JUD HOYT: We have some administrative things today. First of all, there are two handouts-- 34 and 35. We've got the lecture notes. And I'm handing out the solutions to homework number 4. And homework number 4 is going back. So you can pick up homework number 4. It's in the orange folder in the back there.

And some reminders or announcements-- these will appear on the website today. So just to remind you, your written reports are due here on Thursday, December 9, in class. I have come up with a schedule for the presenters, for the people doing oral presentations, that's shown here.

We have two lecture periods that we've dedicated to oral presentations-- the last two lectures of this term. So on December 7, we have four speakers. And these are the four speakers I have chosen for December 7. And your presentation time should be-- of your actual presentation-- should be about 16 minutes.

And then beyond that, we'll leave three minutes for questions and answers period. So it's a short Q&A. And that's for each talk. So I'll make each talk about 20 minutes. And we'll have four people in one period. So that should work out fine.

And the same thing on December 9. These are the four speakers that I have chosen for December 9 and pretty much approved all the topics. If you're not sure your topics are approved, go on the website. And I have an indication whether it's been approved or changed.

If you need help-- you're going to be graded. For those of you who are doing the oral presentations, you're graded on two things-- same as the written, but a slightly different-- is you're graded on the presentation itself, which includes your oral presentation as well as your slides. So you need to provide handouts, the same way that I give out handouts every lecture-- not as long because you're only speaking for 16 minutes.

But if you need any help making Xerox copies-- you don't have access to a copier-- you don't need to pay to make those Xerox copies. My assistant, if you send her the PDF file, or the PowerPoint file preferably, or if you bring her a hard copy of your presentation, she'll make enough Xerox copies for the rest of the class.

You need to give her some time, though, to do that. So make sure you get her that the day before, so she has time. And her hours are 8:00 AM to 6:00 PM-- 2:00 PM, 8:00 to 2:00 so. Make sure you get her that the day before if you want help making Xerox copies.

And the second component of your grade, of course, is on the technical content. So one part of your grade is on the presentation. And that includes the quality of your graphics and the quality of your handouts. And the second part is on the actual technical content, the depth, and to see that you read adequate number of references and that you understand the material and that you can present it.

Same thing for written reports-- the only difference is you won't be graded on your speaking because you won't be speaking. But for the written report, you'll be graded both on technical content, as well as on the presentation. How nicely did you lay out the report? How well is it sectioned? Are the references clear and easy to read? Do you have good graphs? Or are the graphs impossible to read? Things like that-- so both aspects will be part of it. And all of this information will be posted on the web. But this gives you an idea of when you need to prepare yourself.

So yeah, this is where we are right now on November 23. The Thanksgiving holidays coming up on Thursday. We won't have a class Thursday. Today, we're finishing up chapter 10 on etching. And then we'll move on. We have two more formal lectures that I'll be giving. They'll be one on silicides and contacts and novel gate materials, and then on growth and processing of strained silicon and silicon germanium on December 2.

And then you folks will start speaking on the 7th and finish up on the 9th. And your written reports, again, are due in class on that final Thursday, December 9. OK. Anybody have any questions or anything about the final report or the schedule? Send me an email if you have any questions about it.

OK. So let's go on to today's notes. Each chapter now, you notice we're going through a little more rapidly as we finish up. This will be the last formal chapter of reading that you'll have in this class. There'll be a little bit on silicides you'll have to read. And then the last lecture, there is nothing written up on it. So you just come to lecture.

So hopefully, you're finishing up reading chapter 10 on etching. What I'm showing here on the first slide of handout number 34 is kind of a summary of what we talked about last time, as far as plasma etching and the various mechanisms. There are certain characteristics which we said were important-- for example, the pressure of the chamber in which you're etching, the energy of the species. And these end up affecting things like selectivity and isotropy.

At the top of the list of the type of etching are the most physical etching processes here, like sputter etching or ion beam milling. It's pretty much purely physical. There's not a lot of chemical reactions going on. If you go to the very bottom type of etching, which is wet chemical etching, it's entirely wet. There is no physical bombardment.

And everything in between has a certain component of both chemical processes and physical processes going on. And these arrows point to the direction in which the quantity increases. So anisotropy, which is-- so the etch will be more vertical than it is horizontal-- increases as you go from wet chemical, which is very isotropic-- or if you go to sputter etching, it's very anisotropic. So anisotropy increases in this direction.

Selectivity is very poor, or very low, for sputter etching. You can sputter anything. And then the sputter rates are not very different among the elements. Selectivity increases going down. So reactive ion etching is less selective. Plasma etching is more selective. And wet chemical etching tends to be the most selective. Selectivity increases in this direction.

In general, the energy of the process increases going from the bottom to the top. So there's more energetics involved in sputter etching. The particles can come down with hundreds of electron volts and knock off atoms on the surface. So these processes tend to be more energetic. And the pressure of the system, here, tends to be lowest here, and then increasing as you go from here-- this type of sputter etching-- to, say, plasma etching.

So that's just what we talked about last time by way of mechanisms. Today, I wanted to give some specific examples. And I want to emphasize gate etching issues. And I'll talk a little bit about modeling, the mechanics of how people model etch processes.

On slide 2, this is a chart-- or it's half of a chart that I took directly from chapter 10 in your text. And what it shows is a series of common materials, and some typical etchants that are used, and then some comments about each of these etchants.

For example, a very common way of etching polysilicon in a generic way is to use SF6 or CF4. What are the characteristics of that type of etch? Well, it tends to be isotropic or near isotropic and not very good selectivity over SiO2. So hit once you hit the SiO2, you continue etching. So this is a general purpose etch you might use if you're trying to strip off some poly off the back of the wafer or something-- probably not something you might use if you need anisotropy or something like that.

A more anisotropic etch, but also not very selective to SiO2, would be a mixture of CF4 and hydrogen, or ChF3, some of these freon types. You can mix CF4 with oxygen. I think we showed last time an example of the chemistry of this. It's isotropic generally, reasonably isotropic. But it's more selective to SiO2. Again, a lot of times, you're etching poly on top of oxide or some dielectric. And so the selectivity of the etch, of the polysilicon etch, to the SiO2 is very important.

The most commonly used etchant gases today for polysilicon-- and we'll talk about this when we talk about gate etching-- are HBr and chlorine or some combination of HBr, chlorine, and oxygen. These tend to be very anisotropic. So you can get very well-defined sidewalls. And most of them are reasonably selective compared to SiO2. With HBr and O2 combinations, you can get up to 100 to 1 selectivity, depending on the etcher.

If you want to etch single crystal silicon, you use exactly the same etchants you used for polycrystalline silicon. In general, there's not a big difference between the two. If you need to etch oxide, you can etch SF6. Interesting, any of the fluorinated species will help you with etching oxide.

Notice that SF6 was also listed as etching poly. So that means it's not a very selective etch. And isotropic etches include a variety of the freons, like C2F6 and C3F8. They are somewhat selective over silicon. When do we etch oxide? In a Mohs process, for example, we're often etching oxide when we're opening up contact holes. So we have an interlayer dielectric of some sort, maybe 1,000 or 2,000 angstroms we need to get through it. With reasonably good control of the sidewalls, we also need to stop on the silicon. You don't want to etch away your source strains. So anisotropy and selectivity are also very important.

Silicon nitride-- there are some isotropic etches, like CF4 and O2. An anisotropic CF4 and H2-- this one is selective over silicon. So it'll stop on silicon, but not over SiO2. And sometimes, you have stacks consisting of oxide and nitride. And you need to etch one of the two selectivity-- selectively. So it gets a little tricky. You need to pick the etchant and the reactor carefully.

What are some other materials? So those are the more common ones. Maybe a little less common, but things that you do have to etch often in plasmas are, for example, aluminum wiring. Or aluminum contacts need to be etched. It's very common to use chlorinated species, like chlorine or CHCl3, or a mixture of chlorine and nitrogen. Chlorine itself can be somewhat isotropic. And some of these other gases are more anisotropic.

Tungsten can be etched in SF6 or chlorine. So you notice a lot of the metals-- these are all metals-- are attacked by chlorine, which isn't too unusual if you know a common wet metal etchant is hydrogen chloride, HCL. So again, etching titanium or ti-nitride, these are very common etchants-- chlorine or mixtures of chlorine with hydrocarbons.

Photoresist. An important thing, we need to be able to not etch or pattern so much, as you typically want to strip it. So an etching step would include an oxygen plasma, which reacts with the hydrocarbons and removes the photoresist. And the nice thing about O2 plasma is it's very selective. It doesn't etch too much of any-- it doesn't etch nitride, doesn't etch oxide. It will etch a little bit of silicon by oxidation. That is, it'll oxidize a little bit of the surface silicon.

And then if you happen to do an HF dip, that'll get removed. So it has a small amount of silicon removal by an oxidation process. So again, on page three, that's the second half of the chart I took directly from chapter 10 in your textbook of some common etchants. Now, I want to go here on slide 4. I wanted to go through some etching challenges in moss frontend processing from the last five years or so. And I'll talk about a couple of them.

One is, for example, we've talked about in the past DRAM, or dynamic random access memory. There are different ways of storing the charge in that memory. You need to make a capacitor of some sort on which you store charge. And that is your bit. Whether there's charge there or not charged there tells you whether you have a one or a zero.

This capacitor you might think, well, let's just put a capacitor on the substrate. We all know how to make an MLS capacitor. But the problem is to make very high density DRAMs. You can't just make a capacitor flat on the silicon substrate. You can, but then it takes up too much area. So people have all different types of ways. People are thinking of ways of stacking the capacitor on top of the transistor. That's one method, of stacked capacitor topology.

Another method is take the capacitor and bury it down in the substrate somewhere so you can get a smaller area for your DRAM cell. You can pack more memory on a given chip, and then you'll have much higher performance. So this is an example where you might need to etch down into a silicon substrate to make a trench prior to making a capacitor.

This is an example where you have etched a trench that looks like this, lined it with oxide, and then filled it up with polysilicon to make one of the plates of the electrode. The other plate of the electrode could be the surrounding silicon material. And I took this out of an IBM journal of research and development just off the web if you want to look at that in more detail.

So we need to make very high aspect ratio etches for these trenches. And we also need to care about the damage that we do to the silicon because it's going to be part of a capacitor. So we want it to have smooth surfaces and things so that we don't break down the oxide and we get a good electrical performance of the capacitor. In fact, let's see. Let's take a look here at page five. I took this from the 2003 ITRS table 73.

This is just for DRAM trench capacitors. What are some of the technology requirements we're looking at in the long term here, going from 2010 all the way out to 2020 or 2018? What do people expect they might need? Well, here there are some indications here of characteristics. They have things called the trench structure, how it will be shaped-- maybe more like an upside down bottle-- the circumference of the trench, the trench surface roughening factor, the effective oxide thickness that people will be using in that trench capacitor.

Just to give you an idea, in nanometers here, we're talking about several nanometers, maybe two nanometers, scaling even thinner and thinner. The trench depth in microns. This would be a 30 for a 35 femto farad capacitor, just as an example. How deep will we have to etch, reasonably deep-- somewhere between 5 to 7 microns is what people are talking about doing.

Aspect ratio. So remember, we talked about aspect ratio as being an important attribute of our capabilities for etching. Well, so here is the trench depth divided by its width. We're talking about very high aspect ratios that are increasing. And they're also quite large. In 2010, which is a ways off, we're reaching aspect ratios on the order of 100, and beyond that, going even higher.

And you notice beyond 100 the cells are all marked red, indicating in that range, people don't really know yet how to etch such deep aspect ratio, 100 to 1. So new technology is going to have to be developed in order to do that. They have some indications here also of what the upper electrode might be made of. It seems to be metal. The dielectric material, people are talking about using high Ks.

Again, you're trying to store a fair amount of charge in a small space. You can do that with a higher dielectric constant. The bottom electrode could be a combination of silicon and metals, as well. So just to give you an idea of what kind of trenches one needs to etch. And this increase in the aspect ratio is pretty amazing in terms of what technology requirements are going to be imposed on advanced etching techniques.

OK, so that's one example. If I move on to slide 6, we not only have to etch silicon into the single crystal material, but we also have to etch silicon when we're etching the gate. And in fact, etching the gate is becoming more and more complicated. This is an example on slide 6 of etching what we call a compound gate, or a gate that already has a gate stack, or already has several materials that it consists of.

And I apologize this is not a very good SEM. I took this off the website, again, from that IBM article. And what it is on the very top, there is photoresist that has been etched-- that has been patterned, I should say. Then there is a layer of tungsten silicide that has been etched. Then there's a layer below it of polycrystalline silicon that's been etched down to, say, the oxide. It's down to the gate oxide.

And when presented with something like this, you have to think about using two different types. After you pattern your photoresist and you develop it and you're ready to etch it, underneath you have two different types of layers. You need to etch a silicide layer in this example, which was preexisting as a uniform layer. It has to be etched. And then you need to stop on the silicon and then switch gases to etch the polysilicon underneath.

Now, this gets a little tricky. So you have two different etch gases. That doesn't sound too bad. You have an edge gas to etch the tungsten silicide and another series of gases you use to etch the polysilicon. One problem, though, you can come across, for example, tungsten silicides may not be super pure. It may have some oxides incorporated into it.

And these oxides won't etch if the etchant is highly selective to oxide. So if you're using a single etch gas, then it's very selective to oxide, say is what you might want to do for the poly. That same etch gas might not do very well in the tungsten silicide. And what happens when you use an etching gas that is very selective to oxide and you don't do a breakthrough, if you have any native oxide on top of that structure and you don't break through it, you'll end up with something called RIE grass.

And it's called grass because it looks like grass in the SEM. On the next slide, I'll show you a picture of grass on slide 7. This is an example of an extreme case of RIE grass. And you have a flat substrate down to which you have etched. And sometimes it's called micro masking. What happened, an etchant was used to etch this gate, whatever this material is. An etchant was used that had a very high selectivity of oxide.

And there were apparently small oxide inclusions at the surface of the wafer when the etch was started that didn't get fully removed. And so those small oxide areas will act like little micro masks. They were not intentionally put there, but there's just a little bit of oxide here and there. So a breakthrough etch wasn't used. As a result, if the etching you're using is highly selective to oxide, it'll never etch it. And instead, you end up with little posts which look like little blades of grass.

The solution to this is, well, you can say, well, don't use an etch that's so selective to oxide. But if you want to stop on SiO2, you wouldn't want to do that. The solution is the very first part of your etch step should be a breakthrough etch where you make sure you break through all the native oxide and then you continue on down. So it gets a little tricky when you start etching multiple materials or when you have to think about the fact that there may be materials on the wafer you don't know about.

Native silicon dioxide may be there, even if you didn't put it there. It grows there automatically on the wafer. So you have to remember that you do get it off the surface before you use a highly selective etch. Slide 8, what I've done is I've listed some what I call etching challenges in MLS frontend processing as they apply to gate etching and logic devices in particular.

This on the left-hand side is a TEM micrograph. I think I showed this the very beginning and several times throughout the course just to give you an idea of what a device five years ago or so might have looked like. This is a fairly large technology by today's standards, but I think it's a quarter micron or something like that channel-link technology.

Here is a polycrystalline silicon gate. There's a thin layer underneath, which is the gate oxide separates the gate from the channel. And on top there is a layer, which is silicide. Now, that was put there afterwards, in this case. So the gate was first etched. And then metal was put down and then was reacted with the gate and the source drain in a silicide or a salicide process. We'll talk about that next lecture, when we talk about silicide.

So you didn't have to etch through this. This was reacted later on. But what do we care about? Well, in this picture, first we need to control the gate length or the critical dimension, sometimes called the CD. The gate length is defined from this point to this point. So this width or length, so to speak, of the polysilicon needs to be controlled exactly on all areas of the chip and all areas of the wafer.

The sidewall profile. What do I mean by that? Well, this profile, going from the very top down here, has to be controlled. You usually want it vertical. You don't usually want it splayed out like this. You don't want it retrograded like this. You don't want it straight with a little slit at the bottom. You want it perfectly vertical and stop. That sounds kind of trivial, but it's not easy, for various reasons.

Selectivity. We care about, as we're etching this polysilicon, we don't want to etch away the resist that would be masking it on top or whatever mask material we're using, if we're using a hard mask of SiO2. So selectivity to the etch material, to the masking material, is very important. That's one type of selectivity.

Selectivity to the gate dielectric. Down here in this is a dielectric material, probably a high K. That's probably 1 to 2 nanometers thick. We need to stop instantly, immediately, as soon as we hit it, automatically, and not go through it, and not over etch too much. So this is a real issue, getting chemical selectivity at the same time as getting good sidewall profile. I just mentioned the fact that Native oxide is often present on the polysilicon. We have to use a breakthrough etch as the initial step of etching. Otherwise we'll have all kinds of grass problems.

There can be impact of different dopants. If the gate is doped prior to etching, which it can be if you put down in situ doped polysilicon, the etch rate for boron or heavily phosphorus doped poly may be different from-- and they may have different etch characteristics than that of undoped. So you need to take that into account.

And we also mentioned at the very beginning of the last lecture, gate dielectric damage due to the antenna effect. What is that was you have ions coming down, being collected by the gate. Where do they go? Well, they don't go through the field oxide because that's too thick. They get shuttled down the gate and try to-- the current tends to go through the gate oxide, but to be careful that you don't put too much current through your gate oxide and damage it. So those are some basic issues when it comes to poly gate etching.

Here's an here's an example on slide 9 of a gate etcher. It's a little bit older technology now, but at the time it was designed by a company called Lamb to be dedicated to gate etching only. It's called the TCP 9400. It's relatively low pressure, low bias, high density plasma, has a mechanical clamp check. Remember, we talked about the importance of controlling the temperature of the etch. To cool the wafer, it's got helium on the backside.

This is an example of a wafer that's frontend processing only, so no metals are allowed to go into the tool. And you can get an etch rate of what's called the main etch. It's pretty fast. In fact, higher than you probably need in most etches today. The height of the poly today is probably only 1,000 angstroms. And this goes at 320 nanometers per minute. So quite efficient etching for poly gates. And it can give you the kind of profile shape and selectivity that you need.

If we're going to slide 10, this is just an example of a recipe. If we were to use that Lamb, TCP 9400 gate etcher-- I took this off the Stanford website. This is an actual recipe that they have published showing the individual steps that you would need to use, just to give you an idea of some of the complexity of this etch process. Now, each step-- they're numbered here 1 through 9-- corresponds to a different time sequence, a different step in the time sequence for the etch.

So the very first step labeled number 1, you're adjusting the gap which is a parameter inside the tool. The pressure is being stabilized at 13 millitorr automatically by the computer. And you're setting up to flow some of your initial gases. You're flowing C2F6. And once you've stabilized after 20 seconds, you go on to the next step. OK, so the next step is step number 2. You're clamping the wafer mechanically to the chuck to get good control of the temperature.

And you have certain characteristics in that step. So that lasts about 20 seconds. Step 3 is the first etch. Remember we mentioned when you're using a gas that's highly selective, you need to first start out with a breakthrough etch to get through the native oxide. Here you're etching for about ten seconds. You etch just based on time. You're flowing C2F6 all together with-- I guess that's the only thing you're flowing.

And you've got RF power on the top and bottom, 250 and 40 watts. And your pressure is 13 millitorr. So this should break through any native oxide for about ten seconds. And then the next step is 4. You need to stabilize the gas flows before the main etch. The main etch is the etch that's going to etch the polysilicon. So you turn on different gases. You turn off the C2F6. You turn on the chlorine, a mixture of chlorine and oxygen.

That's what we call the main etch that's going to do most of the work in getting you through the polycrystalline silicon. So you're just stabilizing at this point. You notice there's no power on the top and bottom electrodes. So you're not actually-- the plasma is not on. So you're not lighting a plasma. You're just stabilizing the flows through the chamber. And that takes a while.

And this condition in this particular case, stabilization step is about 20 seconds. Just gives you time to get the old gases purged out from the prior etch, which was the breakthrough etch, and get the new ones purged in and establish your pressure of 10 millitorr. Then you actually do the main etch. So here you turn the plasma on by turning the power onto the top and bottom electrodes. So the plasma is now lit, and you're actually etching here for a time of, in this particular example, 60 seconds. So you etch by time. So it's a 60-second etch.

And then you move to the last stabilization now. And you're stabilizing to get ready for what's called the over etch. So there's a breakthrough etch to get through the native oxide. There's a main etch, which does most of the work. Now, the over etch is the amount of time you need, remember, where you've pretty much hit the bottom, but you may have stringers or other things. So the over etch typically has to have even greater selectivity than the main etch.

And so the over etch in this example, you notice they removed the chlorine. So the chlorine flow, which was 40 CCM, has been removed. And you just have HBR and O2. This gives you a much higher selectivity. Maybe not as good of an etch rate, but it does have very high selectivity so that you will not break through that oxide. And then you do the over etch in this step where you light the plasma for 30 seconds using just a mixture of HBR and O2.

You then unclamp the wafer. The plasma goes off. After those 30 seconds, you unclamp the wafer. And then you end and you're ready to unload. So it's all computer controlled. In this example of just etching polysilicon, you have three types of etches a breakthrough a main and an over etch and stabilization steps in between. And this is an example.

And then you come out and you look at your etch rate and your sidewall profile and you see how it looks. Most etches, as we said, end up leaving a sidewall polymer that's either like a hydrocarbon based, or it might be a glass based, like an oxide-based material. That is what's responsible. Remember, it's the inhibitor layer that keeps you from etching sideways. That's what gives you such perfectly vertical side walls. You must have that inhibitor layer.

The question is, what do you do with that inhibitor layer? Now that it's on your wafer, every single gate has some goop, whatever it might be, sidewall junk, polymer, whatever you want to call it. You need to do a post-etch cleanup process after that because that inhibitor layer is there. If you go into a TEM or something, you can see it, what you'll see is-- or maybe even an SEM.

You'll see your polysilicon gate, and you'll see some material maybe that looks like this, some maybe amorphous-looking material. Very thin. It might only be 20 angstroms, something like that. Very thin. Thickness may vary from top to bottom. Whatever. But that's an inhibitor layer. That is what gave you that sidewall. But typically it's removed by either an H up dip or H combined with piranha. If it's some kind of hydrocarbon, you need the extra piranha.

These are typically not things you want to leave on there. They can introduce contaminants in subsequent steps. So you have to do a cleanup process. OK, so that's just an example of using a dedicated gate etcher, how you would etch a polysilicon gate. Let's just talk about a couple of the characteristics people often see. And if you do gate etching, this is what you will be up against.

Here's an example of what's called notching at the bottom of the polysilicon gate. It's probably a little bit hard to see. This is not the greatest micrograph. These are SEM cross-sections. These line widths from here to here is normally 0.2 microns from this point to that point. So it's a polysilicon line. And this is using sort of a standard flow rate. And here's where they modified the flow rate, making a little bit higher.

And if you notice, the gate looks sort of straight, and then it goes down. It kind of goes in at the bottom like a little notch. This is considered to be undesirable in many cases. You typically want the gate sidewall to go straight down. Why would you not want the gate to be notched? Well, there are some examples of what it might do. Let's say your gate did go in like this. So you have it vertical. And then it kind of-- I'm going to make it worse. I'll make the notch look really bad. So it might look like that down here and something like that.

Well, the problem with notching like that is when you go to do your iron implant, depending on what angle-- I'm going to implant the source and drain. Depending on what angle and how bad the notching is, your ion implant might end up looking like this. So there might be a little region on either side where the source and drain extensions don't actually get implanted and where there is no gate.

Now you say, well, I just do a high enough temperature anneal, and then this arsenic will move. And it'll probably overlap. You'll probably be OK. But what if it doesn't? You can't go on probably. And if that notch is not well controlled, you can end up with all kinds of high series resistance in that region. So people typically don't want a notch because it ends up affecting what happens to the ion implant, because remember, this poly that you're forming, this gate is the mask for your next step, which is going to be implant the source and drain extensions.

So this is an example of how they got rid of the notch. This notch is not as bad or doesn't look as dramatic as the one I drew on the board. They got rid of it by just changing a few things. They changed the pressure and the flow rate. And they're able to get a more vertical type of sidewall. So if you have to do polysilicon gates, you'll find yourself twiddling a lot of the knobs-- the pressures, the flow rates, and things like that-- until you get a shape that you're happy with.

Here on slide 12 is another article I took out of literature. It's an older paper. It's from *Applied Physics Letters.* And what they're making a point here has to do with something called gas residence time in the plasma etcher. This residence time concept is also important if you're growing thin films. So it also is important in reactors and things. What it is has to do is the amount of time here in seconds or fraction of a second.

Effectively, that gas, the volume of gas, occupies the plasma etcher itself. And you can see that residence time increases. So it goes up with each one of these lines as they go to higher and higher pressure, or it decreases as they go to lower flows-- as I go to lower pressures. And as I increase the flow rate, it goes down. So you're flowing through faster.

Now, why would we care about the gas residence time? Well, we do care about it because the amount of time the fresh gas is sort of resident has to do with the amount of time the gas is available to actually do etching. And if you have a very long residence time, byproducts can develop. Actually, after all, these reactions are chemical reactions. They evolve byproducts. Byproducts can then sit and develop there for a period of fractions of a second. And those byproducts can then end up affecting the etch or the sidewall polymer formation or things like that.

So residence time ends up being an important parameter. The way people control it is, well, they either up the flow rate, the total gas flow rate, or they change the pressure in some ways in the tool. Here's just an example from that same paper, an example of what happens as a function of the silicon etch rate as a function of flow rate. So here is etch rate in nanometers per minute. And this is a particular silicon etch. What they're using is chlorine, very common.

And this was the so-called their present etching technique they had using a high gas flow rate and a conventional etching technique. So it also, when you change the residence time, increasing it is going to change the etch rate. So in this case, a shorter gas residence time increases the etch rate. I'm sorry, increase the flow means the amount of time any volume of gas is in the reactor is shorter.

And why? Because it increases the availability of neutral etchant species. So you have the unreacted etchant coming in faster. And it also reduces the concentration of the reaction byproducts, because again, these are all chemical reactions. You want for a chemical reaction to go faster, you put in the reactants faster, and you remove the byproducts rapidly, as well.

So not only does it affect etch rate, but changing things like pressure also affect the shape of what you get. Here's an example on slide 14 from that same article of three different pressures. Here we're at 10 tour to try to define a polysilicon gate. I'm sorry, 10 millitorr, 1 millitorr, and 0.1. So we're going down here each case by an order of magnitude and pressure. We're using lower pressure at high flows.

You notice we get more anisotropic etching and less notching, at least when using this particular chlorine polysilicon etch. So here at 10 millitorr, you get a fairly isotropic etch. It's not totally isotropic, but you notice it's etching sideways, as well as vertically. And how can I tell? Because this is my photoresist. And it's undercut on the left and right side. So the line is a lot narrower than what had been designed.

There's a little bit better anisotropy at a lower pressure at one millitorr, but we still have-- now we've developed a little notch at the bottom, which is, again, that's not considered to be good. Finally, we have the lowest residence time, or the lowest pressure here, at 0.1 millitorr. This paper demonstrates very vertical sidewalls and no undercutting. So this would be considered a very good quality etch.

You may be able to make one device with this etch. I'm not saying this etch won't work, but from a manufacturing point of view, how do I control-- how do I control my line width? In my critical dimension, it's not well controlled at all. So this would not be considered a manufacturable type of process. There's another consideration. So we talked a lot about-- we care about the etch rate and the shape of the sidewall because that affects how we're going to implant it and do other things.

A huge consideration anytime you are etching features-- this is not just true of a polysilicon gate, but I'm giving the example for a polysilicon gate, so gate etch is the example-- is the topography. And so what's the problem? Well, highly anisotropic etching. So etching, there's only down. Only vertical. Doesn't etch sideways. It's required to achieve good control-- we just said that-- over the gate length and the shape.

But unless you have a long over etch that you do at the end of the step, you can get stringers left behind. And let's just show an example here. So on the left, I'm showing where we've etched through polysilicon. There was no topography on the original part of the wafer. It was just flat. So I have photoresist and I've etched polysilicon anisotropically. Everything looks very good.

Now, what's happened on the right? Well, on the right-hand side, I had some initial topography. So I had a step that looked like this. So as a result of having that step, the polysilicon overlayer, which is shown by this sort of stippled region, the polysilicon went over everything. So the polysilicon in this region is actually thicker than it is in the flat portion. So by the time I cleared the gate and I just hit this oxide right here, I have left behind a stringer because I haven't etched long enough because the polysilicon thickness was thicker in this region.

And remember, I'm only etching vertically. It's not going sideways. So this stringer is considered to be undesirable because it's now an extra gate electrode, or it's an extra thing you didn't want in your circuit. How do I get rid of it? Well, you say, just keep on etching. Go to do an over etch. So even though I've finished etching the poly gate and I'm done in this region, I'm not done over here. So you keep it in and you keep banging on the top of that stringer until it finally goes down to 0.

So you need a very long over etch whenever you have surface topography in order to get rid of stringers. What's the disadvantage of a long ever etch? Well, you might break through your oxide because you're still beating on that oxide. So unless you have really good selectivity, as you're beating the stringer down and getting rid of the stringer, you could be destroying the oxide.

What's the other bad thing? Well, your over etch may not be perfectly anisotropic. It may be etching sideways, a little bit on the gate. So unless it's perfectly anisotropic, you'll have that. So there's always a balance with over etch. So this is why you will find people say, if you want to get the finest patterns, your wafers should be completely flat. You should have no topography because if you have a completely flat wafer, you never have to worry about stringers. You etch down. You stop on the layer below. And you stop. Period.

Your over etch in that case is only dictated by the non-uniformity in the film across the wafer, as we talked about before around uniforming etch rate. Now we have another requirement on over etch. I have to not only take into account non-uniformity across the wafer, but if I have topography and I have an anisotropic etch, I have to get rid of stringers.

So you will see the topography on modern devices, as they get smaller and smaller, everything has to be smoother and smoother. That's why people went away from locos. Locos creates this big hump on the way. People don't like that. We do shallow trench, which creates a perfectly flat wafer. That way I can pattern a very, very fine line and I don't have to worry about stringers.

Slide 16 is from a different text. It has another example of stringer formation. And maybe this cross-section might be easier for you to see. So here's an example where I have a field oxide. So the field oxide has a certain step height to it, shown here. And I have-- and of course, when you deposit the polysilicon, it gets deposited everywhere on the wafer. So this cross-hatched region is the poly. Then I want a pattern. And this is the photoresist after it's been developed and patterned. OK.

Now, step B here, I put it into the etcher and I use the photoresist as a mask to etch the poly. So this looks good. I have very nice sidewalls and I stop on the oxide. But what happens? Well, in this region right here, look at the thickness of this gate material where it goes over the step. The thickness from this point of the top to the very bottom is almost twice what the thickness is in the flat region.

So as a result, once I've cleared the flats, I still have this little notch of material that's un etched that's sitting there because my etch only etches vertically. If the etch had etched sideways, then you wouldn't have to worry about that notch. So an isotropic etch doesn't tend to give you stringers. Stringer formation is not such a big deal. But if it's anisotropic, it is. So what we do is you continue in step C with an over etch. You continue to beat down on this, and you introduce an extra 50% of an over etch to get rid of the stringers.

If you leave the stringers in place, you'll have shorts on your devices because then you have this poly going all the way around the field. You can short the source drain to the gate, and you can get all kinds of problems. Now, this is drawn by a cartoon, or by an illustrator, so it's very easy. It doesn't show any detrimental effect of the over etch. In practice, over etching to remove stringers almost always has a little detrimental effect.

And so the name of the game is trying to control your topography so you minimize your over etch so you don't have to destroy the structure by over etching too much. So let's look at slide 17. This is an example on that paper, that article from IBM, on the effects of topography on gate conductor yield with respect to shorting. As I mentioned, if you end up with stringers, the gate can be shorted to the source and drain. So here's an example. And they etched here 21 wafers or so-- 20 wafers.

And the percent yield. So if you have a high yield, it means you don't have gate to source drain shorts. If you have a low yield, it means you have shorts between the gate and the source drain. So this wafer had no topography. So they etched the gates on perfectly flat wafers. And the yield was very high. It was like 90% for these five wafers.

Now they introduced-- in the field oxide or somewhere, they introduced a 50-nanometer step. You can see the yield goes down a little bit. Not too bad, though. So they have more shorts. And here if you introduce a 100-nanometer step on this particular gate etch, you get a yield of only 30%. So 70% of devices are shorted. So this means you're getting a lot of stringers forming.

So you can look for stringers visually in the microscope with a TEM or an SEM, or you can look at them with special test structures where you look at your shorting and your yield. So this is an example where you would say for this particular gate etch process that the 100-nanometer step is too much. You need to reduce your step height and not produce such an abrupt step. Otherwise you need to increase your over etch time or do something. So stringers are really a big issue, and it's always a trade-off between their removal and messing up the structure.

Here on slide 18, I just showed you some-- I've shown some ways where we can calculate over etched requirements based on the topography one might have during gate etching. So here's an example on this particular figure where I'm showing you a field oxide. Remember we said the field oxide might have a certain height, which we're calling T sub FOX, the thickness of the field oxide. And it has a certain angle that it's been etched. So it's not perfectly straight up and down. It has a certain angle, theta, that it makes with respect to the substrate. So that's your field oxide.

And this black region on top is meant to represent the polysilicon. So remember, polysilicon is always deposited everywhere across the chip. So the black region is the thickness of the poly. Now, you notice that on the sloped portion, the thickness of the poly from here to here is given by T poly times the cosine theta. So it's thicker here than it is here. It's just by geometry.

So the thicker portion of the poly here means that you need to increase in order to get clear all the poly. And to avoid forming a stringer down here, you need to increase the etch time. And in fact, this formula gives you here, for this particular topography, gives you an example calculation of the over etch here in percent that you would need for this topography. And it's 1 over cosine theta, that quantity, minus 1, times 100%.

So if you have a sidewall angle of 45 degrees-- so this particular example says 45 degrees-- just because of this, we need about a 40% poly over etch time to avoid stringers. So what do people do? Well, either you have no topography, you keep it very flat, you bury your field oxide, and that's called a shallow trench. Or if you are going to have a little bit of topography, you make it-- you kind of ramp the angles and you make very shallow ramps. You don't have anything really abrupt.

So depending on this angle, you have to adjust your over etch time. Otherwise you'll end up with stringers. Sometimes, by the way, here on slide 19, sometimes I've been calling these things stringers. Stringers is the derogatory term. Stringer usually means something that was left behind that you didn't want there. Sometimes a stringer is something we want. In fact, when we want it, we give it a positive name. We call it a sidewall spacer.

So we often use anisotropic etches without an over etch to get a sidewall spacer intentionally. Here's an example of an intentional production of a stringer. We have a substrate. We put one film on. It could be polysilicon that we have now patterned with a very sharp step.

We put a second film on where we adjust the height-- or the thickness, I should say-- of the second film to be about, oh, just the right thickness so that after I've done a purely vertical or purely anisotropic etch, when I'm done, I have a sidewall spacer forming. So this material is left behind. The nice thing about it is self-aligned. So I didn't have to put another mask down and do photolithography to pattern this.

All you have to do is have a step in your film, number one. You put another film down. You clear the whole thing with a vertical etch. And bingo. Right on the sidewall, self-aligned to it, is a spacer layer. In fact, that's exactly how spacer layers are formed in modern MOSFETs. The poly gate is etched vertically. A certain thickness of nitride or oxide is put down. It's etched anisotropically, and you end up with a sidewall spacer.

So stringers, so to speak, are not necessarily bad. Usually the word stringer means something left behind that you don't want. But if you do want it, you call it a sidewall spacer. And the width of the spacer, which is from this point here to that point there that's sort of the lateral dimension, is a function of how long you did etch and how much over etch you did, as well as the thickness of film two relative to the thickness of film one in this example.

So you adjust all of those to create different width spacers. That's a very critical process in modern technology. Now, this slide-- actually, I apologize-- is not in your handouts. It's on page 20. There's two slides here. The next three slides are not in your handout. This one is. It's the next two after that are not. And if you need the up-to-date handouts, they're all posted on the web. So you can get the PDFs yourself.

I wanted to list here on slide 20 a couple of new issues for gate etching and patterning that have come up in the last, say, three to five years that you're going to hear more and more about. In the year 2000, this is what people were concerned about. They were concerned about etching poly, making it vertical, maintaining a critical dimension of, say, 0.18 nanometers-- 0.18 microns, rather.

And they were etching through only polysilicon and putting a silicide strap on top after the gate etch had been completed. So that was what should call the more traditional or conventional technology. What are people doing these days? Well, people these days are not just doing this. For one thing, the gates are a lot narrower. Look at today's gates. People are, instead of doing something like 200 nanometers, we're down to 70 or 80 nanometers. So the dimensions have shrunk quite a bit.

And people are doing what they call fully silicided gates. So instead of just reacting the polysilicon just at the very top of the gate, people are actually putting down enough metal and reacting it at the right temperature that the metal reacts all the way to the very bottom. In this case, on the bottom would be nickel silicide. This was formed by depositing nickel and reacting it at 450 all the way to the bottom.

So there's no polysilicon left because these days, people don't want polysilicon gates. Polysilicon is a great material to etch. It's highly compatible chemically with oxide. But it has gate depletion problems. It doesn't have enough carriers in it. Polysilicon, what's the highest you can dope polysilicon? Well, maybe 1 to 5 times 10 to the 20. So that limits your resistivity and it limits your gate depletion.

People want to use metals. What's the concentration of electrons in a metal? Well, 10 to the 22 or 10 to the 23. It's huge. It's three orders of magnitude higher. So we need more electrons or more holes, more carriers in our gates these days. The way we do it is we take poly and you convert it completely to silicide. This is called a fully suicided process. This is only in research. There's a lot of issues with doing this.

Here's an example where they tried to do the siliciding with cobalt in the upper micrograph, upper SCM micrograph here. It didn't work too well. As you can see, the cobalt went down in. It didn't make it to the bottom. In fact, what happens the silicon got sucked up. And the next time we'll talk about silicide reactions. The silicon was actually moving, and you ended up with a big void. So you had no gate in this example.

So this particular paper was just showing-- I took from a few months ago at the VLSI conference in Hawaii. Nickel silicide works much better for fully siliciding than does cobalt silicide, which leaves behind voids. So there's is an example. Not only are we etching polysilicon gates these days. We may end up siliciding them completely. And that's what people are considering.

There's another type of process which I'm picturing here on the lower left called damascene replacement gate process. And here's an example of a gate where I have ti-nitride. So this material-- well, it's got tungsten on the top. That's the bright material. And this sort of gray material is a thin layer of ti-nitride. And the gate dielectric is hafnium dioxide, which is down here. This is on a silicon germanium type of high mobility transistor. That was also published about six months ago.

Doesn't look anything like a polysilicon gate. So not only are we fully siliciding the gates, but sometimes we're actually replacing the polysilicon with another material, with a metal material. And ti-nitride is a very popular one. So gate etching is evolving from not only having to etch poly and stop on dielectric, but now we may have to etch metals.

So lots of new technologies are being developed over the last two or three years and will continue to be developed to push the technology forward. So it's these next two slides, I apologize, not in your handouts, 21 and 22. But they are on the website, so you can download the PDF file. How did they make-- I just wanted to show you this lower left. How does one make a gate that looks anything like this, just to give you an idea? Well, this is the way you do it.

Again, this particular example I took from IEDM of 2002. People use the ordinary process to make a polysilicon gate. So you put down poly. You etch it. You make your sidewall spacers. So that's just like what we've learned about in this course. And here is a very short channel MOSFET. Here's a very long channel MOSFET. Now, at the end of defining your MOSFETs, you put down this gray material, and you put it everywhere, very thick. So you put down a dielectric. And then you CNP it down.

So you have dielectric that's been CNP'd and flattened. It's been flattened down just so it stops, the polishing stops, just around this height. Now what you do is the polysilicon is then chemically etched and removed. So you etch out with an etchant that digs out the poly. And if you want, you can even dig out the gate oxide and replace it with a high K.

So this is what's called a replacement gate process. So actually, I use poly. To etch your poly, you put down sidewall spacers. You go through the high-temperature anneal. Now, why would anyone want to do this? You put down poly, you pattern it, and then later on you're going to remove it. Sounds kind of crazy. Why didn't you just put the metal down to begin with? But does anybody have any ideas on why you might want to use poly and then stick the metal in at the very end?

Well, there's a couple of reasons. One reason is we know how to etch poly perfectly. We spent the last 20 years getting perfectly vertical sidewalls. So people know how to etch it. That's one reason. The second reason is poly can go up to very high temperatures. What do I have to do in this process? After you put the gate down, you have to implant the source drains, and you have to anneal them, typically at 1,000 degrees these days for 10 seconds, or 1,050.

There aren't too many metals you can pop up to 1,050 and not either melt them-- most metals will melt-- or if you don't melt them, they could undergo dramatic phase transformations. Some of them could diffuse into the wafer and cause problems with deep levels. So metals are not considered very good to have at high temperatures. So this is a way of using poly as a dummy. We know exactly how to pattern it. The only problem with it, it just doesn't have enough carriers.

Well, we'll put it in the process. We won't change the process till the very end. We'll etch it out. I don't know how manufacturable-- I don't believe anyone is manufacturing products. This was just published in research about two years ago. It is kind of a neat idea. So after you remove this and you open this hole, well, then you have an open hole. You can put down your high K, your hafnium dioxide.

You can put down your ti-nitride by sputtering, and a tungsten cap, and then maybe some capping layer of dielectric and CNP the whole thing down. So that's how they got-- in that last picture, that's how they got that replacement gate. And in fact, slide 22 was also not in your handout. But this is what it looks like in TDM a cross-section in a 50-nanometer MOSFET where they remove the poly, they dug it out, they put in this very dark black layer of hafnium dioxide. So that is your gate dielectric.

The ti-nitride is next, a very thin layer. That is your gate material. They then have tungsten that they fill it up with. So they have the tungsten material, which is a little bit more environmentally friendly. Or it's an easier material to work with to fill up the plug. So that's an example where the poly was completely removed. So new types of things have to be done with forming gates these days.

OK, let me move on here. Now on slide 23, I'll move away from the gate etching and talk a little bit, spend a few minutes, about modeling and simulation. There's a lot of similarity between the models that we introduced for deposition in the last couple of lectures that were in chapter 9 and the etching models. And why is that? Well, both deposition and etching use incoming chemical or neutral species and ion fluxes. And they have a lot of very similar processes. The difference is when you're depositing, the species come down. They move around and they stick, and they react in some way that deposit.

When you're etching, they come down, they move around. They have a chemical reaction, but they actually remove material. So one can be considered mathematically the inverse of the other. But because in math, you can just change the sign from positive to negative in a computer program, you can use very much similar methodologies to model etching.

So here's an example of things that might take place. During etching you have ionic species coming down. You have things coming down and being sputtered away. And you may have things coming down and being emitted or desorbed. So a lot of these fluxes, the whole flux equation, is very much analogous between deposition and etching.

Here on slide 24, just like in the case of deposition, the etch rate is proportional to the net flux arriving at a point. Remember we said the deposition was proportional to the net flux, the flux in minus the flux out? Well, the etch rate is often proportional to the net flux. So we have the same kind of situation. You have a wafer, which has a certain patterned surface. You form a virtual dashed plane above it, and you look at the distribution, the angular distribution of the species arriving just above the surface on that dashed plane.

If you have a chemical etching species that is something that's neutral that is not ionized, you assume you have an isotropic arrival distribution. So you put a cosine theta to the n, where n equals 1, for a chemical species. If you have an ionic species, it can be accelerated towards the surface. It can arrive very much with a vertical orientation. And here we prefer etching-- we put in cosine theta to the n where n could vary between 10 and 100, or something like that.

Remember, the higher the value of n, much more vertical the distribution. In the deposition case, remember we had the concept of a sticking coefficient where we use very much that same concept. When we're etching we say an ionized species usually sticks where it lands and it's going to react and etch, while a reactive neutral species may have a low SC value. It might bounce around before it finally reacts and etches off the surface.

As far as the physical component of etching like sputtering, the sputtering yield has the same angular dependence that we used in the deposition case. So again, very similar process going on when we include sputtering. So in slide 25, I'm showing what's called the linear etch model. There have been some very specific models developed. Here I just want to give a general purpose etch model, a couple of them, which can be applied broadly.

The so-called linear model assumes that there are chemical and physical components that act independently. So they're not linked. They're independent of each other. And in this case, we have an etch rate that is given by just a sum of two terms, one term that has to do with the chemical flux, F sub C. That's the chemical flux onto the surface at any point. And another term that is proportional to F sub I, the ionic flux.

In front of each one of these fluxes is a K factor, KI and KF. These are relative rate constants for the two components. Usually you use these as fitting parameters. If I have a large KI, I have a lot of ionic component, whereas if I have a large KF, there'll be more of a chemical component. What is that-- the physical component can be due to purely physical sputtering, for example, like we talked about before.

Or it actually can be some ion enhanced mechanism in the regime where the chemical flux is not limiting the ion etching. So this is a linear etch model with the independent, chemical, and physical components. Here's just an example of some simulations. Remember, speedy SM was a topography simulator, does both deposition and etching. Here on slide 26., I have some speedy simulations using the linear etch model.

And we have A, B, and C. Well, what we're doing-- each one of these contours, again, represents the shape of the etched surface at a given time. So we're looking at snapshots in time. Here I have an etched mask that looks like this. So this is my masking material on the surface. And I'm etching into the silicon, or into the film below.

Here I have n equals 1 for my chemical flux angular distribution. So it's going to give you fairly isotropic type of etches here. Etching laterally, this dimension, just about the same as you etch vertically. So there's an isotropic type of etch. Here's where you have n equals 80 for the ionic flux angular distribution. And you end up with very vertical etching. Basically everything's coming straight down. And etching, there's not much lateral etching.

So this is all chemical etching. The ion flux is zero. All physical etching in this case. And here's a 50/50. I have half chemical and half physical fluxes of both. And you can see it etches a little bit laterally and a certain amount vertically. So you can generate a variety of etch profiles using a relatively simple model with only maybe three or four parameters. And then you can fit this to what you actually see in your actual etches.

The second type of model, which is quite different that's shown here in slide 27, is called a saturation dash absorption type of model because the fact-- very often, this is used when the chemical or the neutral etching and the ion etching or the physical etch components are coupled. So when one of these affects the other, when you have a synergistic effect, which we know is often the case, especially when you have side wall inhibitor layers or something, you need to use this type of model.

So here's an example. Let's say you need an ion flux to remove an inhibitor layer that is formed by chemical etching. So instead of a linear model where these two just added up linearly, these add sort of resistors in parallel, so to speak, or capacitors in series if you're electrical engineer. The etch rate is just given by the inverse, 1 over, the sum of 1 over Kl FI plus 1 over SC FC. So there's 1 over a chemical flux and 1 over an ionic flux.

If you plot this equation, what does it look like? Here's the overall etch rate as a function of-- now I'm plotting as a function of the ion flux times K sub I. So it's a function of this quantity. So it goes up. And the parameter here is the chemical flux, SC. So if either flux is 0, so if SC is 0, I get 0 etch rate, regardless of what my ion flux is. You need both the chemical and the ionic flux.

So if either flux is zero, the overall trade is 0 because both of them are required to etch the material. And you end up with these sort of saturating type of curves. Here's for a chemical flux of 0.5. As I increase the ion flux, you increase the etch rate until at some point it sort of saturates and it's no longer a function of ion flux. And it's limited, then, by the chemical process.

So you have a series of desaturating curves. It sort of looks like transistor characteristics because you have two things controlling it, both the chemical flux and the ionic flux. So basically, this is that same plot I just showed. The etch rate tends to saturate when one component gets too large relative to the other. So you're always rate limited by the slower of the two series processes, either the chemical etch rate or the ionic etch rate.

This is a very generalized approach. Again, it can be broadly applicable, and you can use it to fit a lot of different etch models, again, with relatively small number of parameters, maybe three to four parameters. . On slide 29, I'm showing a speedy simulation, again, where we have equal chemical and ionic components. And you notice what you get is reasonably anisotropic etching. There's a little bit of lateral etching here down at the bottom.

And why is that? Well, because in order to continue etching, you need an ion flux to remove the inhibitor layer or whatever it is. And that ion flux tends to be very vertical because we have a large n value in this case. And so the side wall inhibitor is not removed. The inhibitor layer is not removed on the side walls, but it's only removed on the flat surfaces.

If you change this n direction, you'll start to see a little more lateral etching-- if you change the n number, rather, which the n will determine the angle of the cosine theta to the n dependence. And then you'll get a little bit more lateral etching. So there's an example for equal chemical and ion components.

OK, so I just want to summarize what we've done in chapter 10 on plasma etching. There's a lot of important issues. Some of the key ones we've covered are the selectivity with respect to the layer below you, plus the selectivity with the layer above you, or the mask. The directionality of the etch. How much do you etch vertically versus laterally? And the shape of the profile on the sidewall. There tend to be two components of etching. There's a chemical one that is reactive neutral species or free radicals.

There's a physical component, which is just due to the ions being accelerated towards the surface and reacting or sputtering off components from the wafer. These can be completely independent. And then we use the linear model, or they can act in a series in a synergistic manner in what we would use the last model that we showed.

Topography on the wafer can really increase the selectivity requirements. It makes the patterning of fine features and stopping on thin layers very difficult. Of course, lithography is impacted, as well, by the presence of topography because of the depth of focus. But all modern processes seek to minimize topography. So people have moved away from field oxides that go up or that go down. We want fill oxide that's completely flat with the surface so you see shallow trench isolation.

Special high-density plasma etchers have been developed to use chemicals like HBR, chlorine and oxygen combined together, specifically for gate etching to control the selectivity so you can stop on the gate dielectric and to control the shape of the sidewall profile. At the same time, you minimize damage to the underlying structures and the underlying gate oxide.

The models that I've talked about here for etching are fairly what I would call empirical. If you want more predictive or more physically accurate models of plasma etching-- Professor Sarwin used to teach a class. I'm not sure he's presently teaching it-- really on all the details of plasma etching as an entire field in and of itself. We've only been able to give it two lectures and one chapter of your textbook.

So these models, while helpful to people that we have now, people to use topography simulators like speedy, it's not going to give you the same level of chemical understanding if you had a full-blown model of the etching. Again, it's fairly empirical. We're fitting curves in speedy. And we're fitting the shape. But it gives you some idea of what's going on. If you want more details, there are other more detailed classes you can take.

OK, before we finish up, let me just remind people if you came in late that we have set the schedule. After today is Thanksgiving. That's a holiday. There's no class. Next week we'll talk about silicides and novel gate materials. After that we'll talk about strained silicon. And then you guys start talking.

I've picked the first four speakers here for the reports. You have to speak for 16 minutes, and then I have three minutes for questions and answers on the seventh and on the ninth. These are the people speaking on the ninth. If you're doing a written report, you must bring it to class on Thursday, December 9. That's the last you can hand it in. And this will be posted, this list, on the website. OK, thanks.