
Interactive Level-Set Deformation

On the GPU



Aaron Lefohn

Institute for Data Analysis and Visualization
University of California, Davis

Problem Statement

- Goal

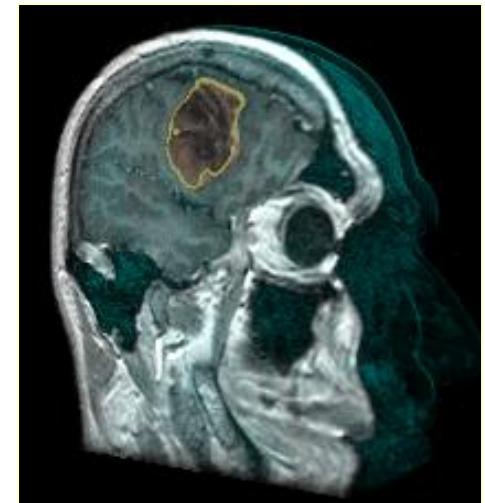
- Interactive system for deformable surface manipulation
- Level-sets

- Challenges

- Deformation is slow
- Deformation is hard to control

- Solution

- Accelerate level-set computation with GPU
- Visualize computation in real-time

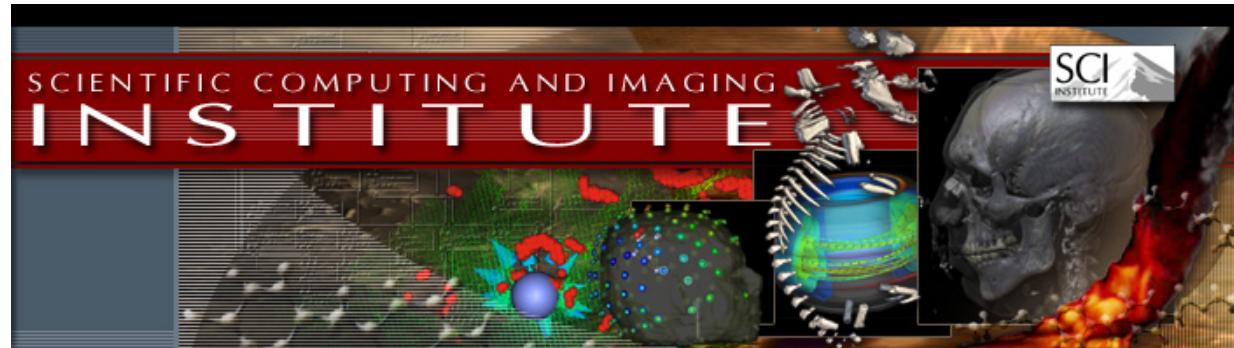


GGPU
Aaron Lefohn
University of California, Davis, USA

J2

20
vis04

Collaborators



University of Utah

Joe Kniss

Joshua Cates

Charles Hansen

Ross Whitaker



GPGPU
Aaron Lefohn
University of California, Davis, USA

J3

20
VIS04

Overview

- Why?
 - Motivation and previous work

- How?
 - Streaming level-set algorithm
 - Real-time visualization
 - Segmentation application

- Does it work?
 - Demonstration
 - User study



GPGPU
Aaron Lefohn
University of California, Davis, USA

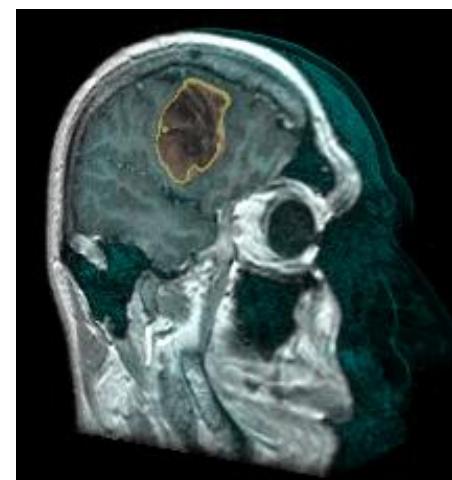
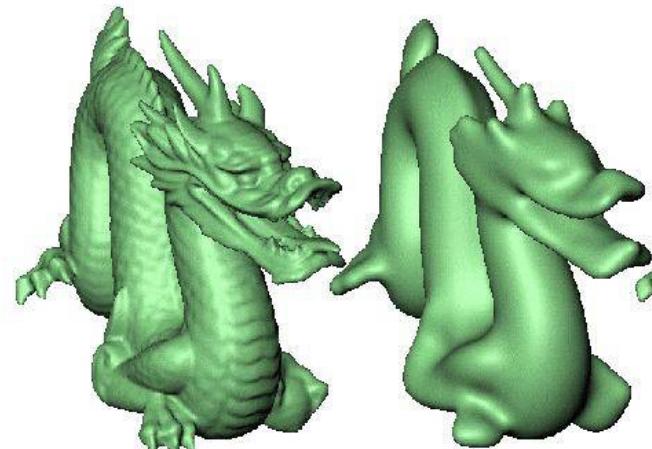
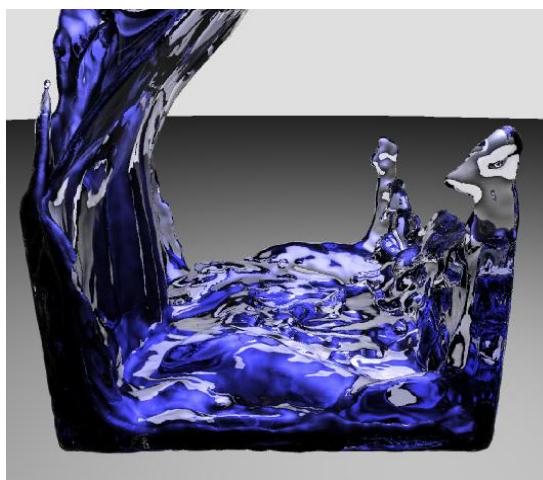
J4

20
vis04

Deformable Surfaces

- Applications of Level-Sets

- Fluid simulation
- Surface reconstruction for 3D scanning
- Surface processing
- Image / Volume segmentation



GPGPU
Aaron Lefohn
University of California, Davis, USA

J5

20
VIS04

Level-Set Method

- Implicit surface

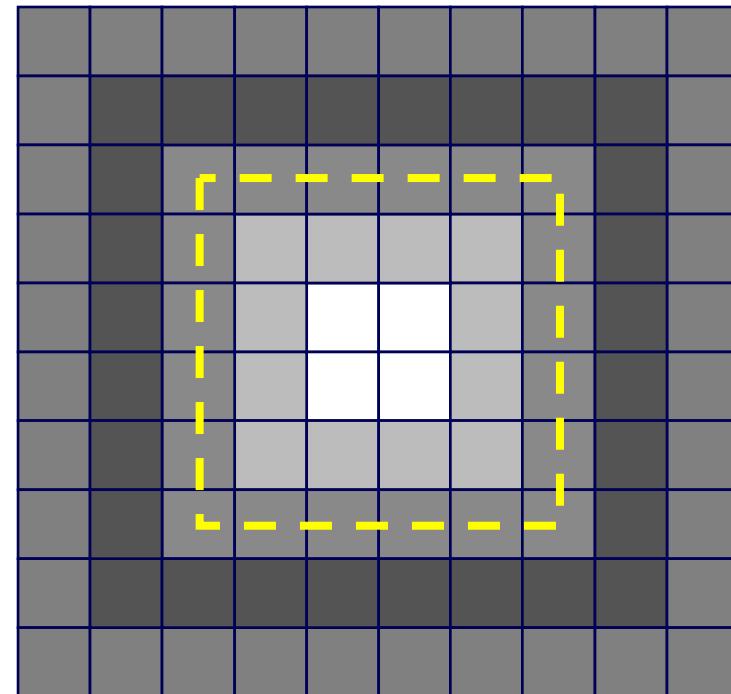
$$\mathbf{S}_t = \{\mathbf{x} | \phi(\mathbf{x}, t) = k\}$$

- Distance transform

ϕ denotes inside/outside

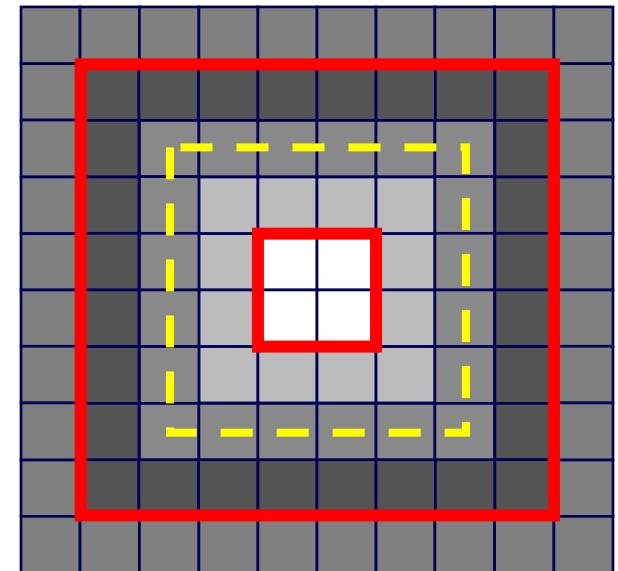
- Surface motion

- $\phi(\mathbf{x}, t + \Delta t) = \phi(\mathbf{x}, t) + \Delta t F |\nabla \phi|$
 - F = Signed speed in direction of normal

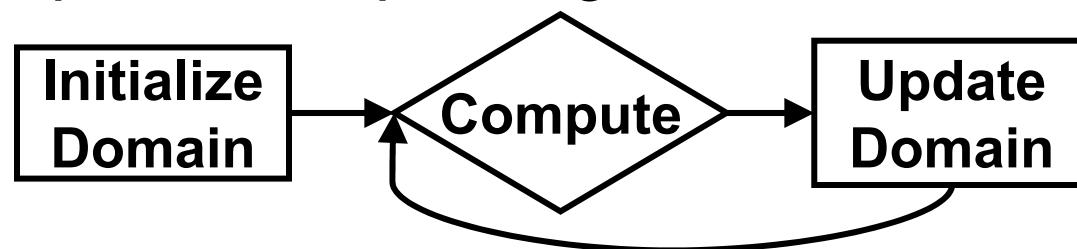


CPU Level-Set Acceleration

- Narrow-Band/Sparse-Grid
 - Compute PDE *only* near the surface
 - Adalsteinson *et al.* 1995
 - Whitaker *et al.* 1998
 - Peng *et al.* 1999



- Time-dependent, sparse-grid solver



GPU Level-Set Acceleration

- Strzodka et al. 2001
 - 2D level-set solver on NVIDIA GeForce 2
 - No narrow-band optimization

- Lefohn et al. 2002
 - Brute force 3D implementation on ATI Radeon 8500
 - No faster than CPU, but ~10x more computations
 - No narrow-band optimization



Overview

- Why?
 - Motivation and previous work
- How?
 - Streaming narrow-band
 - Real-time visualization
 - Segmentation application
- Does it work?
 - Demonstration
 - User study



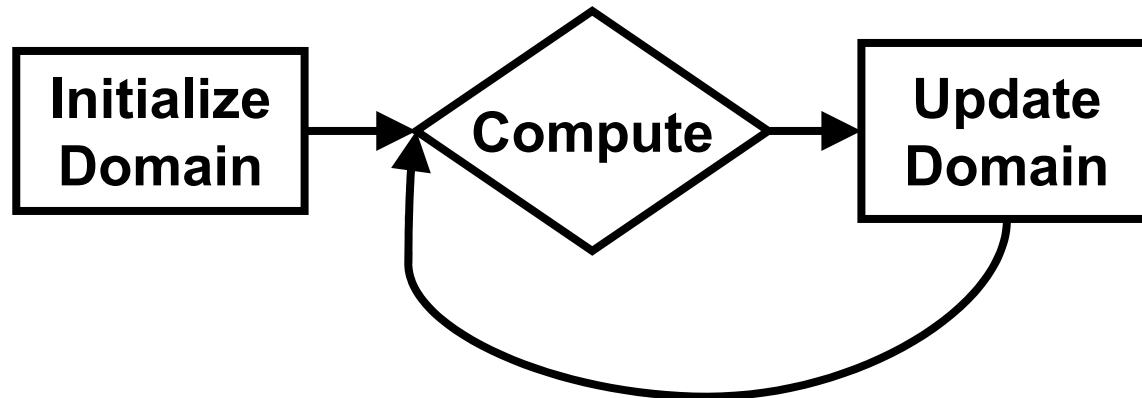
GPGPU
Aaron Lefohn
University of California, Davis, USA

J9

20
vis04

GPU Narrow-Band Solver

- Sparse Volume Computation
 - CPU algorithm: Traverse list of active voxels
 - GPU algorithm: Compute all active voxels in parallel



- Data structures change after each PDE time step



Algorithm Goals

GPU Narrow-Band Solver

- Goals
 1. Leverage GPU parallelism
 2. Perform sparse computation
 3. Minimize GPU memory usage
 4. Fast update of sparse data structures
 5. Interactive visualization



Algorithm Solutions

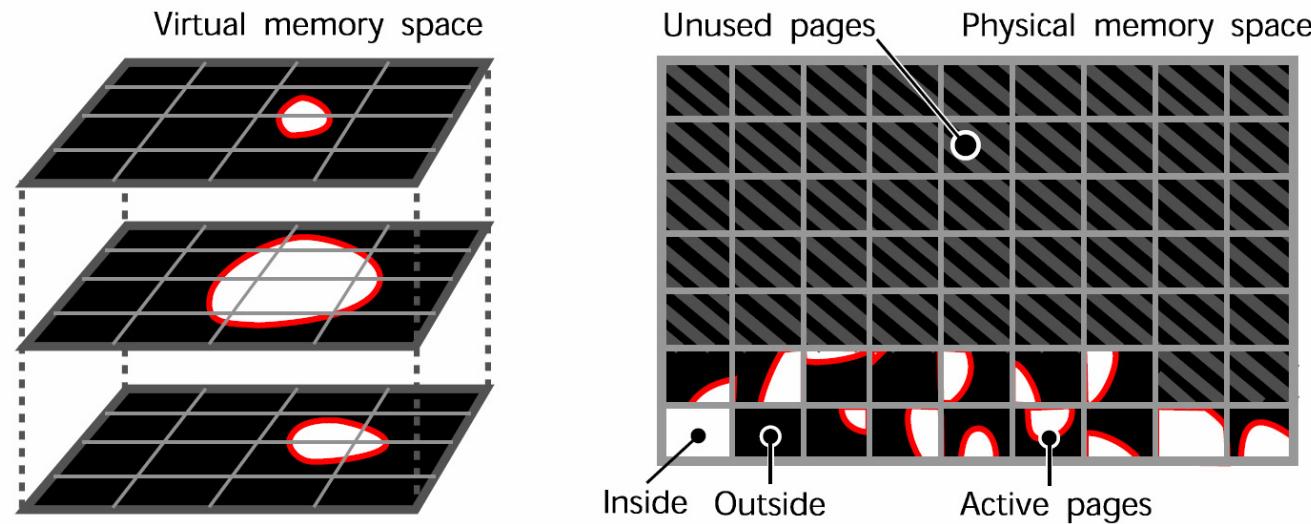
- Pack Active Voxels Into 2D Texture
 - Increase parallelism, reduce computation and memory use
- Efficient GPU-to-CPU Message Passing
 - Fast update of packed data structure
- On-The-Fly Decompression Volume Rendering
 - Interactive visualization without increasing memory use



A Dynamic, Sparse GPU Data Structure

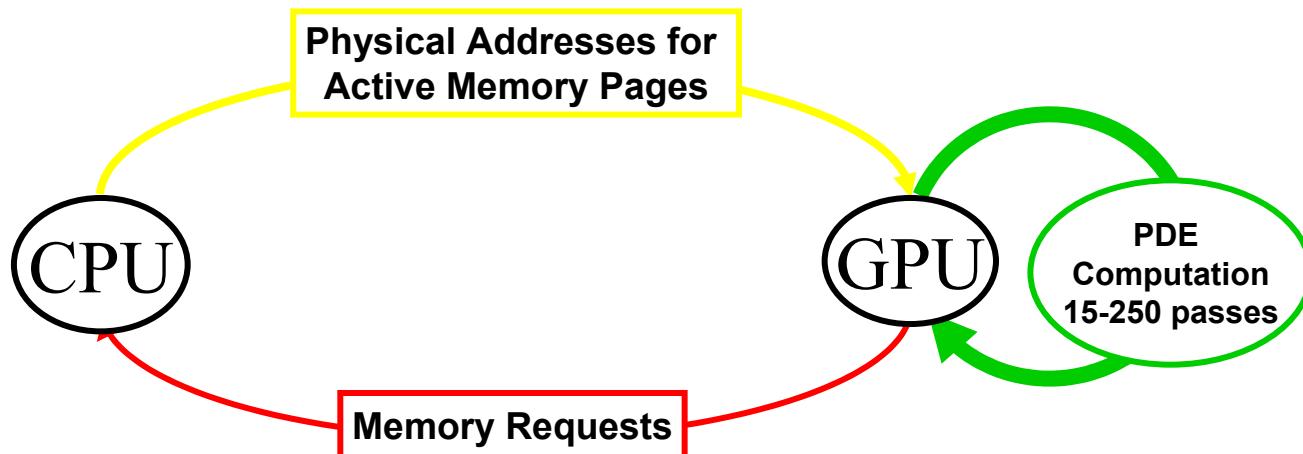
- Multi-Dimensional Virtual Memory

- 3D virtual memory
- 2D physical memory
- 16 x 16 pixel pages



A Dynamic, Sparse GPU Data Structure

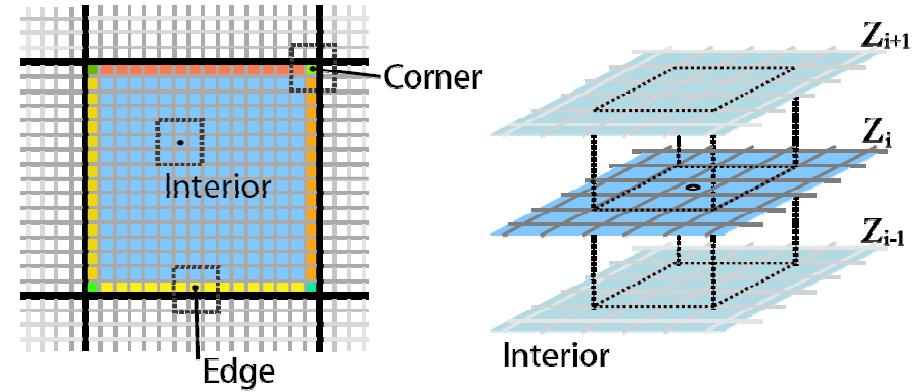
- GPU: Computes PDE
 - Level-set computation (2D physical memory)
 - Issues memory requests
- CPU: Manages memory
 - Memory manager
 - Page table (3D virtual memory)



A Dynamic, Sparse GPU Data Structure

- Problem

- Neighbor lookups across page boundaries
- Branching slow on GPU



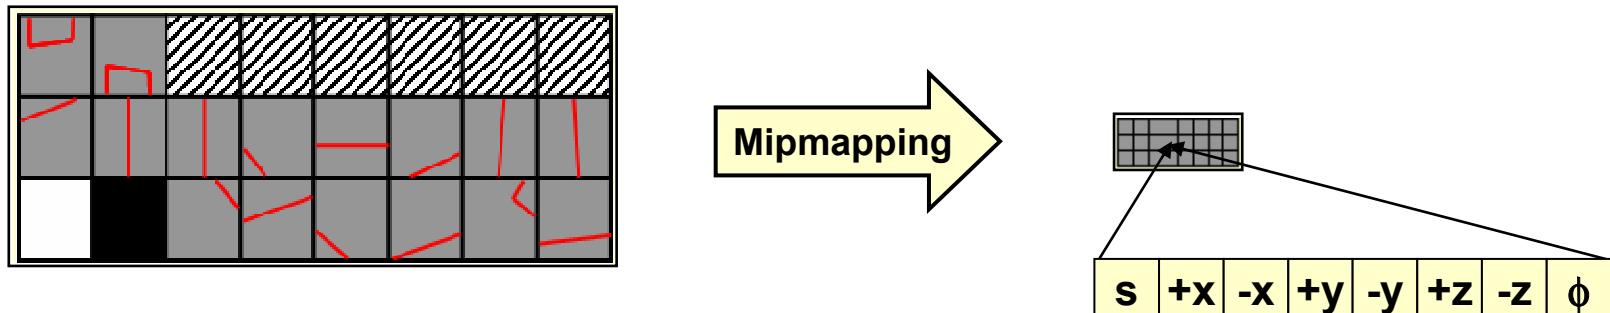
- Solution

- Substreams
 - Create homogeneous data streams
 - Resolve conditionals with geometry
 - Lefohn 2003, Goodnight 2003, Harris 2003
 - Optimizes cache and pre-fetch performance
 - Kapasi et al., Micro 33, 2000



GPU-to-CPU Message Passing

- Problem: Active Voxel Set is Time-Dependent
 - GPU memory request mechanism
 - Low bandwidth GPU-to-CPU communication
- Solution
 - Compress GPU memory request
 - Use GPU computation to save GPU-to-CPU bandwidth



Overview

- Why?

- Motivation and previous work

- How?

- Streaming level-set algorithm
 - Real-time visualization (with Joe Kniss)
 - Segmentation application

- Does it work?

- Live demonstration
 - User study



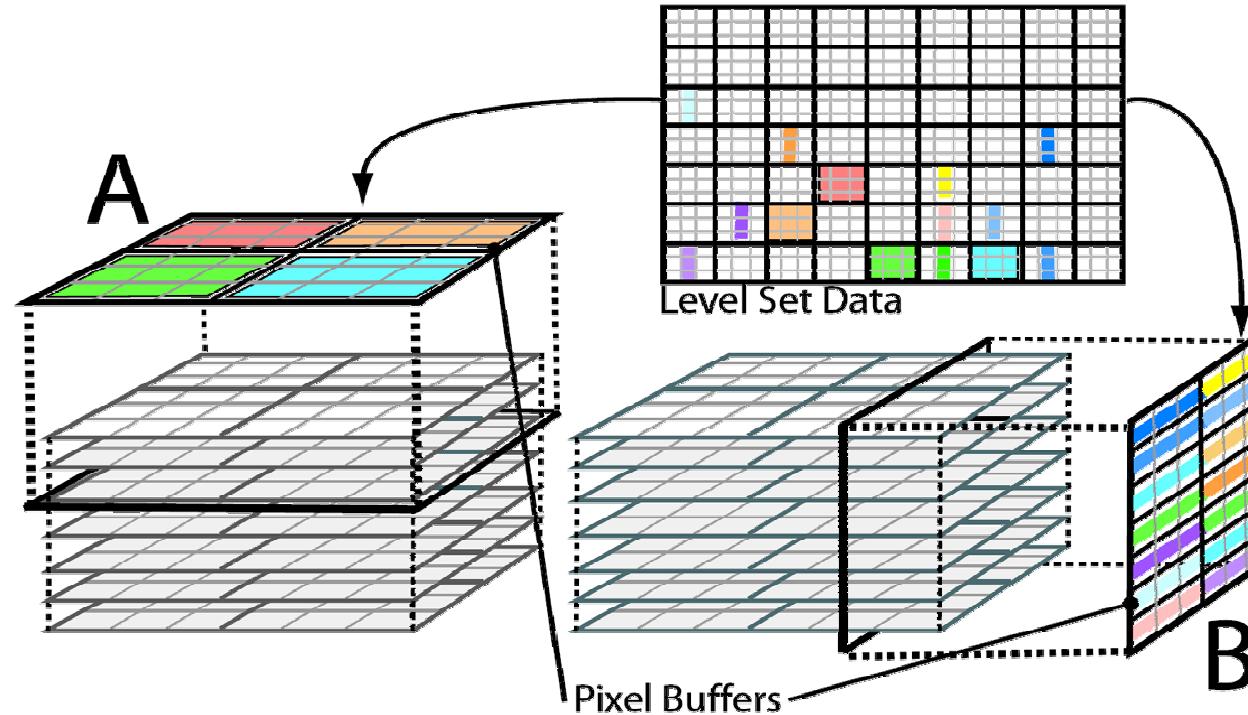
GPGPU
Aaron Lefohn
University of California, Davis, USA

J17

20
vis04

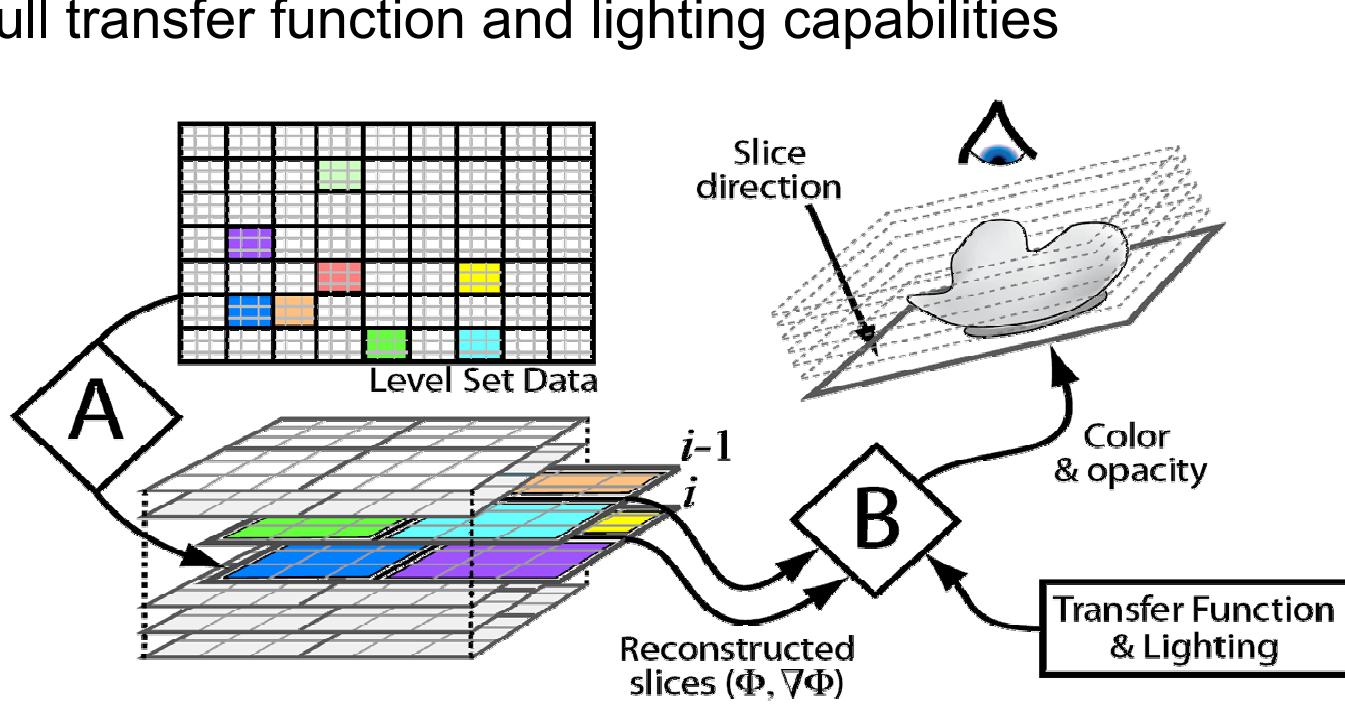
Direct Volume Rendering of Level Set

- Reconstruct 2D Slice of Virtual Memory Space
 - On-the-fly decompression on GPU
 - Use 2D geometry and texture coordinates



Direct Volume Rendering of Level Set

- Deferred Filtering: Volume Rendering Compressed Data
 - 2D slice-based rendering: *No data duplication*
 - Tri-linear interpolation
 - Full transfer function and lighting capabilities



Overview

- Why?
 - Motivation and previous work
- How?
 - Streaming level-set algorithm
 - Real-time visualization
 - Segmentation application
- Does it work?
 - Demonstration
 - User study



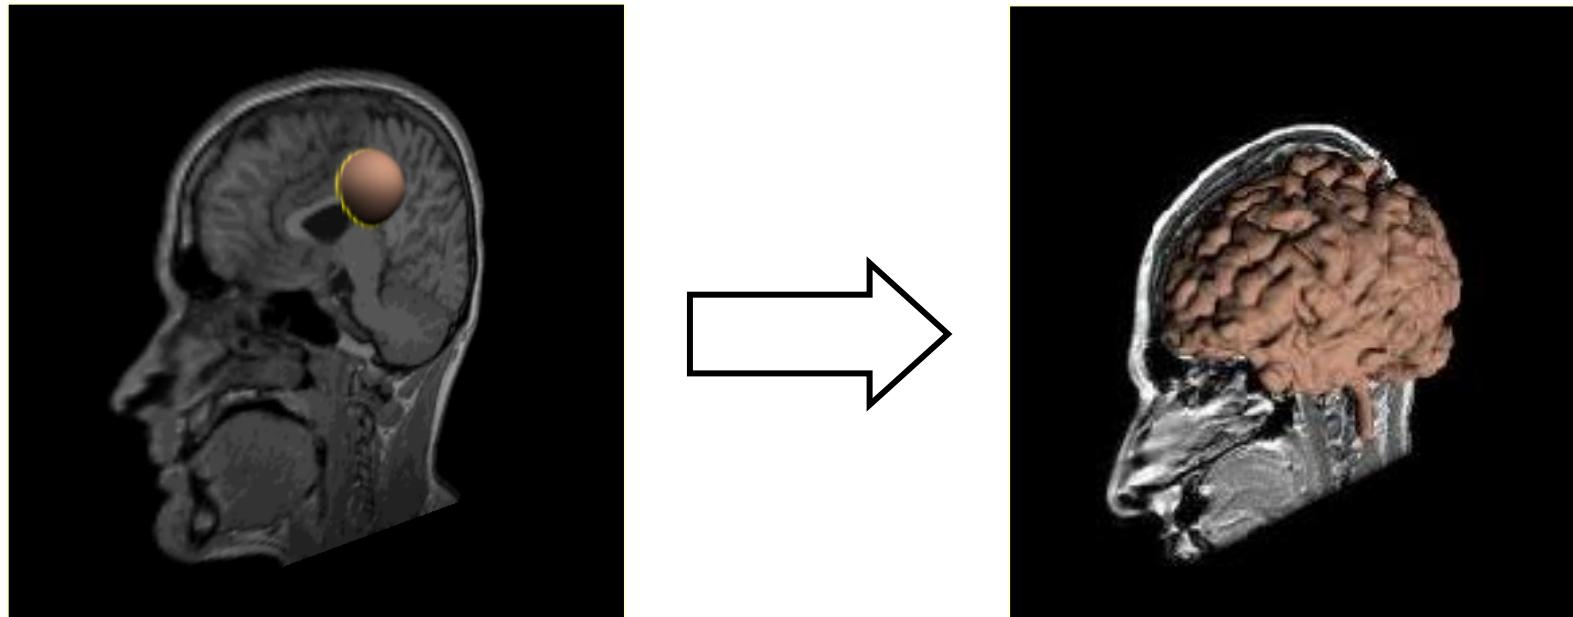
GPGPU
Aaron Lefohn
University of California, Davis, USA

J20

20
vis04

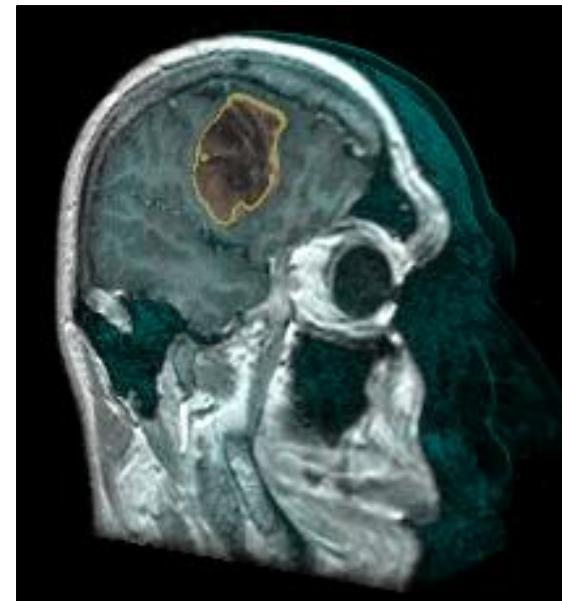
Level-Set Segmentation Application

- Idea: Segment Surface from 3D Image
 - Begin with “seed” surface
 - Deform surface into target segmentation



Demo

- Segmentation of MRI volumes
 - 128^3 scalar volume
- Hardware Details
 - ATI Radeon 9800 Pro
 - 1.7 GHz Intel Pentium 4
 - 1 GB of RAM

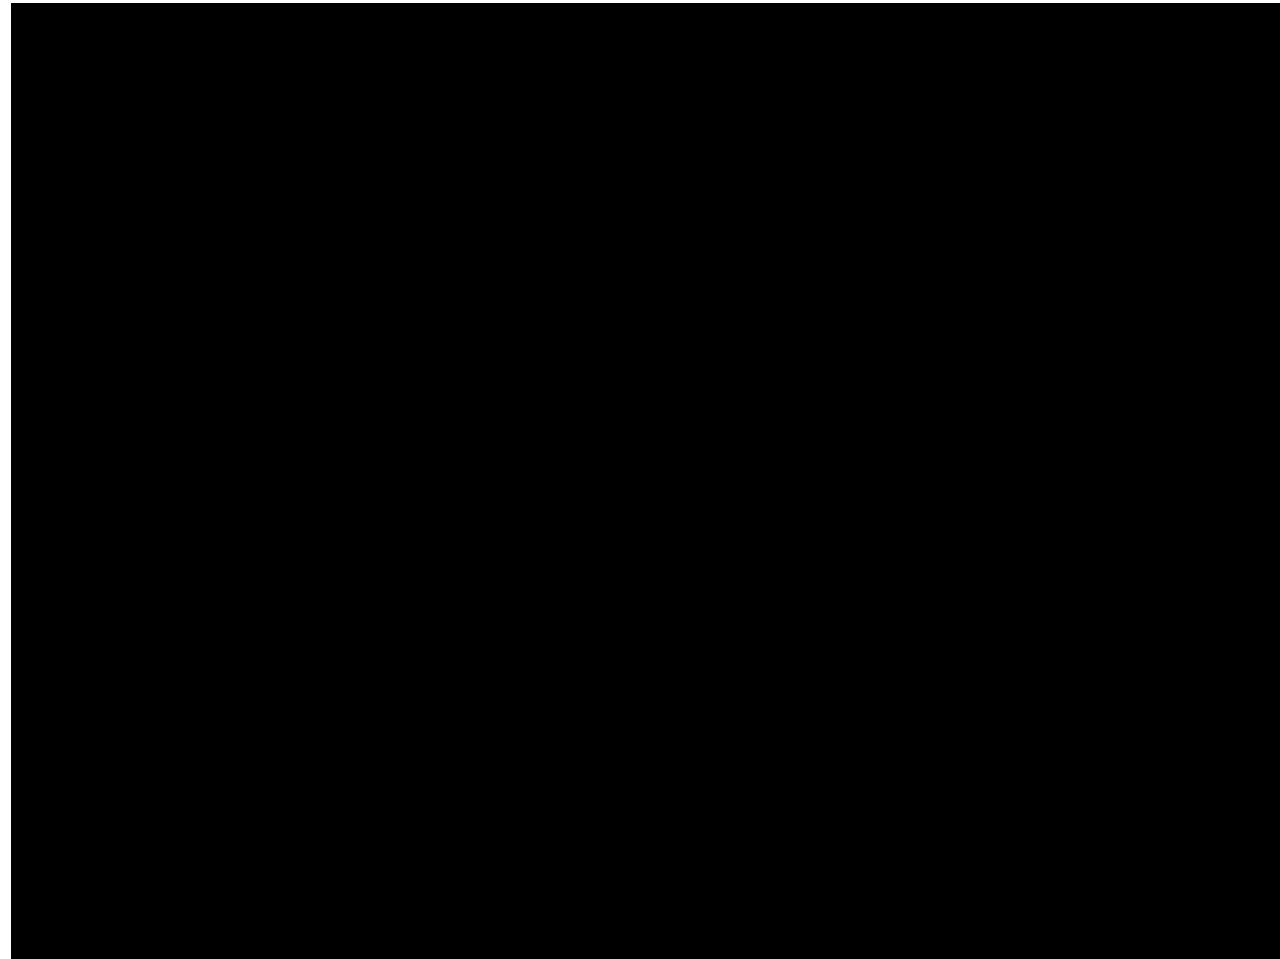


GPGPU
Aaron Lefohn
University of California, Davis, USA

J22

20
VIS04

Demo



GPGPU
Aaron Lefohn
University of California, Davis, USA

J23

20
VIS04

GPU Narrow-Band Performance

- Performance

- 10x – 15x faster than optimized CPU version (Insight Toolkit)
- Linear dependence on size of narrow band

- Bottlenecks

- Fragment processor (~80%)
- Conservative time step
 - Need for global accumulation register (min, max, sum, etc.)



Overview

- Why?
 - Motivation and previous work
- How?
 - Streaming level-set algorithm
 - Real-time visualization
 - Segmentation application
- Does it work?
 - Demonstration
 - User study (with Josh Cates)



GPGPU
Aaron Lefohn
University of California, Davis, USA

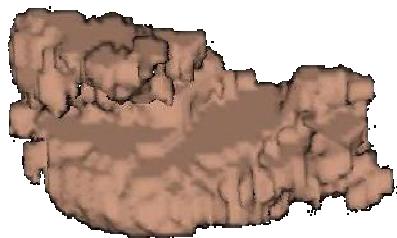
J25

20
vis04

Evaluation User Study

- Goal

- Can a user quickly find parameter settings to create an accurate, precise 3D segmentation?
- Relative to hand contouring



User Study Results

- Efficiency

- 6 ± 3 minutes per segmentation (vs multiple hours)
- Solver idle 90% - 95% of time

- Precision

- Intersubject similarity significantly better
- $94.04\% \pm 0.04\%$ vs. $82.65\% \pm 0.07\%$

- Accuracy

- Within error bounds of expert hand segmentations
- Compares well with other semi-automatic techniques
 - Kaus et al., Radiology, 2001



Summary

- Interactive Level-Set System
 - 10x – 15x speedup over optimized CPU implementation
 - Intuitive parameter tuning
 - User study evaluation

- Algorithm Developments
 - Multi-dimensional virtual memory
 - Substreams
 - GPU-to-CPU Message passing
 - Volume rendering packed data



Future Directions

- Other Level-Set Applications
 - Surface processing, surface reconstruction, physical simulation
- Integrate GPGPU Code Into Open Source Software
 - The Insight Toolkit (www.itk.org)?
- “Interactive Visualization”
 - GPGPU allows for simultaneous visualization and simulation
 - What problems can be solved with “interactive visualization?”
 - What is the user interface for a visualization?



Acknowledgements

- Joe Kniss
 - Volume rendering
 - Tumor user study
 - “Teem” raster-data toolkit
 - “GLEW” OpenGL extension wrangler
- Josh Cates
- Gordon Kindlmann
- Milan Ikits
- Ross Whitaker, Charles Hansen, Steven Parker and John Owens
- ATI: Evan Hart, Mark Segal, Jeff Royle, and Jason Mitchell
- Brigham and Women’s Hospital
- National Science Foundation Graduate Fellowship
- Office of Naval Research grant #N000140110033
- National Science Foundation grant #ACI008915 and #CCR0092065
- **ALIENWARE** Interchangeable mobile GPUs



GPGPU
Aaron Lefohn
University of California, Davis, USA

J30

20
VIS 04

Questions?

For More Information

Google “Lefohn level set”

<http://graphics.cs.ucdavis.edu/~lefohn/>

Journal Papers Based on this Work

Lefohn, Kniss, Hansen, Whitaker, “**A Streaming Narrow Band Algorithm: Interactive Computation and Visualization of Level Sets,**” IEEE Transactions on Visualization and Computer Graphics, 10 (40), Jul / Aug, pp. 422-433, 2004

Cates, Lefohn, Whitaker, “**GIST: An Interactive, GPU-Based Level-Set Segmentation Tool for 3D Medical Images,**” Medical Image Analysis, to appear 2004



GPGPU
Aaron Lefohn
University of California, Davis, USA

J31

20
VIS 04