

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: Good. Well, why don't we go ahead and get started. We're going to be discussing photovoltaic efficiency, measurement, and theoretical limits. And there will be plenty of natural breaks over the course of today's presentation for us to have our debate. This is a fun lecture because we start out by talking about how to measure solar cell device efficiency.

Later, we will discuss the theoretical efficiency limits of solar cells. Why do we focus an entire lecture on efficiency? Well first, as we discussed previously, efficiency is a very strong determining factor for cost. The rationale, again, is that if you have low efficiency, you're going to need more commodity materials to make a given watt peak.

That means you'll need a larger area of solar module to make a certain amount of power, which means you'll need more glass, encapsulance, and so forth. So efficiency is a strong lever determining cost of all downstream components except for the area independent factors like the inverter.

Secondly, efficiency is tricky to measure accurately. That is why there are only a few laboratories around the world, a handful, that are certified to measure solar cell efficiencies. These are the efficiencies that could be reported in, say, the efficiency compendiums, an example of which you've just picked up here as one of the readings.

The reality is that we can measure efficiency or get a pretty close value for an efficiency of a device within our own laboratories. But there are a number of possible errors that can creep up and nip us in the heel if we're not careful. And that's why we spend some time in today's lecture discussing those potential pitfalls.

And thirdly, there are new technologies that are being promised right and left to overcome some of the fundamental limits of traditional solar cell devices, like this one right here. And we have to understand what those limits are so that we can design better ways to overcome them. So learning objectives.

Bit of a small font here. But the idea is-- our very first point is to identify the sources of record solar cell efficiencies to understand where one goes to look them up. How do you find the record efficiency of say, a silicon device or a [INAUDIBLE] device? Eventually, we'll talk about measurement of solar cell efficiencies.

And finally, the theoretical or fundamental limits to solar cell efficiencies. So the key concepts for-- learning objective number one, to identify a source for record efficiencies. My go to place is a Progress in Photovoltaics, it's a journal and in the PV field. And every six months, PiP comes out with solar cell efficiency tables led by one of their editors, Martin Green, professor at University of New South Wales in Australia.

The latest addition that I could find was version 38. I believe this was from July or June this summer. And every six months or so, they come out with a new version. And what you'll find inside of that paper-- this is one of the four handouts that you have today. One of the three articles that you have.

What you'll find inside of a typical Martin Green solar cell efficiency table is a listing of-- in table one, a listing of individual cells. In table two of modules. And within the cells, any new record efficiency will be shown in bold. Over the six months preceding the release of the latest version, there were indeed four record efficiencies that had been made.

And that's pretty impressive advance. Note one thing which we'll come back to later. Note the plus minus appearing here after the efficiency number. Is anybody surprised at that number? Ashley, did you expect it to be that big?

AUDIENCE: No, not that big.

PROFESSOR: Not that big, right?

AUDIENCE: Yeah.

But say the record efficiency of crystal silicon device, 25% plus or minus 0.5. Pretty large delta. We'll explain some of the reasons why that error bar is so large. Another thing to keep in mind. Look at gallium arsenide at 28.1%, just achieved by Alta Devices in March. Keep that number, 28.1 plus or minus 20% in your mind, at least the first one, 28.1

We'll march on to the module efficiency tables right here. So now the gallium arsenide module efficiency record for the efficiency tables right here is 21%, or 21.1%. Crystalline silicon has dropped from 25 to 23 and so forth. And this is fairly typical that record module efficiencies are in the order of 2% to 7% lower than record efficiency cells.

Can anybody guess why that might be?

AUDIENCE: [INAUDIBLE].

PROFESSOR: Yeah. You have a mix of different performers. And when you connect them in series and in parallel, you're going to be limited by the lowest voltage or current, respectively. Yes. OK. So certainly there are mismatches between the individual devices inside of a module.

That's where the majority of that comes from for, say, crystalline silicon or discrete monolithic wafers. But how about for some of these thin film devices? They are deposited using these large chemical vapor deposition reactors, for example, or PVD reactors.

And you deposit a uniform thin film over a large area, use lasers to cut little trenches in the films and discretize the devices that way. So how come there are differences between record cell and record module for thin film?

PROFESSOR: Exactly. So if you have inhomogeneities in thickness or in composition, or even in surface quality from region to region in that large area, you're likely to reduce your

performance.

One analogy, since I know many of you are mechanical engineers, one analogy to this is when you're doing tensile tests with ceramics or brittle specimens at room temperature, and you pull on your specimen, and you obtain a certain fracture stress, you then take that smaller piece, pull again. Now the fracture stress is higher. Pull on that smaller piece that broke off. Yet again, another fracture stress is even higher.

In other words, in that large specimen, there was one point that was extremely weak, another point that was sort of weak, and another point that was mildly weak. And as you increase the size of your specimen, the likelihood of having one of these failure points increases. That's an analogy, let's say, to a large area module, as well, if we could have pinholes or other manufacturing defects inside of a large area module that could reduce the performance locally, and everything is interconnected, it tends to drop the performance overall.

So we have the record laboratory efficiencies. This is another reference source. I am not personally aware where Larry Kazmerski publishes this on a regular basis. I know he maintains this table. And if you email him very nicely, he'll email you back with the most updated version. But I'm not aware of any publication outlet where this is regularly appearing.

But nevertheless, it captures the record efficiencies versus time the same way that you, too, could do if you went to the Martin Green records. And I believe we're at, was it, number 39 now? 38. We're at Version 38. And if you went back in time to all of the different tables and tabulated the results versus time, you would get a plot that looked very similar to this for each technology.

So next up we're going to identify the sources of standard solar spectrum. So we know the record efficiencies. We're taking their word for it right now that they did everything right. Now we're going to learn what that everything is and how to do it right, or at least some of the pieces of doing it right.

We're going to identify, first of all, the sources of standard solar spectrum. So in one scenario you could say, well, let me just go outside and measure the solar spectrum. If you go outside today, it's kind of cloudy. The spectrum would look very different than tomorrow. And so it would be very difficult for you to compare the performance of your device against somebody who might be in Germany or somebody who might be in Finland, or somebody who might be in Brazil.

So that's why we come up with the standard reference solar spectrum. And it's shown right here. The ASTM standard reference spectrum is shown on the nice NREL website for air mass zero. That's in the outer reaches of our earth's atmosphere and air mass 1.5, which is assumed to be a standard in, say, temperate climates.

So there have been numerous revisions to the standard. You might want to use the latest one just to make sure that everything is up to proper spec. The latest standard here, essentially they, in very pedantic detail, walk through all of the possible scattering mechanisms in the atmosphere. And this is their justification for the specific solar spectrum that they're measuring. And so you have a nice explanation in much detail.

Let me just show you once again the standard spectrum. You've seen this before. You've worked with it in homework number, I believe it was one or two. And we see AM0, AM1.5 direct, and AM1.5 global, AM1.5 global being the capture of sunlight from a full hemisphere, direct being looking very closely, in a very small solid angle, directly at the sun on a sunny day. Given these atmospheric conditions with a small amount of atmospheric scattering, obtaining an integrated power density, watts per square meter if you integrate over all wavelengths, it's about 90% of the global full hemisphere measurement.

Any questions so far about this?

OK. So that's pretty straightforward and the measurement devices for the global and direct measurement today. Obviously we didn't have those fancy contraptions 60, 70 years ago in quite the same way, with the same materials and the same

design and the same quality of glass. But we had other measurement devices. And you'll see in a few slides how the solar constant has been varying as a function of time. Granted, the solar output also varies, but our ability to measure the solar output varies as well.

So if you look at the evolution of the solar constant versus time, this is the number of watts per meter squared in the outer reaches of our atmosphere, so AM0. And you can see that our values have changed. So if you were to ask, is this because of the change of our measurement capabilities or because of the change of the solar output, the people who study the sun will tell you that the variation in solar output is expected to be rather small. So the likely origin of this large fluctuation, pre-1960, is most likely due to our ability to measure the spectrum accurately.

So next we'll describe how to simulate the solar spectrum in the laboratory, and we'll describe how a solar simulator works. So great. We have an idealized solar spectrum right here. How do we recreate this in a laboratory environment? How do we obtain a light source that follows this profile exactly?

Well, barring the ability to recreate a small fusion source in the laboratory and the ability to introduce exactly the right Fraunhofer lines into our spectrum, we are not going to be able to reproduce that exactly. But we have several techniques that come fairly close.

So the solar simulator. This is a schematic coming off of the Newport website. Many of the solar simulators in the laboratories here at MIT, as you'll see as you walk around, be Newport Oriole or related brands, we have the light source back here. Note the type of the lamp right here. It's a xenon arc lamp. Mirrors, essentially a series of optics to create the right form factor of the light and ultimately work our way toward a planar incoming beam. And the what's called spectral correction filter, which is right in the middle of the optics train. And that's essentially to correct for variances between the emission of the xenon arc lamp and the ASTM solar spectrum as defined right here.

And we have a shutter, as well, to block the light. For example, if your device or cell

or material is photosensitive, you might want to not expose it for long periods of time. There's also-- when I say photosensitive I don't mean that the device will stop working under light, but that the performance will change. There are materials, for example, amorphous silicon that we discussed, where the performance does change as a function of illumination, cumulative illumination intensity.

So another interesting thing to note here is that we have quasi-planar light coming in at the end of this optics train. But if we're going to be measuring, for example, concentrating solar cell apparatus, it's much more important to have a higher degree of planarity of the incoming light, and then other optics would be used. One might envision, for instance, increasing the distance between the light source and the actual sample, or the optical path length between the collimating mirror and the sample.

So now that we have light approaching our sample right here, so light is incoming on the sample, we have three things, broadly, that we have to worry about. We have to worry about uniformity of the light, the uniformity from small region to region of the illuminated area. We need to worry about spectral fidelity. That means, how closely does the spectrum of our lamp match the ASTM standard solar spectrum? And thirdly, temporal stability. That means if I turn on the lamp this morning and want to take in a measurement immediately, is it stable yet?

OK. Let me wait a half hour for things to-- for example, the thermal loads inside of the system to reach equilibrium with environment. Now I'm going to measure it. What if I come back in an hour and a half? Will I still get the same result? What if I come back tomorrow or next month? Temporal stability is another major concern for solar simulators.

On the right-hand side is just an example of the radiance versus wavelength of a given light source versus time. And you can see there's a new lamp after a certain working period of 1,200 hours.

So we have non-ideal matches, several examples of light sources that don't quite get it right or have several spikes in the output spectrum. So in all cases, here the

AM1 direct spectrum is shown in this dash dot line. And different light sources are shown either solid or dashed lines. So not great.

And then finally, we reach our xenon arc lamp with air mass filters. And the filters are to suppress certain peaks and certain general portions of the spectrum so that we have an approximation of our ASTM standard reference spectrum.

And you notice that it's not perfect. You notice that there are spikes in the output of the xenon arc lamp, as you might expect from the physics involved in the light source. This is of some concern if your device is particularly sensitive to a region, a spectral region where those might be present. So it's not a bad idea to take a measurement of your light source and actually understand how it compares, how it matches up against the ASTM standard spectrum.

In terms of standards, or ranking different types of solar simulators, there are three classes according to the IEC standards. There are, as well, other common standards. There's a standard used in Japan, and the ASTM standards as well.

But let's focus on the IEC 904-9. These are the requirements for solar simulators measuring crystalline silicon single junction devices. So it's a very specific standard. And in a few slides we'll explain what the potential differences are when you're measuring other types of solar cell materials.

We have the spectral match, or spectral fidelity. We have the non-uniformity and the temporal instability in this case, since we're defining a relatively small parameter.

Class A solar simulators have relatively tight specs. But you'll still notice here that the non-uniformity of the irradiance plus or minus 2%, the temporal instability plus or minus 2%. This is where you start to see some of those error bars on the record efficiency measurements, right?

So if a laboratory has a very good handle on its reference solar simulator, it will be able to calculate these effects and estimate what their impact is on the actual solar cell efficiency measure. Note that the temporal instability for the Japanese standard is a little bit more stringent than the IEC test. Minor detail, but depending on who

your collaborators are, where they are in the world, they might be using a different standard than you. Just keep that in mind.

The solar simulator downstairs in Building 35, in the laboratory that several of you have seen yesterday when you did the phosphorus diffusions-- I poked my head in and saw everybody there. So about a dozen of you might have walked past the solar simulator in the laboratory. That's a large area, as in the illuminated area at the working plane is around 20 by 20 centimeter squared. And it is a Class AAB solar simulator. So I believe that would be spectral match, non-uniformity, and then temporal stability.

So you'll typically see solar simulators rated in this way, AAA, Triple A, or Class B solar simulator, or AAB and so forth.

So again, note the significant figures. That's where that comes from. Pretty straightforward.

Next we're going to describe how to accurately measure and report cell efficiency and some common pitfalls to avoid when actually measuring the cells. So this is really directed toward people who are doing active research right now in the field of PV. For those who aren't, enjoy. And we'll come back as soon as this little section is over to the debate. And finally some topics of general interest.

So this is just a small subset of things to keep in mind when you're measuring an actual solar cell. This is by no means a comprehensive list. There are, indeed, people who spend their entire lives optimizing and perfecting the art of measuring solar cells. So the very first thing that you might want to do is have a reference solar cell encapsulated, and then mailed to you after a measurement is performed at a certification laboratory. So this is an example of a very small crystalline silicon solar cell device inside of an encapsulated frame. And the encapsulation is meant to prevent any degradation to the solar cell, as well as damage that might incur during accidental use.

And this cell is called a reference cell. It was tested at NREL. The current voltage,

the short circuit current, open circuit voltage, and fill factor of that device is well-known and reported, and is essentially sent with that device back to our laboratory. Now whenever we want to make a new measurement on any cell that we want to measure inside the laboratory, we'll measure the reference cell first just to make sure that our solar simulator is well-behaved.

When I say well-behaved, what could happen? Well, one of the many things that could happen is that the lamp intensity decreases. That's probably the most common thing that can happen. You'll just have an overall reduction in the output of your lamp, in which case you'll notice a reduction in what cell parameter? Current, voltage, or fill factor?

Current, right? So the light intensity decreases. You'll have a reduction primarily in your current, logarithmic reduction in voltage. So you have your standard calibrated reference cell. This is definitely something that each laboratory that is going to be serious about measuring efficiency should have. This is also an apparatus that doesn't leave our laboratory except for class purposes. It's something we treat very carefully, since I think the cost ran in the few hundreds or thousands of dollars.

Avoid extraneous-- basically, avoid light from coming in from outside of your solar simulator. It is often shown, or the solar simulators are often shown in this manner right here, where you have everything out in the open. That's to show you what's going on inside. When you actually take the measurement, typically you have a small black curtain that's light tight around your apparatus, or maybe even a box.

Ensure 25 degrees C measurement conditions. Remember that the open circuit voltage can change, depending on the band gap. The larger the band gap, the smaller this effect. The smaller the band gap, the larger this effect. The VOC can change as the temperature changes. And so it's important to have a good handle of your temperature. That means that you might, for instance, have active heating and cooling on your chuck. And you'll also, at the minimum, be measuring the temperature on top of the chuck, where the actual solar cell will sit, not far away removed through some layers of insulation away from the chuck.

Even if your temperature is not precisely at 25 degrees C, if you know the temperature dependence of your solar cell, you can correct for it or account for it to pull it back to 25 degrees C.

Next, choose your probe locations judiciously to avoid series resistance losses. Let's think that through for a second. So if we have our solar cell device right here and we're going to be measuring its cell efficiency, if I put one little probe right here in the corner, then the current that's being generated over here has to travel a very large distance through a lot of metal to reach that point.

But if I have probes that essentially will consist of many different individual metal points, and they come down on either bus bar like shown right here-- where you see the green right here, and off of the green are many little probe tips, individual probe tips that will make contact with the bus bar, so you have essentially one coming down right here, another coming down right here-- your series resistance losses will be much less.

But one thing to keep in mind is that when you do have these probes sticking on top of your solar cell device, they will scatter some of the light as well. So there are best practices in terms of what color they should be and how tall they should be, as well.

So choosing the probe location is a large step toward achieving high efficiencies. I don't under emphasize that point. It really does make a huge difference where you put your probe tips, especially for cells with high series resistance, which can be several of the new materials that are being developed.

Account for spectral mismatch between calibration cell and your cell. Let me drive that point home. So I have a crystal and silicon calibration cell right here. I know its spectral response. I know at what wavelengths it responds the strongest. And so it will detect any mismatch between the ASTM standard and the actual light source there where it responds most strongly.

Now let's imagine that I'm designing a new organic material that responds really well. Pick something, the infrared. And so now silicon will stop responding, let's say,

at around 1,100 nanometers. So anything beyond here, silicon won't be able to detect. But let's say my device is very sensitive in that spectral region. My silicon calibration cell says go ahead, take your measurement, everything's fine. It matches your ASTM standard.

But then when I stick my cell in there, all of a sudden I'm getting a super high efficiency. I just made an organic device, and my efficiency's 11%, which would be a world record. I'm ecstatic. I'm really happy until I do the quantum efficiency of both the standard calibration cell and my new cell, and I realize, wait a second. They're nowhere near each other in terms of their responsivity, what region of the solar spectrum they can respond well at.

And the reason I'm bringing that up is, again, to emphasize we have, say, for example, an amorphous silicon device, a gallium arsenide device over here, a SiGs device showing the different regions of the solar spectrum, the solar spectrum shown in this dark orange in the background right here, peaking at around 550.

We can see that the different materials are responding to different regions of the solar spectrum. And crystalline silicon would be in between the amorphous silicon-- actually, it would be a little further over, be starting up at around between 1,200 and 1,100 nanometers.

So we can see how different materials are more sensitive to different regions of the solar spectrum. And there's even a standard test method for determining the spectral mismatch parameter between your device and the reference cell, an ASTM standard for it.

And the typical way to account for this is to measure the spectral irradiance as a function of wavelength of your light source, measure the quantum efficiency, meaning the responsivity versus wavelength of your device, and then measure your calibration cell. And using that math, you can really begin to normalize for these extraneous effects.

Why do I bring that up? I bring that up because there was an example, several

examples, of folks in the literature-- I've taken off their names, so protecting the innocent here-- folks in the literature who report very high efficiency devices, or at the time was a near-record efficiency device. What they showed were the QE curves of the devices, and then the IV curves of their devices.

And you guys, in your homeworks, calculated the short circuit current, which is shown here-- essentially the intercept with the y-axis-- you calculated the short circuit current from the QE. So you know how to perform that calculation.

Well, guess what? Folks at NREL also know how to do that calculation, and a whole lot more. So the folks at NREL did that and said, well, wait a second. When I do the integration of your QE, I'm not getting these short circuit currents over here. I suspect what happened was, in your solar simulator, you were using a silicon reference cell. But your cells were more sensitive to a different region of the solar spectrum-- in this particular case they were more sensitive to the shorter wavelengths-- and you didn't have a properly calibrated solar simulator. So you're over reporting your current outputs. And of course that's very embarrassing for a group.

In this particular case there were merits on both sides. There was a rebuttal to the rebuttal. So it's not a simple black and white case for this particular story, although the logic does fall more strongly on one side. So this is to say, avoid this sort of controversy. Perform your measurements properly, don't over report your efficiencies, and when in doubt, you can always ship your cells to NREL or [INAUDIBLE] or another certified testing center and get a certified cell efficiency. Then you can place your IV curve inside of your publication, and in the little corner over here it'll have the figure of NREL or [INAUDIBLE] and the properly certified information.

Now that can be rather complicated if you're, for example, growing organic materials that degrade quickly. Somehow you have to transport it over there without it degrading. And that's where arranging in advance the transfer of the materials, and perhaps even looking into what sort of transfer chamber you're going to

encapsulate your device in could be of interest.

But at the very least, please, please, please remember spectral response mismatch when you're doing efficiency measurements. Any questions so far?

AUDIENCE: So then does NREL also make non-silicon reference cells?

PROFESSOR: So NREL actually doesn't make these cells. They're local companies that live right around NREL that are manufacturing these and putting them together. NREL will test those cells. Since it takes a few months, typically, to get turnaround on a cell efficiency measurement-- unless you're fast tracked in because your device is degrading or you've arranged in advance and kind of put your place in line-- because it takes a few months turnaround, there's actually a premium on inventory.

And so these companies will manufacture the devices, send them in, get them tested and calibrated, or get the calibrated measurements performed at NREL, and then bring them back to the company and put them on the shelf until you put pick up the phone and call them and say, I'd like a certified cell.

You can actually make your own, too, if you follow a set of standard protocols that you can receive from the folks at NREL. I believe Keith Emery might be a good point contact at first. There are a very specific set of protocols that you should follow if you want to make your own in the lab.

In other words, you have to design the contacts a certain way, the encapsule in a certain way, make sure that the materials comprising the remainder are black so that they don't reflect the light back in, little details that are only gathered by experience.

If you follow all of those parameters, and the folks over there, Keith Emery or Paul [INAUDIBLE] look at it, and they inspect it and say, yeah, yeah, looks good, then you can get that tested and serve as a calibration cell as well. So you can circumvent some of the cost associated with buying a certified cell from a company. But there is a premium to them. They don't tend to be particularly cheap.

So yeah, if you have another material and you'd like to have it calibrated, you can make your own calibration reference standard as long as it doesn't degrade.

Question.

AUDIENCE: Putting the multiple contact in yours while you're measuring the efficiency. Is that very representative of the devices that actually work in the field?

PROFESSOR: So there are a lot of things-- let me spend a minute waxing poetic about the discrepancies between these cell efficiency measurements and what actual cells experience in the field.

We understand the logic behind cell efficiency measurements. We understand we have to have a universal way of comparing a cell in Japan, in Boston, in Freiburg, Germany. So we understand that we need some standard method of cell measurement. This cell measurement proves very useful when you measure the output in terms of peak watts, because then you can multiply by the number of peak hours of sunlight per day and estimate the energy output as a function of location on the earth. So it has its uses.

Now as to its drawbacks. We're measuring at 25 degrees Celsius. These cells are typically operating around 60. We know that there's a voltage drop, primarily a voltage drop with increasing temperature. That's point one.

Point number two, the contacting scheme. Typically on these devices in a module, you'll have soldered contacts in the front, and so they make contact with the entire bus bar, but really, depending on your soldering machine, only in a few locations where there's really a good electrical contact, and on the back in three locations.

So yes, there are differences in how the cell is contacted in real life. Could be one of the reasons for discrepancies between module and cell efficiency, probably one of the minor ones compared to homogeneities.

Why don't we take a quick little pause right here before we dive into describing the efficiency limitations of a typical cell and get some cool demos at the very end. What we're going to do is have ourselves a quick debate. So I'd like to call forward at the

front of the room the representatives of the two teams, one of which is going to debate in favor of the development of novel materials for solar cells. Let me give you an example of one material that's attracted quite a bit of attention, which is pyrite.

I can pass this around as folks are coming up to the front. Here's an example of an iron sulfide-based mineral which is purported to have a very high degree of manufacturability because of the large resource abundance, and also the large refining capacity for the respective elemental constituents. There is a gap between the performance of current pyrite cells and their theoretical record efficiencies, quite substantial, the record efficiency being up closer to 20%, the actual efficiency being down at around two.

That is not unlike, for example, tin sulfide, or other related materials. I'd say copper zinc tin sulfide is probably the most advanced at around 10%, but still about half of its theoretical limit. And so the big question that we're going to debate is, does it make sense to invest a lot of funds to come up with these earth abundant alternatives for our existing solar cell materials, our cadtel, our copper indium gallium selenide, cognizant that the supply of some of these heavier elements, tellurium and indium, is limited. And we may not have enough of these elements in the earth's crust to scale up to the terawatts level.

And so we'll hear two points of view, one in favor of development of new materials, and one against. And just to situate ourselves in a position or a location where these debates actually do happen, you can imagine yourself, for example, in the Office of Science and Technology Policy, which reports directly to Barack Obama. These are about 50 PhDs who are all in an office under the direction of a director, thinking deeply about some of the scientific challenges that our nation faces and the proper scientific response, coordinating amongst many agencies, including the DOD, DOE, NSF, and so forth.

And so you can picture yourselves in a debate, in a lively discussion, all friendly, but really with the potential to influence national policy. Do we direct resource funds to

develop these novel, earth abundant alternatives that we might need in 10 years' time? Or do we focus and allocate resources elsewhere?

So I'll welcome the two participants up to the front.

So folks want some insight into where the Office of Science and Technology Policy actually decided to go. A gentleman who took a version of this course in Berkeley in 2003, I think it was, his name is Cyrus Wadia. Graduated with his PhD from Berkeley and actually went off to join Office of Science and Technology Policy. He took his class project for the PB course, which was analyzing alternative materials and published that paper that you read by Cyrus Wadia on resource abundance that was published in 2009, for that work earned himself a TR 35 reward from *Tech Review*, and then joined the Obama administration's OSTP, and has been developing the Materials Genome Project within OSTP.

The first solicitation for proposals was issued, I believe, a week and a half ago. So that's funneled through the NSF. But it's a larger effort to develop some of these materials involving NSFD, DOD, DOE, and so forth. So it actually did come to fruition through OSTP.

Budgeting is always the big question, though, because that, of course, gets done through the committees in Congress and ultimately reconciled between the House and the Senate, and has to make it through OMB and finally to the individual directorates. So that's how things actually get done.

But it takes the vision of OSTP, sometimes, to drive these larger projects forward. So we'll wish him the best.

We're going to describe the efficiency limitations of a typical solar cell now. And what I'll do is I'll pass around some of these books just to, again, situate ourselves. Did I grab Jenny Nelson's on the way out? Thought I did. Hm. OK. Well. Peter Wurfel and-- oh, there it is. Yes, absolutely.

So there's two books that are entitled *Physics of Solar Cells*. You can see it's a very popular topic. And one is *The Photovoltaic Handbook*. All refer to some aspect of

efficiency limits. And as you look through your sheets, you'll see at the bottom, typically the different pages of each of the books are listed here.

So the very first thing that we should consider before we get into PV technology, we just have some solar device looking at the sun. And the sun is radiating at it at 6,000 Kelvin, 5,800 Kelvin. And it, the solar contraption on the surface the earth, is radiating back at the sun with black body radiation at 300 Kelvin. So much lower power from Stefan Boltzmann's Law we know exactly the amount of power being emitted by that device.

So the two are radiating at each other, and they're in equilibrium. And that results in the very first, we'll call it blackbody efficiency, or maximum solar heat engine efficiency, which would be around 86%. So that's our first fundamental cutoff. Now that 86% is averaged over all wavelengths. There tends to be a wavelength dependence to this as well.

So the next step was to say, OK, well, we know that we can't do better than the theoretical thermodynamic limit of one object looking at another. But what is the actual limit of a solar cell? And this was a question that the first developers of the solar cell asked themselves.

So Prince was one of the three of the team that in 1954 published on the crystal silicon solar cell device, and very quickly followed it up with another article here in JP focused on the theoretical efficiency limit. And a curve was proposed, looking much like this, with two data points for germanium and silicon, theoretical limits, that is, and power density. So not a conversion efficiency per se, but a power density. And of course, assuming a certain input power density, one can calculate an efficiency from there.

And along came-- well, at this time, H-J Queisser had just moved over from Germany to Palo Alto, to The Apricot Barn, and working with Shockley, the esteemed Shockley at the time, to develop a detailed balance model for describing how a solar cell performs, or what its ultimate theoretical efficiency limit could be.

Now what a detailed balance model does is just basically accounting, accounting for all of the photons coming in and out of the device. Through the photons, the electron hole pairs are generated. This is assuming a very high quality material that does not have any form of recombination other than radiative recombination.

So we talked about the different methods of limiting lifetime. Radiative recombination is one of the methods of limiting performance that occurs in high quality materials. If you have a poor quality material, you'll have non-radiative recombination, say, Shockley-Read-Hall recombination.

But in the detailed balance model, only radiative recombination was assumed. That way you can count number of photons coming in and number of photons going out of your device. Furthermore, they assumed that the mobility of carriers was infinite inside of their material. A few laughs over here coming from the folks who are working on organic materials. But it's true. They assumed that the mobility was infinite, such that the separation of the quasi Fermi energies throughout the entire device was equal.

So if you had a limited mobility, if you had a certain higher density of carriers in the front than toward the back, you'd have a difference in the chemical potential through the thickness of your device. They assumed infinite mobility.

So it was a very simple, yet elegantly insightful model that, upon first submission, was rejected. So the first time they submitted this model for publication it was rejected outright. And it took them another several months of edits, probably about a year, and they resubmitted it. And then it became a sleeper paper. It wasn't cited that much. If you go to the Web of Science, for example, and look at this manuscript, you'll see that the number of citations in the early years was rather limited.

Nowadays there isn't a talk about the fundamental efficiency of a solar cell without mentioning the Shockley-Queisser efficiency limit. That's coming directly from this paper right here. And that serves as a motivational story for you. If your paper is rejected, just remember that now you're along with some several esteemed

individuals who have set precedence in the field of photovoltaics and are Nobel worthy.

So that's my personal opinion. If you fall into that category, don't feel discouraged. Regroup and find a way to make your paper better. You have a copy of this manuscript right here, by the way, in your handouts for today.

And interestingly, through their detailed balance model, they obtained an ultimate efficiency versus wavelength, or energy in this case, curve that looks very similar to the curve that you derived in your earlier homeworks, where you just assumed two loss mechanisms. One was non-absorption of light and the other was thermalization of carriers.

So the detailed balance model, the way they reached this point, is different, fundamentally different. The physics that was assumed is basically listed out right here. But the end result is fairly similar to the rough back of the envelope calculation that you performed in an earlier homework.

So again, just to go up, they assumed that the photons with energies greater than the band gap are absorbed, create one electron hole pair. They assume thermalization. So these were your two assumptions that you did in the homework.

Then they do a few additional things. They assume that radiative losses occur within the material, that not every carrier is collected-- that's important-- so you have radiative losses in your device, so that reduces the performance. And then they assume further that there is a thermodynamic loss in their device, in other words, that you don't extract the full band gap of energy, but that there's a thermodynamic loss as you go from band to band, which would be the band gap energy, to the difference or separation of the quasi Fermi energies, which is this $\Delta\mu$ here, which is the change in the chemical potential from the front side to the back side of the device.

So there are some additional loss terms that were included. For a full description of the detailed balance limit, and all the math, I definitely encourage you go to this

website. As well, Peter Würfel's text does a wonderful job of describing this. And of course, you have your paper here. You can read through the original paper yourself, or perhaps suggest it at an upcoming journal meeting for your group.

Here is the calculation of the detailed balance limit for AM0 and AM1.5, essentially as a function of band gap, coming from the PVCDROM. And you can see how silicon and gallium arsenide are pretty close to the theoretical maximum in terms of the theoretical maximum efficiency for a single band gap semiconductor material. So this curve right here, again, is just representing one band gap material, a single band gap material.

So let's take it from here and venture forward into some more realistic performance reduction effects. We can, for instance, take into account recombination mechanisms that aren't only-- for example, that aren't only radiative. We can take Auger recombination. We can take Shockley-Read-Hall recombination into account.

We can also take what's called photon recycling into account. So what is photon recycling? Photon recycling is when you have a radiative recombination event, and that photon gets trapped within the material. It's not allowed to escape. But it gets trapped because, for example, off the top, there's an index of refraction mismatch. Or the angle at which it tries to exit is too oblique, and so you have total internal reflection within your device. And you eventually have reabsorption.

So that's called photon recycling because you have a radiative recombination event. It emits a photon inside of your solar cell device, and that photon bounces around a few more times until it's reabsorbed, generating another electron hole pair. And that is, essentially, the major boost in some of these ultra thin, high performance solar cell materials such as the Alta Devices record efficiency 28% gallium arsenide cell has the ability to do this photon recycling.

Now how do you modify the detailed balance limit to account for finite mobility? I suspect this corner of the room is going to want to hear this.

There is a beautiful piece of work done by [INAUDIBLE]. The PhD thesis is even

more insightful than the manuscript in terms of actually breaking down each individual component. A great deal of modeling went into this, including photon recycling, including some things that are very difficult to model. And the effect of finite mobility was calculated on a modified detailed balance limit model.

And you can see number one, that curve that's starting here at the 20. That is for a mobility that is close to the optimal. If you drop by an order of magnitude, or two orders of magnitude, relative to the maximum mobility inside of a material, your performance begins to degrade considerably.

And now you can understand why silicon, or crystalline silicon, which has, say, hole mobilities somewhere in the range of a few hundreds of centimeters squared per volt second. So keep that number in your mind, on the order of hundreds of centimeters squared per volt hole mobility.

Now you compare to amorphous silicon, which has something in the range of 10 to the minus 3 to 10 to the minus 1 centimeters per volt second. And you begin to see why those materials with poor mobilities are really impacted in terms of their performance. There is a method to calculate the impact of limited mobility on device performance. And if that intrigues you, I would definitely refer you to this manuscript.

One very simple way to think about it is, at least for a band conductor, you would have a certain diffusion length. At least a first order would be the diffusivity times the lifetime and the diffusivity would be related to the mobility inside of your material. So your mobility is factoring into your diffusion length, and the diffusion length is affecting how many charge carriers are collected. So that's a simple way to think about the problem.

A similar related approach that kind of pulls all this together in a graphical form is at the very last page-- some of the very last pages of Peter Würfel's text. And this is going through the efficiency calculation that we just saw, starting from a certain amount of light that's not absorbed. The light that is absorbed-- so that's efficiency loss mechanism number one-- for a thin device, let's say light trapping isn't perfect. We're at around 74% efficiency of light trapping, of absorption of light.

Let's say now that we have a thermalization event that results in about a 33% loss, so we're down to 67% efficiency for the thermalization of charge carriers, just a step. So now our combined efficiency is going to be the multiplicative product of those two. So just the thermalization of charge carriers results in a 33% drop of an efficiency.

Now there's a delta between, say, the band to band and the actual chemical potential inside of our material once we consider the ensemble of carriers, not just those free carriers, but the ensemble of carriers, which is going to dictate the ultimate potential. So we have some carriers that are excited and others that aren't.

These thermodynamic losses results in another 36% drop. And finally, fill factor losses, which represent the solar cell and practical operation. When we have series resistance and so forth, another 11%. So if we multiply these four numbers together, we drop down to 28.

And the beauty of doing breaking things out like this-- and what Peter Würfel does very nicely is dives into each of these in great detail and explains to you exactly how those numbers are derived. The beauty of doing something like this is you can pick the lowest number, say, 64 and 67, and say, I want to work on those. I want to make my PhD thesis about those parameters because that has the biggest impact. And that's what you can do with an analysis like this.

Now you can take this further. You can say, well, these are only four of the many parameters that impact performance. I would like to look at many, many more. And that's what one has done for crystalline silicon, which is arguably the most researched and, from the point of view of physical understanding, advanced solar cell technology. We can see a variety of performance loss mechanisms that have been taken into account inside of the crystal silicon devices.

This was an invited talk by Dick Swanson, the founder of Sun Power, former professor at Stanford, who pulled this together, a very nice presentation. He speaks with authority because Sun Power produces the highest efficiency crystalline silicon

solar cell device commercially.

And interestingly, right here, this is the thermodynamic limit and the calculated limit efficiency versus time. So we start from Prince. That was the very first one. And we have Shockley-Queisser up there, and so forth. So the theoretical limit of solar cell performance has also changed with time. And their calculations vary. So it's important to be able to understand this as well as you go into it.

And here is the actual best laboratory performance. So if we had, for example, a 25% efficient cell over there, it would already be higher than some of the earlier efficiency calculations for crystalline silicon.

What Dick Swanson thinks is the practical limit is right there. That's the Swanson prediction, that we won't get much above 26% for crystalline silicon.

So loss mechanisms visualized right here, and several good readings on efficiency limits.

What I'm going to do is pause here. There is still some material in your text, and some really cool demos that we're going to have to wait until next week to see.