

# Fast Compositing for Cluster-Parallel Rendering

M. Makhinya, S. Eilemann, R. Pajarola



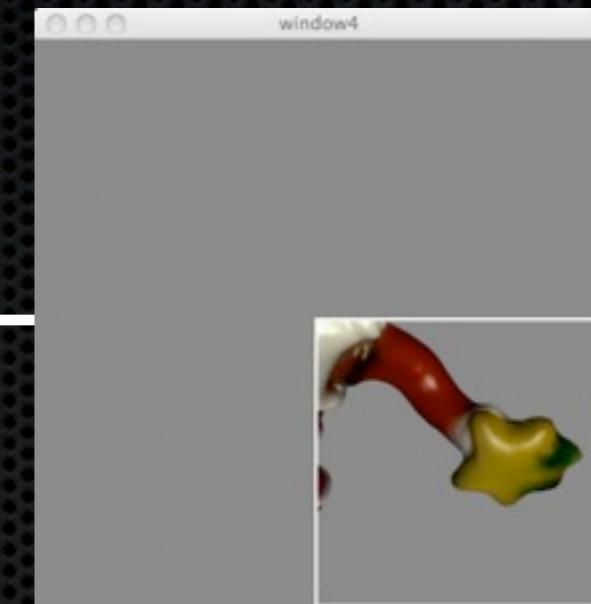
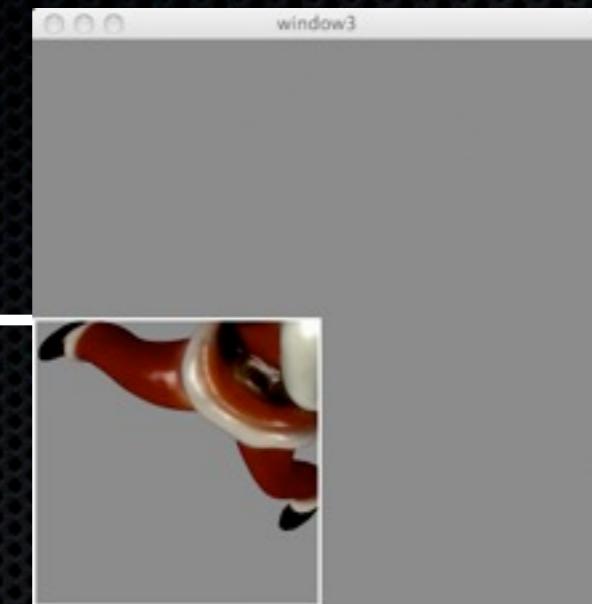
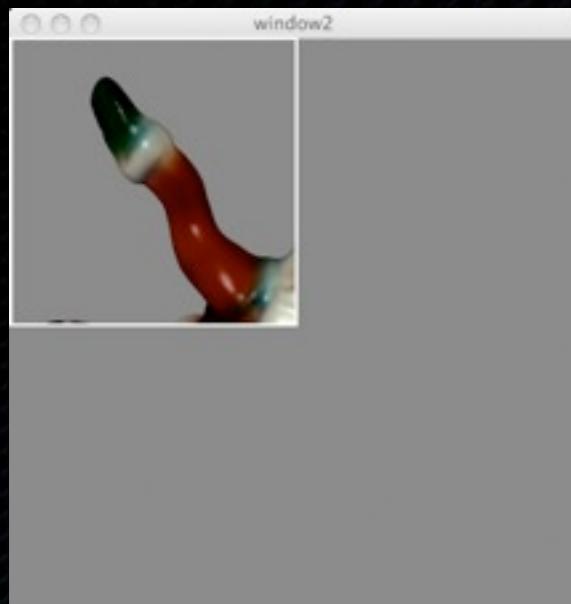
University of Zurich  
Department of Informatics



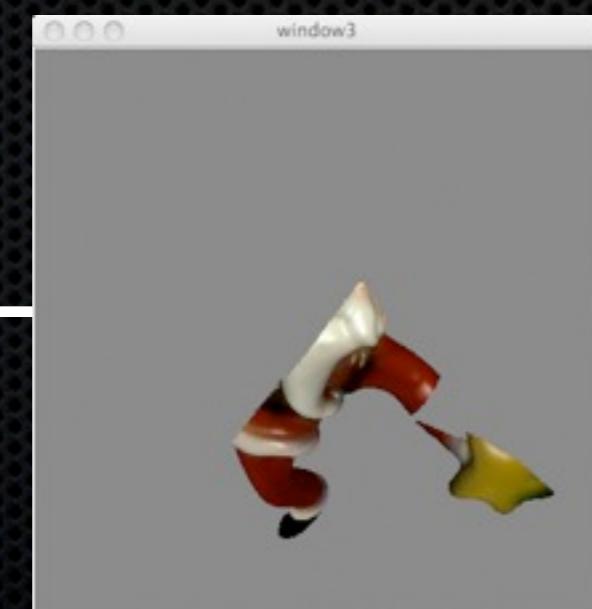
# Content

- Motivation
- Methods
- Experimental analysis
- Conclusion

# Motivation - Focus of Study



Sort-first (2D)



Sort-last (DB)

# Existing Approaches

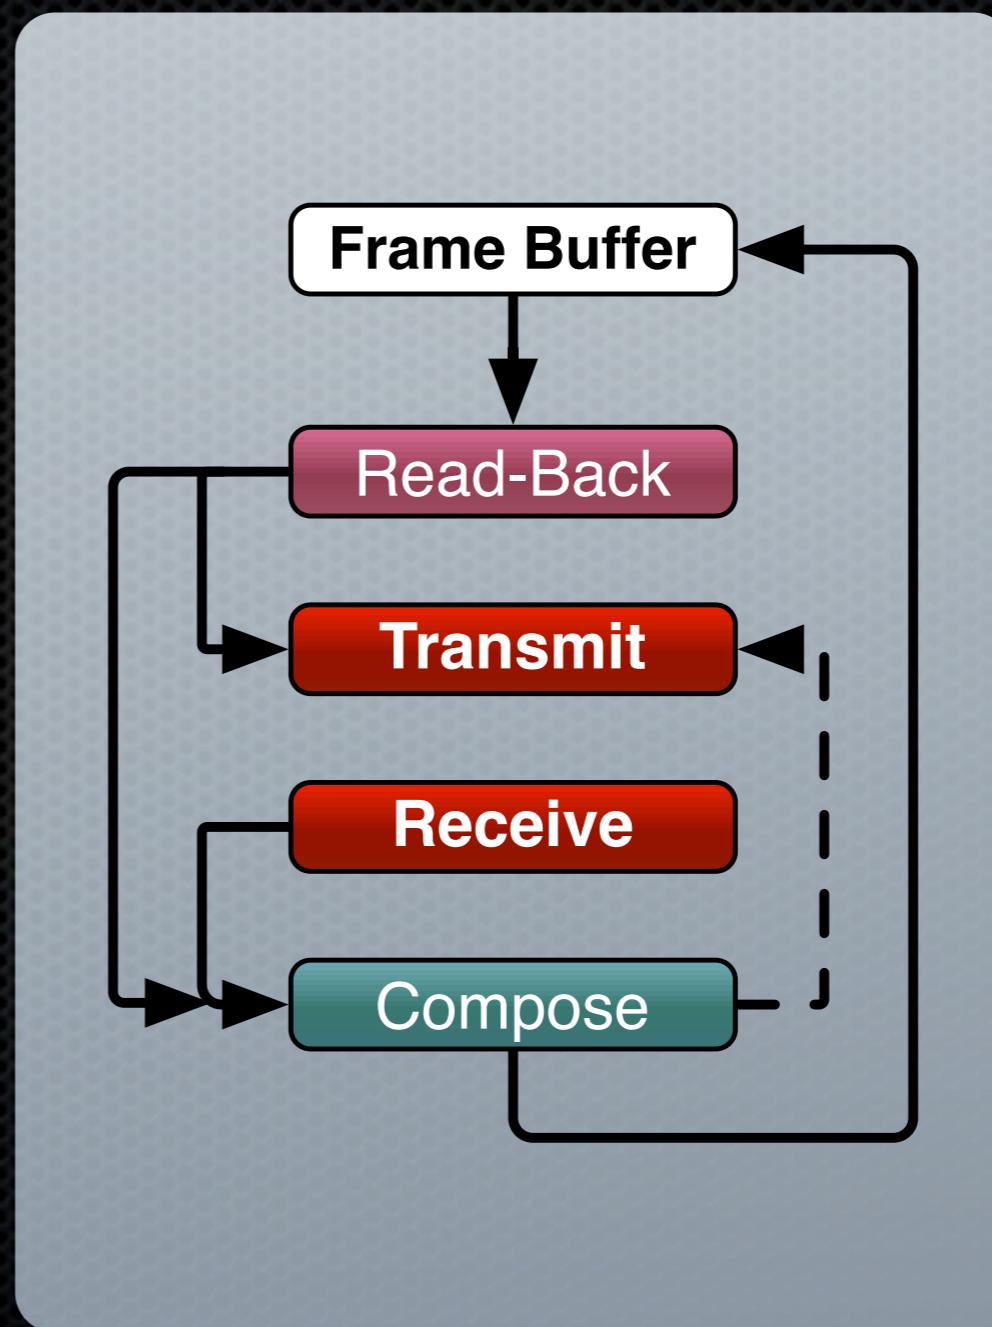
- Read-back / compositing improvement
  - ▶ Less data over network
- Data transmission improvement
  - ▶ Color / depth buffers compression

# Motivation - Other Solutions

- Difficult to integrate in to existent frameworks
- Hardware dependent
- Doesn't necessarily improve performance on modern hardware

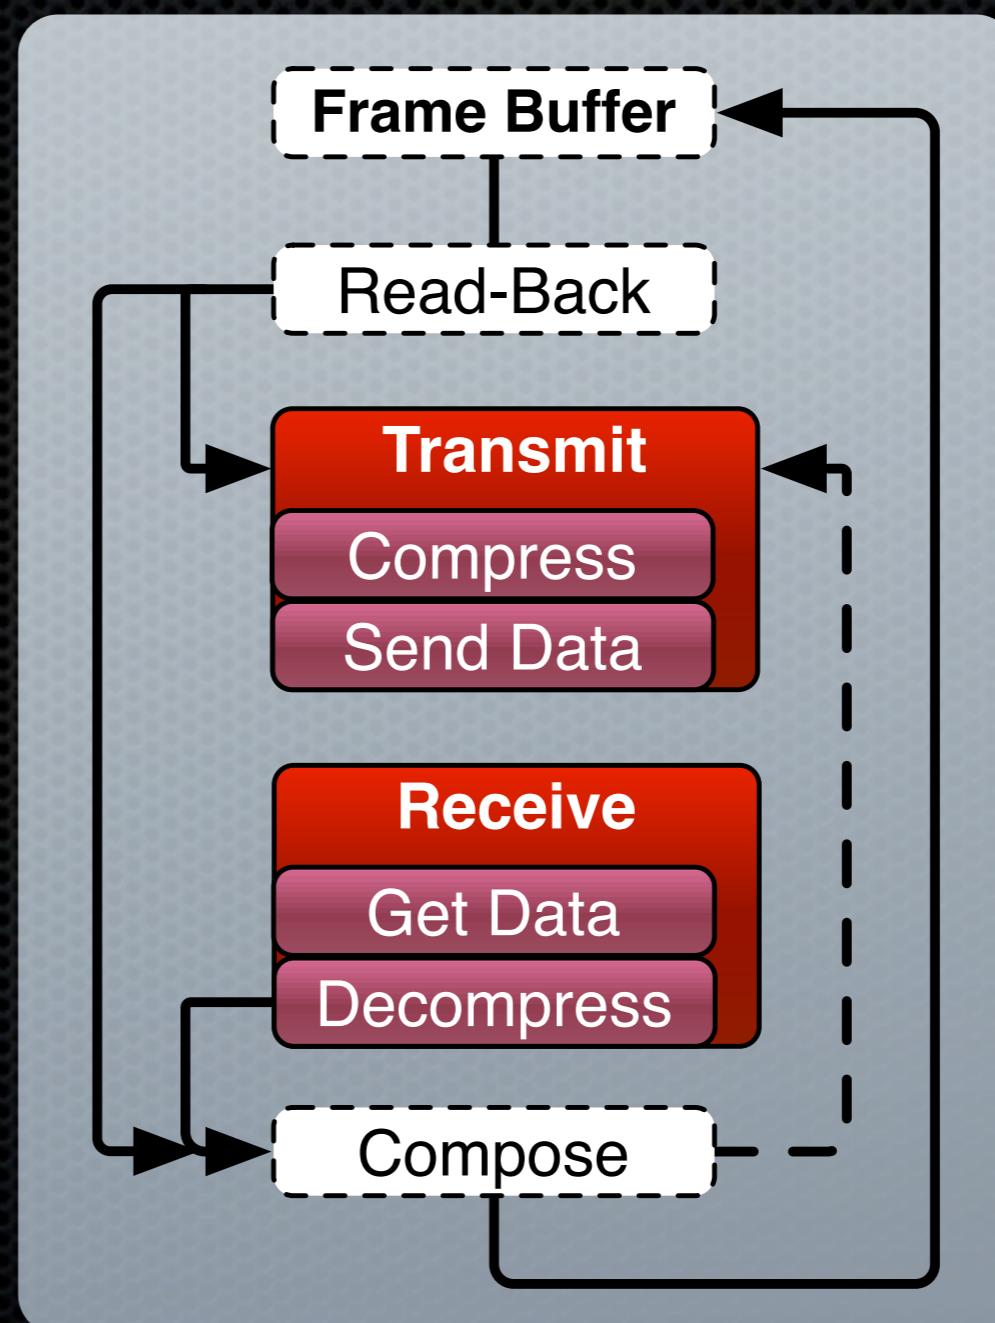
# Methods - Equalizer (EQ)

Typical  
compositing loop  
in parallel rendering

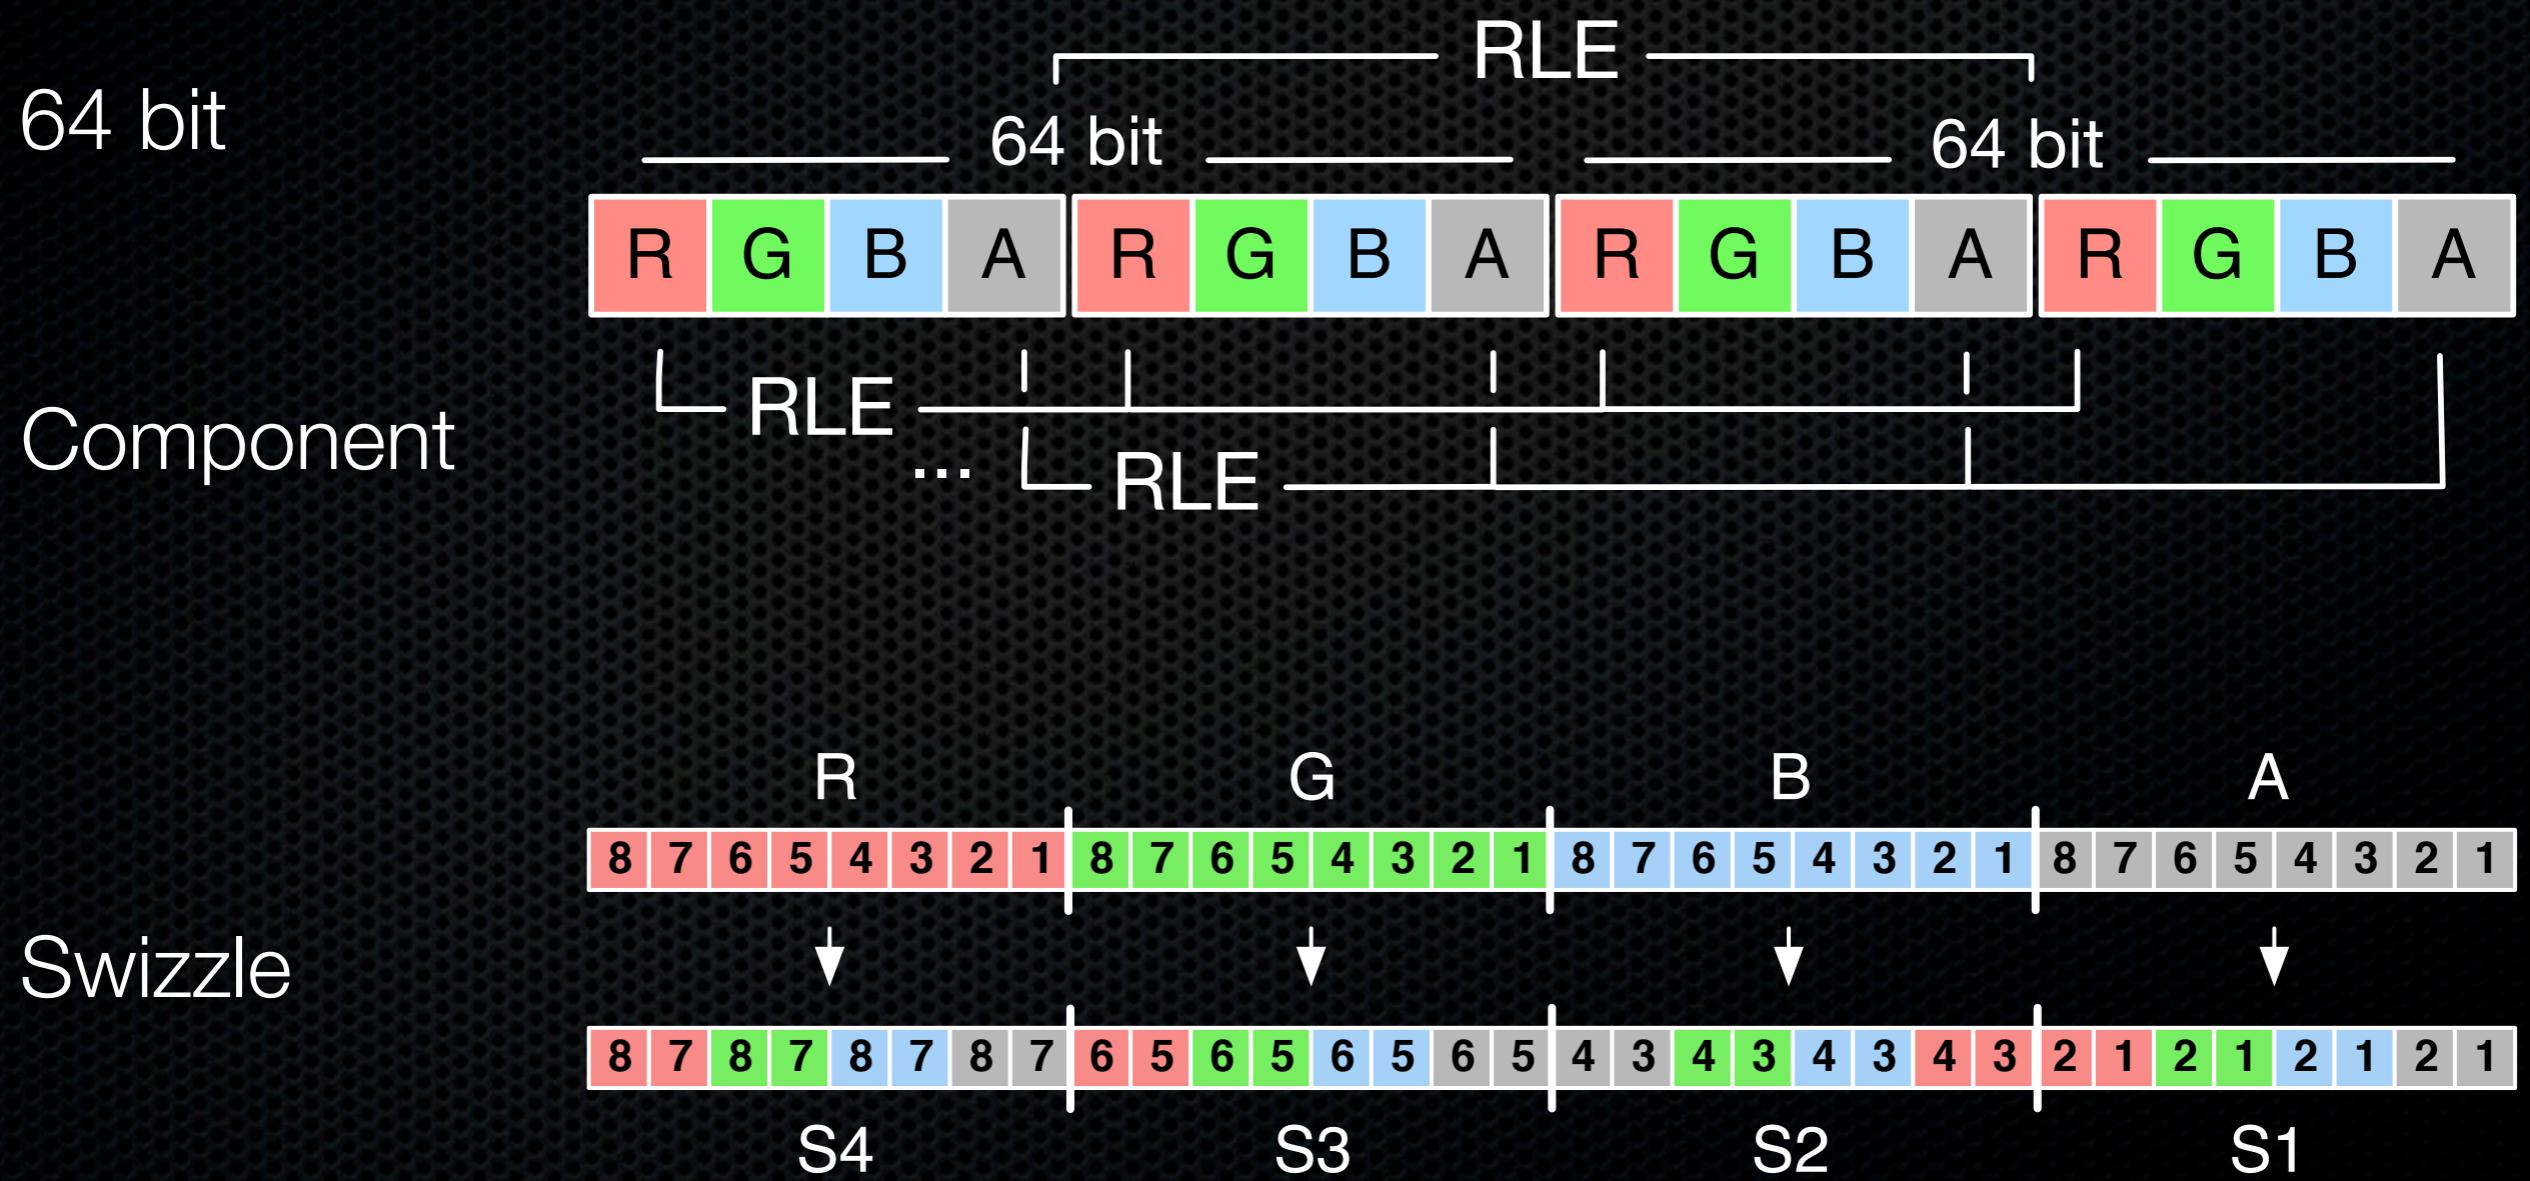


# Methods - Equalizer (EQ)

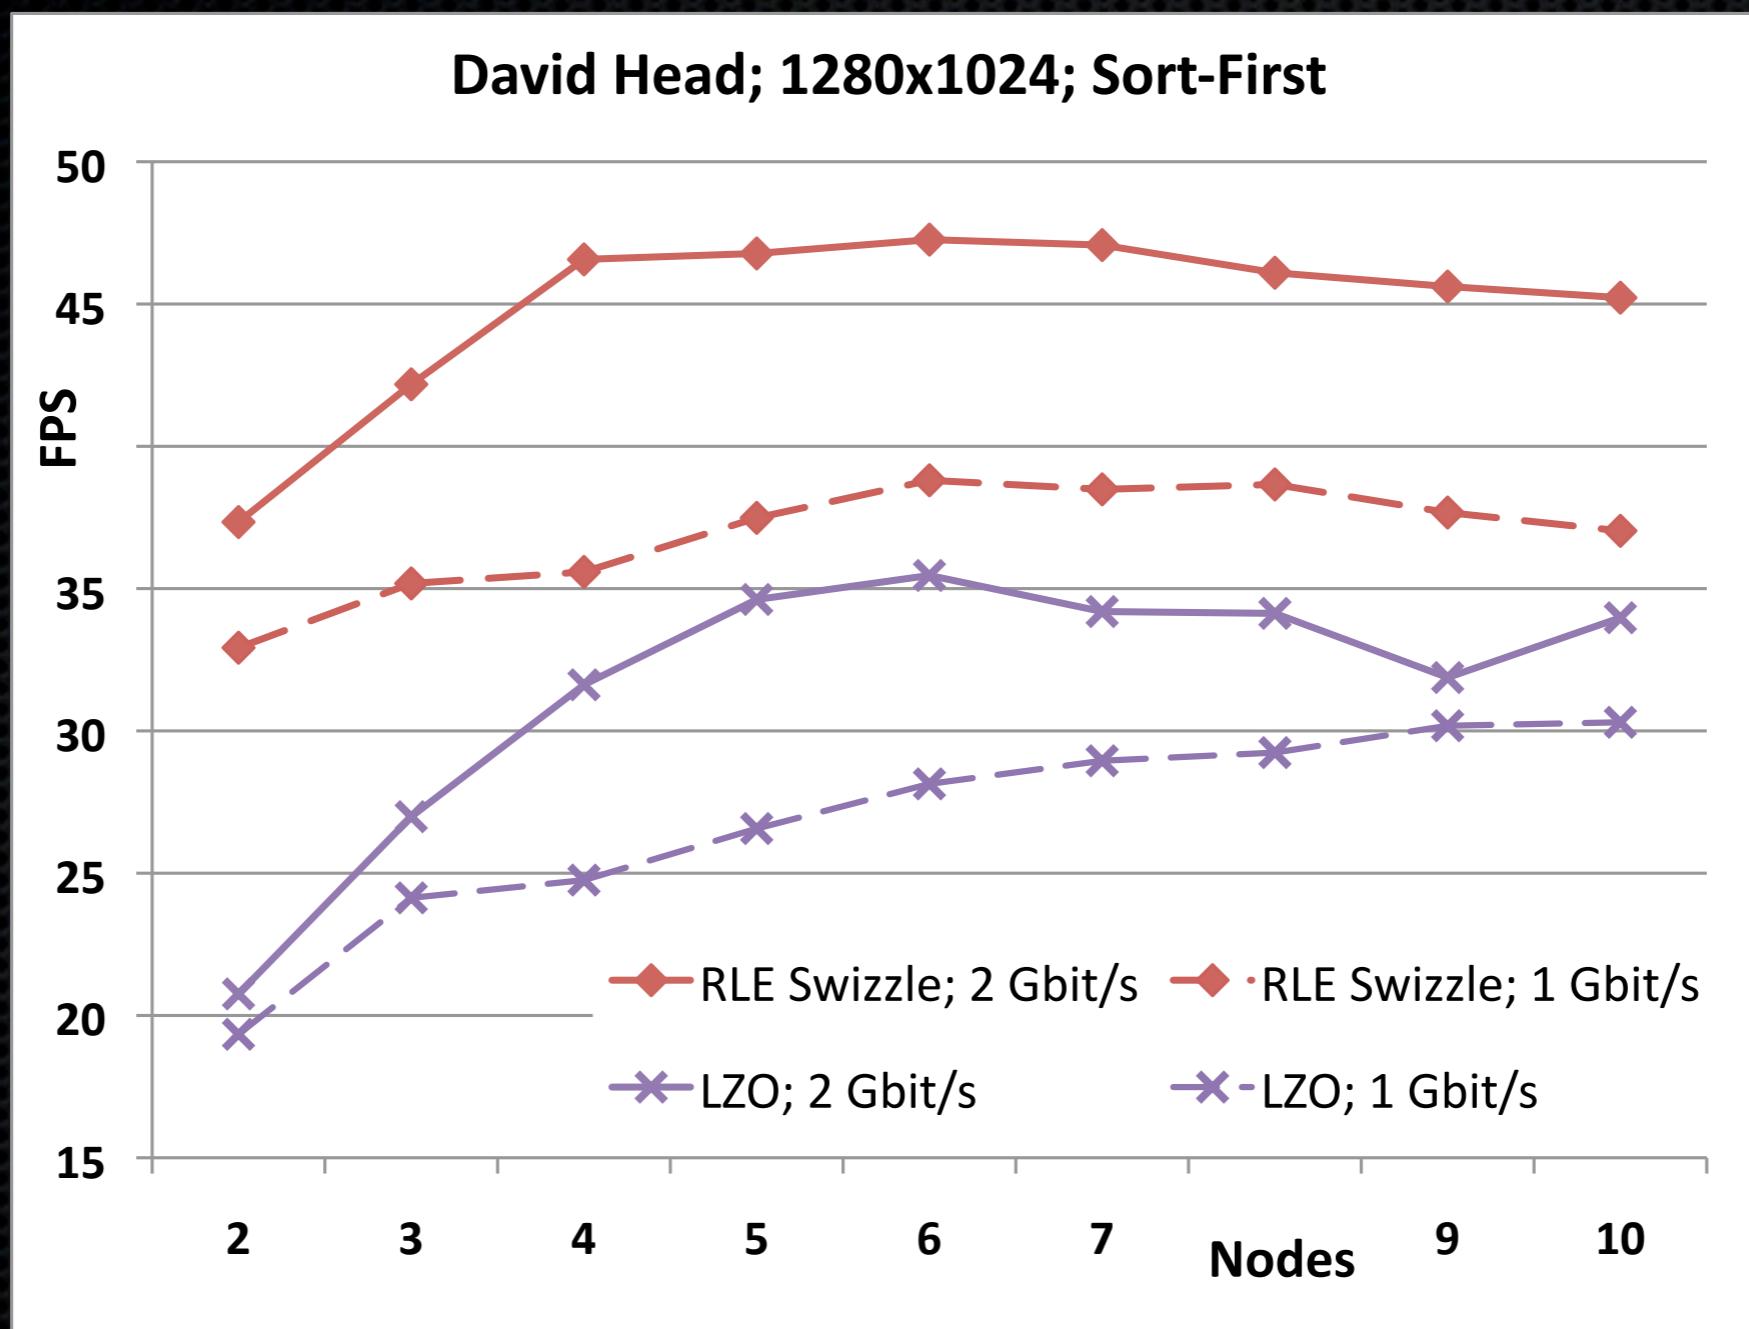
Typical  
compositing loop  
in parallel rendering



# Methods - RLEs



# Methods - RLE vs LZO



# Methods - YUV12

R1	G1	B1	A1	R2	G2	B2	A2
R3	G3	B3	A3	R4	G4	B4	A4



RGB to YUV

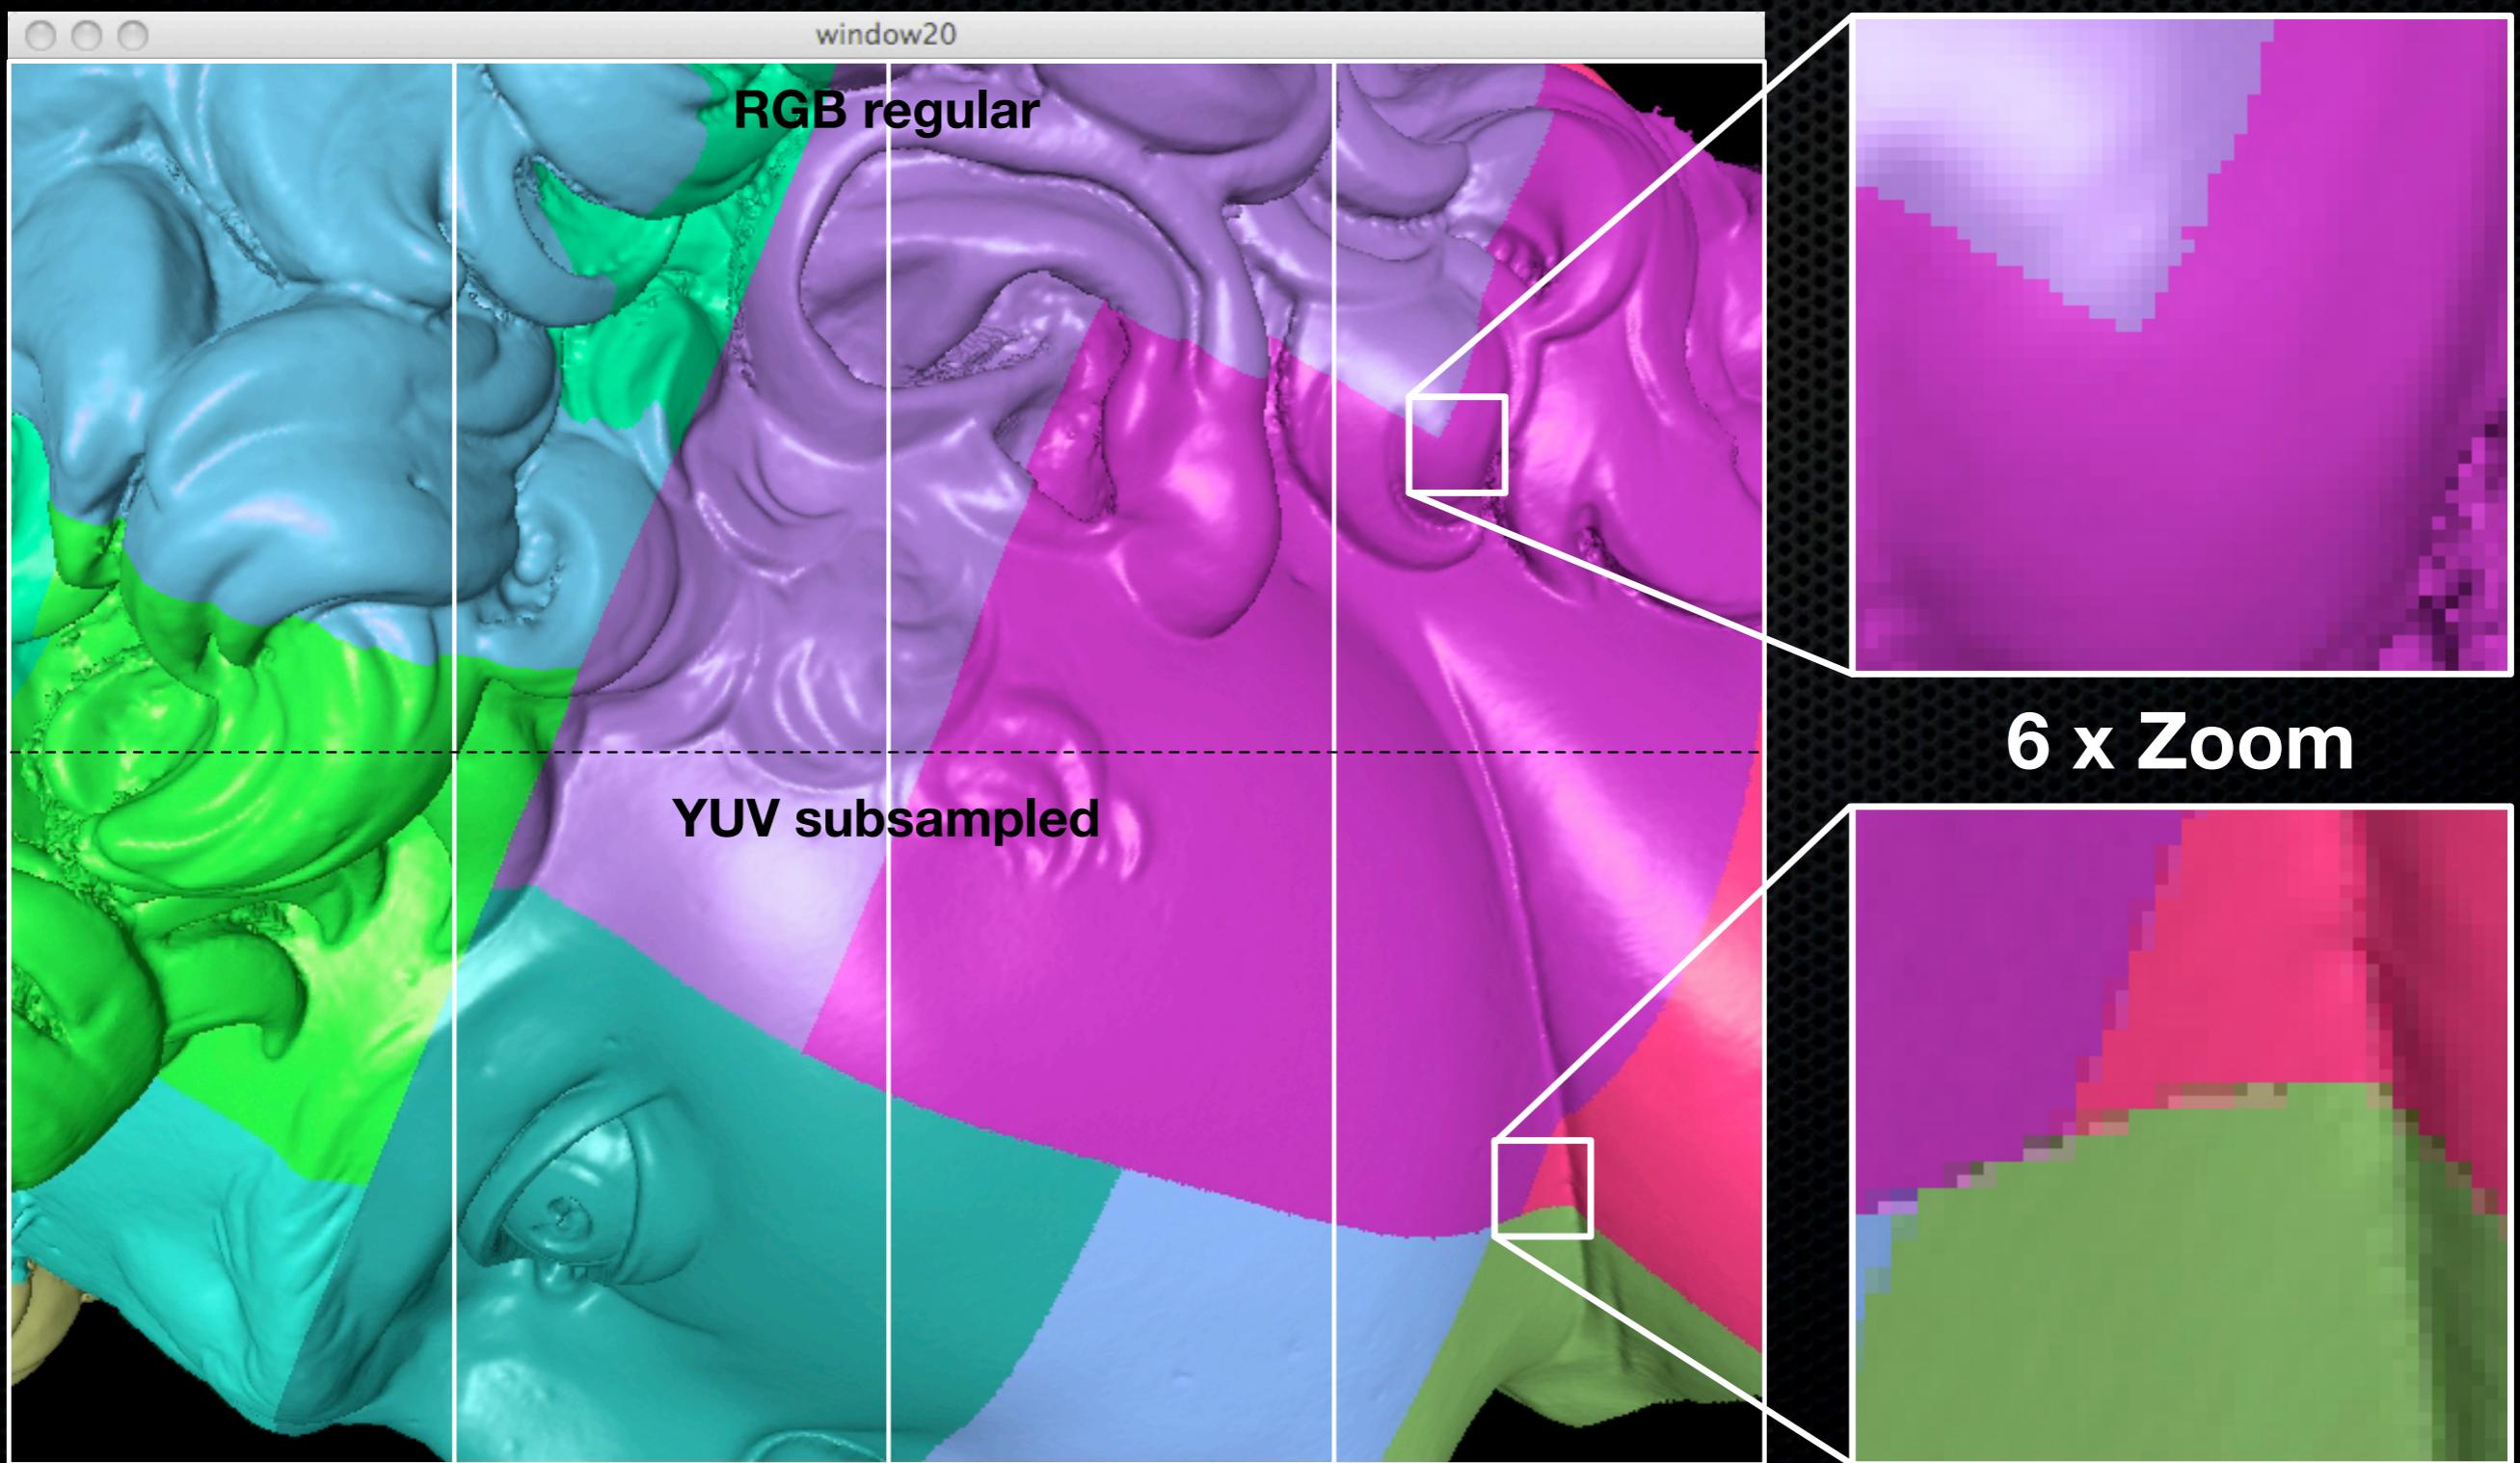
Y1	U1	V1	A1	Y2	U2	V2	A2
Y3	U3	V3	A3	Y4	U4	V4	A4



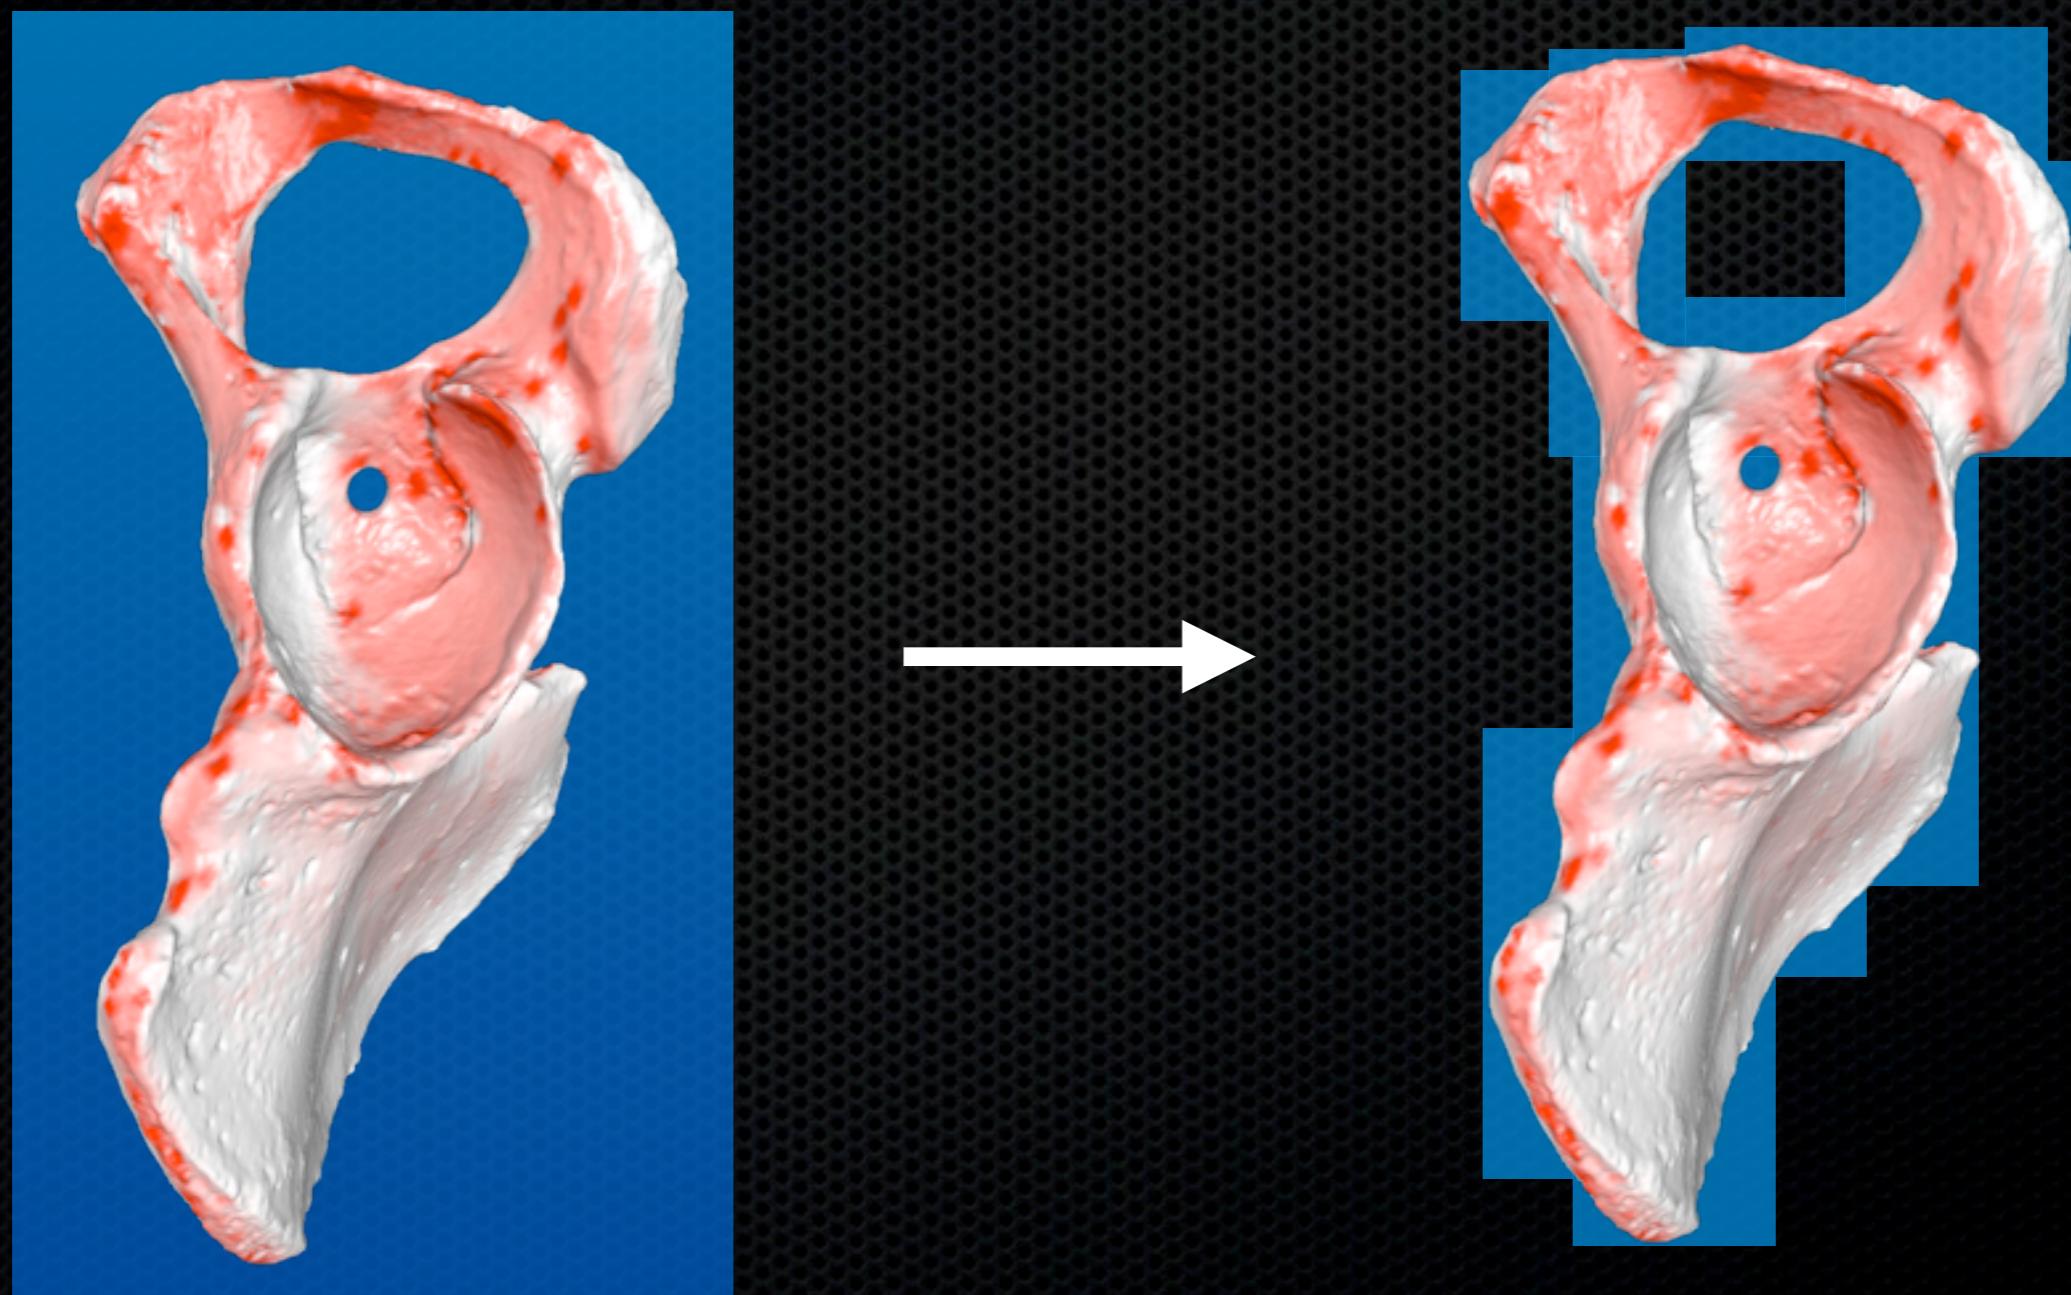
Average

Y1	Y2	Un	An
Y3	Y4	Vn	An

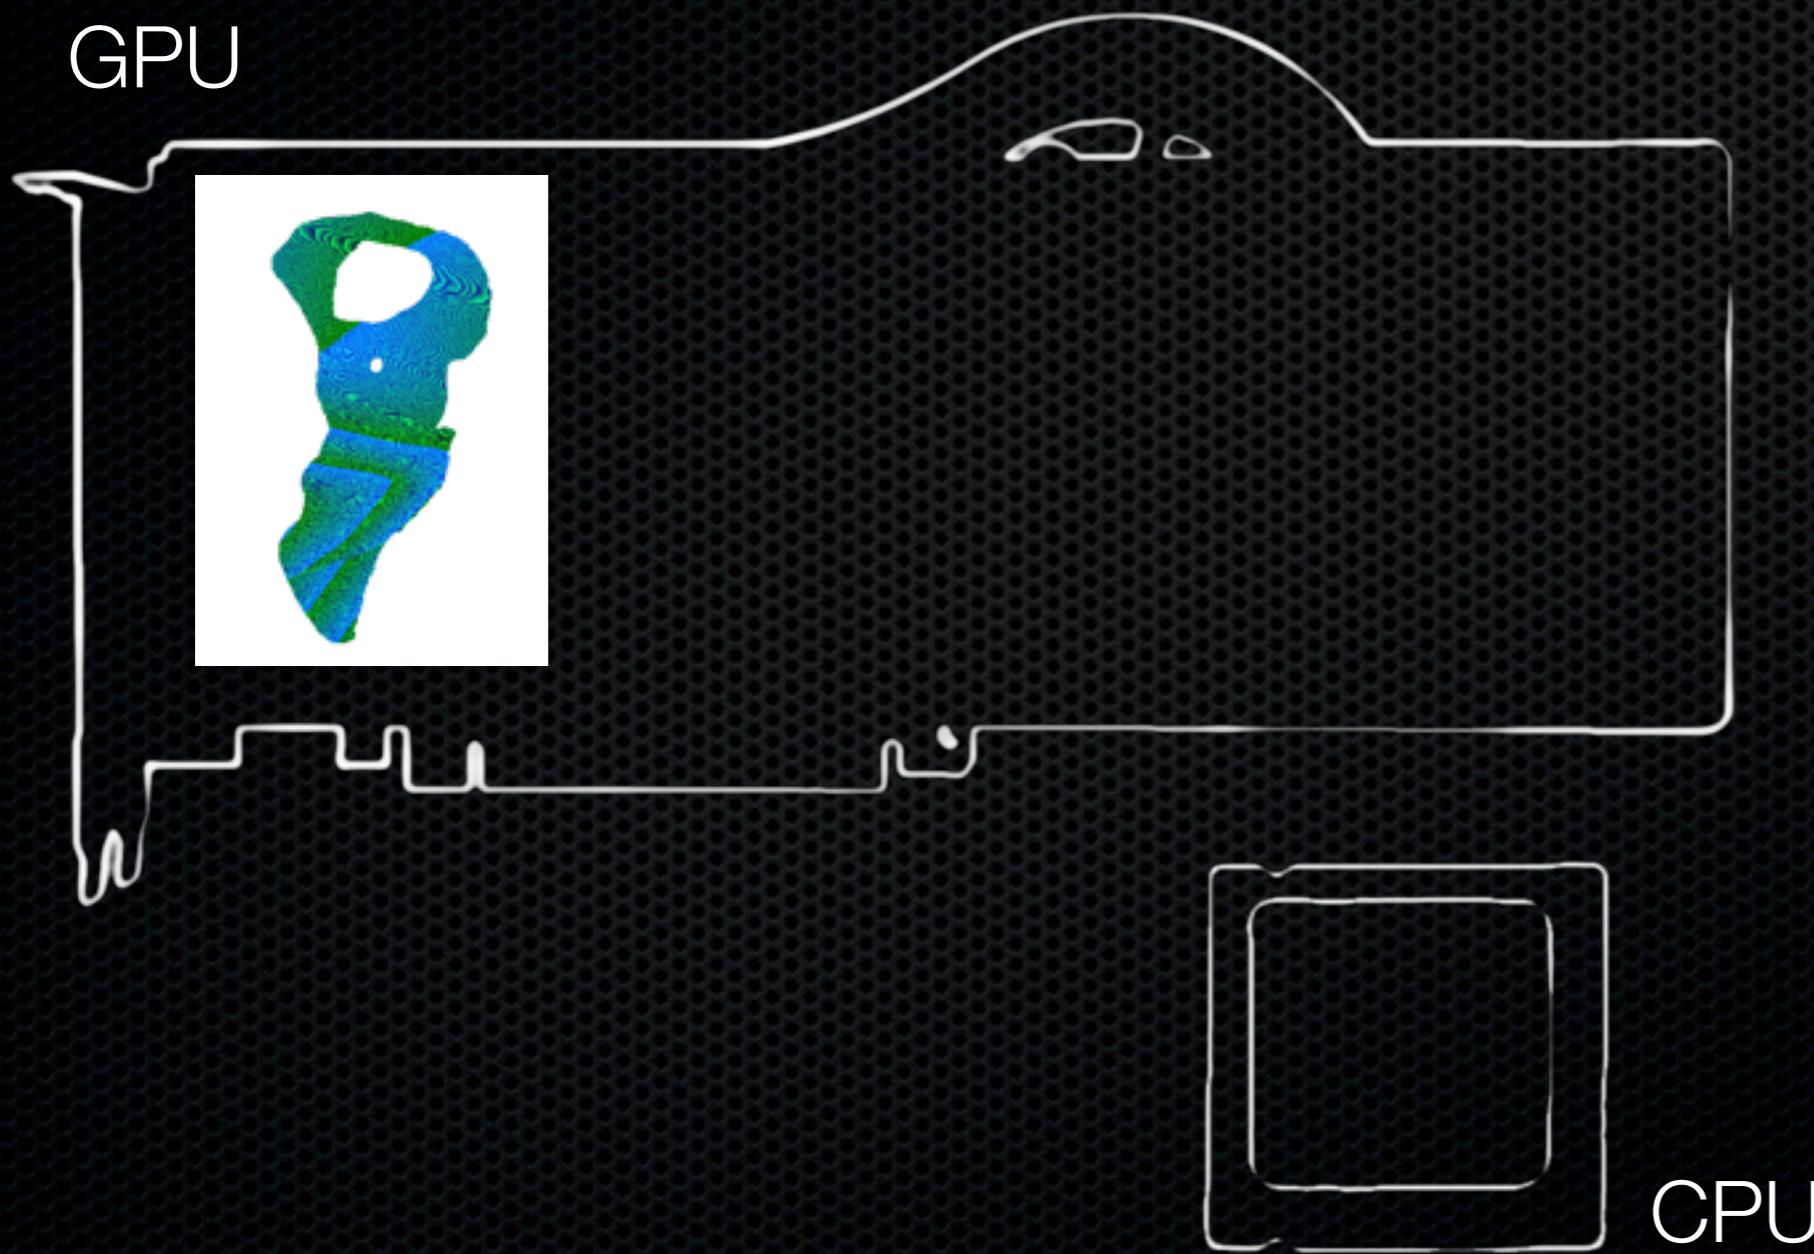
# Methods - YUV12 (Example)



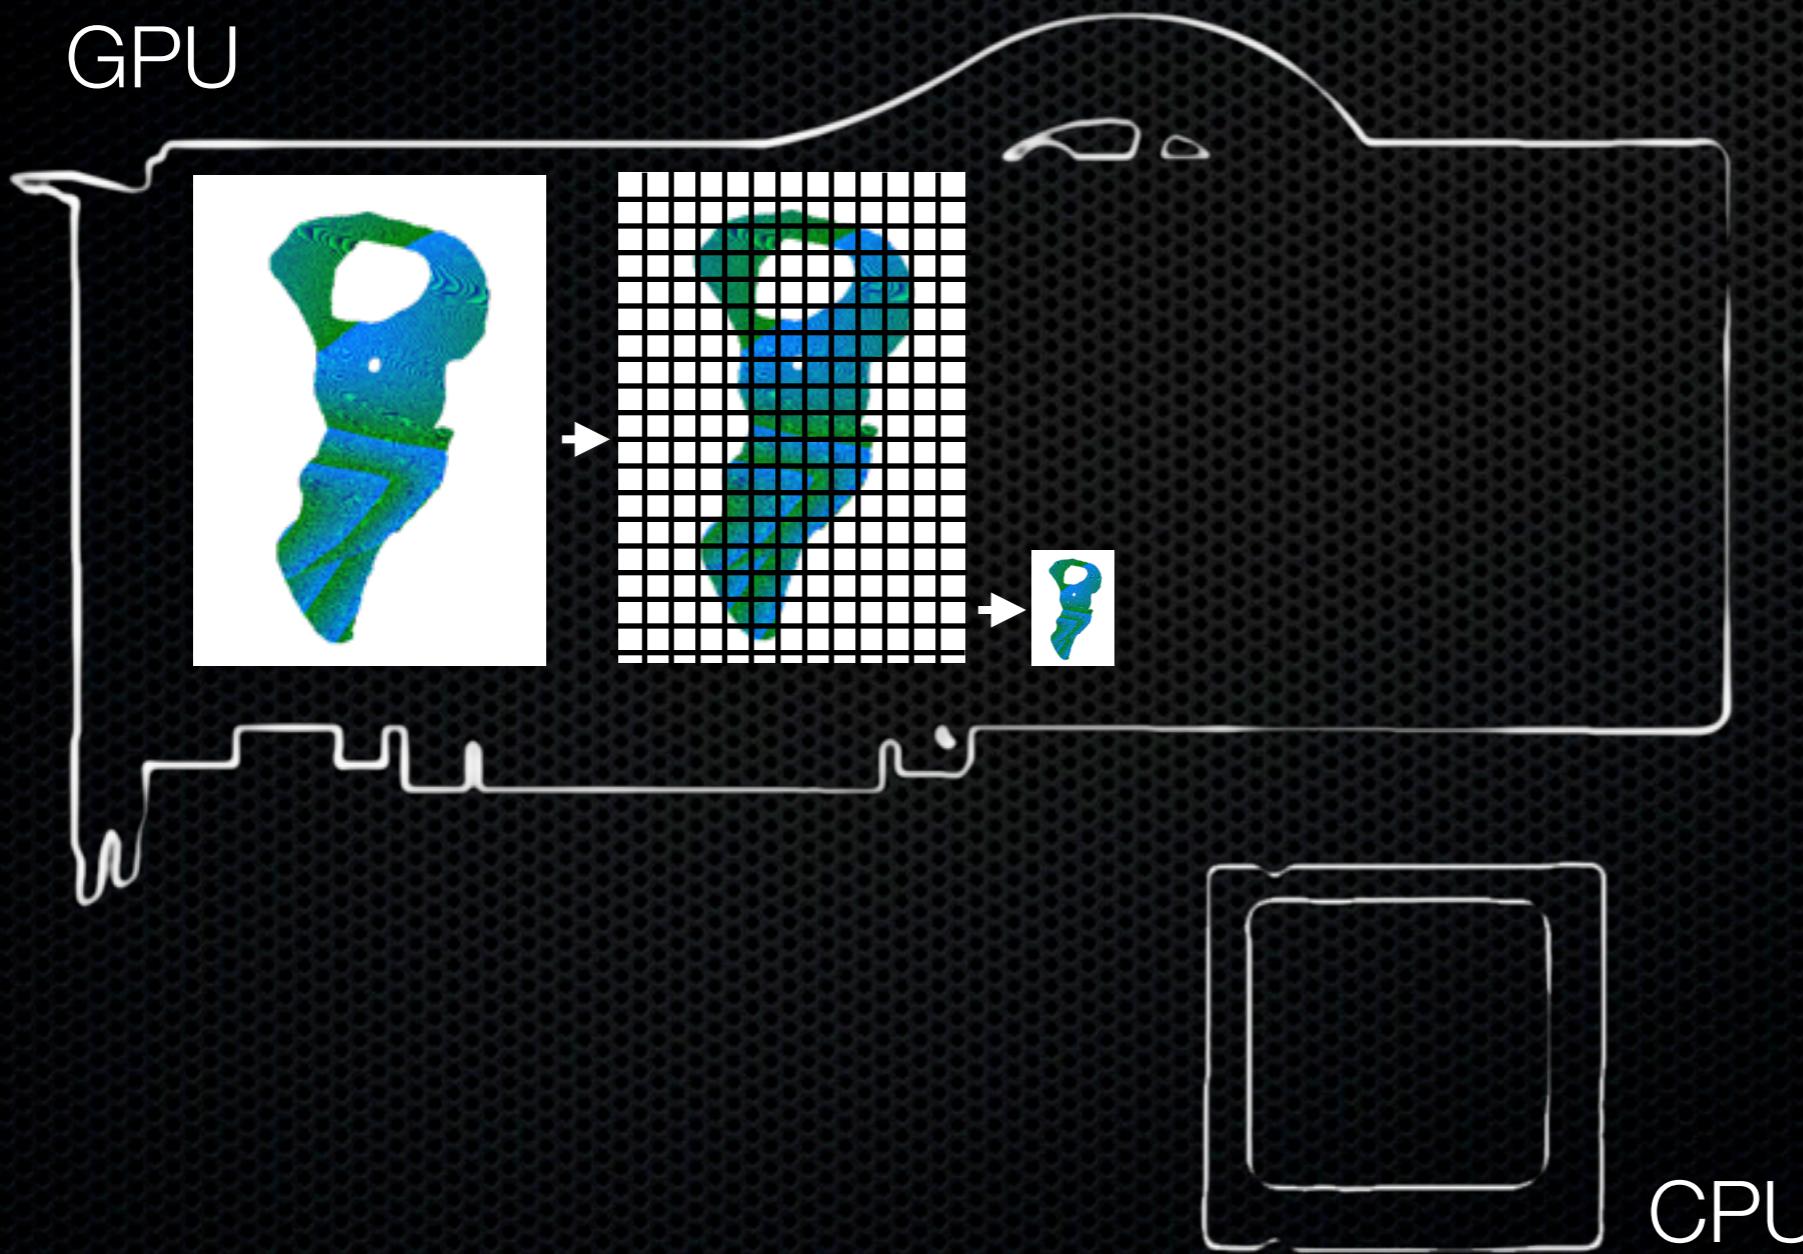
# Methods - ROI (Idea)



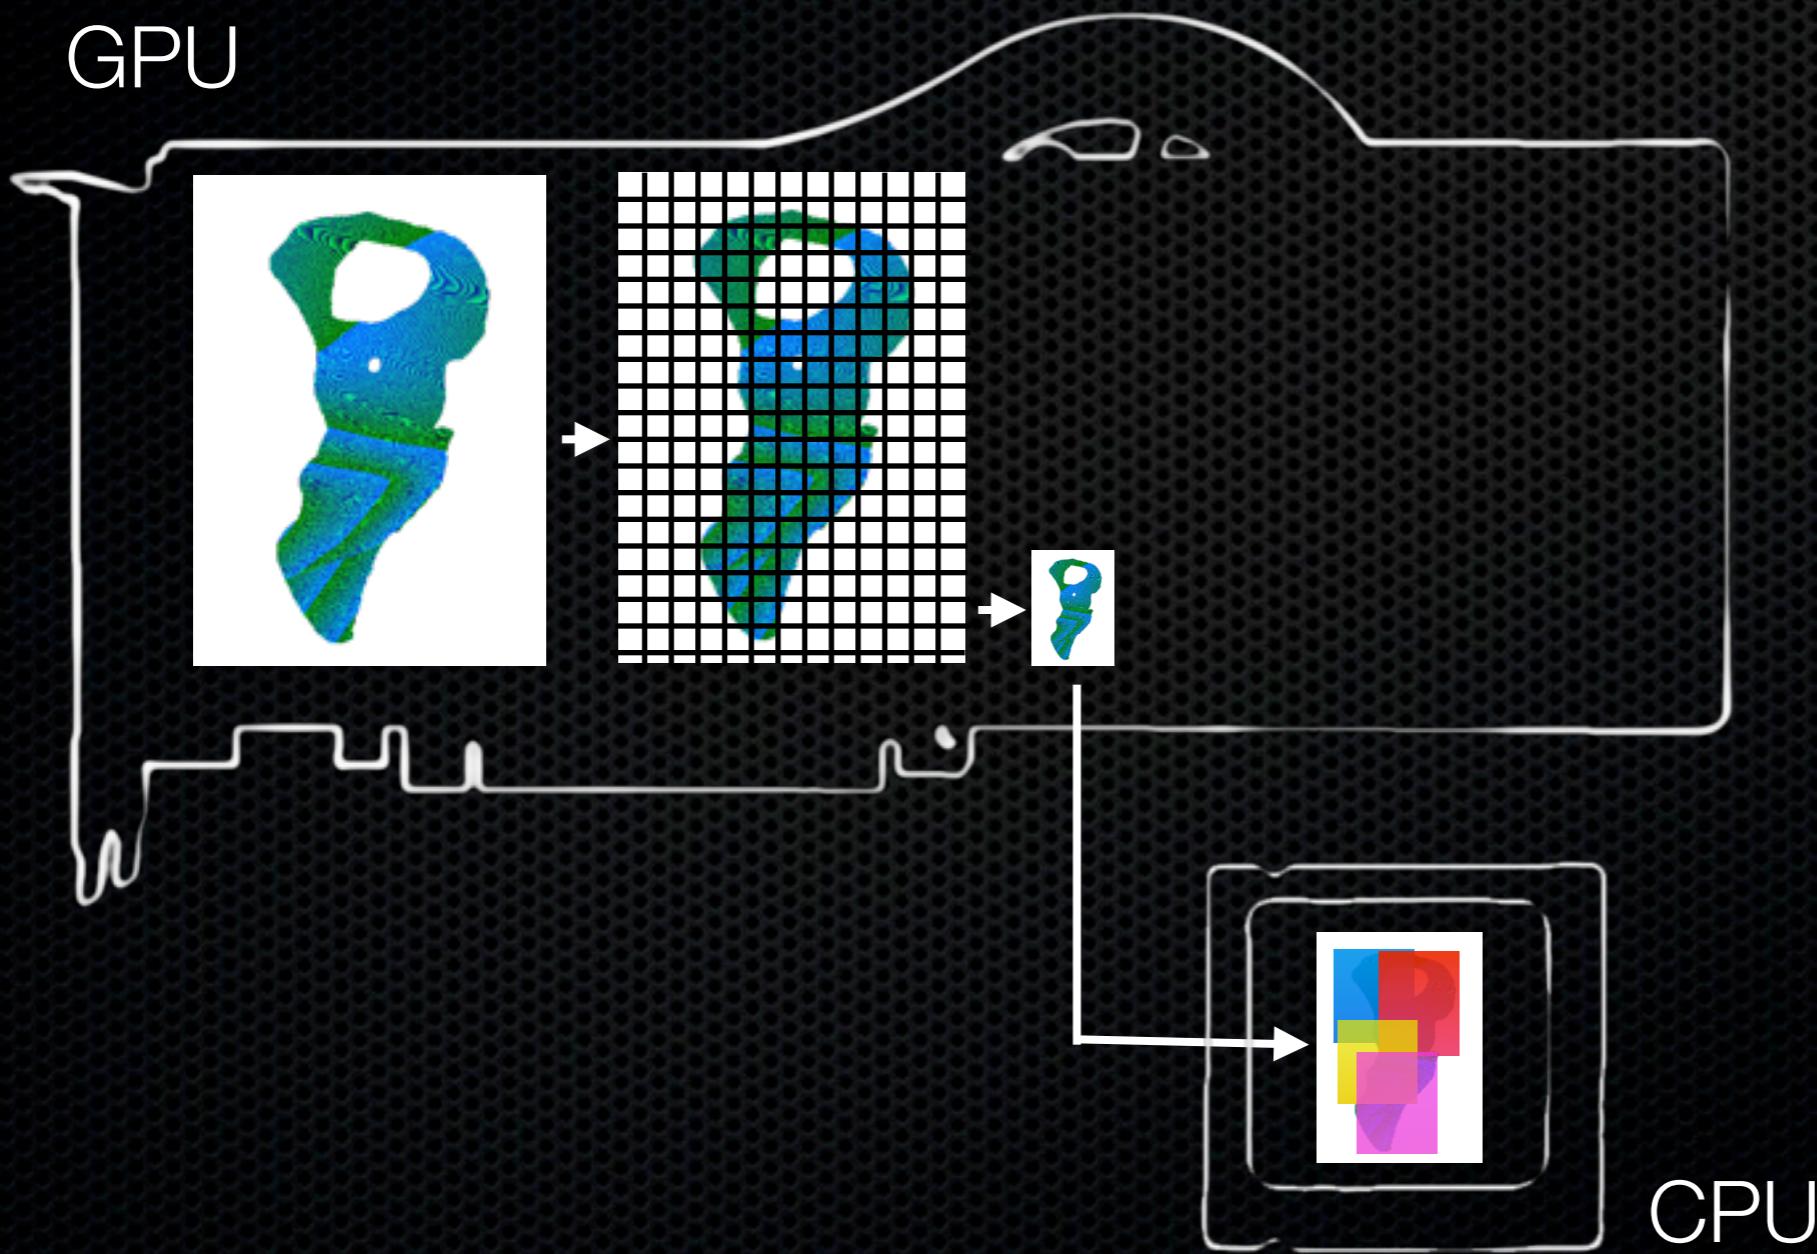
# Methods - ROI



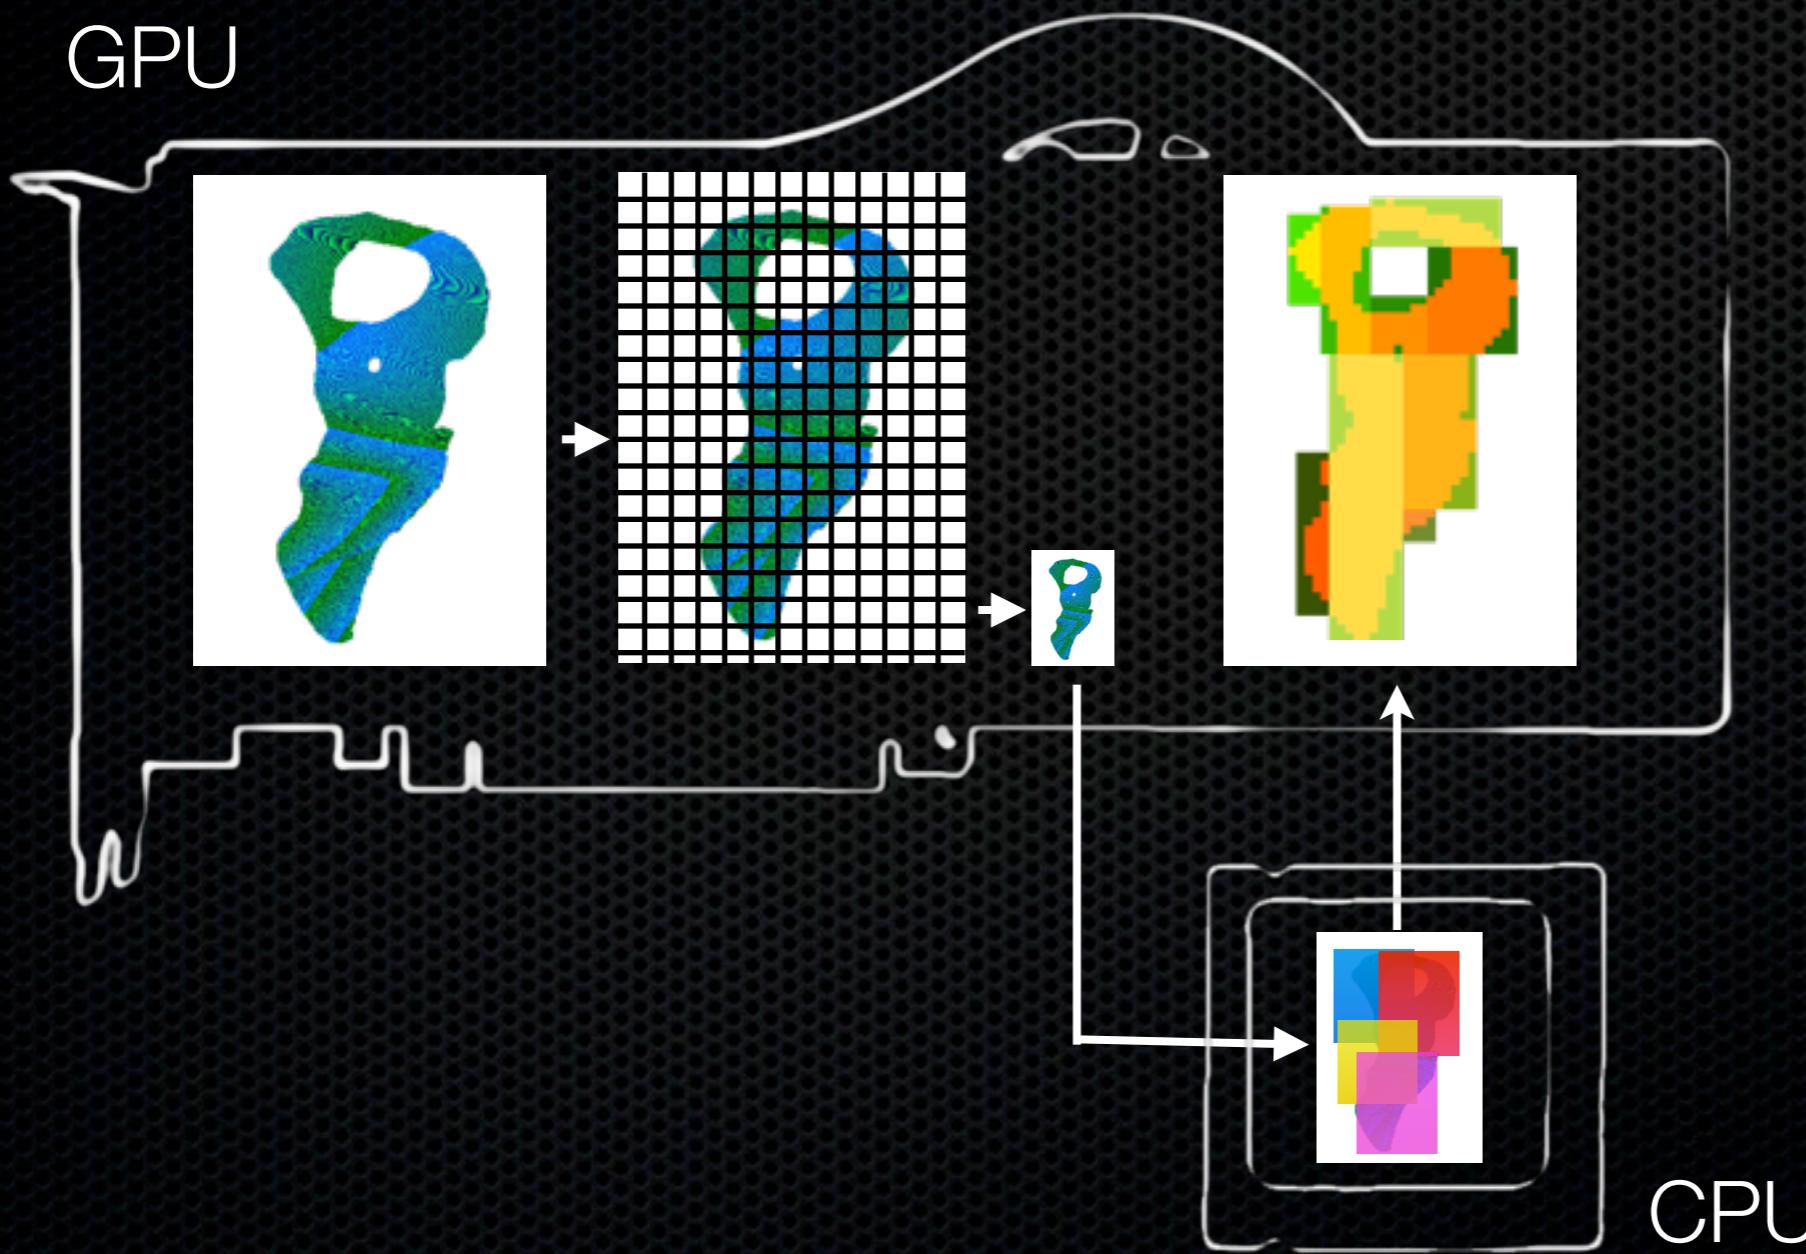
# Methods - ROI



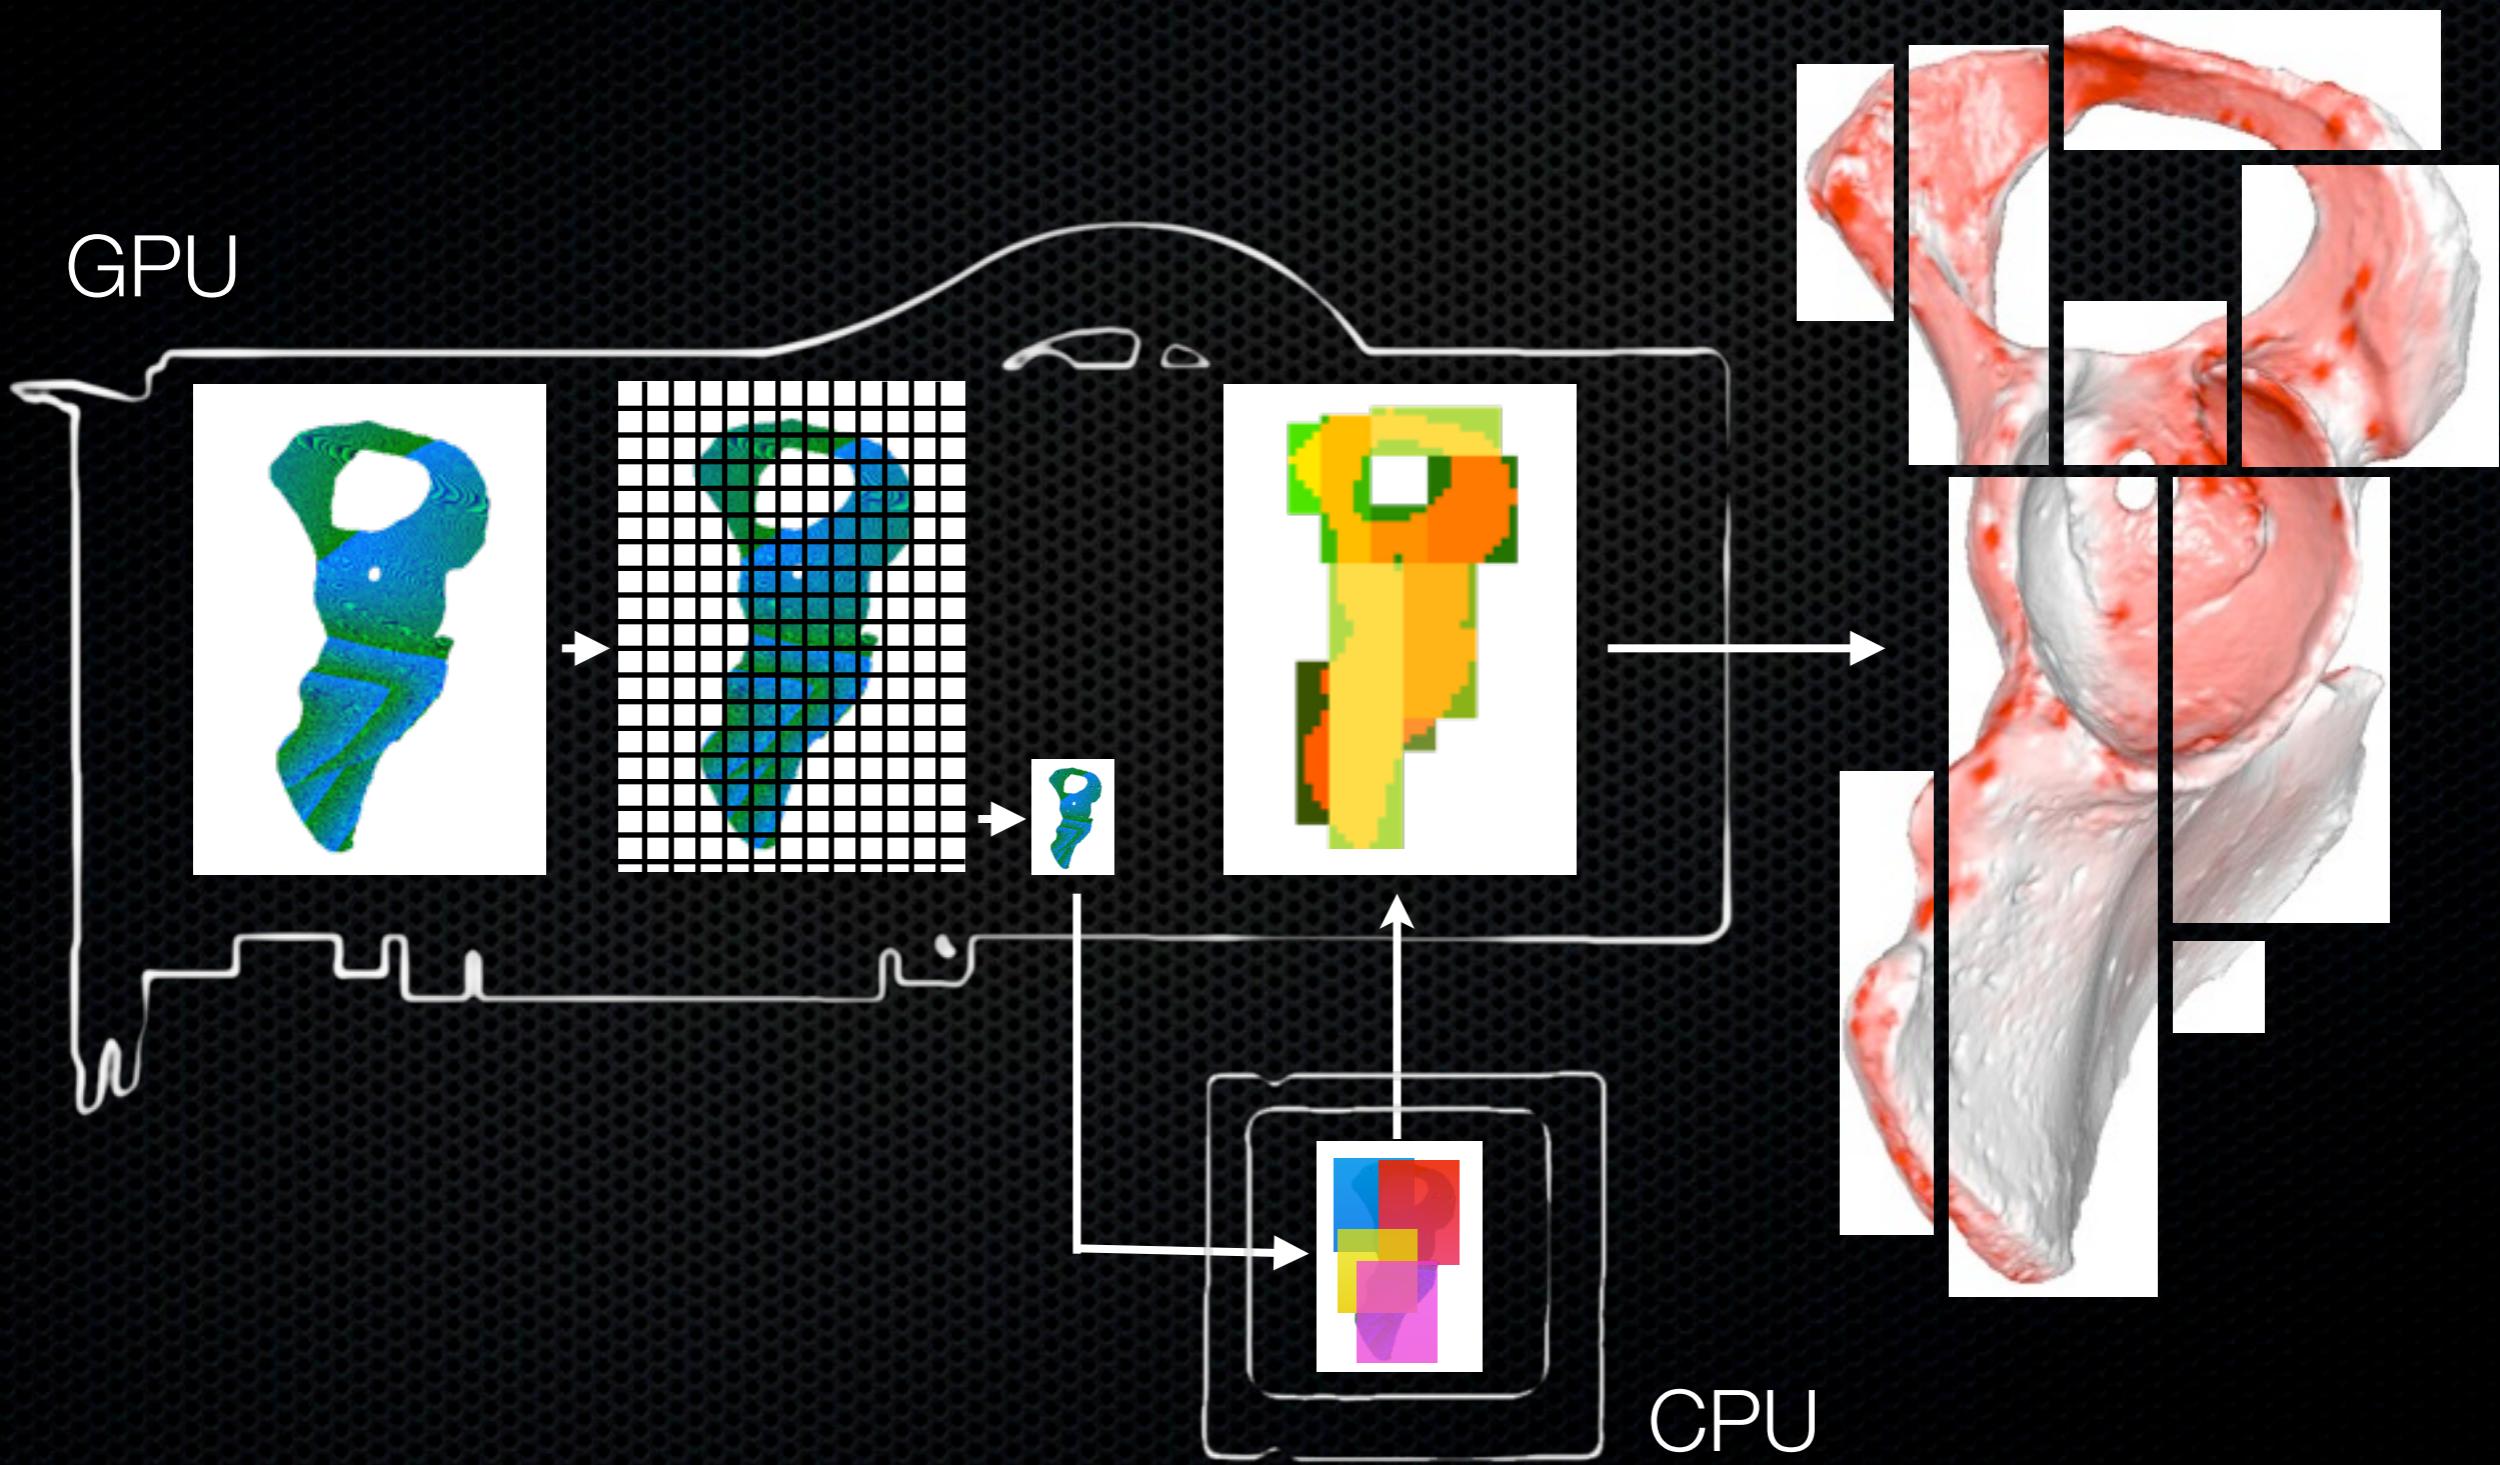
# Methods - ROI



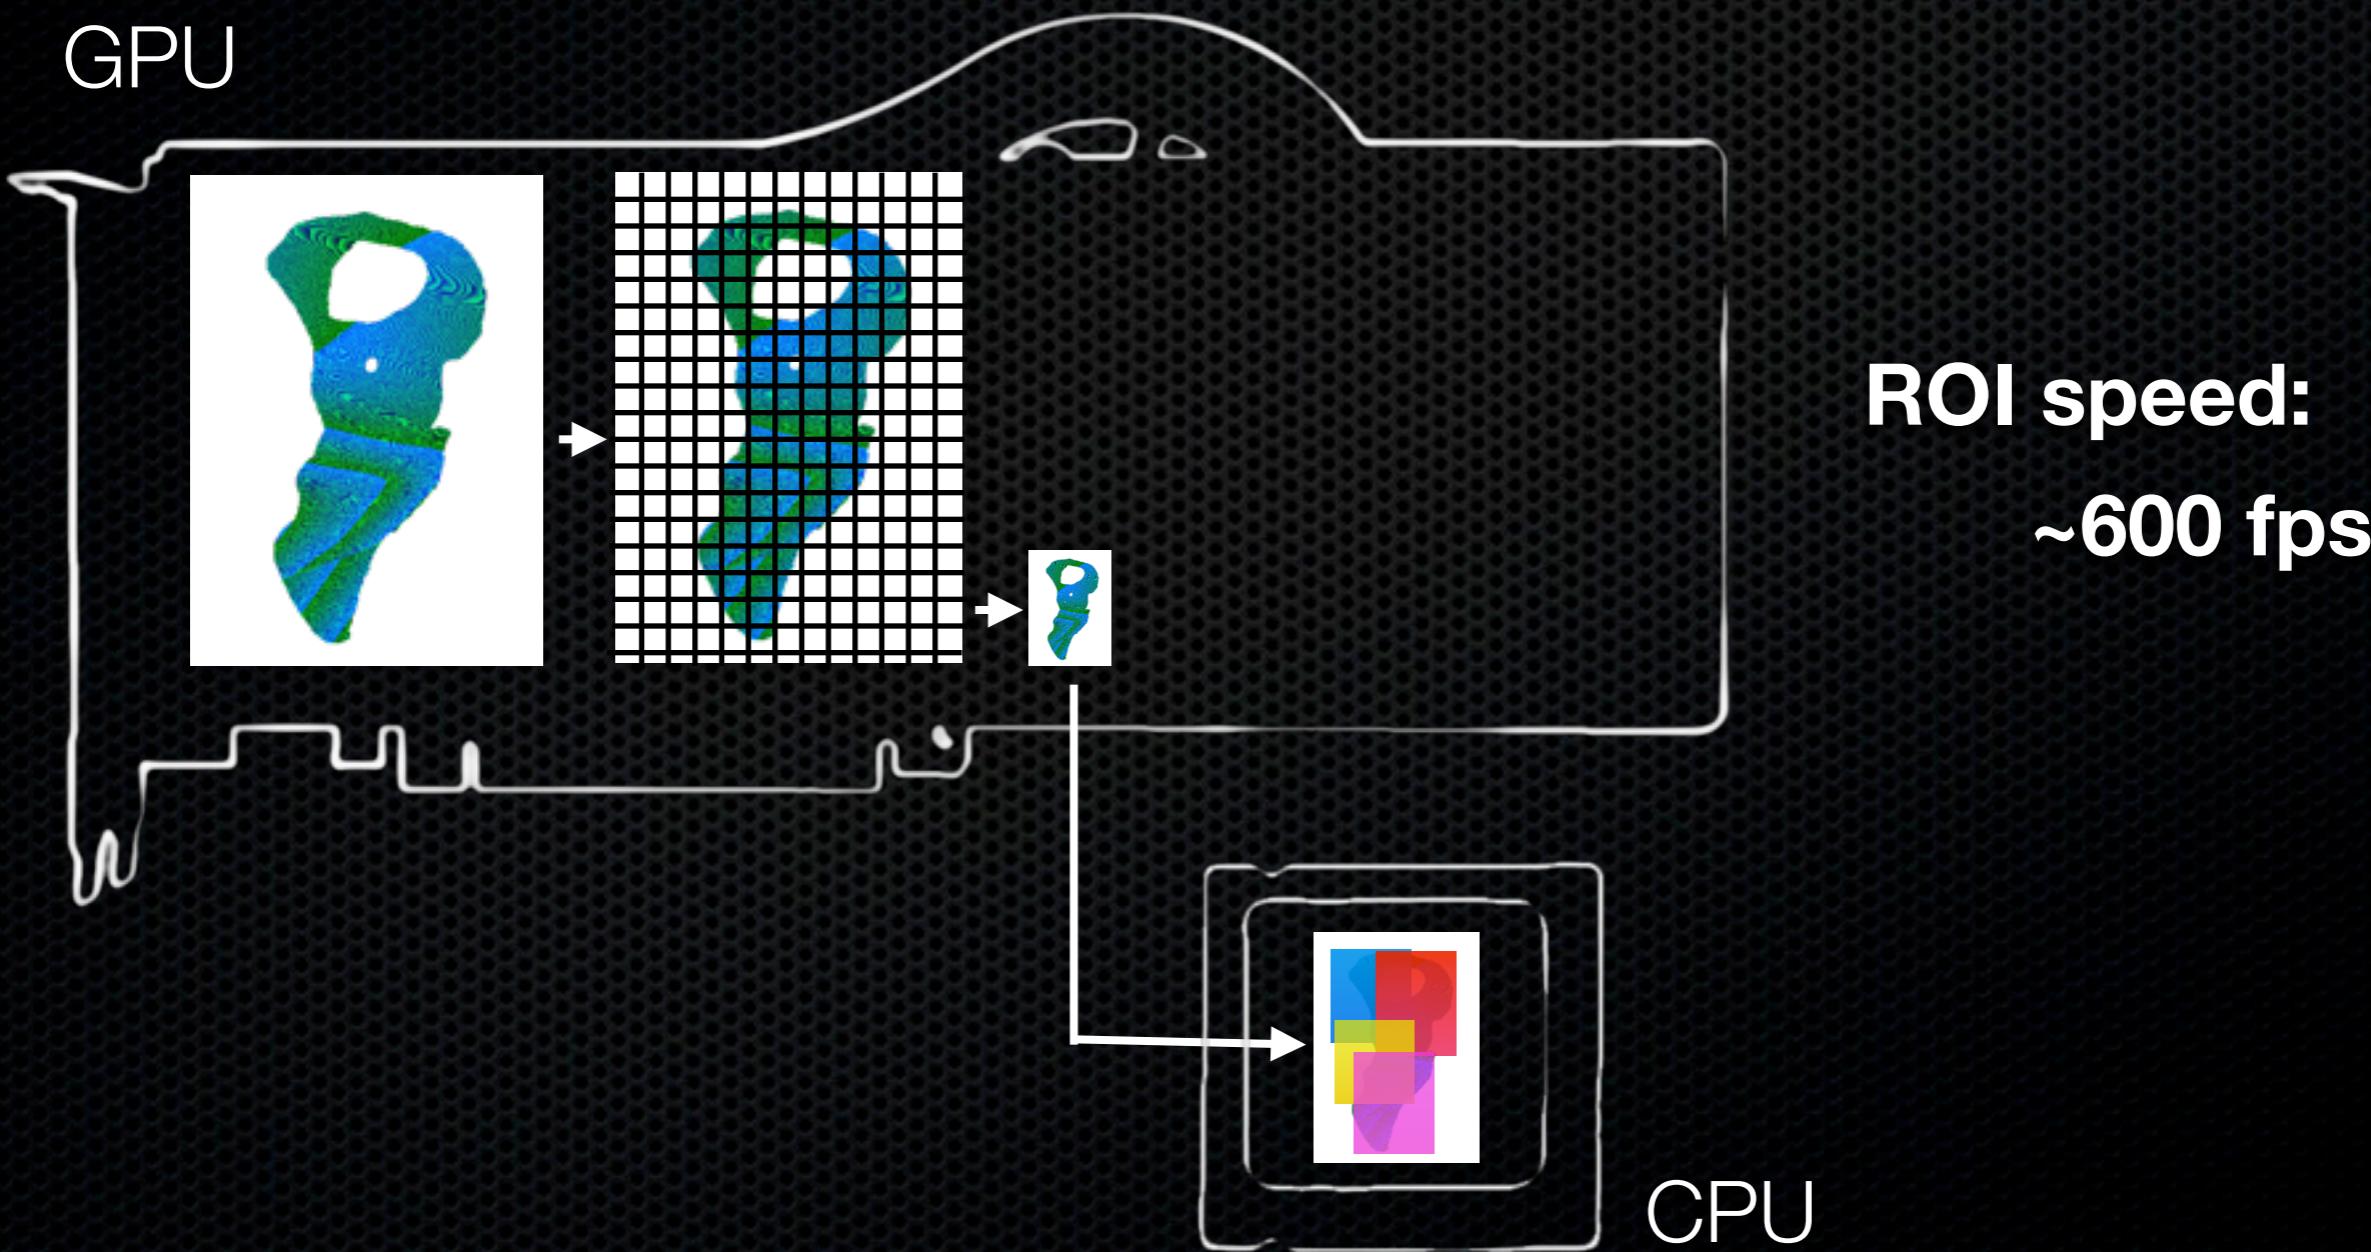
# Methods - ROI



# Methods - ROI



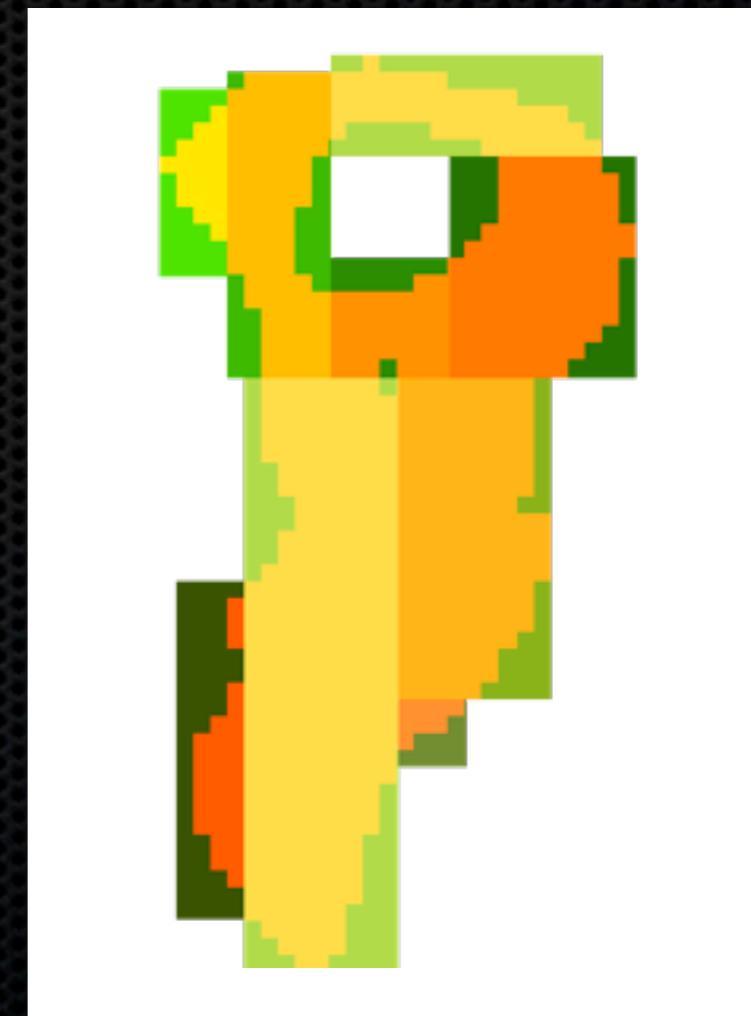
# Methods - ROI



# Methods - ROI

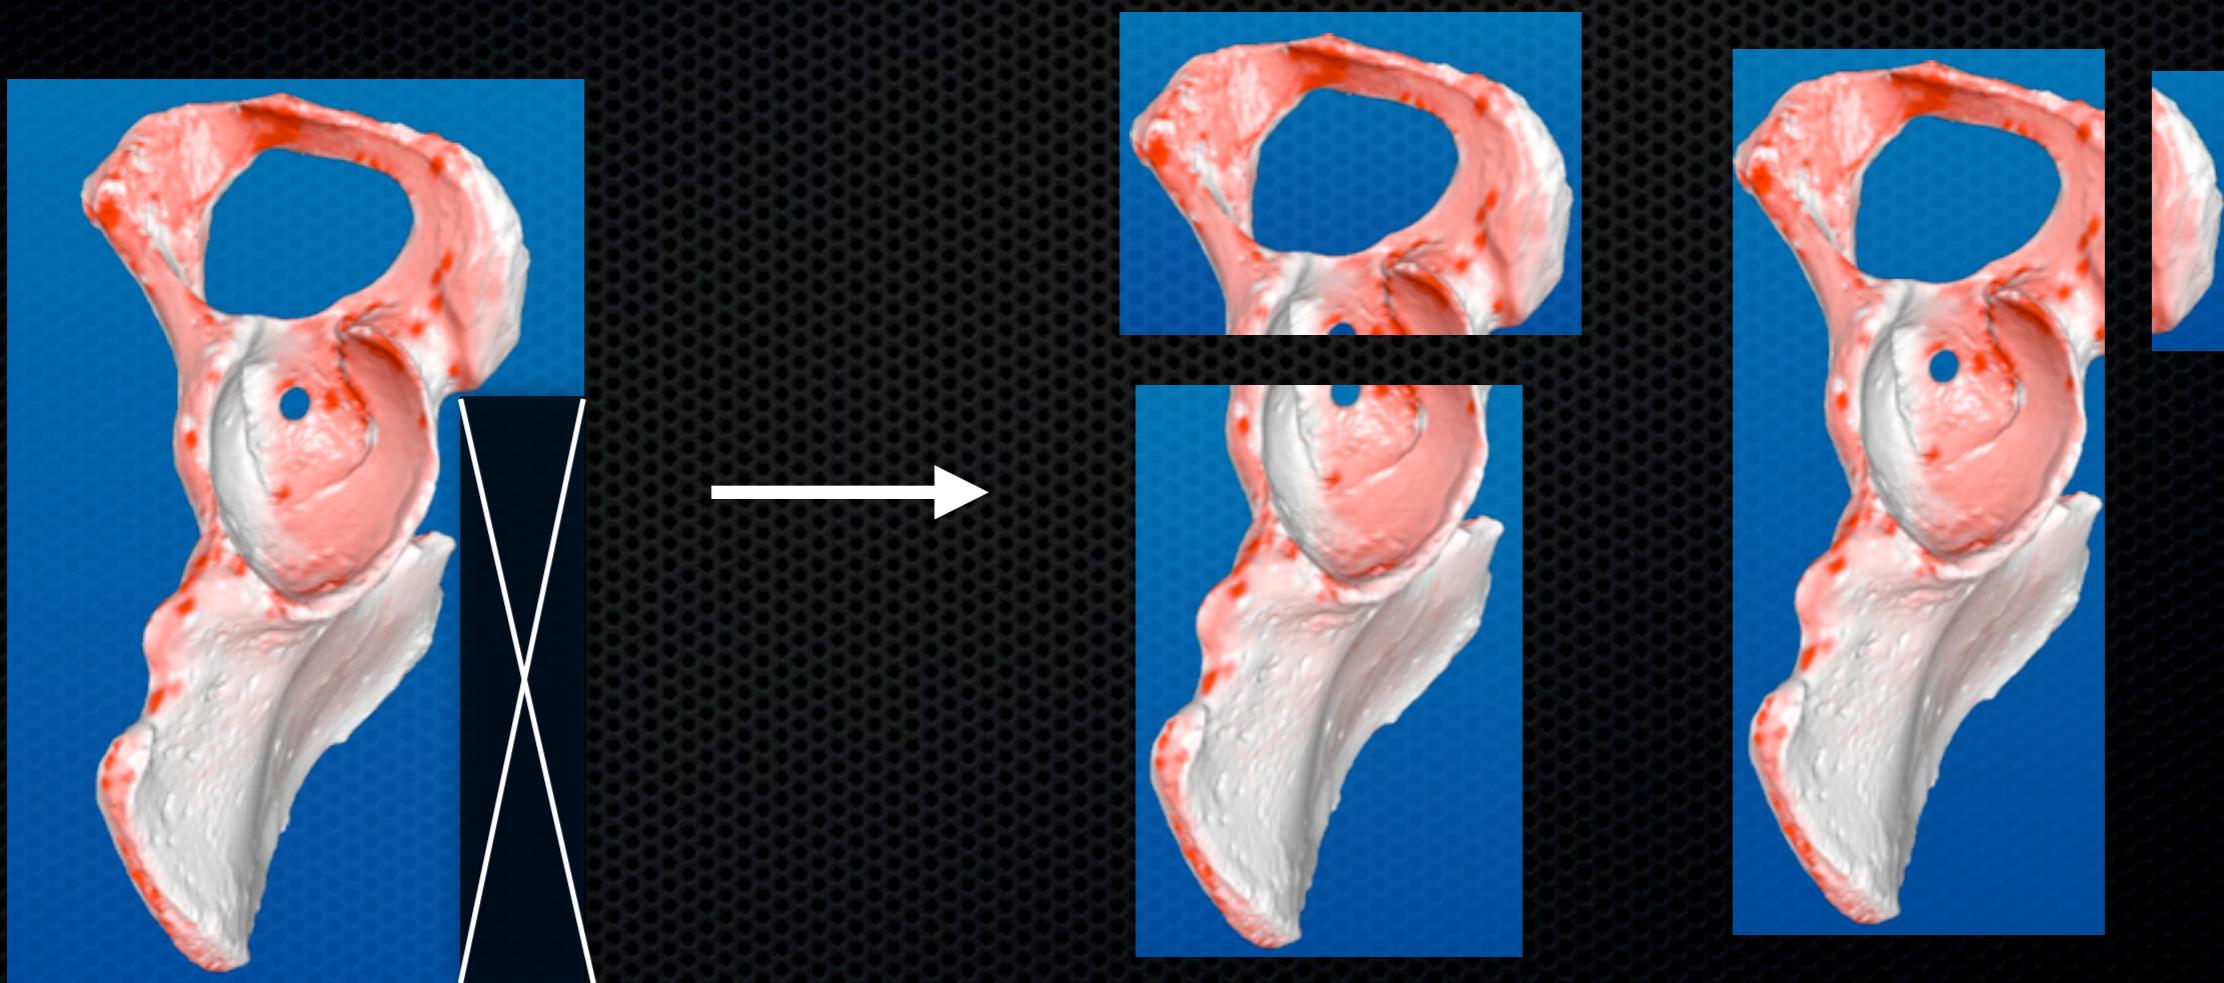
Regions requirements:

- Cover all data
- No overlapping
- Smallest total area

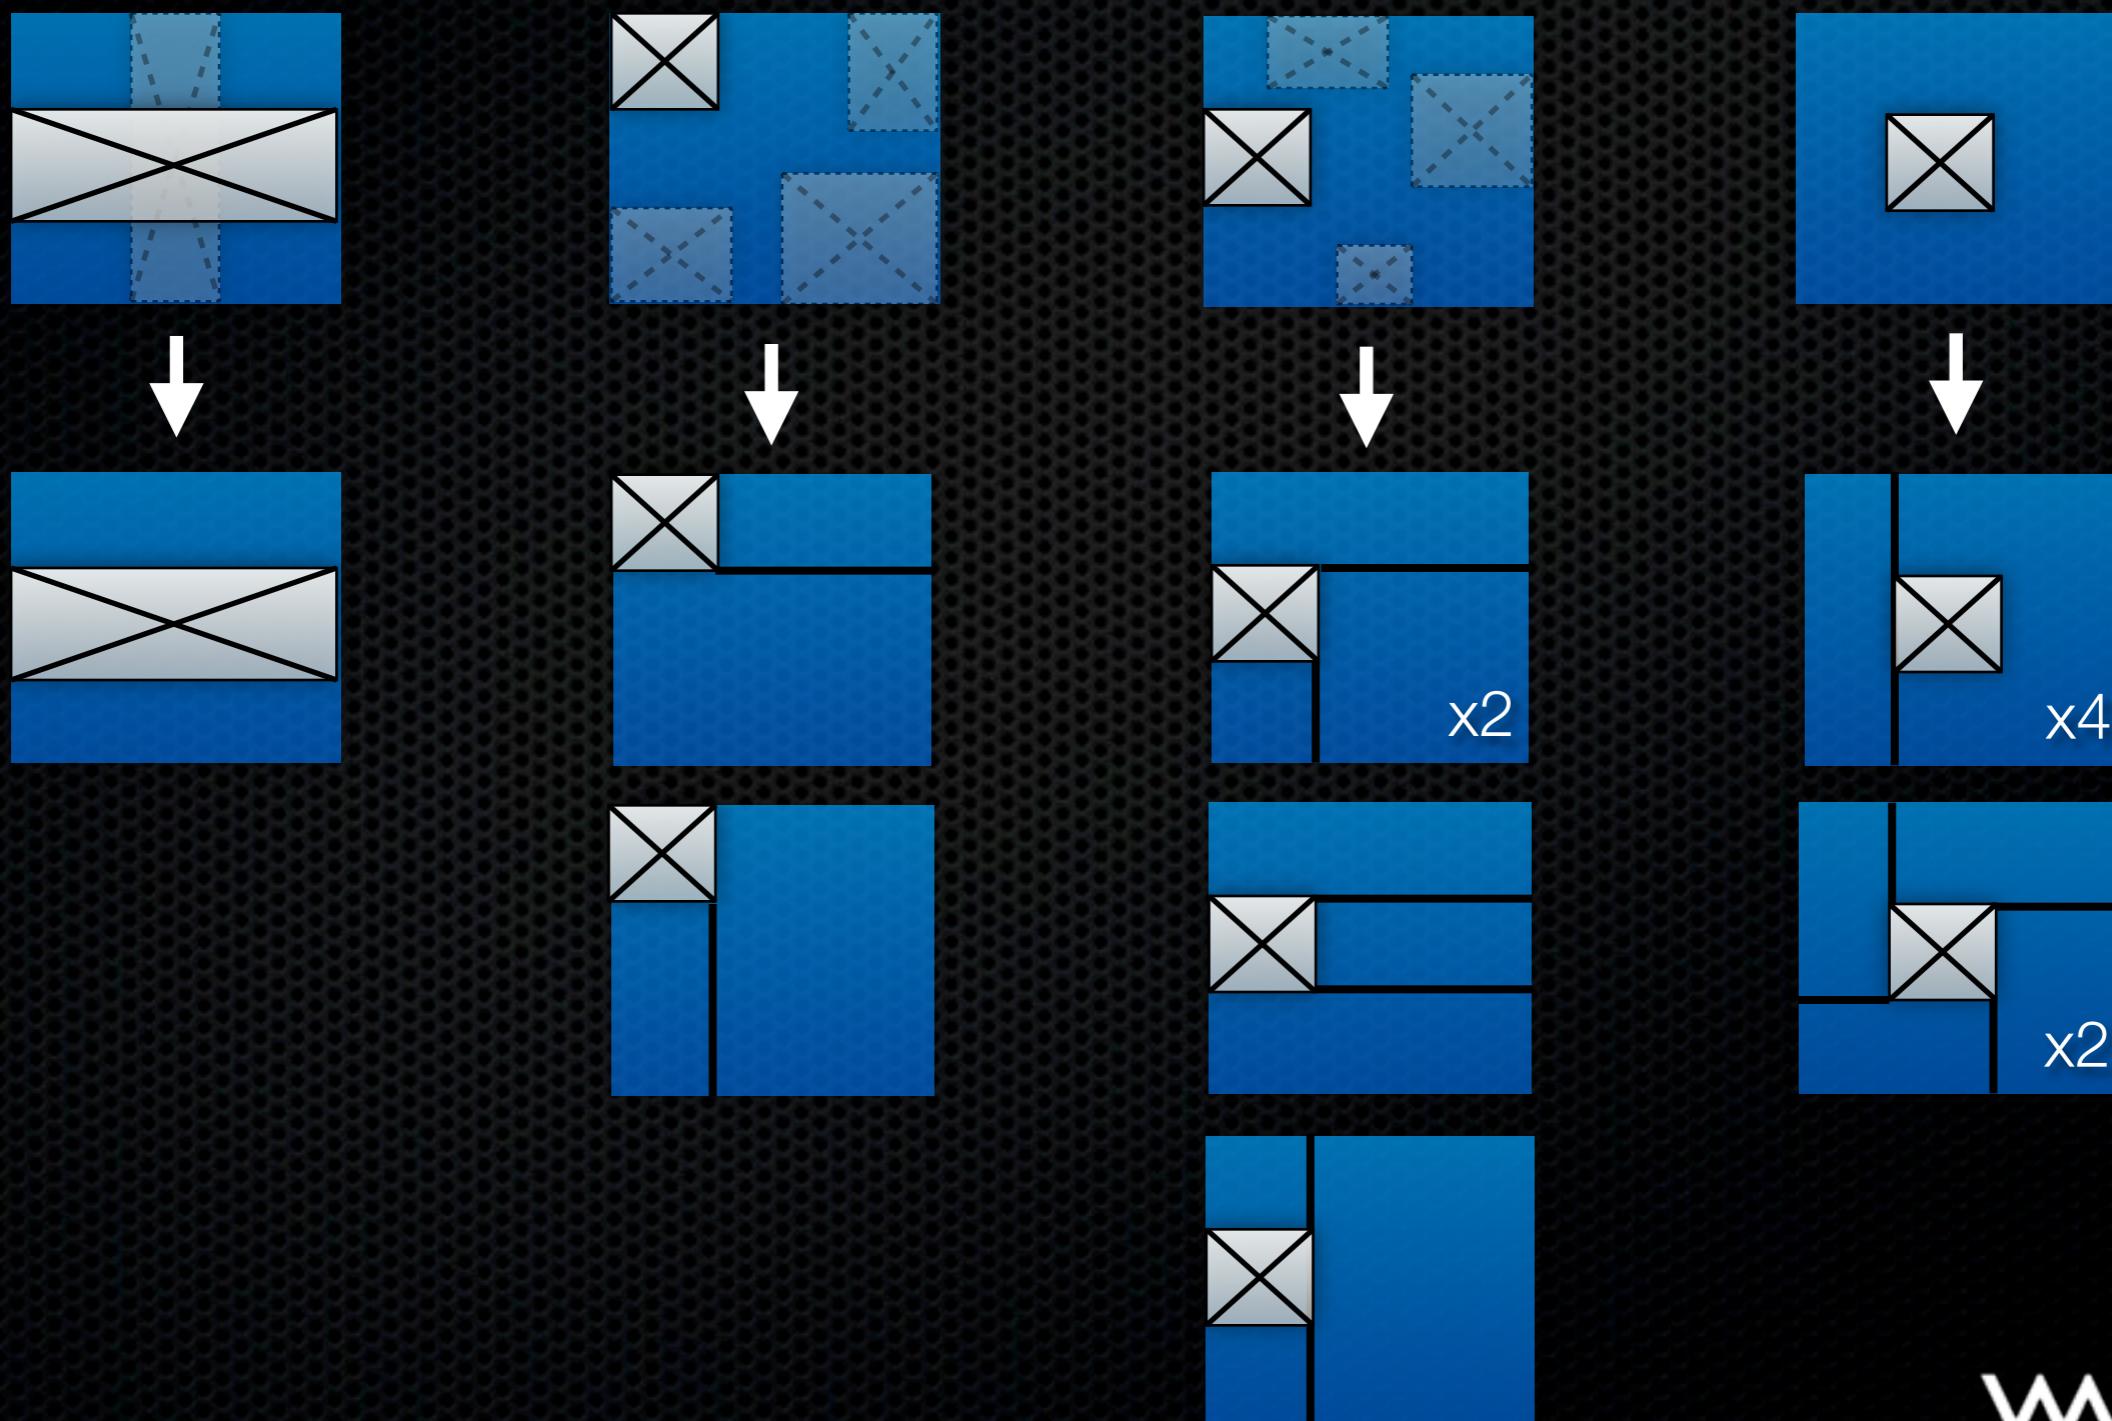


# Methods - ROI

Cut away unused space instead!



# Methods - ROI



# Methods - ROI



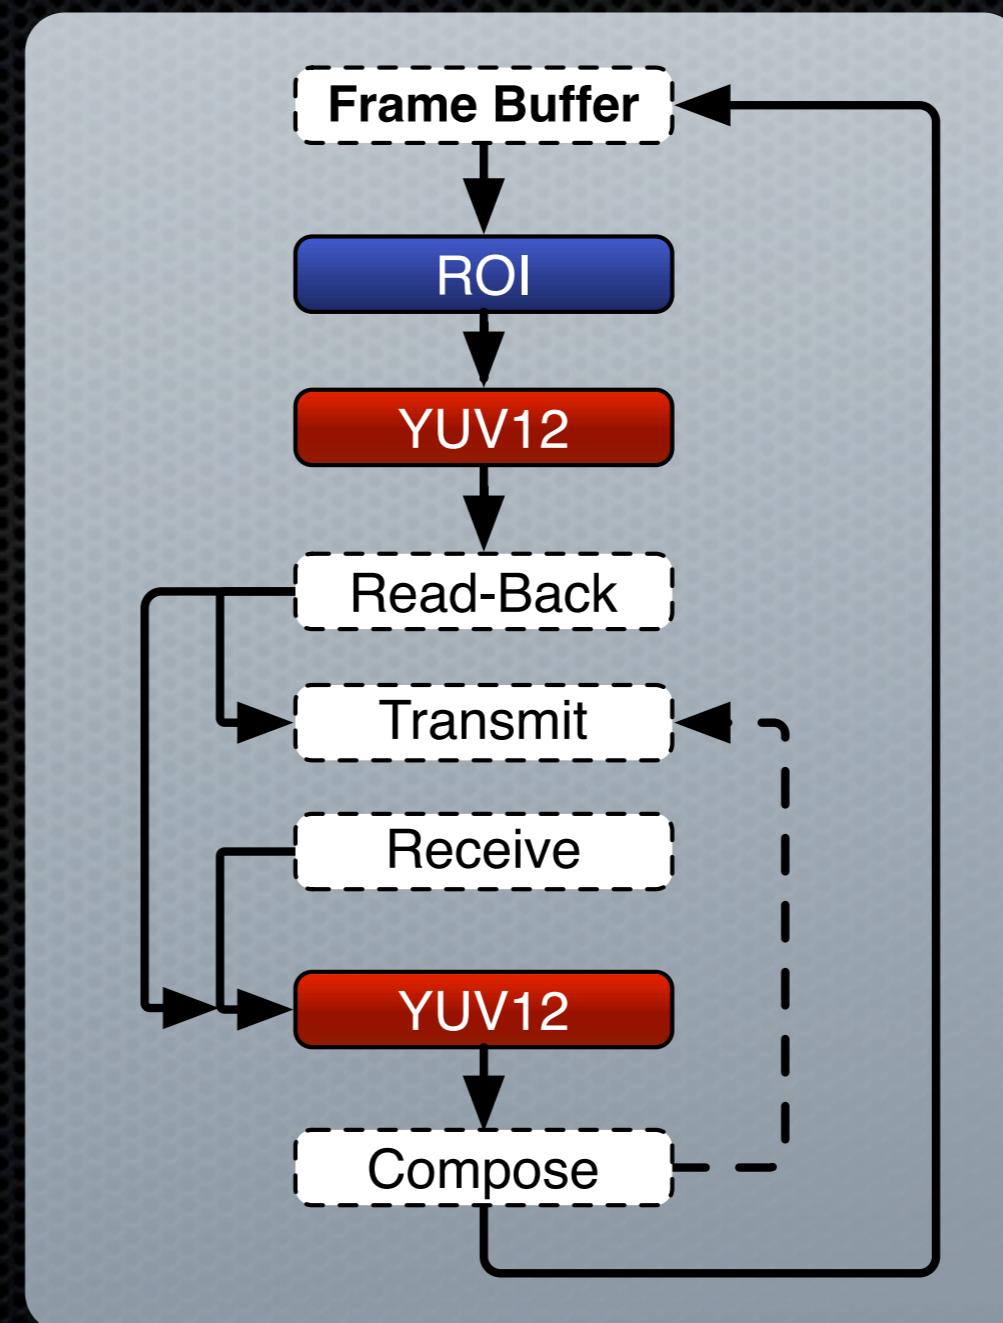
# ROI Demonstration

Demo

# Methods - Integration

ROI and YUV12  
simply integrated as  
separate layers

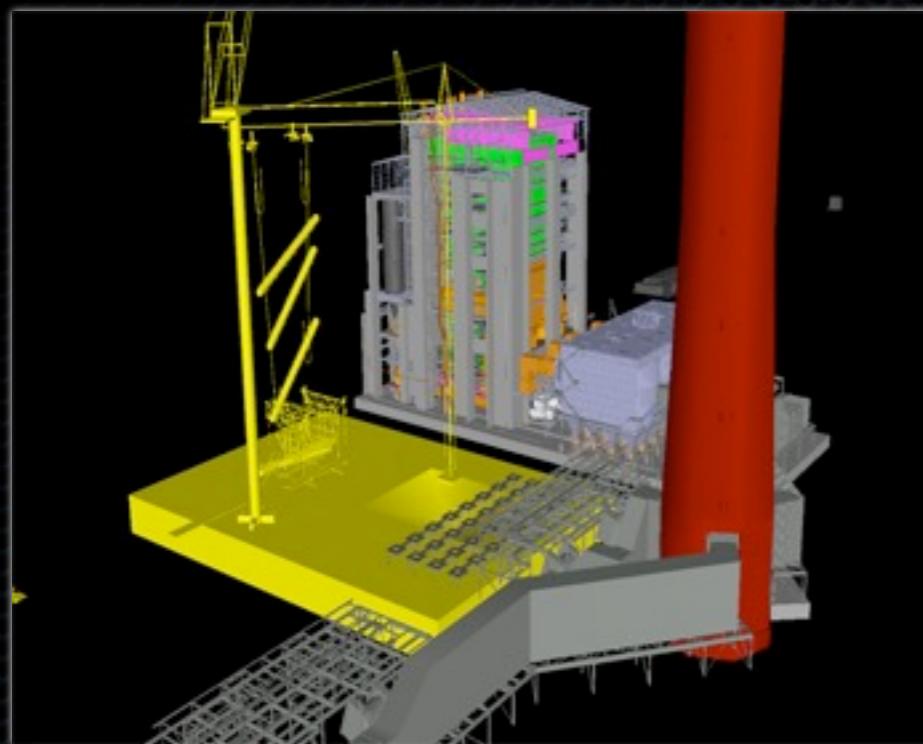
Compatible with  
other compression  
methods



# Analysis

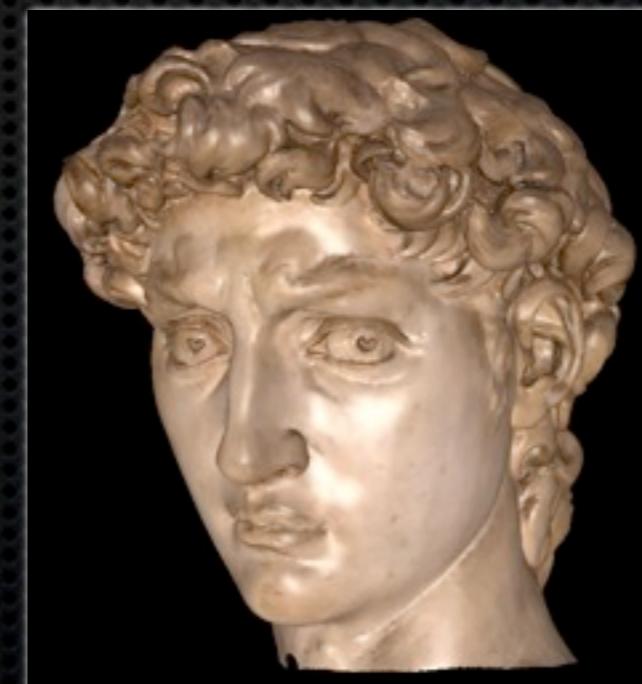
- Setup
- Throughput and rendering limits
- Sort-first performance
- Sort-last performance
- ROI for sort-last

# Models

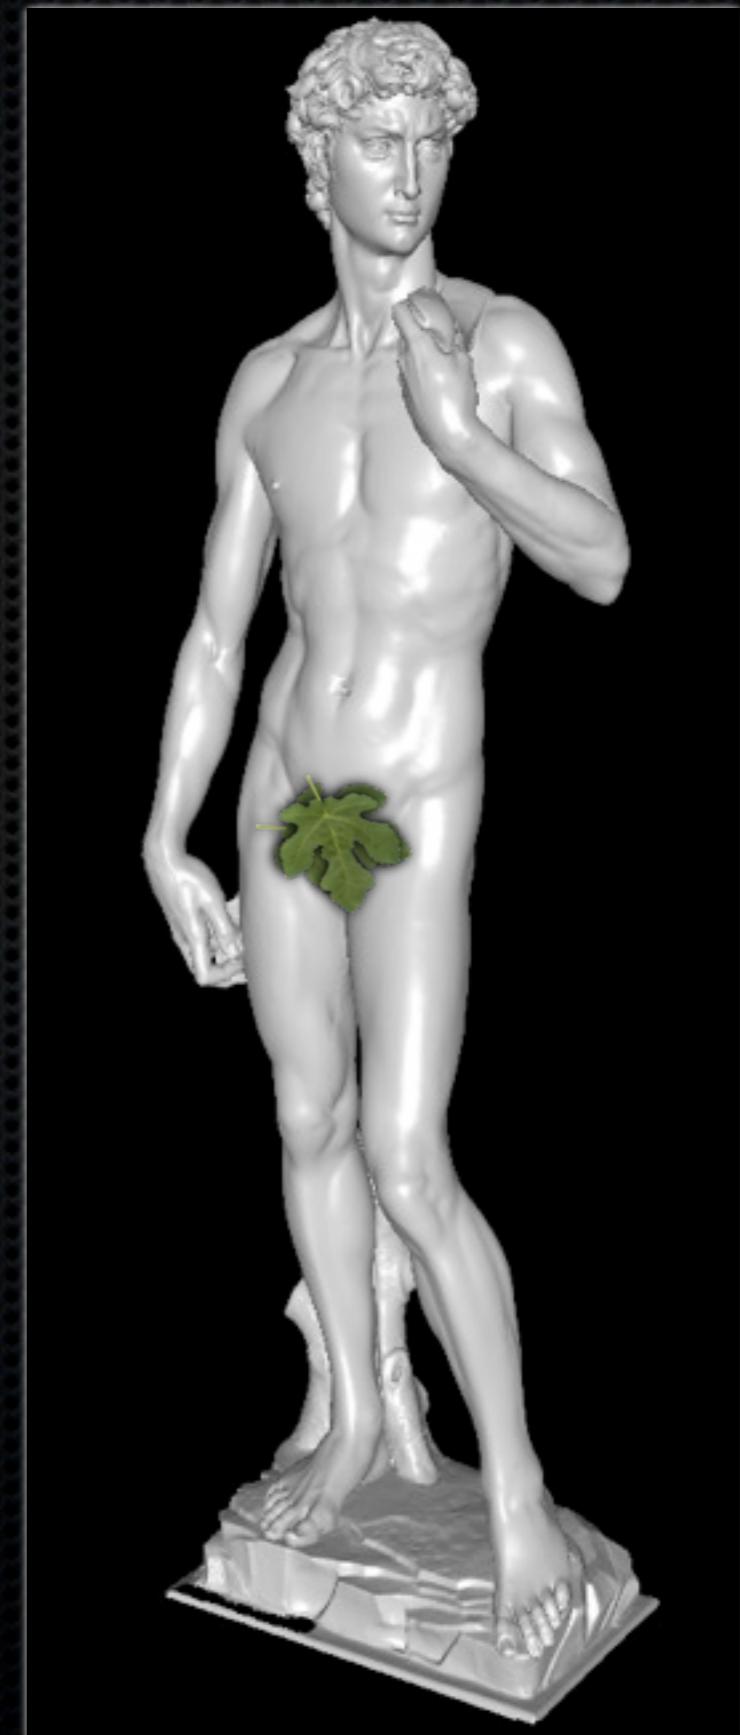


Power plant  
~13 M

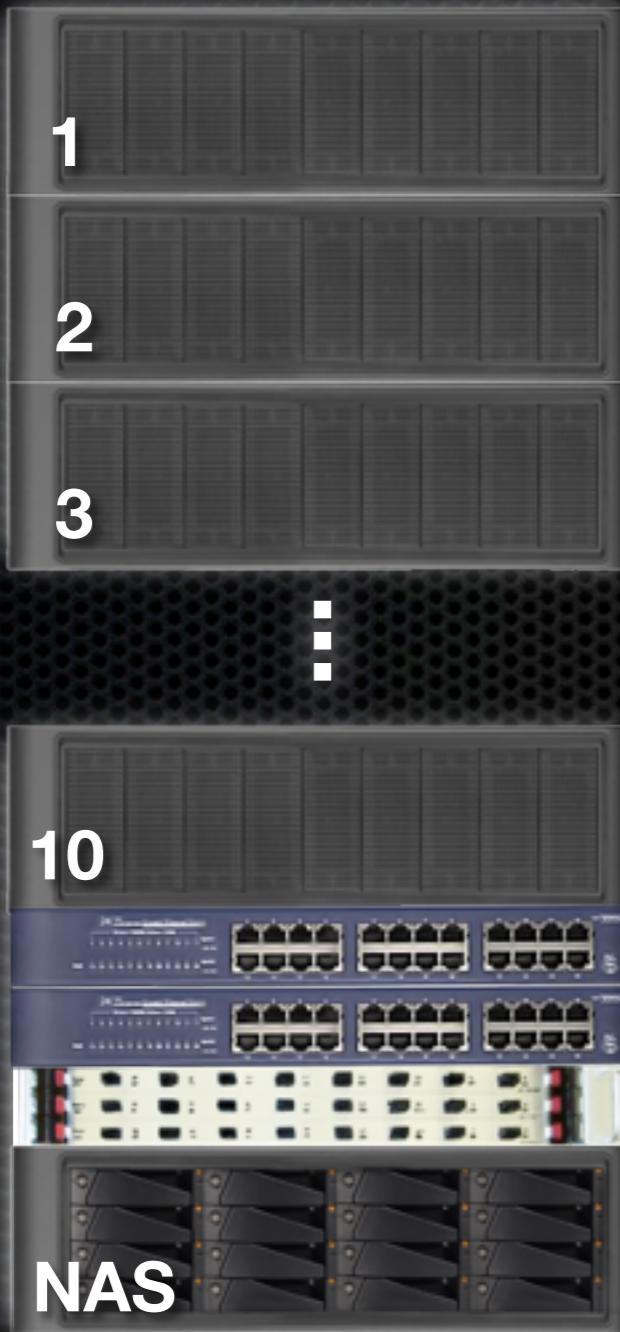
David head  
~4 M



David statue  
~56 M

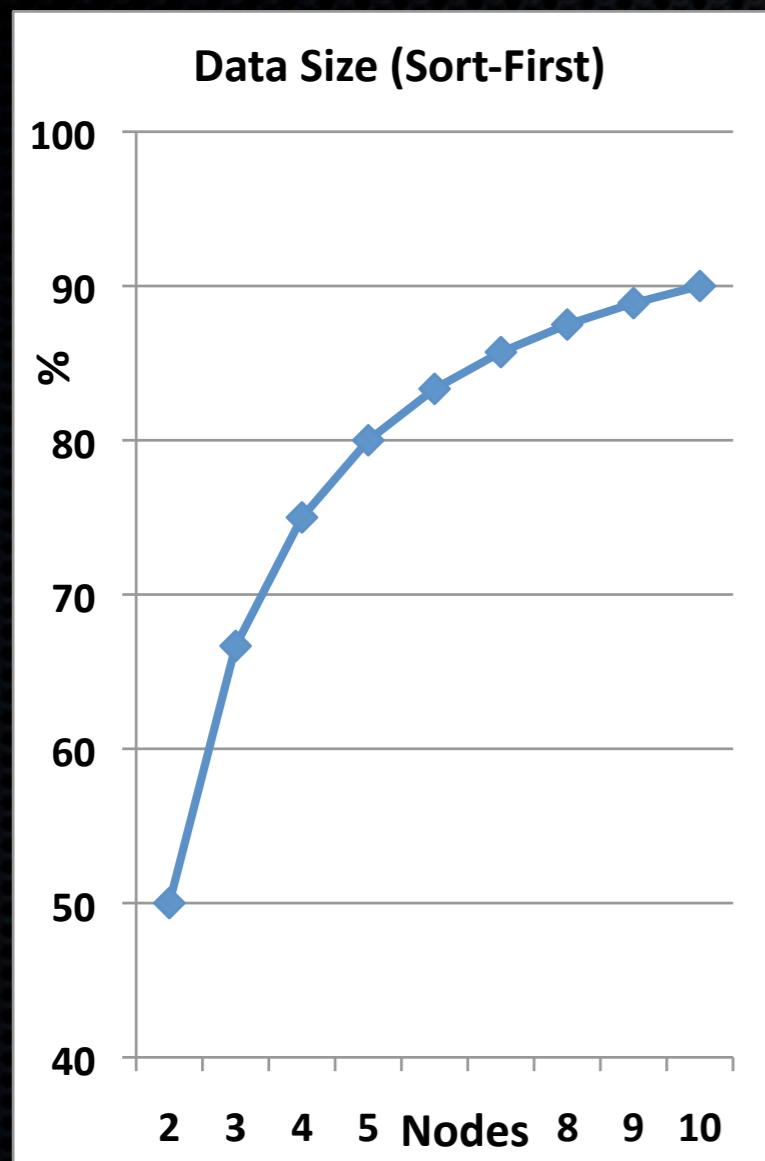


# Cluster Setup (Hactar)

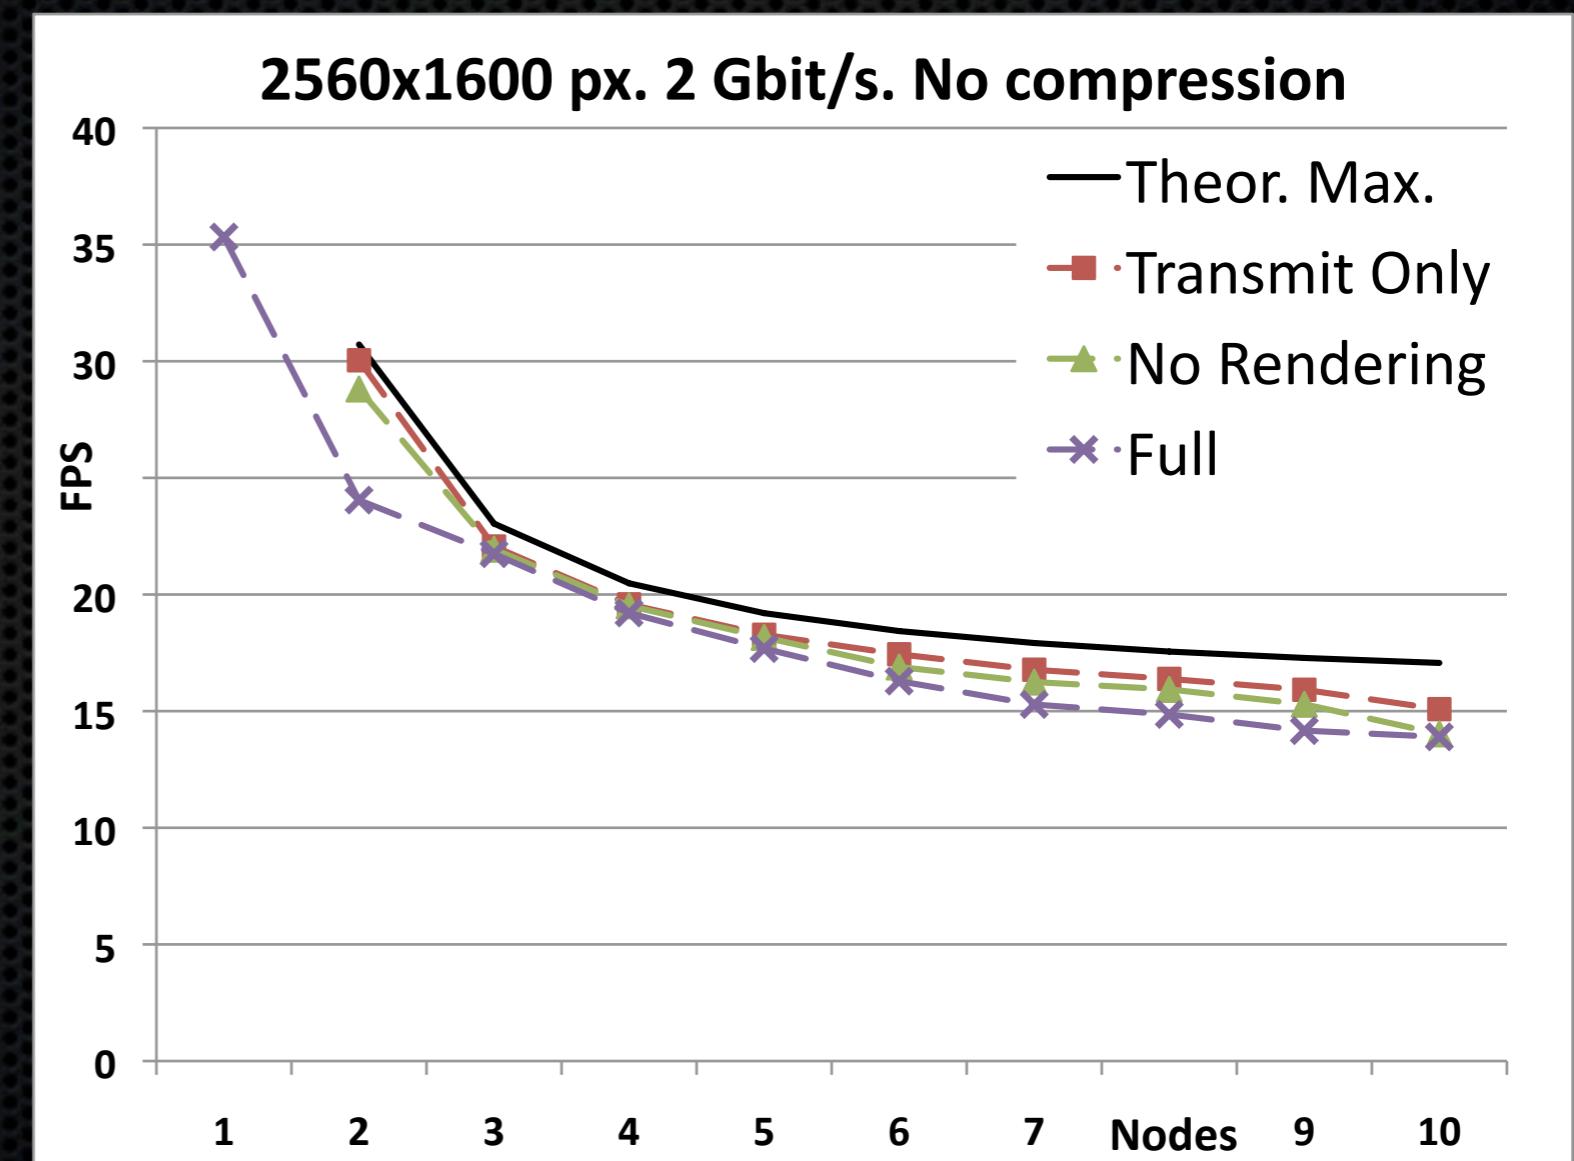


- 2x 2.2GHz AMD
- GeForce 9800 GX2
- 4 GB RAM
- 1 Gbit/s (Ethernet)
- 2 Gbit/s (Myrinet)
- 2 TB (4x 1Gbit/s Ethernet)

# Throughput and Rendering Limits (No Compression)

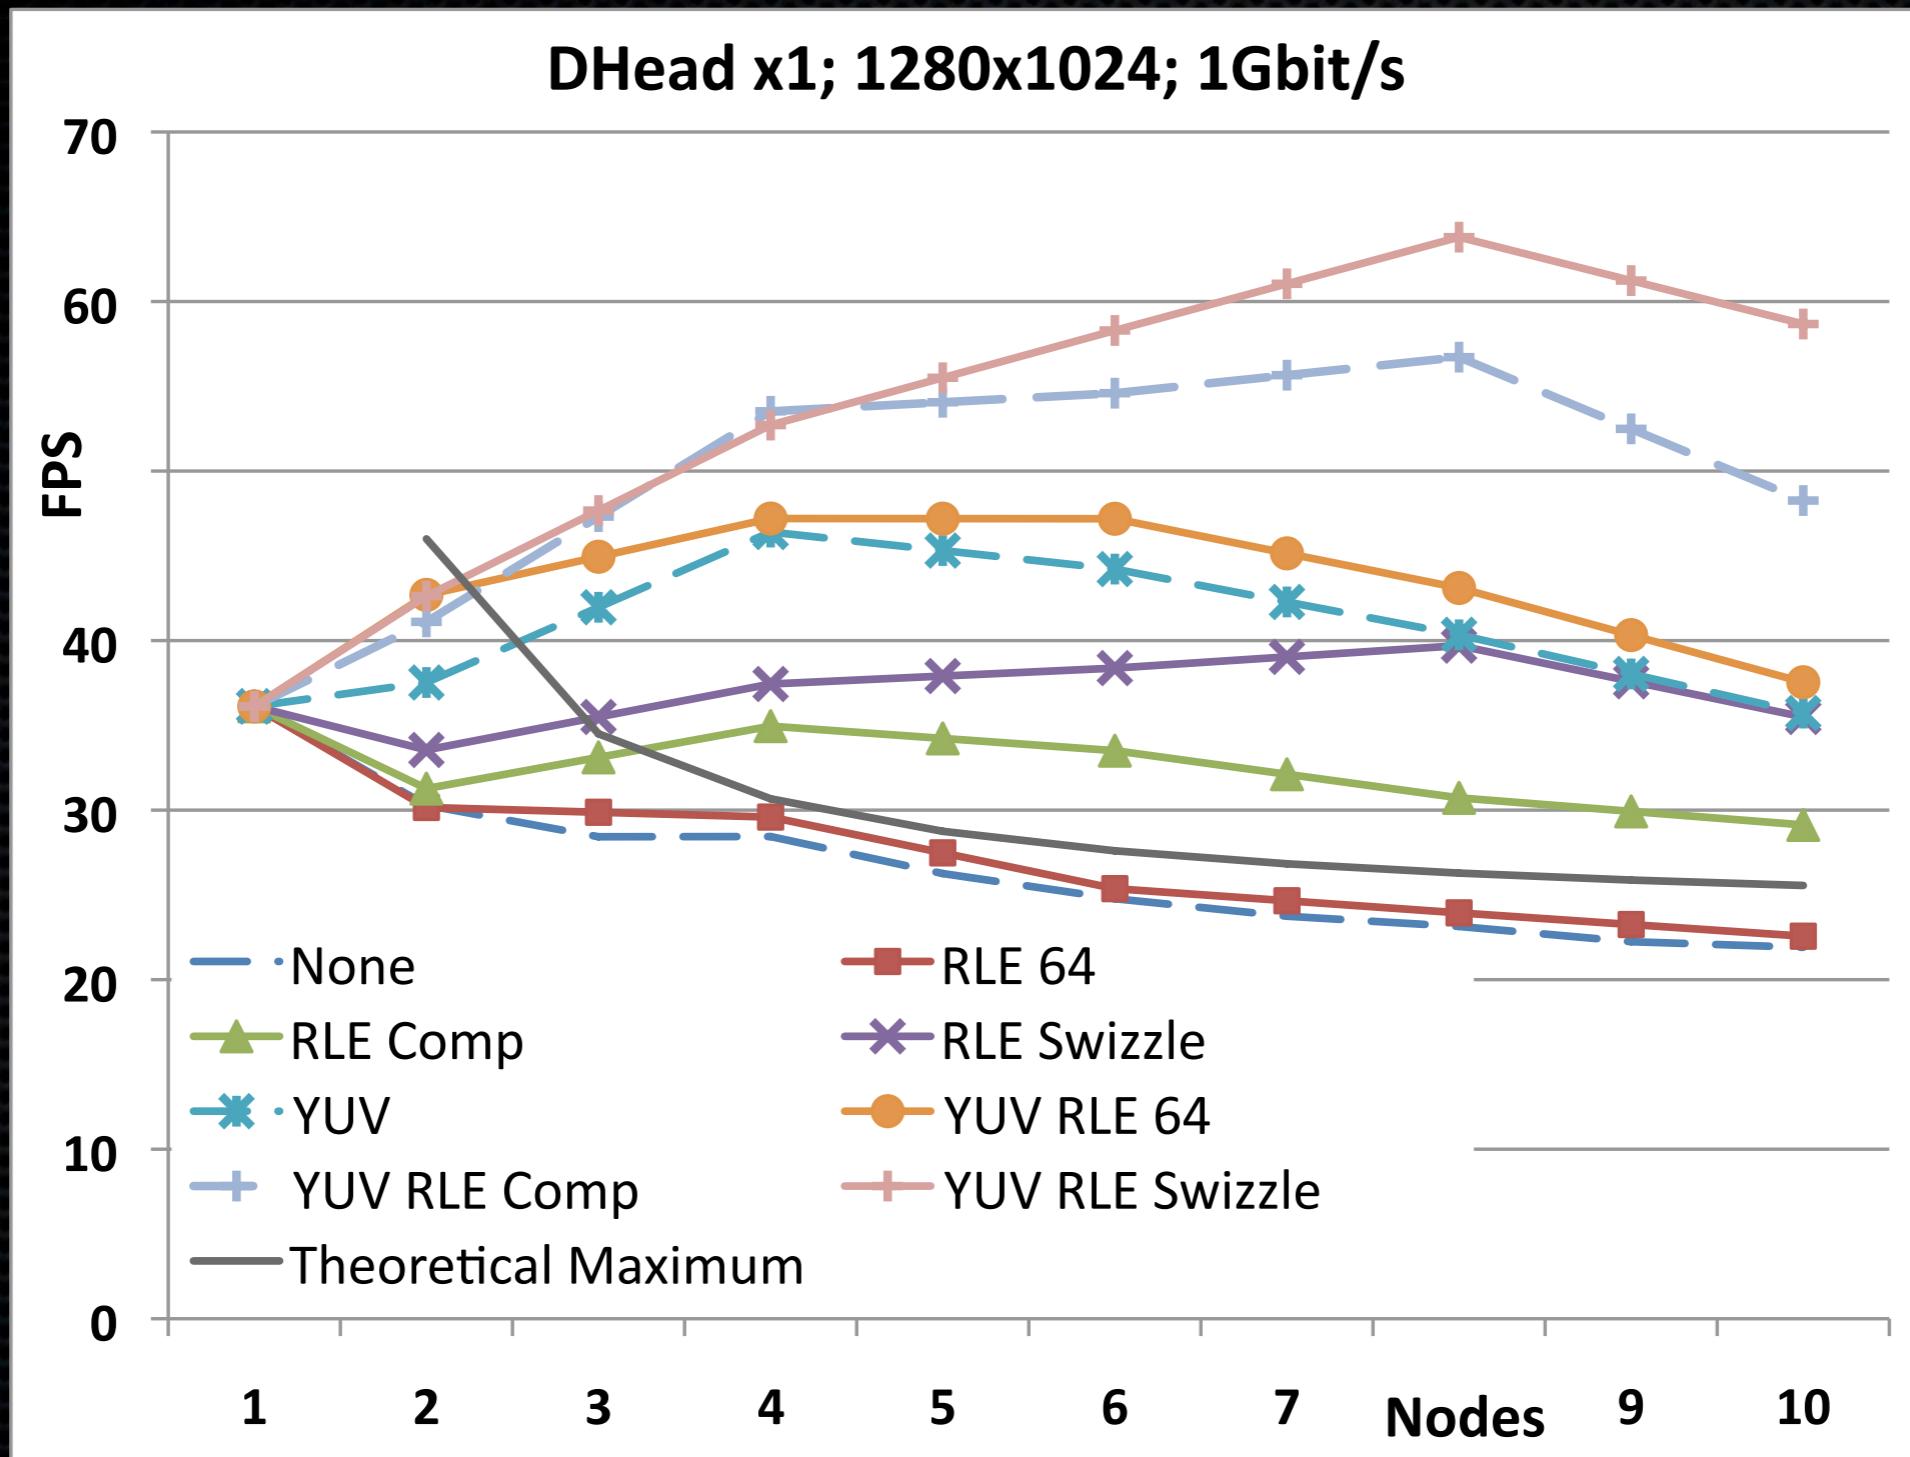


Data to receive

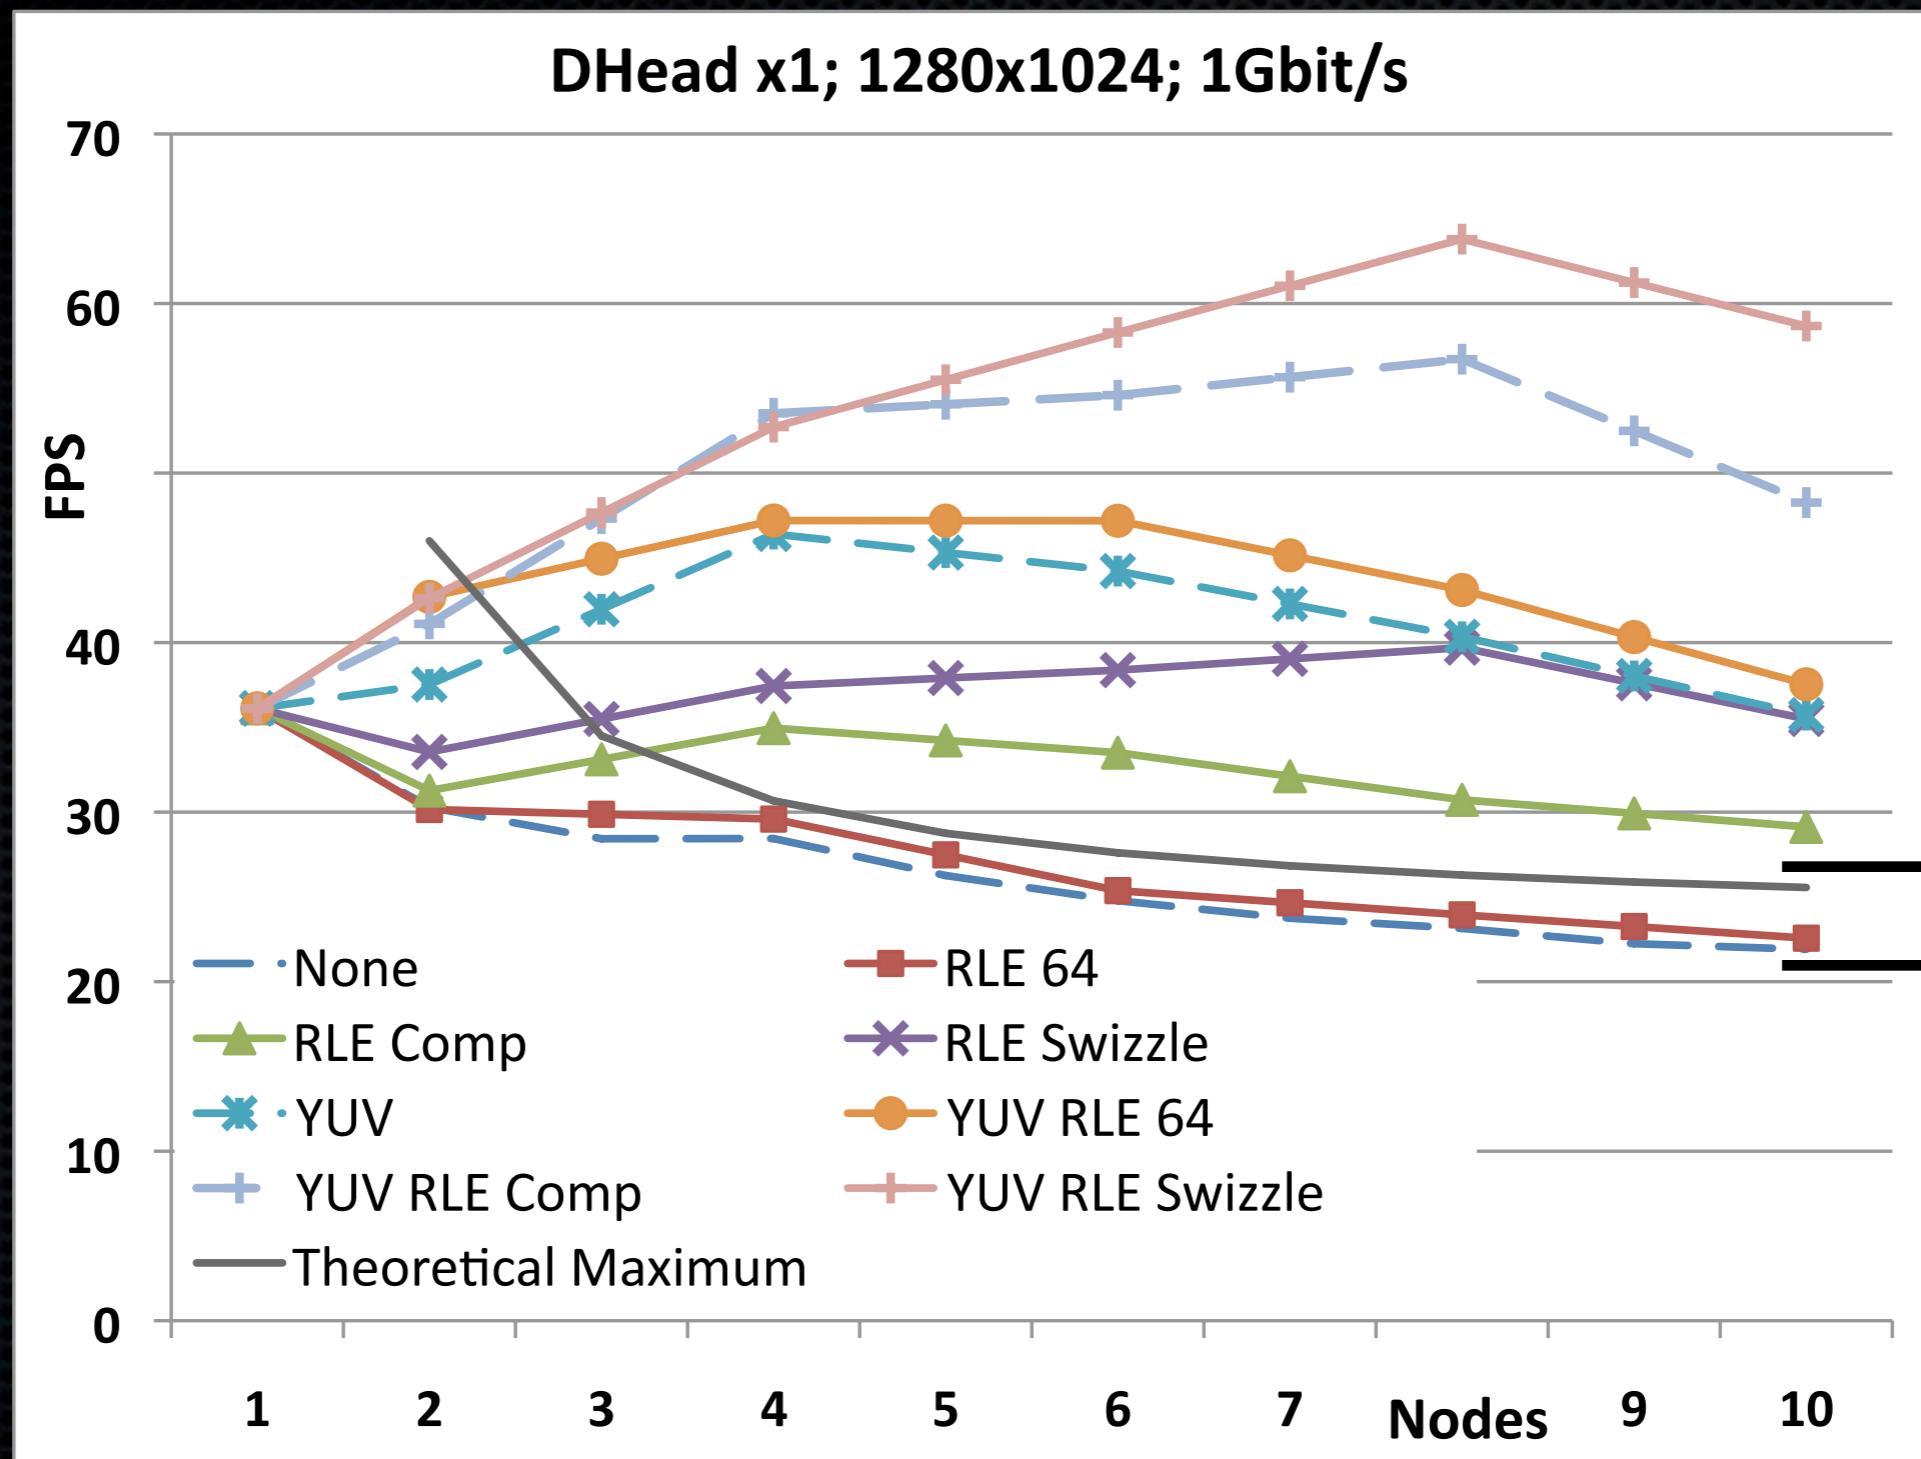


Performance of sort-first

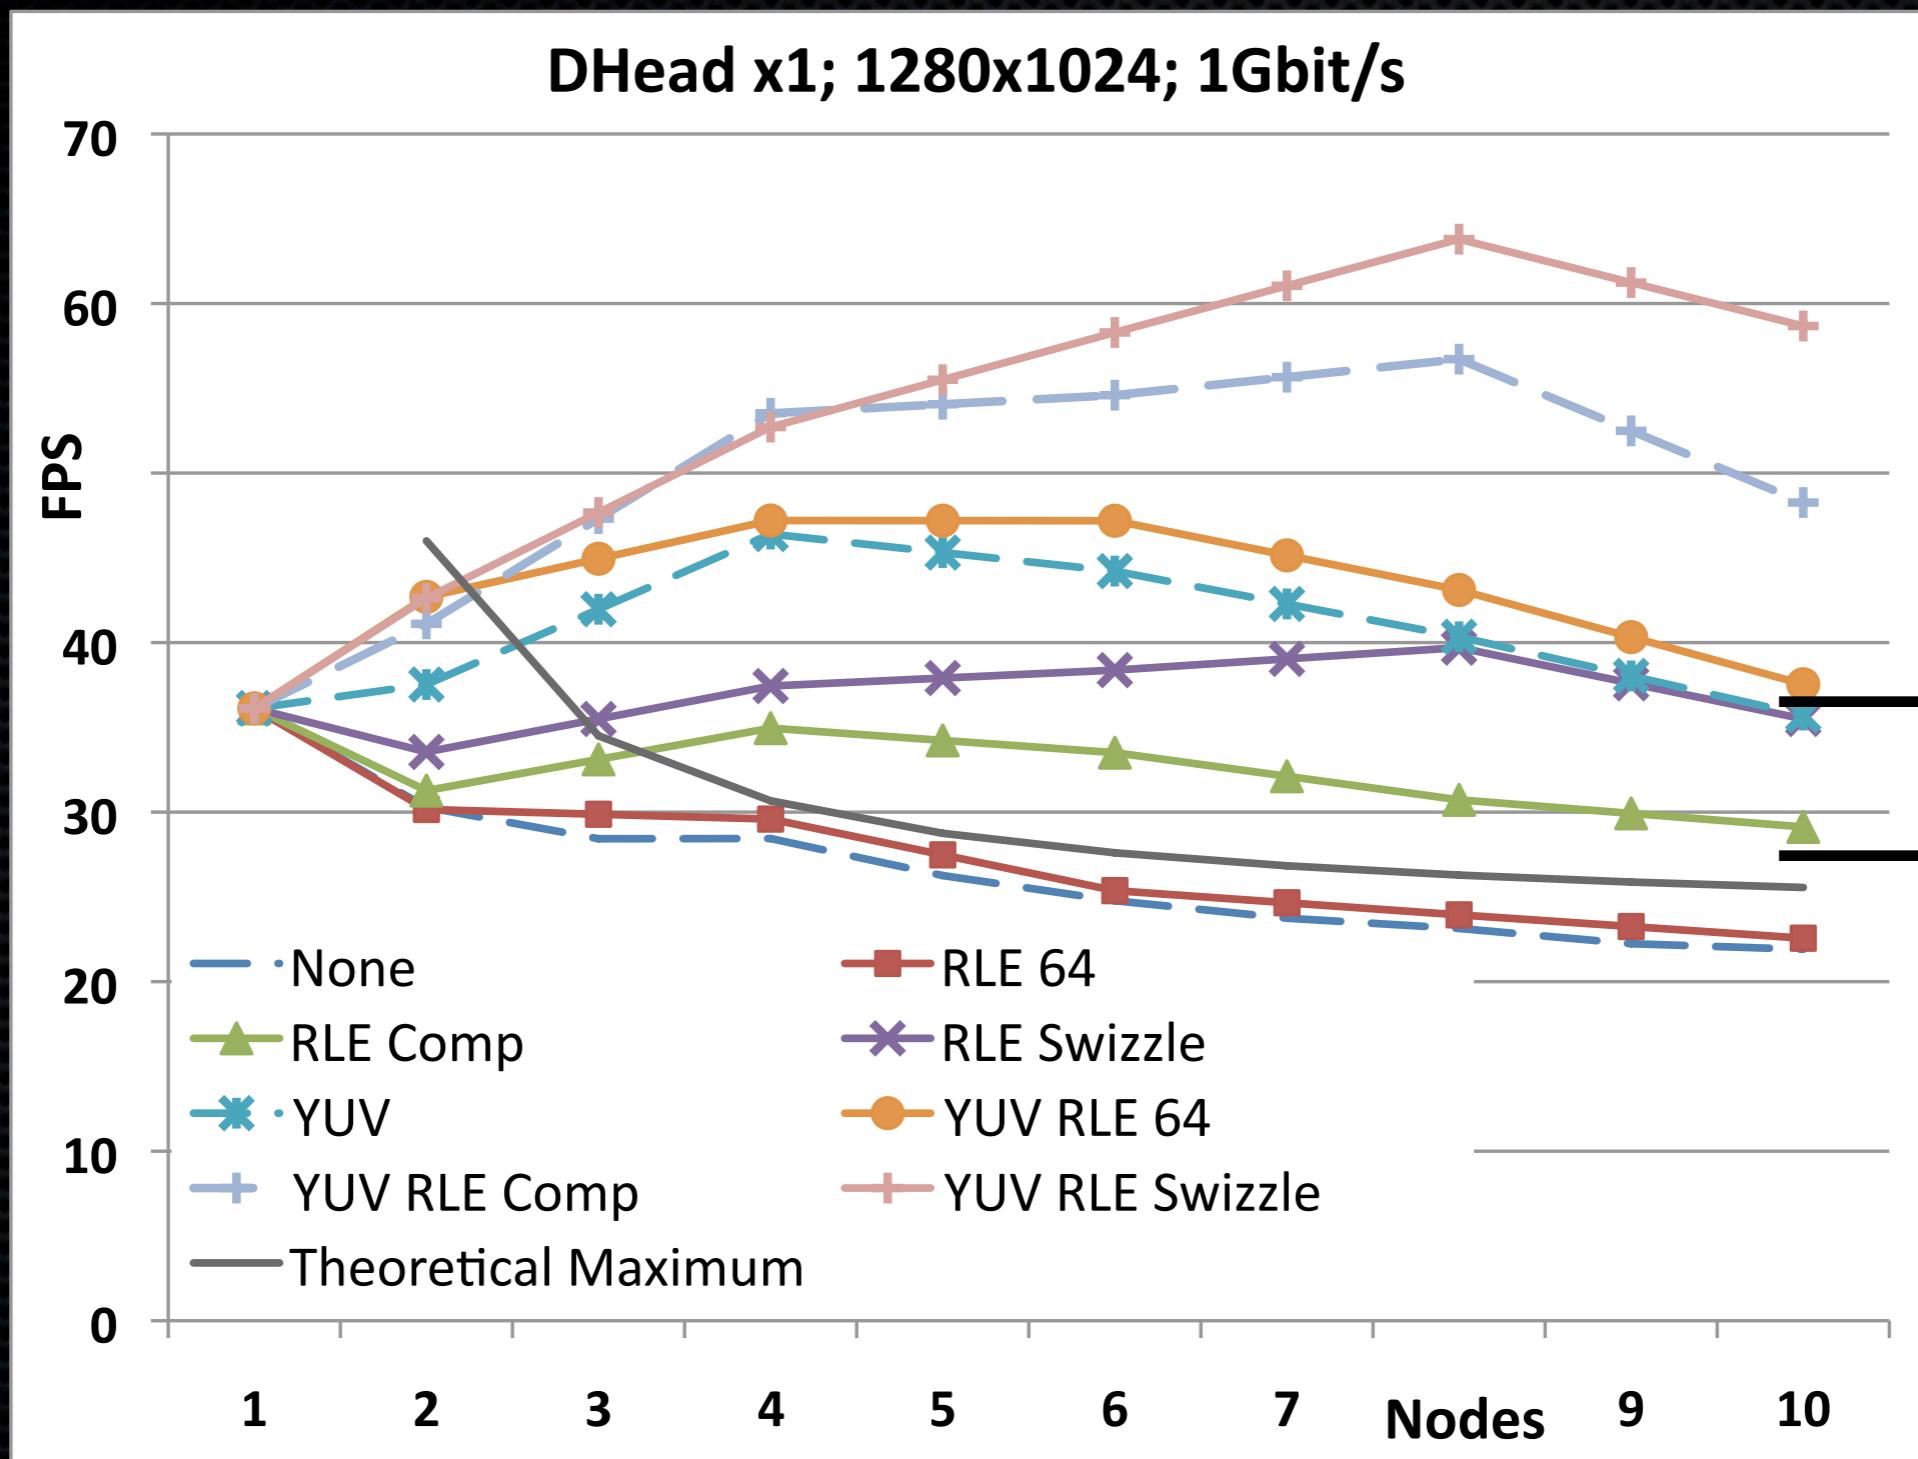
# Sort-First Performance



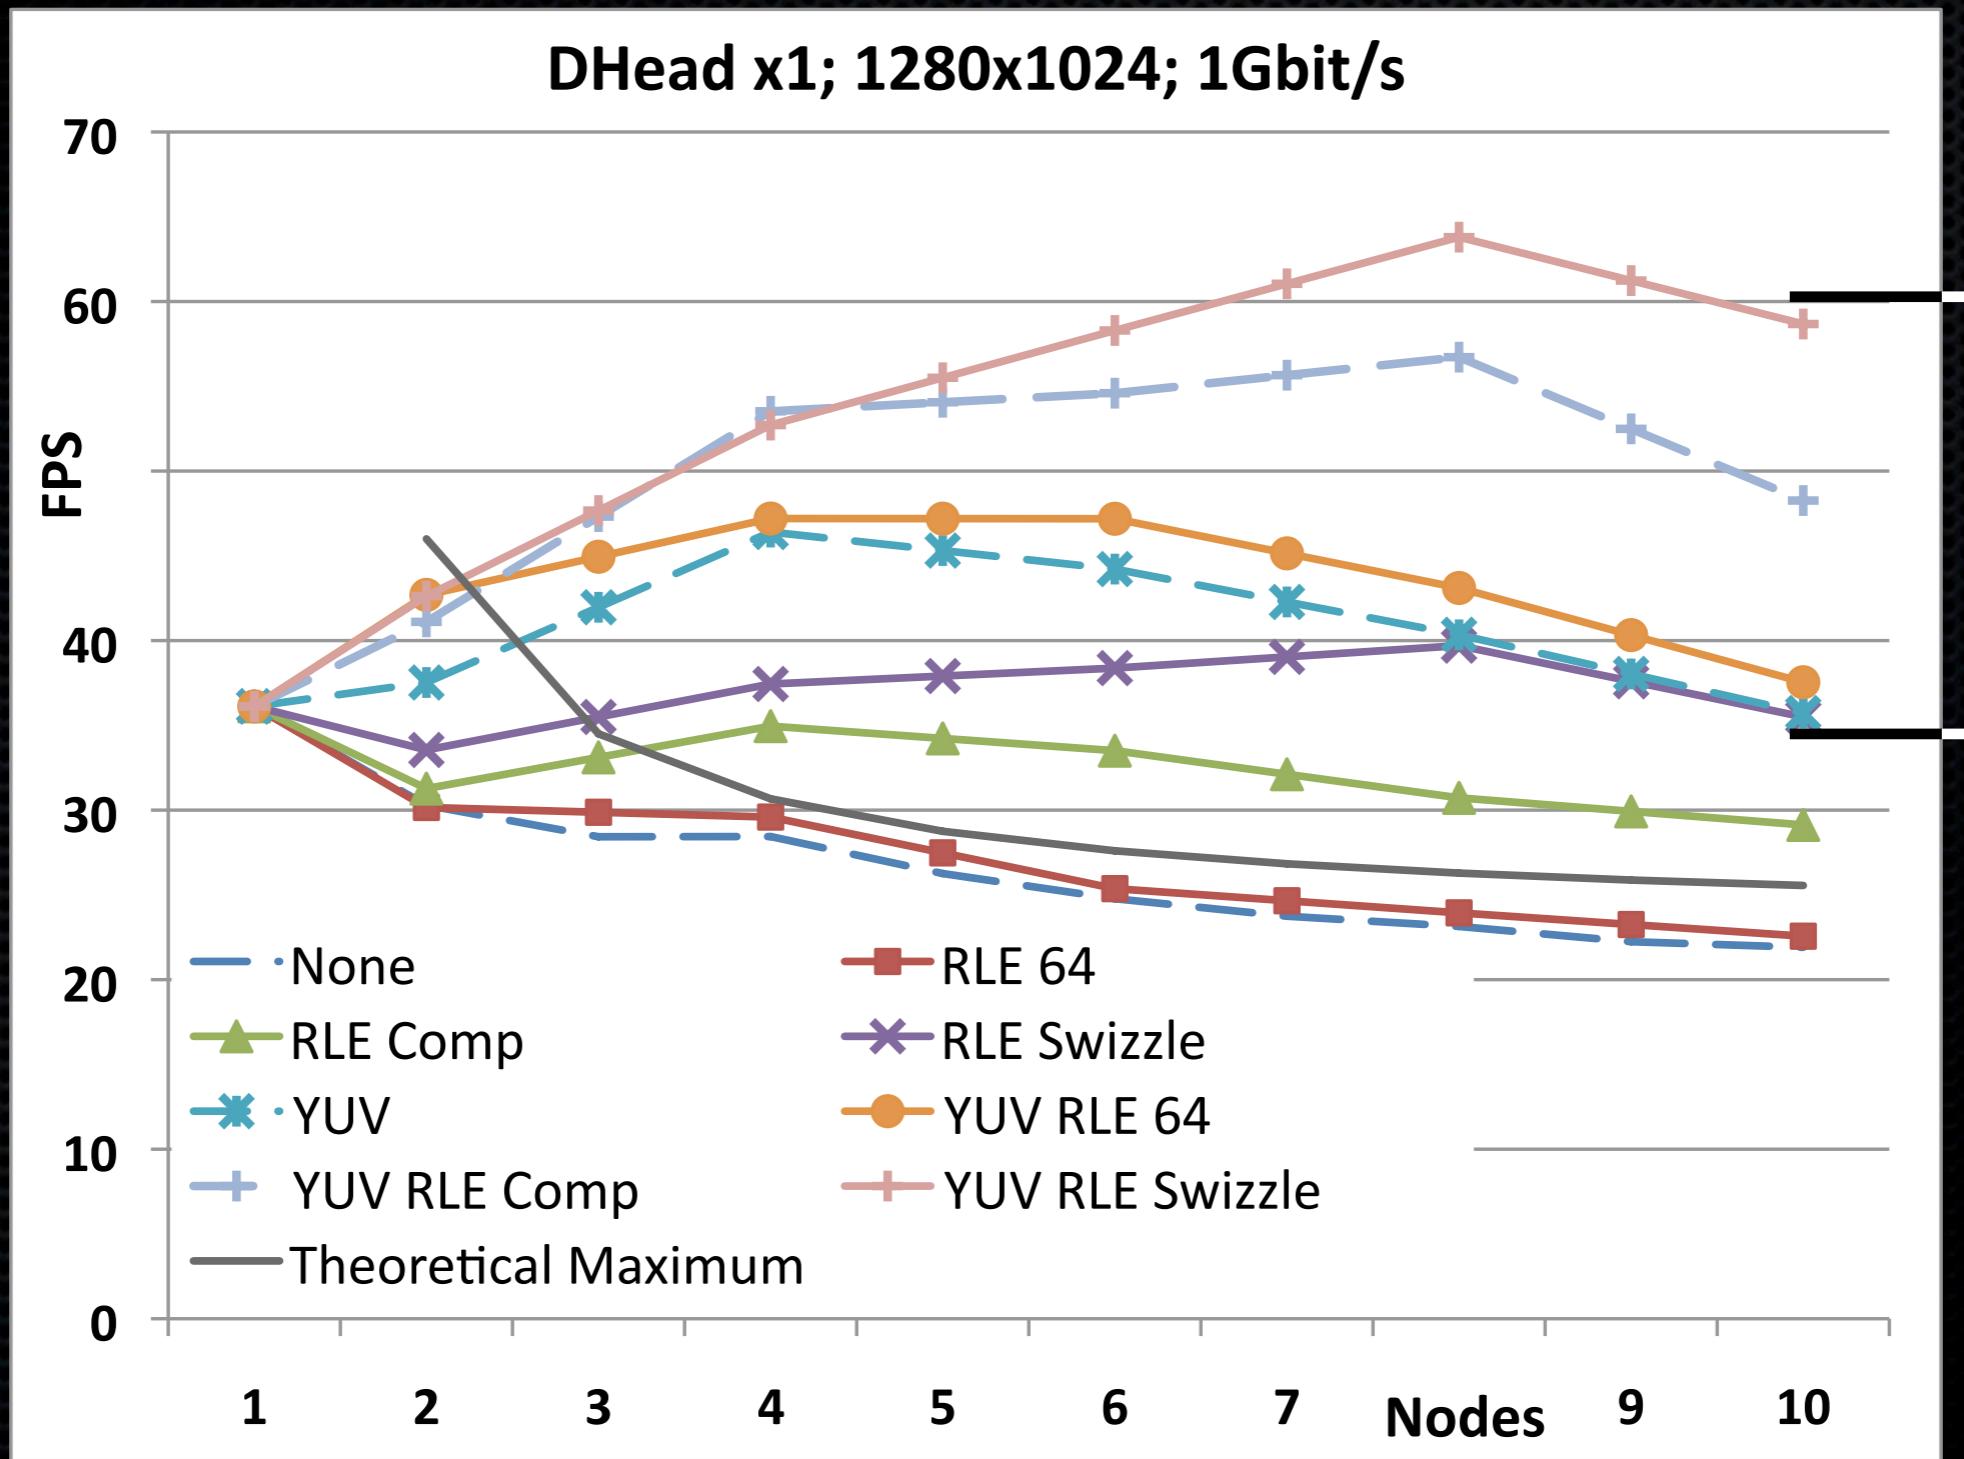
# Sort-First Performance



# Sort-First Performance

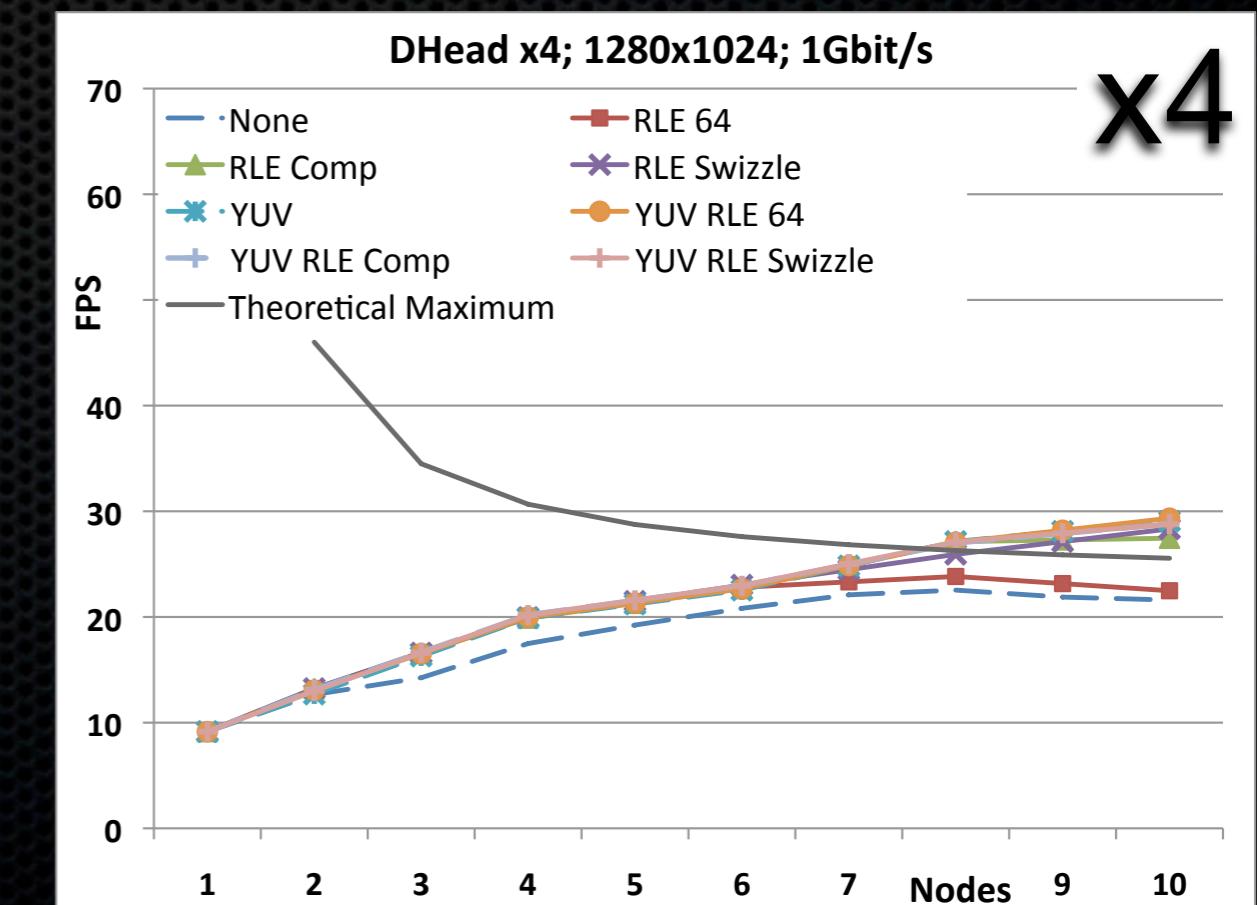
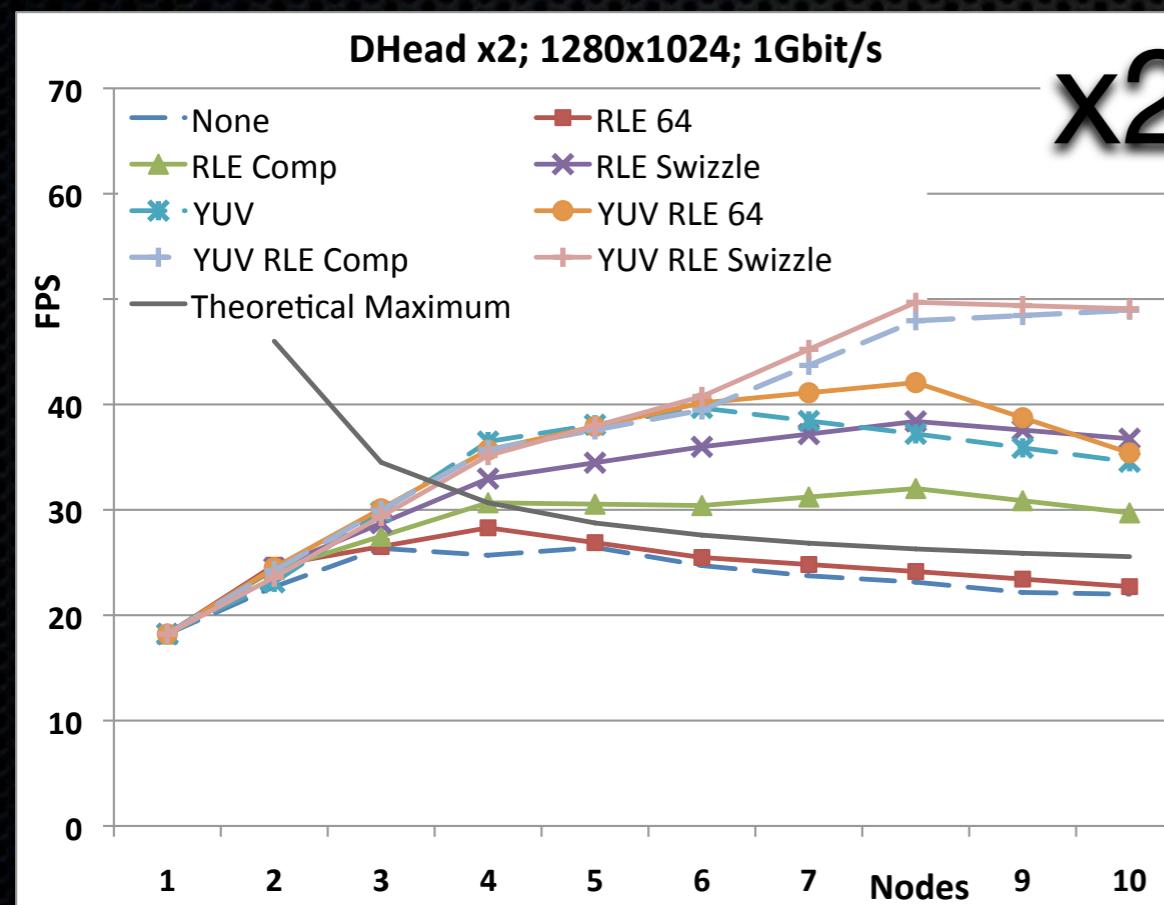
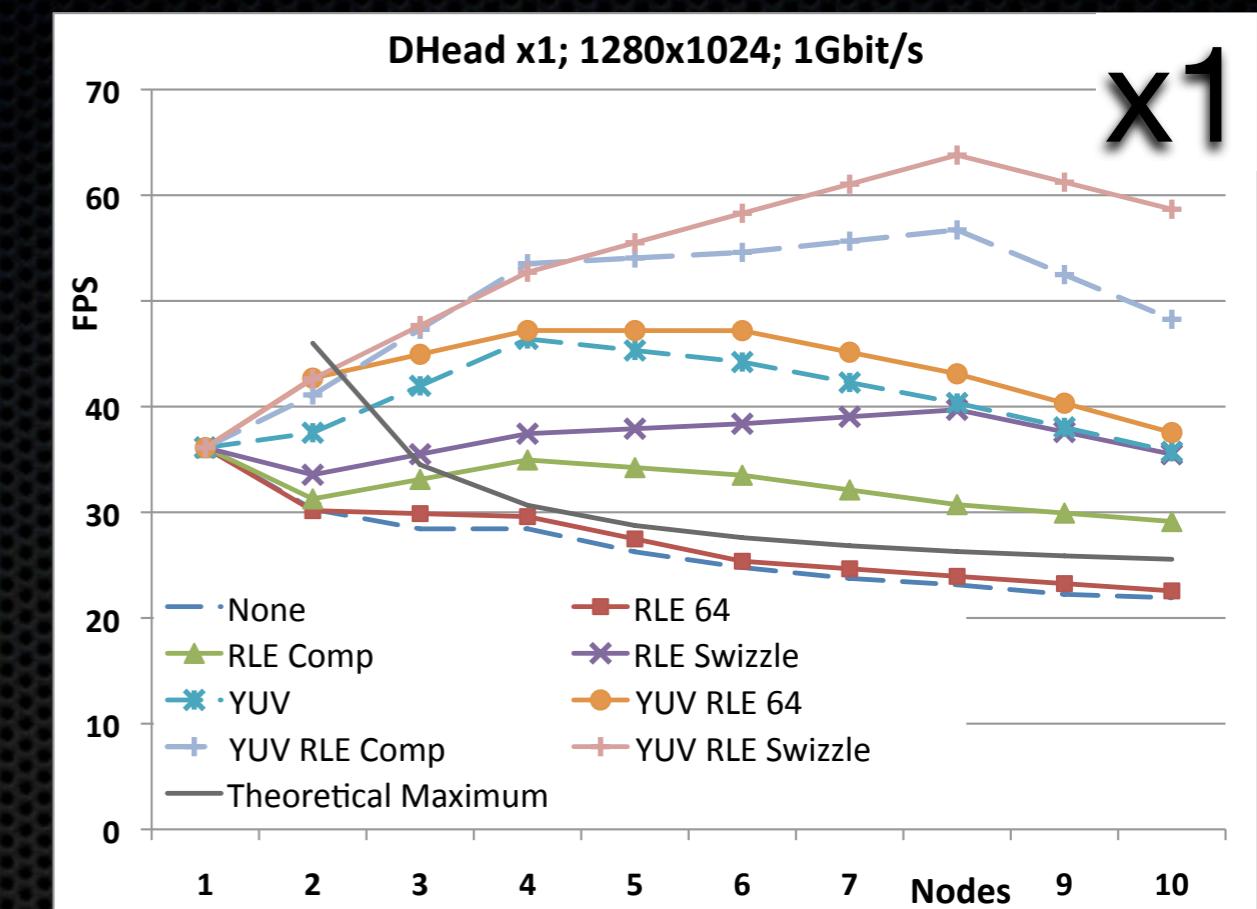


# Sort-First Performance

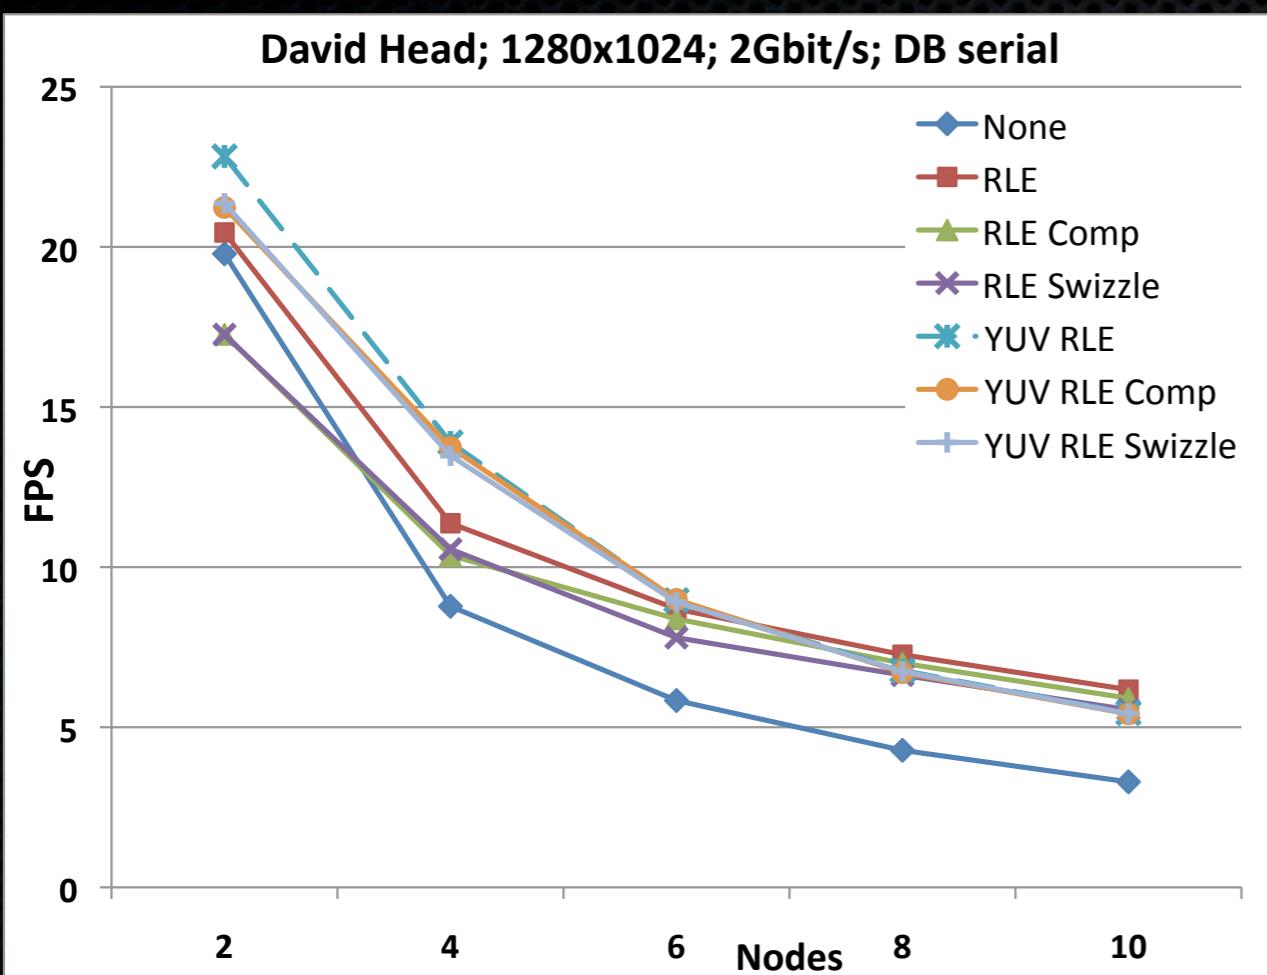


# Sort-First Performance

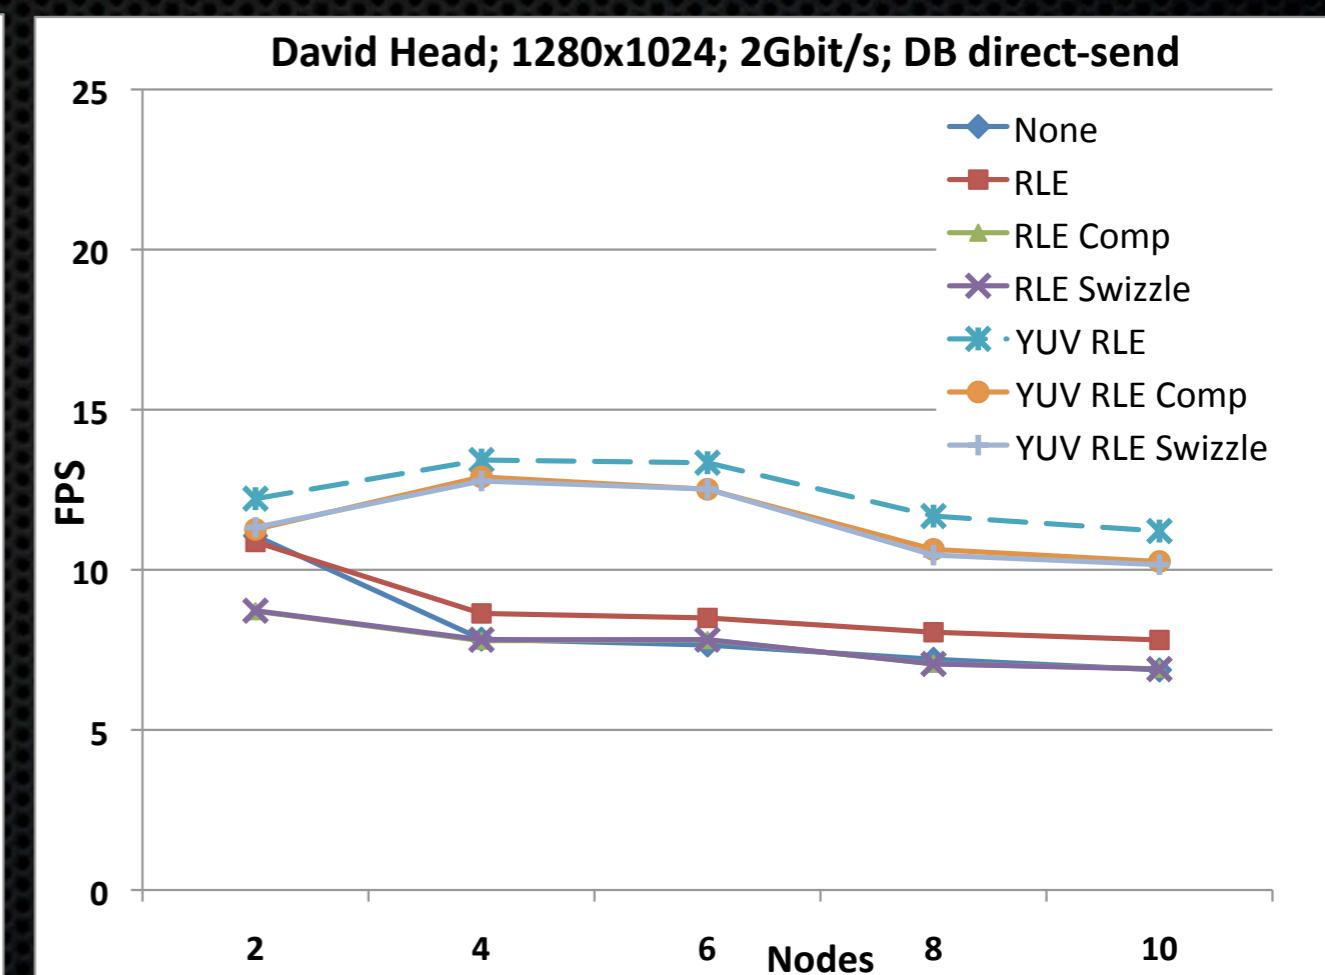
David head model  
rendered 1, 2 and 4 times



# Sort-Last Performance

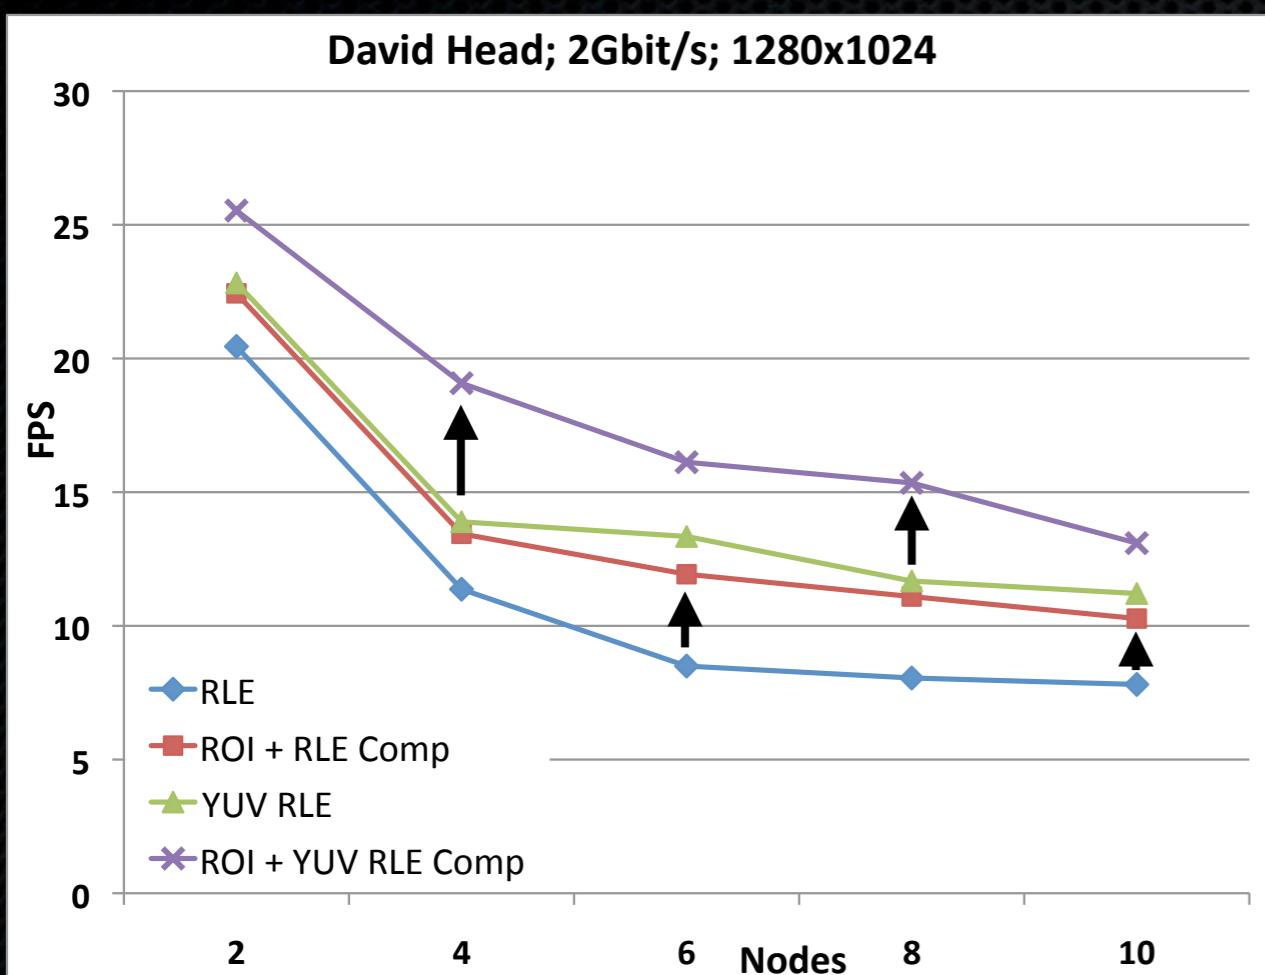


Serial Sort-Last

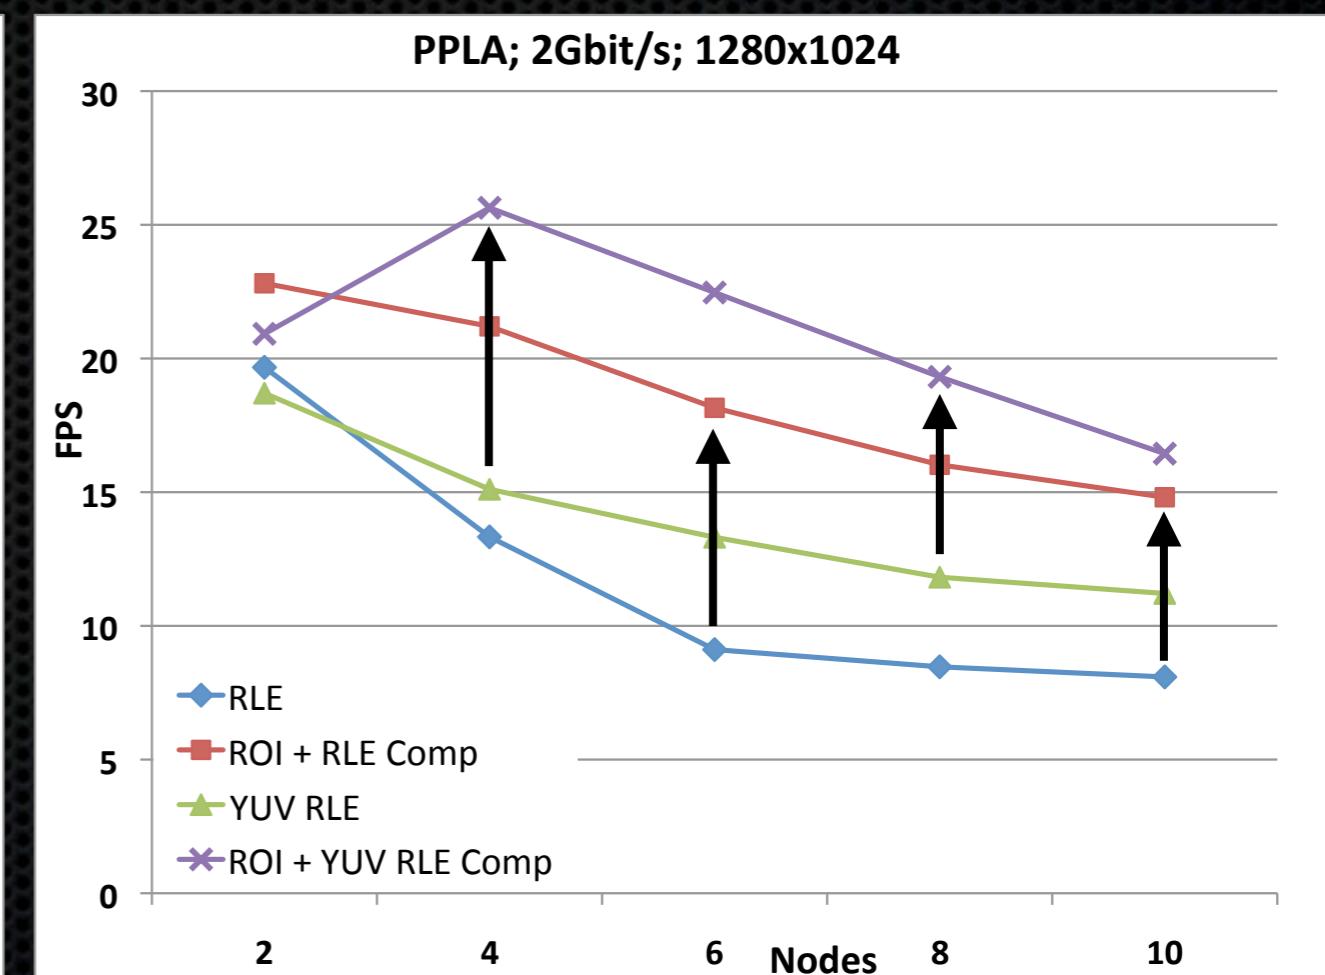


Direct-Send

# ROI Performance

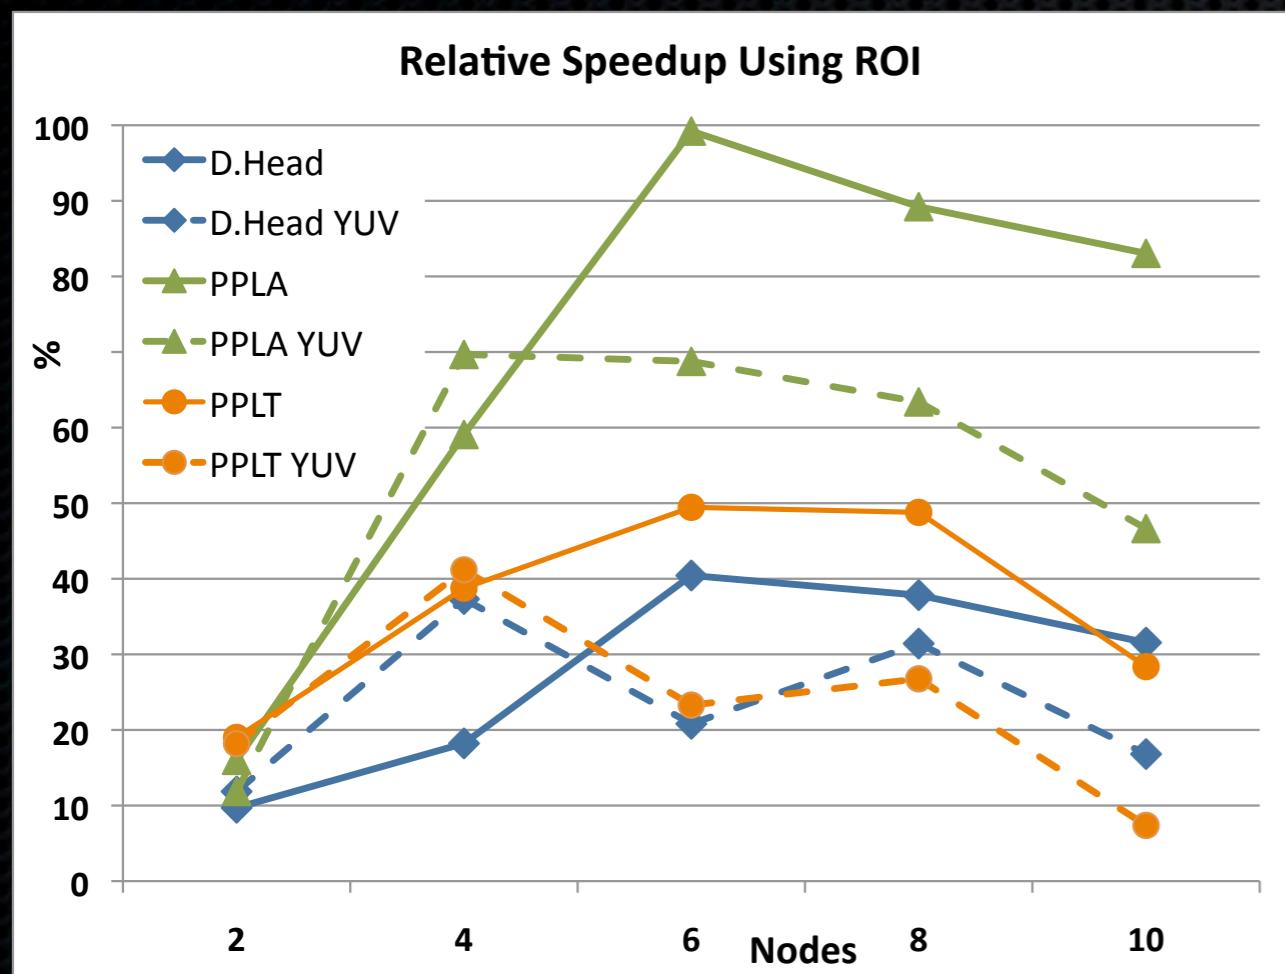


David head

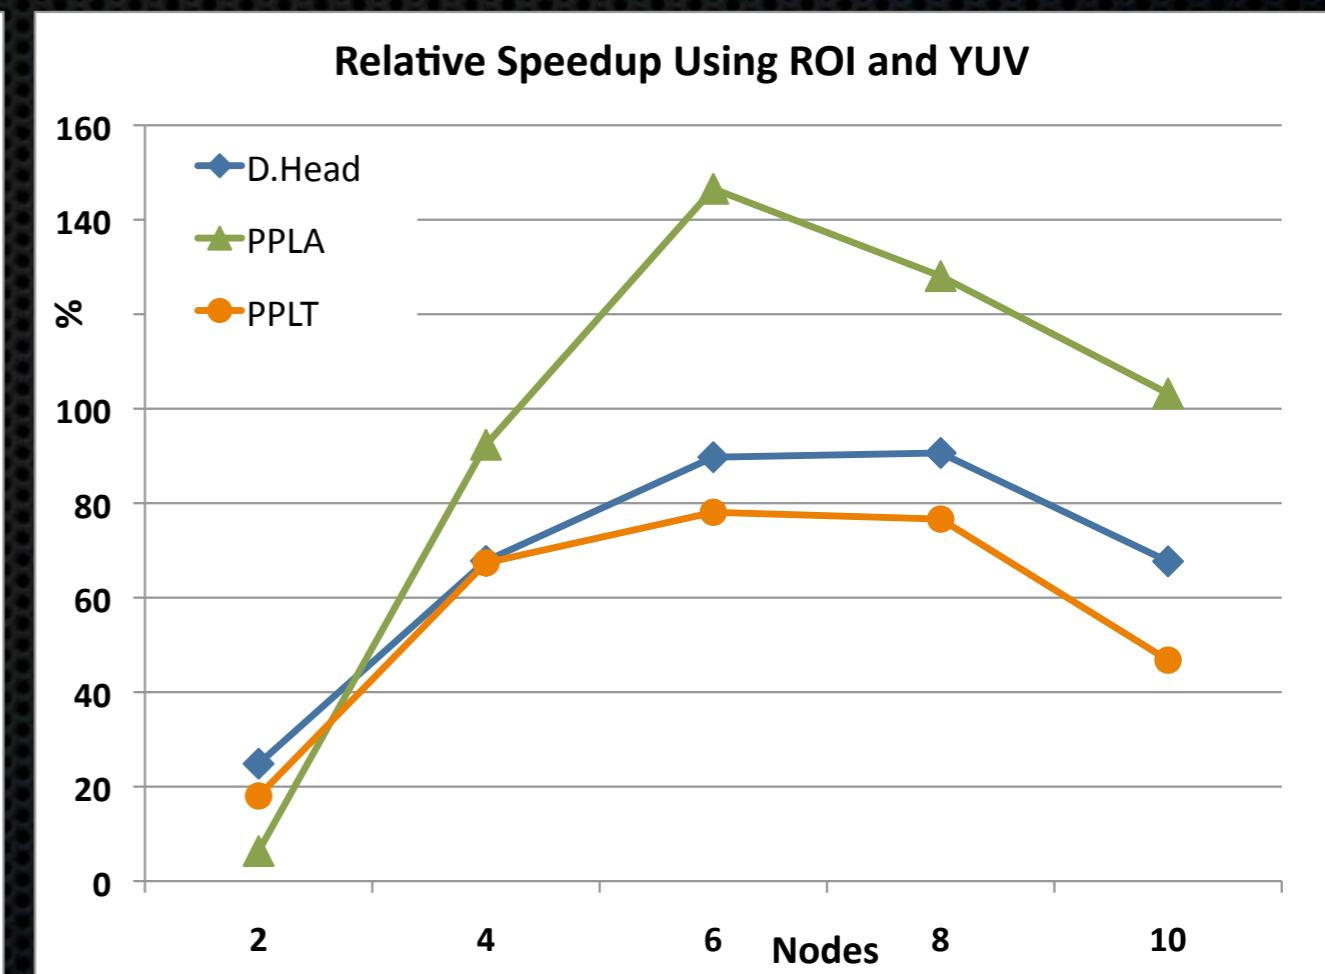


PPLA

# Relative ROI Performance

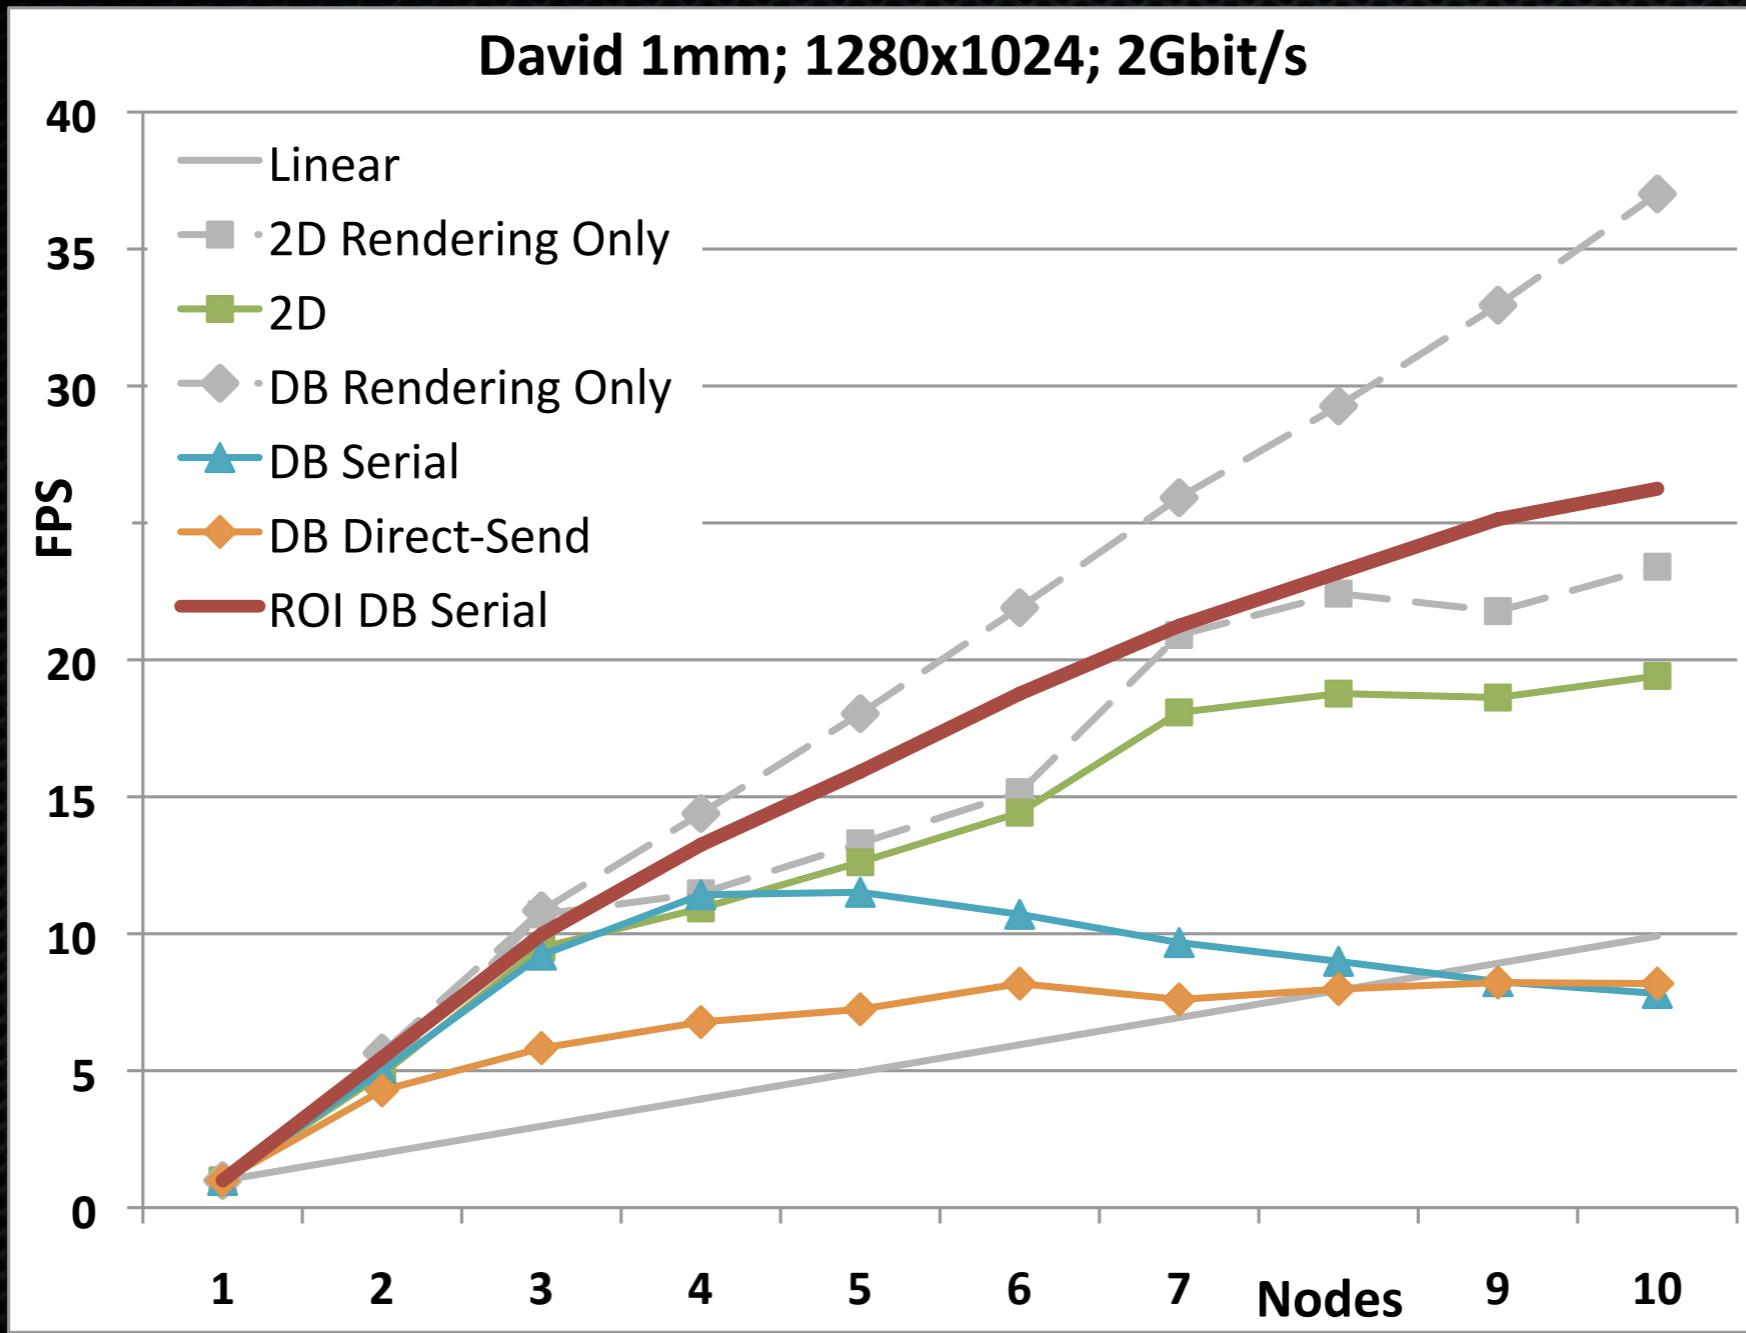


Relative ROI



Relative ROI+YUV

# ROI Best Case



David statue (56M triangles)

# Conclusion

- Several compression techniques are compared in terms of overall rendering performance
- Fast ROI algorithm for sort-last rendering is demonstrated

# Thank You!

## Questions?