SIGGRAPH2012

The **39th** International **Conference** and **Exhibition** on **Computer Graphics** and **Interactive Techniques**

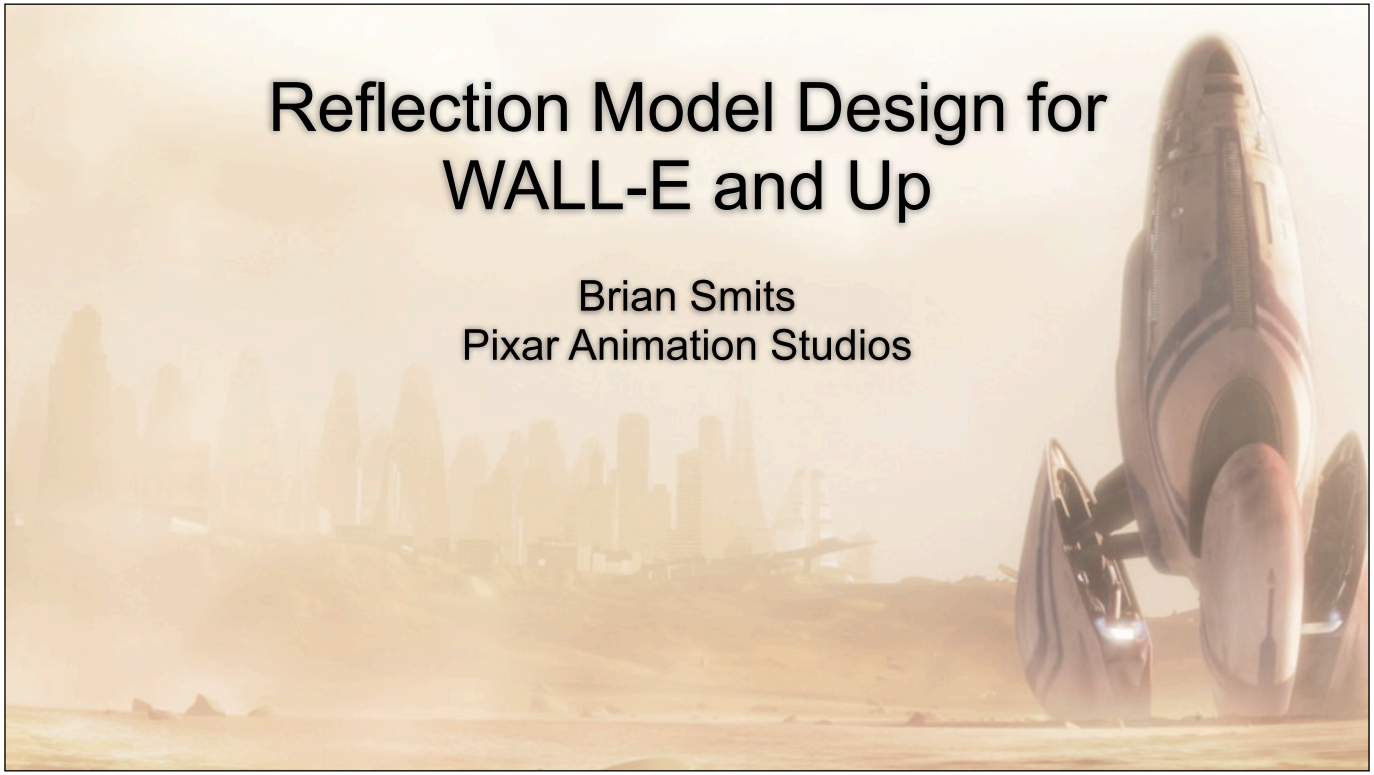
Reflection Model Design for WALL-E and Up
Brian Smits
Pixar Animation Studios

Why Change?

SIGGRAPH2012

- Wanted a grittier more realistic look
  - more ray tracing
  - more metals

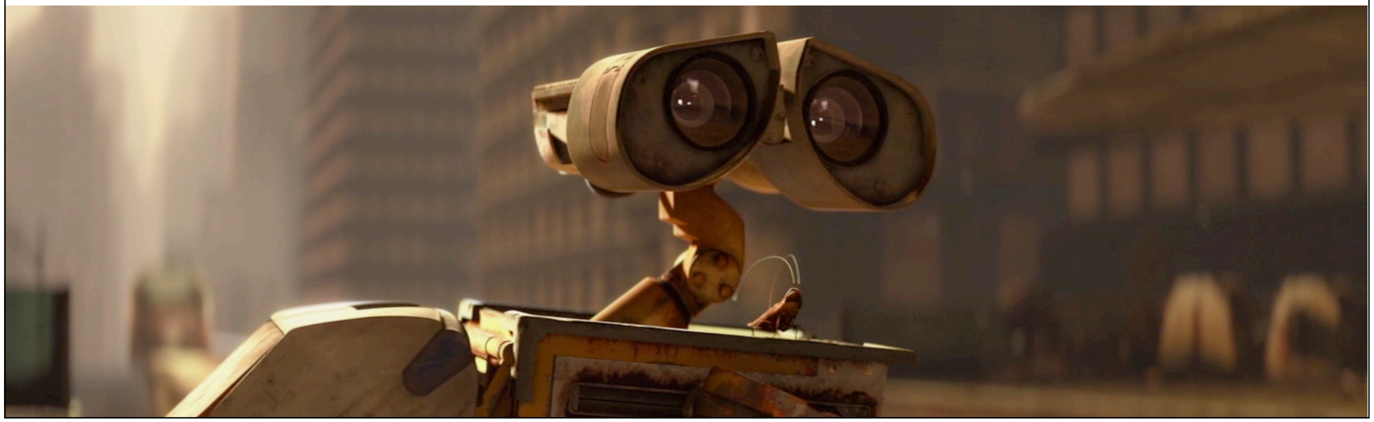

Why Change?

SIGGRAPH2012

- Lots of geometric complexity
- Many locations

Shading Department Changes

SIGGRAPH2012

- Shading more objects in less time
- Tools that increased productivity tended to increase cost of shaders

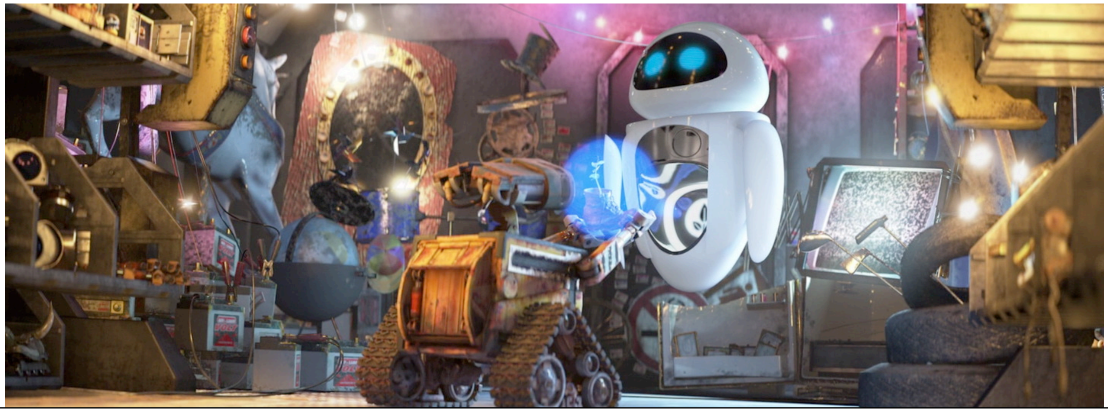

Instead of getting a single object to shade, shading TDs were getting sets to shade.

Lighting Department Changes

SIGGRAPH2012

- More raytracing
  - easier to set up
  - required shading many more points
- Too much time spent adjusting shader parameters
- More GPU relighting

Design Process

- A new BRDF is best developed by working directly with the people who will use it
  - Understand what they need to do and why
  - This was critical for success

Shading requirements are different than lighting requirements, and it's not always possible to satisfy all requirements, so getting the users to buy in to the tradeoffs early is helpful.

Shading TDs care about how easy it is for patterns to drive behavior. Lighters just want it to look good without any tweaking and be easy to tweak if needed.

Weekly (or more) meetings with all interested parties (shading, lighting, technical). Lots of time spent determining exactly how parameters affected the final result, for example overall highlight shape and energy balance. We ended up remapping the roughness the users dialed into the roughness the model took to improve usability. We also ended up with a bunch of emergency overrides to make sure problems could be dealt with very late.

Why Not Pick an Existing Model

SIGGRAPH2012

- We want
  - Broad range of appearances with few parameters
  - Easy to control (directly or with patterns)
  - Ability to cheat
    - Antialiasing done in shader

In many ways the shading TD doesn't care about the choice of the specular lobe or which diffuse model is used, as that is a show wide decision made very early in the production. The size of the highlight at each point on the surface or the amount of diffuse is much more important. Their goal is to use patterns to drive complex surface behavior quickly and easily and make the result look good.

In addition, they need to make sure that as the surface moves further away the patterns they generate don't alias and that the filtered results still work. Ground planes are a good example of something that must filter nicely, but any large object seen at very different distance may have problems. This may involve cheating

In other words we likely want more than one BRDF, plausibly combined, and easily controlled

Goals for a New Model

SIGGRAPH2012

- Fewer and easier parameters
- More consistency for lighting
- Layered shaders are free
- GPU relighting more predictive of final frames

Our old model could easily have 10 interactions with “lights” per layer. Multiple layers could result in 30 to 50 interactions with lights. This was expensive, especially when ray tracing.

Visual features

SIGGRAPH2012

- Highlights
  - Go to white at grazing angle
  - Surfaces have some specular for “rim” effects
- Diffuse
  - Disappears at grazing angles for smooth dielectrics
- Higher dynamic range
  - Pixels more likely to be well above 1

All of these visual goals were things we would get for free by switching to a physically based approach.

This of course meant building a set of lights to drive the reflection model, and an unanticipated consequence was that we also needed to introduce a nice way of mapping over range intensities back down to 0,1. Previously the results of lighting were mostly already in a nice 0,1 range. Now much more of the scene was over range.

Energy Conservation

SIGGRAPH2012

- Helps deliver consistency
- Helps drive realistic partitioning of energy and behaviors when designing the model

While it was important to be at least roughly energy conserving, it was just as important that the principles of energy conservation went into developing the model than that any individual shader maintained energy conservation.   This let a small set of parameters drive more complex behavior.

It was possible to break energy conservation in the resulting shaders, only it took work to break it rather than work to maintain it.

Light can reflect off the interface causing highlights and reflections
Enter the substrate, scatter, and exit causing a Fresnel weighted, view dependent diffuse
Or hit multiple microfacets and scatter incoherently, which we will put into a Lambertian diffuse

We want to control all of this behavior with 3 parameters, color, roughness, and metallic.

Partition Energy

SIGGRAPH2012

- Roughness drives partitioning of incoming energy
- Fresnel effects cause angle dependent partitioning
- Metallic determines
  - specular reflectivity at normal incident angle
  - presence of any subsurface diffuse

Three Lobes

- One energy conserving diffuse lobe
- Two normalized highlight lobes
  - Specular - broad highlights
  - Reflection - sharp, detailed highlights

Specular And/Or Reflection

SIGGRAPH2012

- People used to the distinction
  - lights produced specular highlight
  - ray tracing and reflection maps
- Limit amount of ray tracing
- Blur the lines
  - lights produce both specular and reflection highlights
  - some reflection maps acted as lights as well

Reflections were computed from lights, reflection maps, and ray tracing. Specular only came from lights. We had a light that functioned as both a texture mapped rectangle for reflections and as an area light for specular and diffuse. The texture map had a usable effect on specular and diffuse.

Fresnel

SIGGRAPH2012

- Controls the intensity of highlights and diffuse
- Schlick's approximation
  - Blends surface color to white at grazing angle
  - Generalize this to two colors

$$
\begin{aligned}
F(C_\perp, C_\parallel, \theta) &= C_\perp + (C_\parallel - C_\perp)(1 - \cos(\theta)^5) \\
a + bF(C_\perp, C_\parallel, \theta) &= F(a + bC_\perp, a + bC_\parallel, \theta)
\end{aligned}
$$

We are also going to use heavily the linearity of this Fresnel function to push terms outside the Fresnel function that are independent of the angle into the color weights

Highlight Lobes

SIGGRAPH2012

- Based on dBRDF [Ashikhmin, Premoze]
  - Gaussian
  - Convert cone angle c to power p(c)

$$
\begin{aligned}
G(c) &= \frac{p(c)e^{-p(c)(1-N\cdot H^2)}}{4(N\cdot L - N\cdot V)^2}\\
S(C_\perp, C_\|, c_s) &= F(C_\perp, C_\|)G(c_s)\\
R(C_\perp, C_\|, c_r) &= F(C_\perp, C_\|)G(c_r)
\end{aligned}
$$

We want something simple to work with.
Gaussian lobe for the distribution and a modified shadowing function/normalization term so that G * N.L integrated over the hemisphere is closer to 1 for most roughnesses.
Artists don’t care about the exponent value, they care about highlight size. We also want highlight size to vary nicely (perceptually uniformly) with changes in roughness, so we express highlight size internally as a cone angle and convert to a power when evaluating the lobe.
Cs-perpendicular is a blend between a default dielectric (.05) and the surface color based on the metallic parameter.
Cs-grazing defaults to 1, or white

Can include an anisotropy term as well, with no need to constrain the tangent vector to lie in the plane of N by playing around with trig identities

Partitioning Specular Energy

SIGGRAPH2012

- Roughness r, reflection weight w(r)
- Highlight fade as roughness increases
  - force total energy to be 1-r

$$
\begin{aligned}
(1-r) &= (1-r)\int (w(r)R(C_\perp, C_\|, c_r) + (1-w(r))S(C_\perp, C_\|, c_s)) \\
&= \int (R(C_{r\perp}, C_{r\|}, c_r) + S(C_{s\perp}, C_{s\|}, c_s))
\end{aligned}
$$

We want to combine two lobes, one for specular and one for reflection based on roughness, and we want the total energy to be less than 1-r. We also want to push all of these external weights into the color weights to the lobes. This gives us a simple term for the amount of energy in each lobe, since the lobe itself integrates to 1

Subsurface Diffuse

SIGGRAPH2012

- Non-Metal
  - Heavily affected by Fresnel effects
  - Reciprocity means both light and view matter
- Metal
  - none

$$D_s = C(1 - F(C_\perp, C_\parallel, \theta_V))(1 - F(C_\perp, C_\parallel, \theta_L))$$

We use the same basic Fresnel attenuated diffuse as in the Anisotropic Phong paper by Ashikhmin and Shirley and other places.

C-perpendicular is a blend between a default dielectric (.05) and 1, or white) based on the metallic parameter. C-grazing defaults to 1, or white.

Multibounce Diffuse

SIGGRAPH2012

- Non-Metal
  - as roughness increases, from multiple interactions
- Metal
  - only from multiple interface interactions
- Lambertian
- Use same color for shader writer sanity

$$D_{mb} = C$$

A dielectric when made rougher (say with sandpaper on a smooth dark plastic) tends to get less saturated in color due to more interactions with the interface scattering light back towards the viewer. For a metal multiple interactions should make the diffuse more saturated. For a shader writer who is supposed to make the surface hit a specified RGB value this is undesirable behavior.

Combining Diffuse

SIGGRAPH2012

- Blend based on roughness r
  - Subsurface to balance with the highlights
  - Multi-bounce to include the rest of the energy

$$\begin{aligned} D &= (1-r)D_s + rD_{mb} \\ &= C((1-r)(1-F(C_\perp, C_\|, \theta_V))(1-F(C_\perp, C_\|, \theta_L)) + r) \\ &= F(C_{d\perp}, C_{d\|}, \theta_L) \end{aligned}$$

The first diffuse term gets the same (1-r) weight that highlights did so that they balance each other

Examples

Transmission

SIGGRAPH2012

- Three lobes for transmission
  - transmitted diffuse
  - transmitted specular
  - refraction
- Connected to subsurface scattering

Transmission Examples

SIGGRAPH2012

BRDF Lobe Parameters

SIGGRAPH2012

	Reflected	Transmitted
Diffuse	$C_{d\perp}, C_{d\|\|}$	$C^T_{d\perp}, C^T_{d\|\|}$
Specular	$C_{s\perp}, C_{s\|\|}, c_s$	$C^T_{s\perp}, C^T_{s\|\|}, c^T_s$
Reflection	$C_{r\perp}, C_{r\|\|}, c_r$	$C^T_{r\perp}, C^T_{r\|\|}, c^T_r$

12 colors, 4 floats = 40
20 floats for front side, 20 more for back side
Users never see this. Not a reduction in data, but important for layering.

Additional Parameters

SIGGRAPH2012

- Additional terms to adjust/cheat
  - Scales on Fresnel effect at normal and grazing angle
  - Tint for specular/reflection (iridescence)
- Scaling terms
  - Diffuse, Specular, Reflection
  - Specular and Reflection highlight size
- Corresponding set for transmission

The Fresnel scales could be used to adjust the default reflectivity at grazing angle (effectively an index of refraction term) and the other could be used to dampen the grazing angle intensity

The anisotropy vector could be used to get the wood shading results of Marschner et al.

There were a corresponding set of scaling and tinting controls for transmission

All of these affect one or more of the lobe weights

Examples

Transmission Examples

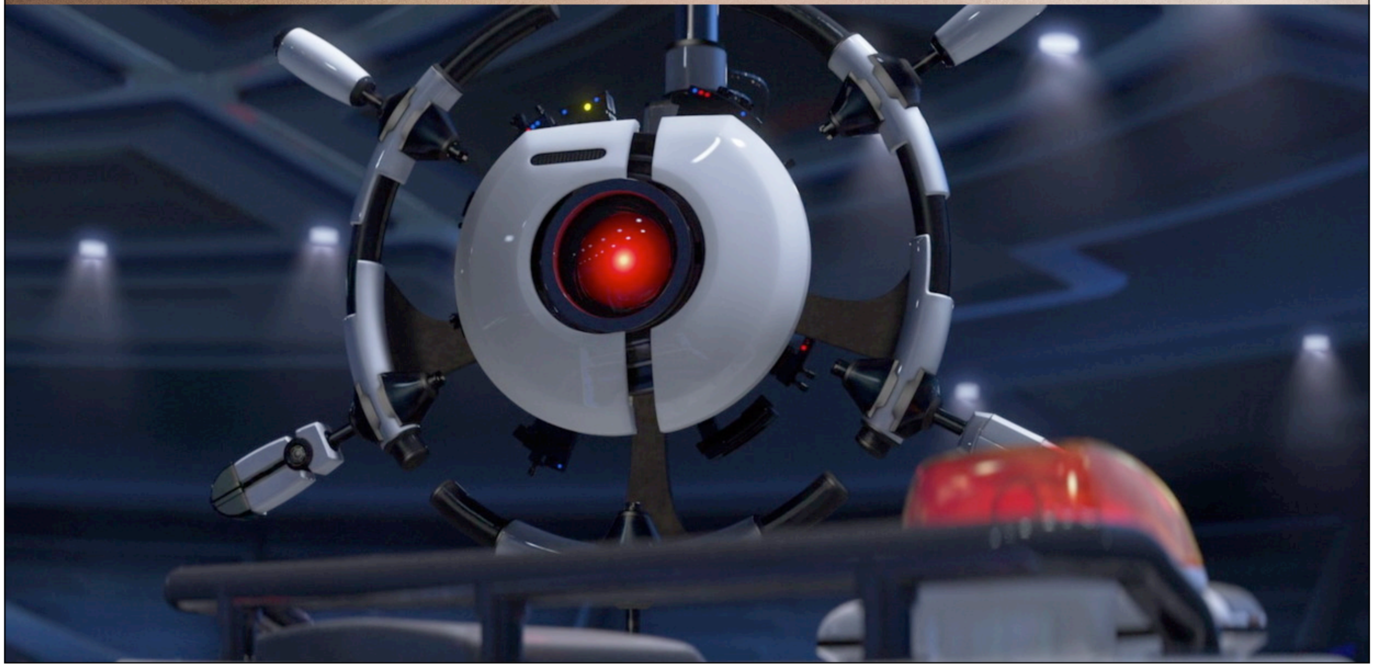

Transmission Examples

SIGGRAPH2012

Could turn all transmitted energy weights into an opacity, allowing cheap thin transparent objects that had a view dependent transmission, since the amount of transmitted light goes down at grazing angles.

Layering

SIGGRAPH2012

- Support many layers
- Give a plausible result
- Start simple, layer based on alpha/presence
  - This is most of the uses of layering

Layering

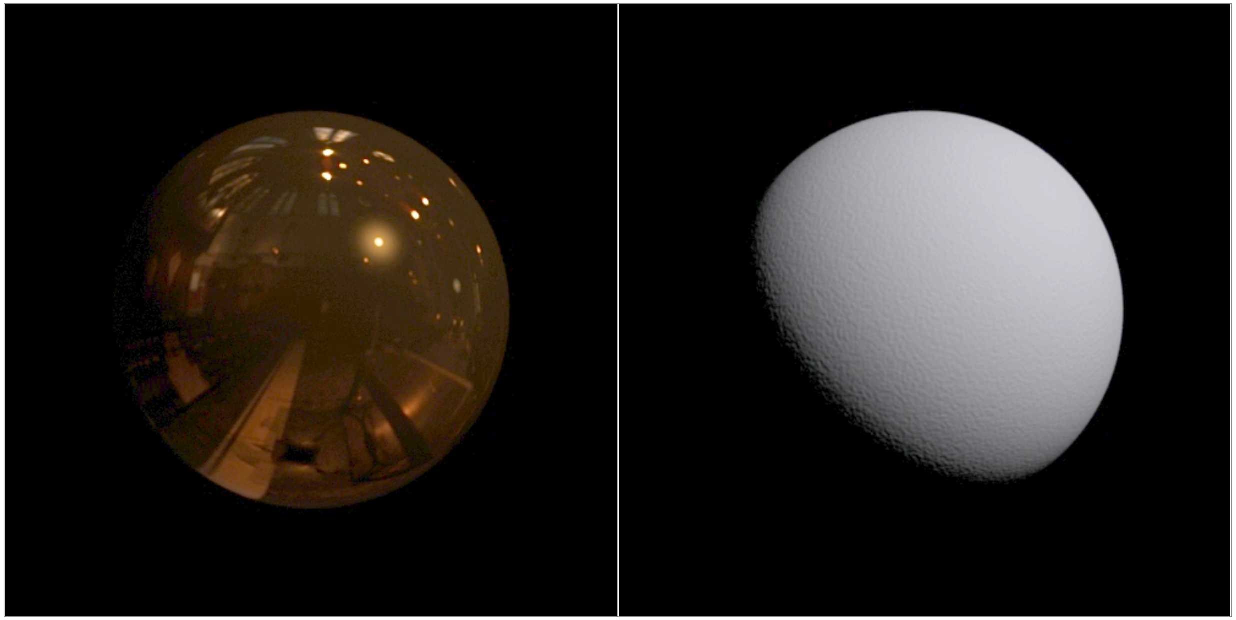

Parameters versus Illumination Results

SIGGRAPH2012

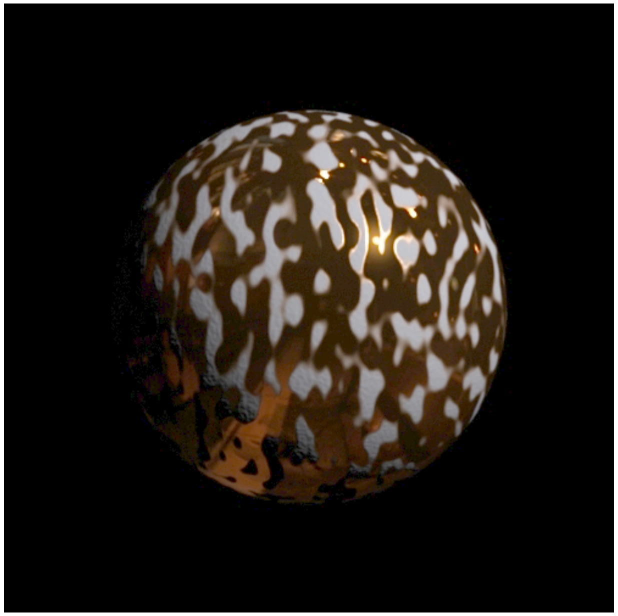

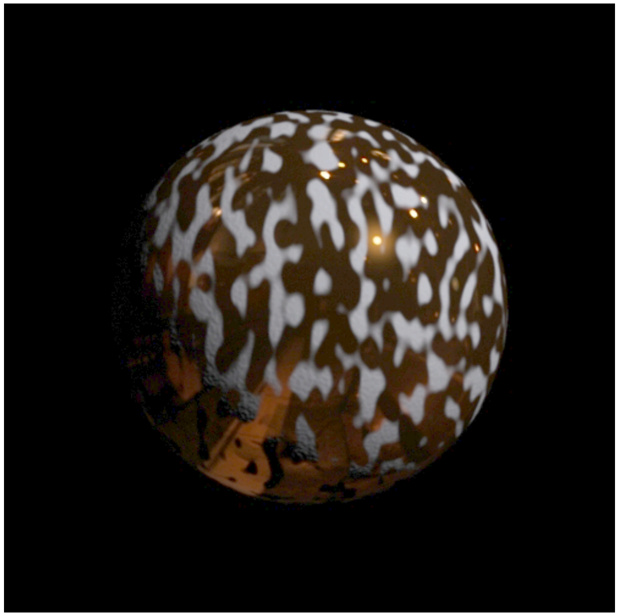

Parameter blending shows artifacts in transitions.
Illumination blending becomes arbitrarily expensive as lighting cost increases proportionally with the number of layers, also, the amount of data needed increases with the number of layers, which is a problem for GPU based relighting
Neither approach is acceptable.

Parameters versus Illumination Results

SIGGRAPH2012

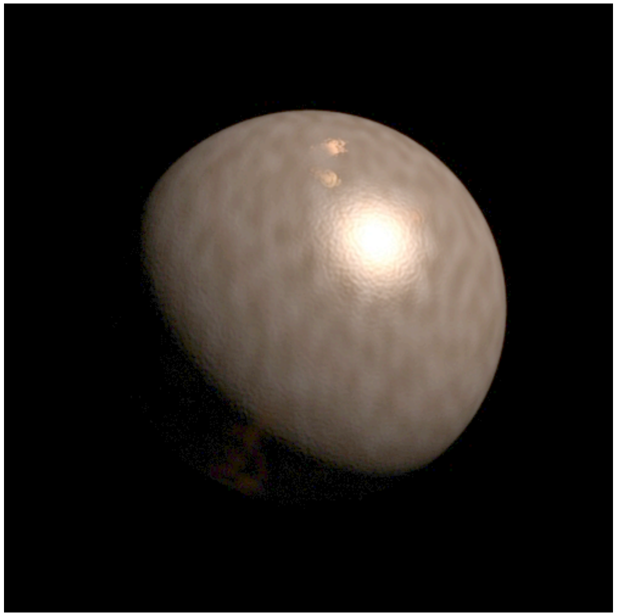

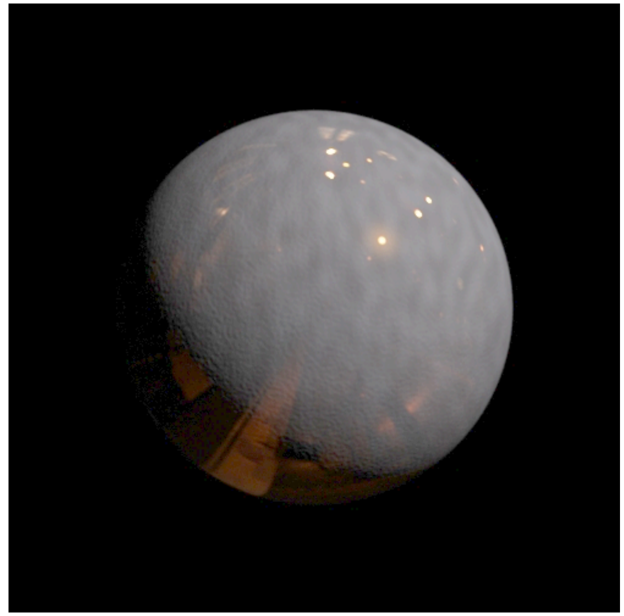

Parameter blending shows artifacts in transitions.
Illumination blending becomes arbitrarily expensive.

Blend Visual Phenomena

SIGGRAPH2012

- Want dust over copper to have a diffuse term and a reflection term, but not a specular term
- Blend weights on lobes and on cone angles
  - Energy weighted blend for cone angles

We can use the opacity or presence of each layer to blend the color weights on the lobes and the cone angles of the lobes. Blending color weights is very close to blending after illumination is done. It is exactly the same for traditional diffuse or highlights with the same cone angles. Blending cone angles is closer to a perceptual blend than blending exponents would be. An energy weighted blending of cone angles is even better as a lobe with a large cone angle but very little energy will have essentially no effect on the results.

Keep Separate Normals

- Do an energy weighted blending for these normals as well
- Might want 3 different normals, we used 2

Layer Based On V Direction

SIGGRAPH2012

- Minimize data and some computation
- Shading artists use dot(V,N) to do layering anyway...

It might be straightforward to work around this restriction here, but it is definitely a bit less straightforward for what we are going to talk about next.

Layering Results

SIGGRAPH2012

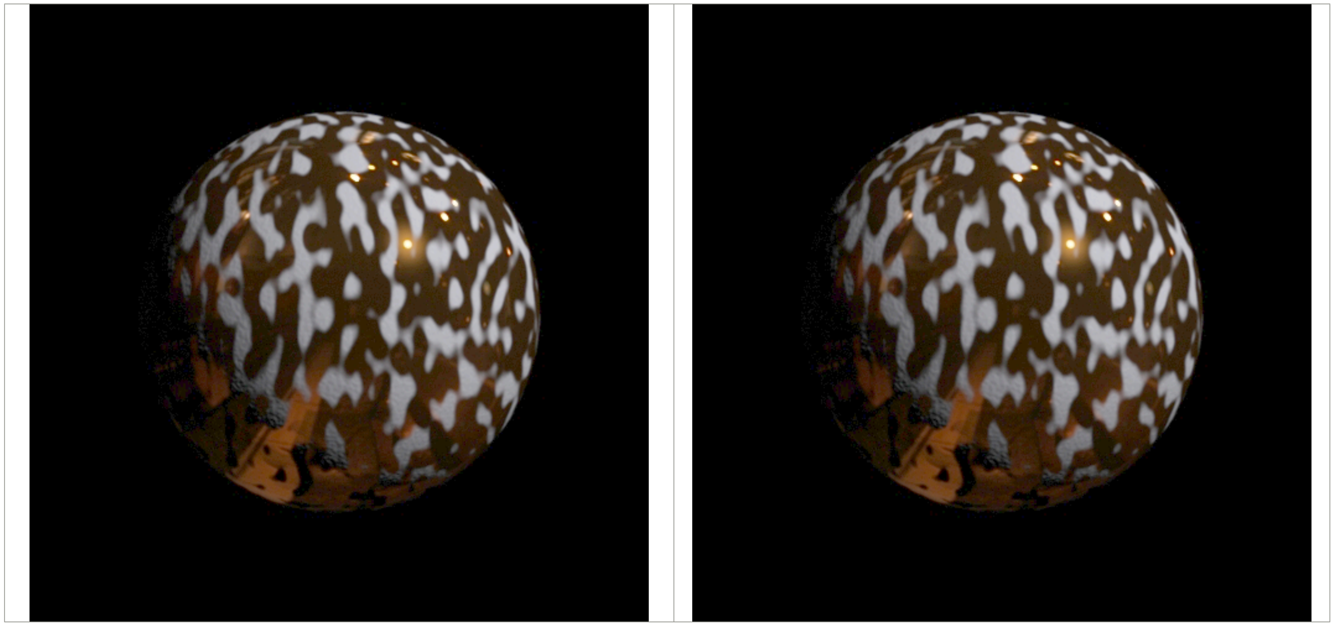

While pattern generation goes up in cost, lighting cost stays exactly the same regardless of number of layers.

We never did this test. The important thing wasn't how closely it matched blending the results of illumination, the important thing was that it acted in a useful and predictable way.

Layering Results

SIGGRAPH2012

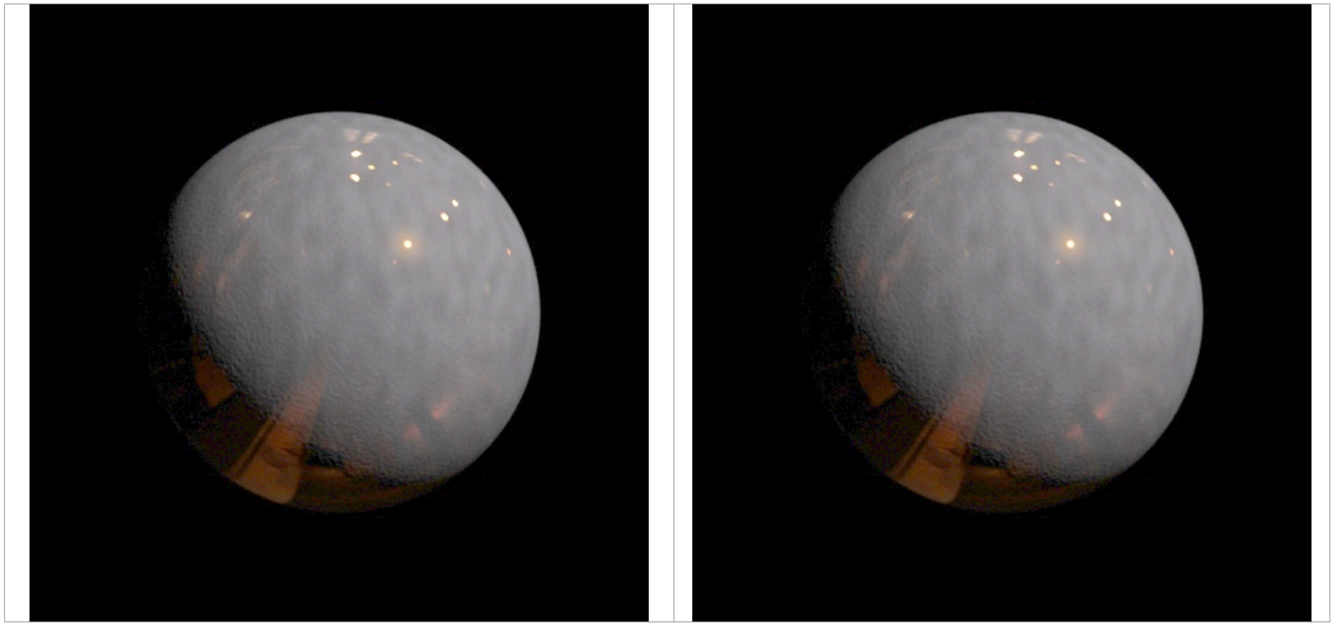

While pattern generation goes up in cost, lighting cost stays exactly the same regardless of number of layers.

We never did this test. The important thing wasn't how closely it matched blending the results of illumination, the important thing was that it acted in a useful and predictable way.

Layering Examples

SIGGRAPH2012

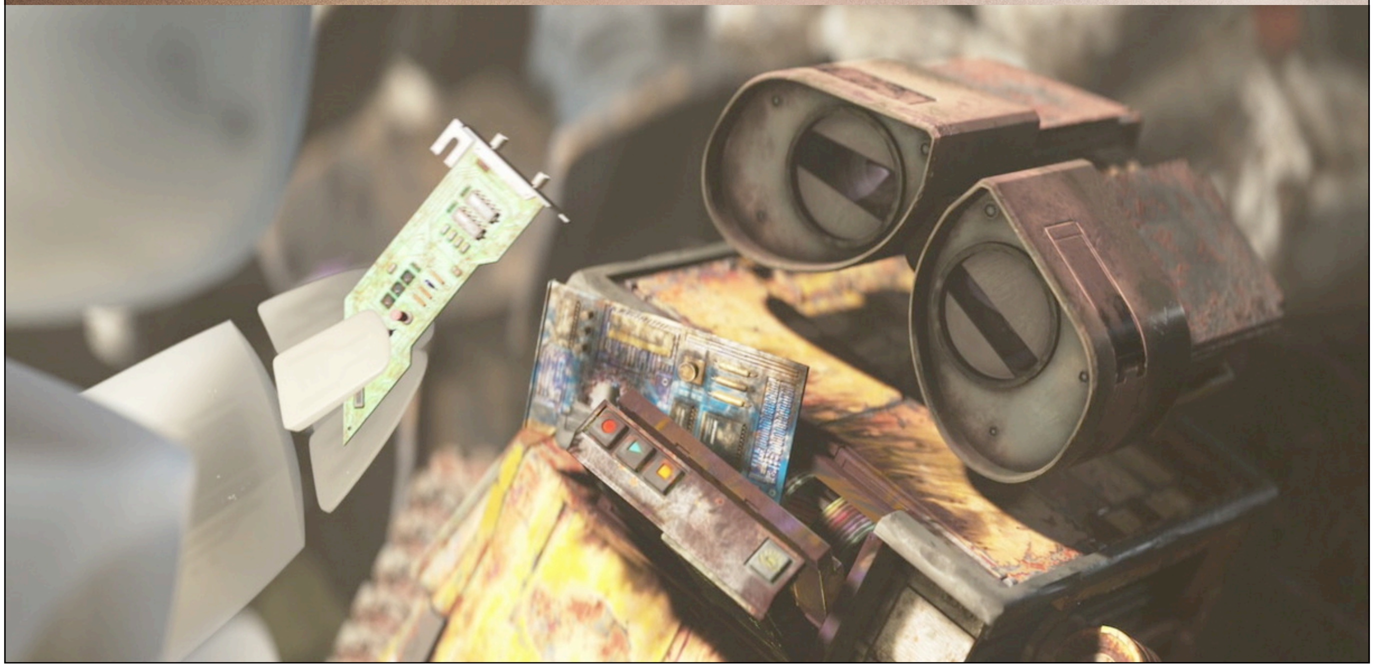

Metal and paint and rust and dirt and oil. Decals often broken out on a separate layer.

Layering Examples

SIGGRAPH2012

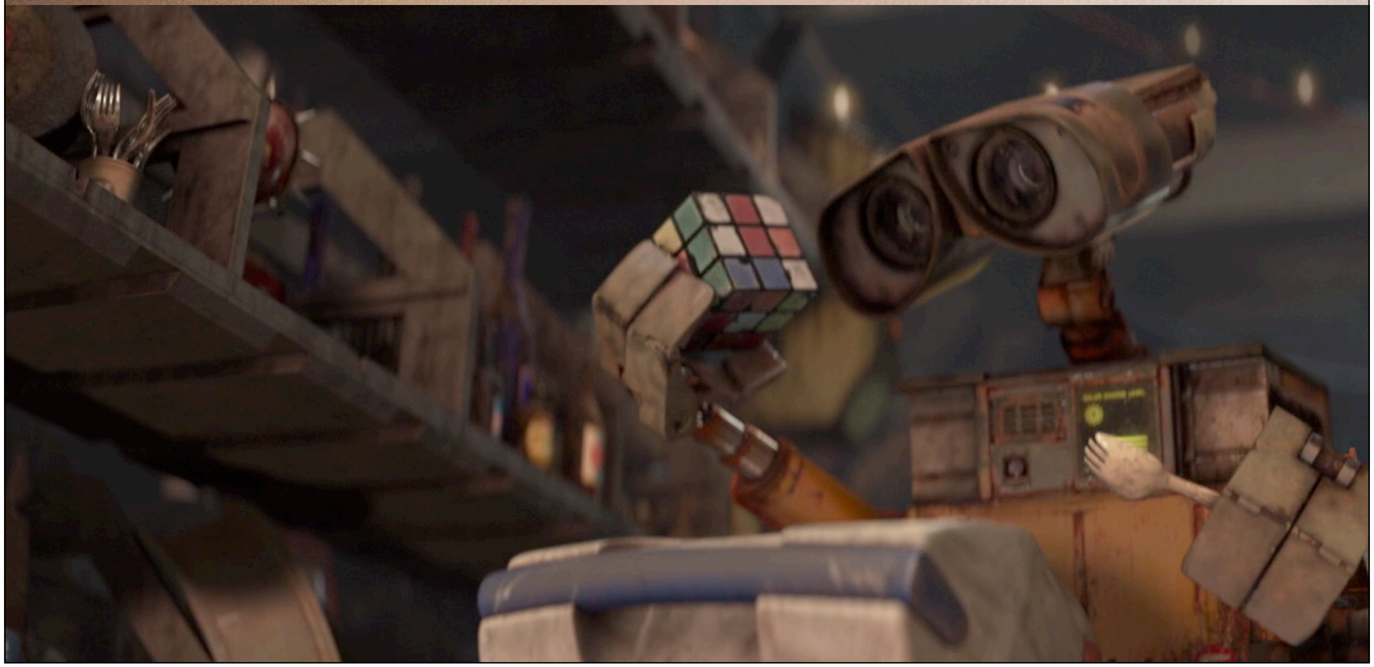

Every object in the Earth portion of the movie had a long history of abuse piled on it that had to come across in the shading.

One Day Later

SIGGRAPH2012

- “How do I put a displaced anisotropic shiny layer under a non-displaced lacquer?”
- “What if the lacquer is chipping, somewhat rough and yellowing with age with some drops of water on it?”

Simple Example

SIGGRAPH2012

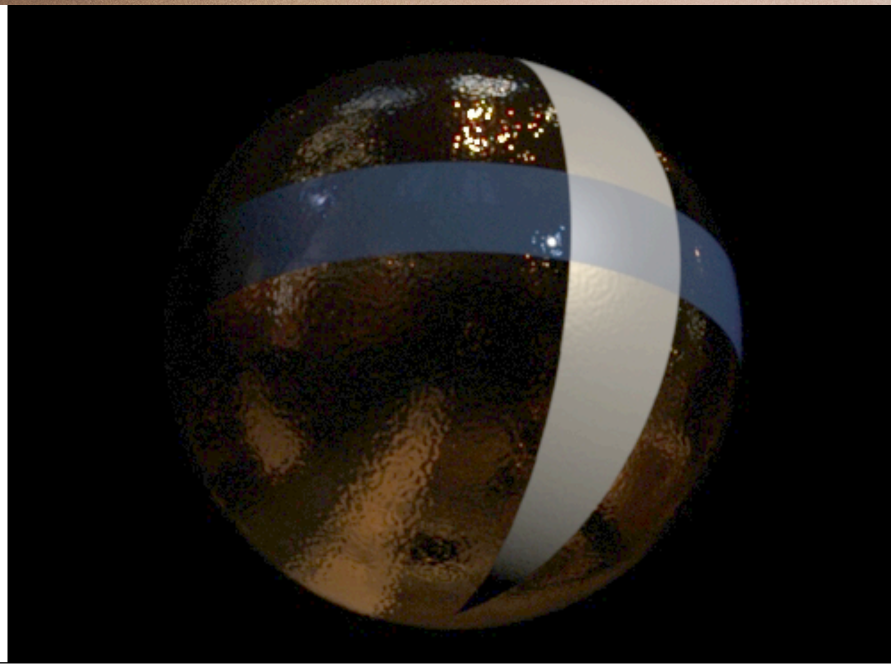

General Layering Idea

- Light goes through each layer based on the transmission lobes of that layer

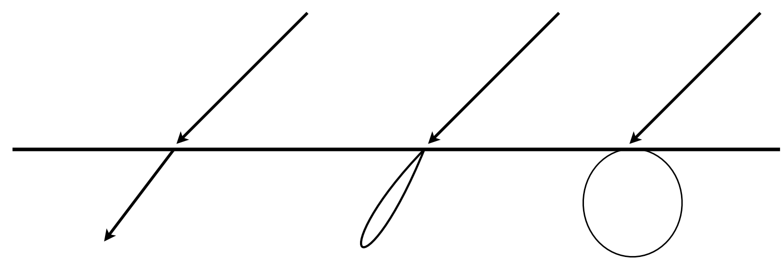

General Layering Idea

SIGGRAPH2012

- BRDF lobes of next layer act as blurring filters on inputs
  - The result can never be sharper than the inputs
- Account for all energy

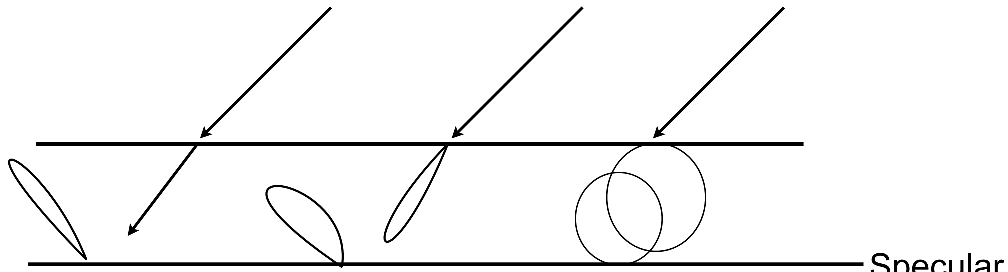

Lobe Interactions

SIGGRAPH2012

- 9 possible lobe interactions
- Each interaction includes the color weights and cone angles for both lobes

Top\Bottom	Reflection	Specular	Diffuse
Refraction	Reflection	Specular	Diffuse
Transmitted Specular	Specular	Specular	Diffuse
Transmitted Diffuse	Diffuse	Diffuse	Diffuse

Transmission Lobe Interactions

SIGGRAPH2012

- 9 more interactions
  - results in 3 transmitted lobes

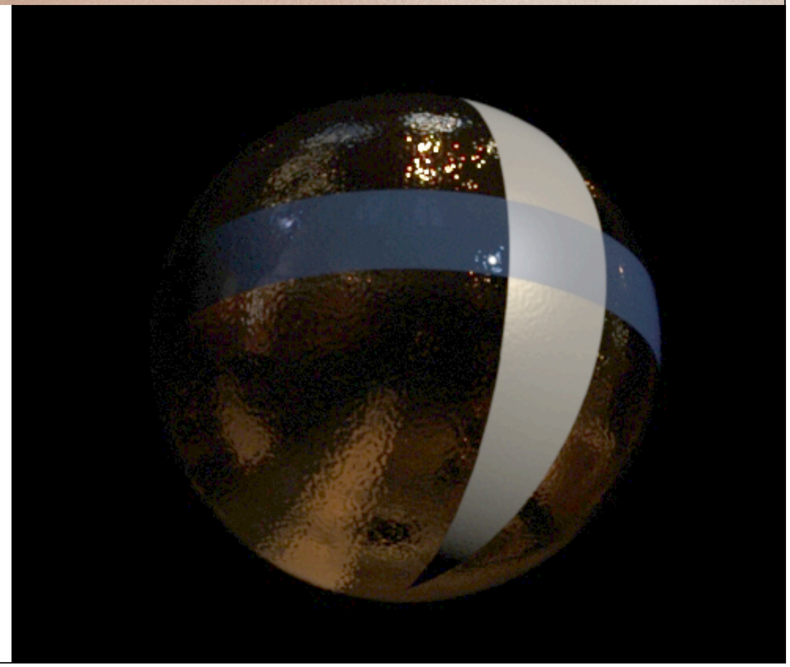

The smooth blue lacquer has it’s own undistorted sharp white reflection, but tints the bumped copper colored specular/reflection from the metal blue. The white rough “matte spray” has a large diffuse component that is partially colored by the layers underneath it. Over the smooth blue lacquer

General Layering

SIGGRAPH2012

- Also needed to include alpha/presence
- Don't include multiple interactions with the same layer
- Introduced an additional specular and reflection lobe when needed
- Powerful, but took some experimenting to learn

Layering Results

Layering Results

SIGGRAPH2012

Layering Results

SIGGRAPH2012

Notice the veins on the leaves are light When lit primarily from the front

Layering Results

SIGGRAPH2012

Now the veins are dark, and the lower frequency internal variation in the leaves is more prominent.

Layering Results

SIGGRAPH2012

Results: It Worked For Shading

SIGGRAPH2012

- Easy to get basic materials
- Flexible enough to hit a broad range of complex materials
- All but a couple models in WALL-E and UP were shaded with it

Hair wasn't included in the model. I wouldn't have done the few true glass objects with it.

Results: Performance

SIGGRAPH2012

- Layering was free
  - Fixed number of interactions with the light independent of layers
  - Fixed amount of data per interaction
  - Pattern generation time could go up
    - Easier to factor shaders into layers than to build up complicated individual layers

Results: GPU Relighting

SIGGRAPH2012

- Fixed data meant inputs to GPU shader and RenderMan illuminance loops were the same
- Didn’t solve all our problems with GPU relighting

For any light that could be fully implemented in CG and for every pixel affected by just one surface the results were exactly the same. Since some lights had some functionality that wasn’t implementable in CG, and antialiasing meant that some areas were blends of multiple objects (hair is the standard example, but for Wall-E it was large amount of small geometry), the end images didn’t always match closely enough for the lighters to have confidence in the results. Subsurface also didn’t show up.

Results: Lighting

- Lighters rarely had to touch the surface shaders
- Lighting seemed to stay on schedule
- More mixed results from Up

Recently someone looked at the number of edits done to materials by lighters. Wall-E had a very significant decrease in surface shader edits. This matches up with reports from lighters on the show that they rarely touched a surface parameter.

What I'd Change : Everything...

SIGGRAPH2012

- Lobe choice
  - Need to be able to sample effectively for Monte Carlo
- Distinction between Specular and Reflection
  - Less fear of tracing rays,
  - more environment lighting means sending rays everywhere anyway
- Early partial evaluation of view direction
  - Not sure how to do layering based on view direction
  - GPU relighting not a goal now
  - View independent baking is very useful for standins

Except the Process

SIGGRAPH2012

- Work directly with the users
- Account for all energy
- Same visual phenomena important
  - Possibly add retro reflection possibility
  - Build in true subsurface directly
- Flexibility is still important
- Even more important to minimize number of lobes

Thanks To:

SIGGRAPH2012

Alec Bartsch
David Batte
Chris Bernardi
Chris Burrows
Sharon Calahan
Danielle Feinberg
Nigel Hardwidge
Stephen King

David MacCarthy
David Munier
Brandon Onstott
Keith Stichweh
Erik Smitt
Colin Thompson
John Warren