

RICK RUNDELL: Good morning, ladies and gentlemen. Welcome to BUILDing the Future of AEC. So I'm your emcee for four terrific presentations today that we've organized around the kind of work that we're going to be supporting in a new facility in Boston that we're calling the BUILD Space. That's what the BUILD is about-- Building Innovation Learning and Design.

I'm Rick Rundell. I'm a senior director at Autodesk, who is running that project. And we've assembled these four presenters. And I will introduce each of them before they present.

And then at the end of the presentation I'll do a quick update on what the BUILD Space in Boston is actually about. And that'll just take a few minutes. And then I think we may have a little time for questions at the end.

But, of course, as always, you can when we wrap up and people are stepping out you can find these guys and catch up with them later. All right?

So our first presenter is Gustav Fagerstrom. He's a registered architect and an associate with BuroHappold in New York, where he leads the computational modeling for structural and facade engineering efforts.

Specializing in design computation, he operates at the intersection of architecture, engineering, and computer science. He's practiced with UFO, KPF, and UNStudio, and has published, given workshops, and been on academic juries at architecture and engineering universities worldwide.

Work of his has been presented at the Venice Biennale, [? Cadria, ?] ACADIA, Fabricate, and the Smartgeometry Conference. Lots of acronyms in that introduction. Anyway, join me in welcoming Gustav.

GUSTAV Thank you. Thank you.

FAGERSTROM:

[APPLAUSE]

Is this a good volume? Yeah? It's not a huge room. Where should I stand? Stand here? Oh, I don't want to block this thing. So yes, thanks so much for having me in this great panel. I'm

really looking forward to seeing what everyone else has put together.

And so I will talk about something that, if you're familiar with it it needs no further explanation-- reality capture. If you're not, its various technologies that help, as the name implies, capturing real data and making something of them.

In the case that I'm going to show you, it's about creating geometry from photos, and taking temperature readings and kind of feeding that back into the geometry.

So I want to start with a kind of musing on the notion of scale. So this was done by the Eames couple for IBM in 1970-something, I seem to recall.

And essentially it zooms out to a very large distance away from the Earth in order to stress, kind of, how tiny we are in the solar system. It's sped up, unfortunately, so if you want to check out the full thing, it's online. And it has a pretty great narrator as well. But it's kind of long when it's in it's full stage.

So here we are at 10 to the 24th power meters, all the way out at the edge of the solar system. Then we go in again and into the kind of negative numbers instead-- so down at the molecular level. And bear in mind this is before they could actually take pictures like this. So this is all kind of '70s CGI.

So the reason why I'm showing this is that I think the point that we're at kind of in history and in AEC is where we have a pretty amazing opportunity to think of data in a kind of scale this way almost. So the kind of ones and zeros of binary, the nucleotides of DNA, or what have you.

And we can kind of choose on what scale to operate-- jump between scales and read data, read capture assets as I'm showing here at certain scales, and kind of put them back in our computers or wherever we want them to go at the scale that fits our purposes.

So architecture and engineering many times is about dealing with real world assets. It can be things like the geolocation of a project, survey data, physical objects, environmental data, and thermal readings-- which is what this focuses on. And then we can go into humidity, wind speed, air composition, and so on.

And all of that is kind of there for the taking. And one step in recording that is you will receive something that often looks kind of like gobbledygook. And its raw data, and what do we do with it? It doesn't have necessarily value as of yet.

So that need some sort of parsing, which is when we kind of work with it-- elaborate it-- so it becomes information from data. And ultimately then, that needs to get into some sort of indexing or cataloging system-- which is, for instance, the way that we many times in digital building practice we catalog, we charter relationships, and we put them in a model for simulation and presentation and other things.

So this particular case study was done in Hong Kong in summer, so it was really, really hot, really, really humid. And the purpose was to kind of sniff out inefficiencies or differences in usage of curtain wall building. Pretty simple geometry case so as to not get bogged down with huge files for reasons of complex curves and such.

So the first part is data acquisition. And this was done by using this rig, which is a Phantom DJI 2 that we fitted with a GoPro camera and some [INAUDIBLE] kind of components for streaming data and reading. And it has its own radio transmitter, and we also have a thermal scanner, which you can see here.

Sorry, I wanted to do this. And we actually did it with two different models. I'll show you kind of what different resolutions can do with this type of experimentation a bit later.

So this little drone gets fitted with this. And you fly it around. And the first step there then, in data acquisition, becomes geometry creation. So let me try and show if this will work. Yeah.

This is the Chinese University of Hong Kong. And the building that you see right in front here is the Yamamoto Building? Is that what it's called? Yeah. It's a science building of some sort. So this is the drone kind of flying from two different departure points to take as full a 360 survey of this building as possible.

And then that gets fed into a photogrammetry algorithm, which-- I don't know if that's provided-- which is called Recap, which you can also take a look at in the Expo space.

So we had that done by Rick's [? eminent ?] calling [? ENQ, ?] who was there at the workshop as well and kind of graciously lent us his fast cloud credits that he has as an employee. You can do this also without that, but it takes longer. So this whole mesh was reconstructed in perhaps 10 hours or so.

And this is the resolution of the texture mapping that you get from just these pictures taken from the drone. And of course, you can see that there are certain things that look like they're

melting.

And it's certainly something where you can go up in resolution if needed. But this is kind of a proof of concept study at this point. So doing it kind of in a fast and furious way in that low cost as well. All the gear is very, very economical.

So this is what we get out of that. And at the same time the drone will have read temperature data from this facade with that thermal scanner that I showed you. And that gets, then, into the other part, which is data management.

So you can think of this geometry creation, you can think of as data management in a way as well-- because we're taking one type of data, which is pictures, and we're making another type of data, which is 3D mesh geometry.

And then we have more abstract data management, which is in this case the temperature readings. So see here. Right. So that was done in this case with the .net program that runs in Grasshopper, which is called Elefront. And it essentially becomes an attribute manager for geometry objects in Rhino.

So with these tools you can control kind of non-geometrical attributes-- things that live in the geometry. In fact, the same thing that you would have in a regular model-- like all your BIM data that you can't see but you have it embedded in the geometry. That's what this does, but with, in this case, [? round ?] objects. But this could really be done in kind of a neat digital environment as needed.

So out of that, then, we create a sort of proto-building information model, containing in this case specific types of information but not all of it. And the way to do that is that the degrees of freedom in flight are the 3 axes here-- the pitch, the roll, and the yaw.

And they need to be correlated with the readings. The readings are based on kind of-- you get the GPS position of how the temperature reading was taken. You know at what orientation in space the drone was at that particular moment.

And that can be related, then, to the CAD space that the model lives in, so as to be able to take those readings and kind of reconvert them into, in this case, color values that are embedded with the tools that I've just showed into the mesh faces.

So this is done with a kind of high quality lens thermal scan that is one pixel. So obviously

there's a question of processing time. And you can see that it kind of paints a very tiny dot per iteration.

Then we also did the run with a 64-- 4 by 16-- pixel array, which is much lower quality. But it covers like it's a broader brush, essentially. So these are two different ways of trying out the same thing. And again, it's a quality versus speed trade off at the end of the day.

So this is an example of what comes out of that thermal scanner. It's just surface temperatures in degrees Celsius. And they are then mapped.

We create a baseline-- maximum, minimum-- for that particular day or that particular experimentation run. And then that becomes a color scale in the software, in this case. And then that gets mapped back onto the model.

So at the end we can get something like this, which is sort of an as-built thermal survey of how the building is performing. What's the surface temperature? What might be the reason for that? You see windows with closed shutters and AC turned off and on.

And then there are also the kind of future step of this which we intend to carry out next year-- is there are some fairly recent research papers where you can use linear regression to estimate the u-value of a facade component without having to take an interior temperature measurement.

And to be quite honest, I don't understand how that works. But there are researchers that have done that and have come up with algorithms for how to essentially get the r-values from surface readings.

So we've been discussing working together with some such people. And hopefully this will become-- kind of in the next iteration it will become data, which is more in line with what we would be looking for as building designers or facade engineers and such.

So finally, project credits. This is by no means my own work. The top row here are the inventors and creators of the software-- not recap, but the rest. And the kind of equipment set up. It's Alan Tai, Roman van der Heijden from Front, Inc. It's an engineering firm.

And then on the lower row here are our excellent student helpers that came from half the world to participate in this. So they should have full credit for this as well. And that's it.

RICK RUNDELL: Awesome. Thank you, Gustav. Our next presenter is Andrew Payne. Andrew is a registered architect and principal research engineer with us at Autodesk. He completed his doctoral degree at Harvard's Graduate School of Design in 2014, received his Master's of Architectural from Columbia in May of 2005, and earned his Bachelor's of Science from Clemson in 2002.

Andrew's work explores embedded computation, parametric design, robotics, fabrication, and 3D printing. He has lectured and taught workshops throughout the United States, Canada, and Europe. He has served as an adjunct assistant professor of architecture at Columbia, and he also sits on the board of directors for the Association for Computer Aided Design in Architecture.

Please join me in welcoming Andy.

[APPLAUSE]

ANDREW PAYNE: All right. Can you guys here me OK? So thank you, Rick. I appreciate the invitation to come and share some of the research that I've been involved with.

Today I'm going to show some of the work that I did as part of my doctoral research, which I just graduated about a year and a half ago from Harvard.

And really it falls under the domain of physical computing, which in the broadest sense of the word is really about creating physical systems that can really sense and respond to the world around them through hardware and software. It's a pretty basic definition.

But when we apply that to the domain of architecture, I think it raises some really interesting questions. Which is to say, what does it really mean for architecture to be considered smart? Right?

We hear smart buildings and smart cities and smart products, but it's really unclear what it means for architecture to be smart. If we look at the literature there's no clear definition.

In fact, if we do actually review the literature, what we find is that most of the research behind smart buildings has revolved around the integration of different building control systems through new novel communication protocol.

So starting in the 1980s, as single-function dedicated systems that really don't communicate well with one another, to fully integrated cloud-based enterprise systems that are available

today. But I think it's interesting that the how and the when things begin to get turned on is still following a relatively antiquated model. Right?

It's a single closed feedback loop where we might have some sensor and some predetermined response to some sort of input that's defined by a system's designer. It doesn't have the ability to learn or adapt to how the space is being used.

And I would question whether or not that should be considered smart or not. In fact, it brings up a question-- what is smartness and how does it relate to intelligence? Intelligence underscores this capacity to learn and understand and reason about how space is being used.

And I like this definition where an intelligent environment is really able to acquire and apply knowledge, and then adapt to how it's being used to improve that experience in that space. Right? That's how, I think, intelligent environments-- what we should be striving for as architects and designers.

So if we put it more simply, physical computing systems for buildings really need to have this capacity to learn and adapt over time. And the other side of that story is that my research was specifically focusing on office spaces.

And we've all been in office spaces where we walk into the space and we have very little control, if any control, over the space and we're often uncomfortable. Right? Many of the buildings that we work in and live in today have been highly engineered.

This is a case study of the Kresge Foundation Building in Troy, Michigan, which is rated as LEED platinum. It was highly engineered, it's energy efficient. And these are the architectural photographs of the space.

But what the reality actually is closer to this. And what we see is that we find portable heaters under the desk and task lights and headphones and desk fans on top of the desk. And what that underscores is the fact that the way buildings are designed are based on this idea of the average user. And really nobody falls under that same profile.

And so the only recourse, because nobody has control over their environment, are to bring in these low-cost, inexpensive devices to augment their environment. But in many ways these are working against the system designer's goals, right?

If somebody's cold, they turn on the portable heater. The thermostat detects rising

temperature and so it kicks on the AC. And so you get this negative feedback loop, which is not only energy inefficient, but also detrimental to the user's satisfaction. Right?

And so I think if we were to also paint with a broad brush about what the future of physical computing systems for building should be, they need to really provide greater personalized control over the environment. So that's sort of the umbrella of where my research was focusing on.

And so now I'm going to sort of explain a little bit-- I've mentioned that physical computing systems are this marriage between hardware and software. And so as part of my research I developed a number of different hardware devices and a software system, which I'll explain in a minute.

So these are just a couple of the pieces. I'll go show quickly a couple of the pieces-- like a low cost wireless sensor module. A lot of these are out on the market, but I designed my own.

It can last for over six months on a single charge. It would track light levels, humidity, temperature, and motion detection. So that gives us an idea of how the space allows us to track information and then store that in a database-- say, every eight seconds was what I was tracking.

I also designed a smart plug and built a smart plug. This solves part of the equation, which was I didn't want to basically throw out all of the existing infrastructure that people have, right?

People already use portable heaters and desk fans and task lights to augment their space. But they have no ability to communicate back with the larger central system.

What this is is basically a smart plug or a smart adapter, which can wirelessly turn on and off anything that you plug into it. So you can plug in your portable heater or your task light, and now that has the ability to communicate back to some larger central building system. Right?

So it takes those everyday objects, but it makes them a little bit smarter by being able to communicate through whatever protocol.

The other side of this is that you have to then communicate with the central system. So I designed a infrared remote that could wirelessly communicate back to a central HVAC system.

So it could control the fan speed, whether it's going to be in heating or air conditioning, and

temperature. It will basically allow you to communicate wirelessly to the central building system. So now your everyday portable heater or desk fan can communicate back to the larger central system.

So those are just a couple of the simple ones. I'll show you in more detail this particular device, which was a smart desk fan, if you will. And I'll go into more detail because I think this is kind of an interesting case study in how we can begin to try to optimize-- a really small scale, but the potential is quite large.

It's been shown in some studies that we can save anywhere between 17% to 49% of cooling cost if we increase the ambient temperature in a room just a few degrees, and use a low power fan to offset that thermal discomfort. Right?

And the key is that it needs to be a low power fan. Generally speaking that's around a 20 watt fan. So this is a six watt fan, so it's about a third of the normal power.

But what it also does is that it's been shown that you can optimize occupant comfort by directing the cool there, where it most affects the user. And for cooling that's been shown to be the face and neck region.

So what this fan does is it actually has an embedded webcam and uses custom facial recognition software that I wrote to direct the cooled air of the fan to the parts of the body that most affect comfort.

So, again, starting like most product designers do-- starting with sketches, moving into more detailed CAD models, actually trying to figure out exactly how this thing is going to be built. So again, I built all of the circuitry involved. These are all 3D printed parts.

And then ultimately I built four of these prototypes which will go on to this user study that I ran, which I'll talk about at the end.

So now I'm going to show just a very quick video about this particular product.

[VIDEO PLAYBACK]

[MUSIC PLAYING]

-Meet the new smart fan. This little guy is designed to be the smartest smart fan ever invented. He has the ability to learn when and where to focus his attention using facial

recognition to direct cool air towards parts of the body that most effect comfort, like your face and neck.

He's also smart when it comes to energy. Studies show that we can save substantial amount of energy if we raise the ambient temperature in a room just a few degrees, and use a low power fan to cool off body.

Yet most fans only have the ability to perform one of two actions-- they can either sit still, or oscillate back and forth. Because most people don't sit still or rock back and forth all day, regular fans waste an inordinate amount of time and energy cooling the area around the actual person. Needless to say, that just blows.

The smart fan uses two small three watt fans, drawing about a third of the power of a typical desk fan. Plus, with his built-in video camera and facial recognition software, he spends all of his time and energy making sure you feel good.

The smart fan also has the ability to communicate with everyday devices, like task lights, personal heaters, and larger building control systems. And as if that weren't enough, he also uses state of the art machine learning algorithms to get better and better at maximizing your comfort. Smart, right?

Welcome to the future of personalized comfort control. Your personal desk fan just got a whole lot smarter.

[END PLAYBACK]

ANDREW PAYNE: OK. So the smart fan and the number of other hardware devices were just part of the equation. The other side of this is software design. And so I'm going to talk a little bit about my setup. I hope I'm not going over my time.

Which, the key here was that I was looking at machine learning algorithms-- algorithms that can actually learn how a system is being used or how a space is being used and begin to adapt.

So after a review of the literature I settled on the idea of using an artificial neural network. And part of the benefits of using a network is that they're incredibly good at being able to be trained. And they don't have to have any prior knowledge of the problem that your training it for. They can represent non-linear functions, which I'll explain more about in just a second.

But I think one of the keys is that they can be adaptive. They can change over time. They can be retrained based on user feedback.

So neural networks in many cases are built on top of the idea of our biological neural networks, where we have a cell neuron, which has any number of inputs that are coming into it, that receive electrical or chemical inputs-- stimulus. And if it reaches a certain threshold, then that cell then fires an electrical signal out through its axons to other neurons that it's connected to.

So if we were to think about this from a computational perspective, we can simulate that by-- we have the equation on the right and the diagram on the left. But essentially we have any number of input nodes on the left, which are weighted. And those just get added together.

And then they get multiplied by some activation function and compared to a threshold value. And again, if that then becomes greater than that threshold value, then it sends a signal out to the next neuron in the list.

So generally speaking, the most common type of artificial neural network-- there's many. This is rapidly advancing. But the most common type is a Multilayer Feedforward network, where you have a single input layer, any number of hidden layers, and then an output layer.

And I'll try to explain a little bit how the process of training a neural network works. There's two phases in training a neural network. Essentially you have the left hand side where you can pass it in any number of inputs. When I say an input, let's think of it as a number. Right?

So we can have any number of inputs on the left hand side. And we can pass it through. And say we're trying to train it to approximate some sort of algorithm. Let's say, for example, we're trying to add two numbers together. We're going to create a neural network to add two numbers together-- or in this case four.

But what we have is that we have these numbers, and those are getting multiplied by those weights. And those weight at first are randomized. So as it propagates through, the value that you're going to get out is completely off, right? You might have two numbers that is 2 plus 3, and what we should get is 5. But what we actually get is something else. It might be 8.

So we have this different discrepancy between what we should have gotten and what we actually receive. So at that point we actually go through a back propagation process that then

goes back and tries to adjust the weight of each of the neurons so that what it gets out will be closer to what it should have received.

So in that case it's actually a process of training the neural network to approximate that function. And so in this case, when we talk about how we would apply a neural network to the AEC for thermal comfort prediction, we need to know a little bit about how we begin to calculate thermal satisfaction within a building.

And the current standard right now is the ASHRAE Standard is using something called the predicted mean vote, which was actually figured out in the early '80s. But what they found out was that six variables really control thermal satisfaction environment.

It's air temperature, mean radiant temperature, air speed, humidity-- and then two personal factors, which is how much clothing someone is wearing, and what type of work they're doing. Are they doing a lot of strenuous labor, or are they just sitting in a chair?

So all six of those factors you form this really long equation. And what comes out of that equation is a value between negative 3 and positive 3, which maps to this thermal sensation scale where negative 3 cool, positive 3 is hot. And really where you want to be is right around zero.

So that's the current state of the art, and what we use to define buildings are based on those values. So what I did was I designed a training interface for the artificial neural network-- which I'll show a video and later I'm going to ask you to click on the video when I play it to pause it at certain points, if that's all right.

So what you're seeing here on the left hand side is that there's a number of sliders. And this is actually using that predicted mean vote algorithm that ASHRAE uses. So you can change any of those variables. And what we find here is the actual value that's coming out using that equation.

So we changed sliders in real time. And we can see that in general the predicted mean vote and the percentage of people who would be satisfied or dissatisfied would be conveyed right there.

What we can then do is actually create what's called a training set. So what we're going to do is create any number of-- can you pause it.

We're basically going to create 5,000 iterations where we're randomly going to create, randomly set, all of those six variables. And then record what the output's going to be. And that's what we're mapping here.

So we're basically writing out a text file that says, OK, we have these six variables. This is the output. These six variables and this is the output. OK. Thanks. This is just mapping what the outcome would be in a vertical scale versus a histogram

Next, what we can do is actually load that into the neural network trainings area-- the middle phase. In this case you have many different options of actually changing the topology of the neural network. You can change many of the different settings-- which, unfortunately, I don't have time to go into. But you can change a lot of the different ways the neural network will behave when it's training.

Now what you'll see here is that once we start the training process-- remember I talked about that you have that discrepancy between what the neural network-- can you pause it? Sorry. What you got out of the neural network and what you should have gotten, right?

That's the thing about training is that you need to know exactly what you should have gotten. And if it's learning over time, what you would find is that that error between those two values is going to go down.

And so what you're mapping here on this graph is over time, it's learning how to approximate that predicted mean vote algorithm. Right? So the error between what it got out and what it should have gotten decreases. OK. Thanks.

So then we can load back in the neural network that we just trained that's based off the ASHRAE standard. And now what you're seeing is using the same sliders-- we can adjust the sliders-- but now the output you're seeing is the output of the neural network and not the actual equation that we got.

Now here is really where the magic happens. We can test the run set up, which is actually-- can you pause that one for a second.

What we're seeing here is that the red histogram is the same histogram that we saw on the training data set. And the blue is actually the profile of the output from the neural network. And if it matches, what you see is that it's a fairly good fit model-- that what you got is pretty close.

But the key here is that-- OK. You can play it now. The key here is that you can actually go back and say, given those current conditions it though I was neutral, but in reality I'm actually cold. I can then retrain the network. And what you see is that the profile shifts.

So what was originally the profile of the ASHRAE standard-- which is the predicted mean vote model, which is based on the average user-- is now being retrained to adjust to the current conditions that said, I am now closer to this cold shift-- because I gave it some feedback that said, under these conditions I'm no longer neutral. I'm more cold And so it retrains the network. I think that's really the key of this whole system.

So I developed on top of this is an online dashboard that would allow a number of users to control all of the different devices in the system, and also see the real time statistics about the space, the temperature, humidity. And then I ran a six week trial within four offices at Harvard.

I used sensors to measure sort of the power consumption and user surveys to measure user satisfaction. And really, overall, what we find is that there was a reduction in power consumption of about 32% across the four offices.

We found that satisfaction with it-- excuse me-- that's the overall satisfaction with the thermal environment improved using the neural network design control system versus the control of the experiment.

The desire for more control actually went down-- which I think is interesting, because in many cases we desire to have more control over the environment. But in using a learning system, that desire went down because people were more satisfied. They didn't need to actually adjust the temperature because it was learning how they were being used.

And then also an interesting-- which has large financial implications-- is productivity. In this case productivity was reported as increased throughout the study. So again, I think I've already gone way over my time. But thank you. And I'm happy to answer any questions.

[APPLAUSE]

RICK RUNDELL: All right. Thank you, Andrew. That's pretty impressive. How do we get one of those fans if we want one?

ANDREW PAYNE: [INAUDIBLE].

RICK RUNDELL: No. All right. All right. So our next presenter is Nathan King. And Nathan is an Assistant Professor of Architecture at the School of Architecture at Virginia Tech. And he's taught at Harvard University Graduate School of Design, and the Rhode Island School of Design.

His background is in Studio Arts and Art History, and he holds a Master's degree in Industrial Design and Architecture, and recently earned his Doctorate of Design degree from Harvard GSD. At the Harvard GSD he was also a founding member of the design robotics group, with a focus on computational workflows and additive manufacturing and automation in AEC Industries.

Outside of his academic work he's also the Director of Research at the MASS Design Group, where he collaborates on development and deployment of innovative building technologies, medical devices, and evaluation methods for global application in resource limited settings.

He consults on the development of research facilities, programs, and software to support the exploration of emerging opportunities surrounding technological innovation in art, architecture, design, and education.

And most recently, he's been instrumental in consulting on the development of our program and a lot of the ideas around the BUILD space that we're developing in Boston. So please join me in welcoming Nathan.

NATHAN KING: Thanks.

[APPLAUSE]

NATHAN KING: So I wanted to talk a little bit today-- instead of showing projects-- show a bit of research surrounding the kind of feasibility and efficacy of additive manufacturing in buildings.

So we see a lot of speculation and research centered around scaling off of 3D printing, different ways to apply printing or other forms of additive fabrication to the AEC industry. But we rarely see an actual conversation around the why and the how and the what are the solutions and the issues and the consequences surrounding that.

So this presentation we'll go through quickly why we consider novel building technology through automation rather manufacturing. We'll kind of segue into the next presentation that's a little bit more on robotics.

The building industry is inefficient. Construction's inefficient. Time delays, cost overruns, labor productivity issues, and so on, affect the building industry in a negative way.

This has been tested through Automation Research kind of through the '80s on how to increase productivity of building and reduce labor on buildings. And it often resulted in sameness. So automation in Japan in the 1980s resultant in sameness rather than variation.

But now in the current kind of contemporary field of computational design, we're often forced to build things that have complex geometry. So another piece of the automation additive manufacturing component, how do we deal with or address complex geometry?

These are some case studies, or really some news blurbs, from the various blogs that kind of talk about these issues in relation to some projects that I'll go into a little bit more detail.

We also have increased performance demands-- so more building parts that need to do more things-- and have demands on environmental performance, structural performance, cost, assembly, simplicity, life cycle, and so on. And often these things are at odds with each other and they force certain levels of customization, especially when considered around complex form.

This typically results in industry segregation. So we have lots of consultancies popping up. In this case, the diagram shows case and the kind of workflow surrounding Louisiana's Sports Hall of Fame, which I'll talk a little bit more about.

But the more parts, the more people have to be involved. The different types of fabrication, the more fabricators have to be involved. Ultimately someone has to control this model. Typically that's not the architecture firm. It has to be a consultant.

All this has to do with additive manufacturing in a way that is pretty intuitive in that the additive manufacturing workflow is typically centralized through some service provider. Geometry goes in, repaired geometry comes back, or a part comes back out.

But that doesn't really work yet in the AEC industry because we don't necessarily understand why or how we would apply these things.

So typical scenarios of the use of additive manufacturing in an actual part production or industrial manufacturing operations is used to shorten part development time-- that's typically through prototyping, rapid prototyping. Decrease manufacturing time, also by reducing tooling

cost and so on.

Part complexity can be increased. Manufacture complex geometries-- so undercuts that you can typically mold to [INAUDIBLE] and 3D print. Multifunctional products-- this is multi-material printing or gradients or different parts, embedded, sensors, electronics, and so on. And ultimately reducing or eliminating tooling.

So it opens up materials and processes that are often for high volume production for low volume production and it allows a certain level of customization. So this is all kind of proven within the industry, but not in AEC.

So adding in manufacturing typically involves the design has everything-- file transfer from a number of different file types to a service provider. And sometimes the service provider repairs the files, or there's a file repair check that then goes to the service provider.

Within the system we have a build material, typically a support material-- depending on the process, a binder. Secondary processing in terms of machining, drilling, sanding, curing. This is all tied into the AM machine or the AM process. Post processing, which is finishing fit. Support material removal, and so on. And ultimately part delivery.

Different typologies within the industry exists. There are service providers. There are in-home offices that own printers that print in house, and so on. And there's different manufactures that work on different levels of functional components.

In 2012-- and this is [INAUDIBLE] report. This is where the data comes from. They produce an annual report on the kind of current state every year of additive manufacturing and the industry. And in 2012, around 30% of additive manufacturing technology was used for functional parts and tooling.

So we're getting to a level now where production of real things and not just prototyping is shown in industry, not yet in buildings.

Some case studies I read through quickly. There's a prosthetic device with a comparison of the normal fabrication technique, which is milling or machining out of a solid block, and then printing through a metal printing process.

So you see lead time shortened more than half. Production volume-- say, this is a low volume production, so depending on the parts-- there are \$250 to \$1,500 around 8-40 pieces. And

then for 3D printing there is a change as you increase part count, surprisingly. But around same part count, you have a cost down almost tenfold.

Lead time is reduced, cost is reduced, and ultimately that's because of material efficiency. So instead of machining out these kind of typologically optimized structural pieces, the printing allows you to do this without the machining-- ultimately providing this prosthetic device for lower cost in shorter time.

Kind of one of the only examples of true mass customization-- in buildings we talk about mass customization a lot. But very few things in buildings are actually mass produced at the level that we would be able to customize.

So outlet covers, carpets, finishes, tiles-- these things are mass produced. Things we talk about as mass customization are really low volume-- low to mid volume.

But this product, Invisalign's a kind of orthodontic appliance. Compared to braces, which are kind of tentatively molded in place. They're adjusted incrementally at the dentist or at the orthodontist. All standard parts assembled in a custom configuration based on the patient.

Invisalign scans a mold that's produced from the patient's mouth, provides a digital model of every stage of the adjustment. And then every month or periodically you get an updated incremental replacement for the mold.

So this is a vacuum form part formed over a printed component. And they're producing up to about 30,000 of these a day. So it's a very high volume, high custom process. Everything that they make is different. Everything they make is high volume.

Optimization of constructional performance, and again, weight. So in aerospace industry the more weight you reduce the more fuel you save. So there's a push for typologically optimized parts through printing. This is kind of one of the big things that we see a lot of.

Optimization of the part of the milled part. This is moderately optimized. It's better than a solid block, but not quite as good as this typologically optimized one. Batch size for both are high. So it's high volume production for both parts.

If we consider the current hinge as the base weight, we lose about 10 kilograms per plane for the printed part-- basically through this material elimination, through the technological optimization process. And 10 kilograms per plane results in a life cycle reduction in fuel

consumption over the life of the plane that's substantial.

Architecture, we deal with the same things. We need to use materials more efficiently. We have increased performance demands. We may not care about grams or microgram reduction of weight of a building, but we need complex geometry. We need multi-performance parts. And much of the same thing that's driving AM and other industries.

A few examples of 3D printed parts on buildings-- all of them are the first. No matter where they are in history, all of these are first ever uses of 3D printing in buildings according to the literature.

This is Zaha Hadid Project, where within the project there's 100,000 pieces of drywall that are curved by hand, plastered over, and kind of form is built up manually by a craft process. 4 parts that couldn't be done that way are across the Zaragoza bridge. There's details here, crazy intersections here, and other form where ultimately all these parts have to come together.

So two different processes were used to produce these four parts. One is an SLA print that was then plastered over, applied to the building. Others are 3D printing molds, like cast parts. But ultimately these are finished pieces. So at the very kind of acute angles or too complex for bent drywall, these things come in to fill the gap.

So this is thought of as an afterthought in my mind. It's applying AM to solve a problem-- a very acute problem-- within the building that just results of kind of the design itself. The design's not necessarily considering 3D printing as a construction methodology. So this is, I think, direct technology transfer-- applying the technology exists to the building.

The same with this project. This is kind of an atrium space on the top of the building, and has a structural kind of connection of all these details. And it needs a cladding. So the aesthetic that was desired was this kind of organic part. And this is a 3D printed plastic part through SLS. It's hundreds of different facets.

If it was to be welded as one part, ground down and then finished, or made by some other method throughout the building hundreds of facets. For this there are 96 pieces that come together in this shell and kind of clamp around the joint. But again it's a finish part of a cladding piece.

So what this means is that we're thinking about 3D printing in buildings kind of down here at

the DFC-- Design for Construction and Construction process-- which is ultimately where direct technology transfer can happen. So we need to make something. This existing process allows us to make it. And so we make it to solve a problem within the project.

But the real potential is to look at earlier in the project in the pre design and schematic design processes what opportunities exist. And consider kind of free form construction or additive manufacturing at this earlier stage in the design. Both are opportunities and both present significant potential.

We don't yet really understand what this one is. This area we're starting to kind of get the picture where we can print parts.

So in the computational design workflow for this Trey Trahan project, the Louisiana Sports Hall of Fame, which was a case study that kind of was supported by CASE and Steve Sanderson of CASE and Trahan architects as part of this additive manufacturing framework development.

We have consultants and relationships all built around this complex geometric interior cladding. These are all stone and cement panels. You see detailing where it embedded lights. So it's a multi-functional assembly.

There's a contractual obligation to a very small surface deviation. So there's the designed surface and there's an installed surface, and those things have to be very, very minimal tolerances.

This resulted in a conceptual design model using other software transferred into a surface design using yet another CATIA-based platform that ultimately ended up as back and forth between multiple softwares, multiple platforms, several different associative modeling tools.

Those files get transferred and ultimately kind of dumbed down for fabrication. Sent to the fabricator. Then the fabricator actually rebuilds the model for the mold for printing. So kind of an insane workflow to create one part. And when you have 1,150 of these parts with almost no tolerances this workflow becomes problematic.

So throughout the production of these pieces, there's around 30% waste or scrap. That ultimately comes from these molds where the mold release falls into the most concave areas and you end up with flat spots. This is something of a process reality that-- an industrial process we understand that.

But as these low volume job shops, we're not really working at industrial mold making techniques. 1,150 molds. Lots of the molds were scrapped. Another important piece is that the molds were used to transport the part to site.

So this is another thing that we have to consider is the transportation of the complex geometry to site. And in this case, it was the molds.

This Reiser Umemoto project, the O-14 towers and other geometric. It's not really that complex in that it's just an extrusion laid by slip forming. But as all these apertures that are put in using wire cut foam parts.

So you can see this drawing. The rebar barely shows up, but there's a really complex rebar pattern around these foam elements. And then the slip form would have to come up and cast around it.

The structural engineer specified the rebar density to be less at the top, greater in the bottom. But the reality of fabrication is that the pattern's exactly the same for every lift because there's no way to organize this amount of rebar onsite.

So one of the questions is would this accommodate 3D printing or additive manufacturing in buildings. And there's some concrete printing techniques that theoretically claim to be able to do this kind of a vertical extrusion.

Then we get a projects like this James Carpenter project, the light well in the Fulton Street train terminal, which is highly successful architectural detail. But how's this assembly? Thousands and thousands of parts, and ultimately over budget.

So this is the drawing from TriPyramid Structures that made this part. This is the assembly process. So this is one node, stainless steel. All the nodes are more or less the same-- some spacing variation. But ultimately it's not a complex geometry problem. It's an assembly problem.

So if you can imagine being on site on a lift assembling this part. And once the cladding starts going into place, you have to have one person on the back, one person on the front, two-sided connection and so on.

From a product design perspective we could design parts that have snap fit, one-sided assembly to try to reduce this labor cost. And we think about that and we're making things like

mice and parts that are produced on hundreds of thousands of parts.

But this particular project is a piece of architecture-- architectural hardware. We don't really think about that assembly time onsite. [INAUDIBLE] labor cost, and really labor productivity over time, show a significant drop in construction labor productivity. And there's even an analysis per country based on how many hours of functional work you get out of a day.

So we have opportunity to really understand what the labor rate of construction is, what assembly detailing is all about. And from other research and kind of longstanding product design research on assembling times we can quantify building assembly cost relative to parts.

So fabricating architecture proposed a reduction theory of part count that starts from many parts to components or sub-assemblies. But through additive manufacturing long term we can think of completely printed product. So all of these parts printed as a single part-- which is the Ford future, the long term.

But in terms of reducing parts, the thing holds true. When this is an example of just a very basic housing-- two-sided housing-- 10 parts if it's a bolted connection, reduced to a one-sided snap fit at two parts.

So if we apply this to buildings relative to actual assembly time. And this is again is based around this Fulton Street detail for the sky net. And again I'm not an engineer, but just schematic design solutions for a snap fit just to understand what assembly time, how this would change the building.

So we have around 40,000 parts eliminated from the assembly process. It's around \$20,000 in hardware savings, resulting in \$25,000 in labor savings when we account for labor productivity. This is lots and lots of days of people not working on a ladder.

There's things not included here in terms of risk reduction-- you know, fewer days on the ladder, fewer days on the lift, less risk on the project. And ultimately we end up with a node part or a printed part that's around \$450, which is not that far out of line with the cost of the node in the first place.

So what I'm proposing is that we take the things from 3D printing or additive manufacturing in industry, apply them to construction but at a little bit of a different conceptual bent to actually get at this first part-- the pre design process.

So looking at instead of just 3D printing things, we're concerned about scale, different build volumes, and so on. When you think about additive manufacturing construction, we're also talking about assembly of components-- bricks, blocks. And ultimately I'd say early buildings were 3D printed in that it's a layer and a layer and a layer of similar units-- blocks, bricks, and so on built together to form whole buildings.

And rather than taking a direct kind of theoretical transfer from industry to apply it to buildings, we should look at this-- where do we apply 3D printing early in the schematic kind of pre design phases to then influence the building design rather than solving problems at the end in terms of technology transfer.

Quick case study that's kind of testing the performance standpoint, actually looking at technology development. And a lot of this research came from a bit of a frustration that we're developing a lot of technology in AEC for 3D printing-- large format concrete printers. Every couple months, we see a new one on the blogosphere.

This project tests kind of a problem-- starts with a problem-based research around and environmental design issue. We love to make glass facades. You can imagine shading a glass facade with this curve. It's ultimately going to generate the demand for different daylighting considerations throughout the facade.

So this is a question we have. The potential for the need for high performance and the potential for custom parts, how does 3D printing or additive manufacturing enable this?

So the goal would be to create a low-cost integrated workflow-- so actually move from the design model to the fabrication process without this kind of cluster of consultants, and so on. And also developing a prototype for a fabrication system.

So 3D printing or additive manufacturing either relies on powder-based processes are kind of inherently supported as they're printing. Other processes have to build support material. But if we think of large format 3D printing for buildings, support material seems to be a catch.

We're trying to reduce material. Using a lot of material to build support material seems to be a problem. So think about flexible tooling.

This is a pen mold. It's robotically actuated to [? grade ?] the form of a [? louver. ?] So it has within it the tolerances or the variation for a curved facade scenario. And we have a post-

processing, a milling process, that comes in after.

So a lot of times we talk about resolution of 3D printing. This really crude clay printer-- it's really a robotic deposition system-- prints on a one-inch bead. So no one could ever sell a printer that has a resolution of one inch.

But three-dimensional rectification, which is standard within the ceramic industry, we mill this out to produce final parts, proof of concept prototype. And ultimately through the kind of printing and rectification process, we get geometrically custom units that then link back to this environmental performance plan and ultimately provide a potential building application through the proof of concept research.

So to run through this, that's the last slide. And that should segue into the next talk on home robotics and customization.

RICK RUNDELL: Thank you, Nathan.

[APPLAUSE]

All right. Our final presenter, except for me, is Nick Cote. Nick received his Master of Architecture from the Rhode Island School of Design in 2015, where he also studied printmaking and illustration.

His training in robotics comes from academic work and research positions concerned with robotics in art and design. He has been a design robotics researcher at Autodesk since June of 2015, when he joined my group as an intern.

And now he's working here in San Francisco with our CTO group. So please join me in welcoming Nick.

[APPLAUSE]

NICK COTE: Can you hear me? Thanks, Rick. So yeah. I'm Nick. So this summer I had the opportunity to work with these guys on some really interesting projects. But it starts in the realm of design robotics just in general.

So when a lot of people get a robot, they sort of wonder what the first thing they should do is. And architects, of course, will try to make a pavilion. These guys-- Nathan King, Center for

Design Research at Virginia Tech, Virginia Tech School of Architecture, couple of master woodworkers over at Rutabaga, and the mass design group-- put this really awesome pavilion together.

It was modeled using a visual logarithm editor. It was structurally tested. It was basically data driven design of a grid shell structure that was created with more than 1,500 totally unique wooden struts, and then about 375 unique steel nodes that connected all of them.

The design for it was set so that you could assemble it very simply using pretty minimal materials-- basically only two materials. Do it on site, ground up, and only need a few people to do it. And this was basically fabricated and assembled over a period of two weeks. It now stands in the Boston Greenway.

So all the parts, like I said, were data driven. They were able to label them and know exactly where they were in space according to the model, and then design them accordingly.

So what you're looking at here is an index that was applied to each of the wooden struts so we knew exactly where to put it when assembling it. The next piece was that we needed a tool that would hold basically the-- so maybe back up a little bit.

Each of those struts needs to be connected by a steel member now. In order to do this, we chose to do it robotically, and created a tool that would hold this assembly as it was being manufactured. First, though, we needed code to send to the robot to make that actually happen.

For most designers and architects, we're not necessarily really competent programmers at the very beginning. And in fact, this is really in the early phases of adoption in architecture schools at all. So we need, sort of, code on the rocks. And that ends up going to visual editors such as Dynamo and Grasshopper-- stuff like that.

So what you're looking at right now is the rapid code that would be generated in order to manufacture a part. The first thing is the tool data and work object data. That's describing how you're holding that part and where in space in relation to your robot the part actually is.

Then you have a series of constants that tell the machine where to go in space. So what you're looking at are coordinates followed by an orientation. So you have to tell it to go to a position and then which direction to point.

The next thing is your main routine. And then, following that in this section of code is a series of instructions that say OK, move linearly to this point. Wait here. And then move the joints of the row up to some other recent position.

So in general, programmer robot has only four parts. It's a series of-- it's a set up, it tells the robot where to go in space, a routine, and then instructions for doing so.

So what I had been working on is a workflow that would do it very simply for essentially non-programmers to work with. It takes those data types. It takes those instructions. Well, this also a data type.

But then it writes that to what's called a PRG file, which is what you were looking at here. That's the language spoken by EBB machines in particular. So if you hear of an EBB, those sort of orange or-- you'll see it in a few of the following images.

This library I created for Dynamo, then was expanded to hold a pretty large variety of other things that would allow you to communicate directly with machine, support a variety of other data types-- such as circular movements and work zones for safety-- a variety of other instructions.

But the key here is that these sort of things enable a variety of motions to be described using the same interface that the geometry itself for the pavilion could have been designed in. So it's closing the gap between, basically, model space and then robot space. So, design space and manufacturing space.

So what would happen is that all the data that goes into the pavilion would then be sent to a CSV file. So we're reading all of the angles for the nodes, passing it through a workflow just like this, and creating points in space that the robot could potentially go to.

We'll skip over this one. So what we are creating is nodes-- really endless, endless ranks of nodes. And this took a really long time to put together. But each of them was welded to a steel cylinder, and there are four flanges welded to that cylinder.

So what happens is that each of those flanges will then go inserted into one of those wooden struts on the main structure. And these are each totally unique. So when I said that each of the wooden struts had their own length and measurements, each of these also is-- there's no duplicate any of them.

So this speaks a lot to the ability for a robot to basically be any tool you need it to be in the given situation. But for mass customization and mass production in this sort of sense, it was the ideal tool to produce this.

So like I was saying, all that data is now generated in the model. It goes into a file that can be read by one of these tools. That file is now read by the machines. So what you're seeing here are the planes in space that we'd like to send the robot to.

The next thing that I was into is where it's being read for the CSV. And then the actual file productions. So what you're seeing here is the name of the file, the node name. So we've got node one-- it's produced here.

And what's interesting is using one of these sliders in this program so you can automate the production of these to just create all-- for our project-- 375 nodes in a matter of seconds.

All those points then get processed through basically something that changes the syntax to robot code. Then we find our machine on the network, and send that code directly to it.

As you'll see in a second, this is a simulation. But it's good for testing the workflow before it happens. But really simply, it pops up and down, twists the assembly so that it clears slag and a person can weld it.

So this is actually what it looks like. That's the node. The person welds. And then it will raise the assembly, rotate it so the next flange can get in place, bring it back down. And then the next flange will be welded onto it.

During a few of the sort of screwups during this project, we had to reweld things. And it turned out that by hand measuring, welding, and producing one of these took about 35, 40 minutes. But doing it with the machine and a person working together, it was only about 3.

OK. So the next step and sort of an obvious thing is that in between the actual production of the pavilion and the robot itself you need good operators. I helped to train some of these people just on using the robot. We had a really, really well-experienced welder teaching basically novices to operate one of these welding machines as well.

So we had people that didn't necessarily have expertise in construction or robotics assembling what is a very complex computationally designed pavilion through a very complex computationally designed process. And it all turned out OK.

What you can see is the robots are mounted beneath the table and that we've got basically a work [? cell ?] on either side of the table. An operator working with a welder on either side made for some pretty awesome shots.

So this now is on the Greenway in Boston. I guess it's heading for some site on the BUILD Space later on, but you can also see a section of it downstairs. That's it. Thanks.

[APPLAUSE]

RICK RUNDELL: All right. Nick, thank you.

So this final project that Nick presented was the first BUILD Grant recipient. So we provided some funding to this group to do work that would be of the type of work that we will eventually be able to support in the BUILD Space when that's finally completed. But we're getting kind of a jump start on the work that we're doing there by funding these research projects.

So I'm going to take just a moment and talk about the project that we're creating that will support the kind of really terrific research that you've seen each of these four presenters present this morning.

The BUILD Space is a physical location. I'll show you some plans here in a minute. But it's intended to facilitate collaboration between industry, academia, and practice, and design, and engineering. So it's targeting the building industry.

We have other facilities at Autodesk that are focused more on manufacturing that some of you may be familiar with in San Francisco and a few others around the company. This one will be particularly focused on the building industry.

So let's take a look at the-- oh, I've got a slide changer. All right. You want to tell me which button actually-- Is it the one in the-- that one. Oh, OK. There we go. All right.

So is anybody here from Boston? Familiar with Boston? OK, terrific. All right. So this is an area of Boston that's been rapidly developing called the Innovation District. A lot of start up businesses are here, as well as some more established businesses are moving in-- most recently Vertex Pharmaceuticals. It's also called the Seaport District depending on who you're talking to.

And we are taking some space in a building here that's a third of a mile long-- it's an old Army

Depot-- for both our offices in Boston. We're moving our offices from [? Waltham. ?] And for this project that I'm talking about today, which is the BUILD Space.

So the building's a very interesting building. It's a third of a mile long. It has a loading dock the length of the building. And it used to function to unload trains. So you'd pull a train up to this building and they would just unload the train into the building. And then eventually the material would end up on ships on the other side. And this was for the Army.

It was acquired a couple of years ago by a developer who's developing it the tenant community around making and design. So it's a terrific fit for us. So you'll see some of the tenants that are in there already.

Our space for the BUILD Space is on the first and second floor opening onto this pedestrian promenade-- that's what the loading dock is turning into. And then the floor plate on the second and third floor eventually.

So here's a quick diagram. This is number 23. This will be our entrance. Offices upstairs on the sixth floor. If any of you visit us in Boston, that's where you'll go find my colleagues. And then on the first two floors we have the workshop that we're developing.

This is what the space looks like now. It's a 21-foot square column grid, which those of you who've done architecture can appreciate how that helps us organize the space. On a good day anyway, that's how we look at it.

It's a very heavy industrial structure. The space itself-- is that the laser? Oh yeah. Great. So the space itself is organized on two levels with heavy equipment on the first floor. Our goal is to be able to do work with any material-- pretty much any material that you would find in a building.

So we have glass, metal, steel. We have a glass studio here. Composite, panel, layup capability, concrete, and ceramics in the back. We have a 5-Axis CNC router, a water jet, and a heavy metal working area. And we're set up to handle panel sizes up to five feet by 10 feet.

And then the heart of the space is just a sort of open fabrication space. So this could be fit out with any kind of equipment that anyone wants to use to build furniture-sized or building panel-sized objects.

The second floor is an open studio space with 50 workstations that can be assigned to

researchers or members of the community or people who want to come in and use the space.

We have schools that are interested in teaching classes out of here. We have start ups that will be housed here. They're doing work in the AEC industry. Maybe even some of the work that you saw here today will be housed there. And then around the perimeter of this is wood shop, metal shop, 3D printing, laser cutting, and a microelectronics area.

The space is also configured so that we can move everything to the side-- everything's on wheels on both floors-- move everything to the side and host an event. And so we have a lot of interest from members of the construction industry in hosting events and providing experiences for their employees in digital fabrication, for example, using some of the equipment that's there.

I invite you to follow us on Twitter, or Facebook, or Flickr. Actually I'm told that our Facebook presence is actually the best one. So I think that's where all the photos are. It's under construction.

It's planned to be opened in May. We're moving the offices in January, but we'll be opening the BUILD Space in May. And I hope that some of you will be able to join us for the festivities at that time.

And that's all I have. So what are we doing on time? 15 minutes? OK. So maybe if anyone has some questions for any of the presenters, we have a few minutes here at the end of the presentation. So I'd invite you to if you had anything that you heard you have a question about, this would be a good time to ask it.

All right. Well, in that case-- did I see a hand trying to go back up? OK. Here. Use the microphone.

AUDIENCE:

This was a question that I had when I was watching your presentation, Andy. A thought that came to my mind was all the personalized comfort zones. We are already sort of doing it in our cars.

And that's why we love our cars more than our office space or even home for some people because we can really do all that customization. This side, left, passenger side a different temperature.

So I'm wondering in stuff making these components sitting on your desk, can they be

integrated into the architecture itself? Because again, it looks like another step involved. Can this be done through design?

ANDREW PAYNE: I think it's a good question. To be honest with you, at least from an office scale, we're seeing a lot of disruption from the residential market-- obviously, in all these learning systems, Nest has obviously been quite successful.

We haven't seen quite the disruption in the larger enterprise level office spaces and so on. There just really hasn't been quite the amount of innovation, I think, that we should be seeing in that space.

To be honest with you, I think, to have the greatest impact would be partnering with someone like a furniture maker like [? Nohl ?] or Steelcase, which outfits a lot of these cubicles and can do it at scale and has the ability to adapt faster.

I think what's one of the biggest questions is how do you begin to deploy this and not sort of have, say, I have to start with a completely new ground up construction that has a completely new system. We ought to be able to retrofit existing buildings but make them smarter.

And I think if you were to actually apply this and try to get some penetration, the way I see it, would be actually partnering with somebody that has the ability to actually roll this out rather quickly.

I haven't actually been in discussion with anybody on this. But I think the best potential in actually impacting the market would be through partnering with furniture makers, honestly.

AUDIENCE: I have a contact if you want to go that route.

[LAUGHTER]

ANDREW PAYNE: Fair enough.

RICK RUNDELL: All right. Any other questions from the audience for any of our presenters? All right. Well then, we'll all get maybe 10 minutes back in our day. I want to thank one more time these four awesome presenters.

I also want a special thanks to Nathan King, who actually pulled a lot of the mechanics of getting this presentation together. So thank you, Nathan. And thank all of you for a terrific presentation, and all of you for joining us. Enjoy the rest of your AU.

