MICHAEL SHORT: I think I might actually use all 16 colors today. Oh no, this is the most satisfying day. Whereas Tuesday was probably the most mathematically intense, because we developed this equation right here, today is going to be the most satisfying, because we are going to cancel out just about every term, leaving a homogeneous, infinite reactor criticality condition. So we will go over today, how do you go from this, to what is criticality in a reactor?

So I want to get a couple of variables up over here to remind you guys. We had this variable flux of r, e, omega, t in the number of neutrons per centimeter squared per second traveling through something. And we also had its corresponding non-angular dependent term, on just r, e, t, if we don't care what angle things go through.

We've got a corresponding variable called current. So I'll put this as flux. We have current j, r, e, omega, t, and its corresponding, we don't care about angle form.

And today, what we're going to do is first go over this equation again so that we understand all of its parts. And there are more parts here than are in the reading. If you remember, that's because I wanted to show you how all of these terms are created. Just about every one of these terms, except the external source and the flow through some surface, has the form of some multiplier, times the integral over all possible variables that we care about, times a reaction rate d stuff, where this reaction rate is always going to be some cross-section, times some flux.

So when you look at this equation using that template, it's actually not so bad. So let's go through each of these pieces right here. And then we're going to start simplifying things. And this board's going to look like some sort of rainbow explosion. But all that's going to be left is a much simpler form of the neutron diffusion equation.

So we've got our time-dependent term right here, where I've stuck in this variable flux, instead of the number of neutrons n, because we know that flux is the number of neutrons, times the speed at which they're moving. And just to check our units, flux should be in neutrons per centimeters squared per second.

And n is in neutrons per cubic centimeter. And velocity is in centimeters per second. So the units check out. That's why I made that substitution right there.

And this way, everything is in terms of little fi, the flux. We have our first term here. I think I'll have a labeling color. That'll make things a little easier to understand. Which is due to regular old fission.

In this case, we have new, the number of neutrons created per fission, times chi, the sort of fission birth spectrum, or at what energy the neutrons are born. Over 4pi would account for all different angles in which they could go out, times the integral over our whole control volume, and all other energies in angles.

If you remember now, we're trying to track the number of neutrons in some small energy group, e, traveling in some small direction, omega. And those have little vector things on it at some specific position as a function of time. So in order to figure out how many neutrons are entering our group from fission, we need to know, what are all the fissions happening in all the other groups?

I've also escalated this problem a little bit to not assume that the reactor's homogeneous. So I've added an r, or a spatial dependence for every cross-section here, which means that as you move through the reactor, you might encounter different materials. You almost certainly will, unless your reactor has been in a blender.

So except for that case, you would actually have different cross sections in different parts of the reactor. So all of a sudden, this is starting to get awfully interesting, or messy depending on what you want to think about it. There is the external source, which is actually a real phenomenon, because reactors do stick in those californium kickstarter sources. So for some amount of time, there is an external source of neutrons, giving them out with some positional energy angle and time dependence.

So let's call this the kickstarter source. There's this term right here, the nin reactions. So these are other reactions where it's absorb a neutron, and give off anywhere between 2 and 4 neutrons. Beyond that, it's just not energetically possible in a fission reactor. But don't undergo fission.

They have their own cross sections, their own birth spectrum. And I've stuck in something right here, if we have summing over all possible i, where you have this reaction be n in reaction, where 1 neutron goes in, and i neutrons come out. You've got to multiply by the number of neutrons per reaction.

For fission, that was new. For an nin reaction, that's just i. But otherwise, the term looks the same. You have your multiplier, your birth spectrum, your 4pi, your integral over stuff, your unique cross-section, and the flux.

And these two together give you a reaction rate. I've just written all of the differentials as d stuff, because it takes a lot of time to write those over and over again. And then we have our photo fission term, where gamma rays of sufficiently high energy can also induce fission external to the neutrons.

The term looks exactly the same. There's going to be some new for photofission, some birth spectrum for photofission, some cross section for photofission, and the same flux that we're using everywhere else. Then we had what's called the in scattering term, where neutrons can undergo scattering, lose some energy, and enter our group from somewhere else.

That's why we have those en omega primes, because it's some other energy group. And we have to account for all of those energy groups. That's why we have this integral there.

And it looks very much the same. There's a scattering cross-section. And that should actually be an e prime right there. Make sure I'm not missing any more of those inside the integral. That's all good. That's a prime, good.

There's also a flux. And then there was this probability function that a given neutron starting off an energy e prime omega prime, ends up scattering into r energy, e and omega. So this would be the other one. And this would be r group.

But otherwise, the term looks very much the same. And that takes care of all the possible gains of neutrons into r group. The losses are a fair bit simpler. There is reaction of absolutely any kind.

Let's say this would be the total cross-section, which says that if a neutron undergoes any reaction at all, it's going to lose energy and go out of our energy group d, e. Notice here that these are all-- these energies and omega is our all r group, because we only care about how many neutrons in r group undergo a reaction and leave. And the form is very simple--

integrate over volume, energy, and direction, times a cross-section, times a flux, just like all the other ones.

Then the only difference in one right here is what we'll call leakage. These are neutrons moving out of whatever control surface that we're looking at. And this can be some arbitrarily complex control surface in 3D. I don't really know how to draw a blob in 3D. But at every point on that blob, there's going to be a normal vector.

And you can then take the current of neutrons traveling out that normal vector, and figure out how much of that is actually leaving our surface ds. The one problem we had is that everything here is in terms of volume, volume, volume, surface. So we don't have all the same terms, because once we have everything in the same variables we can start to make some pretty crazy simplifications.

The last thing we did is we invoked the divergence theorem, that says that the surface integral of some variable FdS is the same as the volume integral of the divergence of that variable dV. So I remember there was some snickering last time, because you probably haven't seen this since, was it 1801 or 1802?

1802, OK, that makes sense, because divergence usually has more than one variable associated with it. I'll include the dot, because that's what makes that divergence. So we can then rewrite this term. Let's start our simplification colors.

That's our divergence theorem. So let's get rid of it in this form, and call it minus integral over all that stuff. Then we'll have dot, little j, to be careful, r, e, omega, t, d soft. So for every step, I"m going to use a different color so you can see which simplification led to how much crossing stuff out.

And so like I said, this board's going to look like a rainbow explosion. But then we'll rewrite it at the end. And it's going to look a whole lot simpler. So now, let's start making some simplifications.

Let's say you're an actual reactor designer, and all you care about is how many neutrons are here. Of the variables here, which one do you think we care the least about? Angle, I mean do we really care which direction the neutrons are going? No, we pretty much care, where are they?

And are they causing fission or getting absorbed? So let's start our simplification board. And in blue, we'll neglect angle. This is where it starts to get fun.

So in this case, we'll just perform the omega integral over all angles. We just neglect angle here. We forget the omega integral, forget omega there. Away goes the 4pi, because we've integrated overall 4pi star radians, or all solid angle. Let's just keep going.

Forget the 4pi, forget the omega, forget the omega, forget the 4pi and the omega, and the omega. Same thing here-- forget the omega in the scattering kernel, forget it in the flux, forget it there, forget it there, and there as well. OK, we've now completely eliminated one variable. And all we had to do is ditch the 4pi and one of the integrals.

What next? We're tracking right now every possible position, every possible energy, at every possible time. If you want to know, what is your flux going to be in the reactor at steady state, what variable do you attack next?

Time. So let's just say this reactor is at steady state. That's going to invoke a few things. For one, it's going to ditch the entire steady state term.

We're going to get rid of all the ts in all the fluxes. This shouldn't take too long to do. I think that's all of them.

And the third thing is if this reactor is at steady state, chances are we've taken our kickstarter source out, because we just needed it to get it going. But the reactor should be self-sustaining once it's at steady state. So let's just get rid of our source term.

I just want to make sure I didn't miss any here. OK, next up, let's go with green. What else do you think we can simplify about this problem?

Well, if you look far enough away from the reactor, we can make an assumption that the reactor is roughly homogeneous. In some cases, it's not so good of an assumption, like very close to anything that has a huge absorption cross-section. Now, I want to explain the physics behind this.

If the neutrons travel a very long distance through any group of materials, then those materials will appear to be roughly homogeneous to the neutrons. If, however, the neutrons travel through something that's very different from the materials around it, then that homogeneous assumption breaks down. So in what locations in a nuclear reactor do you think you cannot treat the system as homogeneous?

Where do the properties of materials suddenly change by a huge amount? Yeah, Luke?

AUDIENCE: Control rods.

MICHAEL SHORT: Control rods, right, so let's say it's bad for control rods. Where else? How about the fuel? All of a sudden, you're moving from a bunch of structural materials where sigma fission equals 0, to the fuel where sigma fission, like you saw on the test, can be like 500 barns, which even though it's got a very small exponent in front of it, 10 to the minus 22 centimeters squared, it's still pretty significant.

So this assumption breaks down around the control rods and around the fuel. But we can get around this. Let's analyze the simplest, craziest possible reactor, which would be a molten salt fueled reactor. It's just a blob of 700 Celsius goo that's got its fuel, coolant, and control rods all built in.

So if we assume that the reactor is homogeneous, which is a pretty good assumption for molten salt fueled reactors, because the fuel's dissolved in the coolant. And it builds up its own fission product poison. So it's got some of its own control rods kind of built in. Usually, we'll have other extra ones too, but whatever.

Then we can start to really simplify things. If we get rid of any homogeneity assumptions, we cannot necessarily get rid of the r in the flux, because even if the reactor's homogeneous it still might have boundaries. So you might be able to approximate it as just a cylinder or a slab of uniform materials.

But if we were to get rid of the r's in the flux term, that would mean that as we graph flux as a function of distance, it would look like that, including infinitely far away from the reactor. Now, is was that true? Absolutely not, so I don't want to leave that up for anyone. We'll fill in what these graphs look like a little later, just leave them there for now.

We can get rid of some of the other r's though, like these cross sections. If the reactor is actually homogeneous, then the cross section is the same everywhere because the materials are the same everywhere. So we can get rid of the r's here, the r's here, and there, and there. And that's it, I think. I don't think I missed any, good.

Next up-- if this reactor is homogeneous, then does it really matter at which location we're taking this balance? Does it really matter which little volume element we're looking at? We say these equations are-- we'll call them volume identical, which means if this same equation is satisfied at any point in the reactor, we don't need to do the volume integral over the whole reactor. It's not like it's going to change anywhere we go.

So forget the volume integrals. Hopefully, you guys see where I'm going with this. And I've never tried teaching it like this rainbow explosion before. But I'm kind of excited to see how it turns out.

So already like 2/3 of the stuff that we had written are gone. What's the only variable left that we can go after? What's the only color left that I haven't really used?

Energy, so we can make a couple of different assumptions. This equation as it is not yet really analytically solvable, because a lot of these energy dependent terms don't have analytical solutions, or even forms like the cross sections. But we can start attacking energy.

Hopefully, this is different enough from white. Yeah, is that big enough difference for you guys to see? Good, OK, we can start doing this in a few different ways.

I want to mention what they are. And then we're going to do the easiest one. So the way it's done for real, like in the computational reactor physics group, is you can discretize the energy into a bunch of little energy groups. So you can write this equation for every little energy group, and assume that along this energy scale, ranging from your maximum energy to probably thermal energy-- 025, let's do this clearly with thick chalk.

There we go. You can then discretize into some little energy group. Let's say that's egi, that's egi plus 1, and so on, and so on. And depending on the type of reactor that you're looking at, and the energy resolution that you need, you choose the number of energy groups accordingly.

Does anyone happen to know for a light water reactor, how many energy groups do you think we need to model a light water reactor? The answer might surprise you. It's just two, actually.

All we care about-- so let's say this would be for the general case. All we care about for a light water reactor is, are your neutrons thermal? Or are they not?

Because the neutrons that are not thermal are not contributing to fission that much. They are just a little bit. And you can account for those. But pretty much, they're not.

Once the neutron slowdown down to get thermal, in the range from, let's say, about an EV to that temperature-- took a surprising amount of time to write with sidewalk chalk-- then you've got things that are about 500 or 1,000 times more likely to undergo fission. And so all you care about is the neutrons are all born. They're all born right about here.

And they scatter, and bounce around. And you don't care, because they're just in this not thermal region. And when they enter the thermal region, you start tracking them, because those are the ones that really count for fission.

And if you actually look up the specifications for the AP 1000, this is a modern reactor under construction in many different places in the world. When you see, how do they do the neutron analysis? Two group approximation. So this isn't just an academic exercise to make it easier for sophomores to understand.

This is actually something that's done for real reactors. So if you ever felt like I'm making it too simple, no, no, no, I'm simplifying it down to what's really done. And I will get you that specification so you can see what Westinghouse says, like, this is how we design the reactor. We made a two group simplification in many cases.

So you can discretize. You can forget it, which we're going to call the one group approximation. Or you can-- let's say two group is the other one that we're actually going to tackle. We're going to do this one, forget energy.

But we're not really going to forget energy, because you can't just pick an energy, and pick a cross-section, and say, OK, that's the cross section we're going to use. If most cross sections have the following form-- if this is log of energy, and this is log of sigma, and it goes something like that, what energy do you pick?

Go ahead. Tell me. Which energy do you pick? Anyone want to wager a guess?

AUDIENCE: The ones before or after the big squigglies.

MICHAEL SHORT: The ones before or after the big squiggles. I don't think that's correct, because if you do it this, then you're going away underestimate fission. If you do it here, you're going to way overestimate fission, or whatever other reaction you have. We didn't say which reaction this is. The rest of you who are silent and afraid to speak up, you're actually correct.

I wouldn't actually pick any single value here. What you need to do is find some sort of average cross-section for whatever reaction that accurately represents the number of reactions happening in the system. And in order to do that, you have to come up with some average cross-section for whatever reaction you have by integrating over your whole energy range of the energy dependent cross-section as a function of energy, times your flux de over-- does this look familiar from 1801 or 2 as well, what's the average value of some function? Little bit?

Well, we'll bring it back here now. So retrieve it from cold storage in your memories, because this is how actual cross sections are averaged. For whatever energy range you're picking-- I'm going to make this a little more general. I won't say zero. I'll just say your minimum energy for your group.

So now, this equation is general for the multi group and the one group and two group method. For whatever cross-section you want to pick and whatever energy range you're looking at, you take the actual data and perform an average for the fast and thermal delineation, where, let's say this is fast and this is thermal, you would have two different averages. Maybe this average would be right there.

You know what? Let's use white so it actually has some contrast. So this would be one value of the cross-section. And maybe the next average would be right there.

So you simplify this absolutely non analytical form of your complicated cross-section to just a couple of values. Maybe we'll call that average sigma fast. And we'll call that average sigma thermal.

So using this analogy and this color, we can then say, we're going to take an average new, an average chi, get rid of the energies, because we can perform the same energy average integration on every quantity with energy dependence. So all we do is we put a bar there, ditch the energies, ditch the energies. And let's just say that flux is going to be what it is.

Same thing here-- yeah, same thing there, and there, and there, there, there, there, here, and here, and here. And there is a cross-section. There is an energy. There is an energy.

There is a cross-section. We don't care about those anymore. And there's a couple of other implications of this energy simplification.

What is the birth spectrum now? What's the probability that a neutron is born in our energy group which contains all energies? 1, OK, so forget the chi, and that one, and that one. And what about this scattering kernel?

What's the probability that a neutron scatters from any other energy which is already in our group into our group, which contains all energies?

AUDIENCE: 1.

MICHAEL SHORT: Yeah, scattering no longer matters when you do the one group approximation, because if the neutron loses some of its energy, it's still in our energy group, because our energy group contains all energies. So forget the scattering kernel. And forget the energy integrals. What are we actually left with?

Not much. There's no green in here yet. Good, because I need to do one more thing. There is no more green. Oh, we did green. We did time.

OK, green, red, orange-- is this the orange I used? Dammit. OK, we use those. Purple, no, we've used it. Oh my God. We've used both blues. Bright yellow. Yeah?

AUDIENCE: [INAUDIBLE]

MICHAEL SHORT: Yes. Chi is the fission birth spectrum, the probability that a neutron is born at any given energy. But because all neutrons are born in our energy group, which contains all energies, then that just becomes one and goes away. There's no birth spectrum, because they're just born in our group. Does that makes sense?

OK, I think I found the actual only color, besides black on a black chalkboard, and white which we already have, that I have left. We also have a slightly darker shade of gray. But I'm literally out. This worked out awesome, because there's one more thing that we want to deal with.

What do we even have left? All right, what is the one term that is not in all of the same variables as the others? That current, that j. What do we do about that?

So we're going-- sorry?

AUDIENCE: The F e to e.

MICHAEL SHORT: The F e to e-- so, actually, I'll recreate some of our variables here, because there's a lot of them. So our F of e prime to e is what's called the scattering kernel. And that's the probability that a neutron scatters from some other energy group, e prime, into ours in e about de.

And chi of e is the fission birth spectrum. And just for completeness, knew of e is our neutron multiplier, or neutrons per fission. And I think that gives a pretty complete explanation of what's up here. So now, let's figure out how to deal with the current term.

This is when we make one of the biggest approximations here, and go from what's called the neutron transport equation, which is a fully accurate physical model of what's really going on, to the neutron diffusion equation. And this is where it gets really fun. You don't assume that neutrons are subatomic particles that are whizzing about and knocking off of everything else.

You then treat the neutrons kind of like a gas, or like a chemical. And you just say that it follows the laws of diffusion. Again, this works out very well, except for places where cross sections suddenly change, like near control rods or near fuel.

But for most of the reactor, especially if we have a molten salt fuel reactor, we can invoke what's called Fick's law. Does this sound familiar to anyone? Fick's law diffusion, 3091 or 5111. It's the change of a chemical down a density or a concentration gradient. So, yeah, you've got the idea.

What Fick's law says is that the current-- or let's say the diffusion current or the neutron current-- is going to be equal to some diffusion coefficient, times the gradient of whatever chemical concentration you've got. Let me put the c in there. So right here, this would be the current. I'll label it in a different color.

This would be your variable of interest. Maybe c is for concentration, or phi is for flux. Oh, that reminds me, where are those bars on our flux? Which term did we do?

Energy-- where's my slightly lighter blue over here? All of these phi's become capital, because we've gotten rid of all the angular, and energy, and everything dependence. Oh, angular dependence-- neglect omega, that should be dark blue.

Omega goes away. And the fluxes become capital. So many terms to keep track of. Luckily, you will never have to. And then the j becomes a capital J.

Did I miss any phi's here? No, because that one was already gone, cool. All right, so we can use Fick's law, and transform the current into something related to flux.

And what we're saying here is that we're getting rid of the true physics, which is that there's some fixed neutron current. And we're saying that neutrons behave kind of like a gas, or a chemical in solution. And so in yellow, we can ditch our current related term, and rewrite it. We don't have any integrals left, as negative del squared phi.

I think the only variable left is r, not too bad. Now, we have a second order linear differential equation describing the flow of neutrons in the system. And so we actually have something that we can solve for flux. I think it's time to rewrite it. Wouldn't you say?

This has been fun. So let's rewrite what's left. Make sure you guys can actually see everything there. We'll write it in boring old white.

So we have no transient dependence. We have left sigma fission, times flux, as a function of r. No source, and we have our neutron nin reactions of-- oh, we forgot our new sigma fission. Then we have our i sigma fission from nin, times flux.

Next term, we have photofission. So we have a new, from gamma rays, times sigma fission from gamma rays, times flux. Next up, we have-- well, last simplification to make. We have scattering. And we have total cross-section.

When we said, forget about energy, and our scattering kernel becomes one-- and that's light blue-- got to make one more modification to this board. Do we care about scattering at all anymore whatsoever? Because scattering doesn't change the number of neutrons left.

So we can then take these two terms and just call it sigma absorption, times flux, because if we take scattering, minus the total cross-section, it's like saying, all that's left if you don't scatter is you absorb. And if you remember, I'll add to the energy pile, we said that our total cross section is scattering, plus absorption. And absorption could be fission and capture.

And capture could be-- let's say, capture with nothing happens, plus these nin reactions, plus any other capture reaction that does something. So we're going to use this cross-section identity right here with a couple of minus signs on it. And say, well, scattering minus total, leaves you with negative absorption, to simplify terms.

I'll leave that up there for everyone to see. So then we have scattering and total just becomes minus sigma absorption, times phi of r. And we're left with-- what was that? Current, that becomes plus. There is a d missing in there, isn't there? A yellow d, minus d.

OK, and that's it. Yes?

AUDIENCE: What is d?

MICHAEL SHORT: d is the diffusion coefficient right here. So we're assuming that neutrons diffuse like a gas or a chemical with some diffusion coefficient. And so we'll define what that is, oh, probably next class, because we have seven minutes. Yeah, Luke?

AUDIENCE: [INAUDIBLE]

MICHAEL SHORT: Uh-huh.

AUDIENCE: [INAUDIBLE]

MICHAEL SHORT: The c right here, that's whatever variable we're tracking. So let's call that flux. Or let's call it n, the number of neutrons, because flux is just number of neutrons times velocity.

So let's say that the concentration was the concentration of neutrons. And we just multiply by their velocity to get flux. So it's almost like we can say that the concentration of neutrons is directly related to the flux. And that way, we have everything in flux. And that's the entire neutron diffusion equation.

Yeah, this is for one group with all the assumptions we made right here, homogeneous. What other assumptions did we make? Steady state, and we already neglected that. And I think that's enough qualifiers for this.

But it's directly from this equation right here that we can develop what's called our criticality condition. Under what conditions is the reactor critical? So in this case, by critical, we're going to have some variable called k effective, which defines the number of neutrons produced over the number of neutrons consumed. And if k effective equals 1, then we say that the reactor is critical.

That means that exactly the number of neutrons produced by regular fission, nin reactions, and photofission equals exactly the number of neutrons absorbed in the anything, and that leak out. So let's relabel our terms in the same font that we did here. So this would be the fission term.

This would be nin reactions. This would be photofission. This would be absorption. This would be leakage.

How many neutrons get out of our finite boundary? And if you remember when we started out, we said we were going to make the neutron balance equation equal to gains minus losses. And through our rainbow explosion simplification, we've done exactly that.

These are your gains. These are your losses. When gains minus losses equals zero, the reactor's in perfect balance. Yep?

AUDIENCE: How does leakage come out to be negative?

MICHAEL SHORT: Leakage comes out to be negative, despite the plus sign here. And that's actually intentional. That's because neutrons traveled down the concentration gradient.

So let's say we're going to draw an imaginary flux spectrum that's going to be quite correct. And I'm doing all of those features for a reason. But let's look at the concentration gradient right here. Leakage is positive when your flux gradient is negative.

That's why the sign is flipped right there. So a positive diffusion term means you have neutrons leaking out down a negative concentration gradient, because if you look at the slope here, the change in x is positive. And the change in flux is negative.

So the slope is negative. Concentration gradient is negative. That's why the sign is the opposite of what you may expect. And the same thing goes for chemical, or gaseous, or any other kind of diffusion.

I'm glad you asked, because that's always a point of confusion, is, why is there that plus sign? That's intentional And that's correct. Cool. Yeah, Shawn?

AUDIENCE: So in that case, if you were to explicitly right out losses, would it be minus absorption, plus leakage?

MICHAEL SHORT: Let's put some parentheses on here, equals zero, and a minus. And when we say plus leakage, we have that plus sign in there. So I'm not going to put any parentheses up here, because that wouldn't be correct.

But what I can say is that gains minus losses have to be in perfect balance to have a k effective equal to 1. Does anyone else have any questions, before I continue the explanation? Cool.

Let's say you're producing more neutrons than you're destroying. That's what we call supercritical. So I just did an interview for this K through 12 outreach program. And they said, should people be afraid when something, quote unquote, goes critical?

Sounds scary emotionally, right? And the answer is absolutely not. If you're reactor goes critical, it's turned on. And it's in perfect balance. That's exactly what you want.

So going critical is not a scary thing. It means we have control. If something goes supercritical, it doesn't necessarily mean it's out of control. Reactors can be very slightly supercritical and still in control, because of what's called delayed neutrons, which I will not introduce today, because we have two minutes.

If a reactor has a k effective of less than one, we call that subcritical. So it's important to note that the nuclear terminology that's kind of leaked out into our vernacular is not physically correct, in the way that it's used. Words like critical are used to incite emotions, and bring about fear. When to a nuclear engineer critical means, in perfect control, in balance, like you would expect, or in equilibrium.

That all sounds kind of nice, makes you calm down a little bit. Yeah, so we can put one last term in front of our criticality condition. We can take either the gains or the losses, move the equal sign and zero over a little bit, and put a 1 over k effective here.

This, then, perfectly describes the difference between the gains and the losses in a reactor. So if the gains equal the losses, then k effective must equal 1. And the reactor has got to be in balance. If there are more gains than losses, which means if you are producing more neutrons then you're consuming, than k effective must be greater than 1 for this equation to still equal zero, because this equation must be satisfied.

So if you're making more neutrons, your k effective has got to be greater than 1. So you have a less than 1 multiplier in front. And on the opposite side, if you're losing more neutrons then you're gaining, your k effective has to be less than 1 to make this equation balanced. Going along with all these definitions right here.

So it's exactly 5 of 5 of. I've given you delivered promised blackboard of Lucky Charms. And we've hit a perfect spot, which is the one group homogeneous steady state neutron diffusion equation, from which we can develop our criticality conditions and solve this much simpler equation to get the flux profiles that I've started to draw here. So I want to stop here, and take any questions on any of the terms you see here. Yeah?

AUDIENCE: Didn't we talk a lot about the different energies, like the one-group, two-group, or the discrete distributed discretized energy groups? So when we're doing the one group, you're actually just treating them fast together?

MICHAEL SHORT: We are. That's right.

AUDIENCE: To know that, like you said, the reactors do two group in the actual analysis.

MICHAEL SHORT: So a lot a lot of reactors, at least thermal reactors where you only care if neutrons are thermal or not, two group is enough. When you have a one group or a two group equation, these are fairly analytically solvable things. You get to any more groups than that, and, yes, they're analytically solvable. But it gets horrible.

And that's why we have computers to do the sorts of repetitive calculations over and over again. Once we've solved the one group equation, I'll then show you intuitive ways to write, but not solve, the equations for multi group equations. Sure. Any other questions?

So like I promised, we didn't stay complex for, long because there's basically nothing left. Yeah?

AUDIENCE: What is the 2n over 2t? Are we saying that?

MICHAEL SHORT: Oh, that's a partial derivative. Yeah, there we go. So this is saying a change in neutron population, or the partial derivative of n with respect to t, because n varies with space, energy, angle, time, and anything else you could possibly think about, equals the gains minus the losses.

I think this is worthy of a t-shirt. If any of you guys would like to update the department shirts to properly take into account photofission external sources and nin reactions, I think it would make for a much more impressive thing, because we kind of printed an oversimplification before. It's too bad. We definitely had room on the shirt. There was room on the sides, and on the sleeves.

Yeah, keep going. It might have to be long sleeve. I think that would be pretty sweet. Yeah, OK, if no one else has any immediate questions, you'll have plenty of time tomorrow, because the whole goal tomorrow is going to be to solve this equation. That's only going to take like 20 minutes.

So we can do a quick review of the simplification of the neutron transport equation, solve the neutron diffusion equation. If you have questions, we'll spend time to answer them there. And if you don't, we'll move on to writing multi group equations.

And also Friday for recitation, it's electron microscope time. So now that you guys have learned different electron interactions with matter, you're going to see them. So we're going to be analyzing a couple of different pieces of materials that a couple of you are going to get to select.

And we're going to image them with electrons to show how you can beat the wavelength of light imaging limit, like I told you before. We're going to produce our own X-ray spectra to analyze them elementally, where you'll see the bremsstrahlung. You'll see the characteristic peaks. And you'll see a couple of other features that I'll explain too. So get ready for some SEM tomorrow.