Terminators and Iron Men: Image-based lighting and physical shading at ILM.

Ben Snow

Industrial Light and Magic

In the Computer Graphics visual effects community, we've learned to live with the accusation that our work looks fake or looks "too CG". In the business of looking at images, everyone is an expert and will offer their opinion.

Of course, we don't *want* our stuff to look computer generated. We spend most of our waking moments trying to get away from that. This course is about some of the ways we do that and I'm going to use some of the film projects I've worked on at ILM to give examples of how. I'll start with a reminder of what we're trying to recreate, particularly in the Visual Effects context. Then I'll give you an overview of how we've tried to improve that over the years, starting with the introduction of ambient occlusion on Pearl Harbor through the energy conserving, image-based, importance sampled approaches we used on Terminator Salvation and Iron Man 2. Finally, I'll talk a little bit about the process (and challenges) of capturing images and environment information on set.

What are we reproducing here?

In computer graphics a lot of the time we're trying to reproduce the reality that the viewer sees with their own eyes in the world around them.

But with visual effects for film, we're really trying to produce *filmed* reality, or more correctly the view of the world that is captured on film or another medium. And it's an important distinction. One of the things we noticed while doing the planes on Pearl Harbor was the way the highlights would flare out in the real footage. But the filmed references and 8 bit log digital scans compressed the dynamic range of the scene, so when we balanced the ambient and reflection lights the highlights tended to turn into mush. On that film we compensated by boosting the brightest values in the environment map. And further we add an additional bloom to the hottest highlights in the composite – trying to reproduce what happened to them in the medium of film.

It's not enough to just match the reality of the lighting of the set. This has to be viewed as very much a starting point. As artists we need to use our aesthetic judgment to make our creatures correctly integrated into the scene (and here the physically based techniques discussed in this course are primary), and also presented in an aesthetically pleasing way the supports the story being told by the shot. It's great when the director of photography has some real reference object like a partial Iron Man suit to light when they're lighting the shot. When that isn't the case its great when they're at least thinking of what the effect would be, and light the scene with that at least partially in mind. In that case our reference material and HDRIs are a good foundation and might give a good result out of the box. But equally often the scene won't be lit with the character in mind, and so our captured environments give us a good foundation

for the reality of the scene but even more than usual need to be complemented artistically to create a truly effective image.

Another thing to remember is that there are going to be crew around when the material is shot. While you don't want the crew perfectly reflected in your lovely Iron Man suit, you don't always need them to clear out either – if Iron Man had been in when they shot the footage, the crew would be in his reflection, but probably made less apparent with some dulling spray.

The 90s – light probes, textures and environment maps

At ILM in the mid '90s during the first wave of post Jurassic Park films, we were using two types of light probes for calibrating our CG lighting back at ILM with the lighting from the location. The first was a chrome sphere for capturing environment maps and providing a 360 degree picture of the light on the subject and the second a grey sphere as an easily matched reference for how the lighting would affect a known object in the scene.

Where possible we still shoot a reflective "chrome" sphere and diffuse "gray" sphere on location, but now we also shoot High Dynamic Range images as well.

CHROME AND GREY SPHERE ON LOCATION (c) 2001 Industrial Light and Magic. All Rights Reserved.

If we can match the grey sphere materials on the computer, we get a way of evaluating how closely our CG lighting matches the lighting on the set.

The chrome sphere gives us a look at that lighting, and we can unwrap it into a spherical environment map that can be accessed with a standard Renderman environment call. We mostly used it for reflections. The chrome sphere gives us a 360 degree view of the scene but the back hemisphere is very distorted at the grazing angle. Also, by the time you shoot it on film and scanned it, you often ended up with an image a few pixels across – way too small to get a ton of detail. One solution to the latter was to shoot still images in four directions with a stills camera and map them on a cube for reflection.

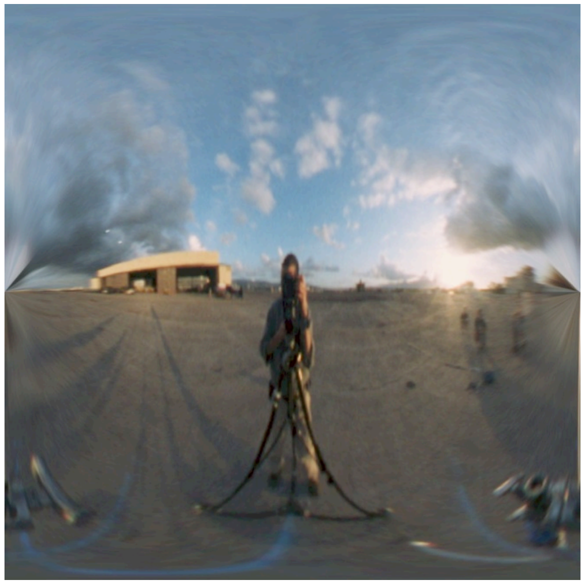

CHROME BALL DERIVED ENVIRONMENT MAP EG

In the early 90s we also relied a lot on painted textures, and fairly simple shaders (straight out of the Renderman companion) maintained by a group of technical artists. We put a lot of faith in our textures, and we had very talented texture artists. The problem with relying too much on textures this way is that you end up with painted in detail like highlights and shadows that doesn't respond correctly when you change the environment, which leads to a bunch of do-over's of look development for different scenes in the movie. So as we improved our look development standards we tried to watch out for that sort of thing.

We were using Cook-Torrance specular functions (and still use a modified version of that in our generic material) but investigated other specular functions like La Fortune. Lighting wise we in general used the standard early '90s CG setup of ambient and reflection environments with spot lights and one or more directional lights for the lighting. People who worked in that environment will remember the evils of ambient and reflection lights which were not correctly shadowed.

Too Much Ambient isn't good for the brain. Mars Attacks!

On various films we tried things like replacing the ambient with multiple spot lights of low intensity or even with their shadows turned down from 100% opacity.

The rock monster from Galaxy Quest. It was difficult to reproduce all the bounce in this environment and still get realistic shadowing. We used a large number of lights of low intensity.

Pearl Harbor and the development of "production ready" IBL.

So, in the rich tradition of visual effects, both Computer Graphics and more traditional, we approach the end of the '90s with a bunch of hacks. Of course ray tracing and global illumination were already in use in production and constant improvements were cropping up each year at Siggraph. But we were, and still are to an extent, in love with the look of our Renderman renders, and were already dealing with scenes of such heavy complexity that ray tracing and global illumination were not practical solutions. At that time we started using Reflection Occlusion and Ambient environments to provide us with a production ready approach to global illumination – a way of providing similar looks to ray tracing and global illumination, but with less expense.

These techniques used a ray-traced occlusion pass that was generated independently of the final lighting. The information in these passes was used as part of the final render calculations in Renderman, but we were able to change our materials and lighting without having to calculate these passes.

Reflection occlusion was developed during Speed II as a way of occluding or shadowing the reflections on the CG cruise ship created for that movie. It was put into full production on Star Wars: Episode 1. It addresses the problem of reflections not being correctly occluded when you use an all encompassing reflection environment. Single channel reflection maps like the one show on the left below are used to attenuate reflection in areas that are either self occluding or blocked by other objects in the scene.

Surface shaders read the occlusion maps and then attenuate the environment to provide us with realistic reflections as shown in the second and third images.

B25 REFLECTION OCCLUSION PASS, B25 WITH REFLECTIONS, B25 WITH REFLECTION OCCLUSION.

Ambient environments is a technique that gave us a way of getting diffuse fill light illumination that was more like what we'd get from global illumination. We introduced this technique on Pearl Harbor. For that film Michael Bay challenged us to come up with a more realistic looking computer graphics airplane than had been achieved before. We also had a need for a computer graphics ships but felt that we hadn't been able to get them to the level of realism that had been achieved with miniatures.

I felt that global illumination was our best bet for the latter, and Ken McGaugh and Hilmar Koch started investigating whether we could get our CG ship model to render in mental Ray. Meanwhile Hayden Landis was working on the airplane solution and felt that the multiple fill lights solution we'd used to approximate more even environmental fill was too cumbersome. He wanted to try using a blurred version of the reflection environment as a big ambient light to provide fill. The problem was that the ambient light wouldn't be shadowed properly but we decided to give Hayden a couple of weeks to try and come up with a solution. Working with Ken and Hilmar, he came up with ambient occlusion, a technique that, like reflection occlusion, uses a pre-rendered occlusion map accessed at render time to give the scene realistic shadowing. In addition, they developed a way to derive directional information so that a given surface point would be illuminated by the most appropriate part of the ambient environment map. Below is an example of the ambient environment render with the different channels used to indicate a light direction, and the final beauty render.

Ambient Occlusion Render, Beauty render with Key Light and Ambient Environment Light only.

Ambient occlusion is achieved by casting rays in a hemisphere around the surface normal for every surface point. The final occlusion amount is dependent on the number of rays that hit other surfaces or objects in the scene. Since the strongest diffuse contribution comes from the general direction of the surface normal, the result is weighted to favor samples that are cast in that direction. But beyond that, we do a first pass to calculate the average direction of the available light arriving at a surface by averaging together the unoccluded rays that hit the surface. This is weighted with the surface normal to bend the direction of the lookup to the appropriate direction in the environment map. We use the term "bent normals" to refer to this effect.

Here are some final images from Pearl Harbor

Reflection occlusion and ambient occlusion were great tools that gave us a definite boost in both realism and speed, and they've served us in good stead ever since. But they're not the end of the story.

Around the time of Pearl Harbor we started seeing some very interesting developments in image based lighting with Paul Debevec's presentations of using High Dynamic Range Images and global illumination to make very realistic renderings. At ILM we developed ways to improve the dynamic range of our composited images with the introduction of the openExr format.

On Hulk we started using Sphereon rigs for capturing high resolution and high dynamic range images of the scene, and developed tools to recreate the set and project that captured image material onto rough proxy geometry for ray tracing and faking diffuse reflections. We also made tools to calculate the location of lights in the scene by projecting rays from the capture location to the brightest points in the projected image. This was quite an expensive process at the time and so wasn't generally used in production at ILM until several years later.

Iron Man and better metals.

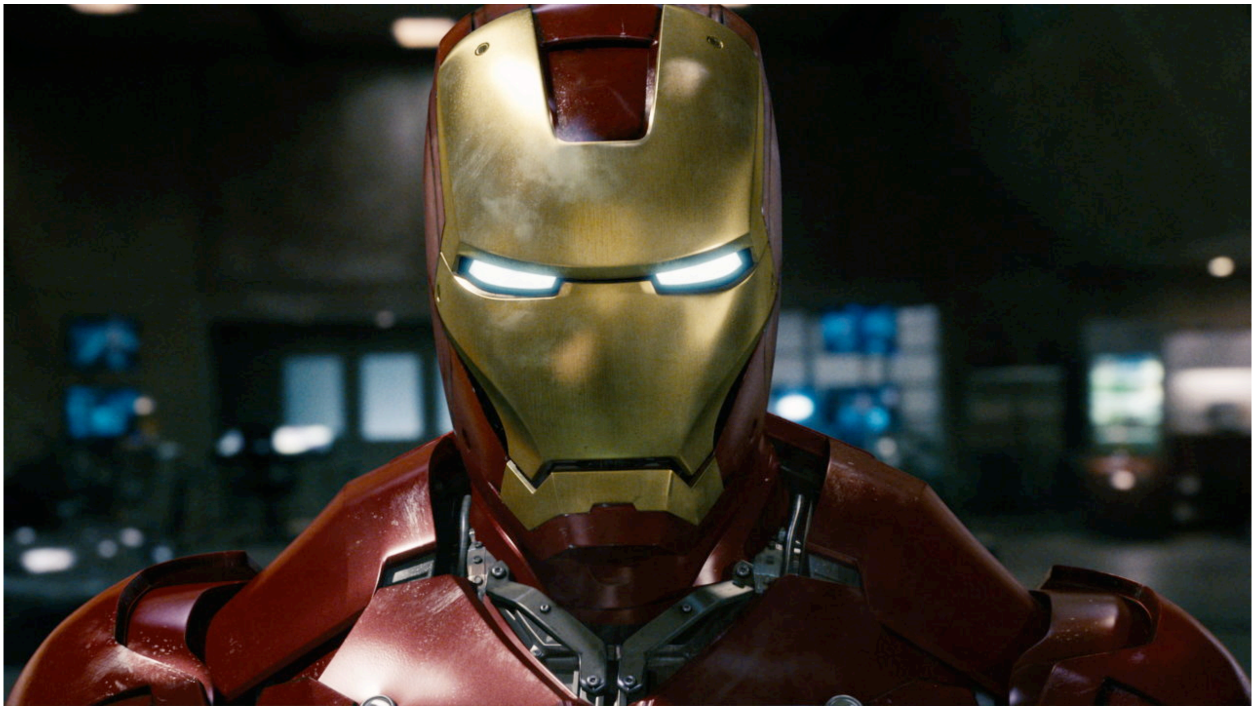

By the time of Iron Man we'd begun to heavily use High Dynamic range images captured on set with digital cameras. We'd also developed a variety of more realistic car finish materials for Transformers, particularly a specular clear-coat ripple and some metal flakes materials. When ILM was approached about doing a test for Iron Man to help make director Jon Favreau more comfortable with the large amount of computer graphics he'd need to use on his film, we were able to leverage the Transformers material to quickly put together a fairly realistic metal suit for the test.

Shot from an early Iron Man test.

The Iron Man armor for the film was originally been discussed as largely computer graphics, because it had to have a lot of additional functionality that would be impossible to do with a practical costume. However, financial reality set in, and the film-makers decided to get several suits built by Stan Winston's company (which later became Legacy effects). The aim was to try and get many of the shots with practical suits, and then use Computer Graphics to complement that and for things like an aerial dogfight, where it didn't make sense to do practical stuff (although we did talk about things like strapping a stuntman in the suit to the side of a helicopter).

The test and what we were seeing from Transformers told us that our car materials were looking pretty good. We'd also settled on a solid approach for acquiring High Dynamic Range images. But I hadn't really been happy with our raw metal surfaces. Chrome was pretty good, but brushed metal and dull metal surfaces tended to rely a lot on texture maps to add detail, and seemed to come out feeling like a fairly diffuse painted oily metal. Also the Winston/Legacy beauty suits were spectacular. They'd actually chromed and gilded the helmets and parts of the suits and then added paint and other finishes making them look great. As we finished photography it was clear our suits were going to have to hold up very closely with the Legacy suits.

Practical Iron Man suit created by Stan Winston studios

One of the things we decided to do on Iron Man was to try and acquire sampled BRDFs for certain materials. The Legacy effects team made some small references spheres for us of the different suit finishes, including the red armor, the gold armor and the brushed metal of the Mark 2 silver suit. We also had them make 1 inch square swatches with the same finishes and sent them off to have BRDFs sampled.

We got back data on the red armor, and some useful stuff on the gold, but the brushed chrome was a problem. It was a little too complex for the agency doing the sampling.

At the same time, the materials we created for using the sampled BRDF were proving to be a limitation. You could plug in the BRDF table and get quite a nice highlight out of it, but it wasn't terribly easy to art direct on top of that, or tweak or even, initially to change the base color. In the end the sampled BRDFs were useful in helping us better understand how the materials responded to light, but not actually used in and of themselves in our materials for the show.

Based on that analysis we came up a new anisotropic specular function to help improve our brushed metal surfaces.

Anisotropic Brushed Metal

The anisotropic specular function we used on Iron Man added separate X and Y specular exponent values along with a direction-of-anisotropy vector (e.g. "scratch direction"). It allowed separate specular exponent values per R,G,B color channel, giving a color fringe used in the second specular highlight in our shaders. This provided the macroscopic "stretched highlights" characteristic of brushed metal seen from a distance. The meshes were UV laid out so that the primary brushing direction was always exactly vertical in the texture, eliminating the need for a direction map. The specular ended up being close to a regular Cook-Torrance for compatibility with the existing controls. Our modifications were for speed, to add the anisotropic qualities and to normalize it. We'd considered Ward and Ashikmin-Shirley's approach but mostly went with out own approach for backward compatibility issues.

The actual brush lines were painted as a standard specular attenuation map (so they mush together and fade out when further from camera), but we also modified the shader to allow darkening the diffuse and reflection components by the specular map to further enhance the visible brush lines.

This was great for the specular highlights from our individual lights, but we'd also been increasing our use of environment lighting based on high dynamic range images, and when the silver suit was in a more environment light based setup we didn't get the lovely bow tie highlights in the reflections of lights coming from in the environment sphere (because at the time that information was treated as separate reflection). We also figured out a way to get anisotropic *reflections* with environment spheres, but it definitely felt like double the work – something better was needed.

We were able to use some better controls to allow selective ray tracing in Renderman that Transformers had come up with (ray tracing still being pretty expensive in that package). We needed to limit the number of trace bounces for speed. However, our materials basically didn't send out a ray if the reflection occlusion pass told them not to. We also set up base colors that were close to the textured colors so that we could turn off some of the texture maps during the ray tracing passes.

Brick Maps for Ray Traced reflections.

In our need for speed we also increased our use of Renderman brickmaps for cacheing surface data for ray traced self reflections. Brickmaps are Pixar's name for a volumetric texture map. Instead of a tradition 2-D texture map projected mapped onto an object, a brick map defines the color at each voxel that was occupied by some geometry when the brickmap was created. Like textures, brickmaps are mip-mapped in Renderman, meaning they have several nested levels of detail as shown below.

Use of High Dynamic Range Images

By the time of Iron Man, use of High Dynamic Range images for environment lighting was fairly well established at ILM. We were using the images mostly as environment maps for reflection and ambient illumination, still much like the unwrapped chrome spheres from Pearl Harbor. For Iron Man we ended up shooting "highish" dynamic range images with 3 bracketed exposures, (0, -3, +3 stops).

In general we created hero environment maps for each set or location, with one or two variations for different lighting conditions that were shared across the sequence. Much of the lighting was still done

with traditional lights even with having the High Dynamic Range images and better tools, although we did start using area lights and reflection cards more heavily.

Look developing the CG suits to match Practical

As mentioned, the aim had been to shoot as much as possible with the practical suits and augment with computer graphics. In the end, by far the majority of the suits were computer graphics – something like 300 CG suit shots vs. 40 practical suit shots. This was because the practical beauty suit was great looking but heavy and difficult for the actors to move in. The lighter stunt suits didn't look as good, even for more action shots. And Jon Favreau was preferring the movements our animators were getting in the animated suits.

It was clear that our Computer Graphics suits had to match the practical suits. With the techniques described above in our toolkit, we were able to get fairly close matches to the practical, which was vital due to the need to cut back and forth.

IMAGE: SIDE BY SIDE CG and PRACTICAL

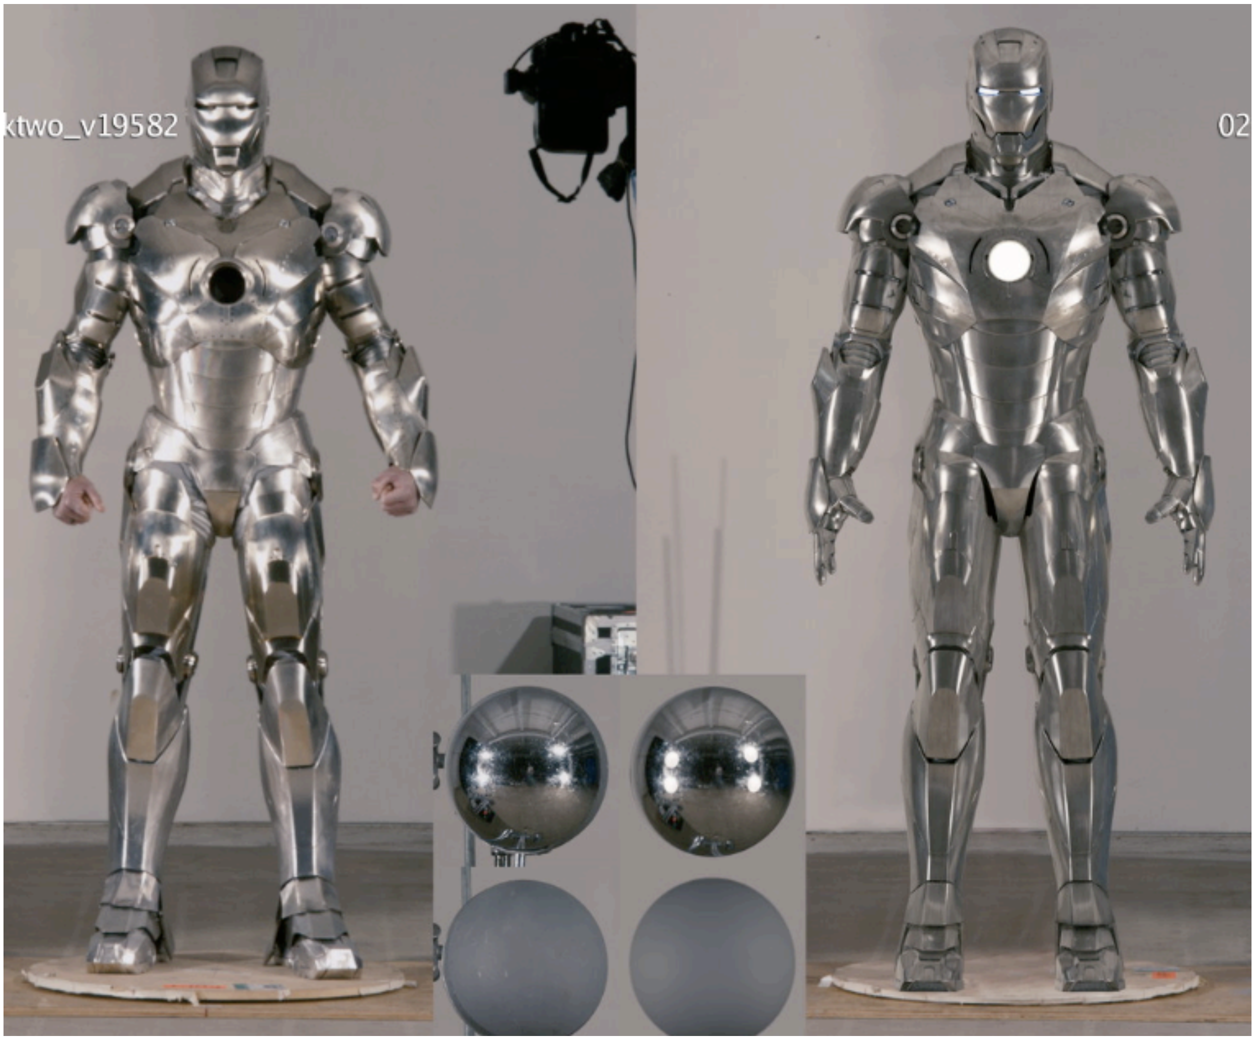

These looked good in the turntables and also in the shots, but a couple of things became clear. Firstly, the shader setup we were using made it possible to create unbalanced lighting that affected different components (ambient, diffuse, specular, reflection) across different areas of a model, and particularly between different models.

Ambient and diffuse are really the same thing, and specular and reflection are really the same thing, but we had different controls for each, and they would drift out of alignment as the look development artist was working on separate parts of the model, and the balance between the components would be inconsistent. In addition, the way we mixed lights and reflection cards to get a good looking image was rather arbitrary and up to the individual artist, and each of those lighting instruments only affected a subset of all the lighting components – the reflection lights wouldn't give diffuse illumination on a surface and wasn't affected by the surface's specular BRDF function, the CG lights gave too perfect and shiny highlights that lacked the complexity of real lighting. The look could change radically depending on the instrument type used by the artist.

I was also pushing to match more of the lighting that I was seeing on set – the bounce cards, larger light sources like Kino Flos etc., that the D.P. Matty LaBatique used to make the practical suits look good. Area lights were still felt to be a little expensive, so some artists used reflection cards instead – again leading to different results. In the end, we tried to coral the lighting by having a sequence lead set up the instruments to be used in a sequence and prove those with a couple of test shots. The other artists would start with that setup. But it was still a time consuming process getting each shot to match and look good, and we felt that a new approach to lighting was needed.

The other issue arising from the large set of controls over the surface and illumination properties was that we'd need to make sure we ran our turntables in a variety of environments reflecting the environments the creature would appear in. On Iron Man we made a strong effort to making sure our main assets would work in all sorts of conditions, but it was sometimes a struggle.

IMAGES: Iron Man in Daylight, Iron Man at night.

EXAMPLE FROM IRON MAN, MARK 2 SILVER SUIT IN TONY STARK'S WORKSHOP

The beginning of the sequence where Tony Stark tests his silver Iron Man flying suit for the first time was a particular challenge showing some of the tools above. We'd shoot footage with the practical Silver Suit created by Winston/Legacy effects for most of the shots, that set a high bar. We knew we'd have to integrate and replace those with our CG silver suit to make the different control surfaces move around as Tony tested the suit.

IMAGE: PRACTICAL SUIT

We started with our "high-ish dynamic range" environment of the workshop.

We created a lighting environment with a combination of lights and reflection cards. The Anisotropic materials described above were key to nailing the brushed metal look.

CG Render of Silver Iron Man suit in the workshop.

Terminator Salvation and more physically correct materials in Renderman.

After Iron Man I moved on to Terminator Salvation, with many of the same ILM people who had collaborated on Iron Man. We were again faced with a bunch of bare, brushed and oily metal and with cutting our computer graphics with practical metallic looking puppets created by Legacy Effects. If anything we had even less of a budget than the already tightly budgeted Iron Man and wanted to leverage image-based lighting for quality, consistency and speed. We wanted to avoid the double work we'd run into with the anisotropic specular and reflections for the silver suit, and we wanted to try and avoid having to rework the materials for the different lighting environments we knew our Terminators would play in.

On Terminator we were faced with fairly dynamic environments filled with practical sparks, fireballs and explosions on the one hand, and harsh exterior desert environments shot in and around Albuquerque in New Mexico. On top of that the D.P., Shane Hurlbut, and director McG had decided on a very harsh and unforgiving digital intermediate process that would great increase the contrast of the image.

Another concern was that unless the computer graphics really matched the dynamic range of what was being photographed we'd be faced with seeing our computer graphics drifting away from the live-action plate – either the black levels would be too bright and flashed in the low end of the computer graphics part of the image, or the bright highlights would drift away in the high end.

In the few years leading up to Iron Man and Terminator, there had been a lot of interesting discussion and presentations at Siggraph and other conferences about physically based materials and lighting, and the move to energy conserving materials, and importance sampled ray tracing. We decided we wanted to use an energy conserving, importance sampled, image-based shading setup.

The goal was a simpler, more intuitive and physically based system of lighting and rendering.

Energy Conservation

The system we arrived at introduced energy conservation into our shaders. This means that the amount of light that reflects or bounces off a surface can never be more than the amount of light hitting the surface. As described in Naty's introduction to this course, in the physical world incoming light can be reflected off a surface, absorbed by a surface or transmitted through it. Ironically, some of our older shading functions created a situation (for example in Fresnel areas) where more light could be reflected off a surface than was being emitted by the light! To make our shading models more physically accurate, we introduced a new, normalized specular function, improved light falloff response, and importance sampled raytracing. The new tools also combined what we used to think of as the separate components of specular and reflection together into one component we called specular. Likewise what we used to call ambient and diffuse now are combined into diffuse. In addition we created new importance sampled image based lights so we don't use separate ambient and reflection environment lights. So instead of having to manually balance ambient, diffuse, specular and reflection in a material we just have to balance diffuse vs specular. Rather than losing artistic controls it makes it a lot faster to hone in on a more realistic and balanced look because the material is behaving more like something in the real world.

Normalized Specular

Let's start with looking at the traditional setup, where we treated specular and reflection differently. The image below shows reflections in a ground plane of three 50% grey spheres hovering above three

different ground planes (apologies that the ground planes themselves aren't clearly shown in these images. In the lower half of the image you see the reflections of the spheres, with each of the ground plane materials having a different reflection blur. As we increase the reflection blur the image of the reflected sphere gets darker because at each pixel we are now gathering light from a wider and wider cone of directions, so the average value of the pixel becomes less.

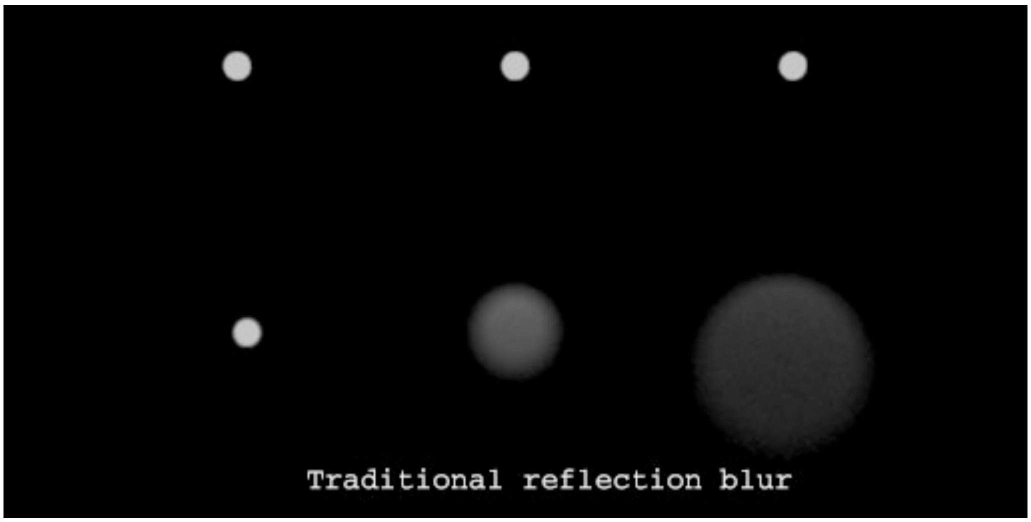

But it didn't work that way with our specular in the traditional ILM lighting scheme. In the image below we have a point light of intensity 0.5 at the position of each of the spheres. These point lights are illuminating three surfaces but this time we're varying the specsize (and not including any raytraced reflection to make the idea clearer). The specular from the point light doesn't get darker as the specsize increases, indeed the specular model we have been using for years at ILM actually gets much brighter with grazing angles so the actual specular values are very hard to predict.

Adjusting the gain of the specular to compensate for different sizes was left to the person doing the look development, so for a model which has a number of different specsizes the way that those surfaces respond to reflections (rougher surfaces should have darker reflections in this example) and the way they responded to specular from spotlights was inconsistent when it should be the same.

So a normalized specular function behaves in the same way that a reflection does. As the specsize increases, the intensity of the specular goes down, and if you looked at a graph of specular response vs reflection angle (ie the specular lobe) then the area under this graph (which corresponds to the spread of energy of the reflection) always sums to 1.0. The lobe is either thin and tall for a small specsize, or short and fat for a large specsize, but the area of the lobe is always exactly the same.

Here is the same render using the new normalized specular model. The specular is the same shape as before but the intensity behaves as it should, with larger specsizes producing dimmer specular reflection.

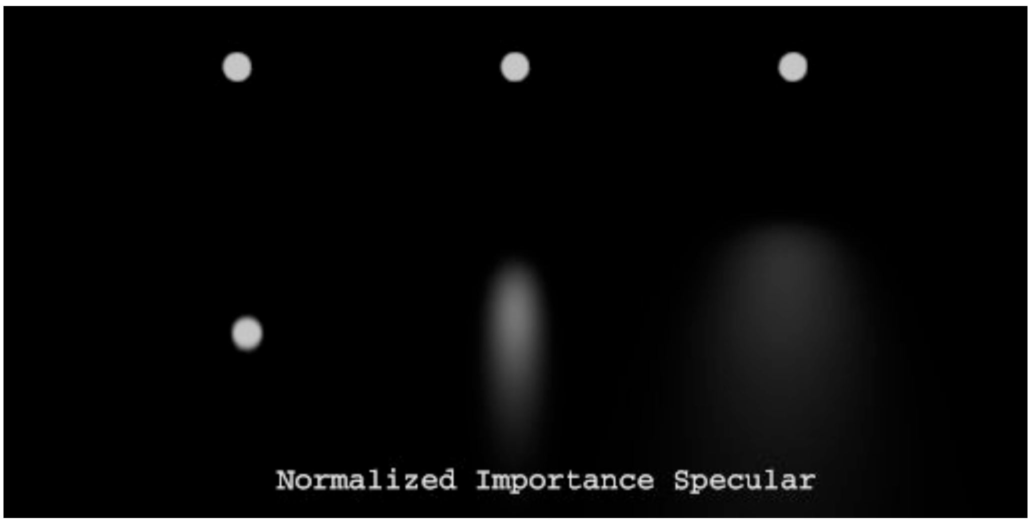

So let's apply this to the look development on a real world model (or real world in terms of a Terminator movie). Below are two images of a mototerminator with varying specsizes for the various different materials but with the specGain set to 1. The left image is rendered with the old traditional shaders and rhe right with the new importance shaders. The look development person would have to go in and start balancing all the specGains by hand but with the new shaders its looks better out of the box, with the duller surfaces like the tires out of the box reflecting light in the correct balance relative to the shinier surfaces.

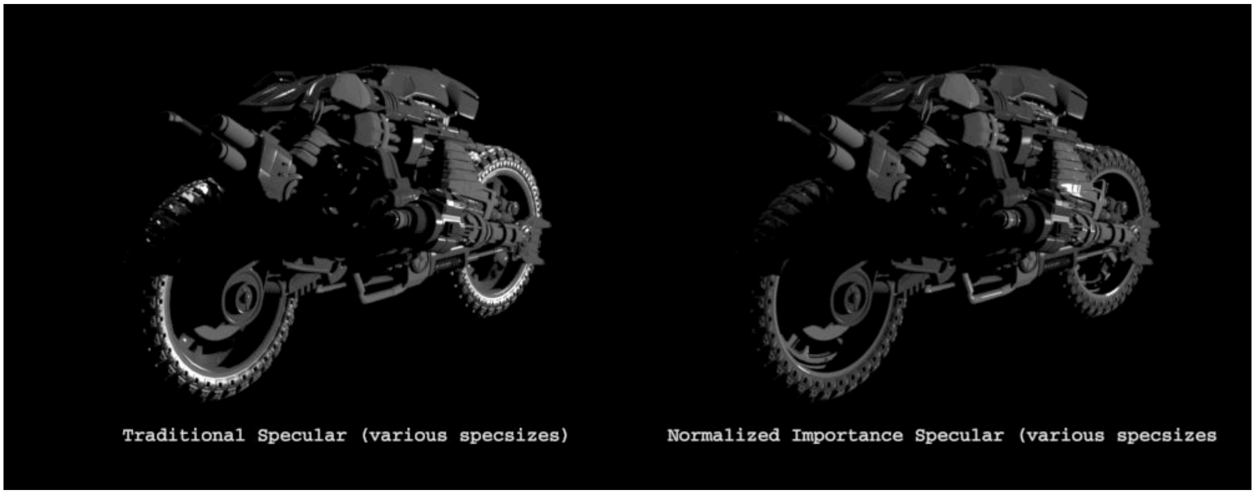

Specular Falloff

Another feature of the new specular is the ability to get more physically plausible specular falloff as lights recede from surfaces.

Traditionally at ILM, our lights defaulted to having no falloff. We had controls in the light that gave us the ability to add a 1/r2 falloff or a smoothstep falloff between a near and far distance or various other ways of defining falloff.   The falloff would affect both diffuse and specular contributions from the light but even after we added separate control of specular falloff it still didn't account for how the roughness of different surfaces should inform falloff.   When you're lighting a single creature in an environment this is not such a big issue, but as we start rendering more all-cg scenes having a physically inspired model can really help in achieving realistic lighting conditions.

The three images below help illustrate this. There are three spheres, each with different roughness. The first is chrome-sphere-ey (specsize 0.003), the middle a glossier metal (specsize 0.1) and the far one has broad specular (specsize 0.2).

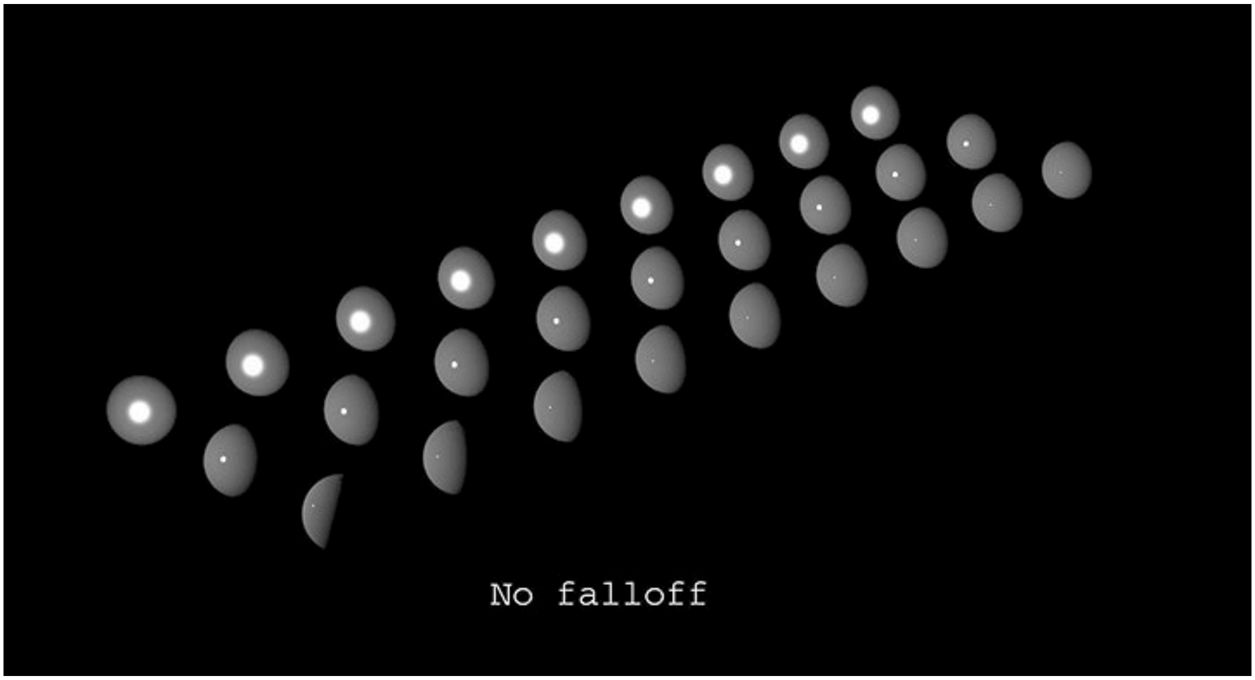

The first render uses the older shaders with no falloff on the light.

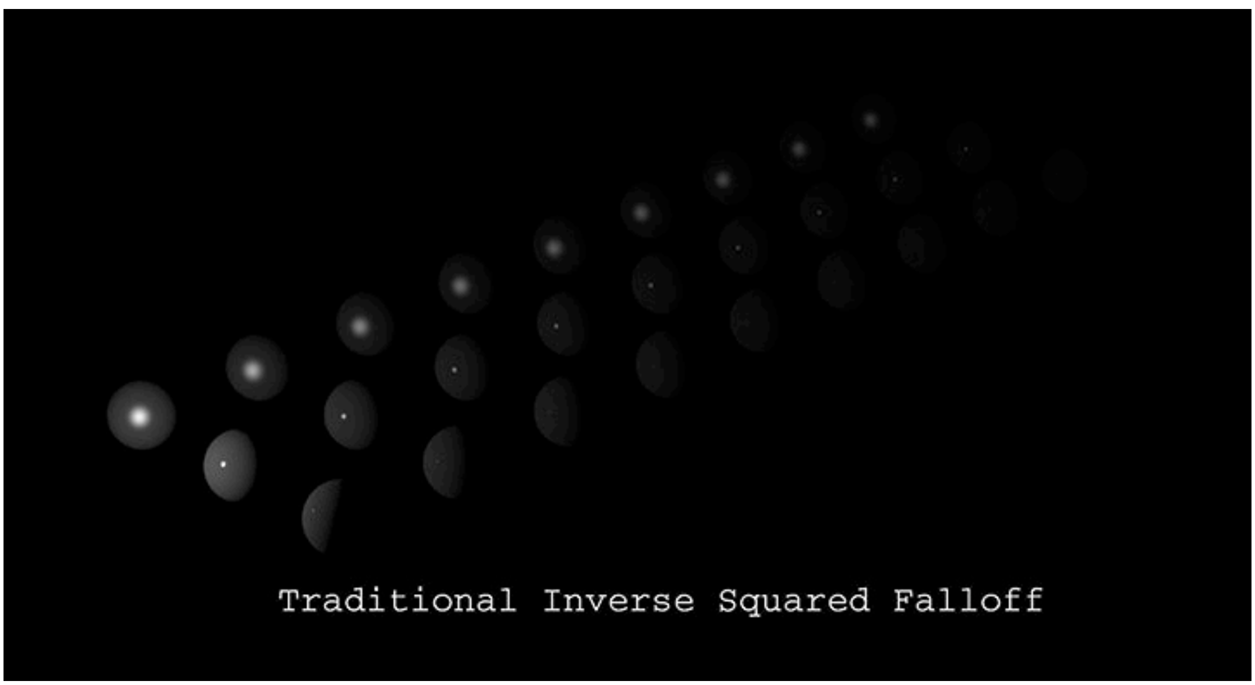

The second uses the older shaders with Inverse squared falloff.

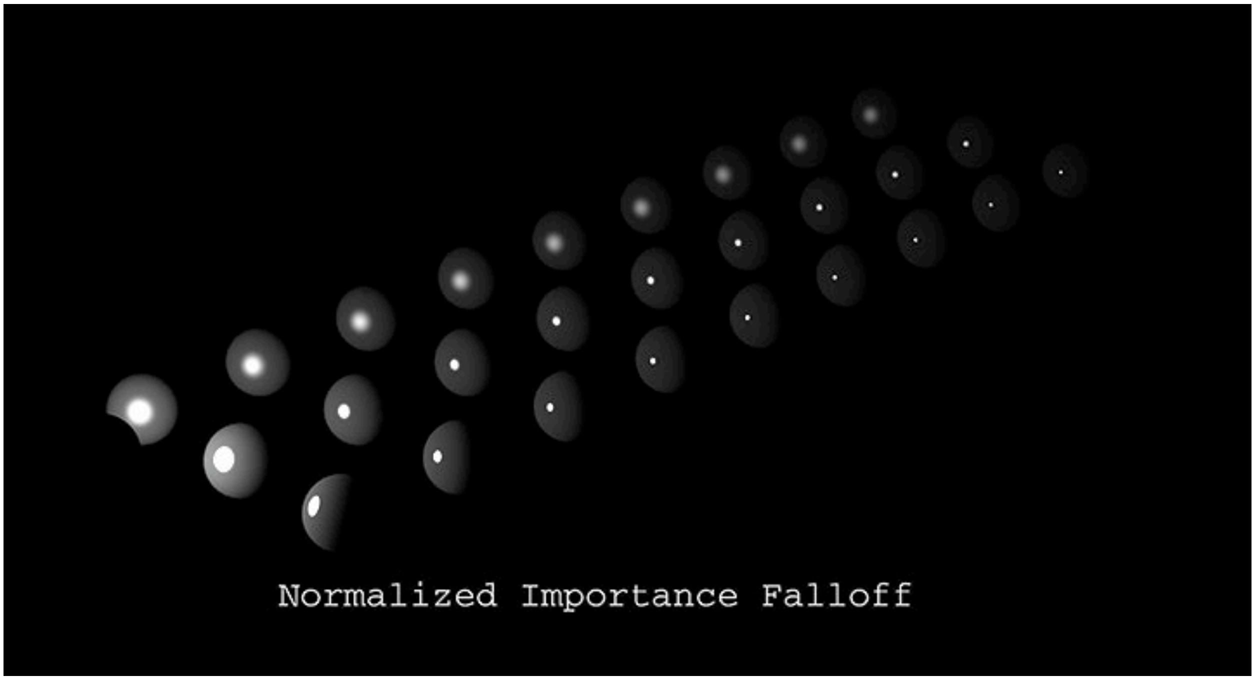

The last render uses a spot light again but with normalized-importance surface shaders. Note how on the rougher surfaces the intensity of the highlight falls off faster. This effect is unachievable with the old shader set.

If you've been used to taking the notion of inverse square falloff as the way things should work, this might seem a little odd. While light does "falloff" in the real world, when you are looking directly at something, or looking at something in a shiny surface as the object gets further away you see more of it. With specular reflection an object or light source will only get darker as it occupies a smaller portion of the specular lobe. For tight speculars, like chrome spheres or mirrors the light source has to get a really long way away before it starts to look dimmer. For broad speculars it will dim much more quickly.

For this to work we had to introduce the concept of a light having a physical size. For distant lights like the sun we define the light size as an angle rather than a physical radius and this angle will be the same for every surface regardless of how far it is from the light.

For the implementation we used on Terminator we gave ourselves the option of decoupling diffuse from this behavior, because some of our team felt there were times we might not want to use falloff. This sparked a holy war amongst our team, since to be physically correct the diffuse should fall off as well. There needed to be a shift in the way we thought about lighting, or more correctly, in the way were used to cheating lighting. As our materials and lights became more physically correct we had to treat them more like real lights, including allowing ourselves much greater light intensities and getting used to the light dimming as our slightly dull objects drove flew or swam away from the light.

Normalized Importance Sampled Raytracing

Importance sampling involves trying to sample your scene as efficiently as possible, by focusing your sampling on the points where you get the most bang for the buck. We do this in two places.

First we importance sample any images used for lighting. We want to concentrate our samples around the brightest points of the image and not waste time on the dark portions that don't contribute much light.

Secondly, when we cast out rays to get blurry reflections, we importance sample the specular function to determine which are the ray directions that will provide most of the lighting contribution, much like the bent normals we'd used back on Pearl Harbor. We also determine the direction which will provide least contribution. Both are taken into account to determine the best possible samples that will accurately capture both the light information and surface response.

A similar technique is used by Sony Picture Imageworks. Refer to the next talk in this course.

Below is a comparison of rendering with our old traditional reflection blur and our normalized importance ray-traced specular reflection.

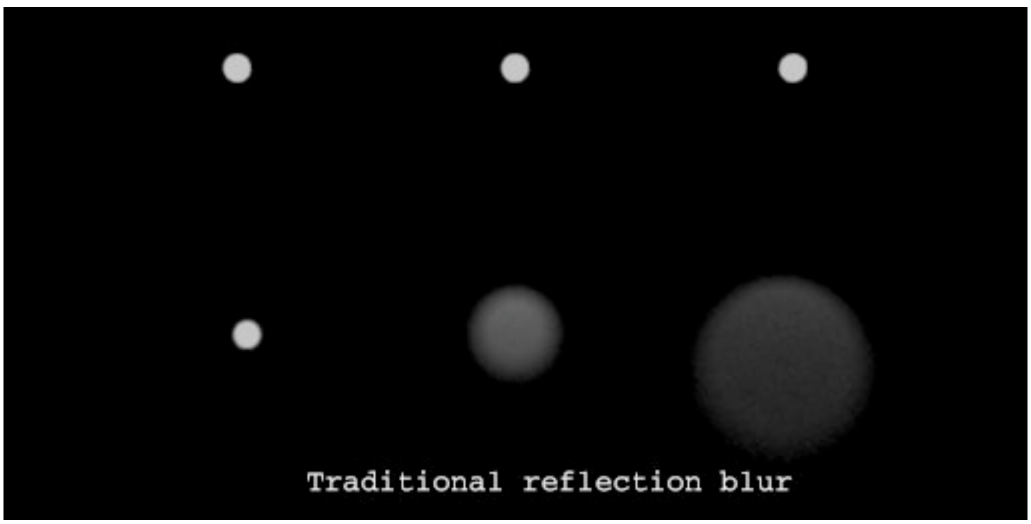

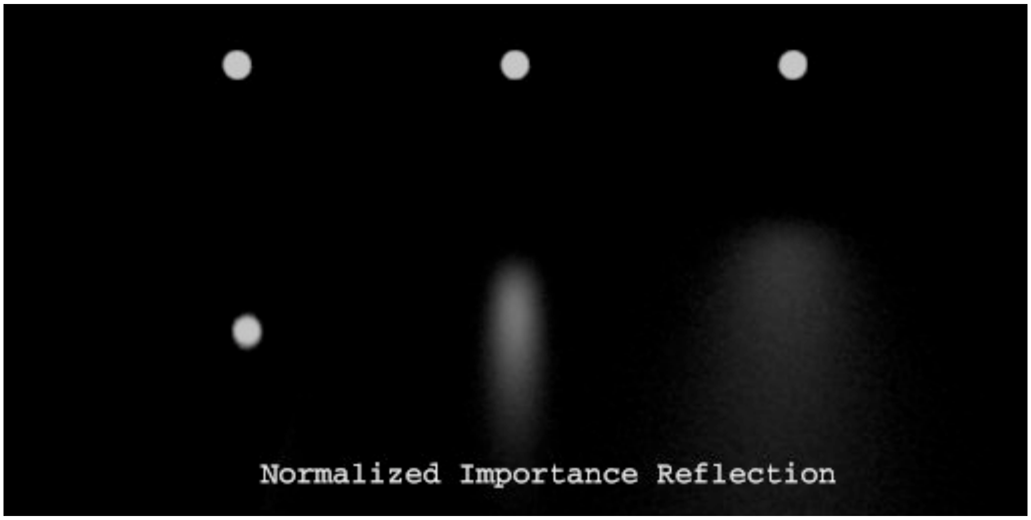

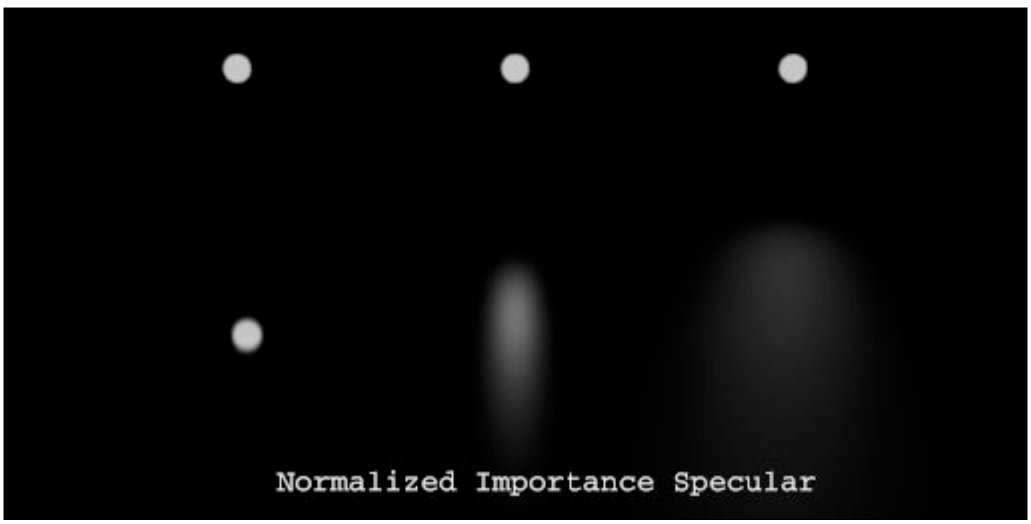

The last image above is the normalized specular image (the second image above) but with the lightRadius set to the same size as the reflected spheres, for comparison with the ray-traced reflection. As you can see from the similarity of these images, reflections from objects, and specular from lights is treated the same way.   This, we feel, is more like the real world.

Importance sampled ray tracing still comes with some cost, and while the time taken to place the lights and get a realistic result has been reduced, the render times have gone up. Sampling strategy is very important, and we were initially getting some rather noisy results.   We introduced sampling controls that increased as the specular sizes needed to increase. We added secondary sampling controls for where a shader is being called by another shader to resolve inter-reflection, indirect diffuse illumination or refraction so by setting these lower than primary values we could avoid an explosion of sampling.

Our preferred importance sampling strategy on Terminator was deterministic texture, which gave us smoother diffuse with fewer samples. However, the deterministic scheme can result in artifacts like

banding in the shadows. We also used stratified sampling, which fixes the latter problem but can result in noticeable aliasing with very bright values in environment textures.

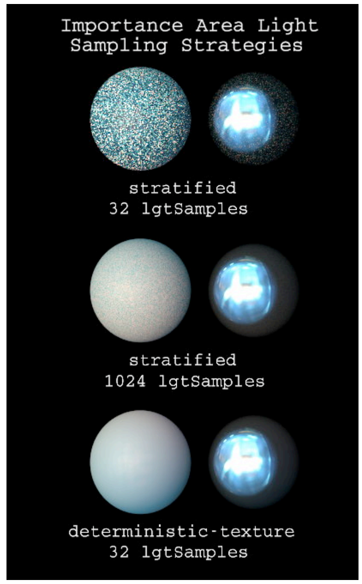

Since Teminator we have improved the texture sampling and on Iron Man 2 implemented a "mediancut-interleaved" sampling method, although the improved texture sampling on that show worked very well also.

For more information on importance sampling, I refer you to Pharr and Humphrey's "Physically Based Rendering" which I believe has a new edition out at Siggraph. There is also a siggraph course on Tuesday at 9am entitled "Importance Sampling for Production Rendering " at which Simon Premoze, who contributed greatly to our new tools, is one of the presenters.

SWITCHING TO THE ENERGY-CONSERVING, IMPORTANCE SAMPLED, IMAGE-BASED MATERIALS

The development of the new materials and lights took a bit longer than expected for our production needs for Terminator Salvation and so we decided early on to start with the traditional lighting and then commit to using the new approaches on a couple of sequences. We left the moto-terminator and giant Harvester sequences as using the traditional approaches and decided to use the new approaches for our skeleton terminators – the old, clunky T-600 that attacks John Connor in the beginning of the movie, and the evolved T-800 skeleton that attacks him in the Terminator factory at the end. This was because the beginning represented a high contrast environment where we had to cut back and forth with a practical Legacy effects T-600 puppet, and that end factory sequence was largely going to be illuminated with flashing lights, smoke and explosions and could benefit from a more image based approach. As work progressed we also made the resistance A-10 Warthog airplanes and the chase down the canyon using the new setup.

Our test case was a CG version of the practical Tow Truck in the film. While we wanted the new tools to give us results that were close to, but an improvement upon the older materials, the goal was to match

reality – or our filmed version of reality – and so we used a real world target for our initial testing, as well as a suite of methodical, more technical tests.

A couple of points here. We put a lot of effort into establishing a fairly solid Image-only based setup for our generic turntable environment (seen in the Truck turntable below). It's particularly valuable to use IBL for look development, because it soley focuses the effort on material setups. With our solid generic environment we felt more confident in being able to transfer an asset between quite different environments without having to greatly re-jig the look. For instance, with a "pure" IBL setup you don't end up baking the response of specific lights in your BRDF parameters. There is a school of thought which suggests that you'll do a better job of judging your material if the lighting field in which the observation is made is complex and high frequency (something you get pretty much out of the box with the real world captured IBL data).

IMAGE: Turntable frame of C.G. Tow Truck

MATERIALS EXAMPLE – INITIAL T-600/JOHN CONNOR TERMINATOR FIGHT

Here is our CG T-600 endo-skeleton. In the end this is featured in all but a couple of the shots in that initial fight.

IMAGE: T-600 turntable.

IMAGE: VE fight before/after.

As we started TD's out on the sequence we found there was a learning curve. Initially, the renders were a lot slower than what they were used to. However, this was offset by not having as many pre-passes for ambient occlusion, brick-mapping etc. so the complexity of the scenes was less. Another issue was that the TDs are so used to a bunch of 3D and 2D cheats to make their shots look good that they found it odd to be working in a more physically correct world where things didn't work the way they used to in the more artificial traditional approach. We almost had a holy war over eliminating the ability to control falloff in the lights, and ended up re-implementing it as described above.

We also asked them to use much more of an image-based approach, and so the environment sphere became a lot more important. The instruments used to add shadows were different – you couldn't just add some barn doors or paint a cuculoris in a slide map – you instead had to use more real-world approaches like adding flags and reflectors. A lot of the time the artists would basically go in and paint the environment map. Shots where the terminator transitioned from outside to inside the downed helicopter during the fight were a problem, due to the need to move between two environments.

In the end for speed we did set things up to allow us to selectively mix in ambient environments as well as pre-computed point-cloud approaches for occlusion and indirect diffuse calculations. It was a bit of a struggle, but we felt the new tools helped the realism and as the tool set matured were definitely the way to move forward.

IMAGE-BASED LIGHTING ON TERMINATOR SALVATION - THE T-800 FIGHT

IMAGE: T-800 in the Terminator Factory

The terminator factory sequence at the end of Terminator Salvation, where the John Connor and Marcus characters face off a charred T-600 endo skeleton was ideal for an image-based approach. D.P. Shane Hurlbutt was lighting the sequence with a lot of pyrotechnic effects creating a chaotic, rapidly changing environment. The issues of specular and reflection being treated differently in the old shader set could create a variety of headaches with all the reflections, but the new energy conserving materials solved this problem for us nicely.

Capturing the environment
This was our first problem. We don't have a good way of capturing high-dynamic range moving images. Our solution was to have the film-makers hold the dynamic effects like explosions and freeze the strobes (where possible) for our HDRI digital stills. Then we'd shoot a chrome sphere with all of the effects. Finally, we shot a whole bunch of the pyro elements on film, both for use as elements to add even more chaos in the 2D composites, and to use as reflection/area lights in the scene.

Creating the sequence foundation
To help smooth over the transition to our new energy conserving shader set, we decided, like on Iron Man, to have our sequence supervisors set up the environment maps and lighting instruments for the sequence. Beyond that, 'though, we had the sequence supervisor actually run a first lighting pass for each of the shots. One of the benefits of the new approach was that the shots looked better out of the box just using the environment maps, and with the lead TDs setting up the shots, we'd always have a high standard first take. It would also mean that if need be we could get a real-looking render out even during animation if we were having trouble getting the animation approved by the director – to remove one more variable.

The sequence lead would start by finding a hero HDRI image set for a shot or group of shots, and have that stitched into a lighting sphere. All the bright lights would be painted out and either images of the lights from the sphere image, or images of the different pyro events we'd filmed separately would be used. The sequence lead would document the different environments and lighting instruments available to the artist.

Choreography of the lights
When the TD picked up the shot, they found a well set up shot file including the environment and many of the lighting tools they needed already set up and running. But they still had to position the lights and start using the running footage to match the chaotic lighting going on in the background. This was greatly helped by a hardware rendering tool Pat Conran wrote for the show called layer cake, which gave fast GL previews of the lighting.

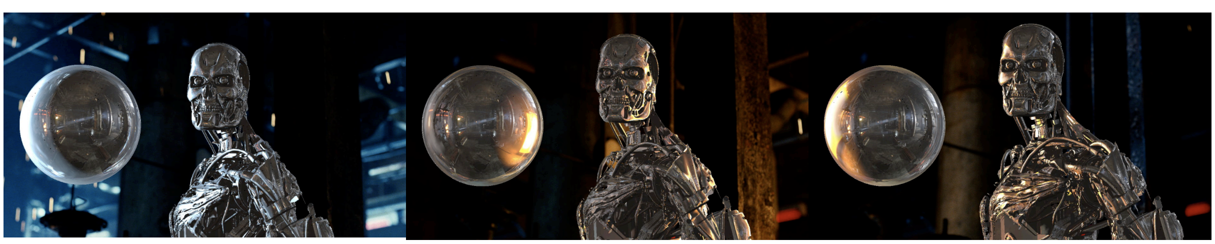

IMAGE: A sequence of frames from hardware preview. Note the element cards in the spheres.

This approach to lighting was new to us all, and as a diagnosis tool we added a feature that rendered a chrome sphere automatically centered in front of the camera for the shot, and baked it into the movie files we used to review the shots in dailies. This gave me a sense, at a glance, of the lighting choreography that the artists were using, and is a tool I really like to have in the tool set.

Iron Man 2: Putting the new shaders and lights to work.

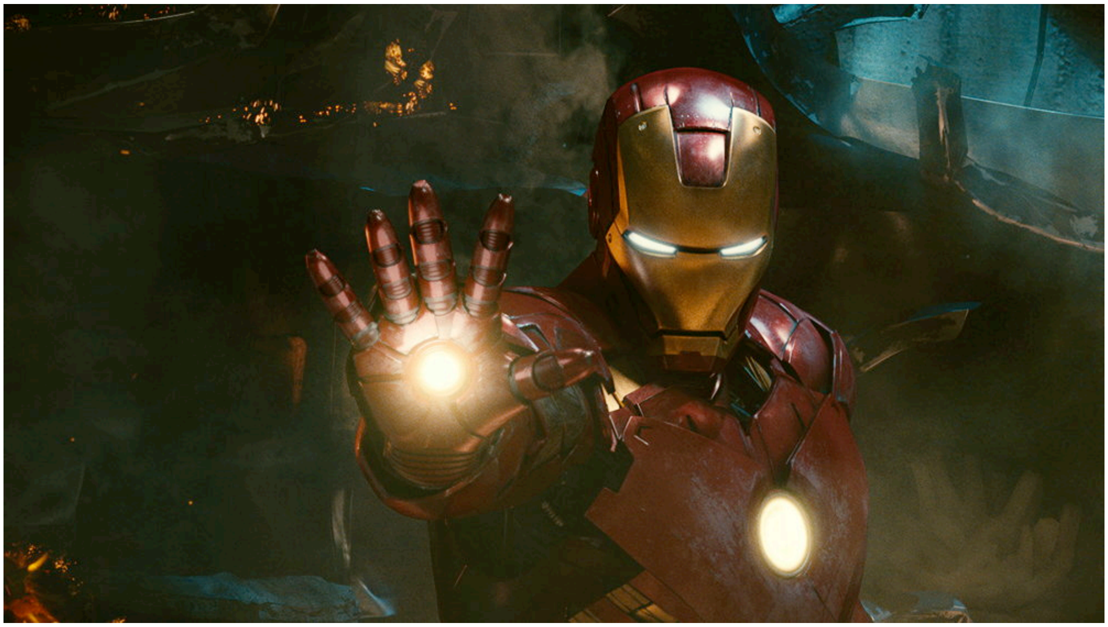

Going into Iron Man 2, we were able to put the new energy conserving, image-based tools to work on the whole production. The shaders and lights themselves had matured, and we'd removed some of the hacks from the transition phase so that, for example, you couldn't override the falloff on lights as easily. The first thing we did was to update the materials for the Iron Man Mark 2 and 3 suits seen in the previous film to our new shading setup.

IMAGE: Iron Mark 2 comparison.

IMAGE: Iron Mark 3 comparison.

New tools to make image based lighting easier

Doug Smythe and the Iron Man 2 team created a set of tools and procedures to create a smooth end-to-end image based lighting pipeline for Iron Man 2.

The first was a tool to organize, stitch and publish the HDR environment sphere images.

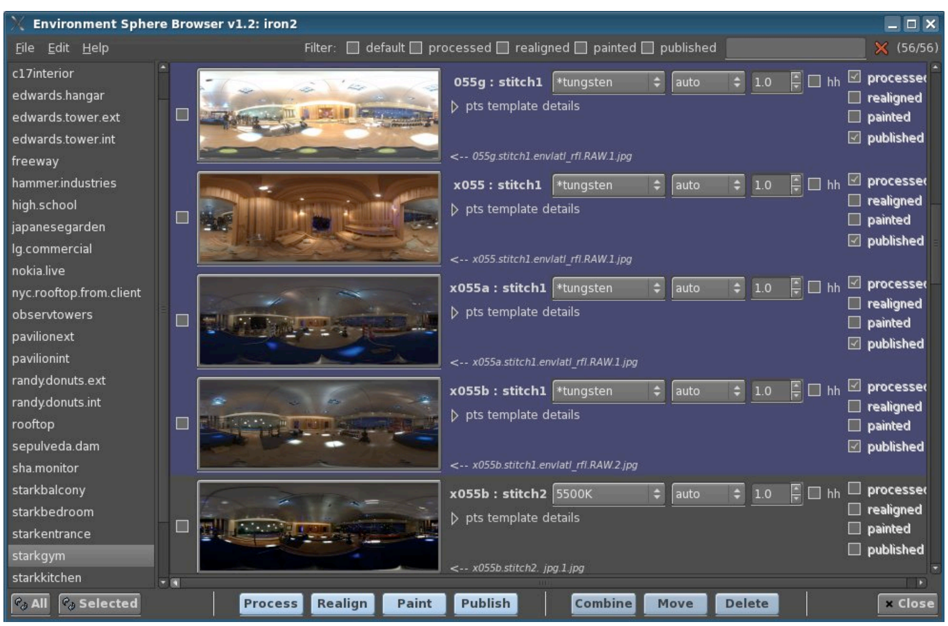

The Environments Browser.

Once stitched we used another internal software tool which we updated to load the env sphere image and semi-automatically create a lighting rig which included textured area lights created for bright areas in the image.

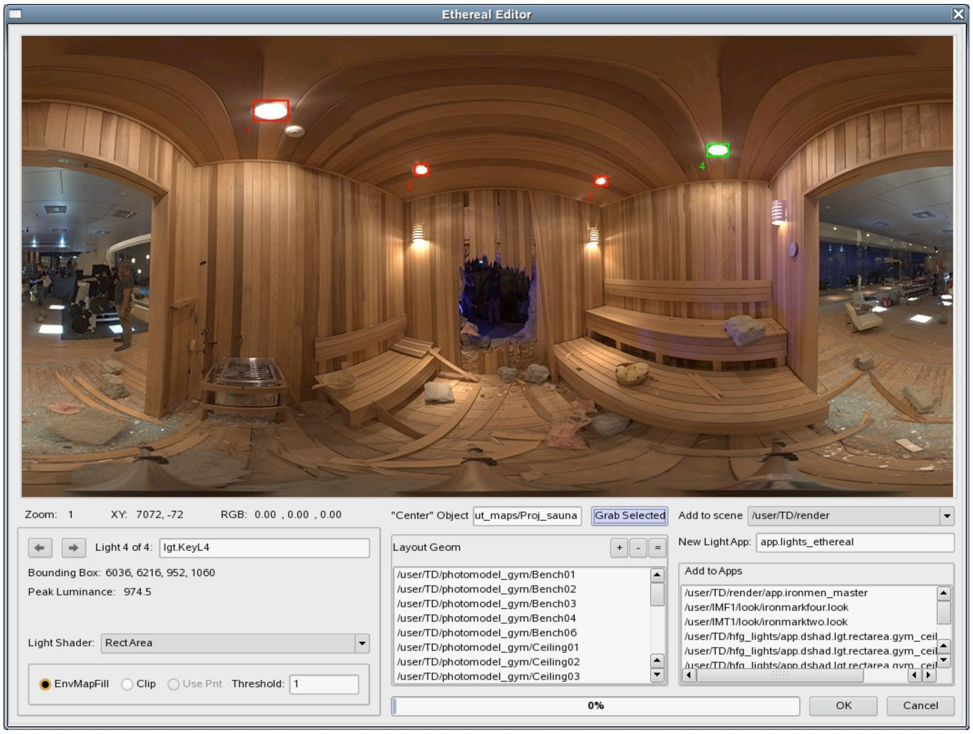

The area light extraction/light rig creation tool. Squares represent lights tagged for extraction.

The artist can manually tag the lights they want removed from the environment map and recreated as area lights. Our matchmovers modeled the whole of the sets in rough geometries, and position the HDR env sphere was photographed from. The light rig tool traced a positon out onto the set to get a rough light position and scale for the area light.

Finally a tool was created to create local environment sphere textures and simultaneously output data in the form of pointclouds or brickmaps, or create textured area light representing geometry that the CG character interacts with so indirect bounce and reflections are locally accurate. Shadowing and local illumination effects from other CG characters could be included in these passes.

Matching the lights that they're using on set

Some of the scenes in Iron Man 2 were lit with the practical suits on set and with the lights were set up by the D.P. with the suits in mind. While this is an ideal case, and certainly such practical references should be encouraged where possible, in Iron Man 2 there were circumstances where the set lighting wasn't perfectly balanced for the suits. In this case we were able to look at times where the D.P. had lit the suits and use images of those lights as lighting instruments in our other scenes.

An example was a small Kino light he'd use to get kicks off the practical suits. It had one warm and one cooler Kino Flo tube in a reflector box a couple of feet wide. We created a CG version of this based on an image of the light to use in scenes like the end battle to add kicks on the suits.

Issues with image based.

In general the new image based tools worked out very well indeed. It certainly made it faster for artists to get a shot looking good pretty much out of the box, and allowed us to focus on the more creative lighting challenges for the show. The image based tools worked particularly well for the daylight exteriors on Terminator Salvation, and for the dynamic factory environment where the lighting was driven by pyro events that we could reproduce on cards. Combined with the normalized importance materials, the image-based approach definitely gives us more realistic results earlier in the rendering process than we've had before. However, there are some issues involved in their use.

- There is a learning curve for artists used to the traditional approaches. In lighting a shot I did for Iron Man 2 I found the new approach to be terrific. I was able to cut the lighting instruments out of the HDR environment capture instruments, position them in the right places and effectively start with a scene in my computer that matched with what I'd seen on set. But I've done a lot of set work and was familiar with the tools used. Some of the artists took to the new setup well, others found adapting to losing some of their previous cheats a little harder to adapt to.
- The better your image collection on set, the easier the path to real-looking renders. We shot environment spheres for most shots (tips on the way to shoot these is below), but in addition I recommend shooting images of the lights themselves used.
- When editing the HDR environment images, make sure you don't lose dynamic range.
- Beware the infinite nature of the lights coming from an environment sphere. If you use an HDRI sphere and a character isn't centered in the location the sphere was shot from for the whole shot, the reflections and illumination can start to play up. That's because the light intensities were correct for that spot, but the light images, being infinitely far away if you're just using the sphere, will remain the same wherever the character is positioned. This was a problem in the kitchen for the House Fight in Iron Man 2, where there were a bunch of recessed lights in the ceiling. These looked great on the practical suit reference pieces we were able to shoot on set, but in our CG environment spheres our lights were not, of course recessed, so when a character moved away from the light he was centered under he was still being illuminated by that light to the same intensity. Using a more hand-built combination of lights and reflectors as we'd done for the comparable Tony's workshop on Iron Man 1 was ultimately a bit more successful. In these cases you can recreate the lighting more correctly in the computer to match the scene rather than relying on the sphere alone. Using projections onto geometry can help as well, but, again, you're limited by the infinite nature of the lights in the source material. In a couple of shots from the Kitchen Fight of Iron Man 2 we actually painted out a bunch of the lights on camera.
- Don't forget the aesthetics. Its great to have your Iron Man look more real out of the box, but remember, the box may not always be lit with him in mind, not with him doing this particular

action. Use cards and reflectors and flags like they do in the real world to shadow offensive lights in the environment or bounce more light into the scene.

Practical tools for recording the lighting in the real (filmed) world

Filmed references

When shooting the scene, even with us using high dynamic range images as our most important tool for capturing information about the lighting and environment, it's still important to shoot visual effects references.

"Setiquette" and relationships on set.

Film shoots are fast and furious, but the crews prefer clear decisions, consistency and predicatability. You need to make sure you have relationships with the right people to get you the information you need, and you want to get people used to what you're going to need to document you environment. Your data gathering crew has to be fast and ready to move as soon as the opportunity to shoot references presents itself.

As a vfx supervisor, I usually communicate with the director and director of photography. However, the most important relationship for getting you the references to help you recreate the reality of the lighting on set is with the first assistant director and their team. You should make sure they're aware of what passes you need as soon as you can. The camera team, particularly the camera operator and first a.c. are important – you'll often want the camera operator to provide a "human motion control" reproduction of his camera moves for a clean plate, and you'll need the a.c.'s help to stay on top of camera information – lens, focal distance, and stereo information etc. Likewise the grip crew who control the cranes and dollies and who can save your butt with a piece of bluescreen thrown up on short notice. The gaffer and his crew are also people you want as allies – you want them to keep the lights and reflectors and diffusers in place for shooting your references and HDRIs so the scene matches what it was. The script supervisor will often keep track of camera information as well, and you can help each other keep track of information. You can keep the script supervisor informed of why on earth they're shooting all these passes which can help you down the track as editorial tries to sort out what to send.

The film set is a complicated place, with an odd sort of hierarchy. You and your crew need to be sensitive and away of what is going on, to stay out of the way as much you can, but to make sure you get the references you need.

What references to get

You certainly need an HDRI that is relevant to the lighting of the scene, and I'll talk a bit more about that in a moment. Beyond (and often before) that, I recommend you also capture the traditional light probe references. Spheres, a Macbeth chart possibly, and ideally some reference related to what you're going to be creating.

Chrome/Gray Spheres

A couple of years back I had a very talented visual effects supervisor tease me for wanting chrome and grey spheres even with shooting the high dynamic range image sets to document the scene. Its true that the HDRI is our most important reference but the chrome/grey spheres offer the following advantages:

- They are on the medium that the plate material is on, and if correctly captured and converted to your image data format, they represent the same colors as that medium. Thus, and this is the important part, you can use them to calibrate the color of you HDRI which is often shot in a different way using different equipment.
- They are a fast way of getting up to the moment information. While it might be hard to perfectly preserve the lighting, position of the performers etc. with an HDRI, you can often quickly get the chrome/grey sphere into place.
- It's insurance. In case you don't get and HDRI or the data gets messed up.
- They are moving footage. You can capture dynamic lighting setups with fire puffers, steam, strobe lights etc.
- They are familiar. More on this below.

There are several types of chrome grey spheres, and you can pay quite a lot for different ones. The one I usually use is a portable half-chrome/half-grey setup that you breaks apart for transport. Its quick and light and lets you do the job efficiently.

I've also used slightly larger separate full chrome and full grey spheres, which are good for multiple camera setups and for moving through a scene, and big ones for long lens and air to ground type scenarios.

Guidelines for shooting the chrome and grey spheres are as follows.

- Get the crew into the habit. I mentioned the importance of predictability and consistency. You want the crew to expect and indeed call for the spheres and digital stills if the shot is a visual effects shot.
- I generally ask for reference passes if I'm adding computer graphics to a shot. Even for bluescreen elements it can be relevant to shoot the spheres at least - for example where you're adding a 3D prosthetic or piece of armor to the person of object you're shooting against blue. The spheres can help in matching the bluescreen lighting to the background (hopefully you've shot the plate first and can use the sphere pass from that to help guide the lighting on the bluescreen element).
- Make sure the lighting is consistent with the plate photography. The references are not much use if the grips disassemble the reflector cards and the electric crew turn off the key light for the reference passes. Ideally everyone should hold their positions for the references as if you're shooting another take, even the actors in some cases. Don't expect a whole lot of success on the latter, but sometimes you can get in there on at least one of the setups and capture a representative HDRI and refs with performers or stand-ins. With some D.P.s this can be a nightmare because they tweak the lighting with every take. At a certain point you have to get what you can.
- Make sure the spheres are as big as possible in frame. If the camera is on a zoom, have them zoom in, but do that AFTER shooting any lens calibration references you might be using. If it's a prime – which is what we prefer for visual effects work – you might need the big spheres
- Make sure the spheres are in the right spot. It's not very useful to shoot spheres that aren't in the lighting that represents where the computer graphics will go. Of course they're not a ton of use if they're tiny in frame either. If it's a relatively consistent lighting environment (eg daylight exterior without a bunch of reflectors etc.) I use two sets – one back where the cg will go and one closer to the camera so its big in frame.
- Hold the sphere out in front of you. You don't want to see the person holding the sphere reflected in the sphere, particularly if you're going to unwrap it for an environment map.
- Don't shadow the spheres with your body – again you want the sphere to be in the lighting the added object would be in – and that probably doesn't include you.
- Make sure the lights go through their performance, ideally with the sphere and camera in a locked position. For the chrome sphere its useful to have it static so you can pull an environment map if you wish.
- Move the spheres through the space if the computer graphics object moving through the space. This is where separate spheres are handy – you can hold them out from your body and apart and offset so they don't shadow one another and the grey doesn't take too much of the real estate in the chrome reflection.
- Camera move is a judgement call. You want some moments of camera and sphere stable, particularly for a chrome sphere if you're pulling an environment from it.

- Be ready, and have your crew ready, at all times when they're shooting VFX shots. You need to get in there as soon as they announce they're moving on. If there is something non camera or lighting related that delays the shooting the A.D. might want to grab the references while waiting. This is not always ideal since things can change but you can always call to shoot them again if things change enough to justify it.

- When in doubt, shoot refs. If the director swears blind it won't be visual effects but you see how it could, its usually worth shooting the refs, particularly if by now your on-set crew is a polished machine. If not, make sure the data gatherer and script supervisor note the nearest set of references that could at least inform this take.

- Don't be sloppy – make sure you and your crew are shooting the spheres properly and that you don't just dash them off, even knowing there is some forgiveness if you're using them to just calibrate the HDRI

How it goes down on set

Before the take. If something in the take is going to adjust the set irrecoverably (eg its going to be blown to bits), you might want to shoot references and even a clean plate before they do the actual effect.

As mentioned, be ready to swoop in with references when the A.D. calls for them. Make sure you tell the A.D. you want references and ideally what they are for every setup at some time before or during the photography.

You need to keep an ear out – ideally you're at or near video village. If the director seems to be happy with the take, get ready to get your team out there with the refs. If you hear the "cut, we got it, moving on" type call, you need to swoop in. Sometimes a call like "spheres " (crews sometimes like "balls" better), "hold the lights" etc. may be necessary. If you're lucky the crew will pretty soon be calling for the spheres themselves. In general, have them shoot them for VFX shots to get into the habit. Don't try and be too nice a guy and overthink it – remember consistency, and also take care in what you shoot – so they know to take care as well.

For complex shots particularly they should be informed ahead of time, but a lot of what we do is on the fly, so you certainly need to make sure they know what you'll be asking for as soon as they start doing takes of a scene. If it's a big special effects pass or stunt, you might request a clean plate or some references before the first take that would mess up the shot, so that you can capture the environment that you'll be adding your computer graphics to in its "before" state – there is often no going back.

Capturing HDRIs on set.

There are several ways of capturing HDRIs on set, but we try to keep a fairly simple approach, mostly in the interests of flexibility and speed. The rules for capturing lighting references apply here.

In theory we're trying to capture the whole dynamic range of the scene, including the brightest lights including the sun. In practice we often end up replacing the brightest lights with lights we have control over, so there is some forgiveness. We've arrived at a resolution of no less than 8k around for the final stitched environment sphere

Not everyone will agree with me on this, and there are several approaches to capturing high dynamic range images on set, including some quite cool semi-automated approaches like the box device created by Hoyt Yeatman.

However, we've found that sending an on set data gatherer out with a simpler setup as described below works very well. It is reliable, flexible and, importantly fast. It is also film set friendly, and sometimes we'll be able to take an appropriate image off to one side or sneak in and grab it between takes, saving further time in the expensive day of shooting.

Recommended equipment:

Canon 1DS mk3 Camera body. This model has a full-frame sensor and with an 8mm lens you can capture a scene with three directions. Other cameras have smaller capture areas and require multiple directions to cover a full sphere.

Sigma 8mm fisheye lens

Nodal Ninja

Tripod

Remote shutter trigger for camera body.

0.6 ND (2-stop) filter for lens.

We setup the camera as shown with the 8mm fisheye lens in the Nodal Nija head mounted on the tripod. Position the lens so it is as close to nodal as possible.

We shoot in Manual exposure mode with manual focus and image stabilization off if the lens or body has it. We remove any extraneous filters or rings from the lens.

In good, not too rushed shooting conditions we'd shoot the following for a direct-sunlit exterior environment:

- 7 exposures, 3 stops separation, center exposure 1/32 sec
- Aperture f/16, ISO 100
- Add 0.6 ND (2-stop) filter to the lens (you don't want this for interior or reference shooting, ‘though.

We adjust to taste in darker situations or situations where we need to move a little more quickly than this would entail (eg the director or AD is shouting at you). Aperture decreases and ISO increases are usually the first things that get changed in the heat of battle.

ACKNOWLEDGEMENTS

I'd like to thank the following people. Pat Conran, Doug Smythe and Hayden Landis all of whom generously allowed me to shamelessly crib from their various show reports, ILM internal documentation and previous courses. Pat's normalized importance documentation was particularly helpful and supplied many of the technical renders in the section on Terminator Salvation. Also Christophe Hery, Naty Hoffman and Philippe Rebours for their valuable advice and suggestions and Lanny Cermak for modelling.

Reflection Occlusion was originally developed by Stefen Fangmeir on Speed II with contributions by Barry Armor, Dan Goldman, Ken McGaugh, Will Anielewicz and Christope Hery.

Ambient environments were developed largely by Hayden Landis, Ken McGaugh and Hilmar Koch with refinements by Kevin Sprout, Dan Goldman and Doug Sutton.

Christophe Hery and Simon Premoze developed the BRDF and anisotropic tools for Iron Man and implemented much of the Normalized importance shaders and lights on the R&D side, with Pat Conran, David Meny, Philippe Rebours making big contributions on the production side.

Thanks to the ILM crews on all the above productions, and the terrific support staff and R&D team.

Finally, thanks to Siggraph organizers for providing this forum for learning and sharing and to my fellow speakers for their enthusiasm and inspiration.