MICHAEL SHORT: So today, I wanted to give you some context for why we're learning about all the neutron stuff and go over all the reactor types that, until this year, the first time you learned about the non-light water reactors at MIT was once you left MIT. I remember that as an undergrad as well. The only exposure we had to non-light water reactors is in our design course, because we decided to design one.

So I wanted to show you guys all the different types of reactors that are out there, how they work, and start generating and marinating in all the different variables and nomenclature that we'll use to develop the neutron transport and neutron diffusion equations. The nice part is now, until quiz two, you can pretty much forget about the concept of charge. So 8.02 can go back on the shelf, because every interaction we do here is neutral, charge neutral.

There'll be radioactive decays that are not the case. But everything neutron is neutral. It doesn't mean it's going to be simple. It's just going to be different. But in the meantime, today is not going to be particularly intense, but I do want to show you where we're going. And this goes with the pedagogical switch that we made in this department starting this year. And you guys are the first trial of this.

We're switching to context first and theory second. I personally find it much more interesting to study the theory of something for which I know the application exists. Who here would agree? Just about actually everybody. OK. Yeah. That's what I thought too. So in the end, we had arguments amongst the faculty about, well, you have to learn the theory to understand the application. And that works really well when you say it behind the closed office door by yourself. But the fact is, I'm in it for-- yeah.

I'm in it for maximum subject matter retention, so in whatever order that works the best. And sounds like, for you guys, this works the best. That's what we're doing with the whole undergrad curriculum, not just this class. So let's launch into all the different methods of making nuclear power, both fission and fusion, and to switch gears since we're dealing with neutrons. I don't know what happened with the-- oh, there we go.

The idea here is that neutrons hit things like uranium and plutonium, the fissile isotopes that you guys saw on the exam, and caused the release of other neutrons. And as we come up with these variables, I'm going to start laying them out here. It might take more than a board to fill them all. And I'll warn you ahead of time, this is the only time in this course that we're going to have V and nu, the Greek letter nu, on the board at the same time. And I'm going to make it really obvious which one is nu and which one is V.

So this parameter that describes how many neutrons come out from each fission reaction we refer to as nu, or the average number you'll see in the data tables as nu bar. And so as we come up with these sorts of things, I will start going over them. And the idea here is that each uranium-235, or plutonium, or whatever nucleus begets two to three neutrons, the exact number for which is still under a hot debate, and I don't think it actually matters, will make a couple of fission products that take away most of the heat of the nuclear reaction.

And I just want to stop there, even though you know there's going to be a chain reaction. And that's what makes nuclear power happen. And we can go over the timeline of what actually happens in fission and what kind of a nuclear reaction it really is. So in this case, this is a reaction where a neutron is heading towards, this time we're actually going to give it a label, a uranium-235 nucleus. And it very temporarily, like I showed you yesterday, forms a compound nucleus, some sort of large excited nucleus that lasts for about 10 to the minus 14 seconds.

So it doesn't instantly fizz apart. There's actually a neutron absorption event, some sort of nuclear instability, at which point your two fission products break off. Notice, you don't have-- let's call them fission product one and fission product two. Notice, you don't quite have any neutrons yet. Neutron production is not instantaneous for the following reason. If you remember back to nuclear stability, when we plotted, let's say, I think that was maybe Z and this was N. And I think this was a homework problem. And you had to come up with some sort of curve of best fit for the most stable combination of NZ for a nucleus.

It was not a straight line. It was something on the order of like N equals-- what is it? --1.0055Z plus some constant, something with a rather small slope. Well, if you have a heavy nucleus, like uranium-235, and you split it apart evenly, let's just pretend it splits evenly for now, you're kind of splitting that nucleus along a rather unstable line. And, as you saw in the semi-empirical mass formula, a little bit of instability goes a really long way towards making the nucleus extremely unstable. So let's say you'd make a couple of fission products that just cleaved that nucleus with the same proportion of protons and neutrons.

How would they decay? Or how can they decay? There's a couple different ways. What do you guys think?

AUDIENCE: It can emit neutrons.

MICHAEL SHORT: It can emit neutrons if it's really unstable, at which point it would just go down a neutron number. Or how else could it decay?

AUDIENCE: Alpha decay.

MICHAEL SHORT: Alpha decay. Let's see, yeah, a lot of those will-- the heavier ones tend to do alpha decay. What would it do at alpha decay? For alpha, I guess it will be going that direction, right? You know what? I'm not going to rule that out yet. So let's go with that. How else could they decay?

AUDIENCE: Through beta decay.

MICHAEL SHORT: Through beta decay, let's say in that direction. Pretty much all these happen, just not necessarily in this order. When you have a really, really asymmetric nucleus, a lot of these fission products will emit neutrons almost instantaneously in the realm of like 10 to the minus 17 seconds, some incredibly short timeline. You will start to decay downwards a little bit. But you're not quite at the stability line, which is why a lot of the fission products then go on. And they deposit their kinetic energy by bouncing around the different atoms in material creating heat. But a lot of them will also send off betas or gammas.

And it may take 10 to the minus 13 seconds for them to whatever the half-life of that particular isotope is. And after around, let's say, 10 to the minus 10 to 10 to the minus 6 seconds, depending on the isotope in the medium, those two fission products will stop. And let's just say that they stop there. So the whole process of fission, it's actually quite a compound process.

First, the neutron is absorbed, forming a compound nucleus. Then it splits apart. Then those individual fission products undergo whatever decays suit them best. And that's the source of the neutrons in fission. Sometimes one of those fission products might be particularly unstable. And it might send off two neutrons. In other cases, though I don't know of one off the top my head, it might be none. But this is the whole timeline of events in fission and the justification for why this happens straight from the first month of 22.01.

And I wanted to pull up some of the nuclear data so you can see what these values tend to look like and also where to find them. I'm going to do that screen cloning thing again. There we go. So I've already pre-pulled up the JANIS library. I've already clicked on uranium-235. Thanks to you guys, I have all the data now on my shirt so you can see a little better. I also have it on the screen.

So let's look at this value right here, nu bar total, neutron production. And I'll make it bigger so it's easier to see. Did I click on the right one? Yeah. So take a look at that. The total number of neutrons produced during U-235, for most energies it's hovering around the 2.4 or so. There's been arguments about whether it's 2.43 or 2.44. And that's a linear scale. That's not very helpful. Let's go to a logarithmic scale.

That's more like what I'm used to seeing. Most of the fission happens for U-235 in the thermal region, in the region where the neutrons are at values, let's say, the cutoff is usually about one electron volt or lower in average energy. And nu bar is fantastically constant at that level. Then as you go up and up in energy, you start to make more and more neutrons. Why do you guys think that would be the case? What are you doing to that compound nucleus as you increase the incoming neutron energy?

AUDIENCE: It's going to have more energy.

MICHAEL SHORT: It's going to have more energy itself. You might excite other nuclear states that can then lead to other sorts of decays or other neutron emission. So to me, that's the reason why, once you hit about 1 MeV, you can start to see a lot more neutrons being given off. The reason we usually treat this as a constant, notice I haven't given it an energy dependence, is because most of the fission that happens is at thermal energies.

For that, I want to show you the fission cross section. There are a lot of cross sections. And it's probably going to be on a different graph, because it's in different units. And this gives you a rough measure per atom, what's the probability of fission happening as a function of incoming neutron energy? At those high energies, you have relatively low cross sections, or low probabilities, of fission happening.

Then there's this crazy resonance region that looks like a sideways mustache. But then as you get down to the lower energy levels, it gets much more, in fact, exponentially more, likely that fission will happen. So almost all the fissioning in a light water reactor, or any sort of other thermal reactor, happens at thermal energies. And that's why we take nu bar as a constant. You don't have to, especially if you're analyzing what's called a fast reactor or a reactor whose neutron population remains fast on purpose.

And so with that, I want to launch into some of the different types of reactors that you might see. And you guys already did those calculations in problem set one, so I don't have to repeat them for you. Let's get right into the acronyms. So if you haven't figured this out already, nuclear is a pretty acronym dense field. Can anyone say they know all the acronyms on this slide? You're going to know about 90% of them in about 90 minutes. So it's OK. Or you'll have seen them at least. Any look completely unfamiliar?

AUDIENCE: Most of them.

MICHAEL SHORT: Most of them?

[LAUGHTER]

Well, let's knock them off. So [INAUDIBLE], last Thursday, already showed you the basic layout of a boiling water reactor, one of the types of light water reactors. And the reason that this is a thermal reactor is because it's full of water. Water, as we saw in our old q equation argument, is very good at stopping neutrons, because, if you guys remember this, the maximum change in energy that a neutron can get is related to alpha times its incoming energy. Or this alpha is just A minus 1 over A plus 1 squared. And I think this would actually be a 1 minus right there.

A is that mass number of whatever the neutrons are hitting. And that one comes directly from the neutron mass number. If you remember, this was the simplest reduction of the q equation, the generalized q equation for kinematics that we looked at. When I said let's do the general form, then OK, let's take the simplest form, neutron elastic scattering. Here's where it comes back. If a neutron hits water, which is made mostly of hydrogen, and A is 1, then it can transfer a maximum of all of its energy, let's say, to that hydrogen atom, therefore, giving the neutron no energy and thermalizing it or slowing it down very quickly.

To show you what one of these things actually looks like, that's the underside of a BWR. Did [INAUDIBLE] show you this before? OK. So you've already seen what this generally looks like. What about the turbine? Has anyone actually seen a turbine this size close up, a gigawatt electric turbine? I'm trying to see which one of those pixels is a person.

I don't see anything person-sized. There's a ladder that looks to be about 6 feet tall, so to give you guys a sense of scale of the sort of turbines that we say, oh, yeah, we draw a turbine on our diagram. Well, it's not actually that simple . These things take up entire hallways, or kind of airport hangar sized buildings. I've never seen one in the US, but I've seen one in Japan. It was a lot cleaner than this. But, otherwise, it looked pretty much the same.

And the way this actually works, for those who haven't taken any thermo classes yet, is this turbine is full of different sets of blades that are curved at an angle so that when steam shoots in, it transfers some of its energy to get the turbine rotating. And there's going to be a generator, kind of like an alternator, to generate the electricity there, which looks to be roughly 100 feet away. Just to give you a sense of scale for this stuff.

As [INAUDIBLE] showed you, a pressurized water reactor is another kind of light water reactor with what's called an indirect cycle. So this water stays pressurized. It also stays liquid, which is good for neutron moderation or slowing down. Because in addition to the probability of any interaction, some probability sigma, if you want to get the total reaction probability, you have to multiply by its number density to get a macroscopic cross section. This is why I introduce this stuff way at the beginning of class, so you'd have time to marinate in it and then bring it back and remember what it was all about.

And so every single reaction that goes on in a nuclear reactor has got its own cross section. We'll probably need half the board for this one. You can say you have a total microscopic cross section. These are all going to be as a function of neutron energy. What's the probability of anything happening at all? And these are actually tabulated up on the JANIS website. So let's unclick that, get rid of neutron production, and go all the way to the top, n comma total.

So all this stuff is written in nuclear reaction parlance, where if you have, let's say, n comma total, that means a neutron comes in, and that's the reaction that you're looking at. So this data file here, once I open it up, will give you the probability that anything at all will happen. You can see as the neutron energy gets higher, the probability of anything happening at all gets less, and less, and less. And it follows the shape of most of the other cross sections. And I'm going to leave this up right there.

You've also got a few different kinds of reactions. You can have a scatter. Let's call that scatter, which we've already said can either be elastic or inelastic. It may not matter to us from the point of view of neutron physics whether the collision is elastic or inelastic. All that matters is the neutron goes in, and a slower neutron comes out. Because what we're really concerned with here is tracking the full population of neutrons at any point in the reactor.

So we'll give this a position vector r, which has just got x, y, and z in it or whatever other coordinate system you might happen to use. I prefer Cartesian, because it makes sense. At every energy going in any direction, so we now have a solid angled vector that's got both theta and phi in it any given time. And the whole goal of what we're going to be doing today and all of next week is to find out, how do you solve for and simplify this population of neutrons?

Make sure to fill that in as velocity. Let's see. Let me get back to the cross sections and stuff. If we want to know how many neutrons are in a certain little volume element, in some d volume, in some certain little increment of energy, dE, traveling in some very small, solid angle, d omega, supposedly, if you have this function, then you know the direction, and location, and speed of every single neutron everywhere in the reactor.

And this is eventually what the goal of things like Ben and Kord's group does, the Computational Reactor Physics Group, is solve for this or a simplified version of it, over, and over, and over again for different sorts of geometries. And in order to do so, you need to know the rates of reactions of every kind of possible reaction that could take a neutron out of its current position, like if it happens to be moving, which most of them are, out of its current energy group. Which pretty much any reaction will cause the neutron to lose energy. What's the only reaction we've talked about where the neutron loses absolutely no energy? It's a type of scattering.

AUDIENCE: Forward scattering?

MICHAEL SHORT: Yep, exactly, forward scattering. So for forward scattering for that case where theta scattering equals 0. Again, you missed. The neutron didn't actually change direction at all. And, therefore, it didn't transfer any energy. But for everything else, for every other possible reaction, there's going to be an energy change associated with it and probably some corresponding change in angle, because a neutron can't just be moving, and hit something, and continue moving more slowly. There's got to be some change in momentum to balance along with that change of energy.

And it might slightly move in some different direction. And all this is happening as a function of time. As you can see, this gets pretty hairy pretty quick. That's why we put the full equation for this on our department t-shirts. But no one ever solves the full thing. What we're going to be going over is, how do you simplify it into something you can solve with a pen and paper or possibly a gigantic computer? But it's not impossible.

So inside this sigma total, we talked about different scattering. And then you could have absorption in all its different forms. What sort of reactions with a neutron would cause it to be absorbed?

AUDIENCE: Fission.

MICHAEL SHORT: Yes, fission. Thank you. So there's going to be some sigma fission cross section as a function of energy. And if it doesn't fizz, but it is absorbed, we'll call that capture. But capture can mean a whole bunch of different things too, right? There could be also a whole bunch of other nuclear reactions. There could be a reaction where one neutron comes in, two neutrons go out, like we looked at with beryllium in the Chadwick paper from the first day or like what actually does exist for this stuff.

So JANIS doesn't like multi-touch, so you have to bear with me on the small print on the screen. But there should be-- yep, here it is. Cross section number 16, there is a probability that one neutron goes in. That z right there is whatever your incoming particle happens to be. And in this case, we know it's a neutron, because we picked incident neutron data. And 2n means two neutrons come out. Let's plot that cross section.

You can see that the value is 0 until you hit about 4 or 5. Oh, it's actually 5.297781 MeV. So that's the q value at which this particular reaction happens to turn on. Might be responsible for a little bit of the blip in the total cross section. So technically, if we were to turn on every single cross section in this database, it should add up to that red line right there. So you can start to get an idea for how much of all the reactions of uranium-235 are due to fission. That's the one we want to exploit.

So let's find fission, right down there. Oh, wow, there's a 3n reaction. I want to see that. That doesn't happen until 12 MeV. Yeah. So neutrons don't typically tend to hit 12 MeV in a fission reactor. So this is a perfect flimsy pretext to bring in another variable. It's called the chi spectrum or what's called the fission birth spectrum. Yeah.

We've already talked about the neutrons being born and how many there were. But we didn't say at what energy they're born. In fusion reactors, this is pretty simple. You've already looked at this case. What is it? 14.7 MeV. That's a lot simpler. That's the fusion. For fission, it's not so simple. For the case of fission, if you draw energy versus this chi spectrum, it takes an interesting looking curve from about 1 MeV to about 10 MeV with the most likely energy being around 2 MeV.

So you aren't really going to get neutrons at the energy required for a 3n reaction in a regular fission reactor, just not going to happen. But it's good that you know that that exists. So let's go and answer my original question. How much of the total cross section is due to fission?

Most of it, especially at low energies. So let me get rid of those 2n and 3n ones, because they're kind of ruining our data. It's making it harder to see. That's better. So you can see at energies below around, let's say, a keV or so, almost all of the reactions happening with neutrons in uranium-235 are fission. This is part of what makes it such a particularly good isotope to use in reactors. The other one is, you can find it in the ground, unlike most of the other fissile isotopes, unlike, I think, any of the other fissile isotopes.

Thorium you got to breed and turn it into uranium-233. I'll have to think about that one. But then you start to look at, what are the other components of this cross section, like zn prime, inelastic scattering, which doesn't turn on until about 0.002 MeV, but later on is one of the major contributors and actually is responsible for-- wait, I've brought this for a reason. --is responsible for that little bump in the total cross section. So eventually all these things do matter.

But let's think about which ones we actually care about at all, because what we eventually want to do is develop some sort of neutron balance equation. If we can measure the change in the number of neutrons as a function of position, energy, angle, and time, as a function of time, and that would probably be a partial derivative, because there are like seven variables here. Before I write any equations, it's just going to be a measure of the gains minus the losses.

And while every particular reaction has its own cross section, there's only going to be a few that we care about. There will only be one or two types of reactions that can result in a gain of the neutron population into a certain volume with a certain energy with a certain angle. And for losses, there's only one we really care about, total, because any interaction with a neutron is going to cause that neutron to leave this little group of perfect position, energy, and angle.

So that's where we're going. We'll probably start down that route on Tuesday, because I promised you guys context today. You've all been to the MIT Research Reactor. A couple of you-- are you running it yet?

AUDIENCE: Yeah.

MICHAEL SHORT: Awesome. OK. Yeah. Yeah, so Sarah and Jared are doing that. Anyone else training or trained? No. I'd say folks are usually pretty scared when they find out MIT has a reactor. And they're even more scared when they find out you guys run it.

AUDIENCE: Yeah.

MICHAEL SHORT: What they don't realize is there's been basically no problems since 1954. The only one I know of is someone fell asleep at the controls once and forgot to push the Don't Call Fox News button, and it called Fox News or something. So there was a big story about, asleep at the helm, ignoring all of the alarms, and passive safety systems, and backup operators, and everything else that actually made sure that nothing happened.

But nowadays, correct me if I'm wrong, you actually have to get up every half hour, reach around a panel, and hit a button, right?

AUDIENCE: No. It's on console, but it beeps at you.

MICHAEL SHORT: Ah.

AUDIENCE: Yeah, it's pretty tiring.

MICHAEL SHORT: So you want to hit it before it beeps at you.

AUDIENCE: It's reminding you to take hourly logs.

MICHAEL SHORT: OK.

AUDIENCE: It does go off every half hour.

AUDIENCE: It is half hour, but you we don't do [INAUDIBLE].

MICHAEL SHORT: Ah, OK, yeah. I'd heard the button's every half hour. Gotcha. Cool. Yeah, so for all of you watching on camera or whatever, just know that these guys got it under control. So onto some gas cooled reactors and to explain some of these acronyms. There are some that use natural uranium, though pretty much all the ones in this country, you need to enrich the uranium to get enough U-235 to turn the reaction on. But you don't have to do that in every case.

And you'll also see these acronyms, LEU, MEU, or HEU, standing for Low, Medium, or High Enrichment. The accepted standard for what's low enriched uranium is 20% or below. An interesting fact, though, you can't have something at 19.99% enriched uranium and expect it to be low enriched uranium, because every measurement technique has some error. And what really determines if it's LEU is when an inspector comes and takes a sample, it better be below 20% including their error.

So you'll usually see 19.75% given as the LEU limit, because there's always some processing error, inhomogeneities, measurement error. Hedge your bets, pretty much. Like in England or the UK, the advanced gas reactors have been churning along for decades. They actually use CO2 as the coolant, which is relatively inert. And they use graphite as the moderator. So in this case, the coolant and the moderator are separate, unlike the light water reactors we have. So this way, the graphite, right here, just sits in solid form and slows down the neutrons, not quite as good as water, but pretty good.

There is an issue, though, that CO2, just like anything, has a natural decomposition reaction, where CO2 naturally is in equilibrium with CO and O2. And O2 plus graphite yields CO2 gas. Graphite was solid. In talking with a couple folks from the National Nuclear Laboratory, they said that 40 years later, when they took the caps off these reactors, a lot of that graphite was just gone with a good explanation. It vaporized very, very, very slowly over 40 years or so due to this natural recombination with whatever little bit of O2 is in equilibrium with CO2 and possibly some other leaks.

I'm sure I wouldn't have been told that if there was a leak. So I'd say the feasibility is high, because they've been running for almost half a century. The power density is very low. Why do you guys think that's the case? Yeah.

AUDIENCE: [INAUDIBLE]

MICHAEL SHORT: Mm-hm.

AUDIENCE: [INAUDIBLE]

MICHAEL SHORT: Absolutely. So well, let's say, you need the same cooling capacity, but you're right. CO2, even if pressurized, is not as good a heat transfer medium as water. Water is dense. It's also got one of the highest heat capacities of anything we've ever seen. The other reason is right here. If you want enough reaction density, then it not only matters what the per atom density is, but what the number density is. And if you're using gaseous CO2 coolant, even if it's pressurized, there are fewer reactions happening per unit volume, because there are few CO2 molecules per unit volume than water would have.

So that's why we pressurize our light water reactors, to keep water in its liquid state where it's a great heat absorber, takes a lot of energy to boil it, and it's really dense so it's a very effective dense moderator. These have been around forever. Let me think. When did Windscale happen? Windscale was also the source of an interesting fire that you guys might want to know about. It's one of those only nuclear disasters that hit 7 on the arbitrary unit scale. I don't quite know how they determine what's a seven.

But there was a fire at the Windscale plant due to the build up of what's called Wigner energy. It turns out that when neutrons go slamming around in the graphite, they leave behind radiation damage. And when my family always asks me to explain, what do you do for a living? And I can only think, well, they don't know radiation damage. They've watched Harry Potter. I'd like to say, radiation, like dark magic, leaves traces. Well, it leaves traces in the graphite in the form of atomic defects, which took energy to create.

So by causing damage to the graphite, you store energy in it, which is known as Wigner energy. And you can store so much that it just catches fire and explodes sometimes. That's what happened here at Windscale. 11 tons of uranium ended up burning, because all of a sudden, the temperature in the graphite just started going up for no reason, no reason that they understood at the time.

It turns out that they had built up enough radiation damage energy that it started releasing more heat. And releasing more heat caused more of that energy to be released, and it was self-perpetuating until it just caught fire and burned 11 tons of uranium out in the countryside. This was 1957. So again, a 7 on the scale with no units of nuclear disasters. Argue it's probably not as bad as Chernobyl, so they might want a little bit of resolution in that scale.

There's another type of gas cool reactor called the Pebble Bed Modular Reactor, a much more up and coming one, where each fuel element-- you don't have fuel rods. You've actually got little pebbles full of tiny kernels of fuel. So you've got a built-in graphite moderator tennis ball sized thing with lots of little grains of sand of UO2 cooled by a bed of flowing helium or something like that. And then that helium, or the other gas, transfers heat to water, which goes in to make steam and goes into the turbine like I showed you before.

So this is what the fuel actually looks like. Inside each one of these tennis ball spheres of mostly graphite, there's these little kernels of uranium dioxide about a half a millimeter across covered in layers of silicon carbide, a really strong and dense material that keeps the fission products in, because the biggest danger from nuclear fuel is the highly radioactive fission products that due to their instability are giving off all sorts of awful, for anywhere from milliseconds to mega years, after reactor operation. And so if you keep those out of the coolant, then the coolant stays relatively nonradioactive. And it's safe to do things like maintain the plant.

Then there's the very high temperature reactor, the ultimate in acronym creativity. It operates at a very high temperature, which has been steadily decreasing over time, as reality has caught up to expectations. When I first got into this field, they were saying, we're going to run this at 1100 Celsius. Then I started studying material science. And I was like, yeah, nothing wants to be 1100 Celsius. By that time, they downgraded it to 1000. Now they've asymptoted it at around 800 or 850 due to some actual problems in operating things in helium.

It's not the helium itself, but the impurities in the helium that could really mess you up. And the sorts of alloys that they need to get this working, these nickel superalloys, like Alloy 230, they can slightly carburize or decarburize depending on the amount of carbon in the helium coolant. Either way you go, you lose the strength that you need. So I'll say feasibility is low to medium, because, well, we haven't really seen one of these yet.

Then onto water cooled reactors. Has anyone here heard of the reactors they have in Canada, the CANDU reactors? That's my favorite acronym. I hope that was intentional. It what?

AUDIENCE: It's convenient.

MICHAEL SHORT: Yeah. [LAUGHS] It's not like the-- well, they're not sorry about anything, but whatever. At any rate, one of the nice features about this is you can actually use natural uranium, because the moderator is heavy water. You have to look into what the sort of cross sections are. Even though deuterium won't slow down neutrons as much as hydrogen will-- where did my alpha thing-- oh, it was right here all along. Even though A is 2 instead of 1 for deuterium, it's absorption cross section, or specifically-- yeah, because it doesn't fission.

Its absorption cross section is way lower than that of water. It actually functions as a better moderator, because fewer of those collisions are absorption. And because you have a better neutron population and less absorption, you don't need to enrich your uranium. You also don't need to pressurize your moderator.

So you can flow some other coolant through these pressure tubes and just have a big tank of close to something room temperature unpressurized D2O as your moderator. The problem with that is D2O is expensive. Anyone priced out deuterium oxide before? Probably have at the reactor, because I know you have drums of it.

AUDIENCE: It's like a couple thousand per kilogram.

MICHAEL SHORT: A couple thousand a kilo, it's an expensive bottle of water. It'll also mess you up if you drink it, because a lot of it, even if it's crystal clear, filtered D2O, a lot of what the cellular machinery depends on the diffusion coefficients of various things in water, those solutes in water. And if you change the mass of the water, then the diffusion coefficients of the water itself, as well as the things in it, will change. And if you depend on, let's say, exact sodium and potassium concentrations for your nerves to function, a little change in that can go a long way towards giving you a bad day.

And actually, we have a little piece of one of these pressure tubes upstairs if anyone wants to take a look. There's all these sealed fuel bundles inside what they call a calandria tube, just a pressurized tube that's horizontal. The problem with some of these is if these spacers get knocked out of place, which they do all the time, those tubes can start to creep downward and get a little harder to cool or touch the sides and change thermal. And now I'm getting into material science. It's a mess.

Then there's the old RBMK, the reactor that caused Chernobyl. You can also use natural uranium or low enriched uranium here. The problem though that led to Chernobyl-- one of the many problems that led to Chernobyl was, you've got all this moderator right here. So if you lose your coolant, let's say you had a light water reactor and your coolant goes away, your moderator also goes away, which means your neutrons don't slow down anymore. That one reaction is messing up. There we go. Which means your neutrons don't slow down anymore, which means the probability of fission happening could be like 10,000 times lower.

So losing coolant in a light water reactor, temperature might go up, but it's not going to give you a nuclear bad day. In the RBMK reactor, it will and it did. And in addition, the control rods, which were supposed to shut down the reaction, made of things like boron 4 carbide, or hafnium, or something with a really high capture cross section were tipped with graphite to help them ease in. So you've got moderator tipped rods, which induce additional moderation, which helps slow down the neutrons even more to where they fission even better. And that's what led to what's called a positive feedback coefficient.

So the more you tried to insert the control rods and the more you tried to fix things, the worse things got in the nuclear sense. And in something like a quarter of a second, the reactor power went up by like 35,000 times. And we'll do a millisecond by millisecond rundown of what happened in Chernobyl after we do all this neutron physics stuff when you'll be better equipped to understand it. But suffice to say, there were some positive coefficients here that are to be avoided at all costs in all nuclear reactor design.

In the actual reactor hall you can go and stand on one of these things. It's a very different design from what you're used to. I don't think anyone would let you stand on top of a pressure vessel. First, your shoes would melt, because they're usually at like 300 Celsius or so. And second of all, you'd probably get a little too much radiation. But this is actually what an RBMK reactor hall looks like for one of the units that didn't blow up. There were multiple units at that site.

Then there's the supercritical water reactor. Let's say you want to run at higher temperatures than regular water will allow you to. You can pressurize it so much that water goes beyond the supercritical point in the phase sense and starts to behave not like liquid, not like a gas, but somewhere in between, something that's really, really dense, so getting towards the density of water, not quite, which means it's still a great moderator, but still can cool the materials quite well to extract heat to make power and so on and so on. Yeah.

AUDIENCE: So supercritical refers to the coolant not the neutrons?

MICHAEL SHORT: Good question. For a supercritical water reactor, it most definitely refers to the coolant. It's the phase of the coolant where it's beyond the liquid gas separation line, and it's just something in between. Any of these reactors can go supercritical, where you're producing more neutrons than you're consuming. And that is a nuclear bad day. But the supercritical water reactor does not refer to neutron population, just a coolant. Good question. It's never come up before. But it's like, should have thought of that.

And so then my favorite, liquid metal reactors, like LBE, or Lead-Bismuth Eutectic. It's a low melting point alloy of lead and bismuth. Lead melts at around 330 Celsius, bismuth 200 something. Put them together, and it's like a low temperature solder. It melts at 123.5 Celsius. You can melt it in a frying pan. This is nice, because you don't want your coolant to freeze when you're trying to cool your reactor, because imagine something happens, you lose power. The coolant freezes somewhere outside the core. You can't get the core cool again. That's called a loss of flow accident that can lead to a really bad day.

And the lower your melting point is the better. Sodium potassium is already molten to begin with. Sodium melts at like 90 Celsius. And when you add two different metals together, you almost always lower the melting point of the combination. In this case, forming what's called the eutectic, or a lowest possible melting point alloy. The sodium fast reactor has a number of advantages, like you don't really need any pressure. As long as you have a cover gas keeping the sodium from reacting with anything, like the moisture in the air, or any errant water in the room, you can just circulate it through the core.

And liquid metals are awesome heat conductors. They might not have the best heat capacity, as in how much energy per gram they can store like water. But they're really good conductors with very high thermal conductivity. They also are really good at not slowing down neutrons. So these tend to be what's called fast reactors that rely on the ability of other isotopes of uranium, like uranium-238, to undergo what's called fast fission.

And I want to show you what that looks like. Let's pull up U-238 and look at its fission cross section. And you might find that it should look a fair bit different. So we'll go down to number 18 to fission cross section, very, very different. So U-238 is pretty terrible at fission at low energies. It's pretty good at capturing neutrons. This is where we get plutonium-239, like you guys saw on the exam.

But then you go to really high energies and all of a sudden, it gets pretty good at undergoing fission on its own. And so the basis behind a lot of fast reactors is a combination of making their own fuel and the fact that uranium-238 fast fissions even better than it thermal fissions. So something good for you to know, even though it's not a fissile fuel, that's light water reactor people talking. You can get it to fission if the neutron populations higher.

Now, there's some problems with this. It takes some time for neutrons to slow down from 1 to 10 MeV to about 0.025 eV. If your neutrons don't need to slow down and travel anywhere, and pretty much all they have to do is be born and absorbed by a nearby uranium atom, the feedback time is faster in these sorts of reactors. They're inherently more difficult to control.

And you can't use normal physics like thermal expansion of things that might happen on the order of micro to nanoseconds if it takes less time than that for one neutron to be born and find another uranium atom. You can still use it somewhat, but not quite as much. So it's something to note backed up by nuclear data.

And that's what one of them actually looks like. These things have been built. That's a blob of liquid sodium on the Monju reactor in Japan. And where I was all last week in Russia, they actually have fleets of fast reactors. Their BN-300 and BN-600 reactors are 300 and 600 megawatt sodium cooled reactors. One of them in the Chelyabinsk region they use pretty much for desalination down in the center of Russia, where there's no oceans nearby and probably dirty water. They actually use that to make clean water. They also use this for power production and for radiation damage studies.

So when it comes to radiation material science, these fast reactors are really where it's at. Yeah, you just noticed the bottom. I went to Belgium, to their national nuclear labs, where they have a slowing sodium test loop. It's not a reactor, but it's like a thermal hydraulics and materials test loop. And I asked a simple question. Where's the bathroom? And they started laughing at me. And they said, we're not putting any plumbing in a sodium loop building. You'll have to go to the next building over.

And that's when I noticed, there weren't any sprinkler systems or toilets. But every 15 or 20 feet, there was a giant barrel of sand. That's the fire extinguisher for a liquid metal fire is you just cover it with sand, absorb the heat, keep the air out, the moisture out, wick away the moisture or whatever else sand does. I don't know. But you can't use normal fire extinguishers to put out a sodium fire.

AUDIENCE: When you said sand, I thought of kitty litter.

MICHAEL SHORT: Ah. I don't know if that would work.

[LAUGHTER]

I guess it's worth a shot.

[LAUGHTER]

With glasses, and safety, and stuff, of course. And the ones that I spent the most time working on, like I showed you in the paper yesterday, is the lead or lead-bismuth fast reactor. This one does not have the disadvantages of exploding like sodium. It does have the disadvantage, like I showed you yesterday, of corroding everything, pretty much everything. And so the one thing keeping this thing back was corrosion. And I say the ultimate temperature is medium, but higher soon. Hopefully, someone picks up our work and is like, yeah, that was a good idea, because we think it can raise the outlet temperature of a lead-bismuth reactor by like 100 Celsius as long as some other unforeseen problem doesn't pop up, and we don't quite know yet.

These things also already exist in the form of the Alfa Class attack submarines from the Soviet Union. These are the only subs that can outrun a torpedo. So you know that old algebra problem, if person A leaves Pittsburgh at 40 miles an hour and person B leaves Boston at 30 miles an hour, where do the trains collide or I forget how it actually ends? Well, in the end, if a torpedo leaves an American sub at whatever speed and the Alfa Class submarine notices it, how close do they have to be before the torpedo runs out of gas?

So what I was told by the designer of these subs, a fellow by the name of Georgy Toshinsky, when he came here to talk about his experience with these lead-bismuth reactors is, there is a button on the sub that's the Forget About Safety, It's a Torpedo button. Because if you're underwater in a lead-bismuth reactor and a torpedo is heading at you, you have a choice between maybe dying in a nuclear catastrophe and definitely dying in a torpedo explosion. Well, that button is the I Like Those Odds button. And you just give full power to the engines and whatever else happens, happens. The point is, you may be able to outrun the torpedo.

And quite popular nowadays, especially in this department, is molten salt cooled reactors that actually use liquid salt, not dissolved, but molten salt itself as the coolant. It doesn't have as many of the corrosion problems as lead or the exploding problems as sodium. It does have a high melting point problem though. They tend to melt at around 450 degrees Celsius. But there's one pretty cool feature. You can dissolve uranium in them. So remember how in light water reactors the coolant is also the moderator? In molten salt reactors, the coolant is also the fuel, because you can have principally uranium and lithium fluoride salt co-dissolved in each other.

And the way you make a reactor is you just flow a bunch of that salt into nearby pipes. And then you get less, what's called, neutron leakage, where in each of these pipes once in a while uranium will give off a few neutrons. Most of them will just come out the other ends of the pipes, and you won't have a reaction. When you put a whole bunch of molten salt together, most of those neutrons find other molten salt. And the reaction proceeds.

And it's got some neat safety features. Like if something goes wrong, just break open a pipe. All the salt spills out, becoming subcritical, because leakage goes up. It freezes pretty quickly, and then you must deal with it. But it's not a big deal to deal with it if it's already solid and not critical.

So it's actually five of. It's zero of five of. I'll stop here. On Tuesday, we'll keep developing the many, many different variables we'll need to write down the neutron transport equation, at which point you'll be qualified to read the t-shirts that this department prints out. And then we'll simplify it so you can actually solve the equation.