MICHAEL SHORT: Today I want to pick up where we left off-- well, to remind you where we left off last time, we were watching videos of crushing things with the explicit purpose of understanding material properties, so that we can talk a little bit about radiation damage and nuclear materials. Since I got more than a few requests to say, what's all this nuclear materials? It's like the biggest research field in the department, and yet, it's not talked about in 22.01. Well, now it is.

So we talked before about all the different stages in radiation damage from creation of single defects to their clustering into things like voids or loops and super structures that have end up having macroscopic effects on material properties, and all of them is due to the production of crystal and defects due to radiation. And to refresh your memory quickly, I'm going to jump ahead to the stress strain curve that we were looking at before we started watching videos of crushing things, to remind you about the different material properties and what they actually mean.

So anyone remember what we mean by toughness in relation to this curve?

AUDIENCE: The area [INAUDIBLE]

MICHAEL SHORT: That's right. The amount of energy it would take to actually cause this material to fail. That's a measure of toughness. How about strength? Remember this curve is stress, which is a force per unit area, versus strain, which is an amount of elongation. The strength of the material is how much stress you can put in until it starts to either plastically deform or it hits its UTS, ultimate tensile strength, where it will just fail.

How about ductility? What do we mean by that? Either intuitively or on the curve, yeah?

AUDIENCE: How much you can stretch it?

MICHAEL SHORT: Exactly. How much you can stretch it before it fails, indicated by this point right here. The strain to failure would be a good measure of activity. And finally, stiffness. What do you think? Stiffness is more of a response function, so it's how much does it deform in relation to how

much stress you put into it. So it's the slope of this part right here, so that's why I want you guys to know that we mean actually different physical things by these properties, which will be important to note when we start to discuss what radiation damage actually does.

So the basic mechanism of radiation damage is like you might imagine. Let's say this green particle is a neutron or a heavy ion or a proton or an electron or anything. Anything that's energetic enough to cause atomic displacement. So as that neutron or whatever enters, it will strike some of the atoms in this perfect crystal, creating what's called a primary knock on atom, or PKA, for short. And then that neutron and the released PKA will go on to hit more and more atoms, resulting in what we call a damage cascade, leaving behind a lot of different types of defects.

We talked about these last time, but I'll just refresh your memory. A vacancy is a type of defect that, well, it's not really a thing, right? It's just the absence of where an atom would have been, but we refer to them as defects of their own that can diffuse and move because let's say another atom moved into the position of this vacancy. Then we can say the vacancy moved to that atomic position.

There's also interstitials, or atoms where they shouldn't quite be. And this leaves behind a whole bunch of damage that we quantify using a measure called DPA, or displacements per atom. It's a simple measure of how many times has every atom left. It's a lot of sight. That's it, though. It's not actually a unit of damage, and I'll be giving a talk at MRS, the materials research conference tomorrow, railing against this DPA unit because I'm going to explain this a little bit right now.

What is a DPA? A DPA measures the number of times that each atom has moved out of its original site, but it has nothing to do with how many times it stays out of its original site because a lot of those atoms will get knocked away and then just move right back, but the DPA part only measure is what we call the ballistic stage of radiation damage.

Let's see if this works. I've just realized I can jump back to a slide without inducing epilepsy. Yeah. So what DPA actually measures is how many times does this process happen? How many times do the atoms get knocked around? But it says nothing about where they end up, and that's the really interesting part about specifically radiation material science. Because let's say one of these interstitials were then to combine with one of these vacancies. It's like they were never there.

Even though they were displaced, and would be counted as part of the DPA, or the radiation damage dose, the net effect on the crystal material is nothing. So let's say what really the DPA is. It's a simple formula that I think you guys may recognize. This look familiar from all of neutronics that we've been doing? It's yet another reaction rate. It's an energy dependent flux times another type of cross-section that we call the damage displacement cross-section, or sigma D, and it's integrated over your entire energy range, and that's all there is to it.

So with what you know 22.01, you can understand the basic unit of radiation damage. As you might imagine, we've had four lectures on neutronics, so if you can understand all there is to know about DPA after four sophomore lectures, it's probably a pretty simple unit. You're right, it is. What goes into this damaged displacement cross-section is also something that might look a little familiar is a cross section that says, what's the probability of some particle coming in with energy E and imparting kinetic energy T to another struck atom?

That comes right from-- remember our treatment-- I think I've drawn this probably 50 times now. Our hollow cylinder treatment of a charged particle with charge little ze interacting with a particle a big ZE at some impact parameter B. We wanted to know well, for all possible approach paths, the area of this hollow circle, or the probability that this particular approach path is taken, is just the area here 2pi b db. With some constants in front of it, which actually is that cross section what's the probability that our particle goes in with energy E and imparts kinetic energy T? It's directly related to that impact parameter B.

And this is the same thing that you're seeing right here. You then multiply by this little function nu of T, which represents the amount of damage, or the number of displacements done, for each one of these reactions. And there are simple models, there are mostly linear models for-- if a particle comes in with energy E, leaves with energy T, how many displacements happen? It's a pretty simple linear piece-wise model, and that fairly well approximates the number of displacements that happen, but I want to get the idea of DPA versus damage. They're two very different things, and they're often equated.

Much like the material properties of strength, ductility, hardness, and toughness are equated in colloquial speech, but that's absolutely wrong. So is the idea of DPA and radiation damage. Because DPA, again, just measures the number of times that an atom is displaced. Damage is some measure of the number of messed up atoms at the end of the game, and they operate in very different timescales.

It takes femtoseconds to picoseconds for a damage cascade to happen. So the DPA is all over in less than a picosecond, but it can take years for all these different defects to diffuse, to cluster up, and to form these super structures, and actually end up causing the damage that can lead to material property degradation. So what sort of factors would affect the speed at which these different defects end up finding each other?

What could you vary about a material or its environment to change the speed of these atomic diffusion jumps?

AUDIENCE: Temperature.

MICHAEL SHORT: Indeed. Temperature. Reading off my list-- well, the whole list jumped up. OK. You got the first one. What were you going to say?

AUDIENCE: I was going to say temperature also.

MICHAEL SHORT: OK. Yeah, absolutely. Temperature determines diffusivities. It also can change phases or crystal arrangements, like for the case of anything iron-based. The dose rate, the rate at which those neutrons come in can change the rate at which the defects cluster up. Chemistry, if you have solute atoms, which I've drawn here. You may have let's say chromium atoms and iron, and the chromium atoms are a little bit bigger.

Defects may be attracted to or repelled to those extra solute atoms, changing the way that they interact with each other. And then micro structure. Things that are bigger than on the order of atoms. Grain boundaries, dislocations, all of those defects that we talked about last time, just to refresh your memory of what those are.

We have been talking about zero dimensional defects like vacancies. We spent a while on dislocations, these one dimensional defects that other defects can be attracted to. We saw an example of a two dimensional defect, known as a grain boundary, where you can see this line between different arrangements of atoms. And there can be three dimensional defects. Like inclusions of some separate face sitting in the material. Like the manganese sulfide we found in the Alcator fusion reactors power rotor.

And all of the presence and density of all those different defects can be quite strongly influenced. Let me start that sentence over. The movement in clustering of those defects can be quite strongly influenced by the presence of all those other defects. So again, the DPA

actually tells us this part of radiation damage, and that's what we tend to simulate with these ballistic binary collision approximation simulations, where we just say like billiard balls, how many atoms knock into each other? What it doesn't tell us is everything else, and it's the stuff that happens here that can tell us will our materials fail in nuclear reactors?

And there's evidence for this. I'm not just ranting against it, no I am, but I'm doing so with evidence, so it's justified. So here's a nice experiment I like to show in every talk for this case. These folks took pure nickel and put it in the same reactor, at the same temperature, and got the same amount of swelling. All the conditions were the same. Same temperature, same materials, same microstructure, same reactor, same neutron energies. Just a different dose rate.

A 30% difference in the rate at which neutrons arrived at the nickel, and they get the same result in void swelling, one of those bad things that happens, at two and a half times the DPA, which tells us that there's a very strong dose rate effect for material damage. So if you want to answer the question, well, how much dose does it take to reach 3% swelling in nickel? Can't answer that question, you don't have enough information.

Even if you say, how much dose does it take with one of the neutrons at 600 Celsius in this one reactor? You can't answer that question. Kind of tricky. And a lot of the rest of nuclear materials data looks something like this. Now, I don't want you to worry about what the axes say. They're not readable because they're not important. What I do want you to know is what's the quality of this data set you see?

Would you be bold enough to draw a trend line through a single data point? No. What about three where it doesn't actually match up with one of them? Or is there any reason why you think they made this parabolic instead of a linear line? I can draw a line that would fit between the error bars of these two right here. So the trick is doing these experiments is extremely difficult and expensive.

So just throwing something near the MIT reactor for a month, because we did this, we took a few hundred milligrams of copper, aluminum, and nickel, threw it in near core position of the MIT reactor, and that cost $40,000, and that did about 0.002 DPA, or about the dose that you'd receive in a normal power reactor in one day. If you want to actually say how long will it take to get materials to the end of their useful life, this tends to be anywhere from 10 DPA in light water reactors, to hundreds of DPA in proposed fast reactors to 500 DPA for

TerraPower's traveling wave reactor.

Now, I don't particularly have-- let's see what's 500 divided by 0.00-- I don't have 10,000 years to wait for the final answer. The best we can do right now is to stick them in a reactor called BOR-60 in Russia. I've actually been there. It's in the very Western edge of Siberia-- I don't know if you could call it that-- in a city called Dimitrovgrad. They have a sodium cooled fast reactor.

For those of you who are wondering when our advanced reactor is going to be built, they are built. Just not in this country, not very much. But Russia's got a fleet of sodium cooled fast reactors that can get you 25 DPA per year. And if your reactor is going to go to 500 DPA, you have to know whether or not your materials will survive, you have to wait 20 years for the answer.

So what investor is going to be like, all right, here's $10 billion, but I can wait 20 years for a return on investment. No. I can wait 20 years to start building the reactor, which means 40 years for a return on investment. Chances are, if someone's got $10 billion to give, they're going to be dead by the time they get a return. So this is a no win proposition.

So what we really need to know is what is the full population of every single type of defect in an irradiated material? That's what I mean by damage. Did I show you guys this movie yet? The orange one? We've talked about vacancies in an abstract sense, but this is a movie of one of them actually moving about on the surface of germanium. So this is a scanning tunneling microscope image-- I think that's what it stands for-- and these are atoms on the surface of germanium.

And that right there, that darker orange thing moving about is vacancy diffusion. It's actually happening. You can see it in real time.

AUDIENCE: Is this real time then?

MICHAEL SHORT: Pretty much, yeah. So I think this was-- yeah-- 30 frames a second, or so. Anyway, I don't remember exactly, but I'd say that's why I always reference everything in the presentations. I encourage you guys to look it up. And then the only reason these slides aren't up yet is because they're 300 megabytes, and I didn't have the bandwidth to upload that from my house. Now that I'm on campus, I can get a 300 meg presentation up there because it's full of movies.

What sort of things could happen to these defects? So radiation produces all these crazy defects, then the DPA description is over. What could happen next?

[INTERPOSING VOICES]

MICHAEL SHORT: Sorry, Jared, and then-- yeah.

AUDIENCE: Material could crack.

MICHAEL SHORT: Material could crack. That would be the worst case scenario, but that is indeed what happens in the end, and I'll show you some pictures of that actually happened. Yeah?

AUDIENCE: You mentioned that displaced atoms can find their way back?

MICHAEL SHORT: Yep. So they could recombine with different types of defects and annihilate each other. If you have a vacancy and an interstitial nearby, the interstitial can plug the hole of the vacancy and your left with another perfect crystal. But now what happens if two vacancies find each other? Then you've got the makings of a void, or we then call it a small vacancy cluster of two vacancies, but it's actually more stable for these vacancies or interstitials to find each other and make these larger defects than it is for them to sit alone in the crystal structure.

So there is a thermodynamic driving force bringing them together, and then as those defects build up, then what Jared said could happen. You could crack the material because it could get weaker, less ductile, less tough. Weaker is the opposite of strong, and what's the other one? Toughness-- oh, and harder, actually. So the origins avoid swelling I'll start with the humble vacancy.

A void is nothing but a bunch of vacancies or a pocket of vacuum or gas in a material, and it all has to start with these single vacancies. As they cluster together, they reached this threshold in terms of free energy where putting a few of them together is not quite energetically favorable, but it's not so unfavorable that it never happens.

So once in a while, you'll get a few vacancies to come together, and that cluster will survive for a little while. All the while, you're making more and more vacancies nearby, and if it gets to a certain size, that free energy goes negative. And when the free energy goes negative, it becomes stable on its own, and then that void will simply continue to grow, and grow, and grow.

And so there's this process of absorption and emission of defects by larger or smaller void, so if you have a whole bunch of voids near each other, some of them can be emitting vacancies, which can be captured by the other ones, and this is part of why they don't all just disappear at once. They have finite lifetimes, long enough that you can build them up to the size where they become stable. Then this free energy eventually curves down, becomes negative, and then they just stick around, and we've actually seen these clusters or voids diffusing, so it's not like vacancies alone are the only thing that moves.

We've actually seen clusters of defects diffusing, mostly in one dimension, but what you're seeing here is a TEM, or transmission electron microscope image, of one dimensional diffusion of a vacancy cluster. That little black blob right there is a pocket of vacuum that's moving back and forth. And if it happens to find another pocket of vacuum, it could then combine to a bigger pocket, becoming a bigger and bigger void.

The other problem too is that most materials generate helium when you irradiate them with neutrons. Did we go over the what's called the N alpha cross-section? Does that sound familiar to anyone? All right, I'm going to pull up Janice like I do pretty much every class. Let me show you what's going on. But this is an important one to note.

Because a pocket a vacuum is not that stable, but if you get a little bit of gas to stabilize that pocket of vacuum, then that pressure differential goes down and that void becomes a bubble, and that bubble is more stable. Good. Because on the last isotope, where I was showing someone [INAUDIBLE] alpha cross-section. How convenient.

So among the millions of cross sections that we've gone through, there is this one right here called Z alpha. So what-- Yeah?

AUDIENCE: [INAUDIBLE]

MICHAEL SHORT: Got to clone the screen. That's right. Thanks. There we go. So there's this one here called Z alpha, which means neutron comes in, alpha particle goes out, alpha particle is just an ionized helium ion, which very quickly pulls in two electrons from anywhere else in the metal and becomes helium gas.

And this cross-section is not zero, especially at higher energy starting around 2 MeV, there is a small, but non-negligible chance that a neutron will go in, and a helium atom comes out. And those helium atoms have nowhere to go, they find the easiest place to sit, that happens to be

pockets of vacuum. Voids.

And what that actually does is stabilizes those voids so the curve I showed you back here. This is the case of free energy for a vacuum pocket of void, and that free energy gets lower and lower as you start to fill that void with gas. So as the voids fill with gas, they become more and more stable, and a lot of materials generate their own gas. So that's that, and they end up forming these bubbles.

You guys remember, last time I showed you a bunch of voids. They look like diamonds all aligned in the same direction. What do you see different here? They're not quite circles, right? But they're kind of round edge squares. They are also all pointing in the same direction. If you look carefully at this one, this one, the one above it. And the big one over there. It's harder to see for the small ones, but for these three big ones, you can see that they're all rotated in the same direction, giving away the crystal orientation of the material.

But the reason that they're starting to swell up from diamonds into bubbles, is they're full of gas and they stabilize the voids. You can also get dislocation buildup. Normally, you would have to deform a material to create and move dislocations, but when you apply radiation, you can just create dislocations. I mean look at that. This is kind of cool. You've got a dislocation source right here. Every one of those lines you see is a dislocation, and you can see it's spiraling out and ejecting dislocations from this one little spot.

Any combination of small clusters can collapse into dislocations or the stress induced from irradiating things can cause more stress that can move more dislocations. You create what's called this network forest of dislocations that makes things a lot harder to deform. So let's see, I want to show a couple more videos of this because it's very clear in some of these.

You can actually see along these lines right here a few different orientations of dislocations, and if you watch up here, you can see some dislocations moving, and then there's a source that emits more right there, and so all the time, you're creating dislocations that are being emitted from different places and colliding with each other. The trick that we didn't talk about yet, is when dislocations from different directions collide, they get stuck. And when they get stuck, they can't move. And when they can't move, you shift the balance from slip to fracture, which means, like Jared said, it's easier to just break something, rather than plastically deform it.

And the effects of this are things like stiffening. An increase in the Young's Modulus. Because

if you remove some of the compliance from the material or make it stiffer by injecting all sorts of different defects, it takes more stress to impart the same strain. That might not be a bad thing on its own. Your materials get stronger, that sounds like a good thing, right? Not always because it doesn't just come as stiffening. But now from an atomic point of view, why does this stiffening happen? Anybody have an idea?

I'm going to jump back to the stress strain curve. So the stress, or the yield stress of a material, is usually defined as this point right here. When you go from elastic reversible deformation to irreversible plastic deformation. If something gets stronger, it means that this yield stress goes up. And if something gets stiffer, it means the slope goes up. These two tend to happen at the same time. So if something is stiffer and stronger, then the stress strain curve would be drawn more like that.

And what actually physically happens at this point? This is when dislocations start to move. Dislocation movement is irreversible. You can't just snap it back when you relieve the stress. And by making something stronger and stiffer, you make it more difficult for those dislocations to start moving. And you can do that by throwing any defect in their way. And since radiation creates pretty much any and all defects, it's a great way to stiffen and strengthen the material.

So one of the reasons things get stiffer and stronger is, remember before we showed you that video of dislocation sliding through material? Folks don't remember. I'll bring it up right now. I didn't see a lot of shaking heads. This one, right? This one right here. So remember, before, we showed you the way that dislocations move. So the ways that materials can deform without breaking.

If you throw anything in its way, you're going to make this process a lot more difficult. And all radiation damage does is throw things in the way. So if you throw absolutely anything in the way from solute atoms to interstitials to vacancies, all of a sudden, it's harder for that dislocation to move because some of those bonds are stretched out, or there's a few extra atoms in the way. And you can then start to create what's called little pin sections called jogs.

Let's say a little vacancy moves over to this dislocation, meaning that it goes up by one atomic position. Then you've got pieces of this dislocation that are not in this preferential slip plane, and they get stuck. So all of a sudden you go from a completely gliding-- or what we call glissile dislocation-- to one that's stuck, or sessine, those are the actual material science words that we use.

And what it ends up leading to is a strong loss in ductility. At the same time as things get stronger and stiffer, they tend to get much less ductile. So what you're looking at here are fuel shrouds. These are fuel boxes that surround the fuel pins in a Russian sodium fast reactor, and usually what you do is you would grab onto this piece right here and lift up to remove this fuel from the reactor during refueling.

What happened here is they grab that, they started to pull, and they heard a little clink, and up came half the fuel box, and the fuel stayed down in the reactor with no way to pull it out. So this is the reason why radiation damage is such an important field of study is you might not know anything has happened until you shut the reactor down and go to take out the fuel and realize that you can't because everything is as brittle as glass.

This is actually a talk I saw yesterday. We had a summer visitor in our lab from Kazakhstan. And most of their radioactive materials came from this reactor that shut down in 1999 on one side of Kazakhstan, and they wanted to transport those materials to the other side.

So they hired the cheapest truck drivers to go on the bumpy roads. And the scientists were freaking out, because only they knew that all of the metal that all those guys thought was going to be ductile like metal was more brittle than glass. And any sort of bump would cause just complete shattering of this metal and catastrophic release of radioactive material.

So this took them-- let's see-- I think Kazakhstan is smaller than the US. So who here has done a cross-country trip? How long it take you?

AUDIENCE: Six days from Seattle to here.

MICHAEL SHORT: OK, and did you stop along the way?

AUDIENCE: Yeah.

MICHAEL SHORT: OK, so this trip took them 13 days because they went slow. And I don't think the roads in Kazakhstan are as good as they are in most of the middle of the country. The coasts are terrible, but the middle is pretty good. So this trip went slow because the scientists said this happens. You should be careful. And luckily, there were no problems and no release of radioactive material. Pretty cool.

What you want to happen is for dislocations to move on the easiest planes. And so what I have redrawn here is, let's say you've got a bar of some metal, some face-centered cubic metal, as

you pull on it, like this, it will actually deform at about a 45 degree angles. You might wonder, why does that happen?

But if you look carefully here, what's the closest packed plane of atoms that you can see? It's not this plane normal to the stress direction. It's like this one or like this one, and so you actually end up getting deformation in what's called slip plains, or the easiest directions for things to deform.

And without going into any of the math or atomistics, I just want to show you some examples pulled out of, again, the fusion reactor. So this is a piece of rotor steel from the same Alcator rotor where we found that inclusion. We were pulling it in this direction, and look what formed. All of these slip bands at 45 degree angles, showing you that just because you pull in something in one direction, doesn't mean it deforms in that direction. It deforms and little slices in the direction that dislocations can move the easiest.

So when you actually pull on something, to show you diagrammatically, it deforms something like that. You get a mixture of bending and rotation to make it look like the bar is bending uniformly straight, but on the microscale, it's not. So this is a pretty slick image of a single crystal of cadmium being pulled in this direction, and you can actually see every plane there, that's a slip plane. That's a plane where dislocations have been moving all the way to the outside of the material, which is pretty cool, and this is the process that you want to happen.

Anything in the way of those dislocations, you don't start forming these slip bands, and you'd make it more preferable that the thing will just break and fracture. To show you some extreme examples of slip, that's when you have to go nano. So these are some pillar compression test that used a focused ion beam, which we will be using to top off our study of electron interactions with matter and ion interactions to take a piece of metal and carve out a little cylinder. And all they did is smash.

They came down and compressively pushed on the cylinder, and look how it deformed. Not in the way you might intuitively expect, if you don't know any material science. So it deformed all on 45 degree angles and very weird compression. Not actually weird, if you know what's going on.

there's lots more neat examples of this. If you don't push too hard, you can actually see these perfectly symmetrical slip planes at 45 degree angles to the axis of compression, and you

know every one of these pillars you see this happening. And this is what you want to happen to nuclear materials because you're really trying to balance this between slip and fracture towards the direction of slip. That means that something will deform a little bit before it just shatters like those channel boxes from the Russian reactor. Any questions on what I said before I go on to the macroscale properties?

All right, and let's get into the real world stuff. What actually causes this embrittlement? Well, there's a few things. Remember we saw videos of those dislocations in a traffic jam. If not, then I'll refresh your guy's memory. It's a phenomenon called pileup, let's see. There is the traffic jam. Do you guys remember this video now-- I know it's been a week, but these dislocations are moving and feeling each other's stress, and so they can't move as easily as they would want to, so you end up with a phenomenon called pileup.

This happens both near other dislocations and near any other defect that gets in the way, like a grain boundary. So for smaller materials you end up with more of this pileup, and they tend to be of a fair bit stronger, and a fair bit-- they can be less ductile, with some exceptions. I won't say they're always less ductile, but if you put anything in the way, notice, this just says barrier. Any other defect can act as a barrier.

And this ends up shifting what we call the ductile-brittle transition temperature. This is the property that people worry about for reactor pressure vessels, because you would want the pressure vessel, which in cases the entire core of the reactor, to always be ductile in the worst possible situation. The worst possible situation is on the absolute last day of operation at the coldest temperature it could possibly be.

As you guys know, or you probably know, when you make something colder, it tends to be more brittle. So who's frozen stuff in liquid nitrogen and broken it before? Good. I'm glad to see a few hands, so what do you guys freeze and break? So just like froze a bottle of Pepto Bismol and shattered it? Nice. That must have been a fun mess to clean up. Yeah, what about you, Sarah?

AUDIENCE: A banana and a coin.

MICHAEL SHORT: Cool, OK. And we're able to break the coin? There you go. So normally you'd be able to bend a coin, or if it's a one yen coin, you can bite through it. Not when immersed in liquid nitrogen. What about you, Charlie?

AUDIENCE: Flowers.

MICHAEL SHORT: Flowers, OK. Classic. So take a rose and shatter it. Yeah. All these normally ductile materials become extremely brittle when they get colder. So do reactor pressure vessels. The problem is, they don't just get brittle when they get to liquid nitrogen temperatures. At the end of their life, they can be brittle at room temperature from radiation damage.

So there is what's called a ductile-brittle transition temperature before a pressure vessels are irradiated, there's about this 50% line, whatever this temperature was where you have let's say, 50%, or 10%, or 0% ductility. Whatever measure you're using, you say, OK, we always want to make sure that it's got a certain amount of energy absorption capability, toughness, at a certain temperature.

And as you irradiate that vessel, this shifts over this way, and this upper amount of energy, this USE, or what we call upper shelf energy, because this looks like a shelf, goes down. What you want to make sure is that this change in ductile-brittle transition temperature never reaches room temperature. You might think, oh, it's OK, reactors run at 300 Celsius, and things are pretty ductile right there. Well, you usually have to shut the reactor down once you're done with it to refuel, and if something goes wrong, there's a pressure spike, you can have a condition called pressurized thermal shock, or PTS.

In that case, you would have a sudden pressure wave from, let's say, from steam explosion or whatever you could have, and you want that vessel to be able to absorb that energy instead of breaking in half, because if you break it half, that's a radioactive release. The way you test ductile-brittle transition temperature is what's called a Charpy impact test. It's probably the highest tech, lowest tech test I've seen.

You simply hit things with a hammer. A very well calibrated, precise hammer. Let me pause so you can see what the sample looks like. You have these little bars with a notch in them. The notch is to make sure that acts as a stress concentrator, and you know where the breaks going to happen. So in a Charpy test, you line up this little sample, and you've got-- actually in your reactors, usually, you've got pieces of the pressure vessel in this form lining the inside rim of the reactor pressure vessel. So every time you refuel your reactor, you take off another few of these out and you hit them with a very well calibrated hammer.

And you can measure by actually turning this dial and letting the hammer turn it as it moves through the material, you can see how much energy was absorbed by the material as the

hammer comes back up. So it breaks right through the material, in this case, it's in a quenched or brittle condition, and for some reason, they have a lot of footage of the guy standing and not breathing, but what I want to show you is what actually happens here. I didn't make the video, I just got it. There we go.

So what you can see is that if the hammer were to move through air with absolutely no drag, it would come back to the zero position. If they had encountered some resistance, like with a piece of steel in the way, it then measures the amount of energy in joules that piece of steel absorbs from the hammer blow. The larger that is, the better. And by doing this test at a number of different temperatures, you can recreate this ductile-brittle transition temperature curve.

So they'll take a few Charpy coupons, they will test them at, let's say, every 25 Celsius, get a bunch of points, draw the line through the points, and decide where is the material brittle. At what temperature will it become brittle? To show you what something looks like when it's not brittle. The same test is done in what's called the normalized condition, where you simply heat the steel to a high temperature, relaxing out most of the defects, and bringing back as much of the perfect crystal structure as possible, which is really good for letting dislocations move through it.

So the same test is done by the same awkward feller who likes to stand there and not breathe, but you'll notice a very different result of this test. Doesn't look like it, but if you actually look at how much energy was absorbed, much, much higher. So something like 18 times more energy, and you can qualitatively see the difference between these two conditions by looking at the fractured surfaces, and this is where it starts to get intuitive.

Something that's ductile would tear more like taffy, where something that's brittle would cleave or break in half much more smoothly, so these are the two pieces of metal that we just showed. You the one that absorbed 180 joules by lots of defamation, and the one that absorbed 10 joules by fracture in a brutal way. This is what you want your reactor vessel to behave like, but the problem with these ductile-brittle transition temperature curves, is this not just this part that you're worried about, it's that part. So even at high temperatures, things get less ductile.

So it's a combination of temperature and number of defects. And if either one of these criteria fails, if you become too brittle at low temperature, or your total ductility at high temperature

goes down too much, that's the end of life of your reactor vessel, and this is one of the biggest problems in life extensions of light water reactors. They were built for 40 years and they originally had license for 40 years.

How many of you guys have heard of the license extensions going on now, to 60 or 80 years? Yeah, so a few of you guys. I've heard heard, why not run the reactor longer. Not build a new one, but keep getting all this clean green cheap nuclear energy? This is why. You have to be absolutely sure that your vessel, your primary containment, will survive. And we're not so sure because well, we jump to the part of the video that's got the Charpy coupons. Those. We ran out.

We only plan to put these vessels in service for 40 years, and folks put 40 years worth of these coupons, plus some extras, in the reactor vessel. Now, in order to prove that it's actually safe to continue operation, you have to have some amount of material to test and say, OK, this vessel is still ductile, it's still going to work. What happened? What do you do when you run out of coupons? Anyone have any ideas, because I'm sure the industry would love to hear them.

AUDIENCE: Would you have to start using material from the vessel itself?

MICHAEL SHORT: You could start using material from the vessel. That's actually what I plan to do too, but with some very strong caveat. So if you were to scoop out a piece of the vessel, you then create a stress concentration. In addition, reactor vessel looks like a gigantic forging of really thick carbon steel with a very thin liner of stainless steel. And the stainless steel is there to prevent corrosion from the reactor water.

That thin, quarter inch bit of stainless steel is what actually saved what could have been one of the worst nuclear accidents in US history, the Davis Bessie plant, where there was a crack in the vessel. Boric acid actually ate through a whole chunk of the pressure vessel, leaving the stainless steel intact. And it's that little quarter inch stainless steel that saved the plant.

But if you were to take something out from the inside of the vessel, the part that gets the most damage, you'd be taking out some of the stainless steel, which is a problem. You could take a piece out from here, maybe the outside, but then you've got a stress concentrator. Any sort of chunk that is missing is where a crack is going to preferentially form, so you would weaken that vessel by taking a piece out. Anyone else have any ideas? Yeah?

AUDIENCE: Is it impossible to just replace the vessel?

MICHAEL SHORT: Yes. A new vessel means a new reactor. So the license for the plant is intimately tied to the license for the vessel. Any other ideas? Yeah?

AUDIENCE: The creation of the vessel, just put extra in there and then take that out.

MICHAEL SHORT: That's what they did, right? So that's why these Charpy coupons were there, but what do we do about the vessels that we already have? Yeah?

AUDIENCE: Can you make Charpy coupon or coupons that are similar to the status of the ones most recently taken out of the vessel, and then just put them in?

MICHAEL SHORT: That's what they're doing. So they're taking these Charpy coupons, which this is bigger than actual size. They break them, so let's say this region's all garbage, and then they cut a little mini Charpy coupons out of the last piece, and they're putting those back in. So that's absolutely right. You've just probably recreated a year's worth of licensing work and ideas and in a class.

But I just want to get back to Charlie's idea because that's what I think has to happen is you'd like to be able to take a piece out from the actual vessel and run a test on it. The only way to do that is to go nano. Take the tiniest, tiniest little piece out, and perform some other sort of measurement. So this is the idea that our group has had in using what's called stored energy of radiation damage, so I don't mind telling you about it, even though it's not like funded or papered yet because it's educational, and it's cool.

So every kind of defect takes energy to create. Defects don't just create themselves. You either have to raise the temperature of a material or in our case, irradiate it. And it's the energy of those incoming neutrons that bounces around different atoms and creates all these different types of defects. So those defects are storing energy in the material. And so if you think about how much energy does it take to destroy something, it would have to be the energy that it's already stored plus the energy that you put into it during the test can reach the failure energy.

What if you could measure the stored energy? What if there was a way to know how many of each of those defects there actually were in a material? We think there is. Well, we know there is. It's called differential scanning color imagery. It's a way of measuring the change in heat capacity of a material, where you take two very small furnaces-- you don't have to put this in your notes, by the way, this is just for fun.

You take two small furnaces, put your chunk of your material on one, and you apply a lot of heat to both of them. And you look at the difference in the amount of heat you have to put in to keep the two at the same temperature. So normally, you would get the heat capacity of a material, how much heat can extort per degree Kelvin. If this material's got a bunch of defects already in it, then you should release that defect energy by heating it and that would take a little less energy to heat it up, but there's a lot of problems with calorimetry, so we're actually using what's called nanocalorimetry.

We're doing this process on nanograms of material and seeing if you can irradiate something and measure its stored energy because if you could, you could take a tiny little razor blade, take out the smallest sliver of the vessel-- smaller than a grain of sand. Not enough to cause a crack-- enough to measure its stored energy. And I want to show you guys what this process looks like. So I'm just going to deviate from the actual lecture, and jump into the topic I'll give tomorrow. I think it's more interesting and more relevant.

There we go. I'm going to skip through some of this stuff, but it's within the last five minutes, I'll try to get through this, see if this is a record. It's what we call the ultimate snipe hunt. Has anyone here been on a snipe hunt? What's a snipe hunt?

AUDIENCE: People bringing you in the woods and tell you you're looking for a bird that doesn't exist.

MICHAEL SHORT: That's right. Pretty much this, right? They say bang a bunch of sticks, get a bag, and go look for snipes. They don't exist, right? They actually do. Snipes are real. You pretty much have to be British to know it, because they hunt them there for sport, and apparently, they're delicious. That's actually where we get the term snipe because the actual size of the sniper compared to the sniper is about that. If you can shoot that bird with a gun, you are an expert marksman and deserve the delicious and tiny treat that you've then blown apart with your bullet--

AUDIENCE: [INAUDIBLE]

MICHAEL SHORT: Yeah, exactly. So you can you know rain bird dust on whatever meal you've already prepared. That's what I like in finding these radiation damage defects too. Because some experiments have been done to plot the number of defects versus their size.

And as the defects get bigger, the number of them decrease. So most defects are very, very small, and it turns out that-- first of all, the resolution of the screen is funny. I think I know how

to fix that. Clone the screen and then jump back to presenter mode. That usually-- that's not what I wanted. That's what I wanted, great.

Most of the defects that cause these reductions in material properties are too small to see, even in the transmission electron microscope. So I don't have to tell you this stuff again, that I don't like the DPA, I showed you that. What we want is that. We took some inspiration directly from the Manhattan Project. Luckily, I have an uncle who works at the DuPont Library and Dupont was quite responsible for the Manhattan Project.

So this memo between Eugene Wigner and Leo Szilard-- one of whom won the Nobel Prize, the other one probably should have-- said, hey, radiation stores energy by neutron collisions like cold working and amorphization. So we've dug up this original memo from the 40s, and said, let's do this for everything because every defect has its unique amount of energy that it stores and creating it in some different amount of eV per defect, and we've done some molecular dynamics simulations to show that this amount of energy stored is pretty universal. When you irradiate something, we predict that it stores about 2% of its energy in radiation defense.

So if you know the number of neutrons that hit, and you know that the amount of energy per neutron, you know how much you're looking for. You know what your signal should be. And to jump through to the whole idea of differential scanning calorimetry, it's like what I drew here, but a lot more legible. You simply heat two materials, one of which contains your sample, measure their temperature, and look how much energy it takes between them to keep them at the same temperature.

We did some of these measurements on a piece of steel from the nuclear reactor, and we got a whole bunch of interesting looking peaks for the red curve compared to the blue curve in the other irradiated conditions, so we think there's something there. So we tried a more controlled experiment irradiating aluminum with helium ions and the accelerator in Northwest 13. And we were encouraged because the initial time that we heated this material, we got some stored energy out of it in some funky spectrum that might tell us what the defects are.

The bad news is we got that with the control heat too, and the really bad news is that when you normalize all these curves, you get something that you can't tell if I drew it or my son drew it. Looks suspiciously like the doodles that he does, not scientific data. And the problem is that DSC, differential scanning calorimetry, induces a lot of artifacts in the signal that we couldn't

separate from the noise.

So our solution was to go nano. To use a nano a DSC, or nano differential scanning calorimeter, that can heat about 10,000 times faster than a traditional DSC. So you can get your energy out from smaller materials way faster than these artifacts can manifest themselves. What we think is going to happen is that every one of these peaks here is going to correspond to one type of defect that's released at a certain temperature.

And by extrapolating, or say, integrating the area under those curves, you get the energy in each type of defect. And by extrapolating to a zero heating rate, you should know which type of defect they are. And if you know which defects you have and how many of each one, you know the full defect properties and material, you should know it's material properties. Because we already know if you have this many dislocations, it's this brittle, the question is how many dislocations.

So we start off by-- we use a different kind of calorimeter. It actually fits on a chip. In fact, there's two on a chip. There's one that we put our material on, and one as a reference that we both put in the accelerator being irradiated at the same time to control for that effect, and this is what they actually look like. The scale bar here is 100 microns, and that transparent spot is a little bit of aluminum that we vapor deposited onto the calorimeter. Right there.

And so the way this process works, is we take our DSC chip, we put a mask over it, vaporize some aluminum to deposit on one of the calorimeters, take the mask away, and irradiate the whole thing, and then finally put it in the nanocalorimeter, and I'll show you what happens slowed down by a factor of 1,000. That pulse right there. That whole thing just went from zero to 450c a millisecond.

And the reason it took a second is because I've slowed down the video by 4,000 or by 5,000 times, and that little pulse of heat actually released some of the defect energy. We were able to see very clearly, the first time we heated the sample, this extra area corresponds to some sort of energy release. We then heated that same sample a whole bunch of times, and made sure that it was always the same, which meant we had a fully relaxed material.

And it shows some sort of a trend. If you note this data was taken like two months ago, it's pretty fresh. Not published yet, so hopefully by the time this video comes out, it will be. And we see some sort of trend between the amount of irradiation it gives and the amount of stored energy you can get out of it. So this is what we hope can be used instead of those Charpy

coupons. We can go much, much smaller and just take out tiny pieces of the vessel and get the same information as you would from a Charpy test but on the nanoscale.

So the question then is, where is the defect fingerprints? Where are those individual defects that we were looking for? Well, I think they're just popping up right here. The reason for that is we picked a very fast heating rate for our experiments, and doing these sorts of measurements on other materials shows that at 10,000 Kelvin per second-- think about that for a second-- 10,000 Kelvin per second.

So in a millisecond, something heats up by 10 K, which is ridiculous. Yeah, something like that. You don't really see any peaks. The heating is so fast that the defects don't even have time to find each other, annihilate, and release their stored energy. So we need to repeat the experiments at some lower temperatures, see what the peaks are, but if you go too low, you end up getting a lot of noise in your signal. So there's going to be some sweet spot that we haven't yet found in order to see this stored energy.

So we're just at the very first experimental stage of trying to see can we extend reactor lifetimes. After doing simulations for a year that probably no one believes until you do an experiment, including me. But for now, it actually shows some sort of a trend, so it's just enough justification for us to buy one of these nanocalorimeters and start looking for real.

So if you want to see, now I've taken you from basic material science, to where's the field going to, how do we keep our reactors running in about two hours. I think that's the most compact introduction to nuclear materials I can possibly give you. So any questions on what you've seen today?

AUDIENCE: Is that roughly the trend you would expect?

MICHAEL SHORT: Yes, I would expect an up. That's the best I can say. As far as is it actually a line? Is it a curve? I am not as brave or stupid as some of the other folks that will draw an arbitrary shaped line through a single data point. So I'm not drawing a trend line yet. Yeah, any other questions? Yeah?

AUDIENCE: Is it-- or I guess you're making the assumption that one little spot in the reactive pressure vessel to say what the rest of the vessel has been exposed to?

MICHAEL SHORT: Oh, not at all. Take a whole bunch. Then, instead of just doing Charpy coupons of one place,

which is what we do now, you can get a map. We don't have that information now, but if you take pieces from all over the vessel, then you get an actual 3D map instead of a single point of saying all right, well, we picked what we think is going to be the worst condition. How do you know? You don't. How do you know for sure? You make measurements like these. Any other questions? Yeah?

AUDIENCE: What would a peak in this graph correspond to in terms of like-- How does it relate to some sort of damage effect?

MICHAEL SHORT: We would expect that a peak would relate to a certain type of defect reaction occurring. When some type of defect gets high enough in temperature that it goes from stuck to mobile, and as that moves, it encounters anything else it will in the material, and will react with all the other defects nearby, decimating the population of that defect and slightly depressing that of the others. And as you go higher, and higher in temperature, the slower and slower defects start to move. Yeah?

AUDIENCE: So you can get rid of the defects by heating it quickly?

MICHAEL SHORT: Mm-hmm

AUDIENCE: Would there be a way to self repair our radiation damage to pressure vessels themselves?

MICHAEL SHORT: That would be awesome. But the properties of the vessel are highly dependent on, not just its composition, but the heat treatment that went to make it. If you heat that vessel, you both remove the radiation damage and remove the strengthening put in by the forging and heating process. So you would have, if-- again, if you let's say, replace the vessel you have a new reactor. If you heat the vessel too much, it's no longer a code stamp vessel. Pretty tricky spot that we're in, huh? But we're trying to science our way out of it.

Well, it's a couple minutes after. I don't want to keep you longer, but I'll open on Thursday with a little story about how mass attenuation coefficients can get you out of apartheid South Africa. I'm serious. And then we'll move into dose and biological effects.