PROFESSOR: All right. Why don't we go ahead and get this started here? We have a cornucopia of different silicon materials out in front here in display, and we'll walk through some of them shortly. What I wanted to do right at the beginning of class was to give a little bit of an update on quiz number two. Some of you have probably seen this already and are aware that on Thursday we're expecting a short little decision tree as to how to process your solar cell to obtain the lowest dollars per watt peak.

So this little exercise-- it will last for about a month-- is coincident with our technology section of the class. So remember, we went through the fundamentals. Now we're on the technologies. And then finally, in the cross-cutting themes. So coincident with the technologies portion is designing your own solar cell and optimizing the dollars per watt.

So this will entail actually fabricating a solar cell, which is kind of fun. And Joe will be your guide throughout this process, so you'll be able to actually take a piece of bare silicon and finish up with a device, a rudimentary device, but something to take a picture yourself, post on Facebook, that sort of thing. You design your own solar cell.

So the idea isn't only to optimize for the performance of the solar cell, but we decided to throw in a little curve ball and design for dollars per watt peak. Now this is a little bit of a contrived exercise since we've arbitrarily chosen what dollars are associated with each different process step, but it's not too unlike what you would face in actual industry if you had real data coming off a production line and knew exactly what it cost for each process step.

So instead of having 30 plus components in a more detailed cost model, we've

decided to simplify it to this little diagram right here. So this is a flow chart for the fabrication process of your solar cell. You'll start with a wafer. It has a certain cost associated with it. You'll have some decisions to make concerning light management, whether you want to texture your front surface or whether you want to leave it bare and reflective like this right here.

So whether you want a reflective front service or you want to texture it, there's a certain costs associated with it. So you can probably go to some online resource, like PVCDROM, and use their simulator or the one you've already constructed for homework number two, and calculate what the predicted efficiency boost should be if you texture your front surface. Keep in mind on this very simple solar cell here, we have no anti-reflection coating. So the texturization is pretty much all you've got for light management.

Next, on the emitter, the choice is whether to make a deep emitter or a shallow emitter. The text goes into that in some detail. But your decision is basically if you make a shallow emitter, you have less Auger recombination in that front region. And so your blue response to the device will be better. But you run the risk when you do your contact metalization of firing through that very shallow emitter and shunting your device.

Whereas, if you decide to go for a deep emitter, it stays longer inside of the furnace because of the phosphorus will diffuse deeper inside of the device. You blue response will be poorer, but you'll have less risk of shunting. So it's up to you to use all of the tools that you've assembled so far to make a value-based judgment whether or not it makes sense to go with this or that as your selection choice.

And finally, narrow and wide fingers, this you can probably guess already pertains to series resistance and shading losses. So these are all representative of trade-offs, trade-offs in terms of the technology and trade-offs in terms of cost. And you have all the tools necessary to calculate or estimate what these outputs should be based on what you've learned so far.

And so by Thursday, what we've asked you to do is to make an estimate of what

technology pathway your company is going to pursue. Remember, you want to optimize the dollars per watt. You want to minimize that quantity, which means you want to reduce the number of dollars you invest in your solar cell. But you also want to increase the watt peak that you get out of it. And so at the end of the day, it'll be a performance/cost trade-off in each of these different process steps right here.

And sometimes it won't be entirely obvious which one to choose because so many factors will converge. And so it'll be up to you to make an engineering decision, a professional judgment, as to which path you should pursue. Since it is kind of-- you know, there's a little element of competition in here, so we decided the dollars per watt peak shouldn't be completely neglected at the end and we all get certificates of merit and all feel good about ourselves. We decided it should be worth some part of the grade, but not such a large portion of the grade that everybody's freaking out and saying, oh my gosh, I don't have the right tools to make this decision. I feel like I'm not being graded fairly.

So the portion of dollars per watt is really only going to be affecting 10% of the final grade of quiz number two. And so it will be based on a ranking system where the highest one will be 100% and so forth. But just 10% of your grade. So it's enough to, I would say, create maybe a sting of the pride if you don't happen to hit the highest performance metric, but not enough to sting the actual final grade of your class, which will be one lumped quiz, quiz one, home works, final, and so forth. Right? Any questions about quiz two so far? Yes, Jessica?

AUDIENCE: I completely understand, but there's even a note in number three under the deep emitter. And you guys say, any numbers you should give as far as [INAUDIBLE] or are they responsible for [INAUDIBLE]. It seems like its lacking some numbers. And I understand optimization, but I'm having trouble putting just how much better. And it say it'll be much more effective if you do etching. Well, how much is much more effective?

PROFESSOR: Oh, the etching for the--

AUDIENCE: For the etching, we gave a rough [INAUDIBLE]. So you can look that up.

AUDIENCE: I did look that up. And for the other ones, is there--

AUDIENCE: So that one, you can get a pretty good estimate for.

AUDIENCE: OK. For the other ones, is there going to be a [INAUDIBLE]

AUDIENCE: In terms of shunting your device, it's really hard to predict the shock resistance. But if you do shunt your device, you essentially ruin it. So I would just take that into account. You're not going to get exact answers. But you can do your best to estimate [INAUDIBLE] resistance from the [INAUDIBLE] spacing and your emitter thinness.

PROFESSOR: Yeah. Believe it or not, you might feel like you don't have the tools right now to get quantitative answers, but you do. You have a number of the tools here to get, say, 90% the way there. And in engineering, 90% of the way there is well beyond what you'll actually face in the field. So that's pretty good.

If you have specific questions about what would be a good resource to look up about this, what would be a good resource to look up about that, send an email. And what I'll do, if I receive something in that nature, I'll respond to the class so that everybody has benefit to that information and no one person is particularly advantaged. So it's worth a try.

If it's something that was just covered yesterday in lecture, I might be a little bit more reticent. But if it is something to the effect of, gee, how would the lifetime improve with these different gettering scenarios, sure, absolutely. We can give you a little hand there. But everything else, you should definitely have that information available so far.

This is meant to be a fun exercise, but also one that illustrates the trade-offs involved with designing solar cells. And trade-offs very similar to this are evaluated on a daily basis in industry, or perhaps not quite as often as they should be in industry. But at some point, they were. And the designer of the manufacturing line made those judgment calls. OK.

So again, the pre-analysis, what is due on Thursday is 20% of the grade. The dollars per watt peak metric, at the end of the day, is only 10% of the grade. It's meant to really serve as a stimulus, a little bit of competition, but not meant to really harm you if you happen to not achieve a good value there.

And this is meant to be an educational mission, so the solar cell efficiency analysis at the end is really heavily weighted. We'll be walking through some of the characterization tools in the laboratory so you can determine what exactly went wrong with your devices and quantify them. And that'll be a real chance for you to get a tutorial of how solar cells are not only made-- you'll be there when they're actually fabricated-- but also how they're analyzed and how they're assessed. So it's up to you to really grab this opportunity.

Maybe if you're working in your own devices and want to bring some of them along, you're welcome to do that as well. We won't take up the time when everyone else is in the room, but we might stay longer afterward and help you walk through the analysis as well.

And what we've done, just to resituate ourselves, we talked about the silicon feedstock, right? So we chatted about how you go from quartz in the ground and the carbon-baring feedstock material to the purified, highly purified, silicon feedstock material. This right here is probably on the order of somewhere between 8, 9, or 10 nines pure, very, very pure material, this Siemens-grade polysilicon right here in my hand.

And you've taken a look at this during last class, so you have a sense of what it is up close and personal. Silicon, in fact, has been so well refined that, for a period of time, NIST, the National Institutes of Standards in Technologies, they were thinking about redefining the unit of mass in terms of a silicon boule, essentially a silicon sphere, that would be polished down to about 4 nanometers mean surface roughness with a very low defect density, isotopically pure silicon to serve as a new standard for mass because it could just be purified so well and because their standard reference units were beginning to shift relative to all the others around the

world, the one in Paris relative to the ones that were stored in Washington and Delhi and others around the world.

The values of the mass were shifting as a function of time when they would perform these round-robin. So either the mass in Paris was changing or everybody else was changing. This obviously was unacceptable for an institute that was focused on standards. And so they decided to reformulate the standard for mass.

I'm not quite sure where that project currently stands. So if anybody has further information about the NIST unit of mass, I'd be happy to hear it. But that gives you an idea of how well-purified silicon can be and how well-controlled it can be as well.

During the integrated circuit fabrication, which uses this ultra purity silicon to produce very nice single crystal wafers like this one right here, the investment per gram of silicon can be on the order a few tens or even low hundreds of dollars per gram of silicon and still turn a profit. But in the solar cell, on the other hand, you can invest, at most, a few tens of cents per gram of silicon. This is because the solar cell has to compete against bulk power. That's its competition coming out of the wall right over there. So the solar cell has to be able to be produced much more cheaply. And as a result, typically thinner wafers are used and less expensive starting materials and faster growth methods, resulting in more defect-rich materials.

So one group decided, gee, the embedded cost in the wafer is just so large, it's just so large that we have to make it thinner. And we have to avoid using these ingots, like this one right here, from which these wafers are sawn. So your wafers are sawn out of the ingot like this, like shown. During the process, about 50% of the silicon is lost to sawdust.

And they said, well, let's develop a better way. Let's extrude the wafers directly out of liquid molten silicon and make ribbons of silicon instead. That way, we don't have the sawdust, and we don't have to have this expensive ingot solidification step.

So ribbon growth has been explored since the 1970s at least. And the advantages

is that you have no kerf loss, in other words, no sawdust, due to wire sawing and, hence, more efficient silicon utilization. Immediately out of the gate, if your wafer yields are comparable, you get about a factor of 2 gain because this wafer right here is about 170 microns thick. And the sawdust is around 170 microns as well. So that's about a factor of 2 if you're able to produce a ribbon of silicon directly out of the melt.

The disadvantage is that traditionally there's been lower material quality and, hence, lower performance because of the thermal stresses during growth of a very, very thin foil or thin fin. The thermal stresses can be larger, resulting in plasticity, resulting in dislocations and other defects that can reduce minority carrier lifetime. And traditionally, there has been as well a higher capex. And a third disadvantage, traditionally, in ribbon has been that the form factor or the shape of the wafer has just been different than the ingot material.

Why is that important? Well, if you're trying to displace the dominant design, the wafer, you would do well to make your wafer the same size and shape as the dominant design. Why is that? Well, if you want to make a cell out of it or solar cell device, you'd want to make sure that you can take advantage of the same manufacturing equipment. And that's just a plug-in-and-play, drop-in replacement.

If you require customization of the downstream components on the cell in the module level, you'll wind up having to invest more money in those processes, which might counteract the advantage that you get out of using less silicon. Yes, Ashley?

AUDIENCE: What's capex? Is it--

PROFESSOR: Oh, capex. Capex stands for capital expenditure, capital equipment expenditure. And that relates to the cost of the equipment that is typically-- well, in the business world, typically one undergoes what's called an accelerated depreciation where you amortize the cost of the equipment over five years but then assume that it runs over a longer period, maybe 7, 10 years or so giving you profit back.

So in layman's terms, what this means is capex is the equipment cost, in other

words. And then you just take the cost of the equipment and parse it out. For each wafer you produce, you allocate a portion of equipment cost to that.

So let's take a little walk through history and go back to some of the earliest methods of ribbon growth. So one of the earliest forms of ribbon growth was the so-called edge supported ribbon, also known as string ribbon. And there were developments of this general technology in different places. Ely Sachs, former professor here at MIT, now founder and CTO of 1366 Technologies just up the road in Lexington, developed the string ribbon material here at MIT in the early 1980s, late 1970s.

And the general idea was to use two filaments like so that would be passed through a crucible. And then the silicon would flow in between those two filaments much like soapy water flows between the little circle when you blow bubbles. So a meniscus would form here and then eventually solidify into a solid piece of silicon, and you'd have edge-supported ribbon, otherwise known as string ribbon because you're using the strings to define the edge of the ribbon.

So I have a wafer here, an example of a wafer here, a string ribbon sample. Oh, here it is. It's hiding from me.

So this is an example of one of those materials. Here we go. And like usual, it's good to handle these wafers with some care almost like a photograph. So here's an example of a string ribbon wafer, one particular wafer that was laser cut out of a growing ribbon.

As you can see, this larger ribbon right here-- these can grow up to be a few meters long. They're rather long. You can pick them up if you have gloves on your hands. And they're quite flexible at that length. You could actually even bend them with a radius of curvature of about a couple of meters.

So the reason you wear gloves, obviously, is to prevent your fingers some soiling the wafer. We talked about sodium contamination and other forms of contamination. Silicon is nontoxic, so it won't affect you. It's really you affecting the wafer, much like

putting fingerprints all over a nice, clean photograph.

So there were similar technologies developed by Ted [INAUDIBLE] at NREL out in Colorado. But the general idea is shown right here. Now some of the earliest edge-supported ribbon samples were developed back in 1970s. It really took a while before they were commercialized in full. And that was done through Evergreen Solar, which was founded in 1994 by Jack Hanoka, Rich Chlebowski, and-- oh, goodness-- Mark Farber. So the three of them a co-founded Evergreen Solar.

And they developed the string ribbon growth process shown right over here. Eventually two ribbons face to face, and now four ribbons side by side. So this was called the Gemini because there were two ribbons face to face. And then eventually, the quad process were four ribbons edge to edge.

And you can see the conventional ingot multi-crystalline silicon. Here, the different steps forming the ingot, eventually slicing, and so forth and the string ribbon process here being much simplified in correspondence. So not only was the process simpler, but you'd use about half as much silicon.

And here's Rick Wallace, the inventor and developer of the Gemini process, up there showing one of these longer meter-length ribbons with some flexibility. So the company had a joint venture with REC and Q-Cells, Norwegian and German companies respectively, to form a factory in Germany. REC would supply the silicon feedstock, Evergreen the growth technology here, and Q-Cell some of the cell fabrication expertise. And very recently, Evergreen Solar encountered some financial difficulties-- we'll get into that during the third section of the course when we talk about cross-cutting themes-- and is in the process of filing for bankruptcy.

So this process-- so Sovello is continuing as its own company, but the Evergreen plant here in Massachusetts in Marlborough, about an hour west of here, has effectively shut down. So that was the trajectory of this particular technology through commercialization and ultimately not making it. If you would like, my personal opinion about why Evergreen never quite took off, yes, there are some technical factors, but as well it failed to grow fast enough to keep up with the rest of the

industry and scale with the rest of the industry. And part of that can be traced back to the mid 2000s when silicon was scarce, the inability to source the feedstock material. Yeah?

AUDIENCE: Excuse me. Can you back one slide?

PROFESSOR: Sure. It takes a while. It's a big file. OK.

AUDIENCE: How do you seal the space between the filaments and the bottom of the crucible?

PROFESSOR: Right there, right?

AUDIENCE: Yeah.

PROFESSOR: Since this is your graphite crucible right here and these are your filaments popping up through the graphite, the beauty is you don't have to seal that. The surface tension of silicon is greater than that of water. So if you've ever filled up water to the top of a glass and seen that meniscus that forms, the silicon meniscus would be even higher than that.

AUDIENCE: Oh, that's cool.

PROFESSOR: Yeah. It's pretty nifty.

AUDIENCE: So I'm imagining just like molten metal. You don't want that spilling out the bottom.

PROFESSOR: No.

AUDIENCE: That's really cool. OK. Cool.

PROFESSOR: Yeah.

AUDIENCE: Are the ribbons a single crystal? Or are there grain boundaries in them?

PROFESSOR: Yeah. So let me show you the actual ribbon right here, and you can inspect it first hand. These do indeed have grain boundaries. So what I'll do is I'll place the ribbon inside of here for ease of carrying around. If you'd like to take it out, feel free. They're more where this came from. So in case there was a little accident along the

way, don't feel too bad.

The growth of an ingot is about one to two days, but you get thousands of wafers out. The growth of a wafer itself-- if the growth rate was around, say, let's pick a number somewhere between 2 and 5 centimeters per minute, then it would take-- let's see, with this, you have a 15 centimeter wafer-- it would take somewhere on the order of four minutes to grow wafer. And you'd have a faster growth of single wafers from the ribbon process, of course, lower throughput. The silicon utilization of the wafer growth process was a lot higher than that of the ingot growth.

Some smart people realized along the way that you could grow these ribbons vertically, but you encountered the following problem. During the growth of-- here you go. During the growth of a vertical ribbons, if this was the ribbon growing vertically-- it should be straight. Apologies. There we go. Let's make sure we're good engineers here.

And so this is meant to represent a growing ribbon. This is the liquid, and this is the solid silicon right here. The growth velocity would be in this direction right here.

So you're growing the ribbon out of the melt. This is your melt. This is the ribbon that's growing up. You're looking at the cross section right here, so looking at the ribbon edge on.

So you're pulling it in this direction, so the growth velocity is here. And the direction of latent heat of fusion extraction-- so you have liquid silicon solidifying here. During the solidification process, there's heat released. And that heat has to be conducted up the solid and then radiated outward from the fin, from this thin ribbon.

So the direction of heat extraction is also parallel to the direction of growth. What that means is the growth velocity will be limited by the speed at which you can extract heat up the ribbon and then radiated outward. So there are many ideas tossed around about potentially growing in media that are able to extract heat [INAUDIBLE] transport. You can use your imagination.

But ultimately, growth continues in air, and you're limited to, at most, around 5

centimeters per minute growth velocity because of the extraction of latent heat. If you try to grow faster than that, you'll eventually just pull the solid off of the liquid. It'll dissociate much like pulling an ice cube off of a top of a glass of water. Surface tension won't be able to hold the two together.

So you have here a conundrum. How do you grow faster? If you want to increase the throughput and instead of spending minutes to grow wafer, you'd like to grow a wafer per second, how do you do that?

Well, one group of folks thought about this a bit and said, well, what if we do this? If we take our growth velocity and in some way, shape, or form now our growth velocity is going to be perpendicular to the direction of heat extraction, what would that geometry look like? And they came up with something that looked a bit like this right here, a horizontal growth mechanism.

So you see the [INAUDIBLE] interface is now at an angle. It's almost vertical at this point, a slight angle. And the pull velocity is almost perpendicular to it. So now, you're able, in theory at least, to grow much, much faster.

This was a schematic of the ribbon growth on silicon process. There's also another company called AstroPower that developed silicon film. It was later purchased by General Electric. So these technologies were developed with the intent of pulling very, very fast. And indeed, you can literally extrude the silicon at around 49 meters per second. But the problem about this is that you wind up with very small grains and very poor crystalline quality when you try to grow at the speeds. And so it winds up being a metallurgical problem of how do you ensure the proper grain size when you're growing using these technologies? So there is some work in that regard, but never really took off in commercial production. Yeah?

AUDIENCE: So does pulling at a lower speed with the horizontal ribbon increase your quality by increasing your grain size? Or is it not really--

PROFESSOR: If you're able to control the nucleation and growth process at the very beginning, theoretically, that could be possible.

AUDIENCE: OK.

PROFESSOR: Yeah, question?

AUDIENCE: You had mentioned form factor for these wafers before.

PROFESSOR: Yeah?

AUDIENCE: So is there like a standard form factor for solar cell manufacturing?

PROFESSOR: Yep. So the standard form factor today is akin to this one right here. It's about a 15.6 by 15.6 centimeter squared lateral dimension form factor for the wafer. And I can pass this one around as well.

This right here is what's called a "pseudo-square." You can see the edges are kind of rounded off. And that's because it came from a CZ wafer like this one. It was just chopped out of it. Let me see if these two are coincidence. It would be a-- oh, yeah. Look at that.

So you can see where the solar cell actually came from. So that's the standard diameter of a, say, linear dimension, usually rectilinear shape, a square. And the multi-crystalline silicon ingot material are typically of this size as well. And you can already see that these wafers that I have up here are a bit small. These were the previous generation size. I believe these are 12.5 by 12.5 centimeter squared.

Most laboratory devices that you and your colleagues will manufacture are on the order of 1 by 1 centimeter or smaller because-- well, because of a variety of factors. One is the transparent conducting oxide as we saw in our homework problem. We're limited in how big we can make the device by the sheet resistance of that transparent conducting oxide.

Another problem that we typically run into is just that we're not able to deposit uniformly over a large area. We don't have a deposition equipment for it in our labs. We're there trying to optimize a new material. We don't necessarily worry about making module-sized devices out of it. Yeah, question?

AUDIENCE: Is there a reason why the form factor is different than that used for device manufacturing?

PROFESSOR: Sure.

AUDIENCE: Like [INAUDIBLE] uses circular wafers.

PROFESSOR: Yeah. So if we were to imagine a bunch of circular wafers inside of this module over here, you can imagine the circular wafers side by side. That was how they were done at once upon a time. Obviously you didn't have 8-inch. It was much smaller. Or a 6- or 8-inch. This would be a 6-inch wafer.

But the wafers were a little smaller, but you still have circular wafers and a lot of dead space in between. So as you can see, because of the rounded edges, the packing density is very low. The equivalent would be, say, oranges at a market where they're all stacked on top of another and you have all this dead space in between. And so the idea was to optimize between the cost of the silicon and the cost of the encapsulant materials by shaving away a little bit of the silicon and losing that-- and perhaps recycling it, to be honest-- and the encapsulant materials, where you have this dead space in between the wafers, a small amount of it, where you have glass and encapsulant but no active device underneath.

Another interesting development, as you can see just from the device point of view-- so this would be an Evergreen string ribbon wafer right here, as you can see. And this right here, a larger area device. Does anybody notice a difference besides the shape? In particular, I lead you to the busbars. How many of those thick, vertical lines appear down the wafer?

AUDIENCE: [INAUDIBLE]

PROFESSOR: This has two, and this has three, right?

AUDIENCE: Yeah.

PROFESSOR: So the busbars-- the optimization of these busbars-- this one has three-- that's

really to minimize series resistance. Because now that I have a larger wafer, you have so much current flowing through it, being generated, that the series resistance through those very thin metal wires would end up resulting in large power losses, essentially heat instead of electricity. And so they added the third busbar, even though it increased the shading, to reduce the series resistance losses. So you can see these optimization problems are used quite frequently in solar.

Let me go back one step. There was an interesting question about could we grow single crystals using the vertical ribbon growth. This is a technology. And I don't know if there are actually any of these, many of these samples left in the world. They're quite rare. So I do ask if you want to come up here, take some care with it.

This is a dendritic web sample. So this technology went out of commercial manufacturing, I believe, in 2005. Must have been. Or 2004. It was developed by Westinghouse, which is used to be one of the powerhouses in solar located in Pittsburgh, Pennsylvania.

They had a very active solar activity. It was a kind of a crucible out of which many solar experts then went into diaspora around the United States and set up their own activities elsewhere. And one of the technologies that they developed was a single crystalline ribbon technology like this right here. And if you look very closely, it really is a single crystal.

The growth methods to make this, though, was extremely intricate. It involved, among other things, control of the temperature, of the liquid silicon to within 1/100 of a degree Celsius at melting temperature, which is an extreme feat of engineering. The uptime of these pieces of equipment, meaning the growth time, was around 50%. And the other 50% of the time, the operators were trying to make it work.

So it grew very, very thin material. It wasn't able to scale to the form factors that we see nowadays. The throughput was quite low. The cost was high. And so it didn't quite make it, but from an engineering point of view, it was a marvel in terms of what they were able to accomplish.

So history of crystalline silicon development is riddled with these technologies that didn't quite make it with these materials that were extremely inventive, extremely ingenuitive. But at the end of the day, the dollars per watt peak just couldn't continue to justify their existence. And there were a number of factors that could contribute to making that happen.

So in terms of wafer fabrication in general-- this includes both the wafers out of ingot materials but also ribbons-- where do I personally see this field going? These are some notes. So in terms of cost, the cost per watt peak can be reduced by using cheaper starting materials. That means instead of using this expensive Siemens poly, perhaps an upgraded metallurgical silicon process. Growing or sawing thinner wafers. Growing, for example, on a ribbon technique.

Sawing, maybe making the saws themselves thinner but more robust so that they don't snap as they're pulling through the material at about 5 meters per second in that slurry with the silicon carbide or diamond grit. Very challenging engineering as well.

This second bullet point right there can be encapsulated in a larger team called improved silicon materials utilization. In other words, the grams of silicon that you use to produce a watt peak of a solar cell. So improving that number right there.

Increasing furnace throughput-- that means increasing ingot size, growth, speed, and so forth. There are many people right now trying to grow these ingot right here up to a ton, one metric ton, so 1,000 kilograms. That would mean for the full-sized wafers, you would have something on the order of 6 by 6 bricks. It's pretty large, a pretty large ingot. Maybe even 7 by 7.

And improving the material quality so that you can improve efficiency, efficiency being a huge leverage over the entire cost structure. Because if your solar cell is able to produce more power, that means that you use less encapsulant, and less material, and so forth per unit power produced, and even less labor to install it and less racking and framing materials downstream. The scaling issues, so polysilicon production is currently-- well, this is higher now. It's about 100,000 metric tons per

year. And about half of that-- well, about a quarter of that, now, maybe a third is for the semiconductor industry, about 3/4 for the PV industry.

The slurry and the silicon carbide grit needed for wire sawing is, at some point, going to become an issue. These are huge volumes of waste that need to be transported through the factories. And of course, the silicon loss due to wire sawing and ingot casting, resulting in only 50% of the silicon here in this ingot being used in the actual wafers to make solar cells.

The technology enablers-- using lower-- let's put it this-- lower cost feedstocks. You can't compromise on quality ultimately, so this is a little bit of a false choice right here. Using lower cost feedstocks produced by the upgraded metallurgical route, for example. Producing and handling thinner wafer and growing faster, larger, higher quality ingots.

And there's a lot of innovation to be had in this space right here. I believe the numbers in the last quarter, start-up companies raised on the order of $250 million from venture capital. And that wasn't including a new $50 million deal that was just announced of a company attempting to produce upgraded metallurgic grade silicon through liquid routes, purification. This was just announced this past week, if you go to Greentech Media.

So there's still a lot of active innovation in this area despite the current market conditions. And those of you who are looking for jobs right now, if you're clever, you'll find them here in this space. Any questions so far about wafers? Yes?

AUDIENCE: Does laser cutting cause as much dust?

PROFESSOR: Does laser cutting cause as much dust? So let's walk through that. If we're thinking about the ribbon growing from, say-- from this ribbon right here, I'm going to extract this wafer. So I need to make an incision horizontally right around this point right here. If you look at the total height, the wafer's around 15 centimeters long. And the laser cut itself is something on the order of maybe, oh-- I'm going to guess-- a few tens of microns, maybe 100 microns in that order. And so that the amount of kerf

loss in that regard would be 100 microns over 15 centimeters, so a relatively insignificant fraction.

If you're trying to chop up this using a laser, yes, then you would have significant losses. But since you're growing that ribbon straight out of the melt, the laser cuts themselves are a very small fraction of the total silicon. Yep?

AUDIENCE: Can the sawdust be collected and remelted then?

PROFESSOR: Wonderful question. Can the sawdust be collected and remelted again? There was a lot of work done to try to figure that out. At that point, the sawdust is mixed with this glycol-based slurry, and with the silicon carbide grit, and with fragments of iron coming from the stainless steel wire, and nickel and chromium and other impurities inside of the wire. And so a lot of the work was focused on separation of those different constituents, shall we say.

And when the silicon prices were very high, maybe in 2007, 2008, when the spot prices were $500 a kilogram, there was a large incentive to use every single drop of silicon you had including separation. But in recent years, the incentive to do that has really dropped. And the one company I knew that had a very active slurry recycling program let it go. So there may be companies out there that are looking into it, but I'm not aware of their activities. OK.

Let's hop forward into cells and devices. So now we've talked about the market shares of different technologies, feedstock refining, wafer fabrication, how we make these wonderful different pieces of silicon. Now we're going to talk about going from a wafer into a solar cell device.

So just to situate ourselves, raw material, silicon feedstock, the module in the system over here. In the middle, we have the wafer to the cell. And this is the portion of discussion forthwith.

Cell processing. Let's have a look at this. Again, it's a very different world now in a cell fab line then it was in the crystallisation environment. So in wafer fab, which means wafer fabrication and the section of the company dedicated to producing

wafers and ingots, it was a little bit more dirty. You had forklifts moving these big crucibles around with chunks of silicon in it, operators coming by with garden hoses and washing down furnaces after they're finished.

Here in the cell fab line, it looks almost more like a clean room. Almost, I say, because these folks aren't in full bunny suits. They're usually just with jackets with booties. Sometimes you see them with hair nets as well to protect from hair and other particulate matter from getting inside of the tools.

But by and large, the wafers are brought in. And either in a series of inline processes-- this is a wafer, wafer, wafer, wafer. So there are four wafers across moving through what looks like an etch tank to do the texturization on the wafers. Whereas in wafer fab, it was pretty dirty. In cell fab, it looks pretty clean.

You have a combination of these inline processes like this one shown here. We have wafers on conveyor belts moving through lines. And batch processes, where little robots [? pick in ?] places, line wafers up inside of crucibles or boats, and insert them into furnaces for batch processing. So this is the crystalline silicon cell fabrication. In on one side go bare wafers like this, and out the other side come fully processed solar cell devices.

So the very first step after wafer sawing is the saw damage etch. After the sawing process, you have subsurface damage, something on the order of 5 to 10 microns deep beneath the wafer's surface. And keep in mind these are only about 170 microns thick. So you have subsurface damage that needs to be removed.

And you can take advantage of the subsurface damage by etching it in such a way that you etch along the damage and form texturization. So it's a bit of a two-in-one here. You clean the wafer, you create your texturization, and you remove your saw damage so that when you lift your wafer, the wafer doesn't break because there's some hairline fracture caused by the silicon carbide grit.

After you have your wafer-- so you start with your p-type wafer. And this represents the cross section of the wafer from the backside of the eventual cell to the front side

of the eventual cell, about 170 microns thick. Wide would be something on the order of 15.6 centimeters in a real device. We're just looking at a small section of it here.

So as we walk through the different steps of cell fab, we'll see them evolve over here. The first step after the saw damage etch is to do what's called an emitter diffusion, to create your p-n junction. Straight out of the box, the p-n junction is created after the saw damage etch. And typically, what we do is deposit a lower resistance or more highly doped-- that's why we have the 2 pluses here, that means very highly doped-- emitter right underneath where the eventual contact metalization will go. That's to reduce the contact resistance. That's to create the tunneling junction between the semiconductor and the metal. OK.

So we have the high resistance emitter over here. This is representative of a shallow emitter. You remember in your quiz two you have this decision whether to take a shallow or deep. This architecture, which is used in industry, actually combines the best of both worlds. It has a shallow emitter over most of the solar cell device to improve the blue response, minimize Auger recombination. But it also has a deeper emitter right underneath the contact metalization to prevent shunting and to reduce contact resistance.

AUDIENCE: I assume we have to choose one or the other.

PROFESSOR: You have to choose one or the other unfortunately. To create this-- it's really to create the combination, what's called a selective emitter. It's an emitter because it's the charge separation portion of the device. But it's also selective in the sense that you selectively place these low resistance portions across in a geometric fashion underneath your eventual contact metalization.

You have at the end of this diffusion process what's called a phosphorus silicate glass etch, PSG, Phosphorus Silicate Glass etch. After defusing in the phosphorus in the gaseous form, what you'll do-- or actually, you'll watch it being done, since it's happening inside of a furnace. This phosphorus-based gas will deposit a thin glassy layer on the surface of the sample, which then needs to be etched off or removed before you can do further processing. So that's what the phosphorus silicate glass

etch is about.

Then there's a nitride or a silicon nitride anti-reflection coating that's placed on the front surface. And as we calculated in lecture number two, this silicon nitride coating is only how thick? About?

AUDIENCE: 70 nanometers.

PROFESSOR: 70 nanometers, right? It's really, really thin. But yet, that's enough to create that quarter wave interference effect that leads to a very blue looking solar cell device. So the reason they looked blue is because of that anti-reflection coating.

We are going to omit the ARC coating in our design for quiz number two. It requires silane gas, which we don't have access to down here in the laboratory. We'd have to go to either [? NTL ?] or Harvard CNS to get that deposited. So because we want this to be a hands-on experience, we don't to take the wafers out of your hands, do some magic off to the side, and bring them back and say, oh, here you go. Because your level of ownership in the process just plummets in order of magnitude in the process. We want you to be able to see it every step of the way.

So we omit the anti-reflection coating in our quiz number two. But in commercial production, that's done. And people pay a lot of attention to that step.

And finally, the metalization is deposited on the sample and fired. Now, the metalization, how is it deposited? We'll see in a few slides what the screen printing process looks like. And then, we'll actually do it ourselves. You'll press the button on the tool and deposit your metal on yourself.

But the metalization is typically deposited onto the devices on the front side and on the back. The front side, you have to line up-- in commercial production-- line up with the low resistance portion of the emitter so that you are able to extract the full benefit from the selective emitter design and not shunt your device elsewhere.

And on the back contact, this is typically aluminum. The aluminum, some of the aluminum will indiffuse into the silicon and create a p-plus region in the back side

here, which is a minority carrier blockade layer. It pushes the electrons away from the back junction and toward the emitter. And so it prevents back surface recombination.

So you see every single little step of the solar cell fabrication. A lot of smart people spend a lot of time thinking about, gee, how do I optimize two or three things at once? Question up there?

AUDIENCE: So both the front side and backside metalization is [INAUDIBLE]

PROFESSOR: No the front side metalization in this case-- thank you for that clarification-- the front side metalization in this case would be silver or silver-based paste. And in most commercial production, this silver-based paste includes metal oxides. It could be glassy frit. It could be lead oxide. It could be some combination of elements. That is able to etch through the silicon nitride anti-reflective coating.

This is only 70 nanometers thick, but silicon nitride is a very strong material. It's a ceramic material. So you have to be able to etch through the silicon nitride and make electrical contact with the silicon underneath.

And some of the earliest screen printed metalization cells that got in the range of 15% or 16% efficiency only made electrical contact about 10% of the silicon. But it was enough to have these percolation paths for current to flow up into the metalization. It's a miracle that it works at all. But it's a very effective, cheap manufacturing process that, nevertheless, is still being used in commercial production today, even among some of the highest efficiency cell architectures.

And so for each of these different processing steps, somebody had to sit there and think deeply about optimization of different functions. The [? aluminium ?] on the backside, somebody had to think about, gee, how do I prevent the wafer from bowing, bowing too much, due to coefficient of thermal expansion mismatch between the aluminium and the silicon? Somebody had to think about, how do I create the right eutectic with the silicon-- the aluminum silicon eutectic is around 577 degrees Celsius-- so that you create a good ohmic contact on the backside?

How do I diffuse in a certain amount of the aluminum to create this back surface field to prevent back surface recombination? How do I get the right back surface reflectance of the light coming off of here so that I have multiple optical bounces through my device and so forth? So a lot of optimization goes into making a solar cell device to get the Liebig's law of the minimum, to get each plank Liebig's law as high as you possibly can so you can achieve a high device performance. So hopefully, this walk through now, you can have an appreciation for the difficulty that some of your colleagues face at solar cell fabrication plants.

Finally, as last steps-- I mean, this is a real miniature cross section in the lateral dimension right here. We only have two contact metalization fingers. If you look at this solar cell device right here, we have several dozen, right? If I were to make a vertical cross section through it, you'd see several dozen contact fingers. But this is just meant to be a caricature.

So on the edge here, we have edge isolation. And what this is doing is preventing shunting pathways from going around to the back. So it's preventing the emitter from being able to make electrical contact to the backside of the device. And this is typically done by inserting a trench, a laser-based trench, just-- gosh, it must be on the order of half a millimeter from the edge.

I'm going to pass this around, this solar cell device right here. And if you look very, very carefully, it's literally a few hundred microns from the edge at most. You may be able to see the edge isolation, the trench that is formed by the laser. But it's very difficult to see. So I'll pass this finished device around as well. And feel free to pick it up and look at the backside and the front side.

On the back, you'll see some silver paths in the middle of all that aluminum. And if anybody has ever tried to solder to aluminum, you know exactly why those silver pads are there. It's so you can solder to them and make contact to the back of one device and contact it to the front of the next. And you'll notice that they're aligned, so the back pads are aligned with the front. So I'll pass this around right here. Yes, Ashley?

AUDIENCE: I assume that in order to go in terms of where you want to put the edge isolation, do you want it as far out as possible so you're not losing that edge part. But you also need to make sure you're actually making a full [INAUDIBLE]. There's some optimization--

PROFESSOR: Exactly.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Exactly. If your laser edge isolation machine isn't well-calibrated, you're losing area, active area, of your solar cell, hence your current output is going to be lower. Because you know your solar cell has a certain current density, a certain, say, milliamps per square centimeter. But then if your area, if you're square centimeter is smaller, because you're cutting too far away from the edge, you're throwing away good material. This isn't the trench all the way through. It's just electrical isolation. So essentially, this material over here still exists. It's still hanging on to the device, but it's electrically isolated. This trench here is only about a couple of microns deep. And you're losing area.

This area over here is not contributing to the photocurrent of your device. Any electron making it up into the emitter over here will just stay there and recombine eventually. It won't be able to be pulled out of the device.

A funny, but true story-- there was a company once that I worked with to solve a problem. And they were getting lower efficiencies in their new process. And they couldn't figure out for the life of them why they were getting lower efficiencies. They checked everything, everything, everything, everything.

And it turned out that they were cutting their wafers to a slightly larger size than they were before. Actually, it was a slightly smaller size, because it was a lower efficiency. And their tester had embedded in it a fixed number for the area. It wasn't measuring the area of each wafer independently.

It just had a fixed number for the area of the cell. And so it was dividing the total current output by a bigger area than what was actually there. And so it was

"measuring" a lower efficiency than what actually the cell was outputting. So again, these geometric parameters can come up and bite you if you're so fixated on the electrical performance parameters.

Testing and sorting. So after you create your device, you have this beautiful solar cell. And just from simple electrical engineering and maybe as an extreme example if you're stringing Christmas lights together, you know that if you have one bad apple, it can drag down the performance of the entire string, right, if you're connecting these in series. And so it makes sense to test each of the cells individually and make sure that you sort them together with their like cousins.

So if you have high performance cells, you bin them all together. And you make models of the high performance cells. These will be high performance modules. The bad apples you put together with the bad apples and so forth. And that way you can extract the maximum value out of the product you've created.

You take your good cells, you put them into a higher efficiency module. It looks exactly the same, but its producing more power, so you can sell that module at a higher price than you would, say, a lower power output module. OK. So that's what the test and sort is all about.

Turnkey solar cell fabrication lines, very common since the mid 2000s. There are companies-- Centrotherm, gosh, [INAUDIBLE], Roth & Rau, others that were producing either turnkey equipment or even turnkey lines for the entire fabrication line. Even a local company, Spire, just up the road here in Massachusetts. These typically consisted of wafer inspection systems on the input side.

You don't want to invest any money in a wafer that's ultimately going to break, so you want to be able to inspect your wafers coming in to make sure that they're high enough quality to be worthy of your cell investment. Next, you have wet processing to do the texturization. That's shown right there. Saw damage texturization.

And these are typically inline tools with little ceramic rollers, some pretty nasty acids being use. Silicon is like a rock. And if you want to etch the rock, you need to have

some pretty strong solutions, some very high or very low pH, very basic or very acidic respectively. And most of the time, in multi-crystalline silicon, we use an acidic solution. It textures the wafer independent of grain orientation. For the single crystal materials, we use a basic solution that is isotropic or anisotropic in nature. It creates nice little pyramids.

So here, you see the wafers being drawn over an etch bath. And very small quantities of liquid are used per wafer in this arrangement. You just coat the wafer's surface, and that's about it.

If you were to do it in a batch mode, you need a big bath like a bathtub. And you dunk your wafers inside of it. So that would be the bathtub. This would be the shower equivalent. So more water efficient. In this case, acid efficient.

And then the cells come out on the other side and go into the emitter diffusion process. And these are a series of furnaces. We'll see one such furnace over the course of quiz two when we make our solar cells.

So this is the phosphorus diffusion furnace right here. The wafers are typically loaded into boats and then inserted into furnace where phosphorus containing gas, POCL3, also called "pah-cul," is flown into the chamber. The chlorine components and the oxygen dissociate from the phosphorus, which is then driven into the wafer. The oxygen reacts with the silicon, creates that phosphorus silicate glass on the surface. And the phosphorus is driven into the solar cell creating the p-n junction, creating your device.

And here's an example of Czochralski wafers being loaded into the phosphorus diffusion furnace and then out again, just showing the degree of automation of some of these furnaces. This showing a stack not dissimilar from the one in the laboratory downstairs in building 35 where we'll be doing our phosphorus diffusion.

So the next-- after we have our p-n junction-- the next step would be to create the anti-reflection coating. And this is done by a process called Plasma Enhanced Chemical Vapor Deposition, or PECVD for short. And in the PECVD process, you

flow in silane gas and ammonia.

Silane is silicon with a bunch of hydrogens, four of them. And ammonia is nitrogen with a bunch of hydrogens. And the nitrogen and the silicon react on the wafer's surface and create the silicon nitride coating. The hydrogens, 90% of it, evaporates off. But about 10% of it hangs around. Between 1% and 10% go into the wafer or stay at the interface there and eventually are driven into the wafer passivating bulk defects. So again, a multitude of different things going on at the same time.

Eventually, the visible effect is that you've created your anti-reflection coating. The wafers go in looking shiny and come out looking blue. But what's happening underneath the surface is that some of the hydrogen is going into the wafer. Hydrogen, the first element on the periodic table, very tiny.

And in the PECVD process, where you have a plasma, you have hydrogen ions, which basically means you have a proton without its electron. And that proton is very fast, moving through the lattice. There's lots of space for it to move through the silicon lattice. And it's also very reactive because it doesn't have that electron.

So whenever it finds a defect or a dangling bond, it'll usually lodge itself there, and attach itself, and passivate that defect. And that's what hydrogen passivation is all about during the silicon nitride anti-reflective coating deposition. These are examples of inline processes for doing an anti-reflection coating.

I believe there are a few different variants of this inline process, one of which is a sputtering mechanism to deposit this anti-reflective coating. Of course, then during sputtering process, you have to worry, as well, about hydrogen. Do you have the benefit of hydrogen passivation? Perhaps not as much, so additional engineering is needed.

But the inline process could be potentially faster and higher throughput than the batch process using the PECVD. So again, manufacturing trade-offs.

Next, we have the printing line and screen printing. So this looks very similar to screen printing for a t-shirt. Here is a t-shirt being loaded into a screen printer. And

here's a solar cell being loaded into screen printer. This is a close up of the screen, of what the screen actually looks like.

Here, the screen, which is comprised of this mesh of metal-- here the screen is bare. And so the metal that's deposited on top can go through those holes in the screen and onto the wafer underneath it. And here, there's a coating, a polymer coating of the screen, which prevents the metal from going through the screen at those places. So it shades the solar cell underneath and prevents metals from being deposited there.

And you have fingers and busbars. And those are eventually the thin little fingers that you see right here going sideways and the vertical busbars that you see going vertically right here. Question?

AUDIENCE: Yeah. So when you do these [INAUDIBLE] emitter [INAUDIBLE] where you have some areas of high resistance emitters and others of low resistance--

PROFESSOR: Yeah.

AUDIENCE: Do you use a screen to shield it. Or do you [INAUDIBLE].

PROFESSOR: Great question. So some of the earliest designs for the selective emitter--if we go back all the way up to here. Yeah. To the selective emitter portion. So the earliest designs used photoresist process. But I would say nowadays, there are a few technology options that are much faster, one of which involves a creation of porous silicon on certain regions of the wafer that you want to etch back and create the shallow emitter leaving the deeper region intact.

And that, you could use a form of shading. You could use a wax even for it. There are a variety of technology options for achieving that goal. But in a sense, many of the selective emitter designs involve a deep diffusion first and then a partial etch back. For example, creation of porous silicon and etching that material away. Ashley, you had a question?

AUDIENCE: Oh, yeah. So what does "turnkey" refer to?

PROFESSOR: Turnkey. Excellent question. So turnkey manufacturing line-- what it refers to is that I'm the vendor of the equipment. In one case, I say, here, Ashley. Here's a piece of equipment. It's going to cost you $1 million. And good luck getting it set up and running. I'm out of here. I'll see you.

AUDIENCE: Right.

PROFESSOR: A smarter company might come along and say, I'm going to guarantee an output from my piece of equipment. I'm going to guarantee that you'll be able to make 16.7% solar cells, 16.7% efficiency. I will send my engineers to your factory, and they will help you get the equipment set up and running. And once it's running up to spec, then they'll come back home, and you'll be on your own. And you'll be able to optimize it further.

And so you walk in knowing that you have this guarantee of a performance. Then you can go to your financing agency. You can go to Joe and say, hey, Joe, give me money for my new factory. I have a guarantee that I'm going to hit 16.7% and have a pathway to get to 17.2. My CTO right here thinks-- she's a really small person, and she has a pathway to get to another 0.5% out of it.

And so you can go to your financier and get money more easily than, say, in the first scenario where you're given a piece of equipment and then the person high tails it out of there.

AUDIENCE: Right.

PROFESSOR: So turnkey refers to the idea that you turn the line on. You essentially turn the key, and you're getting high performance out. In reality, it takes a month or two to ramp up to that point.

AUDIENCE: Right.

PROFESSOR: To get high yields and to get high performance. But you have the support of the company there on the ground helping you achieve that. And the turnkey lines were actually one of the real reasons why technology flowed around the planet so quickly.

Because up until about the mid 2000s, high efficiency cell was limited to a few laboratories and a few companies in the know. But once turnkey equipment manufacturers got into to the mix, they started creating these turnkey lines and selling the equipment around the world and the expertise of how to make high efficiency devices, both the architecture and the processing know-how.

And this is how, within in the last 5 to 10 years, you've seen such an explosion of companies around the globe in all sorts of places that traditionally haven't been experts in solar cell manufacturing suddenly knowing how to manufacture solar cells. It flows. The know-how flows through the equipment vendors.

So finally, testing and sorting. This is the last stage of the solar cell manufacturing process. Here, we see a little pick-and-place. That means a little robot that picks up wafers and deposits them. The simplest incarnation is just suction cup. The more fancy ones involve Bernoulli lifters, essentially pressure differentials pulling wafers up.

So you have wafers being loaded onto a conveyor belt, coming off of one conveyor belt onto another one. And they're moving forward. And what you see right here in very low resolution are two probes coming down. This, evidently, is a two busbar cell, not a three busbar cell like this one.

The probes come down and make contact with the busbars. And the probes have multiple contact points, so the series resistance along the busbars is not affecting your measurement. Cell efficiency measurement is always tricky because depending where you put your probes, your measurements are going to change because of the series resistance. So these probes right here are long, and they contain multiple contact points.

And they're essentially touching the busbars. And light flashes onto the device simulating the sun, so simulating AM 1.5 conditions. And an IV curve is measured, is swept. I can't really tell from the photograph or from the movie right here whether the IV curve is being swept at illumination, meaning you're sweeping your voltage when the cell is illuminated, or whether the illumination intensity itself is used to vary

the forward bias condition of the cell. They could be doing it in one of two ways.

But most likely, what they're doing is they're flashing the lights, measuring the IV characteristic of the device, and then sorting the cell based on that performance. It goes into a computer. Efficiency is calculated, just like you did on your homework. And just like that, it's calculated.

And then, as the cell moves down the line, the robot knows, oh, that's the cell that got 16.6. We put it over here. Oh, that next cell got 16.8. We put it over there.

Some additional companies sort their cells based on color because they want to have the aesthetic appearance of homogeneity within the module. They want every cell to be of uniform aesthetic value inside of a module so that you have a nice, uniform color.

AUDIENCE: Is that considered [INAUDIBLE].

PROFESSOR: Whether or not this module right here is considered uniform or different would depend on you, Jessica. You're the customer, and you decide whether this is good enough for you or whether it's not.

AUDIENCE: It's not.

PROFESSOR: It's not? All right. Well, then we have to work harder. So the customer requirements really drive the industry. So some customers are more discerning. Obviously, if this is going to large field installation, we have big barbed wire around it. Who cares as long as the module's producing high performance?

But if it's sitting on the facade of the train station in downtown Freiburg, Germany, where every single person riding the train, entering the station, sees the modules lining the side of Deutsche Bahn's headquarters, you want to make sure that those look nice. So there are differences depending on where they go and where they're installed.

High efficiency cell architectures. So there are a plethora of different architectures

out there. There are some that, for example, put all their contacts on the backside, so there's no shading. And these are interdigitated positive, negative, positive, negative, positive, negative contacts here. So this is called an interdigitated back contact structure. It's used by the company called Sun Power.

And so there's no metalization loss on the front side. All your contacts are on the back. Because lateral carrier diffusion is involved, meaning the carriers have to diffuse laterally, they don't have to diffuse only one dimensionally, you probably can't use PC1D to model this cell. You'll have to use a two-dimensional device simulation like Sentaurus.

If anybody has any two-dimensional device simulation questions, Ashley right here in the front is our resident expert, so you're welcome to ask her.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. And then there are also other device architectures which we'll get to during our thin films discussions. A couple of ancillary topics, barriers to scale.

This is the size of a 1 gigawatt peak plant manufacturing facility for wafers, cells, and modules. This is a palm tree right here. These are roads. So you get a sense of scale. This is located in Singapore. It's a company called REC that has this factory.

These are 18-wheelers right here that are taking the materials out and selling them to customers. So you get a sense of the scale of these facilities. They're rather big. And if you say, OK, this is a gigawatt fab, but we need to be producing on the scale of terawatts, which are three orders of magnitude larger in area than this, how big is that factory going to be? It's about half of the state of Rhode Island.

Granted, it'll be distributed throughout many different regions, but it's a big, big factory. So one of the interesting questions is, can we produce the silicon in a faster way that involves less area? Because area generally relates to capital equipment costs, not always, but quite typically. If you have a larger area because you need more equipment in there, for more equipment, it's a higher cost.

So can the production costs be reduced by a higher throughput growth mechanisms? So instead of using thin film or crystalline technologies that are currently being used today-- apologies for that. Instead, if we used, let's say, a float glass-like process. So these would be extruded pieces of silicon on some bed of--I don't know-- liquid tin would be for float glass, an equivalent for silicon. You could reduce the area by about two orders of magnitude.

And if you envision instead these high speed printers that print out your reports for your exam or class notes, they're spitting out 55 pages per minute on 8 and 1/2 by 11 inch squared sheets. If instead those were 15% solar cells being printed, you could envision an area the size of five football fields instead.

So this starts opening the mind that, wow, our way of manufacturing these solar cells, this discrete process where it's very segregated-- wafer, cell, and module. Wafer manufacturing almost like a commodity. Ingot of aluminum. The cell like a device. The module-- as we'll see in a second-- like an automobile, an assembly process.

If instead we managed to blend these processes together and reduce the barriers, the discrete barriers between these different processes and reinvent the manufacturing process thereof, we stand to make this a lot cheaper, and a lot faster, and a lot smaller to produce.

We might even have our own solar cell manufacturing equipment mounted on our desk. When we need to print a solar cell device or power something, we can produce it right there. So that's kind of the vision of the future where this might be going and why bright minds like yourselves are needed.

We talked about silver. We know there's a limit for how much silver can be used in the front contact metalization. We're using about 10% of it right now. And if you're looking for environmental impact of crystalline silicon technologies, I've included many different sites right here that talk about the environmental impact of solar cell manufacturing since we have mentioned acids. We have mentioned gases like silane. We've mentioned CO2 production when we produce the wafers.

We'll talk about this later on in the class, but in essence, we're looking at around 1/10 or 1/20 the CO2 intensity of coal. So it's not a zero-emission source to produce that module, but it certainly is a lot less than, say, our fossil fuel sources.

This declining US market share has really captured the attention of politicians lately, the fact that the US used to comprise 75% of the PV production market. This is to produce and manufacture the modules. And today, it's on the order 5%. This is a risen concern within many in the DOE and today's DOE and government. Meanwhile, the market is growing substantially.

And so an open question is, what is the future of US market share? If all goes well, we should have a small Greentech Media article published on this topic probably within about a week or so, so keep your eyes open.

And let me briefly jump into module manufacturing. Do we have a question? Oh, we're all set. OK. I'm going to hop into module manufacturing. It'll be the last five minutes. Just to show you how you go from the cell to the module, it's an assembly process, very, very straightforward.

We have coming in here sheets of glass, encapsulate materials. And we'll be able to see this up close and personal and feel the materials when we go visit Fraunhofer CSE in the first week of November. We have a field trip going up there. That'll be a lot of fun.

And the encapsulants are a lot of fun. They're polymers. They're really tough. You can take the Tedlar back skin, this white stuff here in the back of the device, that white skin right there. That's called Tedlar.

As the name would suggest, it comes from DuPont. It's a polymer. Really, really tough. If try to take some in your hand and try to tear it, it's nearly impossible, even for the strongest people here. So it's impermeable, very strong material.

The ethyl vinyl acetate, or EVA, is a polymer that infuses the glass in the front side with the cell. And with the Tedlar in the back, it kind of forms this sticky, mushy

material when you heat it up above 150 degrees C. And it binds everything together in what's called a laminate. So let's walk through that real quick.

To get to the point of a module, we need to take our good apples with our good apples or our bad apples with our bad apples, essentially the like-binned cells, and start stringing them together. That means contacting the front side with the backside of adjacent cells. So the front of one cell is connected to the back of the next. The front of that one is connected to the back to the next, and so forth.

And they're connected in series in a big, long string. And that's done at this tabbing, stringing, and layup table. Typically, this is done by an automated solder system. I just put the cells together, and it wires them for you. But usually, there's a manual inspection process afterward because sometimes the soldering isn't perfect. A human being is typically there fidgeting through, making sure that everything is primo.

Then we have the lamination process, which takes those strings. They're very fragile at this point. They're just solar cells connected with some solder-coated wire, so they're very fragile at that point. And these are then laid up on the top of sheets of the encapsulant materials and the glass and eventually laminated together to form that nice package.

So at the lamination stage, coming out of the lamination, we'd have the glass on one side, the Tedlar in the the back, and the cells in between encapsulated by the ethyl vinyl acetate, the EVA. And we wouldn't have the frame yet around it. And so the put that frame, we would need essentially a large machine that takes those pieces of extruded aluminum and pushes them together around the edges of the laminate, fixing them on there.

And this is the examples of the tabbing and stringing right there. And let's see, OK. So the framing materials right here are typically done by machines in places with high labor costs. And they're done by human beings pushing them together at regions of low labor cost.

And finally, the junction box is deposited at the end. And the junction box, what it does is it collects the power outputs from each of the cells and very conveniently gives you two leads. So there could be some power electronics inside of here that allows the current to flow around this module if this module's under performing, if it's broken, or if it's shaded. There would be a bypass diode inside of the junction box that allows the power to flow around the module and not get sunk into it.

And it also works to collect the power outputs from all the cells and produces two leads, which can then be conveniently plugged into either adjacent modules, which would be strung in series, or into an inverter, which would then take the DC power here and convert it into AC power for your consumption.

And that is how a solar cell is made. So I welcome you to spend a few minutes at the very end to come up and take a close look at some of these materials. Ask some further questions. And on Thursday, we'll start diving into thin film technologies and talk about how those are made as well.