

JUDY HOYT: Just remind people where we are on the course schedule, we just finished up chapter 9 on deposition and epitaxy, and we're going to plow ahead, moving along to chapter 10, which we'll cover in the next couple of lectures. Chapter 10 is on etching, which is a very important capability for CMOS processing.

And today is the 18th, and I guess homework number 5 is due today. So I brought a new folder here up front, this yellow folder, we'll use to hold your homework number 5. I'm going to have-- homework number 4's been graded, but I need to go through it. So I'll have that back to you next week.

So let's take a look at today's handout, which is handout number 33, and we're starting chapter 10 on etching. We've talked a lot about, now, a lot of different processes. We spent the last three lectures talking about depositing thin films. After they're deposited, typically they have to be etched to make some kind of a circuit pattern or whatever you're trying to pattern.

An important requirement of etching that we're going to talk about is that films need to be etched selectively. You want to etch one film, but you don't want to etch what's underneath it typically, or sometimes, you don't want to etch the mask. So selectivity is very important, and we'll talk about that.

Besides etching thin films, sometimes we want to etch down into the substrate itself for dRAMs these days to-- so you don't use so much chip area. The capacitors, the storage capacitors, are actually sometimes in the trench process. They're actually etched deep trenches into the silicon. They make capacitors to store the memory-- to store the charge and the memory that way. So sometimes you need to etch silicon substrate.

Dry etching, as it's called, or plasma etching, is really critical to CMOS fabrication. I'm going to emphasize dry etching. I only have two or three slides on wet etching. Wet etching is still used in research and development, but it's very rarely used in a manufacturing process, and we'll see a little bit why that is. So the next two lectures, what I want to try to do is give you a basic introduction to etching, mostly plasma etching, some specific examples.

The next lecture, I'm going to have some specific emphasis on gate etching, polysilicon gate etching, since that's really very important to the CMOS process. And we'll talk a little bit, next lecture as well, on modeling, how we model etching processes. Whoops, I forgot to go to full screen mode here. Let me do that.

So let's move on to slide number 2 in your handout. Basic introduction, this is what we talked about in the beginning of the course. Just to remind you, the planar process, what it consists of, usually, you'll have a deposited film here, which is shown in blue, and then to pattern it, you need to put down a photosensitive film, which is called photoresist, that can be put down at room temperature or spun on.

You then expose it through a mask using ultraviolet light, or the right wavelength light. That either causes the resist to stay wherever it's struck by the light or to be removed or etched where it's struck by the light. Either way is a positive or negative resist, and you end up with, here, after developing the resist, you end up with an etch mask.

So I have photoresist shown here in gray, and you're ready now to etch the blue thin film below it, and you do that either in a dry or a wet process. Then you remove the resist, and you're done. So now, you've transferred the pattern from the mask to the resist and from the resist to the film. So we won't talk about the lithographic process of transferring to the masks to the resist. That's a whole class, Hank Smith's course. What we're going to talk about here is how we etch this thin film.

The most important things people care about are selectivity and directionality, and I'll talk a little bit more about what they mean. Selectivity comes from the chemistry. That is, how selective is the etchant to a particular material. The directionality, that is how the etches vertically versus horizontally, that usually comes from more physical processes, like energetic ions bombarding the wafer. And modern etching techniques have a little bit of directionality and a little bit of selectivity, both of those.

Simulation tools are only more recently being developed. They are starting to play an important role, just like in deposition. Again, you will not find etching modeled in Supreme 4. There is no topographic modeling in that. You need a topography simulator like Speedy, and usually, topography simulators do both deposition and etching because they are really inverse processes.

In one case, you have things coming down, reacting and sticking. In other case, you have things coming down, causing a removal reaction, and something coming off, an etching. But they're very much mathematically very similar processes.

So slide number 3 shows a little bit about these characteristics, some of the important characteristics that I just talked about, and illustrate them. Up at the top, we have a situation where we start with photoresist or a mask. Let's say, that looks like these dark regions on top are the pattern or the mask, and we want to transfer that into this hatched region, which is the film we're trying to etch.

After etching, shown on the upper right, what you have is, here, you see the mask, which is retained. Good fidelity, the mask was not attacked, so there's good selectivity. But we have what we call undercutting. You can see what happened is not only did we etch the film vertically, but we also etched laterally on each side, and so the feature size has been distorted. It's not exactly what the feature size was on the mask because of this undercutting.

So that's the top two pictures illustrate undercutting. The second one here, labeled B, illustrates two things, both undercutting and poor selectivity. Up here, I had good selectivity. The resist was not attacked, and the layer below, whatever it might be, was not attacked by the etch. You noticed, it went down, and it stopped.

Here, I have both undercutting and poor selectivity. Indeed, you can see what's happened. I've etched laterally left and right, so my holes end up being bigger than I had wanted them to be on the mask, and I've also attacked the layer underneath, so I don't have very good selectivity. The etchant continues to etch this layer-- it continues to etch the bottom layer as it's in the solution, or maybe in the plasma, and not only that, the etch also attacked the photoresist. And we caused it to get rounded, and that also can distort the feature size.

Usually, what we want most of the time, although it's not always the case, most of the time, we want what we call anisotropic etching. We want the etching to go down perfectly straight. Generally maybe a slight angle, but we want to be able to control the angle in any case, almost vertical.

And usually, we want highly selective, so we'd like to have what we call a selectivity on the order of 25 to 50. So that means the etch would etch this top film, the film we're trying to etch, 25 to 50 times faster than it etches the film underneath it. Both of these, anisotropy and selectivity, are usually desired. Unfortunately, it's very hard to get both, the best of both worlds, but there are systems that try to optimize this, and that's what we'll talk about.

So as we move on the bottom of slide 3, as you move from the left panel A to B to C, what you're moving is from completely isotropic etching, as you can see by the little vectors here, that the film etches-- this particular etching is etching this material vertically at the same rate it is laterally. Here's a little more anisotropic.

The vertical etch rate is faster than it is laterally, but still, there's some lateral etching, and this is completely anisotropic. There is no lateral etch rate, only vertical etching. So you're moving from completely isotropic to completely anisotropic-- I'm sorry, from left to right, so you're getting more and more directionality.

On slide 4, I've listed some things that are fairly obvious to you, but I just wanted to make note of some of the practical requirements on etching. First of all, we want to get the right profile, whether we want sloped or vertical. That's obvious.

We don't want much undercutting. Undercutting or etch bias is bad because the feature size on the mask is not maintained on the wafer, and that's not what the circuit designer designed. So that's not good, and you have lack of control.

You want good selectivity to other exposed films and to resist. You don't want the etchant attacking something that you don't want it to attack. You'd like the etch to be uniform across the wafer, so as you etch a trench on the left of the wafer and the right of the wafer, the trench depth is exactly the same.

This is an important one that people-- number five, people often forget. We don't want to damage the surface, and we don't want to damage the circuit electrically. And we are bombarding, in a lot of these plasma etches, we're bombarding, with hundreds of electron volts, the surface with ions. And they those ions can constitute a current to the wafer, and that current can actually damage sensitive electronic components.

Obviously, we want the etch to be clean. We don't want to introduce a lot of impurities. We don't want to introduce a lot of metals into the wafer, and safe. So here's an example, on the bottom part of page four, what I mean by plasma damage.

Let's just say you have a MOSFET. You have a field oxide FOx here on the left, another one on the right, and this thin region in the middle is meant to represent the gate oxide. And let's say-- and I have a polysilicon gate that looks like this. There's a large contact pad, which we often have perhaps for contact. And then it goes down, and it's a very narrow gate, which is what we were doing today, say 90 nanometer technology.

Now, I go to put this in a plasma etcher. I'm etching something on the wafer, maybe not the gate, but something else. So I have all these arrows going down here. It's supposed to represent ions because I have a plasma current, so I have ions that come down, strike this-- they strike everywhere, but in particular, they strike this large area.

Because this is a conductor, this current, with all these ions-- the flux of ions to this large area, constitutes a current. This current can then flow down the conductor, and where does it go? Well, if this is a thin oxide, it can tunnel through that oxide. So all of a sudden now, I could be forcing a very large current through a thin oxide, and you can end up with charges buried in that oxide and damage.

And this is called an antenna because this thing can collect ions from a large area and funnel them down, acting like an antenna, and you can do a lot of damage. You have to be very careful. Yes?

AUDIENCE: [INAUDIBLE]

JUDY HOYT: Not necessarily. It depends on what I'm etching. This may not be the gate etched. For example, after you do a gate etch, you need to open up-- you may need to open up some other holes somewhere else. This could be happening during the gate etch.

It could be happening during another step, but there are, nevertheless, ions coming in, in this particular example, into the gate. And so this may not actually be a gate etch step. It's just an example where you have the gate exposed, potentially collecting current. It can go through the oxide. So that's something we need to worry about when we have high density plasmas.

Slide 5, so let's just talk about some of the basic concepts, and I'll go through the wet etching relatively rapidly. As I said, there's two types of etching-- dry, which means you use a plasma, and wet, which means you use a beaker or something in a solution.

So you typically, for wet etching, you submerge the wafers in some specific bath. The good thing about wet etching, it tends to be highly selective. It's all based on chemistry. There's no physical components. You have no bombardment, so as a result, because it's based on chemistry, you can select the etchants in the solution that they will etch exactly what you want.

It's usually isotropic. There are a couple of what we call crystallographically dependent etches. KOH, potassium hydroxide etching silicon is an example of that. There aren't a whole lot of cases of that, though. Most of the time, the etches end up being selective.

Here's an example. Here's a wet etching of SiO₂, but by HF. So you add HF to SiO₂. You end up with some extra water, and the SiO₂ goes away. You can etch silicon in a combination of an oxidizer, like nitric acid, and HF. Sometimes, you add acetic, so HF nitric acetic, but here's an example. Etching silicon wet etching with nitric NHF, so the nitric oxidizes the silicon.

The HF strips the oxide, and you end up with overall etching of silicon. HF nitric, though, is not that easy to mask because it also etches any oxide on the wafer. So you have to watch out. You'd have to use nitride to mask it, but those are some examples of wet etches.

Slide 6, let's talk a little bit about isotropic etching and undercut. If we're doing isotropic etching, as I'm showing here in panel A at the top, one is supposed to be my mask. Layer 1 is the mask. Material 2 is what I'm etching. It's isotropic, so it's etching vertically and laterally at the same time. I'm going to have some undercutting.

And in fact, we can sort of mathematically express the amount of anisotropy, how anisotropic is the etch, by defining this number here called a sub f. If you define this anisotropy factor, a sub f, to be 1 minus the lateral dimension that you go, the lateral in a given time, etched amount divided by the vertical. So in my diagram here, this would be 1 minus b over a, so I etch laterally by distance b-- I'm sorry, b over d. Laterally, by distance b and vertically by a distance d, and so that would be the anisotropy factor.

So this anisotropy factor is 0 if you're isotropic etching because then b equals d where r lat equals our vertical, and then af goes to 0. If you're totally anisotropic, this thing basically-- b goes to 0 because you don't etch laterally at all, and this number goes to 1.

Now, typically, if you're etching a film, what's a practical thing you need to think about? Well, you need to know the etch rate. That's fine. Let's say you know the etch rate, and I give you the film things, you can calculate how long you need to etch. You just divide the thickness by the etch rate, and that gives you the time. But that's not the time you should use for the etch in practice.

In practice, you almost always do something called overetching. So you typically leave it in the bath longer than what you would need to etch the film thickness. Why is that? Well, basically, you need to ensure that you actually do completely remove the film because the film-- if I tell you the film is 1 micron, first of all, you shouldn't believe me.

I may not know what I'm talking about. It may not be a micron. And even if I do know what I'm talking about, films in reality are never exactly one thickness across the wafer. There's always variation. So you may have a slight high spot in the center or a high spot at the etch. Who knows?

So you almost always do overetching, and that's what we're picturing in the upper right. What we've done here is, in showing you the contour of what the etched region looks like for different times, so this is after a certain amount of time. This other dashed line is a slightly longer time, and now you've already broken through.

You've hit that, but maybe somewhere else on the wafer, you haven't quite etched the film. So you go a little bit longer, and you end up with the profile shown in the solid line here. So that's because we have done what we call overetching, and overetching has to be done in any practical process, either in research or in manufacturing, to take into account variations in not only the film thickness, but variations in the etch rate. The etch rate is not exactly uniform necessarily, depending on the etcher or even depending on the solution.

For wet etching, the selectivity is usually excellent. I mean, we usually choose our etchants so that we get good selectivity, and again, selectivity is defined as the ratio of the etch rate of the film we want to etch divided by the etch rate of the film we don't want to etch. And chemical reactions tend to be very, very selective.

So just to illustrate this idea of overetch and selectivity, I wanted to go through a quick example that I took from your text, and that's what's shown on slide 7 and 8. Here's an example of a question that you can consider.

You have a relatively simple setup. You have a half micron layer of silicon dioxide on a silicon substrate, and you need to etch it down to the silicon. Now, you assumed that the nominal etch rate-- so you're putting this in some etcher, whatever it is. Let's say it's not HF, but it's some other type of etch. Has a nominal etch rate of $r_{\text{sub O}_x}$, so that's a certain number of microns per minute that you've been told that the oxide etch is in this etchant.

But there is about a plus or minus 5% variation in the oxide thickness. The oxide's not perfect. It was grown in a tool that gives you a 5% variation across the wafer. There's also a plus or minus 5% variation in the oxide etch rate. You're putting it in a solution, or let's say you're using a plasma etcher, it's not perfect, so it varies plus or minus 5% across the surface of the wafer.

So the first thing we want to figure out is how much of an overetch is required, and we usually represent overetch as a percent of the total time. So if I'm going to etch for a minute, to get rid of the film, and then leave it in an extra minute to make sure that I've taken care of non-uniformities, That's called 100% overetch because it's 100% of the main etch time. I'm contributing 100% of that to the overetch step.

So we want to know how much overetch is required as a percent in order to make sure that all the oxide on the wafer that you etched all the way down at every point in the wafer. So let's just go through part a first and how we would think about that.

Well, I've just drawn a little cartoon down here. Basically, this green material is supposed to represent the oxide. It's nominally half micron on the wafer, but at the thickest point, it's thicker than that. It's 5% thicker, so it's 0.525 microns.

So basically, what we want to do is, the overetch has to be done to make sure that you remove the oxide at the thickest point for the slowest etch rate, so you take a worst case. So I take the thickest point is 0.525, and then the slowest etch rate is 0.95 times the nominal etch rate. It's 5% slower than the nominal etch rate.

So the time to etch this, worst case, the thickest region, assuming you had the slowest etch rate, is just this 0.525 divided by the slowest etch rate, so divided by r_{Ox} times 0.95. So the worst case etch time would be 0.553 divided by r_{Ox} minutes.

So now, again, we express the overetch as a percent in time. So that's just going to be the worst case etch time, which we just calculated, divided by the nominal etch time. The nominal etch time would be the time to remove half micron at the nominal etch rate. So the nominal etch time is given by 0.5 divided by r_{Ox} .

So we just take the ratio of the worst case etch, the longest etch, to the nominal etch, and we get a 1.1, so that's about a 10% overetch, which isn't bad. 10%-- to do a 10% overetch these days is pretty darn good. Most etch-- most films have a greater variation than plus or minus 5% in a lot of plasma tools. Plus or minus 5% variation is conservative, so it's not unusual to see overetches of that length.

Now, the second thing-- basically, I'm leaving that wafer in the plasma etcher or in the solution 10% longer than I need to, to remove the nominal film. The second question is, now we know that, what selectivity of the oxide etch rate to the silicon etch rate is required so that, at most, we remove 5 nanometers of silicon?

Assuming we do the overetch that we calculate, assuming we do a 10% overetch in part a, so what we're asking here is this. I'm going to be etching down, and in some regions where the oxide is thinner, and because I have an overetch, I'm going to actually be exposing the silicon surface to the etchant. And it has a finite etch rate, but I don't want to remove-- I want to make sure I don't go into the silicon by any more than 5 nanometers.

So that sets a requirement on my selectivity. So the selectivity is how etch how fast am I etching through the oxide divided by how fast I'm etching through the silicon. If that selectivity is too low, if it's 1 to 1, then during certain regions during the overetch, you'll be zipping through that silicon substrate at the same rate that you're etching here.

So obviously-- and this is exactly the kind of practical calculation you'll need to know. If you have to say, oh, I don't want to attack-- I don't want to eat away the source drains when I'm etching the contact holes. Well, how much selectivity do I need if I'm only allowed to eat away, say, 5 nanometers, something like that, given a 10% overetch?

So if we go to slide 8, uh-oh, we can think through that here in solution b. So what's the most amount of silicon you would remove? The worst case will occur under the thinnest oxide being etched at the fastest rate. That's the worst case, so I have the thinnest oxide on the wafer, and let's say, at that point, it happens to be etching the oxide at the very fastest rate. That gives you the most amount-- the quickest that you'll clear the oxide.

At that very point in time then, the etchant starts attacking the silicon. So the thinnest oxide is 0.95 times half micron, so it's 0.475 microns. The fastest etch rate is 5% faster than the nominal r_{Ox} , so that's 1.05 times our r_{Ox} . So the time to etch through that oxide is 0.475, which is the thickness, divided by the fastest etch rate. So it's just this number here, 0.4523 divided by r_{Ox} . So after that amount of minutes, all of a sudden, in that one part of the wafer, the silicon is exposed.

So now, the time that the silicone is exposed to the etch is equal to the total time, which is this-- what we calculated before, 0.553 over r_{Ox} in part a, minus the time it took to etch the thinnest oxide. So that's the total amount of time then that the silicon itself would receive the etch. Assume it's etching at the fastest rate.

So the time the silicon is exposed is just what we had from part a, the total etch time, 0.553 over r_{Ox} minus the time it took us to get the silicon opened up in the worst case. So we end up with a silicon exposure time of 0.1 divided by r_{Ox} in minutes. That's how long the silicone, in the worst case, would be exposed.

So now, we know what we have to calculate. We want to etch only 5 nanometers, which is 0.005 microns of silicon. That gives me an idea of what the etch rate of silicon, the maximum, can be. It's going to be that 0.005 divided by the time it's exposed, which we just calculated. So the maximum etch rate of silicon ends up-- do the math, ends up to be 0.05 times r_{Ox} microns per minute, so this is always in terms of r_{Ox} .

And so now I can calculate-- given this, I can calculate the selectivity of oxide etch rate to silicon, just like we did before. So I just take this rate r_{Ox} divided by this etch rate of the silicon and divide it out, and you end up with a rate-- the r_{Ox} 's go out, you end up with this ratio has to be 20 to 1 or better. Yeah?

AUDIENCE: [INAUDIBLE]

JUDY HOYT: Which one?

AUDIENCE: [INAUDIBLE]

JUDY HOYT:

Oh, yeah, the time silicon is exposed? All right, well, that's fine. You can do the math to kind of get the idea. But what this says is, any time you need to do an etch, you typically have a maximum requirement on the film below that you can remove, and then you have a certain overetch requirement, which is set by non-uniformities and other things that can be set by topography.

That determines what kind of selectivity you need, so knowing what kind of selectivity you need, you can go out and figure out what etch it is you need to find in the plasma etcher. So typical selectivity requirements, we'll talk about this later next time, but for gate etch, typically you need 50 to 1, so something like that.

50 times faster etching the polysilicon compared to etching the oxide, and it can even be higher than that depending on topography and other requirements. That's a simple example. It just gives you an idea of how one thinks through these types of problems.

So one problem that we haven't talked about too much here, but shown on page 9, slide 9, is something called mask erosion. The fact that, during the etch, the mask itself may get etched, not just the layer underneath. So the mask can be eroded during the etch, and this can happen in both isotropic and anisotropic etching. And so in this example, a, you can see that the mask was eroded by a certain width or etched off Δm here, and that can happen if you're doing isotropic etching or anisotropic etching.

Again, mask erosion is bad because the feature size that you get ends up being different from what you would wanted on the mask, and it can even end up affecting the shape sometimes of the feature that you get. So mask erosion is something we want to avoid.

If you want to know about chemical etchants, you can look in table 10.1 in your text, and there are some common ones listed. There are websites you can go to for common etches, but again, as I mentioned, because they're isotropic, what chemical etches are very rarely used in mainstream IC, maybe for manufacturing some MEMS devices or some very specialized devices, but for mainstream CMOS, you don't typically find--

The only wet etch that you typically do find is HF, but that's not used for pattern transfer. HF is usually used as a dip during the RCA clean to remove a very thin layer of oxide. But that's not patterning. That's just cleaning off the surface.

So we're going to spend most of the discussion of chapter 10 then on this plasma etching, or this gas etching. It's fast and simple in most cases, and the real benefit is it's more directional. You can get anisotropic etching. Plasma etching can have both chemical effects, so very highly reactive species can be created in the plasma, and it can also have ionic effects, which would be directional so a sputtering action as you accelerate ions towards the surface.

So both of these can play a role, and depending on their ratio of the chemical effect to the physical effect, you'll either get less or more anisotropy. Here's an example of a plasma system. This should remind you very much of the plasma enhanced CBD. The only difference is we're not depositing. We're actually removing material.

Here's a typical RF powered system. Here, you have a gas coming in, and you have an etchant gas, for example, could be CF₄ or O₂ or some combination coming in, creates ions in the plasma that are highly reactive that can go down and end up etching off species from the wafer surface.

Here's just an example shown on slide 11. Again, this should be very much analogous to what you read about in chapter 9 on plasma enhanced deposition. Again, we have-- when we have a plasma, the plasma tends to self bias because of the differences between the mobility of ions and electrons. It tends to be self biased to be-- the potential to plasma tends to be positive.

And so also, the smaller electrode tends to have a higher voltage drop near it to maintain current continuity, so the trick in the plasma is you put the-- if you're doing etching, you put the wafer here, the target here, to be the smaller electrode. There's a large voltage drop then. On the electrode towards the left here. This voltage drop will accelerate the ions towards it, and you can get etching.

If you're doing sputter deposition, you do just the opposite. You put the wafer over here on the large electrode. You sputter away the aluminum target, and it goes to the wafer, but I'm trying-- again, now, I'm trying to etch material on the wafer, so you put the wafer on the smaller electrode during RF plasma etching.

Slide 12 shows some typical reactions and species that might be present in a plasma. For plasma etching, a particular example, assuming we're flowing CF_4 which is a very common species used to do etching. Here, there are different types of processes. You can get ionization, dissociation.

So the CF_4 can be banging around into electrons, and you can dissociate it into a free radical, CF_3 , and free fluorine, both of which can be highly reactive. And then this, you can have ionization of those species to create CF_3 ions or fluorine ions, so you have all these different both neutral and ionized and very reactive species present.

Typically, we have something like 10 to the 15th per cubic centimeter of neutrals. 1% to 10% of those neutrals may be free radicals. Remember, free radical is highly reactive. It's got an unsatisfied bond, and so it's a highly reactive species. It's very important for etching. And then you have a certain number of, say, 10 of the 12 or so per cubic centimeter ions and electrons. This isn't a standard parallel plate plasma system.

In a standard system like this, just two parallel plates with an RF voltage, the plasma density, so the number of these reactive species per unit volume, is usually determined by the voltage drop, the sheath voltage, so to speak. So increasing the power of the plasma increases the plasma density. So you more species and a higher etching, but you also get a higher voltage drop between the plasma and the wafer, so you get more damage.

So there's sort of-- in this type of reactor, you don't have as many knobs as you might like, so there's a trade off here. You can etch faster by tweaking up the RF power, but you have to be careful because you're also increasing-- tendency to increase the ion energy. And increasing ion energy may or may not be desirable from a damage point of view.

So I wanted to go on slide 13 and talk about the different plasma etching mechanisms, and there are basically three. There is chemical etching, which just like in wet etches, there are chemicals or species present, which do isotropic etching. They etch the same in both directions. These would tend to be-- typically not ionized, not ions. But they give very good selectivity because they have chemical specificity.

There is physical etching, which is anisotropic. It tends to be ions that are accelerated down towards the wafer, and so they have a directionality associated with them, but less selective. If you're just banging down on the wafer, it doesn't matter what material you bang on. You're going to remove some of it.

Ion enhanced etching is sort of a combination that we'll talk about with our ionic species that are present that also have chemical action. So these are the keys. This is where we get to reactive ion etching. You can get both anisotropy and selectivity to a certain extent at the same time. Most etches today use reactive ion etching or ion enhanced etching, but let's go through these three different mechanisms.

The first is chemical. As you might imagine, it's typically etching by neutral species, free radicals, so for example, you're creating this extra CF_4 , which is very reactive, and CF_3 , which is very reactive, and fluorine, which can then-- this free fluorine can come down, hit the surf-- or come to the surface, react with the silicon at the surface, and create C_4 , which is volatile.

You can add things to this to etch silicon a little bit more efficiently. You can add oxygen, O_2 , and that helps react with the CF_3 and reduces the recombination of CF_3 plus fluorine to go back to CF_4 . So that helps-- adding oxygen tends to push this more to the right, so you get a higher etch rate.

So a very common chemical etch that is fairly isotropic is to flow CF_4 plus O_2 . That etches silicon in a reasonably isotropic manner because it doesn't depend much on ions. It depends just on this surface chemistry.

So as I indicate on slide 14, these processes, these chemical equations, are really just chemical. They are isotropic and selective like wet etching.

So here's an example of, you have some species being created in the plasma. They are transferred to the surface. They adsorb on the surface. We have this free radical. They react with the surface or the film, and then they come off as a byproduct.

Just like we talked about in deposition, we usually characterize the arrival angle distribution of these species coming down the etchants. Usually, there's some kind of cosine theta to the end if it's isotropic, n equals 1, and a certain sticking coefficient usually very low. The surface reaction takes some time, and you can get some desorption. But you can use this type of picture to set up mathematical models of the process.

So that's chemical etching, and again, you're just creating free radicals. It's not that much different, in some ways, from putting the etchant in a beaker of wet chemical. On slide 15, I talk about physical etching, so what is physical etching?

Well, ion etching, for example, on the right, you could have, in the plasma, there are ionic species. The plasma-- there's a voltage drop between the plasma and the surface, so these ions can be accelerated towards the surface. So it tends to be more directional.

There's an electric field or a voltage drop, as we mentioned, across the plasma sheath right near the surface. We often model this as a sticking coefficient close to 1. The ions don't go [INAUDIBLE]. They just come in. You assume they come in. They react, or they come in, and they knock off whatever's on the surface.

Typical, in a physical etching mechanism, you might have species like CF_4 plus. Again, it's not really chemical. It's just coming in as a large massive molecule comes in and can knock something off, or argon itself you can remove by sputtering.

So it's not-- it tends to be not very selective. If you're doing purely physical mechanisms, it tends to remove the photoresist or erode the photoresist. It's hard to get selectivity because most sputter rates of all the elements are about the same. They're not that much different.

They can vary by a factor of 2 to 4, but we just did an example where we needed a selectivity requirement of 20 to 1. And other etches, as I mentioned, need 50 or 100 to 1, so purely physical etching is not going to give you that. But it is very directional. It is very anisotropic. The other problem with it is, of course, you can damage the surface because you're putting very energetic species in [? ion ?] and planting them into the near surface region.

So the third type-- so we did we did the chemical etch mechanisms, the physical mechanisms. The third type is this sort of ion enhanced etching, and it's the most confusing one, but it's also the most important. It's been observed that, a lot of times, that the chemical and the physical components of plasma etching are not independent, that they interact with each other in terms of the etch rate and the resulting etch profile.

And there are a couple of-- many examples of this. I'm just showing one here. This is a particular example of etching silicon, and this is a plot of the etch rate in silicon on the vertical axis in nanometers per minute as a function of time. And there are different gases flowed into the plasma etcher.

In the beginning, you're flowing xenon difluoride gas only. You do have a silicon etch rate, but it's slow. Now, you add an argon ion beam, so you're adding some physical component. You're adding some argon in there, and all of a sudden, the etch rate goes up by a factor of 5 or 10, something like that. So you get a large increase in etch rate when you have both an argon plasma ion beam plus the xenon difluoride.

And then you turn the argon off-- or you turn the xenon difluoride off, and you have argon only. And you get a little bit sputtering, but not much. So really, you only get fast etching, in this particular case, when you have both the ion beam and the xenon difluoride present at the same time. So there's a synergistic effect going on, and this is not uncommon. It's very typical.

We'll talk about some mechanisms. The nice thing about this reactive ion etching, or as people call it, is the etch profiles can be very anisotropic, and you can also get selectivity. So you can try-- you can get the best of both worlds.

How does it work? Well, there's a lot of different mechanisms. I don't think people know a lot of the details, but there are-- people try to model it. But some examples of this synergistic mechanisms are shown below, for example, here on page 17 on the left.

So you have some reactive neutral species that's in there. You have the ionic species at the same time. Well, what might be happening is there could be a chemical etch that's going on at the surface based on, say, free fluorine, or whatever it happens to be. So there's a chemical etch going on, but the etch rate that the chemical equation, the chemical-- the reaction may depend on the bombardment by ionic species. So there's sort of a synergy going here. In the presence of this bombardment, the etch rate goes up.

On the right, this is a little bit easier to understand. You might imagine that, during this etching, very often what happens is inhibitor layers form. So you'll have reactions that take place on the surface, but you have a lot of carbon, the CF₄. A lot of the etchings have carbon in them. They tend to form some kind of polymer that sits-- a monolayer or two of polymers that sit on the surface.

These polymers are there, but they inhibit the next reaction from taking place. In some ways, they inhibit the etch rate locally wherever they exist. But now, if you have ion species around, they can come in and bash up, they can bombard the polymers, and remove them only where they hit right at the bottom of the trench. The inhibitor layer can still be left on the sidewall.

So here's an example where you have both chemical etching going on. You have the formation of inhibitor layer, but only where the ions come down and hit do you remove the inhibitor, and then you get etching only at the bottom of the trench. So here's an example of reactive ion etching.

The presence of these little byproducts, which you might think are, oh, these are just parasites to the whole process. They're the name of the game when it comes to doing reactive ion etching and getting a very anisotropic profile. You want inhibitor layers to form on the sidewall because you don't want any sideways etching, but you don't want them to form on the bottom. Well, they'll still form on the bottom, but if you have ionic species coming down, you can remove them.

So this is how people do reactive ion etching and get such both selectivity, and the reason you get selectivity, because there's still a chemical reaction going on here. So it's not just etching-- it's not just sputtering. It's a chemical reaction and a little bit of removal of the inhibitor together that count.

So as you see on slide 18, this ion enhanced etching or reactive ion etching is hopefully the best of both worlds. You get good selectivity, and you can get good directionality. So although we may not know the exact mechanism, I kind of prefer, in my own mind, to have this inhibitor layer removal mechanism.

The two components, that is the physical bombardment part and the chemical reaction etch part, they act in series. So you get this anisotropic etching with very little lateral undercutting because of the directed ion flux. So RIE is one of the most common forms of etching that you will see today because you can get, to a certain extent, the best of both worlds.

Slide 19, talk a little bit about how to control the shape or what are some of the issues. Even with ion enhanced etching and this chemical etching, the slope of the resulting sidewall is not always perfectly vertical. There's some tricks.

It needs to be adjusted empirically. The ion flux may not always be perfectly normal to the surface. The ion flux may be coming at a slight angle. You could have bowing of-- that can cause bowing of the side walls of whatever you're trying to etch.

You can also get sloped sidewalls when the inhibitors, formed during the etch, have a high deposition rate relative to the etch rate of the inhibitor and the substrate. That's a little tricky, but on the next slide, I'll show you an example of how we get sloped sidewalls. So you need if you change the relative inhibitor deposition and etch rates and substrate etch rate by changing the chemistry or changing the specifics of the plasma conditions, you can control the slope of that sidewall.

So I just want to show, on slide 20, an example of how that works, and we kind of walk through these. There are two cases. One case on the left, I'm assuming-- we're assuming that the inhibitor deposition rate is very fast compared to the etch rate. On the right, the inhibitor deposition rate is relatively slow compared to the etch rate of the film, so let's see what happens in both cases.

If we start on the left, here's my mask. Here's the film I'm trying to etch. Now, just don't get confused here. The inhibitor deposition process and the etching process are happening simultaneously. I'm just showing them in separate views, separate steps, to help us think through it.

So let's just say, in the first time step, the first couple of seconds, the inhibitor deposits everywhere. So you have this inhibitor formation. Now, in the next time step, the inhibitor is removed, and you etch. So here, the inhibitor DEP rate is fast compared to the etch, so I don't etch very much of the silicon. Just a little bit, whereas on the right, I etched quite a bit because this etch rate of the silicon is much faster than it is in this case.

Now, I form, again, the next step, a little bit of inhibitor deposition here on both the left and the right. Here, the inhibitor DEP rate is pretty fast, so I have a thick inhibitor. Now, during the etch part, I etch just a little ways, whereas here, look, I etched a deep ways.

And you continue on doing that, and you can see that, in the left case, you're going to end up with a more sloped sidewall because the inhibitor is going down much faster than you are etching into the silicon. So you end up with this-- there's a certain amount of lateral etching that takes place, and so you end up with this kind of sloped sidewall. , Whereas in this case, where you've adjusted the inhibitor DEP rate, it's relatively slow compared to the etch rate, or the etch rate is relatively fast if you want to think of it that way. You end up with a much more vertical profile.

So the angle of the vertical profile depends on this ratio of how quickly the inhibitor forms the polymer, and how quickly it ends up being etched away, and the etch rate of the silicon, or whatever you're trying to etch at the same time. So very small changes. Usually, you do this by flowing multiple different types of gas and things. Very small changes in those gas flows can then change the angle which you achieve, and obviously, people try to tune this to get the right angle for their particular application.

So that's just an example of how RIE works and how you can control that angle, and that's a pretty critical thing to control these days. So slide 21, now, I want to go through a couple of different examples of the types of etching systems, and then we'll get back more into how we can do the modeling.

Over the years, a lot of different configurations have been developed, some of which make use primarily of chemical mechanisms, some physical, and then a lot of them, primarily, these days, are this ion enhanced or RIE mechanism. The most old fashioned type you will find in pretty much any thin film lab is an old fashioned what they call barrel etcher. It's called a barrel etcher because it's shaped like a barrel, a cylinder.

And you have here one of the outer circumference of the etcher is one of the electrodes, and then you have some kind of shield on the center as well. And you have an RF bias then here between one side of the barrel and the other, and what you end up with, this is purely because of the geometry and the chemistry of what you're using, purely chemical etching.

So you're primarily creating the free radicals that we talked about earlier that participate in these chemical reactions, and barrel electrodes are used today a lot in fabs, but not for anything critical. They're usually photoresist strippers, so a common way to strip photoresist off of wafers is put them in a barrel etcher and put an oxygen plasma.

Photoresist reacts as an organic. It reacts very quickly with the oxygen, and this process strips the resist without attacking much else on the wafer. And this is called ashing, just commonplace. So nothing-- the shape and design are not real critical. It's just a way to strip off certain species such as a photoresist.

The next level of complexity in a fab that you might find is shown on page 22, a system called the parallel plate system operated in the plasma mode as opposed to an RIE mode. So here's a plasma mode parallel plate system.

We have two parallel plates that form the electrodes, one with an RF power input, and then there's another electrode on which the wafers sit. And you put in some gases like CF₄ and O₂. They have roughly equal areas, or maybe the wafer, is grounded in the chamber. It might be slightly larger.

Here the sheath voltage may only be 10 to 100 volts, relatively low, so the energy in which these things come down is relatively moderate, say 10 to 100 electron volts. So there's not a whole lot of ionic component. You're not getting very high acceleration of the ions.

It's primarily chemical, so in a parallel plate system in the plasma mode is pretty much still chemical. The etching tends to be fairly isotropic. It gets us down and sideways about the same, and it tends to be quite selective. Again, it's mostly chemical type of etching, so that's a parallel plate system in a plasma mode.

Now, there's a second way to operate the parallel plate system shown here on slide 23, parallel plate in the RIE mode, Reactive Ion Etching. In order to do this, we need to get more directed etching. We need to have a stronger ion bombardment. We need to have a higher voltage drop between the plasma and the wafers.

So here, you make the wafers electrode much smaller than the outer electrode, than the other electrode. So the wafers sit on the smaller electrode. Here, the voltage drop across the sheet, instead of being 10 electron volts, is like 100 to 700 100 to 500 electron volts, so 10 times higher at least, maybe more, 20 times higher.

So you have much greater energy of the incoming ions. You also have lower pressures to get more directional etching, maybe 10 millitorr. Lower pressures means, what, a longer mean free path, less randomization, so the ions again get accelerated towards the wafer, very directional.

So this tends to be a more physical component than for the plasma mode. There's better directionality. A little bit less selectivity, however, for the RIE mode, but you can-- this is what you trade off. You trade off directionality for selectivity to a certain extent in these systems. Damage is going to be worse. A little bit more damage because you have higher energy ions.

Oh, on slide 24 is just an example of a photograph. It's a very old fashioned, way back 20 years ago, but still found in some labs is the Applied Materials 8100 Dry Etcher. This gives you anisotropic etching, this particular one, of films like silicon nitride, silicon dioxide, and polymer layers.

You can relatively control the oxide slope, the slope of a contact hole, slope angle, by using sidewall polymer deposition, just we talked about in that example of the inhibitor layer. The inhibitor layers are typically polymer, so sometimes people will call them sidewall polymers. The chemistry is fluorine based in this particular unit, so you would use things like CF₃, oxygen, SF₆, anything with fluorine in it, creating very reactive fluorine species.

The inside electrode where the wafers sit is a hexag-- hexagonally shaped. The outside barrel, which is shown here, is the outer electrode. So the wafers sit on the inside electrode, which is smaller, and so you operate in the RIE mode. There's a fairly large voltage drop across the sheath near the wafer surface.

So that's kind of old fashioned these days, although you'll see them in a lot of university labs, and even some fabs. There are RIE etchers here at MIT in the clean room here. The most sophisticated type of system today is a little different, though. It's shown here on page 25.

It's called the High Density Plasma etch system, or an HDP, and what this does is it uses a remote non capacitively coupled plasma source, such as a microwave electron cyclotron resonance or inductively coupled plasma. So the mechanism by which you create the ions, that mechanism is separate from the mechanism by which you bring them to the wafer.

So you have different voltages. Before, if you need-- in the old fashioned here RIE type, basically, if you want to get a higher plasma density and a higher etch rate, you drive up the voltage here, but that also means you have a bigger voltage drop to the wafer. So the plasma density and the damage to the wafer-- or the voltage drop here are directly related. You cannot decouple them very effectively.

In HDP, High Density Plasma, these things are somewhat decoupled, so you have a separate RF source here that biases the wafer. So this separates the plasma power or the density of the plasma from the wafer bias, the accelerating field. And so now, you can get reasonably high etch rates, but not to completely damage the wafer surface.

So on some of the characteristics are listed here on slide 26. You can get very high density plasmas, maybe 100 times higher density than you can achieve. Otherwise, you get faster etching.

The pressures tend to be lower even in the 1 millitorr range, 1 to 10. To get higher ionization efficiency, longer mean free paths, more anisotropic, more directional etching, so things do come down to the surface perpendicular. The nice thing is, about these systems, you get a high etch rate, pretty good selectivity, good directionality, so very vertical side walls.

At the same time, trying to keep the ion energy and the damage low because you have a separate power supply to create the plasma from the power supply that accelerates everything towards the wafer. So if you need to etch deep trenches, this type of thing, high density plasma is perfect. If you are doing optoelectronics, and you want to make waveguides, high density plasma where you need to etch through microns of material extremely vertical, perfectly vertical sidewalls, high density plasmas are perfect for that. Very high etch rate. Good control of the directionality, and you don't damage the surface too much.

Here's an example, actually, slide 10-7, of an HDP system. It's fairly old one now. This was back in the mid 90s, this was popular, made by a company called Lam in California on the West Coast. It was called the TCP 9400. It's, again, a little bit obsolete, but it was developed for high density plasma etching of polysilicon gates, specifically to etch the gate electrode and to stop on silicon.

So it was designed to have very, very-- give you very vertical side walls for gate etching, but not etch too much of the gate oxide underneath, so you get pretty good selectivity. So this etcher tends to be dedicated. In many fabs, they have a dedicated etcher just for gate etching and not for anything else.

So that's the RIE-- example of RIE etchers and high density plasma. They're all over the place in fabs. There is one other type of etching, not that common, but people do it, shown here on slide 28, and that's sputter etching or even ion milling. And this is purely physical. It's highly directional. It doesn't have much selectivity. As I mentioned, the sputter rates of the elements don't vary very much.

There's not a lot of chemistry going on. The nice thing about it, though, is you can etch almost anything. For certain materials, it's actually hard to remove them because they don't always have the right gas. But if you need to get the film off somehow, if you use a physical mechanism, you can sputter it off. Sputter etching almost always uses a heavy molecule or a heavy species like argon, an argon ion.

The bad thing about it, of course, it damages the wafer surface. It can damage devices. It can also do things like produce trenching, which is a non-desirable sort of effect. Here in panel a, I show what trenching is.

So this could be, here, let's say you're masking this, and you're trying to etch down here through this layer. Well, because, if you're using sputtering right near an edge, you can get the ions kind of bouncing off near this edge, and they can then come at a certain angle where they actually produce a little trench right near the edge of a feature. So you get enhanced etching right near the edge of a feature. You don't necessarily want that. You'd like to have the etch be straight across, so trenching is common.

Ion bombardment damage, you can get redeposition, so this can come in here, sputter off this film, and redeposit those atoms down here, where you don't necessarily-- you may not want them. There can be charging that happens if you have an insulator up here. And then as these ions come in, that can distort the ion path, and so your trenches can bow out, for example.

So the reason we bring up these sputter etching is not so-- and ion milling, not because people use them that much, the pure sputter etcher, but in any etcher system where they have a very strong physical component, a certain amount of this sputtering will take place. And if you're not careful, you can end up with trenching and things like that. So this, if you turn up the physical knob too much on your etch, you can start having things like this, which are indicative of that you're having a little bit too much sputter etching going on.

So slide 29 is kind of a bit of a summary of the different plasma etching types of systems. Here at the top are the most physical processes listed. At the bottom are purely chemical processes, so we go from sputter etching and ion milling to-- which is completely physical, to high density plasma, which is physical plus chemical. RIE, a little more chemical. Plasma etching, totally chemical. Wet etching is totally chemical. There's no physical process-- no bombardment inside the beaker.

So if we go from here to here, we think about it, the pressure is going down if I'm going in this direction. Anisotropy, so more anisotropic, it goes in this way. So wet etching is total isotropic. This sputter etching tends to be very anisotropic.

Selectivity actually increases down this way, so wet etching is much more selective. Sputter etching is totally non selective, and the energy of the process tends to be higher here for these. A lot more energy and damage than in wet chemical etching. The energies are just a few electron, or half electron volts, just chemical reactions. Just keep that picture in mind when you're looking at the different types of plasma etchers in the fab and when you choose which plasma etcher you want to use.

This slide just summarizes the different types of processes, all of which can be occurring during plasma etching, depending on the particular etcher that you're using, just as an example. So I'm assuming this top layer is my mask, and this is the film that we're trying to etch. There can be ionic species present, of course, that can cause some sputtering here and trenching.

There are reactive neutral species, like free radicals. They cause the chemical part of the etching right here, which can be isotropic and lead to undercutting. There can be mask erosion because the mask is attacked by the etchant, and that can change the shape eventually of what you end up etching.

Probably the most important-- one of the most important features on this whole slide, you can barely see, I guess we should have highlighted it here, is this little thin slim area called sidewall inhibitor deposition. That's probably the most mysterious and yet the most important part of reactive ion etching today. Without that, we could not do CMOS as we know it today.

The sidewall inhibitor deposition, what is it made of? Well, people don't actually know. Maybe some of the etched byproducts. It actually also depends on mask erosion. It turns out that the rate at which you form this inhibitor depends on the total amount of photoresist exposed that is on the wafer.

So people will develop a process using resist as the mask. It works perfectly. They get nice vertical sidewalls, good sidewall inhibitor formation. Now, they go and they say, I don't want to use resist as my mask. I want to transfer the mask pattern to an oxide, and use oxide as the mask. You might try that. So they no longer have resist in the etcher when they're etching the poly or whatever. All of a sudden, the sidewall inhibitor doesn't form because there's not enough carbon in the etcher. You don't have photoresist.

So sometimes, the presence of photoresist, you want a little bit of mask erosion just to contribute to the side wall inhibitor. The key about the sidewall inhibitor is that it's removed on a flat surface because the ions come down, and they can knock it off. It tends to not be removed on a sidewall. That's why it's called a sidewall inhibitor.

This is a possible mechanism why ion enhanced etching gives you such nice vertical profiles, because of the formation of that sidewall inhibitor. So those are some examples of all the different things that can happen during plasma etching.

Slide 31, I'm listing some manufacturing issues. These are fairly practical things, some general plasma etching conditions. So parameters that you have-- if you're going into the lab, an MTL, and you need to take an etcher, what knobs can you turn that you can directly control in a standard parallel plate system?

Well, you have the RF power. That's generally a knob you can turn. You have the pressure to a certain extent in the plasma etcher. You can control the gas composition, which species you decide to put in and their flow rates. So in a standard parallel plate system, you may have power densities like in this range, 0.1 to 5 watts per square centimeter.

If you want to increase the etch rate, you can turn up the RF power, so you get a higher density plasma. You'll get a higher self bias. That means the voltage drop between the plasma and the wafer or the electrodes will increase. That's going to increase the ion energy at the same time.

An HDP system, you have an extra knob. You have a separate power supply for the plasma and for the wafer bias, so you can get a very high plasma density with a more gentle-- without necessarily getting such high ion energies. Plasma densities are much higher here, up to 10 watts per square centimeter. The power density at the wafer, there, though, and the ion energies, look at them. They can be much lower, maybe only 10 to 100 electron volts here. So faster etch rate with less damage to your wafer in an HDP system.

How about the pressure ranges? These are very important here, shown on slide 32. This reactive ion etching, you typically would use something between-- the reactors are designed for 10 to 100 millitorr, and HDP system would drop that to the range of 1 to 10 millitorr, so maybe a factor of 10 lower.

So if you have a higher pressure, you're going to get more gas phase collisions. This decreases the directionality, or the other way around, HDP has a lower pressure, longer mean free path, more directionality. But you also can pay a price in terms of the plasma density. Usually, if you have a higher pressure, you'll get a higher plasma density up to a certain point.

Above a certain pressure, you get a lot of collisions between the gas molecules and the electrons, and that limits the energy electrons, and that limits the overall ionization rate. So pressure is an important knob that can be used to control not only etch rate in the plasma etcher, but also has effect on the sidewall profiles and things like that, and selectivity.

One thing we haven't talked about much, I've sort of implied implicitly here that, in most people's minds, people think of plasma etching as being a room temperature process, and that is more or less true. But because we don't usually intentionally heat the wafer except certain exceptions, like aluminum etching, where the temperature of the etch system, most of the time, we don't intentionally raise the temperature of the wafer.

However, the plasma puts a lot of energy into the process. You have a pretty high voltage drop. You have a large current flowing in there. V times I is power, so you're dumping power into this little space where your wafers are sitting. What does that mean? V times I power-- that power has to be dissipated somehow, and it comes out, a lot of the time, as heat. So the gas can get hot, and the wafer can get hot.

But we don't usually need to heat it itself in order to increase the etch rate or improve the process. For aluminum etching, they sometimes do heat the system up to maybe 50 degrees C or so, and that's to keep the species volatile because it turns out the species that tend to be formed during the aluminum etch tend to not be very volatile at room temperature, and you'll never get them off. You'll get these inhibitor layers, and the etch will slow down. So to remove byproducts, sometimes, you will see people heating the wafer.

This unintentional heating though, is much more of a problem. The wafers can get up to 100 degrees C without too much trouble. This needs to be controlled very carefully these days because that side wall inhibitor deposition tends to go down as the temperature goes up. So you never want to have the wafer getting so hot, self-heating, that the side wall inhibitor starts evaporating away.

Next thing you know, you lose your side wall inhibitor, and you're-- all of a sudden you have a very poor looking profile, very isotropic. And that can happen unintentionally just because you turn up the power a little bit too much, and all of a sudden, you're putting enough power in. The wafers are heating, or you change wafer size or something in the plasma etcher, and you get too much heating, not enough sidewall inhibitor, and then your anisotropy goes through the roof-- goes to pot, so you're in trouble.

So these days, the trends are to have better wafer temperature control. People have heat removal at the chuck. They have very careful clamping of the wafer, tightly controlled onto the chuck. Sometimes, they have helium gas flowing on the backside of the wafer to try to take heat away, just so you can control it because you don't want this etch-- especially if you're running the reactor through 25 wafers.

The first wafer might be kind of cool, but by the time you get to the last one, the whole thing has heated up. And next thing you know, from wafer to wafer, you're getting a lot of variation in the etch profile. So controlling the temperature is also more and more important these days.

Loading effects, there are things called macroscopic loading effects, and this is because you can deplete the etch and species across the wafer or across the whole etcher if you're doing more than one wafer at a time, so you'll see this. You put one wafer in. You measure the etch rate. Everything looks great. You have it.

Now, you put all 10 of your wafers in. All of a sudden, the etch rate goes way down, so now, you screwed up your etch. So in fact, when you measure the etch rate, you should have put 10 in, 10 dummies, because otherwise, you can have this loading effect, and you don't really measure the etch rate appropriately.

So it's kind of difficult to control. People get around it just by always using the same number of wafers with the same amount of material exposed every time. If you change that, you need to remeasure your etch rate. It's a pain, but it's a practical reality.

That's not that hard to solve. This one, though, on page 35, this micro loading, it's a little tricky. And this refers to the fact that the etch rate can vary over very small distances on the surface of the wafer. And why is that? Well, it's because the density of the open area, the area that you're etching.

The density of the area that's reacting can vary over small distances depending on the chip design. In one area, you may be doing a lot of etching in the chip. Another area, it might be mostly photoresist. So you have this sort of micro loading effect, so you get differences analogous to the macroscopic loading.

Well, what people do is they put in dummy structures intentionally, where you don't really need to do an etch, but you're doing etching of the active area there anyway just to make a more uniform appearance to the plasma of the surface. This one is even trickier, number two. You may have differences in the aspect ratios.

So you may have a portion of the wafer where you're etching a structure that's very closely spaced, and then you'll measure the etch rate here, shown in this SCM micrograph. The etch rate will be lower when you have a higher aspect ratio trench. On the same wafer right next door, a few microns away, the trench is etching faster where the aspect ratio is smaller, so there's a bigger opening between the trenches here. The ions can get down. The etch rate is a little faster.

So right next door, depending on the feature, density, and dimension, you're getting a different etch rate. This is sometimes called aspect ratio dependent etching, or people call RIE lag. This guy is lagging behind this one in some ways, so something to be aware of when you're etching complex patterns.

So why might this happen? Well, you may get depletion or trapping of the reactant species when you're traveling to a bottom of the trench. You can imagine these reactive species have to make their way to the very bottom. They may react the longer the distance they have to go, when it's a very tightly controlled tube, they may not have a chance to get to the bottom.

You may have distortion of the ion paths due to charging. You may even have shadowing effects. The point is that the net result is the probability of a reactive species getting to the bottom of a trench is very narrow, goes down as you make the trench more and more deep and more and more narrow. So that's something we need to take into account when you're doing these kind of dense patterns.

So slide 37 is kind of the summary of this introduction. We talked only a little bit about wet etching. It's primarily a chemistry problem, chemical reactions. Very good selectivity, but it etches the same amount in each direction, the same rate. It's isotropic.

Dry etching, there are two kind of species in the plasma that tend to be important, the reactive neutral species or free radicals, just free fluorine. It's neutral, so it doesn't get accelerated, so it's fairly isotropic, but it's also quite selective.

Ionic species, on the other hand, can be accelerated towards the surface. They tend to have more of a physical mechanism. It can be anisotropic, so they etch vertically much more than laterally, but therefore, sometimes not very selective.

In plasma etchers today, there are different mechanisms. There's purely chemical mechanism, which is based on neutral species. There may be sputtering or physical sputtering by acceleration of the plasma, and the most popular is the sort of ion enhanced etching, where ions enhance the chemical processes, or ions remove these inhibitor layers, which is key to keeping the etch going.

There's a lot of different types of plasma systems, just literally dozens of them. There's barrels. There's the plasma mode, RIE high density plasma, and there's even sputter etchers and lots of variations. So these tend to be purely chemical in the barrel range and purely physical if you're in the sputter etcher. Most of the systems these days operate in between those two limits, a little bit of chemistry going on, and a little bit of the physical etching.

So that's about all I have to say introduction. Please go ahead, read chapter 10. Next time, we'll talk in detail about how you etch polysilicon gates with good profiles and good selectivity. Also homework 5, your last homework is due. You can put that in the folder.