Summary of

M~R to MR:

Getting the Maxwell look in Mental Ray

Summary by Starsoul

.

wedge   Maxwell / Mental Ray Render Settings

Siemens A31 mobile phone –

Maxwell:
- All **Maxwell** materials,
- 2 keylight emitters and
- 1 large fill light.
- Sample level 19.44,
cooked for 8.5 hours, 365x640

Mental Ray:
- All **DGS** Materials,
- 4 MR Area Lights.
- No GI / No FG,
- Min/Max samples 2/3,
600x958, complete in 47 minutes

➔ MasterZap

At minimum, make sure you are working in a proper gamma, i.e. set your **render** gamma to 1.8-2.2... (2.2 is the "theoretically correct" but some people like around 1.8 because it "looks like it has been exposed on film")... and/or use tonemapping.

- Use GI and FG, or perhaps FG multibounce (a bit less tweaking needed, but renders a bit slower).
- Use area lights w. proper physical characteristics (falloff, matched photons)
- Use at minimum a good gamma, or an exponential tone mapper, or fancier tonemapping still. A lot of "the **maxwell** look" lies in good tonemapping.
- Use DGS shader
That's good starter tips, methinks.

➔ floze

Maybe try to
- experiment with very low finalgather radii (0.3/0.3 for min/max with view dependant ON) and
- 2 or 3 diffuse bounces.
- And *use* a proper gamma, Master Zap is totally right about that. Set the gamma in the renderglobal's framebuffer options to 1 divided by desired gamma,
e.g.: 1/1.8 = 0.556
=> But if you have filetextures in your scene, you first need to gamma correct them too in order to get correct results. You can do this by piping them through a gammaCorrect node with gamma value set to the real gamma (i.e. 1.8) again.

➔ wedge

fun test of those **maxwell** type renders...
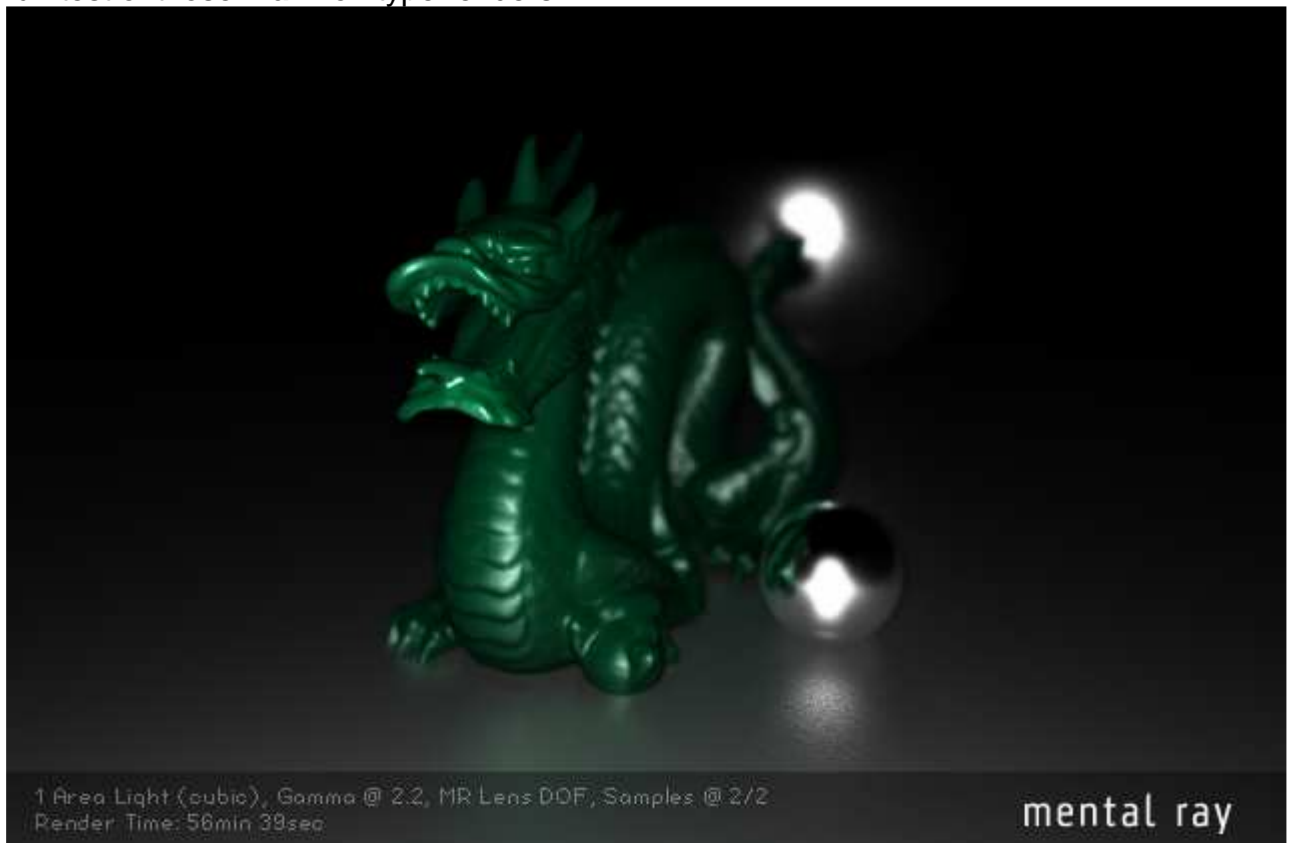
FuriousD   Tonemapping & 1.8 Gamma

I've been doing a bit of tone mapping rendering to exr,
- setting the gamma to 2.2 and
- gamma correcting my textures to .445.

Now I noticed you mention 1.8 ? how come not 2.2 ?
btw im using XSI, maybe thats the difference ?
I've been using the xsi-blog for reference, I'm sure you've seen it.

➔ http://www.xsi-blog.com/?p=45

➔ MasterZap
The "technically correct" is the gamma of your screen, which for most modern
computer screens is 2.2-ish.

- The thing is, stuff that's been to film and back gets some other non-linearities and a
"look". Some people think (fairly arbitrarily, i.e. no "physics" behind it) that rendering
to a gamma slightly lower than 2.2 (around 1.8-ish) "looks" more like this.

- Another important thing not handled by gamma is how overbrights are handled. Just
clipping them to white isn't good, so a more "soft" treatment of overbrights is
necessary (i.e. tone mapping that does "soft clipping").

Rendering to EXR and tonemapping in post is a very good idea.

MasterZap   Gamma Curve

It's a common misconception that TV monitors gamma is somehow a "flaw" to be
"corrected". That's not the case.

It's by design, and for very good reason; The curve created by a gamma of 2.2-ish is
approximately perceptually linear.

I.e. double the value LOOKS twice as bright. It isn't twice as bright, it just LOOKS that
way.

Due to this, it's the perfect way to encode images - no bits are wasted encoding what
our eyes cant' see. (We can't percieve high intensity differences at high total
intensities, but we can perceive INCREADIBLY small intensity differences on low
intensities).

=>   I always bring up this site when on the subject.

wedge   Tonemapping in Post

here's another test of
- area lights with
- fg and
- proper gamma.

this time i used the happy buddha model from stanford, and
➜   instead of correcting gamma with the framebuffer,
        i rendered to HDR and corrected in Photoshop.

but how can I get rid of this **under-glow?**

=> Jozvex
Are you using FG multibounce? I assume you are, and in that case you need to
- set your shadow trace depth on your area light, and
- in the **Render** Settings, to the multibounce depth +1.

EDIT: also, which model should i be using for realistic lights? Quadratic or Cubic?

➜ wedge
i'll try the scene again with higher shadow trace depths. on the area light, is the
setting is under raytrace shadows (ray depth) or under the MR options?

➜ Jozvex
Under the Raytrace Shadows.

Now, there were some difficulties as I expected in translating your shaders and light over to **Maxwell** because the units and colours etc are just completely different, so I just sort of matched as close as I could with guesstimation.

Some info:

Both renders rendered in 15 minutes and reached Sample Level 12.2.
- For the camera I used a Gamma of 2.2,
- Fstop 5.6,
- Shutter Speed 125,
- ISO 500.
- The light is only 20 watts (haha)
because it's quite large and the actual scale of the scene is very small in **Maxwell**. For example the buddha statue is only 1.2 cm wide I think. This was just the default size it started at so I left it.

➜　　The only difference between the two renders is that the one on the right has narrower DOF and has the fraunhofer diffraction (aka Glare) calculation turned on, just to see what it would look like.

➔ wedge
 i imagine that after very many samples the M~R and MR would look very similar...

the biggest differences i can see are in the hardness of the MR render,
and the blue fill color in the indirect illumination,
and overall brightness of the fill illumination.

are you using a physical sky or sky dome? or are you just using a black
environment?

➔ Jozvex
I'm using an
- "area light" the same size and colour as yours, and the
- skydome set to the blue of your room box.

I compared how it looked with the skydome vs the room box and it was pretty much
the same, only the room box version was slower.

➔ wedge
the room box was making my renders act all funny,
- so i replaced it with a very dim blue camera environment.
- i was also using cubic lights at 120,
- so i made it quadratic and
- reduced the power to 30.

here's what i've come up with. i'm pretty satisfied with
- mray area lights combined with
- DGS materials for that nice maxwell look –
➔ just takes a little more tweaking to find the right balance of area light strength, distance, size, and fg environment.

EDIT: rendered to HDR and balanced in Photoshop, render time 45 min 21.58 seconds

Settings:
- All ordinary maya blinn shaders with mib_glossy_reflection attached;
- gamma is at 1.8,
- hardrendered into the 8bit framebuffer.
- *Liiittle* touchup done in photoshop
- (added some vignette and
- sized down the image by a few pixels).

I'd love to share the scene but it's *so* huge.. maybe I'll crunch down the polys tomorrow to get it on a reasonable size.

edit: The scene file with bounding-box-replaced Buddha can be found here:

➔ http://forums.cgsociety.org/showpos...36&postcount=60

floze   Benefit of a maya blinn with glossy reflections attached

Quote:

Originally Posted by wedge
those look great, floze! i have been using pretty much DGS only these days... whats
the benefit of a maya blinn with glossy reflections attached?

Well, imo it's
=>   the highlights of a blinn which implement a fresnel reflectance,
and the scalability of a mib_glossy_reflection which also has a fresnel option
and some (fake-) dispersion that I ususally use these days.
Both together make up a perfect team. 😊

floze   Hdr sky generator for HDR Shop

Quote:

Originally Posted by wedge
anyway, next up on my list: the maxwell physical sky look!

Yeah, that's one big advantage of maxwell I guess.

Banty has written some pretty decent hdr sky generator for HDR Shop:

➔   http://www.banterle.com/francesco/download.html

a quick test with GI+FG 🙂

and a test with Aperture 😳
=>   Addition: www.Lightengine3d.com

here are my attempts at nice outdoorsy look with the
- mib_light_infinite node and an
- HDR skymap to get that fill lighting.
if only mental ray had a physical light shader!

FuriousD   Aperture Shader Node and there uses

Dagon, I love the detail in your renders, down the bottom of the budda
those little bumps look great, special on the Aperture version.
I must play with Aperture, do you know if theres any XSI version ?

Do you use it just for depth of field ? or for exposure as well ?

➔ dagon1978
i'm using just the
- depth of field...
- and adjusting the gamma to 2.2 (0,45)

i think the bump look good because of resolution 🙂 nothing special in my setup,
- i added just a dummy geosphere round the light visible in reflection only

p.s.   IBL hints

great stuff here for the sky... nobody tried the env_sky attached to ibl? look good for
me...
=> 1) and for wedge, the ibl dont shoot direct light, u have to turn on "emit light" for
this, but is very slow, i think it's better to use a direct light for this 🙂

=> 2) So all I did was placing a directional light in the direction where I assumed the
sun was coming from. That was no problem because I used diffuse bounces instead
of photons for indirect illumination.

MasterZap   Why lighting in proper gamma space

Quote:

Originally Posted by Lamster
What puzzles me is that the material is so washed out, even though it arguably does look "Natural". Is it because my lights are too bright?

Partially. A big error is to "light" your model (or even set up your materials) for a wrong gamma 1.0, and expect it to magically look "good" when you put on the correct gamma.

A scene properly lit for gamma 2.2 will look stupid in gamma 1.0
➔   way too dark with way too strong highlights.

So light it in the proper gamma space.

It's also a matter of us computer users being so "used" to the "wrong gamma" look that when it gets fixed we say "hmm, this looks wrong".

And as said above, the "compromise gamma" of 1.8-ish tends to work best to both work "correct"-ish and solve "expectations" with the added bonus of looking somewhat "film-y".

MasterZap   Lighting in proper gamma space - Workflow
Quote:

Or that my material Diffuse settings are set too high? They are at the default value of 0.8. All the file textures are applied to a default blinn for control.

Workflow:
- Fix the light intensity is my first suggestion.
- Tweak the materials for the new gamma. THATS when things start to look "natural".

Also, if you set mental ray's gamma (and not your applications "viewer" gamma) it will automatically de-gamma incoming textures (with the exception of "float" textures, which is assumed to be in linear space already).

If the setting you are changing truly changes the mr gamma setting internally, you should *not* have to "un-gamma" your textures!

Of course, this depends on the app; some app's use mr's build in texture loading, some bypass it, and then all bets are off. I am not familiar enough w. the Maya integration personally to say either way. I know in *max*, the gamma you set in "Preferences" is actually the app's own view gamma, not mental rays internal gamma.

<u>floze</u>   Sky Dome Render with Mental Ray & Free Scene File

Some more skylight fun using one of Paul Debevec's probes:

The probe's sun has been artificially eclipsed for
- finalgathering,
- replaced by a directional light. Still there's enough dynamic range
  to cause some trouble.

➔ Free Scene File by <u>floze</u>
The Buddha is replaced by his bounding box, same shader assigned. Pretty basic
setup, dont expect too much; no custom shaders etc., no worries about your rayrc. 🙂

Attached Files
🖥 <u>happy_buddhaDummy</u>

Additional:
<u>wedge</u>   Buddha Scene File etc.=> http://graphics.stanford.edu/data/3Dscanrep/
I took the original scene down to save some space, but <u>here</u> you can find the original
buddha at the stanford 3d repository, along with the PLY-TO-MAYA plugin so you
can load them up.

floze   Explanation of Settings used in the Free Scene File

Quote:

> Originally Posted by **francescaluce**
> floze, a couple of things about your scene:
>
> .the output is set to hdr, but the framebuffer is not.
> .you dechecked preview.convert tiles.. so even if set to 32bit, maya will use 8bit.
> .zdepth is enabled and hdr does not support it.

Ehm, yeah, the original file that came from wedge was like that, I actually changed it later on but I pulled that scene (online one) from an older version I did. I neither used 32bit nor zDepth.

Quote:

> Originally Posted by **francescaluce**
> .hypershade does not support right to left connections
> so all the shader setup is something queer to say the least.

Errr.. yes it does support left to right connections..?
All I'm doing is
- connecting the eccentricity to reflection spread (which does make sense), and
- the roll off to reflection base weight (which also makes perfect sense).

If you ever compared reflectivity of a blinn to the mib_glossy_reflection you would notice
=> it's harmonizing perfectly this way.

The actual reason I'm doing this connection is laziness (I dont want to go back to the mib_glossy_reflection every time I change my blinn settings).

Quote:

> Originally Posted by **francescaluce**
> .fg accuracy set to 1 ray is ... well I don't know.

Crank it up and watch the rendertimes increasing; does it look like insufficient with only one ray? Havent heard anyone complaining about blotches in my pics yet..?
Quote:

> Originally Posted by **francescaluce**
> .photon reflect/refract set to 64 is useless for this scene.

I set it as high as I can, it does neither hurt processing times (which are negligible for the photon tracing part) nor the quality of the image. What's your problem with a depth of 64?

Quote:

Well, thanks for that flattery. I know those guys pushing buttons randomly, I know I am not one of them. I dont force anyone to use my funky setups, seems like you are?

Btw: the only actual 'bug' in the scene kept undiscovered by your investigation. The photon intensity should be 4 times as high as it is right now, according to how spherical/physical area lights work (physical light = 4pi * maya light). But what the heck.

<u>floze</u>   Hint about setting up Number of Photons (GI) with FG for M~W Look

Quote:

Originally Posted by YourDaftPunk
Floze, those settings are friggin' sweet. Do the final gather settings add most of the points in the rendering phase? I thought the goal was to have most of the final gather points added in the FG pass, not at render time?

JOB 0.4 progr: 19.2% rendered on shawn-c1af26c83.4
RCFG 0.4 progr: computed 2000 additional finalgather points for tile [32 63] x [416 447]
RCFG 0.3 progr: computed 4000 additional finalgather points for tile [32 63] x [384 415]
RCFG 0.4 progr: computed 3000 additional finalgather points for tile [32 63] x [416 447]
RCFG 0.4 progr: computed 4000 additional finalgather points for tile [32 63] x [416 447]
RCFG 0.3 progr: computed 5000 additional finalgather points for tile [32 63] x [384 415]

Also, is precompute photon look up necessary with global illumination? I have scenes with 6million photons (8-10 bounces) where the precomp. takes forever. Seems my plain photon renders come out comparable to the FG+photon in about a fifth of the time.

-shawn

The more photons you add
=> in the less sense it makes to use FG.

Generally it makes less sense to add in fg and photons at all, **the closer you want to get to the 'maxwell-look'** I guess.

The scene I posted was an experiment, if I had to set it up for animation or anything else it would look completely different for sure.

For comparison, here's a really simple to set up brute force **path_material** test:

The rendertime isnt that bad, but consider
=>   it's a very plain and homogeneous setup, and you sacrifice every benefit of your maya shaders.
=>   However, it's probably the **easiest way** to get the maxwell feeling (noise + rendertime 🙂) into your mental ray renderings.

mental ray for Maya
AMD2600+/1GB RAM
~34min

And because I have too much spare time these days, here's another:
- skylight test,
- fg diffuse bounced on **maya blinns**:

mental ray for Maya
AMD2600+/1GB RAM
~1h 25min
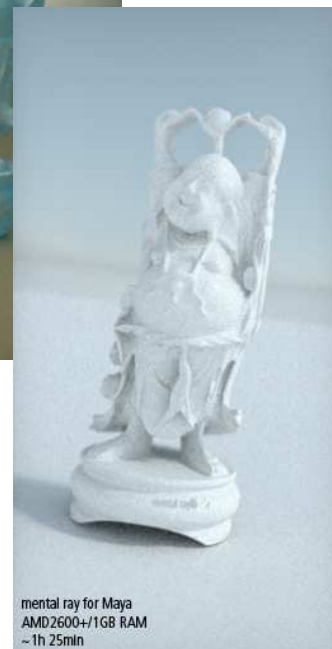
YourDaftPunk   Test Render with Photons only

tried the scene with

- **photons only** and => it renders fast, but you have

- convert the light to a cone to control their scatter to ~180 deg.

It would be nice if I could shoot say a million photons more from the camera- I guess that would be like final gather, but as you said:

"The more photons you add in the less sense it makes to use fg."

6m photons, noFG
P4 3200
8min48sec

<u>Saturn</u>   Test Renders with MR / DGS

This my test in XSI. We should have definitivly a thread dedicated to MR. I used
- DGS material and mib_glossy_reflection.
➔   Trying to respect the DGS rules (D + G + S <= 1) to keep physical
        correctness, but in the same I am using:
- FG and not photon.

Both images use FG.
1st took about 5 min, including FG phase and DOF.
2nd was longer about 20 min because all pixel rendered are glossed reflection. (edit:
once optimised the scene took about 10 min to calculate, I used more simple path for
reflection and FG rays)

I am using a MacBook pro 2.16 ghz / 1Gb Ram running under XP.

When using DGS, I have a doubt about the rules above. I don't see how to respect
when I plug a texture in the diffuse slot. How this rules must applied ?
Is the sum of the luminance I should take care ? Or RGB value ?
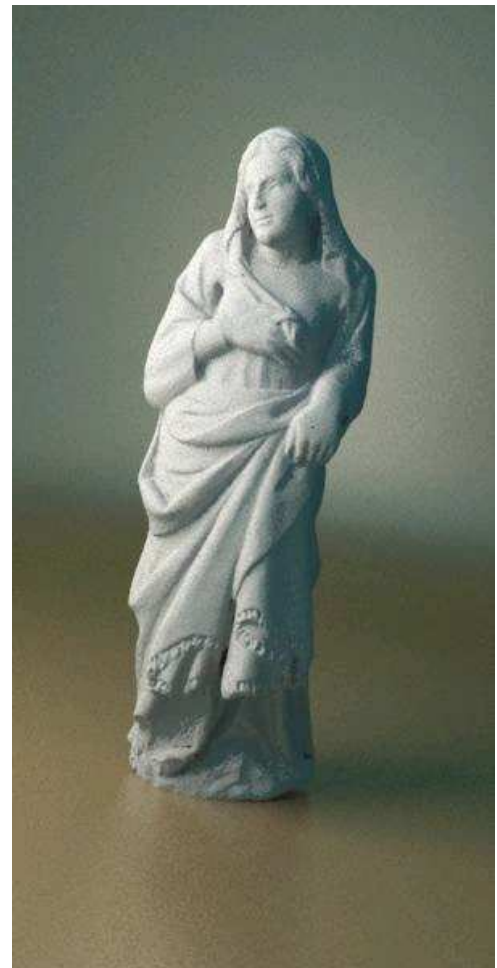
Attached Thumbnails

<u>Cobra1979</u>   Test Renders with FG&GI

Here is my test, really learn a lot of things here thanks you all!

12min FG+GI Athlon 1.8Ghz 728mo ram ➔

← And this one with SSS but need some improvements I think, still working on it.

Wedge   Scene Setup / Aperture Shader

So I am taking a look at Aperture from lightengine3d...

here's what i came up with:

- rendered to EXR,

- no post, except gamma/exposure correction.

- 1 quadratic area light,

- 1 dgs shader,

- 4 aphrodite models.

render time 4 hours 24 minutes for 853x480!

but i think thats because i am on a crappy 3ghz p4 with only 512 mb of ram and was

doing other things on top and i didn't tweak the apertute settings. i bet i can speed

this up substantially.

Mmm bokeh... so starry!

Settings:
- 80mm lens,
- fstop 2.4,
- shutterspeed 1/75,
- ISO 100, gamma 2.2.
The statues are about 15cm tall I think and the
- light is using a flurescent tube preset set to 80 watts.

<u>Chris-TC</u>   Bokeh

That's a very cool looking scene, Joz!

I've set up a similar scene (mental ray and **X-DOF**). In one of my initial tests I used a very reflective material but could not get the Bokeh to look as strong as in your left image.

Can you give some info on your settings? How big are your statues and what f-stop and focal length did you use?

Also, why does your right image not show any Bokeh at all, yet it has depth of field?

Will post more test of this scene tomorrow, including a material similar to yours.

Edit: I had not refreshed the thread, just saw your big image. Looks amazing!
I don't like the unevenly distributed grain though.

PAGE 8

➔ Jozvex
Quote:

Originally Posted by **Chris-TC**
Can you give some info on your settings? How big are your statues and what f-stop and focal length did you use? Also, why does your right image not show any Bokeh at all, yet it has depth of field?

Ok, my settings probably aren't very realistic because I don't have much real-world camera experience.

Settings:
- 80mm lens,
- fstop 2.4,
- shutterspeed 1/75, ISO 100, gamma 2.2. The statues are about 15cm tall I think
- and the light is using a flurescent tube preset set to 80 watts.

About the right image (that used to be there) not having bokeh, I'm not sure. I think it's hard for **Maxwell** to actually separate what light is coming from which bounce and so when you turn off some of the interactions it maybe just has to guess?? I'm not sure. The options are:

Direct light, indirect light, direct reflective caustics, indirect reflective caustics, direct refractive caustics, indirect refractive caustics. And then the fraunhoffer diffraction haha.

I had everything turned off except 'direct light' for that right hand **render**.

<u>Chris-TC</u>   X-Dof compared to md_aperture ?

Quote:

> Originally Posted by **FuriousD**
> Chris: thats really quick 14minutes, how do you find X-Dof compared to md_aperture ?

I never really tried the md_aperture but my understanding is that it gives you less control and that its depth of field slows things down considerably.

X-Dof is insanely fast and gives you very smooth results.
It also works perfectly together with motion blur which is amazing
=>   because proper DOF + motion blur is almost impossible to achieve otherwise, but again I don't know how md_aperture performs.

Thank you for posting the settings Jozvex.
I will try and approximate your fantastic looking result. I love the amount of ultra-sharp detail that's in focus on the first statue. That's just too cool looking!

<u>Chris TC</u>   Mental Ray Test Renders / DOF blurring (HDR) in Post?

As announced, here's my try to get things look similar to **Maxwell**.
Unlike **Maxwell** I had to add some subtle **grain in post**,
but the rendering time was less than 3 hours for this image (and I even overcranked the settings).

Quote:

> Chris, might it be that the blur is being applied on an 8bit image? Bigger question- what programs/plug-ins support post DOF blurring in HDR formats?

No, the blur is being calculated directly in Mental Ray, as a combination of lens shader and output shader. It should do all calculations in floating point.
The main post DOF solution seems to be Frischluft Lenscare, but I believe it's limited to 16 Bit.

Quote:

> Originally Posted by **Jozvex**
> Hehe, even if I had listed the material and light setup it'd still be a bit tricky trying to recreate them.

Oh yes, I'm aware of that. That's why I only asked for the camera settings.
BOKEH Effect

<u>floze</u>   LINK to Bokeh – Explanation

As for the depth-of-field discussion and the whole 'bokeh' story, I'm not sure if everyone knows what's actually meant by this. To talk the same language, here's a helpful link for everyone not being familiar with the japanese word bokeh (heh, not *everyone's* speaking japanese):

➔　http://www.kenrockwell.com/tech/bokeh.htm

So what you did, Duncan, would probably be considered as **'bad bokeh'** (which *does not* mean 'bad image'!*).
- If you turn down the spherical aberration attribute in the md_aperture lens
=> you would get 'better' or neutral bokeh (like Jozvex' maxwell renderings have).

Of course it's always an artistic decision to either have the one or the other, just wanted to point that out and maybe raise another topic where maxwell seems to be rather strong.

* Because some people think the out-of-focuse region of your image should be rather at rest than appearing busy. If you *want* this region to be busy for compositional/taste/artistic/whatever reason it's whole another story of course.

➔ MasterZap
But from the article you linked, the CG people defintion of "bokeh" is 180 degrees the opposite of the photographers.

It's the usual story; the **photographers** considers the **'perfect gaussian blur'** (that is so easy to get out of any renderer) to be "good bokeh"...

....whereas the **CG people** think the 'cute little circle - or better yet, hexagon! - of light that is slightly dimmer in the center' which is such a trivial **artifact** in photography but *so* hard to do in CG is "good bokeh"!

I.e. most people in CG I talk with who wants "bokeh" really WANT what this guy calls the "bad bokeh". That is the desirable goal. To make it look "photographed".

I mean, who wants standard, boring, dull and ugly gaussian DOF? Meh.

yashu   Setting up light Iitensity / Shaders correctly

to get the right intensity plug a mentalray light into the mentalray light shader slot of your light (this will ignore most of the maya light controls) then adjust the intensity on that one and when you got the wanted intensity simply copy it onto the photon intensity value. Note that the intensity of an mr light is the V setting from the HSV color, and that can go up to huge values and not only in 0-1 space.

➔ floze

Quote:

> Originally Posted by **MaNdRaK18**
>
> If You take a look at Floze's "mr physical light tour", You'll see that actual corelation is 4*Pi, not 40. That applies to all light types except flat area light (sqare, disc), where You must multiply Your light's intensity with only PI to get proper photon energy value.
>
> That all being said for lights with quadratical decay.

...and the fact that you need correctly set up shaders. Use dgs if you're not sure about that, or tweak your maya shaders as shown in my guide:

➔   http://individual.floze.de/fileadmi...S_exploited.pdf

Additional CG-Talk-LINKS:

- ➢ **Rendering: a test in interior lighting**
  http://forums.cgsociety.org/showthread.php?t=109035&highlight=Rendering%3A+test+interior+lighting

- ➢ **VRay-like interior renders with mental ray**
  http://forums.cgsociety.org/showthread.php?t=190232&page=1&pp=15&highlight=vray

- ➢ **M~R to MR: Getting the Maxwell look in Mental Ray**
  http://forums.cgsociety.org/showthread.php?t=364576&page=9&pp=15&highlight=maxwell+render+interior