[SQUEAKING] [RUSTLING] [CLICKING]

JACK HARE: All right, so welcome back, been a bit of a break. You may be-- yes, exactly, why are we here? What are we doing with our lives? We're learning about plasma diagnostics. And you may remember, before we went away, we spent a very enjoyable lecture drawing all sorts of diagrams, like this, where we had some energy levels here.

Maybe this is a ionization energy level. This is maybe some lower energy level. And something happened. Photons were flying around. Electrons were flying around. And we said that if we want to understand what the ionization state is in our plasma and also the excitation state, which would be if you have some intermediate level here, which I guess I've already used I. So I'll have to come up with some other notation for it, N, like that.

Maybe we want to know how many electrons are in this intermediate level, in the lower level, in the upper level. Because at the end of the day, in spectroscopy, what we're detecting is photons, which are being emitted. And the photons are being emitted when the electron drops down either from being a free electron into a bound state or between two bound states.

And we went through all of these processes. And we said, if we really want to understand the plasma, we need to model all of these processes. And obviously, that's an awful lot. So today, I'm going to go through. And I'm going to present some simplified equilibria in which we do not need to include all of these processes. And these are the equilibria which we tend to have a better chance of modeling. Anyone remember what this process was called here? This is my personal favorite, when you have two electrons in an excited state, and one of them drops down, and that provides enough energy for the other electrons to ionize itself.

STUDENT: Autoionization.

JACK HARE: Autoionization, yes, absolute favorite. That's crazy. OK, good stuff. So now, let's go through some equilibria most of which do not involve autoionization.

And the reference for all of this-- there are many textbooks on plasma spectroscopy. A very famous textbook is Griem. That's the name of the author. And the title of the book is *Plasma Spectroscopy.* Griem is not actually a particularly good book for learning spectroscopy. It's a pretty good reference. But you will find all of this stuff in here, if you want to know more. Hutchinson has some of it in a slightly more fragmented fashion. So I've been taking most of this material from Griem.

So the most simple equilibrium is an equilibrium where we're in complete thermodynamic equilibrium. If we're in a thermodynamic equilibrium, this implies straight away that we have a Boltzmann occupation of states.

So our Boltzmann occupation of states is something like the occupation of the i-th state, whether that's an ionization state or it's an excited state within an ion, is proportional to exponential of minus the energy of that state over the temperature. And the electrons, and the ions, and the radiation all share the same temperature. So there is only one temperature for electrons, the ions, and the photons.

The other thing about a thermodynamic equilibrium is that all the processes balance. So all of the up processes that excite the system are exactly balanced by their counterparts that de-excite the system. This is how we get into a steady state here. That's actually not true in other equilibria. There can be processes which excite and different processes that de-excite. As long as the rate is the same, you're in an equilibrium. And we'll talk about some of those later on. But for thermodynamic equilibrium, all processes balance with their opposite process.

OK. Now, how often do we think we get a thermodynamic equilibrium in a plasma?

STUDENT: [INAUDIBLE]?

JACK HARE: Maybe in space, probably getting closer. I mean, a thermodynamic equilibrium is an infinite in the system, which is not changing in time. So of course, in reality, we can never realize it. The space plasma, you might get close to thermodynamic equilibrium.

But this also implies that you have to have your photons in thermal equilibrium with your gas, which means that your system is optically thick. And space plasma is-- at least plasma in space, in outer space, is rarely optically thick. Maybe inside some astrophysical body, it could be optically thick. So this we don't really get very often. So this is rare.

OK. Next one I want to talk about is local thermodynamic equilibrium, which I'll write as local TE. So people will often call this LTE, like this. OK. Now, in local thermodynamic equilibrium, we make the assumption that our plasma is optically thin, which means that we can ignore the processes that involve photons being absorbed and emitted in terms of changing the levels. So we have no-- I guess I'll call it radiative processes.

The range of processes still occur. We can still observe the lines coming from them and do spectroscopy on them. But they don't dominate the balance of what determines what states are excited and what are not. In fact, that balance is then obviously dominated by the non-radiative processes. So these are collisional processes. The collisional proceses dominate is.

So this is collisional processes which excite from a lower state to an upper state, but also collisional processes which de-excite from the upper state back down to the lower state. So our occupation of the upper and lower states is set purely by these collisional processes.

However, we don't need to know the cross-section for these processes, which we would do in our more detailed full model. Remember, this sigma ij is the cross-section for a process that takes us from state i to state j. But they are unimportant in LTE, due to the balance.

So we keep the feature of a thermodynamic equilibrium. But the processes balance the opposite process. So the actual cross-section of the process is not important. That's what makes it a local thermodynamic equilibrium.

And if you look at what this means and you work through the mathematics of your population balance, you come up with the famous Saha equilibrium. So a Saha equilibrium, named after its discoverer, has a form that says that the electron density times the ionization density for state z in its ground state over z minus 1 in its ground state-- so this is the density of atoms in ionization state z and their ground state atoms in ionization state z minus 1 in their ground state times by the electron density-- is equal to, most importantly, a factor to do with the ionization potential. So this chi z here is the ionization to go from z minus 1 to z.

So the change in energy there over the temperature. And then there's a few other factors out front that come from the fact that we're going to be assuming the populations are Maxwellian. There's an ion mass temperature over 2 pi h-bar squared, the 3/2. And then there's also a factor that's to do with the degeneracy of these levels. And we talked about the degeneracy as being different ways to arrange electrons and still get out the same energy. So this is gz, for the z ionization state, 1 over the degeneracy of the lower ionization state.

So the Saha equilibrium doesn't tell you about how the excitation states are distributed within each ion, but it does tell you what the ionization is. And so this tells you, for a given temperature, how many particles you expect to be in each ionization state. And because of this exponential factor here, which comes from a sort of Boltzmann-like argument, we're going to end up with most of our particles still being in lower ionization states.

OK. Any questions so far? Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: The radiation doesn't have to be the same because it's optically thin. And we don't care about it. But the ions and the electrons are the same temperature, yes.

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, so this notation here, this lower number here, is like if you have a series of levels within an atom, that atom has a charge of plus z. And then I've got these levels, which I could label 1, 2, 3, 4, and so on like that.

This 1 refers to the energy level with an atom with charge plus z. And the z here refers to that charge. So these are atoms with different ionization states, which means, in general, their energy levels will have different energies. But here, we're taking the ratio of the atoms of ionization state z in the ground state, level 1, and the atoms of ionization state z minus 1. So they have one more electron.

And they're in their ground state, 1. And the energy of that ground state will not, in general, be the same as the energy of the ion with one fewer electron because all the energy levels will shift when we reduce the number of electrons. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah.

STUDENT: [INAUDIBLE]?

JACK HARE: Ah, yes. Yeah, but you could skip that by just taking ratios of this equation with an ionized argon. And then you get out the ionization to get the fifth ionization level. So you don't actually have to do it recursively. You can do it all in one go. You could also do it that way, if you wanted to.

But yeah, in general, these ionization potentials are tabulated. They're things that people know from experiments. It's 13.6 electron volts for hydrogen. It's less for every subsequent element. And so you need to put in the temperature, of course, of your plasma. But then you can calculate this pretty quickly.

So this is a nice and simple-- relatively simple equilibria. And so for example, if you're doing spectroscopy and you see a line which has an energy which corresponds to the z ionization state-- you see a line that corresponds to argon 5+, then you know you have to be at least the temperature that produces argon 5+. You can't be at a temperature where argon is completely neutral or only singly ionized.

So straight away, this sort of equilibrium is like if I see that line, it must be at least this hot. And that's often good enough for spectroscopy. We'll talk about more advanced tactics for using spectroscopy later on. OK. Other questions? Any online?

OK. Next equilibrium is something called corona equilibrium. Coronal equilibrium is called coronal because it was first derived in the case of the solar corona. Sometimes, this is also called a collisional radiative equilibrium. And that name, radiative, is much more descriptive. And I wish we used it more. But you'll still find most people calling it coronal equilibrium instead.

So it's called a collisional radiative equilibrium because we consider the excitation going from a lower state to an upper state only collisional processes. But for any de-excitation going from an upper state back down to a lower state, we consider only radiative processes.

So this is a slightly seemingly contrived setup that actually occurs a great deal of the time. It's a setup in which our system is optically thin. And so we are very unlikely to have a photon wander by and suddenly excite our system because the photons mostly just stream through the plasma without being absorbed. So this is optically thin. The, photons escape.

And yet, it's also so collisionless-- the collisions happen so rarely-- that, although they are responsible for excitation, before you have chance to have a collision de-excite your system, it will spontaneously emit a photon. So it instead will have spontaneous emission taking us back down. So this works very well for low density plasmas. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Yes, exactly. But we're not dealing with any radiative processes that rely on there being a large number of photons like stimulated emission. Those will never happen. It's just imagine you've got your ion floating around. Lots of photons stream by. They will never interact with it. Occasionally, a collision will happen. It excites.

And so now, it's an excited state. Then the clock is ticking. Two things could happen now. Another collision could de-excite it and nonradiatively take that energy away. Or spontaneous emission could occur. And we've chosen-- or this system works in a regime where the density is low enough.

Remember, these processes, the collisional processes, all depend on density, like that. We've chosen a density which is low enough that there will never be a second collision more quickly than we have spontaneous emission here. This occurs a lot in tokamaks and other plasmas. And it's also relatively easy to solve for. So this is a very favorite model in plasma spectroscopy.

OK. So if you have that, we can also write down the occupation of the excited states. So this is some upper excited state, number of atoms in that state, compared to the ground state. Everything is being driven back to the ground state by spontaneous emission. Even if it takes multiple steps, you'll eventually get the ground state. And you'll get there long before any collision occurs.

So we are going to end up mostly in this ground state here. And this ratio here is going to be equal to the number of electrons doing the collisions to pump it up to an excited state times by the cross section to go from the ground state to this excited state averaged over the velocity. So that's our standard reaction rate here.

And it's going to be divided by the rate at which any process takes it from the upper state down to the lower state here. And this could be by a chain of multiple processes. So for example, if this is the upper state, and there's some intermediate state, and this is the lower state, we'd have to add up the processes that go straight down and the processes that go via some intermediate.

And any chain of these processes will happen more quickly than other collisions. So this is the balance between these two here. And we can write it very easily for this collisional radiative equilibrium. But the up processes are no longer balanced by the down processes. They're completely separate. That's a big break from our local thermodynamic and our thermodynamic equilibria. OK. Questions on that? Mm-hmm?

STUDENT: [INAUDIBLE]?

JACK HARE: That won't happen in this case because, before the second collision happens, it will have decayed, right? Because we're saying that these collisions are so infrequent. Yeah, really, we just have collisions going up. If the collision only took us up to here, then, before a second collision could occur, we'd have spontaneous emission and decay. If it goes up to here, then there could be multiple pathways back down, yeah. OK.

STUDENT: [INAUDIBLE]?

JACK HARE: I'm sorry?

STUDENT: [INAUDIBLE]?

JACK HARE: Right. So this is one of these things, like when is this model valid, right? And so then there will be a dimensionless parameter, which will be the ratio of those time scales. And you would like that parameter to be very large or very small, depending on whether you write it one way or the other way up.

What is very large and very small? 10 is a pretty large number. 3 can sometimes be large enough, right? So if you're using this model in a situation where the ratio is only 3, it will work roughly. It won't be completely wrong. There'll be a nice, gentle transition from being right to being wrong. And you will be somewhere in the middle there.

And so a lot of the time, we will use this model. And if you go and look in Hutchinson's book-- I'm not covering it in the lecture today-- actually, there's a limit to the energy levels for which CRE works. And so we find that transitions between the lower energy levels tend to be very rapid. They're very favorable because of the overlap between the upper electron wave function and the lower electron wave functions.

And so you can actually use this coronal equilibrium to describe the first n energy levels. And then you have to use a different equilibrium, like maybe LTE, to describe the upper energy level. So you actually have a transition between which models are valid even within the same atom with the same ionization state. So this stuff can get complicated. Yeah, that was probably too much information, but still. Any questions online?

OK. And then the final thing, after we've done coronal equilibrium, is really just having to actually solve those equations. So this is what we call non-local thermodynamic equilibrium. For many years, I thought there was a hyphen here and this was non-local thermodynamic equilibrium. So we had a thermodynamic equilibrium, which mediated non-locally. It turns out I was wrong. This is just not a local thermodynamic equilibrium, which is a much less descriptive name, but is much more valid. And you'll often see this written as lowercase nLTE.

So in general, you have to solve all of the different processes, exciting and de-exciting. But for a specific nLTE model for a specific plasma you're interested in, you may still be able to restrict what those processes are. So what you need to do is actually consider all the relevant processes because, even in these models, other processes are happening. It's just we think the rate is very small. And so those processes are not relevant.

So if you're doing an nLTE model, you simply have to calculate what processes are relevant and what aren't. So I'll give you an example here, if we're dealing with a photo-ionized plasma. For example, I've got a gas cell, a chamber filled with gas. And I've got windows on it, which are transparent to X-rays. And I have an X-ray source outside.

Those X-rays will stream in. And they will photoionize the gas. The photons will be absorbed. And we'll have photoionization. And then the occupation levels will be set by those photoionization processes. We won't have to consider collisional processes because the photon field, the strength of the X-ray drive, is so strong that the photons are much more likely to be involved than any collisions.

And so for a photonionized system, for example, we can neglect collisional processes. So that reduces the number of equations that you need to solve. And so you may be able to get away with a smaller set of equations. But if you calculate it and all your processes are roughly equal, then you really do have to solve that big equation I put up on the board last time that was effectively 0, because we're in steady state, equals the sum over all processes.

And as I mentioned, people spend their lives writing code to calculate these coefficients and to solve this balance. And it can get very, very messy indeed. But that's all I'm really going to say on nLTEs. So in general, if we're trying to do spectroscopy, there was a workflow that we might need to do. So this is modeling spectroscopy.

So first of all, we're going to model the spectra from a blob of plasma, a little cube like this. It's got some mass density, which is what spectroscopists often work in, rho. And it has some temperature. Occasionally, your spectroscopists might work in terms of ion density instead. Obviously, these are same up to the ion mass here.

Then you need to determine what equilibrium is appropriate. This could be a Saha-type equilibrium or a collisional radiative equilibrium. So maybe you know this by back of the envelope calculations. Maybe you go through and calculate a load of rates and determine what's most important. But you have to determine what your equilibrium is.

So then you need to calculate the ionization states. What organizations are present in your system here? So these are the density of ions with charge z, all the different z's. In some cases, you may be very weakly ionized. You may have a lot of neutrals. And you may have just a few atoms which have an ionization of 1.

In other cases, you may go all the way from neutrals to fully stripped. For something like tungsten, which has 74 electrons, 76 electrons, that's an awful lot of energy. But you may well be able to get a plasma [INAUDIBLE]. So then you'd have to take into account all the other ionization states.

Then you have to calculate the excited states. And you have to calculate the excited states for each nz. So for all nz, you need nz, and then whatever the excitation is. Then you need to calculate the line strengths. So these are these coefficients a, i, j. And you'll remember that these had some hand-wavy relevance to the overlap integral with the dipole operator between the upper and the lower state.

So this is the dipole operator, which is radiating away a photon of energy, h-bar omega ij. And that process has some probability to occur. And so this tells you how many photons with that energy that you're going to see. And you need to do that for every transition between every set of excited states for every ion in your system. So this is hard.

Once you've done that, you can now predict the spectra. So this is predicting the intensity of radiation with frequency omega for the very specific density and temperature that you started out with. Now, if you're very, very lucky, your system is optically thin, so tau less than 1. Life is good.

If tau is greater than 1, now you need to consider radiation transport or the fact that some of the radiation will be absorbed as it moves through your plasma. And actually, for both of these as well, you'll have to consider things like reflections. So the fact you've got metal walls on your tokamak, or whatever it is, means that some light going the other way will bounce back.

I've got my plasma here. Some of the light from this bit of plasma is going to be emitted this way. And then it will bounce back into your detector. And although it's being emitted from here, it will look like it's being emitted along the line of sight here. So if you don't take into account the reflections, you will be thinking you've got lines from one bit of plasma when they're, in fact, from another bit of plasma.

You'll also need to model the response of your detector. But there are no perfect spectrometers out there. And they, in general, have a response curve to them, which is like response as a frequency of omega. They're not flat. They treat different frequencies differently. They absorb infrared photons more than visible photons. And then finally, you need to compare this to your experimental data.

So for example, if you had a spectrum that has two lines in it like that, and you've done all of this process, and you get out a spectrum that looks like this, you'll have a choice. You can either say, it's close enough for today. Or you iterate and you do all of this again.

And you pick a new density and temperature. And you do some sort of gradient descent type thing until eventually your model spectra matches up with your experimental spectra. And then you say, aha, now I finally know the density and temperature of my plasma. It's a fair bit of work to do spectroscopy. Questions? Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, so you can. I will show you some tricks to narrow this parameter space in a moment. This is the most general case. I had whole page of notes on this, and then my iPad hasn't updated them. But I wrote the word line ratios. And I'm going to try and reconstruct the memory what exactly I was going to say. So we'll see how it goes. Yes?

STUDENT: [INAUDIBLE]?

JACK HARE: Well, other diagnostics can be expensive they can be limited to certain parameter ranges. If you're doing astrophysics, you can't do Thomson scattering on your black hole corona. So you're stuck with this sort of thing. Spectroscopy does provide extremely detailed information on all of the atomic kinetics. This is what people call this sort of distribution of different states, the atomic kinetics.

So if you're far from local thermodynamic equilibrium, say you're looking at fast electrons in a tokamak, the spectrum may provide information on those fast electrons that something like Thomson scattering doesn't because Thomson scattering may be as sensitive just to the bulk of the plasma.

So this is a time-honored time honored diagnostic. We have a lot of knowledge about how to do it. It's just that it turns out to be very hard. I've presented this in a slightly flippant way. But this is what you need to do if you actually want to solve all these things.

Some of these steps have already been done for you, right? I mean, a lot of this is now included in computer packages. And so if you know what your equilibrium is-- for example here, this rho t getting out some spectra here, there are large codes PrismSPECT is a commercially available code. You give it a density at a temperature and the size of your plasma and it will do radiation transport through it and tell you what the spectra is.

You run that for a few different rhos and t's and you create a table of spectra for different rho and t, you can very quickly fit that to your experimental spectra. So none of this stuff is impossible. It's just, if you're doing it from scratch, it can be very hard. Yeah, I'll just take Nicola's question.

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, that's a great question. So what we talked about here is that there is some photon with some energy, right? So if we think that that is true, our spectrum should look like a set of delta functions, right? And those delta functions should sit at omega ij.

And I know that delta functions are infinite. But they should have a height that is related to the strength of the transition here. What we'll talk about next lecture is line broadening. So that is when these lines are broadened by other processes. And we may be able to, in some circumstances, learn more about the plasma from the broadening.

For example, if there's Doppler broadening, we'll learn about the temperature. If there's Stark broadening, we'll learn about the density. And then there are other Zeeman splitting and things like that may even be related to the magnetic field. But we'll get on to that later. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: I would say yes, in general. And the worst the case of this is the fact that this is line integrated. So we're going through the plasma here when we're collecting all of that light. And so if the plasma-- for example, rho or t-- changes along that curve, then you can't say exactly how that density changes because there are lots of different solutions that give you the same result. So definitely not unique.

People do inverted spectroscopy and very cool techniques like that, where they measure lots of lines of sight and have symmetry assumptions. But that's quite tricky. Any questions online? OK, I'm going to try and do the line ratio thing. Let's see how this goes.

So you may have a spectrum with some very well-isolated lines. So these lines could be ones that are some distance apart from any other lines. They stand by themselves. It's very clear what their wavelength is or the frequency is. There's no confusing it as anything else.

An example of a poorly separated line would be if you have three lines all on top of each other, like this. Because these lines are probably adding together in some way, you can't exactly tell what the intensity of each of these is. But if you have nice, well-separated lines, then you have some intensity coefficients for this one, some intensity coefficient for this one here.

So this line ratio technique, you need to have strong, isolated, and identified. You need to know what these lines correspond to. So identified, in this case, means we need to know what the ionization state is for it. And we also need to know what upper level and lower level it corresponds to within that ionization state. So we need to know the-- I guess we technically need to know the density of the upper state here.

And so an example of this would be if I have a line that comes from aluminum that is twice ionized and aluminum which is three times ionized or something like that, these two ions will have very strong lines for certain transitions. You'll often have a very strong transition from some upper state, the lowest lying upper state, back down to the ground state because, in general, that's where you're going to end up. And also, in general, you have a large energy gap here. So this is pretty favorable.

So maybe you've got some well-isolated lines that you know correspond from emission from an upper level down to the ground level here. And you can measure their intensity very well. So we now know the intensity of omega 1. And we know the intensity of omega 2. And this is what I mean by strong. You don't want to be doing this on little lines that are down here in the noise because we can't measure their intensity very well.

So what you do with these lines is you say, OK, I know that the intensity of this line-- so the intensity at some frequency corresponding to the transition from the upper state to the lower state here is equal to the number of ions in the upper state times by the probability of transition from the upper state to the lower state.

And then because this intensity often has units of energy here-- this is like a power-- we obviously need to have an energy unit here. So this is h-bar omega upper lower. So if we have two lines here, we can take the ratio of these. So we can have intensity of omega z upper lower over intensity of omega z minus 1 upper lower like that. And

We can see then that we're going to get out a z upper over n to the z minus 1 upper a upper lower a upper lower-- I'll put little superscripts on these, so we remember that these are different coefficients for our different ionization states. And then the h-bars will cancel. So we'll have an omega z upper lower omega z minus 1 upper lower.

But these are things that we can calculate, look up in our atomic database. These are things that we have literally measured because our spectrometer is well-calibrated. This line ratio we've also just measured. So that means we can infer the ratio between the number of atoms in one ionization state and another ionization state.

Here, I've done them as neighboring ionization states. This could be like z minus 2, z minus 3. It doesn't really matter. You don't generally get that because, for a certain temperature, you only have a few ionization states active at a time. And because you've now inferred this quantity, you can then use an equilibrium, such as Saha. And you can get out the temperature.

So you can say your Saha equilibrium says, at a certain temperature, I should have twice as many ions in the lower ionization state as the upper ionization state. That's what I see. Therefore, that is my temperature. So just simply from looking at line ratios without doing all of this complicated modeling, you may be able to get a handle on the temperature.

And that could be a first step that you then feed back into your algorithm where you say, I know what the starting temperature roughly is. So let's only look in the solution space surrounding that. So this is quite a powerful technique. And often, this is as far as people go. They'll do a line ratio. They'll get a temperature out.

And I mentioned earlier the even simpler thing to do is just be like, hey, I've got a line that corresponds to aluminum 3+ ions. That only happens when the temperature in the plasma is like 10 electron volts, maybe 12, maybe 8. But that means my plasma is probably 10 electron volts. Perfect. And that's good enough for a lot of purposes. OK. Questions on this? Yes?

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, so no, so absolutely. I say you're measuring it. But you know it, right? It's close enough. Yes, there may be some Doppler shift involved to it or some of these other line broadening mechanisms, which means you don't know exactly where this is. You may also have uncertainties in your atomic code, especially when you're working with X-rays.

The X-ray spectroscopy models can be quite inaccurate to within 20%. And that's still pretty good. And so you'll say, OK, I thought it was going to be here. There's nothing anywhere near here, apart from this line. So that is probably this line. So making a line identification is very tough.

And especially in a plasma in the laboratory, you think my plasma is made out of aluminum because that's what you've made it out of. But then you forget that the lenses are coated with magnesium fluoride. And so there may be magnesium fluoride absorption lines when the X-rays from your experiment photoionizing.

Or your electrode is made out of copper. And now, you've got copper lines as well. Or someone left their thumbprint on one of the electrodes. And now, you have hydrocarbons and salt. So you have sodium lines on there as well. All of this has happened to me very recently. So actually, it's a nightmare. You think, my plasma is made out of this. And the answer is no. Your plasma is made out of whatever crap there is in the machine, so yeah. OK. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Yes, and preferably negligible, right? So yeah, this is definitely a technique that works best if it's optically thin. You could use it if you have very strong priors about what the density and temperature are so that you can calculate the opacity. But if you have those strong priors, then you don't need to do this technique. You can use it to confirm post-hoc afterwards that your technique was right, yeah.

STUDENT: [INAUDIBLE]?

JACK HARE: Well, I mean, that's really tricky. If you're able to do an experiment where you can increase the depth of your plasma, that would be one easy way to do it. But of course, if you're in a tokamak, you can't just make the plasma twice as big just to see what would happen to the lines. And so in some experiments, you could use that technique.

Otherwise, yeah, you'll have to use maybe other diagnostics to give you ideas of density and temperature that you can then check how optically thick the various lines are, and then iterate from there. Or simulation as well gives us some hints as to what the answer is. Any other questions?

What are we doing? Ah, X-ray spectroscopy, very good. So the reason I'm talking about X-ray spectroscopy in particular is that quite a few people doing plasma physics end up using X-ray spectroscopy. And also, it's got its sort of niche nomenclature that I wanted to introduce you to so that when people start talking about it, you have some vague idea of what's going on.

Of course, all forms of spectroscopy are valid. A lot of people do visible spectroscopy, infrared, ultraviolet. But if you're working with plasmas at fusion conditions, which are like kiloelectron volt plasmas, then you're going to be getting kiloelectron volt photons, right? And so that means they're in the X-ray regime. And X-ray spectroscopy can be a little bit different.

So X-ray spectroscopy has very strong characteristic lines. So these lines tend to have a high transition probability, aij. And they tend to be narrow. So they show up very, very strongly. They're easy to identify. They're easy to work with.

And these lines correspond to transitions often between some of the lowest energy levels inside our atom. So here, this 1 is not referring to some abstract concept here. This is literally n equals 1. So this is the lowest energy level down there, the energy level that your ground state hydrogen atom occupies here. And these are the other excited states.

Because we're dealing with relatively high energy systems when we're doing X-ray spectroscopy, an ion may be almost completely stripped. There may be only two, or three, or one electrons here. And people start referring to these ions as helium-like or lithium-like. So a helium-like, unsurprisingly, has two electrons. And a lithium-like has three electrons.

What they mean is this could be something like tungsten, or aluminum, or anything else like that. But it's so ionized that there's only two or three electrons left. And that means that you can use all of the spectroscopy, all of the quantum physics that you do on helium, which is a relatively simple system to solve compared to something more complicated.

And all you have to do is replace the atomic charge, the nucleus charge, from 2 to whatever your actual nuclear charge is. So you just take z from 2 to-- I don't know-- 13 for aluminum. But it means that you have a pretty good idea of what these energy levels are because the quantum physics of low atomic number elements is pretty well understood.

So these characteristic lines will occur when there is some vacancy in a lower energy level. So for example, if we have an electron up here in the n equals 2 level and there's space in the n equals 1 level to drop down, it will emit a photon here.

Now, why would there be space in this lower energy level? I'd updated my notes and that was also not updated as well. This may be because of collisions, or photoionization, or something else like that. If we're in a very hot plasma, it's very reasonable to have hot electrons flying around that collisionally ionize us or, indeed, just collisionally exciters up to some very high energy level, like 7.

And that leaves a hole here. And because this is energetically unfavorable, to have a lower energy level unfilled, one of these electrons in one of the higher energy levels will very rapidly drop down into here. And these different transitions from down to the ground state from 2 and 3 and so on have different names.

So first of all, we have what are called the K-shell transitions. This is transitions from-- or down to n equals 1. We have the K alpha, which is from n equals 2 to n equals 1. That's the one I've drawn here. You have K beta. That's n equals 3 to n equals 1. And so on, K gamma, all that sort of thing.

Now, it turns out that the K alpha is always the strongest transition. So you'll typically get an electron dropping down from 2 to 1 rather than from 3 to 1 more quickly. And this line will show up more strongly. And it's strong, again, because of this overlap between the lowest energy state and the second highest n equals 2 energy state.

These wave functions look most similar. As you start getting out to high ends, these wave functions start to look really weird. And when you do the overlap integral, it's not particularly large. So quantum mechanically, it's much easier for the electron to transition between these two wave functions here. So the K alpha is always strongest.

We also then have the L-shell. This is transitions down to n equals 2. And you'll have L alpha, which is 3 to 2, L beta, which is 4 to 2. Then you also have the M shell and so on. But we tend to focus on the K shell because these lines are the highest energy atomic lines you can get from a system because you're dropping down to the lowest energy level from somewhere else. So these lines are very bright. And they're very good signatures of what the plasma conditions are.

So just to give you an example, this is energy here and this is intensity. And we have something like copper. The copper spectrum will have some Bremsstrahlung component background here that you can never get rid of. And then at sufficiently high energies, it will have two peaks, like this. This one is the K alpha. This is the K beta.

The K alpha, in this case, is 8.1 kV. And the K beta is 8.9 kV. And these lines are so narrow and so well-defined that they form a very strong spectroscopic signature. So if I'm looking inside my tokamak plasma that I want to be hydrogen, and I look at these high energy kiloelectron volt ranges, and I see peaks there, I can say straight away there is iron in my plasma, there's tungsten in my plasma, there's molybdenum in my plasma.

And of course, all of those impurities are radiating very strongly because these transitions are very strong. And so that really ruins your power balance in the reactor. So a lot of spectroscopy, early spectroscopy work, at least in magnetic confinement fusion, was looking for these signatures of impurities in the plasma here.

OK. Questions on all of this? The nomenclature thing is basically knowing that K means transitions to 1 and alpha and beta mean from the level above and the level above that and knowing that the K alpha is much stronger than the others and maybe knowing that sometimes people are like, this is a K alpha transition in helium, like aluminum. And you're like, ah, cool, yeah. But questions?

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, absolutely. Bremsstrahlung is producing those wavelengths as well, in general, but not lines. So that's why they're so unique. And maybe in some system, you could have synchrotron radiation or something like that. But in general, these are very specific. So yeah.

And they're slightly different energies. So K alpha for helium-like is at a slightly higher energy than alpha for lithium-like. And so you may also then have the-- say these ones are the helium-like. And this is the K alpha and the beta for the lithium-like

And then because you know exactly what these lines are because they're so specific, you can do your line ratio trick that we used before and work out what the temperature is because this is a lower ionization state than this one. Yeah. Obviously, if you are fully stripped, if you have got to a temperature where there are no electrons left in your system, you no longer have the line. You would only get a line from recombination. But the recombination has a spectrum, as we talked about, with an edge, but still a spectrum. yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: No, I believe that is separate. Yeah, actually, it's been a while since I reminded myself what Lyman alpha was. That's the one that's 121 nanometers, right?

STUDENT: [INAUDIBLE].

JACK HARE: Oh, OK. Any astrophysicists? No? OK. Yeah, that's like a UV line. Is there a chance that it is related to these? Yeah, I think it is an X-ray line, but I can't remember The most common one. It makes me think--

STUDENT: [INAUDIBLE]?

JACK HARE: It's called the Lyman series or the Lyman-- there's a Balmer series. What's the Lyman series?

STUDENT: [INAUDIBLE]?

JACK HARE: OK. So maybe in hydrogen, they have separate names that correspond to these. And I think, historically, people were doing spectroscopy on astrophysical bodies, which were the first bodies we could looked at that had these temperatures. So people were not doing spectroscopy on copper straight away.

But if you point to your spectrometer at the sun, then you can get these alphas-- or these Lyman series imbalances. I'll look that up. Good question. Any questions online or comments if you know anything about Lyman series and Balmer series?

OK. We can use this in reverse to design filters for our X-ray systems. So we often want to be able to measure just the light in a certain region here. So this would be like taking a picture of just the X-rays. And so we want to be able to filter out all this other light and just use the light in this region.

So this is looking at X-ray filters. So a filter is a solid object. It's usually a thin foil made out of some metal or some non-metal here. And the idea of the filters is that we still have these different energy levels here-- n equals 1, n equals 2, n equals 3 and 4. But now, this is for z equals 0. So this is for some solid, like a lump of tungsten or something like that.

And then if we have a photon coming in, if there's a possibility that photon has enough energy, h-bar omega, to excite an electron from one of the inner shells, from 1 and 2 from 3, then that process is now favorable. It's likely that photon will be absorbed. If it doesn't match with any of these bandgaps or any of these transitions, then the photon will go through. But if it does match, then it will get absorbed.

So what we see in our absorption plot for a filter-- so this is that quantity, alpha, the opacity, as a function of frequency-- is we see a series of what are called edges here. So this is the L edge. And this is the K edge.

And what you get is your lines corresponding to your frequencies corresponding to K alpha, K beta, all the way up to the highest K, which is ionizing from whatever the highest occupied state is up into the continuum. And same here, you get-- I'm trying to space these parabolically-- L alpha, L theta.

So effectively, you have a sudden increase in absorption whenever it is possible to do this ionization. If you have done the homework yet and you've looked at some of the other filters that you can generate on the Henke website, you will have noticed these really weird absorption features.

Henke normally gives you transmission. And so you will have seen 1 over this. The absorption is 1 over the transmission. But you will have noticed these really sharp edges. That's not a numerical problem. When you first see it, you think, oh, the code is broken here. No, these correspond to these very strong absorption edges here.

So if you have a spectrum that has some lines in it and you put that spectrum through this filter, this filter would cut out, for example, this highest energy line here. And you'd just be left with the spectrum in this region. So this is how you use X-rays to do filtering. Oh, OK. Questions on that?

And the position of these edges depends on the element. So if I use an aluminum filter, or a tungsten filter, or plastic filter, they'll have edges in different places. And so by using different filters, I can image different parts of the spectrum. So we'll talk about X-ray imaging in a moment. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Not at all, yeah, just photo-absorbed, yeah.

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, so that's a bit that doesn't-- yeah, yeah. That's the unique quantum and treating the photon as a wave in order to get-- well, you need a quantum treatment of it in order to get this roll-off here. Yeah, my explanation only covers why there's a sharp edge. I haven't seen a nice intuitive explanation for this roll-off afterwards. But all materials at high enough energies, no matter what, their absorption goes down to 0. So gamma rays will get through anything eventually, Yeah.

STUDENT: Professor?

JACK HARE: Oh, sorry, I'm just going to answer another question, Jacob first. And then I'll get back to you. So the question was, why the edges get taller and taller? I think it's due to the likelihood of this interaction, but I'm not sure, yeah. Sorry, Jacob. Your question?

STUDENT: How do you design a material to specifically change or take out a certain-- like the K edge or the L edge?

JACK HARE: You don't design a material. You have a look and see what materials are available.

STUDENT: Yeah, I mean, how do you choose a material?

JACK HARE: How do you choose a material? Well I choose it by going on the Henke database. And I generate these curves. And I see what they are. And I guess, if you're very experienced, you just know, for example, that this K ege is at 1.6 kV for aluminum. And if you're trying to measure photons of energies below that, then 1.6, this aluminum, is quite a good filter.

But you can't engineer a material to have a K edge. It's an atomic process. So it really depends on the elements. If, for example, you have a filter made out of a mixture of two elements, then it's just the combination of the absorption for those two different elements. So yeah.

STUDENT: And you have to get the L edge to do the K edge, right? Or can you just get--

JACK HARE: Yeah. No, no. The L edge will always be there as well, and the M edge, and so on, but down to lower and lower energies, right? So they may not be interesting for your bit of the spectrum.

STUDENT: All right, so you'll still see the bottom ones. But you just know that they are too low of an energy to--

JACK HARE: If you think your spectrum looks like this and you know that there's no interesting features down here-- it's just Bremsstrahlung with very bright lines up at high energies-- then you don't really care what the structure is down here because when you multiply the transmission via your spectrum, there'll be like little fluctuations down here that you don't care about because your spectrum is going to be dominated by these lines you think.

STUDENT: OK, makes sense.

JACK HARE: Other questions? OK, let's talk about X-ray imaging, so why we might want to use a filter. Obviously, if we're doing spectroscopy, we don't really care about the filters because we can just do spectroscopy. And we can see where the lines are. But if you don't want to do spectroscopy and you want to make an image, you're going to need some of these filters so that you know what energies you're looking at. So this is X-ray imaging.

So one of the facts about X-ray imaging is that there are no lenses. And that is a consequence of the fact that the refractive index is roughly 1 for X-rays. And lenses work on having a difference in refractive index from air. So if your refractive index of all your materials is about 1, you can no longer bend the rays. So the simplest thing that you can do for X-ray imaging is called pinhole imaging.

Pinhole imaging is what the first cameras for visible light used as well. You have your object, which, in this case, I'm drawing as like a triangle with very little symmetry for reasons that might become apparent later on. You have your pinhole, which is a little opening of diameter d in some plate, which is otherwise completely absorbing.

And then this is just geometric optics. So you trace the rays of light from the corners of your object through the pinhole. And if you get it right-- and I'm probably going to get it wrong-- you get an object which is inverted from the one that you started with. So it's upside down. And it's flipped up-down and flipped left-right as well.

That's not usually a big problem. You can usually work out what's going on. The object will also be magnified. So there's some distance, u, between your object and your pinhole and some distance, v, between your pinhole and your detector. It could be film or some camera.

And so this is also magnified by the ratio of v upon u. And you can check that out by drawing some lines for yourself. And you'll see very quickly, from geometric arguments and similarity of triangles, that you get out this magnification here. But this will also be blurred. So you'll not get a perfect image. And it will be blurred by two effects. The first one is geometric blurring.

The geometric blurring is a function of the fact that you have to have some finite size pinhole here. Ideally, this pinhole would be infinitesimally small. And then all of your rays of light would have to go through the pinhole to make it through the detector. In practice, if your pinhole is very, very small, then you don't get any light through at all. And you don't get a very bright image. So you have to increase the size of your pinhole to get enough light through to make an image.

The geometric blurring is effectively-- if you think that your object consists of two little bright blobs, like this-- so you're looking at it. And it's just two little LEDs, like that. The projection of those dots, say a dot here and a dot here, through the pinhole-- I draw one of these for the top side of the pinhole and one of them the bottom side of the pinhole.

So if this is just a missing light uniformly in 4 pi steradian, then it will project a circle through here. So we will end up not with two dots, but with two circles. So this is purely geometric. I'm just using ray optics here. I don't need to think about waves whatsoever.

And the diameter of these circles here is going to be equal to M plus 1 times the pinhole size. So we can change how much this blur occurs by changing our magnification and also by changing the pinhole. And in the limit that our pinhole size becomes very small, our blur also becomes very small. And we can resolve point objects again.

But like I said, we need-- well, if we have a very small pinhole diameter, then we have no light. So there's a trade-off here. And if you do photography and you're aware of f numbers and things like that, some of this may be very familiar to you. But for those of you who don't, I think this treatment is valid.

There's another limit, which is the diffraction limit. And this is where we need to take the wave nature of light into account. So we have our little pinhole again here. But we know, from Huygens principle, if we got some wavefronts coming in like this, every point on the wavefront is a source of spherical wavefronts.

Normally, when there's no pinhole here, that doesn't matter because our next wavefront is then made out of the sum of all of these little spherical wavefronts. But when we put our pinhole aperture in the way, we block out some of those wavefronts. And instead, we start to get diffraction.

And so we start to get spherical waves coming out from this. And so if we initially had a point source, it's now going to be imaged into a blurry source here. And so the-- which one's-- OK, I can't remember. I think it's the Rayleigh criteria. The Rayleigh criteria says that if you have two dots like this, for example, that are some distance apart then we can only distinguish them if their blur is separated and this separation, which sets the resolution, is 1.22 lambda over d.

That's an angle space. So multiply it by v to bring that angle into a real position on our detector here. So this 1.22 comes from zeros of the Airy function. These are actually like little sinks. And we're trying to make sure that the first zeros are separated between these two here. So you may have seen this if you've done diffraction theory before.

So the geometric blur doesn't care about the wavelength of the light. It doesn't care about the energy of the photons. The diffraction blur does care about the wavelength of the light because it's a wave effect here. And so we can make the diffractive blur go away by having a bigger pinhole or by having a shorter wavelength.

So if we go to a bigger pinhole, then our geometric blur will be worse. And so there's actually a trade-off between these two. There'll be a sweet spot that gives you the best resolution, where you minimize the sum of the diffraction and the geometric resolution. But you need to think about this when you're designing pinholes.

OK. Oh, the other thing about this is if you have a source that is emitting lots of light in lots of different wavelengths, then the short wavelengths will be distractingly blurred. And they'll be smeared out over your detector. And the long wavelength won't be. But that means that you'll have this blur over your whole detector from the short wavelengths, which reduces your signal to noise. And so this is especially where you want to have a filter.

So you want to filter out the low wavelengths so that they don't blur out your image. And you want to just have the high wavelengths, like these ones here, make an image just in some of the K alpha emission from the very hot parts of your plasma. OK. Questions? Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Yes.

STUDENT: [INAUDIBLE]?

JACK HARE: Yes, I'm not sure how that comes into this geometric optics framework.

STUDENT: [INAUDIBLE].

JACK HARE: OK. Yeah, I see where that's coming from. But this is for a pinhole. It's always going to be inverted. I guess, theoretically, if you put two pinholes-- no, that doesn't work. Nope, it's always inverted. Yeah, so maybe it's negative in someone's sign convention. But not in what I'm using here because it's always true. Yeah, good question.

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah. Well, what I've missed out here is, in fact, that the transmission of as, you know, a thin bit of metal foil is actually opaque to visible light. So if I put a very thin bit of aluminum foil in there, I will definitely block out the light. And you need like 200 nanometers. So you can just deposit it on some basically transparent plastic. And then you will also get the filtering from the aluminum K edge, for example.

So yeah, if we haven't gone down to very low energies, this is kind of like an X-ray regime treatment of it because, of course, you need X-ray photons in order to even get the L edge. In reality, this goes down and some other stuff happens down at this visible wavelength. Yeah, but any metal that is sufficiently thick-- and it doesn't need to be very thick-- will block all the visible light out. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Yes, you're right. I just constantly get confused between wavelength and frequency. Filter out long wavelengths, which is visible light, generally. OK, good question.

You also do imaging-- and I'm not going to talk about it at length. Pinhole imaging is quite limited. You can also do imaging with bent crystals. So if you've ever looked at diffraction of X-rays from crystals, you'll know that the crystals will diffract light of certain wavelengths. And if you have a bent crystal, you can also do focusing with it.

So people spend a lot of time very gently bending bits of crystal. And that makes a reflective lens. So you can't have a transmissive lens. But you can have a bent mirror, effectively. So there are two ways to focus light. You can have a lens like this that your light goes through. Or equivalently, you can have a mirror, like this, that the light reflects off. And because the mirror is bent, they will get focused to a different point. And as you adjust this mirror, it will depend where this gets focused on.

Now, even a crystal is not a mirror for X-rays. It's a mirror for a certain wavelength of light. And so if you put your detector here, you might have the K alpha focused onto your detector. But for example, the K beta, which is coming in the same direction, might get reflected or diffracted off to a different point, which is off your detector.

So this is a way of doing monochromatic imaging. So effectively, your crystal here acts as both the filter and the focusing element. So you can make beautiful images. You don't need to have a tiny pinhole. So you can get lots and lots of light. And you can get images with just a single wavelength, which can be very powerful for saying this is a very hot bit of my plasma that's producing K alpha. So bent crystal, you get what's called monochromatic imaging.

But bending crystals is extraordinarily expensive. So pinholes are very, very cheap. And you'll find a lot of people using pinholes because they work pretty well. And of course, the bent crystals only work for individual energies. And so if you want to try and measure more broadly what the radiation is from the plasma, you'll miss out on information. OK. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: My last set of notes for today, I promise. Sorry, what?

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, diodes are pretty good, yeah.

STUDENT: [INAUDIBLE]?

JACK HARE: No, no. It's nice silicon, 50 nanometers thick, perfectly good for X-rays up to about 10 kV, so nothing wrong with diodes.

OK. So a lot of these detectors are simply film. If you went to the dentist a decade or so ago, they would have used film in their X-ray things. And in fact, most of X-ray film was developed for medical applications. And film is really good.

OK. So film has a huge drawback. It's time-integrated. So you have no sense of what the time is. You have an integral of the intensity over time. So if you have some filtering, maybe you have still frequency information left in your image, because you're like, oh, this is coming all from the K alpha because I put a special filter in. But you don't have any time. So that could be a really big problem if you're doing something with a rapid implosion, like in inertial confinement fusion.

But it doesn't have to be a big problem. For example, if you've got a signal that is very bursty in time, like an ICF-- so this could be a few hundred picoseconds across. And you have an X-ray diode looking at it. And that diode has time resolution, but no spatial resolution.

Then you could use that diode signal to say, well, all of my really bright X-rays are missing 100 picoseconds. So although this is time integrated, I actually know that the time that this was integrating over was only P to P plus 100 picoseconds. So you can use a combination of diagnostics to localize your signal in time. That might be OK.

Film is really high resolution. Film grains are very small. So it can be much higher resolution than a detector, like something that has pixels. You obviously don't have to scan it at very high resolution. But we can do that. It's also very well characterized.

So although it may not be linear, we know like this amount of darkness corresponds to this much energy deposited. So film is actually really precious. And unfortunately, a lot of the film manufacturers, like Kodak, have stopped making X-ray film because there isn't a market anymore.

I believe Omega has one of-- the Omega facility at Rochester has one of the largest stockpiles of remaining medical film in a fridge because they can't buy it anymore. And no one else has any left. So people keep begging Omega for some of their film from time to time. So it's good stuff. But a lot of people have now moved to something called image plates.

Image plates, at first glance, looks an awful lot like film. It's a bit of plastic, plastic-looking stuff. It's time integrating. Again, there's no electronics to it. What happens with image plates, which is neat, is the X-rays excite it up to a metastable state. So you have the X-rays excited up to this metastable state up here.

And because it's metastable, it doesn't actually decay straight away. This may be-- ah, sorry, yeah it doesn't make sense unless I do it like this. The X-rays excite the atoms in the plastic up to some upper state. And then via some non-radiative decays, because of the peculiarity of the plastic, the system decays down to a metastable state. And metastable, in this case, means that there is no permitted transition down to this lower state here. This doesn't happen. So this is a forbidden transition.

Now, even though the transition is forbidden, there will be some other mechanism, mostly to do with collisions, which will eventually de-excite the atoms in your detector again. So you've got like a few hours to get this bit of image plate out and put it into a scanner.

What the scanner does, very cleverly, is that it uses a laser to excite the image plate back up to some upper level. And then it looks for the fluorescence as it decays back down again. So you can read off the number of stored photons in each little bit of the plastic using this little laser and do this scan. And then you can reuse the image plate, which you can't do with film. So image plate is pretty nice. But you need an expensive scanner. So it has a laser readout.

So both of these are time integrated. You can get time gated cameras. So these are cameras, for example, which have some gate on them, like this. And maybe that gate is 5 nanoseconds or something like that, where you're really getting a very, very fast burst of X-rays. So that avoids this problem. You know exactly when the X-rays were coming out. But of course, a time gated camera is very expensive because now it's got lots of electronics.

These are often things with names like Hybrid-CMOS. These cameras can be very, very expensive. And they also may have relatively low resolution because each pixel costs a lot of money. So if you can get away with using film or image plate, you might be tempted to do that.

Finally, there's something called a streak camera. A streak camera is a slightly odd device. And when you first see a streak camera image, it's very hard to understand what's going on. What you get out of your streak camera is actually something that has got a lot of intensity, position on one axis, and time on the other axis.

So if you have something like an ICF capsule that is roughly circular and you put your streak slit here, initially, if this little circular blob is emitting nice and uniformly, you would see nice, uniform emission in all of x. And as this ICF capsule gets smaller, your slit stays the same. And so the region that's being lit up on the slit gets smaller as well. So you'd see a streaked image that looks like this, where this has got more light on it and this region has no light on it.

So this is a very clever technique for monitoring like the evolution of a 1D profile of intensity in time. And the reason these streak cameras are so popular is that they're extremely fast. And they're actually based on the same technology that our CRT television uses with a beam of electrons being swept across here. They are very old, though. So everyone who has one, they keep breaking all the time. OK. Any questions on detectors? That's all I have for you today. Yes?

STUDENT: [INAUDIBLE]?

JACK HARE: Normally, what you actually have is a bit of film that sits behind the streak camera window. And so the film gets exposed. And it has this pattern on it. And you'll have some timing dots from a laser that will tell you how to calibrate your time axis. And you might have some fiducial that tells you how to calibrate the spatial axis here.

But because it's something like film, it does have a dependence on the intensity. So it's not binary. But the intensity may be very non-linear because to get this image, you're generating an electron beam and sweeping it. And there's, like, magnets, and electron bunching, and all sorts of things like that. So interpreting these in a very quantitative way can be very, very hard. Yeah?

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, so there's something called a-- often something called a microchannel plate, which has some voltage across it. When the photons come in, they ionize. The photoelectric effect, they release an electron from the surface. That voltage sweeps the electron and accelerates it up.

There'll be magnets, which bend the electron for a certain part of the time axis. The electron will collide with the phosphor screen. The electron excites the phosphor. The phosphor decays and emits green light. The green light is recorded by [INAUDIBLE].

Yeah, so they're very fun. But they are very complicated, which is why you shouldn't-- if you see a region that is twice as bright here than it is here, that doesn't necessarily mean that it's emitting twice as many X-rays because it's very, very non-linear. And they're very sensitive to voltage.

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, it's a pulsed magnetic field that's ramping up in time. So as the magnetic field ramps up, it sweeps the electrons to different places. That's exactly how a cathode ray works. And we were able to do cathode ray TVs to make moving images at 30 frames a second, so yeah.

STUDENT: [INAUDIBLE]?

JACK HARE: Yeah, but all the electrons from the photoelectric effect are be coming out with very similar energies. And then they're accelerating. So maybe they've got an energy range of a few EV. But they're being accelerated over a kiloelectron volt, which is also why you should not mess around with a CRT because inside it there is an electron beam that has kiloelectronvolts, yeah. Any other questions? Anyone online?

OK. Next lecture, we are going to talk about line broadening. And we're going to finish off spectroscopy. So see you on Tuesday.