

JUDY HOYT: So before we start with today's formal lecture, there's a couple of things I want to talk about. One is the schedule, which I have shown up here. And you'll notice, or if you remember, I've actually shifted around one of the lectures. Here we are today, I think October 19. And we're supposed to have this lecture on the Suprem4 process simulator. I decided I wanted to do that lecture later, primarily because I want to have talked about transient enhanced diffusion and those effects.

A lot of the examples I show and how to use Suprem relate to transient enhanced diffusion. And I realized I hadn't really-- I wouldn't have introduced it. So we're going to have instead the ion implantation lectures. There's four on ion implantation, two of those are on TED, and then we'll do that lecture. And as I mentioned also because of the Red Sox game last night, I don't have homework number 4 ready to go out to you yet. And so we can blame the Yankees for that. But I'll have it next time.

There is another handout as well before we go to the formal lecture, handout number 20. It's a one-page document. Hopefully, you picked that up. I want you to start thinking about your term projects. And so it's a good time to start thinking about it. What I've listed here is maybe 15 or so potential topics that you might want to use for your term project. This is not an exhaustive list. These are just examples.

So for example, you could study dopant diffusion in silicon germanium. You could read up on that. You could study laser annealing of ion implants for shallow junction. These are all things that we don't get a chance to cover any great detail-- selective epitaxial growth of silicon or silicon germanium. But they are related to front end processing I'd like you to learn more about in depth. So you pick a topic or you make up a topic. Any topic you pick or make up, it needs to be approved by me.

And next time over the next couple of weeks, I'm going to bring in a clipboard. And you'll have a chance to sign up for a topic. And I also want to know whether you want to give an oral presentation or a written report. The written report is going to be-- it needs to be 20 pages or less. And an oral presentation in class is probably going to be about something like 15 to 17 minutes plus questions. And I need to start scheduling that so I need to know whether you would prefer to do your final term project as a written or oral.

And the term project is not meant for you to go read one textbook or one paper and then regurgitate it. The whole idea is you go to a library, you go online, and you read a series of research papers, so maybe some texts, but most of these topics are new enough that you'll have to read research papers, and then write up a report or give an oral presentation about the topic. So I don't want you just to regurgitate one paper.

And then the other thing is if you do decide to do a presentation, you'll need to have handouts for the class, just like the handouts that I give you, not as long. And if you need help making the handouts in terms of getting them xeroxed, my assistant will help you with that. So even if you do the presentation, there'll be some preparation of graphical information. So that's to get you started thinking. And as I say, next time I'm going to bring the clipboard and you can start signing up for topics.

OK. So let's go to the formal lecture for today, which is handout number 19. And we're going to talk-- as I mentioned today, we're not doing the Suprem4 lecture. I shifted the schedule. And we're going to talk about chapter 8. The next four lectures are going to cover chapter 8 so I hope you will read it. Chapter 8 is about ion implantation and annealing.

So we've already talked about introduction to the basic process flow in CMOS. We know how to fabricate wafers and clean them. We discussed in detail thermal oxidation and two-dimensional effects. And we spent the last four lectures talking about diffusion and how we model it. We're going to come back to diffusion after we've talked about ion implantation. But first, I'm going to spend the next couple of lectures talking about the basic concept of ion implant and how it's used, the impact of the crystal structure, there's a phenomenon that's called channeling, how ion implantation is modeled, and damage annealing, and transient enhanced diffusion.

Ion implantation is somewhat unique in this course or in any course on processing. It is probably the only process that we can actually model from more or less first principles. That is, if you look at our models for [INAUDIBLE] models for oxidation, there are a lot of chemical processes involved, and there's a lot of unknown constants here and there. And the models tend to be very empirical.

Ion implantation, on the other hand, is based on fundamental physics, a lot of it. And so it's one of the few processes you can actually model without too many free parameters. You'll see there are some parameters but not as much as other. So for people who like first principles modeling, it's somewhat satisfying.

OK. So let's go on to slide number 2, showing you what I call the old fashioned or traditional approach to getting dopants into a wafer. You would coat it with oxide. You would open up a hole. And we talked about last time, last few times, doing what's called a predeposition.

So that used to be, in predeposition, it was a high temperature in diffusion from the gas phase. Typically, in a pre-dep, you would have in the gas phase some kind of dopant gas like phosphine or whatever. And you would keep the concentration of the dopant at the surface at the solid solubility at that temperature at a very high value. And you would do it for a certain amount of time to put in a certain control dose. And then you would turn off the gas that had the dopant in it. And you would do a drive in, a deep diffusion drive in.

A big problem with this method though, although it was the original method, it was very limited and it's very hard to control the actual dose, the integrated dose of the dopants that were put into the wafer. If you had an extra few minutes in the furnace or the furnace temperature was slightly different, you would get a completely different dose. So this method was good in the old days but it really gave way some time ago to ion implantation as the method.

And the only problem with ion implantation, the good part is it has tremendous control of the dose. The bad part or the issues we have is that it does introduce damage. And we have to anneal that damage and repair the crystal in order to make a good device.

So let's look on slide 3. And from this slide, we can understand why ion implantation was introduced into manufacturing in the late 1970s. And this is an interesting history. And it's also an interesting example because it gives you an idea of if you have a new idea-- let's say you're working here on your PhD and you have a new idea for some great new process.

The new process, whatever it might be, has to really be an enabling technology. It has to do something that you cannot do by any other method because all new processes have some bad aspects to them. They cost money. There's development time. Ion implantation adds damage. Why would anybody want to do it? Well, this is why because it did something that they couldn't do in the past. And basically, in the late 1970s, ion implantation was introduced as the only means to accurately control the threshold voltage of a MOSFET.

Without ion implantation, people couldn't make circuits because they could not control the V_T at the level that they needed to be able to control it. So if you go back and look at this as an equation we introduced a number of lectures ago-- we didn't derive it. I just wrote it up there and you'll have to take it as truth. But what it says is that the threshold voltage of the MOSFET depends on a number of terms-- the flat band voltage. It depends on the doping in the substrate, the oxide thickness through $[\phi_{SiO_2}]$ and things like that.

There's one last term here, which I circled, which is the electronic charge q times ϕ where ϕ is the dose of a dopant that's ion implanted into the surface of the wafer divided by $[\phi_{SiO_2}]$. So this extra term, this isn't here unless you're doing an implant but that extra term can be added in. And if you can very precisely control ϕ , the dose, then you can very precisely control the threshold voltage and set it to a level that you need.

So on the left-hand side, I'm just showing a schematic, just reminding you from the first lecture, what people do for a V_T adjust implant. For example, if you're making an N-MOSFET on the right-hand side, you would ion implant with boron, a very, very thin surface region. You block off with photoresist, the PMOS area. And you would implant boron to adjust the threshold voltage. Typically, doses are pretty low, 10^{12} to 10^{13} , something like that. And energies in this range, we'll talk about dose and energy in more detail in the next couple of lectures.

The nice thing is that this dose can be accurately controlled, and that's because it's based upon integrating a current, electronic charge that goes into the wafer. That's because of the way we do ion implantation. So it turns out the dose is just 1 over the area, where the area is the implanted area, times the integral over time of i . That's supposed to be an i , i standing for the beam current that's collected into the wafer, divided by the electronic charge, q , integrated over time.

So because the ions coming in are charged, they represent a current flow into the wafer. So they're charged, and as the current flows into the wafer, you can measure the current with an ammeter very accurately. And then you can integrate it over time. So you start the implant opening up a shutter, and then you iterate over time. Each charge that comes in gets counted, and then you stop.

So you can very, very precisely because it's based on charge counting techniques, not on some chemistry pushing away from into a hot furnace, hope the temperature is right. Oh, you pushed it for ten seconds too long. What? Pull it out. It's more of a physics-based approach. From a manufacturing point of view, that was an important breakthrough in the fabrication of MOSFETs.

Let's just do an example on slide 4 that illustrates the controls that you can get during ion implantation. So let's say that we need to calculate the dose rate. OK, so that'll be the number of atoms per square centimeter going into the wafer per second or per unit time. I'm going to assume that we have an ion beam current, so that's the beam coming into the wafer. The beam current is 0.1 microamp. OK, and we measure that very accurately. It's a tenth of a microamp.

And we assume we're implanting an area that's 20 by 20 square centimeters, just for the sake of-- we just know that's our area. So basically, I can calculate the dose. And just like that formula we just said, it's 1 over the area times the integral of the beam current, which is 0.1 , over the electronic charge, dt . Now if the beam current doesn't vary with time-- say, it's a constant 0.1 -- then the integral is easy. It's just the current times the total time of the implant divided by the electronic charge times the area for a constant current.

OK, so the dose rate will then be, just take the dose and divide it by the total time. Average dose rate, dose divided by t , it's just the beam current divided by q , the electronic charge, times the area. So you end up with a dose rate. For that beam current, you're putting in about 1.6×10^9 atoms per square centimeter per second. So that's pretty good control.

Let's say, I want to implant 10^{12} atoms per square centimeter. Well, you can just divide that into 10^{12} , and you figure out how many seconds it would take. It's a reasonable amount of time. And then you can adjust your beam current and the time to make it something that's physically realizable. So this has tremendous control, very good accuracy, and it's very reproducible.

It has good dynamic range because you can adjust your beam current down from something like 0.01 microamps if you want to do a very, very low dose. You're putting at a very slow rate. You put the ions in, and you integrate them up. Or you can do it if you have a very high dose. You might have to have a milliamp of beam current going into the wafer. Either way, you get very, very good control.

So let's compare this ion implant case to the case of where I want to use the old-fashioned method. So instead of implanting boron, I'm going to do a boron pre-dip at 1,000. And again, in order to control this, we don't want it to depend on the flow rate of the gas, so we typically do make it so that the surface concentration is at the solid solubility for boron, which is about 10^{20} per cubic centimeter.

So remember, when you're doing a pre-dip, you get a complementary error function profile and that the dose of the integral of that can be given by this equation. So the dose is 2 times the surface concentration divided by the square root of π times the square root of dt , so where d is the diffusion rate, the diffusivity of boron in silicon at that temperature. And so we can look that up in the chart, so it's about 10^{-12} centimeters squared per second at that concentration and temperature.

So in one second, you've already put into the wafer 10^{14} atoms per square centimeter, a very high dose. So there's no way you could really do-- if I want an implant of 10^{12} , there's just no way you can get that low of a dose and with any kind of control by using a chemical pre-dip technique.

So again, this was the big breakthrough in ion implantation. People found they needed to control threshold voltage, and they needed to implant relatively low doses very accurately so they get a very good number for V_T . And pre-dip was not going to be a way to do it, so ion implant filled that need.

OK, so let's go on to slide number 5. Before we get into the details of the physics, we'll just talk about some basic concepts. As I mentioned, it's the dominant method used today to introduce dopants, in spite of the fact that it has a huge amount of damage it does to the lattice. And these are all the good things about it. I've already mentioned control, but dynamic range, a large range of doses. It's essential to maintain or control the V_T .

You can make a buried or retrograde profile. Remember, with a complementary error function, if you plot the shape of the profile as a function of depth-- so this is depth, and this is the concentration-- in diffusion from the surface, you always get a higher concentration at the surface than you do here. But with an implant, you can actually-- if you want, you can make it look like that. You can bury the peak and make it look retrograde. So you have the ability to do that. You have the ability to control that.

It's a low-temperature process. It happens at room temperature, which is an advantage. And there's a wide choice of masking materials. If you're trying to mask a diffusion, well, there's only a few materials you can put down, either oxide or nitride, that might keep the dopants from diffusing through there. If you're trying to mask an ion implant, you can use photoresist, a very simple, low-temperature material.

There are some disadvantages, though. Of course, we mentioned that you are bashing up the crystal, and you're damaging it. And you have to remove that damage by heating it, and the heating causes some problems with diffusion. In fact, there's an anomalous type of diffusion due to ion implant damage that we're going to study called Transient Enhanced Diffusion, or TED, that happens due to the injected defects from the implant.

Nevertheless, people have found out how to model TED, and they found, in some cases, ways around it. And the other disadvantage is you can insulate, and you can have oxide on the wafer. You can tend to charge it up, so you have to be a little bit careful. But there's ways to flood the wafer with an electron gun and all kinds of things to keep your wafer from charging up.

If we go to slide 6, this is a schematic diagram of what an ion implanter looks like, at least a sort of a traditional diagram of how an implanter is designed. What you have here in this big box are two key things. And all of this is in a vacuum system, so this entire thing is being pumped down by a pump. You have an ion source. So that has a gas in it or something that ionizes the species of interest.

Maybe you have a gas in there. You put BF_2 . OK, so BF_2 gas or whatever. And then you're going to ionize, so you have a way with a hot filament of creating ions, and you create B double plus, B plus, fluorine, BF_2 -plus. You create all these different ions, of which there are some that contain boron, which is what you want to implant, let's say, in this example.

The ion source extracts the ions at a given energy, say, from 0 to 30 keV. So we can set the energy of the ions coming out just by knowing this extraction voltage, the extraction that we apply to this grid. So we know the energy of the ions coming out, and then we do a standard-- what's called an E cross B filter. This happens in any kind of mass analyzer. This is what's happening when you do secondary ion mass spectrometry or whatever.

But if you know the energy of the particles coming out, and you bend them in a magnet, you can mass analyze them and only pick those ions that have the right mass. So I can be sure that I just implant boron. I don't want to implant fluorine. I don't want to plant nitrogen or oxygen or carbon or anything else that happens to be in the source. I only want to pick one species. In this case, let's say boron.

So basically, how this works, if you go back to any of your elementary physics class, you know that what's happening here in the ions as they traverse this arc in the magnet, the centrifugal force is balanced by the magnetic force, the v cross B force. So that has to be the case for ions going through this arc. And you also know the velocity of the ions coming into the arc because you know their energy. And that is set by this extraction of voltage, the extraction.

So I have a velocity that I can plug into this upper equation, and that is going to relate, then-- end up relating the mass of the ion that makes it through this resolving slit. That mass can then be related to the energy, which is related to the extraction voltage and the magnet current. So this is usually an electromagnet, which I can adjust the current. So all you need to do to be sure that you're implanting one particular species is, you adjust the extraction voltage, and you adjust the current going through the magnet. And then the only ions that make it are those that you want to implant with the right mass.

So this is a way. And then once you have the ions that you need-- you've extracted out all the others, and you just have the ones you need-- then you can put them through a high-voltage accelerator to get them up to, if you want to implant very high voltages, say, 50 to 100 or 200 kilovolts. You focus the beam, and then there are a series of XMY scan plates where you have an E field that's taking the beam and moving it around like I'm moving this laser pointer right now, scanning it randomly across the wafer.

And that's how you do the ion implant. And then the important thing is that the wafer is attached here to this device, which has a current integrator, which integrates the dose going into it. So it knows when to start and stop the implant. So a very basic sort of physics-type of accelerator setup.

We go on to slide 7. What I just showed was the practical aspects. Let's talk a little bit-- we won't talk about the physics of stopping yet, but let's just talk about what profiles look like. I drew a profile up here that I said is somewhat what an ion implant might look like. Let's think about what that might be.

Imagine we have a surface of a silicon wafer. So the silicon is on the right here-- this is the interface with the vacuum-- and on the left is a vacuum. And you have a beam of ions coming, in this case, right in this red spot. Any given ion that comes in is going to come in and experience a series of collisions. And in fact, many of them are billiard ball-like. It's going to have a series of collisions with atoms in the substrate. You can see it's having little collisions here and there. And it loses energy as it goes in via these processes and eventually comes to rest at some point here, shown by this little blue dot.

So there's a couple of pieces of terminology. We call the total range of this ion as the total path length. It's the sum of all these little lengths. That's the total length it traversed. The total length is not as interesting. What we usually care about is the depth. So the projected range, that's the projection of its total length along the normal to the surface. That's what we care about, how deep did it go. Usually, projected range is very near the peak of the distribution.

So the things you need to remember are, ion implantation is a random process. Ions come in. A whole bunch of them come in, and they randomly hit various things and lose energy and end up in some distribution. So it's a statistical type of process. The ions are pretty high in energy. We're talking about 1 kilovolt at the lowest probably up to maybe a megavolt. That's the range of energies, believe it or not, that people use in making modern devices. They can even get as high as a megavolt.

These hit the silicon substrate, and they lose energy. I already mentioned about these nuclear billiard ball-like collisions, just like on a pool table. There's also electronic drag forces. So there's two mechanisms, and in the next lecture, we'll talk in great detail about the physics of nuclear collisions and the physics of this electronic drag. Just for now, take it as a given that there are these forces, these things that slow the ion down so they eventually end up coming to rest in the crystal.

So if we go on to slide 8, let's take a look at a plot of the distribution of ions implanted into silicon. Here's an example where the beam energy is 200 kilovolts. So you adjust the beam energy so that you have 200 keV ions coming in. And this is a plot of the concentration of the dopant in atoms per cubic centimeter as a function of depth into the crystal, into silicon.

And we're showing several different types of ions. So if you look at this blue curve here, this is antimony. Notice the antimony didn't go very far. It ended up piling up, looks like a Gaussian-like distribution, pretty close to the surface. The black is phosphorus. It went a little further. Arsenic, I'm sorry. The black is arsenic. Phosphorus here is the red. It goes further. And the one that goes the furthest here is boron. That's shown in green.

Boron is the lightest of all those elements, so it kind of makes sense. Boron is lighter. It's going to lose less energy in its collisions, so it's going to travel further overall. And not only does it travel further into the crystal, its projected range here is about 0.5 or 0.55. It's also broader. Its distribution is very broad. Look how broad this is compared to the antimony, which has a very, very tight distribution.

So as a physical process, it depends very much on the mass. We often describe profiles very roughly by a Gaussian distribution. It's not really Gaussian, but for a heavy ion like antimony, it turns out to be reasonably accurate. So the first-order description is a Gaussian, and a Gaussian has two variables. One is the projected range, R_p , which gives the peak point in depth of the Gaussian, and the standard deviation, ΔR_p , which for a Gaussian, is a measure of the half-width basically.

So for a simple Gaussian, we only need two parameters. And in fact, you can write-- on slide 9, I've written down the equation for a simple Gaussian profile with a certain peak concentration. And this Gaussian is centered at R_p . You can see that just by plugging in x equals R_p , this whole thing goes e to the 0. So you get the peak concentration at x equals R_p , and it has a standard deviation, ΔR_p .

The integral of this, if you integrate the Gaussian, that Gaussian's integral is known analytically. So the integrated dose is just square root of 2π times ΔR_p times the p concentration. So this tells you, if I told you I implant antimony at a particular dose and energy, you can look up-- there are tables which you can look up the ΔR_p , and then you can calculate, knowing that if I give you the dose, and you know ΔR_p , you can calculate the peak concentration right away if you assumed it was a Gaussian. So you can do some very simple things with simple hand calculation.

Again, what is q ? It's the dose. It's the integrated number of ions per square centimeter. So literally, in this profile, if I took the area under this curve, under the concentration versus depth, that's the dose the number of ions per square centimeter. And again, we control that by measuring the integrated beam current.

We'll say a little bit more about this. But if I go back briefly to slide 8, you notice some of these-- the antimony is quite symmetrical and quite Gaussian-looking. The boron profile is not so symmetrical. It's skewed. It has a larger sort of tail region or region towards the surface. It's got the skewness to it, and we'll talk a little bit more about why that is and how to model that.

So it turns out the projected range, R_p , and ΔR_p , the standard deviation, of all the common dopants in randomly oriented silicon, so assuming the silicon is amorphous, and there's no channeling going on-- we'll talk about channeling later-- these have all been measured experimentally. They've also been calculated very accurately from first principles. And these theoretical calculations can be done for almost any combination of ion and substrate.

So you could be ion implanting cesium into carbon. And if you know the basic mass and the number of atoms per square or cubic centimeter in the substrate and the energy, you can calculate what the distribution looks like in an approximate sort of-- if you're assuming you have a Gaussian. So all of this has been calculated and measured. In fact, it's tabulated in a number of books.

And it's in your textbook. Rather than tables, what your textbook does-- I took this directly from the text-- it gives you plots. So on the upper left here, this is a plot of R_p . So this is the depth in microns that the ion travels. It's the peak, R_p , as a function of the energy of the implant for the common dopants. So here's boron. Again, you can see boron goes very far. So if I'm at an energy, say, of 80 kilovolts, and I have boron, the projected range is a quarter micron. That's 0.25. For arsenic, at that same voltage or that same energy, rather, it's about 0.05 microns or about 500 angstroms. So it's much, much shallower.

And so you can see R_p has been tabbed, has been is given by these straight lines here. And the standard deviation, ΔR_p , as a function of energy is also given. So again, boron at 80 kilovolts in the lower right, it has a standard deviation of 0.08 microns. So it's going to be much broader than arsenic, which has a standard deviation of only 0.02 roughly.

So you can read these right off the plot. And if you make the Gaussian assumption, all you need are these plots and a simple hand calculator, and you could draw the implant profile roughly for any of these species at a given energy. Now, these also assume that the substrate has been tilted and rotated-- and we'll talk a little bit about this-- to avoid channeling. So this is assuming the substrate is amorphous.

So let's go on to slide 11 for a minute, and let's just think about what happens. So that's the distribution. Now, we want to think a little bit about what happens if I don't scan the beam across the wafer. I was saying, we take this laser beam in a real implanter, we're continuously scanning it to randomize and to implant the same dose everywhere on the surface.

But if we just hold the ion beam at one point, and just look-- let's look at the case where the implant beam is centered on a point in space, x , y , and z , at the 0, 0, 0 position. So if I draw axis x going this way, y going vertically, and z going out of the board or out of the page, at the 0, 0, 0 position, I imagine bringing in ion beam and just holding it there and then looking at this cloud of ions and where they end up, where they stop.

And in fact, people have done this on the computer. It's not that hard to do a simulation called a Monte Carlo. Monte Carlo is just like gambling, just like in-- you send in a certain random distribution of ions. They have known energy but different impact angles, and you simulate these random trajectories for a group of ions. Maybe you could ion implant 10,000 ions at some spot in the wafer. And you see where they end up, what their three-dimensional distribution looks like.

So this particular simulation is what 1,000 atoms of phosphorus that were shot into a silicon wafer at 35 kiloelectron volts. And that's what the cloud looks like. That's where they end up. And in fact, it's kind of an elongated ellipse in the x direction. x is the implant direction, and because a lot of the high energy ions undergo very small angle collisions. So they don't end up going sideways very much. They end up scattering forward and then eventually stopping when they lose all their energy.

Let's go on to page 12. Now, I'm showing some two-dimensional projections. What I showed before was a three-dimensional picture of this elongated ellipse, this cloud of where all the ions ended up. Now, I'm projecting that cloud onto the yz plane here. So this is along the beam direction if you're looking straight. The beam is coming straight in here like this, right in the center. Along the beam direction, it has sort of a lateral distribution that's kind of symmetric, which makes sense.

If we look at the top view, so looking top-down, and the beam is coming in here right at this point, y, in fact, it looks-- you can see the depth distribution. It's kind of peaking right around here at this point, and then it's kind of tailing off as you go in deeper. So you can approximate it by a Gaussian. If we guessed, the center is somewhere around here, around 50 nanometers. That's R_p . And ΔR_p is about 20 nanometers. So that's the depth distribution.

The lateral distribution is pretty Gaussian, and it has a certain straggle, what we call ΔR_p perpendicular. That is in perpendicular to the beam direction. So we have two straddles to really think about or two standard deviations. There's the one in the beam direction, so in depth, and then there's the one perpendicular to that sort of in width. That's the lateral straggle.

Now, I just showed that to give you an example, physically get a handle on what's going on at a given spot. But we don't usually just hold the beam on one spot of the wafer. As I said, we usually-- what you do is you randomly scan the beam across the entire wafer. If you want to mask at given region, you implant through a mask into a window. And the surrounding areas on either side of the window will be masked from the implant.

For example, just showing here on the lower right, I've imagined I put down a mask. It could be photoresist. It could be oxide. Somehow, I've patterned the mask, and I've opened up a window. Now, when I do the ion implant, of course I implant across the whole wafer, so I'm implanting across everything here coming in. But the ions stop in the mask. The only ones that make it into the substrate in this particular example are the ones that come through. Only in this window area does the substrate get implanted.

So in this course, we're interested not just in the depth distribution, but in the lateral distribution of the profile. So I care about this window. I'm interested to know not only how deep do the ions go, but how do they spread laterally. If I'm making the source and drain, I care about how they spread laterally into the channel. That's an important thing to note.

So in the lateral direction, of course, measuring the profile is hard. We saw last time two-dimensional profiles are very difficult to measure. But we actually mathematically model them by assuming that the profile is composed of a product of a vertical and a lateral distribution.

So in terms of the x and y, let's say in this particular example that the x is depth. OK, going deep into the wafer, into the crystal. And the y is the lateral dimension. Then we can approximate the concentration of xy distribution as being some vertical distribution, be it Gaussian or whatever we'll talk about, times a lateral Gaussian. You recognize this as a Gaussian function, $e^{-y^2/2\Delta r^2}$ perpendicular squared. OK, so it's a Gaussian with a standard deviation, Δr perpendicular, in the y or the lateral direction.

So I've written over here on the left what we call a point response function. If I just shoot the beam in at one point, you can see what that looks like. And when you go to do this simulation that's being done at SUPREM-IV, what it's doing is it's taking this point response function and then superimposing that response function across the whole window. And then it can give you, then, laterally what the distribution looks like, not only in the lateral direction, but in the depth direction.

In fact, if you do this properly, and SUPREM-IV does it for you, it leads to some kind of lateral distribution under the mask edge. So that's just to remind you-- I haven't told you anything about the physics of stopping-- just to remind you how these profiles vary in x and y. So let's go on to slide 14 and think about what happens at a mask edge.

And this is very important because you're making a MOSFET, because you're defining the source and drain extensions with some kind of mask. Let's assume the mask is thick enough to block the implant, so the mask might be the gate. You have the polysilicon gate of the MOSFET, and you want to implant everywhere around it so it blocks the implant from hitting the substrate.

So the lateral profile under the mask is going to be determined by what we call this ΔR_p perpendicular, this lateral straggle. And this is an example of looking at two different implants. Here on the left, this is a 35 keV implant of arsenic and a 120 keV implant at the edge of a poly gate. So basically, over here on the right, see this, all these different grains? That's the polycrystalline silicon. That's the gate. This is the gate oxide underneath the gate.

And this is the region to the left of the gate, which we often think of as the source. And you can see this contour defines sort of where the-- the edge of where the arsenic ended up. In fact, this is a cross-section TEM, and we showed something like this last time. This was done by Transmission Electron Microscopy with staining of the junction area.

So they actually take the thin TEM sample, they dip it in acid, and the acid preferentially etches and makes very thin the regions that are very heavily doped with arsenic. So this is how you can see, with your eyes or with a microscope, how you can see where the junction ends up. And you can see this profile right here underneath the mask edge or under the gate, this distance from here to here is dominated by this lateral straggle of the implant, depending on how much annealing you do afterwards.

So as we decrease the device geometry, as we make this gate length shorter and shorter, the implant straggle under the mask becomes a more dominant effect on determining the actual channel length. Knowing the lateral straggle is an important feature of MOSFETs. They not only determine the channel length, but remember, we gave examples that the shape of that lateral profile also determines some of the short channel effects, like VT roll-off and things like that.

OK. Again, covering some of the more practical issues before we talk about the fundamental physics, on slide 15, I've mentioned a couple times this idea of masking and implant. OK? So what do I mean about that? And how would you calculate how thick a mask you might need?

So this is an example of masking an implant. I'm doing an implant, and again, the beam is rastered over the entire wafer, but I only want to implant certain regions, maybe only the P-MOS regions or N-MOS regions. So I have to put some photoresist over a certain region. And the question is, how thick does that material have to be? If it's for resist or if it's nitride or silicon dioxide, how thick does it have to be?

Well, it turns out that a dense material, has a high density, can be a lot thinner than a light material, because a light material allows the beam to go in deeper. So here's an example. I'm showing you a plot here, concentration on the vertical axis as a function of depth. And let's say, this is what it looks like. This profile right here is what the ion distribution looks like. And this depth right here marked x sub m , that's the thickness of the mask.

And you can see what I've done here. Actually, I've chosen a mask that was too thin sort of intentionally. What happens is all this distribution got implanted into the mask, and a little bit of its tail right here, which is hatched by this red hatch region, a little bit this tail actually got into the silicon. Generally not desired, but just to show you as an example. So the concentration, c star, at x_m gives you an idea of the concentration that will now be at the surface of the silicon. And this whole dose is what got into the silicon.

So for masking, what we do is we calculate the concentration profile in the mask material. And the reason we have the star here, and I know that we use star in this course to mean a lot of different things, but this star means you're in the mask material. So whatever that material is, you have to use the projected range, R_p star, of boron in that material or the standard deviation, ΔR_p star, of boron in that material in this equation.

So this is just like a regular Gaussian. The only difference I've done is I've got this represented as R_p star and ΔR_p star. So what I want is this concentration, c star, at this point to be less than some background concentration, whatever it might be. Let's say, your wafer is doped 10^{14} per cubic centimeter.

You want to make sure that the concentration profile in the mask is such that by the time that the profile gets to the edge of the mask, it's down below 10^{14} so you really can't see it. That means you have a good mask. You've designed it to be thick enough. If it's thinner, then you get too much of the material, of the species is implanted into the wafer itself. It's not a good mask.

So let's look at slide 16. And in fact, here what we want to calculate is-- we'll calculate the area of this red region. That's the dose that penetrated the mask and made it into the silicon or into the material below. So we can actually calculate the required mask thickness from this equation that I just showed back on slide 15. I'm going to set this concentration equal to the background concentration, whatever it is in the wafer, solve for that, and that gives me the mask thickness, x_m , in terms of R_p star in the material and the standard deviation, ΔR_p .

So basically, it says the mask should be as thick as R_p , obviously, plus some multiple m of the standard deviation. Typically, m might be five standard deviations of thickness or something like that. And so the dose that penetrates the mask, then, I can just substitute that in. All you need to do is integrate this little red region. It's given by this. It turns out to be a complementary error function. So you can actually figure out how much you made in into silicon. Generally, design the mask so that it doesn't make it in.

And there's a calculation on page 457 of your text, the hand calculation of this type of thing to give you an idea. A lot of times, people use hand calculations or they use SUPREM-IV to figure out, OK, if I have a certain thickness, I have a half micron of oxide, how far is boron going to go into that? And is it thick enough for me to mask my implant? So that's masking in terms of the vertical direction.

How about in the lateral direction? This is a little bit more tricky, and you just have to think here intuitively. This is on slide 17. Implants are, by the way, rarely done perfectly vertically. If this is the wafer, it's very rare that the incoming beam is at 90 degrees. And we'll talk in a little bit about why that is. But it turns out a lot of implants, either for device design reasons or whatever, are done at some tilt angle. The beam is like this. You actually tilt the wafer.

And in fact, halo implants are almost always done at some kind of tilt. And the reason you tilt it in is because you want to, in this case-- this is the gate, and I'm trying to do an implant, and I'm trying to get the ions to sneak a little bit underneath the gate a certain distance, and use the implant to get the ions underneath there. So in that case, I'll have to take this tilt into account when I'm calculating the profile.

And in fact, you can get shadowing effects. This is an example of a 50 KB phosphorus implant, and this was tilted at a very high angle, say, 30 degrees. And this red region, which represents the gate, the polysilicon gate, it ends up shadowing the beam. So you end up with a little space here where the shadow was, so that did not get ion implanted. So if I want to reproduce a symmetric profile on the source side and the drain side, I then need to tilt the wafer the other way and come in at 30 degrees to the normal in this direction for symmetry.

So if you're doing tilted implants, you have to think about how your device is oriented with respect to the tilt direction. And maybe do some other implants to symmetrize what's going on. And again, that's a geometrical effect. It's relatively easy to simulate. This is a SUPREM-IV simulation. You just have to keep track of that.

OK, the next concept I want to introduce is thinking a little bit about, how does this Gaussian-- let's say, you implant a Gaussian profile. OK, here on slide 18. How does that profile evolve when I'm going to have to heat the wafer to get rid of the damage? How is it going to evolve? Well, we first do a very, very simple back of the envelope type of calculation.

And let's look at a Gaussian implant distribution. This is what I just-- the equation 1 in what we just put down for a Gaussian implant. It's characterized by this R_p , projected range, and a standard deviation, ΔR_p . Now, if we go back to a prior handout, 14, when we were talking about diffusion of Gaussians, we actually wrote it like this, the concentration can be written in terms of. And look at the denominator here. You got $4 dt$ in it as opposed to $2 \Delta R_p$ squared.

So in fact, these are both Gaussian functions. They are mathematically equivalent, and I can make them equal if I just set ΔR_p equal to the square root of dt . I can equalize these two arguments. So we can then write down an equation, a very simple equation for the impacts. If I do an implant, and you assume it's Gaussian, and then I anneal it in a case where it's non-concentration-dependent, I expect the final profile to remain Gaussian. So a Gaussian remains a Gaussian, preserves its shape upon further annealing.

And in fact, I can just write down by inspection what should happen to that. If you start out with this distribution here-- this is the ion implant distribution-- what I replaced for the argument in the denominator of the exponential, instead of having just this for $4 dt$, I've got a new standard deviation, which is ΔR_p squared, plus that essentially. And again, where we see the square root of dt , you can also write that out here in the pre-exponential in terms of ΔR_p and dt .

For the very simplest case, you have to do a hand calculation. We know how a Gaussian implant profile will evolve. We're going to go, of course, beyond that in the next few lectures, talk in more detail. Here's an example on slide 19 of just those equations. So this is assumed to be as implanted, the red, assuming a very simple Gaussian distribution with a certain standard deviation, ΔR_p .

And after it's gone through diffusion, it's shown here in the blue. The peak has dropped, and now it has a new standard deviation, which is ΔR_p plus the square root of dt . So this gives us an idea of how we can evolve an implanted Gaussian during a diffusion step.

So let's go on to slide 20 and talk about a real profile. Real profiles are not perfectly Gaussian. In some cases, they're more Gaussian than others, but they're generally more complex. If we look at the profile for a light ion, and light means very light compared to silicon, say, boron, these light ions tend to backscatter. Imagine you have a big, heavy bowling ball, and you're throwing a ping pong ball at it. That's the boron. It's a ping pong ball.

Well, a lot of those boron atoms will ping right back, right? A few of them will scatter in forward, but a lot of them can come right back. And so you can see that if you look at a boron distribution, towards the surface, it's actually skewed. The concentration profile is a little higher towards the surface, as if a lot of the boron ions came in and just got pinged right back. So this blue distribution that represents boron at 38 keV, it's skewed towards the surface. And it relates to the fundamental physics of the scattering.

Heavy ions like antimony, they actually tend to scatter deeper. They're heavier than silicon. They can knock the silicon forward, and they themselves retain a lot of forward momentum because they're like a truck. If the truck hits a little Subaru, kind of tends to go right through the Subaru. So the heavy ions like antimony have this skewed-- or skew deeper into the substrate.

The interesting thing about these two protons, so this is antimony at 360 keV. So we adjusted the energy to give them the same projected range, so you notice they peaked at the same point. 38 keV peaks at the same point as 360 keV antimony, but they have very different skewness. A Gaussian distribution by definition is not skewed, right? A Gaussian, by definition, is perfectly symmetric. That's what the equation says.

So if we're going to represent this, we need a little more complicated mathematical function than a Gaussian. And we'll talk a little bit about that type of function. So here's an example just on slide 21 of some different energy boron implants. So I have concentration of boron versus depth. And here's a 50 keV implant. Interestingly, it looks pretty symmetrical. It's fairly Gaussian. 100 keV.

But as I'm going to higher and higher energies, it's becoming less and less symmetrical. So just looking at this black curve right here at 500 keV, look how skewed that is towards the surface. Whereas, if I had used the Gaussian approximation that you can do with a simple hand calculator, this dashed line is what you would have gotten for the profile. So the Gaussian is somewhat accurate. It gives you a rough idea of where the peak is, but it doesn't represent this forward scattering of the skewness.

In fact, these black profiles are called something called Pearson IV distributions. Those are the solid lines in that plot. And Pearson IV is one of the most popular. It's a statistical function. It's one of the most popular types of profiles to use to accurately represent distributions, range distributions for ion-implanted atoms in silicon.

They work really well. Pearson IV works really well as long, as you don't have ion channeling. It does an excellent job of describing the ion implant profiles. As the name suggests, Pearson IV needs four variables. It has four moments, so to speak. Gaussian only has two, so naturally the Gaussian is more simple.

In fact, on slide 22, just reminding you or introducing you to the idea that an arbitrary statistical distribution can be described by a series of moments. And here are the first four moments of an arbitrary distribution. The first two, we've already talked about. The first moment is called the projected range, R_p . And mathematically, you can derive it or calculate it by-- it's the first moment, so it's the integral over all space of x of x times the concentration, c , of x . So you take the weighted integral, and that is the definition of R_p . You may have seen some of these types of concepts in statistical or statistics classes.

The standard deviation is the second moment. It's actually the integral of x minus R_p quantity squared times the constant concentration, that whole thing divided by the integrated dose. Remember, the integrated dose, q . Just remember that q here is equal to the integral of the concentration-- overall space, dx . So you're always normalizing this.

So the first moment and the second moment, they're used, of course, in the Gaussian distribution. The third and fourth moments are used in the Pearson IV. One is called the skewness. Again, that's x minus R_p cubed, the weighted integral of that, divided by q times ΔR_p cubed. And the kurtosis is defined in terms of the fourth-- the x minus R_p to the fourth power. And this skewness and kurtosis, as they've been defined here, they're actually γ and β are dimensionless, because you notice we're dividing by ΔR_p cubed here and ΔR_p to the fourth.

Just so you know, if you run SUPREM-IV, you won't have to do it. But if you were stuck on a desert island, and all you had was an HP calculator, you could actually-- and your life depended on it-- you could actually calculate the Pearson IV distribution from the simple differential equation. It's the solution f . The Pearson IV is a solution f to this differential equation.

So the differential equation says df by dx^n , where x^n is this variable here, x minus R_p , is equal to this function over here, is equal to x^n minus a quantity times f divided by this polynomial. So it turns out this differential equation can be resolved relatively simply using a numerical method if you know these moments, so if these numbers, a , b_0 , b_1 , and b_2 .

And each one of these variables, like a , a is defined in terms of the skewness, the kurtosis, and of course, ΔR_p , the standard deviation. b is related, as shown here. So all these constants are defined in terms of the moments of the Pearson IV. It turns out that if the skewness, γ , for most dopants in silicon, you can actually estimate the kurtosis, β , from the simple little analytical formula. So in fact, they're related this way. So really, for most dopants, there's only three independent moments, R_p , ΔR_p , and β -- and γ , because you can calculate β from γ .

If you just look at this equation just for a moment-- you don't to solve it-- but just looking at it, what are some properties of this function? Well, the peak, where does the peak occur? Well, the peak occurs wherever a function is set equal to 0. So the peak occurs where this-- or the derivative of the function is 0. So it actually occurs at x minus a , which is R_p plus a . So it doesn't occur directly at R_p . Unlike the Gaussian, the peak of this occurs just slightly to the left or the right of R_p .

So when the skewness, γ , is 0-- so when γ is 0, f simplifies to a Gaussian. So if I just take this and a goes to 0, then I'm going to get a simple Gaussian. There's no skewness. If the skewness is negative, so γ is negative, like in the case of boron, a is going to be positive. So if γ is negative, a positive quantity, and the peak of the distribution occurs deeper than R_p .

So for boron, if you're careful about calculating it in Pearson IV, the peak will just be slightly deeper than R_p . It'll be R_p plus a . When the skewness is positive, like for arsenic, a heavy atom like arsenic or antimony, the peak is going to be slightly shallower than R_p . So again, slightly more sophisticated than a Gaussian is the Pearson IV. And that's how it's actually calculated by the computer.

So these four moments that we've been talking about, they've actually been tabulated. I took this particular table I copied out of a 1980 book edited by Gibbons on ion implantation. And it's just showing as a function of energy and a couple different energies, 20 50, 100, and 200 keV. For three common dopants, it's showing R_p , σ_p -- now, σ_p is the same as ΔR_p . It's just a different notation-- the skewness, γ , and σ_p perpendicular, or ΔR_p perpendicular.

So just looking at boron, again, its skewness is negative, right? All of these γ numbers are negative. That's why it's skewed towards the surface. And so you can look up in tables like this, and if you have your calculator, you can do the Pearson IV equation, and you can calculate what the profile should look like. OK, so that gives you some idea of how engineers use these numbers.

Now, let's go on to slide 25, and let's talk about implants in real silicon. Well, to be a little more careful, real silicon is single crystal, right? We know it has the diamond cubic structure, and it has a regularity to its structure. And in fact, I'm hoping not to make anybody car sick here. But be careful, don't shake your head too fast. But this is a picture of looking down the 1,1,0 direction, the axial channels in crystalline silicon. What do you see in this picture?

What you see, well, you see a very regular pattern of atoms, and you see big openings, wide open spaces where there aren't-- where the atoms at the surface block the atoms below. So you can imagine being an ion going down this crystal. And if you're aligned, if your energy or your momentum is aligned just in the right direction, you can channel right down some of those channels. And you can imagine going down there without losing much energy, without banging into a lot of the silicon atoms because of this regular crystal structure.

And in fact, this is a two-dimensional representation, a cartoon up here of what I'm talking about. Imagine these red atoms are silicon. They're all lined up in a row, and my incoming ion comes in just pointing down that channel. It can just barely skim the edges of the channel, have very few small angle collisions, and be steered a long way before it stops.

That's very, very different than going into a random material where there are silicon atoms all over this page. So ion channeling, because the material is single crystal, is really going to complicate all those equations I just talked about, the Gaussian, Pearson IV. They're all assuming that the atoms of the crystal were randomly oriented. They didn't say anything about going down a channel. The physics of channeling is quite a bit different from what we've just been talking about.

So channeling is good and bad. Mostly, it's bad because it can produce some really unexpected results. As you can imagine, if this is your crystal, and you're looking down the 1,0,0 channel axis-- so these are all the open spaces. If you happen to be aligned, the beam, right to that axis, you can get profiles that go a lot deeper than you would have expected. So this is an example of implanting phosphorus at 40 keV into silicon. And here's the concentration of phosphorus as a function of depth.

And in fact, looking at different doses, the very low dose here is about 1×10^{13} . Look how deep the phosphorus went, all the way down to-- I don't know. It's got the shoulder here. Goes down to about 0.6 or 0.8 microns. As we go higher in dose, notice this. The very highest dose 7×10^{14} , you don't see that channeling tail. Why would that be? Does anybody have any idea why this very high dose implant didn't go in real deeply, whereas the low dose one did?

[INAUDIBLE]

OK. So you have a high dose is coming in. So what is that dose going to do to that beautiful crystal, all those beautiful channels? Smash them up. Basically, just imagine taking that, just throwing a bunch of baseballs at that thing. And pretty soon after you've thrown enough baseballs, that perfectly nice-looking channel, there's a whole bunch of atoms in it because you completely randomized the crystal. And that's called amorphization.

So when you put in enough dose, you throw enough baseballs at this thing that you randomly-- now, instead of having a beautiful crystal structure, the silicon is completely amorphous. There is no structure left. Then you are going in, and now this implant was essentially done, most of it, into a material that didn't look like this. It just looks like an amorphous mess. And that's why you don't get this channeling tail.

This implant was such a low dose, it barely did damage the crystal at all. The crystal maintained its nice structure, and so the ions were able to channel in and go very deeply. So this is bad because it turns out it's very hard to control what this tail looks like. And a lot of times, you don't want such a deep profile. So in fact, in this case, channeling is viewed as something one wants to avoid.

Actually on page 28, here's another example, just as we just saw, of boron. Boron is a famous ion for doing a lot of channeling. It tends to be steered very easily. One problem with ion channeling is that the profile shapes don't scale with those because, again, a low dose doesn't damage the profile much. A high dose here, like 2×10^{15} , does a lot of damage.

And so this green profile is not simply equal to the 2×10^{13} profile, bumped up by 2 orders of magnitude. That's what you think it would be. But in fact, it's got a different shape because as you were doing the implant, you're doing a certain amount of damage, and you're changing the appearance of the crystal and channels. So there's all kinds of issues when you're trying to model ion channeling.

Page 28 actually has a little bit more information about ion channeling. There is something in ion channeling called the critical angle. What is the critical angle, ψ ? Well, it's the maximum angle. It's the largest angle of the incoming beam with respect to the axis of the channel before the steering action of the row is lost. So at this angle, you can just barely come in and be steered and go in and out and in and out.

And if you have a slightly higher angle, what's going to happen is you're going to get knocked out of the channel. So that critical angle is important to know. It tells you how well does the beam have to be aligned to the channel to be sure that you get this channeling effect. And ψ , see here what's called the critical angle, in fact the reference above has a little bit of information on how to calculate it. And these are some examples of critical angles.

Say, along the 1,1,0 direction for boron is 3 degrees. Arsenic is about 4.4. So it gives you an idea, if this is my crystal and I'm coming in, well, if I'm within 3 or 4 degrees of that axis, I'm going to get a lot of channeling. If I tilt the crystal, so I'm outside 3 or 4 degrees-- say, I'm at 7 degree tilt-- most of the ions will come in, and they won't actually be aligned with the channel. So you won't get much channeling. So it kind of gives you an idea of how much accuracy, when you're shooting your ions in, do you have to have to enable channeling.

As we see here on slide 20, people do something called controlled misalignment of the crystal when they put the wafers in the ion implanter. And they don't put them in like this. They try not to put them in so you're looking straight down the 1,0,0 axis. What they do is, they take the wafer and they put it on the plate, and they tilt the plate a little bit with respect to the beam. In fact, typical tilts are about 7 degrees.

And they actually also rotate it. So if I've just taken this crystal here, and in the computer, they tilted the crystal and rotate it. And look, lo and behold, a perfectly crystalline material all of a sudden looks somewhat random to the eye. And so it looks somewhat random to the incoming beam. You see the channels. A lot of them have now been blocked by the tilting and rotation. And when it looks random, that's good. That means you're not going to have a lot of channeling.

You can never get rid of it all together. You can minimize it. Because the problem is, as I come in here, it looks pretty random, but then I might get knocked. Maybe I get knocked by 7 degrees and rotated some of the ions, and they get knocked back into a channel. So you can scatter incident ions back into channeling directions in their first few collisions. So you can't get rid of it, but you can minimize it.

And ions that are not channeled early in the process, as they get closer to their end of range, as they lose energy, the critical angle increases. In fact, I didn't mention that. If I go back one page to slide 28, this was [? Linhart's ?] expression for the critical angle. It depends on some constant, e_1 , divided by e . So this ψ goes up.

As I get slower and slower, then the critical angle goes up, so the probability of channeling is going to be bigger at the tail of a profile. In fact, you see that on slide 30. That's very evident when we actually look at a profile that was ion implanted, where we tried to eliminate channeling by tilting the profile, tilting the wafer by 7 degrees off the 1,1,1 axis. So this is concentration versus depth, and there's a couple of different profiles here.

In fact, what people were trying to do was, they were also trying to cover the wafer with an amorphous material. People thought, oh, well, if I-- why don't I put a layer of something amorphous? That will prevent channeling, right? Because I'm covering up the crystal with these amorphous material. But in fact, it doesn't really avoid channeling. Putting on a little bit of nitride doesn't really help.

You still get this so-called exponential channeling tail. You see, all four profiles have this exponential tail. That's due to channeling. As the ions get to their end of range, these ions down here have lost a lot of their energy. When they get to this point, the critical angle for channeling goes up because it's much more likely they get steered into a channel. And oops, they go a little further. So this tail is called residual channeling. Even an amorphous overlayer will not eliminate it, and tilting the sample will not eliminate. So this tail does need to be modeled if you want to get a very accurate simulation of what your profile really looks like.

So let me just summarize a little bit about-- so far, I haven't talked about the physics of scattering or of stopping or whatever. But we've just talked about how one analytically might treat ion implant profiles. We have a Gaussian distribution, which has two coefficients. It does an OK job. It's not very good, but it's good. It's OK. Pearson IV has four coefficients. It's very good for amorphous, but it doesn't include channeling.

If I take Pearson IV, and I add an exponential tail, so now I need six coefficients, I can include channeling. And it works in some cases. There's something called a dual Pearson, which we'll talk about. That's two Pearson profiles superimposed, has nine coefficients. So look, as I'm going from top to bottom, I'm adding more and more free parameters. Of course, you can model anything. You can model an elephant with enough free parameters.

Legendre polynomials have 19 coefficients. That's getting kind of ridiculous. But there are tables of a lot of these coefficients that are built into textbooks or into SUPREM-IV. And they have a lot of these coefficients as a function of mass, energy, dose, tilt angle, all that.

These analytic calculations are very fast. They're very efficient. But you have to be very careful when you're using them because the coefficients for dual Pearson or coefficients for this Legendre polynomials, whatever it is, a lot of times people are just fitting experimental data, which may have artifacts of its own. Remember, when we talked about [? sims, ?] we talked about a knock-on process.

So [? sims ?] itself will tend to produce tails, just from the secondary ion mass spectrometers, from the analytical technique I used. So you have to be sure that somebody wasn't actually modeling an experimental artifact. Anything based on tables, you have to take with a grain of salt and just be a little careful. Check it out.

Slide 32 actually just shows you some SUPREM-IV, four different models of varying complexity for modeling a particular implant. These are concentration contours. Each color or each grayscale represents a different concentration of the dopant. This is a polysilicon mask edge, so none of the dopant can make it into the mask over here. And it only is implanted on the right-hand side. So it's particularly tricky to get the right profile.

So this particular model, upper left, is assuming you're going into amorphous silicon, so it has a very simple shape. This is assuming the dual Pearson type of model, so it's going to have some skewness. This particular profile is dual Pearson, where there are extra coefficients that take into account tilt of the wafer and rotation. You can see it looks quite different from that.

All three of these here in the upper left corner, these three are what we call analytic. They're based on these equations that we've been talking about. Monte Carlo model is quite different, and we'll talk about that next time. Monte Carlo, literally in the computer, you shoot ions into the crystal, and you follow the scattering processes. And you will see where all those ions end up. A computer can do that. It can sit there all day, and it can calculate the trajectories for 10,000 ions and then plot their statistical distribution.

So Monte Carlo tends to be quite accurate if you have enough physics in. Unfortunately, it takes a long time because the computer has to go in there and follow every ion, see where it ended up. Here's an example, actually, a few words about Monte Carlo. What you do is, you integrate the equation of motion for a representative ion from the point when it hits the surface of the silicon until the point where it comes to a rest, where it stops.

And you repeat this process. You follow the ions for as many times as you can, using random starting impact parameters. And there is this so-called law of large numbers that says, the error, the scatter in what you're going to see, goes like $1/\sqrt{n}$. So if you want three decades of a profile to be accurate, then you need-- with less than 10% error-- you need 10 to the 5th ions. So 10,000 ions, or 100,000 ions, rather, is about the minimum that you need to get a nice smooth profile.

So that's going to take a while. You can try on some of your homework, so you may end up trying a Monte Carlo simulation. You could see. It could take half an hour, it could take overnight simulating one implant, depending on how many ions you need to follow. This is an example of some Monte Carlo simulations that were published back in 1997 for-- here, they're doing relatively low-energy boron. And get that here, 0.5 keV, 1 keV.

And you can see the statistical nature of these profiles. The simulations are these little boxed sort of digitized looking things because they're actually following individual ions and trying to build up a distribution based on that. So the nice thing is you can get reasonable accuracy. The bad thing is it could take you days if you're trying to fit a profile, literally days to do that.

And before we finish up, I wanted to mention there's one other type of calculation that's not very popular but was developed a while back and actually is kind of a compromise between Monte Carlo and analytic. It's something called the Boltzmann Transport Equation, or the BTE, developed by Giles. And he was able to write down a Boltzmann transport equation and then model an ion distribution. Instead of following individual ions, he was following the distribution, the momentum and the energy distribution of these ions.

So there's economy of scale in doing that. And he actually followed the distribution as it changed and as it evolves in the computer. I won't go into the detail, but basically by following distributions instead of following ions, it's much faster than Monte Carlo. But it can have as good accuracy. Unfortunately, it never became very popular, so I don't think it's in SUPREM-IV. There may be some versions of simulators out there out in the world, though, you may come upon that have BTE method for ion implant. I think it's actually a very good method. It's just too bad it didn't catch on in a big way.

So on slide 35, let me summarize. We saw the ion implant is the preferred process for introducing dopants. It has excellent control. It's very reproducible and uniform. The key parameters, energy of the incoming beam, the dose, the tilt of the wafer with respect to the beam, the rotation of the axis of the wafer, whether there are any masking materials on the wafer.

Controlling the temperature of the wafer during the implant is actually somewhat important because it'll control the damage to a certain extent that gets done to the crystal. The dose rate, the beam current, how quickly are you putting it in. Actually, the dose rate determines how much you heat the wafer up, and that's important. So in practice, there are a number of key parameters.

There are tables for the first three moments, R_p , ΔR_p , and skewness, that you can get out of the literature. The most popular analytic equation to use or formulation is the Pearson IV for nonchanneling. If you want to model channeling, you need more parameters than four. Usually, people use dual Pearson and lookup tables. Monte Carlo is slow, but it's probably the most accurate or one of the most accurate. The BTE is another approach you may find in the literature. It's an alternative to Monte Carlo, but it's faster, but not that widely used.

What I didn't get to cover today, we'll spend the lecture next time on the detailed physics of these billiard ball collisions, the physics and modeling of this nuclear stopping, these collisions and this electronic drag force. So we can actually see in more detail how the physics of the implant. Today, I just mostly talked about practical aspects.

OK, if you came in late, a couple reminders. Pick up this handout number 20. This describes your final term projects. And next time, I'm going to bring a sign-up sheet. You can sign up for the topic you want to use or study, and most importantly, whether you want to do an oral report or a written report. And next time, I'll have the homework.