

PRESENTER: So what are we trying to do when we hit render? Rendering is an art, and some would argue that rendering is a religion, and once I learn a renderer and know how to get rendering stuff out of the render, it becomes even more of a religion. But what we're really trying to do is recreate this art of what happens in nature.

So how does it really work? First, it's very rudimentary, there's a thing and this is you, this is your eyeball, and all vision relies on one thing, light. Without light, we would not be able to see. So the rays of light come down from its source, bounce off an object and onto our eyeball, and then it's captured on the back of the eyeball, and then pumped to the brain and the brain creates an image out of it. And really that's not all different from what a rendering is trying to do, it's just that we have, depending on who you talk to, 6,000 years of evolution or 44 billion years of evolution on how the eyeball works, and renders have about 25 years worth of calculations and how they work, and obviously the limitations of hardware.

Renderers do what all software programs do, they take data and they apply math, and I think we forget that sometimes, it is that it's just all math, when you open up Outlook, it's math, when you open up V-Ray, it's math. So what we're trying to do is make the math as easy for the render to work as possible.

So as a renderer, what you're trying to do is find the lowest acceptable results. Because if you had picture perfect every time, it would take hours and, in fact, my friends at Pixar take days to render one frame of their movie. So in Toy Story 3, when they're walking the first time through the day care center, each frame of that took seven days to render, and each frame was one 24th of a second. Their acceptable level of rendering took seven days.

The question you have to ask yourself is, for this image, what is my acceptable level of rendering? Is it the resolution? Is it the amount of reflections? Is it the lighting? All of this stuff is a part of that math. So when we reflect an object one off of another, it has to calculate first what the object looks like, and then calculate what it looks like in the reflection and then calculate so that's three calculations. If you do the Hall of Mirrors, you're talking hundreds if not thousands of reflections, each that need to be calculated. Does that make sense?

So the first thing we're going to talk about is the way renderers work. One is through rasterization and the other is through raytracing, and actually, to be perfectly technical with

you, to make this even more boring, some renders use both at the same time. But let's take a look at the differences. Rasterization starts at the object and I was trying very hard to think of an analogy that would kind of fit how this is.

Think of you and your best friend and a piece of glass between you, and you are an artist and you would like to draw them on the piece of glass. Now, there are two different ways that we could do that, either A, they could sit with a laser pointer and shoot lasers at the glass and wherever the laser went you could trace that from them, or you could sit at your desk and shoot lasers at, no, that wouldn't be good, you get the idea, and then where the laser went it would actually make a mark on the glass. So rays are projected from the object to the camera, or rays are projected from the camera to the object, this is all important because it helps you understand when you choose your renderer.

Here's rasterization. Here's our camera, here's the thing, and that grid pattern that's really dim is the picture we're creating, and from the object raise our cast towards the camera, where they strike the image they put a little dot, they fill in the rest and then that goes to the camera. And the reason that this is really fast is because you'll see that if there is no object out here, no ray. So it really focuses in on the objects that are being calculated.

Here's ray tracing. Here's our object, here's our camera, and that really light grid is where we want, and from there, the rays are shot out to the object and when it hits one, it transfers that on to the image. Make sense? And then it's all calculated inside. Questions? Because as we get into the actual renderers, we're going to start to talk about, this is a ray tracer, this is a rasterizer, so that will be first thing.

In 1974, at the University of Utah, that was where most of rendering technology was born, and one of the dudes that was a part of that, was James Blinn. James Blinn went on to work for NASA and shot rockets, so it took that kind of a brain to conceive all this stuff. But he recognized that as technology advances, rendering times will remain constant, and why is that? Any guesses?

AUDIENCE: Dead bodies.

PRESENTER: Actually, if it were just for deadlines, we'd want to shorten it, right? But what actually happens is, because we are human, for every advancement that we do in speed, we fill it with something else. For example, the very first renderings were in the 72 x 40 pixel range, took 24 minutes, this took one second on my laptop to do. So what do we do? Well, let's make it 720 x

405, that's HD dimensions for a 720 screen at Best Buy. You know what? We walked past those, now those are on the back shelf, and we look for 1080p, and the next one will be 4000p, and then the next one, they're already developing, 8000p. What used to be 72 pixels good enough, we'll now need 8000 pixels good enough.

Speaking of the University of Utah, a lot of people ask me, what's up with the Teapot Mr. Blinn with his coworker Newell they needed an object to render, and so Mrs. Newell had this teapot, and they thought it was very compelling because it had curves for reflection, it had saddle's for where things come together, it had objects that cast shadows onto itself, so another guy that was working with them was Ed Catmull, does that name sound familiar at all? He's the guy that started Pixar, and so when they needed a renderer they created RenderMan and so he used the teapot. And so from there, we started to all identify rendering, so it's this kind of inside joke. It's not that we think that 3ds Max users are making that many tea pots, thank god you've got the teapot button.

But it really does allow you to, in 3-D space, you can tell if you took a box and rotate it around, what's the top, what's the bottom, if you take a sphere of any kind of a normal primitive, you can see that this is immediately identifiable.

So another thing we have to determine is, are we creating photo realistic or non-photorealistic renderings? So here's our little photorealistic guy, and it looks like he could be sitting on our desk, and here's that exact same model now rendered non-photorealistically realistically. And there are some renders that are good at one and not so good at the other, and vice versa. Well, I do design, so I'm probably going to want a photorealistic.

When I was first starting in this business, one of the first jobs that I did, I had to, when the lawyer calls and they take a supposition, I took the brick pattern that was a part of Autodesk Biz and I put it on a wall and I hit render, and then he said, doesn't that look cool, and I said, that looks cool. And a few years later, the owner said, this is a brick pattern that we wanted, and this is the brick pattern you built, we are suing you. So, very often, you can get lost in the materials without seeing the design, and non-photorealistic allows you to present your ideas without those kinds of.

So are you're going to want to do non-photorealistic rendering?

Now that we understand some of the basics, when we talk about how these are all built, there is a number of terms that I want to introduce you to. Shaders, materials and maps. Very often

these are often confused, especially by those that are looking to get into rendering.

What is a shader? A shader is a number of instructions that determine how something is affected, and while all materials are shaders, not all shaders are materials. Ooh, heavy. So here is a car paint shader, and you'll see that the shader for car paint has a lot to do with how I make car paints like flakes, and then here is an ocean shader, and it talks about waves, and you wouldn't apply an ocean shader to a car, and you wouldn't apply a-- well, you might want to play a car paint shader to an ocean. But you get the idea. Shaders can be applied to lights and cameras and other objects. How do we want to affect how they render, where materials, on the other hand, are really based on how an object looks.

There are two things that we worry about when we talk about a material, what is its surface and what is the light bouncing off of it, remember how the eye works? Those are the things that we focus on for materials. Here you can see an object growing from no materials to very shiny and reflective materials.

At this moment, Box is letting me know that he can connect the internet, thank you so much. Is that technical glitch number one? We're not sure.

The, how are materials built? How do we make it shiny? How do we make it not-so-shiny? Remember all that stuff with the eyeball and back in Utah in 1974? They're the ones that took all the science of light and made maps out of it. So here is our eyeball and how an object bounces light, that's its diffuse color, and how it reflects that's its glossiness, and how it transmits stuff through it that's its translucency, and you can see that in the world of rendering all of that stuff still comes through.

Here are all the different maps that you can apply to change the way an object looks in 3ds Max, and guess what? Those are enough. I thought maybe I'd show you. Here are some maps that are applied to an object, this is glossy, so it's very easy to see the difference. Here's glossy. Then you can start to do science computery stuff, here I'm changing the surface of this same object by applying what's called a bump map, and a bump map takes a picture of an object and where it's black it pushes it in and where it's white, it pulls it out, and where it's gray, it leaves it alone. I didn't add any render faces to this, I just changed its surface.

Here's an object with no material, it's more than just a red apple with a bite out of it. It really needs to have a number of different layers in order to make it look like an apple. I wanted to

have fun comparing apples and oranges. This one has those two different pictures applied to the surface, one for bumpiness, one for just the speckled effect, I added translucency and glossiness and you can see, if I'm moving on from teapots to apples and oranges I need all of those things.

So we started off with the whole conversation, how do we see, and that's really where this part comes into. When we talk about rendering, we're really talking about illuminating our objects. So how does a renderer do that? And the first thing it does is it takes from our light source and it throws a beam of light onto an object. Now, this picture you see here, again, goes back to 1970 something, it's called the Cornell Box and why it's used in rendering is because it does a really good job of showing reflections and color bleed and light bouncing around in a space. So I'm going to use that to show you all the different kinds. So this is just light source onto an object, and you'll notice that the ceiling of this box isn't illuminated, there's no light that shines onto it, so what we need there is indirect illumination. The light that strikes an object and then bounces off and hits another one. And when you start to talk about different renderers it's death, how it bounces light around, that really determines what is the difference in the renderers. Does that make sense?

There are some that use radiosity, and what radiosity does is it first goes in and it paints the whole object with a net, and some of the areas it recognizes are very tight, so there needs to be lots of lines there, where some of the areas I can get away with not-so-dense netting, and wherever those two lines cross, when light strikes it, it paints the amount of light on that object. Let me back up and you can see the effect.

And I apologize the projector is not getting the colors and the richness of this on there.

This is called photon mapping, and photon mapping is very similar to the way light bounces around, so rather than a light particle it takes a photon and it starts to spray it and were bounces around, the photon picks up green, and I don't know if he can see but on the side of the sphere there is some light colored bleeding on the side. Another type is final gather, this is what Mental Ray and other softwares use, and what it does is, it says for this pixel right here, what's illuminating that and where could the potential illuminations come from? So it looks at this, it sees direct light from here, but it also sees reflected light from there, and reflected light from there, and that kind of thing, so that's the way final gather works.

Hang with me.

This is irradiance mapping, and that's almost exactly the same as final gather, and then this is light caching. You can see that this one has kind of a bumpy feel to it. And then there's this one called brute force, and brute force is the easiest to understand, that one is just trying to reproduce nature.

The only difference between our renderer and our eyeball is we're receiving billions of light rays. And in order to calculate the billions of light rays that we're actually seeing, it would take us billions of years to calculate all that. So again, I'm trying to find, what would be the number of rays that I can bounce around in here and still get acceptable results? This is 600 rays. From the camera, remember radiosity, 600 rays come out, and they hit 600 places, and what happens to each of those? They bounce off and they hit something else, another 600 rays, for every one and then bounce off and there's another 600 rays in there. So if I say six or seven levels of light calculations, it's got to be a bigillion, right?

Everybody got the light illumination? Yes, we memorized it, OK, good.

There are two types of renderer. One is called a progressive renderer, a progressive renderer will show you everything and continue to refine it, so you can see that it starts off really yucky and then it's getting better and better and better. Let's start over. So it starts off, this is one second, and then it keeps refining. So that's a progressive renderer. Makes sense? So all of the calculations are working on all of that at the same time.

This is called bucket rendering. Bucket rendering starts off and it first goes through and it calculates, and each one of those little squares that you see is kind of where it's focusing its information. And you can see that there are a number of squares going at the same time and this is a little laptop, so it didn't have a whole lot of them, but you can see that there was like 8 or 9 buckets going at any one time, and it just so happens that every one of those buckets is a CPU on my computer. So this CPU is working on that square and the CPU is working on that square, and all my CPU's go, bing! At 100%. Does that make sense?

The last thing you're going to really think about when you try purchasing a renderer or committing to a renderer for a project, is whether it is biased or unbiased. Unbiased is the easiest to understand, there's no controls. It has to calculate every pixel with a ray. So I can't say, you know what, this is blue over here and this is blue over here, just draw the same blue between, that's what the bias is for.

Biased allows me to go ahead and make some adjustments to the way it's all processed for

speed, where unbiased, I can't intervene, so it has to calculate all that stuff out. Of course, unbiased are the most accurate and biased are not. Questions?

One other thing you'll think about or hear about when you go down to the show floor and talk to the rendering guys is, they'll tell you they're a CPU or a GPU renderer. So the CPU is the computer chip, that's the heart on the motherboard, that's what does all the calculations. For ever, the CPU was the thing that did the rendering, but when people ask me, way back then, what kind of video card do I need for my VIS4? I would say, it doesn't matter, the only thing that the GPU does is it shows you the image on your screen.

Well, then they started to figure out that everything I need to show you an image on the screen is exactly the same thing I need to show you an image in a picture. So if I optimize some of my rendering algorithms to the same thing that the GPU does, I can do faster rendering using those, then I can do those. Does that make sense? And the fact that most computers have about 8, I think box technologies down on the floor has one that has 22 CPU's in it. It costs as much as my hou-- no. Whereas I can get a halfway decent video card with 128 GPU's on it.

Now, unfortunately, there's always a good and a bad, or a trade off. For this, I can put in 120 gigs of memory. This one I think currently tops out at 32. I know that some manufacturers have said they're coming out with a terabyte of memory on this, but you talk about your models and textures and all of this stuff that's in a 3ds Max file, it all has to be put on the GPU before it can start to render. So if you have 1,000 tea-pots and each one is really dense, you could fill that all up very quickly.

Any other questions? Before I get to you, the only other thing, one of the things I get asked a lot about is there are manufacturers that do like, let's take NVIDIA, for example, they do a Quadro card and a G-Force card, and we know that the components, the GPU's are virtually the same on both. Now, one is like \$3,000 and one is like \$800. And you're like, well that's simple math, I don't even need a coupon. Why would I buy a \$6,000 card? And it's because it's not the GPU's, it's the way they process. So these G-Force cards are made for gaming, and for those of you who have ever played a video game, you know it's kind of like really intense, and then you get through the level, and then the action kind of calms down. And at that point, the G-Force card breathes for a minute, and it's built to breathe. When you hit the render button, all the GPU's go, OK, I'm rendering, and it could take hours, and then it goes to the next frame. It's built for long-term heavy usage, and GPU rendering doesn't get a break to

breathe until the whole job is done. So I have witness gaming cards smoke when you render out big animations on them. Just be ready to swap out lots. So they are priced for their function. Does that make sense?

Questions?

AUDIENCE: So, are they as powerful? Is it one-to-one? Does one core of the CPU equal one core of the GPU?

PRESENTER: Unfortunately no. No, it doesn't. The question is, is it one-to-one, where one core equals like 10 of these, or if I have 60 of these, does that equal to 4 CPU? And, unfortunately no, because if I am a rendering technology and I write to the GPU, my algorithms are built for that. And if I'm a rendering technology that's built for this.

A couple of things about these two things, very quickly, the benefits, I can very easily add another video card to my computer. I would have to change out the motherboard to add more CPUs, but a lot of CPU rendering is easily added to other computers. So I can take an image and separate it all out to a number of computers through what's called network rendering.

Now, for those that aren't familiar with network rendering, network rendering allows you to take your computer, this is you, and these are all the computers that are laying around the office, for whatever reason. So what I'd like to do is, take the CPU's in my computer and add the CPU's in this computer and that computer, and that computer, and I've even gone ahead and bought a rack mounted rendering computer. Ooh. I need a managing software to do that, and that's the reason that 3ds Max ships with BackBurner. And what happens is that I go to each of these computers and I start, whether or not it's going to be a manager or a server each of those different computers. Makes sense?

Now, BackBurner doesn't render, it's just the managing software, and if you go up to a higher end one like Deadline or Cube, or any of the other managers, it's the same kind of thing. Basically, what happens is, I hit the render button, let me show you. Here in the Rendering dialog box, it says Submit or Network Rendering and it opens up this dialog box, and then you can see here's all the servers that are available. So every computer that I've started that service on, that server program they now appear in this list. And I say OK, send it to all of those.

And what happens is, the manager says "all right kids, let's start rendering". In each of these

computers, it starts up 3ds Max, it sends all the files that are needed for that scene, it loads them up on the max, and it starts to render on this computer. Frame number one, and then it starts on that one, frame number two, and frame number three, and four. But that one only has four cores on it, this one has 22, our different machines are rendering at different speeds, so this one might get done after 30 seconds, the little laptop is taking a minute to render each different image.

They're all collected back together again, we can see how they're doing on any computer on the network with the monitor, I put that one up over there. And we can see how the whole thing is working. At the end, we get individual frames for the animation, and then we used another software to bring them all together, or the RAM player and save them as an AVI, or an MOV, or an MP4, however we want to save that animation. Does that clarify how network rendering works?

The good news is, there are a number of renderers that come with 3ds Max that didn't cost me a thing, so no additional cost to get all of those computers working on my animation.

AUDIENCE: [INAUDIBLE]

PRESENTER: The good news is, all of them do. Rather than submitting a network rendering under the render pull down, I can actually separate out, so rather than frame-by-frame, it'll take that one big image and say, you do the bottom 10%, you do the next 10%, you do the next 10%, and it brings them all together into a picture.

A lot of people when you talk about cloud processing, the very first thing, the very basic application is, well, that will work out great for rendering. And the actual mechanics of cloud rendering are identical to what we just saw with computers inside of our office, the only problem is that it's not supported under the end-user license agreement with Autodesk. So all of your Autodesk software needs to be on a computer that you can walk over to and turn off. All of it. AutoCAD, Revit, all of that stuff.

We're working on trying to figure out not only the mechanics, but the legal aspects of cloud rendering. The good news is, if you look up on labs, we actually have a 3ds Max cloud renderer that you can use. Scanline, Art, Arnold, all for free just to try it out, because it's our job to get the mechanics. The legal department will figure out the license, but we got the mechanics. So if you want to try cloud rendering it's there for you.

That was a lot, but it is the foundation for our next discussion. Are there any questions about GPU's, or CPU's, or biased, or unbiased, or photorealistic or non-photorealistic, or maps, or any of that stuff? Yes, sir.

AUDIENCE: That cloud-rendering you were talking about, are there concerns of certain locations that it has access to for those mats?

PRESENTER: So the question is, when you cloud render, or when you network render, are there concerns about where are you located? With the cloud rendering, that's one of the mechanics that we are trying to figure out, so when you say cloud render, it has to send all of that to another location. When you network render, all of your maps can either be sent with, and the IT department loves that, because it really fries out their routers quickly, or you can centralize them all to a specific location. Does that answer the question?

AUDIENCE: When you are saying [? top runner, ?] you're also including render [INAUDIBLE] 3rd party and render farms like REBUS and things of that nature.

PRESENTER: So the question is, for cloud rendering, does that include rendering farm services, like REBUS and others? And the answer is no. REBUS takes network rendering as you would in your office and creates it in their office. So they have all the computers and hardware, you just send them the file, they have 3ds Max, all that, and they do the management of creating a network, and you get back images. Cloud rendering, my definition is, you set up how many computers you want to work on with Google or Amazon or those kinds of places, and you send it to yours and you pay for the processing power back. Makes sense? Anything else?

So the question about plug-ins for cloud rendering, that's the big mechanics. There are some softwares that don't use plug-in, 3ds Max was built on an open-architecture for plug-ins. So if I bought teapots, you know what I mean, super teapot plug-in, that creates a teapot that actually has like a lid, OK, I'm wearing out my teapot analogy, but they don't have that same plug-in at REBUS or on the cloud, then they can't recreate that teapot or forest, or railing, or some of these rendering technologies are all third-party plug-ins. Just because I send it to Amazon doesn't mean it will be able to render if it doesn't also have this render that goes along with it. Sir

AUDIENCE: So, downstairs I had conversation with someone in NVIDIA, we were talking about the GPU's that they're doing. It sounds like programs like Mental Ray are moving in that direction. Do we need to be concerned about the future of CPU use, also? Or should we really be looking at

GPU's in the future and making that switch and investment in terms of hardware within the office?

PRESENTER:

So the question is, obviously rendering technologies are taking more and more advantage of GPU's, let me see if I paraphrased this properly, in the future, should we be looking to GPU use more or should we still? The answer to that question is yes and no, and I apologize that I wish this class was really definitive, you know what I mean? But it really does determine what is important to you. So if I pick a renderer that is GPU-focused and we will see which ones are written specifically for CUDA core technology or open CL technology, or that kind of thing, then that means that every one of my computers on my farm have to have that technology? It doesn't matter what CPU's currently, as long as it's a CPU that processes windows, it works, AMD and Intel and all of those kinds of things. So there is a agnosticism, is that a word? So, there's agnosticism to the CPU, and like I said, I can use the laptop on the receptionists' desks for rendering at lunch and at night, whereas to put a \$6,000 video card makes it a difficult investment for some companies to make. I really think that as soon as we can get the technology and legal stuff done with all of that, this will all become a moot point and you will just send it up to your private cloud. That would be my vision of the future.

One of our clients in Los Angeles recently relocated very quickly to a new location because their office building that they were rendering in could no longer handle the temperatures that were generated from all the rendering computers nor the power requirements. They were blowing fuses all the time when they would hit Submit a job for rendering. So they actually moved to a new location with new wiring and all that stuff and new ACS just so they could render. That's a considerable investment to a rendering pipeline. And very soon they'll just be able to go, and rather than nine computers, 50 computers, you can have 1,200 working on a rendering all at one time. Animations.

AUDIENCE:

Just talking of that GPU vs CPU also, I think some of us have had a bunch of experience in Unreal, the game industry, some of us are thinking about real-time rendering as kind of being the future of rendering. I know it's not [INAUDIBLE] with accuracy but, I'm just curious if that's something you think that should be [INAUDIBLE]

PRESENTER:

There's a chapter in the hand out book about real time rendering and I just didn't have time to, but there's definitely benefits and pros and cons to that, but the quality definitely comes at a big cost in real-time rendering. For the kind of look that you're looking for. Does that answer that question?

AUDIENCE: Yes. I work at Autodesk, by the way, because I'm 3ds Max [INAUDIBLE] I'm trying to figure out whether I can work on [INAUDIBLE] trying to figure out if real-time rendering is something you want to integrate into cloud service [INAUDIBLE]

PRESENTER: Absolutely.

PRESENTER: For the network rendering, does each server need to have a 3ds Max?

PRESENTER: So the question is, for network rendering, does each computer that were starting up a server, does it need to have 3ds Max on it? Yes. Does it take a license? No. So with one license of 3rd Max, you can put it on 999 computers. And I only say that number because I can't find anybody with more, so we say unlimited.

AUDIENCE: What are your preferences on biased or unbiased rendering?

PRESENTER: No, I have no preference.

AUDIENCE: What are the good cases to use on or the other.

PRESENTER: As I hit a biased renderer coming up, I'll say what it's good for, and then when I hit an unbiased renderer I'll say what it's good for.

Right about now we're all kind of bored, and there was a lot of technical information that hit us, so now we get into the fun part. Before we go into that, let me just poll the audience. How many of you are actually rendering today? How many of you are not rendering and looking forward to rendering in the near future? For those that are rendering, why don't you all stand up? And for those that aren't rendering, stand up too, so you don't feel weird.

So we recently did a survey and talked to everybody who renders. And here is the outcome, and you can't read it. For those of you who are using A360 rendering, please sit down. There's the dude. Released in 2017, with 3ds Max was the Autodesk Great Tracer, anybody using that? Sit down. How about IRA? IRA comes with 3ds max. Nobody? Maxwell? So these little ones are Corona. How about Mental Ray? Scanline? Guess who this is? VRay?

This slice is Mental Ray, it's about 20%. This slice is Scanline, it represents about 22%. This slice is VRay. I apologize, for some reason it worked perfectly 10 minutes ago. I just wanted to show you in a recent survey the renderers that are being used. One of these is Octane and then Arnold is the little sliver there.

So let's go ahead and take a look at the different renderers, and kind of pros-and-cons them

We start with Scanline. Scanline is a rasterizer. Remember what that meant? It projects from the object onto the computer, or onto the image. It requires radiosity in order to do indirect lighting, and it's really good at cel animation, and the reason that it still represents 22% of all the renderers is because it's still used in a lot of cartoons, especially anime and manga out of Japan. So the pros for Scanline is that it can be lightning fast, because it is a rasterizer, if there isn't stuff around the object, it just ignores it and it focuses in on the object that it's working on. There's decades of content, 25 years, this is the first renderer, you can do ink and paint for cel animation, it uses all the standard shaders and materials and lights and it comes with unlimited network rendering with BackBurner and others. If I am doing an interior shot or need to bounce a light around, first I have to do a radiosity solution. 10 years ago that was unacceptably long, now, that image that you saw in the handout, it took like a minute to do the radiosity solution, that used to take 20. Realistic results require considerable preplanning and they probably all need you to take it into Photoshop and fix it.

Any questions about Scanline? It comes with 3ds Max so there's no additional cost.

Released in 3ds Max 2012, I'm guessing, was this thing called Quicksilver. Quicksilver was our first GPU renderer, it's hardware based. As part of that, you can not only use it for realistic renderings, but you can see here, I did use it for a number of non-photorealistic renderings, as well. And then the fun thing about Quicksilver in 3ds Max is, this is a rendering, and this is the viewport, so I can actually set Quicksilver to be my viewport if I want to work in this mode. Quicksilver has lots of non-photorealistic modes, hardware based for fast rendering, the more you use it, the faster it becomes, so the very first time I send all of that stuff to the video card, it has to calculate all of the textures and model and all that. Once it's on the video card, it only looks for differences. So if you're just changing colors or changing reflections, or those kinds of things, it doesn't have to do the whole model and all the lights and everything, just those few little changes and then it starts to render again. It supports multiple material types. Again, unlimited, so all of the ones that ship with 3ds Max will be unlimited. It doesn't take advantage of all hardware configuration, so your hardware card has to use DirectX, not all the shaders are supported, and for me, personally, my computer comes to a complete screaming halt when I hit the render button using Quicksilver. I mean, not even Outlook, not even fantasy football, nothing.

AUDIENCE: Do you get multiple GPUs when you [INAUDIBLE]

PRESENTER: Yes, so there are controls in there that you can actually tell it which card you're using and that kind of thing. Any other questions about Quicksilver? Again, about 1% of you are using this. It is lightning fast.

Introduced in 3ds Max 2016, the ability to send up to our A360 cloud renderer. So you can see that it does a great job of doing very realistic elements and up on the cloud. You can also do illuminins or lightings studies and you can also do panoramic output for VR sets. Lots of good things to talk about there, again, cloud rendering, panoramic renderers, lighting analysis, and it frees up your PC while you render. One of the negatives is it's not using any of the rendering technology that is currently in 3ds Max so there is a conversion, and whenever you convert you get different results. So here's a room that I rendered in Mental Ray, changed it to cloud rendering and sent it, and you can see the differences.

AUDIENCE: Do you have to convert materials?

PRESENTER: So the question is, do you have to convert materials, and it will convert the materials, it's just that, yikes, you might not get the results you're hoping for.

AUDIENCE: Will it to the same with V-Ray?

PRESENTER: So, will it do the same with V-Ray? Some of those aren't supported.

AUDIENCE: Scene converter?

PRESENTER: The question is, what about the scene converter, in 3ds Max 2017 we included a scene converter which will convert those to your V-Ray materials to regular materials and then those will get converted to the A360 materials. But there's no A360 materials in Max.

AUDIENCE: Is there a predictable amount of rendering time?

PRESENTER: Is there a predictable amount of rendering time. And the answer is no, it really is who's using the renders at the time and those kinds of things. So a late night render probably will go quicker than other. Any other questions about A360?

AUDIENCE: Do you have to buy credits?

PRESENTER: Yep. Ding! It's like if you knew. So the question is, do I have to buy credits, and Hi-Res,

illuminants and panoramics all cost me credits. There's a number of credits that come with 3ds Max and after that I have to supplement, if I want to continue. Yes?

AUDIENCE: Is there a video [INAUDIBLE] on each different monitor when rendering on each I did a video [INAUDIBLE]

PRESENTER: So the question is, or there was a concern, you tried to do a number of different images. So one of the things that's not on here, as far as a benefit, is the A360 renderer allows me to send all of my cameras at once. And so I don't know if you checked all the cameras for your different images.

AUDIENCE: No, they were just images, [INAUDIBLE]

PRESENTER: Maybe that's this. Animations are not supported in, so it wouldn't automatically switch materials.

AUDIENCE: What is the difference between the A360 cloud renderer and just [INAUDIBLE] what is the difference in the actual rendering [INAUDIBLE]

PRESENTER: So the question is, what is the difference between A360 and my rendering farm at home. The technology that we have loaded, this is all in Autodesk's building, on our servers, and it has a rendering technology called, I can't remember the name of it, but it's not a render that's native to 3ds Max. Do a little size one and see. That's what I did here, this renderer cost me no additional charge. But it needs to convert all others to this other rendering engine that's not in Max. Yes.

AUDIENCE: I've got a microphone, sorry about that. Do you have to package the whole file or does it support proxies in bins and things like that?

PRESENTER: It does. Proxies and those elements are typically native to our renderer, and so they wouldn't be supported by this other renderer. Any other questions about A360?

Introduced with the newest version of 3ds MAX is the ART renderer, or the Autodesk Raytracer. This particular renderer is the one that you will now find in Revit, in AutoCAD and Inventor, so we wanted to make sure it was in 3ds Max so you always got predictable results. This particular renderer is a CPU only progressive render, it's physically based, so unbiased and you can see that it gives you really good results. Very straightforward controls, physically

based, ideal for Active Shade iterations and it supports all the stuff in MAX. Again, unlimited. It uses iterations to avoid noisy animations. Let me show you the ART renderer. I've disappeared back here behind my computer screen.

Let me open up 3ds Max. Here is 3ds max, and let me show you the art dialog box. I haven't really talked a lot about the rendering dialog box for a while, and we talked about Scanline which is 25 years old, there's lots of documentation about that. Here is the rendering dialog box, I hope you can see. Here's where I choose my renderer, so here's all the renderers that I have installed on the computer, and then here are the controls and because this is a physically-based CPU renderer, you can see that I can basically tell it how long to render, and then the rendering method.

So here I have it set let's set it down to a minimum so you can kind of see and I'll hit the rendering button, and because this is a progressive renderer what I'm anticipating is, it's converting all of that stuff to be rendered and all of a sudden the image will pop on the screen and what it'll do is it'll continue to refine, and refine, and refine through the steps of progressive rendering so all CPU's are working on the entire thing all at once.

Makes sense? Yes. Does it support render Elements yes.

AUDIENCE: So, besides physically-based, does that mean it's a physically-based shader model?

PRESENTER: Yes. So the question is, I say physically-based, does that mean physically-based shader models? It's made to do photorealistic renderings.

AUDIENCE: So, PVR materials [INAUDIBLE] just converting over to that shader model is the same shader model? Or are you using other shader models

PRESENTER: That was a long one. So, do the physically-based rendering shader models, PVR shader models, is that what's using? No, ART is using the shaders that are in 3ds Max to create, it's just that the ones that are supported are the arch and design physical materials, those kinds of things.

AUDIENCE: How would that handle displacement maps?

PRESENTER: How does it handle displacement maps? Exactly as any other renderer would. So you see, the render is done, it still looks like if it needs some time. So when you're doing animations with a progressive renderer, I wanted to just point out that if I'm using different times for this, do you

remember the little laptop in our network rendering? It might only get this much done in that amount of time, whereas a big supercomputer might get twice as much done in that same amount of time and it'll look much better. If I'm combining these images together to create an animation you get that crawl effect of all those little dots because they're not gone.

So what you have to do with a progressive renderer is, you have to actually set it to achieve the same level of quality for all computers. And so, the little laptop might take an hour to get this done and a supercomputer might only take five minutes to get this done, but they'll achieve the same level of quality. So you don't get that crawl in there. Does that make sense? That's perfect science, you're the best teacher in the whole world.

OK. Where were we? ART. So we talked about the pros and cons of ART. Those are all the renderers that Autodesk and specifically that 3ds Max team contributes to. Does that make sense? yes

AUDIENCE: Can you live-preview with ART?

PRESENTER: Yes. So the question is, can I live-preview? I think if there's one benefit that I should have put up there is that it's really good at active shade, and active shade allows you to assign a progressive renderer to a frame, and then as you make changes to your model, it automatically updates in the frame. I'll show you that in another moment.

AUDIENCE: I'm just curious as to what the decision was to make it CPU-based instead of GPU-based.

PRESENTER: Why is it CPU-based? This is something that we acquired with another company, and so it was down that road already. Another reason, I believe, is that we didn't work on converting it completely is because you're seeing that GPU base needs to be kind of optimized for a specific hardware manufacturer, and we love all of our manufacturers. So there are a number of third party renderers unless there are other questions regarding Scanline, ART, Quicksilver and A360.

So there are a number of third party renderers that ship with 3ds Max, and third party means that it's actually created by others. And so we have a licensing agreement with these other rendering guys. You may recognize their names, many of you are using Mental Ray, so Mental Ray has been shipping with 3ds Max since Max 6, I think I said 5 in your hand out, but it's actually 6.

Mental Ray, the guy said Mental Images won an Academy Award for the rendering technology

that they've created in Mental Ray, it is fully functional, fully featured, it is an amazing render you could see that he gets amazing results. Especially when Marianne is doing it. Marianne is the person that did most of those images. So a longstanding heritage with Max, all of the Autodesk material libraries are based on Mental Ray, you can achieve amazing results, it does physically accurate renderings, and again, unlimited rendering with BackBurner.

There's also a standalone version which you can just buy in Mental Ray and put it out on computers and then, rather than sending it through a render, you save a Mental Ray file and you send that to Mental Ray and it will render it. So it doesn't matter if it comes from Max or Maya, it doesn't matter, you can use that Mental Ray to do mixed pipelines.

I think the guys at Mental Images would always agree that when it comes to one of their challenges, it's the documentation of their "how to use Mental Ray" that has always struggled with. If I want to learn how to do better limestone and I go to the Mental Ray help, the very first thing they show me is the mathematical equation and how that-- and it's just like, "you ain't helping me, dude". I just want my limestone to look better. There are some perception problems with Mental Ray, and one of the first ones is speed, everybody thinks that Mental Ray is slow, and I think it really is knowing how to tweak your Mental Ray to get the most out of it. So there are a number of books that show you how to get better speed out of it, but out of the box it feels slower than maybe some of the other renders. And then there is another perception problem that it hasn't been worked on in a while, and recently Mental Images was acquired by NVIDIA as a part of their showcase of what their cards can do, and I think that there may be some truth and some misunderstandings about the development. Any questions about Mental Ray in your box? Yes?

AUDIENCE: [INAUDIBLE]

PRESENTER: So moving forward, the compatibility of Revit, AutoCAD and Inventor to Max, will still be based on the AutoCAD materials. Is that the question? and the answer is yes, will they be Mental Ray-based or ART-based, that is the question mark for the future.

Recently, the guys at NVIDIA also created what's called Iray, and Iray, unlike Mental Ray, which is a CPU-based renderer, Iray is a GPU-based render, because we're a GPU manufacturer. So let's make one that takes advantage of our hardware. Again, you can see the amazing results that Iray can do, thank you Maryann. What a great artist she is. In fact, she did such a great job that NVIDIA used this image for their ads for their GPU's. I'm pointing

to the lower right hand side. So, again, amazing results very quickly with Iray.

Some of the pros, is that it uses all the Mental Ray shaders, so if I've already set up stuff for Mental Ray, it works with Iray. It is GPU and CPU-capable, so I can get the whole computer working on it. Obviously it's optimized for NVIDIA technologies, it's impressive results quickly, and again, unlimited rendering with BackBurner and other manufacturers. Because it is a progressive renderer, we know that we need to use iterations in order to get level of quality rather than time in order to get the same result for animations. It is optimized-- hey, that's the same thing in the pro side. And the only reason I did that is because, what if I have a non-NVIDIA card? All of the capabilities of Iray are not in 3ds max, so they have their own Iray shaders, some Iray specific things, that Mental Images creates that they don't license to us or doesn't get to 3ds Max.

AUDIENCE: Is it limited on map support?

PRESENTER: Is it still limited on map support? So the question is, is it limited on map support? Yes, not all the maps that I showed you earlier are supported. Does it have an issue with normal math? That one's new to me, I can't answer yes or no. Does it? Anybody?

And then, as you saw in the big graph, not a lot of people are using it for some reason, because it is an amazing render. When you install 3ds Max those are all the renderers, and there's one more, this View Renderer, if you're using the View Renderer you have a very specific workflow that you're doing and you wouldn't be in the class, so I get that. Any questions about when you installed 3ds Max you've already got five or six renders?

AUDIENCE: In all future versions, do you think it will support Iray, with 3ds Max?

PRESENTER: Irays is going to be in all future versions of 3ds Max, all I can talk about is today, today, and it is in 3ds max today. There is a little bit of a change, when you look at the big graph of all the different renderers that were on there, you saw the majority were used by one type, so one of the biggest requests that I got as product manager is, I don't use some of those renderers, can I just turn them off and not install them? In the installation process of 2017, you can actually uncheck some of the third party renderers that ship with 3ds Max.

Any questions about that?

AUDIENCE: Does that work for Iray?

PRESENTER: Iray. Yes, it's exactly the same I submit to the network and it fires off Max and brings up Iray, and it's just that, again, I need to make sure I use iterations. Let me just show you. So here is ART. Let me show you in Iray. So organized, it's amazing.

So here is an Iray file. Couple of things that I wanted to show you in this particular file, if we had the chance, was that I used, I went up to the NVIDIA site and I actually got some Iray shaders. So if I take a look in my material editor, you can see that there are Iray shaders in here. Because I installed that, now, this particular shader is optimized for Iray. So I use that on some of the elements, and then, if we take a look at the dialog box, very similar to ART, all I really can choose between is time, number of iterations are unlimited, and unlimited is unique to Iray in that it'll just go until I come back. But there is obviously a point of diminishing returns where I can't tell from four hours to six hours what the defer you've worked on this for two hours and I can't tell the difference, I should have stopped two hours ago.

Again, it's a progressive render, so if I hit the rendering button. It shows me what I did with ART, no, it will refresh in a second. But exactly like the ART renderer, it will show me all of the image and then refine, refine, refine. And if we set that to unlimited, it will refine forever.

AUDIENCE: My question is about network rendering.

PRESENTER: So the question is network rendering. In that dialog box you saw time, you saw iterations, and you saw unlimited. So if I do time, I will get different results depending on the power of the GPUs and CPUs in the different computers. If I use iterations, I will get the same results of quality, machine-to-machine, and if I use unlimited and send it to the network, IT will come down on me so hard.

Any questions about what you're seeing?

AUDIENCE: Is there a way to see how many iterations it went through to decide?

PRESENTER: Yes. So the question is, can I see how many iterations this quality is? If you look on the rendering dialog box back here, it's telling me what iteration it's on.

AUDIENCE: Does it save all the iterations or does only save the last iteration?

PRESENTER: It saves the last. So the question is does it save all the iterations, and no, it keeps refining, so the last one is not saved, it's just all a part of the process. If I tell the rendering dialog box to save it will save the image at that number of iterations, but it doesn't save each iteration.

Back to our PowerPoint already in progress. We have just a few minutes left, so here are the third party renderers very quickly.

This is obviously the one that we saw, the biggest number of users. Vray is used by every industry that 3ds Max is in. So here you can see, the guys at Scanline did this shot in Max using Vray for Captain America Winter Soldier, most of the rendering is done in design, visualization, product animation, video game trailers and cinematics. If you're using 3ds Max, you've heard of Vray. Even our guys in product branding use 3ds Max and Vray to do all of the product shots. I keep talking to Keith Chamberlain, the guy in charge, the story behind these is fascinating.

I really want to somehow get to you this story. This is not just a thing, this is the 3ds Max logo, but what this is, is a ballerina, and if you were to capture her movement over time this is her dance step. And they took that as inspiration to create this form and I think it's a really cool story.

So there's not enough that I can say about Vray, we could spend this whole 90 minutes on it. Amazing feature sets, you get really good results using 3ds Max elements, there's obviously a huge community, there's no lack of learning, their documentation is exceptional and I think it was their documentation that propelled them. If I wanted to learn how to do good limestone, they had an article on how to do good limestone. That kind of thing.

When you come to its concept, there's a lot to it. So if you just want a simple rendering that looks great, you're going to be prepared to learn a lot to get that. And then it's costlier than some and cheaper than others. There's no doubt that its cost is its value, but it's an additional \$1,500 ballpark for your Max. And then, for some reason they ship with a hardware dongle. Welcome to 2002.

AUDIENCE: I think that's a [INAUDIBLE]

PRESENTER: I'm sure it is and nobody knows piracy better than I do.

AUDIENCE: [INAUDIBLE]

PRESENTER: Yes. They're working on that, but when I said, hey guys I want to show Vray, they said, here's your dongle, I was like, still? So there you have it.

One of the largest growing, at 5%, is this Corona renderer, and you could see that corona is

amazing, they call themselves "slightly biased". They focus mostly on unbiased technologies, but they have a few different things that you can adjust for speed. These guys looked at rendering and said, they really want to provide an easy to use type of technology that gives great results. They have a rental type, if you just need to rent it for a few months, or you can do permanent licensing, I'm jealous. And they've really come up with some truly innovative features, but it's a small team working on it. Pixel for pixel is a little slower and they're still working on it.

Any questions about Corona?

AUDIENCE: With a lime?

PRESENTER: With a lime.

AUDIENCE: Does it come with some material sets?

PRESENTER: I'm not--

AUDIENCE: Just answer.

PRESENTER: So the answer is yes.

One that I was excited to learn about for the class was Maxwell, done by Next Limit Technologies. That same image was rendered, and one of the unique things about Maxwell is that in their dialog box, after the image is rendered, I can make adjustments. So here's the Maxwell rendering dialog box, and you can see that as a progressive render, as a GPU-based render, I get a preview as it continues to render, it's like 12%, but if I think, well, that's kind of dark. I can start to adjust while it's rendering some of these things, and it gives me incredible results.

It really is for what I call hyper-photorealistic type of work. It's based on photographics, so you saw the adjustment wasn't brightness or brighter/darker, any of that, it was ISO and those kinds of things. The techniques are based on photographic principles. The rendering engine interface is just another window to manage, so you hit render and the rendering frame buffer comes up and then a dialog box comes up and then the final, so it's a number of rendering dialog boxes all at the same time. Not horrible, but there was just. Accuracy, obviously takes longer, and if you don't know the terminology of photography, then you'll learn that too.

AUDIENCE: I would add a pro that's network rendering, because it's really smooth. So when you wanted [INAUDIBLE] a new [INAUDIBLE] you can have a spare with many processors as you want, with many computers as you want with certain licensing commendations and it will compile the end result at the end. So one might be at a sample of 12, on might be at 18, and it will give you the added-in results. And it's very simple to set up.

PRESENTER: So, Maxwell user commenting that network rendering is also a very big pro for them.

This one is, again, new, Octane, developed by the guys that OTOY in New Zealand. Just a couple of things that I did in their demo set. What's kind of interesting about Octane is that it comes with a separate rendering dialog box. So you can see that I can load all of my stuff without even firing up 3ds Max, I can make renderings in the dialog box itself. GPU-based. They have a material converter to help you with that, distributed GPU rendering, which is kind of interesting, with enough horsepower, if you put a massive video card in your machine, it actually feels like real-time rendering. It's that well-optimized. The licensing, I thought it was kind of wonky, after purchasing the standalone, then you have to purchase the plug-in, then, that kind of thing. Its NVIDIA-based, this is just personal, but I've been rendering for a long time, and it took me a while to understand the terminologies of Octane, lights, cameras, and so on. So that was me.

In two minutes, I'm going to show you the new kids on the block. The first one is the Radeon ProRender, and this is done by AMD. It has VR output, so that's kind of interesting. That one's kind of weird, Radeon ProRender Pro's. You would think that the guys at AMD, who making ATI cards, are like, "I'll show you, here's one based for ours", and instead, they went GPU-agnostic, so I thought that was kind of interesting. It uses both CPU and GPUs, ships with a huge material library. It's not very difficult to learn, and from the start it uses all the native 3ds Max stuff, so all your files all the way back to Max 2.5 work. VR output, and it's free.

AUDIENCE: So is it open source?

PRESENTER: Open CL. So they're still kind of early in development, so there are still some features that didn't work as expected for me, exposure control, physical cameras, etc. If you want to do GPUs, it has to be open CL 1.2 hardware. And again, you guys, we haven't even heard of it.

Redshift is a renderer that's been out in the film industry. If you saw Cloudy With a Chance of Meatballs, you saw Redshift rendering. So now there is a Max version of it. It is lightning fast, it is GPU-accelerated, so it's not just utilizing GPU's, they are accelerated. Production proven,

they update often, \$500 a seat or so, and it recently won a five-star review best in class in 3D World Magazine. So it's still kind of new to 3ds Max, they know that they have some improvements to do with some stuff, again, NVIDIA-based.

And then, the last one I wanted to show you is this thing called Arnold. Arnold was recently acquired by the Solid Angle team, which was recently acquired by Autodesk, so they are now a part of our ecosystem. And again, production proven, if you saw the Martian, you saw Arnold in action. So there is no question that it does amazing work. 90 seconds. So, production proven, it does really good with huge data sets, it is optimized for just tons of information. Artist-friendly, easy to get better than good results immediately, and it uses a lot of, when we built it for Max, we built it with Max, so a lot of the same stuff works Max to Max and it's easy to adapt.

We're still working on it. We just released 0.8 today, so you can download that from the Solid Angel site. Takes a little tweaking and it's definitely pricey. So each individual node on your render farm needs a license of Arnold. So I'm sorry I didn't get to the best practices, but they're all in your book, so downloaded the handout and enjoyed some of the best practices I had for you in there. Thank you so much for attending. I hope I gave you a good insight.