Part 2

Mobile graphics trends

- Hardware architectures
- Applications

Hardware architectures

Brief history of mobile graphics hardware

	Apple (*PowerVR*)	**Samsung (*mostly ARM*)**	**Imagination PowerVR**	**ARM Mali**	**Qualcomm Snapdragon/ Adreno**	**AMD**	**Intel**	**Nvidia**
2007	iPhone (MBX)		MBX SGX535/541 (GLES 2.0)	Buys Phalanx				
2008	Buys PA Semi	Hummingbird (Cortex A8) Omnia HD (TI OMAP 3 & Power VR SGX530)	SGX543 (GLES 2.0, GL 2.1)	Mali 400 GLES 2.0	Buys Imageon (Adreno)	Sells Imageon		
2009	iPhone 3GS (SGX535)		SGX545 (GLES 2.0 GL 3.2)					
2010	A4 (ARM Cortex A8)							Tegra 2 (Cortex-A9, GLES 2.0)
2011								Tegra 3 (Cortex-A9, GLES 2.0)
2012			Series 6XE/XT GLES 3.1 GL 3.2 (28nm)	T600 GLES 2.0, DX9.0	Adreno 530 GLES 3.1+, OpenCL 2, DX 11.2 Vulcan 1.0			
2013	A7/A8 & A8X (GT64XX)							Tegra 4 (Cortex-A15, GL 4.4, 28nm)
2014	A9 (GT7600)	Exynos 5433/7410 (20nm, Mali-T760 MP6)	Series7XE Vulkan 1.0 GLES 3.1 (latter ones 10nm)	T700 GLES 3,1, DX 11.1 OpenCL 1.1				Tegra K1 (Cortex-A15, GL 4.4, 28nm)
2015				T800 GLES 3,1, DX 11.1-11.2 OpenCL 1.2				Tegra X1 (Cortex-A57, GLES 3.1, GL 4.5, Vulkan, 20nm)
2016	Plans to build its own GPU		Apple will no longer require its services in 18-24 months Furian?					Next Tegra generations seem to be for automotive

Architectures (beginning 2015)

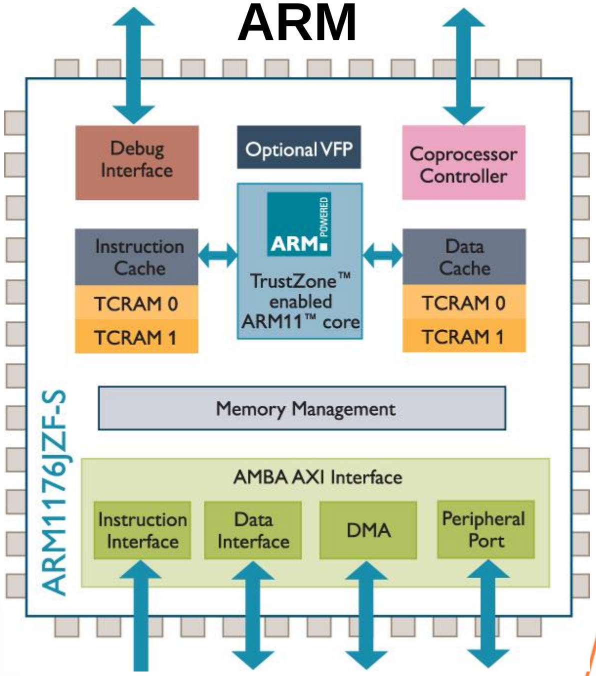
ARM

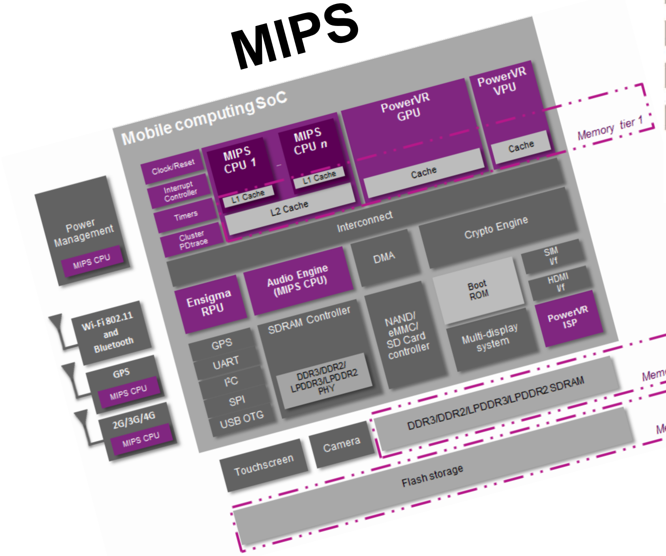
MIPS

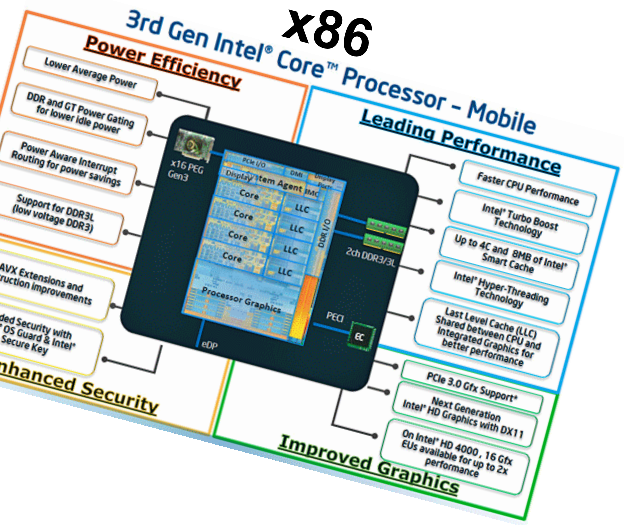
x86

Architectures

- **x86 (CISC 32/64bit)**
  - Intel Atom Z3740/Z3770, X3/X5/X7
  - AMD Amur / Styx (announced)
  - Present in few smartphones, more common in tablets
  - Less efficient
- **ARM**
  - RISC 32/64bit
    - With SIMD add-ons
  - Most common chip for smartphones
  - More efficient & smaller area
- **MIPS**
  - RISC 32/64bit
  - Including some SIMD instructions
  - Acquired by Imagination, Inc. @2014

Architectures – RISC vs. CISC but…

- **CISC (Complex Instruction Set Computer)**
  - Fast program execution (optimized complex paths)
  - Complex instructions (i.e. memory-to-memory instructions)

- **RISC (Reduced Instruction Set Computer)**
  - Fast instructions (fixed cycles per instruction)
  - Simple instructions (fixed/reduced cost per instruction)

- **FISC (Fast Instruction Set Computer)**
  - Current RISC processors integrate many improvements from CISC: superscalar, branch prediction, SIMD, **out-of-order**
  - Philosophy → fixed/reduced cycle count/instr
  - Discussion (Post-RISC):
    - http://archive.arstechnica.com/cpu/4q99/risc-cisc/rvc-5.html

Landscape has changed a bit…

- **Status by 2014-2015:**
  - Intel Atom X3/X5/X7 announced (March 2015)
  - AMD announces Amur / Styx (20nm, Oct. 2014)
  - Nvidia launches Tegra X1 (March 2015)
  - ARM the only EU big technology company
  - Imagination announces Furian (sub 14nm, March 2017) Imagination’s chips are in iPhones & iPads

- **Nowadays:**
  - Intel quits mobile Apr/May 2016
  - AMD cancels 20nm chips (Jul. 2015)
  - Nvidia cancels Shield tablet (Aug. 2016)
  - ARM acquired by Softbank (Sep. 2016)
  - Apple tells Imagination that their IP will not be needed in 18-24 months (Apr. 2017)

Architectures (nowadays)

ARM

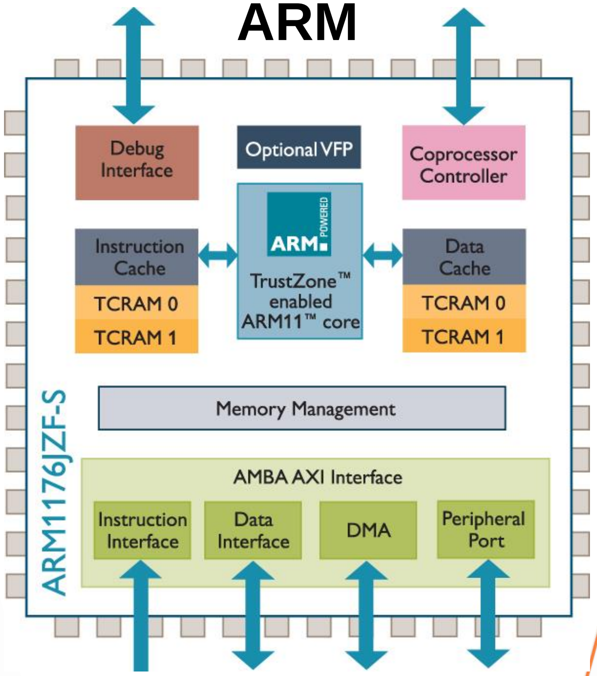

MIPS

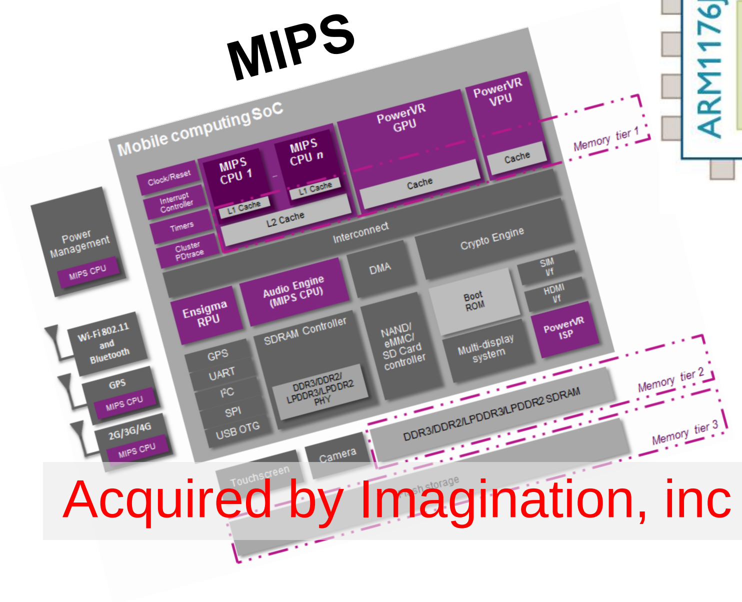

Acquired by Imagination, inc

x86

Architectures – ARM

- **ARM Ltd.**
  - RISC processor (32/64 bit)
  - IP (intellectual property) – Instruction Set / ref. implementation
  - CPU / GPU (Mali)
- **Licenses (instruction set OR ref. design)**
  - **Instruction Set** license -> custom made design (SnapDragon, Samsung in Galaxys, Apple in iPones & iPads)
    - Optimizations (particular paths, improved core freq. control,…)
  - **Reference design** (Cortex A9, Cortex A15, Cortex A53/A57…)
- **Licensees (instruction set OR ref. design)**
  - Apple, Qualcomm, Samsung, Nvidia, AMD, MediaTek, Amazon (through Annapurna Labs, Inc.)…
  - Few IS licenses, mostly adopting reference design
- **Manufacturers**
  - Contracted by Licensees
    - GlobalFoundries, United Microelectronics, TSM…

Architectures – ARM…

- **Suppported on**
  - Android, iOS, Win Phone, Tizen, Firefox OS, BlackBerry, Ubuntu Phone, …
- **Biggest mobile market share**
- **Typically paired with mobile GPUs. Existing brands:**
  - Adreno 4x0/5x0 – Qualcomm
  - PowerVR 8XE (Rogue) – Imagination
  - Mali T8x0/G51/G71 – ARM
- **General strategies:**
  - Cache coherence – week sequential code guarantees on multithreading!!
  - Heavy **dependence on compiler** → optimize instruction scheduling
    - Operation dependencies , loop unrolling, etc…
  - Use SIMD extensions

Architecture types

- **High performance**
  - Premium smartphones & tablets
- **High area efficiency**
  - Medium-to-low smartphones
- **Ultra-low power**
  - Smartwatches

Architectures

Mobile GPU architecture trends

Graphics pipeline trends

- **Tiled rendering**
- **Data (texture) compression**
- **Other optimizations**

Tiled Rendering

- **Immediate Mode Rendering (IMR)**
- **Tile-Based Rendering (TBR)**
- **Tile-Based Deferred Rendering (TBDR)**

Architectures – GPU

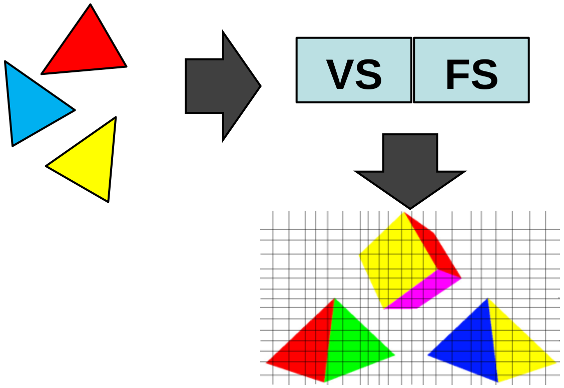

- **Inmediate Mode Rendering (IMR)**
  - Geometry is processed in submission order
    - High **overdraw** (shaded pixels can be overwritten)
  - Buffers are kept in System Memory
    - High bandwidth / power / latency
  - Early-Z helps depending on geometry sorting
    - Depth buffer value closer than fragment → discard

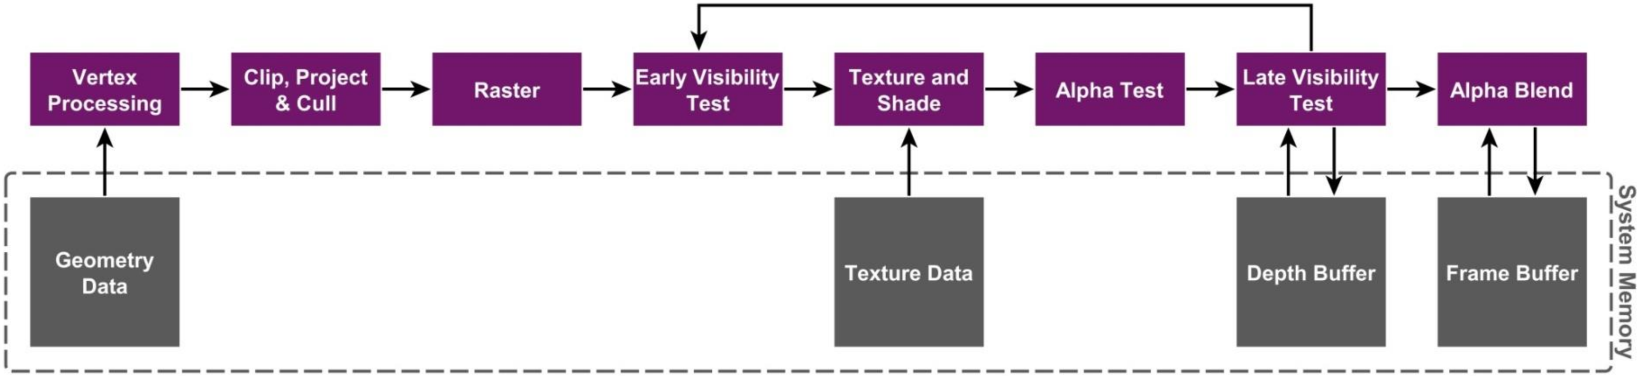

http://blog.imgtec.com/powervr/understanding-powervr-series5xt-powervr-tbdr-and-architecture-efficiency-part-4

Architectures – GPU

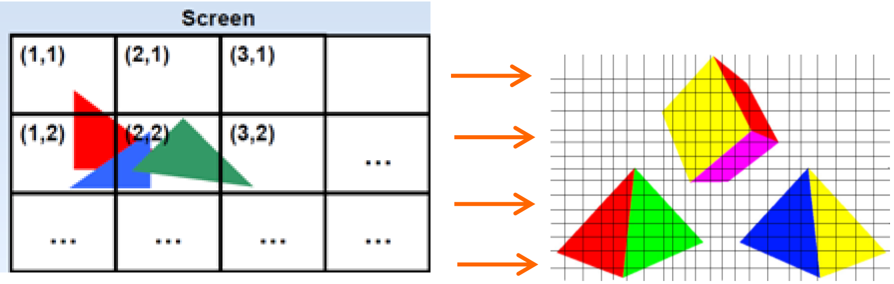

- **Tile Based Rendering (TBR)**
  - Rasterizing per-tile (triangles in bins per tile) 16x16, 32x32
    - Buffers are kept on-chip memory (GPU) – fast! → **geometry limit?**
  - Triangles processed in submission order (TB-IMR)
    - **Overdraw (front-to-back -> early z cull)**
  - Early-Z helps depending on geometry sorting

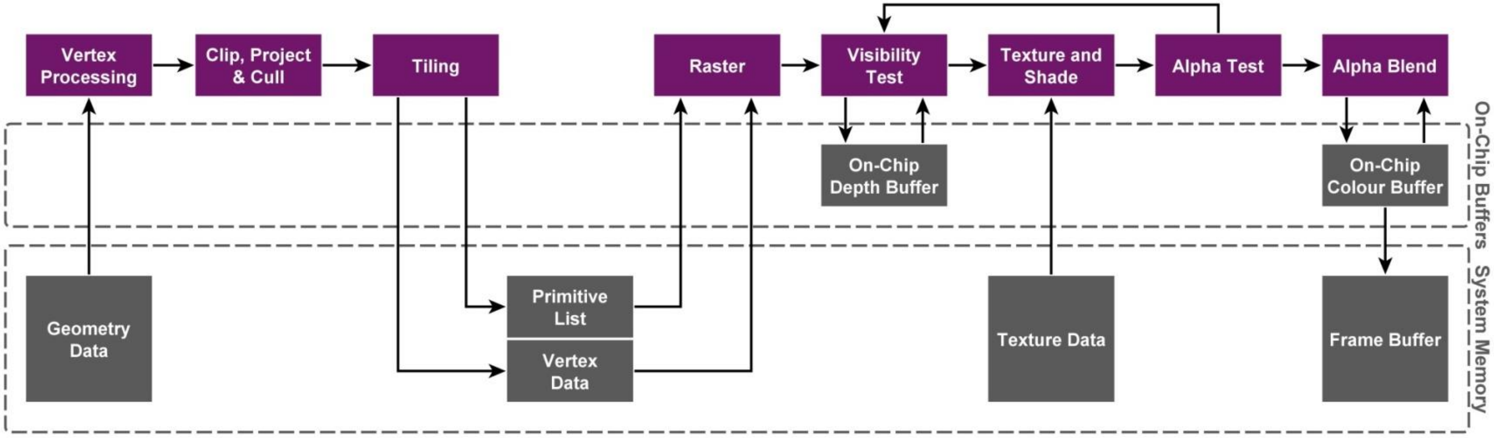

http://blog.imgtec.com/powervr/understanding-powervr-series5xt-powervr-tbdr-and-architecture-efficiency-part-4

Architectures – GPU

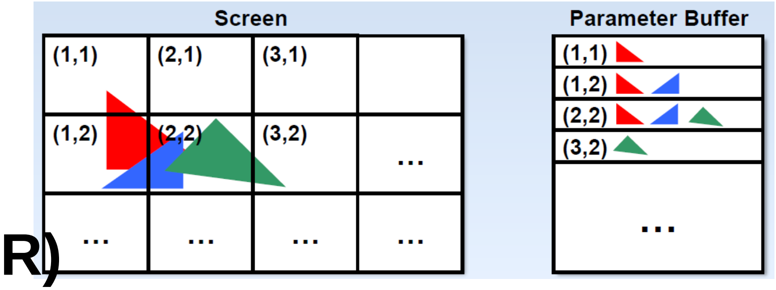

- **Tile Based Deferred Rendering (TBDR)**
  - Fragment processing (tex + shade) ~waits for Hidden Surface Removal
    - Micro Depth Buffer – depth test before fragment submission
      - whole tile → **1 frag/pixel** ☺ → **Limit: ~100Ktri + complex shader**
    - iPAD 2X slower than Desktop GeForce at HSR (FastMobileShaders_siggraph2011)
  - Possible to prefetch textures before shading/texturing
  - Hard to profile!!! ~~~Timing?

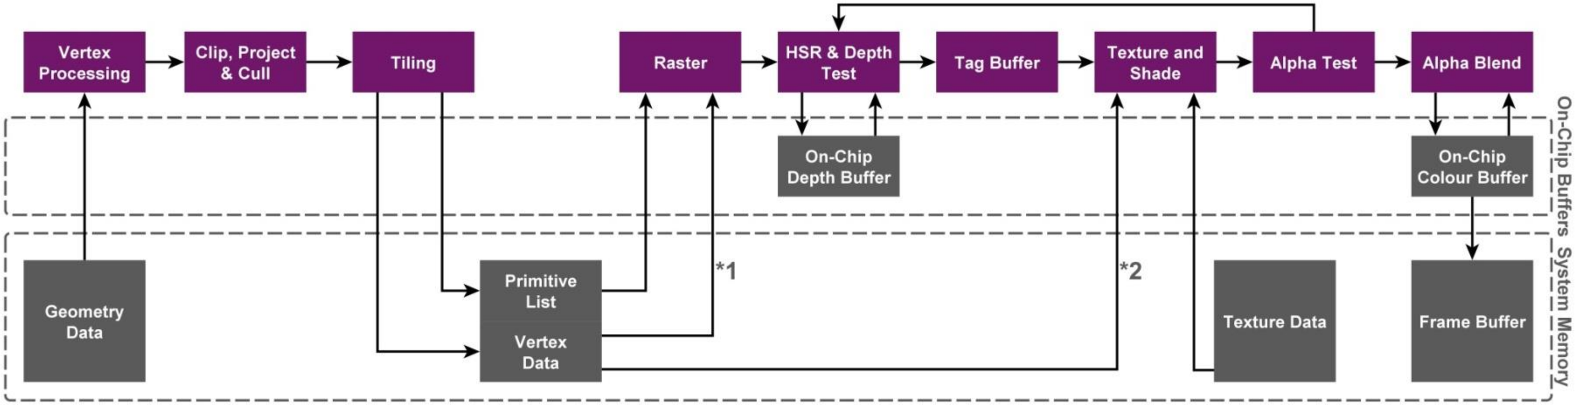

http://blog.imgtec.com/powervr/understanding-powervr-series5xt-powervr-tbdr-and-architecture-efficiency-part-4

Data/texture compression

- ARM's Adaptive Scalable Texture Compression (ASTC) supported by most mobile GPU vendors
- ETC2/EAC standard compression OpenGL ES 3.0
- Compression hardware also present in display hardware
  - – Rendered images stored and transferred to the display in a compressed
    - Saving bandwidth

Other optimizations

- Deferred shading
- Primitive elimination
- Skipping updates to pixels that do not change
  - ARM memory transaction elimination

Trends

- Specific hardware for ray tracing
- Learning libraries & hardware (e.g. Qualcomm's Fast CV, Nvidia's CUDA Deep Neural Network)
- Skipping updates to pixels that do not change
  - ARM memory transaction elimination

Applications

Applications

- **Wide range of applications**
  - Cultural Heritage
  - Medical Image
  - 3D object registration
  - GIS
  - Gaming
  - VR & AR
  - Building reconstruction
  - Virtual HCI

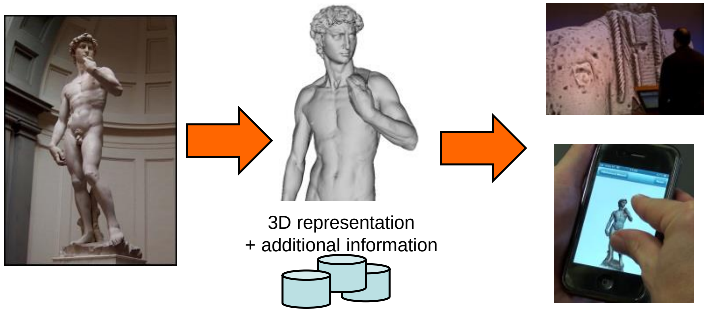

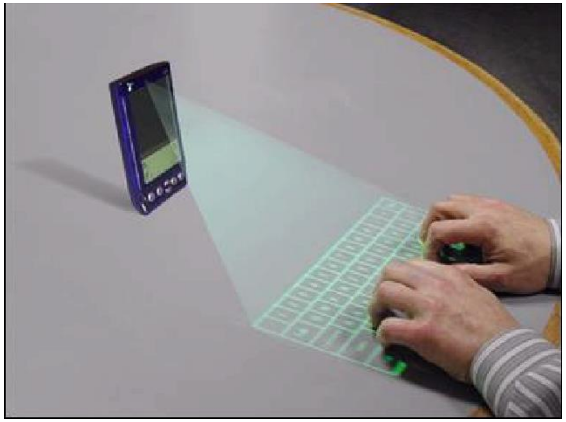

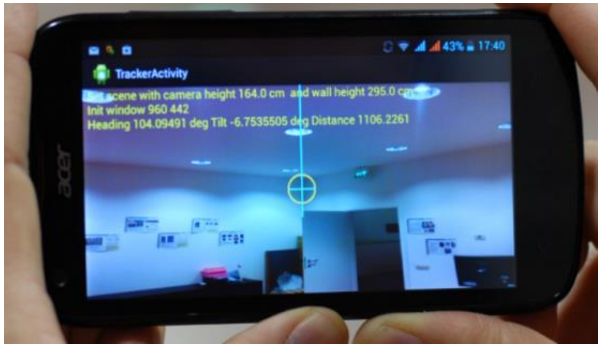

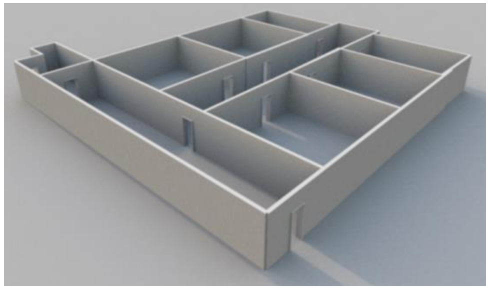

Mobile 3D interactive graphics

- **General pipeline similar to standard interactive applications**

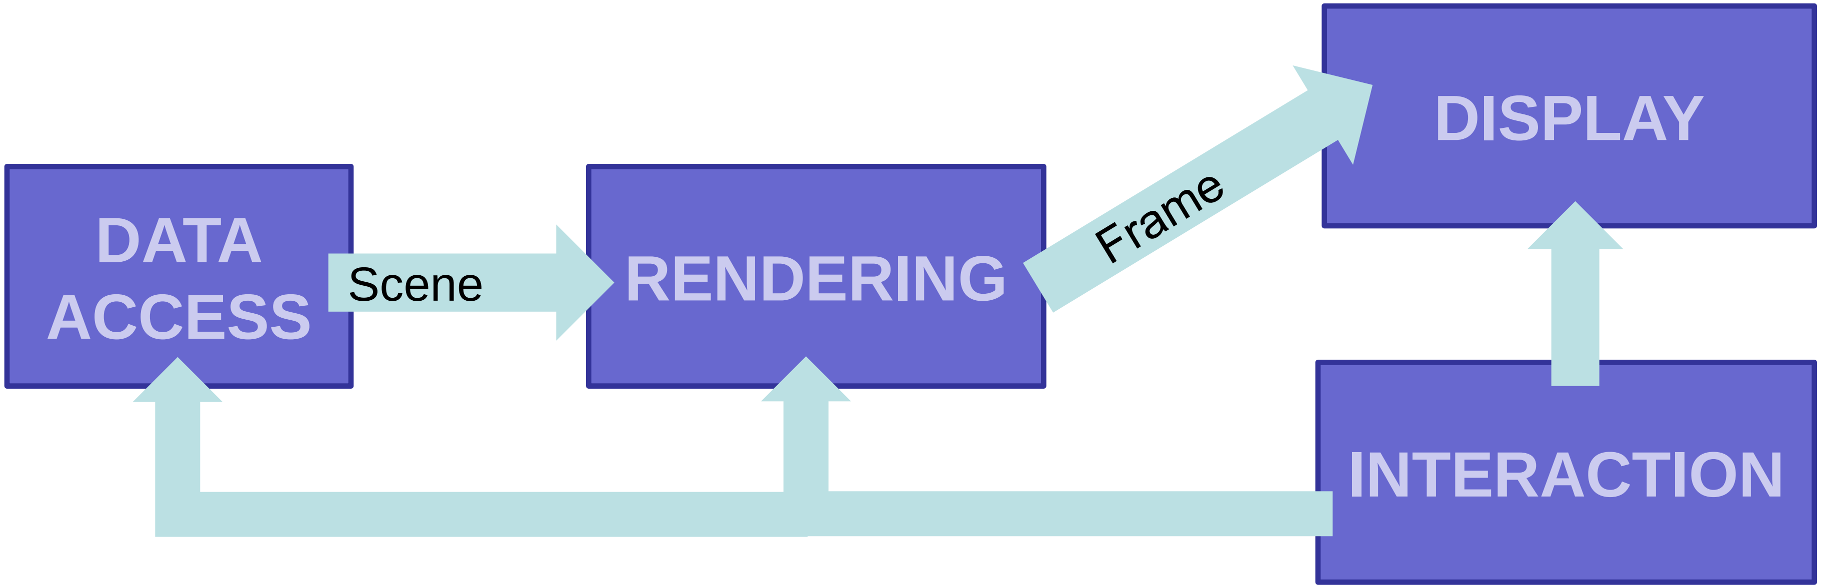

Remote rendering

- **General solution since first PDAs**

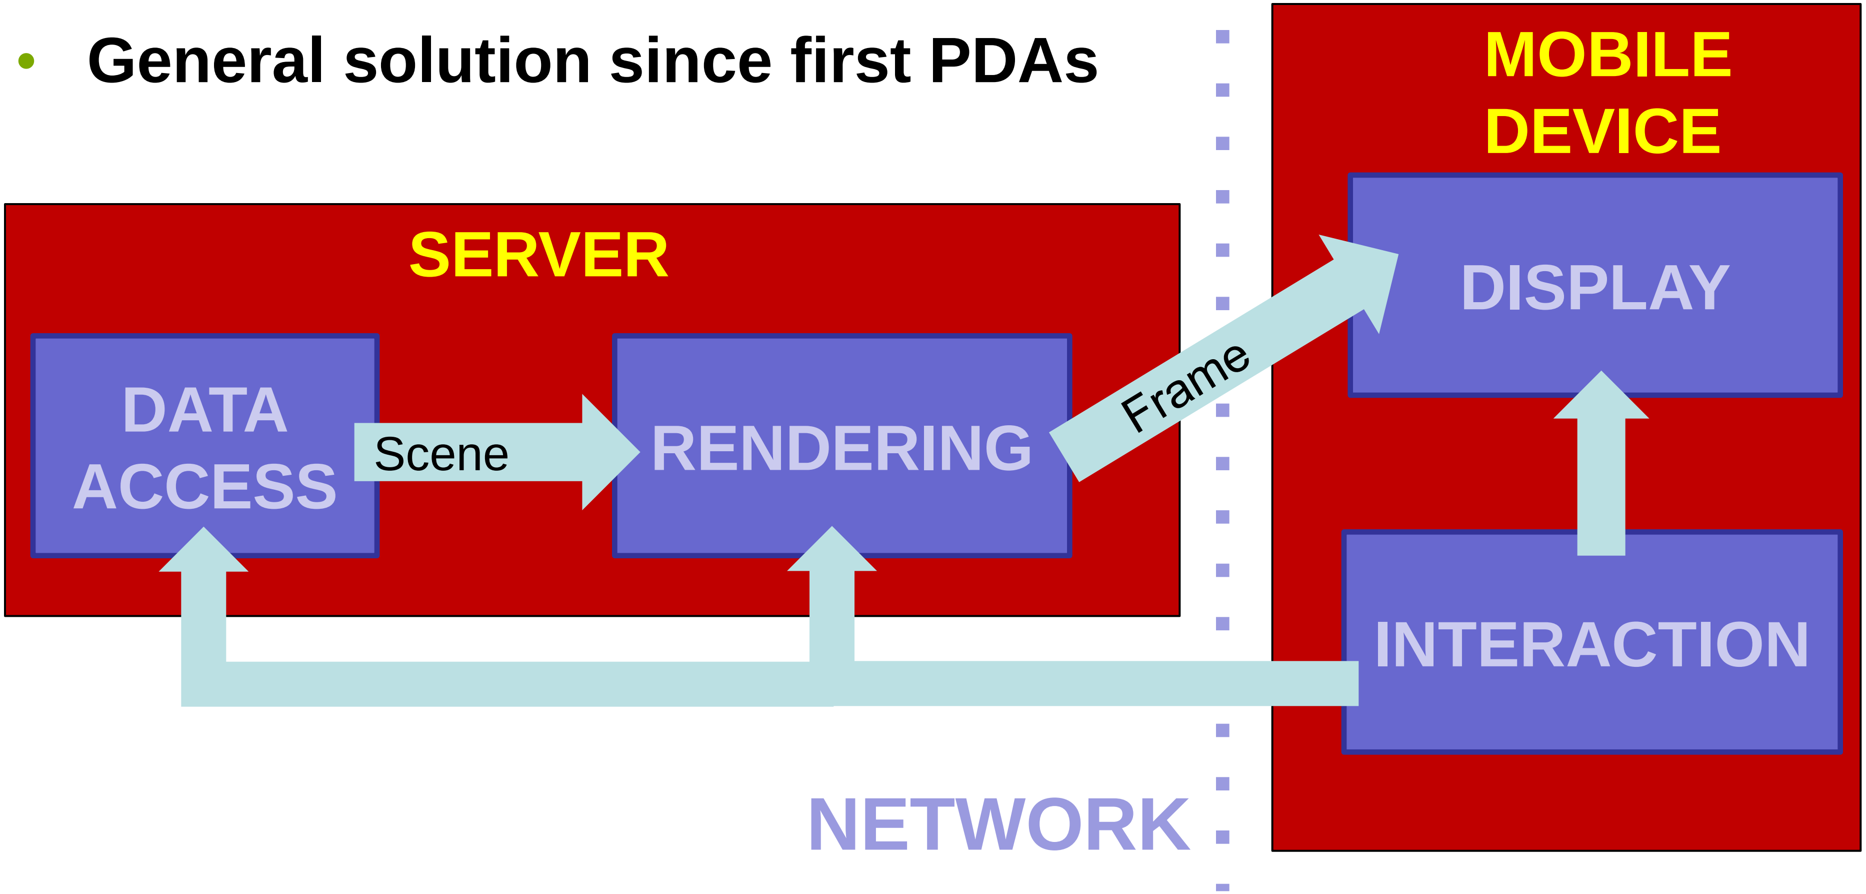

Remote rendering

- **3D graphics applications require intensive computation and network bandwidth**
  - electronic games
  - visualization of very complex 3D scenes
- **Remote rendering has long history and it is successfully applied for gaming services**
  - Limitation: interaction latency in cellular networks

Mixed Mobile/Remote rendering

- **As mobile GPUs progress...**

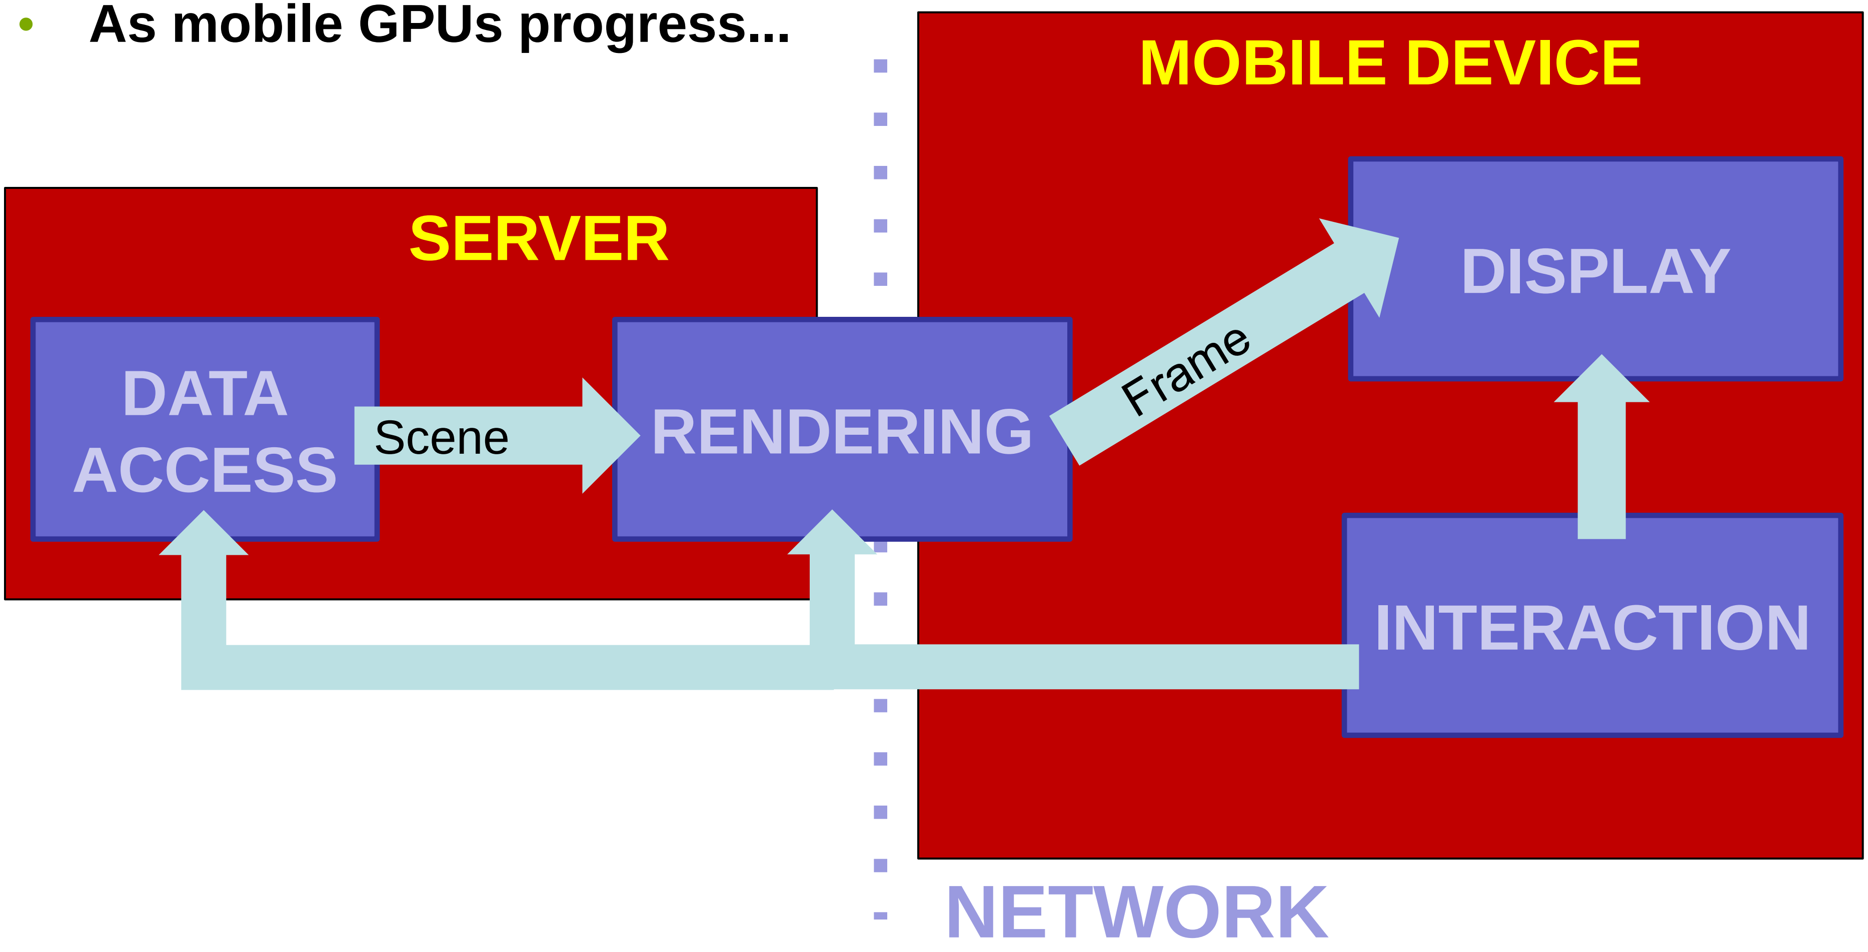

Mixed Mobile/Remote rendering

- **Model based versus Image based methods**
- **Model based methods**
  - Original models
  - Partial models
  - Simplified models
    - Couple of lines
    - Point clouds

Eisert and Fechteler. **Low delay streaming of computer graphics** (ICIP 2008)

Gobbetti et al. **Adaptive Quad Patches: an Adaptive Regular Structure for Web Distribution and Adaptive Rendering of 3D Models.** (Web3D 2012)

Balsa et al.,. **Compression-domain Seamless Multiresolution Visualization of Gigantic Meshes on Mobile Devices** (Web3D 2013)

Diepstraten et al., 2004. **Remote Line Rendering for Mobile Devices** (CGI 2004)

Duguet and Drettakis. **Flexible point-based rendering on mobile devices** (IEEE Trans. on CG & Appl, 2004)

Mixed Mobile/Remote rendering

- **Model based versus Image based methods**
- **Model based methods**

Point clouds organized as hierarchical grids. Tested on PDAs

- Point clouds

Duguet and Drettakis. **Flexible point-based rendering on mobile devices** (IEEE Trans. on CG & Appl, 2004)

Mixed Mobile/Remote rendering

- **Model based versus Image based methods**
- **Model based methods**

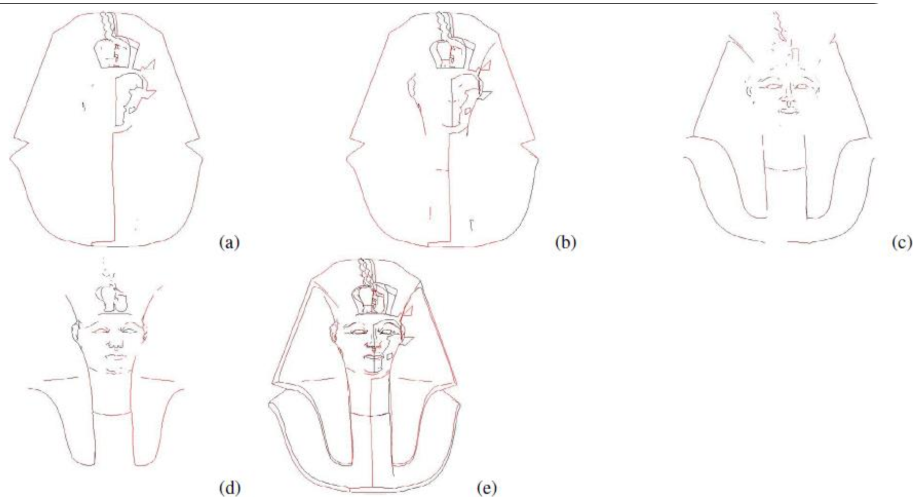

Transfer couple of 2D line primitives over the network, which are rendered locally by the mobile device

- Couple of lines

Diepstraten et al., 2004. **Remote Line Rendering for Mobile Devices** (CGI 2004)

- Point clouds

Duguet and Drettakis. **Flexible point-based rendering on mobile devices** (IEEE Trans. on CG & Appl, 2004)

Mixed Mobile/Remote rendering

- **Model based versus Image based methods**
- **Model based methods**
  - Original models

Eisert and Fechteler. **Low delay streaming of computer graphics** (ICIP 2008)

Intercept and stream OpenGL commands
Better performances with respect to video streaming
Limitation: clients need powerful GPU

Mixed Mobile/Remote rendering

- **Model based versus Image based methods**
- **Model based methods**
  - Original models
  - Partial models
  - Simplified models
    - Couple
    - Point c

Eisert and Fechteler. **Low delay streaming of computer graphics** (ICIP 2008)

Gobbetti et al. **Adaptive Quad Patches: an Adaptive Regular Structure for Web Distribution and Adaptive Rendering of 3D Models.** (Web3D 2012)

Balsa et al.,. **Compression-domain Seamless Multiresolution Visualization of Gigantic Meshes on Mobile Devices** (Web3D 2013)

Mobile

on mobile devices (IEEE Trans. on CG & Appl, 2004)

More details in Part 5

Mixed Mobile/Remote rendering

- **Image based methods**
  - Image impostors

Noimark and Cohen-Or. **Streaming scenes to mpeg-4 video-enabled devices** (IEEE, CG&A 2003)

Lamberti and Sanna. **A streaming-based solution for remote visualization of 3D graphics on mobile devices** (IEEE, Trans. VCG, 2007)

  - Environment maps

Bouquerche and Pazzi. **Remote rendering and streaming of progressive panoramas for mobile devices** (ACM Multimedia 2006)

  - Depth images

Zhu et al. **Towards peer-assisted rendering in networked virtual environments** (ACM Multimedia 2011)

Shi et al. **A Real-Time Remote Rendering System for Interactive Mobile Graphics** (ACM Trans. On Multimedia, 2012)

Doellner et al. **Server-based rendering of large 3D scenes for mobile devices using G-buffer cube maps** (ACM Web3D, 2012)

Mixed Mobile/Remote rendering

- **Image based methods**
  - Image impostors

    Noimark and Cohen-Or. **Streaming scenes to mpeg-4 video-enabled devices** (IEEE, CG&A 2003)

    Lamberti and Sanna. **A streaming-based solution for remote visualization of 3D graphics on mobile devices** (IEEE Trans. VCG, 2007)
  - Envir

    Bouque **mas**
    for mob
  - Depth

    Zhu et al. **Towards peer-assisted rendering in networked virtual environments** (ACM Multimedia 2011)

    Shi et al. **A Real-Time Remote Rendering System for Interactive Mobile Graphics** (ACM Trans. On Multimedia, 2012)

    Doellner et al. **Server-based rendering of large 3D scenes for mobile devices using G-buffer cube maps** (ACM Web3D, 2012)

Image representations are created by the server, and warped in real time by the client to account for user interaction

Mixed Mobile/Remote rendering

- **Model based vs Image based methods**
  - Constraints: rendering quality, bandwidth, interactivity

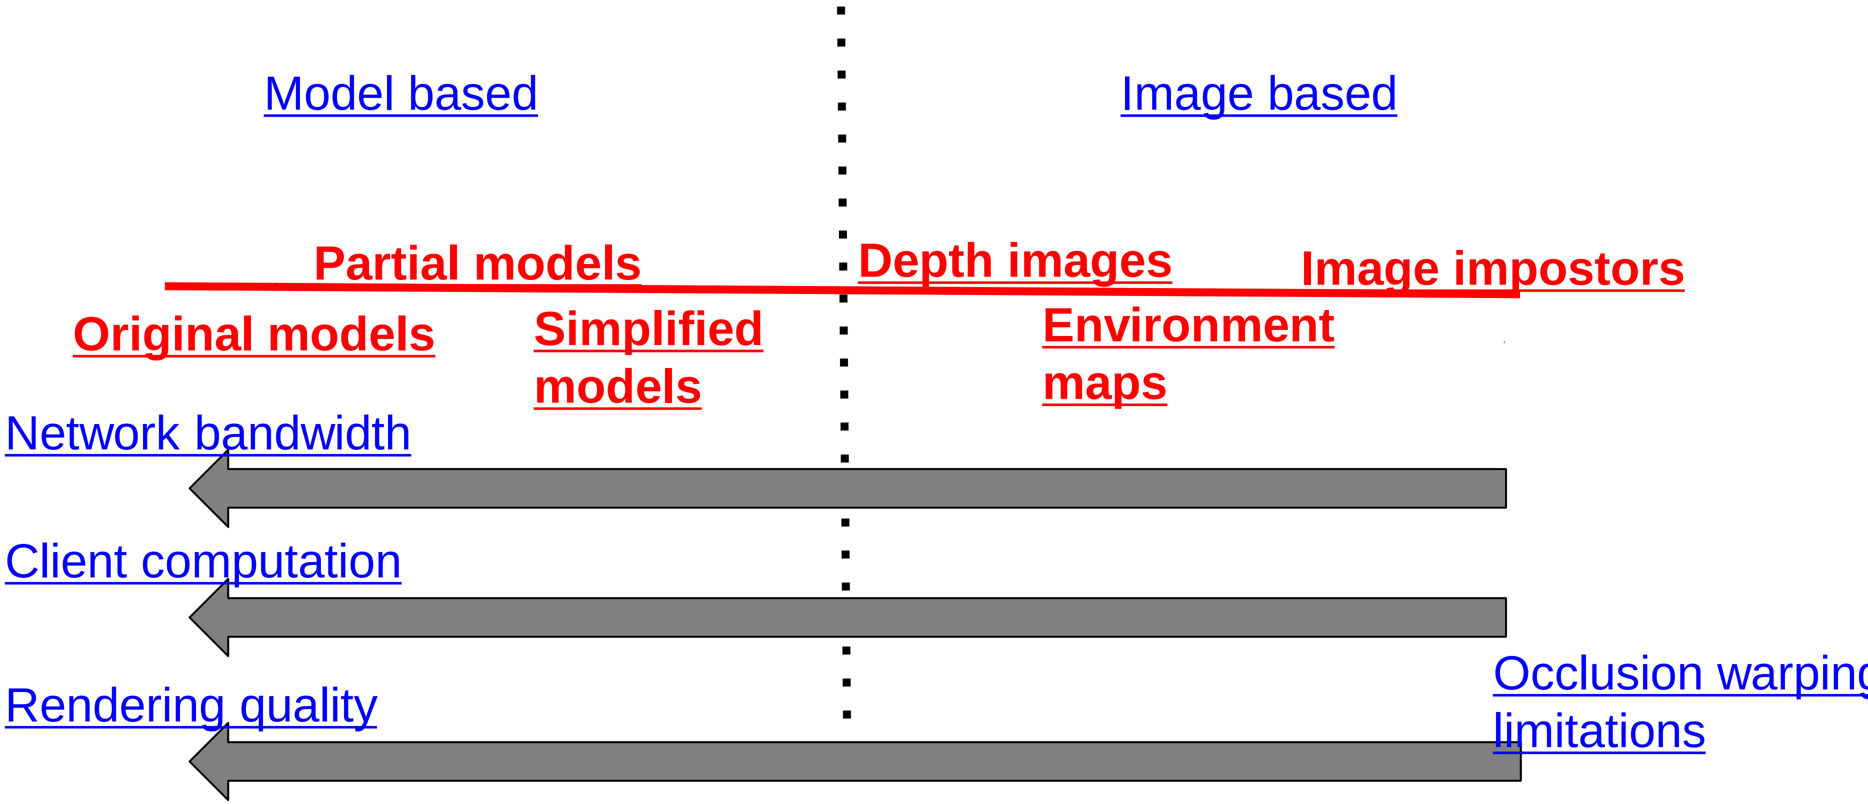

Mobile visualization systems

- **Volume rendering**

Moser and Weiskopf. **Interactive volume rendering on mobile devices**. Vision, Modeling, and Visualization VMV. Vol. 8. 2008.

Noguerat al. **Volume Rendering Strategies on Mobile Devices**. GRAPP/IVAPP. 2012.

Campoalegre, Brunet, and Navazo. **Interactive visualization of medical volume models in mobile devices**. Personal and ubiquitous computing 17.7 (2013): 1503-1514.

Rodríguez, Marcos Balsa, and Pere Pau Vázquez Alcocer. **Practical Volume Rendering in Mobile Devices**. Advances in Visual Computing. Springer, 2012. 708-718.

- **Point cloud rendering**

Balsa et al. **Interactive exploration of gigantic point clouds on mobile devices.** (VAST 2012)

He et al. **A multiresolution object space point-based rendering approach for mobile devices** (AFRIGRAPH, 2007)

Mobile visualization systems

- **Volume rendering**

Moser and Weiskopf. **Interactive volume rendering on mobile devices**. Vision, Modeling, and Visualization VMV. Vol. 8. 2008.

Noguerat al. **Volume Rendering Strategies on Mobile Devices**. GRAPP/IVAPP. 2012.

see section 5 for details

Campoalegre **ualization of medical volume models in mobile devices**. Personal and ubiquitous computing 17.7 (2013): 1503-1514.

Rodríguez, Marcos Balsa, and Pere Pau Vázquez Alcocer. **Practical Volume Rendering in Mobile Devices**. Advances in Visual Computing. Springer, 2012. 708-718.

- **Point cloud rendering**

Balsa et al. **Interactive exploration of gigantic point clouds on mobile devices.** (VAST 2012)

He et al. **A multiresolution object space point-based rendering approach for mobile devices** (AFRIGRAPH, 2007)

Mobile rendering

- **Nowadays...**

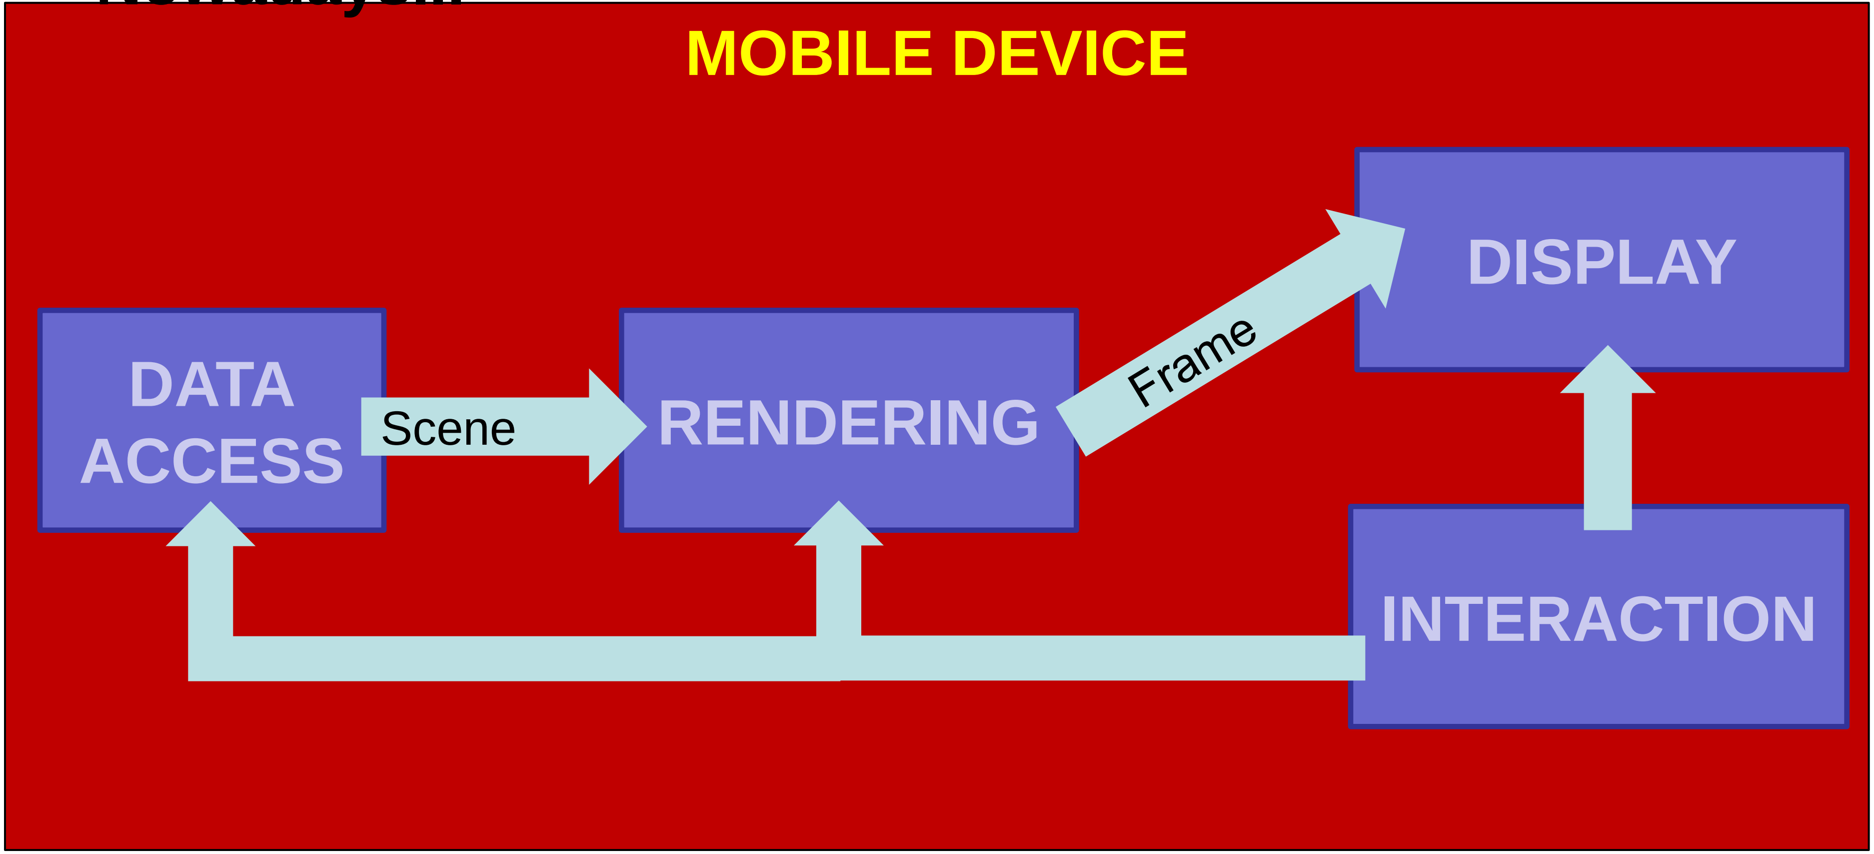

Mobile rendering

- **Or better...**

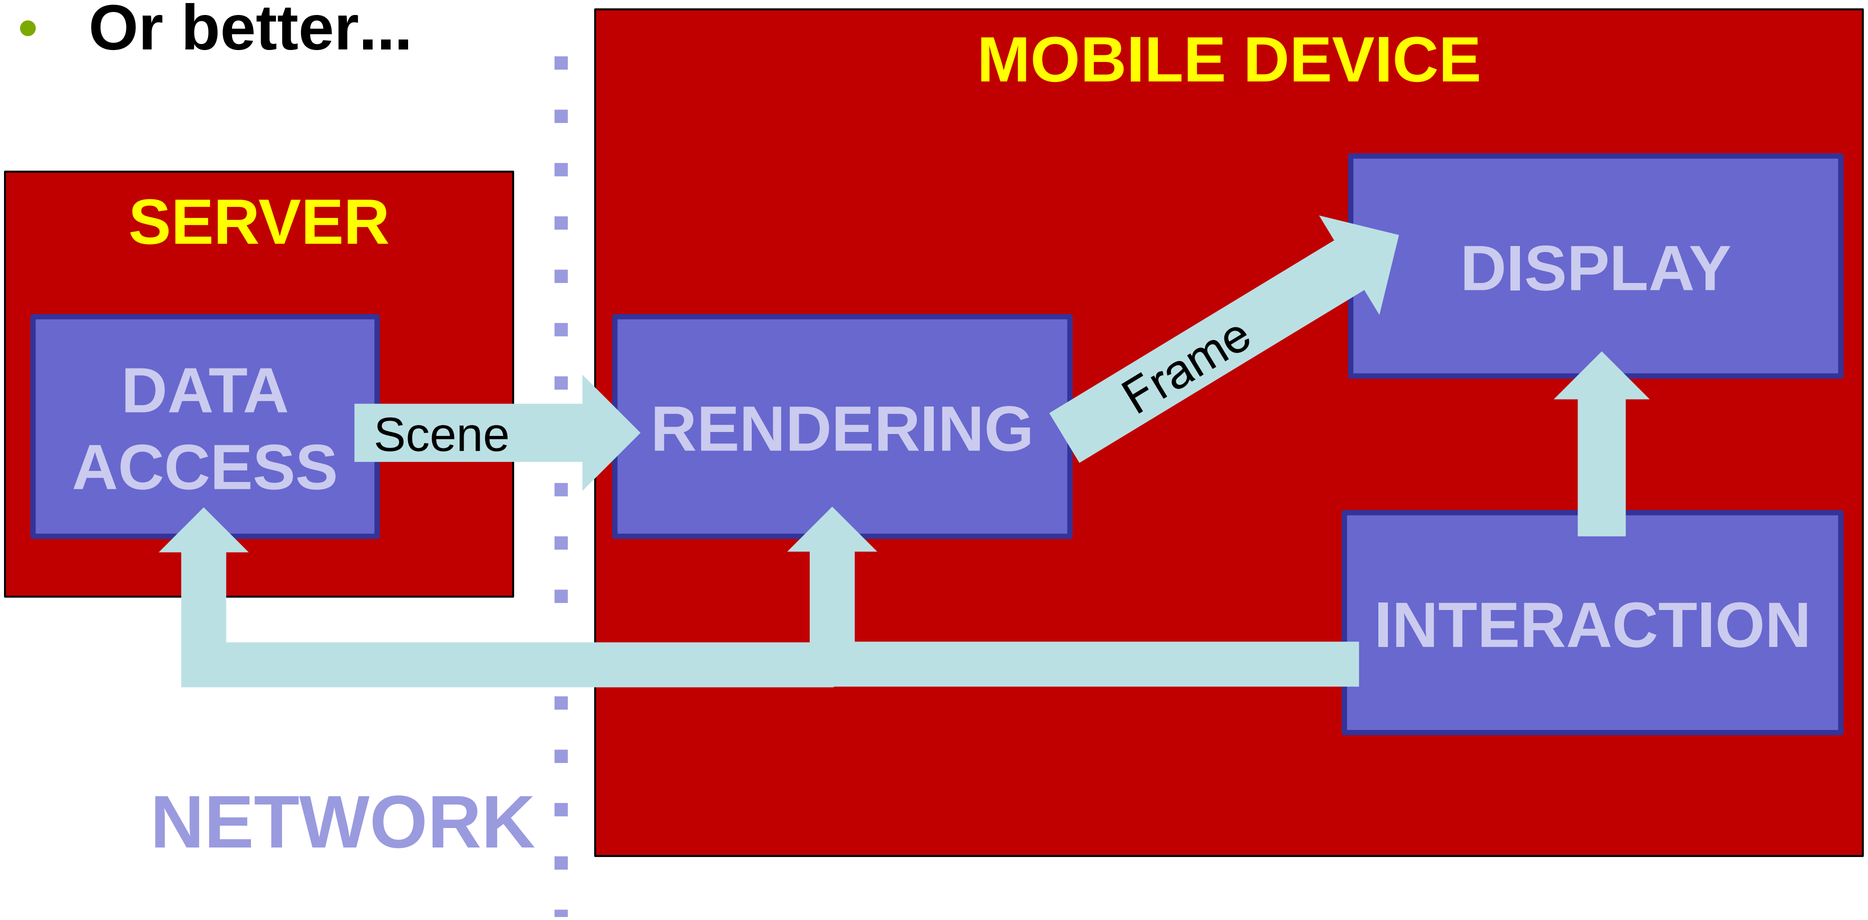

Mobile rendering

- **Or better...**

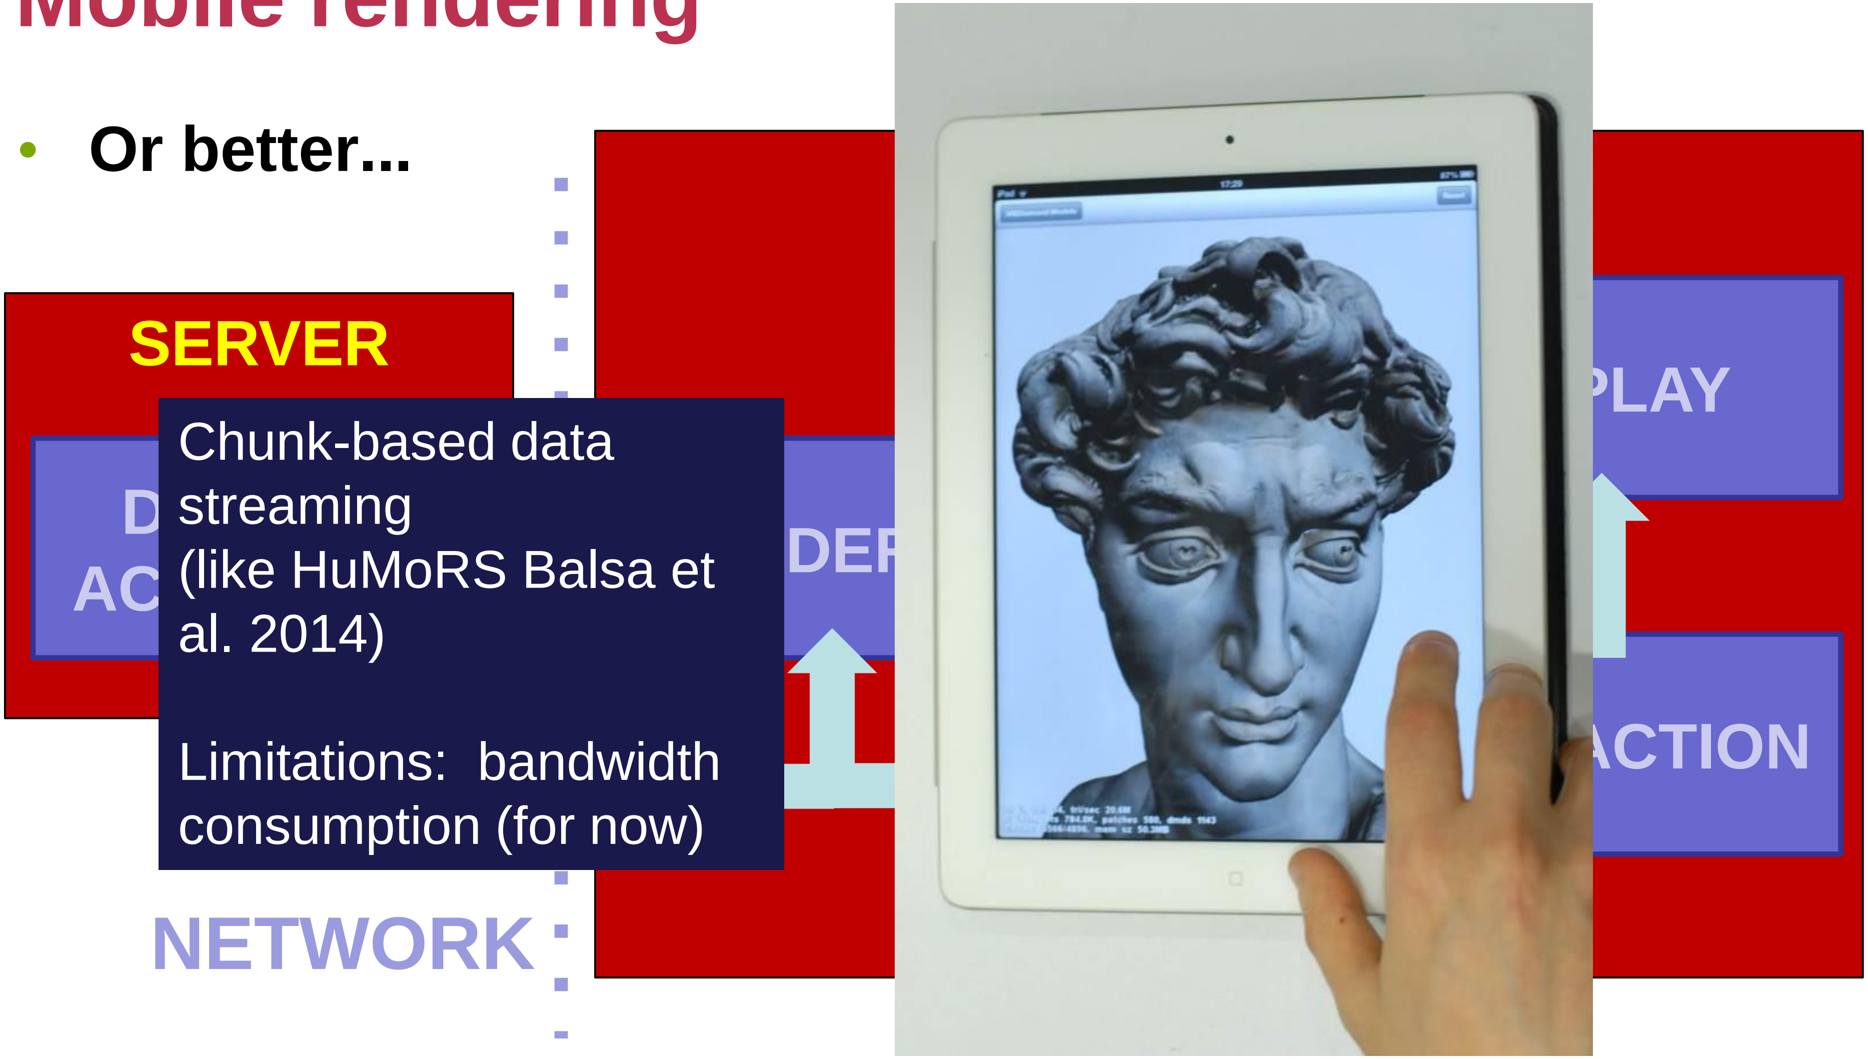

Chunk-based data streaming
(like HuMoRS Balsa et al. 2014)

Limitations: bandwidth consumption (for now)

Mobile rendering with capture

- **Exploiting mobile device sensors...**

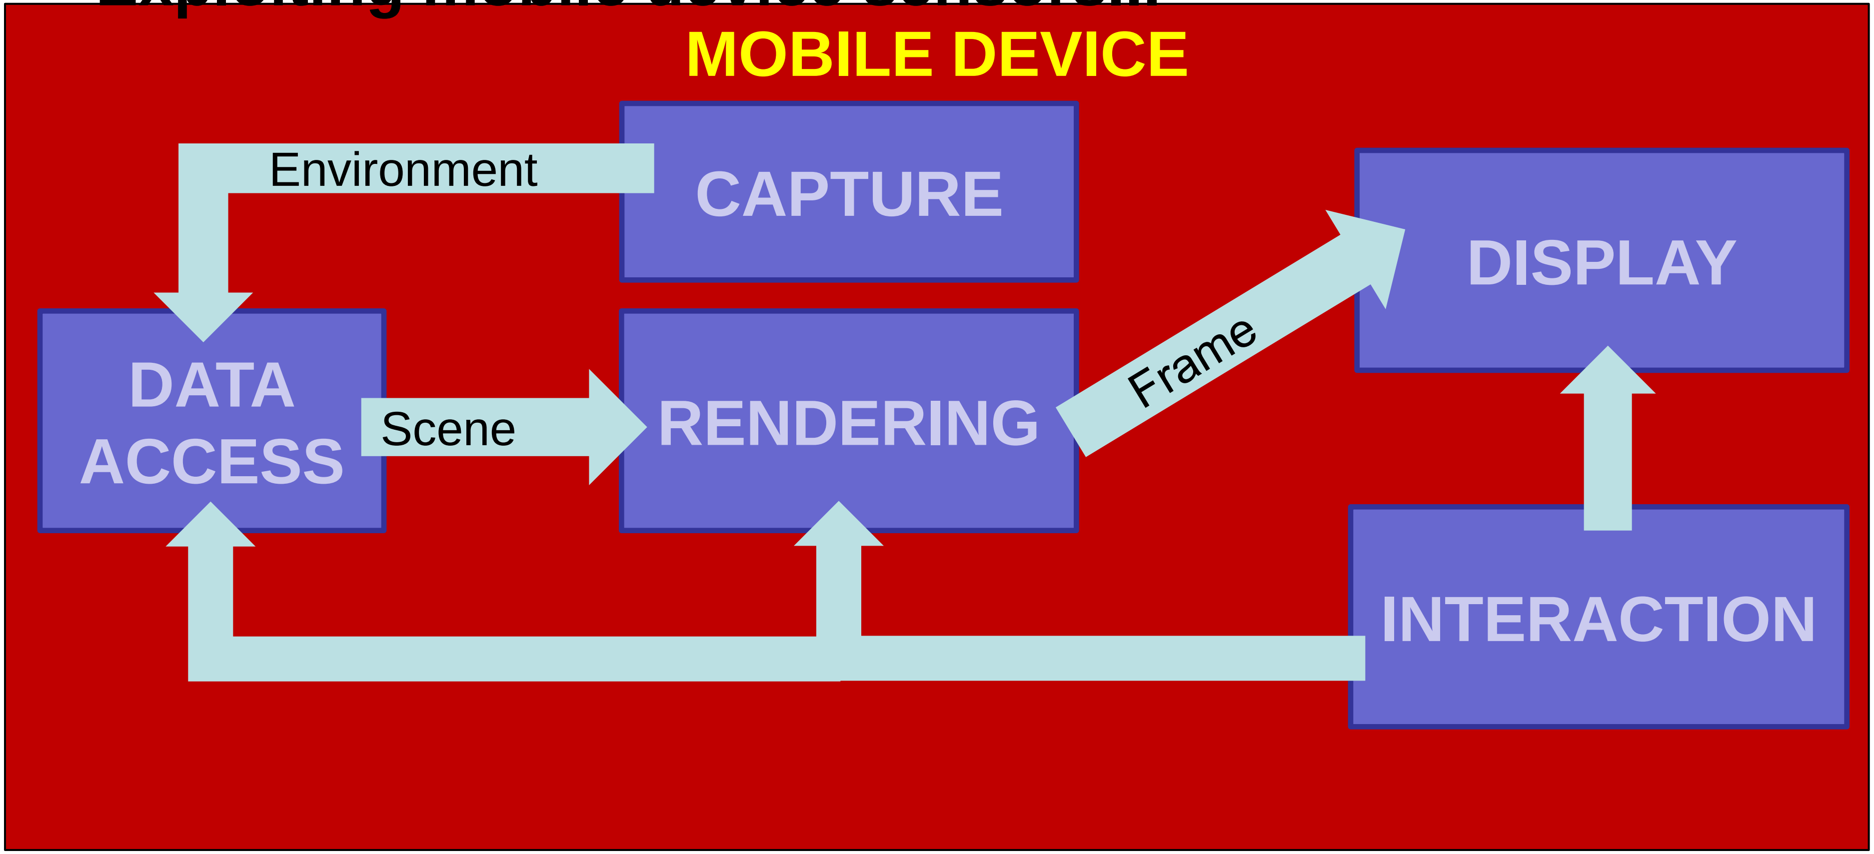

MOBILE DEVICE

CAPTURE

Environment

DATA ACCESS

Scene

RENDERING

Frame

DISPLAY

INTERACTION

Mobile rendering with capture

- **Exploiting mobile device sensors...**

3D scanning with mobile phone
Kolev et al, CVPR 2014
ETH Zurich

Kolev et al. **Turning Mobile Phones into 3D Scanners** (CVPR 2014)

Tanskanen et al. **Live Metric 3D Reconstruction on Mobile Phones** (ICCV 2013)

Mobile rendering with capture

- **Exploiting mobile device sensors...**

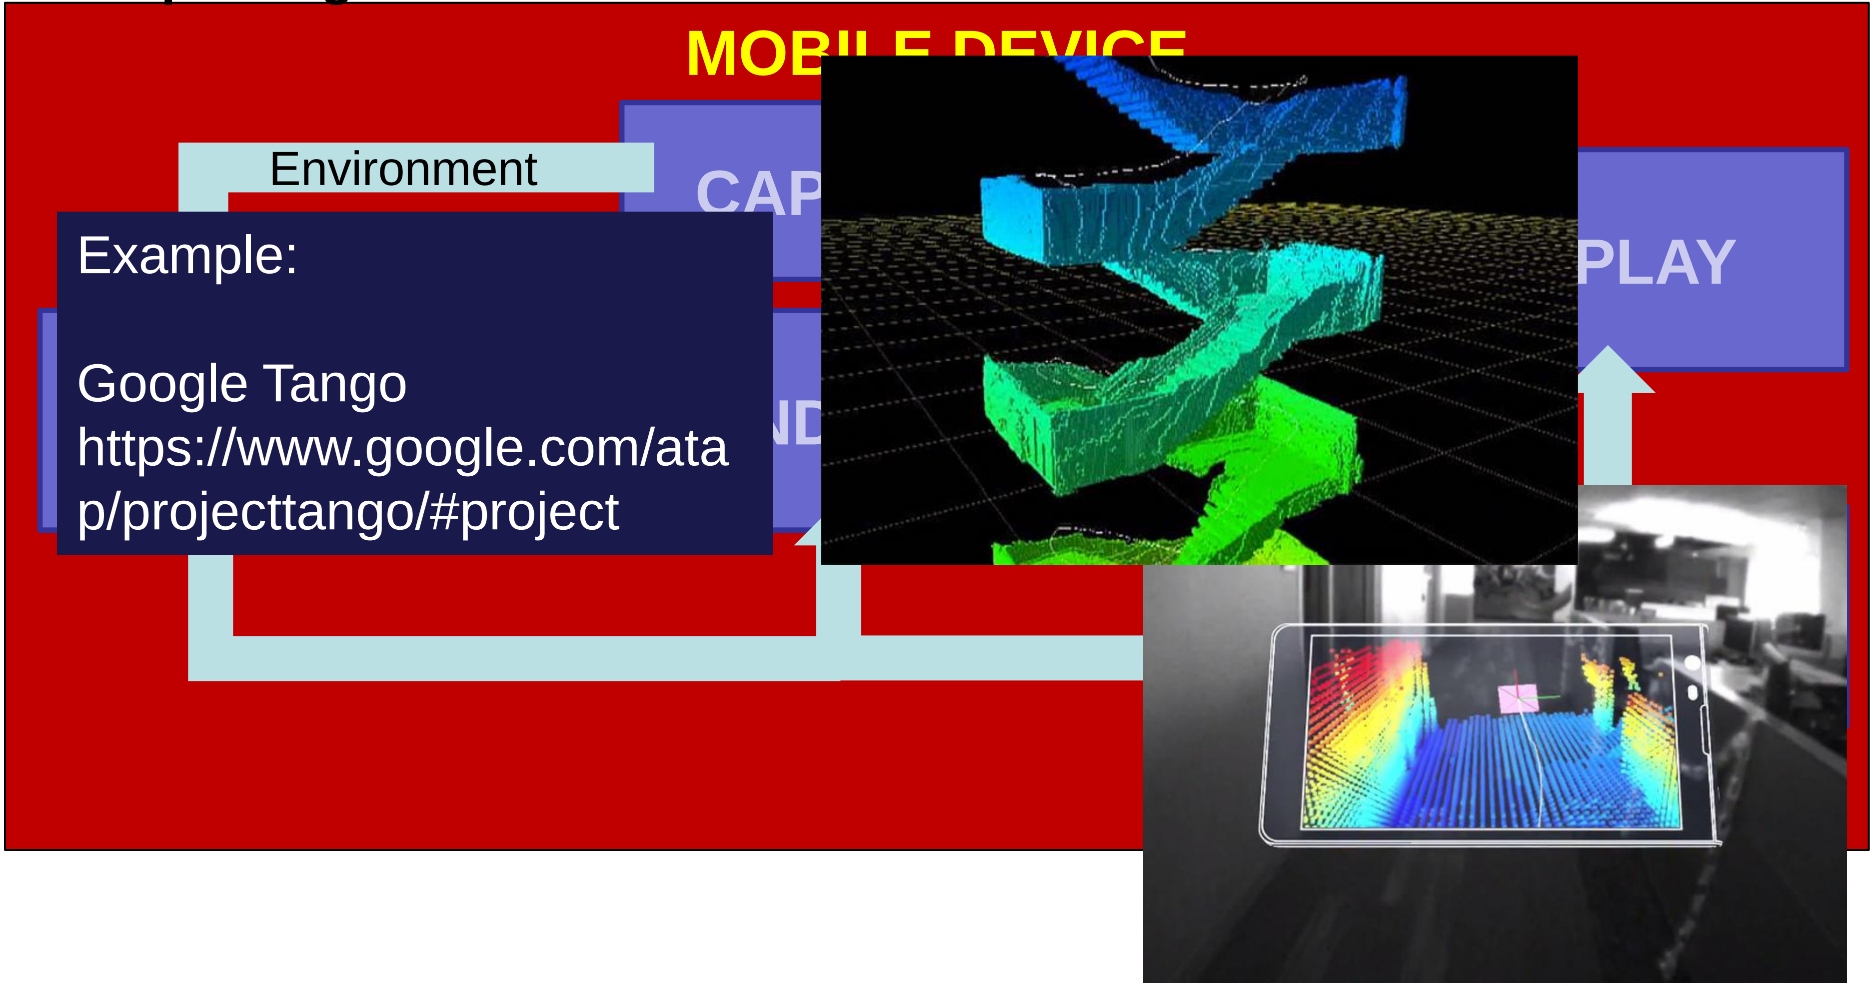

Example:

Google Tango
https://www.google.com/atap/projecttango/#project

Mobile rendering with capture

- **Exploiting mobile device sensors...**

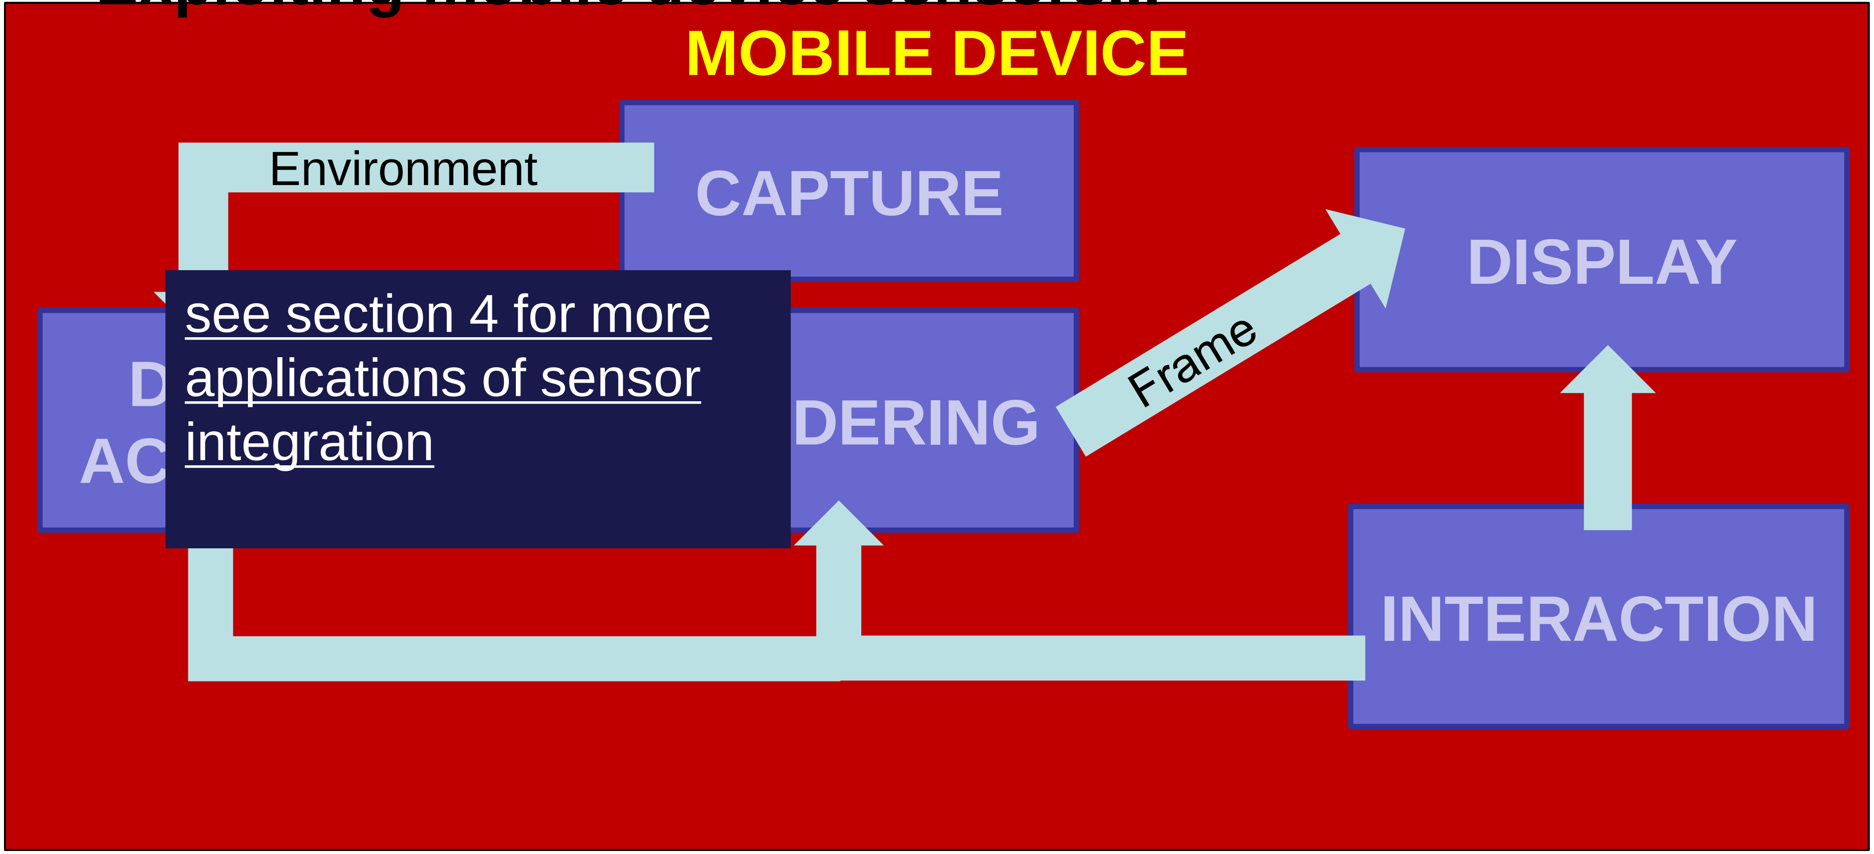

see section 4 for more applications of sensor integration

Trends in mobile graphics

- **Hardware acceleration for improving frame rates, resolutions and rendering quality**
  - Parallel pipelines
  - Real-time ray tracing
  - Multi-rate approaches

SGRT: Real-time ray tracing

- **Samsung reconfigurable GPU based on Ray Tracing**
- **Main key features:**
  - an area-efficient parallel pipelined traversal unit
  - flexible and high-performance kernels for shading and ray generation

Shin et al., **Full-stream architecture for ray tracing with efficient data transmission**, 2014 IEEE ISCAS

Lee, Won-Jong, et al. **SGRT: A mobile GPU architecture for real-time ray tracing**. Proceedings of the 5th High-Performance Graphics Conference, 2013.

Adaptive shading

- Triangles rasterized into coarse fragments that correspond to multiple pixels of coverage
- Coarse fragments are shaded, then partitioned into fine fragments for subsequent per-pixel shading

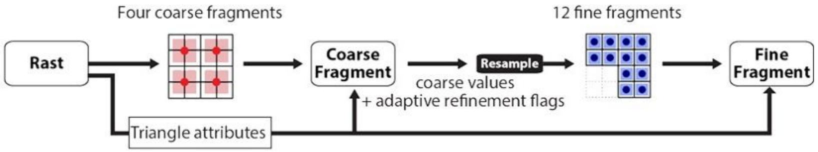

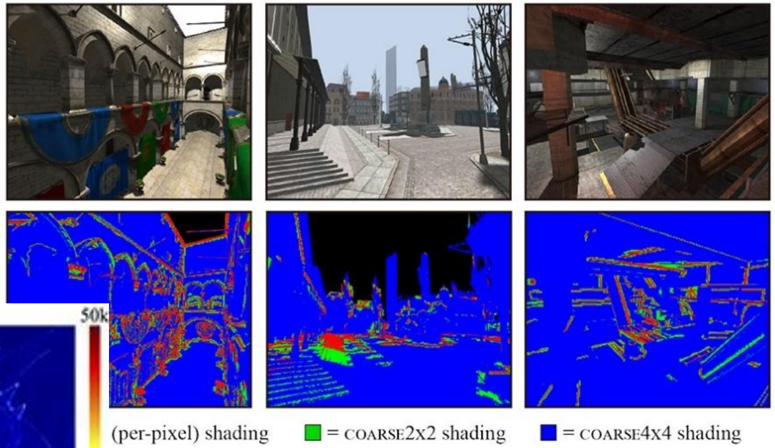

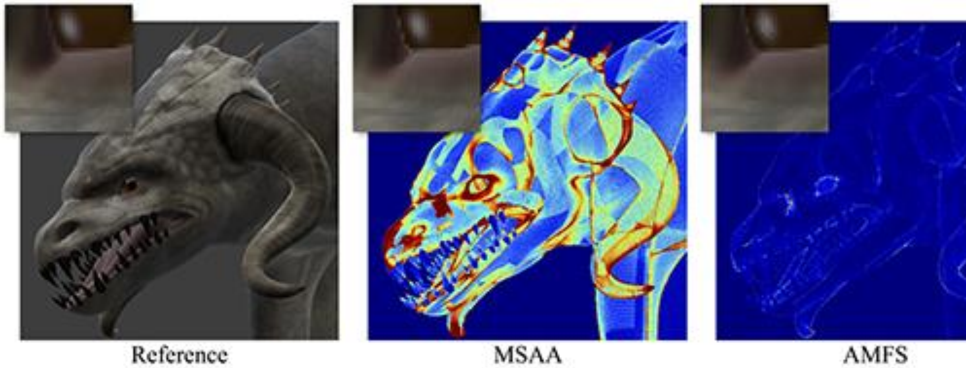

He et al. **Extending the graphics pipeline with adaptive, multi-rate shading**. ACM Transactions on Graphics (TOG) 33.4 , 2014.

Clarberg, Petrik, et al. **AMFS: adaptive multi-frequency shading for future graphics processors.** ACM Transactions on Graphics (TOG) 33.4 , 2014.

Won-Jong Lee, et al. **Adaptive multi-rate ray sampling on mobile ray tracing GPU.** In SIGGRAPH ASIA 2016 Mobile Graphics and Interactive Applications (SA '16).

Mobile rendering with capture

- **Exploiting mobile device sensors...**

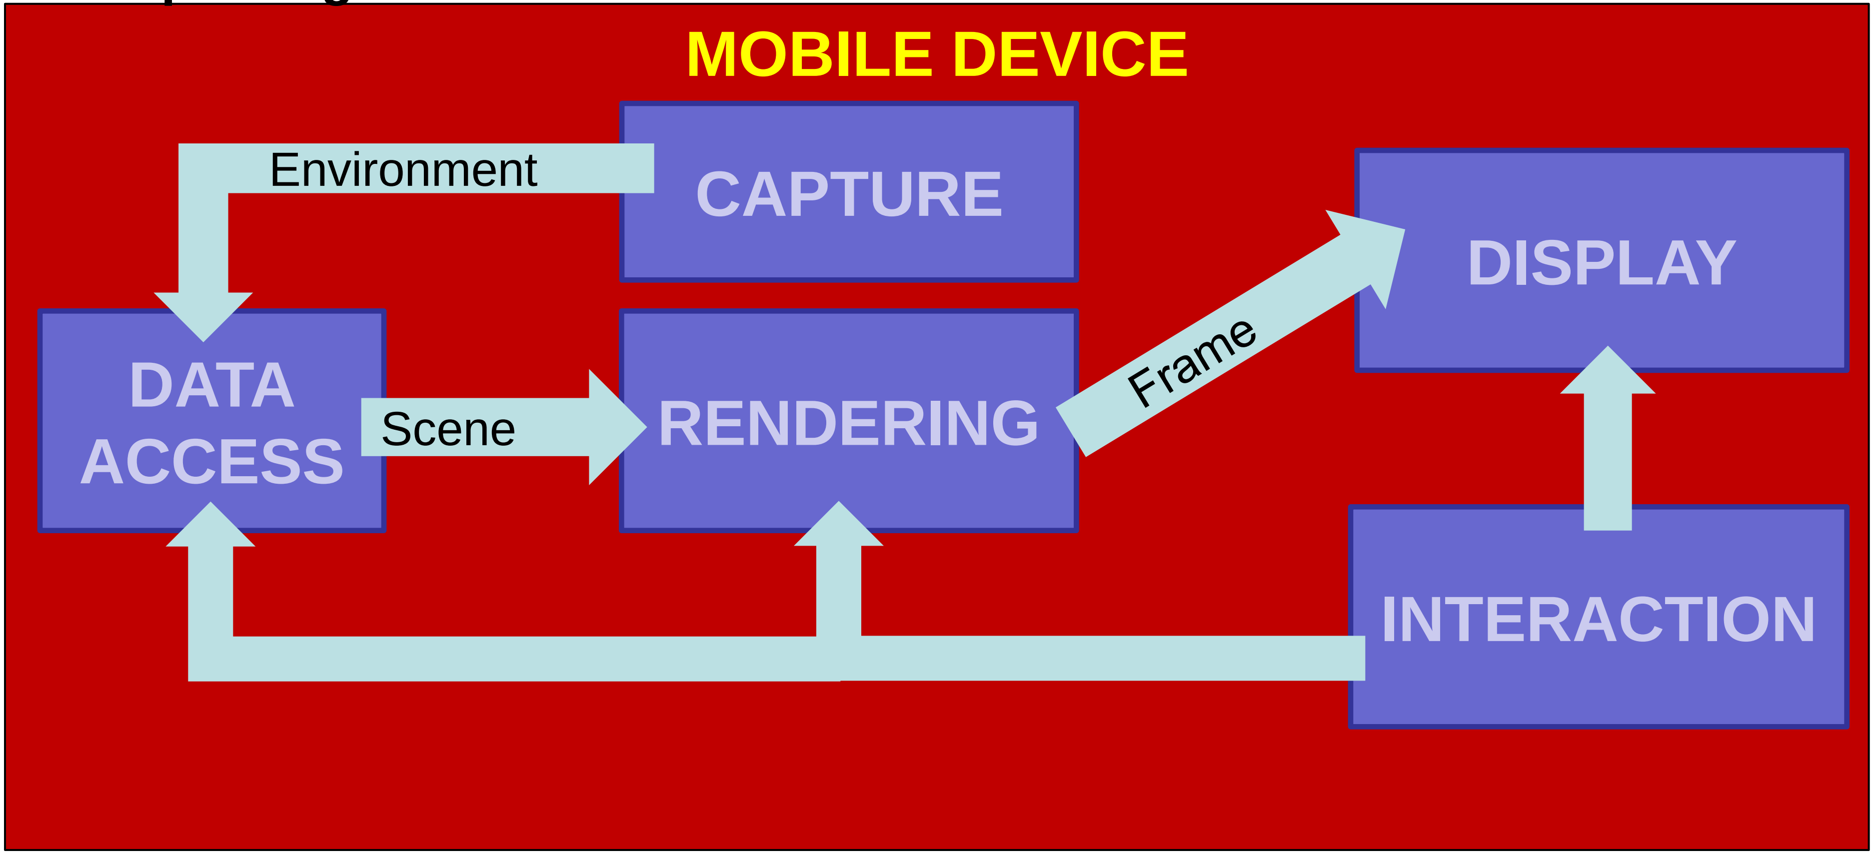

MOBILE DEVICE

CAPTURE

Environment

DATA ACCESS

Scene

RENDERING

Frame

DISPLAY

INTERACTION

Examples: Physical simulations

- **Framework for physically and chemically-based simulations of analog alternative photographic processes**
- **Efficient fluid simulation and manual process running on**

Echevarria et al. **Computational simulation of alternative photographic processes.** Computer Graphics Forum. Vol. 32. 2013.

Examples: Correcting visual aberrations

- **Computational display technology that predistorts the presented content for an observer, so that the target image is perceived without the need for eyewear**
- **Demonstrated in low-cost prototype mobile devices**

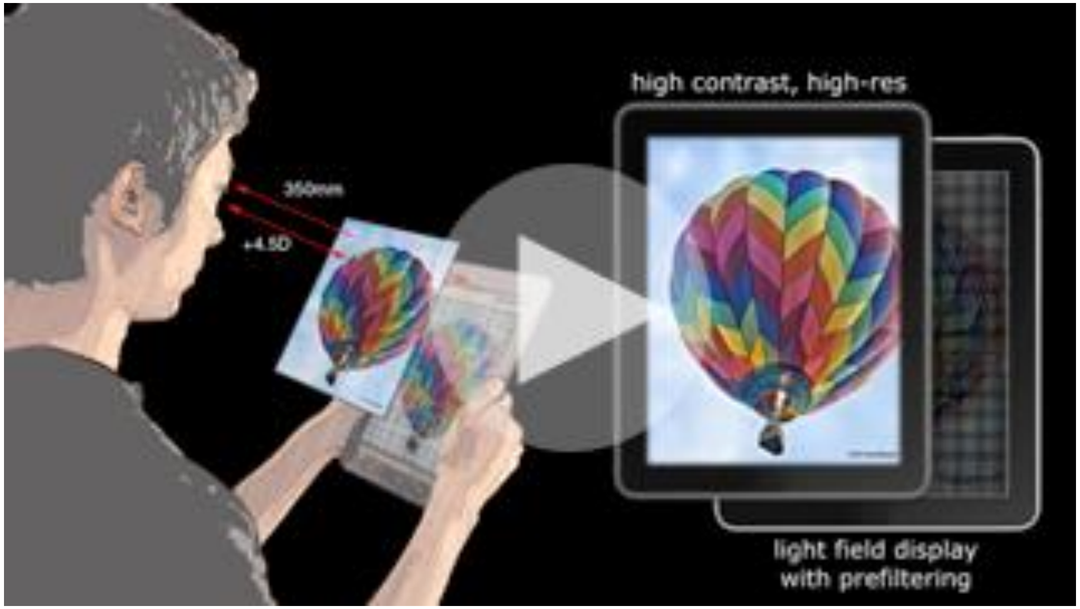

Huang, Fu-Chung, et al**. Eyeglasses-free display: towards correcting visual aberrations with computational light field displays**.ACM Transactions on Graphics (TOG) 33.4, 2014.

Conclusions

- **Heterogeneous applications**
  - driven by bandwidth and processing power
- **Trends**
  - desktop software solutions tend to be ported to the mobile world
    - gaming
    - modelling and 3D animation
    - complex illumination models
- **Sensor integration open new scenarios**
  - examples: live acquisition, mHealth (using sensors and cameras for tracking and processing signals)

Next Session

GRAPHICS DEVELOPMENT FOR MOBILE SYSTEMS