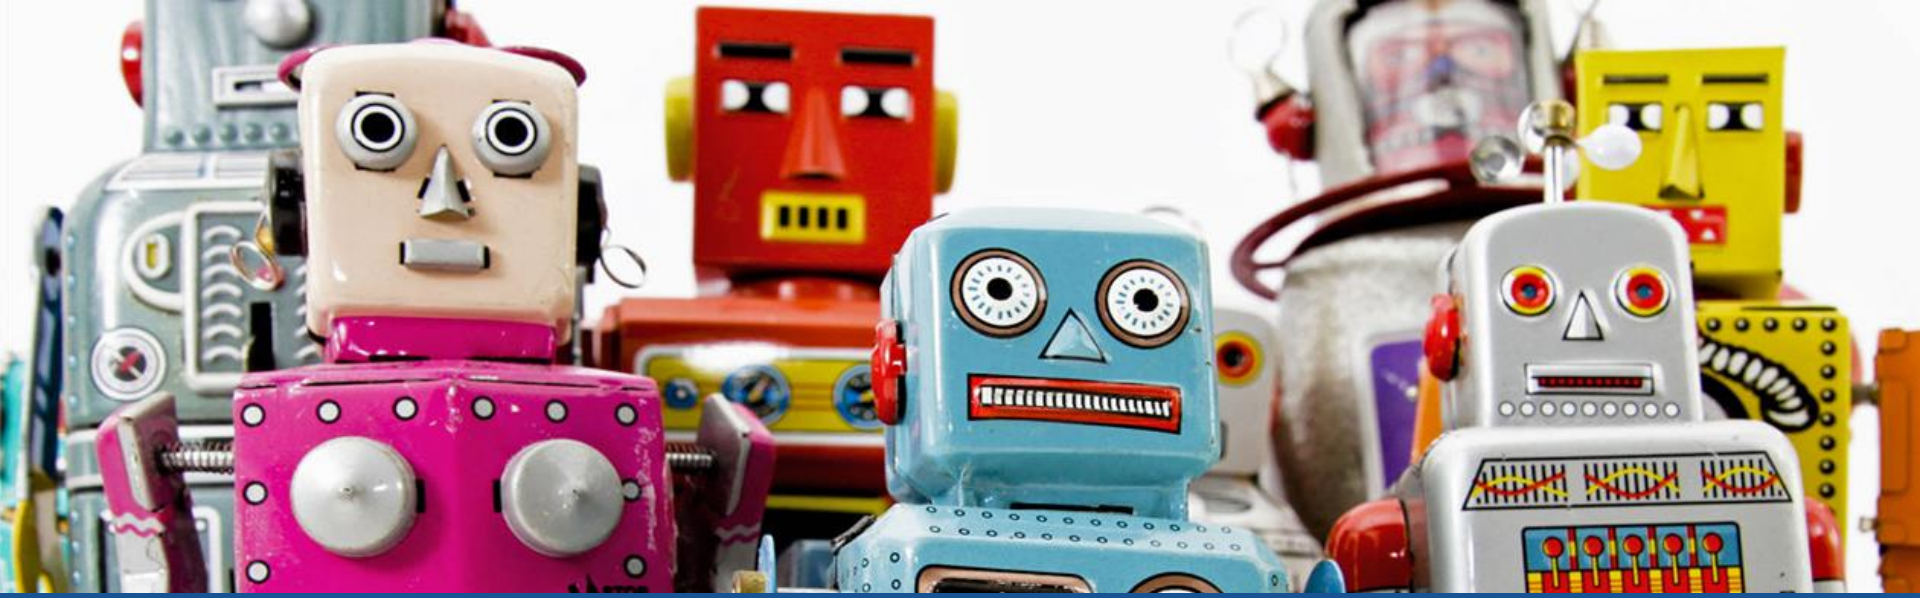

Movidius VOLA

HW Acceleration for Volumetric Applications

David Moloney PhD, CTO Movidius

Introduction

- 3D sensing becoming more and more pervasive
- Started with Kinect
- 3D mapping devices like Project Tango
- Next step from active to passive sensing
- From depth you can reconstruct volume
- And volumetric representations allow many applications

Memory Efficient Volumetric Data

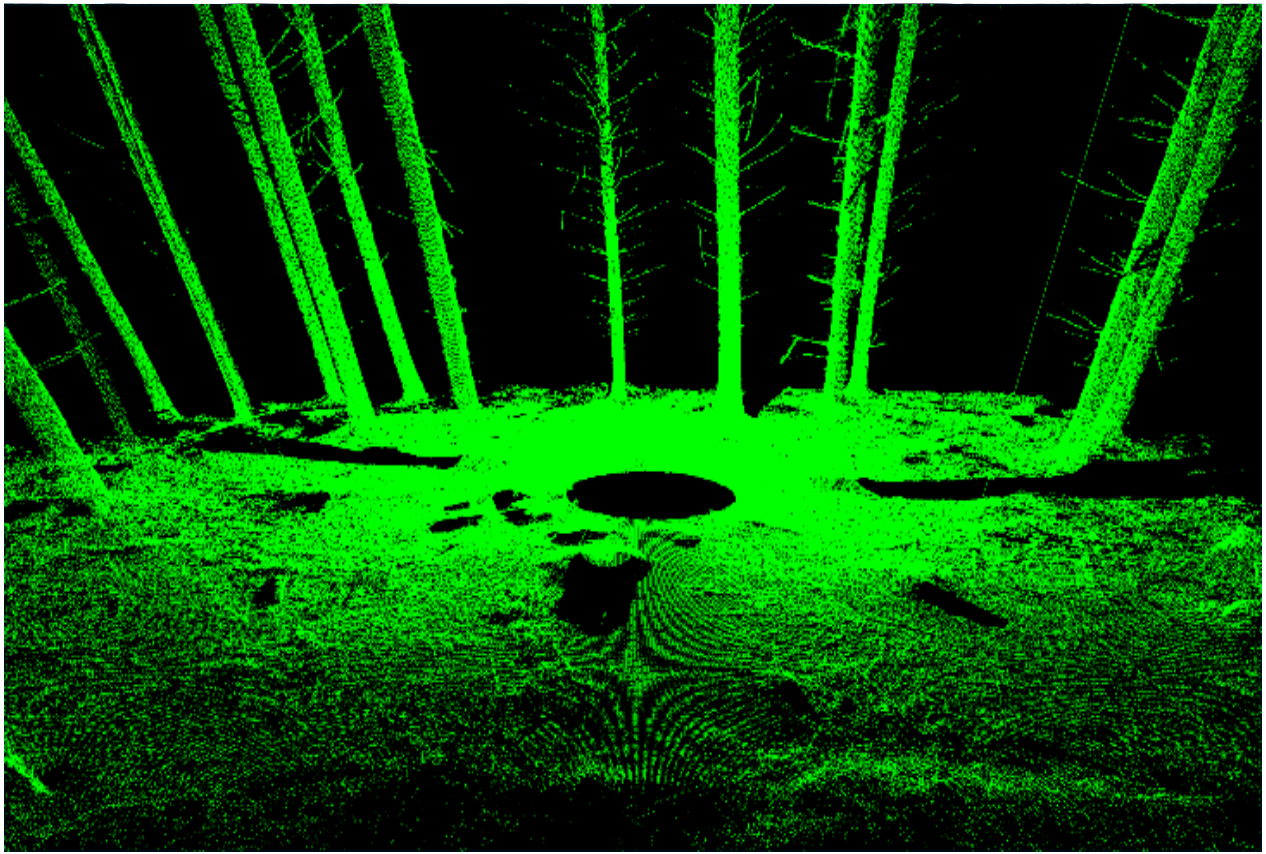

SLAMbench (Platform-Independent Kfusion)

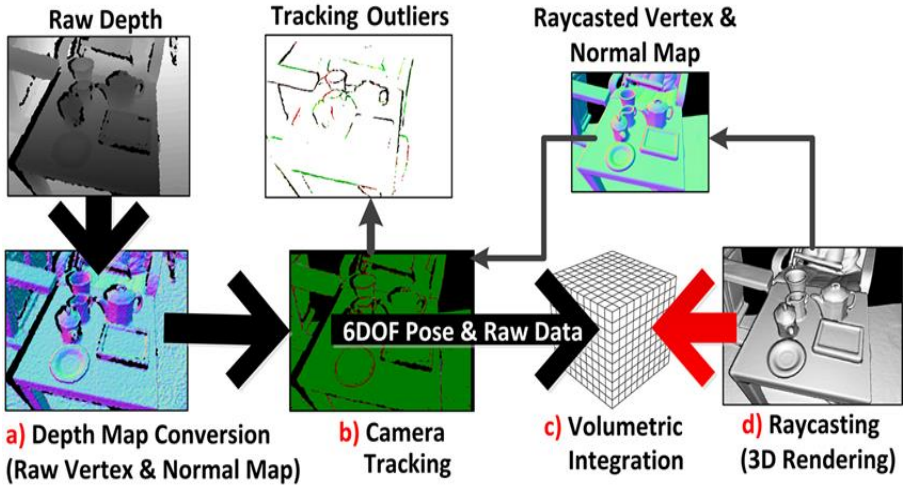

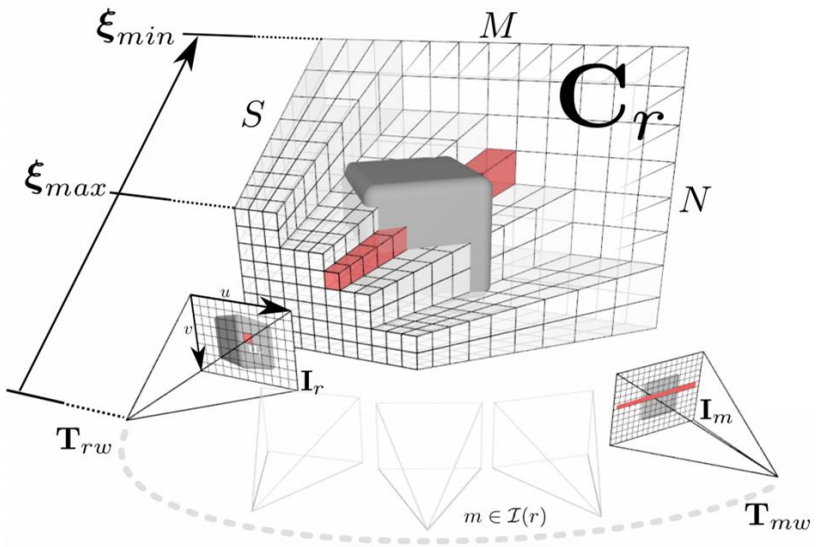

- Kfusion fuses depth frames (Kinect) into 3D geometric map
- Uses voxel grid of TSDFs to represent 3D surfaces
- 3D surfaces recovered by ray-casting at TSDF zero-crossings
- Localisation estimates location and pose

SLAMbench on Myriad2

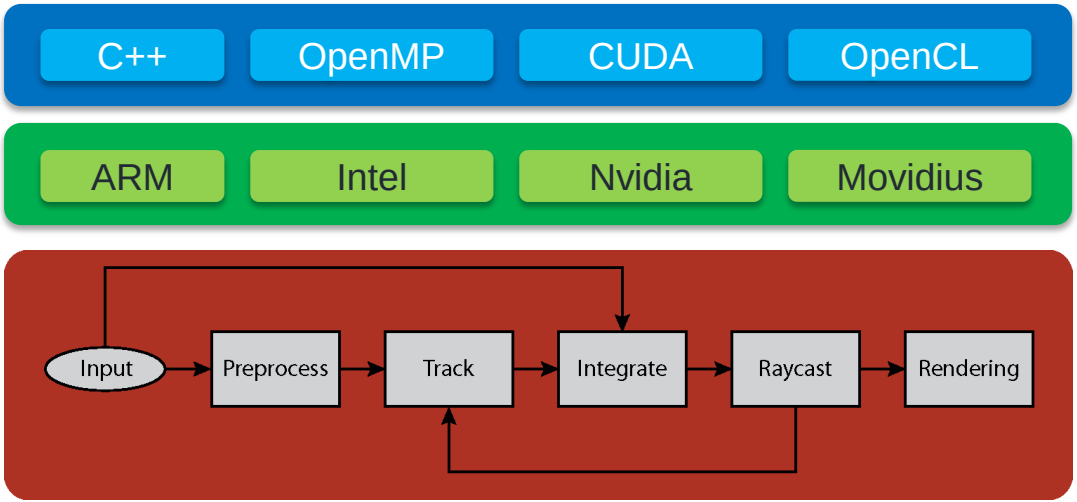

	ICL-NUIM Living Room, traj2: synthetic (QVGA)								ASUS Xtion Pro - real (QVGA)							
	128^3				256^3				128^3				256^3			
energy/frame	CPP	OMP	CUDA	OCL	CPP	OMP	CUDA	OCL	CPP	OMP	CUDA	OCL	CPP	OMP	CUDA	OCL
TK1	4.36	2.01	0.28	0.00	6.21	3.12	0.44	0.00	3.27	1.71	0.22	0.00	4.37	3.08	0.33	0.00
TX1	3.10	1.52	0.17	0.00	4.53	2.47	0.24	0.00	2.29	1.23	0.17	0.00	3.14	1.77	0.21	0.00
XU4	3.43	2.72	0.00	0.48	5.38	5.19	0.00	0.76	2.54	2.50	0.00	0.45	3.74	3.23	0.00	0.65
MA2150	0.59	0.00	0.00	0.00	1.52	0.00	0.00	0.00	0.47	0.00	0.00	0.00	1.03	0.00	0.00	0.00

SLAMbench on Myriad2

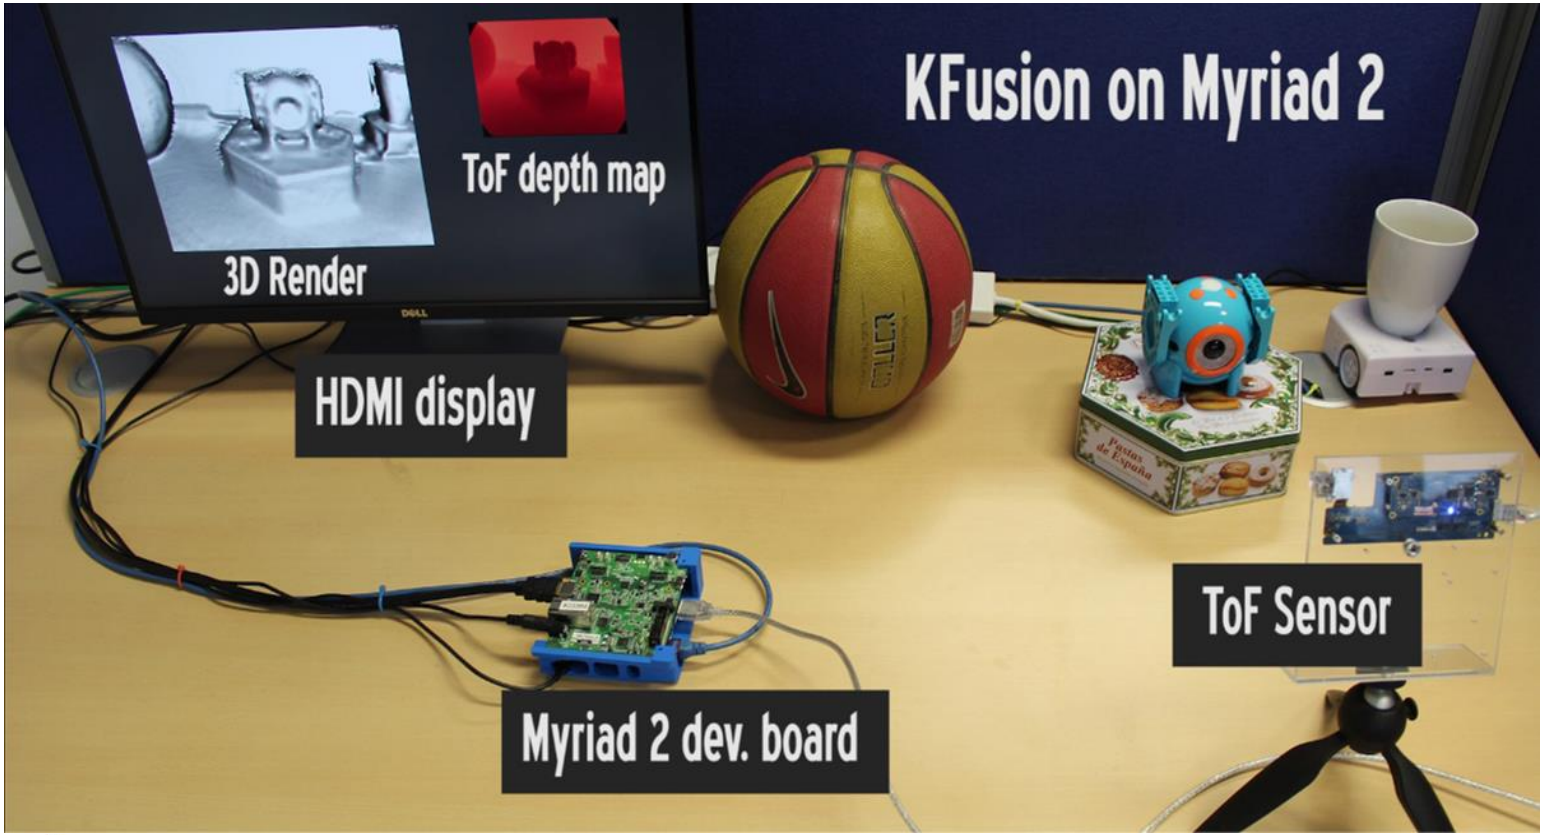

SLAMbench on Myriad2

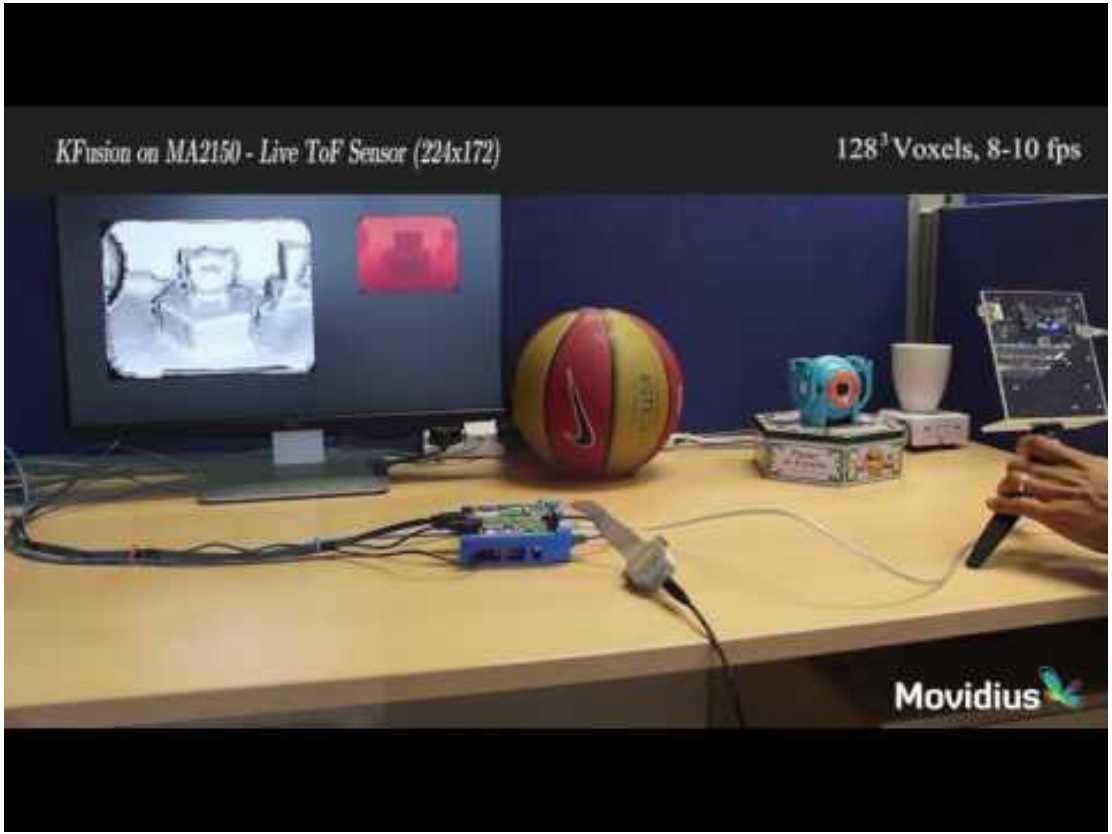

Dense vs Sparse Volumetric storage

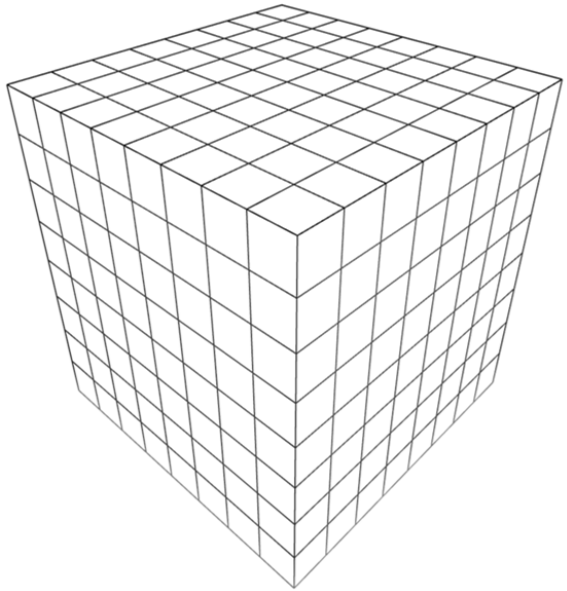

Dense 512MB storage for 5^3 m volume in SLAMbench using 32-bit TSDF per voxel

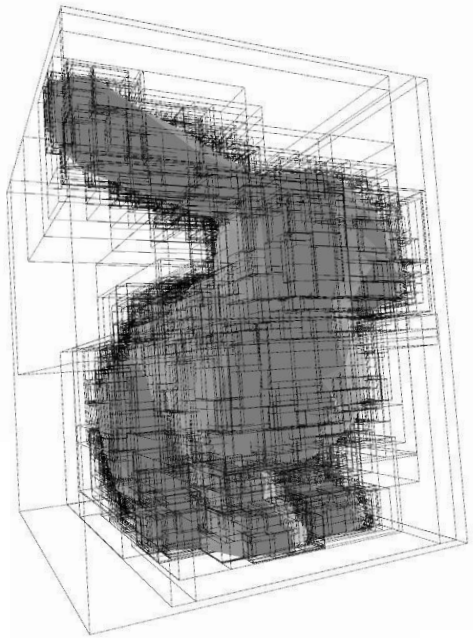

Octree storage – only store 2.5D manifold (no empty space)

Sparse Voxel Tree LoD

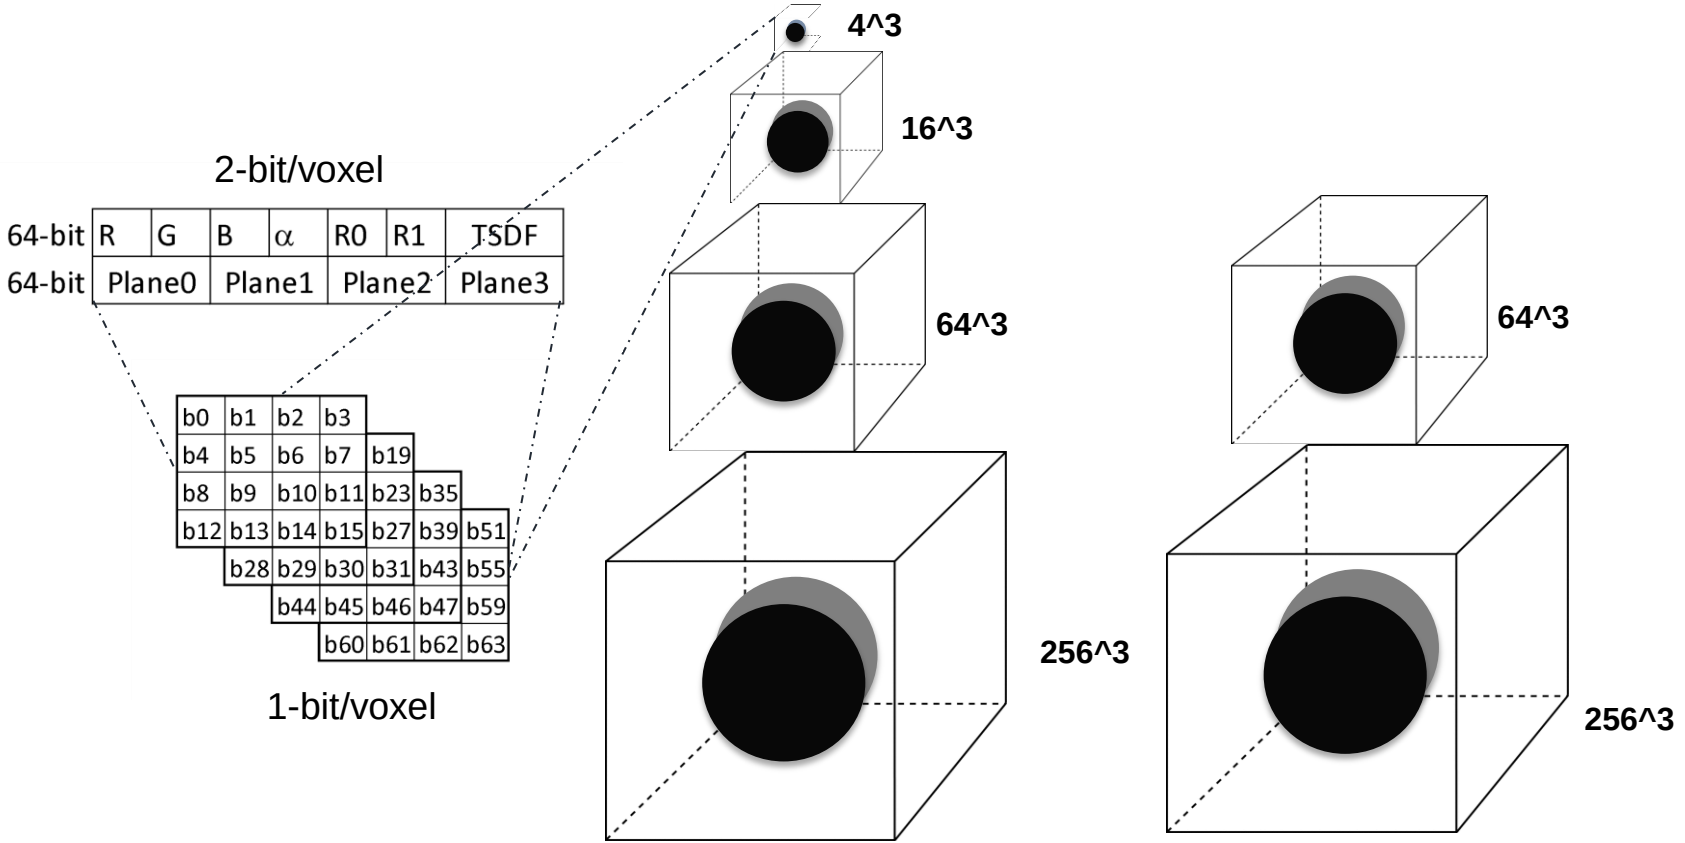

SLAMbench

VOLA Bit Per Voxel Map

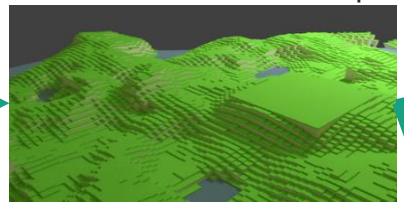

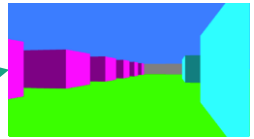

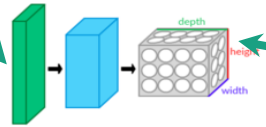

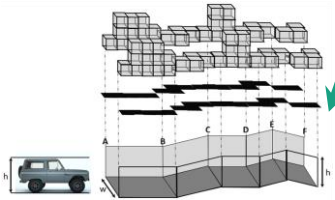

point
insertion/deletion/testing

Collision detection

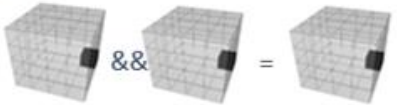

point
deletion

2D projections

2/3D histogram pyramid

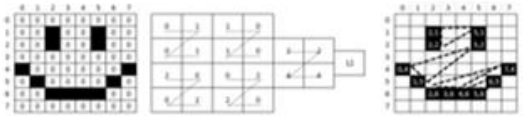

2/3D path-finding

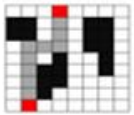

Volumetric Data Sharing

Platin Cement works

Map
https://goo.gl/iKXhQo
800 x 600m (0.5km^2)
raw .obj file 32MB

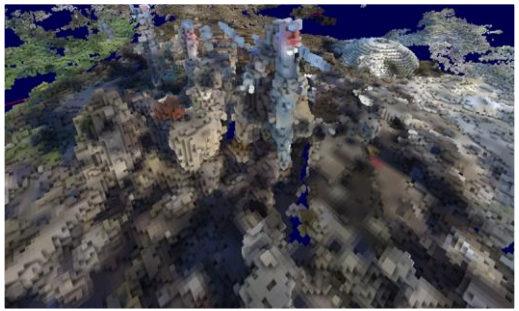

128x compression

SfM video
https://www.youtube.com/watch?v=MZ583jQZSR4

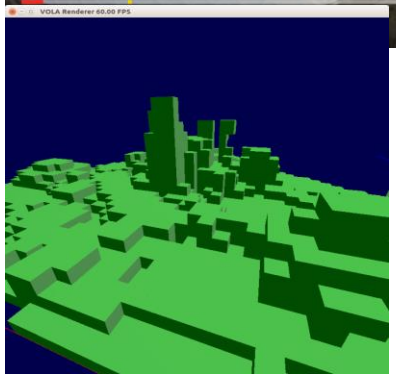

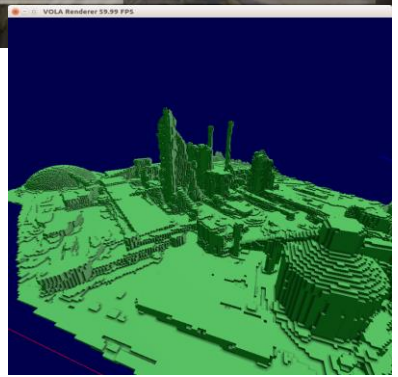

Movidius

Volumetric Data Sharing

VOLA Volumetric CNN

- CNN to identify objects from their VOLA volumetric representation
- Allows objects to be located and marked in VOLA
- Proof of concept using 3D letters ala LeNet achieves 82% accuracy

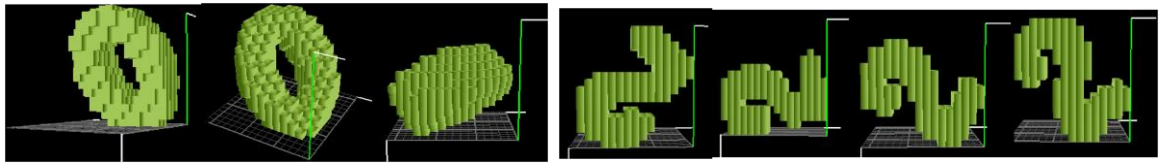

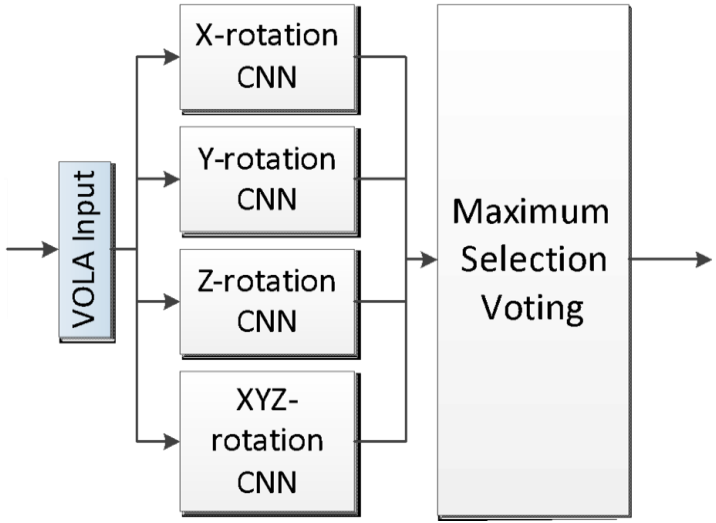

3D font CNN	Test Accuracy
X-rotation	82.95%
Y-rotation	84.70%
Z-rotation	77.33%
XYZ-rotation	82.43%
Average	81.85%

VOLA JiT Audio Models for AR/MR

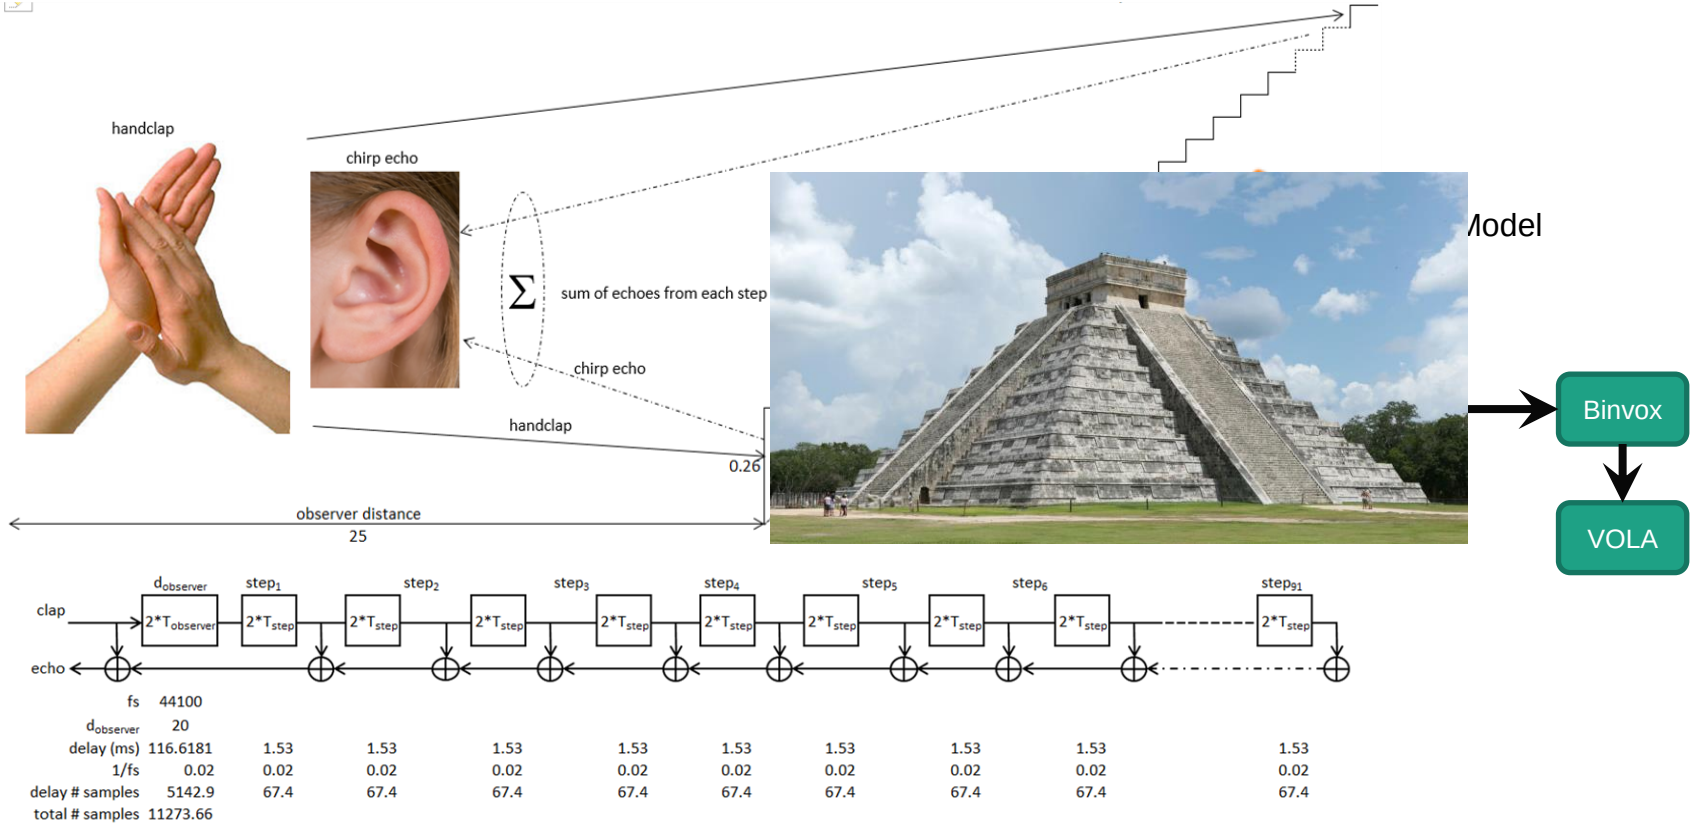

fs	44100									
$d_{observer}$	20									
delay (ms)	116.6181	1.53	1.53	1.53	1.53	1.53	1.53	1.53	1.53	1.53
1/fs	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
delay # samples	5142.9	67.4	67.4	67.4	67.4	67.4	67.4	67.4	67.4	67.4
total # samples	11273.66									

Crowd-Sourcing Volumetric Maps

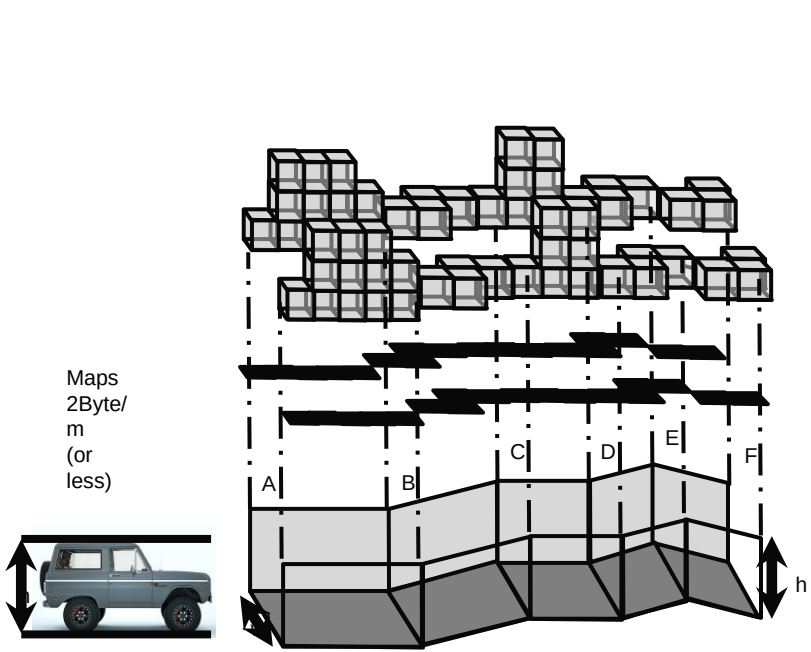

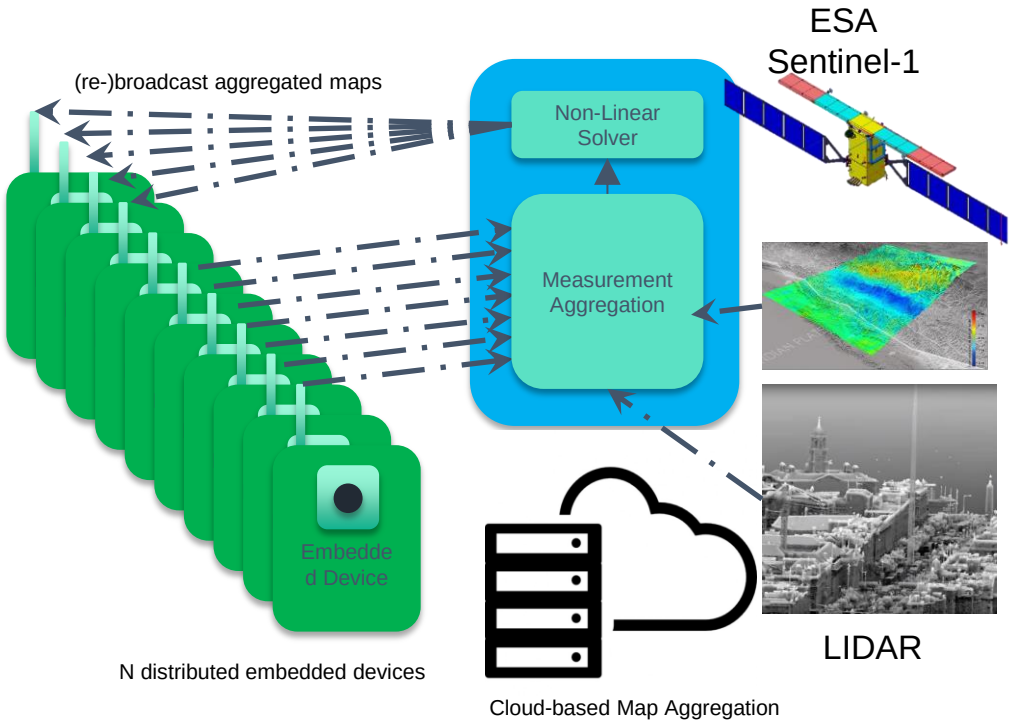

SHAVE ISA Volumetric Data Accelerator

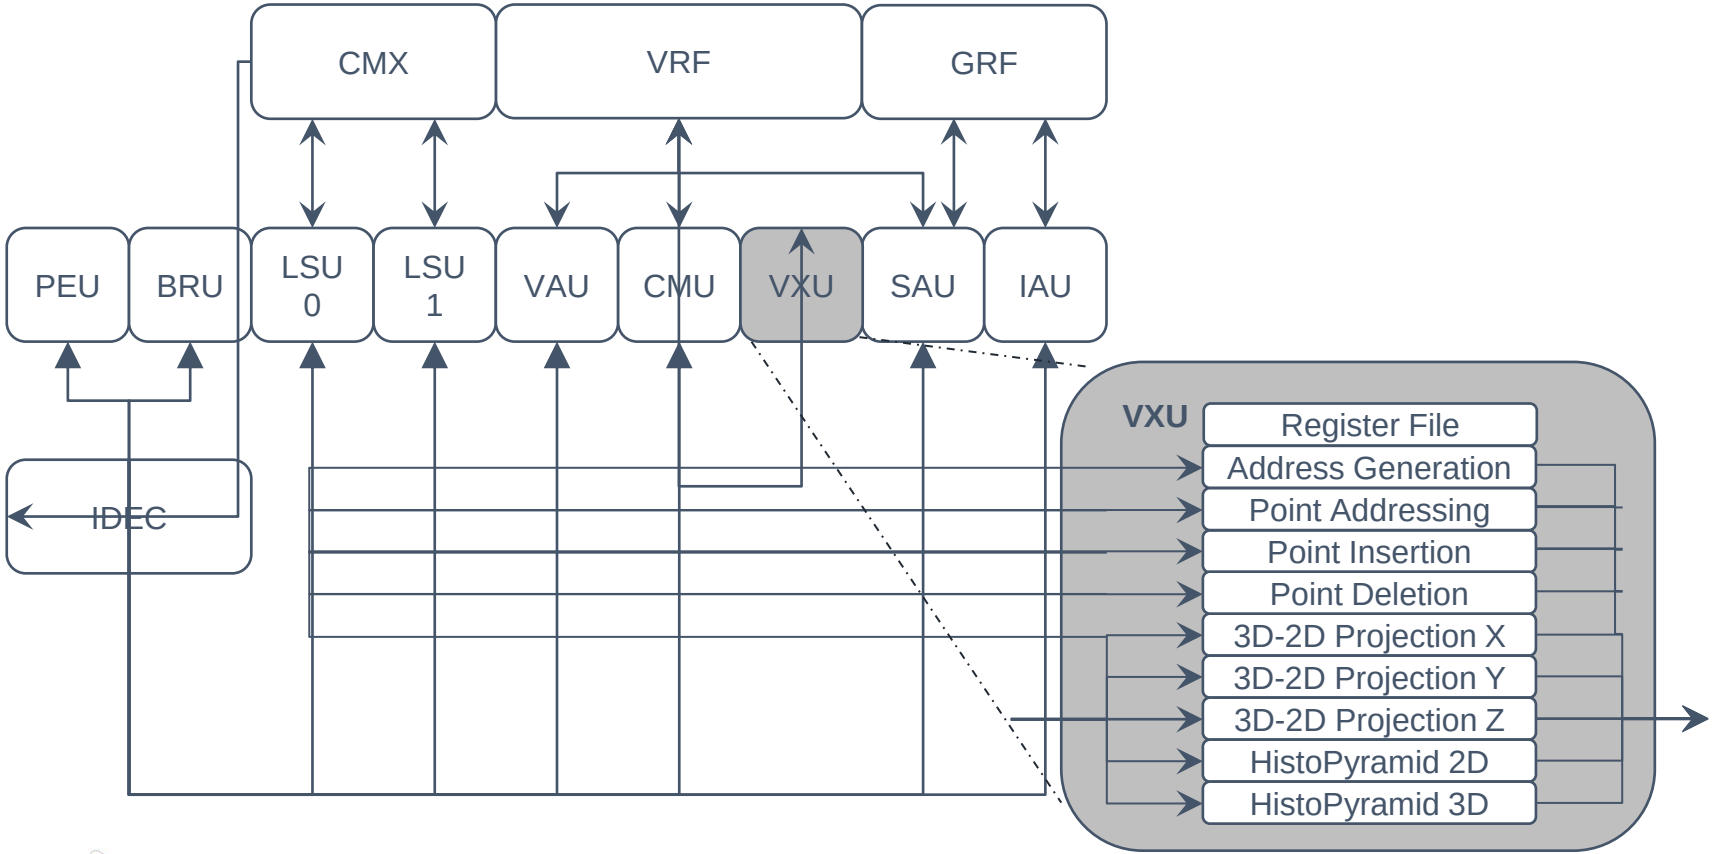

Voxel Insertion/Deletion Logic

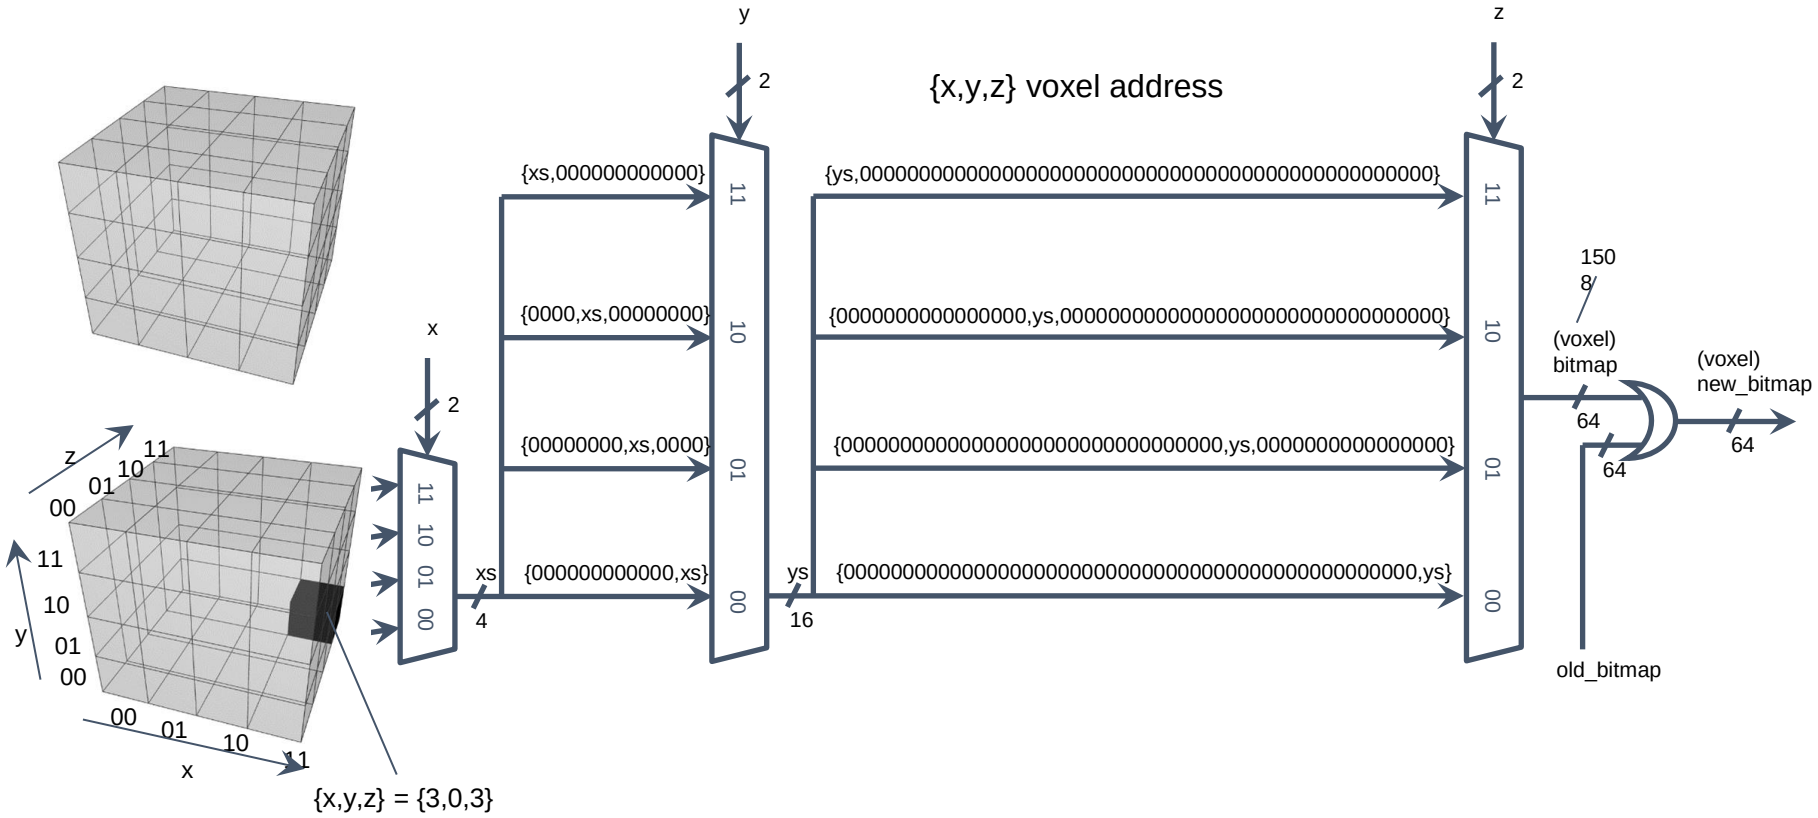

Conclusions

- Volumetric applications can run efficiently on embedded platforms
- Optimal data-structures can allow 128x reduction in RAM requirements
- Bit-per-voxel Octree allows compact interchangeable format for M2M
- Two bit's per voxel allows colour and other information to be stored per sub-volume