

MITOCW | MIT6_774F04_lec07.mp4

JUDY HOYT: The curvature of the surface, and they want to look at the oxidation rate as a function of how curved the surface was. So this is a top view or a bird's eye view here of his structure in part A. But look on slide 3. Part B you're looking at a side view here. So you can see he had etched this cylindrical structure in the top. It's got a circular pattern. After preparing this surface, he then put it in the furnace and grew this thermal oxide on all those free surfaces, which is shown here in blue. That's the silicon dioxide.

So that was a real experiment. And then, just like we saw in the previous photo from Marcus and Shang, he put down-- he deposited polycrystalline silicon, and that's simply used for contrast during the SEM analysis and to protect the oxide. So he has got this polysilicon over everything, so from a top view all you would see is the poly. And then to reveal this oxide and what it looks like looking down, he then lapped or polished off the top, the poly.

So he's about halfway down the pillar, partway down the pillar. And then a top view is shown here in the lower right. You can what he would see. You'd see this polysilicon, which is, again, just a contrast medium. But you would see this white ring here is the silicon that was initially oxidized, and the blue ring all around it is the oxide that grew on the inner part of the cylinder and the outer part of the cylinder.

And then you can vary the cylinder size. He can create a whole lot of these cylinders on the same wafer. And look at how the oxide grows, what its oxidation rate is as a function of the radius of curvature. So it's a very nice scientific experiment to look at these dependencies. If you go to slide number 4, these are some actual-- the upper left is a scanning electron micrograph. Not transmission, but it's an SEM, so we're looking at backscattered electrons.

And there is some contrast associated with the different materials. And this is sort of a typical result of a plan view SEM. And an artist's drawing of what is seen in the SEM is shown down here. So just looking first at the SEM, this wall right here-- so now we're looking down on the chip or on the surface of the wafer, and he's etched a wall. Right?

Looking down along certain directions. This is probably along the 110, the 110 equivalent, 110 planes. He's also etched these regions here. Look at on the upper right. This is a cylinder of silicon. This darker region is the SiO₂, and surrounded by that is polysilicon. If you want to look down here, it's a little bit easier. Everything's labeled. The thing you see right away is that the oxide grown is thinner for both a concave and a convex corner.

So here's an example right here where the silicon was rounded, but it's in a concave corner. And the thickness of the oxide, which is this dark region, is much thinner here than it is on a flat surface here or here. So it's thinning up when you're oxidizing a corner, and even a convex corner as well, the oxide growing is a little thinner.

In fact, the effect seems to be a little more pronounced for this concave corner where it really looks like there's not much oxide grown here at all. So let's say we're just talking about the oxidation of the cylinders. Let's say just the silicon cylinder, this pillar right here, and there are differences in the oxidation rate along the different directions. What would you think? People have any suggestions? Anybody have any ideas? If I take a cylinder of silicon that was etched-- the wafer was originally 100 and I have a cylindrical surface.

Let's say I'm growing in the thin oxide regime. Why do you think you might get different thicknesses going around the perimeter of the cylinder? What's one reason? Any ideas? How about thinking about what's the orientation of those planes as I make a cylinder. If I etch a cylinder, I'm exposing all different planar orientations. So locally, at each point in the cylinder, there's sort of a different crystal plane that's exposed.

So you might expect, if you're in the thin oxide regime, at least there might be some differences due to the fact that I'm looking at different orientations. So that's one effect. Any other ideas on any other potential effects that you might think about? Etching a small cylinder compared to etching a flat wafer. Orientation is one.

All right. Well, let's go ahead and see, then we'll give the answer away in the next couple slides. In fact, you go on to slide 5, and it doesn't completely give you the answer but it shows you at least the results that Gao got. This is a plot from his work, and the y-axis is the normalized oxide thickness. So this is normalized thickness to the thickness he grew on the flat surface of the wafer. So he measured the thickness on the flat part and then he measured the thickness on these cylinder sides and he just divided.

On the flat part he grew about 500 nanometers or half a micron of oxide on all surfaces. And by the way, a flat surface here-- he's plotting this normalized thickness as a function of $1/r$, where r is the radius of the cylinder. So if you're on a flat surface, the radius of curvature, the radius is infinity. So $1/r$ is zero. So right at this point zero, this corresponds to-- and all the curves converge, that means you're oxidizing a flat surface.

As you go out to higher and higher, $1/r$, that means that r itself is getting smaller so you're getting a tighter and tighter-- a smaller diameter cylinder just so you get yourself calibrated. And he's got two different types of curves here, convex radii and concave. If we just go back one-- whoops. One way, one slide to slide number 4, here's a concave cylinder.

So you're oxidizing on the inside of a surface, like you're in the cave here so you're oxidizing. And here's a convex surface. You have silica and you're oxidizing on the outside of the cylinder. So you can do it either way. You can imagine just geometrically, knowing that oxide has to expand, it would be different if you're oxidizing inside a cave, in a concave area, versus outside on the outer surface. And in fact, you see this. Look at the concave radii, all these dashed lines compared to the solid lines.

The normalized oxide thickness at a given temperature, say here at 900 degrees, is a lot thinner for the dashed compared to the solid. So a concave cylinder, he found a lot less lower oxidation rate, oxidizing the inside of the cylinder as opposed to the outer surface. And in addition to you see that difference right off the solid versus the dashed, the other thing is look at it as a function of temperature. 1,200 degrees, it's pretty darn flat. 1,100, almost flat.

So at those temperatures, the oxidation rate doesn't seem to depend very much on the curvature of the cylinder at all. It's oxidizing equally rapidly. But at low temperatures, so should I go down to right between 1,000 and 900, you get a big retardation in the oxidation rate. So there's a much more pronounced effect when you're at low temperatures. That's kind of consistent what we saw in Marcus and Shang, the 950 degrees C pillar oxide had a much more of a corner effect.

So this is scientifically in agreement with that qualitative result. OK. So this was the actual data that he got. And given this data, then he had to come up with some kind of explanation. That's what we're going to talk about. What are the physical mechanisms to explain his results? The first one we already said because we know it from last time, relatively simple, is the fact that the crystal orientation of the silicon is changing along the surface of a cylinder or any surface that's not flat.

And we know in the thin oxide regime, that impacts the oxidation rate. So at sufficiently low temperatures or when he was growing the thinner oxides, that's going to have an effect. Now, of course, it's said that on the flat regions he was growing 5,000 angstroms. That's not particularly thin. So depending on the temperature which he was growing that at, you probably don't expect a huge effect in his results from crystal orientation. But nevertheless, you have to take into account.

If you're oxidizing a trench and you're growing a very thin oxide, the orientation will impact the rate. We saw that last time. These other things, though, are more telling and more having to do with what he actually saw. The two dimensional diffusion of the oxidant, you need some kind of numerical technique to solve the diffusion equation in multiple dimensions. So if I just go back a couple of slides-- let's go back to slide now, say, number 4 and look at oxidizing this concave corner.

As the oxide grows, the oxygen or the water, whichever is he's using as the oxidant, has to diffuse in two dimensions through this structure. So any changes due to orientation or whatever in that thickness in which it has to diffuse through are then going to affect-- eventually will affect the diffusion through that. So we have a two dimensional problem now. Before we had a simple one dimensional. Deal-Grove we just had to diffuse from the top surface to the interface.

Well, we have to do the same thing but now in a two dimensional problem. So as the shapes change, we'll see diffusion of the oxidant will take effect. Let's go back then to slide number 6. The third one and probably one of the most important in explaining Gao's results is the stress due to the volume expansion. We know that oxide layers that are formed on silicon are under compressive stress. OK?

So they're kind of being compressed to get onto the wafer. Even when they're in the planar case, right? Just on a planar silicon wafer surface without any curvature. These stresses can be increased quite a bit on a curved surface, because the volume expansion is going to be confined dimensionally. So here's a way of picturing this on a curved surface. So the central region that's colored in gray, let's say that's my silicon pillar.

Let's say I've grown a certain amount of oxide already that's on this outer region. That's the oxide that I've grown, and now I'm just growing a little bit of new oxide. Remember, in oxidation the reaction takes place at the interface, not at the surface. So to grow this little band here of new oxide, let's say a few angstroms, 10 angstroms that I'm growing, the oxidant has to fuse into the silicon surface, has to react at the surface, and then the new oxide has to push out on the old oxide to make room for itself.

It's got to push up and push down and consume some of the silicon. When I was on a flat silicon wafer, the oxide could just push up and push out a little bit at the edge of the wafer if it had to. But I'm not. I'm on a curved surface here where the oxide is completely attached to itself, so to speak. So it has to actually expand along the direction of the circumference in this direction.

So it has to do work, more work on the oxide above it than it would have to do on a planar surface, because it's really confined all the way around that surface. Really, the curved surface has got to flow in order to get this to happen So you have the solid material this glass that has to essentially flow. So there's a certain amount of stress built up as a result of that. The deformation of the surface can be quite large. Forget about the cylinder for now.

That was Gao's experiment. Let's do something a little more realistic, like LOCOS. What is LOCOS? Remember, LOCOS is you start with a silicon wafer. You pattern part of it with silicon nitride under which the oxidant-- through which the oxidant cannot diffuse. So you're locally oxidizing. You're not oxidizing over here on the right in this lower right picture. You're oxidizing over here, only over here on the left. So you're locally growing a thick oxide.

But look at the surface. The top surface of this oxide here. Remember, at one point when it first starts out, it's relatively flat. This top surface has to stretch, in order to maintain this shape, by as much as 15%, 20%. So the oxide is itself having to stretch and deform. All that involves doing work and stress buildup. Even, for example, also the nitride itself is pushing down on the oxide. In order to lift this nitride layer, the oxide has to do work.

So there's definitely stress buildup in these regions near corners or on curved surfaces, and that slows down the oxidation rate. So how does it slow down the oxidation rate? Well, let's look, for a moment, at slide number 7. These stresses can impact a couple of things, both transport of the oxygen through the oxide, like a diffusion process, and the interface reaction rate.

And here's an example again. This is showing a small portion, a quarter or a section of a silicon surface that's curved because it's been etched or whatever, and the oxide that's already been grown here is on the outside, this annulus shown in the medium gray. This little white annulus, this little white region, the nascent oxide-- nascent meaning new. It's the new little piece of oxide that's grown. This little growing oxide has to do mechanical work against the oxide above it, because it has to push it out, and it's constrained.

So this mechanical work actually can modify the energy, the activation energy of the process, the amount of energy that it takes to do that process at a given temperature. So what we write is EA, the activation energy for the surface reaction rate under the stress case. When we're oxidizing in a curved surface, it's going to be equal to EA on a flat surface, whatever that might be-- say two electron volts-- Plus some extra terms.

And there's a term that depends on the normal stress. σ_n is the normal stress. Normal, so it's in this direction, in the radial direction. The σ_n times some fitting parameter V_R plus σ_t is the stress in the tangential to the growing surface, so along the circumference, times some other fitting parameter volume V_t . So then k_s in the curved or stress case is just the ordinary k_s times these two exponentials-- again, this is an activation energy-- minus $\sigma_n V_R$ over kt and a σ_t term.

So basically the surface reaction rate is going to go down. The amount it goes down will depend exponentially on these two stress terms, and these are fitting parameters that we put in-- essentially you can put in the model. And so you can see then how you can introduce a temperature dependence to this effect. Depending on how strongly temperature dependent is will depend on the numbers you put in for V_R and V_t .

OK. So let's go on to slide number 8, and that was the surface reaction rate. You can intuitively understand how that surface reaction might be affected by the fact that it takes more energy. You have to do work in order for that reaction to take place that you didn't have to do on a flat surface. How about the diffusivity and the solubility? Well, how could that be affected by stress?

Well, this is meant to be a highly schematic picture of a ring structure that would exist in silicon dioxide where these black regions of black circles are the silicon. We have oxygen here as well taking up a certain amount of space, say, in the bonds. There is some sort of region in here where there is sort of free space or interstitial space where the oxidant can diffuse through. And let's say p is a hydrostatic pressure in the growing oxide, so there's a certain amount of pressure in the solid.

Well, as you put the solid under pressure, it reduces that. You can imagine it reducing the interstitial space in the network. So you have this network. Although it's somewhat random, it has this structure, and inside the network there's this open space. As you put the whole solid under pressure, but that bond's compressed, the amount of interstitial space in the network is going to be reduced to a certain extent.

And if you're oxidant, your water molecule, your oxygen molecule has to diffuse through that interstitial space. You can imagine the diffusivity or the diffusion rate is going to go down. As I compress on the oxide, it makes it harder for the molecules to diffuse through. It might even affect the solubility, as you might imagine, hypothetically, that the amount of water vapor molecules or oxygen you can stuff in there maybe would be affected by how much pressure this solid is under.

So again, mathematically we can do what we did before. We write the stress diffusivity as the unstressed-- or not unstressed, but the normal planar diffusivity through the oxide times something that goes exponentially it depends on the pressure, which is the hydrostatic pressure is due to the oxide stress, times a constant, which is a fitting parameter V_D over kt . So again, D will go down as you go under stress, diffusivity. And you might hypothesize the same thing for the solubility, a similar dependence.

OK. Let's go on to slide number 9. So just as shown in your text, I'm just summarizing then the impact of stress on our basic oxidation parameters that we know from our Deal-Grove model. The parameters from Deal-Grove, $k_{\text{sub } s}$, the reaction rate D and C^* all get modified in a way that depends on the stress in an exponential fashion. So it's very sensitive over kt . So the stress effect will be different at different temperatures.

So these parameters, these V_R , V_t , V_D , and V_s , they're reaction volumes. They are designed in SUPREM-IV. They are put in as fitting parameters. So they have some default values, but people typically modify them as needed to fit the shape that's observed experimentally from LOCOS or from a pillar or from your structure.

In practice, this number $V_{\text{sub } s}$ ends up being zero, which essentially people believe that the solubility is not dramatically itself impacted by the stress. And so this extra term here, $e^{-\frac{P V_s}{k t}}$ equals to one. And so people generally believe it's the $k_{\text{sub } s}$ and the D that are being modified by the stress.

OK. So let's go on to page or slide 10. There's one more parameter that we need because we said the oxide is a glass and it's flowing, and the glass flows differently or its ability to flow depends on temperature. And that is expressed through a viscosity parameter, η . And in fact, the stresses that are in the glass or in the oxide are high enough that the viscosity itself, its rate of flow actually is a function of the stress in order to model it accurately.

So this is an equation that's been derived and that is used in SUPREM that gives reasonably good agreement with the experimental observations. So the viscosity as a function of stress is equal to the stress independent viscosity η , but it's a function of t , temperature, times these sort of parameters. A shear stress in the oxide, so it's actually linearly dependent on the stress in the oxide. So if this goes up, the viscosity goes up, and something that depends on the hyperbolic sine.

So that's just an equation that people have derived more or less empirically. So we have these parameters that are changing with stress. So let's go on to slide 11. You can probably get the impression by now there's no way you can easily do these calculations by hand. Deal-Groves, piece of cake. You sit down. You just simply solve or integrate the equations. They're all analytic. These are non-analytic.

The shape of the growing oxide changes with time. $k_{\text{sub } s}$, D , and η all change with stress, and the amount of stress is changing with time as the oxide has to push on a thicker and thicker thing above it. So with all this time dependence, you really need a numerical simulator. You need a computer to do these calculations for you and to integrate them in time. This is exactly what has been done in SUPREM-IV, which can model the shapes of oxide on curved surfaces.

You'll have a chance to do this on one of your homework problems, hopefully. And this is a reasonably common thing in IC technology, that you don't typically always grow an oxide over the entire surface. Very often there are patterns on the surface. In fact, the very first step is to isolate the active area from the non-active region, and it's often done by the process called LOCOS we talked about the first lecture. Poly buffered LOCOS, slightly fancier version, and shallow trench isolation.

So here's an example of LOCOS. Remember, you start with a silicon substrate. You put down a thin oxide, which is a pad oxide, which reduces the stress on the silicon due to the silicon nitride, and you put down silicon nitride and pattern it. Silicon nitride is a perfect or reasonably good mask, as oxygen cannot diffuse through it. So the oxidant diffuses in here and only oxidizes on the left side.

So this is what the local structure just in your mind, as a starting structure before you do oxidation, that's what it looks like. Let's go on to slide 12 and see what that LOCOS structure now looks like in a simulation, the SUPREM simulation, illustrating the stress effects. So there are two simulations that are shown here and let's start on the left. And what this is, the y-axis is in microns so that's vertical in depth from the silicon surface, and the x-axis is in microns as well.

This point right here corresponds to where the mask edge is. So we're doing a two dimensional simulation. And what you see is this upper region is the silicon nitride, as you can see, and the yellow region underneath is the silicon surface. Look how the silicon surface is shaped in this case. And the field oxide is out here. Now an important parameter that people care about, if you're scaling silicon circuits, is you want to be able to put the devices closer and closer together, as close as you can get them.

So in order to be able to do that, you need to grow the field oxide where you want it, and you don't want field oxide in the active region. Any part of the field oxide that encroaches or pushes its way into the active region means that it's scaling the active region in which you can make your device so you don't have as much control. You don't get a nice, well defined active region.

So in this model, you can see what's happened is the field oxide has encroached in. The mask edge start here at x equals zero, it's encroached in from x equals zero by about-- oh, about, say, half a micron. The oxidant has diffused in here and has oxidized underneath the nitride. So you have this region where the silicon is sloped and the encroachment distance, I would quote here as being half micron. So your active device really starts here in this region.

So this is the model run, though, without stress. And people did these models and they found, oh, it doesn't really agree with what we get experimentally when we grow LOCOS. We do a cross-section and we look at it. It's not what it looks like. In fact, what they found it looks more like what's on the right. And what's on the right is the SUPREM-IV model, but this time turning on the stress dependence. OK? So you have the diffusion rate through the oxide, the surface reaction rate are now all going down as a function of stress.

And where is their stress? Well, there's a lot of stress in this region here, in this corner region, because the oxide has to push up against this nitride film. So what you see here is look at the amount of oxide that was grown here in this bird's beak area. There's quite a bit of it has grown, the thickness here. Compared to, in the bird's beak area, when you turn the stress model on not so much oxide has grown. OK?

And as a result, with stress, the amount of encroachment, which I'm marking between these two lines now here, is quite a bit less. It's about a factor of 2, maybe. Here it encroached-- I'm sorry. I said 0.5. It actually encroached by 0.6 microns. The encroachment here, according to my eye, is about 0.3. So it's about half. That's because stress was included in the model. This turns out to be a lot closer to what people actually measured when they do LOCOS oxidation.

That's for LOCOS. You say, well, not too many people do LOCOS in manufacturing. Well, there is still some LOCOS or LOCOS-like structures being done. People nowadays do shallow trench isolation or STI. Well, we still get corner effects and stress effects at STI. And I've taken this from an article or a short course given at IBM back in 1998. And what is shallow trench? Remember what we have.

We have an active region here shown in the center under the gate, and isolating the devices we have field oxide, which is this thick oxide on either side all the way around it. In fact, goes all the way around. We've cut a cross-section here. And we do this by first etching a trench in the silicon and then doing some oxidation. We do a little thermal oxidation of that sidewall.

And the engineers, the electrical folks, are very interested in the exact shape of this corner, because it turns out, depending on the exact shape of that corner and how thick the thermal oxide is, you'll get electric field spikes or build ups right there, or you'll get a region where the oxide is weak. And tends to break down and when the gate goes over that, the device can have breakdown effects, or you can actually cause unusual transistor effects just based on the shape of this corner and this oxidation.

So there are two different simulations shown here. Up on the upper right is without any stress. And you see the thickness of the oxide as you go around the corner is reasonably uniform. It doesn't change too much. With stress effects, look at the corner. The oxide thickness is quite a bit thinner here. It's thinned up in the corner region, compared to on the flat surface or on the vertical surface.

And this is a problem because that oxide gets too thin, you're going to tend to have breakdown effects or it's going to affect the electrical properties of this device. So it seems maybe to your eye initially, oh, what's the big deal of that? But when you go to simulate how these devices electrically behave, it makes a big difference. So the stress effect here comes in because, again, you have this curved surface that needs to stretch, and there's stress buildup whenever you have a corner.

That tends to lower the oxidation rate. And that's also been observed experimentally. So let's go on to slide number 14. Stress also has another interesting effect. Let's say you're not doing shapes. So you don't do LOCOS, you don't do shallow trench. You're a very simple process person. You only oxidize flat wafers. OK. Fine.

You can still see stress effects in the form of the history, of what the history of the oxide itself has gone through, what thermal history. Because we know even planar growth has some intrinsic stress during the oxide growth. The oxide is under a certain amount of stress even when it's planar. It's just under more stress when it's curved.

But this stress can relax upon annealing. At a high enough temperature, above 900, 950, the glass can flow a little bit and the stress will relax. So when people measure the Deal-Grove linear parabolic rate constants that describe oxidation in this intrinsic stress state, the 1D stresses are already accounted for. OK?

But let's look at the experiment that's shown up on this slide. Let's say we do a two step oxidation. And you decide, OK, I take two wafers. I put one in the furnace at 1,100 degrees and I grow a certain thickness of oxide. Say, 1,000 angstroms or whatever. And I put another wafer identical in a furnace at 800. I grow the same thickness, 1,000 angstroms.

So I have two oxides, both the same thickness, grown at different temperatures. However, they look the same but they're not exactly the same. The intrinsic stress in this 1,100 degree C oxide is going to be less than the intrinsic stress in this 800 degree oxide because the higher temperature, the oxide can flow more. There isn't as much built in stress. OK?

So you say, well, this is under more stress than that. What's the difference? OK, maybe you'd see a little less, little difference in wafer curvature. That's a way you could measure it. All right. That's fine. Now you go to the next step in your process, says, I want to grow another oxide underneath these two at 800 degrees.

Now, if I hadn't said anything about stress, you'd say, well, Deal-Grove would say I have 1,000 Angstrom oxides here, 1,000 here. I put them both in the furnace. You just use Deal-Grove, you would get the exact same thickness on both wafers. Deal-Grove doesn't tell you anything having to do with the history of how you grew the prior oxide layer. It just depends on-- remember, just the $t_{sub i}$, the thickness, the initial thickness. It cares nothing about what temperature that oxide was grown at.

In reality, though, if you go to look at the wafer that comes out of the furnace after this 800 degree C, these two wafers will have different thicknesses. And so how can that be? They're both in the same furnace. They're sitting right next to each other. They had the exact same thickness to start with. Why would it matter? Well, actually the wafer oxidized at high temperature is going to grow faster compared to the 800 degrees, because it has lower stress levels.

So it has a little bit more open network. So you can diffuse through that oxide more rapidly, and it's a k sub s . Its surface reaction rate is going to be a little bit faster because it doesn't have the stress in it. So you can see these differences even if you're not doing shape surfaces, just in planar surfaces. Now we chose a very extreme example. 1,100 degrees, which is quite hot, above the flow point of the glass at 800, which is quite a bit different.

Maybe you typically wouldn't do such a wide range, but it really does tell you that the history effects can be important when you're trying to do accurate modeling. And in fact, most simulators don't take this into account, so it's something you'd have to account for, either empirically in your process or calculate it yourself. That's sort of an interesting effect that goes beyond shape effects.

So let's go on to slide 15, and we're going to talk a little bit about the evolution of what's called recess LOCOS. Before, earlier today I talked about LOCOS and we just said we put down a mask, we have a flat surface to begin with, and it becomes a little strangely shaped as we oxidize it. Remember, we had quite a bit of encroachment. Even 0.2 or 0.3 microns of encroachment is too much.

And not only that. If you look back-- let's see. I'm going go back a couple slides to, say, slide 12. When I take the nitride off and I go to use this as a processing my chip, the surface is not very planar. So I have to do photolithography now on this surface. It's not very planar. I have this big oxide hump sticking up here in the field, and the field is much higher here. Look at the field. It's up here at 0.3 microns above the original surface.

So I got this big mountain sticking up over in the field. Photolithography doesn't like to be done on surfaces that are rough or surfaces that have big steps. The depth of focus from your camera or whatever, when you're trying to focus down, the smaller the feature size you're trying to focus on, the depth of focus becomes more and more difficult. And so you the lithography tool might not be able to focus on this valley when it's well focused on the highest part of the mountain.

So there's a big requirement with litho today, forcing us to make all our chips so we can keep the surfaces as flat as possible in every litho step. So a step like this that might have been acceptable, 0.3 micron step years ago when people were making large devices and the lithography tools didn't use really high power lenses, now we're using a high power lens, we cannot tolerate such a big step. So there's always a move in silicon industry to keep the surface as flat as you can. LOCOS doesn't do that.

So before people invented shallow trench, STI-- which is, again, a very planar process because you use a polishing process to polish things down at the end-- they did something called recess LOCOS. So this was an attempt-- well, you say, well, you have this big mountain growing up on the side. Let's give it a handicap. Let's etch the silicon in the recess to begin with so that as the oxide grows, when it's finished, it'll end up being pretty flat.

And so you try to compensate for that. And this is an example of how that works, and I'll show you it doesn't come out exactly as flat as you want. But let's say this is a starting structure. So I have nitride over here, little pad ox underneath, and I've etched this recess over on the left. The recess is maybe 0.3 microns deep, and I'm doing that so that the field oxide will be recessed and I'll get a flatter surface. And then I go after certain timesteps.

So this is after I've grown a little bit of oxide in the center here, and you can already see, even for a thin oxide, immediately some corner effects. The oxidation rate is a little slower here and here compared to on the flat. And on the far right, you can see I'm starting to-- the corner effects are still evident and I'm starting to develop this bird's beak thing. The stress of a nitride is causing the oxidation rate here to be also a little slower, and the fact that it has to diffuse through.

So let's go to the next time step. This is sort of a time evolution, which is easy to do in the computer. Here's the oxide again after a certain period. Once more. And finally, after 90 minutes, the structure is almost flat. So I achieved what I wanted in that in the flat regions out far away, in the field region here, the surface of the oxide is at zero or close to. Not zero, but it's at a height that's just about equal to the surface of the silicon. So that's flat.

Once I strip the nitride off and etch it, I still have this little bump except for the bird's head. And if you look at this, you can see it looks like a bird's head and a bird's beak and the bird's eye would be right about there. So that we really can't get away from, and that's because the oxidant diffuses underneath that nitride, causes some oxidation, and the exact shape of the bird's head depends on, of course, what kind of stress model you use.

But it's better than it was before. The lithography has a flat surface here, a flat here, but it's got this little bump here, which is not so easy to get rid of. So the next technology people develop is shown on slide 17. People thought of this other effect, other idea called the Swami process. This was done at Hewlett Packard, and this was even a little more sophisticated.

They were trying to keep the oxidant from getting around. Remember, if we go back here-- let me go back for a second. The starting structure here, the problem is that at this corner, the oxygen can diffuse straight in and, right off the bat, start oxidizing, and that's going to form where the bird's eye or the bird's head ends up. So they thought, all right, let me put this nitride mask all the way down this the edge, and that'll keep the oxygen from going up there and reduce that amount of oxidation.

So this was the Swami structures. They're very similar to recessed LOCOS. You create the recess, but they create an oxide mask here, and they also put a paddocks-- I'm sorry. A nitride mask along the edge of the recess, and even at the very bottom, at the foot. So trying to minimize the oxidation in the active device region by masking the etched sidewall.

But they want to do this in a way that you don't build up too much stress, because after all, the oxidant will find its way over here and start growing in oxide. If the nitride is too thick here, it can build up too much stress. If you build up enough stress in the oxide, you can actually crack the silicon or introduce dislocations. They're really not cracks, they're dislocations, but that can be a problem.

So you notice they did this in a way that the nitride is a little thinner along the sidewall and down here so that the flap can be lifted without too much stress being induced. So we're going to take this structure now, put it in the furnace at 1,000 degrees, in moisture or water vapor for quite a few minutes and see what happens as we oxidize it. So here on the left on slide 18 is the starting structure. And now look after 450 minutes of oxidation.

This is the simulation. Unfortunately, I'm sorry about the colors, but I think you can see this top layer is the original nitride, silicon nitride. Look at the flap. The flap that was along this sidewall has been lifted up because the oxidant has moved its way under there and pushed it up, and then it's been-- it's quasi vertical over here. The flap has been pushed out by the oxidant that was coming under here from the left and pushing on it.

So the thin flap lifts up during oxidation, but it's still reasonably flat surface. Not perfect, but they're trying to reduce the amount of bird's beaking and bird's head that formed. That's the simulation. And in fact, on the lower left, on slide 19, the upper diagrams are the same as what we just saw. But if we turn to slide 19 in the lower left, they have actually-- SUPREM-IV form can actually be used to calculate stress contours, so it's actually calculating the stress in the oxide of course at each point, and at each point, assigning an oxidation rate.

That's how it knows how to calculate the shape. There is some stress buildup or concentration here here in this yellow region, so these different-- I don't have stress numbers associated with this, to be honest. But if you run SUPREM-IV, it'll tell you in each color what that corresponds to what stress level in the oxide. On the lower right, which is kind of interesting, this is an actual-- from the HP process, this is the actual experimental data.

So they grew this. They made this Swami process, and then they did a polished cross-section, scanning an electron micrograph. And this is the act of silicon region afterwards. This region right here is the oxide and the nitride. And you can see if you compare the simulation upper right to the lower right, the actual data, it's not that far off. It reproduces some of the key features. You can see this little divot here, this little V shaped well.

Well, that's also seen in the simulation. So the stress effects did a reasonably good job of getting the shape right here at the corner is not perfect, but it's close. The actual device shown here, the actual data seems to have a little bit more of a bird's head, then. Maybe the stress effects weren't perfectly modeled, but it's pretty darn close. People spend a lot of time calibrating their stress models and their stress coefficients to try to get close to what's measured experimentally.

OK. Let's go on to slide 20. What I just talked about is specifically related to seamless processing devices, trying to make the active area. By definition, when you're doing that, you're doing a patterned oxidation, and you have stress effects. Here's a neat application of stress effects that has nothing to do with CMOS. For those of you who are interested in nanotech or nanostructures, on page 20, this is an application of stress oxidation effects to form nanowires.

I took this from a PhD thesis of Harvey Lu at Stanford. He graduated in 1995. And what he was doing, and a lot of people since then have done the same thing, he was first etching a very small, tiny pillar in silicon. So a tall pillar. Say, a couple microns tall. But where the starting diameter of the pillar would maybe only be a couple hundred angstroms. Say 200 angstroms.

So he starts with this little needle of silicon. And you can do a 200 Angstrom pillar or a 300 Angstrom diameter pillar with a very good electron beam lithography machine. At least in those days, that's about as small as people could get-- a couple hundred angstroms. You patterned the dot and then you etched the silicon into a nice pillar. And then he took that pillar, those pillars, and he put it in the furnace and oxidized it.

And the idea was he wanted to get the pillar diameter down. He could not lithographically make the pillar any smaller, but he wanted to look at quantum effects, current conduction, in this very thin wire. He wanted to get it thinner somehow. And an obvious way to make it thinner is oxidize it, right? As you oxidize, you will consume from the top and from all the way around the sides of the little cylinder, you will consume the silicon by the SiO_2 production.

And then you can strip the oxide at the end if you want and make electrical measurement which turns out not to be so easy to contact a pillar that's only 20 angstroms. A little minor technical problem, but he was able to do it. And in fact, this is a cross-section TEM, a transmission electron micrograph, of one of Harvey's pillars. And what you can see here is two nanometers in diameter. So it's 20 angstroms. So it's a really tiny wire of silicon.

And this amorphous looking material all the way around it on the sides is the silicon dioxide that he grew. This particular one was oxidized in dry ox at 875 for about 10 hours. And you can actually see the lattice fringes in the single crystal silicon that remained. What was really neat is what Harvey found is, depending on the temperature that he did this, he could oxidize that pillar and it would stop.

The oxidation rate would go way down when he got down to a certain diameter pillar. And so that's quite convenient because then you could use this to form reasonably-- let's say lithography is not perfect, right? Some of the dots came out when he patterned at 200 angstroms, 250, maybe, 210. Something like that. Litho is not perfect. But he could get a reasonably wide range in the starting pillar diameter.

The final pillar diameter, if he oxidized it long enough, was pretty darn uniform. Which was kind of nice, because you have a process now that's physically slowing down as it gets smaller. And so they all end up-- even though you have a large distribution to begin with, they all end up very close to a single diameter. If you're trying to make some device or some interesting quantum thing, that was a very nice side effect.

It's a side effect of, he believed, the stress effect. That there was stress built up in the silicon and in the oxide in such a way that the oxidation rate, as the pillar got thinner and thinner, the oxidation rate would go down, and it ended up being somewhat self limiting at a certain diameter pillar. It's not exactly the same kind of effect that's going on perfectly as LOCOS, but it's related.

It gives you an idea that you can use these stress effects in processes to create structures that you wouldn't be able to otherwise create. If he didn't have this effect, he'd have a very good chance of really ending up with very few pillars that were of similar diameter. OK. So just as a point to make. And if you're working on quantum devices, you'll see a lot of work where people do shaped oxidation. They don't just make pillars.

They make small structures with electron beam. They pattern a silicon on insulator layer, and they take advantage of corner effects to create funny looking shaped structures that are used for quantum devices. So shaped oxidation stress effects are used all the time today in research other than just people doing CMOS devices. So what I want to go on to beyond the stress is to talk about now a more atomistic or microscopic or point defect based model.

So far we've been using macroscopic models. We write chemical equations. We talk about diffusion coefficients. We're not really looking at an atomic scale of what's going on at the interface or in the oxide. There's also a very atomistic picture of the process that turns out to be very important. It's important because it explains certain physical phenomena that are critical to fabrication such as oxidation enhanced diffusion, oxidation retarded diffusion, and these things that tell us that the oxidation process itself dramatically changes the point defect distribution in the near surface region.

Most of the ideas are related to the fact that there is volume expansion going on and the need for free volume at that interface. And so point defects play a role in that volume expansion. So let's go on to slide number 22. This is a kind of ugly looking equation. I don't really like to use it, so you don't have to look at every single term and understand it perfectly.

But it sort of shows qualitatively what we're talking about. This is an equation we talked about in chapter 3 when we were discussing SiO₂ precipitation and using gettering. But the same equation essentially applies here when we're oxidizing a silicon surface. And what it tells us is that we have a certain number of silicon lattice sites, the silicon sub SI, combining with a certain number of oxygen species and combining with a certain number of vacancies.

Notice that vacancies are being consumed in this process, forming SiO₂. Vacancies at the surface are being consumed, presumably, because we need space for that oxide to move, to push around a little bit. So consuming a vacancy will help you find extra space, essentially, at the interface. So I consume a vacancy on the left. I form SiO₂, and I actually form-- on the right side I also create these interstitials, a certain number of interstitials. Plus stress, of course.

We talked about stress. So this oxidation reaction happening right here at the borderline between the pink and the white region here at star, they can consume vacancies or they can generate interstitials. So I've got a vacancy going in, interstitial going out either way in order to provide the volume needed for the reaction. So we need room to put these oxygen atoms at that interface. And that room comes from vacancies and also picking out silicon atoms from the lattice, creating interstitials.

So let's go on to slide 23. This generation of interstitials is extremely important because it is used to explain the non-local effects of oxidation, such as oxidation enhanced diffusion, which we'll show, or oxidation retarded diffusion. That's because the interstitial can diffuse very far from the interface and change the diffusivity of dopants. So this shows how these models can actually be used. At this interface, about 1 in 1,000 of the oxidized silicon atoms is injected into the bulk.

So for every thousand atoms that I oxidize, one atom goes into the bulk that wouldn't be there ordinarily if I were not oxidizing. So it doesn't sound like a lot. Only one atom out of 1,000 goes into the bulk of the ones out of the surface. OK, that doesn't sound like a lot. But their impact is huge in oxidation enhanced diffusion, oxidation retarded diffusion, and the growth of stacking faults. So it's a very important effect.

And on slide 24, we have a picture from your text, a very schematic picture of what's going on when we do oxidation. So this is a LOCOS structure you recognize by now. Here's my nitride on the left. There's no oxidation taking place on the left. I am actively oxidizing on the right. You can see this thick layer of oxide that I've grown. At the surface here there are a number of processes taking place. This G is meant to represent the generation of interstitials or the injection.

It's a certain rate of generation of interstitials that get injected and go into the bulk. Here they are. Here's an interstitial. There's also a flux of interstitials from the bulk to the surface, and they can recombine. So interstitials can be generated during oxidation. They can also find their way to the surface and they can recombine there. The interstitials can recombine in the bulk with vacancy. They can find a vacancy and go away.

They can diffuse over to this surface, which is a non-oxidizing or inert place in recombine, or they can go into the bulk. And here, here's a buried dopant marker layer, which started out looking like this dark green. It was very narrow. And lo and behold, underneath where you're doing the oxidation, the boron diffuses a lot, so it becomes very wide. So they come in here and they disturb the boron diffusion. They enhance it.

And so the junction here, the width of this boron profile is now quite wide underneath the oxide where it was growing. Under the nitride where there was no oxidation taking place I have inert diffusion. There's very little broad motion. And these stacking faults, which are in the substrate. Let's say we have a way of introducing stacking faults. They are found to grow. They actually get longer under the region where the oxidation is. So these are all indirect means that people use to understand what's going on in the point defects in this model.

So let's go on to page 25. I mentioned about the stacking fault. I don't want to go through detailed crystallography, but stacking faults occur when you have an extra layer of atoms. So if I look here on the lower left, it has a certain stacking sequence. I insert this extra layer of atoms. It has a certain layer and a certain length in the crystal right at this point.

So it's not a complete layer that goes all the way throughout the crystal, but it doesn't-- so it doesn't belong there. But I can grow this fault. I can make it longer within the crystal by adding extra silicon interstitial atoms, people believe, on either edge of it. So stacking faults, people use their growth rate as a measure of the injection rate or the injection of silicon interstitials if they grow or shrink.

On slide 26 here's some actual photographs or optical micrographs. The surfaces were etched. Remember, we talked about defect etching. You can put it in a solution of acid that will preferentially etch where the crystal is imperfect. And in fact, these are stacking faults. There are pictures of them. They have a certain size here. These are all-- scale bar up here is 10 microns is shown up here. After 10, 20, 30, 40, 50 all the way down to 60 minutes, these are all at the same magnification.

These stacking faults, these crystal defects have actually grown quite a bit at 1,200 degrees. This is during oxidation, during high temperature oxidation at the surface. These were down under in the bulk, so people believe that that oxidation was injecting interstitials and causing these defects to grow. In fact, we go on to slide 27, this is some data from the literature on the stacking fault length here on the y-axis as a function of oxidation time at different temperatures.

Here 1,200 is up at the top. This is 1,050. And so you can see the stacking faults increases, this length increases with time of oxidation, and the growth rate is faster for higher temperatures. So if we go here, 1,050, 1,100, 1,150, the growth rate is faster. Of course, the oxidation rate is going up as well. So this gave people maybe a hint that the faster you oxidize, the more interstitials that you're injecting.

The oxidation induced stacking fault growth rate is smaller for lower partial pressures of oxygen. So again, slower oxidation rates. This is an idea that the rate of oxidation determines the rate of injection of interstitials into the bulk. So let's go back. Let's go on to slide 28. This is the diagram I just went through. These are the different processes we talked about that contribute to this oxidation enhanced diffusion.

So it's this generation at the interface where I'm oxidizing and the recombination, the balance of G and R, generation and recombination, determine what your net flux of interstitials is into the bulk away from the surface. It's this generation rate G that's directly proportional to the oxidation rate. So if I oxidize faster at a given temperature by upping the partial pressure, I will generate a larger rate of injection of interstitials and I'll get more boron diffusion in the Berry layer.

And by the way, this Berry layer can be 2, 3, 4 microns away. It doesn't have to be close. The interstitials diffuse very quickly and very rapidly, so this is the problem. Oxidation can have effects at a distance. A micron or so away can greatly impact what's going on in the chip. So they diffuse away from the interface, and they enhance the boron diffusion rate.

Here's an actual simulation of oxidation enhanced diffusion. What happened was the surface started out by being implanted uniformly with a certain dose of boron. So I have a very thin layer of boron up near the surface, and then we grow a padocks and we grow a nitride layer on the left. So on the left there is going to be no oxidation taking place. We put it in the furnace at 850 for an hour, and you see an oxide has grown on the right, not on the left, which makes sense. But look at the boron. The junction depth for the boron is much deeper here.

You can get the junction depth, say, from, let's say, the region between the green contour and the blue. The junction depth is a lot deeper underneath the place where the oxide grew. So this tells us away that we are getting oxidation enhanced diffusion of that boron. This is something that's been simulated. Doping effects on oxidation.

So besides the oxidation injecting interstitials, people also noticed that at low temperatures, depending on the wafer doping, they would get a different oxidation rate. So this is kind of interesting. If you look at A, A is a very light-- so I'm plotting oxide thickness versus time. A is lightly doped silicon, kind of the 15th. F is 10 to the 20 or 3 times the 20. Much more heavily doped with n type, with phosphorus.

So especially at low temperature-- at high temperature, not so much of a dispersion, but at low temperature, a large dispersion. So this needs to be accounted for in our models because depending on you have lots of different doping concentrations on the surface of a wafer when you're making a chip, you'll get different oxide thicknesses. You need to take this into account. People, in fact, not only found the thickness to be different. They extracted the B and the B over A parameter.

So as a function of boron doping-- so this is the oxidation rate, or these rate constants B/A and B as a function of boron concentration. Oh, I'm sorry. It's actually phosphorus doping in this reference. The B/A parameter is the one impacted. So above a certain concentration of dopants, here above about 10 of the 19, B/A takes off. So it's going up quite rapidly. The diffusion through the oxide doesn't seem to be affected.

You can kind of imagine that I've got doping at the interface. You expect there's some effect of that doping on the reaction rate. And in fact, if we to slide 32, people had an idea of what it might be. Maybe it's a chemical effect of the dopant, perhaps, but people liked this model better. People liked the idea that the vacancies available in the near surface region are important to the oxidation process.

So anything that caused vacancy concentration to go up might cause the oxidation rate to also go up. But what do we know about highly doped regions from your homework that you did? When you increase the doping in a certain region, the total number of vacancies go up, right? Because you move the Fermi level and you create charge vacancies and you count up all the vacancies and they go up. So remember these equations you used on your homework that you're handing in today.

If I'm intrinsic, I'm here. My Fermi level is at mid-gap. I'm dominated by neutral vacancies. They have a certain concentration. 10 to the 13 or whatever. As I move the Fermi level up so I make it more and more heavily doped, I create a lot of v^- and v^{2-} . So I could have a hundred times the number of vacancies at the surface. Oxidation needs vacancies to take place. That's the hypothesis.

So people explain the enhancement of the high doped oxidation case due to extra vacancies that are around at the surface. That's one potential explanation. And in fact, I just reminded, you did this in your homework. You did some calculations on slide 33. The vacancy concentration that's been calculated for you here as a function of doping concentration, look where it takes off.

Of course, somewhere right around 10 to the 19 it really starts to take off. That's exactly where Hogg found that the B/A coefficient took off is right where the vacancy concentration is going up as a function of doping. So let's just say on slide 34 that these vacancies are available to provide sites, these extra vacancies, for the oxidation reaction.

Then we can imagine writing B/A equal to something, some number R_1 where R_1 has to do with all the mechanisms other than vacancy driven processes. Obviously oxidation is not only just controlled by vacancies. It's controlled by a lot of things. But there must be some term added in this $k \times V_{\text{total}}$, CV_{total} . So the term on the right represents the vacancy driven process.

It's, I believe, to be directly proportional to the concentration of the total number of vacancies in all charge states. So doing that, you can then rewrite the B/A parameter to be B/A in intrinsic material, lightly doped, when you're not extrinsic, times this 1 plus some exponential type of factor-- so this is empirically determined here as a function on temperature-- times something that goes like CV over CVI . So the concentration of vacancies in this heavily doped material divided by that intrinsic material minus one.

So if I'm intrinsic, this goes to zero and this term goes away. As this increases, this ratio goes up above one. This term starts to kick in. Interestingly, though, when we think about this now, this model for the vacancy dependence is only going to depend on the electrically active concentration. Because remember, we said we move the Fermi level up.

So if we add phosphorus to the wafer, it's only the amount of phosphorus that contributes electrons that's going to move the Fermi level around. So people had an idea on how to test this. What if I compensate the wafer? OK. Well, actually let me go on. I'll talk about that next. Let me just show slide 35 just to show, compare boron and phosphorus.

We know from our calculations, actually, that n type regions, just because of the point defect statistics, n type regions tend to have higher charge vacancy concentrations at a given doping than p type, than boron. And in fact, what we see is that the B over A parameter is more impacted in n type silicon, very heavily doped, than it is in p type. That's the difference between the right and the left hand side on page 35.

That's some hint. What I started to talk about, which I think to me is a little more convincing, that this is really maybe is due to vacancies. If you look at the oxide thickness as a function of the electrically active phosphorus-- and this solid line is for uncompensated. What is uncompensated? I just add phosphorus at certain level. Here I add it to 10 of the 19, 5 to the 19, 10 of the 20, and 5, 10 of the 20. And I don't add any other dopants. No boron.

As the phosphorus goes up, the vacancy concentration goes up and so does the thickness. Now the dashed line is material that's been compensated. What does that mean? Well, what that means is when I add 10 of the 20 phosphorus, I add 10 of the 20 boron. There's still 10 of the 20 phosphorous there and 10 of the 20 boron, but electrically I haven't added any extra electrons or holes because they compensate each other in terms of dopants.

So the electrically active concentration has not gone up when I add as much boron as I add phosphorus. You add equal amount of p and n type doping. I still end up with fairly-- I end up with material that's compensated. I haven't moved the Fermi level. I haven't created a lot of excess charge defects. And in fact, he found that there wasn't that much increase in the oxidation rate when he added a lot of phosphorus but compensated with boron. Equal amount of boron.

So if it were a chemical effect, just the amount of phosphorus atoms-- let's say a phosphorus atom chemically causes oxide to grow faster. It wouldn't matter whether you added an equal amount of boron. But if you add an equal amount of boron, you don't change the Fermi level. So this is an argument that says to me that it really is the electrical concentration that's changing. It's the movement of the Fermi level. It doesn't prove it's a vacancy effect, but it gives a little bit more weight.

Here on slide 37, it's just a simulation of this type of effect. What we have on the left hand side is local oxidation where it's taking place here at 800 degrees for 30 minutes and you're seeing get a little thicker oxide or get a five times thicker oxide growing in this heavily doped region. So this region over here on the right is heavily phosphorus doped. So it was I implanted prior to oxidation.

So you need to take this into account. If you're making a device and you implant a region for a source drain and then you go to oxidize it, it's going to have a different oxidation rate compared to those parts of the chip that are being oxidized that don't have high doping in them. So in SUPREM you can easily take that into account. And the last thing I want to cover is here on slide 38, we talked about oxidation.

It has all these point defect effects. We also talked a couple last time or the time before about interface charges that come along with oxidation. And this model on slide 38, is while we're talking about atomic level things, this is an atomic level cartoon or picture of what is actually causing the fixed charge and the density of interface states. And what it supposedly is here, in the near-- here's a miamorphous SiO₂ up here and my crystal silicon.

In the near interfacial region, there are believed to be a few interstitial silicon atoms. Remember, some interstitials get injected into the bulk. Some go up, out. There's a few that get stuck near the interface. This little extra silicon interstitial with a positive charge on it. It's not bonded to oxygen atoms, so has an excess positive charge. That unbonded silicon interstitial, the extra dangling silicon interstitials that are there that are positively charged are believed to be responsible for QM.

Remember, fixed charge doesn't change with bias. The density of interface states that change is believed to be associated with dangling silicon bonds right at the interface, and we can passivate these because they're not in the oxide. They're right at the silicon dioxide surface. We can passivate these by adding hydrogen.

So when you build a CMOS or you build a MOSFET, the last step is 450 degrees in hydrogen ambient. The hydrogen diffuses in. It satisfies these dangling bonds. It bonds here and it covers up that charge. The hydrogen can't do anything about this silicon. It's unoxidized. It needs oxygen atoms. So that's why when you do a forming gas anneal here, the last step, you can reduce the density of interface states.

The fixed charge is a little different. To reduce it you don't anneal on hydrogen. This is so-called Deal triangle. What you do is you anneal the oxide after oxidizing it in an inert ambient at high temperature. So the idea would be that you try to get some of those silicon interstitials that were stuck there. They're unoxidized. You give them time to diffuse away from that interface so you can get rid of that fixed charge.

That's exactly what you will find. If you run an oxidation program in the furnace, what people typically do, they grow an oxide, 800 or 900 for an hour. The last step of the program is to let the wafer sit at 800 or 900 in the tube with no oxygen flowing. Just argon or nitrogen. And that last half hour anneal is the QF anneal. So if you grew in dry oxygen at 800, an oxide, and you didn't anneal it, you might have on 111 silicon about 6 times to the 11 for your QF per square centimeter.

If you anneal it, you can get it down at 800. If you anneal it in nitrogen or argon, you can get it down to about one times 10 to the 11 on 111 silicon. 100. All the numbers will be about three times lower. If you anneal it at a higher temperature, it'll also come down. So the whole idea is to anneal that in inert ambient at a relatively high temperature right after your oxidation.

OK. This is the last slide. Just to summarize chapter 6 and what we've talked about today. We know have to control the oxide thickness to atomic level for a thin gate dielectrics. There's a basic Deal-Grove mechanism which involves diffusion through the oxide, reaction at the interface. This doesn't work perfectly, but there's been a lot of situations that people can correct the Deal-Grove model.

This non-planar oxidation is impacted by the orientation of the surface, two dimensional diffusion through a shaped oxide, and the stress effects. Stress retards or slows down the oxidation by slowing down the interface reaction $k_{sub a}$ and slowing down the diffusion through the oxide. And these silicon interstitials are injected into the bulk during oxidation, and we're going to come back and revisit them in the next few lectures or over the next several weeks because they're going to dramatically impact diffusion of boron and other dopants.

So that's about what I have. I believe, for today, your homework is due. You can bring that up right now and put it in this orange folder. If you didn't pick up your homework last time, it's in the back with the handouts. I'm going to be out next week. Your TA, Maggie-- Maggie, you want to raise your hand in the back? Is going to give lectures as usual. So we'll start diffusion on Tuesday. And on Tuesday, then your new homework goes out.