


Reducing Aliasing Artifacts through Resampling

Alexander Reshetov





I will be presenting... Yet Another Post-Processing Antialiasing Algorithm.

A lot of such algorithms were proposed in recent years, so not only I will have to explain what I did, but also is why.

Presenting...

Full Disclosure:
"Rippa Yappa Pari" is
"Paris is Splendid Indeed"
in Japanese

3/44

In the interest of full disclosure, "Rippa Yappa Pari" is "Paris is Splendid Indeed" in Japanese and you could even find this song on YouTube.




Unless we're talking about some fancy analytical methods, rendering is all about sampling. I am not talking about this one, but one you could read about in Wikipedia.

It is All about Sampling

sampling is the reduction of a continuous signal to a discrete signal

WIKIPEDIA



5/44

There are 2 reasons for this:

- Target devices are discrete
- It just makes life easier (similar to wine tasting)

To compare subsamples, we use a novel symmetric angular measure in a variety of applications that assess the difference between adaptive tessellation).

CR Categories: 1.3.3 [Computer Graphics]: Picture/Image Generation

Keywords: antialiasing, post-processing effects, deferred shading

e-mail: alexander.reshetov@intel.com

1 Historical Background

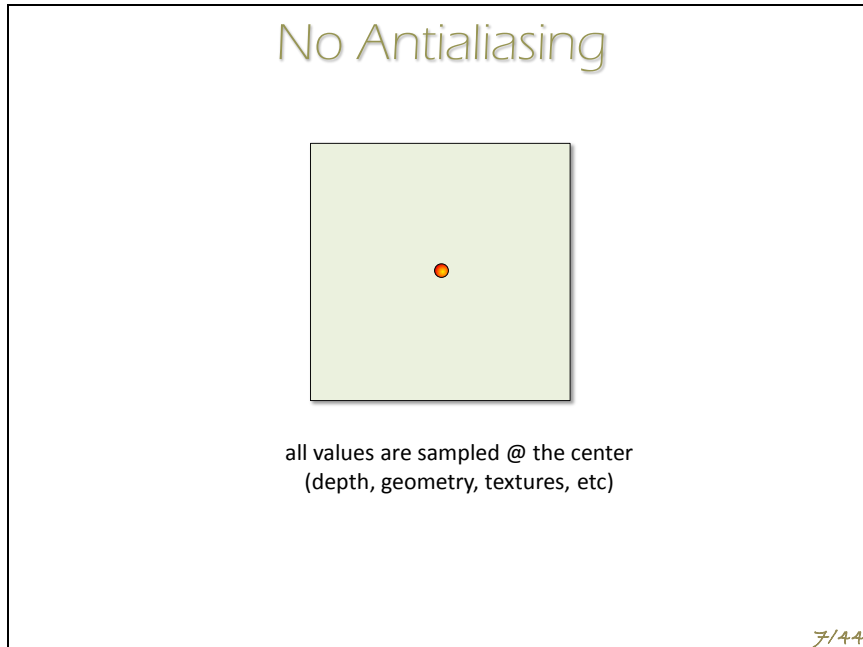
Rendering is about sampling.

2 Antialiasing through Resampling

During the pre-processing pass, we compare each subsample pixel with the pixel center, setting up a bit mask ('1' for similar for different subsamples). This mask, in effect, allows retrieval of precomputed filter coefficients. To fully utilize hardware tessellation

So, my original idea for the paper was just to say it in the introduction, and then proceed with the main idea presentation.

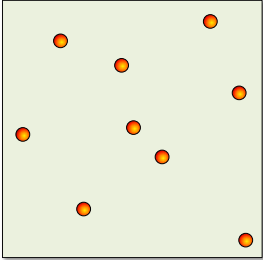
But, probably, it would've been too extreme. So let's talk about prior art now.



I will introduce key ideas of the new approach after briefly describing the existing techniques and their limitations.

Sampling everything once per pixel is, obviously, the fastest way, but quality will suffer.

Supersampling

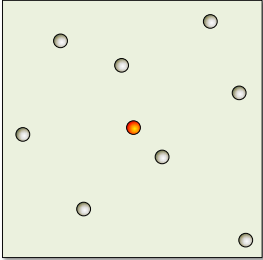


values are sampled @ multiple locations
(and resulting colors are averaged)

8/44

Au contraire, supersampling allows the best quality but rather expensive.

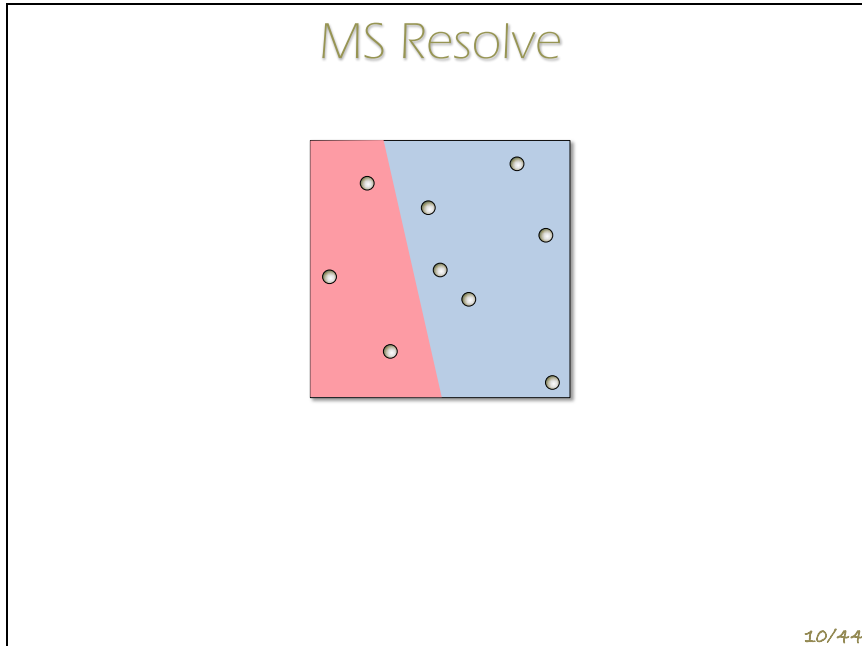
MultiSampling



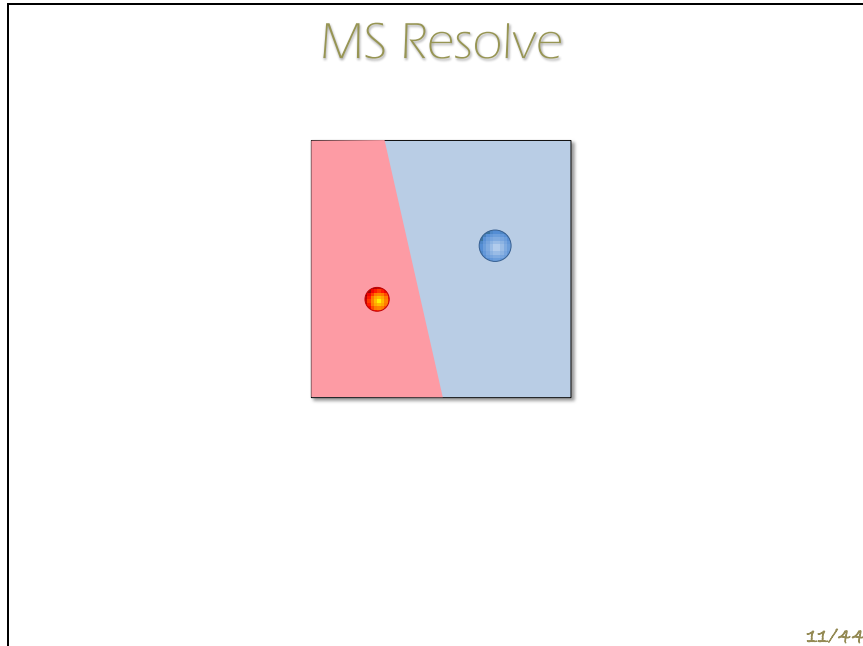
- Depth is sampled @ subsamples
- **Everything else** @ the center

9/44

Multisampling is an attempt to get all the advantages of supersampling at a fraction of the cost. It works nice for direct shading pipeline when the final pixel color is resolved through hardware acceleration.



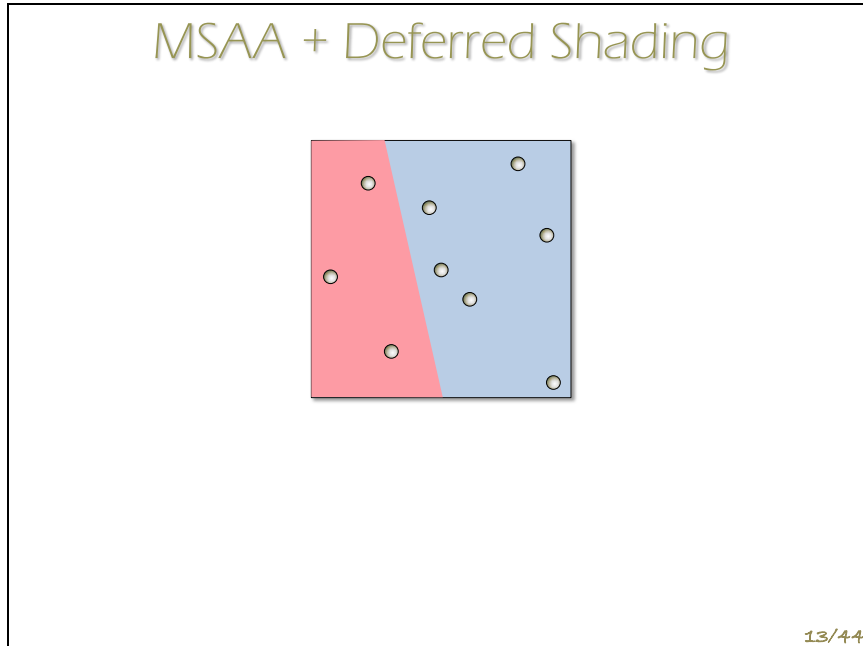
So if there are only 2 primitives overlapping a pixel, subsamples, obviously, will be split into 2 groups and




shading will be done only once per group.



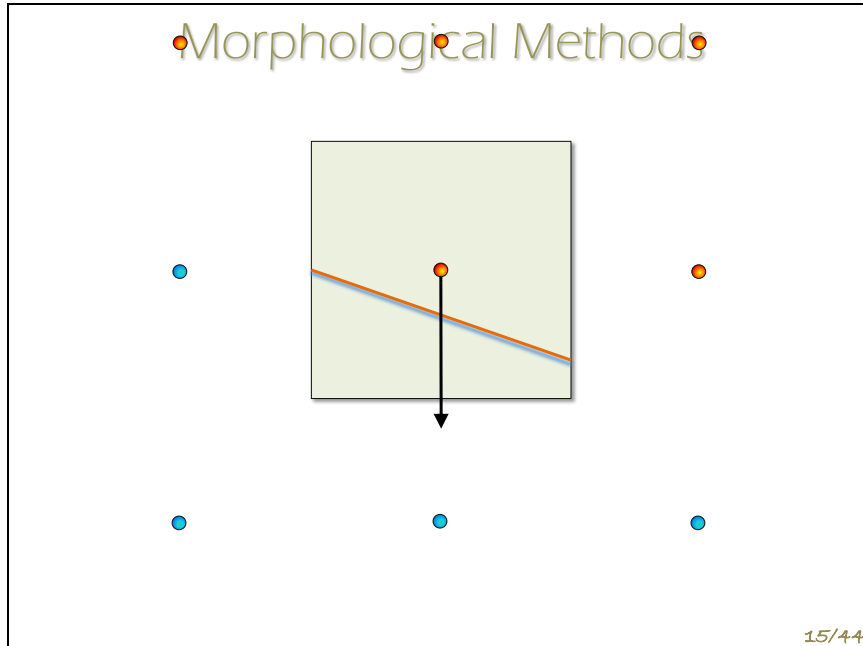
The final color will be computed by blending these 2 colors with weights proportional to coverage.



For deferred shading though all subsamples have to be explicitly written to output buffer for later processing, so it is not very efficient.



The reason for this is that resolve and lighting computations are noncommutative.

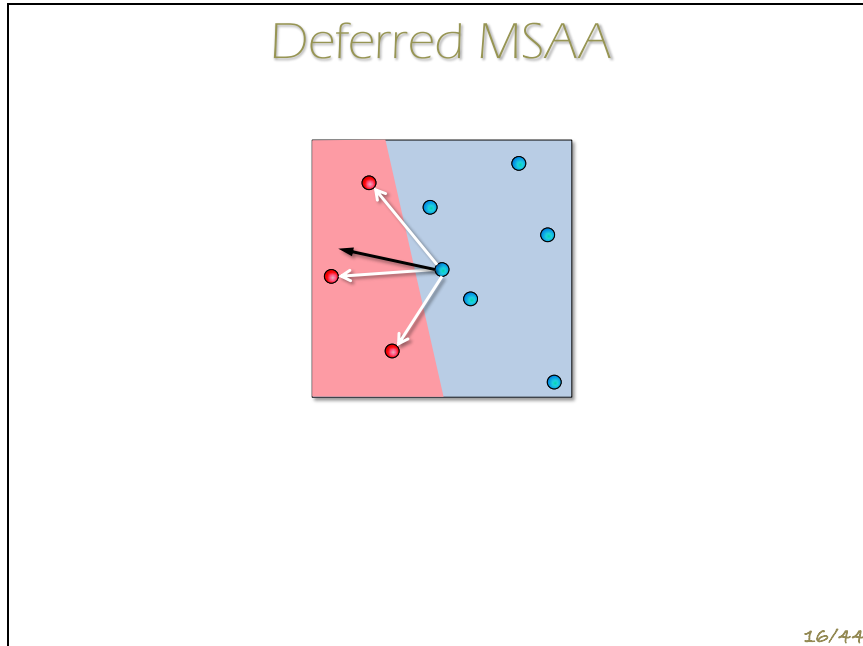


In morphological methods colors are computed once per pixel, and additional knowledge is obtained from analyzing non-local neighborhood, allowing to reconstruct plausible object silhouettes and then improve the quality through linear interpolation.

Accordingly, such methods are a good match for deferred shading, and, actually, for any other shading, since they are independent from the rest of the pipeline.

For this reason, such methods become popular in last few years. Who would knew?

Still, there are some inherent quality limitations, especially in temporal domain.

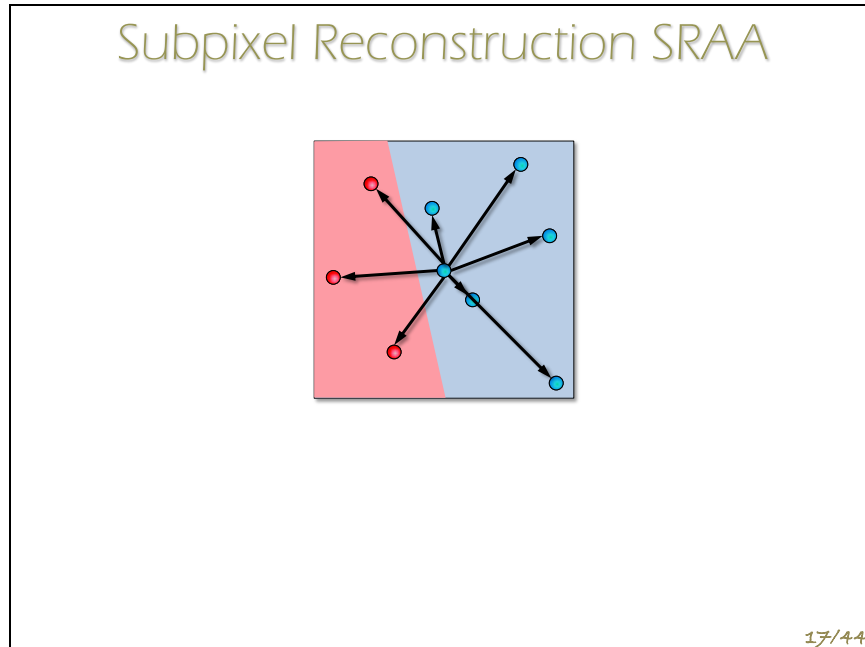


It is possible to improve quality by considering additional geometric samples, while still executing expensive shading computations once per pixel.

The technique I describe today falls into this category.

Another approach to this effect is deferred MSAA, proposed by Matt Pettineo. It works by averaging all subsamples which depth is significantly different from the depth of the central sample and then bilinearly interpolating the image in the direction of the average vector.

It is a very simple approach, but it doesn't scale well with number of samples and it is possible to come with a better direction of the resampling vector.



In sub-pixel reconstruction, all processing is done implicitly, applying cross-bilateral filtering to upsample per-pixel colors from geometric data.

It is a very interesting idea, but it substitutes a linear block filter with non-linear one in which coefficients are proportional to geometric difference between subsamples, expressed numerically.

AA Taxonomy

antialiasing method	quantity					
	depth	coverage	geometry	shading values	storage	BW
no antialiasing		x				
multisampling antialiasing MSAA [Ake93]						
coverage sampling antialiasing CSAA [You06]						
supersampling antialiasing SSAA [Lei80]		x				
MLAA [Res09, BHD10, Per10, Bir11, JME*11], FXAA [Lot11], SMAA 1x [JES*12]		x				
a directionally adaptive edge antialiasing [JYP09, Joh12]	x					
geometric methods [BWG03, CD05, Mal10, GG12, Per12]	x	x				
edge blurring: directionally localized DLAA [And11]; normal filter NFAA, screen-space SSAA [Uni11]		x				
temporal reprojection [NSL*07, YNS*09, Kap10]		x				
spatial/temporal supersampling + morphological antialiasing SMAA 4x [JES*12]		x				
deferred MSAA [Pet10]		x				
subpixel reconstruction antialiasing SRAA [CML11]		x				
surface based antialiasing SBAA [SV12]						
resampling antialiasing RSAA (this paper)						

Legend:
sampling rate per pixel

- x:**
not used
- once
- some
- many
- all
- ∞

18/44

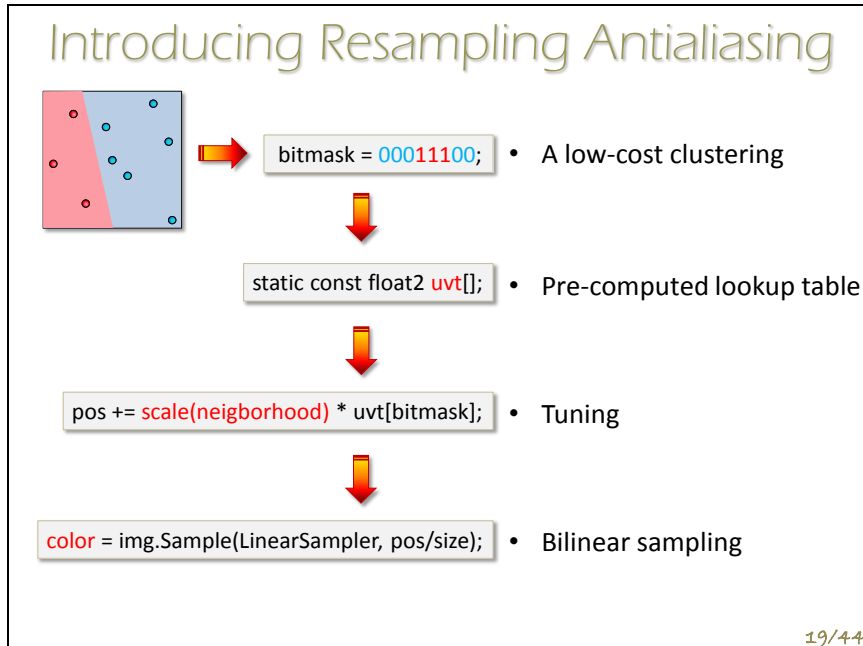
You could read more about various post-processing antialiasing methods in the last year's Siggraph course. I have also included all those algorithms, and a few new ones, into this convoluted table. To my surprise, reviewers did not object to this table excessively, so it is now official.

Please don't try to read too much into this table: it is more like an art rather than a science.

Basically, I am using shades of gray to indicate typical sampling rates for a number of entities required for antialiasing computations.

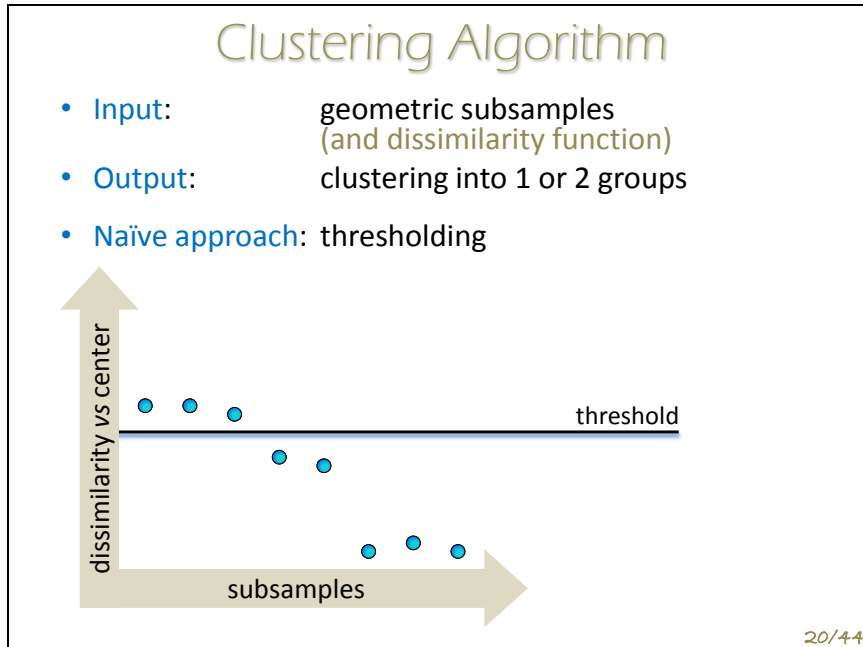
A white color stands for a single sample per pixel and then it scales up to a black color for many subsamples. It is not quite 50 shades of gray, but a lot.

This table just illustrates the fact that all of these values can be sampled at different rates. And it also improves the Hirsch citation index for all these papers.



The algorithm I present today requires splitting all subsamples in pixel into 1 or 2 groups. The resulting binary mask is then used to retrieve filtering coefficients from the precomputed lookup table. All this processing is done inside the pixel. There is also an optional tuning step, adjusting the coefficients by considering local neighborhood. The last step of the algorithm is just bilinear interpolation.

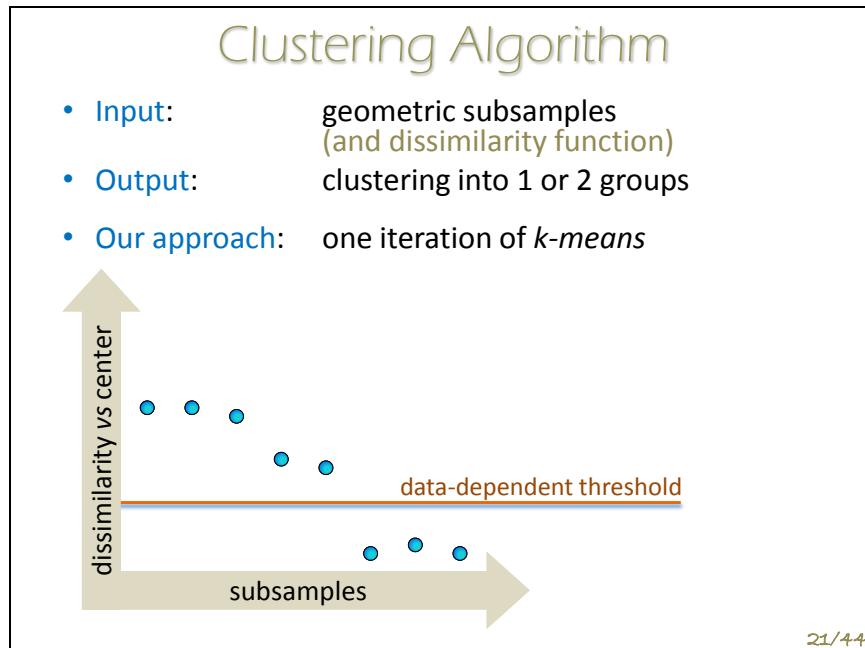
Except this, all other steps are original and I will describe them now.



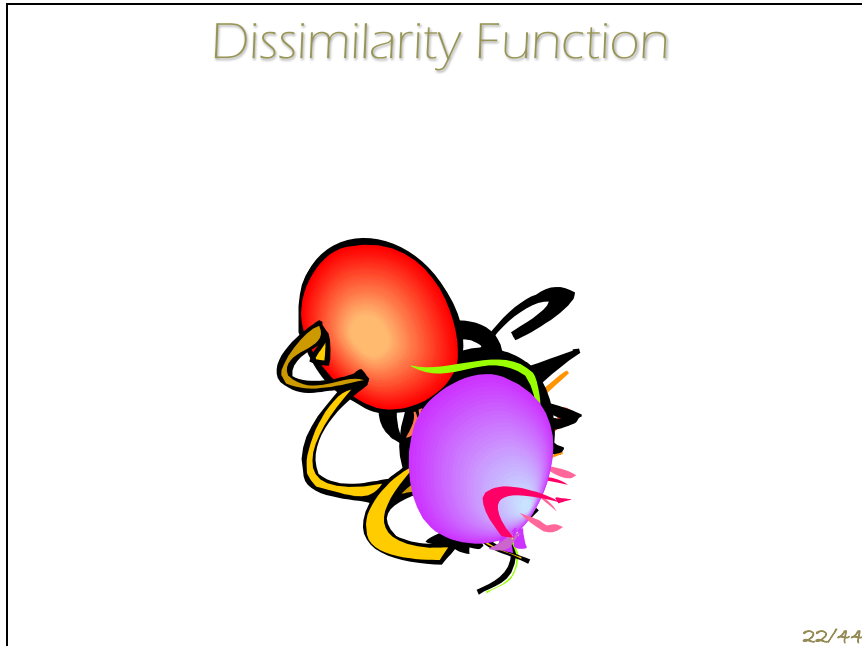
The clustering is based on considering the dissimilarity between all subsamples and the center of the pixel, expressed numerically.

Naïve approach would be to just split everything by comparing these values with a predefined threshold.

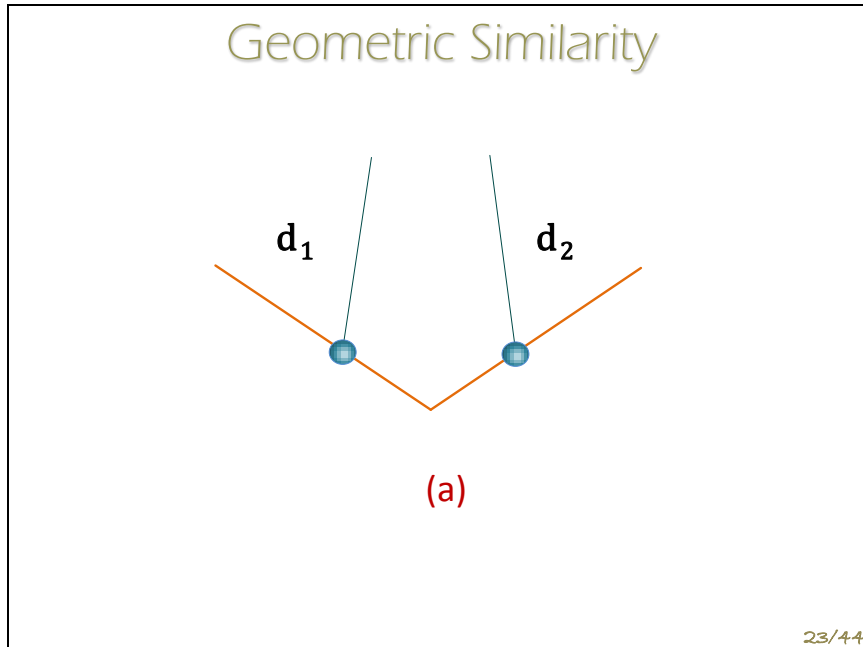
It is not the best possible approach, as illustrated on this figure.



This is one of the most typical problems in science, and there are multiple ways to do it. One of the most widely used is *k-means*. It is good, but expensive, requiring multiple iterations. So what we do is restrict ourselves just to one iteration.

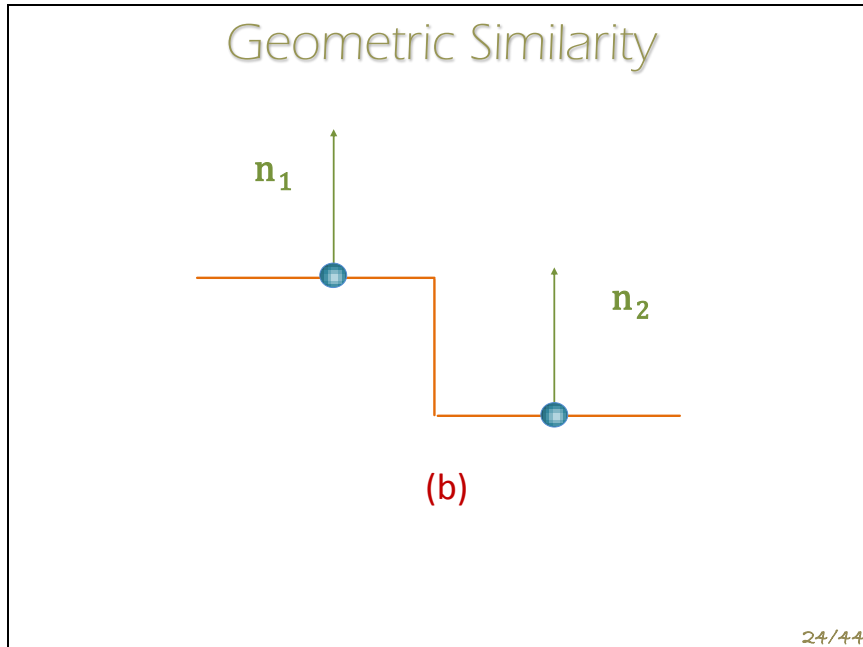


To cluster subsamples, we need a mechanism to measure dissimilarity of 2 subsamples... without using colors.

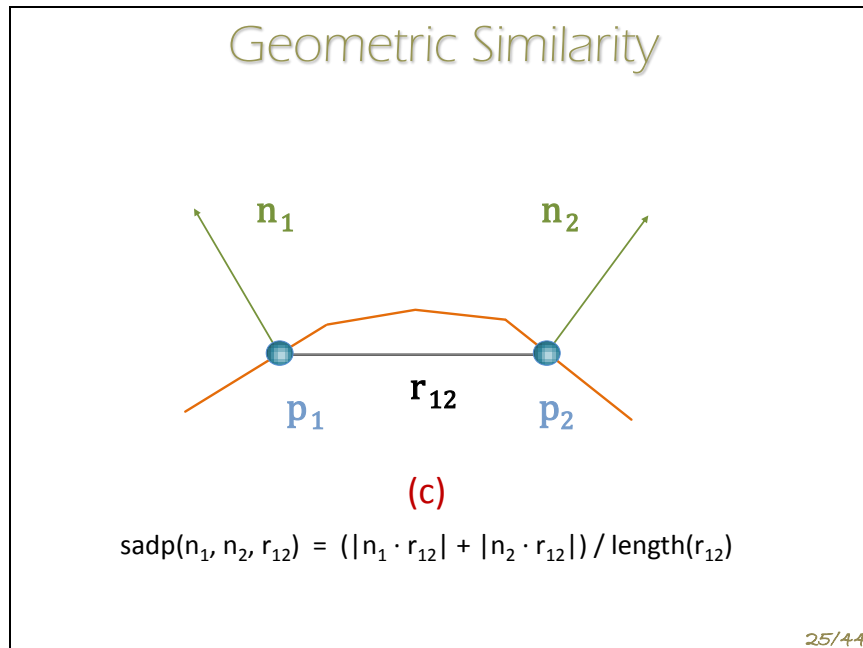


But we still can use geometry.

One popular approach is to compare depth of two samples, assuming that if depth is significantly different, subsamples belong to different objects. It is very simple, but may fail in corners, as shown on this image.



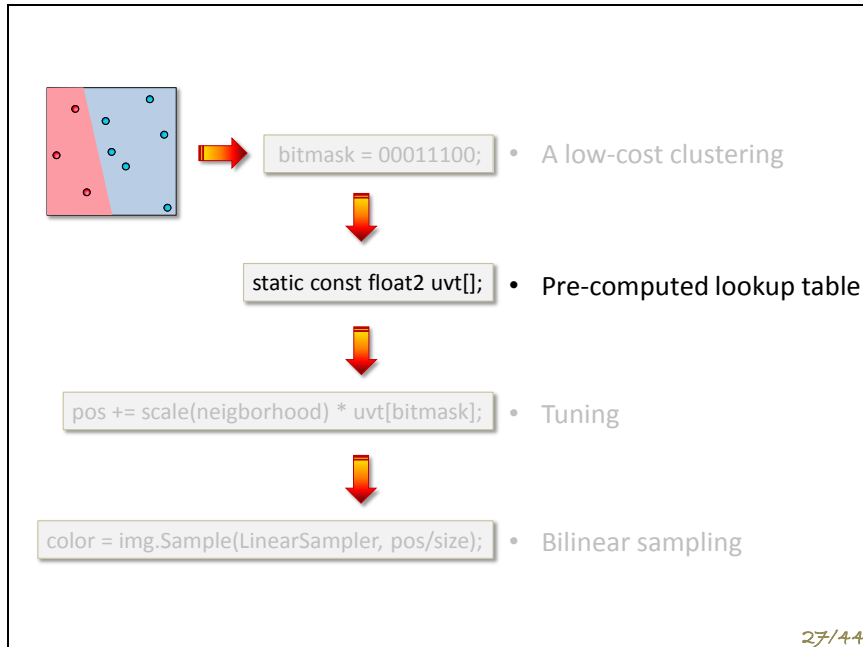
Corner detection could be resolved with a measure that uses subsample normals – for example, their dot product. However, this measure will not be able to distinguish spaced-out almost-parallel surfaces.



And finally a new measure which we propose. It is a sum of the absolute dot products of two normals and a normalized vector from one subsample to another. It addresses problems of other measures; it is symmetric and does not depend on camera position.



We call it SADP and, since it can be used in other computer science disciplines, I created this clipart to help you remember it better.



Now we have an index, so we need an array to index into.

Lookup Table (for filtering offsets)

- Why?
 - The set of all bit combinations is finite
 - The solution can be precomputed
 - Saves run-time computations and
 - Allows to be thorough
- How?
 - just learn it...

28/44

The set of all bit combinations is finite, so it is possible to precompute all possible solutions. It will save run-time computations, and will allow us to be thorough.

We still don't know how to solve this problem, but we can learn it.

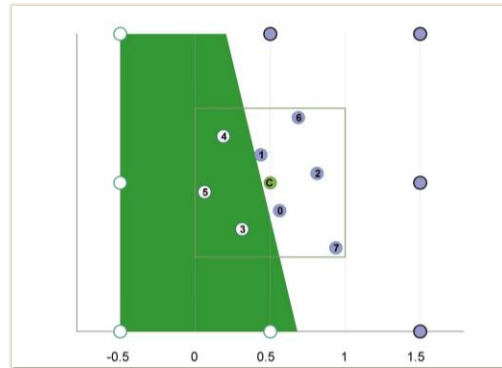
Lookup Table (for filtering offsets)



29/44

Not quite like this but close.

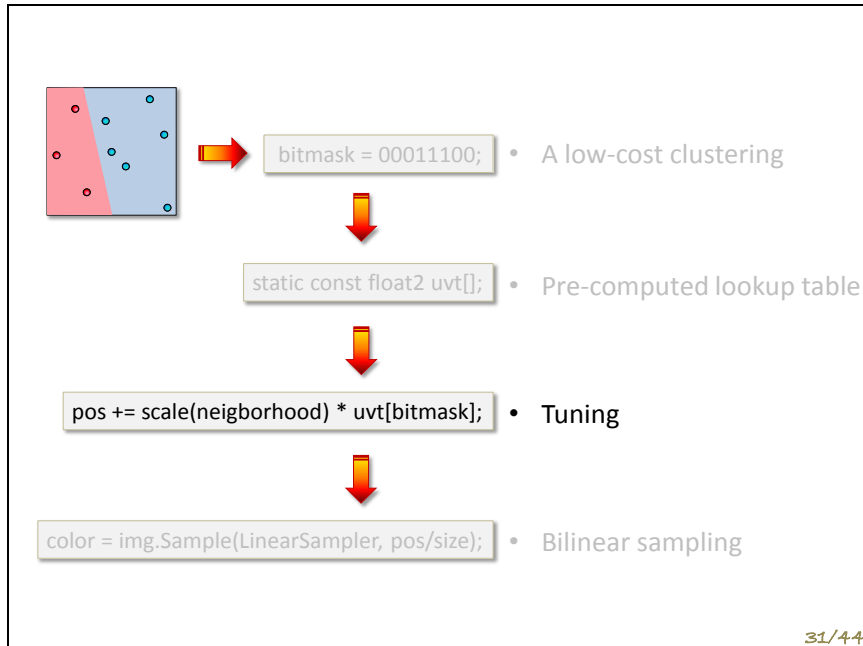
Using Machine Learning



$$uv_{opt} = \underset{uv}{\operatorname{argmin}} \sum_{n \in \text{sample set}} (color_n - color_{uv})^2$$

30/44

We use a simple training data set, for which anti-aliasing can be done analytically, and for each bit combination find offset vector which minimizes average error for the training data.



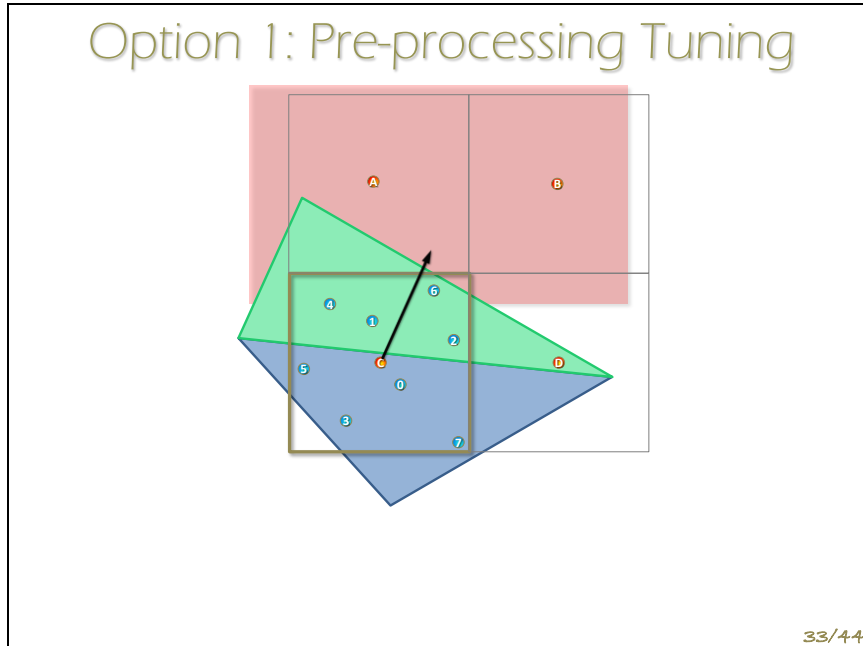
This will result in a good solution which can be further improved by looking into local neighborhood.

Why We Need It

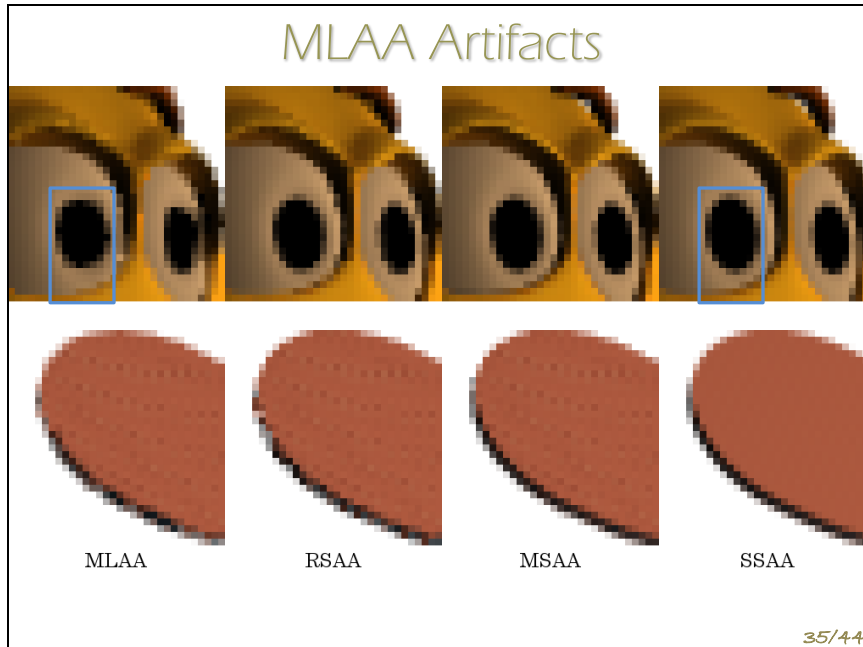
$\text{dissimilarity}(\text{green}, \text{red}) > \text{dissimilarity}(\text{green}, \text{blue})$

32/44

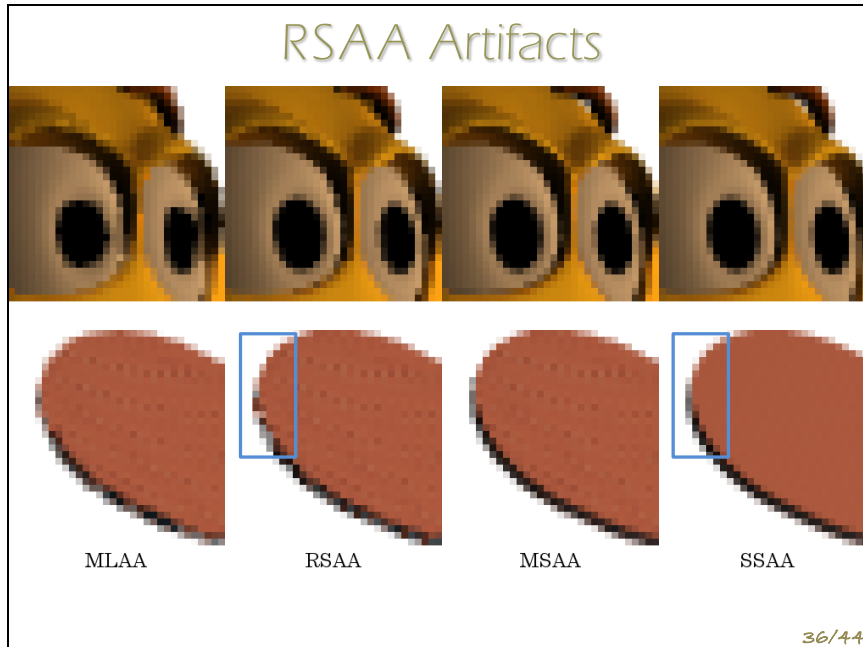
Let's talk why it may be necessary. We have 2 small triangles overlapping a given pixel. By looking on the pixel subsamples, we decided to interpolate the final color at the position defined by the black vector. We were hoping that samples in this direction are representative of green subsamples, which is not true. It will result in blending of blue, red, and a little bit of green, which is not what we need.




We can fix this situation by using SADP measure. One way is to compare green subsamples with closest outside samples and keep only ones which have their matches (it will be subsample 2). This requires extra work, which is unnecessary if pixel is not resampled.



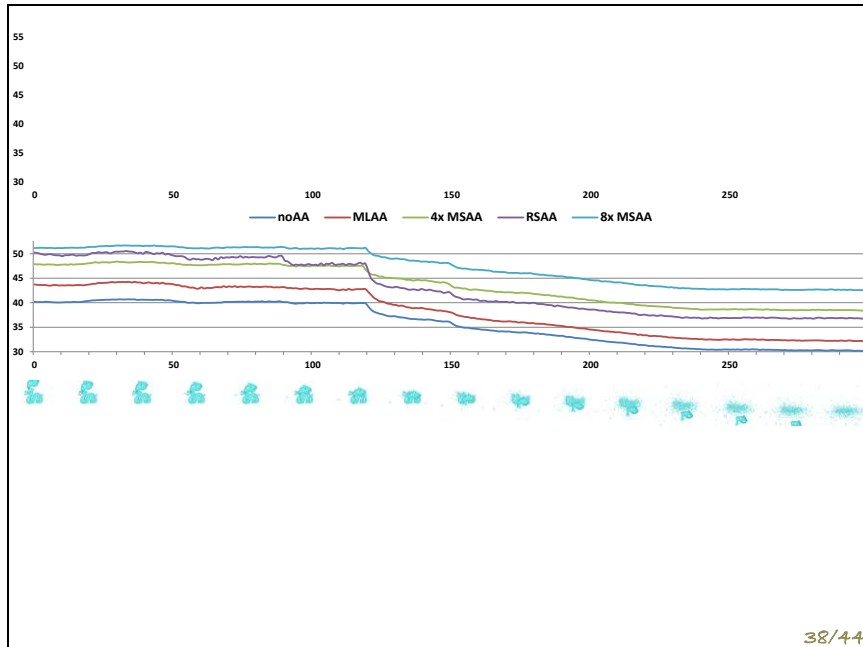
Let's talk about artifacts. This picture shows 4 antialiasing methods side by side. Let's first look on MLAA issues. The right pupil looks OK by itself but it is very different from the correct shape. It is because MLAA hallucinates silhouettes due to lack of any more detailed data. It is bad in spatial and especially temporal domain, resulting in flickering.



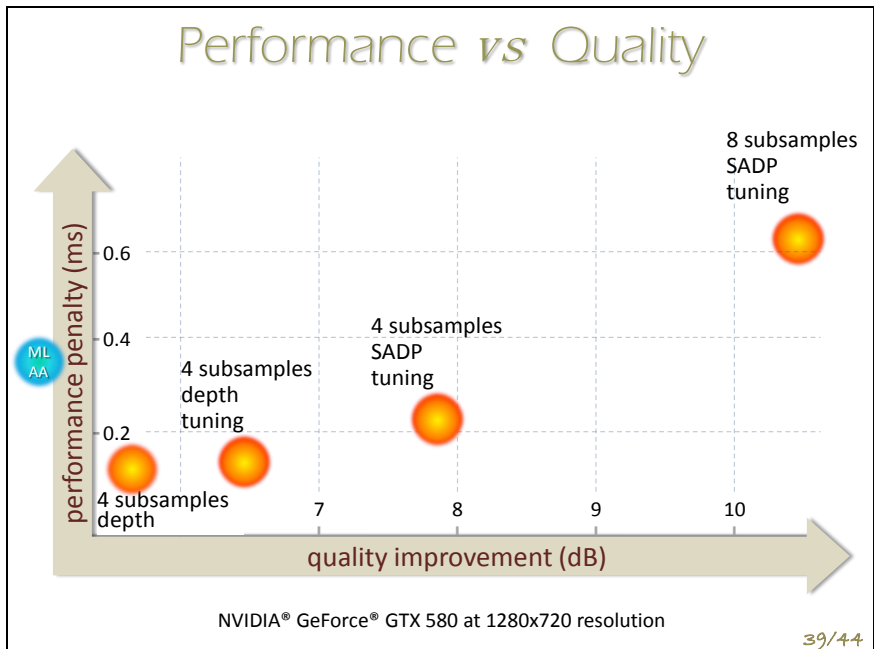
RSAA problems are most noticeable when its assumptions are broken, and there are more than two distinct surfaces overlapping a pixel.



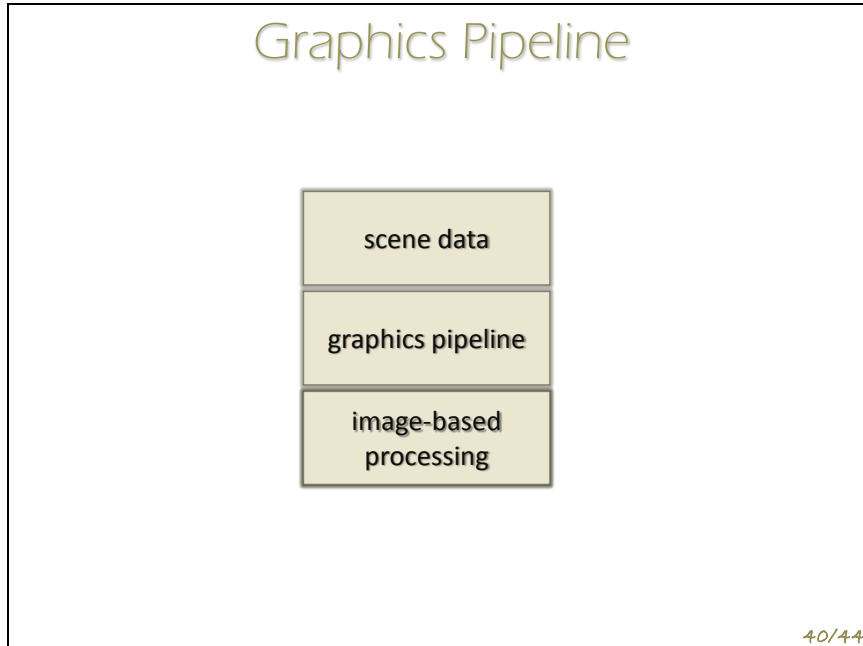
Overall, RSAA allows about 10 decibel quality improvement, compared with 5 for MLAA. Generally, RSAA is a little bit worse than MSAA which uses the same number of subsamples.



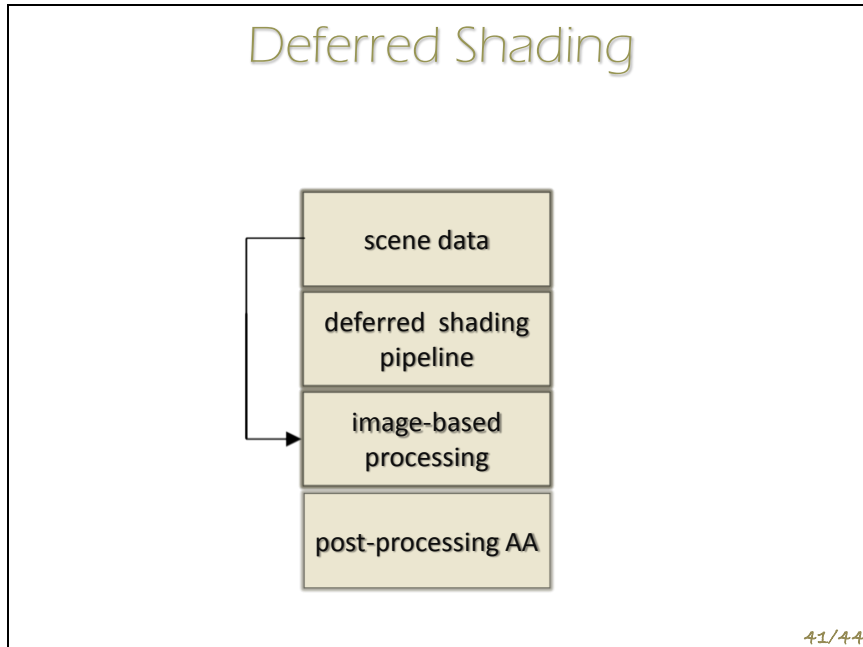
There is one noticeable exception. When there is lot of incoherent subpixel geometry, RSAA quality nosedives, while still exceeding one for MLAA. Just by looking on this chart, it is possible to identify the moment of collision.



At highest settings, RSAA allows about 10 decibel quality improvement, taking 0.6 milliseconds on 580. If we use only 4 subsamples, quality will go down by 2 decibel, saving about 0.4 milliseconds. We will not save much by using depth test instead of SADP, but will lose in quality. And, it doesn't make much sense avoiding tuning either. The performance data in this chart excludes the time required for computing and storing per-pixel normals, which we assume could be a part of the deferred shading pipeline anyway. My implementation of this feature takes about 0.25 ms. And, for comparison, MLAA will be right about here at blue blue (using [Jorge Jimenez](#) implementation).

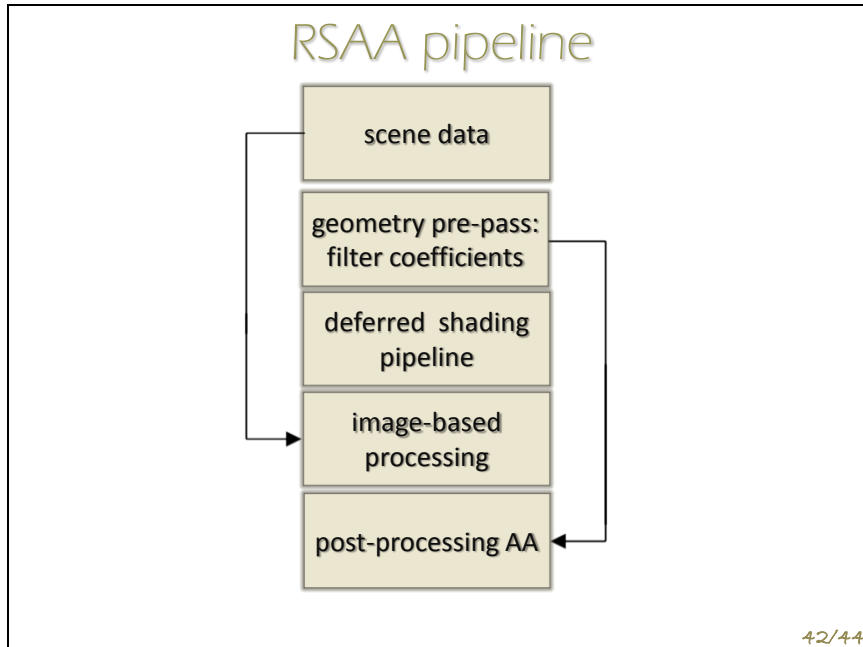


On the other hand, MLAA is completely orthogonal to graphics pipeline. During rendering, a 3D scene description is converted to 2D image, as alluded on this chart.



In deferred shading, some of the processing, well, is deferred till the last stage, allowing greater artistic freedom in scene composing and lighting.

Post-processing antialiasing techniques are a natural match for deferred shading, since they are completely orthogonal to the rest of the pipeline.



Technique, which I presented today, introduces an additional pre-processing stage, at which filter coefficients are computed. These coefficients are then used at a post-processing stage to antialias the final image.

Conclusion

- 8X increase in the geometry sampling rate allows noticeable improvements in quality at a reasonable cost
- RSAA components
 - SADP measure
 - Fast adaptive clustering
 - Machine learning approach
 - Tuning of filter coefficients
 - Example of Direct X pull mode attribute evaluation


43/44

It is indeed possible to improve quality by considering additional geometric samples but still shading only once per pixel.

RSAA includes some features which could be used in other computer science disciplines.

I did not discuss the last component here and ordinarily I wouldn't advice relying on my Direct X code, since I am opposite to an expert in this area. However, there are simply no other examples of using the pull mode evaluation on the Web and it is a rather kinky concept.

Acknowledgments



Thank you!