



Video Streaming for Foveated High-resolution Rendering



Master's thesis, Marc Aurel Kastner
Supervised by M. Stengel, Prof. M. Magnor



- Panoramic scenes create enormous FOV
 - Video data continuously increases
 - Resolution: 2K → 4K → 8K → ...?
 - Frame rate: 25 fps → 60 fps → ...?
 - Unable to process *full* high resolution frames
- Need for high resolution with low bandwidth

Source: Flickr
Grzegorz Rogala



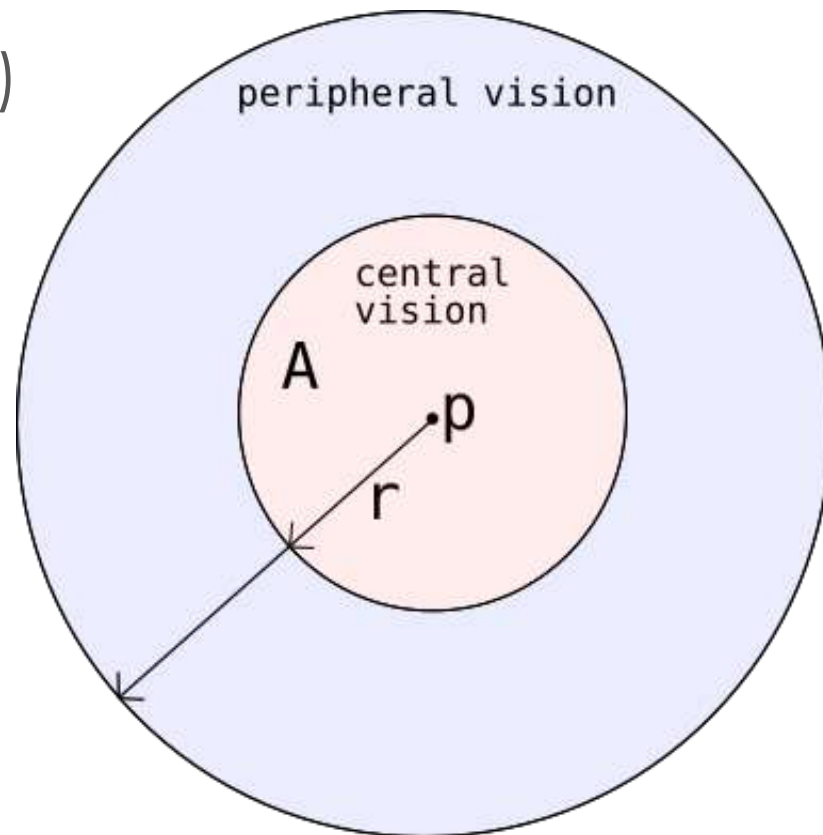
Foveated imaging



Source: Wikipedia



- Simplified model:
 - Aubert / Foerster (1857)
 - Linear fall-off until 20°
 - Then, strong drop
- Still commonly used
 - Conservative
 - Simplicity





Foveated MPEG (Geisler et al. 1996)


- Static / Passive
 - Saliency estimation
 - Molding acuity into data
- e.g. sport & news clips
- Prone to error
 - Users might look in different directions
 - Not well optimized to one users' view

Foveated 3D Graphics (Gunther et al. 2012)

- Dynamic
 - Pupil Tracking
 - No data modifications
 - e.g. rasterization, raytracing
 - Adapts to users' eye
- Make this for videos



- Recent VR often use Head mounted Displays
- HMDs allow pupil tracking
 - Active gaze estimation
 - Knowledge about users' field of view





- Decrease bandwidth
 - Use gaze estimation
 - Filter resolutions simulating acuity
 - Optimize streaming
- Restraints
 - Dynamic
 - Performance
 - Perception



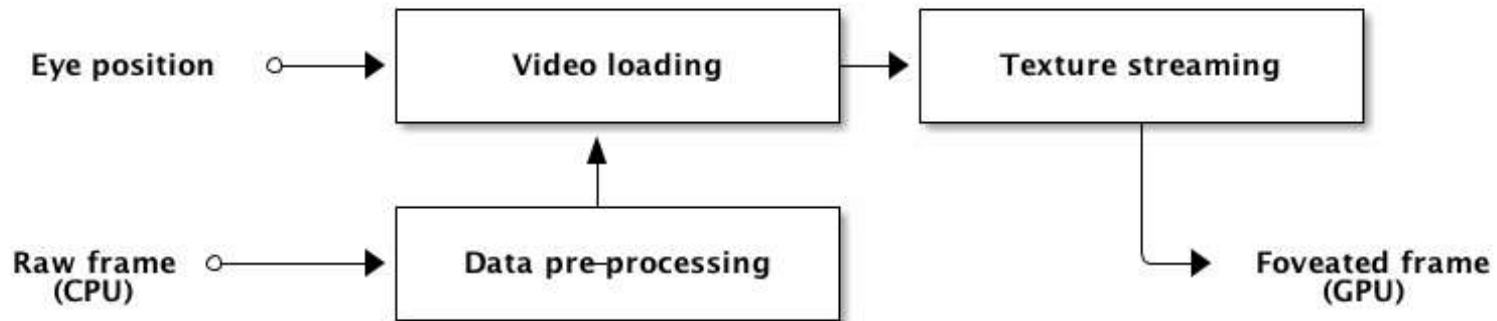
Source: Wikipedia



- Motivation / Idea
- Related work
- Design
- First approach
- Second approach
- Evaluation



- Three main pillars
 - Pre-processing
 - Video loading
 - Texture streaming

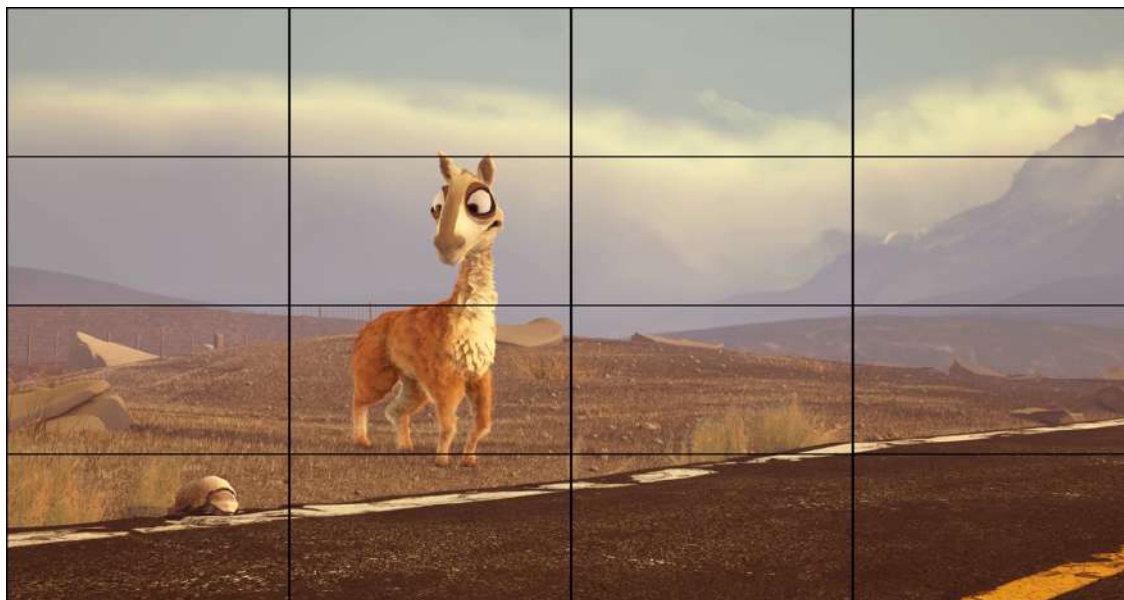





- Video formats don't allow region or pixel access
- Impractical to decode *full* frames
 - Access smaller videos
 - Have multiple quality versions
- No relation to eye position
 - Dynamic




- Uniformly distributed
- Spatial sub-division
- Used for
 - Ray tracing
 - Video tiling







- Idea
 - Pre-process scaled down variants of textures
 - Save different resolution levels
 - Reduce stress on GPU, simplify filtering, avoid moiré patterns, ...
- Often used in video games



Source: Wikipedia




- Idea
 - Use video tiling
 - Create multiple mip-map levels per video tile
- Allow region-based loading






- Wide-screen videos
 - Arbitrary codec
- Acuity distribution:
Hybrid!
 - Grid-based loading
 - Sampling-based streaming





- Depending on eye position
 - Select suitable resolution
 - Threshold mip-map levels with distance
- Grid 4x4, 3x Mip-map
 - 48 video files
 - Load frames as needed





- Seeking is major bottleneck
 - Every video file is open in parallel
 - An eye movement triggers wall of seeking
- Solution
 - Multi-threading model
 - Sometimes use *wrong* mip-map levels
 - Didyk et al. (2010): Retina takes 60 ms to adapt
 - Allow system to keep up



- Two masks
 1. Acuity mask
 2. Interpolation mask
- Allows minimum bandwidth
- Linear reconstruction with look up tables

ID	X	Y
0	15	12
1	46	56
2	89	96
...

Acuity mask

0	0	1	1
0	2	3	1
4	5	6	7
4	4	7	7


Interpolation mask



Frame reconstruction

ID	X	Y
0	15	12
1	46	56
2	89	96
...

0	0	1	1
0	2	3	1
4	5	6	7
4	4	7	7



Streamed to GPU




every frame



0	0	1	1
0	2	3	1
4	5	6	7
4	4	7	7

once



Shading




- 360 degree panorama
 - Pictures per frame
 - JPEG/PNG
- Sampling may introduce visible flickering
- Acuity distribution
 - Fully grid-based
 - Radial blur in post-process reduces peripheral frequencies






Background: Cube mapping

- Six single textures for 360° surface
- GPU: One cube map per mip-map
 - One Grid per cube side



Source: David J. Eck




Source: Emil Persson



- For every frame:
 - Calculate view vector
 - Decide *mip-map*
- Update only relevant region per frame
- Acuity data as look up table




+




+



=



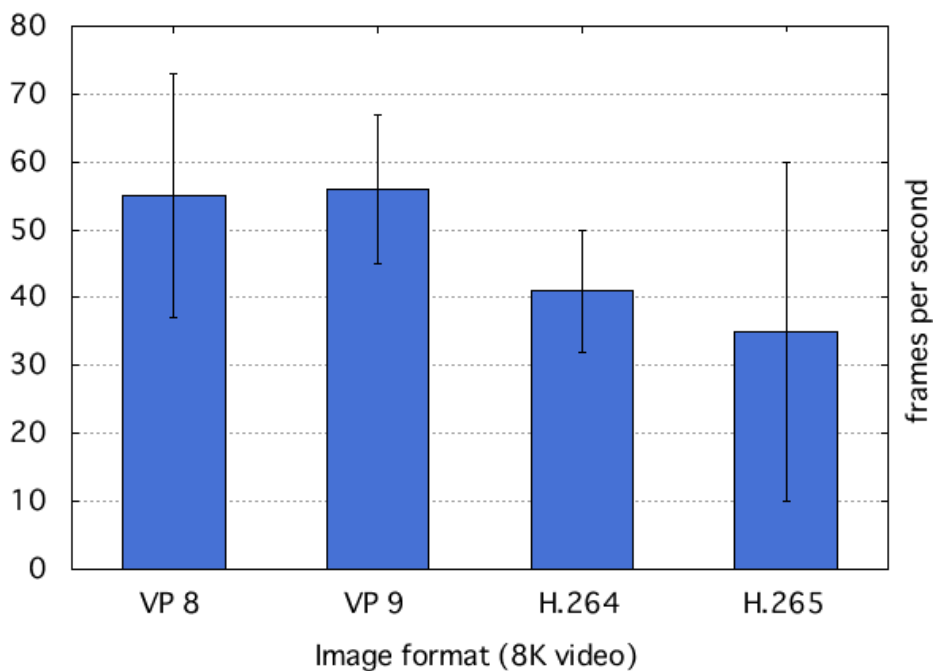






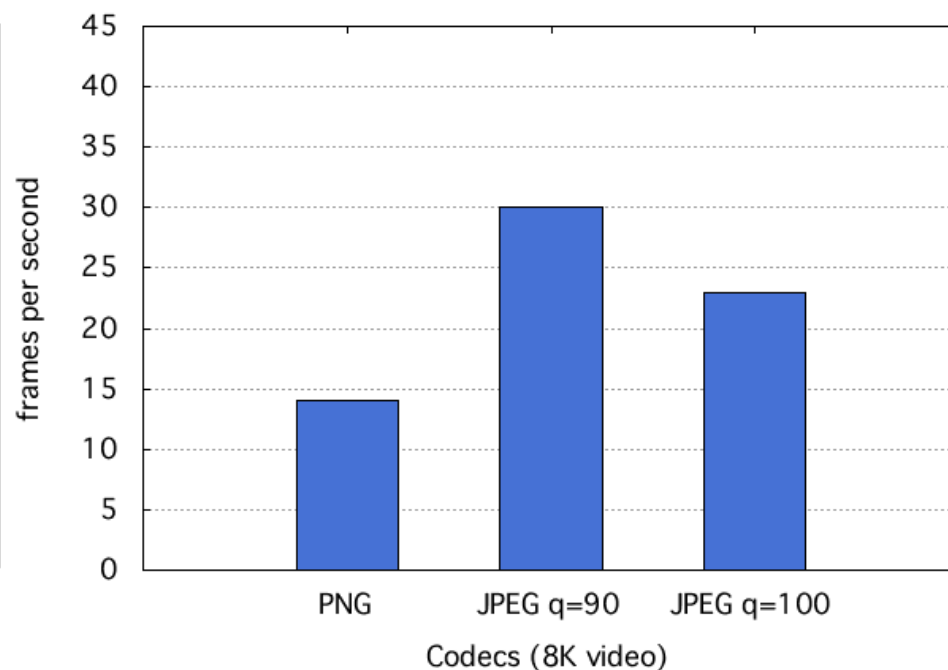
Foveated Video Approach 1

Loading performance for different codecs



Foveated Video Approach 2

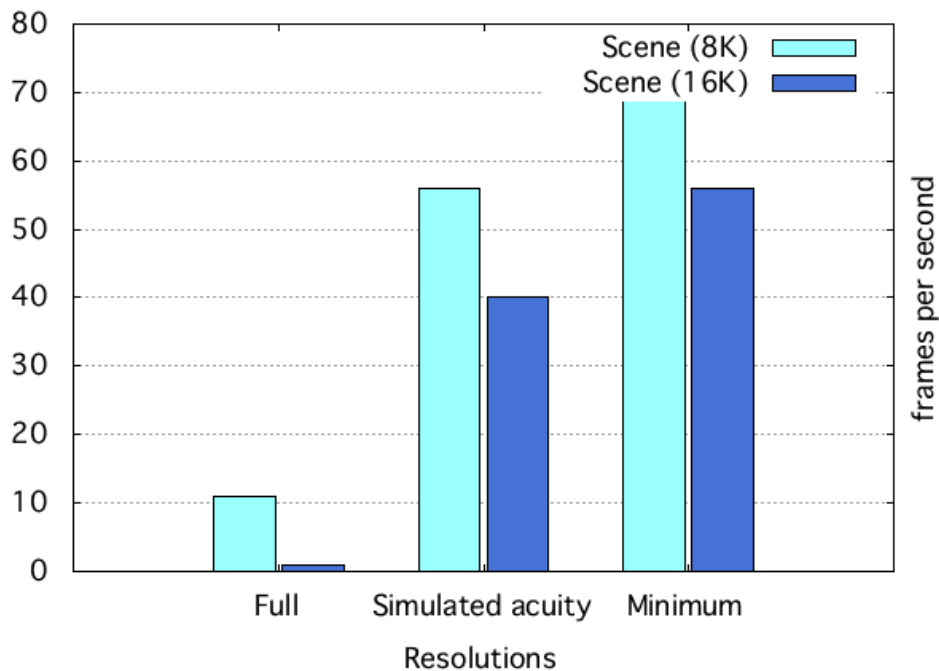
Loading performance for different image formats





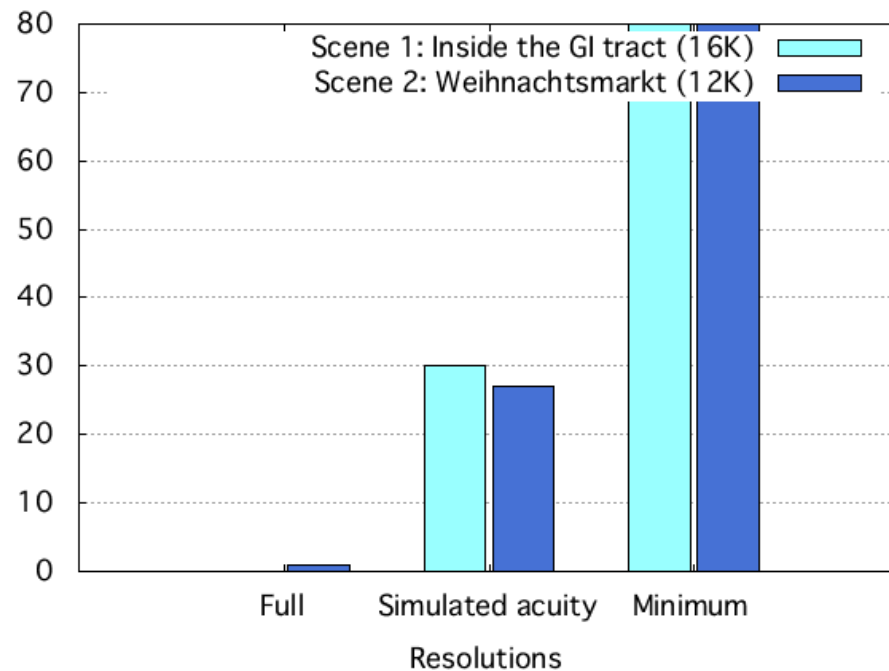
Foveated Video Approach 1

Performance with different resolution settings



Foveated Video Approach 2

Performance with different resolution settings





Wide-screen videos	Transferred pixels	Percentage
Sampling-based acuity (1080p acuity mask)	690.926	2.08%
Minimum resolution (smallest mip-map 240p)	129.600	0.39%
Full resolution (8K wide-screen)	3.177.600	100.00%

Panoramic videos	Transferred pixels	Percentage
Grid-based acuity (best case)	1.773.229	2.89%
Grid-based acuity (worst case)	2.789.232	7.74%
Minimum resolution (smallest mip-map 120p)	212.064	0.59%
Full resolution (12K panoramic)	36.000.000	100.00%



- Exploit human visual system
- Replace bandwidth with storage
 - Pre-processing is cheaper than bandwidth
- No alternative: Common video codecs very slow $>4K$
- Potential for various future apps (VR, mobile, internet)





Thank you for your attention.



<http://graphics.tu-bs.de>



Appendix



- Slide 5
 - Wilson S. Geisler et al. “Implementation of a foveated image coding system for image bandwidth reduction”. 1996.
 - Brian Guenter et al. “Foveated 3D Graphics”. In: ACM Trans. Graph. 31.6 (Nov. 2012), 164:1–164:10. issn: 0730-0301.
- Slide 16
 - Piotr Didyk et al. “Apparent display resolution enhancement for moving images”. In: ACM Transactions on Graphics (TOG). Vol. 29. 4. ACM. 2010, p. 113.
- Full bibliography in thesis p. 69-76



- Slide 2
 - Grzegorz Rogala, Creative Commons
https://www.flickr.com/photos/grzegorz_rogala/5097827282/
- Slide 3
 - JeffPerry, Public domain
https://en.wikipedia.org/wiki/Foveated_imaging
- Slide 12
 - Mulad, Creative Commons
<https://en.wikipedia.org/wiki/Mipmap>
- Slide 20
 - Emil Persson, Creative Commons
<http://www.humus.name/index.php?page=Textures>
 - David J. Eck, Creative Commons
<http://math.hws.edu/graphicsbook/c5/s3.html>
- Evaluated scenes
 - Appendix Slide 34
- Everything else created by myself or ICG



- Hardware

- MacBook Pro
- 2,5 GHz Intel Core i7 Haswell CPU
- NVIDIA GeForce GT 750M 2048 MB
- 16 GB DDR3 RAM
- SSD


- Software

- SDL 2.0
- OpenCV 3.1
- FFmpeg 2.8.6
- STB_Image 2.8
- PIL 3.0

- Mac OS X 10.11.3
- C++ '14 (Main)
- Python 3.5.1 (Pre-Proc)



- Scene 1: Caminandes
 - Wide Screen
 - Original 4K, upscaled 8K/16K
 - Blender Foundation (Creative Com.)
 - <http://www.caminandes.com>
- Scene 2: Inside the GI tract
 - Panoramic
 - Original 4K, upscaled 16K
 - HybridMedical (Creative Commons)
 - <https://www.youtube.com/watch?v=upSnH436ya8>
- Scene 3: Christmas market
 - Panoramic
 - Original 12K
 - Captured with GoPro (Matthias)





- Grid
 - 8x8 for widescreen 16K
 - 4x4 everything else
- Mip-map levels: 3
 - Full resolution
 - Mid resolution = 1 / 4 times full
 - Low resolution = 1 / 16 times full




Approach	Data	Percentage
Acuity mask	2.76 mb	0.69
Interpolation mask	44.00 mb	11.06
Total used by approach	46.76 mb	11.75
Total used by classic video	398.00 mb	100.00

Approach	Data	Percentage
Cube map (high resolution)	1152.0 mb	0.69
Cube map (mid resolution)	72.0 mb	11.06
Cube map (low resolution)	4.5 mb	11.75
Total used by approach	1228.5 mb	106.64
Total used by classic video	1152.0 mb	100.00



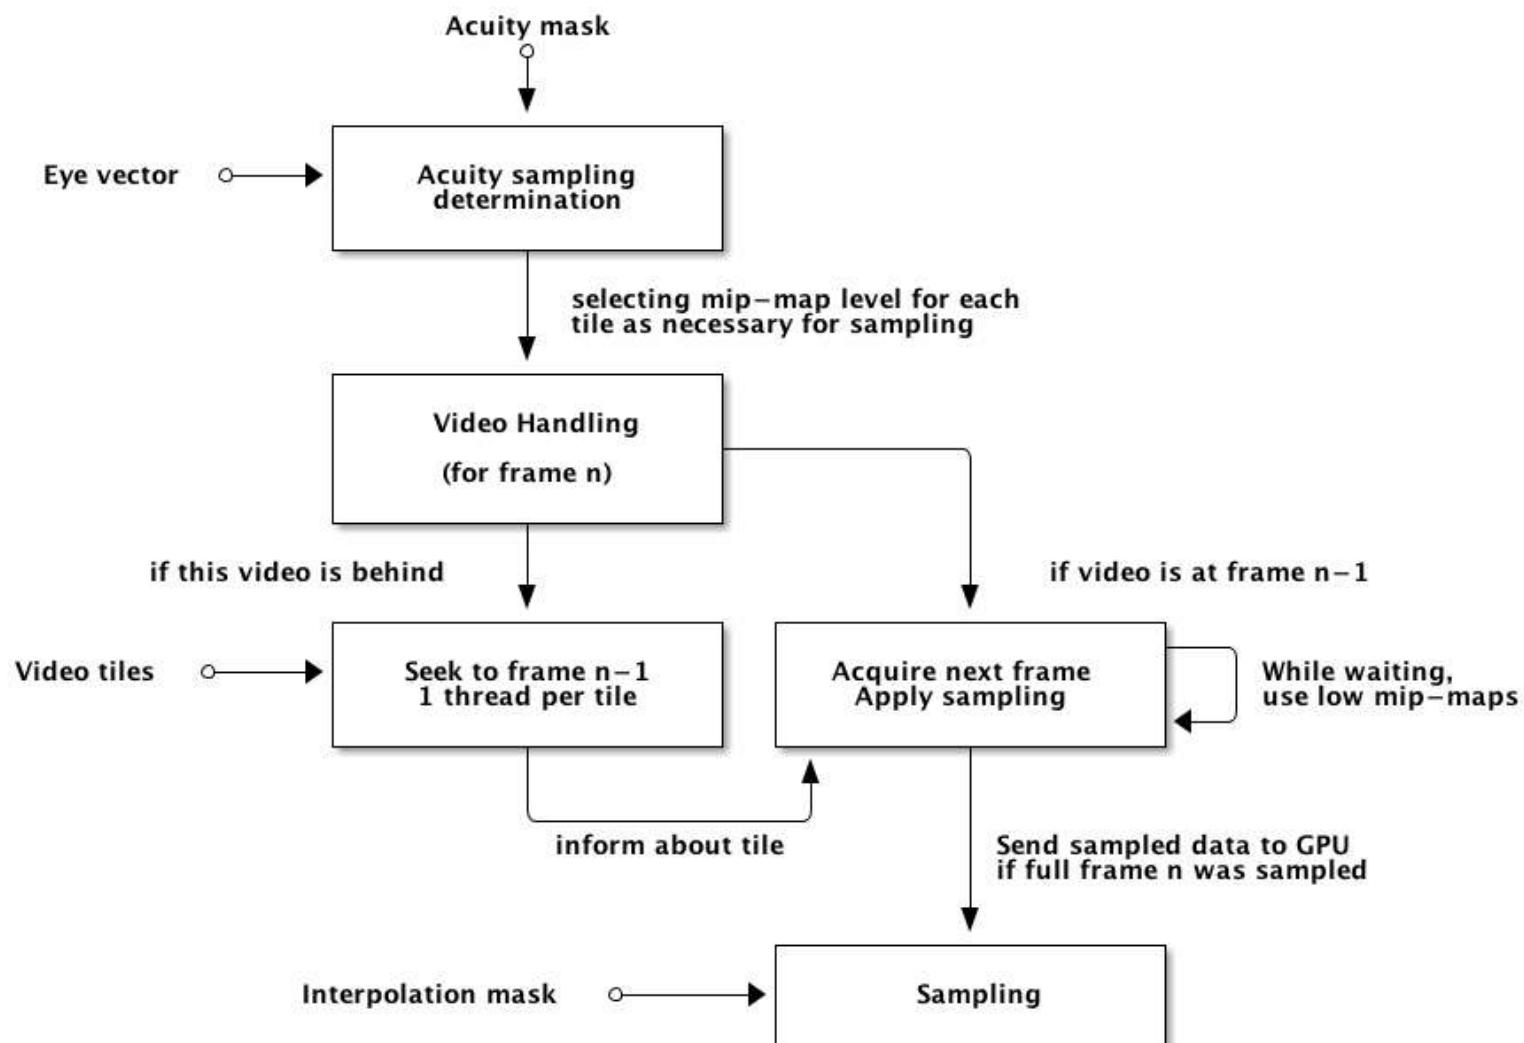
Approach	Time spent	Storage used
Framework 1: VP 8	3 h 58 min	3237 mb
Framework 1: VP 9	8 h 17 min	2294 mb
Framework 1: h.284	0 h 36 min	1923 mb
Framework 1: h.285	1 h 30 min	104 mb
Framework 2: PNG	50 min	1608 mb
Framework 2: JPEG q=90	17 min	816 mb
Framework 2: JPEG q=100	21 min	2500 mb

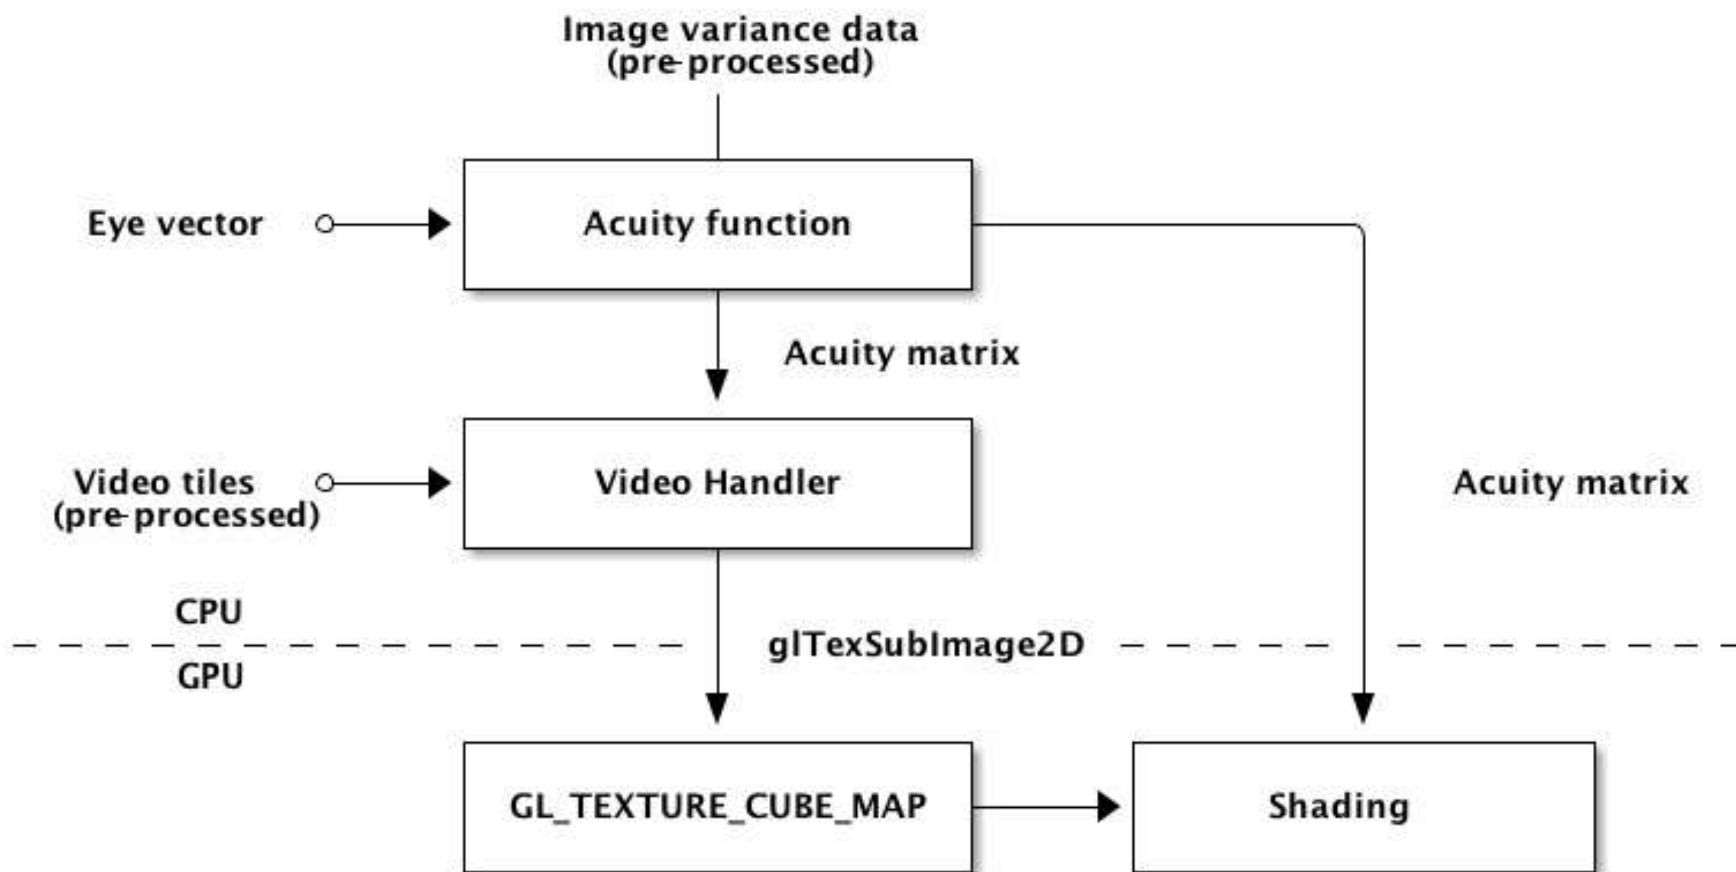
Approach	Mip-map 0	Mip-map 1	Mip-map 2
Framework 1 (Scene 1)	1885 MB	365 mb	46 mb
Framework 2 (Scene 2)	721 MB	81 mb	14 mb
Framework 2 (Scene 3)	3520 MB	463 mb	92 mb





- Saliency-based approaches
- Post-processing
 - Radial blur filter / Gauss blur
- Data
 - Bigger Grid size
 - Higher mip-map levels
 - Less distance between mip-map levels







```
1 side, axis calculateCubeside(vec3 eye):  
2     switch eye:  
3         eye.x > 0.5: return +X, X  
4         eye.x <= 0.5: return -X, X  
5         eye.y > 0.5: return +Y, Y  
6         eye.y <= 0.5: return -Y, Y  
7         eye.z > 0.5: return +Z, Z  
8         eye.z <= 0.5: return -Z, Z
```

```
1 vec2 calculateGridPosition(vec3):  
2     side, axis = calculateCubeside(eye)  
3     otheraxis1 = (getAxisUnequal(eye, axis).a + 0.75) * 2  
4     otheraxis2 = (getAxisUnequal(eye, axis).b + 0.75) * 2  
5     return (otheraxis1 / 4.0, otheraxis2 / 4.0)
```




```
1 void selectMipmap(vec3 eye):  
2     setquality(ALL, lowest)  
3     // Grid-based acuity  
4     vec2 fovealCell = calculateGridPosition(eye)  
5     setquality(fovealCell, maximum)  
6     foreach(neighborCell : adjacent to fovealCell):  
7         setquality(neighborCell, halfway)  
8     // Variance-based acuity  
9     foreach(noiseCell : cell near maximum(color variance)):  
10        setquality(noiseCell, halfway)
```



- Idea
 - Add analysis to pre-processing
 - Color variance (RGB/HSV/LAB)
 - Hard edges (FFT)
 - Temporal differences
 - Set mip-map level not only based on visual acuity but also above analysis
 - For example:
 - use mid resolution instead of low resolution if salient feature in tile



- Current implementation: "Debug"
 - Eye position = Mouse pointer
 - No stereoscopic rendering
- Port to HMD: No obstacles
 - Code is standard C++ '14
 - All cross-platform libraries
 - Stereoscopic rendering presumably no major influence on performance



- Hardware
 - GPU least busy component
 - Due to parallel video access SSD necessary
- Pre-processing necessary
 - Mostly time consuming
 - Overhead in storage
- Grid



- Perceptual user study to verify results
- Bandwidth allows new applications
 - Mobile devices
 - Remote streaming
 - Internet/Cloud?
- Approach
 - Other video codecs
 - Container format, storage optimizations
 - Add saliency-based approaches (variances, edges)