The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

MIKE SHORT: All right. So I am super excited about today because this is, in my opinion, the highest point of the apex of the course where we're going to put together everything you've done so far and start to explain things like Bremsstrahlung, radiation damage, X-ray spectra that you get in a scanning electron microscope, and actually find a way where cross-sections are areas.

Remember before, I told you guys cross-section is measured in barns, in centimeters squared? Think of it kind of like an area. We'll actually be able to show mathematically that some of them do derive from actual areas. So it will make a lot more sense. It's not just an abstract concept.

What I want to do is a quick review of the ionization and excitation collisions that we did last time. Remember we had this imaginary hollow cylinder where we said that there is some ion traveling in this direction with charge, let's say, z times the electron charge and it's colliding, kind of, with an electron somewhere else separated by some impact parameter b. And let's say this hollow cylinder head a shell of thickness db.

And we started off with this situation where we wanted to say-- we can find the y momentum as the integral of the y force. We did something. And one of the intermediate steps we came up with was that our energy in part of the electron is p squared over 2 times mass of the electron, which came out to 2z squared e to the 4th over the mass of the electron impact parameter squared velocity squared.

Then we multiplied by the electron density in the material, which was the number density of atoms times z, the number of electrons per atom, times the area, the cross-sectional area of this hollow cylindrical shell, which came out to 2 pi b, which is the circumference of that circle right there, times db dx. I'm just going to leave that there for now, and we'll come back to it in a second.

I will mention, though, that at the end, we came up with our stopping power expression. Let's call this ionizations that came out to 4 pi, this constant k0 squared, which comes from the Coulomb force law, little z squared, big Z, e to the 4th over mass of the electron velocity squared-- I'm running out of room-- times log of mass of the electron v squared over this mean excitation energy, which comes out to around, let's say, 10 to 19 electron volts times z.

And what this meant was that if we graph stopping power as a function of energy, there's a couple components to it. One of them is this roughly 1 over energy component. So let's say there's a component of it-- yes, I love when I can do that-- that actually-- that follows 1 over e. And there's this logarithmic component that goes that way.

And if we sum them together, we ended up with a much higher stopping power at low energies. But around this local minimum of about 3 times the mass of the electron c squared, it starts to increase again. And that's due to this. How many excitations can you make as a function of energy? So if you think about it, this is kind of like an energy term. How much energy do you have divided by how many ionizations can you make? So it's kind of an intuitive result in that it's how much energy you've got versus how much energy it takes for a single unit process. And that's all mediated by this 1 over e term.

The last thing that we said is that this curves back down here. And that's because at around 500 times this mean excitation energy, which as you can see is around 1 keV times z-- it's not a very high energy-- you start to get charged neutralization. The reason for that is that for-- in order for this formula to work, we have to assume that the deflection is really, really, really small, infinitesimally small. If it's not, and if the ion manages to catch that electron, it's not going to lose energy by undergoing a Coulomb interaction. It's going to lose energy by absorbing that electron and neutralizing.

And so for really low energy, the ions are moving so slowly that they can capture some of those electrons. And it gets less and less effective at stopping the material. And this led to-- let's see-- this, we're distance, and this, we're dt dx times the energy required to make an ion pair, which we'll just call this i. We ended up with a curve that looked something like this. This right here we call the range. We just gave that the symbol r.

So that's the review of last time, but I left some extra space for some reasons which will become clear in about half an hour. Right now I want to take you through a little bit of a whirlwind in terms of Bremsstrahlung. I'm not going to derive anything about it. We're just going to go over what the cross-sections and stopping powers look like, why they take that form intuitively. But we're not going to go through a rigorous derivation because, well, we simply don't have time. But I do want you to know what sort of things exist.

So for Bremsstrahlung, better known as braking radiation-- who's actually heard of this before? Does anyone know roughly-- let's say if we were to say some cross-section for Bremsstrahlung. Let's call it cross-section radiative. Because in this case, we're actually talking about the charged particle radiating away photons.

So let's say we had some nucleus of charge big Z, and we had some ion-- we'll look at a different color-- some ion of charge little Z moving towards it. It will either be attracted or repelled depending on what the sign looks like. And sometimes you might get radiation of a photon. Let's call it of energy E equals hc over lambda. That'll become important in a sec.

So this radiative cross-section-- I don't care about-- particularly care about the exact form, but what is it proportional to? What sort of factors do you think would make the emission of that photon more or less likely? Yeah?

AUDIENCE: The energy of the charged particle.

MIKE SHORT: Sure, the energy of the charged particle. So you can-- there's one expression I like you. Can only take my money until you take it all. Well, the same thing goes for the energy of a photon.

You can't radiate any more than the energy of the particle coming in, so there's going to be some maximum energy which is going to correspond to some minimum wavelength, which is-- let's say if we just put our initial energy in here, that gives us our minimum wavelength right there. That's true.

And it can radiate at any energy smaller than that or any wavelength larger than that. I'm basically saying the same thing. Because if this particle started off further away and felt less of a deflection, it could still emit a photon, but of a longer wavelength. Don't have room to draw enough wavelengths, but I want to make sure that's physically accurate.

So it's actually going to be also proportional to z of the large nucleus. The stronger the pull, the more of that breaking radiation you're going to see. It's actually proportional to z squared. What this says is that heavier z materials produce a lot more of this breaking radiation. And it's also proportional to the inverse of the mass squared. What this says intuitively is a heavier particle deflects less and emits less of this breaking radiation.

Hopefully this makes a lot of sense to you guys, that the stronger the pull of the nucleus, the more deflection you're going to get, and the more Bremsstrahlung you'll get. The larger the mass of the incoming ion-- let's say that's mass m-- the less deflection you'll get because the less-- what is it? The less momentum transfer you can apply with the same force if you've got a heavier particle. Does this make sense to everybody?

So what this says is it's really important for high z materials. And we'll see a little bit why. If I actually write the full expression for stopping power-- and I promise to you I'm not going to derive it because we don't really have the time for that. It's proportional to the number density. You should always think there will be a number density and stopping power, because the more atoms there are, the more they stop things. It's just directly proportional. Times that kinetic energy plus mec squared. And again, this is not something I want you to memorize, but it is something I want you to be able to decompose and explain why the parts are there.

Let's see. Times some radiative cross-section where this sigma radiative is some constant cross-section. This ends up being about 1/500 barns times z squared times this parameter b, which if you see in the reading is actually given just-- this b scales roughly with the atomic-- I'm sorry-- yeah, with the proton number of the material. So you can see that in here, in the stopping power is actually directly the cross-section. So the components of a stopping power-- there's going to be some probability of interaction, and there's going to be some energy transfer part.

This is an interesting result to show. This is why I wanted to just write the Bremsstrahlung stopping power. Because in here, actually, is the cross-section times some other stuff. Pretty neat result. And so now I want to show you how our cross-section is actually contained in the ionization stopping power.

So let's bring this back down for a sec. You can think of the likelihood that the ion comes off with any particular energy to be directly related to this impact parameter. Because as we saw, the final expression for stopping power is directly related to this parameter b. We ended up integrating over all possible b to get the total stopping power in the material.

But if we didn't do that integral-- if we stopped, let's say, at this stage in the game, and we said, all right, well, the stopping power at some fixed b actually depends on the probability that that particle enters into this cross-sectional area right here, this 2 pi b db, that right there is the cross-section for scattering as a function of the incoming energy and the outgoing energy.

This is one of the coolest parts, I think, is that there is an actual area right here. It's the area of a hollow circle is the actual cross-section for scattering with a given ingoing and outgoing, outcoming energy. And then the rest of this stuff-- if you pull all this together, if you take a microscopic cross-section times the number of particles that are there-- because this is your atomic number density, and these two together are your electron number density. Like we talked about before with reaction rates and cross-sections, this thing right here is your macroscopic cross-section for an incoming and an outgoing energy contained directly in the stopping power formula, because then we integrated over all possible cross sections, which means all possible outgoing energies for a given incoming energy.

And so the last bit, the way to link these two together, which is why I left a little bit of space right here-- we know right now, because I wrote it up there, that the scattering cross-section as a function of the ingoing an outgoing energy per unit energy is just the area of that hollow circle. So let's divide everything by dt. We end up with the total formula for cross-- for the scattering cross-section.

We don't know what this b db dt is. What's the differential probability between impact parameter and outgoing energy? We don't quite know, but we can express this as a change of variables that we do know. We do know-- if we have a certain impact parameter, we know what the scattering angle is going to be. There is a well-known relation for that. And there's a derivation in the book and in another book that I want to point out to you guys by Gary Was called *Fundamentals of Radiation Material Science* on page 32. If anyone wants to see the derivation from which this result came from, you can head right there, and it's free on MIT Libraries.

Meanwhile, we do know our relation between this impact parameter and the angle, and we do know a relation between this angle and the outgoing energy. And this is where some of the hard sphere collision stuff comes in. What's the maximum amount of energy that a particle can impart to another particle in some sort of a hard, sphere-like collision?

Let's take the easy example. If the two particles have equal mass, how much energy can one particle impart to another?

AUDIENCE: All of it.

MIKE SHORT: All of it, right? It can impart a maximum of, let's say, this incoming energy ei. As those mass ratios change, you can impart a maximum-- let's say your maximum becomes what's called gamma ei where this gamma right here is 4 times those two masses multiplied over the sum squared. The full expression-- I'm going to use a different color because I'm running out of space here. The full expression for t is actually gamma ei over 2 times 1 minus cosine theta.

The two intuitive limits from this is if theta equals pi, then this here equals 2, and t is our t maximum, just gamma times ei. If theta equals 0, that whole thing equals 0. And in our case of forward scattering, no interaction occurs, and the energy imparted is 0. But the important part here is we have a direct relation between t and theta and the angle, which we can put in here. And we actually have a direct relation between b and theta, which I wrote down so I wouldn't forget.

So we actually have-- our impact parameter is classical radius of the electron times cosine of angle over 2. You don't have to know where these came from, but the point is we have a relation between B and the angle. We have a relation between angle and energy. So we can just do a change of variables to get our final cross-section. And it ends up being, I think, pi times radius of the electron squared over gamma.

So in this way, we can go from known relations between each of the variables and an actual physical cross-section that has units of real area down to an energy-dependent form for this cross-section, which I think is pretty cool. And then this is the one that you would see tabulated in the JANIS tables, like the energy-dependent cross-section for the S reaction or the scattering. So this is one of my favorite parts of this course, because you can see how cross-sections really do follow directly from areas.

So now for the other part, now that we've got this Bremsstrahlung stopping power and we've got our ionization stopping power, it's useful to find out which one is more important when. To do that, we just look at their ratios. So if we look at the dt dx from ionizations over the dx dx from radiative energy transfer or Bremsstrahlung-- let me make sure I get this one right. It's proportional to z times mass of the electron over m squared times t over 1,400 rest mass of the electron.

So what this tells you here is that-- I'm sorry. I think I have those backwards, because radiative should get more important at higher energies. So what this tells you is for higher z materials, Bremsstrahlung becomes more dominant, and for higher energies, Bremsstrahlung becomes more dominant. So if we want to generalize our stopping power curve from just ionization to everything-- so I know I had another color. No. I ran out of colors. I need a fourth one.

There's going to be some Bremsstrahlung component that starts to get more and more important with increasing energy. And so then if you extend this curve, you're going to radiate more and more and more power the higher energy you go-- not from this component or that component of ionization, but from radiation or Bremsstrahlung. So this has some pretty serious implications to answer questions like, how do you shield beta rays?

Does anyone have any idea? Based on this formula right here, what would you use to shield beta particles and not irradiate the person standing behind the shield? Let's ask a question everyone knows. What do you use to shield photons really well? Lead, tungsten, something with high z.

Because as we saw from before-- I'm going to steal a little bit of the Rutherford stuff. If you graph the energy versus the mass attenuation coefficient, you get a curve that looks like this, but everything increases with increasing z. You get more mass attenuation with increasing z. And also, denser materials tend to be higher z. Is that what you want to do for beta particles?

You say-- so Monica, you're saying no. How come?

AUDIENCE: Don't you just want something with a low cross-section?

MIKE SHORT: That's right. Well, you don't necessarily want something with a low cross-section, or else it might not shield at all, but you are on the right track. You can actually look at the difference between these stopping powers, and cross-sections are embedded in there. So I think that answer is pretty much correct.

But also, you're going to get more Bremsstrahlung or more breaking radiation in higher z materials. So if we actually look at what thresholds does this become important in lead, this ratio is about 1 at around 10 MeV, which means that you lose an equal amount of energy to Bremsstrahlung as ionization at 10 MeV for electrons.

In water, this ratio is about 1 at 100 MeV. So what this says is if you want to shield electrons or beta particles safely, you actually have to use lower Z materials because they won't make much Bremsstrahlung. But because, like Monica said, then the cross-section is lower, you actually have to use more. So you don't have a choice. You can't just use less high z material. Because while you will stop more of the electrons, they will create more x-rays in the process. And those x-rays are highly penetrating, as we know from these mass attenuation curves.

Once you get to high energy, this is-- these are logarithmic scales, so let me correct those and say these are log of e and log of mu over p. It gets millions of times less effective at shielding high energy photons. So that's one of those really important things to note is if you're designing shielding for something, and there are electrons involved that are even around 1 MeV or so, you can't just use high z materials to shield them, or you will create more problems than you solve.

That's a pretty important implication. It's quite important for what's called betavoltaic devices. It's kind of a sidetrack, so I'm going to stick it on a board that'll be hidden soon. Has anyone heard of a betavoltaic device? Anyone? What are they?

AUDIENCE: It's like a beta source that emits electrons onto a semiconductor [INAUDIBLE].

MIKE SHORT: Yeah, it's a beta battery. All it is is, let's say, some pieces of silicon, some circuit that grabs the power, and a beta emitter. And these beta particles directly hit the silicon, and the movement of those betas constitutes a charge. And it's direct-- it's direct conversion of radiation to electrical energy. They're not very high power, but they last for a very long time.

How long? Around a few half lives of that beta decay. So for most of these beta emitters that have half lives in the realm of, like, 10 to 1,000 years, you can make a microwatt battery that could last for millennia. This could be pretty useful. Let's say if you wanted to have some secret sensors in a naughty country like North Korea, you could drop these tiny little beta particles that would just-- betavoltaics that would just trickle charge a battery, make a measurement of-- I don't know-- radiation level, or weight of the dictator, or whatever you happen to want to measure, and send that off once a month or once a year with no need for external monitoring.

Or let's say you're designing a mission to land on a comet, like the Rosetta Philae Lander, and your radiothermal isotope generator is going to burn out in, let's say, 10 or 20 years. You might not need that much power just to measure temperature, or light levels, or something else, or a gas that you might want to know what's there. But you have to choose your beta isotope wisely.

If you want to make these things in a little chip-- and they actually have been commercialized in a chip that's about that actual size using about two curies of tritium. Anyone have any idea why one would choose tritium?

AUDIENCE: It's got a short half-life.

MIKE SHORT: Yeah, it's got a short half-life, so you can get a lot of power out of it. That's one of the two correct reasons. And what is the other one? Lets see who's memorized there KAERI table of nuclides. What do you think its beta decay energy would have to be for this not to blast anyone in the vicinity?

AUDIENCE: Low.

MIKE SHORT: Very low. Why do you say that?

AUDIENCE: [INAUDIBLE] they don't penetrate all the way through the [INAUDIBLE].

MIKE SHORT: That's true. Their range is much smaller. But the range of all betas is pretty low in materials. But the answer lies right here-- less Bremsstrahlung. Lower energy betas give most of their energy off in ionization rather than by radiating Bremsstrahlung.

So you can have a device with two curies of tritium, which if that's released to the outside world, that's bad news. That's something that you might have to report. But as long as it stays contained in this device, it does not have enough energy to produce many x-rays from Bremsstrahlung. And therefore, it does not require an enormous amount of shielding.

So you can't just pick a 1 MeV beta emitter which you might get a lot of power out of, because it's also going to be a big, crazy X-ray source that you wouldn't want in a cell phone or a sensor or some other device you might put in your pocket, or even 20 feet from you. Cool.

So that's the idea behind Bremsstrahlung. There's a little bit more I want to tell you about, and I'll save that for the sidetrack board. We use Bremsstrahlung in a lot of really interesting applications, including cyclotron, one of which we just took delivery of here at MIT, or a synchrotron. And I'll just briefly explain how these work.

In a cyclotron, you've got two D-shaped magnets. They actually call them dees because we're so creative in naming these things. You inject some source of charged particles, and there is some electric field lines across these two dee magnets. And what this says is that in between the magnets, the particle accelerates.

And inside each magnet, the path curves. And it accelerates some more. And it's moving even faster, so it takes longer to curve. Than it moves even faster, and it takes longer to curve, and so on and so on, until it finally shoots out the side.

And so this is one way that you can have an extremely compact-- and I'm talking like garbage-can-sized-- accelerator that brings things up to about 13 MeV. That's the one that we've got in the basement of Northwest 13. The problem is every time these particles bend, they send off photons, what's known as cyclotron radiation. And the higher energy that is, the more intense that cyclotron radiation gets.

So you've got this garbage-can-sized device with a little hole right here, and it's just blasting out photons in all directions in this one plane-- let's just call it the plane of death-- which you don't want to be in, which is why this the thing is behind 4 feet of concrete shielding, and in the middle of a room, to help-- that 1 over r squared keeps your dose down. But we actually use this plane of death in a synchrotron.

What it is is it's a circular accelerator. It's not quite circular, so let me correct my drawing a little bit. There are straight segments, and there are slightly curved segments. But it pretty much looks like a circle if you look at it from high up enough. In each of these curved segments, there is a bending magnet. That's my best drawing for a magnet.

And what this does is it continuously changes the path of these charged particles going through usually electrons. And you end up with intense beams. Let me use a different color. You end up with intense beams perpendicular to the original path before it went in that bending magnet of synchrotron radiation.

So it's kind of like a gigaelectron volt spinning ninja star of death, except at the end of every one of these stations, you have what's called a beam line. Because there's 60 or 80-odd of these beam lines coming off with, let's say, 80 kv and below Bremsstrahlung x-rays, you can use those for a whole lot of different analysis techniques. You can simply irradiate things. You can send those x-rays through a monochromator to select only one wavelength, and then use that wavelength to probe the structure of matter down to the atomic level.

There's actually one of these just down in Long Island. About a 2-and-1/2-hour drive from here, there's Brookhaven National Lab. And they just opened up the National Synchrotron Light Source, or NSLS version 2, where they can actually measure distances with single nanometer precision. So inside this beam line is a bigger room which is encased in another room which is encased in another room. And the whole point of that is for vibration and temperature isolation.

So they maintain this entire room to within a speck of 0.1 Celsius. And it's the least vibrating place, probably, in the US. I don't know about on the planet. But it's got basically no vibration. So the atoms are effectively standing still except for their normal vibrations in the material. But there's no source of external vibration.

And the cooling has to come in through these convoluted channels so as not to blow on the sample, so as not to make any convection currents or temperature changes. And they can actually probe the structure of matter with single-nanometer precision using these synchrotron x-rays all produced by Bremsstrahlung. So it's not all bad. You can use Bremsstrahlung for good.

Then there's a little bit-- I have to hijack a little more area from Rutherford scattering. You might think about, well, what is the actual spectrum of this Bremsstrahlung. Well, you can look to see what's the probability that an atom enters into any of these concentric, hollow circles. It looks to be less and less likely that you're going to enter through one of the center rings and more and more likely that you're going to enter through one of the outer rings.

If you start farther away, there's less of a pull to change the path of that ion or electron, and the Bremsstrahlung is going to be lower in energy. This is actually described by what's called Cramer's law, which says that the intensity of the Bremsstrahlung as a function of wavelength scales with some constant k, and that constant scales with-- surprise, surprise-- the atomic number of the material times some lambda over lambda minimum minus 1 times 1 over lambda squared.

And what this says is that there's some minimum lambda or some maximum energy that you can impart to this Bremsstrahlung, which again, you can only take some energy before you take it all. And there's going to be some sort of a fixed minimum lambda if we draw this intensity. And I graphed this on Desmos just before coming here, so I know it looks something like this where that right there is lambda minimum.

It's taking more area. If you then change variables from lambda to the angular frequency where, if you remember, the energy of the photon is just h bar times that frequency-- so it's kind of like converting into energy with just a tiny, little constant in front. And I mean really, really tiny little constant.

You end up with an energy relation that looks like some maximum angular frequency or some maximum energy. And this is kind of a simple, linear-looking relation, this 1 over energy relation minus 1. So if we graph energy versus the intensity of the Bremsstrahlung, you end up with a curve something like this where your max energy is the same as your incoming particle energy.

Now, who here has done any sort of X-ray or SCM analysis before? You have. So can you tell me, is this the Bremsstrahlung spectrum that you tend to see?

AUDIENCE: Well, I've done [INAUDIBLE] analysis with imaging.

MIKE SHORT: OK. Have you ever gotten a regular, old X-ray spectrum to see what elements are there? Can you draw what one looks like?

AUDIENCE: Maybe.

MIKE SHORT: You want to try? They're all the same. So if you remember any particular one, you're correct. Yep. There's some peaks. And then what does this background stuff look like? Yeah. There's some noise and junk on the back of it, right? So this is actually correct. Thank you.

And what you actually see here is a bunch of characteristic peaks. These will maybe be like the L lines and the K lines for one element or another, these characteristic X-ray peaks, on top of the Bremsstrahlung, the breaking radiation which constitutes the background here. And what you actually see-- I'm just going to draw the background curve under Julia's curve here-- looks something like this. What happened to the real spectrum? Why don't we observe what actually exists? There are a couple of reasons. Does anybody have an idea?

So let's take this to the extreme. Why don't you think you would observe physically-- and this is when we actually get into the real world-- any x-rays with energy in, let's say, the eV range if you were to try and observe any x-rays at all? This is where we actually get into what do these detectors look like. So there will be some active piece of your material if this is your detector. This is most definitely not Rutherford scattering anymore.

And there's got to be some window. We can make it as thin as we possibly can. And they make it out of the most X-ray-transparent structural material that they can, which tends to be beryllium. So beryllium has got an atomic number 4. It's the first and lightest element that you can make structural anythings out of.

So if you want to protect your detector from, let's say, air or something-- if this were full of air, it would absorb the x-rays, so you want there to be pretty much nothing. You can put a very thin, seven-micron beryllium window in front. But the problem is we've already got one of these mass attenuation curves. And when you get down to these energy levels, you attenuate everything.

So the lower energy your Bremsstrahlung is, the less likely you're going to see it. So even though this is the actual Bremsstrahlung spectrum, this is what we observe. And I haven't finished grading the tests yet. But I like I promised, for the two folks who do the best, I'm going to ask you to bring something in for elemental analysis. This is precisely what we're going to see.

You're going to see this Bremsstrahlung which is not the actual spectrum coming out, but this has to do with the absorption of x-rays in the detector window, as well as some self-shielding. If we're using a scanning electron microscope, which is nothing more than an electron gun, and you're firing electrons to some distance in the material where they'll then interact and send off x-rays, you've also got this part of the material to contend with, some self-shielding.

So not only do the x-rays all have to get through the detector window to be counted-- so the high-energy ones, which we'll have with small wavelength, get through, but the low-energy or long-wavelength ones might get stopped. You also have to get out of the material itself. The electrons don't just produce x-rays in the outer atoms of the material. They go down a micron or two. And then the x-rays that are produced in those interactions have to get back out again.

So it's interesting. It's kind of like the inverse photoelectric effect, right? In the photoelectric effect, photon comes in, electrons come out. In a scanning electron microscope, electrons come in. Photons come out. Many of them are these characteristic x-rays. Because now if we start to review what sort of interactions are possible when we fire electrons into material, we've just gone over Bremsstrahlung.

And we know that with higher and higher energy electrons, you're going to get more and more Bremsstrahlung. But you're not going to see the actual x-rays produced at low energies no matter what, because this isn't just a system on paper. It's real life. And you're going to get characteristic x-rays that come from energy transitions.

So if you fire in an electron, and you happen to undergo one of these ionization collisions, you might just knock an electron out. So let's say an electron comes in, knocks an electron out. Then another electron fills that shell, giving off-- in this case, it would be a k alpha or a shell 2 to a shell 1 X-ray the way I've drawn it, which is why Julia has got everything right on the spectrum here.

There's the Bremsstrahlung, and then there's these characteristic peaks. The background is due to radiative stopping power, and these characteristic peaks give away some of the ionization stopping power. And so all in one spectrum, you can see just about everything going on in this material.

The last thing that you can't see that I would be remiss if I didn't talk about it as a radiation material scientist is radiation material science, which really is concerned with mostly Rutherford scattering, Rutherford or hard sphere scattering.

This is the last of the major interactions between charged particles and matter that concern us. It's not really in your reading except for, I think, being mentioned once, because they didn't seem to think it's important. But I happen to think it's extremely important, because this is the basis behind radiation damage.

In all of these collisions right here, you have some sort of displacement of electrons. And those electrons can get ionized, and other ones will fill them back in the holes and whatever they'll do. But at no point were nuclei displaced.

You can transfer a lot of energy without moving any atoms around. But when your energy starts to get lower, you end up with a new kind of stopping power-- let's call it nuclear-- which scales with, as always, a number density times pi. Let's see.

Little z, big z, e the 4th-- everything looks pretty similar so far, except for now we've got the energy of the incoming material, and now we have a mass ratio. Because in this case, you're actually undergoing some sort of a hard sphere collision between one atom and the other times the natural log. This is going to look awfully familiar.

It ends up being some energy term over-- actually, let's just go with-- yeah, gamma ei over some new energy. This thing right here is called the displacement threshold energy. And it ranges from about 25 to 90 eV, but it's usually 40 eV.

And what that is-- it's the max-- the minimum amount of energy that has to be imparted to a nucleus smack head on in order for it to move from its original atomic position. And that's what's known as a hard sphere type collision. Or in this case, it's just like all the other qequation-looking scenarios that we looked at before. So let's say the little nucleus goes off, and the big nucleus goes off. This should look familiar by now because I've harped on it probably too much.

Now what I want you to consider is this big nucleus had a position. It liked where it was, and now it's been knocked away. What's left over is an atomic vacancy, which is the most basic building block of radiation damage. So sometimes it's neat to look at the ratio of-- let's see. Make sure I get this ratio right. Ionization on top. To look and see when is ionization versus radiation damage actually important.

And the ratio scales with 2 times the mass of the nucleus over MeZ times their respective natural log threshold things. And that's the ionization potential and log gamma ei over ed. So what this says is that for higher energies, ionization is more important, and for lower energies, nuclear stopping power or radiation damage is more important.

If we graph these two, let's say, on a log e graph-- and let's say we have our nuclear stopping power in blue and our ionization stopping power in green-- we end up with curves that look something like this. That's our nuclear. That's our electronic or ionic.

So what this actually says is if you fire high-energy neutrons or high-energy protons into a material, it's ionization that does most of the damage at high energies until you slow down to around like 10 to 100 keV level, curiously very similar to this 500 times i bar, the mean ionization potential, at which point Rutherford scattering or hard sphere scattering becomes the dominant mechanism.

And so what this says is if we want to draw a picture of what radiation damage looks like-- let's say we had a proton that we're firing into a material, and it hits some atom that we're going to call the PKA or Primary Knock-on Atom. That PKA then becomes-- let's say it was nickel. It's like a nickel plus 26 ion because you've knocked the nucleus out of its electron cloud, effectively, and it's now flying out through the material.

That proton might go off to do more damage somewhere else, but it's not actually the protons or the incoming particles that do the bulk of the final radiation damage. The radiation damage is mostly self-ion radiation. Even though it all starts with the incoming particle, nothing would happen if the incoming particle didn't show up. Most of the final results of the damage are from these heavy ion collisions. And so we actually talked a little bit about-- I think we talked about when this ionization starts to pick up for electrons versus heavy ions.

If you think about when electrons start to radiate away most of their energy, taking it away from radiation, it's like 10 to 100 MeV. What would be the case for a heavy ion, like even a proton? Well, what's the only thing that changes when you change from an electron to a proton here?

AUDIENCE: The charge.

MIKE SHORT: Well, yeah, the charge is the opposite sign, but of equal strength. But what else in this formula?

AUDIENCE: [INAUDIBLE]

MIKE SHORT: That's right. So for heavy ions, for even things like protons, you need to go at approximately 1,837 squared more energy than the electron. So we're talking in the gigaelectron volt to teraelectron volt range for ions. So this is why Bremsstrahlung is not important for any sort of ion interactions unless you are a high-energy physicist and you're working in the GeV or gigaelectron volt in the upper range.

So we like to say in the radiation damage field if you want to know the total stopping power from all interactions, you have to take into account the ionizations. I'll just make that a minus sign for the symbols. The nuclear and the radiative.

For most radiation damage processes except for high-energy electron radiation, we neglect that. The reason is that the radiative to ionization stopping power is pretty close to zero. It's like-- even at 10 NeV, it's like 1 over 2,000 squared-- or 1 over, I guess, 4,000,000. It doesn't matter at all.

With heavier ions, it becomes even less of an issue because you can deflect a heavier ion less with the same Coulomb. And so what ends up only mattering is the ionization stopping power and the nuclear stopping power. It's this nuclear stopping power that leads to collisions. And it's like two of five of.

So I want to stop here and answer any questions. And I'll hijack a bit of the neutron discussion on Thursday with some review of this and filling in the last gaps of radiation damage. So anyone have any questions from today? Yeah?

AUDIENCE: Can you repeat what you just said about why the radiation term goes away?

MIKE SHORT: Yeah. The radiation term goes away because of that.

AUDIENCE: And that's under the assumption you're working with a proton or a heavier--

MIKE SHORT: Yeah. If you're working with an electron, then it actually does matter. If you're firing 10 MeV electrons into something, you must account for the radiative stopping power, because there's a lot of it. At 10 MeV, there's as much radiative as ionizing, and there's basically no nuclear yet.

But for anything heavier-- even muons, which are approximately 237 times heavier, or protons, which are approximately 1,837 times heavier, it totally doesn't matter because it scales with the mass ratio squared. It might be 267 for muons. I forget that middle number. But still, 267 squared is a pretty big number.

Was there another question here? I thought I had seen a hand.

So remember, you guys had said you want to see some radiation material science or radiation damage. This is where it comes from. This is why I love teaching graduate radiation damage and 22.01 at the same time-- because they're the same thing. Except you guys get the derivations, and in the grad class, I say I assume they know it. And then in the homework, I find out they don't. But it doesn't matter because they're supposed to. At least you guys will, so you've got the power. And knowledge brings fear, as I like to say.

OK. I'll see you guys on Thursday when we'll wrap up a little bit more radiation damage, because I can't resist. And then we'll start moving into neutron interactions, which is kind of taking a step down from here, because there aren't really any electronic interactions. But because we can deal with enormous populations of neutrons, things are going to get messy. Have you seen the equation shirts that we have here, the neutron transport equation shirts? Yeah. We're going to derive that on Thursday.