

The logo for VIS 2019 features a stylized blue network of nodes and connections on the left, followed by the text "VIS 2019" in a large, blue, sans-serif font. The "VIS" part is in a lighter blue shade, while "2019" is in a darker blue.

# VIS 2019

## Interactive Visualization and On-Demand Processing of Large Volume Data: A Fully GPU-Based Out-Of-Core Approach.

---

Jonathan Sarton - Nicolas Courilleau - Yannick Remion - Laurent Lucas

CReSTIC – Université de Reims Champagne-Ardenne – France

ICube – Université de Strasbourg – France



**3DNEUROSECURE**  
calcul intensif et simulation numérique

Université

de Strasbourg

# Introduction


---

# Background and motivations

## Large volume data, how to

- interactively visualize them
- process them on-the-fly ?

→ interesting to use **GPUs** !



# Background and motivations


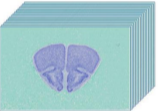
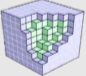
## Large volume data, how to

- interactively visualize them
- process them on-the-fly ?

→ interesting to use **GPUs** !


### Issue : memory occupation

- Large datasets
- $\gg$  GPU and CPU physical memory !
- Interactive manipulation complicated

| Domain/Application   | Data size                    |
|--|------------------------------|
| Mesh voxelization                 | ~ 100 GB                     |
| Histology<br>Electron microscopy  | ~ 100 GB<br>to<br>several TB |
| Regular 3D grid                   | And beyond                   |

→ **Elaborate out-of-core algorithms**

# Out-of-core data access





GPU data cache

+


Octree

Or



Multi-resolution  
Page Table



Gigavoxels




[Crassin et al., ACM SIGGRAPH i3D, 2009]



[Hadwiger et al., IEEE SciVis 2012]

# Out-of-core data access

Better for very large volume !!





GPU data cache

+


Octree

Or

Multi-resolution  
Page Table



Gigavoxels




[Crassin et al., ACM  
SIGGRAPH i3D, 2009]



[Hadwiger et al., IEEE SciVis 2012]

# Data representation and storage

- **Multi-resolution:** to choose the desired level of detail  
⇒ Reduces the amount of data




# Data representation and storage

## 3D mipmap

- **Multi-resolution:** to choose the desired level of detail
  - ⇒ Reduces the amount of data
- **Bricking:** Volume subdivided into small bricks (e.g  $32^3$ ,  $64^3$ ).
  - ⇒ Allows the out-of-core approach


## Data compression with LZ4 algorithm

- Loss less
- Good compression ratio
- Real-time decompression







# Multi-resolution, multi-level page table hierarchy



# Multi-resolution, multi-level page table hierarchy




# Multi-resolution, multi-level page table hierarchy



- One page = 3D coordinates of the bloc in the next cache level + one flag:
  - Mapped
  - Unmapped
  - Empty
- Implementation: CUDA  
3D Textures
- Cache replacement algorithm:  
Least Recently Used (LRU)

# Virtual addressing



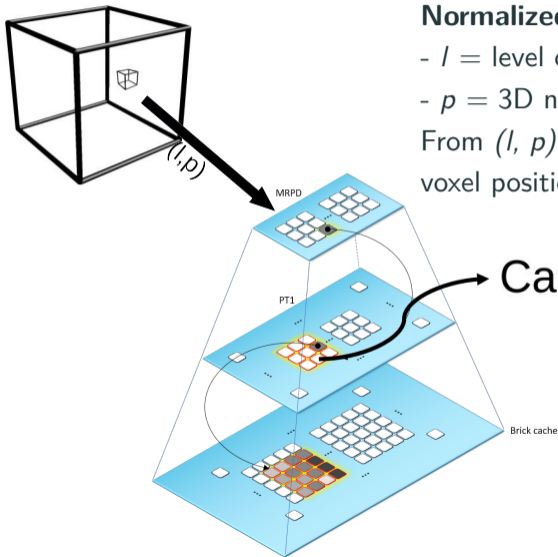
**Normalized volume navigation**  $\rightarrow$  address  $(l, p)$

-  $l$  = level of detail

-  $p$  = 3D normalized position  $(x, y, z) \in [0, 1]^3$

From  $(l, p)$  address, we get the corresponding 3D voxel position into the brick cache.

# Cache miss



**Normalized volume navigation**  $\rightarrow$  address  $(l, p)$

-  $l$  = level of detail

-  $p$  = 3D normalized position  $(x, y, z) \in [0, 1]^3$

From  $(l, p)$  address, we get the corresponding 3D voxel position into the brick cache.

**Cache miss !**

**How to allow on-demand processing of  
any part of a large volume during its visualization ?**


1. Cache usage updates
2. Brick requests management

## **A GPU data structure fully managed on GPU**

### **Advantages**

- Avoids many data transfers between CPU and GPU
- Take advantage of the massively parallel environment of GPUs
- Free the CPU for other eventual processing

# Brick request management on GPU




- Size = number of bricks in the **multi-resolution volume**
- Marked with a timestamp



## GPU → CPU communications

A simple list with the requested brick IDs



## GPU ← CPU communications

Only the bricks ! (With CUDA Zero Copy)

**Model in action: interactive  
visualization & on-demand  
processing on GPU**

---



# Out-of-core virtual microscope

## Virtual microscope

2D multi-resolution visualization of a high resolution image stack.

Interactive navigation:

- move and zoom in a slide
- navigate through the volume from slide to slide



# Out-of-core virtual microscope

## Virtual microscope ...


2D multi-resolution visualization of a high resolution image stack.

Interactive navigation:

- move and zoom in a slide
- navigate through the volume from slide to slide

## + on-demand processing

Region-growing from a voxel selected by the user in the screen space



## Virtual microscope ...


2D multi-resolution visualization of a high resolution image stack.

Interactive navigation:

- move and zoom in a slide
- navigate through the volume from slide to slide

## + on-demand processing

Region-growing from a voxel selected by the user in the screen space




**Cache miss due to processing outside the screen space !**

## Out-of-core virtual microscope

Electron microscopy dataset

$4096 \times 3072 \times 2130$  8bits  $\approx$  **27 GB**

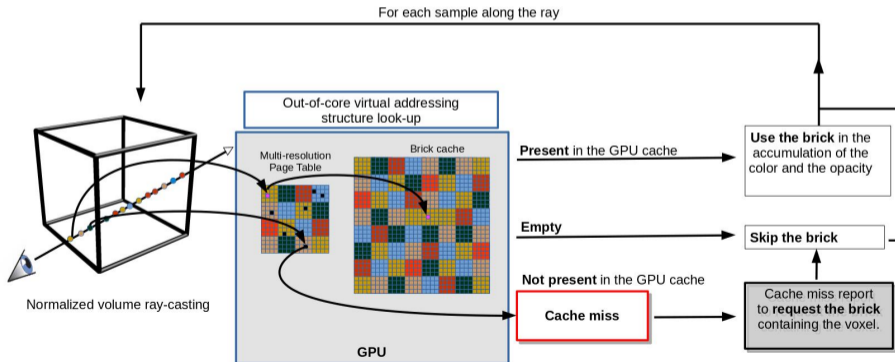
Rendering performance:  $\approx$  **250 FPS**




# Out-of-core Direct Volume Rendering

## Ray-guided approach

- Intuitive visibility selection: no additional culling calculation
- Intuitive out-of-core integration: only load visible bricks on GPU cache




# Datasets



Primate hippocampus

Light sheet microscope

$2160 \times 2560 \times 1072$  16bits  $\approx$  **12 GB**



Mouse brain


Histological scanner

$64000 \times 50000 \times 114$  RGBA  $\approx$  **1.5 TB**




# Performances – frames frequency


On a single workstation  
NVidia GeForce Titan X



Primate hippocampus




Mouse brain




# Performances – frames frequency


On a single workstation  
NVIDIA GeForce Titan X



Primate hippocampus



Mouse brain



- Primate hippocampus ( $2160 \times 2560 \times 1072 \approx \mathbf{12\ GB}$ )
  - Brick size:  $64^3 \implies \approx 27000$  bricks (7 LOD)
  - One virtualization level  
→ Need **1.2 MB** on GPU
- Mouse brain ( $64000 \times 50000 \times 114 \approx \mathbf{1.5\ TB}$ )
  - Brick size:  $64^3 \implies 3.13$  million bricks (10 LOD)
  - One virtualization level →  $\approx \mathbf{63\ MB}$  needed on GPU
  - Two virtualization levels →  $\approx \mathbf{13\ MB}$  needed on GPU

## Conclusion

---

- Out-of-core data management: multi-resolution multi-level page table hierarchy
  - Entirely managed on GPU
  - GPU – CPU communication reduced
  - Good rendering frequency even for very large volume of data ( $>$  TB)
  - Weak GPU memory and computational footprint
  - General purpose context : interactive visualization & on-demand processing

The logo for VIS 2019 features a stylized network of blue nodes and connections on the left, followed by the text "VIS 2019" in a large, sans-serif font. "VIS" is in a teal color, and "2019" is in a dark blue color.

# VIS 2019

## Interactive Visualization and On-Demand Processing of Large Volume Data: A Fully GPU-Based Out-Of-Core Approach.

Jonathan Sarton - sarton@unistra.fr

Nicolas Courilleau - nicolas.courilleau@neoxia.com

Yannick Remion - yannick.remion@univ-reims.fr

Laurent Lucas - laurent.lucas@univ-reims.fr