SIGGRAPH2011 VANCOUVER

Filtering Approaches for Real-Time Anti-Aliasing

http://www.iryoku.com/aacourse/

SIGGRAPH2011 VANCOUVER

Filtering Approaches for Real-Time Anti-Aliasing

Morphological Anti-Aliasing

Alexander Reshetov
Intel Labs

alexander.reshetov@intel.com

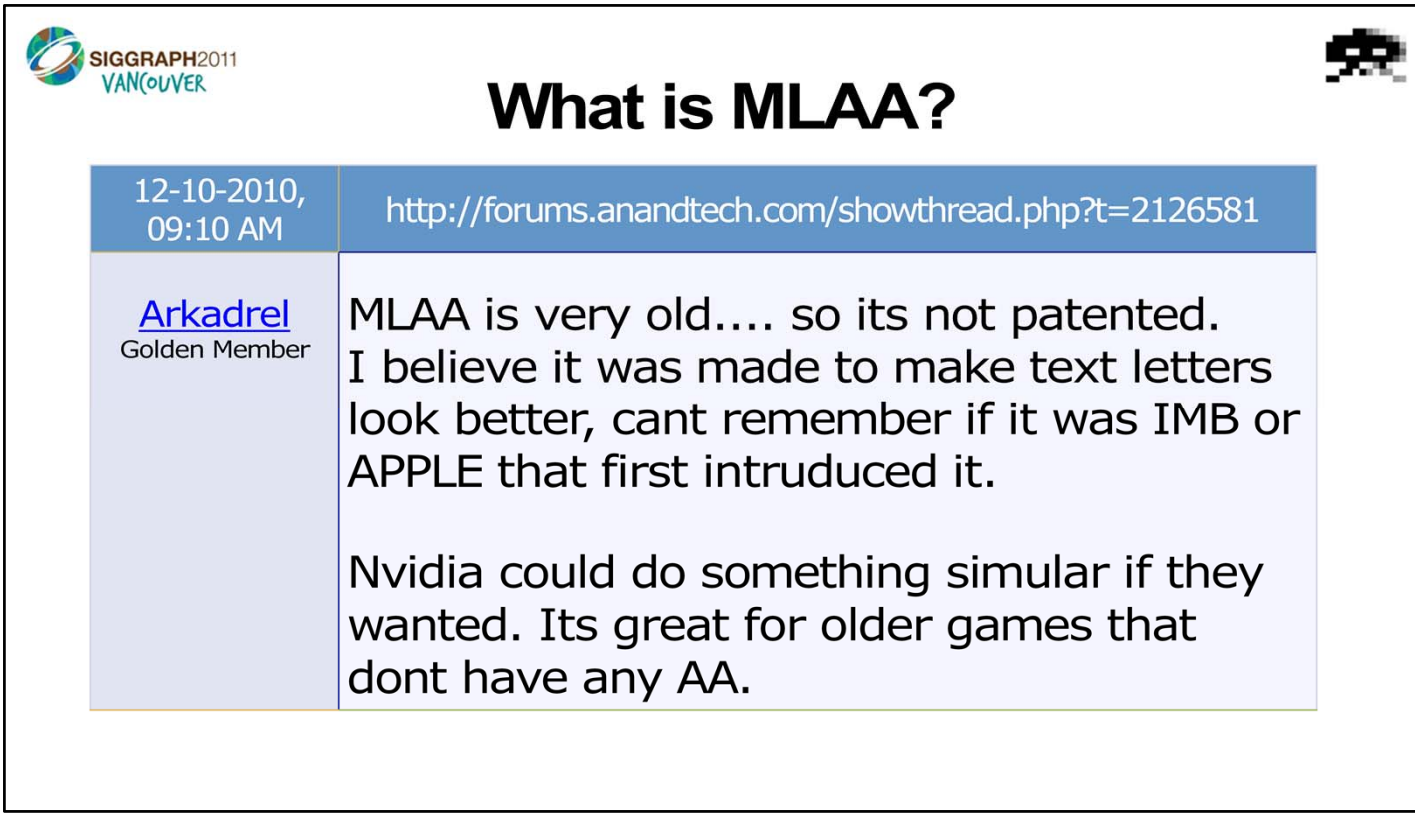

This is a quote from one of the AnandTech forums (orthography preserved).

Well, there is some truth in this statement. M.L.A.A. is, indeed, not patented. Unfortunately, it makes small text look terrible. And, obviously, Nvidia wanted, and did something similar.

Arkadrel, if you're out there, sorry for mocking your sincere posting. For legal reasons, I have to apologize every time I am showing this posting.

SIGGRAPH2011 VANCOUVER

This talk: MLAA in retrospect

Let's talk about what MLAA actually is.

I will describe the original morphological anti-aliasing algorithm from a two-year perspective.

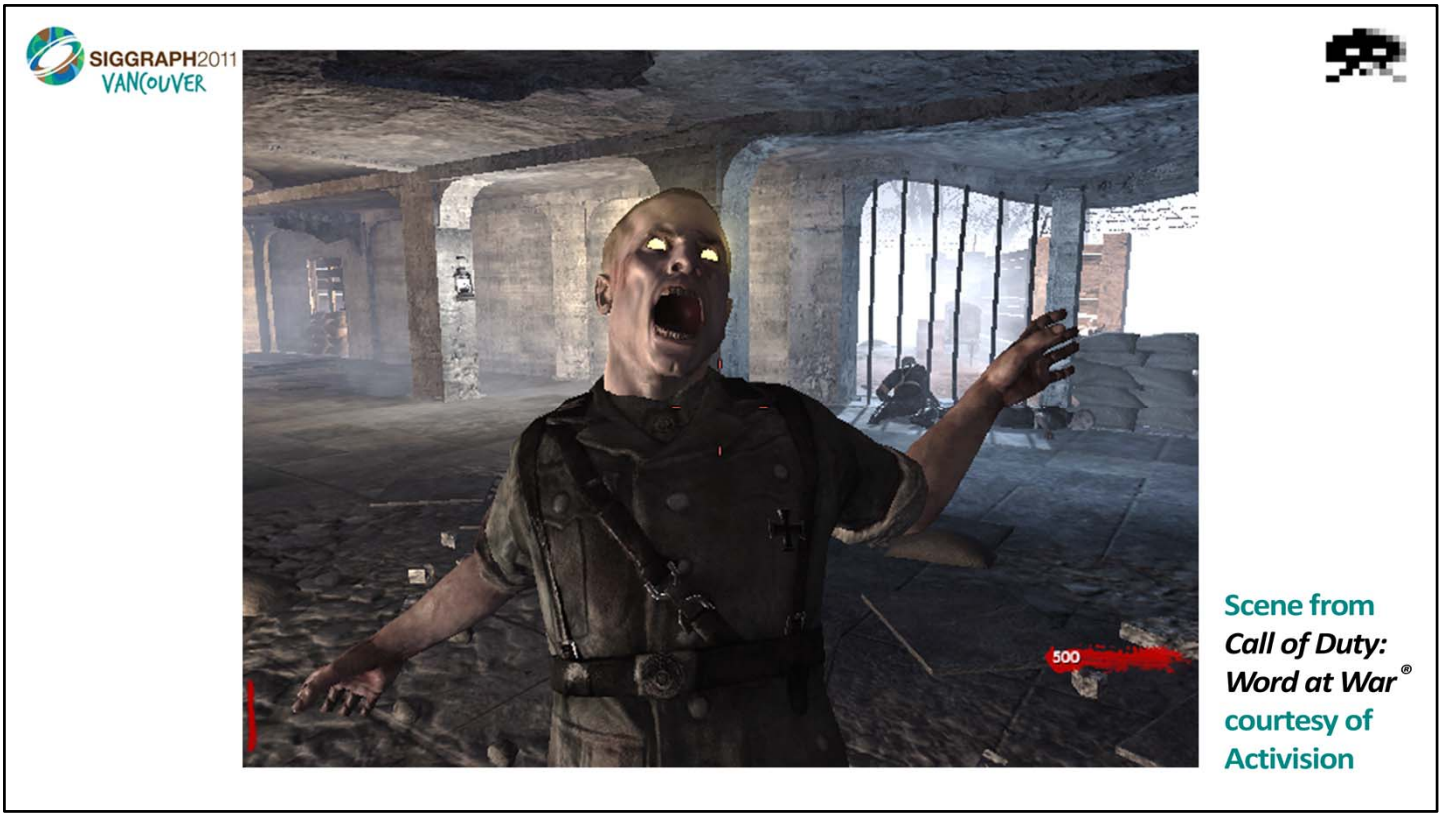

This might look like a very nice picture…

but if we zoom in, we will see that it is actually ugly. All these zigzag edges... you would never see this in a real life.

On the other hand, we can still recognize eyes, the general face oval, and other features.

SIGGRAPH2011 VANCOUVER

The Plan

1. Somehow find silhouettes in images
   (and hope that it will correspond to real objects)
2. Blend (*aka* filter) colors around the silhouettes

So this is our plan. We want to:

- somehow find silhouettes in images

and then

- blend colors around the found silhouettes.

SIGGRAPH2011 VANCOUVER

Meaningful similarities between…

- post-processing antialiasing
- super-resolution
  - Raanan Fattal. Image Upsampling via Imposed Edge Statistics. Siggraph 2007.
- computer vision

➡ recovered (aka *hallucinated*) silhouette edges are used for image enhancement / recognition

Silhouettes come handy in many computer science disciplines.

In all these methods, plausible silhouettes are first *hallucinated* and then used for image processing.

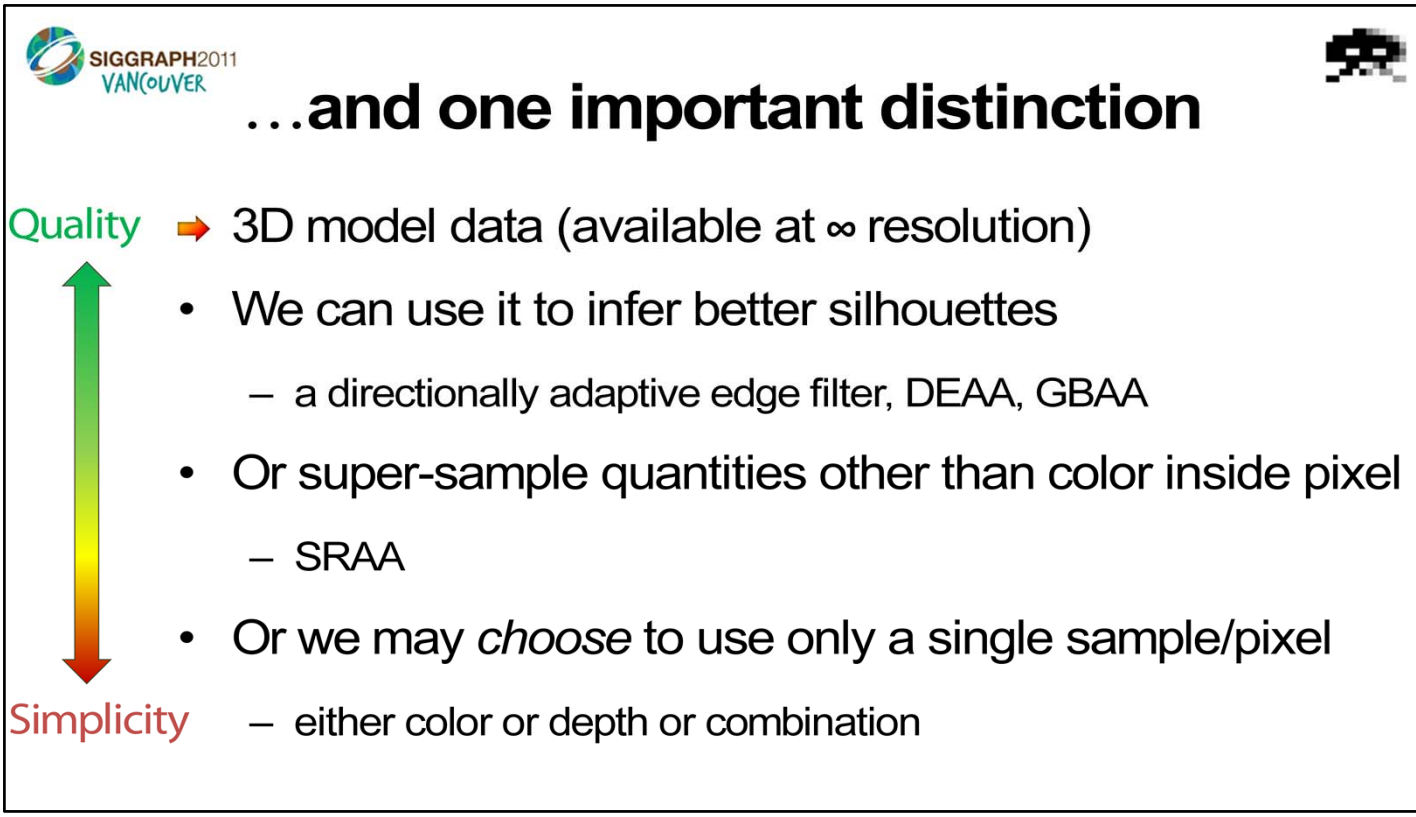

At the same time, in rendering, the exact silhouette can actually be computed from the available 3D data. Such approaches will be presented later in the course.

If we ignore all this additional data and just go with a single color sample per pixel, we will get the simplest possible and the most universal algorithm, but the price for this is quality.

SIGGRAPH2011 VANCOUVER

Why (we hope) it will work

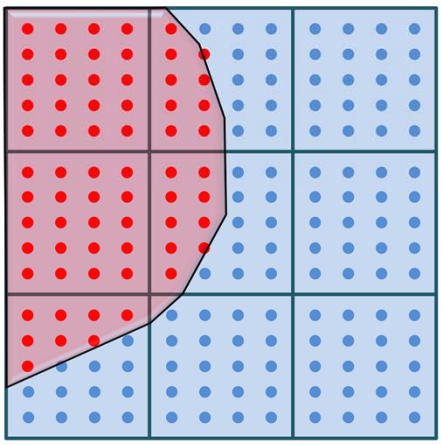

Super-Sampling Anti-Aliasing:

1. sample each pixel

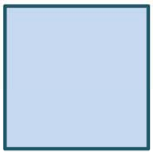

2. average computed colors

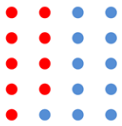

Best quality, of course, could be achieved with super-sampling, which is the gold standard for anti-aliasing, since it emulates integration processes in a camera or a human eye.

SIGGRAPH2011 VANCOUVER

Simplifications

- For pixels with 2 distinct sampled colors, integral can be approximated with area computation:

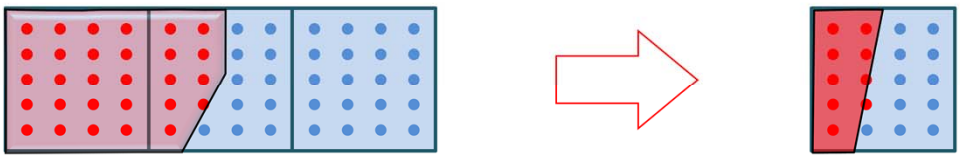

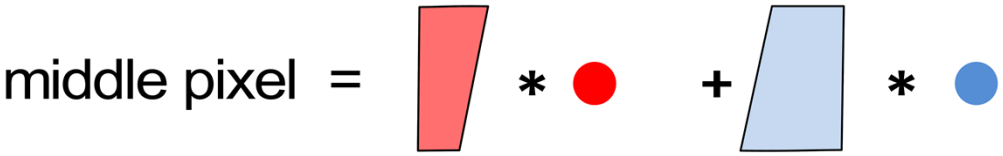

(comes from the left pixel, — from the middle one)

For some pixels though the sampled colors will not be significantly different; and we will just do unnecessary work super-sampling such pixels.

The simplest non-trivial case is a pixel with 2 distinct sampled colors. It usually happens when a silhouette line goes through the pixel.

For such pixels, the integral can be approximated with area computations.

It might even work if we know only a single color per pixel and somehow able to estimate silhouette lines.

SIGGRAPH2011 VANCOUVER

It was done before…

- For a very simple content, pixel *art scaling* algorithms may work
- Developed in 80's to allow original low-res computer games run on better hardware (Wikipedia)

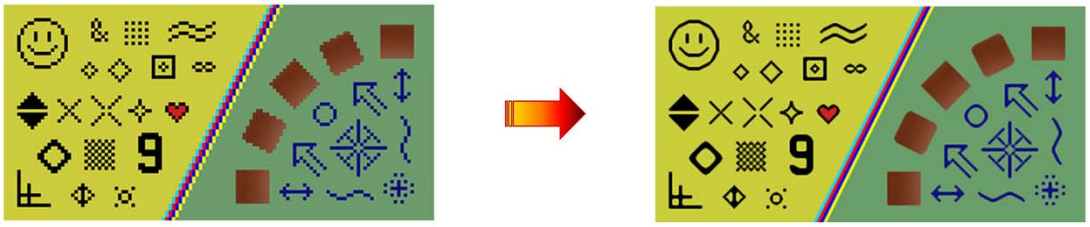

- (see also Johannes Kopf, Dani Lischinski. Depixelizing Pixel Art, Siggraph 2011)

It was done before. In early 80s pixel art scaling algorithms were developed to allow older computer games run on higher resolution displays.

These algorithms use a small neighborhood of a current pixel to compute pixels at higher resolution. It works because the number of potential patterns is limited for older games.

Recently, Kopf and Lischinski come with a clever dipixelizing algorithm, which uses higher order curves.

SIGGRAPH2011 VANCOUVER

What we need

Boolean data (which pixels are different)

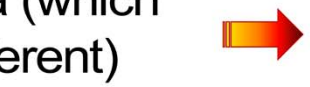

continuous silhouette lines

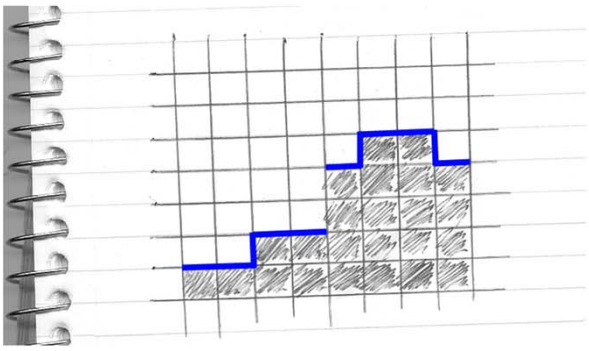

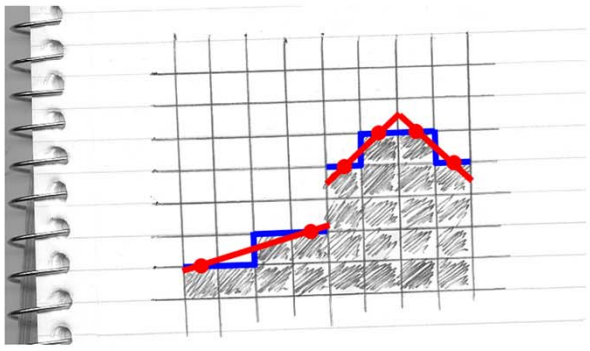

Their algorithm is too expensive for a real time.

We take an inspiration in pixel art algorithms but expand them to allow a general content considering non-local patterns.

We will recreate silhouettes using only Boolean data describing which pixels are different.

This is not the only possible approach, but its advantage is simplicity. This Boolean oracle can use any input data such as color or depth, but its output is either true or false for any given 2 pixels.

SIGGRAPH2011 VANCOUVER

How to decide if pixels are different

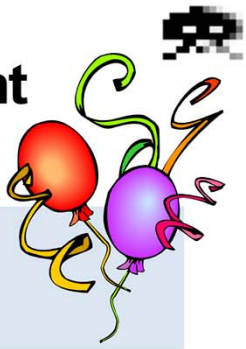

Method	Features
threshold for each color channel	• ≠ human vision • issues with illumination changes near silhouettes
luminosity [ITU-R BT. 709]	• false negatives
Non-linear thresholding (in GOW)	• good detection over the whole range • requires artist's adjustment
depth only	• choosing a scale is difficult • problems with corners
depth + color + object's id + …	• perhaps, the best one (if data is available)

The simplest way is to use only color data, separately for each channel. With this method, it could be difficult to choose a threshold, since the human vision is non-linear.

Small values will result in needless silhouettes, while bigger values will not allow finding all differences.

Another possibility is to quantify luminosity (for example, at 10%), based on ITU recommendations which were designed to match the human vision. With luminosity alone, though, false negatives are possible.

Non-linear transformations may be used as well, to provide a good detection over the whole range of colors, as was done in the game *God of War*.

If only depth is used, it can also exhibit undesirable traits, but, it seems a combination of depth, color, and object ids is the best choice, but it will be application specific. Jiminez's MLAA uses this approach.

The pixel discontinuity data does not necessarily lend itself to easily traceable silhouettes in all cases.

If we go back to our fella, his eyes are easily identifiable.

At the same time, he definitely has some dental problems, perhaps from an unhealthy diet.

General rule is that the bigger differences between pixels the more reliable the silhouette extraction is.

This is good, as it allows to eliminate the most noticeable aliasing artifacts.

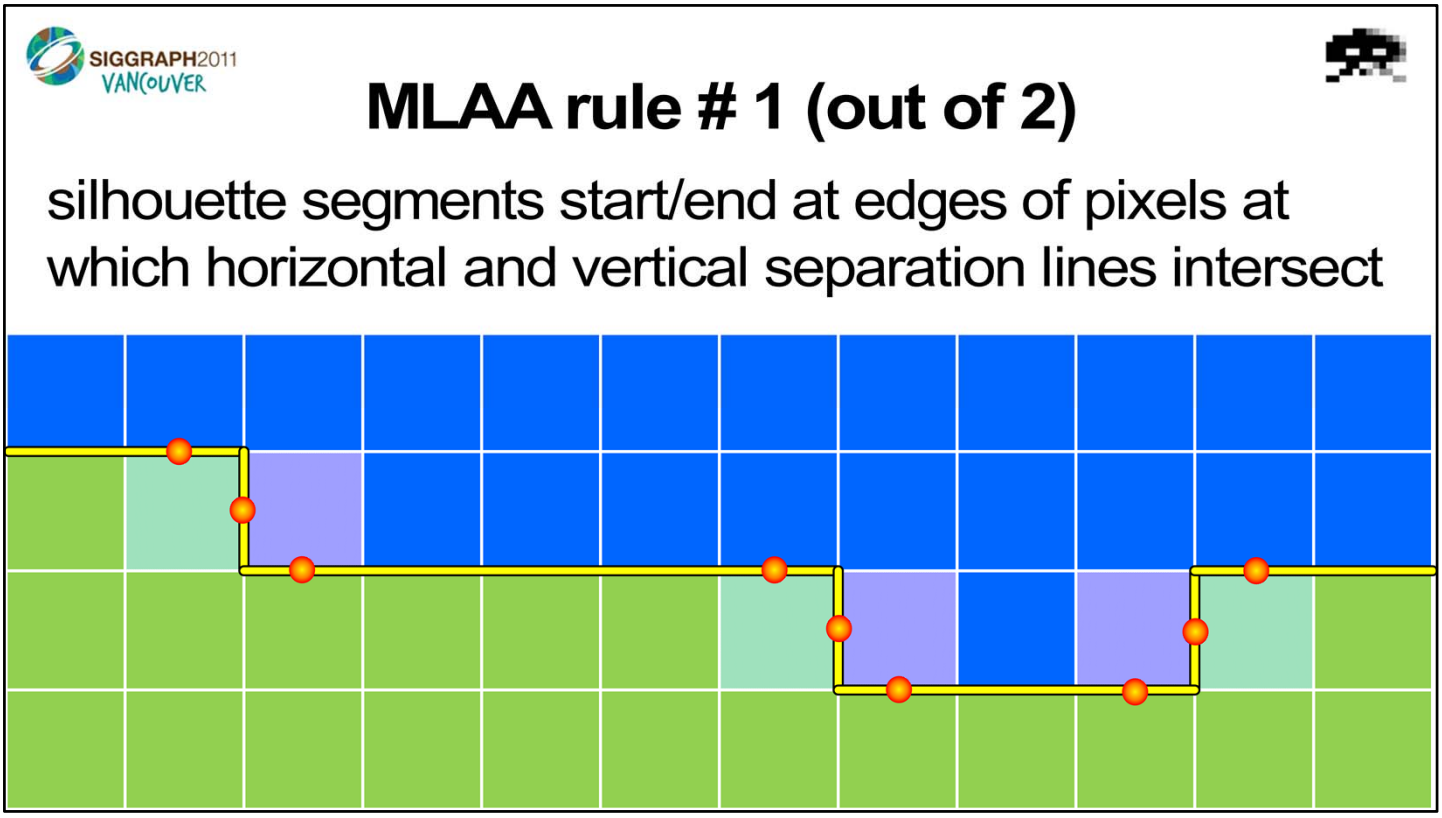

We will have just two rules for extracting silhouettes. On this image, yellow line separates different pixels.

Some pixels will border both horizontal and vertical separation lines (all such pixels are shown with stripped shading).

In M.L.A.A., silhouettes could start and end only at edges of such pixels.

A simplest way will be to use half-edges though in 09 H.P.G. paper some fancier way to find intercepts was presented, using color balance.

SIGGRAPH2011 VANCOUVER

MLAA rule # 2

for each separation line

- look at all start/end points on adjacent orthogonal lines
- choose the longest segment

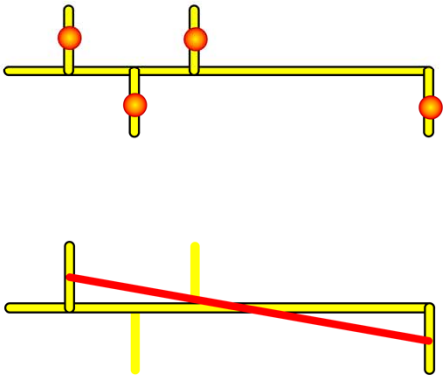

Once we find the end-points for silhouette lines, we have to connect them.

And for each separation line, we choose the longest possible silhouette fragment.

This is a difference between original M.L.A.A. and Jimenez's M.L.A.A. in which the first found silhouette is used.

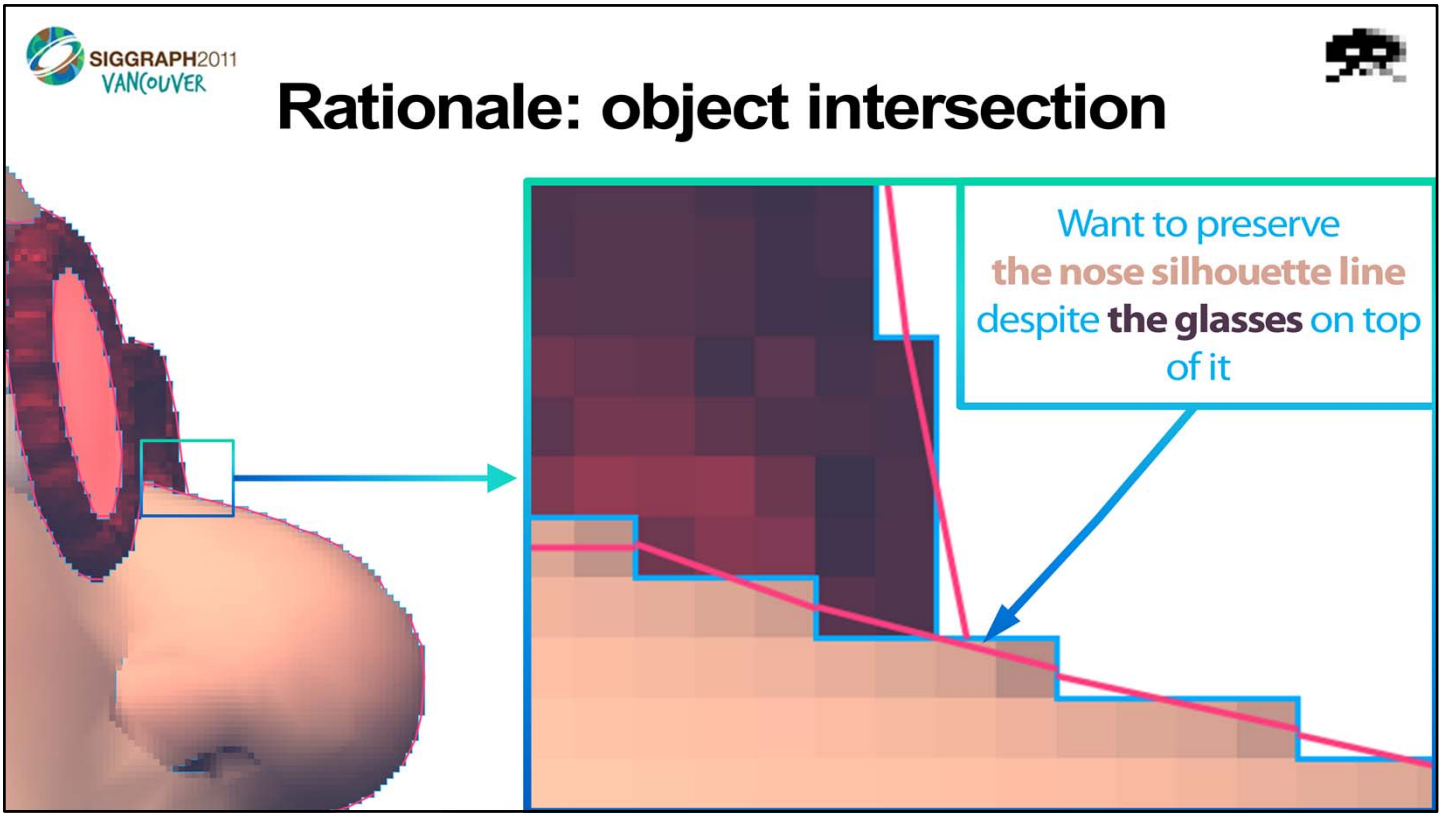

Rationale for this is to preserve the object silhouette despite other objects intersecting it as shown on this Edgar image.

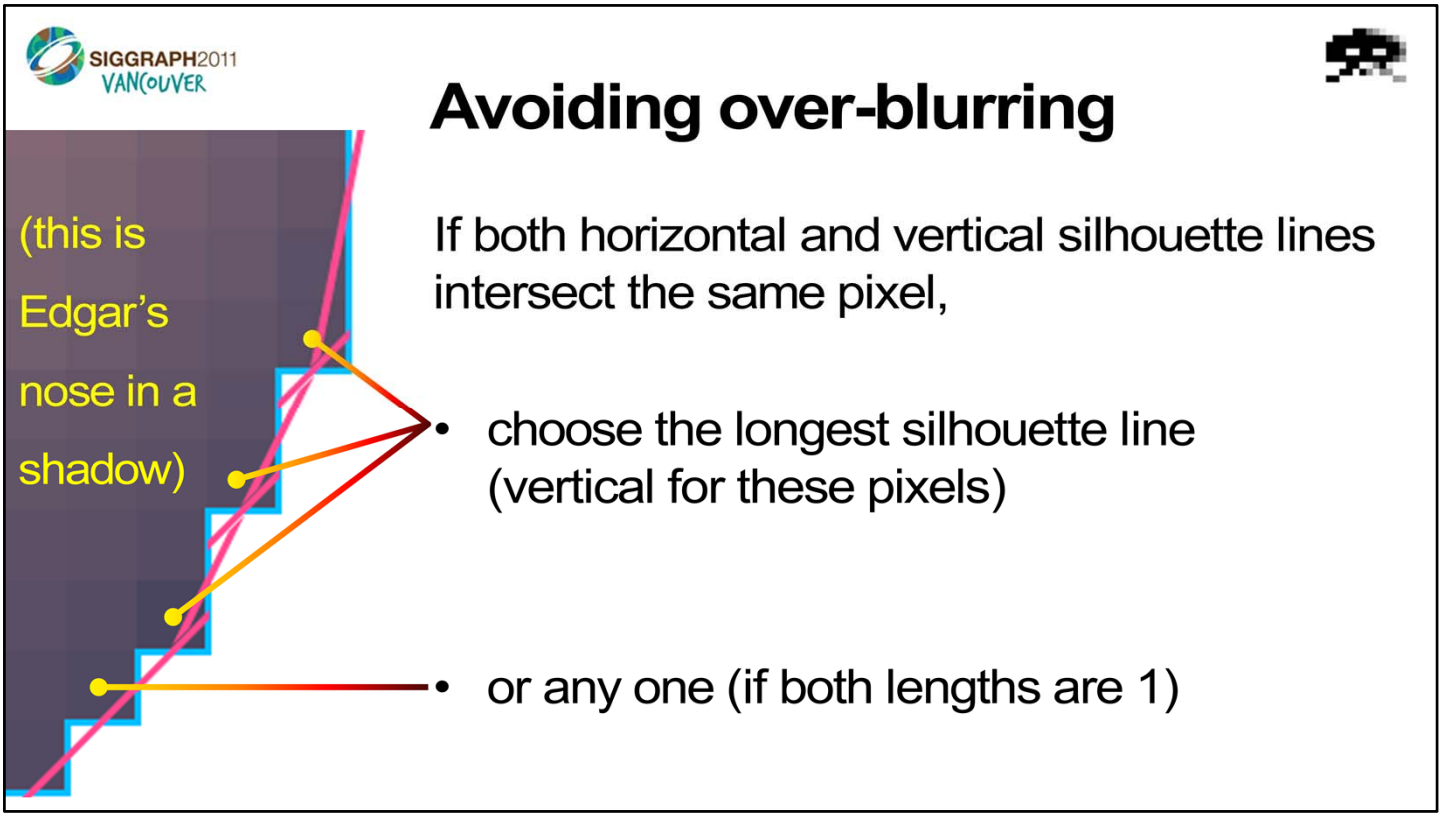

If both horizontal and vertical silhouette lines intersect the same pixel, we could choose the longest silhouette line to avoid over-blurring.

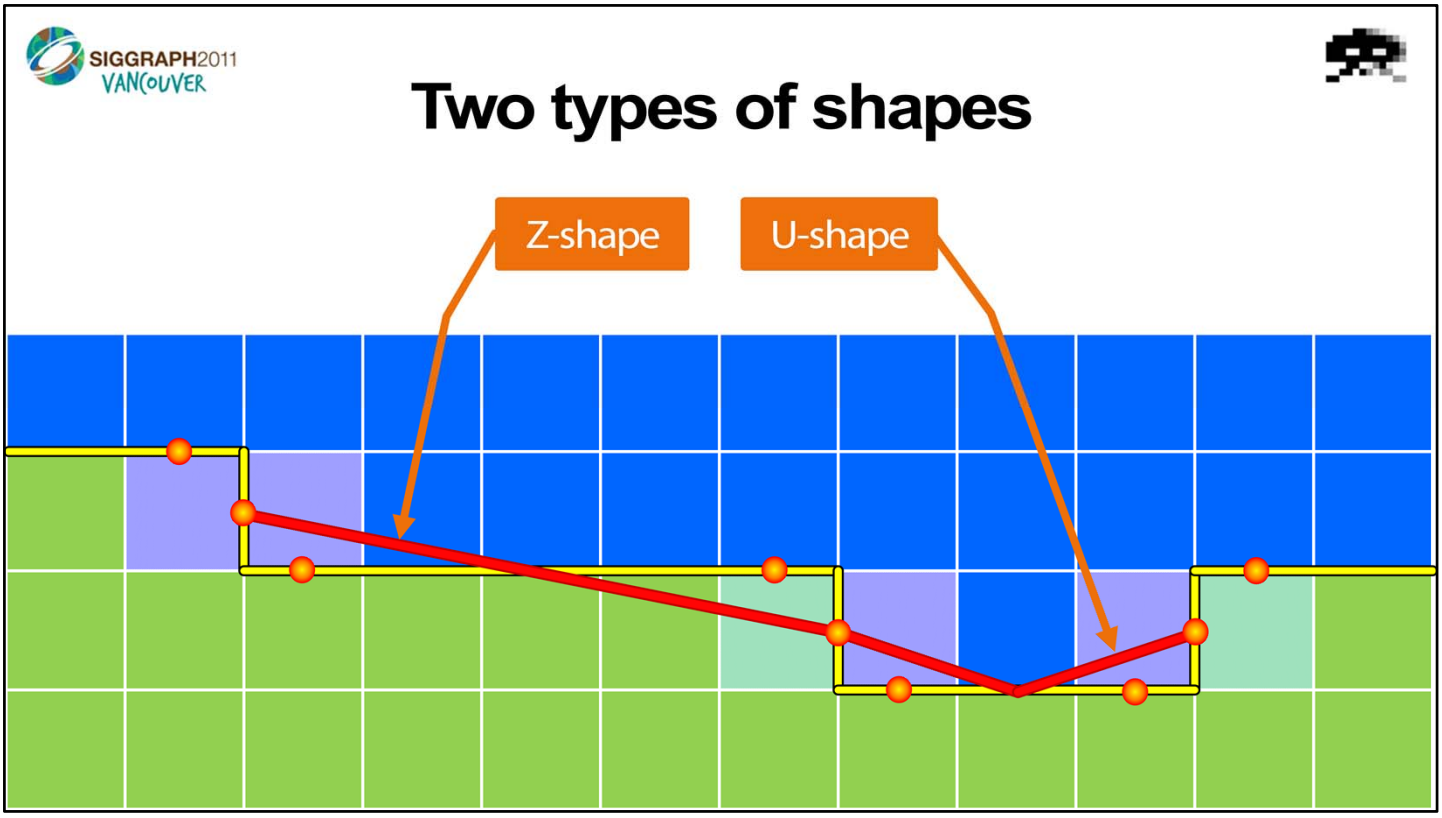

And, finally, if the silhouette end-points are on opposite sides of the separation line, we will get a so-called Z-shape. Otherwise, a U-shape will be created, which consists of 2 linear segments.

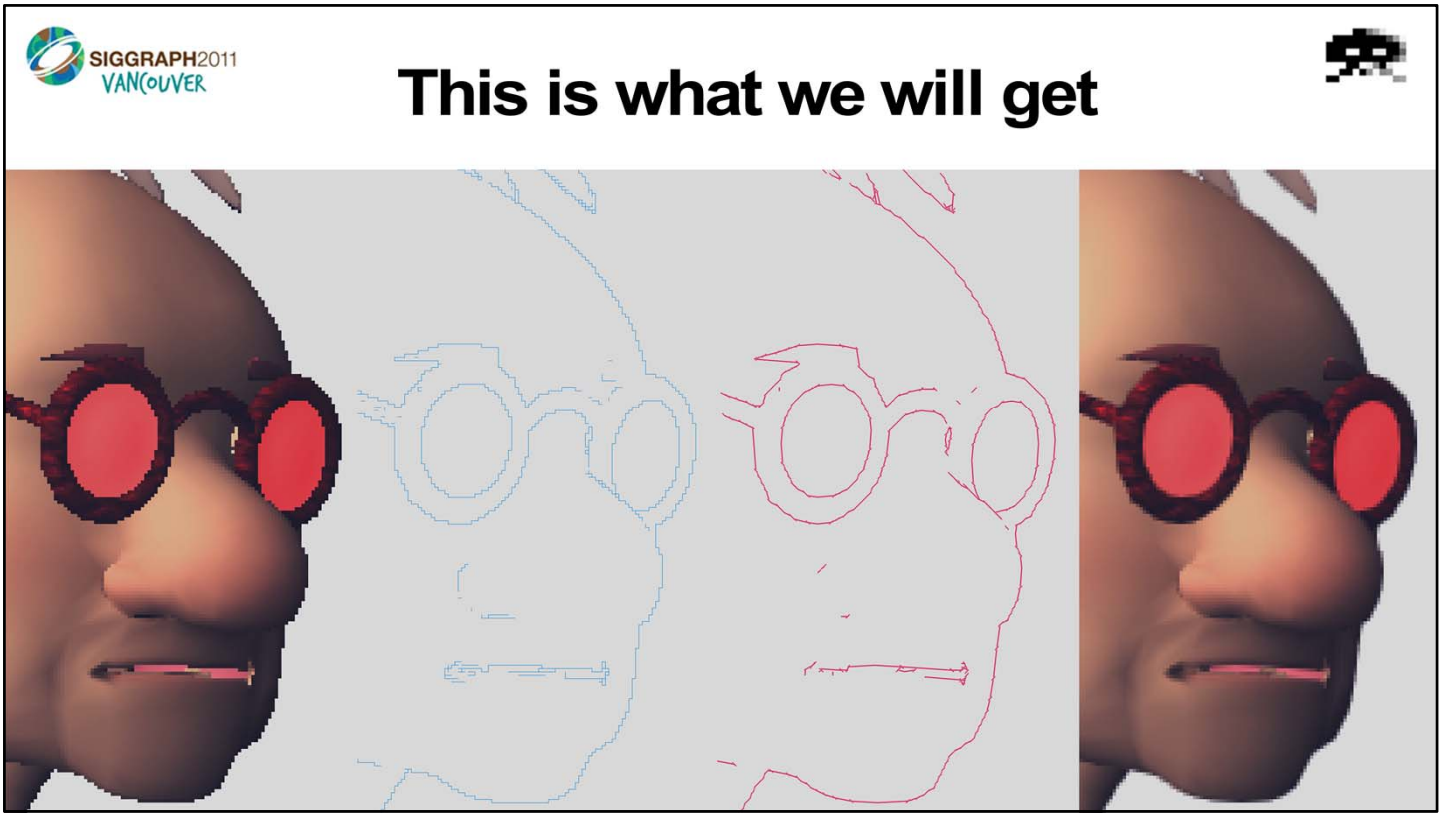

So, putting everything together, from left to right: original image, separation lines, reconstructed silhouette, and anti-aliased image.

SIGGRAPH2011 VANCOUVER

MLAA in a one sentence

- (1) detect all pixels that are different from neighbors to
  (2) approximate silhouettes and then
  (3) filter colors around these silhouettes
- Steps 1 and 2 allow innovation and differentiation
- Step 3 *seems* to be OK in RGB space (without gamma)

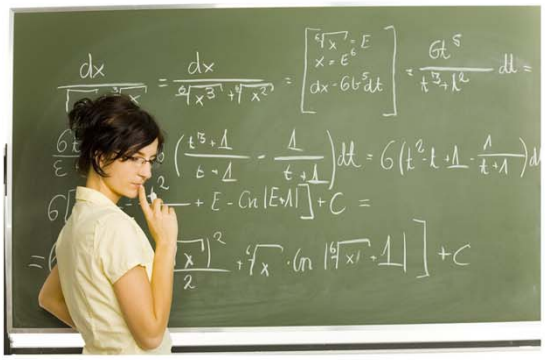

In a nutshell, M.L.A.A. uses silhouettes hallucinated from discontinuity data, to filter colors around such silhouettes.

Other rules are entirely possible. Most will result in similar silhouettes in good situations, and we don't have enough information in bad situations anyway.

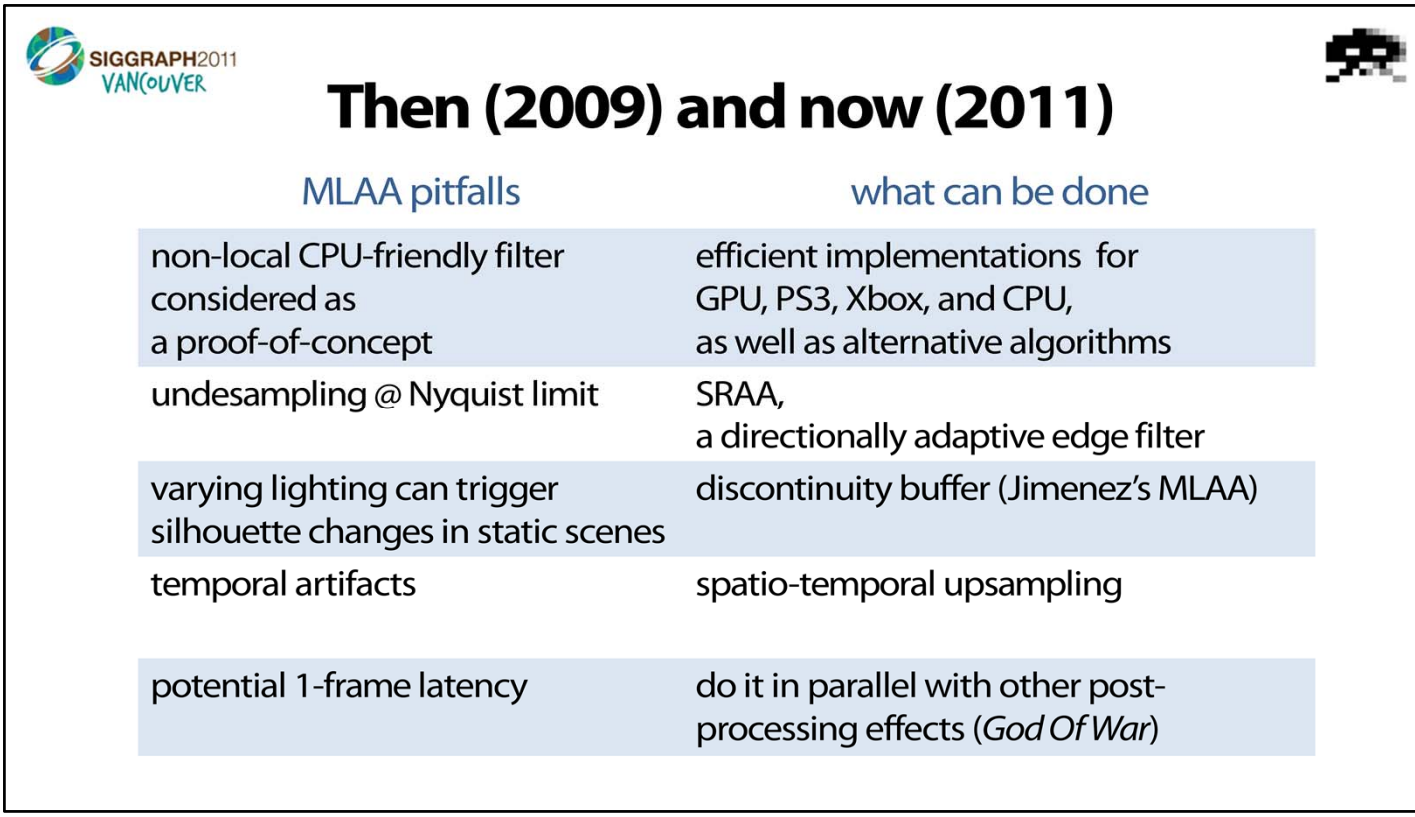

Original M.L.A.A. approach, when it was presented at H.P.G. 09, was more like a proof of concept.

In fact, I didn't even mention deferred rendering, but talked about Leonardo da Vinci, Kazimir Malevich, and Georges Seurat.

Now, most of the original M.L.A.A. shortcomings are addressed, one way or another.

And there is even a new CPU version, executed in about 6 ms, implemented by Alexandre De Pereyra from Solutions group at Intel.

He did a thorough job optimizing and commenting the code, so if you want to look on the code which you can understand, download this version.

SIGGRAPH2011 VANCOUVER

Timeline for 2020?

- AA Naming Guide on AnandTech:

  27 entries for major variations

- Historical perspective: Z-buffer

  killed all *other* invisible surface remc

  algorithms…

- Hardware AA was unable

  to do it (yet?)

Anti-aliasing Naming Guide on AnandTech has 27 entries (and some techniques presented today are not even included yet).

It spells both problems and opportunities.

It is not even clear whether post-processing anti-aliasing is for a long time, or more clever hardware multi-sampling will have an edge.

SIGGRAPH2011 VANCOUVER

So the question is…

- Will retina displays (~300 dpi) kill all AA?
  - (*it will be exciting*)
- Bottom line: it seems that post-processing AA algorithms have matured in time when
  - resolution is good enough to alleviate certain artifacts
  - But not too high to forget about AA at all

One of the major questions is what display resolution is good enough?

This was my original slide, discussing these issues, and it is on the conference DVD. But, for some reason, I did not receive a complimentary iPad.

SIGGRAPH2011 VANCOUVER

So the question is… (*amended*)

- Even 300 dpi are not enough to forget about AA
  - People evolved to notice discontinuities @ higher frequency than eye's resolution (hyperacuity)
- You can read more (see the course web site)
  - John Hable's blog
  - David Luebke's The Ultimate Display

So this is a revised slide. As you see, even 300 dpi are not enough. It is a fascinating topic, you can read more about it on the course web site.

Next talk: Jorge Jimenez on Practical MLAA

this one is MLAAsed ⇨

⇦ this one is not

(if you can read it, you can see it)

This concludes my presentation, thank you very much.