

# VoxelCache: A Cache-Based Memory Architecture for Volume Rendering

---

U. Kanus, G. Wetekam, J. Hirche

WSI / GRIS

University of Tübingen

[urs@gris.uni-tuebingen.de](mailto:urs@gris.uni-tuebingen.de)

WS  
GRIS



- Motivation
- Previous work
- VoxelCache
  - Overview
  - Voxel Memory Layout
  - Block Pre-fetching
  - Cache Implementation
- Results
- Conclusion



- High quality Volume Rendering for large datasets ( $>> 512^3$  Voxels) still requires special purpose hardware.
- Developing ASICs is too expensive and takes too long (at least for us)
- Reconfigurable hardware devices (FPGAs) are large enough and fast enough to implement high performance graphics hardware.



Memory is the key issue for volume rendering systems

→ need a memory system for Volume Rendering that:

- Combines regular memory access patterns and ray-casting
- No replication of volume data
- Allows for efficient caching of voxels
- Can be used with different types and arrangements of external memory
- Can be easily implemented on FPGAs



- Motivation
- Previous work
- VoxelCache
  - Overview
  - Voxel Memory Layout
  - Block Pre-fetching
  - Cache Implementation
- Results
- Conclusion



- VIRIM [Guenther et al. 94]:  
Eight memory modules, interleaved
- VIZARD [Knittel 97]:  
Host memory + on-board cache, lossy block compression
- VolumePro [Pfister et al. 99]:  
One memory module per pipeline, interleaved blocks (skewing), not suitable for persp. proj.
- VIZARDII [Meißner et al. 02]:  
Four memory modules + FIFOs, no on-chip cache
- Many other proposals (see paper)

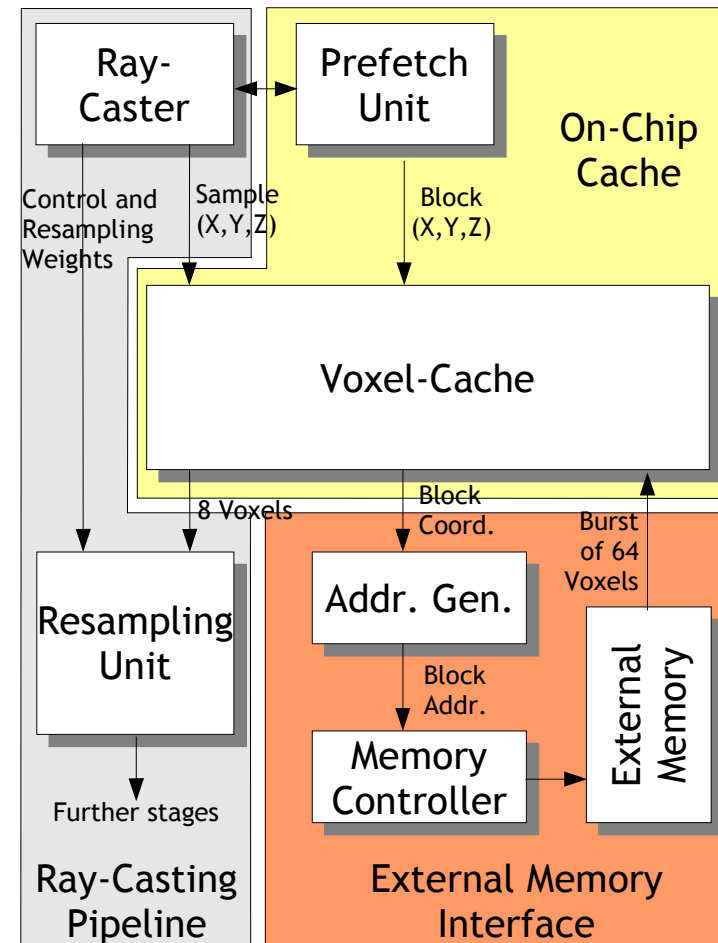


- Motivation
- Previous work
- **VoxelCache**
  - **Overview**
  - Voxel Memory Layout
  - Block Pre-fetching
  - Cache Implementation
- Results
- Conclusion



System split into three parts:

- VoxelCache on-chip Cache
- External memory interface
- Ray-casting pipeline

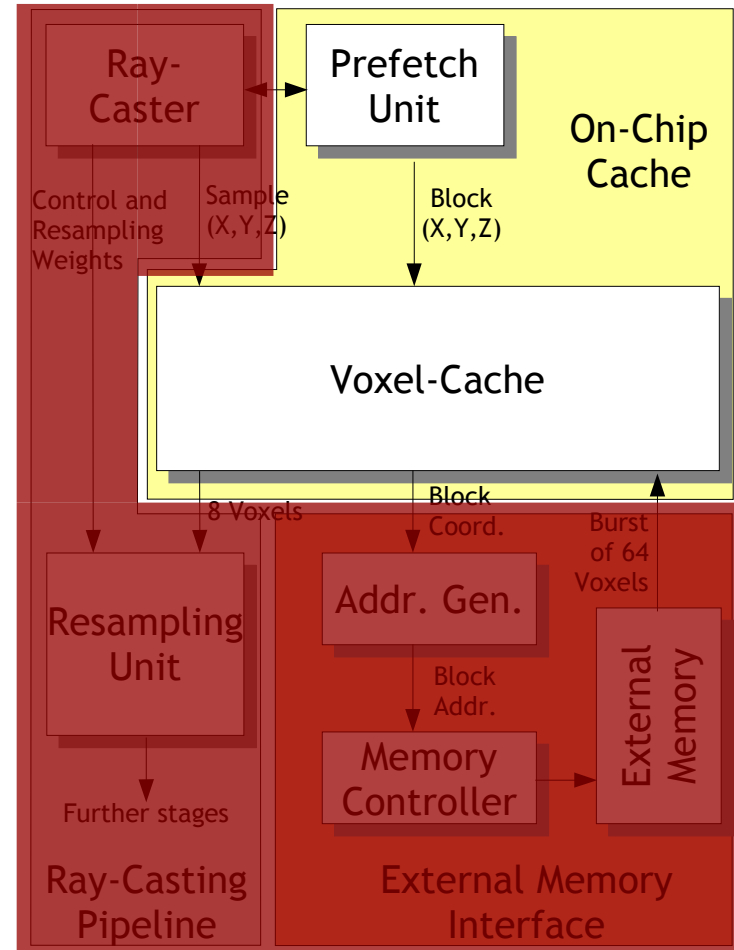






System split into three parts:

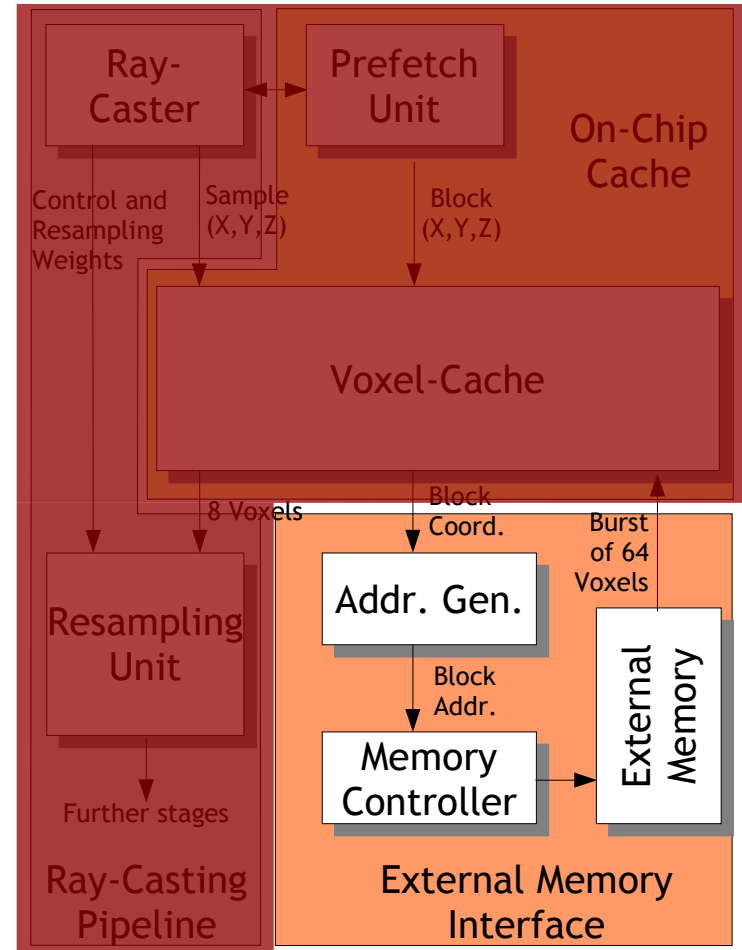
- VoxelCache on-chip Cache
  - Provides voxel neighborhood for resampling for a given sample position
  - Generic interface to external memory
  - Block prefetching to minimize cache misses
- External memory interface
- Ray-casting pipeline





System split into three parts:

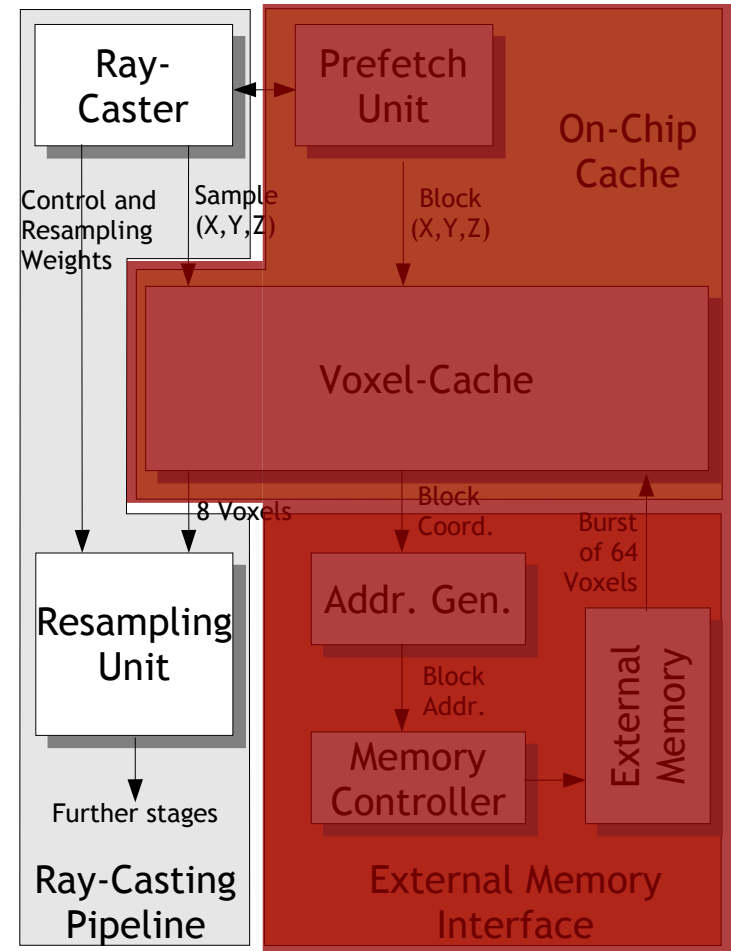
- VoxelCache on-chip Cache
- External memory interface
  - Translates block coordinates into memory addresses
  - Burst transfers of 64 voxels at a time
- Ray-casting pipeline





System split into three parts:

- VoxelCache on-chip Cache
- External memory interface
- Ray-casting pipeline
  - Standard ray-casting pipeline for perspective and parallel projections

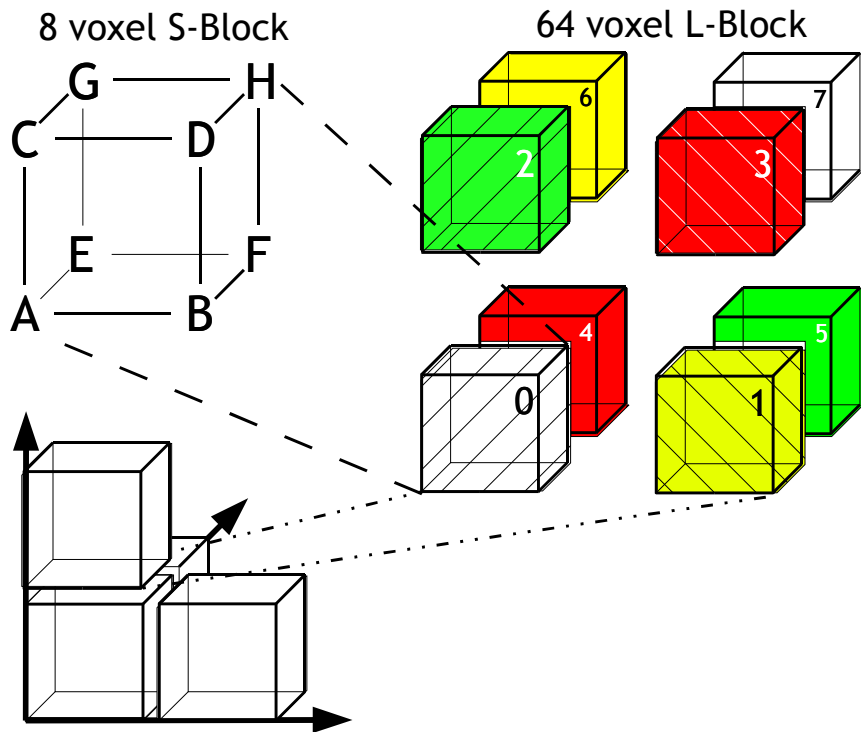





- Motivation
- Previous work
- **VoxelCache**
  - Overview
  - **Voxel Memory Layout**
  - Block Pre-fetching
  - Cache Implementation
- Results
- Conclusion



## Block hierarchy





## Cache memory




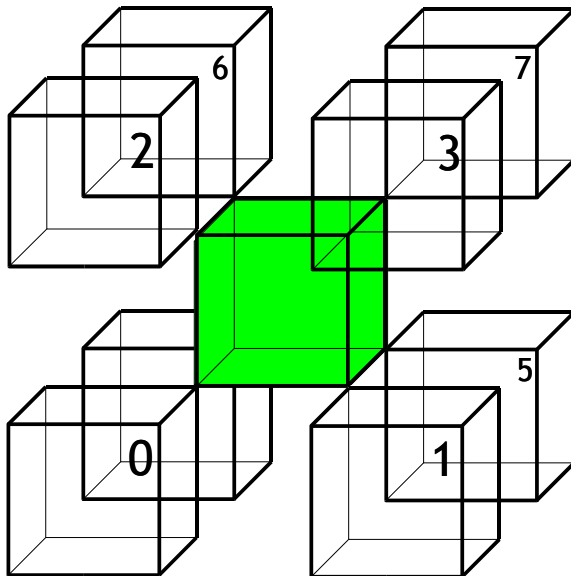


## Accessing voxels from a single S-Block



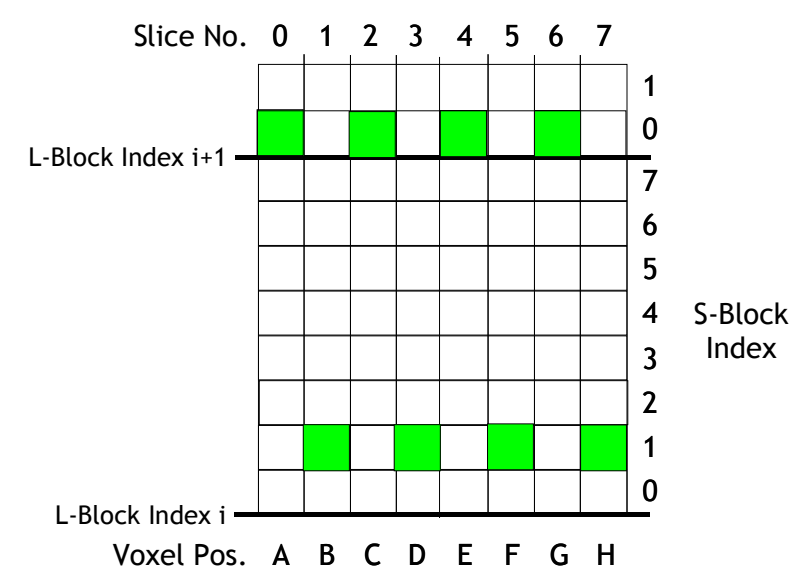



## Accessing voxels from several S-Blocks





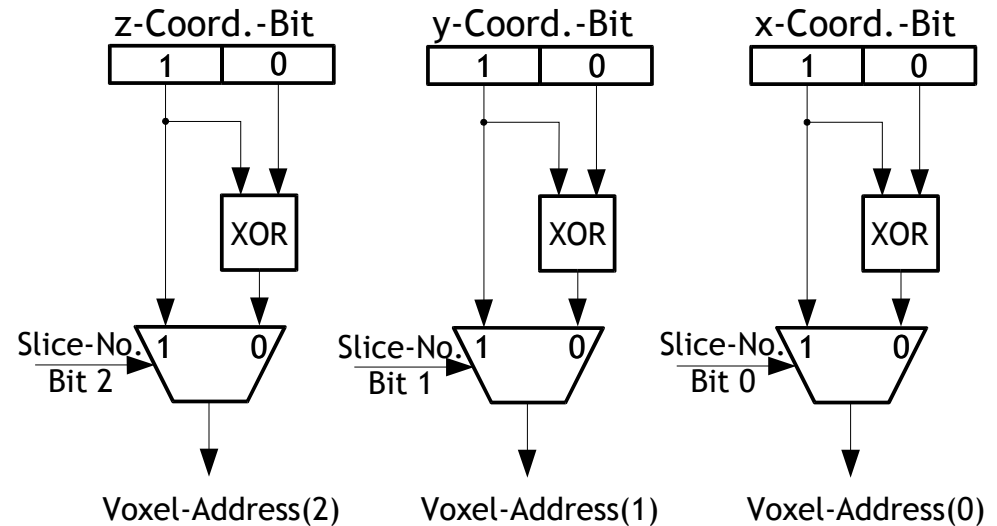
## Accessing voxels from several L-Blocks







- 3-bit voxel-address inside L-Block
- Each cache slice needs one address generator
- Simple logic (3 XOR + 3 MUX per address generator)



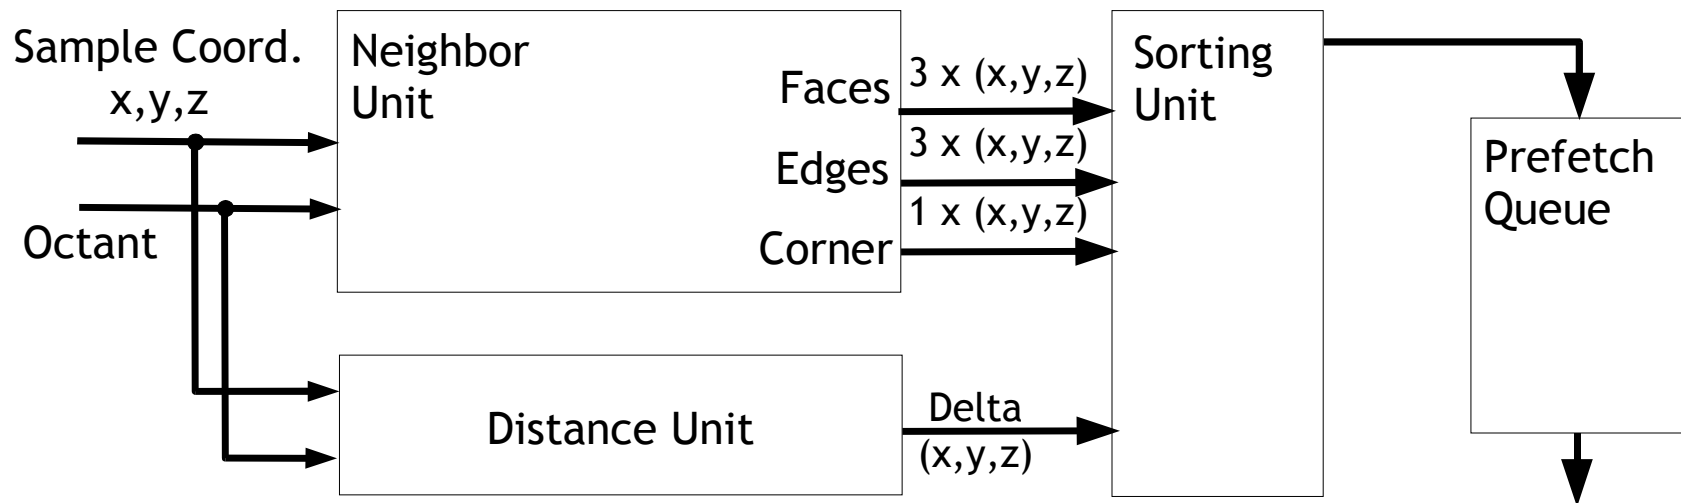


- Motivation
- Previous work
- **VoxelCache**
  - Overview
  - Voxel Memory Layout
  - **Block Pre-fetching**
  - Cache Implementation
- Results
- Conclusion



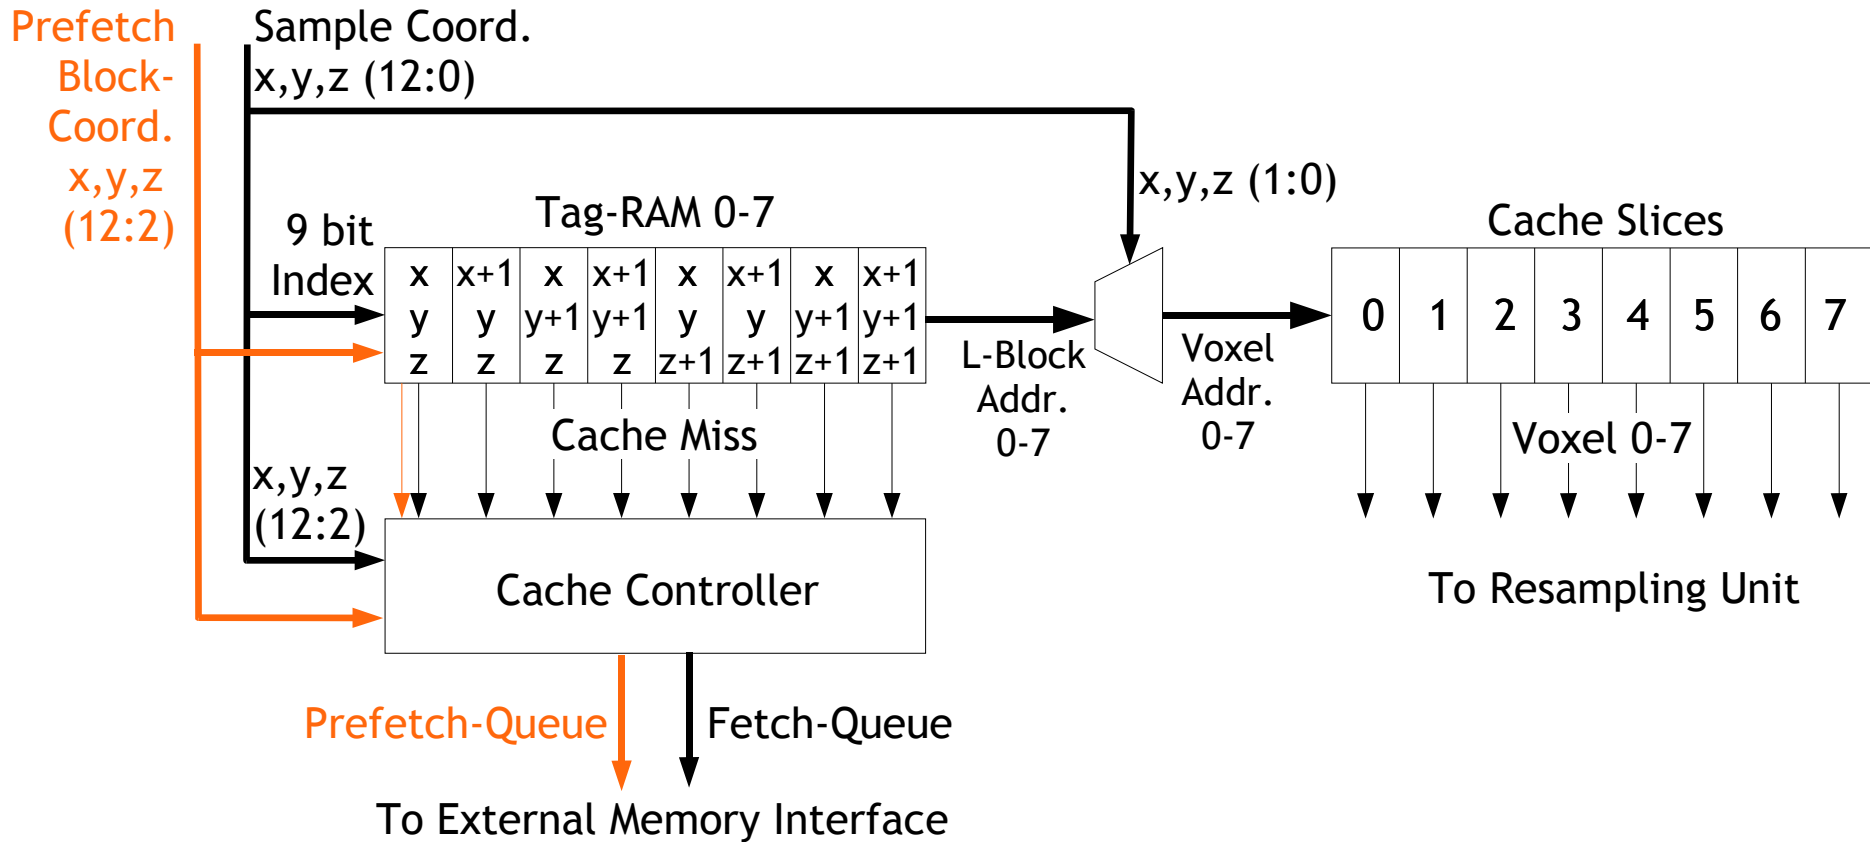
Pre-fetching of L-Blocks based on:

- Ray-direction (octant)
- Sample position in current L-Block





- Motivation
- Previous work
- **VoxelCache**
  - Overview
  - Voxel Memory Layout
  - Block Pre-fetching
  - **Cache Implementation**
- Results
- Conclusion





Cache characteristics:

- Tag-RAM direct mapped (9 coordinate bits)
- Arbitrary L-Block allocation in cache memory
- L-Blocks replacement in FIFO order, unless L-Block to be replaced is required for the next sampling neighborhood



- Motivation
- Previous work
- VoxelCache
  - Overview
  - Voxel Memory Layout
  - Block Pre-fetching
  - Cache Implementation
- Results
- Conclusion

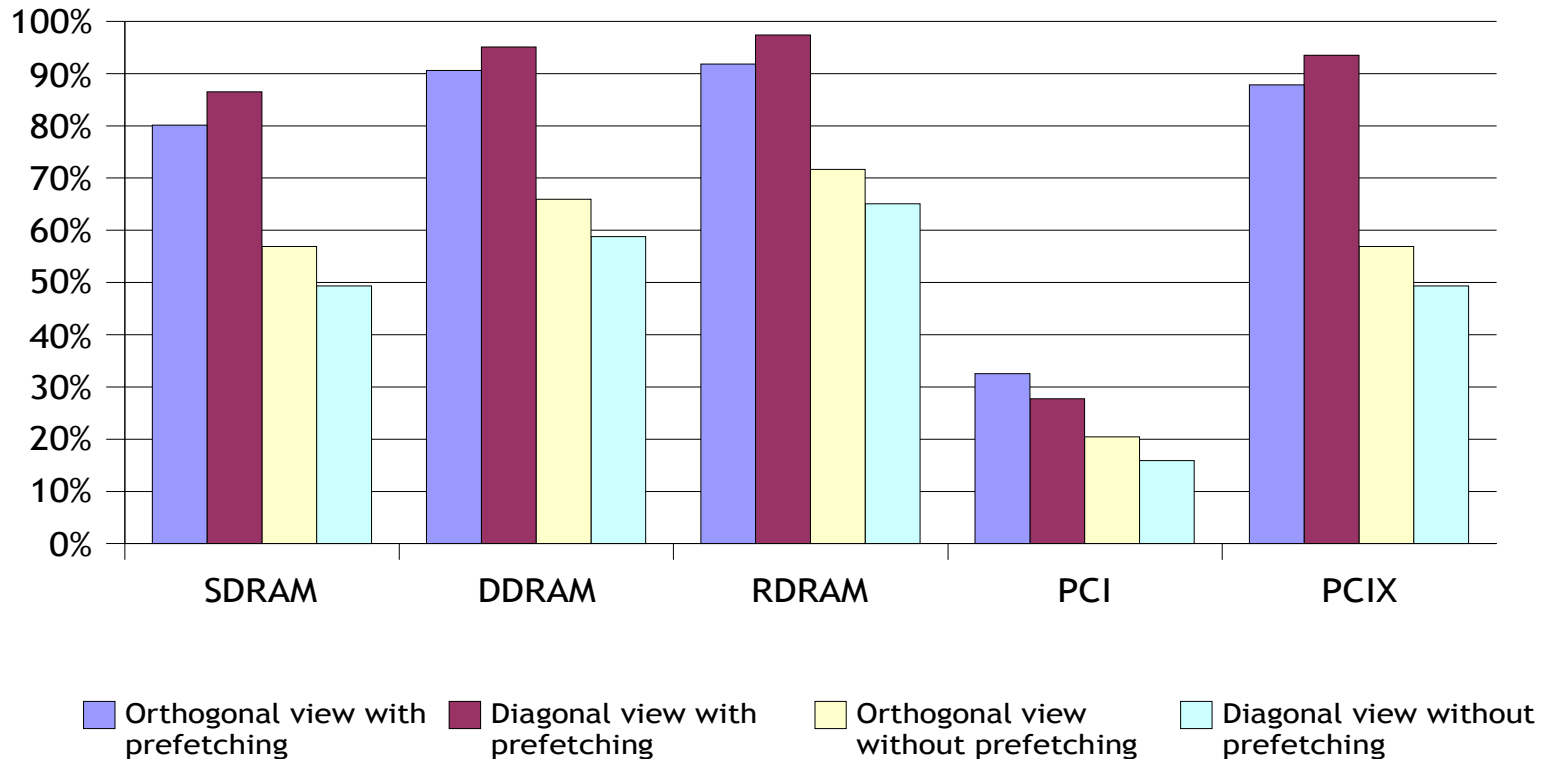


- Cycle accurate C++ simulation embedded into software ray-caster:
  - Assuming a single rendering pipeline running @ 133 Mhz
  - Different memory types defined by latency and bandwidth
  - Cache size of 128 L-Blocks
    - fits well into target FPGA architecture
  - Dataset size  $136^3$  voxels
    - software simulation is time-consuming
    - For size  $>128^3$  an orthogonal ray doesn't fit into the cache anymore




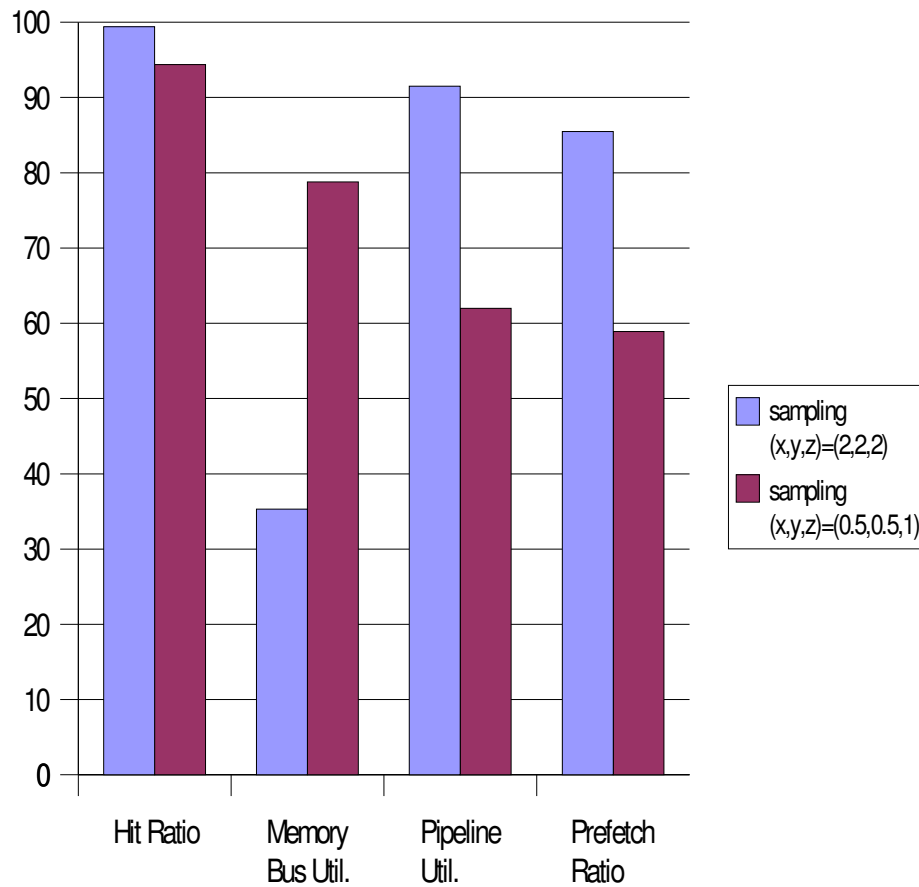


## Pipeline Utilization





## Cache performance vs. Sampling (DDRAM):





- Motivation
- Previous work
- VoxelCache
  - Overview
  - Voxel Memory Layout
  - Block Pre-fetching
  - Cache Implementation
- Results
- Conclusion



- VoxelCache is an efficient memory architecture for Volume Rendering
  - Hit ratio  $> 98\%$  in most cases, increases with sampling rate
  - No data replication
  - Sustained performance even for random access patterns, few pipeline idle cycles
- VoxelCache is designed for:
  - implementation on reconfigurable devices
  - easy adaption to different types of external memory



Future extensions to VoxelCache:

- Integrate VoxelCache with space leaping
- Explore block compression schemes
- Efficient real-time volume data update



Thanks for your attention ...

... Questions ?