#### Motion Planning & Physically-based Modeling Using GPUs

#### Ming C. Lin

University of North Carolina at Chapel Hill http://gamma.cs.unc.edu/voronoi/planner http://gamma.cs.unc.edu/PIVOT http://gamma.cs.unc.edu/PD http://gamma.cs.unc.edu/CULLIDE http://gamma.cs.unc.edu/ICE



# Organization

- Motion Planning
- Proximity Queries
- Phyiscally-based Simulation

# Organization

- Motion Planning
- · Proximity Queries
- Penetration Depth Estimation

# **Motion Planning**

Given the initial and goal configurations, the geometric description of the robot/agent and the environment, find a viable path if one exists.



# **Applications**

- Assembly Planning
- Car Painting
- Waste Cleaning
- Character animation
- Virtual human or artificial life
- Molecular docking
- Surgical Planning
- Virtual prototyping
- Pipe layout

#### **Two Classes of Planning Algorithms**

- Criticality-based
  - Complete
  - Computational expensive
  - Difficult to implement
- Random Sampling
  - Probablistically complete
  - Simple to implement
  - Fast in practice for most scenarios
  - Performance degrading on narrow passages

# **GVD:** Maximally Clear Points

Points with the largest distance values to nearby obstacles



#### Planning Based on GVD & MAT

- O'Dunlaing, Sharir and Yap [1983]
- Canny and Donald [1987]
- Latombe [1991]
- Choset, Burdick, et al. [1994-1999]
- Vleugels and Overmars [1996,1997]
- Guibas, Holleman and Kavraki [1999]
- Wilmarth, Amato and Stiller [1999]

..... and others .....

#### **Basic Ideas**

- Use of rasterization graphics hardware for real-time motion planning in dynamic environments
- Discrete GVD of any primitive types and approximate distance function with bounded errors
- Simple and "implementation-friendly" acceleration techniques for Voronoi-based robot motion planning
  - Construct Voronoi roadmap
  - Biasing path selection
  - Quick collision rejection test
  - Sampling near medial axisReduce search space
  - Establish milestones

# Outline

- Review: Computing GVD with Graphics Hardware
- Motion Planning with GVD
- Implementation & Results
- Summary

#### Voronoi Diagram

Given a collection of geometric primitives, it is a subdivision of space into cells such that all points in a cell are *closer* to one primitive than to any other



#### Accelerated Computation of GVD [Hoff et al. SIGGRAPH'99]

- Compute a discrete Voronoi diagram by rendering a three dimensional distance mesh for each Voronoi site.
- The polygonal mesh is a bounded-error approximation of a distance functions btw a site and a 2D planar grid of points.
- Each site is assigned a unique color & distance mesh rendered in that color.
- For each sample point, the graphics system computes the closest site and the distance to that site using polygon scan-conversion and the Z-buffer depth comparison.







# **3D Voronoi Diagrams**



Graphics hardware can generate one 2D slice at a time



# Outline

- Review: Computing GVD with Graphics Hardware
- Motion Planning with GVD
- Implementation & Results
- Summary

#### **Basic Approaches**

- Depth Buffer providing distance function and distance gradient (finite difference)
- Color Buffer building Voronoi graphs
- Combination of Both
  - Compute weighted Voronoi graphs
  - Voronoi vertices used for milestones
  - Weighted edges used for selecting paths
  - Distance values for quick rejection

#### Voronoi Boundary

- For each voxel in discrete GVD, associate a color corresponding to an object ID and a distance value to this obstacle.
- To extract the boundary, use continuation method similar to isosurface extraction -- starting from a seed point and step to next by bracketing boundary curves in a 2x2x2 region of sampled points.



#### **Constructing Voronoi Roadmap**

- Identify Voronoi vertices
- Extract Voronoi boundary
- Build Voronoi graph
- Select path based on edge weights
- Incrementally construct Voronoi roadmap
  - local planner (PFF) between milestones
  - quick collision rejection test & exact CD

# Voronoi Roadmap/Graph for Dynamic Environments



#### Outline

- Computing GVD with Graphics Hardware
- Motion Planning with GVD
- Implementation & Results
- Summary

# **Planning Using Voronoi Diagrams**

- Use the Voronoi graphs to find the near optimal path (in terms of both path length and clearance to nearby ostacles)
- The graphs are updated on the fly using GPU
- Use the potential field approach to orient the robot
- Applicable to both static and dynamic environments

[Hoff, et al, ICRA'00]

#### **Real-time Motion Planning : Static Scene**



Plan motion of piano (arrow) through 100K triangle model



Distance buffer of floorplan used as potential field

# **Real-time Motion Planning : Dynamic Scene**



around moving furniture



Distance buffer of floor-plan used as potential field

#### Voronoi-Based Sampling

- Use hardware accelerated computation of generalized Voronoi diagram (GVD) for PRM with Voronoi-based sampling
- Classify narrow passages to select different sampling strategies
- Quick collision rejection tests using hardware computed distance functions
- 25 to 1000% speedup over uniform sampling on preliminary benchmarks

[Pisula, et al, WAFR'00]

# 3D Benchmark (I)





# Voronoi-Based Hybrid Planner

- Use Voronoi Graphs as the estimated path
- Use the curvature of the path to orient the robot
- When a collision occur, use random sampling to correct the estimated path

[Foskey et al, IROS'01]



# **Constraint-Based Motion Planning**

- Transform a motion planning problem into a constraint-based dynamic simulation
- Use hardware accelerated computation of distance fields for potential functions and quick collision rejection
- Applicable to both dynamic and static environments with moving obstacles and multiple collaborative agents

[Garber & Lin'02]





#### Outline

- Review: Computing GVD with Graphics Hardware
- Motion Planning with GVD
- Implementation & Results
- Summary

#### **Summary**

- Techniques to exploit graphics hardware for *real-time* motion planning of a rigid robot in 3D
- Simple and easy to implement
- Applicable to both static and dynamic environments
- · Extended to Voronoi-based sampling for PRM
- Core of the Voronoi-based hybrid-planner
- Central to Constrained Motion Planning

# Organization

- Motion Planning
- Proximity Queries
- Physically-based Simulation

**Proximity Queries** 

*Geometric reasoning of spatial relationships among objects* (in a dynamic environment)



# Applications

- Dynamic simulation contact force calculation
- Haptic Rendering -- restoring force computation
- Computer Animation *motion control*
- Motion Planning *distance computation*
- Rapid Prototyping *tolerance verification*
- Virtual Environments -- interactive manipulation
- Simulation-Based Design *interference detection*
- Engineering analysis testing & validation
- Medical Training contact analysis and handling
- Education *simulating physics & mechanics*

# Our ApproachHybrid ApproachObject-spaceObject-spaceProximity QueryOraphics-HardwareOunded-ErrorAccelerationBounded-ErrorAccelerationBounded-ErrorApproximationBounded-ErrorAccelerationBounded-ErrorApproximationBounded-ErrorApproximation

#### **Related Work**

- Proximity query algorithms
  - Exact, object-space
  - Approximate, image-space
- Distance fields







Hoff99

Sutherland74, Catmull75 Rossignac92, Myskowski95, Baciu98

Lengyel90, Hoff00

Goldfeather86, Rossignac90, Stewart00

Algorithms simplified and accelerated using graphics hardware:

- Visibility:
- Intersections:
- Robot motion-planning:
- Voronoi diagrams:

#### Advantages:

- SimplicityBounded errorLinear complexity
- Robustness

# **Distance Fields**

Distance fields give distance to nearest object from any point: Sethian96

- Fast marching method
- Generalized Voronoi diagrams
- Adaptively-sampled distance fields
- Distance fields for penetration computation





Hoff99

Frisken00

# **Distance Fields**

Voronoi diagram computation using graphics hardware (Hoff99)





Depth buffer Result after compositing

Render polygonal mesh approximations of object distance fields Color buffer

distance fields using minimum depth test

# **Our Goal**

Proximity query algorithm with the following properties:

- Performs all proximity queries
- · Handles non-convex primitives
- Requires no precomputation or complex data structures
- Fast and efficient
- Robust
- Portable
- Bounded error

No previous algorithm found, even in 2D!







#### **Distance Field and Gradients**



Draw polygon to encode negative sign in a buffer



Central difference to compute gradient at a pixel













# **Separation Distance and Direction**



Distance at closest point = min separation distance Gradient at closest point = separating axis direction























#### **Demonstrations**

- Simulation test cases:
  - Rigid and deformable objects
  - Different contact scenarios
- Collision response in 2 simulation strategies:
  - Non-penetration constraint, backtracking
  - Unconstrained, penalty-based
- In each demo:

  - No precomputationInteractive frame rates
  - Bounding-box intersection localization
  - Pentium III 800, nVidia GeForce256



# **BUMPER CARS**



- Convex, rigid objects
- Less frequent contact
- Non-penetration constraint, backtracking







Performance					
Average Total Per-frame Proximity Query Times					
Demo	Objects	Lines	GeForce2	InfiniteReality2	ATI Rage Pro LT
Мар	6	719	0.281ms	0.901ms	0.434ms
Gears	13	391	0.015	0.026	0.064
Links	15	440	0.020	0.052	0.038
Cars	18	266	0.007	0.026	0.015
Mont	2	200	1.020	2 260	2 000

Performance timings for dynamics simulations. The number of objects, number of line segments, and the average total time in milliseconds to run proximity queries on all objects in the scene per frame is reported. Timing data was gathered from three machines: a Pentium-III 933MHz desktop with a 64Mb GeForce2, a SGI 300MHz R12000 with Infiniteality2 graphics, and a Pentium-III 750Mhz laptop with ATI Rage Pro LT graphics.

# **Effects of Changes in Error Tolerance**

Effects of Error Tolerance on Performance of Wavy				
Error	GeForce2	InfiniteReality2	ATI Rage Pro LT	
d/4	0.710ms	1.270ms	5.560ms	
d/2	0.315	1.000	1.850	
d	0.211	0.930	0.895	
2 <i>d</i>	0.176	0.879	0.631	
4d	0.165	0.876	0.535	

The effect on performance when changing the distance error tolerance *d*. We used proximity queries on the **wavy** demo with no collision response. The error determines the number of pixels used in the image-based operations. Systems with low graphics performance are more directly affected by the choice of *d* (see ATI Rage Pro LT); however, as the error is increased there is less dependence on graphics performance and the faster laptop CPU overtakes the InfiniteReality2 system.

# Recent Results: Extension to 3D

#### Real-time queries: Deformable models

- Non-convex, complex 3D deformable objects
- ~20000 polygons each cylinder
- Continuous contact
- Unconstrained, penaltybased
- Intersection region in blue
- Distance field gradients
- in red and greenNo pre-computation
- [Hoff, Zaferakis, Lin, Manocha'02]



#### **Real-time queries: Deformable models**

Non-convex, complex 3D
 deformable objects

3D intersections

- ~20000 polygons each cylinder
- Continuous contact
- Unconstrained, penaltybased
- Intersection region in blue
- Distance field gradients in red and green
- No pre-computation
- [Hoff, Zaferakis, Lin, Manocha'02]



3D distance field

# NURBS Torso Models [Seeger et al.'00]

Sample Slices from the Visible SURFdriver Surface Human CT Data Set Reconstruction Progra

Skeleton

Reconstruction Program



NCAT Phantom

(anterior view)

3D NURBS Organ Models



Lungs

an Models

#### Liver Stomach Spleen Kidneys

20

# 4D NURBS Respiratory Model



# **Proximity Queries between Organs**







Right Lateral View

#### M. C. L

# **Modeling the Heart Structures**

Heart exterior

Contact computations between Inner chambers



# Proximity Queries using Graphics H/W

- Performs all proximity queries
- Handles general non-convex & deformable
  primitives
- Requires no pre-computation or complex hierarchical data structures
- Fast and efficient

#### **Some Observations**

- Appropriate for most of applications where numerical accuracy is not the key
- Running time related to the rendering time
- Error bound provides smooth dial between performance and level of approximation
- Varying the max depth in the hierarchical localization allows balancing load between CPU and graphics

# Definition

#### • Penetration Depth (PD)

- Minimum translational distance to separate two intersecting objects



#### **Motivation**

- Contact Handling in Rigid Body Simulation
  - Penalty-based method for contact resolution
  - Time stepping in general simulation framework





#### Motivation

- Time stepping method using PD
- 1. Compute the penetration depth

•

2. Estimate the TOC by interpolation in time domain



# **PD** Applications

- Rigid body dynamic simulation
- Robot motion planning for autonomous agents and animated characters
- Haptic rendering
- Tolerance verification for CAD models

#### **Previous Work**

- Convex polytopes
  - [Cameron '86 '97], [Dobkin et al. '93], [Agarwal et al. '00], [Bergen '01], [Kim et al. '02]
- Local solutions for deformable models – [Susan and Lin '01], [Hoff et al. '01]
- No solutions existed for non-convex models

# Overview

• Preliminaries

- Minkowski sum-based Framework
- Basic PD Algorithm: A Hybrid Approach
- Acceleration Techniques
  - Object Space Culling
  - Hierarchical Refinement
  - Image Space Culling
- Application to Rigid Body Dynamic Simulation

#### Overview

#### • Preliminaries

#### -Minkowski sum-based Framework

- Basic PD Algorithm: A Hybrid Approach
- Acceleration Technique
- Object Space Culling
- Hierarchical Refinemen
- Image Space Culling
- Application to Rigid Body Dynamic Simulation

# Preliminaries

• Local solutions might not have any relevance to a global solution



#### **Preliminaries**

- Minkowski sum and PD
  - $-P \oplus -Q = \{ p q \mid p \in P, q \in Q \}$
  - PD := minimum distance between  $O_{Q\cdot P}$  and the surface of  $P \oplus Q$







#### **Preliminaries**

- Decomposition property of Minkowski sum - If  $P = P_1 \cup P_2$ , then  $P \oplus Q = (P_1 \oplus Q) \cup (P_2 \oplus Q)$
- Computing Minkowski sum

   Convex: O(n log(n))
   where n is the number of features

  - Non-Convex: O(n<sup>6</sup>) computational complexity
    In theory, use the convolution or the decomposition In propertyIn practice, very hard to implement

# Overview

- Basic PD Algorithm: A Hybrid Approach

# **PD Algorithm : A Hybrid Approach**

•  $P \oplus Q = (P_1 \oplus Q) \cup (P_2 \oplus Q)$ 

# **PD Algorithm : A Hybrid Approach**

*P* ⊕ *Q* = (*P*<sub>1</sub> ⊕ *Q*) ∪ (*P*<sub>2</sub> ⊕ *Q*)
– where *P* = *P*<sub>1</sub> ∪ *P*<sub>2</sub>
Precomputation: Decomposition

# **PD Algorithm : A Hybrid Approach**

•  $P \oplus Q = (P_1 \oplus Q) \cup (P_2 \oplus Q)$ - where  $P = P_1 \cup P_2$ • Preservent time Decomposition

Precomputation: Decomposition

• Runtime:

- Object Space: Pairwise Minkowski sum computation

# **PD** Algorithm : A Hybrid Approach

•  $P \oplus Q = (P_1 \oplus Q) \cup (P_2 \oplus Q)$ 

- where  $P = P_1 \cup P_2$
- Precomputation: Decomposition
- Runtime:
  - Object Space: Pairwise Minkowski sum computation
  - Image Space: Union by graphics hardware

[Kim, Otaduy, Lin & Manocha, SCA'01]

# **PD** Computation Pipeline



#### **Convex Surface Decomposition**



- [Ehmann and Lin '01]
- Decompose an object into a collection of convex surface patches
- Compute the convex hull of each surface patch

#### Pairwise Minkowski Sum

#### • Algorithms

- Convex hull property of convex Minkowski sum
  - $P \oplus Q = ConvHull \{v_i + v_j | v_i \in V_p, v_j \in V_Q\}$ , where P and Q are convex polytopes
- Topological sweep on Gauss map [Guibas '87]
- Incremental surface expansion [Rossignac '92]

# **Closest Point Query**

#### Goal

- -Given a collection of convex Minkowski sums, compute the shortest distance from the origin to the surface of their union
- An exact solution is computationally expensive ⇒
   Approximation using graphics hardware

# **Closet Point Query**

- Main Idea
  - Incrementally expand the current front of the boundary



# **Closest Point Query**

- 1. Render front faces, and open up a window where z-value is less than the current front
- 2. Render back faces w/ z-greater-than test
- 3. Repeat the above m times, where m := # of obj's



# **Closest Point Query**

- 1. Render front faces, and open up a window where z-value is less than the current front
- 2. Render back faces w/ z-greater-than test
- 3. Repeat the above *m* times, where m := # of obj's



#### Overview

- Preliminaries
- Minkowski sum-based Francowork
- Basic PD Algorithm: A Hybrid Approact
- Acceleration Techniques
  - -Object Space Culling
  - -Hierarchical Refinement
  - -Image Space Culling
- · Application to Rigid Body Dynamic Simulation

# Motivation

- PD is shallow in practice
- Convex decomposition has O(n) convex pieces in practice
- $\rightarrow$  Culling strategy is suitable and very effective

# **Object Space Culling**

#### • Basic Idea

- If we know the upper bound on PD,  $u_{PD}$ , we do not need to compute the Mink. sum of pairs whose Euclidean dist is more than  $u_{PD}$ 



# **Object Space Culling**

#### • Basic Idea

– If we know the upper bound on PD,  $u_{PD}$ , we do not need to compute the Mink. sum of pairs whose Euclidean dist is more than  $u_{PD}$ 



# **Object Space Culling**

- Basic Idea
  - If we know the upper bound on PD,  $u_{PD}$ , we do not need to compute the Mink. sum of pairs whose Euclidean dist is more than  $u_{PD}$



# **Image Space Culling**

- Rendering only once for the Minkowski sums containing the origin
- Refine the upper bound for every view frustum
- View frustum culling









Models	Tri	Convex Pieces	PD w/o Accel.	PD w/ Accel
Touching Tori	2000	67	4 hr	0.3 sec
Interlocked Tori	2000	67	4 hr	3.7 sec
Interlocked Grates	444, 1134	169, 409	177 hr	1.9 sec
Touching Alphabets	144,152	42, 43	7 min	0.4 sec

# **Hierarchical Culling Example**



# Implementation

- SWIFT++ [Ehmann and Lin '01]
- QHULL
- OpenGL for closest point query

#### Implementation

void DrawUnionOfConvex(ConvexObj \*ConvexObjs, int NumConvexObjs)

void Davoname giClearDeph(0); giClearStencil(0); giClearGL\_COLOR\_BUFFER\_BIT | GL\_DEPTH\_BUFFER\_BIT | GL\_STENCIL\_BUFFER\_BIT); giClear(GL\_COLOR\_BUFFER\_BIT | GL\_DEPTH\_BUFFER\_BIT | GL\_STENCIL\_BUFFER\_BIT);

glEnable(GL\_DEPTH\_TEST); glEnable(GL\_STENCIL\_TEST); for (int i=0; i<NumConvexObjs; i++) for (int j=0; j<NumConvexObjs; j++)

glDepthMask(0); glColorMask(0,0,0,0); glDepthFunc(GL\_LESS); glStencilFunc(GL\_ALWAYS,1,1); glStencilOf(GL\_KEEP,GL\_REPLACE,GL\_KEEP); ConvexObjs[]].DrawFrontFaces();

glDepthMask(1); glColorMask(1,1,1,1); glDepthFunc(GL\_GRATER); glStencilFunc(GL\_EQUAL,0,1); glStencilOp(GL\_ZERO,GL\_KEEP,GL\_KEEP); ConvexObjs[j].DrawBackFaces();

#### **Accuracy of PD computation**

- Our algorithm computes an upper bound to PD
- Image space computation determines the tightness of the upper bound
  - Pixel resolution
  - Z-buffer precision
- In practice, with 256×256 pixel resolution, the algorithm rapidly converges to the PD

#### Overview

- Preliminaries
  - Minkowski sum-based Framework
- Basic PD Algorithm: A Hybrid Approach
- Acceleration Techniques
  - Object Space Culling
  - Hierarchical Kefinement
  - Image Space Culling
- Application to Rigid Body Dynamic Simulation

# **Application to Rigid Body Simulation**

- Interpenetration is often unavoidable in numerical simulations
- Need for a consistent and accurate measure of PD
- Penalty-based Method
  - $-F = (k \cdot d)n$ 
    - d: PD, n: PD direction, k:stiffness constant

# **Application to Rigid Body Simulation**

- Time Stepping Method
  - -Estimate the time of collision (TOC)
    - *s*: separation dist before interpenetration
    - *d*: PD, *n*: PD dir after interpenetration
    - $v_s$ : relative velocity of closest features
    - $v_d$ : relative velocity of PD features

# **Application to Rigid Body Simulation**

-x(t): 1D distance function between closest features and PD features projected to PD direction

n

$$-x(0) = s, \ x(T) = d$$
  
$$-dx/dt(0) = \mathbf{v}_s \cdot \mathbf{n}, \ dx/dt(T) = \mathbf{v}_s$$

-Compute the roots of x(t) = 0

# **Rigid Body Simulation Demo**



Example 1 Average complexity: 250 triangles 60 Convex Pieces / Object



Example 2 Average complexity: 250 triangles 200 letters and alphabets

# **Rigid Body Simulation Demo**



# Summary

- First practical PD algorithm using a hybrid approach
- Acceleration techniques:
  - Object space culling
  - Image space culling
  - Hierarchical refinement
- Application to rigid body simulation

#### **Proximity Queries**

#### • Object-Space Methods

- Considerable pre-processing for hierarchy construction
- Difficult to achieve real-time performance on complex deformable and/or breaking models

#### • Image-Space Techniques

- Primitive rasterization ↔ sorting in screen-space
- Applicable to interference tests

#### **Problems with Image-Based Methods**

- Closed models
- Frame buffer readbacks slow
   NVIDIA GeForce 4, Dell Precision Workstation with 1Kx1K depth buffer taking 50ms

#### Goals

- Interactive Performance
- Complex objects
  - Large number of objects
  - High primitive count
  - Non-convex objects
  - Open and closed objects
- Non-Rigid Motion
  - Deformable bodies
  - Changing topology

# Overview Potentially Colliding Set (PCS) computation Exact collision tests on the final PCS





# Potentially Colliding Set (PCS)

# Potentially Colliding Set (PCS)



# Comparison

- Object Level Pruning ↔ Broad Phase
- Sub-object Level Prunning ↔ Narrow Phase



# **Visibility Computations**

- Lemma 1: An object O does not collide with a set of objects S if O is fully visible with respect to S
- $\rightarrow$  Utilize visibility for PCS computation



# **PCS Pruning**

Lemma 2: Given n objects  $O_1, O_2, ..., O_n$ , an object  $O_i$ does not belong to PCS if it does not collide with  $O_1, ..., O_{i-1}, O_{i+1}, ..., O_n$ 

 $\rightarrow$  Prune objects that do not collide

DCC	<b>n</b> •
PUN	Pruning
~	

$$O_1 \ \ O_2 \ \ \dots \ \ O_{i-1} \ \ O_i \ \ O_{i+1} \ \ \dots \ \ O_{n-1} \ \ O_n$$

**PCS Pruning** 

$$O_1 \quad O_2 \quad \dots \quad O_{i-1} \quad O_i$$



# **PCS Computation**

- Each object tested against all objects but itself
- Naive algorithm is  $O(n^2)$
- Linear time algorithm
  - Uses two pass rendering approach
  - Conservative solution























# **PCS** Computation

$$\bigcirc_{1} O_{2} O_{3} \dots O_{i+1} \bigcirc_{1} O_{i+1} \dots O_{n-2} O_{n-1} O_{n}$$

$$\downarrow$$

$$O_{1} O_{3} \dots O_{i-1} O_{i+1} \dots O_{n-1}$$

# Sub-Object Level: Overlap Localization

- Each object is composed of sub-objects
- We are given n objects O<sub>1</sub>,...,O<sub>n</sub>
- Compute sub-objects of each object O<sub>i</sub> that overlap with sub-objects of other objects

# Sub-Object Level: Overlap Localization

- Our solution
  - Test if each sub-object of  $O_i$  overlaps with sub-objects of  $O_1, ... O_{i, l}$
  - Test if each sub-object of  $O_i$  overlaps with sub-objects of  $O_{i+1}, ..., O_n$
- Linear time algorithm
- Extend the two pass approach



Sub-Object Overlap Localization: First Pass

**Render sub-objects** 

$$O_1 \ O_2 \ \dots \ O_{i-1} \ O_i \ O_{i+1} \ \dots \ O_{n-1} \ O_n$$













# Sub-Object Overlap Localization: Second Pass

Render sub-objects

$$O_1 \ O_2 \ \dots \ O_{i-1} \ O_i \ O_{i+1} \ \dots \ O_{n-1} \ O_n$$

Sub-Object Level Overlap Localization

$$O_1 \ \ O_2 \ \ \ldots \ \ O_{i-1} \ O_i \ O_{i+1} \ \ \ldots \ \ O_{n-1} \ O_n$$



# **Visibility Queries**

- We require a query – Tests if a primitive is fully visible or not
- Current hardware supports occlusion queries – Test if a primitive is visible or not
- Our solution
  Change the sign of depth function

	Visibility (	Jueries	
	De	pth function	
	GEQUAL	LESS	
All fragments	Pass	Fail	
	Fail	Pass	
	Occlusion query	Query not supported	
Exampl	les - HP_Occlusion_t	est, NV_occlusion_query	

# **Bandwidth Analysis**

• Read back only integer identifiers – Independent of screen resolution

# **Optimizations**

- First use AABBs as object bounding volume
- Use orthographic views for pruning
- Prune using original objects

# Implementation

- Dell Precision workstation
- Dell M50 Laptop

[Govindaraju, Redon, Lin, Manocha, GH'03]



# **Test Models: Environment 2**



# **Test Model: Environment 3**



- 250K Dragon 35K Bunny

# **Test Model: Environment 4**



- Breaking dragon 250K triangles
- Bunny 35K triangles

[Govindaraju, Redon, Lin, Manocha, GH'03]

#### Live Demo of Environment 4

• Dell M50 laptop, 2.4GHz Pentium IV-M CPU, NVIDIA Quadro4 700GoGL GPU, 1GB memory running Windows XP







# Advantages

- No coherence
- No assumptions on motion of objects
- Works on generic models
- A fast pruning algorithm
- No frame-buffer readbacks
- Applicable to deformable & breaking models

# Limitations

- No distance or penetration depth information
- Resolution limited accuracy
- No self-collision detection

#### Summary

- First collision detection using GPU
  - Applicable to polygon soups, open boundary, etc.
  - No assumptions on motion
  - Allowing change of topology and dynamic geometry
- Unified approach for object & sub-object level culling
- Linear time PCS computation algorithm
- No frame buffer readbacks

# Organization

- Motion Planning
- Proximity Queries
- Physically-based Simulation

#### **Visual Simulation of Ice Growth**





# Motivation

• "We built a particle system that produced the effect of fingers of ice ... which was tricky, since the particles didn't do that automatically"

- Arnon Manor, Cinesite

# **Our approach**

- A method for modeling and simulating visually complex structure of ice both geometrically and optically
  - Physical Simulation using Phase Field Methods
  - Simulation Acceleration using GPUs & Banded Opt.
  - Controlling Complexity
  - Post-Processing for Rendering

[Kim and Lin, SCA'03]

Phase Field Methods  
• Temperature PDE  

$$\frac{\partial T}{\partial t} = a^2 \nabla^2 T + K \frac{\partial p}{\partial t}$$
• Phase PDE  

$$\tau \frac{\partial p}{\partial t} = \nabla \cdot (\varepsilon^2 \nabla p) - \frac{\partial}{\partial x} \left( \varepsilon \frac{\partial \varepsilon}{\partial \theta} \frac{\partial p}{\partial y} \right) + \frac{\partial}{\partial y} \left( \varepsilon \frac{\partial \varepsilon}{\partial \theta} \frac{\partial p}{\partial x} \right) + p(1-p) \left( p - \frac{1}{2} + m \right)$$



# **GPU Acceleration**

- Fields map easily to GPU
- Framework of [Harris et al. 2002]
  - Fields sent as textures
  - Fragment program runs PDE per texel
     ~100 lines of Cg!

# Performance

- Only unbanded currently possible
- No early exit yet on GPU

		GPU(Hz)	Speedup
	250	624	2.50x
	25	236	9.44x
256x256	8	67.47	8.43x
	3.5	17.67	5.05x
	1.08	3.77	3.49x



#### **Future Work**

- Extension to motion planning of deformable robots and other types of constraints
- PD computation for rotational movement using GPU
- Continuous Collision Detection using GPUs
- Physically-based modeling of paint media, phase transition (melting, solidification, etc)

#### Collaborators

Bill Baxter Tim Culver Mark Foskey Maxim Garber Naga Govindaraju Kenny Hoff John Keyser Theodore Kim Young Kim Dinesh Manocha Miguel Otaduy Charles Pisula Stephan Redon Andrew Zaferakis

#### Acknowledgements

Stephen Ehmann Stefan Gottschalk Sarah Hoff David Hsu Eric Larsen Jean-Claude Latombe Jean-Paul Laumond

# Acknowledgements

Army Research Office DOE ASCI Program Intel Corporation National Institute of Health National Science Foundation Office of Naval Research University of North Carolina at Chapel Hill

