Computational Electromagnetics
Prof. Uday Khankhoje
Department of Electrical Engineering
Indian Institute of Technology, Madras

FDTD: Materials and Boundary Conditions
Lecture - 13.20
MEEP : FDTD in action

(Refer Slide Time: 00:14)

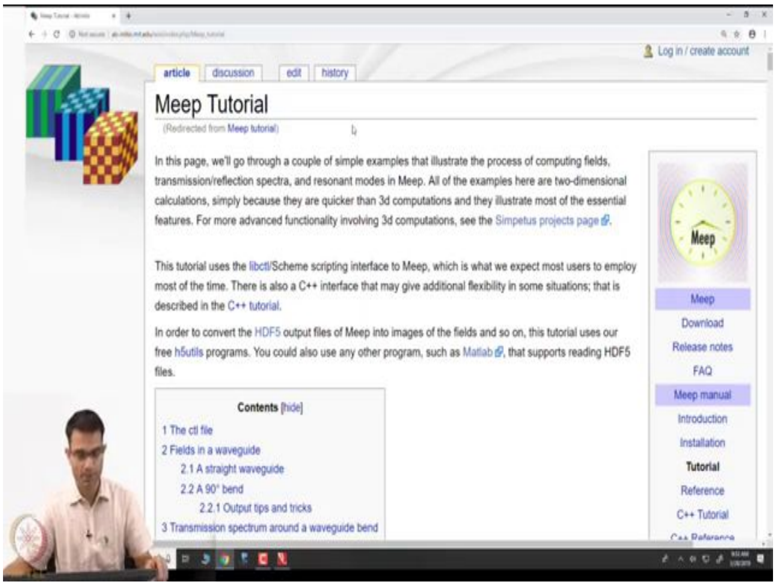

So it is called Meep the text is readable yeah ok.

(Refer Slide Time: 00:24)

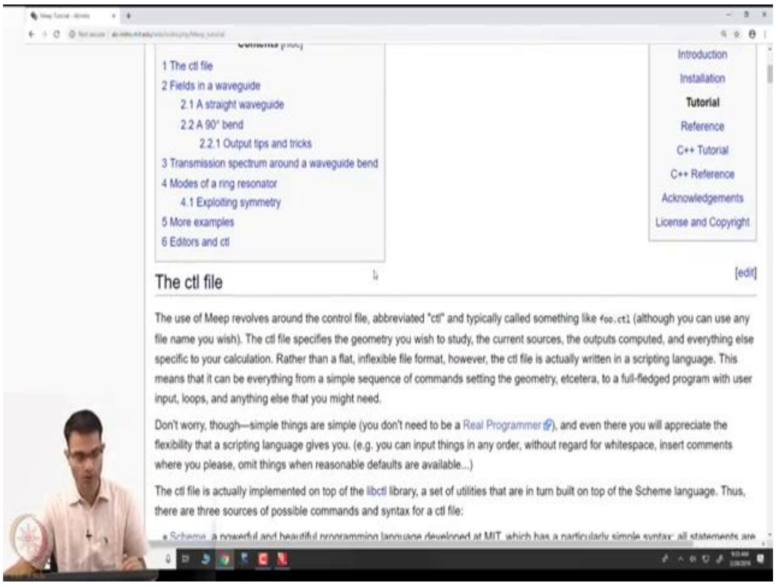

So, I would not read through all of this I will just working through this we will also get an idea of how to what are the steps in running an FDTD program ok.

(Refer Slide Time: 00:31)

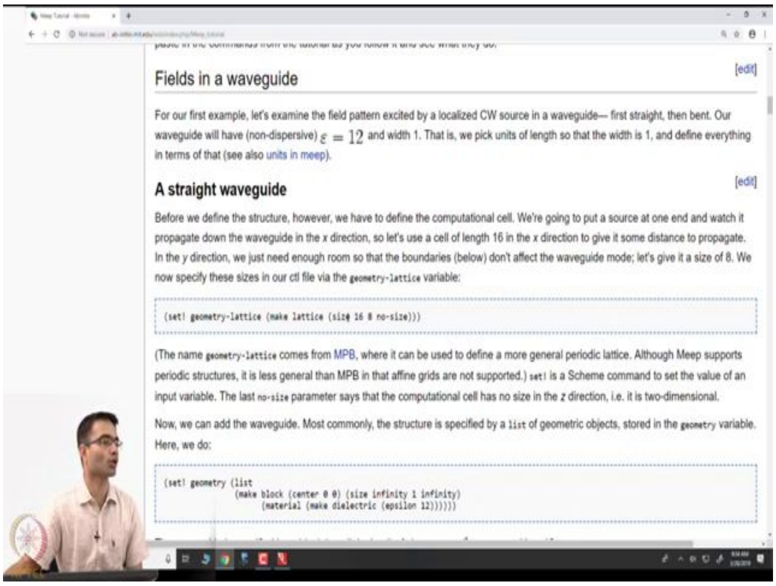

So, let us go to. So, the first example they have is for example, fields in a waveguide. So, imagine there is waveguide I want to send wave through very simple we know analytically in these cases we can even analytically calculate what the fields look like, but that is not our objective let us see what this simulation does.

So, the first line over here is. So, this is as I mentioned is written in a language called scheme which is a slightly unusual language. So, we would not get into syntax, but what is being set over here this size is 16 8 no size this is telling you the size of the computational domain. So, 16 units in x, 8 units in y and no size in z means its 2D simulation that is all that it means.

(Refer Slide Time: 01:19)

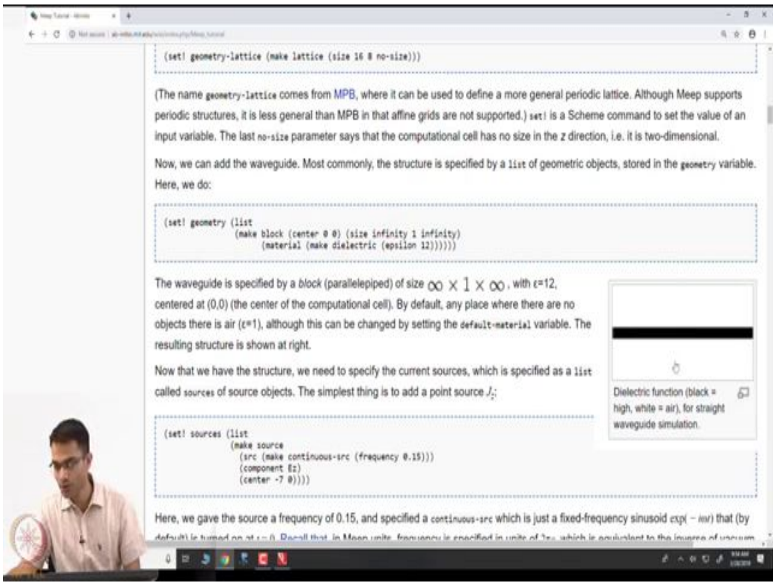

Next I have to make the object. What is the object? It is a waveguide. So, a waveguide as shown in this figure over here is this black strip over here. So, it is like a slab waveguide right. So, I have x is along the waveguide and then y is the other axis right. So, here the waveguide is being made. So, it is being made as a block of some ε. So, $\varepsilon = 12$ has been given ok.

The width is given as 1 unit and infinite in the other dimension so, fairly intuitive right. So, the width of this black strip is 1 it has epsilon equal to 12 and it extends forever of course, it extends forever because I have defined the computational domain about we have 16 units over here right. So, in practice it is going to be 16 units right.

(Refer Slide Time: 15:16)

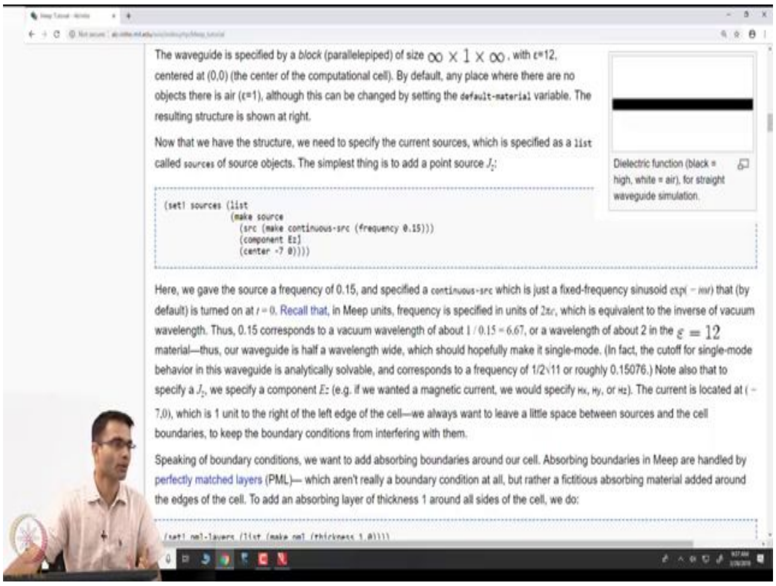

Now, by default everything else is vacuum that is what you would expect ok. Next comes a source I need to put a source to excite some field in the waveguide how will I get the wave in there. So, what have they done? They have made a $e^{j\omega t}$ kind of a field ok.

They have given some frequency in some normalised units 0.15 that is the frequency and they are saying that is E_z polarisation over here and they have given some centre for it ok. So, centre remember our size was 16 so, -8 to 8. So, at -7 that is at the almost at the left most part I am putting E_z polarised source and what kind of a source? Continuous wave it is not that source that starts in time and that decays in time continuous wave the dielectric is the waveguide this black strip over here is the waveguide why would not the wave go out ok? Good question. Why would not the wave go out? So, I am exciting it within this and I mean. So, this is something which is you should have studied in your undergraduate course.

Student: (Refer Time: 03:26).

Ha no. So, the field does leak out it is not the case that the field is strictly confined within the dielectric the field does leak out exact we will see it is not confined purely inside some of it does leak out. We do not have to think worry about how much is undergoing total internal reflection the Maxwell's equation will take care of it.

Student: Only that.

Yeah some will be some will guide it some will get will leak out ok. So, we have we have specified the source what else do we need to specify boundary condition right, if I do not specify boundary condition my simulation will be wrong.

(Refer Slide Time: 04:10)

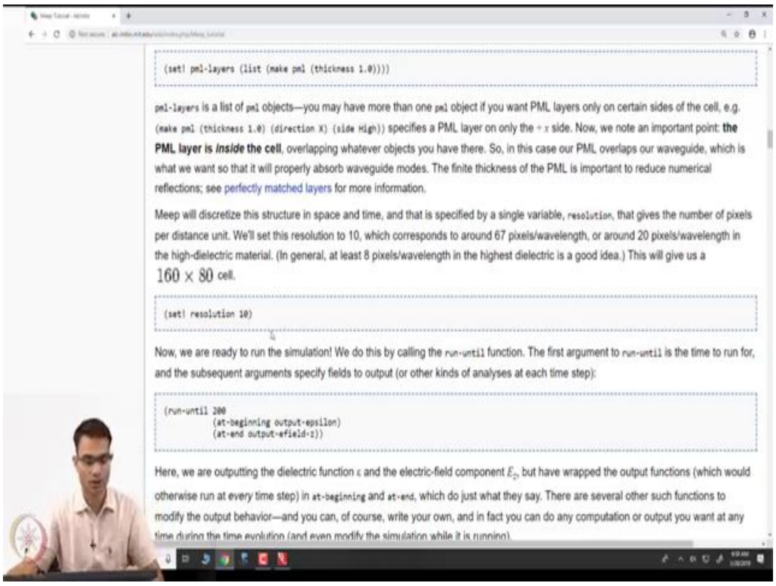

So, they have specified here PML this is just one line specification set a PML of thickness 1. So, by default what it will do you look at this computational domain over here it will put like a picture frame of thickness 1 all around ok.

That is why the source was put at the -7 not at -8 because you do not want source in the material you wanted just outside ok. So, the PML is from -8 to -7 and the source is at -7 ok. Now remember your Δx and Δt is yet to be specified right. So, here we set for example, this set resolution so, that setting Δx in some normalized way.

The Courant parameter is a default parameter that is already fixed in the simulation you can change it if you want, but the they give you a default which will be a stable number. So, fixing Δx is good enough Δt will get fixed according to the courant ok.

(Refer Slide Time: 05:15)

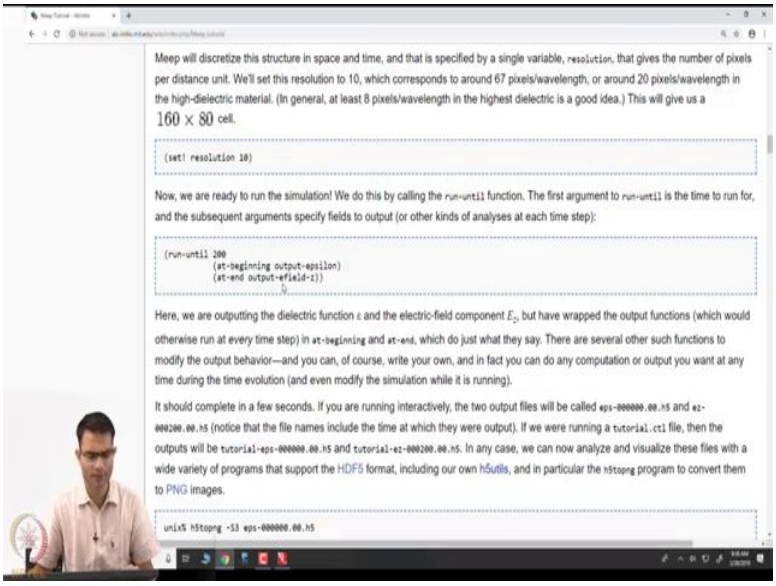

Then finally, you have to tell the simulation that how long do I wanted to run? So, you say 200 time units some more commands are given to output the fields and all of those things ok.

(Refer Slide Time: 05:27)

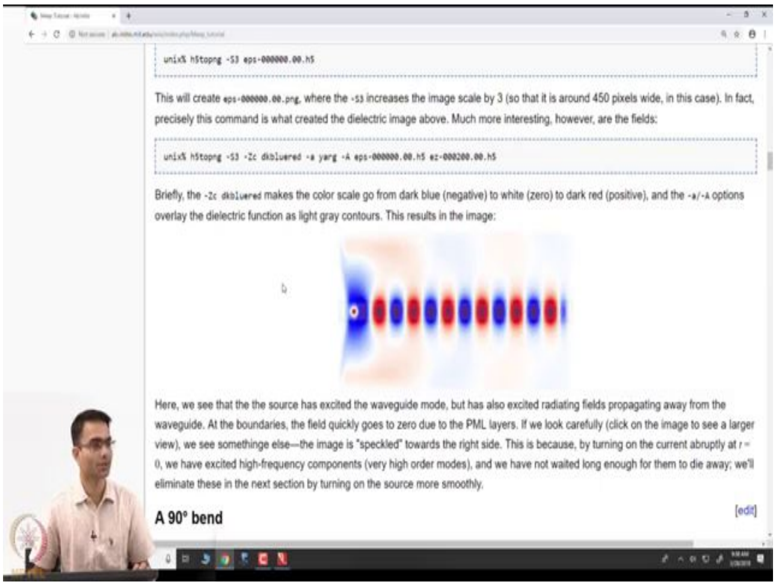

This is what the output of the simulation is at the end of 200 seconds this is what the fields look like right. So, the source was over here and looking at this can you say that there is definitely a wave guiding happening clearly right the field is very nicely confined to the waveguide, but not strictly some of it is leaking out no problem right it is not and this is not a lossy material, it is just $\varepsilon = 12$ there is no imaginary part. So, this field will go on for a very long time and distance right. But you can see its leaking out, next to the left hand side interface you can see there is something strange happening over here that is because there is a.

Student: PML.

PML over here so, its absorbing it almost perfectly, but I mean in practice some of it will not get absorbed right. So, they make a point here that because, they turned on the field abruptly right. So, the simulation started at t-0 and the next instance there was a $e^{j\omega t}$. So, we know that that as far as physics is concerned there is an abrupt turns on the field.

So, there will be some high frequency component right. So, that is what we say towards the right hand side over here the field looks a little what they called speckled you can see the field over here it has some the shading is slightly different from at the very beginning here is the very nice homogenous dark red here is a here there are little white dots inside the red. So, that they says coming from the high frequency components that came as a result of turning the field on abruptly.

Student: (Refer Time: 07:08).

A good question what is the red and blue? So, red will correspond to maximum positive blue will correspond to maximum negative right. So, it is like it is a wave that is travelling right. So, alternate maxima and minima in a wave that is what is happening.

Student: That thing in the.

That thing in the.

Student: (Refer Time: 07:27).Yeah, white it is 0; yeah white it is 0. So, that is what they called this dark blue red that is the you specify the legend for the colour. So, that is how they do. So, typically when you see a simulation of red and white then you should have assumed that one is maximum positive, one is maximum negative and white is 0 this dot that is the source.

Student: Dark.

The dark thing that actually that that is like the darkest red and this is the darkest blue that is how? I mean that is the shades of blue and red shades of not blue and red shades of dark blue and dark red. So, dark blue red is one legend blue red is another legend we can choose ok. So, people use different things for enhancing the contrast.

Student: Ok.

So, this is the example which you could have done analytically also right, but the power of the FDTD is that you can do things which are analytically not possible.

(Refer Slide Time: 08:25)

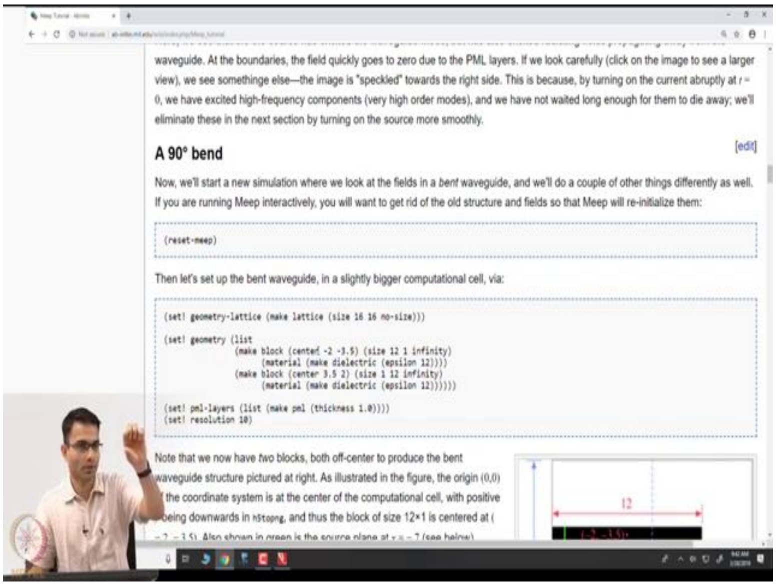

So, for example, they consider a 90 degree bend in the waveguide which you would not be solve analytically right. Now, what do you expect physically to happen I take a waveguide its guiding a wave I bend it by 90 degrees what would you expect will happen physically?

Student: That will leak it out.

Some of it will leak not it will its not like a pipe of water that I bend it and all the water goes out some of it will leak out that is what we will expect. So let us see what if our intuition is confirmed over here. Now same thing they say make the computational domain as before 16, 16 and you have to make two blocks; one like this and the one the 90 degree bend.

(Refer Slide Time: 09:03)

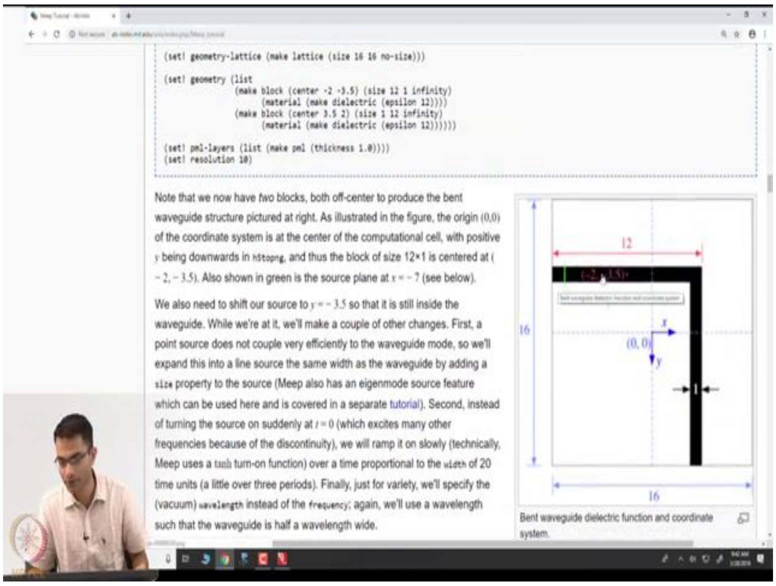

So, they have shown it over here. So, one this way and one this way and they have done it by specifying two blocks same material, but this coordinates are different PML resolution everything as before ok.

(Refer Slide Time: 09:17)

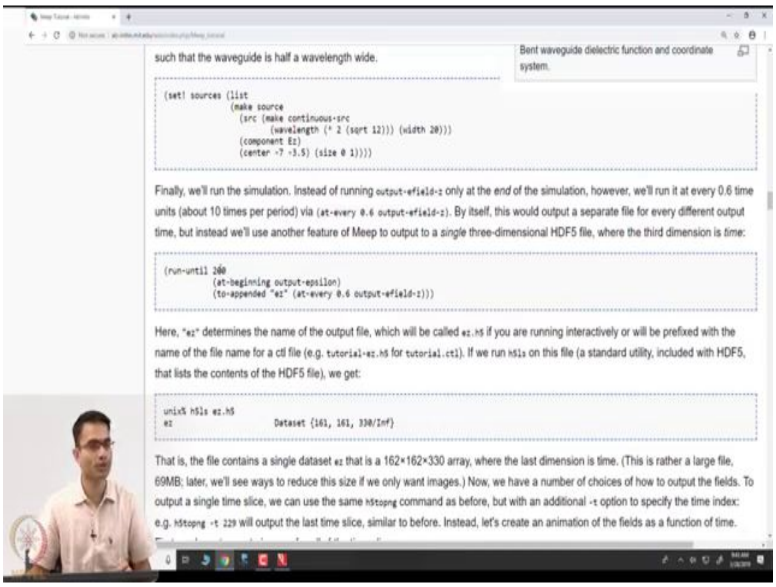

Now, to avoid that speckle phenomena what are they doing instead of turning the source on abruptly they are making it no they are still using a what are they will continuous source wavelength width so much yeah. So, what they are doing now is they are this instead of turning a source on suddenly at t=0, we ramp it slowly over a time proportional to the width of 20 time units. So, this width is equal to 20 time units is giving you the turn on time of the source and the wave length is given to you over here.

Yeah. So, technically Meep is implementing this turn on function using tan hyperbolic function and as I mentioned there is a open source we can go and look at the code and see all that information.

Student: Ok.

And as before run it for 200 time units. So, let us see what the output looks like. So, there is how the output looks like.

(Refer Slide Time: 10:26)

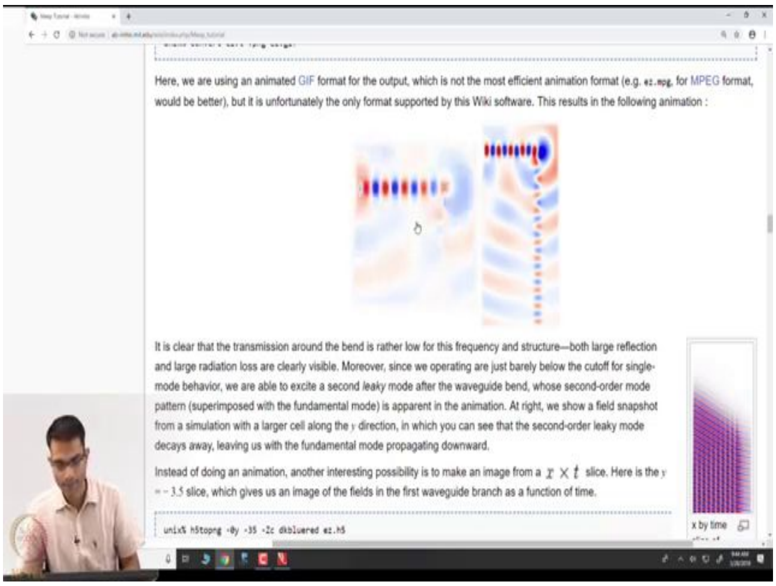

So, yeah this we can see this is the animation at every time instant you can record the field. So, you can see the field is going now it hits the bend and it begins to come into the lower part of the bend, but most of it is getting leaked out at that interface ok

So, that is another very cool aspect of FDTD. If I store the E_z at every time instant I can put it all together and make gif file. So, what they are showing over here is a gif file we will see the animation. So, then you can optimize various parameters if you want to minimize the leakage and all of that we would not go into that.

(Refer Slide Time: 11:13)

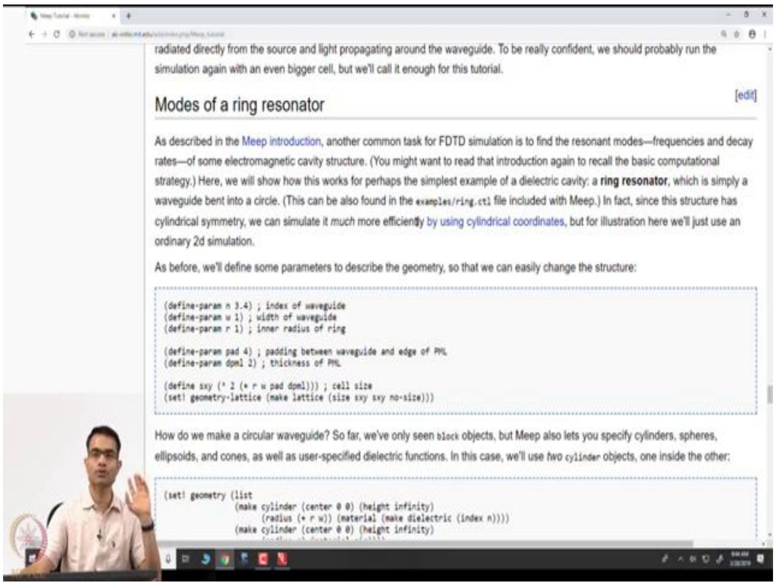

Then the final example I want to show you is modes of ring resonator ok. So, you all have heard the word resonator it means that its a structure which is selective in frequency it allows some frequency and it disregards the other frequencies. Now, supposing you want to design a resonator to work at some frequency how will FDTD will help us in that.

Student: Ok.

(Refer Slide Time: 11:43)

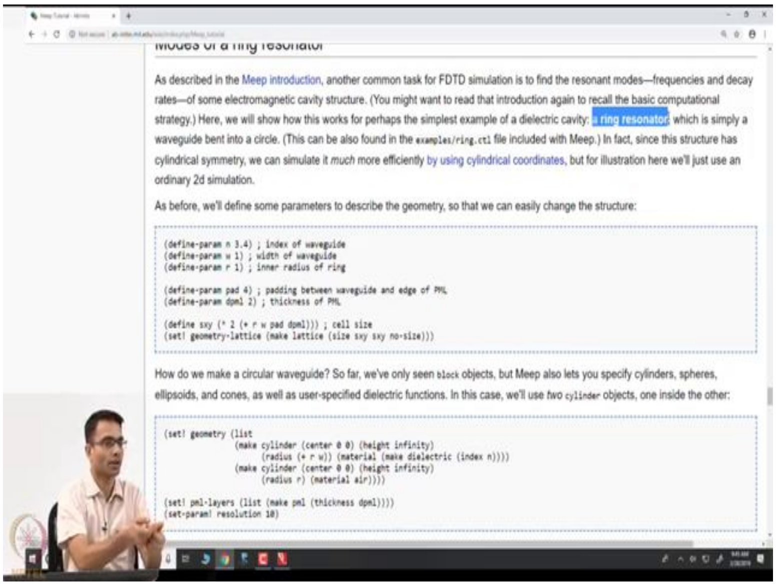

So, the let us look at the geometry over here. So, what have they done this they made a ring resonator ok. So, what is the ring resonator? Take the waveguide from before join it into a circle ok. So, we know that you know integer multiples of the distance around this are what will be the supported modes ok. So, here we want to confirm that intuition over here now hm. So, they have a you know the width of the waveguide radius all of these have been specified.

Yeah some, but $2\pi r = n\lambda$, but λ in the medium.

Student: Ok.

(Refer Slide Time: 12:18)

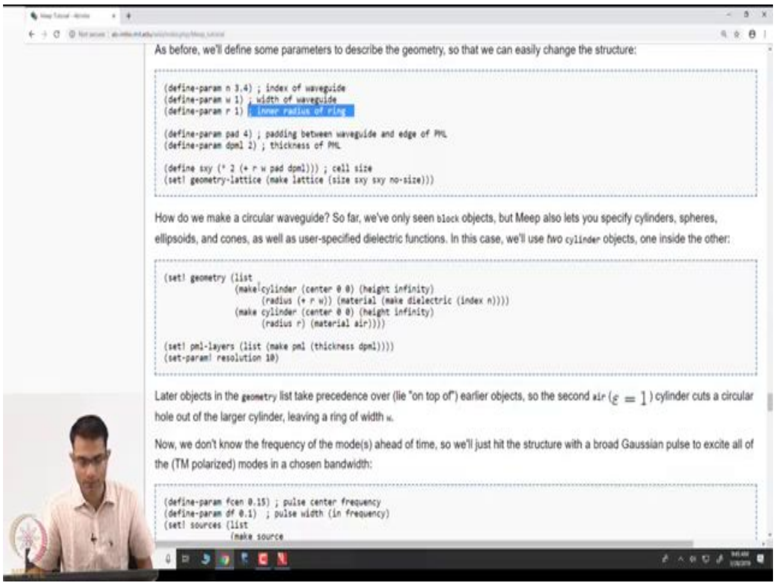

So, they making a cylinder over here all of this is what you would do to make the object ok. Now the question lies what should I if I excite this structure with the wrong frequency what will happen?

Student: Nothing will happen.

Nothing will happen right the fields will not decay I mean the field will not decay I mean the field will decay off it will not be a resonate phenomena, but once I designed a structure with some epsilon if the and if the geometry slightly complicated I would know what the resonate frequency is. So, how would I find that resonate frequency?

I can in FDTD I can specify a source after that what supposing my task someone gives me structure and says find out the resonate frequency of the structure.

Student: Excite.

Excited with all frequencies right; so, that is well all frequencies is technically impossible white band right. So, the step 1 is excite it with the white band ok. So, here they say they are making a pulse kind of a frequency. So, they are saying a centre of the pulse is 0.15 in normalised units and a df right and like a like we have discussed earlier they use a Gaussian source. So, they give a Gaussian source where they specify the centre frequency and the width.

Student: (Refer Time: 13:38).

Yeah so, they work these are physicist they work in a strange set of coordinate units where speed of light is 1 ok. So, everything is normalized very nicely to 1 right.

(Refer Slide Time: 13:53)

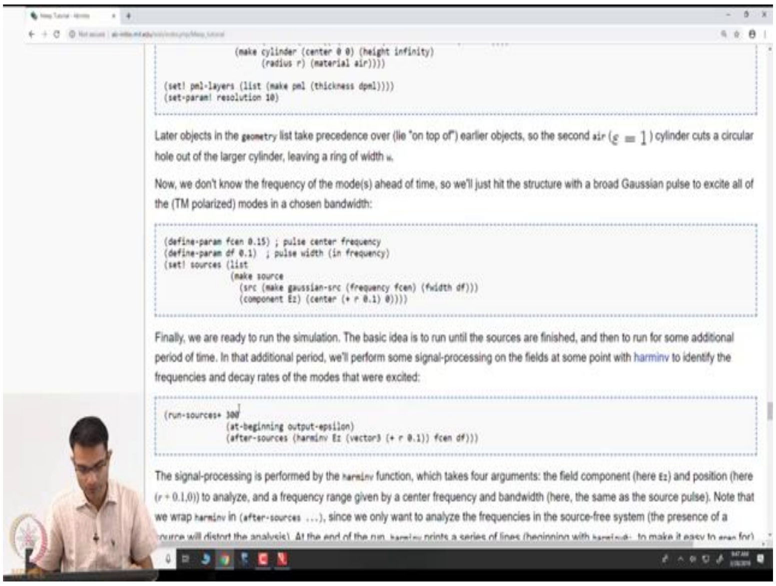

So, it takes a little getting used, but once you get to used it is very powerful everything is normal yeah speed of light is 1 wave length is 1. So, yeah whatever then you run it for some time 300 ok. So, now, let us imagine we have done the simulation. How will you interpret the output? Output will be fields that all times and in positions now what.

Student: (Refer Time: 14:20).

No, but we are exciting into the wideband, but I am not getting frequency information in my FDTSs output is what E as a function of space and time below frequency. So, what do I have to do?

Student: (Refer Time: 14:36).

That is yeah. So, he is saying go around the structure and find out where the nodes are that is a very crude way is there a smarter way. Find out the Fourier transform I have the time domain in the delta calculate the Fourier transforms that is the much smarter way it will give me not just the frequency peaks, but also how resonate they are how sharp they are right. So, these guys have a inbuilt function called harm inverse this is right this is like the inverse Fourier transforms.

So, the simulation itself says that run your simulation for the 300 time instance at the end of it after sources get over run the inverse Fourier transform at this position and some centre frequency is given to you ok.

(Refer Slide Time: 15:25)

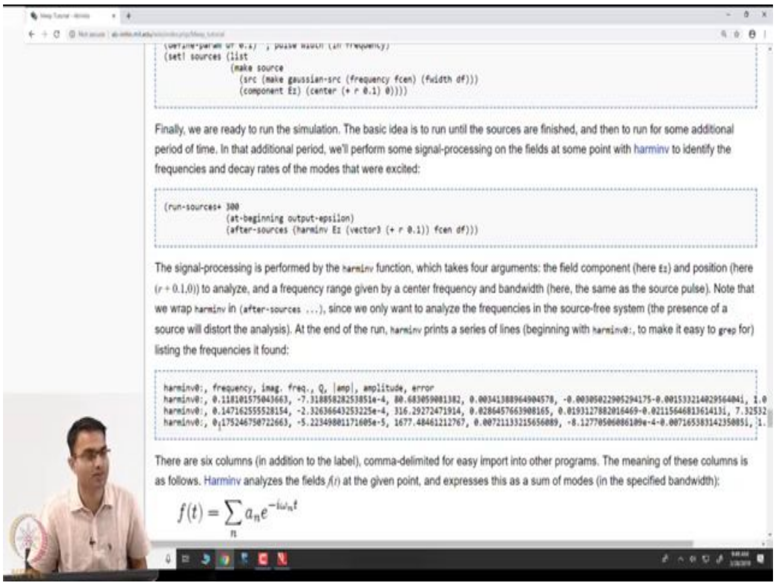

So, again this is something that they have coded and output comes in text. So, you can see that when they do this the peaks are extracted out. So, you get 3 peaks. So, that is the frequency and the Q the quality factor right that tells me if you have studied this before the quality tells me how good the resonance is.

So, the Q factor also comes out. So, for example, the first mode have the Q of 80 then 300 and then 1677. So, 1677 corresponds to the Q which lasts the longest ok.

(Refer Slide Time: 15:55)

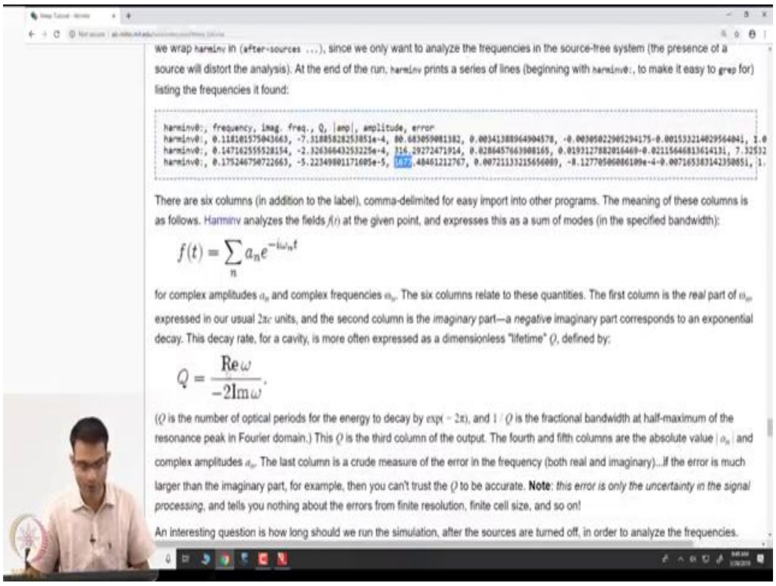

And so, there is some theory of a inverse Fourier transform, but you all know that.

(Refer Slide Time: 15:59)

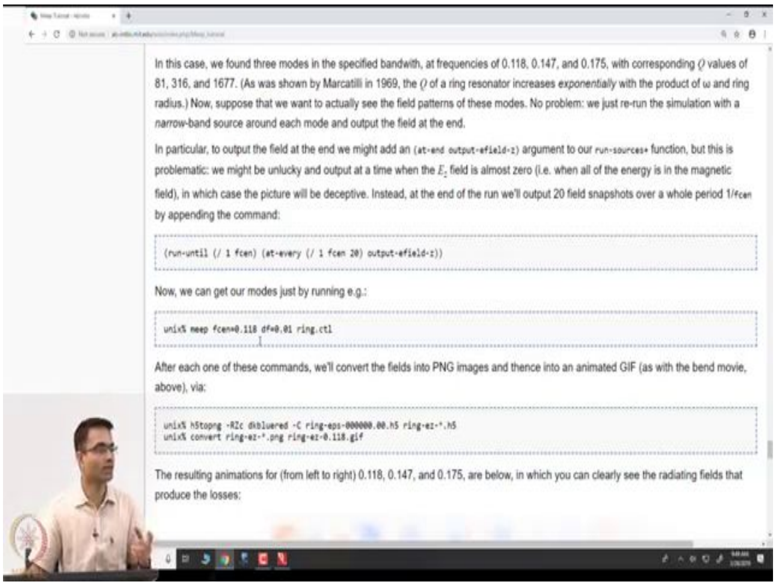

And so, the next thing I could do is pick each one of those frequencies and re run the simulation with the narrow band at those frequencies that will excite only that mode and not the others right.

(Refer Slide Time: 16:15)

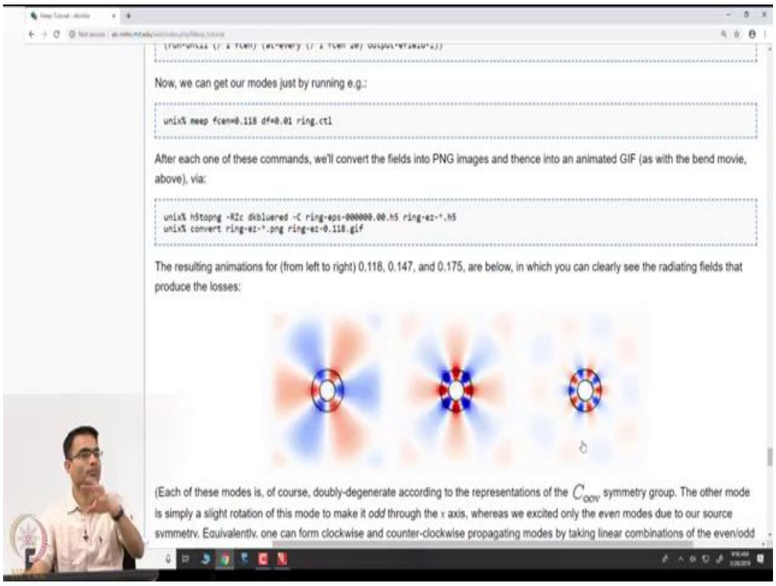

So, what they are showing here are those three modes. So, the first mode the second mode and the third mode and these are again gif files produced. So, this is you can see this had the lowest Q right.

(Refer Slide Time: 16:30)

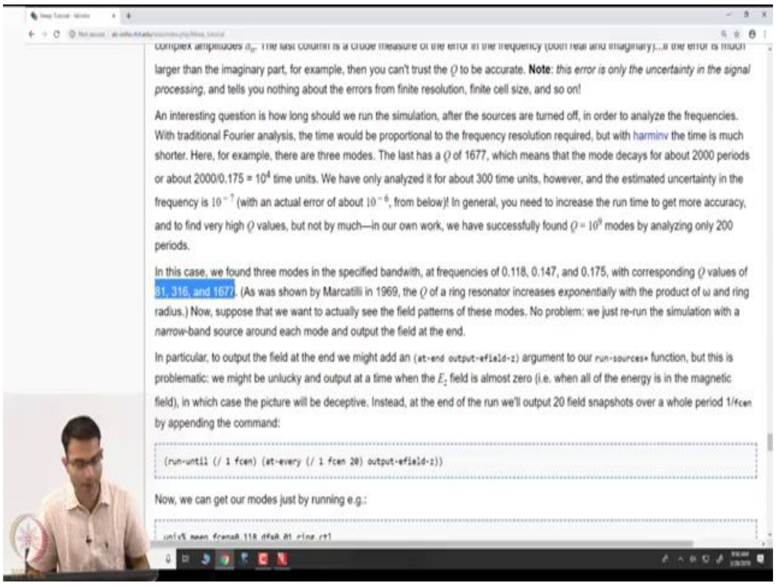

What were our Qs 80316 and about 1600. So, this is the lowest Q and that is you can see that a lot of the field is leaking out right outside the waveguide outside this resonator lot of field is there as the Q becomes larger and finely larger over here. Here the field is much more strongly concentrated inside the resonator.

Student: Right.

So, I have made a structure and I am able to identify these are the resonant modes by using simple FDTD and inverse Fourier transform signal processor ok.

(Refer Slide Time: 17:03)

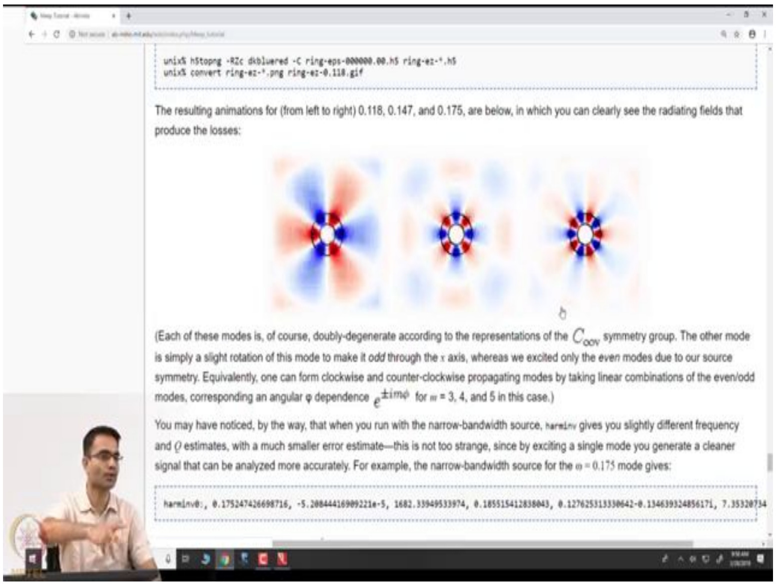

So, that is to sort of show you what should you say in practice how would you use FDTD and is very very powerful design any structure excite any fields right. And, they have lots of other solid examples for photonic crystals and other structures like that and its easiest to run it on a Linux machine. So, let brings to an end this demo work on FDTD.