REUBEN: I'm Reuben. I am currently a senior in Chemistry. And I am a UROP in Bob Sauer's lab over in the Department of Biology. As you guys know from this class, one of the modules that you cover is ClpXP, the bacterial analog of the proteasome. And today I hope to tell you about some of the cool work using modern biophysical techniques that people in Bob's lab have done to gain some really, in my opinion, interesting insight into the actual mechanical mechanism of this motor protein ClpXP as it unfolds and degrades protein substrates.

Before I get started, I want to say, please stop me at any time to ask questions. If I say something that doesn't make sense, please let me know, and I will try to clarify. So today I'm mostly going to be talking about modern single molecule methods, which are the cutting edge in biophysics. And so the first question I wanted to ask you guys is, what are the possible advantages of looking at molecules one at a time rather than performing some sort of bulk assay? Do any of you have any ideas?

AUDIENCE: Maybe the average of a bunch of different molecules isn't truly representative of what a single molecule looks like.

REUBEN: Yeah. So the average is the average. But the average as a statistic has some sort of weakness. It obscures a lot of the vagaries of behavior that may be lost in that average. So it's worth saying that in classical biochemistry you are looking at a lot of molecules simultaneously. In one microliter of one micromolar solution, that's 6 times 10 to the 11th molecules. So taking some sort of bulk measurement averages across all of these molecules. And as you said, you can lose a great deal of information about the variation and the dynamics of these molecules.

So another thing that you can lose, in addition to information about variation, is you can lose information about the procession through some sort of complicated biological process. So let's say that you are studying ClpXP as it recognizes, unfolds, and degrades protein substrates. If you were looking at, for example, 6 times 10 to the 11th molecules of ClpXP sitting in a test tube, then you're smearing across all of these different time points of ClpXP, where you have

some ClpXPs that are unfolding substrates, some ClpXPs that are not bound to substrate at all, some ClpXPs that are translocating substrate. So any sort of readout that you have loses this information about the different states that ClpXP can inhabit. And also, it's very difficult to gain detailed kinetic information about how ClpXP would, for example, transit between these different states if all you can see is a bulk average.

So another way of putting this is that a bulk measurement like fluorescence, you're looking at the unsynchronized activity of these molecules. But by studying single molecules at a time, you can do some sort of post hoc synchronization to actually gain insight into the kinetics, or the order of the different states during some sort of biophysical process.

JOANNE STUBBE: You can do better than that, even in bulk you can synchronize, right?

REUBEN: Yeah. So there are certainly tools like stopped-flow.

JOANNE STUBBE: Which is what we talked about in the first couple of recitations. That's why I brought it up.

REUBEN: Yeah. So stopped-flow is a fantastic tool, which has been used really successfully to study a lot of systems, particularly, I think, the Rodnina lab, which I think you guys covered, used stopped-flow very effectively to study the early stages of translation. But I think that stopped-flow is good for looking at some early stages of processes, but for trying to track a long, complicated process across a significant period of time, stopped-flow is not really an ideal technique.

Because what happens is that the rates of transitions between different processes-- let's say that there is a single rate constant. The actual time that an individual molecule spends in a state before switching to the next state is stochastic. It's an exponential decay process based on the rate constant.

So if you are looking at a bunch of rates all put together, and you start a stopped-flow experiment, it's true that for the first couple of seconds all of the molecules are synchronized. But over the rest of the experiment, they blur out all the way across the time space. And so you can really lose a lot of detailed information about the kinetics. Whereas if you're looking at molecules one molecule at a time, you can track an entire process.

JOANNE STUBBE: But again, the issue where you can only look at them in one way--

REUBEN: Oh, absolutely.

JOANNE STUBBE: --ways. And so, if you don't know anything about the bulk system-- I hate it when people are touting single-molecule, which I think is very powerful, but the fact is that you have many ways to look at reactions that can't be carried over into the single-molecule venue. So that's the caveat.

REUBEN: I completely agree. I was actually just about to, after I finished this slide, ask what were the disadvantages of single-molecule.

JOANNE STUBBE: All right, sure.

REUBEN: And you raise a very good point. No, I always appreciate comments from the audience--

[LAUGHTER]

--and from the faculty in the audience. I don't know. It's true that there are some measurements you can't make in single-molecule. But it's also true that there are some measurements, such as the measurements in this paper, that you can't make in bulk. So you can measure some unique properties of molecules-- particularly molecular machines such as force generation and processive motion-- using measurement techniques such as the optical tweezers, which I'm going to talk about a lot in this lecture, that there's no bulk analog of a measurement like this. So JoAnne mentioned a couple of the possible disadvantages of single-molecule studies. Do any of you have any other ideas of potential disadvantages that might make sure that you want to do some bulk experiments?

So the really big one is that single molecules, there aren't great ways to look at them. There's single-molecule fluorescence, there's optical tweezers. After that, you know, I'm not really sure what else there is. And both of these techniques study from a significant noise issue. You can really only look at one thing at a time. And a single fluorophore, for example, in a single-molecule fluorescence study, is not very bright. You have enormous noise coming from your instrumentation. And it can be very difficult to tangle out the signal that you're trying to go after from your noise.

Another major issue is that many of these experiments are actually very difficult to run, and it can take, for example, months and months and months to do single-molecule experiments, which, you know, a bulk protein degradation experiment, you can mix some proteins and run a gel, and actually learn a great deal about the behavior of the proteins. But all the same, there are some measurements, such as the measurements in this paper, that you cannot take

without single-molecule measurements. So I hope that today I will be able to tell you about some modern biophysical techniques, particularly optical tweezers, so that you can all learn about some of the latest tools in the biophysicist's toolkit.

So I am going to talk about optical tweezers today, but I did want to mention that there is another single-molecule method, single-molecule fluorescence. And this is actually what I do in my experiments. That's also very useful for looking at the dynamical properties of complicated biological processes. So single-molecule fluorescence basically involves tracking the fluorescence of single fluorophores over time, and quantitating it, and then studying the dynamics of the switching between the different states of the fluorophore.

Almost all of these experiments use a physical property called FRET, which is energy transfer between two nearby fluorophores, the efficiency of which is distance dependent. I believe that you guys have learned about it. It's not that difficult to understand how you could go from doing FRET in bulk to FRET in single-molecule. If you haven't heard about FRET, I'm sure you--

JOANNE STUBBE: We'll hear it in the last recitation.

REUBEN: I'm sure you will hear about it very soon.

JOANNE STUBBE: [INAUDIBLE].

REUBEN: Yeah. So this is a very useful technique for answering some questions, but in some ways it's a little bit easier to understand than optical tweezers. So I figured I should talk about optical tweezers today, because I think it's also kind of cool that you can do the experiments that are in the paper today. So optical tweezers-- the main technique used in the Olivares and Sauer paper that I'm talking about today. And the general idea of optical tweezers is that you can use the momentum of light. Because as you guys know from quantum mechanics, light has momentum to trap certain types of particles within a beam of light and directly apply forces to them and directly make unique measurements about distance using this tool.

So with optical tweezers, you can measure nanometer motions at sub-millisecond time resolution. And in a couple of slides I'm going to tell you how that works. And you can also directly apply force to probe mechanics and biochemistry. So you guys have probably all heard about optical tweezers, even if you haven't heard about them in the case of studying single molecules of protein.

So this is a really cool video that I found. These are 12 beads, all trapped in their own optical trap. And some graduate student wrote a program to steer one of the beads around the other beads. So the system that we used in this experiment, that I'm going to tell you about today, is basically the same as this system, except you're not looking at 12 beads at a time. And each of the beads is conjugated to some sort of molecule of protein.

So the general setup in single-molecule optical tweezers experiment in biophysics is to take a bead-- usually the bead is made of polystyrene, because it happens to have some nice properties-- and to conjugate onto it a single molecule of protein. In this case, they conjugated onto it a single molecule of kinesin, which is a motor protein.

JOANNE STUBBE: [INAUDIBLE] like polystyrene shrinks and it traps that use polystyrene columns. Is that--

REUBEN: These beads are functionalized on the surface, and then covered with streptavidin, which I'll talk about in just a second. The beads don't really shrink very much, because you're not really putting them under pressure in any of these experiments. And so what they did in these experiments is that they took kinesin, and they put it on the surface of a bead. And then they took some microtubules, and they put their microtubules on the surface of a cover slip. And they brought the kinesin on the bead close to the microtubules on the cover slip, and they added ATP.

And kinesin is a molecular motor which walks along microtubules. So what happened is that kinesin grabbed onto the microtubules, and it began to bind and hydrolyze ATP. And it began walking. And it dragged the bead with it as it walked along the microtubule. And with an optical trap, you can actually detect how far from the center of the trap a bead has been dragged with extremely high, less than one nanometer, precision.

So they were able to tell that this kinesin was taking very discrete eight-nanometer steps as it walked along the microtubules, which could be related very effectively to the size of these tubulin domains that the kinesin was actually making contact with. So using this they could, for example, study the average time in between kinesin steps as a function of the concentration of ATP. And they could begin to ask some questions about how the binding and hydrolysis of ATP actually led to kinesin's processive mechanical motion.

So before I tell you more about the biophysics, I just want to tell you a little bit about how optical traps work. Because I think that it's not particularly intuitive before you hear about it.

Although, in the scheme of things, they're actually not that complicated. So I'm going to tell you three ways of thinking about optical tweezers, all of which are basically the same, but use different approach to look at the behavior of the system.

So I'm just going to start by pretending that light is just a ray, not a wave, because that's a pretty good approximation in a lot of circumstances, including this one. So the way to think about optical tweezers is to imagine two rays of light, one coming from each side of a lens. And both of these rays, they have momentum, because they're light. And to think about these rays of light interacting with a dielectric particle such as a polystyrene bead, which has a higher index of refraction than the surrounding media, which is usually water.

So what that means is that when the light hits the bead, the angle of the light will change. Basically the bead will deflect the light. And so what that means is that the bead is actually changing the momentum of the light. So what we know from Newton is that, if the bead is changing the momentum of the light, then the light must be exerting an equal and opposite force on the bead itself, opposing the momentum change of the light itself.

So we can see here that this ray A coming from the left is being deflected up and to the right. And this ray B coming from the right is being deflected down and to the right. So we can think about the force on the bead from each of these rays being equal and opposite to the direction of the momentum change of these rays. And say that ray A is imparting a force on the bead which is pointing down and to the left, and ray B is imparting a force on the bead which is pointing up and to the left.

So we can basically sum both of these forces and find that, in this system here, where the center of the bead is moved to the right of the true focus of these two rays of light, then the light is imposing a force on the bead which pushes it back toward the center. And it turns out that, whenever the center of the bead is not in the same spot as the center of the focused light, the light will impart a force on the bead, pushing it toward the direction of that center, no matter how you move it around within the trap. So do you all understand the ray approach to understanding how optical traps work? Cool.

So I'm going to give a slightly more realistic picture. So as you guys know, light is not a ray. Light is a wave. And what that means is that you can't focus light into just a point. Instead, you can focus light into a small volume which is diffraction-limited in size to approximately a diameter of one-half the wavelength of the light, which is often called the confocal volume.

So to think about the shape of the confocal volume, it's not quite a Gaussian, but it's very similar to a Gaussian. So it's useful to think about Gaussians, to say that, you know, it's easy to imagine a one-dimensional Gaussian or a two-dimensional Gaussian. At a confocal volume of light, what you actually have is a three-dimensional Gaussian in which each point in space has a scalar property, which is the intensity of the light at that point in space.

And if you move from the center out in any direction, it basically decays according to this Gaussian curvature. So there's a gradient, a Gaussian-shaped gradient of light intensity going out in every direction from the center, if this is like the optical axis, and say this is like the x direction, and this is the y direction. I don't know if my laser printing is particularly helpful in this situation.

So what you can think about is, say that from your beam of light in this Gaussian shape you have rays of different intensity. I've only shown two here, but try to imagine these rays coming from every part of this light beam, all being focused here so that rays from the center have a lot of intensity, therefore they have a lot of momentum. There's a lot of photons there. So they impart more force per deflection than rays coming over from the side. So if you imagine rays coming from this whole Gaussian volume here, then it's not too difficult to see that whenever you move the bead from the center, there's going to be a restoring force pushing it back to the center.

And the last explanation for how optical traps work is my favorite, but it's a little bit more difficult to understand, even though it's actually classical rather than quantum. So what we know-- a dielectric particle such as a bead, what a dielectric actually is, a dielectric is a material which is polarized by an electric field. So polystyrene is polarized quite a lot by an electric field. And what that does is that, if you have an electric field, say, that has a negative charge here and a positive charge here, then the dielectric will be polarized in such a direction that it is opposing that electric field. And the electric field within the dielectric is lower than the electric field outside of the dielectric.

And it turns out that, for very small particles, in this case particles that are actually smaller than the wavelength of light that we are actually using to trap it, you can approximate a dielectric material very well as just a simple dipole whose dipole moment is facing in the opposite direction of the electric field that it's actually sitting in. So think about what light actually is. Light is a rapidly oscillating electric field that also happens to have a magnetic component.

So the particle that you're studying basically has a rapidly oscillating dipole, which is always opposing the direction of the magnetic field. And so we know that there is a cost in potential energy to separating two charges. And putting a dielectric material in between that charge separation decreases the potential cost of this charge separation. So if you have a gradient of electric fields, as you do at the center of a confocal volume, the dielectric will be most favorable if it is in the place where it is opposing the strongest electric field. And if you move it from the center, then it will oppose a weaker electric field.

And we know that force is the gradient of potential. So the dielectric particle will feel a force driving it toward the strongest area of electric field, which takes place at the very center of the confocal volume. And so it will always be restored toward that spot. So this is kind of an intuitive explanation, but it takes a little bit of thinking about. If you have any more questions about this, I'm totally happy to answer them either now or after the talk. So do you all sort of feel like you understand how you can use light to trap a dielectric particle such as a polystyrene bead?

AUDIENCE: What wavelength of light are you talking about? Like, is this monochromatic?

REUBEN: Yeah. So in this case we're using monochromatic infrared light. The wavelength which is chosen for these studies, it turns out to not damage proteins very much compared to other wavelengths of light, and you're using it at very high intensity. And it's pretty easy to make very high intensity infrared lasers. But you could use, you know, the thing that really matters is what the index of refraction is at the wavelength that you're studying. So as long as you have a big difference, as long as you have a dielectric particle which has a higher index of refraction than the surrounding media at the wavelength that you're studying, then you can use any wavelength of light for this sort of experiment.

So I'm not going to go too much in depth into the actual apparatus for an optical trap, but I do want to just give you the schematic and talk about the very basics. So the general idea is that you have a trapping laser, which is being focused right here on the specimen stage, basically just a microscope where you actually have the bead trapped. And then you have a second detection laser, which is much weaker, 100-fold weaker than the trapping laser, and is at a different wavelength. And this is being focused onto the same volume.

And you have a quadrant photodiode, which basically gives off a different voltage in response to a different location of the trapping laser on it. And this is how you actually measure the

where the bead actually is within the center of the trap. So if the bead is moved out of the center of the trap, then this detection laser is deflected slightly, because it's also refracts through the bead. And so that shifts the centroid of the detection laser as it's read out on this position-sensitive detector. And so you can take this reading 100,000 times a second. And that's how you can get this very detailed, one-nanometer level of resolution, which describes where the bead has been pulled by some biological molecule out of the center of the trap.

So one other thing I should talk about, I want to give a couple of examples of things that people have used optical traps for. But it's worth mentioning that most experiments with optical traps are no longer done with just one trap. They're now done with two traps in which you have two different beads, each trapped within their own optical trap. And you do this because it basically mechanically isolates the system.

So normally you can have vibrations just in the room that would shake your specimen stage. And it turns out that the distances that you're measuring are so small that these vibrations dramatically increase the noise of the system. But if you have both beads trapped within their own optical trap, that provides significant damping from any sort of mechanical vibration. And you can measure the bead-to-bead distance using this apparatus.

And so here, this is from Stephen Block's group. They actually were able to measure single base pair incorporation into a growing mRNA chain as RNA polymerase walked along a template. So it's not perfect. It doesn't look totally linear. But it's pretty incredible to see two-angstroms resolution. And from this they could learn a great deal about the kinetics of RNA polymerase stepping along a template. And they could even keep it in registers so that they could know which base is being incorporated, which change in distance here represents which base incorporation, so they could study the kinetics of how sequence effects the rate of insertion.

Another use of optical traps which has proved very fruitful for some people is to, you know, you can use them to measure forces that molecules apply on a bead, but you can also directly apply forces on molecules by turning the system around. If you move the center of the trap a little bit, then you can apply a force on a molecule which is, for example, bound in the trap, and see how a change in the force on the molecule will change the bead-to-bead distance, basically get a force extension curve. So here they trapped a protein domain in between two beads with a long DNA handle, which basically is designed to be a relatively rigid linker that just keeps the two beads relatively separated so that the optical traps don't overlay at all. And

they applied force.

And what you can see is that as you increase the force by-- in this case, they moved one of these beads just a little bit to the right. You get a little bit of stretching, which is partially in the DNA and partially in the protein domain itself. But then once you reach a certain threshold force, which in this case is 17 piconewtons, they actually unfold this protein domain. And you get a sudden jump in bead-to-bead distance as this now unfolded polypeptide no longer provides any sort of force pulling it back together.

And then if they release the force and allow the molecule to relax, they actually get a different relaxation curve than their original extension curve, because this represents the relaxation of this unfolded polypeptide rather than the refolding of this domain. And another cool thing that they did in this experiment is they brought it right up to this unfolding force, and then they left it there. And it turns out that at this unfolding force, this protein rapidly unfolds and refolds, because it's a very small protein domain, so the protein folding is extremely quick. So this is the rapid transition between folding and unfolding which is taking place at this transition force. So they could use that to study the kinetics of the protein folding under force for this system.

So today, the paper that you guys read for this recitation is about ClpXP, which is this protein that, as you guys studied in class, it recognizes protein substrates that have been tagged for degradation. And it uses mechanical energy released by ATP binding and hydrolysis to unfold folded protein domains and translocate the unfolded protein substrate through an axial pore into an associated sequestered peptidase, where the protein is rapidly chewed up into small peptides that are then recycled into their constituent amino acids.

So I have another figure, which someone else in my lab made more recently, that illustrates a little bit better the actual unfolding and translocation steps, which were the steps that we studied in this paper. So the general idea is that ClpX will translocate a folded protein domain until it's reached a point where the folded domain is too big to fit through the axial pore. And then it will apply what we call power strokes, which are somehow related to ATP hydrolysis and binding, which yank on the ssrA tag and pull the protein down in a repeated attempt to unfold it.

And you can say that there's some sort of strained protein structure. It turns out that most of these attempts do not actually successfully unfold the protein. But some proportion of these attempts-- maybe ClpX gets a particularly good grip. Maybe there's some sort of transient

thermal destabilization in the protein itself. The protein successfully unfolded, and then it continues to take these small steps, these small power strokes along the unfolded polypeptide, translocating it into the peptidase for degradation.

So in the experiments that Adrian actually did for this system, what he did is he attached ClpXP to a bead using streptavidin and biotin, which is, I believe, an interaction you guys are familiar with. And he attached a multi-domain protein substrate to a DNA linker so that the ssrA tag was at the end of this substrate. And he was able to move the two beads together so that, at some point, the ClpXP bound to the ssrA tag. And then he could record the bead-to-bead distance during the unfolding and translocation of this substrate.

So what he saw is that, when ClpXP was successful in unfolding a protein domain, there was a sudden jump in bead-to-bead distance as the folded-- the whole experiment is taking place under a small amount of force. So as the polypeptide has now unfolded for that small amount of force, it's quickly pulled pretty taut. And then ClpXP translocates that substrate, now unfolded, through its axial pore, decreasing the bead-to-bead distance.

And it continues this translocation until it reaches the next folded protein substrate, at which point it stops. It can't translocate any further. And it begins to attempt to unfold the substrate again. And eventually, after some sort of dwell that we call a pre-unfolding dwell, it will be successful at unfolding the substrate. There will be another jump in bead-to-bead distance, and the process will repeat. Go ahead.

AUDIENCE: So you said that it was held taut. Is there an additional force that it experiences because it's being held taut, or is that accounted for?

REUBEN: It's being held a little bit taut. I shouldn't say "taut." There is a small amount of force on the system. You know, often we record in the range of about five piconewtons, which is just required for the optical trapping apparatus to actually work, but should not have a significant effect on the behavior of the system. So I should have mentioned it back at the kinesin. You can understand the stall force for a molecular motor, which is the force that it takes to restrain it from taking additional steps. And it turns out that the stall force for kinesin is about seven piconewtons. The stall force for ClpXP is about 25 piconewtons. So it's dramatically higher than the small amount of force exerted on it by the trap. So it shouldn't have a significant impact on the results. What it just means is that, once you go from folded to unfolded, since there is nothing holding the unfolded structure in a coil, even a small amount of force will pull it

out a little bit and cause this increase in bead-to-bead distance.

JOANNE STUBBE: So I have a question. When you're developing these methods, like when we used to teach [INAUDIBLE] DNA polymerases, how you put the complex onto anything.

REUBEN: Oh, I have a slide on that next.

JOANNE STUBBE: OK. Because to me, that's the key thing. And the question is, in this particular experiment, how many iterations did they have to go through before they figured out length, attachment, all of that stuff?

REUBEN: So fortunately, other groups have worked out many of these issues for other protein systems, so we were able to adapt those--

JOANNE STUBBE: And you would be able extrapolate--

REUBEN: --relatively easily.

JOANNE STUBBE: from one? Because with the polymerases, you couldn't do that.

REUBEN: Other people had studied these AAA's, particularly helicases, as well as various nucleotide translocases that are actually, in mechanical activity, somewhat similar to ClpXP. So we could basically adapt their systems without significant trial and error, which meant that we could get closer to the biology, or closer to the biophysics quickly, which was very nice. But other groups certainly spent 20 years getting these methods to actually work well.

JOANNE STUBBE: Are there any generalizations that came out of that optimization?

REUBEN: Yes. So there are a couple of generalizations which came out of that optimization. The first is that biotin-streptavidin is a really useful interaction. Basically all modern biophysical techniques use biotin-streptavidin or other basically binary, basically permanent interactions for these sorts of biophysical studies.

JOANNE STUBBE: But biotin-enhanced streptavidin has four binding sites. Are they using a streptavidin with one binding site?

REUBEN: No. You give the-- so when you're actually taking a bead and adding ClpXP to it, what you do is, you add a very small amount of ClpXP. And then once you've added it, you add just straight biotin to fill the rest of the binding sites. Because one thing that you never want is, you never

want two ClpXPs on one bead to engage two substrates on another bead, because that would just totally screw up the data that you're recording. So you make sure that it's very, very sparse labeling of molecules of ClpXP on the surface. You know, I think we probably use femtomolar ClpXP during the actual labeling of the bead process.

Another major takeaway from these experiments is that, for these dual bead experiments, it's very important to have a DNA linker. So the distance here is not really representative, but this DNA linker actually goes over here. Because you really want to maintain bead-to-bead separation on the order of more than a micron, so that your two traps don't overlay at all. Because this leads to much clearer readouts on your position-sensitive display, allowing much better data acquisition.

JOANNE STUBBE: So again, just technically, I'm interested in this--

REUBEN: Oh, sure.

JOANNE STUBBE: So when you have very low concentration of anything, usually you have a lot of problems, because the stuff sticks to everything. So how do you avoid, you know, just inherently, especially proteins. I mean, I've dealt a lot with proteins. The more glue you get, the worse-behaved they are, because they stick to everything, even when you modify the containers with different kinds of polymers and stuff like that. So--

REUBEN: Yeah, so that is a major issue. In this case, it's not a particularly big issue, because you're really-- all you need is one ClpXP on the bead and one molecule of substrate on the other bead. So it doesn't really matter what happened to most of the molecules that you add to your mixture as long as you have one active molecule here and one accessible molecule there. So it's probably true that there are ClpXPs which, at this very low concentration, are bound non-specifically to the surface of the bead, and probably bound non-specifically into the actual cover slip.

But it turns out that those molecules have no signal in this process. So it's not as big of an issue. And there's no concentration dependence that you're trying to measure using these very low concentrations. So even if 99% of the stuff you add binds non-specifically, it's just not that big of a deal. For single-molecule fluorescence it can be a really major deal, and so you have to work really hard to basically [INAUDIBLE] all of your surfaces. So-- no, go ahead.

AUDIENCE: I guess it doesn't matter as much for this system since you have two optical traps you're only

measuring a one-dimensional distance. But for optical traps, do you just measure displacement, or can you know two-dimensional, like which direction it was displaced?

REUBEN: You can tell which direction it's displaced, because actually that little picture-- basically it will deflect the beam of light onto this position-sensitive display.

AUDIENCE: But the voltage is dependent upon both x and y directions when working with that?

REUBEN: I can tell you more about the apparatus later. It's dependent on both x and y, yeah. It's called a quadrupole detector. So it basically has four different diodes, and you look at the relative ratio from these four different detectors. That tells you where it is-- basically a four-pixel camera, I guess.

So the actual substrate, the attachment onto the DNA, it's via a sort of chemical biology which is worth knowing about, which is something called the HaloTag. So it turns out that there is this enzyme which, if you make a couple of mutations to ruin it, it will make a covalent bond with haloalkanes. Basically any haloalkane that you add to your reaction mixture, this enzyme will actually form a covalent bond with the haloalkane in the active site. I forget what the original function of the enzyme is. It's some sort of halogenase or something.

So in this case, we synthesized this long DNA linker such that it had a haloalkane at one of its ends. And then we added this halo domain at the N-terminus of our long substrate such that it would form a covalent bond with the DNA which had a biotin at the other end. And then we added that to a bead. And it turns out that these beads are slightly smaller than these beads. So when we actually start these experiments, we trap a big bead in one beam and a small bead in another beam, and we basically move them very slowly with respect to each other until you eventually get an attachment between a ClpXP and a halo domain. And then you begin the actual recording of the experiment.

So I have a nice animation of this actual unfolding process. So this illustrates the effects of an unfolding under force. So once ClpXP successfully unfolds a domain, you get this big jump in bead-to-bead distance, which decreases during translocation and then increases again during the next unfolding step. Alex?

ALEX: Why do you use a DNA linker?

REUBEN: We use a DNA linker to keep the Y-DNA instead of something else. It's really easy to make a long piece of DNA. You know, this is like a micron long. There are very few proteins that are

well-behaved at that sort of length. DNA is very--

ALEX: But why not work with PEG or something?

REUBEN: --well behaved. Stiffness. PEG is very floppy. DNA is relatively stiff. There's still some sort of floppiness, which you account for using what they call a worm-like chain model, which if you have any questions about that, ask me afterward. But DNA is much stiffer than just a single PEG, and also very strong compared to a single PEG. So again, what you see is unfolding translocation, and then unfolding again at the next substrate. So do any of you have questions about the readout that we get when we're actually looking at acquiring the data for this paper? Does this make sense to you guys? Cool.

So I'm going to get into a couple of the figures of the paper. This paper is very difficult, and has a lot of data in it. But some of the figures I think are relatively easy to understand. So I believe you covered these different mutations of titin, which change the mechanical stability and change the rate of ClpXP unfolding in degradation. So using these mutations, we were actually able to directly investigate the unfolding strength of ClpXP.

So this titin I27 domain, there are a couple of mutations you can make right by the C-terminus of the domain, where you basically switch a valine or something for a proline, which significantly decreases the mechanical stability. And we made these multi-domain substrates with these various mutants of titin inserted. And then we investigated the time of the pre-unfolding dwell before this big bead-to-bead distance for each of these unfoldings events before a domain.

And what we saw is that, for a wild-type titin, the most stable, we saw relatively long pre-unfolding dwells. Whereas for the V15P titin, which has the intermediate mechanical strength, we saw shorter pre-unfolding dwells. And for the V13P titin, we saw the shortest pre-unfolding dwells, which basically are the flat periods in between these jumps. So it turns out that-- so you can quantitate the length of these dwells by basically plotting the number of dwells of a certain length, the frequency of dwells of a certain length or shorter, versus that length on the x-axis.

So this is basically making a cumulative distribution plot of the dwells. And it's pretty easy to show that, if you're looking at a single exponential decay process such as a single kinetic step, then this plot should have an exponential shape where it's fit very well by a single exponential decay process. And so you can look at the rate of that exponential to actually gain a lot of

insight into the rate of that particular kinetic step.

So for the wild type titin, we found that the half-life for this unfolding event, which you can extract from this exponential decay rate, was about 55 seconds. Whereas for this V15P titin, which we knew is less mechanically stable, it was about 17 seconds. And for this V13P titin it was about six seconds. And we saw that these plots are not perfectly fit by a single exponential, but they are fit very well. And this sort of indicates that there is one rate limiting step, one major kinetic step in the actual process of unfolding, which is what we assumed, because we believed that these domains unfolded very cooperatively.

And so basically this suggests that a single successful power stroke by ClpXP is responsible for the unfolding of these domains. So do any of you have any questions about how we measured unfolding in this circumstance?

AUDIENCE: Well, just looking at the fits, it almost looks like it decays slower than a monoexponential decay. Right? Do you have an explanation for why that is like that?

REUBEN: No. There aren't that many data points here. And some of these data points out at these high distances are a little bit hairy. I would say there's quite a bit of error in biochemistry, and my guess is that the quality of these fits is like 99.9. And without dramatically more data, it's not very useful to speculate about what this other exponential might be. So I'm not going to say. I honestly have no idea.

JOANNE STUBBE: But I think that when you're looking at data like that, sometimes it's better to look way down here so you can see the quality of the fit, so you can look at the quality of the fit. It's hard to see it on that graph.

REUBEN: Oh, sorry. Yeah.

JOANNE STUBBE: [INAUDIBLE] would mean a lot of the data points. And then you look at the fit.

REUBEN: Yeah.

JOANNE STUBBE: I think the data looks [INAUDIBLE].

REUBEN: Yeah. Given-- I mean, this data took Adrian like a year to acquire, because you get so few data points per experiment. With the error, this is pretty darn good for this type of setup. It may be that sometimes it takes two very rapid power strokes, and that might add some longer

exponential decay process. But it's very difficult to say.

So we know from the crystal structures, as well as from biochemical studies, that ClpXP has these pore loops along its axial pore, which are actually directly involved in the mechanical activity of the protein. So we believe that there are these RKH loops, which are on the top face of ClpXP. And these loops make this initial interaction with the ssrA tag. And then there are pore-1 loops within the pore, and we believe that the pore-1 loops are the loops that are basically the levers of mechanical force application to a substrate during unfolding and translocation.

So these pore loops, you can imagine them undergoing a conformational change, where they grab onto an unfolded polypeptide and basically are translocated downward in some sort of nucleotide-dependent manner, dragging the substrate with it. And if the substrate is unfolded, then it comes along. If it's folded, then this can be either a successful or an unsuccessful unfolding event.

And then there are also these pore-2 loops, which we think are involved in holding the substrate at the bottom of the pore as the pore-1 loops are reset for the next power stroking event. And we've done a lot of experiments which try to explain how these power stroking events are related to the actual process of ATP hydrolysis to ADP and phosphate, and then phosphate release, and then ADP release. And we have some evidence which suggests that phosphate release after ATP hydrolysis is the step which is most intimately coupled with this particular conformational change that applies to mechanical force. But it's--

JOANNE STUBBE: So how do you look at that kind of measurement?

REUBEN: The way that we've measured this most effectively is to do single-molecule experiments of the sort I'm going to show on the next plot with either excess phosphate or excess vanadate. So the idea is that after phosphate leaves, there is some sort of a rate constant, which describes the time it takes for this conformational change to occur. So if you have phosphate or a phosphate analog that rebinds before that conformational change can occur, then that should increase the dwell time.

JOANNE STUBBE: That's only one phosphate binding site? I mean, you have six-- how many-- you have multiple subunits. And depending on whether the other nucleotide is around, you only have one binding site if you have a huge excess of Pi, or is it--

REUBEN: We think that the important phosphate-leaving step is taking place when there is ADP bound. So in this subunit where you have ADP bound and also phosphate bound, and then that phosphate leaves, we think that that is the step that's related to the actual mechanical motion. But we're not going to claim that that is the step, we're going to say we think that's the step.

And our evidence supports that versus any of the other steps such as ATP binding or ADP release, because you can show that adding a lot of competitor ADP or changing the rate of ATP does not have the effect on this dwell time in between steps that you would expect if those were the steps that had this central mechanical role. But the evidence is not totally compelling. But just imagine that something related to this process of ATP hydrolysis is basically causing this pore loop to be translocated downward.

So what we could do is-- I showed these plots where you see this unfolding and then the slow translocation. You can actually, because these optical traps have such high resolution, you can actually monitor the individual steps that ClpXP takes as it walks along a substrate, which I actually think is pretty awesome. So you can fit these steps to a model to get cleaner quantitation. And then you can quantitate the length of these steps.

And you can see that ClpXP takes steps that are a range between about one nanometer and about four nanometers. And we quantize this. We say that ClpXP can take steps between one and four in multiples of one, because we know from the crystal structures that this particular pore loop translocation downward, it moves about one nanometer.

And so what we've found is that you can look at the order of these step sizes. And you find that ClpXP, it rarely takes-- so the order of steps, which you can see A, it's taking a one-nanometer step, and then a two-nanometer step, and then a one-nanometer step, and then a couple other one-nanometer steps, and then some twos and then some threes. And you can see that for B and for C, as well.

We found that the order of these steps is relatively random. It's very difficult to use one step to predict what the next step is going to be. So from that, we think that there is a significant degree of stochasticity in some aspect of the mechanism of this enzyme, determining possibly which subunits are hydrolyzing ATP to power a power stroke.

But we say that it's not completely stochastic. The events are not all equally likely to occur after each other event. So for example, after a four-nanometer step, we very rarely see a second four-nanometer step or a three-nanometer step. We are much more likely to see a

one-nanometer step or a two-nanometer step. Whereas it's much more likely that after a one-nanometer step, we might see a longer step.

So we have come up with a couple of complicated kinetic models that can explain some of this data. And we're not saying that it's correct. What we're more saying is that this is the sort of thing that you have to think about when really asking deep questions about the mechanical activity of this machine. So a model which explains this behavior where you don't see several long steps in a row is to say that these steps, which are quantized according to this one nanometer, what they actually represent is cascades of steps, possibly a two-nanometer step being two very quick steps in a row, or a four-nanometer step being four very quick steps in a row, which occurred too quickly for our instrument to actually catch them.

And you can say that each step is basically controlled by nucleotide hydrolysis or by phosphate release in a single subunit. And say that possibly a four-nanometer step involves phosphate release in four subunits, sort of one after the next. Maybe they're all contiguous, or maybe they're not. We don't really have any evidence going either way. And so you have boom, boom, boom, boom, in a cascade of four very quick steps. Whereas a two-nanometer step involves hydrolysis in two subunits, and a one-nanometer step involves hydrolysis in just a single subunit.

So the reason that we're attracted to models such as this is that after you have one of these four-nanometer steps, you have basically lost your phosphate, or lost your ATP, or whatever, in four different subunits. And so it takes time for ATP to bind and possibly to be hydrolyzed again in these four different subunits. So it's unlikely that you're going to have the time for four ATPs to bind and be hydrolyzed before your next power stroke. Whereas it only requires binding in one subunit to power a single-nanometer step.

So this model, we don't actually have much evidence that suggests that it has to be sequential orders of subunits going around the ring like that. We've chosen this counterclockwise direction completely arbitrarily. The reason that we think that it's limited to four-nanometer steps rather than five-nanometer or six-nanometer steps is that we know that ClpXP, even at saturating concentrations of ATP, basically never binds more than about four equivalents of ATP per hexamer. And we also know that there are these U subunits, these unloadable subunits, which are actually not competent for ATP binding.

So possibly a four-nanometer step could involve phosphate release from all four of the

subunits that have ATP bound. But then the step stops when this cascade event, this cascade of conformational changes, reaches an unloadable subunit, which has nothing bound. So we are trying to add more to this mechanism, but the behavior of this protein is very complicated.

So I should just say, I didn't actually record any of the data in this paper. I don't work on optical traps at all. If you have any questions about the experiments that I do, which more directly investigate the model I just showed on the last page using single-molecule fluorescence, feel free to ask me anytime. These experiments were done by Adrian, who's a postdoc in the lab, who's probably going to start as a professor at UC San Francisco next year. This is the Sauer lab. It's a great place to Year Up. If any of you are looking for a new Year Up, come say hi. So if you have any other questions about optical traps or ClpXP, feel free to ask me.