[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: So we're very fortunate today to have Dr. Maria Gatu Johnson join us to give a guest lecture on neutron diagnostics. Maria is a principal research scientist at the PSFC. She did her PhD working on neutron spectroscopy on jet, on magnetic confinement. But now she works on inertial confinement fusion. And in particular, her work on the magnetic recoil spectrometer on NIF was a key to understanding the record neutron yield that we got last year in the ignition shots. So we're very fortunate to have Maria here, and looking forward to hearing her talk.

MARIA GATU JOHNSON: OK, thanks, Jack. So an hour and a half on nuclear diagnostics. For ICF, that's kind of a lot to squeeze in. So I did a bit of a sampling. I want you guys to ask questions as we go through. If there's anything you're particularly interested in, stop me and we'll talk a little bit more about it, try to touch on the key things.

Also, there's a lot of familiar faces today. So a lot of you know details of what we'll be talking about today, in some cases, better than I do. So then feel free to chime in as well. So with that, we'll get started.

So yeah, as Jack said, we're going to be talking about nuclear diagnostics for ICF plasmas. This, in particular, illustrates three facilities where we do ICF work-- the National Ignition Facility in Livermore in California, the Z facility at Sandia National Labs in Albuquerque, New Mexico, and the OMEGA laser in Rochester, upstate New York. And this is actually a picture of the instrument that Jack just talked about in his introduction, the magnetic recoil neutron spectrometer installed on the NIF target chamber, which is the blue, circular, spherical part that you can see there in the background.

But I thought, as part of this, we'll touch a little bit on-- if I can get this to move forward-- some comparisons to diagnostics for magnetic confinement fusion as well. So these pictures here, that's the inside of the jet [INAUDIBLE] time of flight neutron spectrometer at the [INAUDIBLE] which I did build for my graduate work. So we'll talk about that one a little bit as well when we get to that point.

Brief outline-- find this a little weird between switching the computer and being behind the camera. We'll start by talking about implosion parameters and nuclear signatures. What are we actually looking for in the ICF implosions? What do we get from the nuclear diagnostics? So that's kind of the broad overview background.

And then we'll go into a little bit more technical stuff-- the nuclear diagnostics that we use. And I've divided it into neutron activation, neutron spectrometry, neutron imaging, charged-particle spectrometry, touching a little bit on other charged-particle magnets as well, and then finally, reaction-rate history. And finally, if we have time, I'll just spend a few slides discussing the impact of nuclear measurements on the ICF program that is in particular.

So let's start with the implosion parameters and nuclear signatures. So the nuclear emission from an ICF experiment carries information about the state of the fusion fuel. And this is actually what really excites me about nuclear diagnostics is that they carry information directly about what's happening in their reactions. They're the products of the reaction. So they know exactly what's going on. We can count the number of nuclear products, and that gives us a measure of the number of fusion reactions that happen.

We can look at the energy spread of the fusion products, and that gives us a measure of the plasma ion temperature. We can look at the energy upshift of the fusion products to fuel velocity. Was there a comment online? Can you guys hear me OK?

AUDIENCE: We can hear you great, thanks.

MARIA GATU JOHNSON: Great. And we can look at a scatter or downshift of the nuclear products to study area of density, which I'll discuss what that is, of the compressed fuel and shell. We can also take an image of the nuclear emission to study the spatial burn profile, and we can look at the temporal evolution to determine how nuclear burn evolves in time. And actually, if you look at this, quite a few of these are also relevant for magnetic confinement fusion. We can obtain similar information by looking at the nuclear emission from a tokamak, for example.

The exception is the area of density, which is a specific quantity to ICF. So I want to spend a few slides talking about why that's important. As I think all of you know already, ICF uses the inertia of a dense shell to confine the plasma before it blows apart under its own pressure. We can express a confinement time in terms of the sounds, cs here. This is kind of illustration of how it works. And if we take a--

[INTERPOSING VOICES]

AUDIENCE: Should I turn this one round? And then you've got one fewer computer to look at.

MARIA GATU JOHNSON: Trying to switch slides here. OK, that's good.

AUDIENCE: There we go.

MARIA GATU JOHNSON: We can take a mass average of the local confinement time. And that gives us the confinement time as the radius over the sound speed with a 1/4 of r to the 4 in that. So I'm not going to go through that integral in detail, but it's a very simple one. Yeah, so this is the confinement time.

OK, so then if we look at the standard, number density times confinement, and plug this expression from the previous slide in for the confinement time, it gives us this expression. So we see that number density times confinement time is a direct function of [INAUDIBLE], which is this area of density that we keep talking about, which is essentially a confinement parameter in inertial confinement fusion. And that's why we care about it so much.

Oh yeah, even highlighting it there. High areal density-- or rho R, which you should refer to it as-- is required for a significant fraction of the fuel to burn before it disassembles under its own pressure. We can express the burn fraction, fb, as rho R divided by rho R times 6 grams per centimeter squared. That's at an ion temperature of 30 K. You can derive it at different [INAUDIBLE] temperatures. So that's kind of high compared to what we usually operate at. And this expression actually is derived from the fusion burn rate integrated over the confinement time.

If we throw in some numbers on this, we can throw in that we want a burnup fraction of 25%, which really would be required for high gain. That means we need a rho R of 2 grams per centimeter squared. OK, just to put this in perspective, the best performing NIF implosions to date have had a burnup fraction of about 5%. So this is really high performance. And then if we assume solid D-T, D-T ice has a density of 0.25 grams per cubic centimeters. Then we find that for ignition, we need a fuel mass about 1/2 a kilogram. Again, this is with this expression, solid D-T ice.

So then the question is, as I'm sure you've seen this 1,000 times before in ICF presentations, can we really work with 0.5 kilograms of D-T fuel in a laboratory? Of course, the answer is no. That gives us a little too much yield, which is not quite what we want. And that motivates-- that in order to achieve the required areal density, our confinement parameter, i temperature and confinement time without destroying the lab, we need to compress the capsule.

We actually have to compress it quite a lot. We get these rough parameters, starting from about a 2 millimeter size capsule. Get down the radius of 30 to 50 microns, density like 700 grams per cubic centimeters from the 0.25 that we started with. And an areal density of 2 grams per centimeter squared, temperatures from 5 to 40 keV, and confinement times from about 20 to 200 picoseconds.

So looking at these numbers, these inclusion parameters really set the requirements on the different measurements. Typically, we want to achieve a 5% to 10% accuracy on these kinds of numbers. So that really tells you how well we need to be able to make these measurements.

OK, so then we have a lot of different products that we can work with. And they really carry a wealth of information about ICF implosions. First of all, we have our primary products. All of you know that we primarily work with the D-T reaction, which gives us the alpha particle and the neutron [INAUDIBLE].

PROFESSOR: I don't think so unfortunately.

MARIA GATU JOHNSON: OK, keep jumping.

PROFESSOR: Yeah, sorry.

MARIA GATU JOHNSON: You could use a board eraser, but that's only going to give you an extra foot or so. So it's probably not worth.

MARIA GATU JOHNSON: OK, so that's the primary one we work with. If you look over here, that gives us the yield, the ion temperature, the areal density. We can also use it for yield versus time and use it to infer the confinement time and the radius of the capsule. There's also another branch, which gives us a gamma particle, which is actually quite useful. And we'll look a little bit at that as well later in the class. In addition to D-T, we can also look at primary products from the D-D reactions. We have two, one that gives a [INAUDIBLE] proton, one that gives a helium 3 and neutron.

And actually, a lot of-- many confinement experiments have been working primarily with D-D to date. So you can use that neutron for a lot of measurements as well. And then the helium 3, we also work with in a lot of surrogate experiments where we don't want to do D-T and it's a different thing. And then actually, it's quite similar to D-T. We also get alpha particle, get a proton instead of a neutron with pretty high energy.

OK, so those are the primaries. Counting the number of primaries to generate gives us a direct measurement of how many fusion reactions we had, obviously. So that's a quick way of getting the yield for inflation. Not necessarily always easy, but basically that's how it works. We can also have secondary products, which I actually won't be spending much time on today. But if you're interested in that, I'm happy to answer questions afterwards. So that's when one of the fast products produced in a primary reaction goes on to react with a thermal fuel and to give a broader energy spectrum of the reaction products.

And then we have knock-on reactions. We'll spend a little bit of time looking at how this works. So that's when the fast neutrons born in the D-T reaction hit one of the fuel ions to give faster fuel ion and a scattered neutron. We actually use this quite a lot. And we can also have similar reactions for the alpha particles scatter on the fuel ion and give [INAUDIBLE] fuel ions. I won't spend much time on this today, but this is the signature that they can use to look at what the alpha particles are doing in the plasma, in particular, when these fast neutrons react again, in a tertiary reaction, which I think the example's in here-- yeah, to give these fast neutrons, which we call alpha knock-on neutrons.

Think some of you have heard about that before. OK, yeah, so this is actually pretty cool, these knock-on reactions are the easiest way of measuring rho R. So we look at some examples of that as we go through.

OK, we already talked about this-- counting the number of emitted primary fusion products in a set. Solid angle and scaling it up to 4 pi gives us a measure of the total yield. So this is basically the yield reaction. I forgot to put the y in front of it. But if you do the integral over volume and time, the densities of the reactants-- typically, it could be [INAUDIBLE] and [INAUDIBLE] for D-T plasma-- and the reactivity-- and this is just a Kronecker delta. So if it's 2D, you get a factor of 1/2.

That gives you the [INAUDIBLE]. And the reactivity is the integral over the reactants of the fuel ion distributions, the cross section. And then you can get-- this is actually an example. Reactivity is calculated according to that formula. That's written up in this paper, which if you haven't seen this paper before, this is really a key reference to go to if you want to look at how likely a reaction is.

Obviously, the key is the most probable. Then we have the 2D reactions. We have a probability. And you can go to lower probability, [INAUDIBLE].

OK, yes. Go ahead, John.

AUDIENCE: How do you deal with the fact that you're technically not-- so, when you do this calculation of integrating over 4 pi, you're assuming some spherical symmetry. But we know for a fact--

MARIA GATU JOHNSON: Great question. So in ICF, actually, you can typically assume 4 pi. There are some variations, which we can look at-- will look at later in the talk. But in magnetic confinement fusion, you don't have 4 pi symmetry. So then you have to correct for it.

And actually, this leads to the work that you are doing. This is an example from a paper in 2010 by Sjostrand, where he's using a neutron spectrometer to infer the total yield of neutrons from the [INAUDIBLE] tokamak. But he needs to correct for the emission profile by using the profile monitor so he can scale that single line of sight [INAUDIBLE] neutrons to what the [INAUDIBLE] emission would look like. Great question. OK.

OK, so, thinking about what happens to the particles after they're born, most nutrients escape an implosion and can be counted. But some of them scatter. We'll talk more about that later. But for charged fusion products, like from the helium 3 reaction, stopping in the assembled fuel has to be considered. So the nutrients, they might scatter on their way out, lose some of their energy. But most of them escape directly. Scattering probability is relatively low.

Any charged particles are going to lose energy as they traverse the assembled field. And in fact, an example of what that can look like, we have the helium 3 critical spectrum be born in 14.7 MeV. This is the lower [INAUDIBLE] implosion, lose just a little bit of energy on its way out of the capsule. For high convergence implosions with high rho R, the ions will be fully stopped and can't be counted. So then we really can't rely on measuring primary fusion products, charged fusion products, for those experiments.

And actually, it turns out that ion stopping plasmas is a rich research topic that our group has been working on. I have a few example references there, and there are a lot of other references. OK, let's see.

Yes, so coming back a little bit to the knock-on reactions. So the neutrons, when they scatter off the fuel ions, they will also up scatter the fuel ions and energy. So fuel ions that are starting the thermal distribution, or basically, coal, can be up scattered in much higher energy by these 14 MeV neutrons.

The energy of the ion can be calculated according to that equation, where A is the mass number of the scattering nucleus. And theta is the scattering angle. En is the neutron energy. And this is the examples of what it can look like. So this is the scattered spectrum for tritons, deuterons, and protons. And basically, the energy of a deuteron tritonal proton that's scattered in this way is going to depend on the scattering angle. So this is what you get if you integrate over all scattered [INAUDIBLE].

OK, and we can use this fact. We often measure these products. And we use a number of those products to infer areal density because the areal density depends on the ratio or scale, basically the opposite. The ratio of the number of knock-on ions to the neutron yield will be a function of the areal density. Does that make sense? The more fuel that's assembled, the more of the neutrons are going to scatter.

So we can use that relationship to infer what the areal density of an implosion is. But this only works up to a certain areal density. Like, the knock-on deuterons, for example, would be fully ranged out [INAUDIBLE] 200 milligrams. And we talked about before, it might be at the order 2 grams per centimeter squared. So this is really relatively low performing implosions that we're looking at.

So we often look at the neutrons instead. And we can derive the energy of the scattered neutrons to see that the neutrons carry information about the symmetry of the assembled plasma. Because again, the scattered neutron energy will depend on the mass number of the scattering nucleus and the scattering angle. So we find, if we do that math, that for a detector at a set angle [INAUDIBLE] implosion, neutrons that end up at the detector in a certain energy range are going to be sampling a certain part of the shell of the implosion.

So at the NIF in particular, we often infer that compression by looking at energy range 10 to 12 MeV, which is a really clean energy range of the neutron spectrum. I think I might have an example later where you can see. There are no other sources of neutrons that might contribute in that range. So by looking at the number of neutrons in that range, you really get a measurement of only the neutrons that are scattered.

OK, so if you do that, if you look at the range from 10 to 12 MeV, the neutrons are scattered out from tritons. It's part of the shell. Neutrons are scared of the [INAUDIBLE] part of the shell. And it's a little bit different because deuterons and tritons have a different mass. And you can also broaden the range, and actually, we'll look a little bit at neutron imaging. Neutron imaging typically looks over a broader neutron energy range in order to get enough statistics in the sample the broader part of the shell.

So OK, this looks nice and simple. In reality, typically, the source of [INAUDIBLE] is significantly broader. So this smears out. You're not going to have all the neutrons coming exactly from the center. But this is the basic idea.

The unscattered energy spectrum, the neutrons that go straight out, like the green arrow here that don't lose energy on the way out, will carry information about ion temperature and velocity of the assembled fuel. And actually, in principle, this will be true of any fusion products. When they're born, they will be born with this information. The problem is, if it's a charged fusion product, it's going to lose a lot of that information as it's losing energy on the way out of the capsule. It's going to be a lot harder to infer that information.

This is the expression for the energy of a single fusion product. Then we can make the moments of this energy over the distribution of reactant plans to find that the width of the neutron spectrum is proportional to the ion temperature. There's a small ion temperature related peak up shift of the spectrum. And the peak will be shifted to the direction of flow of the emitting plasma.

There's actually one piece of key information that we've gotten from neutron data at NIF. They found early on that the capsule, when it was pushed through the lasers, it ran off in one direction, basically, based on the neutron spectrum being shifted, which we had to correct for because that prevented efficient conversion of the compression energy into thermal energy of the capsule.

So this is a non-relativistic expression. Really recommend you read this paper to get the relativistic math for how this works. I'm not going to go through it today, but this is an excellent reference which everyone that does neutron diagnostics for ICF uses all the time. And it's actually, originally, comes from the magnetic confinement community. So it's another example of the connections.

OK, so I think by now you've understood that the neutron spectrum provides information on areal density, iron temperature, and yield. And this is an example of what a neutron spectrum can look like. So the primary [INAUDIBLE] here that are unscattered [INAUDIBLE], the width of that primary spectrum is related to the ion temperature of the plasma. We can also-- not illustrated here-- but we can have an upshift that's related to the velocity.

And then by counting them, we get the yield, scaling up to 4 pi. And then by taking the ratio of the neutrons in the down scattered range, the neutrons in the primary range, we get a measure of the areal density. And actually, we often talk about a down scatter ratio rather than an areal density because we can measure the number of neutrons in this range compared to the number of neutrons in this range. And, yeah, go ahead.

AUDIENCE: How do you distinguish between neutrons that have down scattered within the fuel versus neutrons that have down scattered in the lab?

MARIA GATU JOHNSON: OK, so that becomes a technicality of the instruments. You have to collimate them really well in order to look at that. And it turns out the way that ICF is set up, the capsule is at the center of a large chamber. So the room return from the back wall of the chamber becomes a much, much smaller fraction of the neutrons that go down your line of sight.

So basically, the way it's set up for-- Chris is actually looking at exactly this problem-- for the magnetic recoil spectrometer, which we'll be talking about later, the foil is 26 centimeters from target chamber center, back almost 5 meters away. So the solid angle for scattered neutrons coming down the same line of sight is just so much smaller. It becomes negligible.

For the [INAUDIBLE], I don't think it's been [INAUDIBLE] I don't think it's been looked at in detail. But what they do is they take reference implosions so that they know if there's no assembled rho R, they know what the background in that range is. Yeah.

And then Chris is looking at the concept of putting numerous foil really far away. And I'm making him do simulations to see if that's going to work, or if we're going to have a problem with the returns [INAUDIBLE].

AUDIENCE: Why do you cut off the down scattered region at 4 1/2 MeV?

MARIA GATU JOHNSON: OK, so that's actually just kind of random. What we typically do when we do this is we [INAUDIBLE] This is not the best spectrum to look at. We're gonna see if you can see this on the camera. We look at DSR, or down scattered ratios, we call it. And that's the integral in the [INAUDIBLE] sort of see it? Integral in the 10 to 12 MeV range divided by the 13 to 50 MeV range. [INAUDIBLE] just get a quantitative number.

So that's this range here divided by that range here. And the reason we look at that range is because we have contributions from neutrons from the T2 reaction contributing up to 9 1/2 MeV. We have the D-D neutrons contributing here. You can kind of see that peak here. And there's also multiple scatters that kind of break the correlation between rho R and the number of neutrons that you get. So it's the cleanest reading to look at in the spectrum. So that's how that works. Any other questions?

OK, so I mentioned that we can also use the fusion products to look at the spatial emission. So we can take images of primary and scattered neutrons, for example, to provide information on the burned region, size R, and also the thickness of the high density shell. So in this case, this is actually a reconstruction, taking primary images in the 10 to 12 MeV range, down scatter images in the 6 to 12 MeV range, and then doing a fluence-compensated image, which gives us this artifact here, which gives us the picture of the neutron source, which is the primary neutrons and the high density shell, which is scattered neutrons.

And we can measure the nuclear reaction rate to get information about the confinement time and the bang time of the implosion. So what this is here is-- we call it the Lagrangian plot where you follow this simulation. You follow the same fluid element as a function of time. And then you see that the red is that interface between the capsule shell and the gas on the inside. This is for a gas-filled implosion example rather than the [INAUDIBLE] ETIs.

You drive it with the lasers. You get ablation of the surface material, which is why some curves are going off, and the other curves are compressing inwards until you get convergence. The shift you will show in particular moves inwards. The rest of it is converging. Get a little bit of burn here when the shocks hit the center, and you get more burn here where the capsule is at peak convergence, when it's maximally heated.

And then you can measure the emission history as a function of time. And you can see this shock burn and this compression burn. And this particular example [INAUDIBLE] implosion. So gas [INAUDIBLE], so then you often get both of these components. And then with the implosion, you're going to have very little shock. And you can have a lot more compression where we would be completely dominated by the [INAUDIBLE].

OK, so those are some examples of the parameters we're looking for. So with that, I plan to go into more about the technical detector details. So any questions before we move on? Actually, I have no idea how I'm doing on time.

PROFESSOR: Oh, you've got half an hour-ish.

MARIA GATU JOHNSON: That should be good.

PROFESSOR: Yeah.

MARIA GATU JOHNSON: OK, then let's jump into it. So the first one I thought we'd talk about is nuclear activation diagnostics. So they're typically based on indium 115, copper 63, or zirconium 90 isotopes for measurement of primary D-D or D-T neutron yields. If you look at D-D first, that's what we use indium. When a D-D neutron hits the [INAUDIBLE] indium, we get a isomer and the scattered neutron. The threshold for that reaction is about 1 1/2 MeV.

And this isomer state will decay, emitting gamma. But it has to have about 4.5 hours. And this is the gamma that we count to ensure how many reactions happen. On omega, we use copper to measure the D-T yield. And again, it's copper 63. And then the neutron [INAUDIBLE] neutron means copper will be an end-to-end reaction. So with copper 62 and two neutrons, threshold is about 11 MeV. We'll look at the shape of the cross-section. I think it's on the next slide.

And then what happens is that copper 62 is radioactive, will decay to nickel 62. And the half-life for this is 9.8 minutes. And what we actually count are the gammas here as well. Zirconium [INAUDIBLE] zirconium 90. We again get an end-to-end reaction. Threshold is about 12 MeV in this case, which means we're really narrowing in on the primary neutrons at 14 MeV. This is what we use at the NIF. And again, zirconium 89 is not stable. The end product that we get is gamma 909 keV, which is what we're counting.

And if you look at-- OK, so this first plot has the indium and zirconium reaction cross-sections. So you can clearly see why we use zirconium for D-T. The threshold is at about 12. Really covers our primary D-T [INAUDIBLE]. It's also really sharp though, which is actually a useful tool because if the peak is shifted up or down, it's going to impact what you're counting, which means you get an impact of velocity [INAUDIBLE].

You so you can see differences around the implosion. And then indium, on the other hand, is a really broad cross-section, which actually makes it a really blunt tool. If you want to use it to measure D-D, it's by far the easiest. In a pure D2 implosion, we don't have the down scattered D-T neutrons [INAUDIBLE]

In principle, you can use a cocktail of different nuclear activation detectors to piece together information about the full neutron spectrum. And here, we have some examples of parts of the spectrum that can be of interest. So in the D-T implosion, this is what the D-D spectrum look like. You actually have some of those secondary neutrons that we talked about before that are at 14 MeV. You have the primary D-Ts at 2 1/2 MeV, down scattered D-Ts, and then just a little bit of scattered in between.

And that we can get at with the indium in principle. We have the T-T neutrons, which have a peak at about 9 MeV and go from 9 1/2 all the way down to 0. Those you can also attack a little bit. And then, really, the primary thing we're looking at is the D-Ts, which some of the up scattered then-- and we looked at this before-- where the neutron has hit a fuel ion, giving it a lot of energy. That energy in turn reacts to produce another neutron. That's when we get these really high MeV tertiary neutrons, 15 to 30 MeV. And that's actually, in many cases, also really interesting measurements [INAUDIBLE].

Oh yeah, that completely fell off. There's another reaction here that has this cross-section here. It's an isotope of carbon, but honestly, don't remember which one. So then you can really focus in on just those highest energy neutrons. And this actually also shows you-- we have-- it's carbon-12. [INAUDIBLE].

Copper you kind of see is this orange line here. And we have zirconium as the red line. So zirconium is a much sharper threshold of 12 MeV. Copper starts already at 11. So they have a little bit different sensitivity to [INAUDIBLE] neutrons.

At omega, copper activation is used for measurements of the primary DT neutron yield. We have, basically, a little retractor tube that allows a puck to be inserted and then dropped after a shot by pushing a button. But it's still very manual.

So many times, I've been up there for a shot, and this old Russian guy, Vladimir Glebov, pushes the button, gets the black disk, and then runs it over to the counting detector in a different lab. And it's using sodium iodide detectors to detect the gammas in coincidence [INAUDIBLE].

And it's actually quite useful because then you go down to really low neutron thresholds. So cryogenic implosions at OMEGA produce upwards of 10 to the 14 neutrons, but you can measure neutrons down to 10 to the 7, which means you can look at experiments where you're not producing many neutrons at all and still know what you produce.

On the NIF, zirconium is the primary activation element used. And it's used routinely for measurements of primary DT neutron yield. In fact, from the high-performing implosions, there are two measurements that provide the yield that's then reported out. One is the zirconium nuclear activation, the other one is MRS.

So it's implemented in a number of different versions. We have the Well-NADs, which is kind of the go-to reference. It's inserted to 4 meters from target chamber center. There's three different pucks that sit very close to each other so you can compare the numbers from the three and make sure you're not making any mistakes.

And then you have the Snout-NAD, which you can insert much closer. And actually, it's more common to vary the elements in these packets and have the cocktails to look at different neutron interactions. And then finally, you have the Flange-NADs, which sit on the outside of the chamber. And there's a large number of those attached in different positions around the chamber to look at symmetries.

The zirconium detectors are transported-- well, actually this is kind of modified to the Flange-NADs-- we'll talk more about later-- are now counted in situ at the NIF, but they used to be transported. Zirconium detectors are transported to the Lawrence Livermore National Lab NAD Data Analysis Facility, which looks like this with some really old hardware but still does its job.

And then yeah, so the Flange-NADs has actually been converted, fairly recently, to 48 real-time zirconium nuclear activation diagnostics, or RT-NADs, that are permanently installed-- semi-permanently installed on the chamber with a lanthanum bromide detector counting the activation from those continuously. And you can see the peaks when there is an implosion.

But recently, it used to be 48, which is great. You can really look at the symmetry of the emission, which-- I think I'll get to this on the next slide, but the symmetry emission tells us about the aereal density symmetry. So we could look at the symmetry emission. The reason high-yield implosions have started killing these detectors, so we're now down to 21 version, and we actually have to remove them before every high shot. So they're no longer permanently installed on the chamber.

Yeah. So this is what I was trying to get to. The nuclear activation diagnostics often show large low-mode aereal density asymmetries in NIF implosions. So this is when you have this network of 48 detectors that all provide a measurement of the yield above the threshold of 12 MeV. If you're starting to see variations in that above-12-MeV threshold, that means that the birth distribution is uniform in 4 pi. So that means something must have happened on the way out, where some of them, more than others, were scattered on the way out, which means that the aereal density is not symmetric around the implosion.

And actually, turns out the typical scattered neutron fraction is about 20%. You can look at them, the number density times the cross-section times the shell thickness, you get just a rough measurement of how many neutrons will scatter on the way out. It's about 20%. [INAUDIBLE] frequently see variations of plus/minus 8% in the unscattered neutron yield, which means it's a large aereal density variation from one side of the capsule to the other. And this has also been a really useful diagnostic tool in figuring out what's going on with these implosions as we're trying to improve them further, make them perform better.

And yeah so I've mentioned before that you also have an impact on peak shifts here because the cross-section is higher and higher energy. So if you have a flow where the [INAUDIBLE] runs off in one direction, you can have upshift of the peak in that direction, downshift of the peak in the other direction. It turns out it's a smaller effect than the rho R asymmetries, but it's significant enough where it has to be corrected. First, you have to measure that directional flow and correct the distribution for that effect as well.

And we use low aereal density gas-filled DT exploding pushers to set the baseline variations, basically, as-- losing the word-- baselining. There's another word. Maybe it'll come to me later. OK. Did that all make sense, nuclear activation detectors?

PROFESSOR: So why is it that you use copper on OMEGA and zirconium on NIF?

MARIA GATU JOHNSON: Just historical reasons.

PROFESSOR: Oh, OK. Is one better than the other or that can [INAUDIBLE]?

MARIA GATU JOHNSON: So I actually-- the guy who runs the neutron diagnostics at OMEGA now would like to start using zirconium instead. And if I remember correctly-- yeah, the reason for that is the longer half life. It's really hard to work with a 10-minute half life. You have to really run to get to that detector fast enough, whereas zirconium, with the three days, is a little bit easier. And also, the threshold's actually better at 12 compared to 11. Yeah?

AUDIENCE: For looking at the upshift from zirconium, do you just compare that to a baseline where you have a more uniform emission profile versus a non-uniform one, where that steepness of the cross-section actually matters? Or do you compare it to a baseline cross-section that's flatter?

MARIA GATU JOHNSON: I'm not sure if I fully understand the question. But so what you're doing is you're looking in many different directions, and you compare the results in different directions. But you know--

AUDIENCE: So but you just-- you know your baseline just by assuming a 14.1 MeV uniform profile?

MARIA GATU JOHNSON: OK, so there's a couple steps to this. If you get a map kind of like this one, I mentioned you have to correct for the velocity, which is the peak shift. We actually don't get the velocity from this diagnostic. We get it from the neutron spectrometers. So if you have neutron spectrometers in six lines of sight, you can measure there. You don't just measure a number above the threshold, you actually measure the neutrons because you know what the option is, which means you can infer the actual 4 pi velocity vector. And then you can correct this for that. Yeah. Neutron spectrometers.

Any other activation question? OK. Then with that, let's go into neutron spectrometry. This is kind of touching on exactly that point. We do have a large suite of neutron spectrometers on the NIF. These five, the blue ones here, are all based on the same technology. They're neutron time-of-flight spectrometers, which we'll discuss in detail. And then this one in red is the magnetic recoil spectrometer, which I already mentioned a couple of times.

There's actually one more that's not included on this cartoon, which is also a neutron time-of-flight spectrometer based on a different detector technology that's fielded together with the neutron imager on roughly this line of sight. It's been a few years since it was working, but we're trying to bring it back to resolve some [INAUDIBLE] everything. So total of seven neutron spectrometers on the NIF that provide good implosion coverage together.

And this similar setup exists in other ICF facilities. Like on OMEGA, for example, think there are six now that run on DT in different lines of sight. So you can compare the results in those six lines of sight to, again, infer the flow vector in addition to measuring the ion temperature and the ion temperature variations. And fewer of them work on DT but still enough to get a good coverage.

OK. So let's start with the magnetic recoil spectrometer. So again, this is what I've been working with since 2010, so happy to take any questions on this one. And Chris is working on a very similar concept. Sean's kind of working on a similar concept, too, for Spark. So you guys are very familiar with this already.

But for those who have not heard about it before, the way it works, you will have the neutrons emitted from target chamber center, that little blue dot. A fraction of those neutrons will reach a conversion by the plastic conversion foil, 26 centimeters, in the case of NIF, from target chamber center. And this is actually deuterated foil. So the neutrons that interact with the foil, some of them will knock out deuterons.

Forward-scattered neutrons will reach this magnet, which is outside of the target chain wall. It's just a vacuum in between. So they all reach the magnet. And then there is momentum separated in the magnets [INAUDIBLE] different physical location of the detector, right, depending on their energy. Then you use that to reconstruct a recoil deuteron energy spectrum. And then from that spectrum, you can infer what the incident neutron spectrum must have looked like.

We use deuterated plastic, in particular because the detector we use in this instrument is CR-39, and it turns out the deuteron tracks are much, much easier to distinguish above background compared to using protons. Yeah. And we also-- there's a number of detectors with [INAUDIBLE] the sodium hydroxide, scan them in microscopes after the shot, and then stitch the data together.

Looking at-- zooming in on the foil here, so the neutron will hit a deuteron in the foil. And then the recoil deuteron energy will depend on the incident neutron energy, and then also the energy loss that the deuteron has from its place of birth until it hits the back of the foil. So a neutron born at the start of the foil would come out with a lower energy than a neutron born at the end of the foil. And that has to be considered in the analysis of the data.

We can look at, also, a couple of other aspects of this. So when I have-- we used to want a high efficiency to be able to count all the neutrons that came out. Today, we're actually running into saturation problems instead, so we don't really want this to be so high anymore.

But the efficiency of the MRS can be back of the envelope calculated as the foil solid angle times the number density for deuterons in the foil times the foil thickness the n, D differential cross-section for forward scatter, and the aperture solid angle where the aperture is opening in front of the magnet.

And you can throw some numbers on that. This is an example from the NIF. The foil solid angle will be the area of the foil divided by the total sphere at that distance where the foil is sitting. And then number density, we calculate based on manufacturer's specifications. Foil thickness is measured.

This differential cross-section in the forward scatter direction is roughly this number. Aperture solid angle take the area of the aperture and just divided by the foil aperture distance squared. And this is a correction for the fact that the aperture actually isn't sitting straight. It's tilted in front of the magnet. And then this is the [INAUDIBLE] of the [INAUDIBLE] foil which you also have to add as the correction.

This gives us a rough number. In reality, this is not what we use to get the yield number out. We use MCP simulation. Actually, I thought on the way over here, I should have included a slide on MCP because we use MCP a lot as a tool in understanding the response of the detectors that we're looking at.

Even for the nuclear activation detectors that we looked at before, to know how many neutrons that they see and how many might be scattered before they hit the nuclear activation detector, we also have to use Monte Carlo neutron transport tools such as MCMP. So it's not just building detectors. There's a lot of modeling that goes into this as well.

OK. We can also look at what we expect for the resolution. What that is, basically, is if you have monogenic neutrons emitted from target chamber center going through that whole system, then we're going to end up with a wider spectrum than just monogenic on the MRS. So we look at, how wide would that spectrum be assuming we had monogenic neutrons to understand the broadening-- the instrument of broadening.

So we can look at that as three components. We have a broadening effect due to the foil thickness. And this is, again, where deuterons are born on one end or the other, and they're going to lose energy as they go through, which gives you a broadening. We're going to have broadening based on the scattering geometry. We have neutrons that hit the foil head on, deuterons that go straight out. But then we also have neutrons that hit one edge of the foil and go at an angle. And that gives us a broadening effect.

And then finally, we have some broadening depending on the ion-optical properties of the magnet. And that leads us to a total broadening. And it actually turns out that these are counter-related. So if you want higher efficiency, you have to add more foil material, which also enhances the broadening. You want a narrow broadening, you want high efficiency, so it becomes an optimization problem. In the design of a magnetic recoil spectrometer, you have to balance efficiency and resolution against each other.

This is an example for the thin foil magnetic [INAUDIBLE] spectrometer at JET, which is very similar to MRS and maybe even more similar to what Sean is working on for Spark. So this is looking at proton collimator radius and foil thickness, and seeing how varying those two will impact the efficiency and resolution, trying to find some set points that would be ideal operation for getting the signal you need and still having a good enough resolution to make the measurements you want.

And that, actually, we can take a look at what that looks like at JET. This is the MPR magnetic recoil proton spectrometer. In this case, we do use protons not deuterons, because they use scintillator detectors they can count protons without any problems. This is what it looked like before the shielding was added. This is the magnet housing. And we add a lot of shielding to prevent contributions from scattering neutrons that you don't want to look at.

And it's the same concept as for MRS. You have the neutrons emitted from the plasma over here. You have collimators, so they only look at a certain fraction of them. The foil is inside the magnet housing here. The second collimators [INAUDIBLE] the protons are born in foil here. And then the protons are momentum analyzed in the magnet [INAUDIBLE] in a different physical location on the detector array.

In this case, it's an electromagnet rather than a permanent-- maybe I probably forgot to mention that for MRS. But for MRS, it's a permanent magnet. An advantage with an electromagnet is that you can tune it so you can operate at different energies.

The MPR, in particular, can be tuned to operate either for 14 MeV DT neutrons or 2 and 1/2 MeV neutrons. And we've used them with the JET tokamak with an oblique angle like that. If you look from the top, you can see how it traverses all the way through the plasma.

PROFESSOR: Why is the oblique angle used?

MARIA GATU JOHNSON: I'm not sure. That's an interesting choice. I think it might have actually been to maximize efficiency because you see more of the plasma that way.

AUDIENCE: OK.

MARIA GATU JOHNSON: It's a question for, Johan, though. He was involved in the actual building of this system. OK. And then so just look at what a recoil deuteron energy spectrum can look like, this is an example from the NIF. It's a really old example.

But you have the primary DT peak, and then you have the down-scattered neutrons. And then from this, you infer the total neutron yield by scaling for height, the aereal density by comparing the number of neutrons here-- deuterons here and deuterons here and the ion temperature from the width.

And what you do when you analyze this kind of data is you take a model neutron spectrum and fold it with the instrument response function simulated using GM-4 or MCP or a combination thereof and get it on the deuteron energy scale, and then adjust the ion temperature the peak position, which is related to velocity, the amplitude, which is related to yield, and the rho R, which is related to the down-scatter relative to primary diffraction. So you get those numbers out of the analysis. Yeah?

AUDIENCE: You're going past just fitting a Gaussian to this peak. You're using a more sophisticated analytic model?

MARIA GATU JOHNSON: So actually, in most cases, I fit a Gaussian to the peak and then just have a second component that accounts for the scattering. But yeah, there are some slightly more advanced models as well.

And for magnetic confinement fusion, you'd typically have to have more advanced models because you have fast ions due to heating that contribute to broadening of the peak. And to resolve those, you have to have a model for the beam-thermal reactions or the beam-beam reactions that have slightly different shapes. And you can find the relative contributions of those by fitting those different shapes to the peak.

Other questions? OK. So with that, let's look at the neutron time-of-flight technique. So for ICF, this is actually simple because all the particles are assumed to be emitted at the same time. The burn time is so short, order of 100 picoseconds, so you can make that assumption. So then you really only have to measure the neutrons as they arrive on a scintillator at a set distance, d, from the implosion. And you use that time to infer the neutron spectrum.

Yeah. And I mean, this in particular, it's already been converted to temperature expression. But really, what you're looking at is the neutron energy just on a time scale. Yeah. And the same as for MRS, the ion temperature is determined from the width.

And actually, so on OMEGA, the nTOFs are still used as the yield measurement, too. On the NIF, we gave up on that a while ago because it's so hard to know how the gain of the electronics [INAUDIBLE] drifts in time. So what we do is we calibrate it relatively routinely to the nuclear activation detectors and MRS. And since it's cross-calibrated, we no longer use it as the absolute yield number.

nTOF detectors are also used to diagnose aereal density. And I think already mentioned at some point that we use that by comparing to what we call zero rho R implosion, where we really just put DT gas in the really thin shell, drive it really hard directly with a laser, so we know there's no aereal density or negligible aereal density, and then we can look at the difference between zero rho R implosion and one with significant rho R to [INAUDIBLE] the rho R's.

And yeah, this is identifying on the time scale what the 10 to 12 MeV neutron energy range will be. The detectors at the NIF-- I think the closest one is 18 meters from the implosion, and the furthest one is at 27 meters from the implosion. This is what the original equatorial nTOF looked like. It's since been upgraded to look more like that. That actually looks like this, but it's kind of hard to tell.

The way it works now is instead of having these PM tubes directly attached to the scintillator, which is in this volume here, like [INAUDIBLE] you have the photomultiplier tubes facing. So implosion-- I'm the implosion. Detector is over here. The photomultipliers are facing this way so that neutrons hit the scintillator, and the light is collected in this direction, so you have [INAUDIBLE] contributions from scattering in the detectors themselves.

And there's four photomultiplier tubes at each scintillator, so you can have different settings on different photomultiplier tubes to optimize them to look in different parts of the neutron energy spectrum. There's collimators on the way. This is an example of how neutrons come through a wall collimator, with this detector.

In this case, you can also see that both neutrons and gammas coming through the collimator hit this high-sensitivity and fast detectors, which we call a spec detector, which is used to measure rho R. Also recently installed these quartz Cherenkov detectors to just a really thin rod in the same line of sight, which you can use to look at both the neutron and gammas. And these are actually more optimal for the velocity measurement because you get a really precise measurement of the primary structure from those.

Yeah. And it's similar at OMEGA. So this is at OMEGA. It's below the target chamber center in kind of a basement, which we call LaCave. It's this large detector which is a liquid scintillator material. It's quenched xylene, which allows you to have a really fast time response. So it falls off as quickly as possible after a primary peak, which makes it easier to measure the down-scattered neutrons.

Actually, another detail-- on OMEGA, the nTOFs times do not measure down-scattered neutrons in this energy range, because the rho R is much lower at OMEGA than on the NIF's much lower-power laser. So instead, we're using the backscatter edge of n, D backscattering. So neutrons that hit the back of the implosion scatter off of tritium and reach the detector on this side, which gives us an edge at 3.4 MeV, which is much easier to distinguish on OMEGA.

And so that's done with this detector. Again, for photomultiplier tubes, they're optimized for different ranges of spectrum. And look closely, you can see there's two detectors in front here-- one Cherenkov detector as well-- that thin rod-- and this pattern detector, which is one of the primary ion temperature detectors. So it has much better resolution than with large xylene detector. Any questions about that? Ben?

AUDIENCE: Is the thickness of the detector a significant source of uncertainty? I guess I'm guessing that the thickness of these detectors versus the length of these beam lines is really small to be negligible. But is that a source of uncertainty?

MARIA GATU JOHNSON: So it depends on what you mean with "uncertainty." I mean, so definitely, you get a broader spectrum from this large detector than from those thinner ones in front, which reduces your resolution, so it's harder to measure ion temperature. So that's why, like in this case, this detector is optimized for the rho R measurement, and this one is optimized for the ion temperature measurement.

This one needs high efficiency to get that weak component of down-scattered neutrons. This one needs high resolution to measure the peak accurately.

AUDIENCE: So similar to MRS, it's a trade-off between efficiency and resolution?

MARIA GATU JOHNSON: Yeah, yeah. Mm-hmm?

AUDIENCE: Someone asked at APS, and I didn't know why-- what are the benefits of MRS over the nTOFs if they all give temperature, rho R, [INAUDIBLE]?

MARIA GATU JOHNSON: So OK. My perspective, the primary benefit is having more than one technique because you really need to know independently what you're measuring. You can compare the results from both. And many times over the years, as one technique started drifting, and then we figure out what's going wrong by comparing with the other technique. So I think that's the primary advantage.

You can also say that one is that the MRS gives you the absolute yield. It's calibrated from first principles compared to cross-calibrated to other detectors. Yeah. But I think it's really important to have both.

PROFESSOR: But this data is available directly after the shot, whereas the MRS is a while.

MARIA GATU JOHNSON: That's true, that's true, which is a huge, huge advantage. And yeah. And this could be scaled up to rep rate as well because you can make sure you can analyze it quickly after a shot, whereas MRS would CR-39 indefinitely.

[LAUGHTER]

Well, Chris is working on electronic detection, so we'll get there. Other questions?

OK. So then the magnetic confinement fusion equivalent-- we do have a neutron time-of-flight system here, too. But here, it becomes more complicated because we can no longer assume that all the neutrons are emitted at the same time. So here, we have to have two sets of scintillators. We have a start scintillator, which we call S1 here, and a stop scintillator, which is S2.

This is what it actually looks like in real life. There's a collimator through the floor here. The detector sitting in the roof lab above the JET tokamak. So the neutrons come through the collimator in the floor, hit the start detector first and then the stop detector.

We have the start detectors layered to allow us to count at a higher rate. There's five layers in there. And then the stop detector is divided into 32 segments to actually-- that's more of a resolution [INAUDIBLE] thing because you want to know where the light is coming from in order to be able to measure the [INAUDIBLE].

And replacing them on the constant time-of-flight sphere so that you can compare the data from all the different detectors and stitch to make one spectrum. Yes, Kai?

AUDIENCE: How do you know if a neutron hits the S2 is the same one that you just measured at S1?

MARIA GATU JOHNSON: Great question. So you don't. And that's where this comes in.

AUDIENCE: Oh, OK.

MARIA GATU JOHNSON: So we look at data from the different scintillators in coincidence. So you take-- what this example is, is all its events that you can get in the S1 detector on the top. You get a lot more events in S1 because it's closer and it's directly in the beam of neutrons from there. The S2, the beam actually goes through in the hole in the center, so it doesn't hit the S2 directly. The S2 only sees scattered neutrons. But that means you get a lot fewer events in S2's.

And what you do is you go through and look at coincidences between the two detectors. And actually, what you get is you get all the coincidences. You get the true coincidences, and you get background random coincidences. And you have to subtract that back out. But when you do that, the peak will appear because that's then the true correlation that it will be the same between all neutrons [INAUDIBLE]. You were saying?

So in this case, the flight time for a 2 and 1/2 MeV neutron between S1 and S2 is about 65 nanoseconds.

AUDIENCE: But the beam drift time or the amount of time that the neutron spends moving between S1 and S2 is very short, because looking like these detectors, they're very--

MARIA GATU JOHNSON: 65 nanoseconds.

AUDIENCE: So does this function for DT neutrons, or is this [INAUDIBLE]

MARIA GATU JOHNSON: So it's best for DB. So this, it will [INAUDIBLE] they show up at 27 nanoseconds. The time resolution isn't anywhere near as good simply because the flight path is shorter. If you wanted really good resolution, you'd have to make the flight path really long [INAUDIBLE].

But of course, another difference here is here, we didn't really make that point with the inertial confinement fusion nTOFs. But what we're looking at here is running the scintillators in current mode. We're just opening them up, looking at the signal current as a function of time.

In this case, we're looking at individual pulses from single neutrons interacting with the scintillator. So we can divide it into-- we recorded over the entire duration of the pulse and we can reconstruct neutron spectra for any time interval we want where we get enough statistics. So we can actually look at the time evolution of the neutron spectra this way.

And at ICF, we simply get too many neutrons at the same time so it becomes complicated. But Chris has spent quite a bit of time trying to figure out how to do that, too. OK. So that's all I plan to say about neutron spectrometer. Any more questions before we move on? OK.

OK, I think I have just a very short section on neutron imaging. So what we do in neutron imaging is we use a pinhole or aperture close to the implosion and the detector really far away to obtain good magnification. This is actually very-- we have the source again here, which is a target chamber center, which is where the neutrons are emitted. You have a lined aperture, which can either have penumbra, which will encode the signal, or a simple pinhole, where you have a direct correlation, basically, between the neutron emission and opposite, kind of inverted to the detector.

The magnification will depend on the pinhole standoff distance and the detector distance. OK, so I took this slide from somewhere else, and I don't know what numbers they actually threw in there. Oh, assuming a magnification of 200, which you would get depending on what L1 and L2 values you have. If you have 5 microns at a source, it's going to be 1 millimeter at the detector, which magnifies your radius. I talked about before, that implosion would be 30- to 50-micron radius. You magnify them so you can separate individual features much easier on your detector on the outside of the chamber.

And actually, for the NIF in particular, I think typically the aperture is about 20 centimeters from target chamber center. The detector is 28 meters away. Yeah, this is what-- ha, actually, those are the exact numbers. So this is what it looks like. You have the NIF target chamber over here. That's target chamber center. The aperture is fielded right here, 20 centimeters from target chamber center.

And then the neutrons that are selected by the aperture will travel through this collimator structure all the way to the detector back here 28 meters away. And so the neutrons-- this is just a [INAUDIBLE] it looks like. The neutrons come out here through the line-of-sight collimator. This is that other nTOF detector I talked about. You can kind of see that it's a different technology. It's flat plastic scintillators with photomultiplier tubes.

But then the primary-- so that that's just another neutron spectrometer. The primary for the imaging system is the scintillating fiber array, which is fiber coupled to a camera and it's done this way in order for you to be able to gate the camera and get two snapshots. So you can get the primary neutrons and then the scattered neutrons at a later time.

There's actually also-- this is the original NIF neutron imaging system. There's two more now, so we can look at symmetries around the chamber. One of them only has image plate detectors, which I planned to bring image plate, but I forgot. But it's--

PROFESSOR: We discussed it, actually. We talked about X-ray diagnostics, so we talked about it a little bit. Yeah.

MARIA GATU JOHNSON: Yeah, so you know you can't time gate on image plates. So then you just get one image. OK. And I wish I had a better picture. But these pinhole apertures are actually extremely complicated.

So they're made of gold. They're about that long. And you have to make pinholes that are precise all the way through. It's too hard to drill them circular, so they make them triangular instead. And then they're tapered to minimize scatter. So basically, you have an opening, and then the neutrons that go through that opening are all going to be captured at the back end. None of them are going to stop because of the taper inside.

It's still a really hard machining problem to get those even triangular pinholes precise all the way through. We need gold because, as we talked about, neutrons have a pretty low likelihood for interacting in matter. And you want to only select the ones that go through the pinhole. You don't want all the ones around to also make their way all the way back to the detector.

Yeah. So that's fun. And these are examples of the penumbral apertures where the information will be decoded in the penumbra. You'll also get straight through neutrons in [INAUDIBLE] that you can't use to infer anything about the shape of the implosion. So this kind of aperture array is used on all three lines of sight now, but there's development going on to make it coded aperture, which is supposedly going to be simpler and thinner and easier to manufacture. We'll see if that actually works out.

See. Oh, yeah, so these are examples of primary and scattered neutron images. And again, they're routinely obtained. And all DT NIF implosions in these days, it's even in three lines of sight. And there's a lot of work going into tomographic reconstructions to make sure we understand the full emission region.

And we have our equivalent in magnetic confinement fusion, which John is working on. And we call them neutron profile cameras. So it's, again-- the idea, again, is to probe the shape of the neutron source distribution.

But it's much bigger here and more complicated. So instead of that pinhole array that's trying to reconstruct the 50-micron spot, we're looking at a much larger emission. And we do it by using a number of different lines of sight and counting particles along this line of sight, and then try to reconstruct the full emission [INAUDIBLE]. And John can answer a lot more questions here.

OK. So with that, I'm going to jump right into charged-particle spectrometry. And I think I touched on this in the beginning, but it's routinely used to diagnose low to medium rho R implosions. And they can be filled with deuterium, DT, or D helium-3 fuel. And the reason we don't typically use it for high rho R implosions is, again, because the charged particles stop in the assembled fuel, so they don't become indicators on the outside.

OK. This example is showing the proton from the interaction for 14.7 MeV downshifted to about 11 MeV You can look at this downshift to infer the aereal density, which in this case was about 84 milligrams. And you can-- actually, let's see. Yeah. I have an example, one with nice [INAUDIBLE] spectrometer.

The cool thing is you can do this thing in a number of different locations around the target chamber to look at symmetry again to see if things are compressed uniformly or if there's non-uniform issues. So this is a super simple spectrometer. We call it wedge range filter spectrometer. It's just an aluminum filter that's shaped like a wedge. Pass this around. You can see it. So the way this is fielded, it's got a bunch of holes in front.

And you know the holes are used to register where the detector is fielded behind that wedge. You know the thickness of the function of precision on this. And based on, actually, the hole size of the CR-39 detector as a function of precision relative to those holes, you infer the proton energy spectrum.

And it's a little bit more complicated than it sounds because you need to know the diameter versus energy response of CR-39. And that varies from piece of CR-39 to piece of CR-39. So you have to come up with a pretty intricate method for inferring that from the data. But that's done, and these are routinely used to measure rho R from the helium-3 gas-filled implosions in many different locations around the target chamber.

Pass that around. It's a small 5-centimeter round packet. So you really can get it in many locations. And this is also the detector material that's used for MRS, as we talked about before, CR-39 plastic. So the wedge range filter spectrometer is one example that you can see here, too.

We also have charged-particle spectrometers which are very similar to MRS. It's a magnet outside of the target chamber wall. The difference is we don't have a conversion coil, so we're just looking at charged particles directly from the implosion. And you can actually-- this is another advantage of MRS. You can also run MRS in charged-particle mode for experiments where we're not interested in the neutron spectrum. And then you can look at the charged particles that come directly from the implosion.

Yeah. And this is an example of looking at that symmetry and how it can vary around the implosion. This is an insertion module on the NIF, where we can field, actually, up to six of these wedge range filter spectrometers on a single insertion module. There's four insertion modules that have capability of fielding these. So you understand we can field a lot from one implosion. On OMEGA, we can field up to seven on one implosion in different directions. And in each direction, you can add a few more if you want. So a lot.

Yeah. So in this case, these are fielded at 50 centimeters from implosion. These are fielded at 10 centimeters from implosion [INAUDIBLE]. So we can look in different directions. And this is actually a really old example from the paper by Johan in 2004, where he's fielded protons with [INAUDIBLE] in different locations around the OMEGA target chamber [INAUDIBLE] look in different [INAUDIBLE] spectrum.

And if you think back to this, for D3, you often get these two peaks in time, a sharp peak and a compression peak. So you can also say something about the evolution of the experiment. Here, you have the time evolution. You see the small, sharp peak and the larger compression peak. And then you can also see that in the energy spectra, you have less range down-- more range down compression profiles. So you can tell the difference in aereal density between those two types of implosion. It's kind of neat.

OK. I feel like I'm running out of time, so I've got to speed up. We already talked about image plates, so don't need to talk about that. CR-39, I kind of touched on.

This is an example of what it can look like after we've etched it in sodium hydroxide at 80 degrees Celsius for of order hours. If we put it on one of these microscopes, step over, and take pictures for roughly 400-micron frames, the microscope automatically picks up tracks which are due to particles interacting in the CR-39, and records their roundness or eccentricity and the track size. And then we use that to reconstruct whatever information they wanted from that diagnostic.

OK. So then, spend a few minutes on reaction-rate history. So there's a number of ways to do this. And again, so we're looking at really short burn, order of 100 picoseconds. We can field a plastic scintillator really close to the implosion and combine it with the streak camera to measure their reaction-rate history.

This is done at OMEGA. So this picture is of the OMEGA target chamber. This is the laser. So this is not part of the diagnostic. This is how we drive the actual implosion. This is what the detector will look like. So it will be fairly close to the target chamber. There is the plastic scintillator. The light from the scintillator will be coupled through an optical light path [INAUDIBLE] camera to record a streak image. And this is where the burn history is encoded.

And that scintillator will have a rise time of about 20 picoseconds but a fall time of 1 and 1/2 nanoseconds. So the information, really, is encoded in the rising edge of the signal. So we have to unfold it. But when you do that, you can get the burn history as a function of time for [INAUDIBLE] compression.

And this is used a lot for neutrons on OMEGA in the neutron temporal diagnostic. There's also the particle temporary diagnostic or the particle and X-ray temporal diagnostic-- similar concept, where you can tweak that scintillator setup in the center to have a number of different channels, some of them optimized for X-rays, some optimized for protons, some optimized for neutrons, depending on how you filter them and what neutron density focus you put behind [INAUDIBLE] on the streak camera and reconstruct it after. So that's actually a really neat diagnostic that's useful for a lot of things.

I told you, promised you, early on that we'd get back to what we used to gammas for. So one cool thing about the gammas is they don't have the same time dispersion that neutrons do. So the neutrons, when they're emitted from target chamber center, they're going to disperse in time, which is really why we can use just the neutron time dispersion for neutron time-of-flight. So if you put the nTOF detector 20 meters away, look at the neutrons, we measure the energy spectrum, not the emission history.

For the gammas, they don't disperse in time, so we can have a gamma detector relatively far away and we still retain that time history. So that's the cool part. We have the lower probability for getting gammas. I think we saw the branching ratio is about 10 to the minus 5 of the neutrons.

But there is enough of them to count. And by counting the gammas-- in this example, we're using the gamma reaction history detector, which is based on converter gamma rays that are converted to electrons. And electrons generate Cherenkov light in this gas cell, and then it's detected as a function of time. And that's how you infer the bank time or burn history. You get both from this measurement.

This is what this detector looks like on OMEGA, and this is what it looks like on the NIF. And on the NIF, there are four different channels, which you can vary in gas pressure set at different gamma detection threshold. You don't get spectral information, but you can set a threshold and then compare the results from the four.

And actually, it even made it into a movie, this detector. This is the gamma reaction history detector. Hans Hartmann, who built this detector, was really proud to be able to take his kids to this movie at the movie theater.

[LAUGHTER]

OK. So coming back to magnetic confinement fusion, again, here we typically use fission chambers to measure the nuclear reaction rate. And the way this works is you have the fissile material, you get the fission products going into the fuel gas. They ionize-- or, sorry-- yeah, they ionize the fuel gas, and then you get an electric pulse that goes out. And you measure that as a function of time.

And here again, you need MCP model in order to determine what your measured signal actually means in terms of nutrition. These are often also in-situ calibrated, where you move the source around inside the chamber and see what the signal looks like on efficient chambers on the outside.

OK. And I think have five more minutes, Jack, right?

PROFESSOR: Go for it.

MARIA GATU JOHNSON: So we'll take a few minutes on the impact of nuclear measurements on the ICF program at the NIF. And I think we've really touched on this throughout the talk today, right? The nuclear data have been essential for guiding the initial experiments to ignition.

This is the time axis of the yield from the experiments from 2010 through now to 2024. And you can see that this is a logarithmic scale. It's obviously increased a lot over that time frame. And MRS has been part of it from the beginning.

I started here at MIT in August 2010. I think the first data from the NIF came back from MRS a week-- two weeks after I started. They shipped it back in this huge moon lander looking container. It was, like, octagonal box with lots of cool packs to keep this [INAUDIBLE] cold. And we had to work day and night to etch and scan it and turn it around. And that's when we're at this yield level, right? Not registering on the scale.

And then we've been working our way through up to the regions where we actually have target gain. And we've looked at many of these diagnostics today that have been essential for [INAUDIBLE] experiments to ignition. We looked at how we get ion temperature, hotspot velocity, fuel density or aereal density, and yield from the neutron spectrometers.

We looked at how we get the burn width and bang time from the gamma reaction in-situ detector, which is related to the confining time. We also get neutron yield from the activation detectors as well as the map of the fuel uniformity from the real-time activation detectors. And we use neutron images to get the hotspot and fuel shield shape.

And really in particular, these two have been essential for identifying those asymmetries. And seeds to asymmetries have been really hard to eliminate along the way to get there.

OK. This is an example that Johan put together two years ago. On August 8, 2021, an implosion experiment at the NIF ignited and generated a then record neutron yield of 4.5 10 to the 17 1.35 megajoules. This is the MRS spectrum from that particular experiment. We were so excited about that experiment, which really, internally in the community, that was ignition. And I'll explain why in the next slide.

But so this explains what I was talking about before. We have a model neutron energy spectrum, which is a Gaussian with a width governed by the ion temperature, the mean energy determined by the birth energy plus the peak shift, which is related to the velocity. And then we have this component, which is related to the aereal density. We vary those parameters to fit into our mesh and reconfigure an energy spectrum, and then we get a best fit neutron energy spectrum which explains what we actually had.

And so if we look, in particular, at this implosion from 210808, many of the key nuclear observables point to this implosion being in a fundamentally new regime. We saw how the ion temperature took off. Earlier implosion under the same campaign had 5 keV. Now we measured neutron average ion temperature 10 keV-- a dramatic step up.

We saw how-- the burn history is a little bit hard to see, but what happened actually is the burn history peaked later and got narrower. And what this is saying is that the yield took off during compression. So where the yield would have previously tanked, you start having it climb more and more instead. And then it becomes a really narrow burn because it's burning on as the explosion is exploding rather than as you're compressing.

And that, you can see in the neutron images as well. These are neutron images and as well as one X-ray image or the top row from two predecessor implosions. This one, you can see how it gets a lot bigger. And this is, again, because it's going on the expansion.

Yeah. And then the other one-- well, you're actually measuring-- this is not the best spot to show it, but what we're finding is we're actually measuring a lower down-scatter ratio. And the reason for that is we're now probing-- the density ratio that we're measuring is probing the timing the implosion after peak compression. So even though the actual down-scatter ratio-- the actual compression is the same, we're seeing fewer scattered neutrons because the peak compression is before the probing. Does that make sense?

OK. And yeah, this is just another illustration of the same point, where you really see the temperature climb or jump. This is the fusion yield on the y-axis. This is an inferred hotspot mass and hotspot energy, also jumped for this one implosion. Neutron radius increased and burn rate decreased, really showing that we're in a new regime.

And that was this one couple years ago. So definitely done some better ones since then. And there's another one from October 29 that's not yet on this chart, too, which is the second best performing ever, so falls right between these two. And we're at 4 with gain over 1 at this point.

OK, I think that's all I had for today.

PROFESSOR: Thank you very much. Any other last questions?

MARIA GATU JOHNSON: Yeah?

AUDIENCE: What's next to try to go even higher on the gain? Because I think a lot of the changes were capsule quality, et cetera. And is there anything that you're seeing in your new neutron data from this new regime that's guiding further changes?

MARIA GATU JOHNSON: Good question. I don't think there's anything super obvious right now. Part of it is pushing for bigger implosions, which is not directly related to the neutron data. Yeah. Yeah, no, I don't think there's any defect signatures right now that we're going after.

AUDIENCE: Are there any open questions that you see in the neutron data [INAUDIBLE]?

MARIA GATU JOHNSON: Yes, there's a big one. Great question. So actually, I talked a lot about peak upshifts and how we infer velocity from that. So there's two aspects to that. There is the peak shift that's different in each of the different lines of sight that's showing us in which direction the implosion's taking off in.

But also turns out that there is a uniform [INAUDIBLE] that's the same in all lines of sight that's anomalous, that's not explained by the ion temperature and not explained by the direction of velocity. And right now, it looks like the only way to explain that upshift is by non-Maxwellian effects in the fuel line velocity distributions. It's not clear why those would arise.

So that's definitely a big outstanding question that we're looking at, which excites me a lot. I like puzzles.