

Visualization Insider

An In-depth Look at Image Sampling



Introduction

This article was written to clear up some misconceptions about the V-Ray image sampler and was a direct response to a thread posted on the Chaos Forum ([click here to view thread](#)). In this thread, several users found the explanation of the image sampler in the article [Critical V-Ray Settings – Part I](#) very helpful but were confused about why certain settings were recommended. In this article we (3DAS) wrote,

By default, the Min and Max values are 1 and 4 respectively, which provides decent quality for test purposes. However, these values will often need to be raised for production renders. Values such as 3 and 6 are typically good settings for most production purposes. Caution should be taken in using higher values than 6 because of the possibility of excessive render times. Additionally, it's a good idea not to separate the Min and Max values by more than 2 or 3. Doing so can result in a wide range in quality of detail because it leaves V-Ray more room to guess incorrectly.

In the Chaos thread, someone wrote...

"ok after extensive testing, that qmc setting does not make sense - its faster to do min 1 max 6 and the quality does not suffer which is what the document suggests"

Another wrote...

"I totally agree about with that as well. I was using Min 1 max 10 with the to get my glossies nice and smooth, and to give the QMC sampler enough room to sample adaptively for good results. After reading the PDF I tried Min 4 Max 6 settings and it increased my rendertimes with no visible difference in quality"

Followed by many more like...

"But then I started thinking (must remember to engage brain before fiddling) and looked again - so the pshot is the guys at cgarchitect need informing they are wrong on that point! - though I guess it depends a bit on what their qmc adaptive settings were"

Many good points were made in this thread but a lot of assumptions were also made. Hopefully this article will help clear up the subject of image sampling. **I would argue that image sampling is the most important and least understood feature in V-Ray. If you don't conduct good image sampling, having extremely high quality settings in other areas of the program (like irradiance maps and the QMC sampler) will do you absolutely no good.**

Which Sampler....What Values?

The first decision that needs to be made in area of image sampling is which sampler to use. The [Critical V-Ray Settings](#) article discussed some considerations that could make this decision a little easier, but since this article is being written to address the questions asked about the QMC sampler, here is a quote from that previous article...

The Adaptive QMC sampler works best in scenes with a large amount of detail or blurry effects. As a general rule, we suggest using this sampler in scenes where a high level of detail is needed anywhere, but not more than 75% of the rendered image. As mentioned before, when the combined area of an image needing a high level of detail is more than 75%, the Fixed sampler is probably the best choice simply because the small amount of time saved with the Adaptive QMC sampler is not worth the risk of having some areas under sampled. In either case, you should definitely experiment with several configurations to determine the optimal mix of quality versus render times.

These are general guidelines which, like most guidelines, can not be assumed to work for all situations. While they are what we believe to be good guidelines, a more important consideration is the amount of sampling that should be conducted. In the aforementioned thread a question was posed...

When using Adaptive QMC, is it better to use Min/Max values of 3/6 or 1/100?

Well I would simply argue that a Min/Max value of 3/6 would be a better choice more often than 1/100, but realistically you can't answer the question with a simple yes or no, because each scene is unique and each scene would handle these configurations differently. It all boils down to which configuration provides the best mixture speed and quality. In order to fairly compare the different samplers and different configurations, you really need to know what every setting does and how they work. You can't simply throw values in for one sampler and call it better than another because it resulted in a faster rendering and a cleaner image. If you don't compare apples to apples, any assumptions made are invalidated. The [Critical V-Ray Settings](#) article provides a good foundation for understanding what image sampling is all about, but it is not by any means all inclusive. What I want to do here is provide an in-depth look at a few really important concepts, and how image sampling affects the speed and quality of a rendering.

The Impact on Noise

Image sampling has a dramatic impact on two important components of image quality - detail and noise. If you want to increase the detail or reduce the noise in an image, no area of V-Ray will have as great an impact. In the 'Universal Settings' topic Vlado posted ([click here for post](#)), he mentioned that for image sampling, a good choice would be QMC with a Min Subdivs of 1 and a Max Subdivs of 100. I assume that these settings were provided as a general guide for new users, but unfortunately it's new users that most apt to suffer some highly undesirable consequences of such a configuration. Specifically, if you use Min/Max values of 1/100 with certain other settings set improperly, not only might you bring your renderings to a grinding halt, you can actually induce more noise in your image rather than reduce it and you can also degrade the detail in an image rather than improve it. In this article I'll show how easily all of these easily things can happen, but the point I want to make here is that this kind of configuration gives V-Ray a lot of opportunity to negatively impact your work and setting up your image sampling in this way is an unnecessary risk. There's no kind of image or animation that I can think of that would need anything even remotely close to 100 subdivisions per pixel (or even 10 subdivisions for that matter), so why take a chance.

So besides excessive render times, using a Min/Max setup of 1/100 can cause problems with detail and noise. In the first part of this article, we'll look at how image sampling impacts noise in

a scene and then we'll see what affect it has on detail. For both of these image qualities, I'll use a different scene, both rendered with a Core 2 Duo processor. If you want to follow along in V-Ray with this first scene, open the file named interior.max.

The image below shows a simple interior scene where a few V-Ray plane lights were used as well a direct light outside used to simulate the sun. To make the image render faster, I removed all furniture from the living room in the foreground and all the cabinets and appliances in the kitchen (in the background). I also used just a few grayscale materials with default settings to keep things simple.

Notice at the bottom of the image that this scene was rendered using fairly low quality GI settings and with the default noise threshold value of 0.01. I think it's important to mention here that it seems many users falsely believe GI settings to be the main culprit in poor image quality when more often than not the first place they should look for problems is with image sampling. You can achieve high quality renderings using relatively low GI settings and to demonstrate how little an impact the GI solutions have on the image quality for this scene, I'm using fairly low settings.

Notice also that I'm using the Adaptive QMC image sampler along with the default Min/Max values of 1/4. Normally, I would use the Adaptive Subdivision sampler because of its ability to quickly render scenes with large areas of smooth surfaces, but because this article is written to answer questions about the Adaptive QMC sampler, I will be using the QMC sampler in this article.

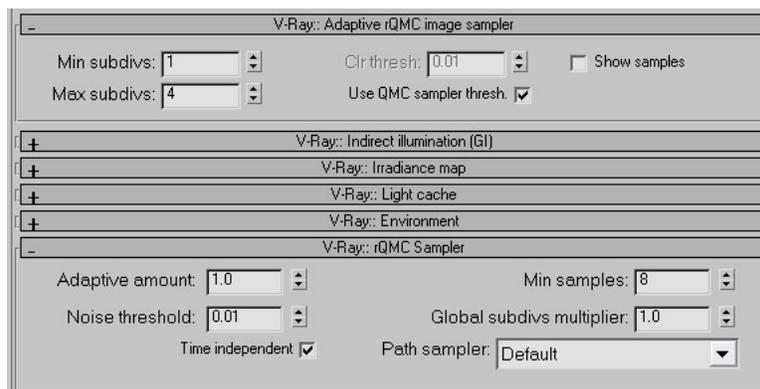
Notice that the image took 2min24sec to render using QMC with the default Min/Max values of 1/4. Notice also the high level of noise plaguing the image, especially in the kitchen area.



Now if we make the simple change of increasing the max value to 100, the image below, shows what we get. The rendering time increased significantly and there absolutely no decrease in noise.



If you're wondering why there is no discernible improvement in the image quality, it's because the Min and Max values are not the only variables dictating how much image sampling should be conducted. In the image below, notice that the QMC sampler has an option called, **Use QMC sampler thresh**. This option is enabled by default, and when it is, V-Ray looks at the **Noise threshold** value in the QMC Sampler rollout to determine how closely colored adjacent pixels should be. If this option is disabled, the QMC Sampler will instead look to the **Clr thresh** value, previously ghosted out, to determine how closely colored adjacent pixels should be.



Now since the Noise threshold value has a global impact on the noise values for all blurry effects within V-Ray, such as the noise in area shadows, blurry reflections, QMC GI, etc, it is often necessary to reduce this value below the default 0.01 value. The biggest problem new users have in doing so, however, is that it's linked by default to the image sampling being conducted and reducing this value can send render times through the roof...if that is, you set your Max subdivs value too high as is suggested in [Universal V-Ray Settings](#). So what happens is the Noise threshold value tells V-Ray how accurately you want the pixels to be colored and the Max subdivs value tells V-Ray how much work it's allowed to perform to achieve that accuracy. To show just how big of an impact the Noise threshold value has on image sampling, I'll reduce the value in this scene to 0.005, and without adjusting any other settings, notice that the rendering time doubled to 5m35s. What's important to note here is that there are no blurry reflections, there is no depth of field, and overall there are very few blurry effects that the noise threshold is affecting at this point. If this scene contained a large number of blurry effects, the rendering time might have quadrupled or increased ten fold. So while lower the noise threshold is a nice way to improve image sampling, it really shouldn't be directly tied to it. Instead, I recommend disabling the Use QMC sampler thresh option, and using the Clr thresh value to specify how accurately to color the pixels in your renderings.



Now going in the opposite direction with the noise threshold, the image below shows what happens when the noise is raised to the Max value of 1.0. Notice, that even though I've specified 100 Max Subdivs, the noise threshold is the real determining factor in how much noise the scene contains. So even though the image sampler is standing by and ready to work hard to get the pixel colors correct, the noise threshold value tells V-Ray that the accuracy needed is very low, and therefore, the image sampler has very little work to do. In fact, if you render this scene again with the Show Samples option enabled in the Adaptive QMC rollout, you'll see only one color...a solid blue color indicating that each pixel was sampled the same amount – once.

Notice also that there are two forms of noise in the image. A very tightly spaced noise consisting of pixels with vastly different colors and a much larger noise appearing as splotches on the walls. Since the noise threshold affects all blurry effects, and since GI is a blurry effect, it should be no surprise that the GI in this scene is causing a great deal of noise everywhere. So if you ever encounter noise in a scene, a clue as to the source might be the size of the noise. Large noise appearing as splotches will usually be a sign of poor GI while small noise will usually be a sign of inadequate image sampling.



To further show just how important the noise threshold value is and how it can overshadow the Min/Max values, notice how little noise the image below contains even though it was rendered with a Min/Max value of 1/1. Because each pixel contained only one sample, there is a noticeable jagged edge to many surfaces, however, since the noise threshold is so low it also resulted in very little noise.



Now, these examples help illustrate the importance of both the Min/Max values as well as the noise threshold value, but they aren't really good examples of what settings to use in a real scene. As was stated in the [Critical V-Ray Settings](#) article, the default values of 1/4 are good for test renderings but you will usually need to adjust these values higher. But instead of just arbitrarily applying a high Max value like 100, I would instead recommend increasing the Min value first. The image below shows what happens when we increase the Min value to 2 and leave the Max value at 4. Notice that there is a significant reduction in noise from the 1/4 configuration and only a modest increase in rendering time of only 5 seconds...certainly a worthwhile change. This is most easily seen by comparing images in the RAM players. Notice also that even with this increase in subdivisions, it's still substantially faster and contains much less noise than the 1/100 configuration.



Now, since I still have noise, I need to decide whether to I want to increase the Min or Max values or decrease the noise threshold. First I'll try to increase both the Min and Max values to 3/5 and the image below shows that I get a noticeable reduction in noise but also a noticeable increase in render times.



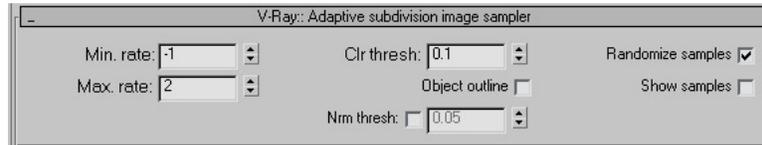
But instead of doing away with this increase in the Min/Max value, I'll instead keep it and try to reduce the render times a little by decreasing the noise. As stated earlier, I don't want to decrease my noise threshold just to improve the image sampling so the first thing I should do is disable the Use QMC sampler thresh. value. When I do I can now try a color threshold of 0.1 without having it impact other blurry effects. The image below shows the result of using a Clr thresh. value of 0.1. The rendering time is reduced by 23sec without any discernible increase in noise.



I want to see if I can improve the image at all by increasing the Max subdivs value by itself, and the image below shows that when I increase this value to 6, there is no change to the image or the render time. What this tells me is that the Clr thresh. value is now what's holding back the quality of the image. In other words, no increase in the Max value would help the image because the Clr threshold is too low.



As I mentioned earlier, the Adaptive Subdivision sampler is a better choice for a scene like this since it has large areas of smooth surfaces, little detail and few blurry effects. But if you switch to the Adaptive Subdivision sampler and open its rollout, you see something interesting. Notice that the Clr threshold value is set to 0.1, which is a fairly high threshold value for image sampling. So while this is the preferred sampler for this type of scene, it's partially because of this value being set so high, which can be done for the QMC sampler. If you were to use this sampler with a more reasonable Clr thresh. that you typically use with QMC, you'll notice that the Adaptive Subdivision sampler loses its edge and no longer beats the QMC sampler.



What makes this sampler really good for these kinds of scenes is its ability to conduct sampling at a rate of less than 1 sample per pixel (i.e. 1 sample every 4, 9, or 16 pixels). Notice also a **Nrm thresh** value which is locked by default to the Clr thresh value. This is another way that the Adaptive Subdivision sampler can speed up things because it can dictate the amount of sampling conducted in adjacent surfaces with varying normals, something the Adaptive QMC doesn't allow.

The Impact on Detail

Up to this point, the discussion on image sampling has been confined to its effect on noise in a scene. However, it has an equally important impact on the detail brought out in a scene. To show this, here is an image of a building rendered with very high quality settings that looks horrible only because of image sampling. Notice that even with high quality GI solution settings, a noise threshold value of 0.0, and detail enhancement enabled in the irradiance map, when we use only one subdivision for image sampling, we lose a tremendous amount of detail.



Going back once again to the 'Universal Settings' discussion, if we use Adaptive QMC with a Min and Max value of 1 and 100, respectively, the next image below shows what we get...a much better image rendered with more detail in less time.



But just like I mentioned earlier, this is not usually the best choice, even when care is given to adjusting the noise or color thresholds properly. This is demonstrated in the next image where the only thing changed is the QMC Min/Max value. Changing the Min/Max to 3/6 caused the rendering time to increase 5 seconds but if you compare both images in the RAM player, you'll see a subtle but noticeable improvement in the sharpness of the image where the 3 and 6 values were used. This is because every pixel was subdivided at least 3 times, unlike the in the image above where some were subdivided no more than once.



But now look what happens when I drop the noise threshold value (or color threshold) to 0.003 for both a 1/100 and 3/6 configuration. The 1/100 configuration resulted in a much longer render time with no improvement in detail. And while reducing the noise or color threshold further won't affect the 3/6 configuration, it will make the rendering times for the 1/100 configuration sky rocket.



Summary

Hopefully after reading this article you have a better understanding of how to use image sampling in V-Ray. While these scenes demonstrate that the image sampler plays an important role in the noise in a scene, a good deal of thought and experimentation should be made to determine the optimal mix of Min/Max values, Clr thresholds, and GI quality. If minimizing render times is important to your workflow, then you should spend a good amount of time determining the thresholds at which one of these 3 variables no longer affects render time and image quality, and the other 2 take over. For example, lowering the low Clr threshold value will do you absolutely no good if your GI solution doesn't conduct enough sampling of its own. Likewise, increasing GI quality will do you absolutely no good if your image sampling ruins its results. Above all, remember to not jump to the assumption that a poor image is due to the quality of a GI solution, because image sampling is usually more likely to blame