A Hardware Architecture for Multi-Resolution Volume Rendering

Graphics Hardware 2005

Overview

- Motivation
- Compressed Multi-Resolution Representation
- Hardware Decompression
- Memory Managment
- Raycasting Pipeline
- Results and Conclusion

Motivation

Visualization of large volume data sets

Scientific simulation

Medicine

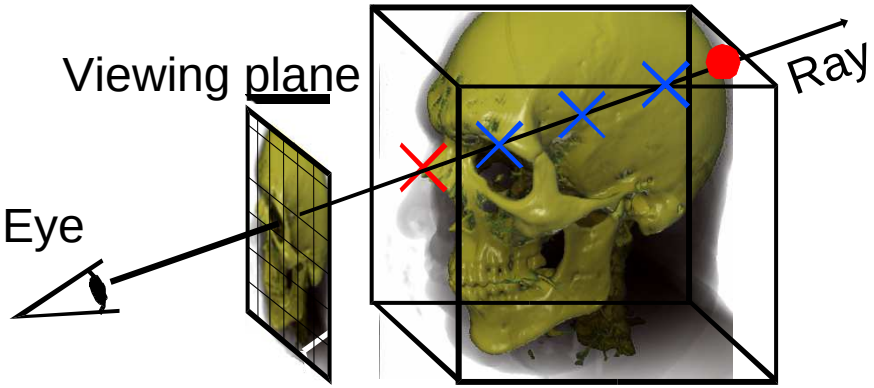

- Entry point of the ray
- Sample point
- Exit point of the ray

Dedicated hardware for volume rendering

VolumePro, VIZARD II...

Able to render >512^3 datasets with large dedicated on-board memory and/or swapping via the host bus.

Introduction

Problem:

Large size of actual datasets leads to problems with memory consumption and bandwidth.

Solution:

- Support for multi-resolution (wavelet) hierarchy
  → Reduction of the memory and calculation complexity: O(n) → O(log n)
- Hierarchy compression
  → Reduction of the needed bandwidth
- Flexible access to the volume data
  → Virtualized memory access (VoxelCache)

Volume raycasting with different level-of-details

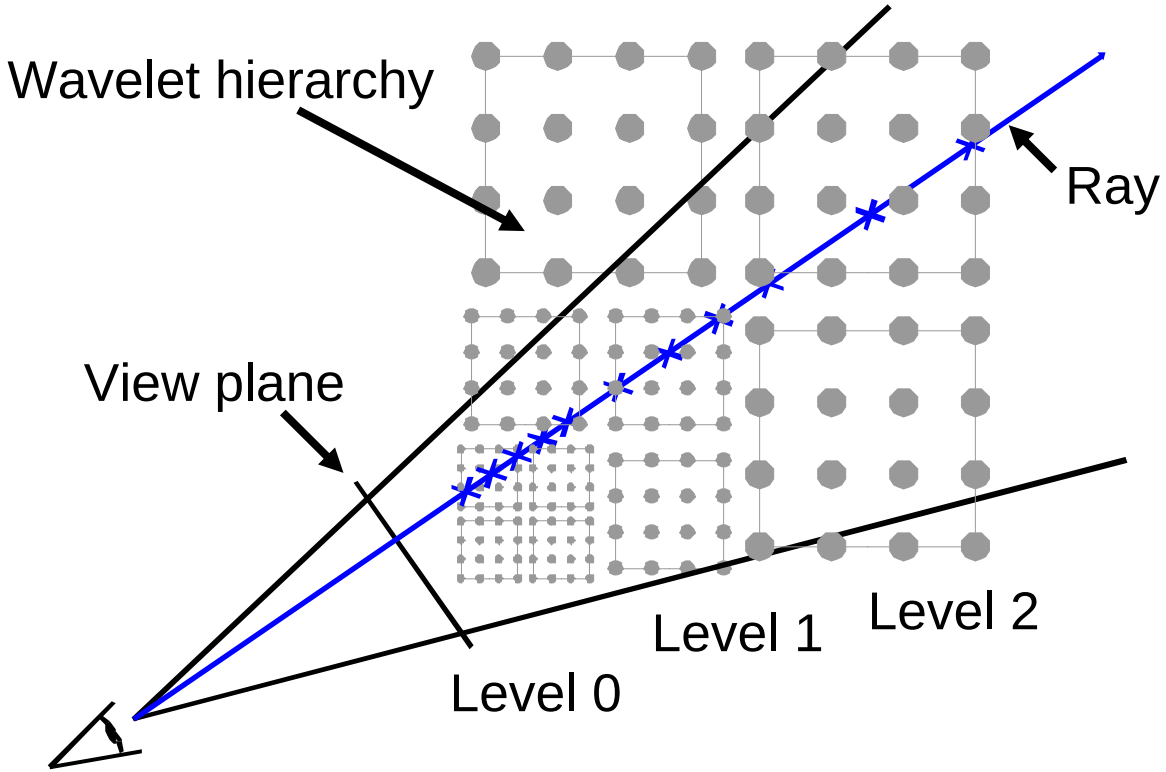

- Voxels closer to the viewer are in a finer resolution
- Voxels far away are in a lower resolution
  → Less memory needed

System Architecture

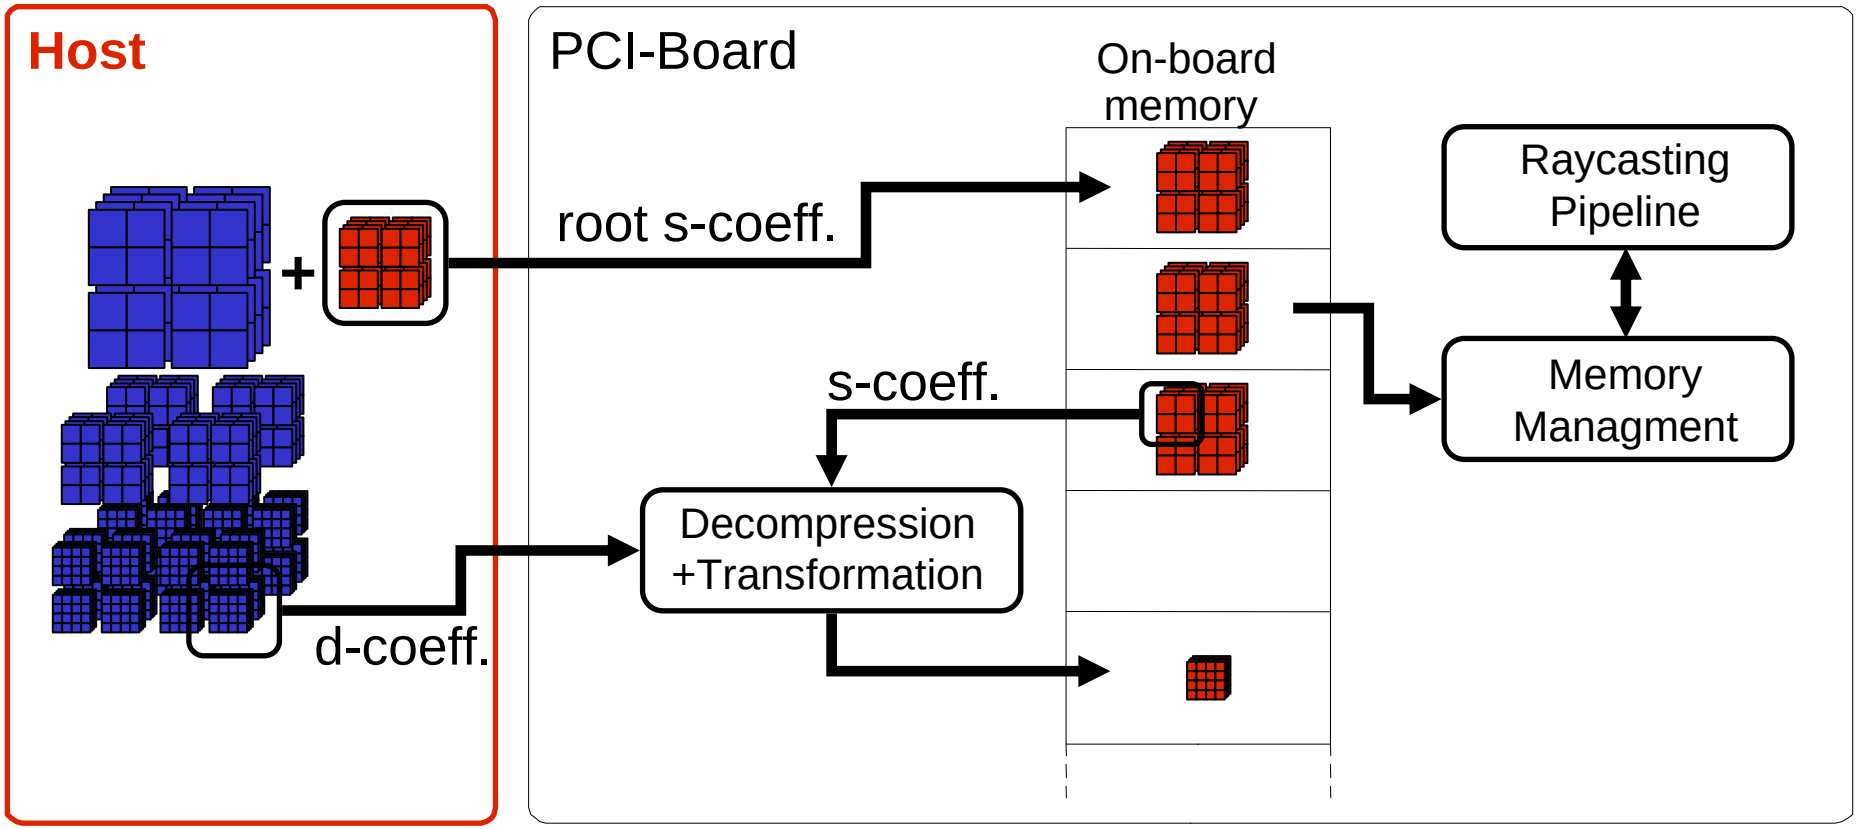

Wavelet Hierarchy

- Linearly interpolated spline wavelets
  → high compression ratio
  → local support of 3 or 5 respectively

- Integer lifting transformation
  → Fixed point arithmetic (16 bit)
  → Low costs for transformation
  (only 3 additions and 1 shift)
  → “In place“ – no additional memory needed
  “Non-standard” decomposition in 3D

see Stefan Guthe, 2001

Wavelet Hierarchy

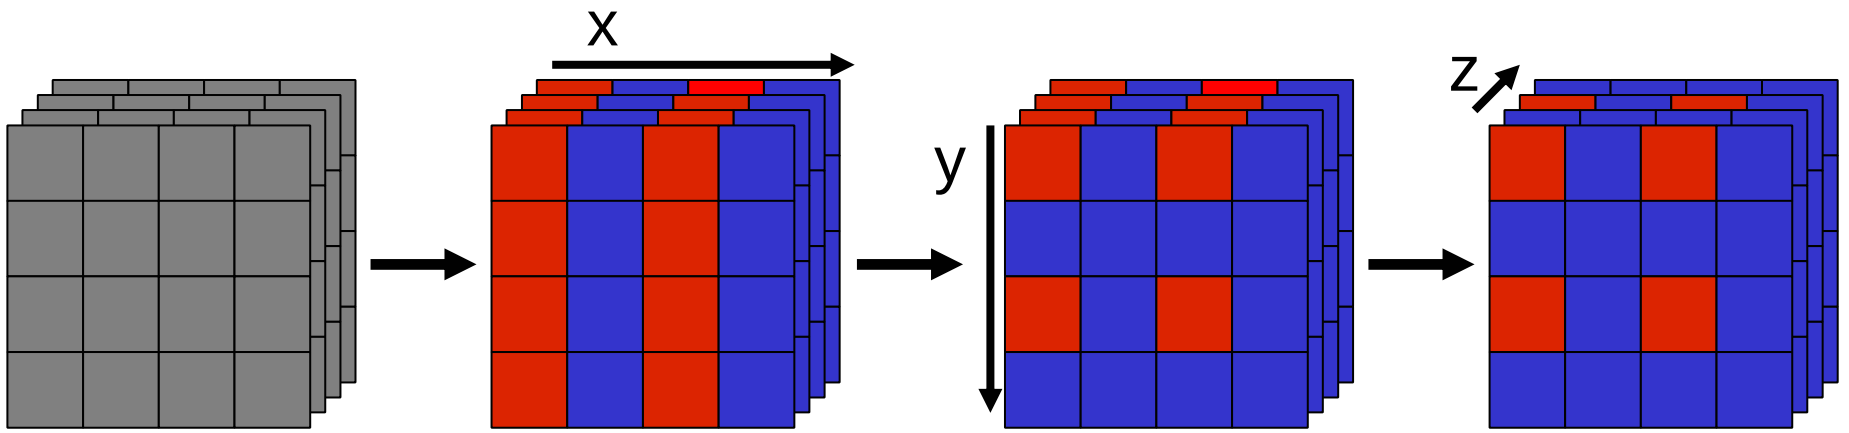

lowpass coefficients (s-coeff.)

highpass (detail) coefficients (d-coeff.)

raw data

Hierarchy Construction

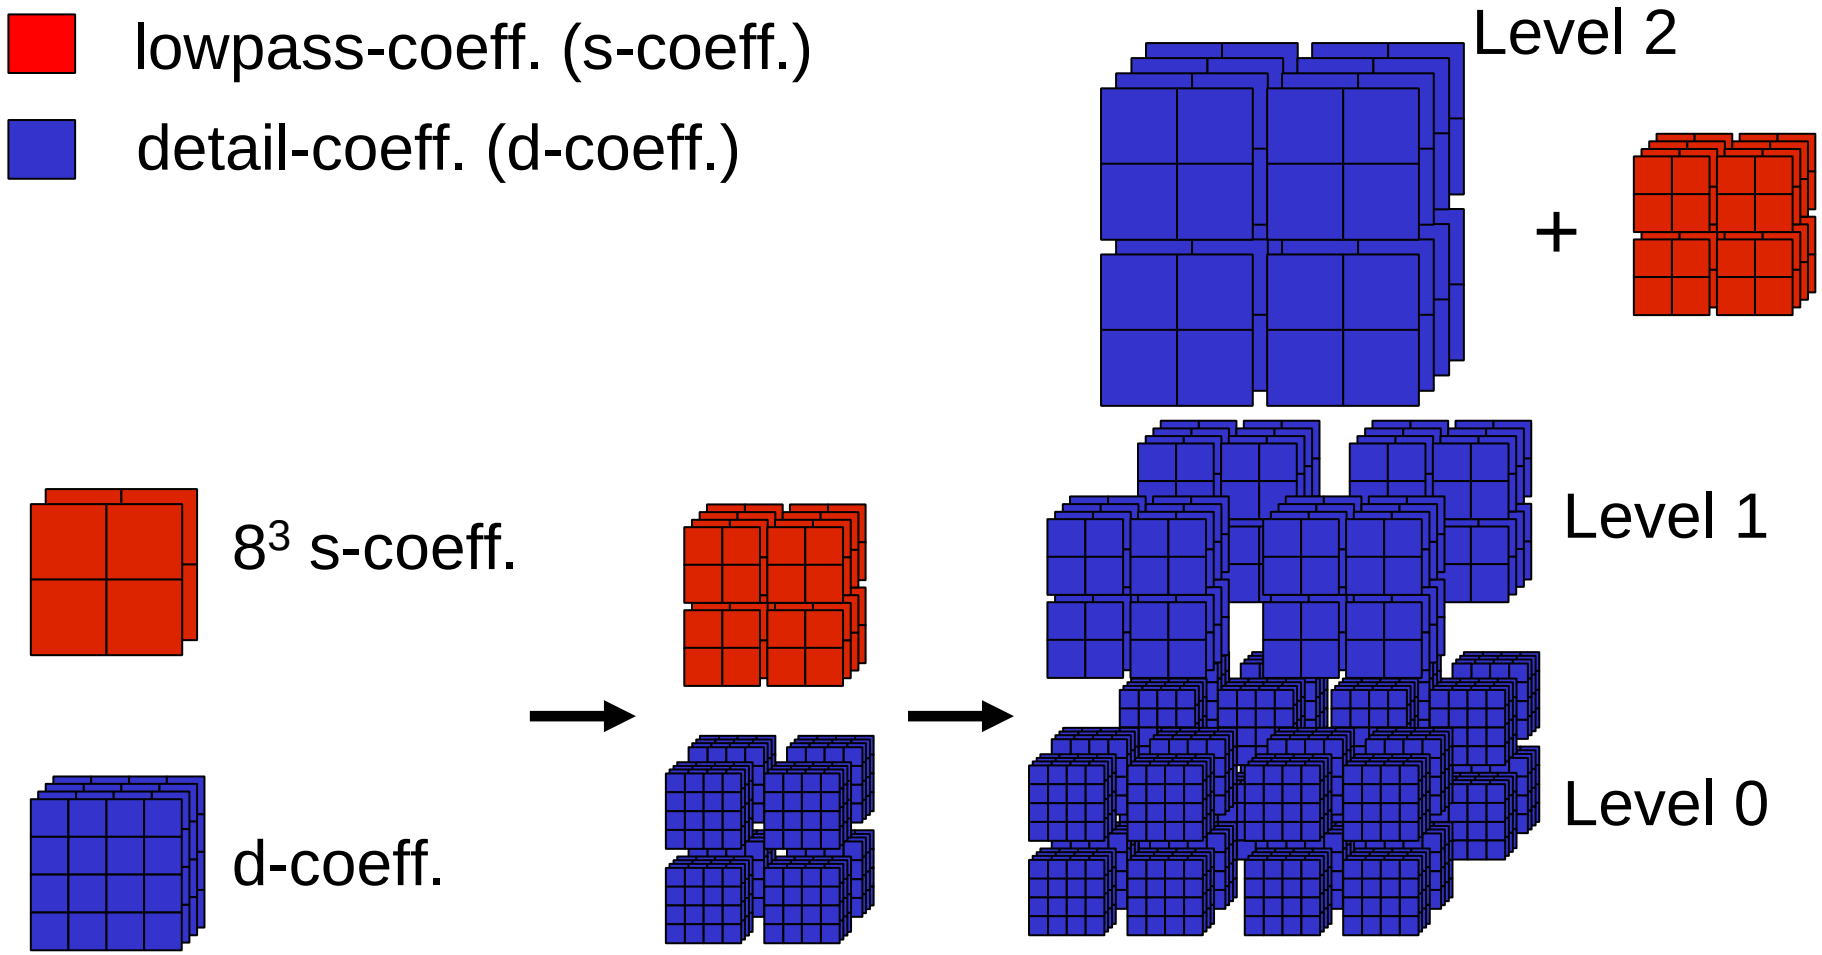

Fixed Huffman Code

Additional coding of the coefficients:

- Series of n (≤ 128) zero coefficients are coded in 1 Byte:

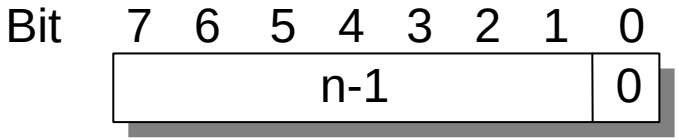

- Coefficients !=0 are coded separately:

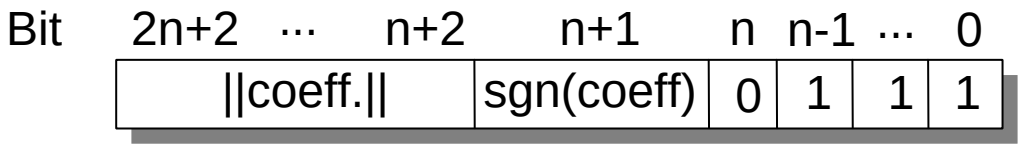

Volume Data Compression

- Volume data is compressed on the host (preprocessing)

→ No major calculations on the host during rendering

→ Host manages octree hierarchy with compressed blocks

→ Significant reduction of the bandwidth between host and dedicated rendering hardware

System Architecture

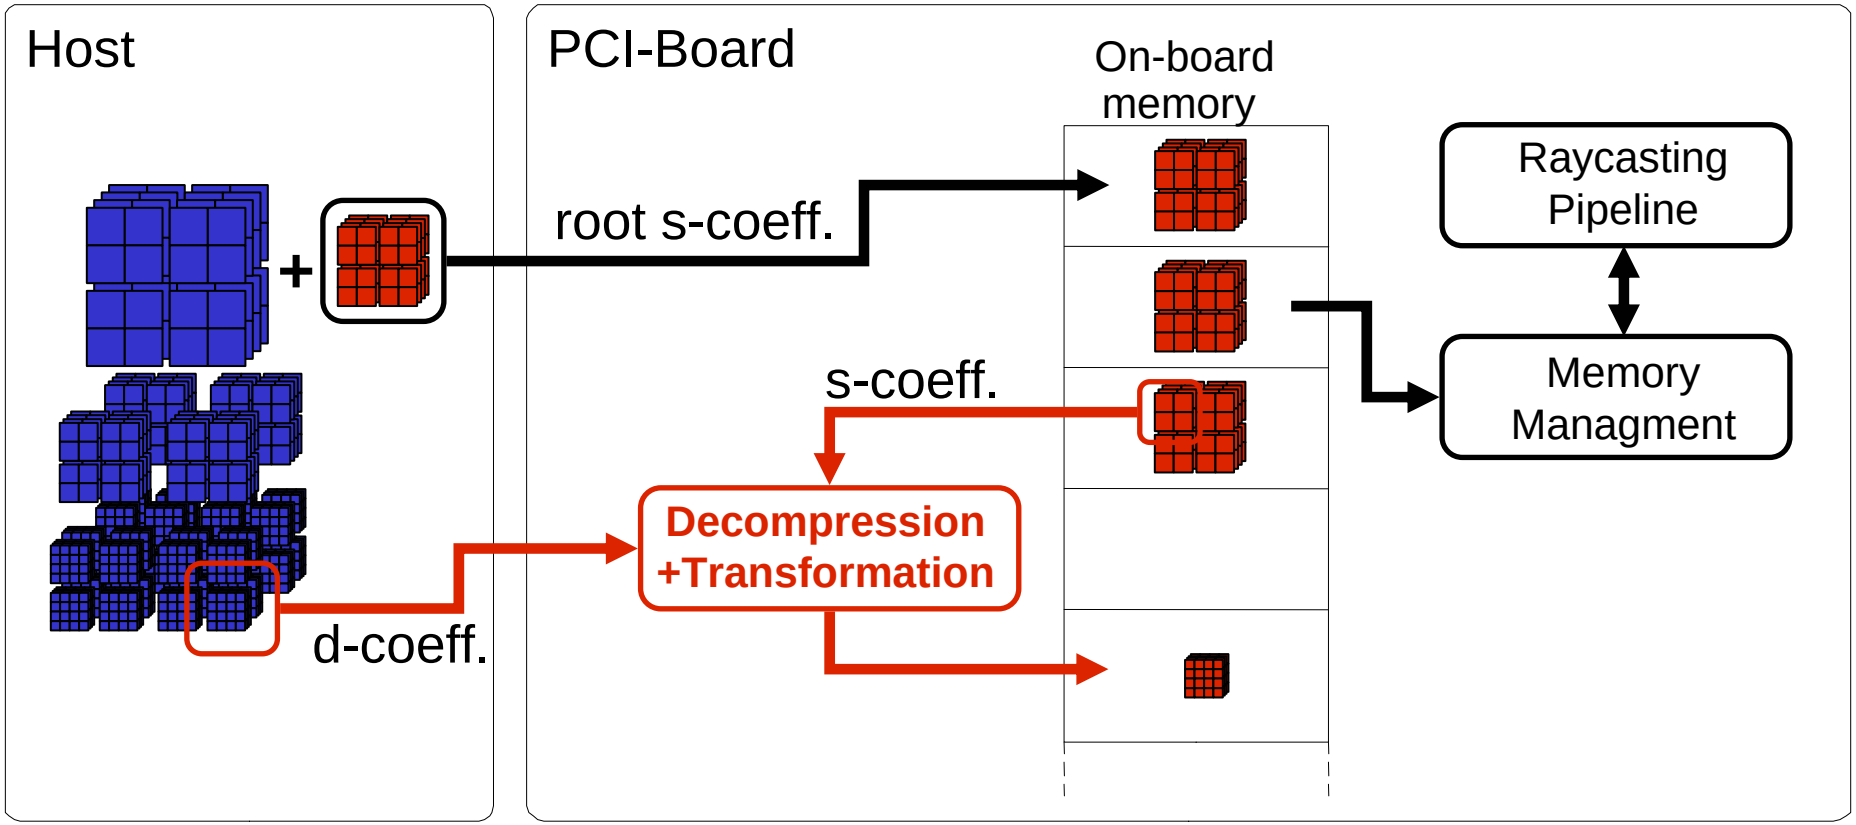

Fixed Huffman Code Decompression

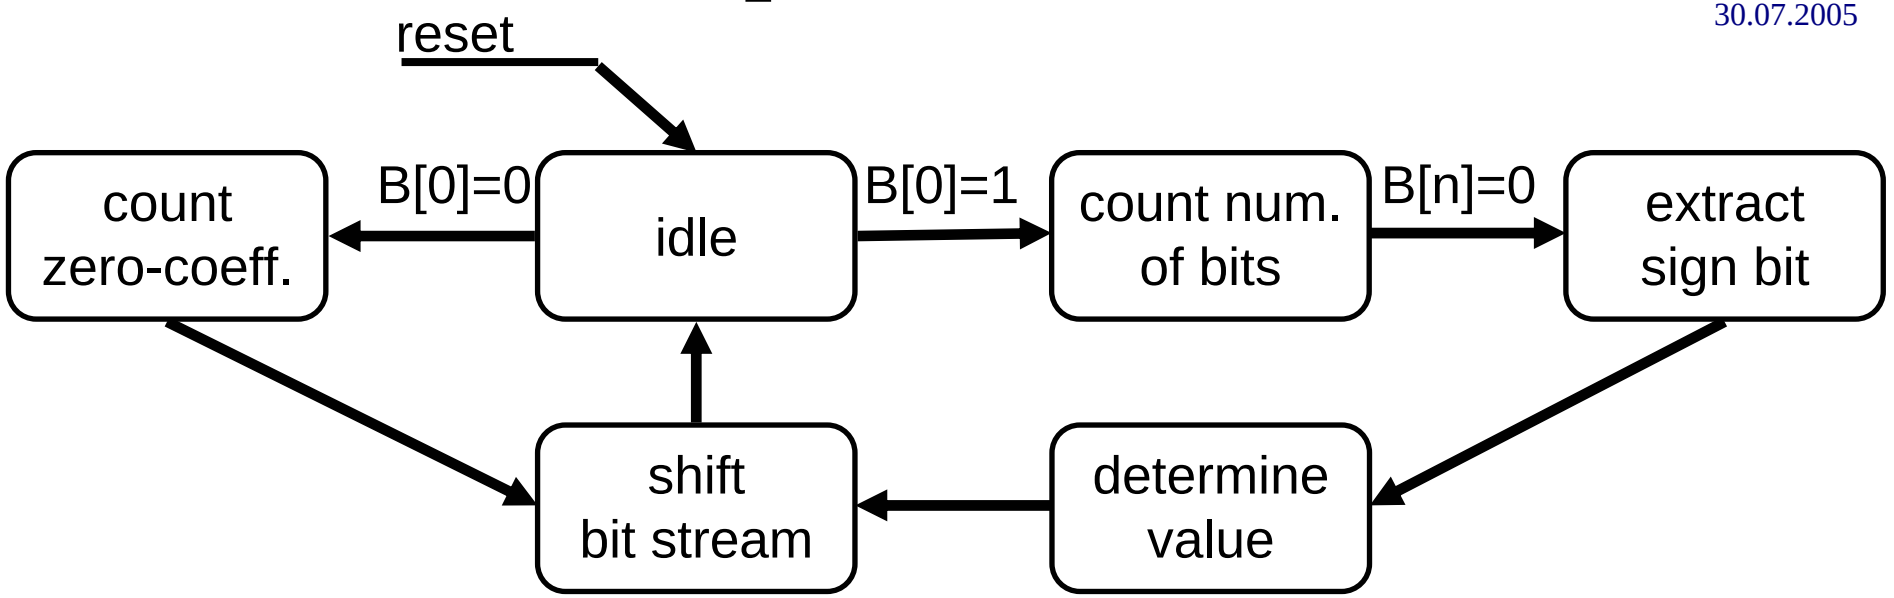

- No expensive arithmetic
- No memory consumption for a codebook

- Compression throughput (10 units): 296 MB/s
- Performance: 160 MHz (post synthesis)
- On-chip resources: 5% logic

Wavelet Transformation

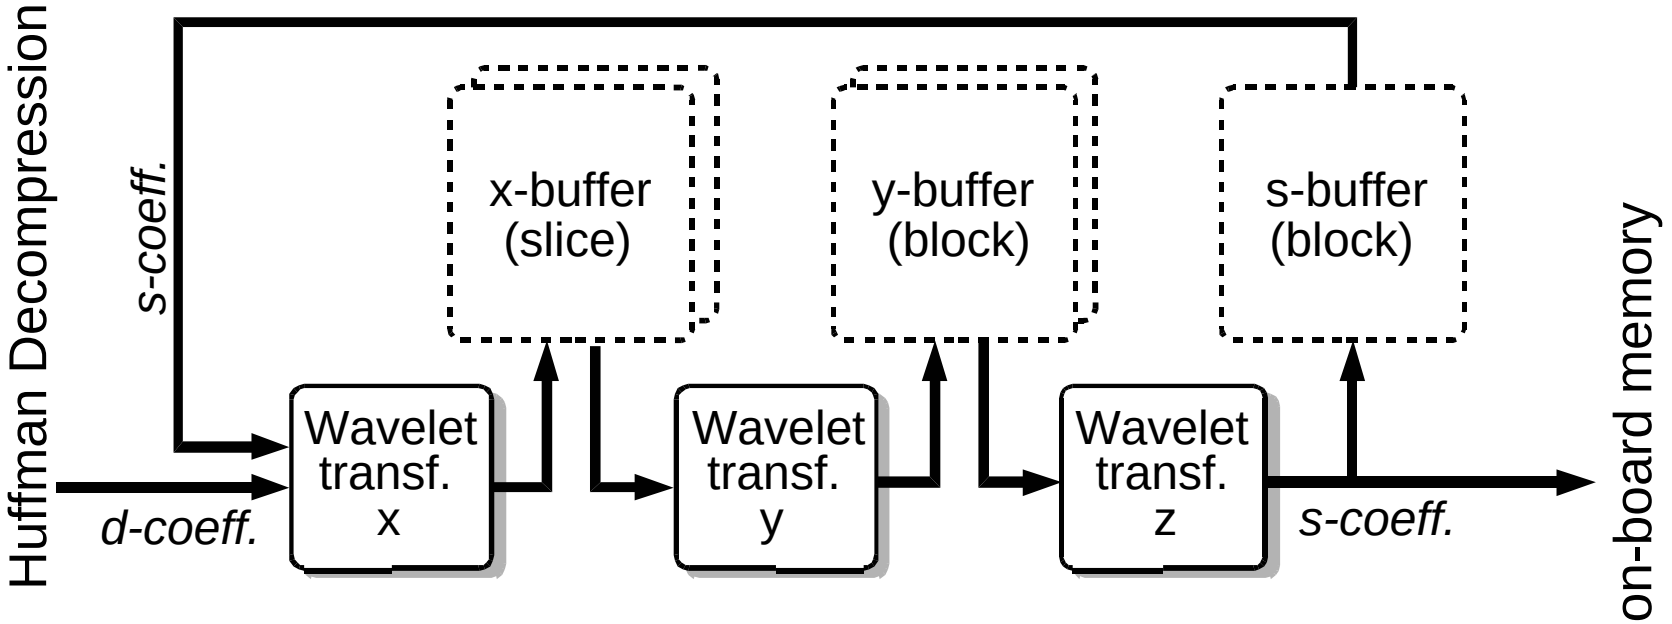

- Transformation throughput:
  40.000 blocks/s ≙ 300 MB/s
- Clock speed:
  165 MHz (post synthesis)
- On-Chip resources:
  4% logic,
  12% memory

System Architecture

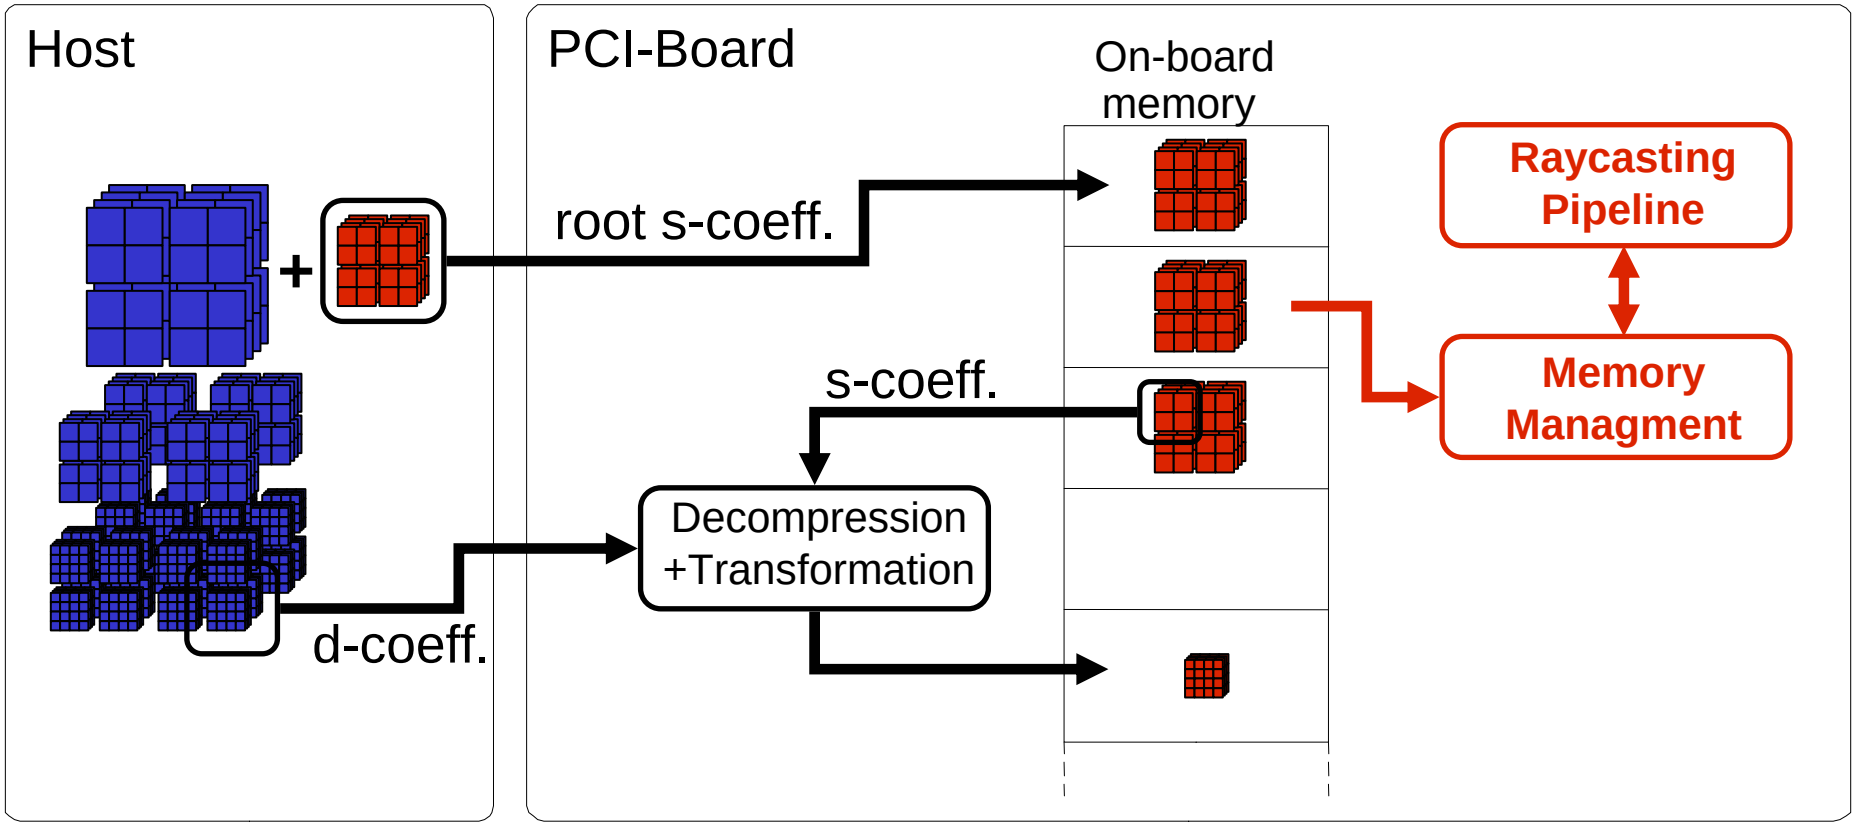

Virtual Memory Access

Support for very large datasets with limited memory resources on the hardware.

The multi-resolution scheme requires a dynamic managment of the on-board voxel memory.

→ VoxelCache:

- Separation of the rendering pipeline and memory
- Efficient caching of voxel data
- Abstraction of memory implementation → Easy change of the used memory type

Memory Access

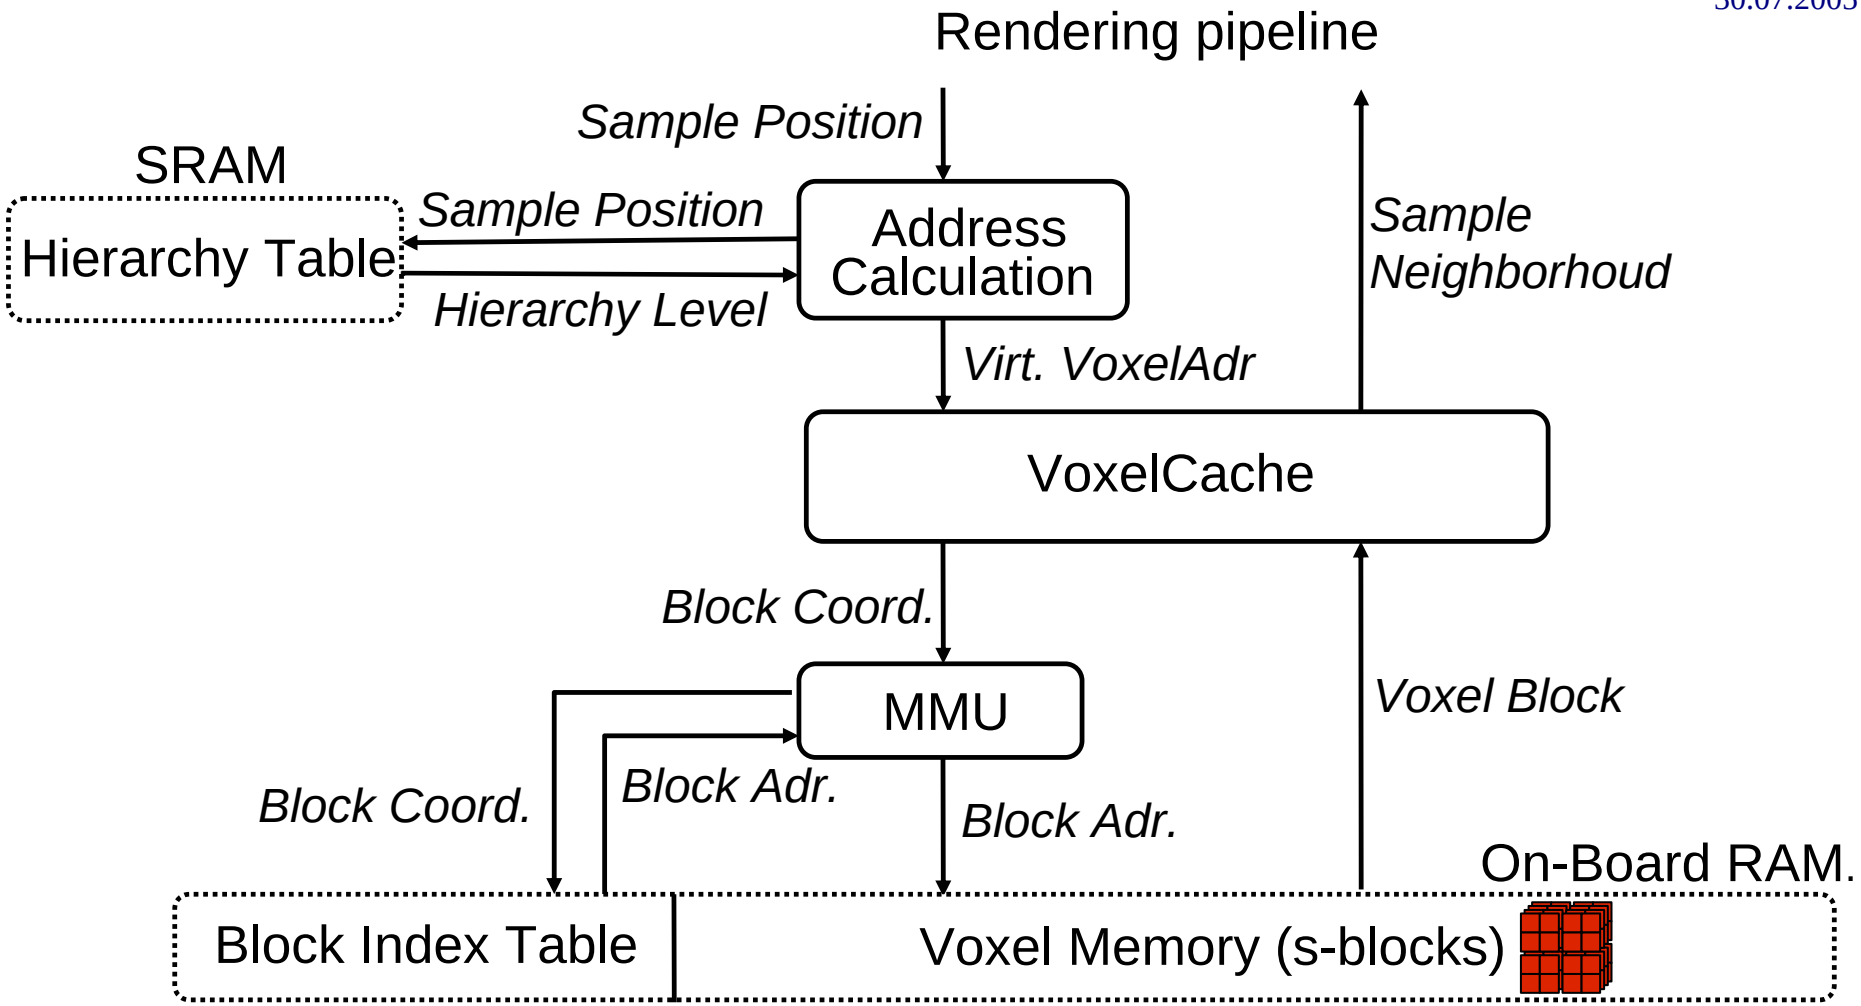

VoxelCache Evaluation

- Cache-Hitrate: 98 %
- Memory bus utilization: 4.5 %
- Utilization of the rendering pipeline: 90 %
- On-chip resources: 5 % Logic
  10 % Memory

Memory type: DDR-RAM

Oversampling: x 1x, y 2x, z 4x

Ray-Casting Pipeline

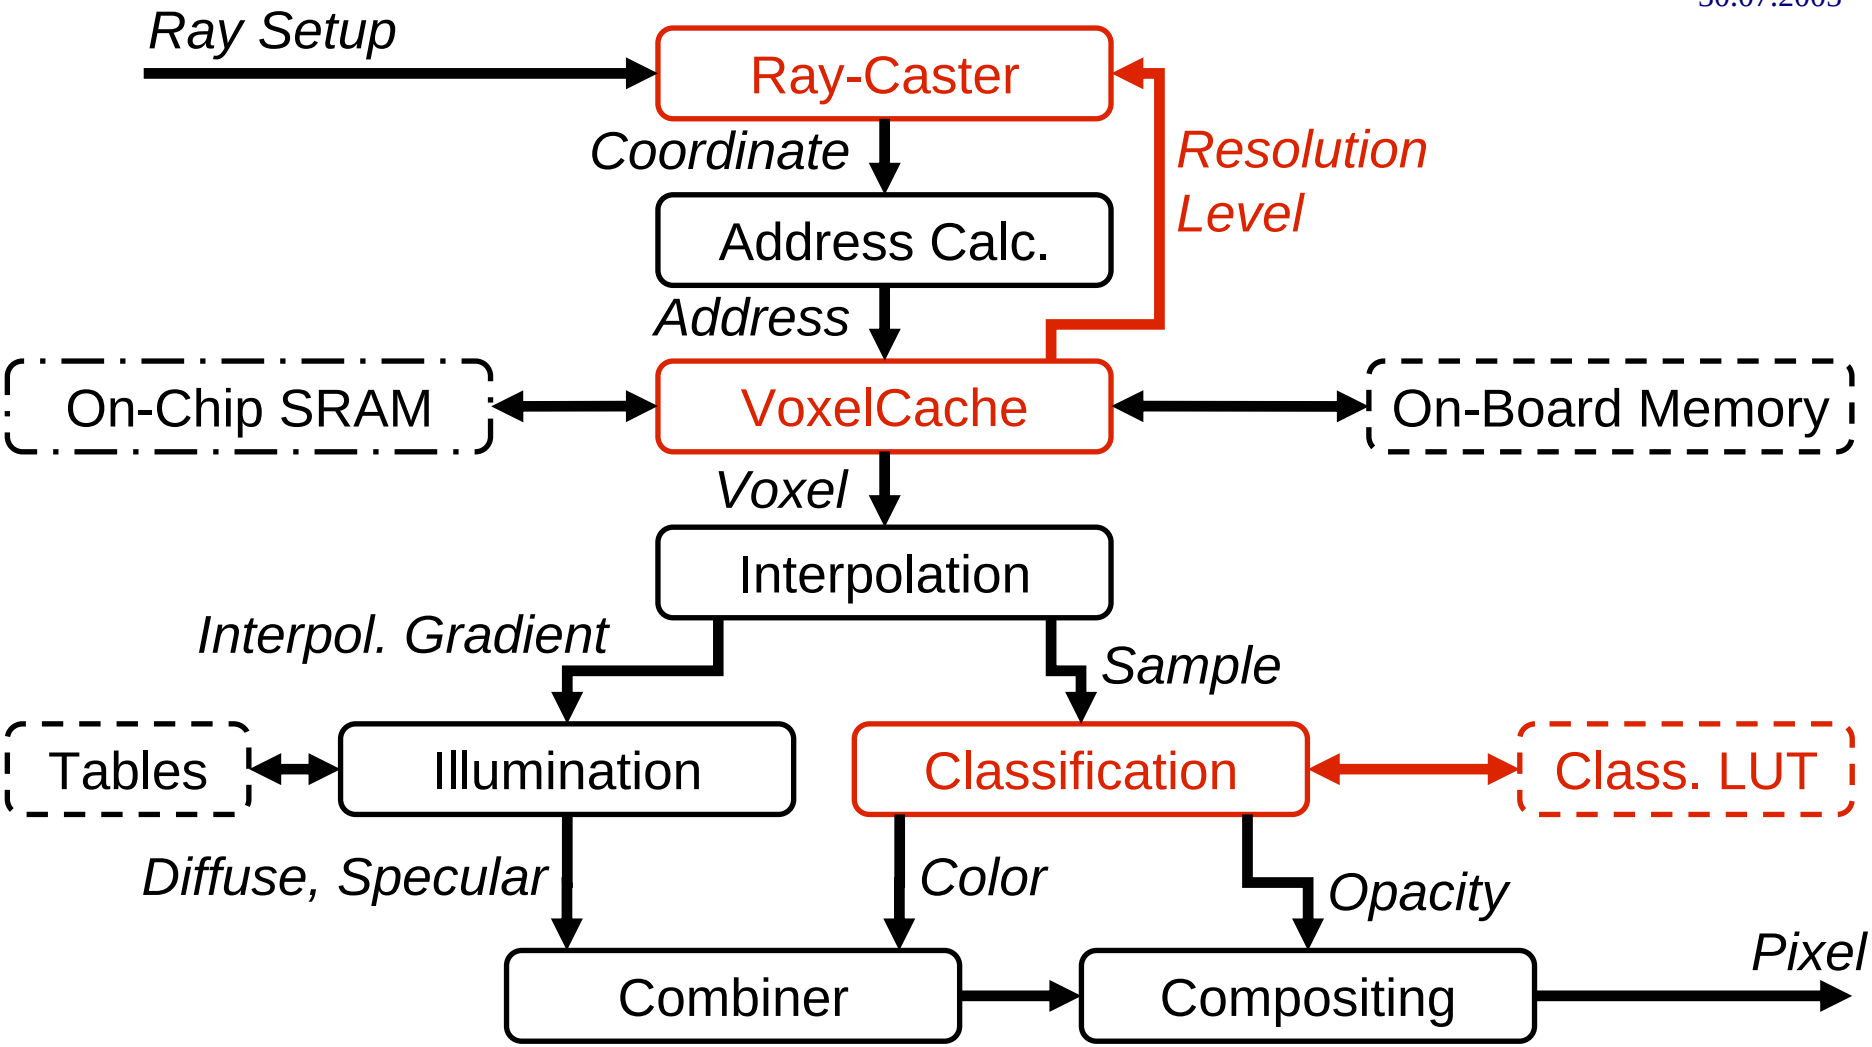

Adjustment of the sampling distance between different hierarchy levels: $\vec{p}_{j+1} = \vec{p}_j + 2^h * \vec{i}$

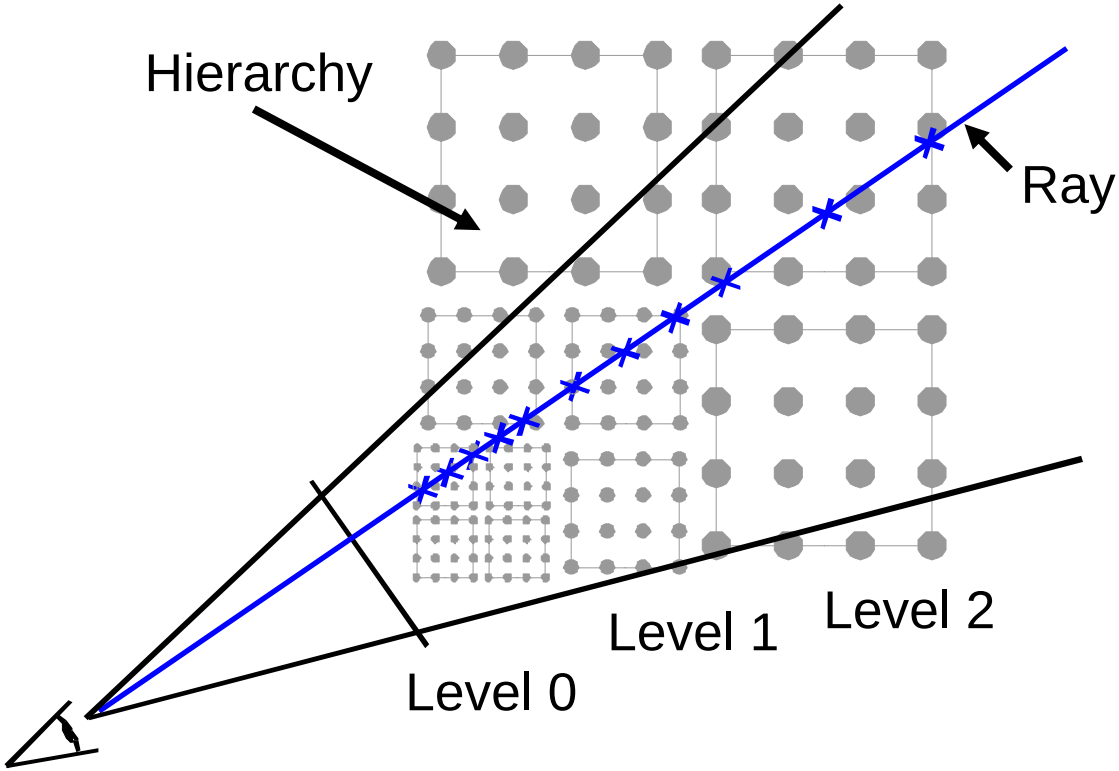

$\vec{p}_{j+1}$: next sample

$\vec{p}_j$: actual sample

h: resolution level

$\vec{i}$: user defined distance

Experiments

Target platform:

Avnet Virtex-II development Kit

- Virtex-II 2V4000
- 512MB on-board DDR-RAM (block index table, s-blocks)
- 1MB SRAM on an additional expansion card (hierarchy table)

→ 2 raycasting pipelines + 1 wavelet decompression unit OR
1 raycasting pipeline + 2 wavelet decompression units

→ Maximum size of the dataset: 2254^3 voxels
(block index table: 5.33 MB → max. 64,000 blocks)

Experiments

Christmas tree (512x512x999)

Size: approx. 500 MB (lossless compression, 16 bit/voxel)

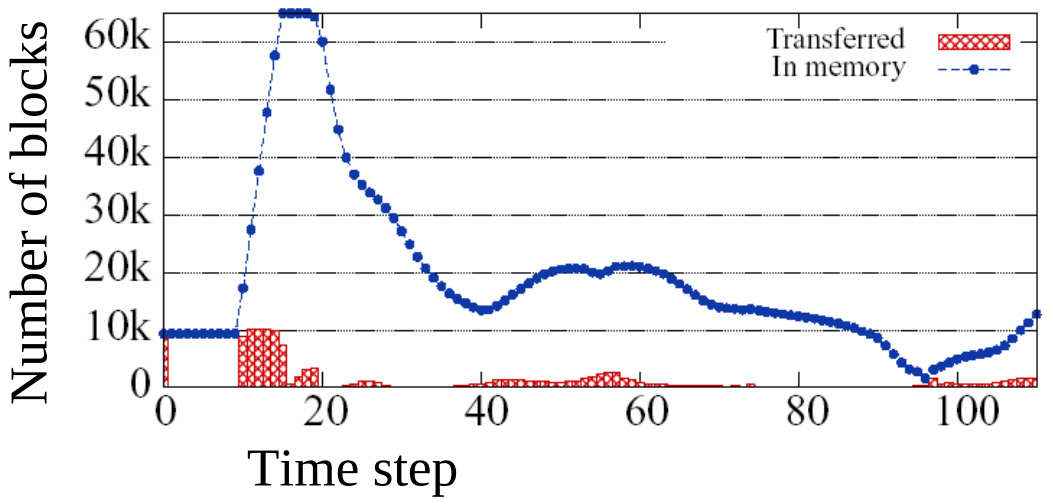

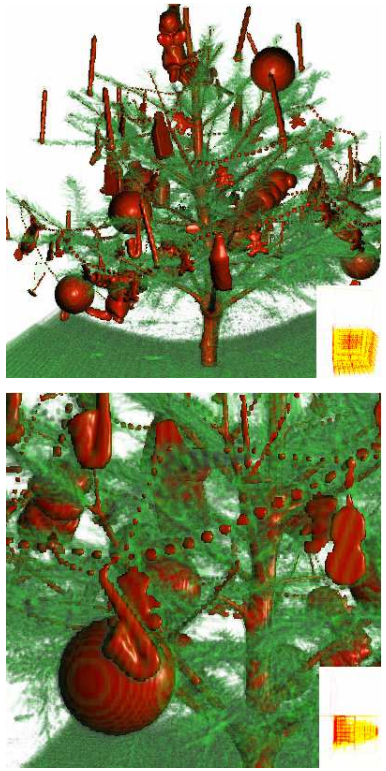

Visible Human – Male (2048x1216x999)

Size: approx. 9 GB (lossless compression, 16 bit/voxel)

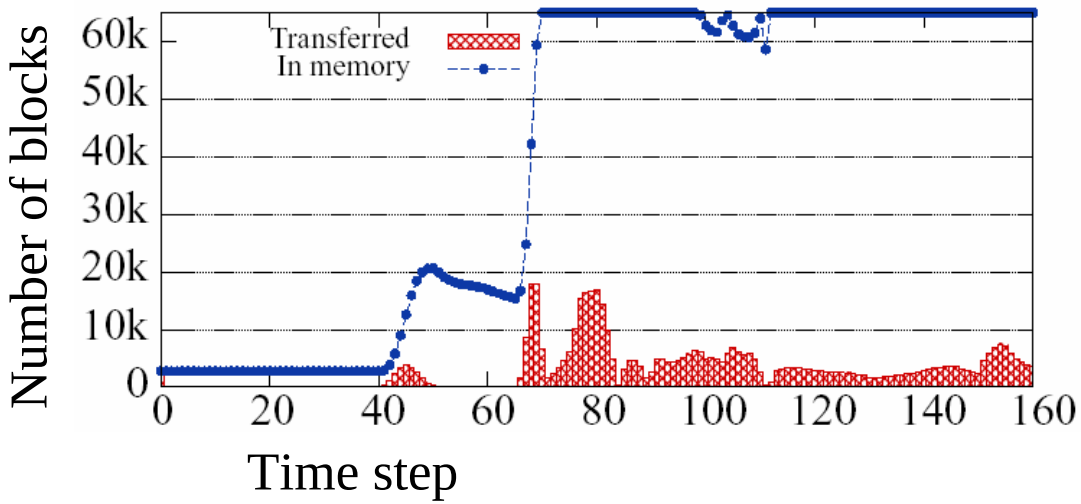

6-8 frames/sec

Conclusions

- Integration of a wavelet-based multi-resolution model in a volume rendering hardware
- Virtualization of the memory interface
- Reduced requirements to memory size and bandwidth
- High bandwidth of decompression due to a full implementation in hardware

Thanks for your attention...

...Questions?