master structure. This coupling here corresponds also to master degrees of freedom, and similarly, this coupling corresponds to master degrees of freedom.

Here we have the displacement vector, the incremental displacement vector. And let's see here that these master degrees of freedom of course correspond to these element entries here. These master degrees of freedom correspond to these entries, and these master degrees of freedom correspond to these entries.

Now what we can do is condense out all the linear degrees of freedom effects. And that is achieved as follows. We recognize that the total effective stiffness matrix corresponding to time t is made up of this part here, which is the linear stiffness matrix plus the total mass matrix. I should point out here that this K matrix contains all the linear contributions, where that M matrix here corresponds to all the contributions from the masses at the master degrees of freedom and the substructure degrees of freedom. Because we're assuming here that the total mass matrix is linear, we can put all the mass effects right in there.

And then since this is not the total effect of stiffness, we add the effect of the non-linear element stiffnesses, right here. And in this particular analysis, this might be,

for example, the non-linear springs down there. The non-linear springs modelling the foundation on the structure.

Well, this can be written in this form. And if we perform static condensation on this combined matrix, we arrive at this matrix here. Notice here we have a condensed degree of freedom vector carrying only the master degrees of freedom, and here we have the entries corresponding to the master degrees of freedom in the $t\hat{K}$ matrix. We put a c on there, meaning condensed. After static condensation, this is a matrix that has been reached. These black squares here correspond to as the black squares that you saw earlier. Of course, they have now been changed a bit because of the static condensation, because the effect of the linear degrees of freedom have been carried into these entries here.

The major steps in the solution are as follows. Prior to the step-by-step solution, we establish the \hat{K} , the linear effective stiffness matrix containing K plus 4 over $\Delta t^2 m$. And this \hat{K} matrix is then used to statically condense out all internal substructure degrees of freedom to obtain this \hat{K}_c , c standing for condensed matrix.

We note that the actual effective stiffness matrix including the non-linear effects is then obtained by taking the linear effective stiffness matrix and adding the non-linear stiffness effects to it. Notice, once again, this matrix is obtained from this matrix here, as I mentioned earlier. We use this matrix to statically condense out all the linear degrees of freedom.

The solution is then obtained in each time step as follows. We update the condensed matrix \hat{K}_c -- there should be no t here-- \hat{K}_c for non-linearities. That means we're adding the $t\hat{K}$ matrix to it, corresponding to the non-linearities that have occurred in the system. We establish the complete load vector for all degrees of freedom, and condense out the substructure internal degrees of freedom.

This is important to recognize. We have to calculate the complete load vector because of the inertia effects on the internal substructure degrees of freedom. The inertia forces on the internal substructure degrees of freedom carry forces to the

master degrees of freedom that we have to take into account in each time step and in each equilibrium iteration. We then solve for the master degrees of freedom, using the small system size, the \hat{c} , the velocities, accelerations, and calculate all the substructural degrees of freedom-- velocities, displacements, and accelerations. These are required here to calculate the load vector corresponding to all degrees of freedom.

Schematically, the solution procedure, therefore, proceeds as follows. For each time step and each iteration, we start with the displacement velocities and accelerations at time t . We calculate the effective load vector, corresponding to time t plus Δt , the way I've been talking about it earlier. And then we condense that load vector, as shown here, to the condensed load vector, \hat{c} meaning condensation. We then, from this load vector, calculate with the effective tension stiffness matrix, the t plus Δt U_c vector, and this vector is used to calculate the displacement, velocity, and accelerations of all the degrees of freedom at time t plus Δt . Of course here I'm showing only the process for a typical time step. If you iterate, as we usually do in non-linear analysis, of course, then you would have to also introduce the iteration counter here in these vectors.

Notice that using this solution procedure, we get the same results as if we had not performed the substructuring. In other words, there is no additional assumption that we're introducing into the solution. You get the same results. However, the solution cost is reduced, because we are dealing with one substructure, a typical substructure, only once regarding the factorization-- the LDL transpose factorization and static condensation that has to be performed.

Let us now look at some example solutions. And the first example that I'd like to consider with you is the solution of the wave propagation in a rod. Here we have a uniform freely floating rod, and we apply to the left end of that rod a force that varies as shown here as step load of 1000 newtons. The material data for the rod are given here. We want to perform a linear elastic analysis, but although this is a linear analysis, we still will learn quite a bit from it regarding non-linear analysis.

We want to consider the compressive force at a point at the center of the rod, right there. And the exact solution for this force at that point is plotted here. You have, once again, apply 1000 newtons, and here we have the time scale. Notice up to $1/2 t^*$, there's no force at point a, because the applied force in fact has to propagate through the rod. Then we have the full force up to $3/2 t^*$, and then no force again, where t^* is the time for the stress waves to travel through the complete rod. This is the analytical solution, well-known, well-established, and we want to compare with this solution when we perform our finite element analysis.

For the finite element analysis, we use a mesh of 2-noded truss elements, as shown here, and we measure the force at the midpoint of the total structure by the compressive force in this element. We use in this analysis first essential difference method, and recognize that the time step that we use has to be small or equal to the critical time step, as we discussed in the previous lecture. And that critical time step is given by L_e/c . Please refer to as the previous lecture notes, and you would see this formula there, which is nothing else than t^* divided by the number of elements. We recall that if Δt is greater than $\Delta t_{critical}$, we would produce an unstable solution.

We also need to deal with the initial conditions for this problem. And we recognize that if the initial displacements are 0 and R is given, then we can calculate the initial acceleration, as shown here. Using now a time step equal to the critical time step, we obtain the correct result, and I pointed that out in the earlier lecture already, that we in fact obtain the exact solution for any number of finite elements to model the rod. And here is a finite element solution, compared with the exact solution. A wonderful solution, right on the exact solution.

Using, however, now a smaller time step-- and this is frequently of surprise-- using a smaller time step, we get a much worse result. Here we use a time step of $1/2 \Delta t_{critical}$. Half the time step that we used before. And just solution looks actually very bad right here.

Frequently one sort of intuitively thinks, well, if I use a smaller time step, I should get

better results. My costs go up, and surely I should get better results. Well, it turns out that using the explicit time integration scheme for this type of problem, this wave propagation problem, with a smaller time step, more cost, higher cost, you get worse results. That is really something to keep in mind.

Now let us consider the trapezoidal rule. The trapezoidal rule-- we discussed this method in the earlier lecture. And I'd like to refer you once again to the notes corresponding to that lecture, where we have identified that a stable solution would be obtained with any choice of Δt . Remember, this is a linear analysis, and the trapezoidal rule is unconditionally stable.

We could use a consistent or lumped mass matrix. And in this particular case, we use a lumped mass matrix, because we also like to compare our results with those that we obtained using the central difference method.

If we use the trapezoidal rule with the time step that we employ in the central difference method-- in fact, the critical time step that we used there-- and we also use the initial conditions computed as shown here, the way we did it for the central difference method solution, we find that the solution is very inaccurate when compared to the exact solution. Here, of course, we also use the 10-element mesh.

If we use 0 initial conditions for the trapezoidal rule, we get this solution. Not much different, really, from the solution that I just showed you. So whether we use 0 initial conditions or the initial conditions corresponding to the equation $\ddot{u} = R$, which we have seen on the previous view graph, makes very little difference when using the trapezoidal rule. It makes a lot of difference using the explicit central difference method. And for that reason, in fact, we have to use the proper initial conditions obtained from $\ddot{u} = R$ at time 0 when we use the central difference method.

Using now twice the time step of what we just used, namely, Δt equal to 2 times Δt , critical of the central difference method, this solution is stable, but very inaccurate. It is stable because we're using an unconditionally stable operator. We discussed it in the previous lectures. Note that this is the solution we're obtaining

when compared to the exact solution.

Now we look at the solution where we have half the time step off the critical size, corresponding to the central difference method. And you see here, the solution, again, very inaccurate when compared to the exact solution, but of course stable, because the method is always stable for any time step Δt . The question now could, of course, be, and should be, what happens if we take more elements to discretize the whole structure? Here we use the 10-element model. Why not use more elements?

And here we now look at the solution obtained with the essential difference method when we use 100 elements to model the total bar. Here the exact solution in red, and that exact solution is reached if you integrate the equilibrium equations with this time step. We discussed that earlier. That exact solution is always obtained for any number of elements that you use to discretize a structure, provided you use Δt equals $\Delta t_{critical}$. If you then use a smaller time step, in fact, Δt equals $1/2 \Delta t_{critical}$ was used here, we obtain a solution that is not the exact solution, but as you can see here, it oscillates for this time step this amount about the exact solution.

So we have much the same observations that we have seen already earlier, namely, when the time step is decreased in the central difference method solution-- from the critical time step, here we see a decrease-- the solution that one might expect to be better actually deteriorates. Of course, it couldn't be better than the exact solution. But we see here a deterioration, and a quite significant one.

If you use the trapezoidal rule, this is the solution that you obtain with the time step equal to the critical time step of the central difference method. Again, here, we're using the 100-element mesh. The solution is certainly better than what we have seen for the 10-element mesh, but it does not give us the exact solution-- in other words, a [UNINTELLIGIBLE] wavefront as we would like to obtain.

Of course, this is a very difficult problem to solve, because we have infinitely steep wave fronts here and there, basically, to model. And an overshoot would obviously

obtain it, but we would like to have said overshoot and undershoot, with this oscillation here as small as possible, of course, with the solution scheme that we employ.

There is one more important consideration that I'd like to bring to your attention. Here we use a 2-dimensional element mesh to model the bar. Notice these are 4-node elements, plane stress elements here to model the bar. And in this particular case, we use the central difference methods. The thickness is given here, the material data are given here. For this solution, notice that this distance here is half the distance that we have here. And that is something that leads to a phenomenon that one has to be aware of.

Notice our the critical time step that we identified for the one-dimensional wave propagation here would, of course, be calculated depending on this length. However, we have to allow for the fact that there is also a possibility of wave propagation this way, and the actual critical time step that we have to use in this problem is given by this length, by the shortest element distance. Please refer to the previous lecture, in which I discussed this already. You have to use the shortest distance between nodal points within an element to determine the critical time step.

So this relationship here does not hold anymore, because this distance is less than that. And if you were to use Δt equal to t^* over 10 elements, which was a critical time step that we used when we solved the truss element model, then we would find that the solution diverges. In fact, in element 5, which was the element that we always used to measure our stress, we find that τ_{zz} , the absolute value of τ_{zz} , is greater than 1000 newtons divided by the area at this time. Notice that it was 1000 newton that we applied. So a stress that should be 0, really, because τ_{zz} should really be 0, it's the stress in the transverse direction, the wave travels this way, and τ_{zz} is acting that way. This stress should be 0. We find, actually, that it grows, and it grows because we are using a time step that is larger than the critical time step that should be used for this problem.

Let us now look at another example, a second example from which we can also

learn quite a bit. Here we have a very simple 3 degree of freedom system. These are the degrees of freedom, x_1 , x_2 , x_3 . Notice we have masses m here, and we have a linear spring connecting these masses, a non-linear spring connecting these masses, and a linear spring connecting this mass to the rigid support.

For this problem, the linear spring has its distance here, and the non-linear spring is characterized by this force displacement relationship. Notice there's a kink going from 1 to 100. So the spring gets suddenly, at a particular distance, or rather deformations, I should say-- displacement x_2 minus x_3 -- at a particular deformation, the spring becomes 100 times as stiff as prior to that point.

Well, we can calculate a critical time step based on this spring stiffness here, the linear part, or rather, this spring stiffness for small displacements. And that critical time step for this problem is given as 1.11 seconds. As long as the response of this system is such that we would, in this non-linear spring, never go around this kink, we would deal with a linear system, and this would be the proper critical time step to use.

However, as soon as we are going around this kink here, due to the fact that the forces that I applied, or the initial conditions that we are applying to these degrees of freedom, are large enough, as soon as we go around this kink, our new critical time step is based on this spring stiffness, and if this given right down here. This is a new critical time step. So if we now perform essential difference method solutions of this problem, and if the solution brings us into this regime, it is very important to use, as a critical time step, 0.14 seconds. It is very important, even if we're only in this regime for just a few steps.

You see, what might happen in the response is that most of the time, you're in this regime. But then all of a sudden, for just a few steps, you're getting into this regime, and because of this high stiffness, you're thrown back into this regime. If your time step that you use in the time integration is then larger than that value, you are having an instability for just these few time steps, and that introduces an error in the solution which, if you don't have a comparison with an exact analytical solution, will

of course be an error that you not necessarily will see, directly. And that's one of the important points.

You see, in linear analysis, when you use the essential difference method in linear analysis with a time step that is larger than the critical time step, you will, after a few time step integrations, see that the response becomes very large, and the instability is obvious. The instability is really obvious when you apply the central difference method with a time step that is larger than the critical time step to the solution of the system.

However, in non-linear analysis, as I mentioned in the earlier lectures, and you are now well aware of it, in non-linear analysis, the frequencies change. The critical frequency being the largest frequency in the system changes. The smallest period of the system changes, as we are seeing here in this view graph, also.

And it is important that your time step is always smaller than the critical time step based on the highest frequency that you can ever encounter in the solution. If that is not the case, then you might have only an instability for a few time steps, and the solution may not have time to blow up, or give you very large displacements. So the instability will not be obvious. Instead, the instability over these few time steps only introduces an error in the solution, an error of which you really don't know the magnitude. And after these few steps have been completed, you're using now a time step that is again stable, giving a stable solution.

That error which we have accumulated, of course, is carried forward, but you have not seen a blow-up of the solution. And that's one important consideration when you apply the central difference method in non-linear analysis. Of course, the important point, the sort of ground rule that you have to observe, is that Δt must always be smaller than the critical time step of the mesh where that critical time step changes throughout the solution.

Let's look at this example here. What happens if we violate that condition? I show you here the solution with Δt equals 0.15 seconds, compared to the solution with Δt 0.05 seconds. Notice that the response of the right math here for both of

these time steps is slightly different, but we can say that for this time step, we don't get displacements that are very much different than we obtain with this time step.

If we look at the response of the center mass, there are more differences in displacements. And if we look at the response of the left mass, in fact, we have even larger differences. But the important point is that if you only have this solution here corresponding to the time step 0.15 seconds, which is larger than the critical time step, you would certainly not expect that this solution was unstable. It in fact was only unstable for a very few time steps. There is an error accumulation, fairly large here. In fact, very large. You can see that these displacements are only about 50% of those displacements. But you cannot see directly an instability here if you don't know this solution.

Let's look at the force in the center truss. Here we list the time, here we list the solution for Δt , 0.05 seconds. And here we list the solution for Δt 0.15 seconds. And you can see here that at the beginning, the forces compare quite well, and then the error becomes larger and larger. In fact, here we have an opposite sign in the force. This being here, say, the accurate solution with a small enough time step, and this being, of course, a grossly inaccurate solution. But it is not evident that this is a grossly inaccurate solution if you don't know this solution. And that's the point I wanted to make with this example.

The next example that I like to discuss with you is the analysis of a tower. The tower is shown here, subjected to a pressure that is induced by a blast. Here we have some geometric and material data regarding the tower. The applied load on the tower is shown here. Notice it's a load that we assume to be instantaneously there, and then to decay, as shown here, over 200 milliseconds. The purpose of the analysis is to determine the displacement and velocities at the top off the tower, to determine the moments at the base of the tower. And in this particular analysis, we use the trapezoidal rule and the lumped mass matrix to solve the problem.

We must make a number of decisions here, and two decisions are most important. We must choose a mesh that is, in other words, we have to establish an appropriate

finite element discretization for the tower, and we have to choose a time step for the time integration. The choice of mesh and the choice of the time step are two very closely related items, and that's what I would like to discuss with you in this example in more depth. Of course, the time step and the mesh must both depend on the loading that is supplied.

Let's look at some observations. The point is that the choice of mesh will determine the highest natural frequency, and of course, the corresponding mode shape, that is properly, accurately represented in the finite element analysis. The more elements you use, just as a rule of thumb, the higher the frequency that can be properly represented, that can be accurately picked up, with the choice of mesh we're using. The time step that you're using will of course determine how high a frequency you're integrating accurately. In other words, you want to use really enough elements to have a mesh that is accurate, and you want to use is small enough time step to integrate the frequency of the mesh that you need to integrate in order to have an accurate response prediction.

I used the word "accurate," and we have to now of course look closely at what I mean by that. It is most effective, really, to choose the mesh and the time step such that the highest frequency, accurately integrated, is equal to the highest frequency accurately represented by the mesh. To choose the mesh and the time step, we look at the applied loading, in other words, the physics of the problem, and we represent this applied loading as a Fourier series. If we do so, that Fourier series will display the important frequencies that are to be accurately represented by the mesh.

Here for this particular loading that we're looking at, we did a Fourier analysis, and you can see that with this Fourier analysis, we include terms now in our first analysis, in case one, up to f_n 17 hertz, and in case two, in our second analysis, up f_n equal to 30 hertz. If we truncate our Fourier representation of the loading at 17 hertz, we obtain this approximation here to the actual applied loading. If we truncate our Fourier representation at 30 hertz, we obtain this Fourier approximation to the actual loading. As you can see, surely, as expected, if you include more terms in the

Fourier approximation, your approximation to the actual loading becomes better.

Now having identified that we want to perform the analysis with two meshes, first there's a mesh that gives us frequencies accurate up to 17 hertz, and then with a mesh that is accurate up to 30 hertz, we have to look for those meshes. And we do so by establishing three meshes that are shown here. Here we have a 30-element mesh. Notice that one element spans from the bottom of a floor to the top off the floor, and across from one column to the other. A 60-element mesh, in which we simply have subdivided each element here in this more coarse mesh into two elements, and a 120-element mesh, in which we again have taken one element here and represent it as two elements. We use these meshes here to calculate frequencies. And on this view graph here, you see the frequencies calculated for the 30-element mesh and the 60-element mesh.

Notice mode numbers in this column. 30-element mesh results here, 60-element mesh results here, in each case, hertz. We notice that up to this red line, the difference between the frequencies is less than 1 hertz. In other words, the difference between the frequencies obtained from the 30-element mesh and the 60-element mesh is less than 1 hertz, and that is also right at the 17 hertz level. So we can say that the 30-element mesh really represents the frequencies up to 17 hertz quite accurately.

You should also remember, please, that since we are using a lumped mass matrix in the analysis, there's no bounding theorem saying that the frequencies of the 60-element mesh must always be below the frequencies of the 30-element mesh. For a larger mesh, we can actually have frequencies that are smaller or larger than for the more coarse mesh. That of course also applies for when we compare the 60-element mesh result with the 120-element mesh result later on.

So we say that for the 30-element mesh, we have accurately represented now the frequencies up to 17 hertz, and therefore, there is the mesh to use. We calculate the time step via is this relationship here as being 0.033 seconds, being the cutoff frequency, and Δt then is given here as 0.0017 seconds. This mesh now should

be integrated with this time step, and then we can say is that the mesh up to 17 hertz represents the frequencies quite accurately, and we have integrated the frequencies up to 17 hertz quite accurately, and therefore, we have an accurate solution of the problem up to 17 hertz-- and what I mean by this "up to 17 hertz" is that also in the Fourier representation, we in effect only have represented the load by the Fourier representation up to 17 hertz, in effect. Of course, in the actual step-by-step solution, we also will carry terms above 17 hertz. But these terms are assumed not to contribute to very much to the response, because we are not picking them up very well in the mesh and in the time step selection.

If we proceed now with the same ideas to also look at the 60-element mesh, then we perform there also the frequency analysis for the 60-element mesh and the 120-element mesh, and we find that up to 30 hertz in the 60-element mesh, we gain represent all the frequencies accurately. Notice that above the red line, below these frequencies here, the difference between these frequencies is less than 1 hertz. Above this line, the difference between the frequencies is larger than 1 hertz, and that is our cutoff line.

Well, if we want to perform now the direct time integration of the 60-element mesh, we proceed as for the 30-element mesh. We calculate our cut-off frequency, our Δt , the proper time step, 20 steps for the highest frequency or the smallest period that we want to integrate accurately. That's the rule of thumb we're using. And with this time step, we perform the analysis.

Notice a smaller time step would be accurately integrate frequencies which are not accurately represented by the mesh. Therefore, there is no point choosing a smaller time step. A larger time step would not accurately integrate all frequencies which are accurately represented by the mesh. It is not good to choose a larger time step, because your mesh is chosen to represent the frequencies up to 30 hertz quite accurately, and to get all of what you can get out of the mesh, so to say, you want to take a time step that is in accordance selected with the mesh. And that times step selection gives you this result here, 0.003 seconds.

Let's look now at the actual response that is calculated. Considering the horizontal displacement at the top of the tower, denoted as u , here, we find that in fact the 30-element mesh and the 60-element mesh give virtually the same response. If we look at the horizontal velocity at the top off the tower, we find that there are some small differences between the 30-element and 60-element mesh. If we were to look at the accelerations, these differences would of course be even larger. But the purpose of this analysis was to calculate the displacement and velocities at the top of the tower.

Let's look at the moment, which was also one requirement for our analysis, the moment right down here at the bottom off the tower. And here we see that the 30-element mesh gives the smooth result. The 60-element mesh gives the zig-zagged result. But overall, really, basically a good response prediction already with the 30-element mesh.

If we look at the displacements of the tower, we have amplified these displacements here. The 30-element mesh at 200 millisecond looks like that. The 60-element mesh. looks like that. At 400 milliseconds, the 30-element mesh looks as shown here. The 60-element mesh looks as shown there.

Regarding these analyses, we now can summarize some comments. We saw some high-frequency oscillations when we looked at the moment reaction response of the 60-element mesh. However, this high-frequency response, which actually can be identified to be 110 hertz, is probably inaccurate, because our mesh did not represent frequencies accurately up to that level. The obtained solutions for the horizontal displacement and the velocity at the top off the tower are really virtually identical for the 30- and 60- element mesh.

Let us now look at another example. This is the example of a pendulum. A large displacement analysis of this pendulum is performed, and basically what we're looking at here is, if I can just have your attention here, is at the pendulum that originally is horizontal, it's pinned right here. And we let this pendulum go through very large motions, and pendle around as shown here. It always should come back

to the horizontal position, unless we have damping in the system. Since we don't have damping in the system, the pendulum should just keep on going forever.

We have a frictionless hinge right here, and this is pictorially represented here on the view graph, a frictionless hinge right there. The pendulum is modelled by a truss element of stiffness EA . The length is given here. The mass is concentrated right there. So there's no distributed mass here. All of the mass of the pendulum is concentrated here, half of it here, and half of it there.

The initial conditions of the pendulum are listed here. The initial angle is 90 degrees theta, so the pendulum is up here, as I explained already. And initially, it has 0 velocity. And now we let the pendulum go, and as I said, it should spring up and down here, of course, under the gravity loading.

The dynamic response that we want to calculate is obtained using the trapezoidal rule. And we are using full Newton iterations to establish the equilibrium during every time step. The convergence tolerance that we use, and that I have to refer you now to our earlier view graph regarding convergence tolerances. The tolerance that we're using here is ETOL equal to 10^{-7} . We talked about what ETOL is in an earlier lecture.

In our first analysis, we chose the time step to be equal to 0.1 seconds. And in that case, the following response is obtained here. We expect that solution is shown as a solid line, and these circles here show the calculated finite element solution. Notice 90 degrees theta here minus 90 degrees here, and then we return to 90 degrees.

The solution looks quite fine, looks stable. And in fact, if you look at the textbook, you would find that there I'm giving the solution up to here. Everything looks good. However, if we plot the strain in the truss, we find that at larger times, the strain oscillates very much, and in fact, this oscillation grows. This clearly signifies an instability in the solution process. The instability is unchanged when we tighten our convergence tolerance or when we use the BFGS method instead of the full Newton method.

We should recall here that the trapezoidal rule is only unconditionally stable in a linear analysis. Here we are talking about a non-linear analysis, and you cannot select any time step magnitude Δt and expect a stable solution. So what we have to do now is to use a smaller time step, and we choose Δt equal to 0.025 seconds, and repeat the analysis. Θ degrees here as a function of time. And once again, a very nice, smooth response solution, starting with 90 degrees, going to minus 90 degrees, coming back to 90 degrees, and so on.

Of course, even with a larger time step, you got a smooth solution for θ . What really told us about the instability in the solution was our look at the strain in the truss. So we look now also at this strain, and in this particular case, with this time step, we see a smooth solution for the strain. You should please observe that the scale on the strain is here different from the strain scale that we had in the earlier view graph. If you refer back to the earlier view graph, you see that we had a different scale there. But we have a smooth solution here, in fact, a solution the way we would expect it.

If we don't iterate, then we obtain a solution that looks as shown here. In other words, instead of getting this oscillatory response that we would like to obtain, we obtain really a highly damped solution response. This, of course, shows the importance of iteration. We have to iterate in the solution of a dynamic problem. It's very important to integrate in order to establish equilibrium for every time step, for every discrete time that we're considering. In this particular case, we see that we obtained a highly damped solution. In other cases, the solution might blow up. So it is important to iterate in the nonlinear dynamic solutions when you use an implicit time integration scheme.

Although the solution obtained with the equilibrium iteration is highly inaccurate, the solution is, of course, stable. And that is also seen when one looks at the strain plot. Notice here we're plotting strains in the truss element as a function of time. No blow-up of that strain, no violent oscillation. It is, in fact, going up here, and then becoming this rather small oscillating, as shown here. So we don't have an unstable solution, but a highly inaccurate solution.

Let's look at the next example now, the analysis of a pipe whip problem. Here we have a long pipe, dimension given here and here, that is subjected suddenly to a concentrated load at its end. There is also, below the pipe, a restraint. That restraint is only active once this gap is closed. Notice there is a 3-inch gap here between the pipe and the restraint.

Of course, this restraint, in an engineering design, is installed to prevent violent motions of the pipe at this point. What our objective here is, to determine the transient response of this pipe. Notice that under this very large load, the pipe very rapidly becomes plastic. And the difficulty of this problem really lies in that the pipe going plastic becomes very flexible here, but the restraint suddenly hitting the pipe, or the pipe hitting the restraint, rather, means that a very large difference is suddenly introduced at this end.

The finite element model we used is shown here. 6 summation beam elements that of course can go plastic. There's more of the elastoplasticity in the system properly. And a truss element from this node to that node. This truss element contains a 3-inch gap.

The material properties for the pipes are listed here. Notice the material property corresponding to the elastoplasticity. We assume no strain hardening and for the restraint. Again, the restraint also is modeled as an elastoplastic material. The analysis performed, in this particular case, using multiple position with two modes and direct time integration. We want to compare the response that we predict using the multiple position procedure, which I discussed earlier in this lecture, with the response that we obtain using a time integration analysis.

For the time integration analysis, we use the trapezoidal rule. And in fact, even in the multiple position, we also use the trapezoidal rule. And for the mass representation, we have chosen a consistent mass matrix. In this particular case of convergence tolerance, ETOL was used with this value down here.

If we perform the eigenvalue solution of the initial system, in other words, of the

elastic pipe in the initial state, we obtain these natural frequencies, 8.5 hertz and 53 hertz. These are also the mode shapes. And of course, these frequencies will change as time progresses in the elastoplastic response, and also the mode shapes will change, particularly once their restraint is hit. However, what we are saying is that we believe we can capture the elastoplastic response of the pipe, approximately, by using this mode shape and that mode shape as generalized displacement patterns to represent the total pipe response. Of course, there are the heavy approximations that we reduce all of the pipe degrees of freedom to adjust 2 generalized degrees of freedom.

To integrate a response, we have to choose a time step. And here we have Δt 1/20 of the cutoff period that we want to use here. Notice that we have only two modes in this problem that we want to consider, with the frequencies that I showed you just now. And this equation gives us, as a time step, 1/1000 of a second. Of course, once again, this estimate, or this choice of time step, is based solely on looking at the linear response of the pipe. In other words, assuming that the pipe remains elastic and that the restraint is not there either.

If we look at the results obtained from these two analyses, we find that the multiple position response for the displacements looks as shown here, the direct time integration response looks as shown here. Notice that this is here the distance of the gap, so at this particular time, the gap is closed and the restraint comes into action. We note from this graph that the multiple position solution is really very close to the direct integration solution for the displacement at the tip off the pipe.

Now let us look at the moment at the built-in end of the pipe. And here we have the moment response plotted as a function of time for the direct integration solution, as shown here, and for the multiple position solution, as shown here. We notice that the maximum moment occurring at the built-in end in the pipe is almost the same when predicted by the multiple position solution and the direct integration solution. There's some small difference, but the difference is certainly not very large.

However, if we look at what happened at earlier times, we see that there is a large

difference between the results obtained using the multiple position solution and the direct integration solution. The reason is that the response here at earlier times is very complex, and the multiple position solution cannot pick that complex response up very well.

However, it's one important point you should keep in mind. And that is, if we had used more modes in the multiple position solution, we would come closer to the response obtained, predicted with the direct integration solution. In fact, if the number of modes used in the multiple position solution is equal to the number of degrees of freedom in the finite element mesh, then the multiple position solution gives you the same response as you calculate with the direct integration solution. I mentioned that earlier in the lecture already, and the reason, of course, is that in that case, all you would have done is a change of bases from the finite element displacements to the mode shape bases.

So if you use, even with this initial stiffness matrix that you have used to calculate the mode shapes, even with that initial stiffness matrix, if use all the modes, as many modes as there are degrees of freedom, in this multiple position solution, you would get the same response prediction with the multiple position solution as we're seeing here, with the direct integration solution. And that means we can be quite assured that as we increase the number of modes in the multiple position solution, we will get closer and closer to this complex response that is predicted using the direct integration solution.

In general, one, of course, does not know the error when you use the multiple position solution, and there must always be a question of how large is the error, for example, that we're seeing here. And that can only be assessed by a more accurate solution. However, I believe that's the multiple position solution is of much value in the initial design analysis stages when you want to obtain fairly inexpensive solutions for the response, and you want to just understand how the system acts that you are analyzing and designing. The multiple position solution really yields relatively inexpensive response solutions of the system which may be very helpful in the initial design stages of that structural system.

This completes, then, the discussion of our examples for which I have prepared view graphs. I'd like to now share with you some further experiences that are summarized on this slide. And let me go over here, where I have the first slide projected. I should point out, before we discuss the information on these slides once more, that we will go fairly rapidly through these slides, and that not all details pertaining to the solutions that we are looking at here are given on these slides. Please refer to the papers-- the references are given in the study guide-- in order to find all of the details pertaining to the solutions that I'm now discussing with the slides.

Here we consider now the analysis of a control rod drive housing, which is modelled using four beam elements and point masses. Notice that the structure is subjected to motion up here. The non-linearity in the system lies in the fact that there's a gap between the tip of that beam here and the spring on the right-hand side and on the left-hand side.

On the next slide, we see now the solutions obtained. We obtain a solution using direct integration, and using a multiple position analysis, including two modes. These two solutions are compared with a solution that was reported at some earlier time by Peterson and myself.

Here you see the comparison of the results obtained. Notice there is a difference here between the solution that we reported earlier, that Peterson and myself reported earlier, and the solution that we generated now. But that difference is due to the fact that Peterson and myself used, at that time, a different solution scheme. The response solution that we obtained with the multiple position analysis is very close to the response solution that we calculated using direct integration.

Now the next analysis that I'd like to discuss this you involves the analysis of a spherical shell subjected to a uniform pressure. And that pressure is suddenly applied as shown down here. Notice this is a step pressure constant in time. The material data of the shell are given here. We want to calculate the elastoplastic response. And some geometric data are given here.

Notice the shell was idealized using 10 8-node axis-symmetric elements. It's a spherical shell, so we use an axis-symmetric model to model the shell. We use Newmark integration with delta and alpha as given here. We chose these values because Nagarajan and Popov had chosen these same values, and we wanted to compare with their solution. We used 2-by-2 Gauss integration and a consistent mass matrix for the problem. The time step was 10 microseconds, and we used the total Lagrangian formulation in the analysis.

The next slide now shows the response that we predicted. Here we see the ADINA solution, a dashed line, compared with the solution by Nagarajan and Popov given by the solid line. Notice we are plotting here the apex displacement of the shell as a function of time. Now in this particular case, we use the consistent mass matrix, 2-by-2 integration, and specific solution parameter, and we want to do now change these parameters in order to see the effect of these parameters on the solution that is predicted.

The next slide now shows again the apex displacement of the shell as a function of time, but using once a consistent mass matrix, and once using a lumped mass matrix. And in this particular case, we use the trapezoidal rule, which corresponds to the Newmark integration with delta equals 0.5 and alpha equals 1/4. Notice that there are only slight differences between the lumped mass solution results and the consistent mass solution results.

Another parameter that one might want to vary is the integration order. Here we have the response, using 2-by-2 integration, which we saw already earlier. Actually, we didn't quite see this response earlier, because earlier we used different delta and alpha parameters. Very close to these parameters, however.

Anyway, here we have the 2-by-2 integration results, and then we have the 3-by-3 integration results and the 4-by-4 integration results here. Very small differences between these solutions, but it's interesting to perform these solutions, and really assure oneself that there are small differences. We know that shells are rather sensitive structures, and we want to make sure about the modeling that we're using

for the representation of these types of structures.

The next slide now shows the analysis of a problem for which we had experimental results to compare with. Here we have a pipe, a very rigid pipe, and a flexible pipe, a nickel pipe here. And these pipes, both being straight, are totally filled with water. At the end of this pipe here, a pulse gun sets off a pulse, and that pulse travels through the water, all the way along. And of course, there's an interaction between the pipe and the water. In other words, it's a fluid structure interaction problem.

This pipe here is very rigid, so that we model it as infinitely rigid. You will see that just now on the next slide. Notice here we have pressure gauges indicated. The pressures and the strains in the pipe, in this nickel pipe, the flexible pipe, were measured by the Stanford Research Institute. Here we have the nickel properties, we assume an elastoplastic material, and the water properties.

On the next slide now, we see the pressure pulse as a function of time. Of course, this was given to us for our analysis. And this slide now shows the finite element model that we use. We used an axis-symmetric model. The center line is right here. And notice that the water here in the rigid pipe was represented by 81 elements. Notice the rigid pipe was modelled as infinitely rigid, because these are rollers on an infinitely rigid surface, so to say. Notice here is a nickel pipe representation, and here we have water elements again, and a thin layer of water elements here.

On the next slide now, we see the response predicted by ADINA and the experimental results, the ADINA response shown as a solid line and the experimental result shown as a dashed line. Of course, the experimental results reported by the Stanford Research Institute. Notice we are plotting here the pressure as a function of time for one of the pressure stations. And here you'll see the same information for another pressure station, and for another pressure station, and one more pressure station-- the pressure stations that I pointed out to you earlier.

Notice that here finally we compare the hoop stress predicted by ADINA with the experimentally observed hoop stress as a function of time at one of the stations

within the pipe. Notice that for the hoop stress here, we have four experimental results corresponding to the measurements that were taken at the north pole, the south pole, east pole, and west pole on the pipe. In the ADINA solution, of course, we use an axis-symmetric model, and that axis-symmetric model only gives one hoop strain. And that hoop strain, as you can see here, lies quite nicely within the experimental results obtained by the Stanford Research Institute.

I should also point out that the comparison of the ADINA results with the experimental results is really very good, except that we don't pick up these high-frequency oscillations that have been observed in the experiment. Well, the ADINA finite element mesh that we used in this analysis is not able, really, to represent these high-frequency oscillations, and of course, that is certainly one of the reasons why we could not hope to pick these frequency oscillations up.

This, then, brings me to the last slide that I wanted to share with you in this lecture, and also to the end of this lecture. Thank you very much for your attention.