

JUDY HOYT: So we're moving along pretty quickly with the first part of the course here. I'm hoping you'll be reading Chapter 3 of the textbook at this point, which is on-- covers crystal growth, wafer fabrication, and some of the basic properties of silicon. I'm going to finish up Chapter 3-- the lectures on Chapter 3 today.

And then I'm going to start Chapter 4, so you can start that as well. Chapter 4 is relatively easy reading, not tremendously mathematical. It's about how we-- although it's a very important topic. It's about how we clean wafers and how we get our impurities in wafers.

So the last lecture, let me just remind you what we talked about. We talked about cubic crystal structures. We discussed wafer manufacturing, basic characteristics of Czochralski wafers. And we'll continue on with that a little bit today. We presented a very simple mathematical model differential equation for-- based on heat flux for modeling the Czochralski growth process.

And, at the very end, I had three slides on the different methods of fabricating SOI, silicon insulator. And, in fact, if you're interested in SOI and getting into it more detail, that's a topic you might consider for one of you-- someone to do for your final project. There's a lot of literature out there on how SOI is really made. I only gave you three slides. I'm sure the class would like to hear more about it if you want to do that as your research topic.

Today I want to talk about more specifically the mathematics and the statistics of point defects in silicon, a little bit more detail about carbon and oxygen. They are so critical to Czochralski silicon and its properties, and then introduce Chapter 4.

So let's go on to the second slide, page slide two. We're going to just-- since we're going to talk about modeling point defects, let's just define them once again. We did this last time, but this is a very schematic ball and stick two-dimensional diagram of what a lattice might look like, let's say, highly, highly simplified.

And the vacancy is shown here in this region where a silicon atom is completely missing from a lattice. You just pulled it out. And there are broken bonds associated with that.

An interstitial is a little bit more complicated. There's actually-- there's two different types of interstitials here. There's one which is shown right here. And this is a pure interstitial. It's actually an unbonded extra silicon atom, like you just stuffed an atom in there in an interstitial space. It's not bonded to anything. That's the pure interstitial.

Some people, the defect experts, also refer to a type of defect which is an interstitial type called an interstitial sea, by analogy to a vacancy. And what it is, it is an extra atom, but it's actually two atoms that are sharing a lattice site. So these two are together. And there is some bonding to neighbors as opposed to a pure interstitial.

A lot of people believe that the interstitial sea is actually the most likely configuration of the interstitial defect in silicon because it takes a lower energy to create it. Just to share a lattice site isn't that hard compared to stuffing an unbonded atom in there.

But the truth is, the distinction is really not important in process modeling. For the purpose of this class, for modeling processes, there's two types of defects. We think of the i-type, the interstitial type. Exactly how it's configured, we're not sure. We'll just use the symbol I, capital I, to refer to an excess silicon atom.

And these point defects, these V and I, they turn out to play really fundamental role in a lot of processes. They control diffusion. When we get to Chapter 7 and 8, we're talking about diffusion, extremely important. The oxidation, the point defects play a very important role activating dopants.

So there's a lot of importance. And that's why we're going to spend a little time today on understanding the statistics of vacancies and interstitials.

So let's go on to slide 3. If you're interested, I've listed here five papers, all of them fairly old but classic. Shockley, the Shockley, who also did the transistor studies, had a paper in the late '50s on the first understanding of statistics for vacancies and their charge distribution.

Shockley and Moll, also famous names, Watkins, did some of the very first studies where people actually detected vacancies in the lattice using a type of technique called EPR, electron paramagnetic resonance. So, if you're interested in that, you can go back to some of the original papers. People actually have some evidence that these things do exist.

Let's go on to slide 4. And we're going to talk about modeling these native point defects. Basically, the presence of native point defects, that is, vacancies and interstitials, it minimizes the free energy of the crystal. It increases the configurational entropy.

If you want a derivation of that, you can look at Mayer and Laus, that's one of the textbooks I referred to earlier in the class. Their 1988 text has a little derivation on that. But, basically, it seems physically intuitive that we can write the following expression for the concentration of neutral vacancies and neutral interstitials.

And the notation here is a little bit cumbersome, c meaning the concentration. The star here, the superscript star means-- the asterisk-- that we're in equilibrium. We're going to talk a lot about in this class when we're not in equilibrium, when we have injected excess vacancies by various processes, ion implantation, or whatever. But, for now, we're going to stick to equilibrium.

This subscript tells you the type of defect, vacancy or interstitium. And then, this little zero means neutral-- uncharged. Well, as we'll see today, vacancies and interstitials can exist in charged states. But, for now, if we're just talking about the neutral and the uncharged vacancies or interstitials, we can write them in this type of expression. They have a temperature dependence.

They depend on a constant, n_s , which is the number of lattice sites. We know that. That's $5 \cdot 10^{22}$. And they depend exponentially on the formation entropy of the defect, which is this s_f and on the formation enthalpy of the defect over kt . So we expect to see this thermally activated behavior. So the concentration of these things is going to go up exponentially as I increase the temperature.

In general, down here, the equation at the bottom is true. The concentration of neutral interstitials is not in equilibrium. It's not necessarily equal to that of vacancies. It may be easier or harder to form. So don't think that they have to be equal to each other, because there's a lot of different ways to form these two different types of defects.

Let's go on to slide 5. In fact, let's talk about a couple of different ways we can take a perfect crystal and form these defects. The simplest way-- and if the way you would probably imagine, if I asked you is, well, how do we create an interstitium and a vacancy?

Well, it just pull a silicon-- rip it out of its lattice site and stick it somewhere, an extra atom, then have a vacancy where the atom was left behind-- where the space was left behind, and I have an interstitial. In fact, this pair right here, this vacancy interstitial pair, is called a Frenkel pair. Incidentally, this can happen.

People believe it costs a fair amount of energy. That is, it's somewhat hard to do because you have to break all four bonds because you're in the bulk of the crystal. Remember, in silicon, every silicon atom is four-fold coordinated. It has four nearest neighbors that are with covalent bonding.

Obviously, the Frankel process is going to create one interstitial for every vacancy. So if Frenkel were the only mechanism by which we could create point defects, then $c_{sub\ i}$ would be equal to $c_{sub\ v}$ in equilibrium.

But it turns out that's not the case. It turns out surface processes are a dominant way of creating these point defects. For example, we can have generation and recombination of vacancies or interstitials at the surface. And they may create only one type.

For example-- and this is, again, a schematic. This is meant to be now a cross-section of a lattice. So imagine up here, this is the top of my lattice. Above it is free space below it is the wafer. And so, this is the top row of surface atoms. So I can take an atom off the surface.

And it's not for a full bonded because, by nature, the surface is a point of discontinuity. There's only a couple of bonds for an atom at the surface. I can-- doesn't take much energy to pull that out and stuff it-- and put it in the bulk and create a silicon interstitial.

So there's a lot of different ways. The surface generation and recombination is very common. The crystal has a number of different ways to achieve different concentrations of vacancies and interstitials, even in equilibrium.

And when we change the temperature of the crystal, the concentration of the vacancies, the concentration of interstitials will change by the creation of either Frenkel pairs or atoms moving in and out of the surface and interstitials diffusing in.

So we'll go on to page 6, or slide 6. It turns out that the equilibrium, no one's ever really measured exactly the equilibrium concentration of v and i by any direct process, particularly at the silicon process temperatures. The truth is, the concentration is so small compared to numbers-- physically reasonable macroscopic numbers, like the number of silicon atoms.

The number of silicon atoms 5, 10 to the 22. We're going to see that vacancy concentrations are in the 10 to the 12 range, one part in 10 to the 10, much, much lower. And so, it's just very hard to measure them directly.

But people have estimated them. And the way they've estimated them is going to-- a little bit circular argument. But it's going to be come out in this course, by fitting impurity diffusion data. In other words, people hypothesize that the diffusivity of a dopant requires the formation-- the existence of certain vacancies.

And then, they see diffusivities increasing and decreasing according to temperature or charge states they say, oh, the concentration of vacancies must be going up by this much. So people measure diffusivities and infer concentration of point defects. It's a little bit circular but, in any case, that's how a lot of people have estimated these things.

So don't expect to find anywhere in the literature some perfect equation. There's a lot of controversy. But there's some agreement these days that these are rough numbers. So here's a rough equation for the concentration of neutral interstitials in equilibrium. It goes something like this.

Again, it's exponentially activated. As I increase the temperature, I get a lot more-- I get more interstitials exponentially. But if you calculate it, put in room temperature here, you calculate it out, it's close to zero. I mean, there's just-- there aren't very many at room temperature of these things. The concentration is quite small.

At 1,000 degrees, that's where we might be doing an oxidation. We put a wafer in a furnace and oxidize it at 1,000 degrees or something, or just anneal it. At 1,000 degrees, it's something like in the range of 10^{10} to 10^{12} to 10^{14} , still very small compared to 5×10^{22} . Nevertheless, it can have a huge impact on the dopant diffusion.

Typical doping levels are in this 10^{14} to 10^{15} to 10^{20} range. And, again, the point defect concentrations are smaller than these. This will have some important implications. What it means is that the doping, the dopant concentration, the concentration of arsenic or boron, is what sets the electron concentration and the hole concentration in the crystal, not the vacancies or the interstitials.

Let's go on to slide number 7 and a little bit now about-- so that all was about neutral. So, as we were saying, as we increase the temperature, the concentration of neutral vacancies and interstitials goes up exponentially.

But it turns out, these defects can have charge states as well. In fact, they can be singly charged, positive or negative, or doubly charged, double plus or double negative vacancies. They've been identified experimentally using things like this electron paramagnetic resonance.

And particularly when people bombard the lattice with high-energy electrons and then they see what kind of vacancies they can see by these techniques, and they do see them in various charge states.

Interstitials are believed to have similar charge states, but they really haven't been measured-- the charge states of the interstitials are more difficult to measure, probably because the interstitials diffuse throughout the lattice much faster than the vacancies. We're going to see the interstitial diffusion coefficients themselves for excess silicon atoms. They can go whizzing through the lattice. So they're even harder to see.

So let's look at another, again, a very two-dimensional schematic way of looking at a vacancy. That was what I'm showing here at the bottom of this page. The silicon atoms are these round dots in black. And these are supposed to represent bonds, these ellipsoidal things, it's supposed to be a covalent bond. Each bond should have two electrons associated with it.

Remember, silicon is in the fourth column of the periodic chart, so it has four valence electrons. So this atom here can be completely satisfied with four covalent bonds. All of its electrons valence electrons can be bonded.

If I rip out a silicon atom here, the vacant site, it's going to distort the nearby lattice. It's going to have some change in the bonding configuration. And this causes a local change in the energy band structure. For those of you who've had solid state physics, you know about energy bands. It's going to cause extra states, actually, in the energy gap.

And we're going to call these deep levels. And we'll show you how we represent them-- are going to be introduced by the presence of these defects. These deep levels exist in the bandgap, and they're split off. They're not in the conduction band or in the valence band. So having a vacancy or an interstitial creates a localized deep level in the bandgap.

So let's take a look at page 8 from your handouts, or slide 8. And this slide shows very schematically or approximately the location of where people believe they can see energy levels exist in the silicon bandgap. Now, hopefully, you've read through Chapter 1. If you haven't, you can, right after this lecture, you can run out and read Chapter 1.

And so, in that chapter, it talks a little bit about the silicon bandgap. This line up here is meant to represent the conduction band energy. Above this there are lots of free electrons. This line down here at the bottom represents the valence band energy, the top of the valence band.

In between there's a forbidden gap. So the carriers cannot have energy-- free carriers don't exist with energies in that energy range. And we know this bandgap is something on the order of 1 electron volt, 1.1 for silicon.

These lines here, the line marked v -double-minus, the line marked v -minus, v -double-plus, and v -plus, they represent the positions in the energy band-- in the forbidden gap of these defects. These are localized states that are formed by the fact that we're ripping electrons-- we're taking silicon atoms out of the lattice. We no longer have a perfect lattice. But we're putting an extra interstitial silicon atom in the lattice. So and these are in various charge states.

This level here, in the mid-gap, it's right in the middle of the bandgap is marked e_i . That's the intrinsic energy level. So that's, by definition, that's pretty much close to halfway, right at the mid-gap point between e_c and e_v . And this line right here, marked e_f , is the Fermi level-- the Fermi energy.

And, again, hopefully you've read Chapter 1. Fermi energy is a concept that comes in. It basically is the energy level at which the probability of finding an electron is $1/2$. That's the exact definition.

The main thing you need to know about the Fermi level is, if the Fermi level is at mid-gap, the semiconductor is intrinsic. There are very few electrons and holes. The only electrons and holes we have are those created by thermally breaking a bond and ionizing an electron from the valence band and bringing it up to the conduction band. So the intrinsic concentration is going to be relatively low.

As I add n-type dopants, the Fermi level moves up. In fact, this picture is for where I've added a certain type of n-type dopant and it moves up towards the conduction band. The higher it goes towards e_c , the more free electrons we get in the crystal, the more free electrons I have in the conduction band. If the Fermi level moves down towards e_v , I have more holes.

So you can think of it as the position of this level in the band gap going up creates more electrons. As more holes are created, it goes down. It's just a way in energy space of keeping track of those densities.

So let's take-- let's just convince ourselves that it's possible for these vacancies to exist in charge states. And so, here's a very crude picture of a vacancy-- of a neutral vacancy. So what's happened here is we've taken the silicon black atom out in the center. And we've taken out all four of its valence electrons. They've been pulled out.

So what happens is, well, again, each one of these little ellipsoids is supposed to have two electrons in it. And this is supposed to represent one dot, one electron from this atom. And this represents one electron from the atom down here. So there's a covalent bond here. There's some kind of a bond. There's a bond here. So this is all satisfied.

So there are-- this area of the crystal is net neutral. There's a vacancy there, but there's no extra electron or absence of an electron. So we call-- we can imagine this in your mind as a neutral vacancy configuration. Again, it's very crude but it gives you some pictorial idea.

In the center, at the bottom, this picture represents a single negatively charged vacancy. So what we do is we have a silicon vacancy. We ripped out a silicon atom. But we left behind an extra electron or an extra electron found its way there, however you want to say it. And here it is. Here's an extra electron sitting here.

And so, then if you count up the electrons here associated with this region, we have one extra. So this has a single negatively charged net in our charge state. So this can represent a v -minus pictorially.

And, similarly, we can take out a silicon atom and we can pull an extra electron out. So not only the four valence electrons, but let's take even a fifth one out and remove it. And now, this region has a net positive charge associated with it because we have the nuclear-- we have the charge of these silicon atoms, and we've unbalanced them. So now we have a v -plus.

So it's certainly possible, you can imagine, depending on this vacancy region, how many electrons are sitting around it to have these different charge states. And people have detected these unpaired spins of these electrons, the fact that there's an extra unpaired spin here and there's lacking unpaired spin here, they've been detected by this technique called EPR.

So we believe that these different charge states can exist. Let's go on to page 9 now and talk a little bit more. We want to talk a little bit more about the statistics. And, again, I'm showing the energy band diagram. This time, on page 9, we're showing the energy band diagram where the Fermi level is right at mid-gap.

So the material is said to be intrinsic. And so, the only electron and hole pairs that are created are those created equally when we break bonds and the Fermi level is at mid-gap.

Let's say we're at 1,000 degrees Centigrade with a semiconductor. You look up in Chapter 1, there's a simple expression for the dependence of n_i , the intrinsic carrier concentration on temperature, and it's about 7×10^{18} to the 18.

So there's that many intrinsic free electrons and that many intrinsic free holes in the material just by virtue of the temperature. The temperature allows a lot of bond-breaking, a lot of electrons to be released from the valence band up and energized into the conduction band.

So the Fermi at this temperature, the Fermi level is-- for any doping level less than 7×10^{18} , the Fermi level is at mid-gap. Now, we have to recall something about how donors and acceptors levels are ionized and whether they're occupied by electrons or not. And, again, if you've had solid state physics, you'll recall this. Otherwise, it's discussed in Chapter 1.

But a shallow donor or acceptor is usually ionized. And what do I mean by a shallow donor? I didn't include a picture of that. I should have. Let me just check out here if this chalk box will open which, of course, is-- let's just-- what do I mean by a shallow donor?

If this is the energy bandgap and the forbidden gap, so to speak, and those are the valence and conduction bands, a shallow donor is like arsenic. Remember, we had arsenic or phosphorus. And it's shallow because it's energy level-- it's a deep level, or it's a localized state also. But it's very close to the conduction band, say within a few tens of millivolts, 0.03 electron volts.

And so, what happens is that those donors can very easily donate electrons. They're easily ionized. And they donate a free electron to the conduction band. That's how they dope the semiconductor. Each one, each donor, by putting 10 to the 18, each one donates its electron to make a free electron, and it's very easy for that process to happen.

So they're usually what we call ionized. That is, the electrons have been donated and they leave behind a net charge.

If the level is deeper, that is deeper in the bandgap, say, like v -double-minus, or v -minus, let's say, down here, this is where my v -minus level is, its energy. So it's much further away from the conduction band.

Deep donors are only ionized when f is below them. So the Fermi level would have to be below here. The Fermi level is below here, as I'm showing here. Then this is ionized and it's put an electron up here. So we have to remember these rules for deep donors. We don't usually think about them for shallow donors because they're all ionized because they're so close to the conduction band.

A deep acceptor, on the other hand, is ionized only when the Fermi level is above it. So these are rules from solid state physics and statistics. And if you know them, it helps. Otherwise, you have to take them as rules for the purpose of this course.

But if we look at-- if we're in intrinsic material and we look at all these levels, these are deep acceptors and deep donors. In this particular case, the Fermi level, which is right here, e_i , is above the vacancy donor level. So it's above these guys. So these will not be ionized.

And same thing here. The Fermi level is below these vacancy acceptor levels. It's down here. So these levels won't be ionized. So what this means is that, in intrinsic material, so it's intrinsic at that temperature, the neutral vacancy, which does not depend on the Fermi level position, that's going to be the dominant vacancy charge state. Its v_0 is what dominates when you're in intrinsic material.

As I now take this Fermi level and move it up, we'll see that these guys will dominate. You'll get more and more of these charge vacancies. And if I move it down, we'll get more and more of these. But, if we're neutral, if we're intrinsic, then it's the neutral vacancy concentration that dominates. And we'll show some equations to actually how to calculate that.

So let's move on to slide 10. And now I'm doing a different case. Notice, I've moved the Fermi level away from the mid-gap. It's no longer in e_i , it's up here. So I have doped this with arsenic or something. I've made it n-type. So I've moved the Fermi level closer to e_c . The closer I move it to e_c , the more n-type I make it, the more electrons we have.

So I'm going to assume the material is extrinsic. What does that mean? That means the n-type doping concentration is greater than n_i . That's the definition. So if I were to dope this with arsenic to, say, 10^{19} , or 10^{20} , I would be well above-- orders of magnitude above n_i . And so, it's extrinsic. The electron concentration is controlled by the donor concentration, not by the temperature.

So, in this case, I've doped it. The Fermi level is here. It's now-- the E_f is now above the E_v -minus level. So E_f is up here above the E_v -minus level. So it means this level is occupied with an electron. So it's acting like an acceptor. And, in fact, E_v -minus is going to be, in this particular case for this position of the Fermi level, it will be the dominant vacancy charge state.

So as before, when I had it mid-gap, the dominant charge state was neutral. As I move it up, all of a sudden, now I've got lots of E_v -minus. If I move it up even further, above this E_v -minus, E_v -double-minus, I'll have lots of those.

So as I move the Fermi level up and down, just imagine, with doping, I can control the concentrations of these types of vacancies, these charge vacancies, either E_v -minus or E_v -plus. If I move it down in p-type material, all of a sudden, I get lots of these, E_v -plus and E_v -double-plus.

So just by changing the doping, I can create more vacancies and more interstitials. So it's not just temperature, because I'm changing their charge-- the charge populations.

So let's take a look at page 11, or slide 11, where we actually do this quantitatively. Now, I've been talking very qualitatively. I say, we move the Fermi level up, we get more of E_v -minus and E_v -double-minus.

Well, there's actual simple mathematical equations that Shockley's paper first wrote down on these charge point defects. And, in fact, they obey exactly the same statistics as are described in Chapter 1 of your textbook for shallow donors and shallow acceptors.

In your textbook, in Chapter 1, I think there's a simple equation that was written down, just gives the electron concentration n is equal to some constant, which we happen to call n_c . And we said, it depends exponentially in Chapter 1, if you go back and read that, on the distance between the conduction band energy level and the Fermi level. That's how we-- Chapter 1, there's a simple derivation of this relationship.

So the electron concentration, as you take this shallow donor and move it around a little bit in the bandgap, you'll change the electron concentration exponentially according to the distance between E_c and E_f .

You may see this-- if you take solid state physics-- people rewrite it sometimes to say that n -- there's another way to rewrite this. And instead of referencing it to the conduction band energy, sometimes people reference it to mid-gap, $E_f - E_i$.

Either way, it's essentially the same relationship. Just says, the electron concentration goes up exponentially depending on the distance between the Fermi level and either the conduction band or the mid-gap point.

The exact same thing-- so this is the concentration of electrons. This is, instead, the concentration of these charged vacancies obeys a very, very similar relationship. So, in equilibrium, the concentration of the E_v -plus, these are single positively charged vacancies, depends on the concentration of neutrals. All of these depend-- the pre-exponential is the neutral concentration-- concentration of neutral vacancies-- times exponential, in this case, of the distance between E_v -plus and the Fermi level over kt .

Similarly, n_i depends on the neutral concentration times the exponential of $E_F - E_v$, just going back. So we're just looking at-- it's just depending on $E_F - E_v$ so it's this distance right here that tells you what the concentration is going to be.

So it's just a simple exponential type equations. You can write down very similar and analogous equations by replacing all the n 's by i 's for interstitial point defects, assuming that they have energy levels in the mid-gap or in the bandgap.

So let's just go on to Chapter-- to page-- slide 12. Here's a picture where I've actually filled in some numbers. People think they know these numbers. They're not known all that accurately. But it's believed that the double negatively charged V^{2-} level, deep level, is about 100 millivolts, so 110 millivolts, 0.1 eV, below the conduction band.

And, similarly, it's believed that this V^{-} level is close to mid-gap. It's about 0.57 electron volts or so below the conduction band. People have done studies where they feel that they can measure these numbers.

And so, as I move the Fermi level around, again, as I move this up, I'm going to change the distance between the Fermi level and these deep-- these energy level positions and therefore change the concentrations of these charge defects.

Yeah, for people who've had solid state physics here will be saying, oh, dear. Oh, wait a minute. We have a circular problem here. If the vacancies are charged, then how do we calculate the Fermi level position? So it can become a very complicated problem. It turns out, it's very simple.

In fact, the concentration-- remember, we said the concentration of these charged vacancies is always in the 10^{12} to 10^{14} range, much smaller than the concentration of electrons due to the dopant atoms. Our doping is usually always over 10^{14} . We typically dope 10^{15} to 10^{16} orders of magnitude above that.

So, in fact, the Fermi level position, what we do is we set the Fermi level just by the dopant concentration, just by the amount of arsenic or the amount of boron. That fixes Fermi level. And then we calculate where the vacancies are just based on those equations on the previous page.

So there's a lot of simplifying. And, in fact, if we go to page-- I just gave away the answer. But, on slide 13, that's exactly what we're saying.

The concentration of these neutral vacancies and interstitials is so much smaller than either n_i at typical processing temperatures or doping concentrations, that we can do a very-- the amount of electrons and holes that are actually bound to these guys is always negligible compared to the total number of electrons or holes present in the silicon that are there either due to thermal activation just due to n_i or due to the extrinsic doping.

So here's how we simplify the charge defect calculations. First, you're given the donor or the acceptor concentrations in the crystal. That is a given, typically. And we calculate the Fermi level in the usual way from solid state physics. We just ignore the fact that the vacancies and interstitials exist. So we use these simple-- we use these equations.

In fact, here's the equation I wrote up on the board earlier. n equals $n_i e^{-\frac{E_f - E_i}{kT}}$ to the exponential of the distance between the Fermi level E_f minus E_i of the intrinsic level over kT , and a similar expression for hole density. So we use these to figure out-- and charge neutrality-- to figure out the position of the Fermi level knowing the doping.

And then, once we have the Fermi level position, we can just calculate the concentrations of charge defects using that value of E_f and the equations from Shockley on the different-- on the exponential dependence of the vacancy concentrations. So it's actually not that hard.

So let's go on to slide 14. So another thing to realize-- this is kind of obvious, but sometimes we forget it-- is that as the doping changes, as dope the crystal more heavily, let's say, and, therefore, the Fermi level changes, not only do I change the charge states, but I, actually, the total number of vacancies changes. It's not just that I'm taking a certain fixed number and changing their charge distributions, how they're charged.

So thinking about it this way, this is the expression we wrote here, this equation, for the neutral vacancy population. It's only a function of temperature, not doping, or not the Fermi level. So that's very simple. It's just an exponential dependence on temperature.

But the total vacancy population is the sum of the neutral vacancies plus the sum of adding in all the charged defects, c_v -minus, c_v -double-minus, c_v -plus. If I increase all these charge defects in their concentration, then the total concentration has to go up, basically, just thinking about it that way.

So, in fact, as I thinking about it-- you never thought about this-- when I'm-- as I add arsenic to a crystal, as to silicon, I make it more n-type, I'm increasing the amount of electrons, free electrons, that are in the crystal. That's why we add dopants.

I'm also increasing the total number of vacancies, their concentration, in the crystal, because I'm making it a lot easier-- it's much more favorable to form vacancies because all those extra electrons are available. So I can form extra c_v -minus and c_v -double-minus.

So as I move the Fermi level up and down, I'm actually creating extra point defects. That's important, because, if I move the Fermi level up, I create all these extra vacancies, all of a sudden, dopants that rely on those vacancies to diffuse, they can diffuse a lot faster.

And, in fact, we'll see that in the chapter on diffusion. We'll see what so-called Fermi level affects. People move the Fermi level up and down. When it's up really high, all of a sudden, the [INAUDIBLE] of arsenic goes way up. And the explanation is related to the fact that you've just stuffed the crystal with a lot more vacancies. Yeah.

STUDENT: We've lost the original [INAUDIBLE].

JUDY HOYT: Mm-hmm.

STUDENT: And the on surface [INAUDIBLE] about. Is this those dependencies still?

JUDY HOYT: Yeah.

STUDENT: [INAUDIBLE]

JUDY HOYT: We are talking about going back to this concentration right here?

STUDENT: Or just in, I guess, the original slide-- page 4.

JUDY HOYT: OK, let's go back all the way to slide 4. This concentration here, ah, how it depends-- this refers to-- yeah, you would have to change the formation enthalpy and entropy. But what we're talking about here is the total concentration of vacancies created by whatever processes.

We're not exactly saying, in this explanation, how they were actually created. We're just going to assume that these relationships exist. If you wanted to know it in detail, you can probably break it out into different origins of vacancies and interstitial, but we don't take that into account in any of these equations, the actual mechanism by which they're formed.

So we lump all of this into a single-- when we wrote down this equation right here, we lumped it into a single activation energy, all the different processes, basically, whether it's a surface process-- again, this is an equilibrium.

Whether it's a surface process, or a bulk process, by Frenkel pair which it was created, this equation doesn't tell you. It's probably the sum of a lot of different processes. So it's an approximation.

Any other questions about vacancy interstitial populations? It's kind of weird. And if you've had solid state physics, it doesn't seem that strange. If not, you should definitely go back and read Chapter 1 to remind yourself of these statistics.

Let's go back-- jump back to slide 15. And I pulled this-- this slide 15 shows a table from Chapter 3. I pulled this right out of your textbook. And what it is, is a table of these energy level positions as best as they're known today.

And first thing you'll notice is, in the vacancy levels, there's the numbers are all printed down. The interstitial levels, there's a bunch of double question marks and single question marks. So the truth is, we really don't in the case of interstitials very well exactly where these are.

The interstitials, no one's ever pinned them down to be able to do electron paramagnetic resonance. They move around too much. So the vacancy levels, we have some ideas. Even these, I would take to a certain extent with a grain of salt. But they give me rough ideas.

And the other thing you need to know-- so you know these distances now. So you can plug all those into the equations. The question is, where are these energy levels referenced to? In fact, we're going to reference them to the band edge.

So as I increase the temperature, remember, the bandgap shrinks, I'm going to say that this distance between the band edge and the vacancy level maintains a constant, even though the bandgap is shrinking. Again, and that's an assumption. But that's-- we're going to assume the defect levels track their respective band edges.

So these v -minus levels track the conduction band, and the v -pluses-- if I go back to another slide where I have a picture-- jumping back for a moment to slide 12. These guys here, the v -minus, are going to track the conduction band.

So this distance, 0.11 electron volts and 0.57, is going to stay fixed even as I squeeze the whole bandgap due to - remember, when you increase temperature, the bandgap goes down, but these distances are going to stay constant.

And, similarly, these are going to stay constant with respect to the valence band edge. It's an assumption. How justified it is, it's not really clear. But that's an assumption that people make when they actually plug in into the Shockley equations.

So let's just do an example since we've thrown a bunch of the Shockley equations out there. And that's shown on page-- slide 16 of the handouts. And I pulled this right out of the textbook. So you can go back and reread from the text if you'd like to do that.

Here's a picture of a slide-- a piece of silicon. And this region up here in the upper right is been doped very heavily with arsenic. So it has a donor concentration of 5 times 10^{19} per cubic centimeter. And the rest of the crystal is boron or p-type-doped at very lightly, at 10^{15} .

And we take this pn junction, and we heat it up to 1,000 degrees in a furnace. And the question is, let's sit down and calculate, using the Shockley equations, the concentration of all the different charged states of vacancies and all the different charged states of interstitials. And let's just see, in those two different regions, what the concentration of these point defects is, here and here.

Well, first of all, we go to 1,000 degrees, and you look up, in Chapter 1, you notice, the bandgap has shrunk. At room temperature, the energy bandgap is distance from E_c to E_v , that's 1.1 electron volts. As we go up in temperature, the bandgap actually shrinks. And there's an expression for that in Chapter 1. In fact, it's shrunk down to 0.78 or so electron volts.

As a result, of course, it's easier to ionize an electron, or to bring-- create electron hole pairs. So the intrinsic carrier concentration is now pretty high. It's 7 times 10^{18} at this temperature.

Does anyone remember, if you've taken 6012, at room temperature, what is n_i roughly? 10^{10} , right. So it's gone up by 8 orders of magnitude, almost 9 orders of magnitude, just because I've increased temperature.

So those of you who work in electronics, at room temperature, n_i is 10^{10} , get that out of your head. n_i is a function of temperature. And it's a pretty big number at these processing temperatures.

So, in fact, in the p-type region, the semiconductor at this temperature is intrinsic because n_i is many orders of magnitude larger than the boron doping concentration. So in the p-type region, the Fermi level is at mid-gap. I swore I put it here, E_f equals E_i . It's at this mid-gap position because the material is intrinsic.

In the n region, actually, the donor concentration is still pretty high, 5×10^{19} is still much greater than n_i . So it's extrinsic. And we use the usual equation, which we just showed on the prior pages, to get the Fermi level position.

In fact, this is the equation, n equals $n_i e^{(E_f - E_i) / kT}$. I know the temperature. And I know n_i . So and I know n is just going to be the donor concentration, 5×10^{19} .

So I have everything I need to find E_f minus E_i . And, in fact, if you plug in a calculator, this E_f minus E_i distance here, from here to here, that's 0.2 electron volts or 0.21 roughly. So I can place the Fermi level on the p-type region as a mid-gap. And, here, on this diagram, this position right here, this dashed line, shows the Fermi level in the n-type region.

And then, we've just drawn in the bandgap, the band edges here, here, and here, E_c and E_v . And each line here represents-- we know these energy levels. They're positioned numerically. This is 0.11 electron volts for the V^{-} . This is 0.57.

So we can draw this quantitatively, this diagram, in the p-type region as well as in the n-type region. And, here, we're showing mostly the vacancies just for illustration purposes and, here, mostly the interstitials. But you could write down the interstitial's levels in the p-type region as well just to show the picture.

So if you sit down and then you plug all these numbers into the Shockley equation, this is what you come up with and on slide 17. And, again, this table is taken directly from your textbook. And all it is, is tabulating, in the p-type and the n-type region, the concentration of all these interesting quantities.

For example, let's just take the p-type region. The doping, we were told, is 1×10^{15} . That was a given. n_i is 7×10^{18} . That's the same in both sides because that's just due to the temperature.

The neutral vacancy concentration, which is calculated from that simple exponential we gave earlier, depends only on temperature, and it's about 5×10^{13} . And it's the same in the p and the n region. Again, the neutral vacancies doesn't depend on the Fermi level.

Here's where the differences start coming up. In the p-type region, look at the concentration of V^{-} . It's about 2×10^{14} , almost an order of magnitude larger than the neutral. So here the neutral and the V^{-} are dominating here in red.

Everything below on the chart, V^{+} , V^{\bullet} , these are all very small, small numbers. And that's because of the Fermi level position basically. And, again, the concentration goes-- depends on the position of the Fermi level respect to these vacancy levels. So here these guys dominate, the V^{-} .

In the n-type region, what's in red, I've highlighted what dominates. Actually, now, both V^{-} and V^{+} in the heavily doped n-type region dominate. And not only do they dominate, look, they're so much bigger.

If you were to add up all these numbers, look at the total vacancy population now is well above 1×10^{15} , much larger than it is in intrinsic material where it was here in the 10^{10} to the 10^{14} range.

So, again, that just shows, in the n-type region of the crystal, there's a lot more vacancies total. So, in that region, if you have a dopant that's diffusivity depends on the vacancy population, it's going to go up its diffusivity.

Clearly, the different charge states dominate in the different regions. And I'm not going to go-- we won't introduce right yet, but when we get to the Chapter 7 on diffusion and Chapter 8 on ion implantation and even in the oxidation models, these numbers, these concentrations of point defects will be very important.

For now, you just get an intuitive feel. As I go n-type, oh, OK, the v_{-} and the v_{--} go way up. As I go p-type, these-- well, this is intrinsic. It's not actually p-type. In p-type, if I made it heavily p-type, these numbers would go up, the v_{+} and the v_{++} concentrations.

So I'm going to leave the vacancy. When we switch to slide 18, I'm going to leave the vacancy and interstitial. You'll get a chance to look at some of those in your homework. But that's an example from the text. I'm going to leave that for now. And we'll, as we go along each process we talk about, we're going to come back to the importance of vacancies and interstitials.

There are two other things about the crystal that are very important that I wanted to talk about and discuss that's discussed in Chapter 3, and that is oxygen and carbon. We said yesterday, or on Tuesday, that the Czochralski growth process inherently introduces oxygen from the quartz crucible and a certain amount of carbon. And we can't get around that.

And these are typical numbers, oxygen about 10 to the 18, carbon concentration about 10 to the 16. And they both have some important effects. For example, for oxygen, these are three effects that are listed.

This is a simple-- or it's an equation you learn from 6012, or if you've taken an electrical engineering class. If you haven't, just take the equation as given. It's equation for v_{th} , the threshold voltage, as a function of various physical parameters.

The v_{th} , that is very important to the circuit designer. It tells you, and to the device person as well, it tells when you apply that gate voltage, the inversion layer forms and you start getting conduction. So circuits behave very badly if your threshold voltage varies from device to device all over the place.

There are-- all these depend on things you can control, like the doping, and the wafer, and things like that. But this one term, this last term, $q q_m / C_x$, where q sub m is the concentration of mobile ions, be it, say, sodium, potassium, calcium, whatever, that exists in the oxide are right at the interface between the oxide and the silicon.

And this, we write as the number of charges per square centimeter. So it could be the sodium concentration in sodium atoms per square centimeter. Well, you just plug some numbers in here, put in an oxide thickness of 10 nanometers, which is reasonable for these days. Then I calculate a v_{th} change or a shift of 0.1 volts by having q_m be just about 6 times 10 to the 11, not very much. That's only 10 ppm, 10 parts per million.

But, already, that amount, 6 times 10 to the 11 sodium atoms, can shift the threshold voltage in one device to the next if it varies by that much, by a tenth of a volt. And that's a big deal to either a device person or a circuit person.

So you really need to-- that's why you'll see the concentrations of these things in the ITRS have to be 10 to the 10 or below because they want it to be a small. They don't want any chance of a 100-millivolt threshold voltage shift. That's just way too much.

So that's one reason for people making logic devices. People in memory devices are even more paranoid, even more scared of contamination. And here's an example, number two on slide 25. This is schematically a dynamic random access memory.

So this is a DRAM that all of us have in all of our computers. And what it consists of is a transistor and that access this-- basically this charge stored on the analyst capacitor. If charge is stored on there, then it's a certain logic state, a one, if there's no charge, it's a zero.

Well, it turns out we have to refresh this periodically. We have a certain amount of charge stored in that node, but it leaks out over time. And, typically, the DRAM has to refresh that bit of information every few milliseconds. But if it's leaking out really fast, it has to refresh it more often. And you suck down your battery power and then the chip doesn't work very well.

So, it turns out, a refresh time of several milliseconds requires a lifetime. This is a measure of the purity of the silicon, a generation lifetime, that goes to τ that goes 1 over a number σ , which is given by this number, times the thermal velocity, which is also a constant, times n_t . n_t here is the concentration in per cubic centimeter of traps.

So if I know σ , typically the cross-section for a typical deep level is this number, 10 to the minus 15 , I know the thermal velocity, I can solve for n_t . This requires that I have less than 10 to the 12 per cubic centimeter or 0.02 part per billion of these traps in the semiconductor.

So it's extremely important. If you have more than that, your charge is going to leak out faster and your DRAM is not going to work. So it's very important to keep these impurity concentrations low for that reason.

So what are these traps associated with? Well, it turns out that trap density scales with the density of certain bad actors. And we'll talk next time about who these bad actors are in silicon impurities, such as heavy metals, gold, transition metals like copper, iron, nickel. A lot of these turn out to be traps, deep levels, that destroy the lifetime.

But, anyway, just to give you a rule of thumb, you need to be in the ppb range or lower. In fact, it turns out the requirements are a little more severe than that simple equation I just did because these elements also have a very bad tendency to accumulate in heavily doped regions, like in the pn junction or source drain, they tend to be-- we'll talk about why they have-- they like to be in those regions. So that's the worst place for them.

So that's why we even require lower carbon concentrations-- or iron. Back here on page 25, I said 10 to the 12 per cubic centimeter. In fact, people, if you look at the NTRS or ITRS, it'll ask for much less than that.

So the truth is, it's impossible to keep our wafers that clean. No matter how careful we are in the lab or in the process fab. So what we do, what people do in manufacturing anyway, is to use a process like gettering, so that it will try to remove these unwanted iron, gold, and copper from the regions of the wafers where the actual devices are located.

So this is a special type of process we do, usually done at the beginning of the process. We do something to the wafer-- or maybe in the middle of the process-- to try to attract these to another location.

So we'll talk about those three types, three levels. First is just the clean room, how do we keep the clean room clean? Second is how we clean the wafers? And third is how we do gettering? So, for today, I'm just going to do this level one control.

And hopefully some of you have seen clean rooms. But, basically, what we do in a clean room is we try to keep the air free of particles. And we do this in practice by these tremendous filters, high efficiency filters, that really trap all particles. So the air coming into the clean room comes from the top.

These high-efficiency filters remove all of the particles, and the air goes out the sides. And it keeps washing the air continuously. We have this constant flow of air taking particles and removing them down to the floor so that they don't get on your wafers.

And, in fact, we actually quote the class of a clean room, where the room gets cleaner as the class goes down. So this is a plot of the total particles per cubic foot in the air of a clean room as a function of the size of the particle.

And the line, each line is parametrized according to a number. That's the class clean room. So, for example, here in MTL, the clean room is nominally rated, at least the integrated circuits lab, at about class 100. So that's this line right here.

So a class 100 says, here we are at a micron-sized particle, you can read off roughly the number of particles per cubic foot in the range of 50 or something like that. Gives you an idea of the number of particles per cubic foot.

If you have more than that, then you move to a class of 1,000 or 10,000. So that's a dirtier room. The higher the class, the lower the number, like a class one, that's the best-- that's the cleanest type of clean room.

Here's an example I just-- here on slide 28-- of just some pictures of different clean rooms. And, basically, we clean the factory environment-- I mentioned these HEPA filters, these high-efficiency particle filters, that recirculate the air constantly. So it's being filtered all the time.

We put people in little suits called bunny suits. Here's an example of a student in a university R&D lab wearing a bunny suit. So you try to cover yourself.

People are the worst source-- the biggest source of particles. We shed skin constantly. We can't help it. It's part of our natural thing. Even if you don't have dandruff, it doesn't matter. You're shedding constantly small particles. So we try to keep as much as possible our clothes contained in these clean suits.

This is in a university-style R&D lab. In a manufacturing, here's an industrial fab. This is a picture I pulled off the Intel website a while ago. Actually, we sometimes put the workers in spacesuits so that not even their face is exposed so they can't spit particles, so they can't breathe particles onto the wafers. So here's a worker in a spacesuit. And he's breathing through a tube.

We, also, besides the air, we have to use ultra-high purity water. We have to filter the water so there's no impurities in the water, no ionized, no sodium, no potassium, no particles. And all the chemicals we use are produced in special chemical factories where they reduce the amount of sodium, potassium, iron, and copper, all that from the chemical.

So everything is ultra-high purity, all the gases are purified. So everything that goes onto that wafer surface is purified. And we use a lot of protocols in manufacturing.

Oh, by the way, these days, in 300-millimeter fabs-- so this is a 200-millimeter-- 300-millimeter, which is the next generation, fab they're going to actually even reduce the number of people walking around, even in spacesuits.

And what they do is they use robots. Here's a robotic. I took this off of a commercial website. This is a track near the ceiling of a 300-millimeter fab. And these little pods, these little square or cubic pods, hold wafers in them. And these travel. These pods travel on the track. They automatically go down to a piece of equipment, open up, and the wafers get loaded.

So the wafers actually never get exposed to the clean room environment. They're kept in pods the entire time. They go from a pod into a vacuum system or pod into another tool. And so, they minimize the exposure this way. So we even have now a lot of robotics just to clean up the process even more.

Slide 29 I talk about-- we'll continue this in more detail next time. But I talk about wafer cleaning. We do introduce onto the wafer a certain amount of impurities ourselves.

In the photolithography process, remember, is that organic material, photoresist, it turns out, is a contaminant. But we have to use it because it happens to do-- it does photolithography. But we have to remove it every step because we certainly don't want it on the wafer when we're doing any high-temperature processing.

So after we do resist strip, we typically then have to clean the wafer using a special clean that we'll talk about next time in RCA clean. And then we do the process step.

So the process of semiconductor fabrication consists of pattern something, transfer the pattern, get all the resist off, do a clean, and then maybe go into a high-temperature step, and repeat the whole thing again. So you can have tens, five, 10, who knows how many cleans that go on depending on the number of mask levels and the number of high-temperature steps.

So why are we trying to remove these things? And, I already indicated, and what we're trying to remove these organic films or trace organics? All of the organics are bad because they can cause leakage in the gate. And they can act like a mask. They prevent you from cleaning the underlying surface. So you need to get organics off. Typically, a first step of any cleaning is try to remove organics.

Particles, well we already-- that's fairly obvious. They can cause defects. Particles very often carry metals in them because you get particles from the wafer handling equipment. Most equipment's made of stainless steel, which is what? It's iron and nickel. So it's a very-- particles are bad because they tend to carry metals into the wafer.

Alkali metals, I've mentioned, sodium, potassium, calcium, these move easily in oxides. They can shift the threshold voltage, like we saw in our calculation, and they can cause reliability problems.

And we already mentioned the transition metals having bad properties. They reduce the carrier lifetime. They reduce the mobility, which is an important property. They're also bad because they can diffuse very easily into the wafer. And they can roughen the surface. So these have a very, very bad-- among CMOS people gold, copper, and nickel, iron, very bad reputation.

So, these, we take major organic contaminants, like photoresist, and we typically remove them in a solution of sulfuric acid or oxygen plasma prior to do any real refined cleaning, like RCA clean. And then we remove trace organics for the RCA clean along with metals. So that's the philosophy. And we'll talk about that next time.

This is just a picture of an RCA cleaning bench. And the next time we'll go through how it works. It's a standard process. And, again, it removes trace organics, not heavy organics, trace heavy metals, and alkali ions. And there's very specific chemicals that are used, which we'll discuss next time too-- on how that's done.

So let me just summarize what we talked about-- a wide range of things. We talked about native point defects, these vacancies and interstitials. They can be either neutral or charged. The charge defects obey the same statistics as shallow donors and acceptors.

So as we move the Fermi level and we dope the crystal n-type or p-type, I can change the number of the concentration of these defects. The good thing, though, is they're still low enough in number that I don't need to use them when I'm calculating the Fermi level position. So that makes it easier.

There are both v-type vacancy type and i-type defects. And we'll talk about how they play a very important role in understanding the processes. And I introduced the idea of contamination control. That's in Chapter 4. And we'll discuss that in greater depth next time.

OK, so that's all I have for today. And, on Tuesday, remember that your first problem set, your first homework set is due. Thank you.