OLIVIER DE WECK: So let's get to work. The V-Model, well, believe it or not, but we've arrived today at the bottom of it, the pointy part of the V. We're arriving at the bottom of it, and the bottom of it is basically you driving down the design in all its glory and all its details. And we call this design definition, and then I want to talk about multidisciplinary optimization, including what happens at the CDR, the critical design review.

Yeah, go ahead.

AUDIENCE: All right, I do have a quick question about A4 based to this. Do you want a CAD drawing or a hand drawing?

OLIVIER DE WECK: I'm fine with a nice dimension hand drawing, but we're gonna talk about CAD next week. So I do think it would be good to have a CAD drawing, but you know CAD drawings can come in very different levels of detail, right? So I think it would be good to have an idea, if the PDR package-- most of the PDR packages that you see that have been submitted in prior years at CanSat, they do have a CAD drawing in there for their sort of suggested configuration. And you know, it is useful because you get to see whether the packaging concept really works. You know, the sketching, you can get away with a lot just sketching. So my answer would tend to be yes, but if absolutely, you don't feel like you need it or you don't want to do it, then a nice hand drawing with dimensions will do as well.

OK, so let's talk about the bottom of the V. The other thing I'm gonna cover today is this idea of multidisciplinary design optimization, and so where does that fit in? So you remember we talked about concept synthesis, concept generation, screening, and selection, and then at the PDR stage, you know, we picked one architecture. So today's topic is really how do you get this architecture, which is shown as this square here to an optimal? And I put optimal in quotes here because we're gonna argue what that means, but you've refined the design, such that you can go present it at the CDR.

So you take this box here, which is your concept, and you open the box-- you can see this on the right-- and you really dive into the details of the various subsystems and components and how they're connected. And what you're trying to do is fine tune your design. So this is shown here schematically. Let's say there's a design variable X1 and you want to find where is the

optimal point? This point here seems to be optimal to minimize metric three and that's the design that you're gonna propose. So today's topic really is about fine tuning the design and coming up with a CDR level design, that's what we want to talk about.

And so what I'll cover is different topics. So first of all, the NASA view of this. It's called the design definition process. Then I'll talk about MDO, multidisciplinary design optimization, the use of so-called CDFs, concurrent design facilities, and then finally, what happens at CDR.

So the general idea of the design solution definition process is that you start with something like this on the left. This is a product breakdown structure, essentially your system architecture, your system decomposition. And on the right side, you then translate this to a detailed design, OK? So for a launch vehicle, we have, you know, the flight segment, this is what-- actually it's gonna fly and it has a payload element, the spacecraft bus, and then the launch and the launch accommodation. So that's the general idea is you have high-level breakdown, but then you need to design every piece. You need to have chosen your materials. You need to have chosen the dimensions. You need to understand the interfaces. And for every part, you also need to decide what? What's a big decision that you need to make when you go to the detailed design, beside, you know, material, dimensions, software? What is always a big decision to make when you do a detailed design?

Go ahead.

AUDIENCE:	Maybe the interfaces between the two.
OLIVIER DE WECK:	Yep, interfaces, all that. What I'm trying to get at is COTS. Who knows this acronym, COTS? Let's see if the, EPFL, anybody know COTS? What does it mean? I think we can't hear you right now.
AUDIENCE:	Commercial off the shelf.
OLIVIER DE WECK:	Commercial off the shelf. So you have to decide in the detailed design are you going to pick components from a catalog, or are you going to custom design these components, right? And so one of the things at the CDR, that I always ask the question-- so you present your CDR design. What percentage, what fraction of your parts at the lowest level are COTS and what fraction did you custom design?

Mike, you've done a lot of these unmanned aerial vehicles and so forth. Can you talk a little bit about, you know, what's the typical fraction of COTS, and how do you make that decision

when you design, you know, a flying thing?

AUDIENCE: It really depends on the scale and how robust. The problem with COTS in the unmanned world is that you can easily go to hobby grade parts. And so if you want to get into more of the commercial, kind of more robust, almost more like defense kind of, like, robustness and you know, durability, you really have to start going to custom designed parts. Because you know, if you can do a whole design space of, you know, all the COTS parts, but you know, they're always just kind of, you now, Chinese-made electric motors and batteries, and none of them that will perform to, like, commercial spec, but it just depends, really, what you need for the vehicle.

OLIVIER DE WECK: I see. So what you're saying is there's a gap. You can get a lot of COTS stuff at the hobby shop level, right? So they're pretty affordable but low reliability. Or you have really high end, very expensive, you know, low production--

AUDIENCE: Like traditional aerospace grade, yeah.

OLIVIER DE WECK: But there's not much in the middle. Is that what you're saying?

AUDIENCE: Yeah, for sure. Yeah.

OLIVIER DE WECK: OK, good. So you know, that's the reality of detailed design that you're gonna face. And surprisingly, in kind of the formal system engineering methodology, there's not a lot of guidance on this. There's not a lot of, you know, guidance to say what's the optimal fraction of commercial off the shelf parts? This is really something you have to work out yourself, and depending on, you know, whether you're in aerospace or medical devices or automotive or making toothpaste, there's-- actually, yesterday there was also a big case study on diapers. Diapers is a $3 billion business, believe it or not. And they were going for-- their solution A was to go for a COTS solution, you know, a lot of vendors. And then they said, no. It's just gonna give us, sort of, average, mediocre performance. We want something much better. So they went for a much higher custom design solution. Took them took them a while to do it, but it was worth it in the end.

So this is a big decision you have to make. So the point here being, at the end, you have a very detailed design. You have blueprints. You have something you can actually build.

So you know, this is a list of all the things you need to do. You define the solution space, develop your alternatives. All the things we've already talked about at the conceptual level, but you do this at every-- you do it at the subsystem level, and even at the component level. What are the alternate designs? Cost performance and schedule. Select your design solution, drive it down to the lowest level of detail, and then identify the enabling products, which would be like support equipment and things like this.

So the best illustration I can give you why this matters is these two pictures I'm gonna show you. The top picture that you see here, this is a picture that was drawn in the '70s to sell the space shuttle program to US Congress. This was the vision, right? This is the orbiter, and you can see it's sitting there in this pristine hangar, right, very clean, shiny floors, maybe four or five people. The payload bay is open. There's, like, one or two carts of ground support equipment. You just load up your new satellite, close the doors, and off you go for your next mission. That's a reusable vehicle.

That's what we wanted, and then what we actually got is this. You can't even see the orbiter, right? But I promise you, it's in there. It's hidden in that scaffolding. So this was the-- it's retired now, but this was the orbiter processing facility at the NASA Kennedy Space Center. And the reality was that to turn around an orbiter between flights was almost like rebuilding it. And the two subsystems that required the most person hours of work were the main engine-- the main engine basically had to be disassembled and inspected and so forth-- and what was the other system? What's that?

AUDIENCE: Heat shield.

OLIVIER DE WECK: The heat shield, the TPS, thermal protection system-- if you've ever seen-- who's seen the shuttle close up? OK, about half, half here. At EPFL, who's had a chance to see the space shuttle, the orbiter, close up? Anybody? Go ahead.

AUDIENCE: One or two.

OLIVIER DE WECK: One or two. So those of you that have seen it, what did you notice about the heat shield? Did you notice anything about the protection, the thermal protection system? Who's seen it?

AUDIENCE: We do not hear very well here.

OLIVIER DE WECK: You can't hear me?

So the point I want to make here is that the thermal protection system, every tile is different. Every tile has a different shape, a different curvature, and there were good thermodynamic reasons, aerothermodynamic reasons for that, but in terms of actually designing, manufacturing, and turning around the vehicle, it was a nightmare, right? So that it turned out to be not a very reusable vehicle in the end, and why? Because a lot of the detailed design decisions ended up being critical, in terms of not really giving us what we wanted, in terms of macro level properties of the vehicle. So detailed design really matters.

Did you get that point, at EPFL? Is it clear?

AUDIENCE: It is.

OLIVIER DE WECK: OK, thank you.

So here's some-- I'm not going to go through this in detail, but here's some flow diagrams, process flow diagrams, from the System Engineering Handbook that essentially describe, you know, how this works. So your inputs are your baseline logical decomposition model, which is essentially your system architecture, right? And then your baseline derived, your detailed technical requirements. You go through all these processes, through all these steps, and then you end up with a whole bunch of outputs, which are more detailed requirements. The specifications, do you remember, we teased apart requirements and specifications? Requirements are what the system should do. Requirement specifications are how the system's actually been designed, right, blueprints. So that's specifications at the subsystem level. And then you get other things, like a verification, validation plan, and then logistics and operating procedures. So very, very detailed work.

The other thing that's important is all the design considerations that go into making detailed design decisions. So here's a fairly long list of them, and you can group them into performance, availability. This is more related to reliability. So you could have a system that has high nominal performance but is poor reliability, right, so it's not available very much. That's not good. We want high performance and high availability. That gives you your technical effectiveness.

And then how do you actually operate the system? So operations, maintenance, logistics, so that's process efficiency. Together, that gives you system effectiveness. And then, of course,

you have the total cost of operating the system-- lifecycle cost, total ownership cost. And then system effectiveness and cost together, give you what's called here affordable operational effectiveness.

So I would argue, in terms of the space shuttle program, it had high system performance. It had OK system availability. They had-- definitely there were some reliability issues, but it was manageable, but it had very poor process efficiency-- operations, maintenance, logistics very, very difficult. So system effectiveness wasn't that high, and as a result, the lifecycle cost was also very high.

There's a very famous paper, where after the shuttle program was ended, a lifecycle cost analysis was done, and it turned out that one launch cost about $1.5 billion. So during the program, people would say, oh, one launch is maybe $500 million, you know, the variable cost, and the rest is fixed cost. But once the program is done, it doesn't matter anymore whether it was fixed or variable cost. All you know is what was the total cost of the program, how many launches did we actually do, and you just take the ratio. And it was $1.5 billion per launch of the shuttle. And so I would say that's way off from what the original target was.

Yeah, Veronica?

AUDIENCE: Does that number account for crew and crew training?

OLIVIER DE WECK: Good question. I think so. I think it basically-- whatever was appropriated by Congress for the shuttle program, and I think that would include crew operations, as well, yeah. So I'm happy to point you to that paper, actually.

Yes?

AUDIENCE: It also accounted for R&D costs as well, right?

OLIVIER DE WECK: Yes, that's correct. It includes everything, all in. R&D costs-- R&D took about 10 years and then operations was about 30 years.

OK, so you see-- and of course, within these are budget limits. So part of the reason why the shuttle didn't have that high process efficiency was because the R&D budget was capped pretty tightly. So rather than optimizing, you know, for maintainability and turnaround, they said, we need to have the system performance. We need to have reasonable system availability, and then the rest, basically, is we're gonna deal with it later.

So what effectively happened is they capped R&D cost, and as a result, the lifecycle cost ballooned. I would argue if they had spent, you know, two or three more years optimizing, you know, maintenance and turn around, really making it reusable, the lifecycle costs could have been much better, but that's not how it worked.

So you have a lot of these competing requirements. And so I want to just ask you this-- this is the first concept question today, OK? So imagine you're-- this is a very detailed design. This is just a bracket here, OK? And this bracket was designed so it has, like, a tip load. Think of this as a cantilever beam. So it's attached on the right side or left side here. You have a load, and then this is your displacement under a given load.

And these three brackets all have the same mass. They use the same amount of material, OK? So you can think of this as you get from the very simple design to the more intricate design, you have more design freedom. So you get better performance. You see the displacement goes down, but you also have more complexity. So it's more difficult to optimize. It's more difficult to manufacture and so forth. And so we can actually show these three choices here on a manufacturing cost, per unit cost, versus the displacement, structural displacement under load chart.

So our simple design, our one bar design, has a displacement of 2.5 millimeters under load, and it costs about $2 to make. The two-bar design has much better performance, like 0.75 or something like that, costs about $3 to make. And then this 17-bar intricate design has 0.6 millimeters, so the best structural performance, but its cost is about $8 per piece.

So my question to you is which of these three designs would you select and why? So if you could go to that link, enter your answer, and then we'll look at the distribution.

So we have 39 responses. That's good. Ah, interesting. OK. So one of you said you're gonna go for the super-duper 17-bar design. Most of you, 74%, said the two-bar design. And one of you said one bar. And then eight of you, 20%, said you're not sure. I guess you need more information, right?

So who picked the one-bar design, the simple design? OK, so please explain why you picked that.

AUDIENCE: Well, for my application. I'm planning on mass producing them, and it's for something that doesn't need the degree of displacement that the other ones are doing, and so I'm really only

seeking to minimize manufacturing cost, and the one bar is the easiest to manufacture and also gives me room down the road to iterate if it's not working.

OLIVIER DE WECK: OK, so So what would be an example, industry or application?

AUDIENCE: Like a children's toy, perhaps, where you don't need very strict requirements on displacement, something that you're mass producing.

OLIVIER DE WECK: Something mass produced, you know, for sort of lower requirements. OK, I think that's a valid answer. Most of you said the two-bar design. So who said two bars? Let's see at EPFL, who picked the two-bar design?

AUDIENCE: Sure. It gives you the biggest biggest bang for your buck. For not that much more cost, you get a much more sturdy design and that's compared for the 17-bar one. Really, it's still pretty sturdy, but you're paying much less, so it seems optimal.

OLIVIER DE WECK: Yeah, so if you think about this in a Pareto sense, I mean, they're all Pareto-optimal, but this one is at the knee of the curve, right? So you're getting, you said bang for your buck or value. If you did, like, a displacement or one over displacement, I guess, per manufacturing cost, this one would do pretty well. So that's also a very sensible answer. Who picked the 17-bar design, the really intricate one? Who picked that?

Oh, I think I might have picked that when I tried out the form. [LAUGHS] I think so. I can't remember now. It's been already a while ago. So if, indeed, I was the one, so how would I justify this? OK, so here this bracket will go-- this will go on the Mars 2020 Rover. And the Rover will see a lot of different thermal gradients. You know, it has to be very precise because it holds scientific instruments that need to be, you know, pointing very accurately. Even though the mass here is identical, but I think we could shrink this. You know, we get the best performance structurally for the mass constraint that we were allocated. And cost, you know, this is all taxpayer dollars, so this is a government program. So we can spend the money. Of course, I don't mean that really, but it is a high-performance application. Therefore, even though the benefit is only small, that small benefit is worth it when you look at the whole mission. So we're we're going to pay for that, and we're going to go for this high-performance design.

Do you see how that works? So the right answer-- and I think a lot of you said, you know,

you're not sure, and it's because-- who said you're not sure? Nathan.

AUDIENCE: I said not sure because you don't know all the specifications or what the tolerances are or any other information that would lead you to make that decision.

OLIVIER DE WECK: Right, so you know, what are the requirements, right? We don't know the requirements here. Are the requirements shall or should requirements? If it's a shall requirement, what it means is there would be some hard limit here, right? The displacement shall be less than one millimeter. If that's a hard requirement, then this solution falls out, does not satisfy. And if it's a should requirement, then it's a goal that you're gonna trade off.

The point I want to make here, this lecture is about detailed design. This is detailed design, right? And you're gonna make hundreds or thousands of detailed design decisions just like this bracket and together, they're gonna give you, you know, the performance cost of your whole system. And so every detailed design decision you make needs to be informed by the requirements but also the trade-offs between these properties. Does that make sense? OK.

So and that leads us-- and so you want to make sure you optimize your system. You want to make sure you get-- and I liked, Katya, the way you said it-- get the bang for your buck, right? You want to make sure you don't have a lot of inefficiencies in your design. So that's what MDO is about.

So let me go to MDO here, what it is. So there's a technical committee in AIAA, the American Institute for Aeronautics and Astronautics, that was founded about 20 years ago. So it's defined as an evolving methodology, a body of methods, techniques, algorithms, and related application practices for the design of engineering systems, coupled by physical phenomena and involving many interacting subsystems and parts. So the main things you need to do MDO are a mathematical model of the system, analysis capability, and in some cases, approximation concepts. So what that means is not-- if you put all this together, it could be very, very slow and computationally heavy. So you might approximate some of these things, and I'll explain a couple examples. Sensitivity analysis, optimization processes, and a human interface-- those are the main ingredients of MDO.

Who's seen this chart before? OK, about three or four of you. Who's seen this chart before at EPFL, this particular graphic? Anybody seen this before?

AUDIENCE: Apparently, no one.

OLIVIER DE WECK: OK, so let me explain. This is a cartoon, OK, and I've experienced this myself working at McDonnell Douglas in the '90s. So each of these airplanes shown here is what the airplane would look like if that particular discipline or group could call all the shots, OK? So let's look at a couple of these.

So the controls group is all about controllability. So you see the flight control surfaces are very big. Now, the problem is the linkages, you know, these control linkages are on the outside. That's really bad for aerodynamics, but the controls people say, I don't care about drag. I need to make sure my control surfaces are deflecting appropriately. Here's the hydraulics group, right? It's all about the hydraulic pump, the filters, and getting hydraulic pressure to all the different subsystems. You know, the big engine here, the power plant group, the engine is the most important thing.

The loft group-- loft means-- lofting is surfaces. You're dealing with the shape of surfaces. So let's keep this simple. So we have a little, you know, balsa airplane with little straight surfaces, very easy for loft. Production is kind of similar. The aerodynamics group wants a very nice high aspect ratio wing, very thin, very low drag. Problem is, there's no room for passengers or cargo here.

You see this, the stress group? They have a nice wing here in the shape of an I-beam. I-beams are great, you know, this is good if you're a civil engineer but not so much-- none of these things would work on their own. So what we really want is something like this, where all these considerations are blended and optimized so we get a high-performance system not just a high-performance subsystem. Said another way, a high-performance system is not the same as just a sum of individually optimized subsystems. That's, in a nutshell, what MDO is about.

So one of the readings for today, post readings, is this paper, 6A. And this is based on a big workshop we had in Germany about five years ago. That's actually longer. It's probably seven years ago by now, but the paper came out five years ago called, "MDO, Assessment and Direction for Advancement." And this is a pretty international group that put this together.

And what we do in the paper, we describe the history of multidisciplinary optimization, where it is today, where it's heading, and some of the key milestones. So this chart is essentially a kind of history of MDO. I'm just gonna point out a few things. The roots of MDO are in structural optimization. So you saw the bracket with all those-- so there's actually a field called structural optimization or topology optimization, and it's very mature now. So you can buy, in some of the

CAD packages, there's some specialized tools for optimizing structures. And when you look at it, it looks like Swiss cheese with big holes, an Emmentaler, big holes. You don't need a lot of structure that doesn't do work for you. So a lot of these structures look like webs, and the reason is because the load path has been highly optimized.

For example, if you look at the A380-- I don't have a picture right here-- but you can look at an A380 wing spars and different structures. They have more holes than structure because they've been highly optimized. So that's the roots of MDO, starting in the 1960s.

Then, I would say, in the '80s, mid-'80s, we started having more complex decomposition techniques, and I will describe one of these to you today. Optimization algorithms were started to be built into mainstream programs like Excel, MATLAB, Mathematica. And then more recently, you know, in the 2000s, we had commercialization of multilevel algorithms, and the one that I'll describe to you is called BLISS.

So let's talk through a very specific example, and then one of those techniques. So here's the example. You're designing a wing, and in designing a wing, the two main considerations are structural and aerodynamic. So here's a simplified wing. It has a wing area, it has a sweep angle, it has an aspect ratio, and so forth. And there's loads acting on the wing as you're flying because of the pressure differential between the upper and lower surface. So there's some net force, P, acting on this wing. And as a result of that, the wing will warp. So you will have some twist angle in the wing, and that's just due to the structural deformation.

So that twist angle and the displacements, in turn, will change the drag that the wing experiences. So the lift and drag of the wing will change as a function of twist angle. But as the lift and drag change, the loads, of course, then change as well, and it feeds back into the structure. And in this case, what we're looking for is maximum range. We want to maximize the range of the airplane.

And this is a sort of simplified version of the range equation, the so-called Breguet range equation. And you can see that the structure influences the range in two ways-- directly by weight, the actual weight of the structure enters here twice in the range equation, and then indirectly through its stiffness, right? The softer it is, the more displacements, the more you affect the drag.

So it's not so clear then what is an optimal wing? Is it the lightest wing possible? But the lightest wing possible will have a larger twist and may have more drag, which will reduce your

range, right? So should you optimize the wing for lightness or for this minimum displacement or some combination or mix of the two? And the answer, of course, is it's a mix of the two. There will be a sweet spot somewhere, you know, between a very stiff and very heavy wing and a very, very flexible, lightweight wing. There will be some optimal mix of the two, and that's the key here.

So let me discuss with you briefly-- this is just one of the methods. There are several methods for doing multidisciplinary design optimization, but this is one of them called BLISS, Bi-Level Integrated Systems Synthesis. And the general idea here in this method is that we have a lower level optimization of each subsystem but it's coupled through a system level optimization. And what you want is, you know, an optimal design at the system level, not just optimal subsystems.

So the context here is we're gonna design a high speed, a supersonic business jet. It's kind of shown in this gray picture. And we have four major subsystems we're gonna consider-- propulsion, aerodynamics, structures, and then the performance, which, in this case, is range, OK? And what's important here in decomposing the problem is the x's and the y's. Now, the x's are the design variables. Do you remember, we talk about the design variables, the knobs you can turn as a designer?

And so these here are-- the upper ones, we call them local design variables, you know, like the throttle setting, the tail sweep, the wing moment arm, the thickness of the wings, the taper ratio. So the idea of a local design variable is that design variable only is used in one of the subsystems. It's local to that subsystem. So you basically can choose that in each of the subsystems, and it doesn't really affect directly any of the other subsystems-- indirectly, yes, but not directly.

And then we have the y's here. The y's are the outputs from one subsystem to another. So drag, lift, NZ is your maximum load factor, your range, your specific fuel consumption, your wing twist, and so forth. So those are outputs from a subsystem or discipline that either matter at the system level or are required as inputs by another subsystem.

And then finally, the last group here, these are known as the shared design variables, so the wing aspect ratio, the altitude, mach number, wing sweep. Those are basically required by more than one subsystem, so two, three, or four, in this case. Then they are considered as shared variables, and they would be considered system level variables because they matter

for more than one subsystem.

So what this particular method does is essentially decompose the problem into the various subsystems, so in this case, again, aerodynamics, propulsion, structures, and range. And these subsystems are sending this information to each other and are all tied together at the highest level through the shared variables, the shared design variables that are shared by at least two subsystems. So again, local variables are only unique to a specific subsystem, and then for the wise, these outputs, internal outputs, there's a distinction between the Y stars, which are coupling variables that are input to a particular subsystem and then the Y hats are coupling variables that are output from a particular subsystem.

So let's zoom in now on just one of the subsystems, let's say, aerodynamics, and see how this works. So here we're zooming in on one subsystem, and that's the aerodynamics. And the idea that you can actually optimize just the aerodynamic subsystem by itself. So what are the inputs into it? We have aspect ratio, wing twist, and W. So what is W? And this is key to this particular method-- the W are the weights, the weights that you give to the outputs. So what you're optimizing here is some objective function, f, shown here below, which is W1, Y1, plus W2, Y2 plus W3, Y3. So you're either maximizing or minimizing that. And what is this here? It's lift, drag, and the lift to drag ratio.

So what's really cool about this method is that what it essentially is telling the subsystem is how to properly weight these different outputs in order to optimize the subsystem, taking into account the system level objective. So the weights are not pre-defined, but the weights themselves are a design variable.

And so you can see the formulation on the right side. So given Q, which is your shared variables, the Y stars, which are the inputs from the other subsystems, and the weights, minimize your objective function by varying your local design variables subject to a bunch of constraints, OK?

And so what is one of the consequences of being able to formulate the problem this way? So this could be a very high fidelity, aerodynamics model, computational fluid dynamics, or it could be what? Or it could be a lower fidelity code, right? As long as the inputs and the outputs are the same, you can now plug in, you know, different fidelity levels of analysis, you know something simple or something based on experimental results, you know, panel method, very high fidelity CFD. This becomes very modular.

The other thing you can do is you can pre-compute a lot of these responses, right? So for a given set of inputs, you can pre-compute over a whole range of input variables what the response of that subsystem will be, and you can store that information and then use what are known approximation methods, like response surface models, to save that information, such that when you do the full system optimization, you don't have to rerun, you know, the detailed codes every time. You can just pull from the approximation you've pre-computed already. So that's one of the really powerful things about this method.

And that's what's shown here, is you substitute-- you have a series of approximation models, one for each of the outputs of the system. You see that? So these little wiggly things here are essentially an approximation of that input/output relationship at the subsystem level. Could be an empirical function, a neural network, a response surface method, or you know, simplified engineering analysis. So you do a lot of this work offline ahead of time for each of the subsystems.

Does that make sense? Any questions at EPFL? Is that clear how this works, at least in theory?

AUDIENCE: Everything's fine, thanks.

OLIVIER DE WECK: OK, so what we have now is essentially these subsystem approximations, and now we can use them for system level optimization. So now let's look at the system level optimization. So that's essentially given approximation models for optimized subsystem outputs, we want to minimize some objective function. So F, here in this case, which is a function of our shared variables, our relevant outputs, Y star, and the weighting factors, right, the omega-- not the omegas, the W's, which are the weighting factors given to each of the subsystems. And we of course, have to satisfy constraints at the system level.

So let me show you an example then of what this looks like. So here's our supersonic jet. This is our cycle zero, our initial condition. Here's the planned form of the jet, and I'm going to show you down here range and just one of the design variables, or a set of design variables, which is our wing box sheet thickness. So you go through this, and it converges after 10 cycles already. So this is what our business jet looks like after 10 cycles. And you know, you see there's a pretty big difference.

So the fuselage is much longer. This by the way, this is the same in ships, right? Your speed

that you can reach, your max speed is driven by your fuselage length, your water length. You can see the wings have a pretty high sweep angle. This is a supersonic jet. And if we look at the details, we see that the range here, the range is about 3,000 nautical miles, and what you see is that the exact range that can be achieved and the approximate range merge. So the approximation models that you have in this method are only approximate. There's some error. And you know, as you run the model several times, several iterations, the difference between the approximation and the actual high fidelity result gets less and less, so that it's actually a feasible design.

And then over here, we see the different thicknesses for the inner wing box, the mid span wing box, and the outer wing box. And you can see that early on in the first four or five six iterations, things are kind of moving around, but then they're settling out. And so for this maximum objective of range, you can see that the inner, mid, and outer wing box thicknesses are quite different.

So that's the general idea here is that this result-- this is a fairly detailed design was produced based on a system level optimization that included consideration of aerodynamic structures, propulsion, and so forth. And there are several other methods, but this is just one of the methods that's been actually quite helpful.

The challenges now in MDO, in this multidisciplinary optimization are different trade-offs. So one trade-off is between fidelity versus how expensive is it computationally. So fidelity is over here. So high fidelity means, you know, 3D finite element model analysis, computational fluid dynamics. Then you have sort of intermediate fidelity or low fidelity empirical models. And then on this axis, we have just doing trade studies, meaning you have sort of a baseline design and you just explore from that baseline design. You do limited optimization or iteration or full MDO, exploring the full space.

And traditionally, you had to choose, am I gonna use high fidelity models? But then I can only look at a few points. So then the question is can you do better? Is there a better solution out there? Or you can do lower fidelity models and explore the design space very widely, and then the question is, can you believe the results because you've done this with approximate models? So the real challenge here is we would like to do high fidelity and explore the full design space. So that's one of the big research areas in MDO is high fidelity but also very computationally efficient.

And then the other challenge is quite similar. So it's again the fidelity here, but then should you just focus on one subsystem? How many of the subsystems, including, you know, the financial, operations, manufacturing considerations, so how many of those subsystems do you include in your optimization? And so you know, ideally, you'd like to do the complete system, but then you may have to sacrifice on some of the fidelity. So that's the breadth versus depth trade-off.

OK so any questions about-- I know this was very fast. I do not expect you, by the way, for assignment four or five to do MDO on your CanSat. This is really more to give you a feeling for, you know, how far things have come and some of the computational frameworks that exist to get you to a really fine tuned optimized design. If you are interested in this, you know, we have the spring class. And you know, I hadn't offered this class at EPFL in the past, but we actually do offer a distance version with the SDM program. So there may be an option to do this at a distance as well. We'll discuss it.

But any questions about MDO? What it is, how it works, at least in principle? Has anybody done MDO kind of work? Sam, go ahead

AUDIENCE: I just had one question about if-- so here, we're looking at optimizing on one attribute, so range, but if we have multiple attributes, it's still up to someone to decide how important those attributes are in the end.

OLIVIER DE WECK: That's correct. That is exactly right, unless you have a way to combine all these into one objective function. And to be blatant about it, to be-- you know, in the commercial world, that's actually clear. It's profit or net present value. So if you're designing a commercial product in a competitive industry, you can actually pretty easily combine all this into a long-- you know, as long as you include sort of the key drivers of, you know, your manufacturing costs. If the product is not as optimal and not as performant, well, probably, you're not gonna be as successful. Your market share might be lower. So in the commercial world, people are doing MDO with profitability as the ultimate objective function.

Now, for situations where it's not about profit necessarily, but it's about a bunch of other considerations, then you can't easily combine it. You can still do utility. You can do a utility, you know, combine it in utility. But I agree with you, I think a good practice is then, in that case, to do the multi-objective, and then you're gonna have a bunch of solutions that are all non-dominated, right? And then it's a human decision making process to pick among those.

Any other questions or comments? What about at EPFL? MDO, anybody have experience with

this before, multidisciplinary design optimization?

AUDIENCE: OK, no one, again.

OLIVIER DE WECK: OK, that's all right. So just to tell you, this is really a big deal in practice. So you know, the structural optimization has been around, that's pretty standard practice, and you know, optimizing control systems. So we've optimized individual subsystems for a long time now, at least two or three decades, but optimizing at the system level is relatively new. And I think the best companies, the best organizations that are out there, they're really investing in this technology, and they're quite good at it. And that's where you can really distinguish yourself from the competition.

So let's talk about concurrent design facilities. And so I'm actually going to use some material here from EPFL because they have a very nice CDF, and I want to give credit here to Dr. Anton Ivanov at the Swiss Space Center for some of this material. And actually one of the papers is also related to this. Anton used to be at JPL. He worked at JPL for close to a decade before going to EPFL.

So what is a CDF, a concurrent design facility? It's essentially an environment where engineers of different disciplines come together to perform, you know, a system engineering study for a project. Key elements that you need is the team, the process, the physical environment, which includes the A/V and the software, and then also some way to do knowledge management, meaning capturing the results of prior studies and ongoing studies. And I'll show you some examples of CDFs.

We also do this in academia. There's challenges for doing this in academia, a short learning curve, you know, synchronizing this with the academic schedule, and then a lot of turnover, of course, as students graduate and so forth. So there's challenges.

And in practice, what we essentially have is, in a CDF, we have a set of models that are linked to each other. And these models can be, for example, here's some Simulink models shown. Here, there is some Excel models. You can have CAD. And one of the tricks is, you know, to link this together in a way that you can quickly execute different designs and look at the results from that.

So in industrial settings, probably the most famous CDF is the one at JPL, at the Jet Propulsion Laboratory, called Team X. Team X has been around for about 20 years, and they

just had an anniversary where they celebrated their 1000th study that was done at Team X. A typical Team X study is about a week. So you do a little bit of preparation, but people work together in the facility for about a week, and at the end of the week, you have basically a conceptual design pretty well worked out. And the results have been impressive.

So studies have shown that the cost estimates that Team X came up with were within 10% of final mission costs, for those missions that were actually built and flown. So of the 1,000 studies that Team X did, a handful of them were actual flight missions. But every proposal that is written starts now as a Team X study. So it's really a big deal.

The European Space Agency ESTEC has a similar facility in Noordwijk in the Netherlands. As well, all the ESA projects are going through their CDF. So one of these is [INAUDIBLE] which is a space-based observatory. EPFL is involved in that project. That went through the CDF as well. Most NASA centers and then commercial applications in medical devices, painting, shipbuilding. And like I mentioned in my introduction, even consumer packaged good companies are now building their version of a CDF, such that they can, you know, very quickly, within a couple weeks, come up with a new proposal for a product and all the major aspects, you know, the underlying physics and chemistry, the packaging, the business case, everything is considered at the same time.

Yes?

AUDIENCE: I was wondering why-- what makes the cost estimate so accurate with this method?

OLIVIER DE WECK: So there are three major ways of doing-- oh, Veronica should answer that question, since she's launching on cost modeling for her research.

There are three major cost estimation methods. One is bottom up, you know, you have the product decomposition, and you add up all the costs. One is called costing by analogy, where you look at a prior mission or a prior product and then you look at the deltas. And then the third one is based on so-called CERs, cost estimation relationships. It's a parametric method. You can also blend these. And you know, I think one of the reasons why it's accurate is because-- there are two reasons why, I would say. So you always have a cost share. One of the subsystems is costing. So you have really people who been doing this for a while. They know these models, and so there's some legacy information. And the more missions we fly, the more products we make, the more robust, essentially, these cost estimates become.

And then the second, I think, is because people that typically work in a CDF, they're non-advocates, meaning it's kind of neutral ground. So rather than pushing a particular idea or trying to be overly optimistic, you know, usually in a CDF, the calculations that are done are pretty non-biased, and that's not always the case. So that's one of the advantages of the CDF is you're getting-- you know, people are trying to say, but really, let's make it look a little better. No, no, we can't do this, because the model says, you know, it is gonna cost $280 million to do this satellite. And that's based on, you know, within prior models, and if you overlay different models, you know, with some error or confidence. That's the right answer. So it's these two reasons. There's experience, there's data, and non-advocacy, really taking a kind of neutral, unbiased position.

So benefits are, you know, improvements of quality, quick turnaround for ideas, better cost estimates, and then increased creativity and productivity in a company.

So let me give you a couple examples here. This is a CubeSat. So this is the ability to design quickly new CubeSats. This picture here I'm showing you is the Swiss Cube that was launched in 2009, still operating. I think it's the longest operating, continuously operating CubeSat so far. There's like over six years of operational data now. And so you know, for the next generation, which is shown here, how would you do a Cubesat in such a facility?

And there's three steps. You define your decomposition level. So you have your system level, your subsystem level, and then at the lowest level it's called equipment. These are the actual components. And you have your product decomposition here, and then for each of these components, you have different sheets that give you either component choices or give you different parameters like mass power with margins that allow you to then choose from this like a menu and pull together your overall design. So that's a Cubesat example.

But I want to talk to you a little bit in more detail through a design of a suborbital space plane that was done in that CDF, and this was a project known as K1000. So some of you may have heard of the Hermes, the Hermes space plane. This was a project that the European Space Agency had starting even in the '80s, and it wasn't built, but there was a lot of studies done. And this is essentially a Hermes-like vehicle but for space tourism, that's the basic application.

And so you know, how would you design a space tourism vehicle in a CDF? Well, you have to start with requirements, right? Without requirements, you're just-- it's like, you know, setting-- my analogy for designing without requirements is like taking your ship and going out to sea

without a map, without a chart, with no destination port in mind. You're just drifting in the ocean. You've got to have requirements.

So here are the high level requirements for this particular vehicle. Level one requirements-- you have to reach an altitude of at least 100 kilometers over sea level. Why is that? Where does the 100 kilometers come from? Why is it not 80 or 90 or 150?

AUDIENCE: That's where space is defined.

OLIVIER DE WECK: Right, and what happens to people who go above 100 kilometers?

AUDIENCE: They get their astronaut wings.

OLIVIER DE WECK: They get their-- you have gone to space, and like, you know, you can claim that you've gone to space if you go above 100. If you don't, if you are at 80 or 90 you don't. That's kind of internationally accepted. So that's really where that requirement comes from.

The second one is a zero g phase. You should experience weightlessness for at least several minutes. Who's done a parabolic flight before here? The zero g, you know, the Vomit Comet? Anybody? Wow, we have to get you guys out there.

At EPFL, are you still with us?

AUDIENCE: We are.

OLIVIER DE WECK: Has anybody done parabolic flights? I know there's a flight campaign coming up in Switzerland later this year?

AUDIENCE: Yeah, it will be next year, I guess, but unfortunately, no one has done this right now here.

OLIVIER DE WECK: All right, guys, we have to fix this. This is unacceptable. I mean, we have to work on this. So it's very cool to do these parabolic flights, but you only get about 20 seconds, right? 20 seconds at the top. It's like riding a big roller coaster, basically. And then you go down, and then you suffer, you know, 1.7 g's or whatever. Then you come back up. But we want several minutes.

And then the next requirement is that the passengers should carry-- there should be six passengers carried on this vehicle. Those are the top level one requirements.

Lower level-- level two requirements-- safety. Don't exceed six g, positive six g's. The vehicle must be controllable at all times. So when you're out of the atmosphere, you have to be able to control the vehicle, that has certain implications. The customer experience-- you should be able to see the Earth's curvature and atmosphere, that thin layer, at all times. The spacecraft impact on the environment should be as small as possible. And then the mass budget, the spacecraft mass should not exceed 11.6 tons with propellants. Where do you think that comes from? What kind of a requirement is that? What flavor of requirement is that? Yeah?

AUDIENCE: Transporting it for proper [INAUDIBLE].

OLIVIER DE WECK: Right, so in this case, it's actually a piggyback, right? You either have a vertical launch or you have a first stage, an airplane from which you are dropped right? So that's sort of an interface requirement.

OK, so once you know your requirements, you can really start developing your models. So here's a model of essentially the kinematics of the vehicle. So you have your environment, your atmospheric model, atmospheric density, pressure, mach number, speed of sound, a gravity model. Then computing the forces and torques acting on the vehicle itself, looking at the derivatives, integrating that. This is essentially the equations of motion of a vehicle in its environment that it's operating in, and then updating that with a certain time step, delta T.

So you can zoom in on one of these, for example, this force model, computation of the forces. And there's two frames. There's the Earth reference frame and, then there's the body centric frame of the vehicle itself. So you have, in this case, the aerodynamic coefficients. You have the forces, lift, drag, pitching moment, and then you have the actuation. This is assumed, you know, the pilots are here in control. The rocket engine itself, secondary propulsion, the rudders, and then attitude thrusters. You need those especially-- you know, once you're out of the atmosphere, the rudders-- the aerodynamic surfaces are no longer effective. And we knew this in the '60s. Vehicles-- like the X15 is very famous. Well, it had thrusters that you could you could actuate once you're out of the atmosphere. So you know, you can go pretty detailed in this.

You would then design the vehicle in a particular way, you know, particular engine, wing dimensions, aspect ratio, a particular flight profile, and then that would be a typical result that you would get in a CDF. There's two curves shown here. The blue curve is essentially altitude and kilometers. So you see there's the 100 kilometers. We need to get above this, right? This

was our requirement. So we go above 100. That's good. We satisfy that requirement. So looks like a parabola, and then we have a-- we glide here, essentially, to the ground. The whole mission takes about 900 seconds, 15 minutes.

And then the yellow curve here is your load factor, your NZ, your vertical load factor. And you can see that's a little noisier as a curve. So we have higher acceleration as we accelerate to altitude. And you can see that roughly here, roughly in the middle of the parabola, the load factor starts to-- you're basically in free fall. You're basically ballistic, and the peak acceleration is about 4.3 g's here, so we satisfy that requirement. And the whole reason we're doing it is for this period here, that zero g experience. So when the load factor goes to zero, or near zero, that's when you get weightlessness. And it turns out for this particular flight profile, we would experience 184 seconds, about three minutes of weightlessness. This is pretty similar to the flight profile that Virgin Galactic is pursuing, by the way.

So you know, looking at this result, you would say, well, I don't know about the mass here, but from, you know, reaching 100 kilometers keeping accelerations below six g, this looks like a feasible profile.

Of course, the other thing that we care about is the view, since this is for space tourism. We want to make sure that, you know, that the experience is great, not just from an experiencing weightlessness, which you can do in a closed can but you actually see something. So this was essentially coupled with a virtual reality visualization. So you can see, you know, the landing here. You can see these three views here. And you can see the vehicle itself. So the pilots are up front, and then here is our six passengers. Every one of the passengers gets their own window.

So this is the view that you would get from the front window. This is the view from the middle window. And then this is the view from the back window, from that last window here, which is closer to the back, to the engines. So what would you, looking at these pictures-- you know, this comes directly out of the simulation and the CDF-- what would you say looking at this? Anything you notice?

EPFL anything you guys notice? Look at the view from the front, the middle, and the back window.

AUDIENCE: From the front, looks better. I guess it will be more expensive, again, for the customer.

OLIVIER DE WECK:	Yeah, so that's exactly what I was asking. So it looks like either you need to redesign the vehicle so that the view is the same, more or less, from every window, or you need to start charging differential prices. Because clearly, from the back window, half your view is obstructed by the wing structure, right? Would you agree with that? Would you charge different ticket prices for this ride, depending on which seat you get?
AUDIENCE:	I guess that's offer and demand?
OLIVIER DE WECK:	Supply and demand, yes. That's right. Go ahead.
AUDIENCE:	It could also merit a change in flight plan. If your structure blocks the downward view, you might want to consider flipping over if like the Earth is what you're trying to see from above. So it may have broader implications for the mission.
OLIVIER DE WECK:	OK, so the point here, is there a way to operationally improve the experience, keeping the vehicle design the same? Absolutely, but what that means is, you would have to put the operation-- so the orientation of the vehicle, you'd have to put that into your design vector if you want to change it.

OK, so I don't know. I think this is pretty cool. So you can actually do a lot in this kind of CDF type setting. And by the way, of the things that spun out of this project is this company called S3. Their headquarters in Payerne in Switzerland. This is actually-- I did a lot of my military service at that airport. Their goal is not space tourism, you see there's no windows on this vehicle, but deployment of small satellites, custom deployment for small satellites with a reusable vehicle.

So you know, I think it's a very interesting concept, but the question is, is it feasible, right? Is it feasible-- the aerodynamics, the economics, structurally, engine reuse? There's a lot of-- this is a great picture, but you know, the devil is in the details. So really working out the details of such a design, that's what a CDF is for, and you can really go deep.

OK so any questions about this? I'd like to do a little partner exercise here. Any questions? So let's do a partner exercise, and the question is, what are your experiences with concurrent design facilities so far? Whether you've done an internship at ESA or JPL, or some other place that has a facility like this, or even if you haven't been there, kind of speculate and discuss what it would be like. Maybe, you know, what would be your favorite chair? Would you pick

propulsion or costing or flight trajectory? You know, just sort of speculate and discuss with your partner, so for which project or application did you use it? What went well? What did not? What could be improved? And again, if you don't have much experiences, just think ahead what it would be like. What would you like to get out of it if you have no experience in it?

AUDIENCE: So there is at ESA, a concurrent design facility, and I was talking with Katya that maybe also United Nations can be considered, like, the facilities there can be considered a concurrent design facility. I don't know if it worked, this approach.

OLIVIER DE WECK: Do you mean in Geneva, United Nations in Geneva or London or New York, or do you mean in general?

AUDIENCE: Exactly.

OLIVIER DE WECK: But do they actually have a physical facility with models and screens, or are you saying this more metaphorically?

AUDIENCE: Metaphorically more.

OLIVIER DE WECK: OK, but you know I-- go ahead. Go ahead.

AUDIENCE: Yeah, we're referring to the layout, to the site, we're looking at the ESA layout, where you have one big room, where you have all the people coming together, all the different models coming together. So that's like all the different experts. Then you have subrooms for different project rooms, and then you have, like, a big conference room. That's kind of the structure of the layout of the facilities.

OLIVIER DE WECK: But you know, maybe a place like the UN could really benefit when they deal with a refugee crisis or with, you know, climate change negotiations. I like this idea. Great suggestion.

What about here at MIT? Any experiences, thoughts? Yeah, please go ahead.

AUDIENCE: Yeah, AFRL, we have a rapid innovation officer. He's like a general equivalent as a civilian, a LOC BOS, and he basically is tasked with breaking all the rules in AFRL. Whereas most projects take years to complete, he fields questions from the field, from the theater, from war fighters, and they have an urgent need. And he gets them a solution in under 12 months. So he pulls in all these experts from industry, academia, wherever. He'll lock them in a room, and

they don't leave until they come up with a solution. And he has the money and the influence, the backing by the AFRL Commander to get it done.

OLIVIER DE WECK: And what does that-- I mean, if you can share, roughly what does that facility look like roughly? Does it have, like, models and computers and screens?

AUDIENCE: Just white walls, like, whiteboard walls, and they will just be covered with everybody's ideas by the time they leave there.

OLIVIER DE WECK: Oh, I see. So it's not as much computationally supported. It's more of a human process.

AUDIENCE: That's right.

OLIVIER DE WECK: Cool. OK. Great. Interesting. Any other experiences? Anybody been at JPL or seen a commercial facility in a commercial company that does this kind of work?

Please, go ahead, Justice.

AUDIENCE: Well it's less of a commercial-- I was interning at the Aerospace Corporation this past summer, and I was working in their concurrent design center. And it uses a similar layout as JPL's Team X's design, where it's basically using models. That they have, like, a group of engineers who sit at desks, and there's multiple engineers per subsystem, and they just discuss. They're given a mission objective, and they discuss how they can best design the system based on that. And they model it and optimize it in the subsystem, and then they feed it back to a systems level design.

OLIVIER DE WECK: Did you participate in that?

AUDIENCE: Yes.

OLIVIER DE WECK: So what did you think of it?

AUDIENCE: A lot-- it's a lot to take in at first just trying to understand how each of the models work. You're not necessarily supposed to work on each model, but as an intern, they wanted me to kind of see what each of the models would do and not necessarily exactly how the optimization techniques work because I didn't necessarily understand it, but just more of have an idea of

how it actually works as a whole.

OLIVIER DE WECK: OK, good. Anybody else at EPFL want to share any experiences?

AUDIENCE: Apparently not, but just a problem with the sound. We didn't hear anything on what the guy before just told us.

OLIVIER DE WECK: OK, Justice, did you push the mic button? OK, yeah, so basically, he worked-- let me try to summarize. He worked at the Aerospace Corporation. He interned at the Airspace Corporation. Was it in Virginia or California? California, in California. And Aerospace Corporation is like a think tank, and it was a very similar set up with different subsystem experts, different models. And Justice was saying one of the big challenges is so much information, taking it all in and understanding all the models and coupling, and so forth. So and I agree. It takes a while. If you haven't been working in this environment, it takes a while to really get used to it.

So the last thing I'll say here is that-- well, two things. One is-- and they're linked-- one is that these CDFs are not cheap. If you really want a high-functioning CDF, you have to have people that are dedicated to this. You know, the software licenses expire. The technology gets old. You need to not just use it, but you need to constantly refresh a CDF, keep it up to date, and it just doesn't happen automatically. You've got to have budget for that. You have to have people who are dedicated to updating the models, updating the facility, and so forth. So it's definitely something that needs to be invested in on a regular basis.

So what happened to us here at MIT in air astro, actually, when I was sitting where you are sitting as a grad student, we created a CDF in the mid '90s. And you know, it was just like the one at JPL. We had desktops, you know, with labels-- thermal, propulsion. And we used it quite actively for maybe five, six, seven years, but we didn't put the budget into it every year. We didn't have a dedicated person who was really. looking after it. So what happened to our CDF is that entropy took over, and it sort of got more and more dated.

And when you go and look in this room now, it kind of looks like a regular conference room. So all the desktops that we had are gone. I guess they'd be pretty old by now. And there has been a pretty strong debate among the faculty whether actually we need this. Should we recreate one? And the argument goes as follows, the proponents say that a CDF is really something special. When you enter in the CDF, you know this is a special place. This is a

facility that's designed specifically for doing rapid studies. All the information you need is there. The walls are plastered with reference information. You know, really, that's what it's for.

The counter argument is that this is no longer the mid-'90s. Now, everybody has laptops. Everybody does mobile computing. All the models and data should be on the cloud, not sitting anymore in a dedicated server facility, but it's sitting on the cloud. And therefore, you know, physically, a facility that's dedicated is no longer needed. What you do is you get together in a more kind of flexible ad hoc way. Everybody brings their laptops. You have the big white walls, like you said. And you can just create a CDF on the fly, right? Everybody brings their computing resources. Stuff is on the cloud. And then you get a lot done, and then you disband again. You clean the white boards and you're done. Therefore, no longer needed a dedicated CDF. Those are the two countervailing opinions right now.

My view is, I think a CDF still has value today. I think it is valuable as a dedicated facility, but it needs to be updated. So yes, the models should be on the cloud. Yes, you can bring your laptops, but it is still valuable to have dedicated computing resources, again. So I now think a hybrid model is actually the right way to go. So we'll see where it goes, but it's a pretty active debate. There have also been, you know, theses written on CDFs, comparing different CDFs and so forth.

So take a look at the Anton Ivanov paper in the reading and just think about this. I think this is a big deal, and it's going to be even more important than the future, whether you're doing space missions, aeronautics, whether you're designing commercial products, you know, or in the military, these facilities that allow you to make these design decisions in a very interactive, very quick way are gonna be more and more important in the future, and this is a big part of system engineering.

Any final thoughts on this before we move on?

OK, so here some lessons learned from the EPFL CDF. So this is particularly in the university environment, so you know, giving access to a wide body of students. You need to adapt this to the university's schedule, the learning curves, the analogy of a smoke-filled room or war room, optimal team size. You know a CDF with 50 people on it is not as effective as one with, you know, seven plus or minus two. And then distributed CDFs-- they don't typically work as well.

So I almost wore my-- I have a shirt that's almost the same. Anybody know who this is or

where this is from? This is Sheldon Cooper in the *Big Bang Theory,* his virtual reality presence. So the point here is that virtual reality is OK. So you can imagine a CDF where half the people are in the room and half the people are dialing in remotely, but it's never quite as good as everybody is there physically. So human interaction is critical.

OK, the last thing I want to talk about today is the CDR. This is a big, big milestone. This is a very big deal. The critical design review, the main purpose of it is to approve the final design in all its glory, in all its details. So generally, the way we talk about it is, if you pass the CDR, you've got the green light to quote, end quote, cut metal. Of course, you know, our systems have software and hardware. So the cut metal is essentially a euphemism for you can go ahead and manufacture the system now.

CDRs are typically the biggest review. The largest number of people you will have at the CDR. You know, if you had 10 people at the SRR and 20 people at the PDR, you're gonna have 50 or 100 people at the CDR. Because any question that comes up on any of the parts-- say, is this really right? Well, what could be the impact of this detail on the rest of the systems? Somebody should be there who designed that part or who can answer the question at least.

And CDRs are typically longer than one day. In fact, for a big project, a CDR can be a whole week of reviews. So people get really tired. Big CDRs are very, very detailed. You really have to drink lots of coffee and stay awake.

So here's the description of what a CDR is. One of the guidelines is that approximately 90% of the engineering drawings are approved and released. So it's OK to have a CDR if you have not 100% of everything done, but you need to have at least 90% complete. You can't do a CDR, with, you know, half the drawings done. And it happens at the end of phase C, the final design phase C.

For very large projects, you can have sub CDRs. You'll have a system level CDR, and then you can have sub CDRs for every major subsystem in your project. So an example, I put a link here. This is from late 2014, the CDR for the James Webb Space Telescope that I mentioned earlier. They passed the last major element level critical design review, which was a big deal. So launch is planned for 2018.

But for big projects, what I'm trying to say is, there may be more than one CDR. There may be the master CDR for the whole thing and then sub CDRs for different subsystems. Now, if you don't-- well, let me get to this on the next chart.

Yeah, go ahead.

AUDIENCE: Are the, like, CDR slides for NASA projects, are those available anywhere as a public entity and can you get access to any of that stuff?

OLIVIER DE WECK: It depends. Not always, because especially if commercial companies, like, you know, Lockheed-Martin or somebody is doing, like, the spacecraft bus, they would consider those details proprietary. So in general, I would say no. But for CanSat, actually, you've probably already found them. For CanSat, you can actually find a lot of this. So prior PDRs and CDRs, you can see that. I encourage you not to look too much. You know, you kind of want to come up with your own thing. But you know, for projects that don't have a commercial interest or proprietary technology, it's often available.

So to show you just-- this is a table from the Handbook, page 178. This is the entrance criteria and the exit criteria for CDR. So entrance criteria means what do you have to have in order to be allowed to go forward with a CDR? And it's mainly three things-- successful completion of the PDR, and responses made to all the PDR RFAs and RIDs. RFAs and RIDs are essentially deficiencies, open items that came up in the PDR. So you have to close them out. A CDR agenda, and then number three is all the technical work packages. And so you see, it's a huge list here-- baseline documents, specification, fabrication, assembly and integration, technical data package, acceptance criteria, verification and validation plan, launch site selection, and operations plan-- super, super, super detailed. This is a lot of information. And at the CDR, you're gonna go through all this. So you can imagine what that means.

And then success criteria, on the right, is, you know, in order to successfully pass the CDR, these are all the things that you have to basically have. You have to have approval for all these things. Now in reality, just like a PDR, at a CDR, things will come up that you missed or that people have concerns about. And so you will make a list of those, and then you have to address these action items coming out of a CDR. If those action items are serious-- if they're really serious, you could fail the CDR and say, you know, this is not gonna work. We should not proceed. And they send you back. Or they cancel the project. That's happened also.

More likely, is that they'll ask you to do a so-called delta CDR. So they'll have you, you know, it's like retaking an exam or something. They'll have you come back, and you'll have to do a CDR again but only on the delta, the portion that raised concerns at the main CDR. You don't have to do everything again. That's called a delta CDR.

OK, any questions about CDR? This is really the-- this is the big one. This is the big milestone. And obviously, you know, the concept or the architecture you chose at the PDR turns out not to be a good one, it's very difficult to make changes to that here, right? You're kind of locked in to whatever design or architecture you chose at PDR. That's why PDR is also so important, even though you have less detail.

Yeah, Sam, go ahead.

AUDIENCE: Do you have to have contracts in place already for the, like, manufacturing and operations?

OLIVIER DE WECK: Oh, absolutely. I mean, basically green light CDR, the next day, you start manufacturing. So you know, every portion, all the components, all the manufacturing, all of that has been bid, selected, contracted. Everything's in place contractually, absolutely.

OK, any questions at EPFL about CDR? Who has actually gone through a CDR on your end? Has anybody done a CDR for any of the projects? OK, go ahead. What was it like?

AUDIENCE: I did also a CDR because I had the PDR last time at the [INAUDIBLE] program. And one comment I want to make. So the teams were allowed to fail the PDR and still fly. If they failed the CDR, then they didn't fly at all. So it was very strict and demanding

OLIVIER DE WECK: Yeah, so would you say that the CDR is more strict than the PDR?

AUDIENCE: The CDR is more strict.

OLIVIER DE WECK: Yes, I agree. I agree. Definitely. OK, anybody else? Are any of you working in systems other than aerospace, like, I don't know, medical devices or you know, scientific instruments?

So the reason I bring this up is because you know, medical devices have to be approved by, here in the US, it's the FDA, the Food and Drug Administration, and then in Europe, there is equivalent, you know, European level thing. It's very, very detailed. It's really not that different. So presenting your medical device design to a panel of experts-- doctors, you know, people in medical safety, regulators-- it is like a CDR. You basically are going through a CDR when you're trying to get your medical device approved for sale and use. So it's not often called a CDR, but it effectively is a CDR.

So this is, especially in industries where there's a lot of money at stake and people, you know, safety, people's health is at stake, this is common practice.

OK, so let me summarize here. So this is the detailed design phase we discussed. It's very important. Basically, you take your PDR level design and you go define all the details in full maturity. You create design documents and models. And you know, as we talked about last time, we're sort of in this transition phase of producing documents versus producing models. Typically, we have both now. Examples would be detailed bill of materials, computer-aided design files-- we'll talk about CAD next week, computer-aided design-- software and control system definition, user interfaces, et cetera, et cetera, et cetera.

Multidisciplinary design optimization is a very powerful concept to optimize at the system and the subsystem level and trade off between disciplines and objectives.

The CDFs are the concurrent design facilities, and they've now become standard practice both in aerospace and product design companies. If you have not had the chance to experience and work in a CDF, I really encourage you to do this. I think it's a life-changing experience. It really-- it's not easy, because like, you know, Justice said, there's a lot of information to absorb. But once you've done this and experienced it, you realize if you don't have a CDF, if you don't design products this way, you really should.

And then CDR is the big milestone, the last gate before, quote, end quote, cutting metal.