JUDY HOYT: Doing right now is hopefully, you're all reading chapter 6, which covers the process of thermal oxidation. Last time, we talked about general introduction to thermal oxides in silicon CMOS. And we also talked about the metal oxide silicon or semiconductor capacitor and how electrical simple CD measurements can be used to understand the quality of the oxide and the charges that exist in that structure.

This time, we're going to cover a basic model called the Deal-Grove. This is after the two people, that's Bruce Deal and Andy Grove, who came up with this linear parabolic oxidation model. It really remains the basic model used today. And then we'll talk about examples on how to use it. And we'll finish up with other models of planar oxidation that are very useful in the very thin oxide regime.

So I'm just curious, how many people here have known or not known but have heard about the Deal-Grove oxidation model? So about maybe a little less than half or so. OK, so that works reasonably well. We go through it because it's the basic model. And then everything after that, a lot of things after that tend to be corrections to the Deal-Grove model.

So let's go on to slide number 2. And that shows the basic schematic setup for how Deal and Grove derived their model. What it shows you is turned sideways is you have a gas stream here on the left that's on top of the wafer. There is a thermal oxide being grown here in this white region, the silicon dioxide. And the pink on the right shows this meant to represent the silicon wafer.

And just to give you an order of magnitude scale. The silicon wafer is thick. Of course, it's hundreds of microns. This oxide that we're talking about growing is on the order of 0.01, 0.01 micron being about 100 angstroms or 10 nanometers up to maybe a micron. That would be the range of reasonable thicknesses of which you would grow.

Once you start growing more than a micron, as you'll see today, if you want to grow one or two or three microns by a thermal process, it gets pretty darn slow. The kinetics are such that it becomes really not very practical. So you don't see too many people growing oxides that much thicker than a couple of microns. If they need oxides that thick, they do a different process called chemical vapor deposition.

So the basic idea of this model, though, is that we have a gas stream here. And there's a concentration C sub G of the oxidant. Let's say it's oxygen. Could be oxygen, could be water vapor. But there's some concentration in the main gas flow , in the main gas stream. As you get closer to the surface of the wafer, which has oxide on it, it has some concentration C sub S that's right at the surface of the wafer.

There may be some boundary layer or whatever that causes a gradient to cross here. But right at the surface of the wafer, it's at C sub S in the gap but in the gas phase. In the solid, right inside at the very surface in the solid oxide, this oxidant has some concentration C sub O right at the surface, or C 0 if you would like.

And then there is some gradient of the oxidant, say of the oxygen molecule. There's a gradient throughout the oxide. And at the interface, it has C sub I, the I standing for the interface. It has that concentration.

And the basic chemical reactions that are happening now at the interface, not at the oxide surface, are shown here. Silicon plus oxygen going to SiO2 or silicon plus moisture, H2O, going to SiO2. And there in the equations that we're going to derive, the simple model, there are three fluxes involved.

There's a flux F1 of oxidant that has to be transported in the gas phase from the major portion of the gas phase in the center of the tube to the surface of the wafer. That's one flux. There's a flux F2 in the oxide. This is the diffusion of the oxidant through the oxide. And there's finally a flux F3, which is what's happening right at the interface as the oxygen is incorporated into the growing oxide.

So let's go on to slide 3. Again, on the upper diagram is exactly what I've just shown you. For those of you who have forgotten, on the upper right, I'm reminding you what the definition of flux is. What is a flux? It's the number of particles. Doesn't have to be, in this case, I'm talking about oxygen molecules or water vapor molecules. But it's the number of particles crossing a unit area-- so a unit area in this direction here-- per unit time.

So there's a flux, in this case, going from the left to the right. In fact, there's three different of them. So what Deal and Grove did was write down three very simple equations, first order flux equations, which show these three series parts of the process. Again, they're all happening in series. So there's a flux, F1. They wrote down using Henry's law.

So Henry's law just says, in the gas phase transport, if you have a concentration difference or a gradient here from this point here in the center of where the gas flow is to the surface, C sub S, that you can write the flux as just some constant, some Henry's constant, a constant number times the difference in the concentration between the surface and the where the concentration is CG. So that's a simple equation from Henry's law.

A second equation they wrote down was diffusion through the oxide. And this comes from Fick's law, Fick's first law. Hopefully, some of you have seen this, maybe not in this particular form. When we talk about diffusion, it'll become very obvious. But this is Fick's first law for diffusion through the oxide. The flux F2 is just a constant number, which we call the diffusivity, D, times the gradient of the diffusion, DN by DX, where n is generally just the concentration of whatever is diffusing.

And you can write that. If this is a straight line, writing the gradient of a straight line is very simple. It's just C at the surface, C 0 minus C interface divided by the thickness of the oxide. So it's a simple equation to write diffusion through the oxide.

And finally, F3 is the flux at the interface. And this is just a surface reaction rate. So we write this final flux as K sub S, a constant, which has to do with how rapidly the reaction is taking place at the interface, times the concentration of oxygen or water vapor right at the interface. So three equations for each of these fluxes.

So we go on to slide 3. And what they're assuming is steady-state condition. And what we mean by steady state is things are not changing with time, these profiles. So if I just go back one second to slide 3, is that this profile is, we're not in the transient phase. This profile has now been established as a linear profile in a steady state. It's not changing with time. You've established a constant profile.

So in those conditions, then, we can write down that the three fluxes in series are all equal to each other. We don't have any buildup or any sync. So F1, the flux throwing through the gas, has to be equal to F2, the flux diffusing through has to be equal to the reactive flux.

So in that case, then, it's relatively simple. If you sit down and you equate those three equations, you can derive this formula here, equation number 4, for the concentration of the oxygen molecules of the water vapor C sub I at the interface. You can write it as some constant, C star, which we'll talk about, where C star is the solubility of the oxidant in the oxide divided by a three-term expression, 1 plus the ratio of K sub S to H.

So that tells us something about the ratio of the surface reaction coefficient to Henry's coefficient H, plus k sub S X node over D. This is involving a diffusivity term.

And so what we are generally going to do in this, we can actually simplify this equation by recognizing that we actually can neglect F1, the gas phase transport, which is usually a very good approximation. It's very fast. So if I go back to this slide number 3 briefly, this flux, this is happening much, much faster than these two processes, than the diffusion through the solid or then the reaction, as it happens to turn out in this case.

That won't always be the case for other processes. When we talk about epitaxial growth or chemical vapor deposition, it won't always be this case that you can ignore the gas phase transport. But as it turns out, the surface reaction is slow enough, and the diffusion through the oxide is slow enough that this is always very fast. And so the beautiful thing about oxidation is you can ignore gas phase transport.

And that's why you can design oxidation tubes to be very simple. You just stand all the wafers up and flow the oxygen through. And you don't have to care about exactly how the gas flows. When we talk about other processes like epitaxial growth, no such luxury by any means. In that case, this flux is extremely important. And in fact, it controls the process many times.

So but we're very lucky in oxide growth that these two processes F2 and F3 are slow. So we can ignore F1. And that just ends up simplifying the equations to a certain extent. And so if I rearrange this a little bit, we get equation number 5 here for the-- we're really interested in C sub O here, the concentration at the surface of the oxidant.

So let me combine now this equation number 4, which tells me what the concentration of water or oxygen is at the interface with-- I'm going to combine this equation with equation 3, just going back again, which just tells me that the surface reaction rate flux.

And we get something like this that says that DX by DT-- so that's the rate of change of the thickness of the oxide-- DX by DT, we know we can write it as a flux divided by a concentration. So N1 here, does what anybody have an idea of what N1 is, just from dimensional arguments?

I want to get DX by DT. Flux tells you the number of atoms per square centimeter per unit time. So N1 has to have the units of the concentration per cubic centimeter, basically, so that it works out, even just to work out dimensionally. Of course, I don't have any chalk. But that's what happens.

But if you work it out dimensionally, flux has a number of atoms per unit time per square centimeter. This is the number of atoms in the oxide. Per cubic centimeter, you end up with something like looks like centimeters per unit time, or length per unit time. So this is how quickly the oxide is growing, the growth rate or the velocity, which is what we're interested in, the growth rate.

Can be written like this. It's the product of K sub S, the surface reaction rate, times the C star. C star is the solubility of the oxidant in the oxide. So that's as much as you can put in at that temperature, OK?

You might have more in the gas phase than that. But that's how much will go in at that temperature. C star divided by this denominator, 1 plus K sub over H plus K sub S X naught over D. So that's telling us the rate of growth of the oxide. Now, we have it, in terms of certain fundamental physical parameters.

So let's go on to slide 5. That was a simple, if you want, differential equation, DX by DT, just says how rapidly it's--- what is the velocity of the growth. And we integrate that. If you look in the text, it has a little bit of more detail on how you do the integration. But you integrate it between 0 and some time and between some initial thickness X sub I. And your final thickness or the thickness at any point in time, X0.

So here, X0 is the thickness of the oxide right at the time, T. And XI is whatever thickness you started. With could be 0, whatever you thickness you started with at the start of the oxidation process. So by doing that integration, you end up with that time, if this can solve for time, is a two-term expression, something that involves X squared, thickness squared, and something that involves a linear term, thickness to the first power.

So this is the origin. The origin of this parabolic/linear model. Here's a parabolic term. It goes like the thickness squared. Here's a linear term. And the multipliers are B. Here, this is the multiplier of the X0 the parabolic term. Looks like this. It's it looks like the diffusion coefficient through the oxide times the solubility in the oxide, C star, divided by N1.

So the parabolic process has to do with diffusion through the oxide. The linear term, the term that multiplies X0, can be simplified here to something that looks like C star K sub over N1. And again, I'm ignoring this 1 over H because this is going very rapidly. So H is very large compared to K. So I can ignore this term.

So this is the linear rate constant, which tells you something about the surface reaction rate. So linear rate constant has to do with the reaction. Parabolic has to do with the diffusion. Notice that they both depend, however, on C star, which is the solubility of either the oxygen or the water in the oxide.

So let's go on to slide number 6. These rate constants, B and B over A, have some physical meaning. We just talked about that, diffusion of the oxidant and interface reaction. And we find experimentally that we can write them as Arrhenius relationships. They're exponentially dependent on E to some activation energy, E to the minus E1 over KT or E to the minus E2 over KT times a prefactor.

The prefactor here is C1. For B and B over A is C2. And these rate constants have been studied for many years. People have measured them or tried to extract them experimentally. And this is a chart, a table, I took directly out of your text, which shows the ambient. So that's the type of gases environment that you are in, the wafer is in, and the rate constants, B and B over A, and how to calculate them.

So in dry oxygen, so you're flowing O2 in the tube. But there's no moisture around. To calculate the B rate constant, you need to know C1. And you need to know E1, the activation energies, wet O2 and H2O and moisture environment. So, by the way, just be careful when you're going to use these numbers.

These are for 111 silicon oxidation. So that's the oxidation of a wafer with orientation surface of 111 at one atmospheric pressure. If you want to know the values for 100 silicon, then you take the C2 values, these B over A, the rate constants, and you divide by 1.68. So it's slower for a 100 interface.

And if you go back and look at your crystallography and look at the number of atoms per square centimeter, it actually makes some sense because on a 111 interface, we have the largest number of atoms per square centimeter available for the oxidation. So it's going to be the rate constants are going to be faster.

But notice, the activation energy doesn't change as you change orientations. It's all for the rate constant B over A, the surface reaction rate, it all has an activation energy of close to 2EV. So the activation energy for B over A is roughly constant. It just depends on the surface reaction rate.

For B, interestingly, the activation energy for B varies with the process. So depending on whether you're in a dry O2 here, you get 1.23 EV. Here, this is 0.7 EV. And for moisture and for water, it's 0.78 EV. So it actually varies with the process. And that makes some sense because what did we say B depends on? It depends on the diffusion of this species through the oxide.

Well, a water molecule, if that's what's diffusing the oxide, is going to diffuse at a different rate or maybe have a different solubility than just oxygen molecule. So we do see a difference in the activation energy for B. And on slide 7, obviously, you can do this yourself. You just sit down with those equations on a semi-log plot.

So this is a log scale on the y-axis. And this is 1 over T or a 1,000 over T. This is a standard Arrhenius plot you've had to make. So on a semi-log plot, any kind of exponential relationship will appear as a straight line because it's log paper.

And these are the different constants. Here we are for water is the B constant, the B over A, rather. And this is the B constant. And these are for 111. So you just literally plot what values are on the table. One thing you see right off the bat, just from this plot, these two lines on top that are faster, that are larger, are for oxidation in water vapor as opposed to dry.

So right off the bat, you know the rate constants are higher for water than dry. So if you want to grow faster in oxide, you would grow it in a moisture environment, not in dry O2. And that's typically the case. People grow thick oxides using a wet environment, thin oxides using a dry environment.

OK, let's go just go on to slide number 8 and just reformulate this. You can do some mathematical manipulations on the law, the linear parabolic law, as I showed it earlier. I've just repeated that equation, where on the right-hand side, we've solved for T. On the left-hand side is the oxide thickness. There are alternate ways to express this law. And in fact, here's a reformulation of it shown in the second equation on this slide.

What you see is what we've done is now, we've expressing it as X0 squared over B plus X0 divided by B over A is equal to T plus tau, where what we've done is tau now is defined as the time that it would have taken, the effect of time, to grow the initial oxide thickness X sub I. So this is just obtained by a mathematical manipulation of the first equation.

So if you set tau equal to XI, the initial squared, plus A XI over B, just mathematically, you do that, then you get a mathematical equivalence of the second equation to the first equation. Sometimes, you also want to explicitly have an equation that shows you what X0 is, the oxide thickness. That's not what's shown here. On the upper two equations, X0 is implicit because it appears here in two different terms.

If you want to solve explicitly for the oxide thickness, you can see this middle equation, X0, it looks like this. It's got a square root dependence on time. And where, again, tau has the same definition, XI squared, it's just that time to grow the initial oxide, the effect of time in the Deal-Grove model.

OK, let's go on to slide number 9. And again, this reformulation, I'm just repeating the equations from the previous slide. So you don't have to remember them. Again, XI and tau account for any oxide that was present at the start of the process. When you first push the wafer in, you might have had other processes that have taken place. And they might already have some oxide present there.

So we can take this-- we can take this equation and differentiate it. And let's say we're interested in the oxidation rate, not the absolute thickness versus time. But how is the rate changing? How fast is it growing? Well, this is the oxidation rate. By differentiating, you can see, it's given by B divided by this quantity, 2X0 plus A. So this gives you an idea of the rate, how rapidly is it growing?

So these equations are very general. And what they describe as the growth kinetics for oxidation. When you're in these following conditions, we have to remind ourselves that it only applies for certain basic case. So this means when it's planar oxidation, so I'm just oxidizing a flat surface, we'll talk next time, when you have a surface that has structure on it, the equations need to be modified. It's planar and it's unpatterned. So this is just an unpatterned surface.

This assumes that the silicon is lightly doped, that you can use these numbers. Heavily doped oxidation, we'll talk about next time. A simple ambient such as oxygen, dry O2 or water, is assumed. If you're going to add other gases, nitrous oxide and other things that people do to grow high K, these particular equations don't apply exactly.

And these equations work when the final oxide thickness is larger than about 20 nanometers, about 200 angstroms. They work quite well. But so that's a certain number of limitations, still very useful because what we're going to do from here is now make corrections. Right off the bat though, when you see this 20 nanometer number, what do you think of?

I mean, gate oxides today, just from your reading of the ITRS and high performance devices, what are they? Are they less than or greater than 20 nanometers? Less than, significantly less than for most high performance devices.

So this tells us off the bat, if we want to use Deal-Grove to model a standard gate oxide thickness that someone might be growing 100 angstroms or less, it's not going to be perfect. We're going to have to make corrections to this. And we'll talk today about those thin oxidation models.

But let's go on to slide 10. And again, just to make a different way of looking at these same equations, I've repeated them again at the top so we don't have to memorize them. But what we can do now is to take two limiting cases. One limiting case I'm going to say is when the oxide is relatively thin. So X0 is relatively small. But it turns out this term here, the linear term X0 divided by B over A is going to dominate over this X0 squared. This is going to become a smaller term.

So we can ignore this first term. And we end up that X0 is proportional to B over A times time. So it's essentially, in thin oxide, we have the equation is linear in time. So it's just we have a constant oxidation rate. And the thickness just increases linearly, according to this Deal-Grove model.

For thicker oxides-- and we'll define later what we mean by thicker-- but at some point, this term, this X0 squared over B, starts to dominate over the linear term. We can ignore this linear term. For long times, where very thick oxides, it goes like X0 squared is a constant times T. So it's basically a parabolic term. So the oxide thickness starts increasing like the square root of time. So initially, it's going up linearly. And then when it becomes thick enough, it starts increasing more like a square root kind of function.

And at that point, it's thick enough that diffusion through the oxide is really starting to limit you. Mathematically, if you prefer to think mathematically, you're going to make yourself a plot. This is actually a log scale, a log plot, log on the Y, log on the X. And it's kind of normalized. What is a plot of is X0 divided by A over 2. So we've sort of normalized the Y-AXIS and T plus tau divided

by this constant, A squared over 4B to make it in unitless. And what you see, this solid line here is equal to a plot of X0 squared plus AX0 equals B times tau time T plus tau. So that's the exact Deal-Grove model. At short time, you can see it approaches or low thicknesses it approaches this linear, rate constant, this linear equation. Remember, on a semi-log plot, this has a certain slope.

And then for longer times, much higher thicknesses, it approaches a straight line, this dashed line that has half the slope. Again, linear here versus square root has half the slope. Different way of looking at it-- I actually like you to think a little more physically, not just look at the equations and say, which term is big and which is small? In fact, it's much easier to remember the limiting cases in a physical sense. And that's what's being shown on slide 11.

So if you look at the case of a thin oxide, it's relatively thin, what's happening is physically, is diffusion. This is so thin, the diffusion process happens very rapidly compared to the rate at which those atoms can react at the surface. So when it's thin, the surface reaction is what limits. It's the slowest step. Any of these serial processes always end up being limited by the rate limiting step is the slowest step.

So clearly, when you're thin, getting through the oxide takes no time at all. What's limiting you is the rate at which you can react. And we know that is a linear process. It just depends on K sub S. Now let's say when you're thick, if you have a thick oxide, well, then the diffusion process is what's rate limiting you. The reaction at the surface can be relatively fast compared to the diffusion rate.

So diffusion is rate limited. And that process ends up depending-- has a square root type dependence on time. So in both cases though, and just for these little diagrams I drew, C star, remember, is the solubility of the oxidant and the oxide. So in the thin case, you diffuse very rapidly through. You have almost no gradient in the oxide. You're getting through very rapidly just by virtue of it so thin.

In a thick case, C star is a solubility. That's what you have at the surface. At the interface, it's almost going to 0 because again, you're rate limited by the rate of diffusion. You have a large gradient. So this flux through the oxide is being driven by this gradient. So those are the two limiting cases. One depends on diffusion. The other is limited by surface reaction.

Slide 12 is just a little example of if you're an experimentalist, how do you sit down? And how do people in the old days used to do this? Well, it's pretty simple, experimentally. You just oxidize a bunch of wafers at the same temperature for different times.

And you plot the oxide thickness measured versus time. You'll see this linear dependence at first. And then it'll start to bend over where it's square root. Easy way to do it is to plot X instead off X0 instead as a function of time, but T time divided by X0.

And you'll find out that the slope, if you just manipulate those simple linear equations, the slope is the parameter B. And the y-intercept is the parameter minus A. So for the Deal-Grove model, it's very simple experimentally if you do this kind of plot to extract those two parameters.

These are some actual calculations of the oxidation rates in dry on slide 13 and dry oxygen using the Deal-Grove model. This is for 100 silicon now. So 100 silicon is much more likely what you're going to be using in a silicon manufacturing process or probably in your research. So it's a plot on the left axis linear scale oxide thickness as a function of time.

And what you do see is a linear regime here, say, when you're at thin oxide, say below 0.1 microns or maybe half of that, 500 angstroms, the lines look quite linear. And then they start bending over. So they're followed by this sort of parabolic behavior where they grow less rapidly. And you can also see, look at the thicknesses, as we mentioned before. Dry oxides, even at high temperatures of 1,100, which is quite high for a furnace.

You're really limited. For practical times, you don't want to be leaving the wafers in there any more than a few hours, four hours, six hours, whatever. You're limited to things on the order of 1,000 to 2,000 angstroms, from a practical point of view, where people grow dry oxides. Of course, you could leave it in the furnace for days. But that's not very economical in terms of use of the furnace, use of electricity, and all the resources.

So if you want to grow an oxide in this thickness range, you use a dry oxide. If you want to go well above that, you should be going to wet oxidation because the oxidation rate constants are so much higher. In fact, on the next slide, slide 14, you can see the oxide thickness versus time for wet ambient, for moisture ambient using Deal-Grove.

And now, look at the time of say two or three hours, instead of growing 1,000 angstroms or so, at 1,100 degrees, you're growing micron here. After about two hours at 1,100 degrees in wet ambient, about a micron, so much faster than dry oxidation, orders of magnitude faster.

And as it turns out, the physical reason why is that the solubility, C star, of water vapor is much higher than the solubility of oxygen, of the oxygen molecule, for various reasons. Just because of the configuration of the molecule and the size of the interstitial spaces, water vapor is just easier to get in there. So it has a much higher solubility. So that whole concentration can be elevated, and you can oxidize it at a much faster rate.

So if you need a thick oxide, you use moisture. Yeah, question.

STUDENT: When B of K is given, so all of the other side, where is the oxide that doesn't react? What would it be then?

JUDY HOYT: OK, so the question is going back to, let's say, slide 11 maybe of these limiting cases. What happens in a real case, if you have an oxidant diffusing through the oxide? And does it all react? Is there any unreactant species at the interface? Basically, it does have to react.

And then if the reaction rate is what's limiting you, then you're sort of in this regime. But there will be a certain amount of oxidant that doesn't react. And that can lead to perhaps excess charges and things like that, some of the more realistic aspects of the oxidation process that you're left over. So it can affect the quality of the oxide, to a certain extent.

And what you generally find is that wet oxides-- you don't see a whole lot of people growing gate oxides by wet. Well, because they want thinner. But even in the old days, when a gate oxide was thicker, 500 angstroms or whatever, you didn't see people growing them in a wet ambient. And the oxidation rate is a lot faster. The quality of the oxide, the number of interface states and things like that, tends to be higher for a wet oxide. So typically, what you do if you're trying to grow a gate oxide, is you grow it dry.

Or if you want to make it thick, and you have practical limitations, you grow a dry step, a wet step, and then a dry step at the end. So the interface was always grown in a dry step to try to reduce the interface state density. So if you need a gate oxide that's 1000 angstrom, you do what they call a dry-wet-dry.

So it does matter if the oxidation rate-- being too high is not the greatest thing in the world in terms of the device properties. I think was there a hand in the back? Or did I miss a question? OK. OK. Good.

So those are the wet oxide kinetics. Let's go on to slide number 15 and just show some examples. Now right off the bat, I'm unfortunately showing an example to which the Deal-Grove model doesn't perfectly apply. But we'll apply it anyway. And next time, we'll talk about why it's not perfect. But there is a process called-- just to make this a little more interesting-- called recessed LOCOS local oxidation, which can be used to provide this lateral isolation between a device A and device B.

We need to grow an oxide here. And it results in a more planar surface than standard LOCOS. Remember, standard LOCOS look like this. When we're done, and you strip off this nitride up in the upper right, you have a very non planar surface introduced. You've got this region up here, where the oxide pushed up. And so the surface is not very planar.

Recessed LOCOS, so that's the basic idea is you etch a trench first. So here, I've etched in the silicon. Around the silicon nitride, I've etched a half micron deep trench. And now, we assume I'm going to subject that wafer to thermal oxidation. Well, on the left and the right side of the trench, that should say silicon nitride.

Sorry, this left and right side should both be silicon nitride. There's no oxidation taking place. So maybe correct this handout on the right, where it's pointing to SiO2, that should say silicon nitride, Si3N4. So we're only oxidizing, we're assuming, in this trench region. Now, this schematic is highly schematic. It's not including any bird's beak effects, which we'll talk about where they originate from, or any two-dimensional oxidation.

We're going to do a very simple calculation of how long would it take to fill that to make a planar surface, assuming I was half micron down. Yeah.

STUDENT: After doing that, will it [INAUDIBLE]

JUDY HOYT: There is in between. Yeah, it's just the arrow got displaced. This arrow is actually pointing to the top region, which is the silicon nitride. There is, as was just pointed out by someone in the class here, this hatched region is meant to represent the SiO2, what they call the Pad Ox, or the pad oxide, and if it was stressed, that's induced by the nitride. So there is, in fact, a thin pad oxide underneath the silicon nitride.

But what's preventing the oxidation from taking place is this silicon nitride. SiO_2 itself won't prevent oxidation. The oxidant will diffuse right through it. The nice thing about silicon nitride is that it's hard to get oxygen and moisture to diffuse through it.

So it can act as a mask or as a blocking layer. So you don't oxidize underneath it, the first word. So this oxide that's shown here underneath it happened before the nitride was put down. OK, good point.

So let's go on to slide 15. And here's just an example of a question. So we're assuming we've etched half micron down into the wafer below the original surface prior to oxidizing. And a simple question might be, how long do I have to put it in the furnace at 1,000 degrees, assuming water oxidation in H_2O to produce a planar surface?

Again, we're going to ignore any of the bird's beak effects or the stress effects that we know are there. We'll talk about those next time. So but just to make it a little more interesting, we want to fill this up to the surface, the original surface, and then make it planar. So we have to be a little careful in doing this because remember, as I'm growing this oxide right here, I'm consuming the silicon underneath it. So you're consuming silicon underneath it. But the oxide is also growing up at the same time.

So it's kind of starting here, very thin. And it's kind of going like this. At some point, it will be just flat or close to being almost flat with the original surface. So how do we figure this out? Well, you make yourself a little schematic diagram as shown here in color in the lower left. And in a rough sense, for every micron of silicon I consume-- so as you start oxidizing, I'm going to start bringing this trench point lower.

Every micron I go down, 2 and 1/2 microns, or 2.2., I'm sorry, microns of silicon is grown, of silicon dioxide, SiO_2. So you can write this mathematically. If I say Y is the thickness of the silicon that was consumed, then 2.2 times Y is X_0. That's the thickness of the SiO_2. So that's just due to the volume expansion when you grow oxide. So we can have that as one equation that describes what's going on.

Slide 17, so we remember that equation. But then we also know that the total thickness of oxide grown, in order to satisfy this requirement that we fill the trench, well, we want to a specific case where X_0 is equal to 1/2 micron, which is the original trench depth, plus Y, which is the thickness of silicon consumed.

So by looking at this geometry, so the distance from here to here is half micron. That was my original trench. Plus the silicon is going to get consumed down a little bit. And that, we're calling the variable Y. We don't know that number yet. So we have two equations here. This X_0 equals 0.5 micron plus Y. And we also have the constraint that 2.2 times Y is X_0. And that's just from the volume expansion that we know.

So if we can simply now just equate these two equations, we have two equations. And we're going to set X_0 in equation 1 equal to X_0 in equation 2. And then you can simply solve for Y. So Y turns out to be about 0.4 micron. So that says, we need to consume about 0.4 microns in order to just fill the trench up to the original surface.

So I'm going to actually go down here. From the original trench depth, I'm actually going to consume 0.4 microns down into the wafer. So X_0, then the oxide thickness must be 2.2 times that. So it's about 0.9 microns. So we really have the answer. We know we need to grow at 1,000 degrees, 0.9 microns.

We were told we can do it in a wet ambient. So you can calculate B and B over A using the table in table 6.2 of your text. So you know this is the Deal-Grove equation. You can just calculate B and B over A and solve for the time. Or if you want to go, on the next slide, I'll show you use the graph.

So on slide 18, these are just the oxidation kinetics or oxide thickness versus time in a wet ambient. And I was told we were at 1,000 degrees. So that's this line right here. And we know we need to grow a 0.9 microns. So that's this black line, horizontal line. And you can get it roughly here. It's just 3.8 or so hours, depending on how accurately you can read off the plot.

If you want to get a little more accurate, you go ahead and calculate B and B over A and plug in the equation. But you can always check the plot. That way, you have a sanity check that you made didn't make some calculator mistake or type in the wrong number. So that seems like a reasonable example of how you would do a simple calculation.

So let's go on now to slide 19. So that was a Deal-Grove model. Immediately, at the same time Deal and Grove published it, they realized that there was a major problem with the model, even when it was first proposed. It doesn't correctly model what they knew to be the thin oxide growth kinetics. And you can see that in this plot.

This is a plot I took out of Andy Grove's book published quite some time ago, but still a famous book, oxide thickness in microns as a function of time. This was for dry oxide grown at 700, very low temperature. And these little bullets or these open symbols are what they measured as the oxide thickness.

And, in fact, what they noticed is that the model itself, which would be in this range, the dry oxides grow much faster for thicknesses below about 20 nanometers, or 0.02 microns, than predicted by the linear parabolic model.

So the linear parabolic model would have had them sort of going like this. And here, these oxides are already going up much faster. So there's a rapid initial oxidation phase, which does not take into account when you have a linear model in the beginning and then a parabolic model. So people knew right off that this really wasn't going to model the initial oxidation kinetics very well.

So if we can go on to slide 20 then, since that time, there have been a number of people have made suggestions to explain why the initial kinetics are faster than what Deal-Grove model would predict. There isn't really one model that's widely totally always accepted. There's a couple of them that are used in simulators today. So if you go to the Supreme Four Manual, and you read up about what thin oxide models are available for you to use in your numerical simulations, here's an example of three of them.

There's the Reisman model, Han-Helms, and the Massoud model. So I want to go through these three there maybe the most popular or well known of all the proposals. So on slide 21, we have the Reisman model. And they actually proposed a relatively simple power law.

But it seems and some people regard it as more of a mathematical fit that quote unquote, fits the data for dry O2 over a wide range of thicknesses. And so their power law is actually looks different from the Deal-Grove model. They say the oxide thickness X0 is some constant A times T plus TI. TI is the initial oxide thickness or time corresponding to the initial oxide to the B power.

Or you can rearrange it so it looks something like this in terms of A and B. Look at this compared to Deal-Grove. You can see it is actually quite different, although they still have a couple of constants. They have a lowercase a and a lowercase b for a given set of process conditions. So similar to Deal-Grove, you have two constants. There are different numbers.

This actually works. The fit mathematically fits better in the thin oxide regime. The problem is it's not as appealing in the sense it's not as quite as simple as Deal-Grove. But their physical reasoning was that the interface reaction controls the oxidation process at all times. This is what they thought. But they believed that this volume expansion, the need for the volume expansion and the flow of the oxide, the interface, control the growth kinetics, the rate.

And A and B were believed to be some time-dependent flow parameters for the oxide. So not as obvious or as simple as appealing as the simple Deal-Grove model. But they also found that mathematically, it just makes a better fit. So people do sometimes use this type of model and then fit their own data with their own A and B constants.

So that's the Riesman model. Han and Helms came along and got additional data. And they actually said this. They actually said that there are two in parallel, a couple of different oxidation processes taking place in the model. So perhaps, they said, not only O2 molecules, but oxygen atoms can diffuse through the oxide, and maybe in parallel.

Or maybe there's a diffusion of oxygen vacancies. And so you have something diffusing in and something else diffusing out. And nevertheless, they said there are two species that you need to take into account. And they have associated reactions at the interfaces. So it looks very much like Deal-Grove. In fact, this is Deal-Grove. If you express Deal-Grove, remember, in terms of an oxidation rate, it's just B divided by 2X0 plus A.

This looks a lot like Deal-Grove. But they've added a second term. And this second term here is to correspond to that second oxidation process happening in parallel. And they propose that all rate constants would be of an Arrhenius type nature, based on the data that they fit. It's just like Deal-Grove. But, in fact, so you say, well, you should double the number of constants in Deal-Grove. So instead of two, you should have four.

In fact, what they found, they only really needed three parameters to fit the data. So they needed a B1 associated with this process and B2 associated with this process and B2 over A2. So they found A1 equal to roughly 0.

So those of you who are skeptical in the audience will say, well, OK, Deal-Grove had two parameters. And Han and Helms have three. So if you add more and more parameters to a fit, you can always do a better job mathematically of fitting data. Everybody kind of knows that from your mathematical experience.

So you could be skeptical and say, all right. They just add another constant. And perhaps that's true because the truth is, there isn't really any definitive data saying what the species that are diffusing through there. But nevertheless, they did come up with a model that's reasonably simple that does a better job in the thin oxide kinetics.

So actually, let's compare these two models just on slide 23 with a Deal-Grove model. So the x-axis here is oxide thickness in microns. And this is time. I'm sorry, the y-axis is thickness. The x-axis is time. And this is plotted for atmospheric pressure dry O2 at 800 degrees.

And so let's look at the different models. This line right here, unfortunately, they're all the same color. I apologize. But you can see where the arrow is pointing to the Reisman model looks something like this. Han and Helms, amazingly, different mathematical equations, but they look pretty darn close. They're not exactly on top of each other in this thin oxide regime. But they're pretty darn close.

The Deal-Grove model with tau equals 0 looks like this and with tau equals 8 looks like that. Doesn't actually approach either one of them very well until you get into the thicker regime up here. But in this thin regime, say below 200 angstroms or 0.02 microns, Deal-Grove doesn't really approach either one. And Han and Helms and Reisman do actually approach the data quite well. They will converge though when you get to oxides thicker than 200 angstroms.

There's one more model that people actually are very fond of. And I'll talk about that now. That's shown on slide 24. This was the Massoud model. This was published quite some time ago at the electrochemical society. And the neat thing about Massoud's experiment was they actually measured in situ. In the furnace, they actually measured the thickness of the oxide growing as a function of time.

So the other people like Deal-Grove and Han and Helms, all those people, they put a couple of wafers in and take them out, put them in for different amounts of time. And so they get a few certain data points. How many wafers can you put in? Maybe five or 10 in an experiment. So you can't really measure the kinetics in great detail if you're just pushing and pulling wafers out of the furnace and just measuring them outside.

What Massoud decided to do to study the very thin regime was actually set up-- and it was non-trivial. He hit a diffusion furnace. He had special windows placed into it where he could put laser beam. So you could put a laser beam coming in, laser beam coming out. And he could do in situ measurements by ellipsometry of the oxide as it was growing in the tube. It doesn't sound like a big deal.

But it turns out, it's pretty challenging because there's all kinds of thermal expansion going on. And the wafer has to sit still and not move while this laser comes in, hits, and bounces off because in the ellipsometry, the angle of incidence is very important and all that. So designing the equipment was like a huge part of this PhD thesis. But the nice thing is, it did produce a lot of data. Look at all these data points. They're all basically give you a nice looking curve. Instead of one point here, one here, and one here, he published in situ measured oxidation rates.

Now, this is the rate. So this is not the thickness. This is angstroms per minute as a function of the thickness grown. And indeed, look below 200 angstroms. All of these rates are much larger than they are where they reach steady state, they reach smaller rates. And these are at three different temperatures, 800, 900, and 1,000. And in the thin oxide regime, the rate itself appears to be going up exponentially as you get to thinner and thinner or shorter and shorter times.

So interestingly, initially, it appears to be, in fact, an exponential process with time. So it's growing very rapidly. And then it starts to converge to a process that looks more linear and then eventually, parabolic. Only with that kind of special apparatus could one get enough data point density to really see that though.

So this is probably the most extensive experimental study that has been done. And what he did was he took Deal-Grove, just as it came, right out of the box, and added one more term. But the term was exponential in thickness. So the oxidation rate had the Deal-Grove B divided by 2X0 plus A plus a term that went like C times exponential of minus X0 over L.

So what is this? Mathematically, this is a decaying exponential. When L gets very large, L is much greater than 200 angstroms, this number goes away, right? You can always make it go to be very small compared to this term. So he wanted it all converged to Deal-Grove at thick enough oxides. But for thin oxides, this is going to dominate, this C term.

And in order to do that, he had to choose L to be a certain number so that it had the right decay length. And L, he found experimentally, it should be about 7 nanometers so that the second term will decay rapidly enough by the time you get to 20 nanometers or so.

In fact, this model does agree very well mathematically with experimental data. It's actually very easy to implement in simulators like Supreme 4. And so it has been found to be implemented in Supreme 4. And in terms of physical explanations, there's a few that have come out. But none has really been clearly identified as being the right reason.

So for now, we use it as sort of a mathematical model or an empirical relation. It's very handy. And it can be reasonably accurate. So it's appealing in that sense. In terms of the physics and the chemistry, it's not a completely obvious.

OK, so those are the three thin oxide models. And hopefully, you'll get a chance to study some of those a little bit more in one of your homework problems. So now let's vary something else. All we varied so far is the following. We varied the ambient. We said it could be dry O2 or water and temperature. What else can you vary? Well, it turns out, people often do vary the pressure. That is the partial pressure of the oxidant.

If you want to grow a very thin oxide, and yet, you want to grow at a higher temperature because you want to get different electrical quality of the interface, you can't grow too fast. So you might want to slow it down by diluting. You can just take the oxygen and flow mostly nitrogen, and have oxygen flow be a very small fraction of the gas stream. That's a way of diluting the gas stream and growing at lower partial pressures.

So just as a reminder, Deal-Grove predicts that the oxide growth rate should be directly proportional to the pressure of the oxidant. C star, effectively, if you can think of it, should depend on PG where PG is the bulk gas pressure. So basically, if I were to lower the gas pressure, you would expect that you could lower the oxidation rate.

So just to be a little more specific, we go on to slide 27. What people have found, though, is it's not completely linear. Experimentally, if you lower the partial pressure of the oxidant, in fact, B and B over A do scale like P, where P would be the partial pressure, for water oxidation.

And the rate constant B does scale like the pressure for dry O2. The one rate constant that's anomalous is the linear rate constant. The B over A for dry O2 doesn't go quite like exactly like pressure. It goes like P to the N, where N is close to 1 but not really 1. It's somewhere between 0.5. So it's a power law between 0.5 and 1.

So this suggests that at least for oxygen, the case of S maybe has some dependence on pressure that's nonlinear. So these are some empirical models that you'll find in Supreme 4, for example. The linear rate constant may go some number B over A to the I times pressure where I here refers to the intrinsic values at 1 atmosphere. N here might be 0.7 to 0.8.

And you notice, you specify the pressure P. So maybe you say at half an atmosphere or a tenth of an atmosphere. But we're not changing-- think of it this way. You're not actually changing the pressure in the tube. The pressure of all the gas in the tube is still an atmosphere. What we're doing is we're flowing mostly nitrogen. And we're putting a small amount of oxygen in the gas stream.

So let's say I wanted to oxidize at half an atmosphere. Well, I just cut the flow. I use equal flows of nitrogen and oxygen. So we're diluting the pressure, basically. We're diluting the oxidant in the gas stream.

So here's an example. Suppose I dilute it by a factor of 10 to 1. So I'm flowing 1,000 SECM of inert gas like argon or nitrogen. And I'm flowing only 100 SECM of oxygen, dry O2. This is what we would get as the oxide thickness versus time.

And the solid lines were if you were to do the very brute force simple idea that it was all scaled like P. And it's just down. The rate constants are down by a factor of P, where p here would be a 1/10 the pressure of atmospheric pressure. The dashed line is where you use a little bit more sophisticated model where, in fact, the B parameter scales like pressure. But the B over A parameter scales like pressure at the 0.8, so just slightly.

So in fact, the kinetics would be slightly faster, the oxide would be growing slightly faster, than you would have predicted using very simple dependence on just saying it depends strictly on P. So it gives you a little bit more accuracy using some of these more empirical models. But again, the idea here is, if for some reason, you need to control the oxidation rate better, you want to slow it down, you just dilute the oxygen in the furnace.

Oh, I just showed some specific examples here of trying to read off what the difference would be. You'd make an error of about a 100 angstroms if you use the simple assumption compared to what's a little bit more accurate here at 1,000 degrees. OK, so let me just get through. That's the standard kinetics on slide 29.

Let's go on to slide 30 then. OK, so what I just talked about was trying to dilute the oxidant so you can grow slower at a given temperature. On slide 30, we're doing just the opposite now. It's not so easy to do from a physical point of view. But you can imagine growing at higher than 1 pressure. Let's say the pressure tells you what? It tells you the number of atoms or molecules, species hitting the surface per unit time.

So if I double that pressure, I can double the amount there hitting that surface. I can double the flux to the surface. So you expect the growth rate to go up. It's just not in your intuition is how I could do that. Here, you have a tube that's made of quartz. I have it open on both ends to the ambient. And I'm flowing some gas to it. How are you going to get a higher pressure? We're all at atmospheric pressure. It doesn't work. But you can do it. You can build a system where outside that quartz tube, there's a big envelope made of stainless steel like a diver's tank.

They have training tanks for deep sea divers. They put these guys in these tanks. They seal it up like a submarine. And they can stuff more and more gas in there, and you can increase the pressure on the diver. So that's exactly what people do with high pressure oxidation. It's a very special type of equipment. It's not something you're going to find in a standard fab. They are available though. They've been commercially made. And people call it HiPOX, the High Pressure Oxidation System.

So here's an example of how the kinetics would go up. And this I took from Mayer and Lau's book, 1990. What they're showing here is on a log plot here, a semi-log, so it's log on the Y, linearly on the X, oxide thickness versus time. And this is 100 in steam. And they're looking at versus temperature-- I'm sorry, looking at a couple different curves here at different pressures.

So at atmospheric pressure, it looks like this. At 5 atmospheres or 10 atmospheres, it looks like this. Here we are 20 atmospheres. Look how much faster we can grow. So this can be used to increase the oxidation rate at low temperatures. If you need to grow a thick oxide thermally, you need to have the properties of a thermal oxide, which are quite unique. Its density is well defined.

It has much better electrical properties than you would get by depositing an oxide. It actually consumes silicon, which is not true of a deposited oxide. Then you have to go to a hypoxic process. Or let's say you have some structure on the surface like silicon germanium or some material that can't go to high temperature because it'll cause the material to diffuse or to relax, or you're trying to limit the amount of diffusion in the substrate.

Well, and you need to grow a thick oxide, you can do it in a hypoxic machine at 800, where you never would be able to do grow a half a micron at 800. But you can do it in HiPOX at 20 atmospheres. So for materials and processes that are getting limited by temperature and they have to go lower and lower in temperature, HiPOX is a way of increasing the oxidation rate.

So in fact, here are some data shown on slide 31 on HiPOX I took from Simon Z's book. And what it shows is the measurements on the y-axis, the parabolic rate constant-- so this is the B parameter-- as a function of 1,000 over T-- so it's a semi-log kind of Arrhenius plot-- in steam. And this has been measured for a couple of different orientations of silicon.

And what you see here is this solid line, which has two different activation energies. There seemed to be a break point here somewhere around 950 degrees but actual activation energy close to one electron volt. This is the B parameter. Here at 5, 10, 15 and 20 atmospheres, this is what the B parameter goes up by. And in fact, I marked this blue line, which is at 800 degrees.

And in fact, I'm reading off the chart here, the B parameter is almost exactly a factor of 20 faster at 20 atmospheres than it is at 1 atmosphere. So we can go up to 20 times the oxidation rate at 20 atmospheres at 800 degrees. So it does scale reasonably well, this parameter does, with pressure at higher pressures.

On slide 32, I'm actually showing you-- remember, I said, you can build this complex diver tank looking thing with a great big steel vacuum system that can hold the pressure and pressurization system. I've taken this from Simon Z's book. He actually talks about how you would do a HiPOX run, if you had such a thing. Just so you can think about it, what happens is the wafers get loaded here at atmospheric pressure. Down here, this solid bar, it means you're at atmospheric pressure.

Obviously, you can't pump up the diver tank until the hatch is closed. So you load them in atmospheric pressure. And you have a sealing step. And here, you're starting to increase the pressure, again, using O2. You do some purging. And then at some point, what you do is you bring this outside steel shell and the tube up together in pressure. And you go up to, say 10 atmospheres.

So your pressurizing in moisture. Then you do your oxidation. Then you depressurize and bring it down to atmospheric pressure and unload. So it's perfectly possible to do this. Obviously, if you try to do it with a quartz tube and put end caps on it and pressurize it, of course, the quartz would explode all over because it's not a strong material. But if you have the stainless steel shroud all around it, then the differential pressure across the quartz is very close to 0. Or it can be just maybe one atmosphere.

We know it can withstand 1 atmosphere. But that way, you keep the quartz from exploding all over the place. So it's a special piece of equipment.

OK, so then let's go on to slide 33, so HiPOX, we can do low pressures. You can do high pressures if you have the right equipment. It turns out, simple oxidation also isn't always done just with oxygen or water vapor. Small concentrations of certain other species are often added.

For example, HCl is sometimes added to the O2 because it turns out, people have found by adding that, they can reduce oxide defect densities or they can reduce contamination levels. In the liquid form, HCl is very good at etching off metals from the surface. So perhaps, it has a similar role when it's in the vapor phase.

So sometimes, people add a small amount of HCl to the furnace. It can produce chlorine, which can react with maybe trace metals by this kind of process that's shown here. So it's not unusual to see, in manufacturing processes oxidation, in not just O2 but HCl, O2 plus HCL.

So these bullets here in the middle of the slide are just some generic kind of observations that have been made on these different ambients. Well, just studied, in the Deal-Grove, we know in detail that you can get 20 to 50 times faster than dry O2 just by putting in wet O2 in steam. So we know that. 3% chlorine in the ambient can increase the growth rate by 20% to 30%.

So also, HCl is a way of, at a low temperature, is boosting the growth rate. The name of the game these days is keep the dopants from moving. But keep the growth rate high enough to give you a reasonable thickness. So you can add a little chlorine to speed up the growth rate. You can also add NF3, another commonly used species. A small amount of that will increase the growth rate by two to three times.

Again, this is adding in a lot of new chemistry. And there's no way in an AB initio, in a first principle sense, to really calculate what those rates. Are so in simulators like Supreme 4, if you go sit down and tell it, oh, I want to grow an oxide in chlorine and oxygen. It doesn't simulate the chemical reaction and all the kinetics and, come on, spit out a number. It has a lookup table of B and B over A values that have been measured in the literature for commonly used ambients like chlorine and maybe NF3 and a few others.

So it looks up in the table what B over A and B should be. As a function of the pressure here, this pressure, delta or chi, these are some functions of the concentration of this additional species that's in the furnace. Could it be HCl or whatever. And they've been determined very empirically. So don't expect first principles models, but maybe some mathematical models would be in Supreme for different ambience.

But again, they've all been calibrated to somebody's experiment. So they may not exactly agree with your experiment. You might have to calibrate it yourself.

Here's an interesting example on slide 34. These are some kinetics in-- HCl has been used for quite a while. It's even old fashioned now. Something that's a little bit newer is people oxidizing an ambience like NO, N2O, and even implanting nitrogen. Remember, we talked about oxynitrides are the first form of high K.

Oxynitride, which is a mixture of silicon, oxygen, and nitrogen, has a higher dielectric constant. You can make it a little bit thicker. And so you get less gate leakage. So gate oxides of oxynitride were are popular for people to study. But here's an example of trying to fabricate an oxynitride where the kinetics are really crazy.

They're not at all linear parabolic. And in fact, this is a plot here on the y-axis is oxide thickness or oxynitride thickness, if you want, as a function of the nitrogen dose. So here, the substrate was implanted first with nitrogen-- and we'll talk about ion implantation-- to a certain number of atoms per square centimeter. And they're plotting the oxide thickness here at different times, so say, for 25 minutes as a function of dose.

And you can see, depending on how much nitrogen you implanted, the thickness actually can go down. So the oxidation rate can depend on exactly what is at that surface. Was it silicon pure? Or was it silicon that had been ion implanted with something? So this needs a qualitatively very different model from Deal-Grove because you need to specify something that the Deal-Grove doesn't even treat. Deal-Grove assumed you had a perfectly pure silicon wafer that you were oxidizing with just silicon atoms in it.

Once you start implanting nitrogen or some other species, all bets are off. You're not going to widely find simulators that can support this type of model. If you want to do this, you're going to have to model it yourself by taking data from the literature and fitting it or maybe doing your own experiments.

So let's look at slide 35. There's another parameter we need to vary here. You can imagine varying. And this is a relatively simple one. And that's the orientation of the crystal, the face of the crystal that you're oxidizing. We know that the B parameter, which depends on the diffusion through the oxide, is independent of orientation. B over A, which is the reaction rate at the interface, of course, it's going to depend on the number of atoms available per square centimeter to react. So it does depend on orientation.

So and these are the dependencies, just to give you rough numbers or relationship. 111 is the fastest. It has the greatest number of atoms per square centimeter. 110 is second. And the slowest is 100. And in fact, these are the relationships. If we take the 111 rates and we divide them by 1.68, we get the 100 rates. And the relationship between 110 is in between.

And again, it's related to the silicone atomic density, the number of atoms per square centimeter that are available. And you can believe that because in order to have the surface reaction take place, silicone bonds have to be broken. So and the number of atoms available is going to be important in that case.

Oh, slide 36 is an interesting example of how even, you might say, well, I only use 100 wafers. What do I care about oxidation on 110? Well, you will care because you often etch a trench and then need to oxidize the trench. And of course, the trench has different orientations of its walls. So this is a Supreme 4 simulation of a trench that was etched into a 100 wafer.

So 100 means what? It means the surface of this wafer, this plane, at the very surface, is 100. The 100 direction or 01 is pointing up. But if once I start etching into it, the planes that I expose will be different orientations, depending on how I design my trench, if I etch into the crystal.

So for example, I etch a trench here down into the crystal. The bottom face is 100 because the bottom face is the same as the top face, the face of the wafer. But the sidewall, this face right here, the sidewall that I've exposed, depending on how I oriented this trench, which is assumed to be square or rectangular, the sidewall face could be a 110, depending on how I orient it with respect to the flat.

So you can immediately expect a different oxidation rate on this face right here from the oxidation on the bottom. And this is an example of a simulation that was done 30 minutes in 900 degrees C in water. And you see the thicker oxide that grew on the sidewall. That's pretty much-- in this bulk region of the sidewall-- that's entirely due to the fact that it's a different orientation.

So this simulator took that into account. It's a little bit thicker here than it is on the bottom. What is this, this weird looking corner effect? That's not orientation so much. This actually shows a two dimensional effect we'll talk about next time, which is a stress effect. When you're oxidizing a corner, again, the oxide has to expand. In a corner, it's a little bit hard to do some expansion. We're just going to turn out, the oxidation rate in these sharp corners, it's going to end up being slower due to stress. So that was taken account into the simulation.

So let me just summarize what we have from today. The basic growth mechanism is transport through the oxide by diffusion and then chemical reaction at the interface as opposed to the silicon doesn't come out of the lattice and go up to the surface. We generally have a linear/parabolic law when we have planar oxidation. Next time, we're going to see there are a lot of different cases where it's non planar and it's not just simple linear parabolic.

The Deal-Grove model isn't completely accurate in the thin oxide regime. But it can be corrected by using Han and Helms or Riesman or you can use Massoud. And they all can be designed to converge to Deal-Grove for when you're at thinner and thicker oxides, say above 200 angstroms. And the growth rate varies with orientation. 111 is the fastest and 100 being the slowest.

So that's mostly what I wanted to cover for today. I have a couple of things for you though. Your homework is going back. So you can come up and grab your homework. It's been graded, up here, homework number one. And the solutions are also up front. So if you can come on up and take the homework and the solutions for homework number one, that's good.

If you weren't here last time, you didn't get to see the Intel wafer, I'm going to hand it back after today. So you're welcome to take a look at it. Oh, these are not in any particular order.

STUDENT: [INAUDIBLE]

JUDY HOYT: What?

STUDENT: People extracted the strained silicon?

JUDY HOYT: Yeah, people tried. Jeff tried. It seems to be pretty darn close, a change of 1% in the lattice constant because I really haven't changed the number of atoms per square centimeter. What does change is as soon as you hit the silicon germanium, oxidation rate goes way up.

STUDENT: [INAUDIBLE]