

MITOCW | MIT6_774F04_lec13.mp4

JUDY HOYT: It's very simple, but doesn't fully capture all of the reality. The simplest analytic expression that captures profiles reasonably well is the Pearson four, which has four moments, and those moments are generally tabulated in tables. We also presented the idea of Monte Carlo or numerical solutions, for example, Monte Carlo simulation, which is quite accurate. And another thing called the Boltzmann transport equation solution.

The other thing we mentioned is that there's this process called ion channeling. Ion channeling itself is quite a challenge in calculating profiles. Calculating the profiles pretty accurately into amorphous solids, or amorphous materials, is reasonably routine. But once you have to model ion channeling, it gets a little tricky.

Today I want to cover some things that we didn't get a chance to talk about. We never even talked about, last time, the physics of the modeling of the ion implantation process, that is the physical energy loss mechanisms of the ions as they traverse. We want to talk about that. Once we've covered that, we want to talk about damage. How do we model the damage that takes place in the silicon substrate, and how do we anneal it? And then I'll give a very brief introduction to transient enhanced diffusion. The next lecture, next Tuesday, is going to be dedicated to talking about TED, but I'll at least briefly introduce the topic. OK. Let's go on to slide number 2.

The right hand side of this slide is just a schematic picture of a cartoon where you're shooting an ion into a silicon surface. And all these little open symbols here represent silicon atoms on some kind of a lattice. You see an ion, which is this dark bullet here. It comes in and it hits the target. And when it does, it may have some kind of a nuclear collision here. It may collide with the target atom and be deflected. There'll be some energy loss there.

Then the target atom itself is going to be recoiled and will have some energy imparted to it. But that deflection is going to then knock the ion a little bit into a different direction, change its direction. And interestingly, you see it going down a path. And right here, after this collision right here, interestingly, the ion got knocked into a channel, is what we're supposed to represent. You see, its direction got changed by this nuclear collision and now it's working its way down a little bit of a channel.

Still losing energy, perhaps at a slightly different rate. And still with each collision creating a silicon recoil that's been knocked off its lattice. So just a very schematic picture. So again, we're bombarding the wafer. The energy of this incoming dark-colored ion is typically anywhere, well, some of the lowest energy implants being done today are one kilovolt, or maybe even slightly less. That's very low, but that is being done. And some of the highest energy implants for making, in bipolar they make selectively implanted collector regions that are quite deep. They can work in the 1 MeV range.

So there's a very wide range that people are eyeing at playing. Typically, implants are between something like 10 to 100 keV, but people do use the full range. Now look at the binding energy of a silicon atom on the lattice. That's only 15 electron volts. So you can imagine that we're coming in with a species that has thousands of times more energy than the binding energy. So obviously, it's quite possible to have billiard ball-like collisions, where you knock these silicon atoms off of their lattice sites.

So the ions collide, and they collide elastically. And they have, as a result, we have ion deflections, which I just showed you here, where the ions change their direction. And when you collide, of course, you lose a certain amount of energy. The ion loses some energy, and it displaces silicon atoms, which we end up calling recoils. In addition to this collision, this sort of pool table type of collision of billiard balls, the ions can also suffer sort of an inelastic drag force from the target electrons. And so this leads to electronic stopping, where you lose ion energy and it actually heats the lattice, to a certain extent. Because they're in this median that has electrons in it.

Eventually, this ion and all the ions come to rest after they've lost all of their energy in both collisional processes and in this drag force. And channeling, as I mentioned, we talked about last time. Along certain directions, ions can travel in the crystal with very few collisions and little drag. So they can go deeper than they would otherwise, and that's tough to model. So if we go to page 3, or slide 3, how do we model these range statistics? Well, what we do is we write down the total energy loss during an ion trajectory, and we write it as a linear sum.

So we treat these two processes independently, nuclear losses and electronic losses. So we write the rate of energy loss as a function of distance, as this equation, dE/dx . It's a negative number, because you're losing energy. N , where n is the target atom density, atoms per cubic centimeter, times the sum of these two quantities $S_{sub n}$ plus $S_{sub e}$. Again, n is the target atom density. $S_{sub n}$ is a function of energy. It's the nuclear stopping power. And it has units of energy times area. So it's eV dash centimeter squared is a typical unit. And $S_{sub e}$ is the electronic stopping power, again eV centimeter squared. you can see the units work out.

If S has units of eV centimeter squared, you multiply it by atoms per cubic centimeter, you get a certain amount of eV per centimeter, or energy loss per unit length in going into the crystal. The thing to remember is that these functions, these are in general these stopping powers are going to be a function of the energy. That is, the rate at which you lose energy is a function of how fast you're going, of how much energy you actually have.

If we know $S_{sub n}$ or $S_{sub e}$ as a function, and we know $S_{sub e}$, the electronic stopping as a function of energy, then you can simply compute the range by doing this integral. So you integrate. The range would just be the integral from 0 to E_0 . Whatever your incoming energy is, your ion-implanted energy, say 100 kilovolts of dE divided by the sum of the two stopping powers. And here again, the density is a constant number. It's just been pulled out. So you can find the range mathematically if you know the physics of these two stopping powers.

So that's where a lot of the modeling has been done in understanding the physics of $S_{sub n}$ $S_{sub e}$. And we'll first talk about, on slide 4, we'll talk about nuclear stopping. Above a certain energy, so it's about half a kilovolt, or 500 eV. Now be careful, because we're approaching this range for some of our very lowest energy implants, making very shallow source drain extensions. People are using quite low energies. But if you're above this, it's reasonably valid to model nuclear stopping as a classical two-body collision between a silicon atom that's sitting still and an incident ion that's coming in with some velocity.

So it can be modeled as, just like on a pool table, two balls colliding, using the conservation of energy and the conservation of momentum principles that you learned in your basic physics class. Now, why does this have to be above a certain energy? Why is that the case? Well, it turns out, if the energy is lower than this, or significantly lower, then what happens is, as this incident ion comes in, instead of looking like a billiard ball, a hard sphere model, this ion has time to sort of hang around the target atom, and the interaction is not just like a Hartsfeld collision. You have other types of interactions. You can have multiple interacts with a lattice, maybe, in the sense of it's going slow enough that as it passes by the nucleus, it has different types of interaction. So if your energy is low enough, these models break down. But the good thing about ion implantation, most of the time you're well above that threshold.

So, as you remember from your basic physics class, we have an incident ball coming in, hitting another ball that was originally at rest. This incident ion is going to be scattered at some angle. And you can figure that out once the interaction potential, V as a function of r , r being the distance between the two bodies. Once that's known, then you integrate that along the path of the ion. You can calculate the scattering angle for the collision, and you can do it once and then look in a lookup table.

So you apply basic physics of conservation of energy and momentum, you can get these numbers. Again, the interesting thing about ion implantation, of all the processes we talk about in this class or anyone does, in IC Fab, it's the only one where, from first principles physics, you can actually predict something. You don't have too many free variables the way you do in some of these chemical processes.

So what it boils down to as far as nuclear scattering goes, is figuring out what is the appropriate interaction potential, the nuclear scattering potential? Well, if you think about these two bodies, an ion and a silicon atom with the ion approaching it, the first thing you might think about is, because you know the nucleus is charged, and it has a certain charge, which is the charge on the electron q times it's z number.

OK. So you might imagine that the simplest type of potential interaction between these two charged nuclei is just the coulombic potential, which is just given by the z numbers multiplied together divided by r . So it's a $1/r$ type of good, old-fashioned Coulomb potential. That's a little bit too simple, though, because the atom is just not a nucleus. An atom has an electron cloud around it. So the electrons around the target atom nucleus actually ends up screening the core potential of the nucleus, to a certain extent, from the incoming ion. So it's not, this is a little bit oversimplified.

So what people do is they put in a generalized function, f , which they multiply by the classical Coulomb unscreened potential, and then this becomes the screened potential, v as a function of r . F is some generalized function of r over a . Again, r is the distance between the two nuclei, and a is a thomas-fermi parameter, which is related to the Bohr radius. It's some function of the Bohr radius times a function of the z numbers. So it gives you, a is some measure of the, if you want to call it the size of the atom, in some way. And r is, of course, the distance between the two, the colliding ion and the silicon atom.

So a very common potential that's often used is an exponential function. You assume that the electrons exponentially dampen out the Coulomb potential, and the people use what's called a thomas-fermi potential. So this function here, shown on the third equation on slide 5 is the thomas-fermi potential for ion implantation. So it's got a $1/r$ dependence times an exponential e to the minus r over a , where again, a is this thomas-fermi screening parameter. So that's potential. That's one that was often used in modeling of ion implantation.

So it turns out that when the screening function is, instead of using exponential, if you use a screening function f just to be a/r , then it turns out that if f is just a/r , and this just the potential just goes like $1/r^2$. That's another way of doing the screening function. In that case, the nuclear stopping power can be approximated by a constant. So it turns out that this number, and if you want to see that derivation, you can go to this article by Gibbons back from the proceedings of the IEEE, oh, back in 1968. And in fact, they derive that it can be a constant number. That depends just on the z number of the ion, and the mass of the ion, and the z number and the mass of the substrate.

So this is a simple equation you can use if you're stuck on a desert island. You need to know nuclear stopping power, you can just program this into your calculator. And I just did a simple calculation for phosphorus, which has a z number of 15 and an atomic mass of twice that, or roughly a 31. Silicon has a z number of 14, a mass of 28 amu. You plug these numbers in and you do the calculation, I get an S_n number of about 550 kiloelectron volts per micron. So that's how much energy I lose, if phosphorus going in, assuming it's a constant linear number, 550 keV per micron, just gives you a rough idea. And that's again, assuming a screening function is just a $1/r$ type of screening, due to the electrons around the nucleus.

And in fact, if you go on to slide 7 of your handout, I took this plot from Mayer and Lao's book. I referred to Mayer and Lao earlier in the course. In fact, handout one has the full reference for his textbook. And what he's plotting here is dE/dx . So again, that's a keV per micron. So it's a stopping power. Or, if you want, you can change both the energy unit and the length unit. It's equivalent to eV per nanometer. I prefer to think of keV per micron, but it's just either way. And it's on a log log plot, and the x-axis is the energy of the ion.

And there are two different types of stopping powers shown here. Right now, all we've talked about is nuclear, but let's focus on this dashed line here for phosphorus. That starts up at 10 keV and it's about 500 or so, a little over 500 keV per micron. And then, as you can see, in reality, the nuclear stopping power given by this dashed line is actually going down as you go to higher energies. But lo and behold, the number we just calculated in our simple calculator, again, assuming a thomas-fermi screening function of $1/r$, comes out at low energies to give you pretty close, asymptotically, to what people calculate with more sophisticated screening functions.

Indeed, my calculation here of 550 keV per micron at 10 kilovolts is pretty darn close to that dashed line for phosphorus. So it's interesting. And you can check it out yourself for arsenic. Again, you know the Z number and the mass, you can check it out to see how accurate, at low energies, it asymptotically approaches. So what do we notice? Well, obvious, arsenic is a very high Z number, a high mass element, and it's nuclear stopping power is quite large. It's on the order of 1,000 keV per micron, something like that. Phosphorus is being less, and boron nuclear stopping even lower, an order of magnitude lower than for arsenic.

The other thing on this plot, besides these nuclear stopping powers, the ones S_n are nuclear, these straight lines that increase with energy, these are the electronic stopping powers, and we'll talk about that next. This is the type of stopping that's due to the drag force of the electrons in the substrate. So let's go on to slide 8 and talk exactly about that, this sort of non-local. Why is it non-local? Well, the nuclear stopping power is local in the sense of, you're having a collision. The ion is coming very close to the atom and the substrate. You're having a deflection. So it's an actual physical collision, and they interact by this coulombic force.

This is a non-local phenomenon called non-local or local electronic stopping, and there are two ways of viewing it. You can imagine some kind of drag force caused by the fact that I have a charged ion coming in, going at some velocity. So it's as plus, arsenic plus, or boron plus. Boron minus is in this sea of electrons, the sea of electrons where they come from. Well, every atom in the crystal has electrons associated with it. And in fact, there are covalent bonds.

There are four covalent bonds for every silicon atom, with eight electrons participating per atom. And these covalent electrons, essentially, can produce a retarding effect on the ion going through it. So you can think of it as sort of a retarding field. Imagine the silicon or the substrate that's some kind of dielectric medium. You have an ion coming through it at some velocity. This blue circle was meant to represent, say, the boron or the arsenic ion. And there's this retarding electric field.

So the interesting thing, the thing you need to note, the main point is, this is a dissipative loss mechanism, but it doesn't change the direction. The electrons are way too light. They're very, very light. They can retard or slow down the ion, but they're not going to change its direction. There's no deflection as a result of this. So, unlike nuclear scattering. So that's one way of thinking of it as a drag force in a dielectric medium. The other way is you can imagine a quote unquote collision. Now, just be careful in the word collision, the use of it, with electrons around the atoms transferring some momentum to those electrons from the electrons from the ion itself.

So you can imagine, here's this target atom. It's got this electron cloud around it. And it transfers momentum and actually results in locally slowing down this ion, basically, reducing its velocity. Again, there's no change in this mechanism, no change in the direction of the incoming ion when this process happens. So both of these don't change direction, and both mechanisms are related to, or involve, the speed or the velocity of the ion.

OK. So if we go to the next slide on slide 9, to first order, people have found they can write the electronic stopping power as some constant times the velocity. So this drag mechanism, the keV loss per micron, due to electronic stopping, increases directly proportionally to the velocity of the electron. And what is the velocity? Well, the velocity just goes like the square root of energy, right? So you can write electronic stopping as some constant, k times the square root of e . And in fact, k has this, very roughly, you can approximate k by about 2×10^{-14} square root eV centimeter squared, as shown up on the top of that slide. So this is an approximation, but to first order, it seems to work.

I should note here, I'm noting in the upper right corner, that a lot of the improvements in ion implant modeling over the last 5, 10 years, or however long, have actually come from a better, more accurate treatment of electronic stopping. Nuclear stopping is very much nuclear physics, and that's been known for a long time now. But this electronic drag force is a little bit more mysterious, a little bit more difficult to model.

And this is where a lot of the improvements have come in the literature. And particularly, this affects light ions like boron, for which most of this stopping is actually electronic. Boron isn't very heavy. It doesn't experience that much nuclear stopping. Most of what slows down a boron ion implanted into silicon is electronic stopping. So boron profiles have become more accurate, more accurately modeled over the last x number of years, because people have a little better models for electronic stopping.

So if we look at this plot on slide 9, this is the total stopping power. It plots both the nuclear and electronic. So again, the units are keV per micron in the vertical axis, and keV on the y-axis. And these are the same plots I just showed you, they're just different color. Here's arsenic nuclear stopping, phosphorus nuclear, boron nuclear stopping, and here is the electronic stopping. And again, it's just basically proportional to the velocity. There's only a small dependence, and we don't even show it here, on the type of ion. So we're just writing this as this black line.

So interestingly, what you see, an interesting energy to point out is the energy at which these lines cross. And that's called the critical energy, $e_{sub c}$. And at that point, the nuclear electronic stopping are equal. Beyond that, the electronic stopping has actually taken over. It's much larger. Because again, the nuclear stopping power is going down as you go faster, or increases as you go slower. But the electronic is going the opposite direction. So boron at 17 keV, above that or below that energy, basically, nuclear stopping will be important. Above that energy, which is a lot of our implants, boron is pretty much being stopped by electronic stopping.

Phosphorus is about 150. So any energy less than 150, you're going to be dominated by nuclear stopping. Above that will be electronic. And arsenic is almost always dominated, at least in the beginning of its path, by nuclear stopping. The interesting thing about this, though, think about an individual ion. As an ion comes in, of course, it comes in with a lot of energy, 10 keV, 100. But by the time it stops, it's got 0. So every ion has to traverse. As I first come in, let's say I started coming in at 10 keV, or rather, 100 keV, and I'm right here.

So as I'm coming in here for boron, you'd say at 100 keV, well, electronic stopping is dominant. That's true. But as the boron ion slows down, it walks down this curve and walks up this point. And when it gets to very low energy just before it stops, in fact, nuclear stopping always takes over. Because at low energies, you really can have a lot of billiard-ball-like collisions. So that's why at the end of the range, back at the depths of the implant that's near the end of range, there's a lot of nuclear stopping that goes on. And we'll talk about what impact that has for damage profiles.

And here's just an example on page 10 on different type of plot. That was a log log plot. I've plotted these nuclear and electronic now on linear axes, just so you get a feel for what they look like on linear. So this is dE/dx . And this is, well, it's actually a square root of energy. That helps linearize the, this straight line, of course, then, is the electronic stopping. And this other line here, that it peaks at some energy, E_1 , and then starts to decrease down, that's the nuclear stopping. So on a linear scale, you can actually see what it looks like. And again, I took this from Mayer and Lao's textbook.

In general, the nuclear stopping dominates, as we said, at low energy towards the end of the range. And that's the location in the substrate where the nuclear collisions are going to produce most of the damage. So we call that end of range damage. At very high energies up here, where I'm up here, particles travel very quickly. They have less time to interact with the nucleus. So nuclear stopping is not as important. They have less interaction time, and so they tend to be dragged down. Nuclear stopping is going down with energy, where electronic is increasing.

OK, so that gives you an idea of some of the physics of the loss mechanisms. So let's go on to slide 11. So let's say I have these nuclear stopping powers as a function of energy, and I have the electronic. How do I get from that to a calculated profile? Well, we know we just simply need to do this integral if you want to compute the range. Or you can, directly from Monte Carlo simulations, you can actually simulate that billiard ball and those electronic stopping processes.

Many years ago, three folks, Lindhard, Scharff, and Schiott, actually took these equations, put in nuclear and electronic stopping powers, and computed the moments of distributions, given $s_{sub n}$ and $s_{sub e}$. And in fact, that so-called LSS Theory generated these tables, I think we showed these tables last time for common dopants in silicon, generated tables of the first few moments. Maybe the first three moments, that generated r_p , this Δr_p , or the standard deviation σ , and the skewness γ . They are tabulated and they were calculated originally by LSS, or they've been fit to experimental data. But you can calculate these moments from first principles physics estimates of $s_{sub n}$ as a function of energy and $s_{sub e}$ as a function of energy.

OK. So that's basically how people do the calculations. What it boils down to, the calculations boil down to how accurately do I know the nuclear stopping power as a function of energy in electronic stopping? If I know that, I can use a computer, you can regenerate the old LSS statistics, or you can actually use a computer to do Monte Carlo simulation. Because once you know those stopping powers.

Given that, OK, so we can figure out where the ions end up in the lattice. Now we want to know about, what kind of damage does this incoming ion coming in at a certain energy do to the lattice? So imagine I have 30 kiloelectronvolts arsenic which is coming into the silicon lattice. And if you look that up in your table in your textbook, you'll find it has a range of 25 nanometers, or about 250 angstroms. So in 250 angstroms, actually, you can figure that out, that is equivalent to roughly 100 atomic planes. It is about 0.25 or 2 and 1/2 angstroms per plane, interplanar spacing.

So you can imagine this arsenic ion is coming in. It goes through about 100 atomic planes. And how many, so in fact, you can think of this 30 keV arsenic ion is coming in like this. This squiggly line is meant to represent its path. Here's where it ends up. And you can think of a cylinder around that, sort of a cylinder that represents the region, which is damaged, because it does a lot of nuclear stopping and you have a lot of recoils generated until it comes to stop.

And in fact, there's a simplified formula, the k_p formula, that people use to figure out, roughly, the number of displaced particles. So that'd be the number of displaced silicon atoms created by an incoming ion. And basically, in simple, you can imagine it might be related to the energy of that ion, which is 300 kiloelectron volts, so that's 30,000 electron volts, and divided by the displacement energy. Well, we use a factor of one half in front of that. But the displacement energy, we said, was 15 electron volts.

So, just order of magnitude estimate, by comparing the incoming energy to the energy it takes to displace an atom from the lattice, we're talking about maybe 1,000 recoils of silicon are created by this one arsenic atom. That's quite a bit. So you're essentially doing little ion implants inside the substrate every time an arsenic atom deflects silicon. It creates what they call silicon knock-ons. The silicons are knocked on, and they themselves have a certain energy, and they can do more damage. So it creates this sort of cylinder-like region of damage, just due to one ion.

And in fact, that was a very schematic cartoon that we just drew with the computer. This is a little more sophisticated on slide 13. These are actually what's called molecular dynamics simulations, where people build in the computer. And of course, this is done at Lawrence Livermore National Lab, would have huge supercomputing capability. They build up a model of the silicon lattice, and they actually model the physics of an incoming atom or ion. Here's a 5 keV boron ion coming in. This ion implanted into the substrate. And they actually follow a whole bunch of ions and actually look at what happens to these displaced recoils.

And each little snapshot picture here, I apologize, it didn't come out very well when we copied it, is a snapshot in time. So the first little cube is supposed to represent a region of the silicon where, after 0.1 picoseconds. So this ion has come in, and just at the time of 0.1 picoseconds, there's this little cloud here. You can see this little cloud. That cloud represents all the displaced silicon atoms. So they're actually following all these atoms displaced from their lattice sites. And then at five times that time, so at half a picosecond, look at how that cloud has grown. It's gone in deeper, and it's expanded sort of laterally.

And then, interestingly, at 10 times that again, roughly, at 6 picoseconds, the cloud looks a little bit smaller. It's reached some size, not that much different from its first size, but why? Well, what's happened is that some of these silicon atoms that were knocked off have found that vacant lattice sites. So you're creating ions of silicon atoms that are displaced, but they find holes in the lattice where they can sit. And so this becomes sort of what the off-site silicon atoms look like after about 6 picoseconds. So people actually do try to simulate this from first principles physics as well.

So that's a more modern kind of calculation. On slide 14, I'm showing you, maybe, a little bit more traditional or old fashioned calculation, what people used to do in the 1980s. In fact, this was taken from a book called The Handbook on Semiconductors on Ion Implantation, published back in 1980. But what it is, rather than that following every atom and figuring out what people do, is they found a way to calculate the energy deposited into nuclear processes. Again, know the nuclear stopping power as a function of energy.

So as we integrate those equations in depth, we can figure out, at any given depth, we know the energy of the ion, we can integrate that up and figure out how much nuclear stopping is taking place at that depth. And therefore, you can think of how much of the energy loss in keV per micron or eV per angstrom is there at a given depth. And here's a picture of that calculation or a plot, its so-called damage density. So it's the amount of energy that is lost at a given depth due to nuclear processes, and it has the units of eV per angstrom as a function of depth.

And this is for a particular ion implant. This is boron 11 being implanted into silicon at 100 kiloelectron volts. And you can see the damage energy. The damage energy density has a certain profile. And this solid line represents a calculation where it includes the silicon recoils, that is the silicon ions that are generated. And the dashed line is when we have only the primary ion. So there's not much difference in this case. But here's an example of boron silicon. It's a relatively light ion. And the interesting thing that it's plotted here, is as a function of x over r_p . So remember, r_p is going to be roughly close to where the peak of the boron. So the boron would peak here at 1.

So the range of the boron atom is actually much greater than the range of the average silicon recoil. That's because the boron is light, so it can't push the silicon that much deeper in. So actually the damage density for silicon recoils doesn't contribute that much. So you can see that these two, dashed and the solid, look pretty much the same. But interestingly, where does the damage peak? Pretty close to r_p . In fact, it's just at maybe 80% of r_p .

And that's a good rule of thumb you can use if you want to know, all right, I'm going to ion implant something into a silicon substrate. Maybe I don't know how to calculate the damage energy density but I can know how to calculate the profile of that ion. Where is the maximum damage? Well, it's close to r_p , but in fact, it's just shy of r_p . So if you want to do a lot of damage, or usually you don't, or you want to minimize, you can figure out, just below r_p is where you're going to peak in terms of doing the most damage to the substrate.

Here's an example on slide 15, a different situation. Now this is the heavy ion. The same kind of calculations published in Gibson's book, Handbook on Semiconductors Ion Implantation Volume. Again, damage energy density in eV per angstrom, again, versus x over r_p . But look at this. This time, the solid line, now again, includes recoils and you get this sort of distribution. The dashed line is for the damage done just by the antimony itself. And you get a very different damage distribution. So what this says is for a heavy ion like antimony, antimony is big enough that it can impart a lot of energy to silicon ions, or the silicon recoils. And the silicon recoils transport a lot of that damage energy deeper into the substrate.

So if you don't take it when you calculate damage, and you have a heavy ion coming in, if you don't take into account the silicon recoils, you're going to get pretty somewhat inaccurate distribution of where most of the damage is done. So you really need to use the solid line and include the silicon recoils in your damage, because the silicon atoms then go on and damage further the lattice. They impart energy to other silicon atoms. This is for 100 kilovolt antimony. But again, look where the peak is. The peak and the damage still occurs at about close to 80% or so of the projected range of the antimony ion. So that's a handy rule of thumb.

Now, what can we do with a damage energy profile like this, or damage density versus x , or x over r_p ? Well, if you know how much damage in a given angstrom, how much energy in a given angstrom will tend to produce amorphization, you could actually use this plot to figure out which part of the lattice is which depth, or amorphized. And so here's an example. What I'm doing here, I'm taking this calculation, which we showed for boron a couple of slides ago, and I'm assuming that there's a threshold for amorphization.

Now this has been both calculated and people have tried to measure it. I'm assuming the threshold for amorphization at room temperature is 6 electron volts per angstrom. So people have kind of looked at this either theoretically or by looking at actual amorphization zones, amorphous zones, and said that when you get above about 6 electron volts deposited into nuclear processes per angstrom of depth, the silicon has so much damage that it goes amorphous. So it's lost its crystal structure. In fact, if I draw a line here at 6 and I see where it cuts this profile, that would say that this substrate, if you're implanting boron at 100 kilovolts, I don't know what the dose was in this particular case. But it would be amorphized between about half of r_p and up to about r_p .

So in that region between half r_p and r_p , between these two lines, it's predicted according to this model, and again, it's going to be very sensitive to your amorphization threshold, this region in here will be amorphized. Everything else will be heavily damaged. So you can see what boron tends to do. Because of its distribution, it tends to create a buried amorphous layer with heavily damaged single crystal silicon on either side.

And I just want to say a little more about amorphization, shown here on slide 17. As we've talked about before, if you give it a high enough dose, you have enough of the crystal is displaced, it becomes completely amorphous. It loses all of its crystal structure. There's no more long-range order. At this point, we have a random arrangement. And the damage accumulation is saturated and really can't talk about damage anymore. Once you've amorphized, there are no lattice sites, so you can't really knock somebody off a lattice site because there is no lattice site. It's no longer a lattice. But it just gives you an idea.

This is taken from your textbook. This is a cross sectional TEM, so this is transmission electron micrograph pictures of an amorphous layer formation with increasing implant dose from left to right. So this is at very high energy. It's 300 kilovolts silicon being implanted into a silicon substrate.

So the interesting thing is, the ion now is silicon. It has the same mass as the substrate. It's relatively light, though. So here at 1×10^{15} , what do I see? Well, here's the surface of the silicon substrate. And you see there's this band, this region, where there appears to be darkness. There seems to be a lot of damage, a lot of stuff going on here. Dark contrast, you can imagine, might be associated with some kind of dislocation loops or some kind of damage in the substrate.

Here's a band here. Now at $1 \text{ and } 1/2 \times 10^{15}$, there's a band, which really has, it's buried and it's lost its crystalline structure. If we were to Zoom in here and do a selected area diffraction on this region, we'd find it's not a crystal. There will be no diffraction Laue spots. Whereas if you go down here, you would see crystalline.

Here's that amorphous layer at 2×10^{15} . It's increased in its width. It still hasn't reached the surface, though. And at 4×10^{15} , it's increased in depth, but again, it hasn't quite reached the surface. 5×10^{15} , almost. And finally, at 10×10^{15} to the 16^{th} , that's a very large dose, you've created an amorphous layer all the way from the surface down to some depth. I don't know what this is, half micron or whatever, and completely amorphized from the surface down.

And this is all at room temperature. It turns out, it's much easier to form an amorphous layer at low temperatures, 77 Kelvin. So you will sometimes see people specifying their implant energy. They want to use a special ion planter that does, if you're doing some of these materials experiments, that can hold the wafer at 77 Kelvin. That means the damage is much more stable. Remember, we were talking about that simulation from Lawrence Livermore Labs.

You saw that damage cloud and then a lot of recombination by diffusion and things. Well, if you're at low temperatures, you have a lot less recombination, so the damage tend to be more stable, especially with the light ion. So if you want to amorphize silicon with a light ion like boron or silicon, you pretty much, to do it reliably, you pretty much have to hold the wafer at low temperatures like 77 Kelvin. If you have a heavy ion that does a lot of damage, like arsenic or antimony, you can easily amorphize a layer from the surface down at room temperature.

OK. So let's go on to slide, that's a picture of damage. Let's go on to slide 18 and talk about damage annealing. So we have to do annealing after we do an ion implant, because we've bashed up the crystal. So what do we want to do when we anneal? We want to remove the primary damage caused by the implant. You want to put all the dopants onto substitutional sites, that way they can be donors or acceptors. We try to make the crystal as perfect as it was when you first started. You never really get it quite as perfect.

We'd also, in doing that, we want to restore the electron and hole mobilities, and the carrier lifetime, hopefully, to what it was. And you want to do all these things without really having much dopant diffusion. So it's a tough job. It's a tough job. And this is a model for damage annealing, relatively simple, but that was published back in 1991, that is very famous now, by Martin Giles, which is called the plus-one model for residual damage. Kind of a funny name.

But what Martin did was he said the following things. He said that most recoiled silicon interstitials-- so again, we're going to call this guy who's recoiled, who ends up in an interstitial space, an i, or a silicon interstitial-- most of them will find a vacancy. And they'll recombine very, very rapidly, either during the implantation process or in the first few seconds of annealing, a lot of recombination takes place. In fact, he calculated the distribution of remaining recoils after ion implantation, without even any annealing. And what he found, there's a net excess of vacancies near the surface and a net excess of interstitials towards the bulk. But there's still a lot of them recombine.

So this is a calculation from Giles, and he's got concentration of interstitials and vacancies as a function of depth. This solid line, or this line on top here, that starts at around 10^{20} to the 20 or mid 10^{20} and then goes down, these are the total interstitials and vacancy concentrations he calculated. They're almost on top of each other. On this log scale, you can't tell the difference. So you create a tremendous number of interstitials, tremendous number of vacancies, but most of them all recombined. And in fact, what he's plotted here is the net interstitials.

So that's the number of interstitials minus the number of vacancies at a given point. So that's how many are left over. So the net interstitials look like this. In fact, if you plot the net profiles here in the near surface region, you have net vacancies. So there's extra vacancies. Down in deep, you have net interstitials. But the total numbers of these net are much, much less than when you do the subtraction of the interstitials from the vacancies, than the total interstitial or vacancies that were actually created. So that's this point. To within three orders of magnitude, a lot, most of these guys recombine.

So to first order, what he said is that you can imagine that all of the original damage recombines, and leaves behind only one interstitial created for every phosphorus atom, or every ion-implanted ion that is ion implanted. So you can imagine. Let's say you are implanting a dose of phosphorus. In this particular example, my guess is 80 keV, 10^{13} to 10^{14} . He was saying that there's one interstitial created for every phosphorus incoming. So he would say there's 10^{13} to 10^{14} per square centimeter silicon interstitials created. So it's incredibly simple. You just take the dose of the atoms that are implanted, of the ions that are implanted, and say that's the number of interstitials per square centimeter that I create. 10^{13} to 10^{14} .

It's called plus-one because, and intuitively what he's saying is, if your implant is totally activated, so every implanted phosphorus ion finds a substitutional site, well, where did that silicon go? It had to knock. That silicon is somewhere as an interstitial. OK? And that's all he said. He said, assuming your annealing is good, you know what you're doing, to first order, so this is to get an order of magnitude estimate, the number of interstitials I create is equal to the dose of what I implanted. So that's pretty simple. I implant 10^{15} , I have 10^{15} excess interstitials now in my crystal. And that's the so-called plus-one model.

This just tells you a little bit more about how accurate that assumption might be. This is on slide number 19. I took this from your textbook. It's a little bit about damage evolution in time. So what it's a plot of is the annihilated interstitials, so the recombined ones, and vacancies per implanted ion. OK. And it's kind of here as a function of time. So if you look at, these are Monte Carlo simulations basically of interstitial and vacancy recombination. And if you go after a short time, really, only excess interstitials remain. And these can end up forming clusters.

But basically, this bulk and surface recombination take place on a very short time scale. So look at the bulk recombination. Very quickly after here, after only, I don't know, 10^{-6} seconds was at a microsecond. There's a lot of bulk recombination that's taken place. Surface recombination, for vacancies, takes a shorter time. And then finally surface interstitial recombination can take place within 100 to 1 second. So you have a lot of interstitial and vacancy recombination, that was the point of what Giles was saying, so that only plus 1 interstitial excess interstitial remains.

Now, what happens to those 10 to the 15th? If I ion implant 10 to the 15th phosphorus atoms, I have 10 to the 15th excess interstitials. The name of the game is, what happens to all these interstitials? Well, it turns out they coalesce. They get together into little defects. And these little defects are called, in curly brackets, called $\{311\}$ defects. And we'll talk in detail next time about $\{311\}$ s.

But the interesting thing about these little clusters, these $\{311\}$ s, they are defects end up being stable for very long periods. So the interstitials come, we have all these excess interstitials, they end up getting together and forming little defects, which then can stay around for 10 seconds, 100 seconds, maybe even minutes and hours, depending on the temperature. And it's these little defects these $\{311\}$ s, that end up being responsible for the process of transient enhanced diffusion that we're going to model in the next lecture.

In fact, on slide 20, just by way of introduction. These are pictures of $\{311\}$ defects. The right hand side is an actual high resolution cross-section transmission electron micrograph. And I can tell it's high resolution, because you can see these little dots, and they form in planes. Well, these little dots, each dot represents two silicon atoms. And they're all lined up in planes, because you're looking at the crystal planes in the silicon lattice.

And you see this little funny looking thing right here that's tilted at an angle. It goes from this point here, this dark band of contrast, to this point here. That is a ribbon-like region that is identified. That is the $\{311\}$ defect. And that's an actual cross section TEM micrograph on the left. It is meant to represent a color cartoon of what this thing actually looks like. The axis here, along the defect, is represented by this vector here, that turns out to be, this vector turns out to be in the $\{311\}$ direction. That's how these things got their name. They were tilted in the $\{311\}$ direction.

The long axis of the ribbon, this is a ribbon-like defect, the long axis points along 100 direction in the crystal. So the long axis here, you can't really see. It'll be into the board or into the slide. These little round things here, these circles, represent interstitials, and they're little dimers. They come in pairs, and they line up along this direction. Could be 100 angstroms long. And it's got some sort of width to it, which you can see here. And it has a certain capture radius. Within a certain region, it will capture these interstitials and form this little cluster.

These $\{311\}$ defects, they flow in pretty quickly, in a matter of a second or less. But they anneal out in timescale of minutes, at moderate temperatures. And as they anneal out, as this ribbon heals, it spits out silicon interstitials. And it's these silicon interstitials that are coming out of this $\{311\}$ ribbon that lead to the transient enhanced diffusion effect, these excess silicon interstitials.

It turns out that below a certain dose, or below a certain damage value, these $\{311\}$ defects can dissolve completely. So you can dissolve them. You can completely get rid of them. And at that point, when they're completely gone, you have no more TED after that amount of time. You go back to normal diffusion. Above a certain damage level, actually, it's a little more complicated. They actually get together, they turn into stable dislocation loops, which are more difficult or sometimes impossible to remove. That's like end of range damage. But there's a certain region in which you can remove all the $\{311\}$ s, certain amount of damage.

OK. So when we're going to talk about $\{311\}$ s and their annealing kinetics in great detail next lecture, when we go through the whole kinetics of transient enhanced diffusion. In the meantime, I want to go on and continue to talk about other aspects of annealing. So we know we create all these interstitials. They get together in $\{311\}$ s. We know if we go to high enough dose, remember I showed you a picture that we can amorphize the crystal, well, let's say I do create an amorphous layer at the surface. How can I get rid of that amorphous layer and restore the crystal to perfect, or to some kind of crystallinity?

Well, this is exactly what happens. I'm showing here cross-section TEM images. This layer is originally amorphous from the surface of the wafer down to some depth, and we're regrowing it. So we're annealing the wafer at 525. The initial implant was quite high. It's 200 kilovolts, 6 times 10 to the 15th of antimony. Again, antimony is a very heavy atom, so it amorphized from the surface down. And we're looking here, initially, at 0 minutes. So guess that kind of got cut off of the slide, but this should be here at 0 minutes, so just when you first start.

After 10 minutes of annealing at 525, look at the interface between the amorphous and the crystal substrate has moved up. So the amorphous layer is actually grown, or regrown, by a process called solid phase epitaxy. So there's no melting going on, but in a layer-by-layer method, the atoms in each layer are taking the template from the substrate and rearranging themselves back into a single crystal form. And this amorphous layer is progressing up from the depths towards the surface, you can see it as time goes on, in kind of a linear fashion in an epitaxial growth.

So we call it solid phase EPI as opposed to vapor phase. Vapor phase would be if I were injecting xylene or something and growing EPI. Here I just have a solid solution that's regrowing at 525. When you're done-- or, well, this isn't quite done. At 20 minutes it's almost done. Maybe half an hour-- you would see the entire amorphous region will have been regrown. It'll all be single crystal. And you'll be left with something though.

There's a band of residual damage that occurs called end of range damage, or EOR. And these are defects that are just below the amorphous crystalline interface. So you get an idea of where they are. They're always just below that. And those are pretty hard to get out. Those generally don't ever go away. You have a certain amount of that, for such a high dose, that you have to live with. You have to find a way to deal with it.

This slide on page 22 actually shows you some data from a book by Gibbons and Sigman, called Laser Annealing of Semiconductors. So it's about laser annealing. The book is. But this particular chapter is about solid-phase EPI or solid-phase regrowth, just in a furnace. And what they're plotting here is the regrowth rate. So this gives you the number of angstroms per minute as that amorphous layer traverses from down below up. The number of angstroms per minute that it goes, as a function of inverse temperature for silicon and under different conditions.

So let's take a curve here. Let's take this one curve right here that is called 100 silicon undoped as a function of temperature. And you can see, for undoped silicon, say at 525, the regrowth rate is about 20 or 30 angstroms per minute. So as a constant rate, these people have measured this, there's a technique called Rutherford backscattering, ion channeling, you can actually measure this rate going up. And you can see how fast it goes. Gives you an idea of how that amorphous layer regrows. And that's at 525. So you can see, amorphous layer regrowth can happen at reasonably low temperatures. You don't have to go too hot.

Interestingly, though, there are lots of different curves here. If you go to the 110 silicon, so if you buy a wafer that is 110 surface instead, you've got a different atomic density on that 110 plane, and in fact, it regrows slower. The same activation energy, e_a of 2.3 electron volts. Interestingly, that number is kind of intriguing. That's also what people believe 2.2 is roughly the bond breaking energy. And so imagine that interface between the amorphous and the single crystal, you probably have to break some bonds there, initially in order to realign the atoms.

So SPE is pretty fast. It depends on orientation, and it depends, what's really interesting, it depends very much so on doping. So for example, here's silicon that's doped with arsenic. You can be quite a bit faster, 3 to 5 times faster, or phosphorous, or boron. Boron is the fastest, it looks like. If you accidentally put oxygen or nitrogen into your sample, look what happens to the regrowth rate. You can go down by almost an order of magnitude. So impurities have a big effect.

And argon would even really do it, really frustrates the regrowth. And so it's important to know what's in your crystal, what your ion implanting will have a big effect on the regrowth rate. Interestingly, when you do this, most of the dopant atoms, like arsenic, phosphorus, boron, are incorporated as that amorphous layer moves up. They're actually mostly incorporated on substitutional sites, even at low temperatures. So it turns out, if you can create an amorphous layer, this is one of the best ways of activating a dopant, is to create an amorphous layer in silicon.

Now this is the way silicon anneals. Other semiconductors, gallium arsenide for example, is just the opposite. If you amorphize it, God help you trying to get the dopants to get activated. It's not so easy. But silicon has this beautiful property that as that layer moves up, a lot of the dopant atoms get forced onto substitutional sites. And then the donors and acceptors are already activated at low temperature, to a great degree. So that's a nice property. Assuming you fully amorphize the crystal.

OK, let's go on to slide 23. I mentioned that life is not perfect. Even if you amorphize, you always end up with this end of range damage. And here's a plot on the left of, let's say this is the surface of a silicon wafer here, going down in depth. And here's my concentration. And let's say I have implanted something, maybe it's arsenic, I don't know what. And it has this ion implanted profile. It looks like this.

And there is, the maximum damage for this particular profile is just below the amorphization threshold. So this region from here to here, from the surface down to, this dashed line is called the amorphous crystal interface, that's all amorphized. So that's going to regrow layer by layer and give you very good crystal quality, if I just go back a couple of slides here. Again, it was amorphized to this depth. If you look at this depth when you're done, there is no residual damage. It's below the initial amorphous crystal interface by some distance, that you end up with it's end of range.

And it kind of makes sense, because at this point, you have deposited-- averted this depth enough energy to amorphize the crystal. Below that, you've deposited a huge amount of energy, but not enough to amorphize. So you've really bashed it up, but it's not amorphous, so it can't regrow by solid phase EPI. Instead, it creates dislocation loops, and that's what we call the end of range.

So if I go back to the slide number 23, just below the amorphous region of this end of range damage. And here is my ion implanted arsenic. OK. So. Now, let's see. I'm using this to make an np junction. So what I need to do, if I'm going to make an np junction, what I typically want to do is just to diffuse the n type dopant a little bit deeper than the end of range damage. So we tend to try to get a little bit of diffusion after we've done an ion implant. Sometimes, that can be helpful, because why would that matter? Well, here's my end of, I've done an overlay.

This is a cross-section TEM micrograph. And here's my end of range damage, this little black dislocation bar. And imagine I diffuse this arsenic just a little deeper. So now I have an n plus region that goes just below the end of range damage, and then below that I have p. Well, this is n plus, this is p. The depletion region, again, always tends to extend, if you remember your np or pn junction physics, the depleted region is going to be in the lightly doped side, where the yellow is. So the yellow is depleted of free carriers.

And you know, in the depletion region, that's where you get the most amount of recombination of electrons and holes to create leakage in a pn junction. So the depletion region is the last place you want damage or defects, because in the depletion region, electron hole recombination is very efficient. So what we do, what people do, this is kind of cheating, but they cover up the end of range damage with a heavily doped region.

So if you see people making source drain junctions, they often diffuse it just a little ways beyond the EOR, and then they get diodes that are pretty good. Reasonably good. If they don't diffuse it at all, they say I'm going to have a perfect anneal, no diffusion, in fact, then the n plus junction might have been back here, and then the EOR might have been in the yellow region, in the depletion region, at the end of range damage, and you could get a lot leakier diode.

So there's a lot of optimization when people activate source drains, or something where they care about leakage. There's a lot of optimization in just exactly how much damage dopant diffusion do I allow, and how much time do I allow to try to get rid of these loops? So it's kind of an interesting, damage annealing is an interesting combination between knowledge of p n junction physics, and electronic recombination mechanisms, and knowledge of the crystal structure and what kind of damage you've created. And you need to optimize that, because you'll never get rid of all damage, to a certain extent, for every single implant. Some implants are easier to anneal than others.

OK. So that was in terms of, I spoke there quite a while about trying to get rid of damage, because I wanted to get good pn junctions. I want to get low leakage. OK. That's only one requirement, is low leakage in the crystal. I also want to activate the dopant. You want to ion implant 10 to the 15th arsenic atoms per square centimeter, you'd like to get 10 to the 15th electrons per square centimeter. You'd like to have all of that activated. Or boron. That's the whole reason you're putting it in. You don't care about boron in the lattice. It's the holes that it produces, or the electrons that arsenic produces.

Well, it turns out, if you amorphize the substrate, again this only applies to silicon, gallium arsenide is just the opposite. But if you amorphize the silicon that we mentioned, that a solid phase epi is an ideal way of repairing the damage and also getting the dopants onto substitutional sites. If you don't amorphize, so you do a lower dose that does a lot of damage but doesn't create an amorphous layer, so you can't have SPE, it turns out that activation is a lot more complex, because you create defects that are somewhat stable over time and temperature.

So here's just an example, of a plot of the annealing characteristics of arsenic and silicon I took out of your textbook. And the y-axis here is the fraction active. So I want to be up to one. If I'm at one, that means every single arsenic atom that I implanted contributes electron and they're all electrically active. And this is as a function of annealing temperature.

And there are different doses shown here. So the red curve assumes I've amorphized the silicon. And look at this. Very interesting. At 550, 525, I can activate for $1e-15$ implant, I can activate almost 80%. That's pretty good. And if I increase the temperature up to 800, it doesn't actually activate much more. So if I want to do a really low temperature anneal in silicon, I can do that with a $1e-15$ implant. Except for 20% of my atoms, I get them active.

But let's say I only implanted $1e-14$ instead. At that same temperature, only 10% are active. Very few. So if I have an intermediate dose $1e-14$ might be more typical dose, say, for the source drain extensions, might be $1e-14$ That's a miserable dose, because it's pretty high to do some damage, but it's low enough that it doesn't amorphize. And so to activate the extensions here, this $1e-14$ dose, you really need to get up to 750, 800 or whatever to get 80% activation.

So we don't just think, oh for any implant I want I just do a certain anneal and I'll get the same results. It really depends on the amount of damage I did, and whether I was below or above the amorphization threshold. Here, if you go to really low doses, it's a piece of cake. Again, here at 10 to the 12th it's not bad. Because you did so little damage even at 600, 700, you can get pretty close to 80% to 100% activation. So the tough spot is this intermediate dose regime which is exactly the dose regime where we're ion implanting source drain tips or source drain extensions. And those need to be active as full as possible.

So unfortunately, it does make our life a little more challenging. That was arsenic. That actually isn't that bad. Even more complicated is boron. And again, I took this plot from your textbook. It's the same kind of plot, now I'm on a semi-log plot, so I've got a log scale on the y-axis and temperature of annealing on the x-axis. And different doses again. So for boron, if you have a low dose up here, 8×10^{12} , again at 550, 600, I can get, oh, maybe 70, or 50% to 60% active. Not that great, but not terrible.

Interestingly, at the very high dose regime, 2×10^{15} , if I kneel at very low temperatures, say 500, I get a reasonable active, well it's not very good. It's only 1 in 10. But it actually goes down. It actually goes down in a certain temperature range between 550 and 700. So you're actually reverse annealing. You're taking the dopant atoms off site. And then finally, it takes up again. When it takes off again, when you get up above about 750 or 800.

So to get full activation of boron, you really need to be in silicon. At almost any dose, you need to be in the high temperature regime, 900 or above. And that's a real problem. That's why you will hear people say, oh boron is hard to activate. In general, it requires higher temperatures or some kind of more sophisticated annealing technique than arsenic. And that's because of the type of damage it does. It does not tend to amorphize, unless you go to really high doses. It's very hard to get an amorphous. And so you cannot take advantage as easily of the solid phase epi regrowth.

Now, one thing you can do. You say, OK, boron doesn't amorphize. What could I do? People do something called pre-amorphization. So before they implant a boron profile, they might implant a high dose of silicon, amorphize the crystal to a certain depth, then put the boron in by implant, and then regrow the whole thing by SPE. And you can cheat that way, in some sense, or it's an extra step, but then you can get a little better annealing characteristics. But this is assuming you didn't do any pre-amorphization. This is just, implant the boron at this dose, and try to anneal it.

So what is this reverse annealing behavior? Well, it's probably, oh, there's a typo here. It's thought. There's a T missing. It's thought to occur. Because there's some kind of competition here between the native interstitial point defects, the things that you're creating-- remember, you create more point defects as you raise the temperature-- and the boron atoms themselves for the lattice sites. Because again, we did not amorphize the crystal. We just created a lot of damage. So there's a sort of range where you just don't want to be annealing. Here at this dose of 2×10^{14} , you want to take the boron up higher. The problem with going to higher temperatures, of course, is you get more diffusion.

So let me go on to slide 26 and just say a few words. We're going to spend the entire next lecture on this, on something called transit enhanced diffusion, and this is sort of exemplified here. Here is an implant that was done, say, of it could be boron or arsenic. It doesn't really matter at this point. But let's say it's boron. And we see two different anneals. The blue curve I did 1,000 degrees C for 10 seconds, and I got that diffusion profile after I implanted the boron.

And now the red curve is 2 minutes at 800. And so, interestingly, you would think, well, I'm annealing it at 800. That's a much lower temperature. The boron diffusion coefficient is down by two orders of magnitude at 800. So even though I'm annealing it for a slightly longer time, it's certainly not 100 times longer or 200 times longer. So you would think that the 800 degrees C profile should be shallower than the 1,000.

So TED, unfortunately, is this strange effect, and we'll talk about how this can happen, is that even at lower temperatures you can even get deeper junctions than you can get at higher temperatures. So this really annoyed people and was quite confusing and it needs to be explained by the implant damage and its effect on how the boron diffuses. And in fact, TED today is pretty much the dominant effect that determines junction depths and shallow profiles. It's not so much ordinary diffusion or concentration dependent diffusion. Well, for arsenic, maybe, but for boron, everything is. You cannot possibly model a shallow junction unless you know how to model TED.

And if I just go on to the next slide, just to show you, if I had given you that problem and told you nothing about TED, and you would have looked up in your book, say, at 800 degrees-- which is somewhere right around here, here's the boron diffusion coefficient, right at this point-- and that at 1,000 degrees it's up here and the difference between the two is more than two orders of magnitude. So it's almost a factor of 1,000.

So clearly, if you're diffusing at 800, it should be a lot less, a lot shallower junction, even if you're doing it for two minutes versus ten seconds. Two minutes is 120 seconds, so 10 into that. So I'm doing it 12 times longer. Fine. But the diffusion coefficient is 1,000 times, should be 1,000 times slower at 800, but it's not. In fact, it's 1,000 times faster. It's a lot faster. And that whole effect is called trans enhanced diffusion. It results from all the interstitials that were implanted, that plus-one interstitials giving rise to {311}s and all that. And so the next lecture, we're going to spend trying to understand these profiles and their time dependence.

So let me summarize on slide 28. Basically, so far, in this lecture, we said that we can separate the energy loss processes completely independent of each other. To first order, we have nuclear processes, which are these billiard ball collisions, and we have electronic drag force. The nuclear stopping dominates when you have a heavy ion over most of its path. For a light ion, nuclear stopping dominates only at the end of range, at very low energies. It's the nuclear stopping that damages the crystal by the creation of silicon recoil. So it knocks silicons off of their lattice sites and they knock other atoms off. And that creates what's called a collision cascade.

The damage profile peaks right near r_p . If you want to calculate that, you can just, maybe 80% of r_p gives you a rough idea of the depth of damage. For heavy ions, the damage is more stable. That is, it doesn't tend to anneal out in the ion implanter at room temperature. And there's a tendency for a heavy ion to form an amorphous layer from the surface all the way down to some depth. That amorphous layer can be regrown relatively easily in silicon. Not true in other semiconductors, however. It results in relatively efficient dopant activation.

For light ions like boron and phosphorus, below a certain dose, it is difficult to produce an amorphous layer at room temperature. And as a result, the activation of the dopants is a lot more complex. It's very dose dependent, and it can be temperature dependent in an odd fashion. There is this plus-one model for residual damage that says that there's roughly one excess interstitial for every primary ion that I ion implanted. So it's relatively easy order of magnitude estimate. And this is after the initial vacancy and interstitial recombination is taking place, which only takes a few fractions of a second.

These excess interstitials, however, they cluster into little {311} defects. Those defects dissolve at a relatively slow rate on the order of ten seconds to minutes, and it's those little, that {311} evaporation or dissolution, that gives rise to the time dependence of TED. So by understanding {311} defects to a certain extent, we can understand the time dependence of TED and, in fact, by understanding {311} defects, we'll be able to come up with a model that could predict this kind of strange looking behavior.

OK. So that's what I wanted to go through today. There were three handouts. Make sure you've got them. Today's lecture notes, homework number four which is due on November 2, and the solutions to homework number two. The solutions have been on the web for a while. I just forgot to hand out the paper copy. OK. So we'll see you next Tuesday. Hopefully we get to watch the World Series.