

***USING GPUS FOR  
COLLISION DETECTION***

**AMD  
Takahiro Harada**

---



- GPU architecture updates
- Problem description
  - Why collision detection is different from other physics problems
- GPU collision detection
- Heterogeneous collision detection

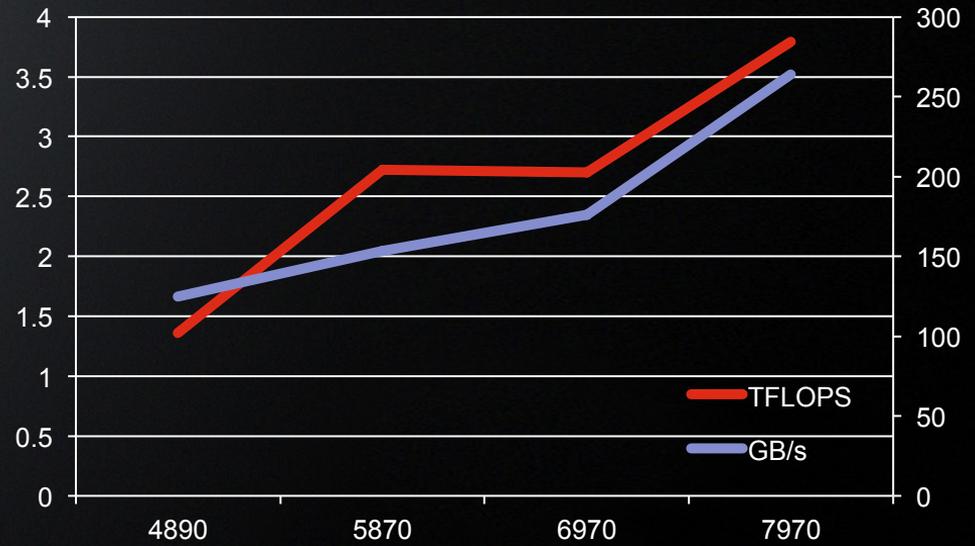


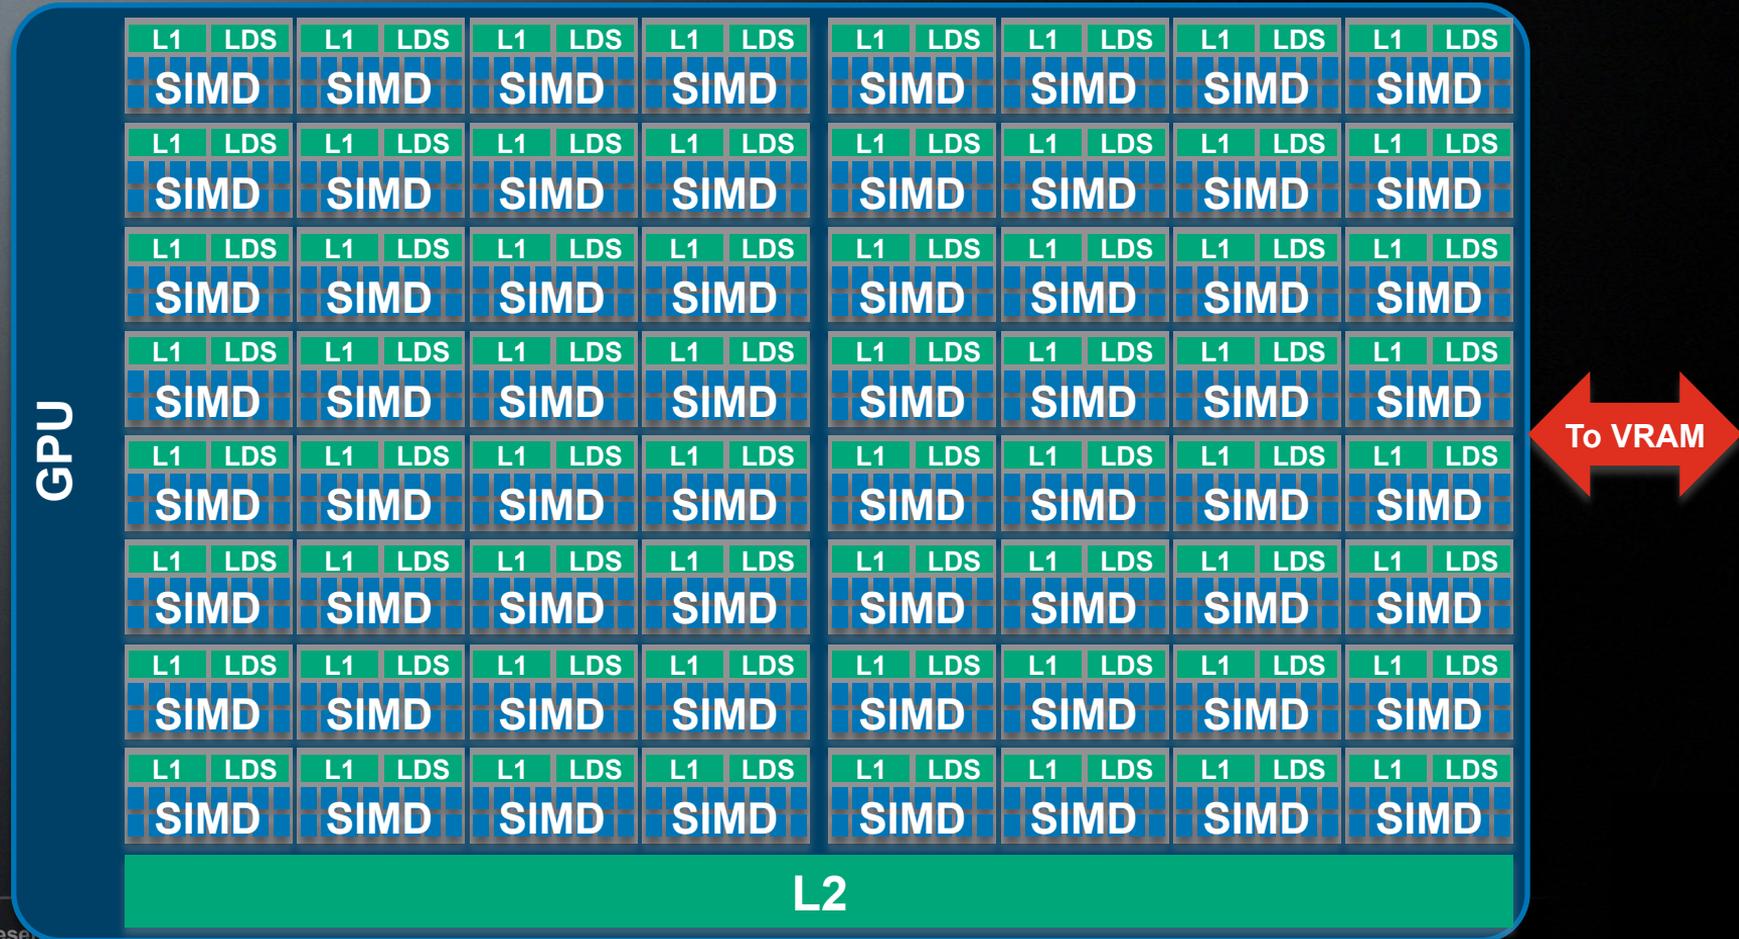
# *GPU UPDATES*



- High ALU performance (FLOPS)
- High memory bandwidth

- Radeon HD 7970
  - 3.79 TFLOPS
  - 264 GB/s
  - 32x4 SIMD engines
  - 64 wide SIMD





# PEAK PERFORMANCE IS HIGH, BUT



Eurographics 2012

Cagliari, Italy

May 13-18



33<sup>rd</sup> ANNUAL CONFERENCE OF THE EUROPEAN ASSOCIATION FOR COMPUTER GRAPHICS

- GPU used to be only good at simple computations
- GPU is improving architecture
  
- GPU performs good only if used correctly
  - Divergence
- Prefer simpler algorithm and data structure
  - ⇔ Linked list traversal

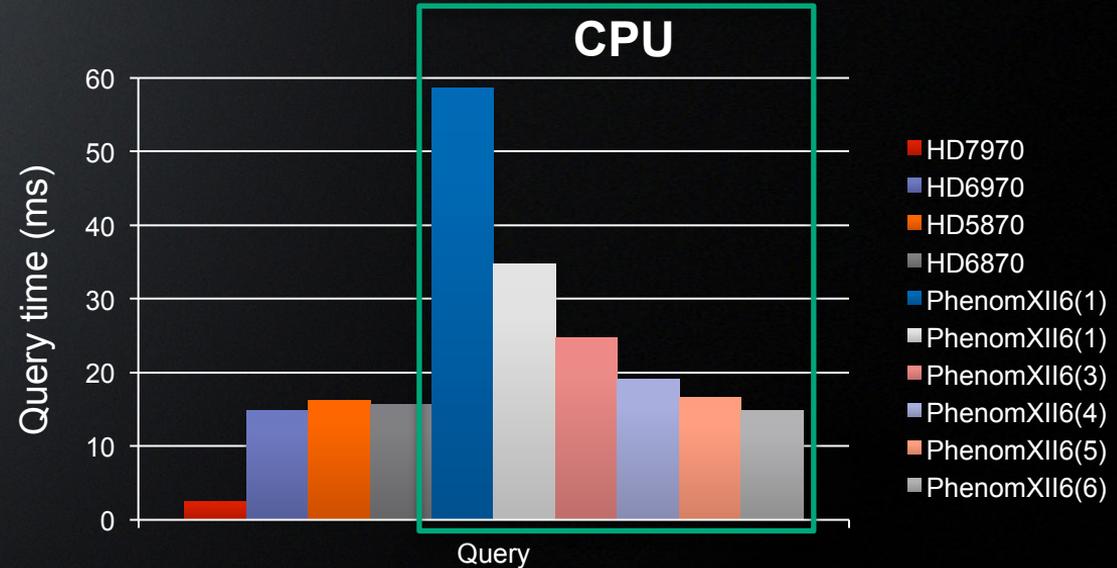


- Low ALU / Mem ratio
- Pointer chasing
- Traversal

```
Node* node = root;  
do  
{  
    process( node );  
    node = node->next;  
}while( node )
```

- CPU is better for this

## Spatial query using a hash grid



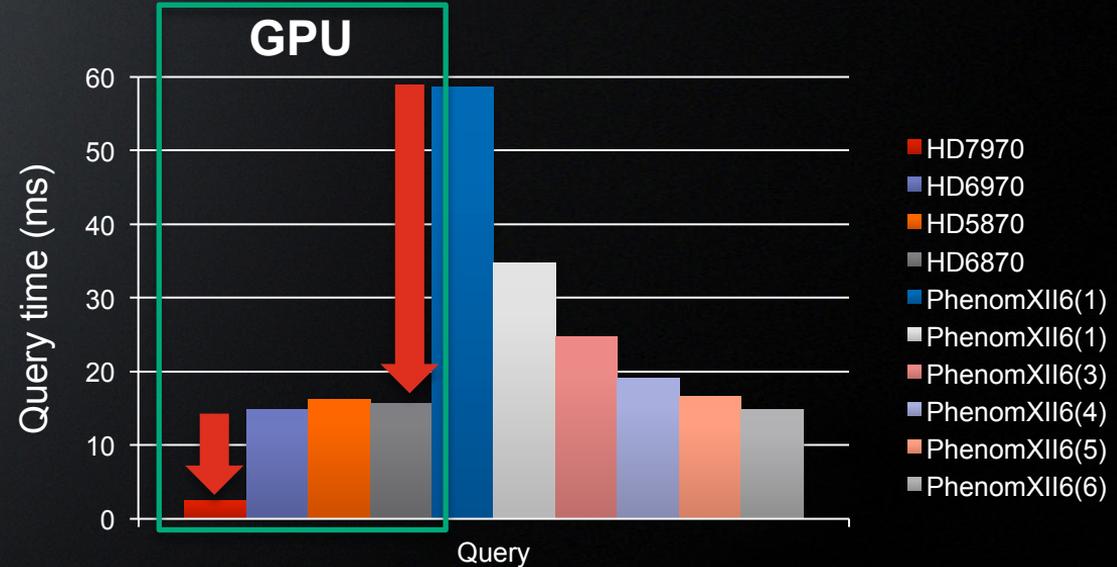


- Low ALU / Mem ratio
- Pointer chasing
- Traversal

```
Node* node = root;  
do  
{  
    process( node );  
    node = node->next;  
}while( node )
```

- CPU is better for this

### Spatial query using a hash grid

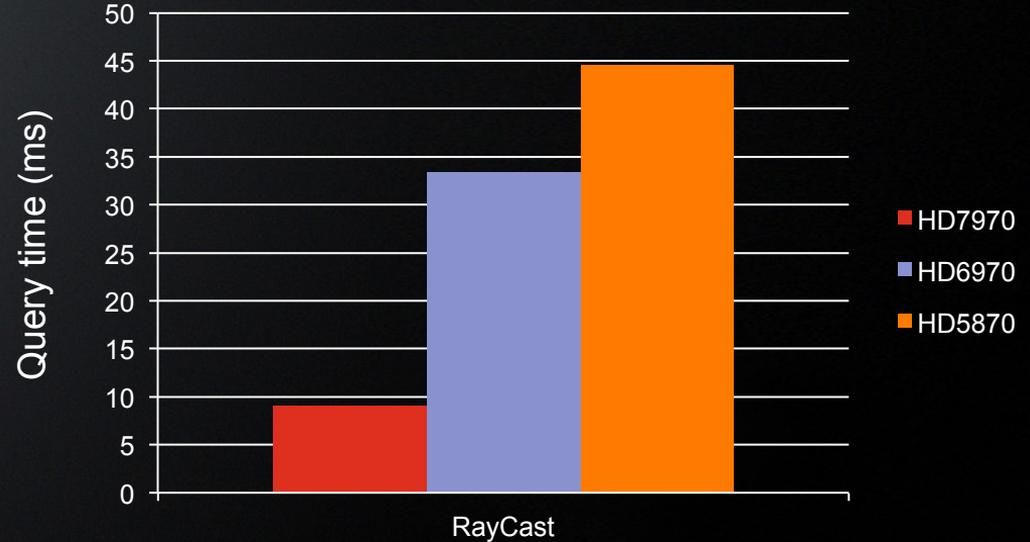




- Cannot explain only by raw performance increase
  - It is less than 2x.
- Key architectural changes of HD7970
  - Non VLIW
    - Increases ALU efficiency
  - Improved memory system
    - RW cache
  - Faster atomic operations
    - On chip atomics
- Improves computation performance in general



- Ray casting
  - Ray tracing
  - Visibility query
- GPU outperforms a multicore CPU
  - Test uses 6 cores



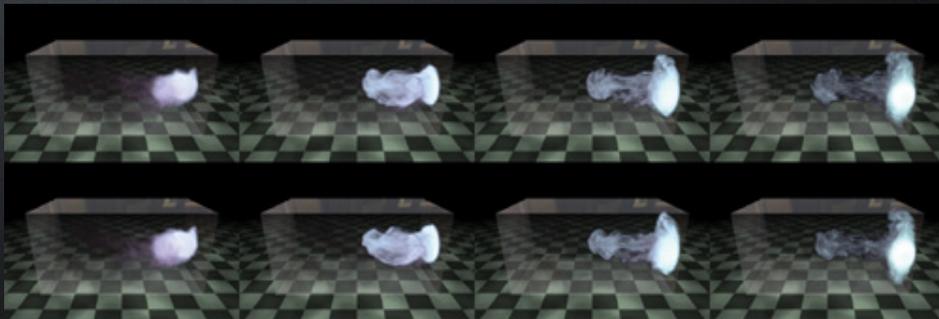
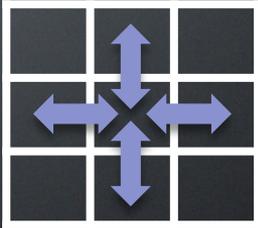
- GPU is not only adding more ALUs on a chip
- Architecture has been improved



# ***GPU COLLISION DETECTION***



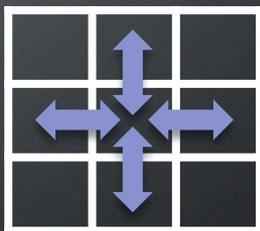
- GPUs were used for physics simulation with fixed connectivity between entities
  - Grid based fluid simulation
    - Entities: Cells
    - Connections: Cell adjacency



Crane et al., GPU Gems3 (2007)

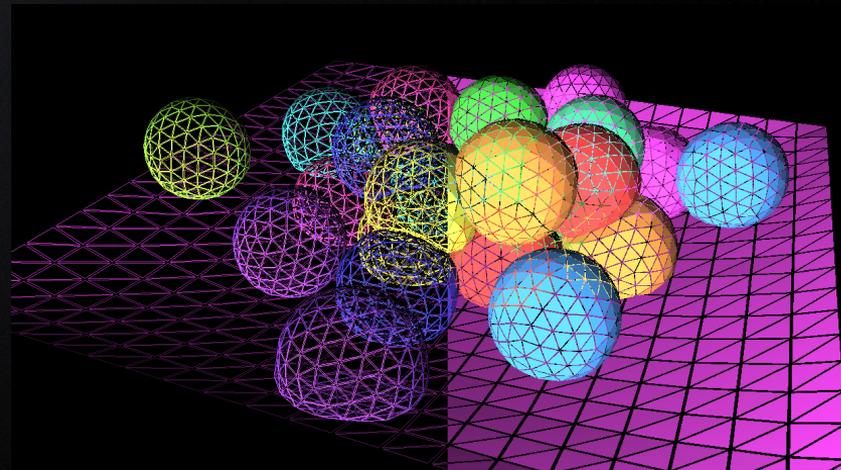
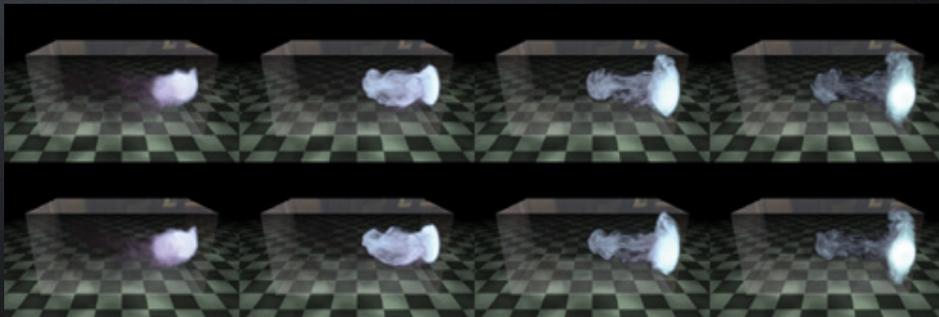
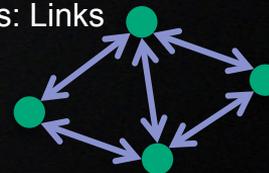


- GPUs were used for physics simulation with fixed connectivity between entities
  - Grid based fluid simulation
    - Entities: Cells
    - Connections: Cell adjacency



## – Cloth simulation

- Entities: Vertices
- Connections: Links



Crane et al., GPU Gems3 (2007)

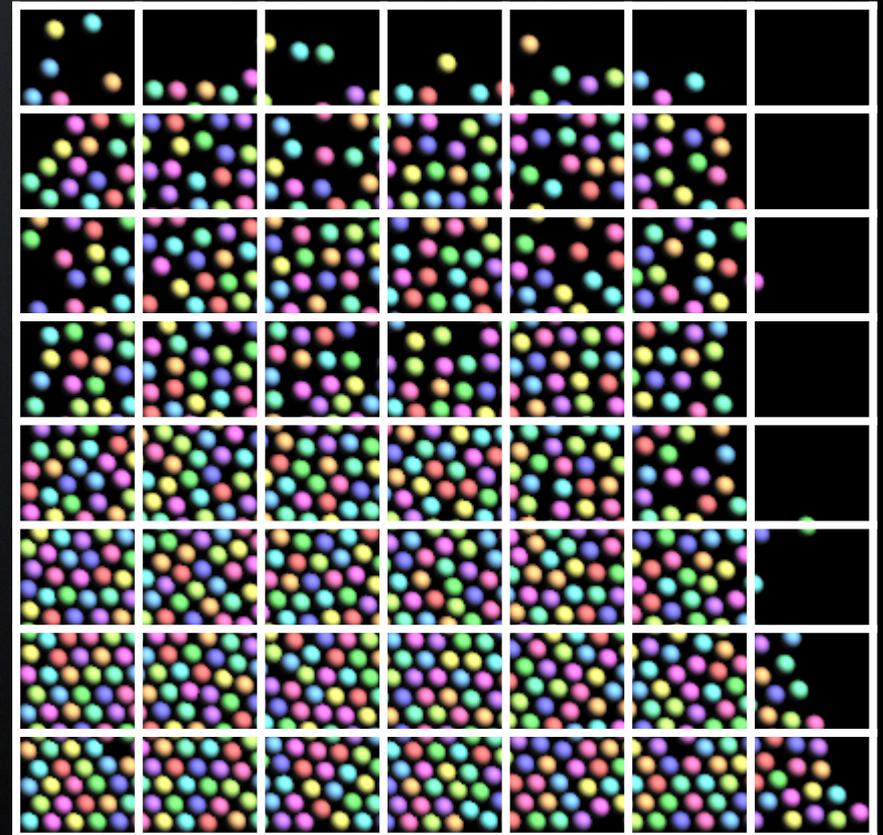


- Rigid body simulation <- Dynamic problems
  - Bodies move
  - Entities: Rigid bodies
  - Connections: Colliding pair



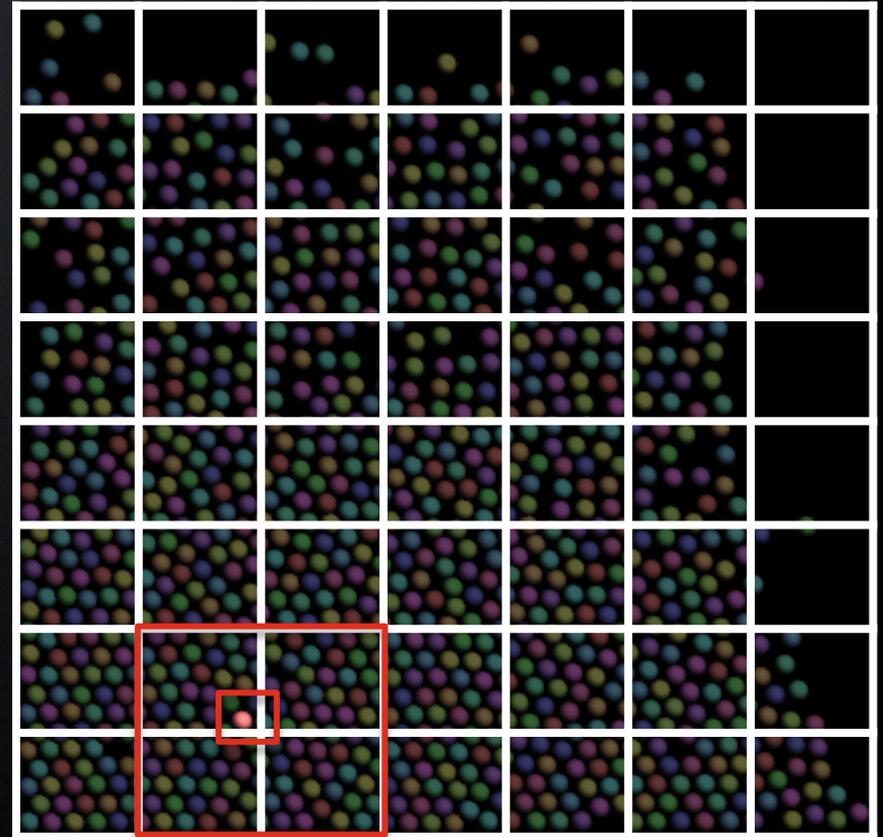


- Simple data structure
- Data is stored per cell
  
- Build
  - Overlay a grid
  - Find cell to which an object belongs
    - Uniform -> cell location can be calculated
  - Store reference to the cell





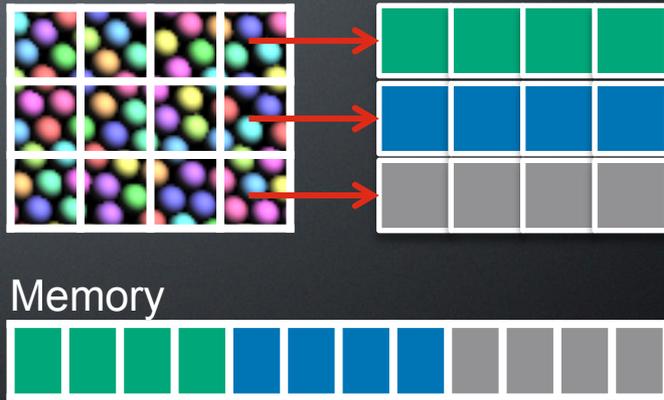
- Collision query
  - Find possible overlapping cells
  - Look up objects stored in the cell





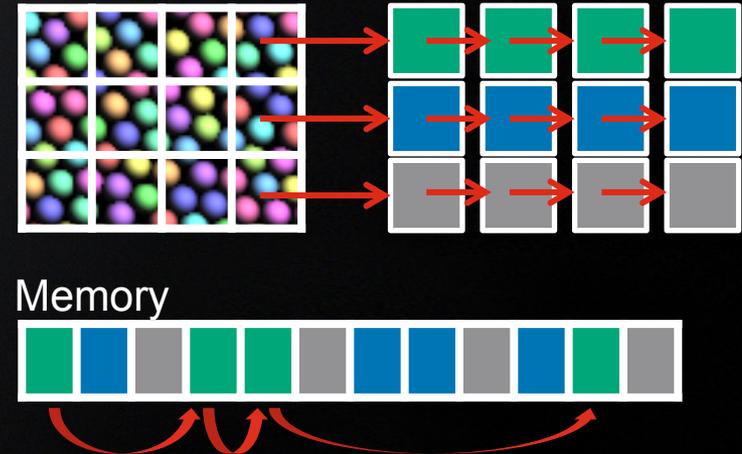
- Linear array

- An element stores data



- Linked list

- An element stores data and link





## Algorithm

- Linear array
  - Fixed sized buffer
    - Each cell has a counter
    - Pros: simple
    - Cons: large memory, fixed entry

## Enabler

- Pure graphics GPU
  - Tests in graphics pipeline



## Algorithm

- Linear array
  - Fixed sized buffer
    - Each cell has a counter
    - Pros: simple
    - Cons: large memory, fixed entry
  - Sort based
    - Sort by cell index
    - Pros: memory size
    - Cons: need a global sort

## Enabler

- Pure graphics GPU
    - Tests in graphics pipeline
- 
- On chip Local memory (LDS)
    - Fast sort (Radix sort)
    - Sort was slow (e.g., bitonic sort)



## Algorithm

- Linear array
  - Fixed sized buffer
    - Each cell has a counter
    - Pros: simple
    - Cons: large memory, fixed entry
  - Sort based
    - Sort by cell index
    - Pros: memory size
    - Cons: need a global sort
- Linked list
  - Each cell has a pointer to data
  - Pros: better for large problems
  - Cons: memory access pattern on traversal

## Enabler

- Pure graphics GPU
  - Tests in graphics pipeline
- On chip Local memory (LDS)
  - Fast sort (Radix sort)
  - Sort was slow (e.g., bitonic sort)
- Fast atomics
  - Intensive atomics was too expensive
  - Traversal was too slow

# WHICH ALGORITHM IS THE BEST??



**Eurographics 2012**

Cagliari, Italy

May 13-18

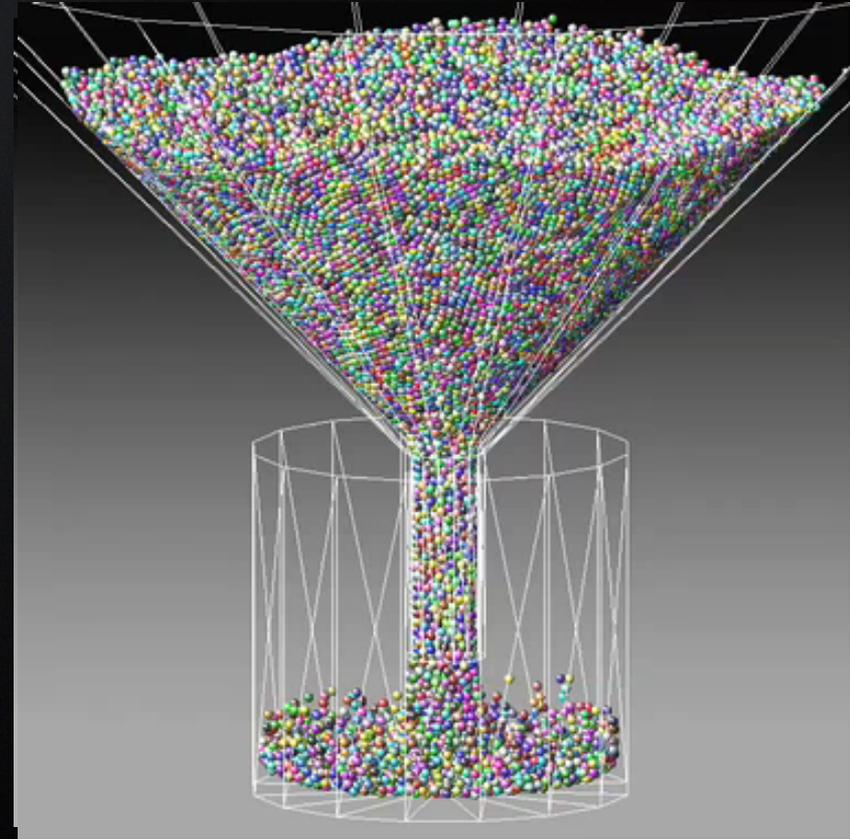


33<sup>rd</sup> ANNUAL CONFERENCE OF THE EUROPEAN ASSOCIATION FOR COMPUTER GRAPHICS

- Best algorithm depends on architecture
  - GPU or CPU
  - Different generation GPUs
- GPU is still changing architecture
  - Today's best algorithm isn't always the best tomorrow

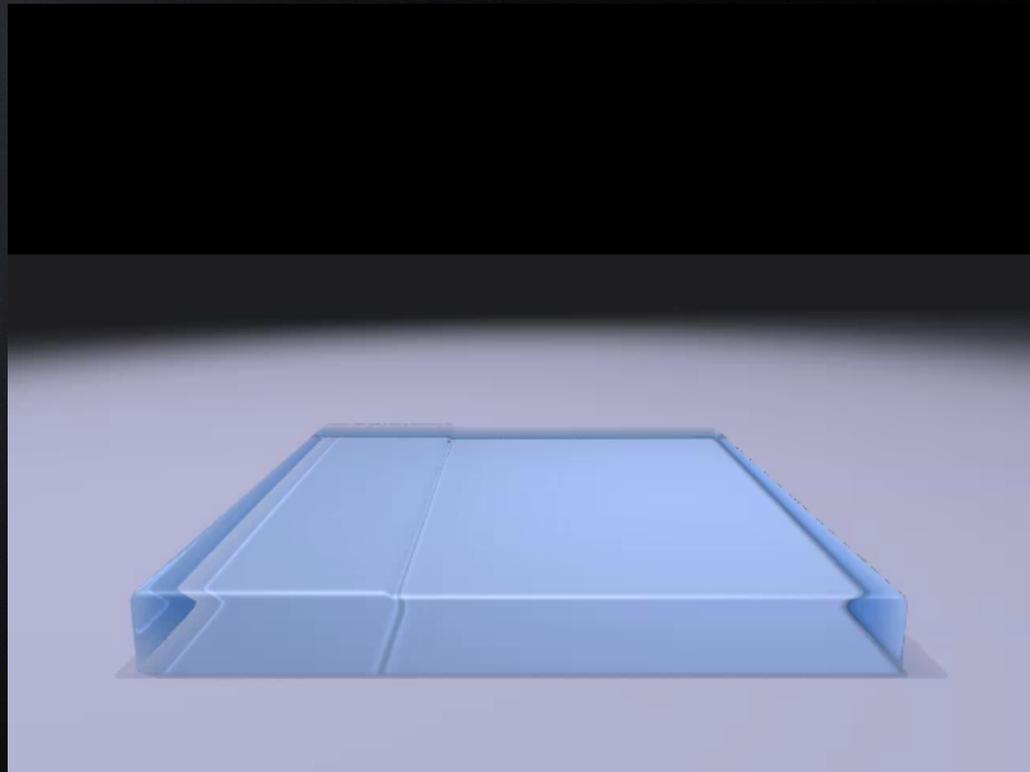


- Distinct Element Method (DEM)
- Interaction force
  - Damped spring





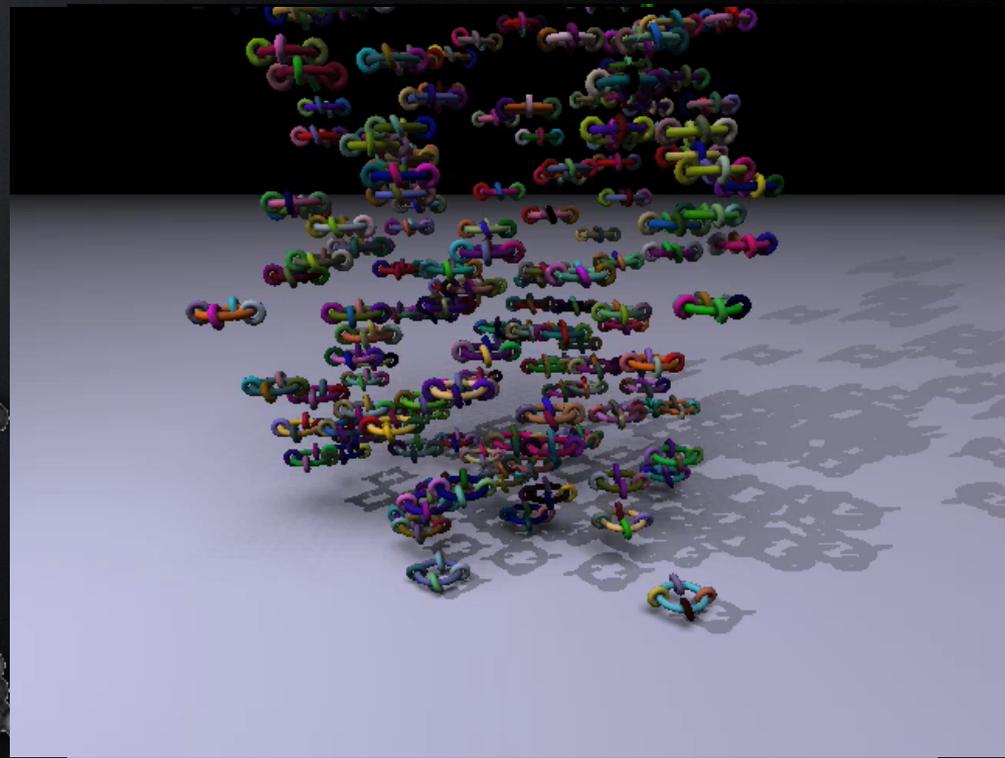
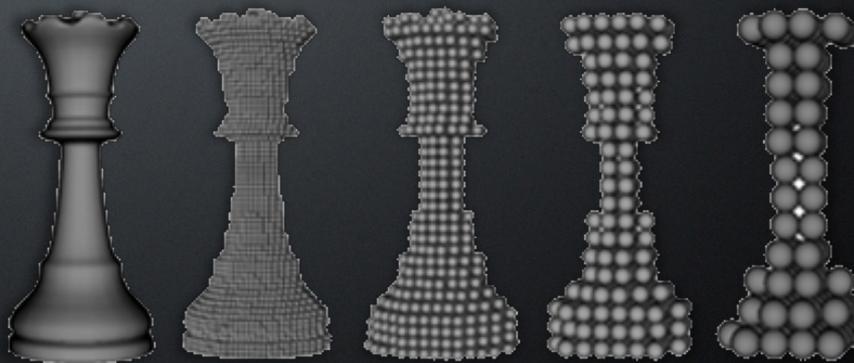
- Smoothed Particle Hydrodynamics
- Solving the Navier-Stokes equation on particles



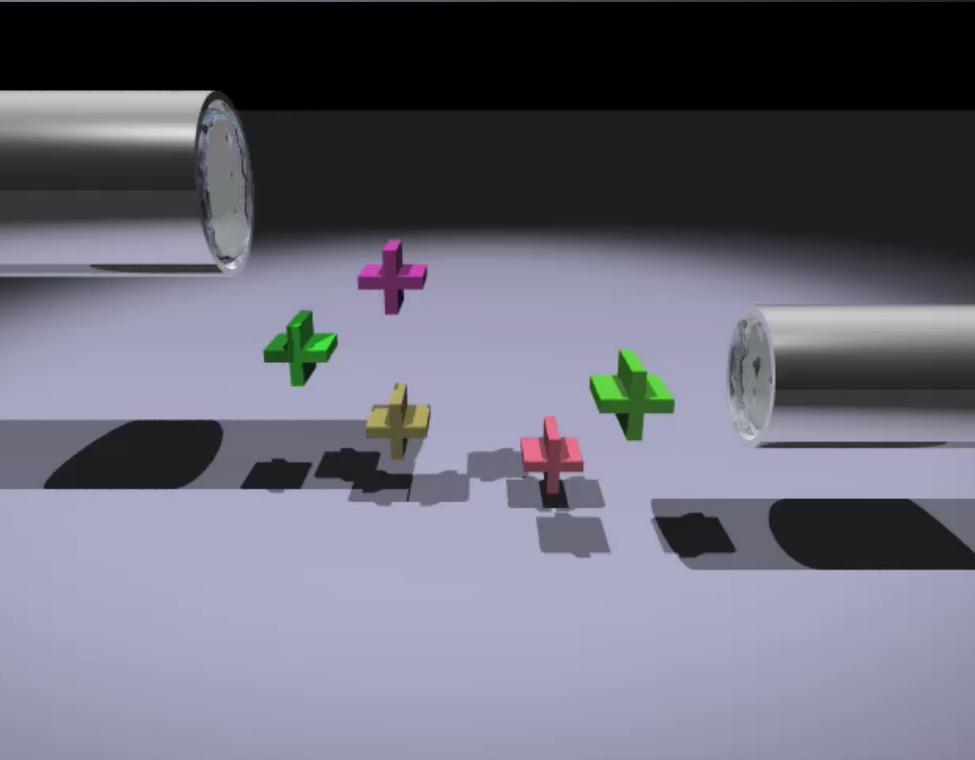
Harada et al., Smoothed Particle Hydrodynamics on GPUs, CGI (2007)



- Represent a rigid body with a set of particles
- Rigid body collision = Particle collision

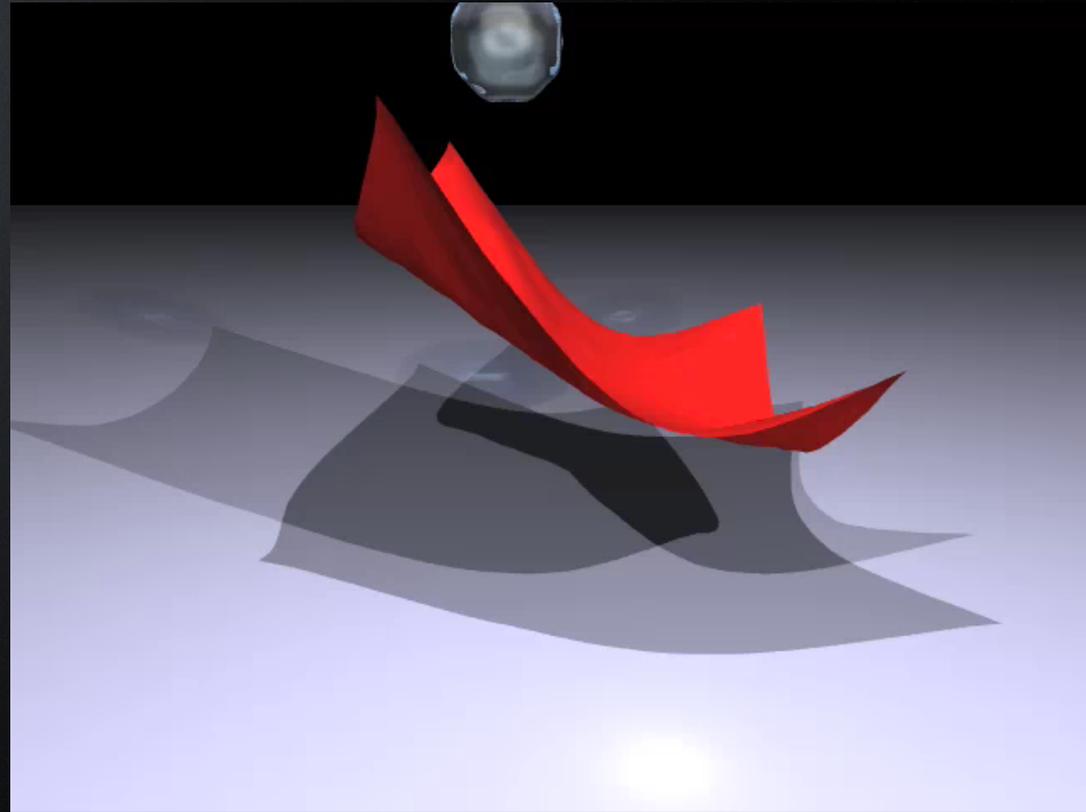


Harada et al., Real-time Rigid Body Simulation on GPUs, GPU Gems3 (2007)





- Particle v.s. triangle mesh collision
- Multiple uniform grids
  - 1<sup>st</sup> grid: particles
  - 2<sup>nd</sup> grid: triangles



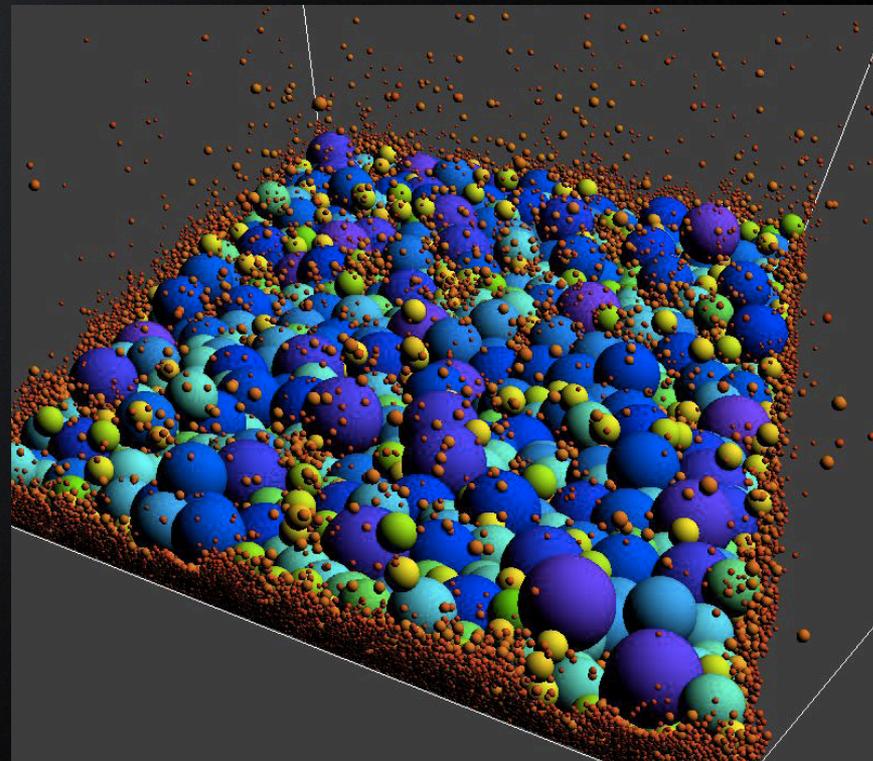
Harada et al., Real-time Fluid Simulation Coupled with Cloth, TPCG (2007)



- Grid
- Sweep & prune
  - Sweep
- Bounding volume hierarchy
  - Build
  
- Rigid body simulation
  - Broad-phase collision detection
  - Narrow-phase collision detection



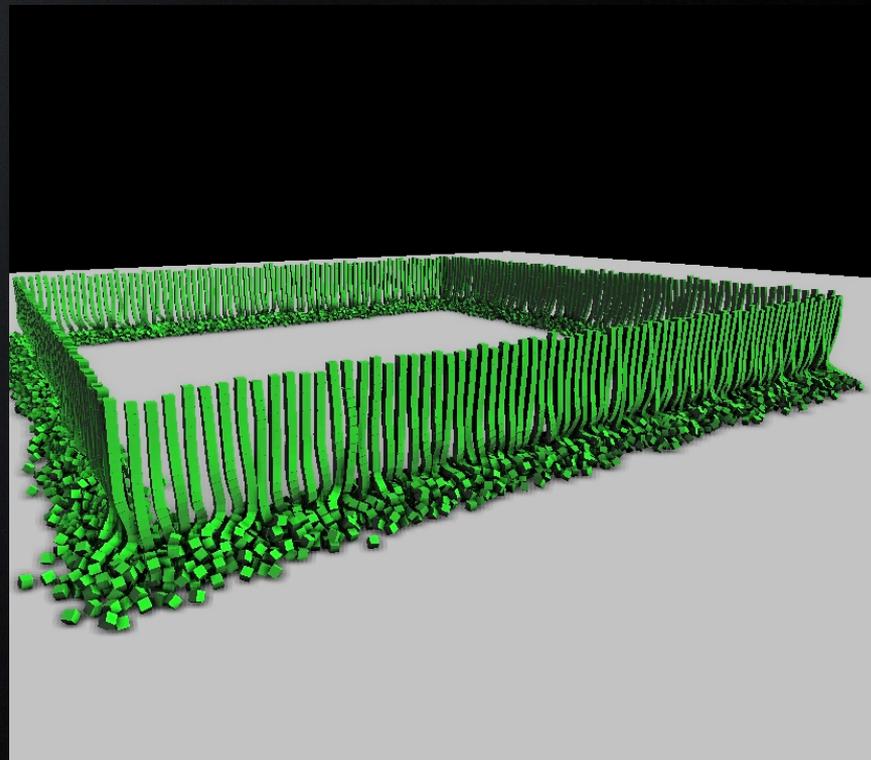
- Sweep & prune
  - Sort end points
  - Sweep sorted list
- Optimizations
  - Split sweep of a large object
  - Workspace subdivision
  - Cell subdivision
- Detail
  - Young J. Kim's presentation



Liu et al., Real-time Collision Culling of a Million Bodies on Graphics Processing Units, TOG(2010)

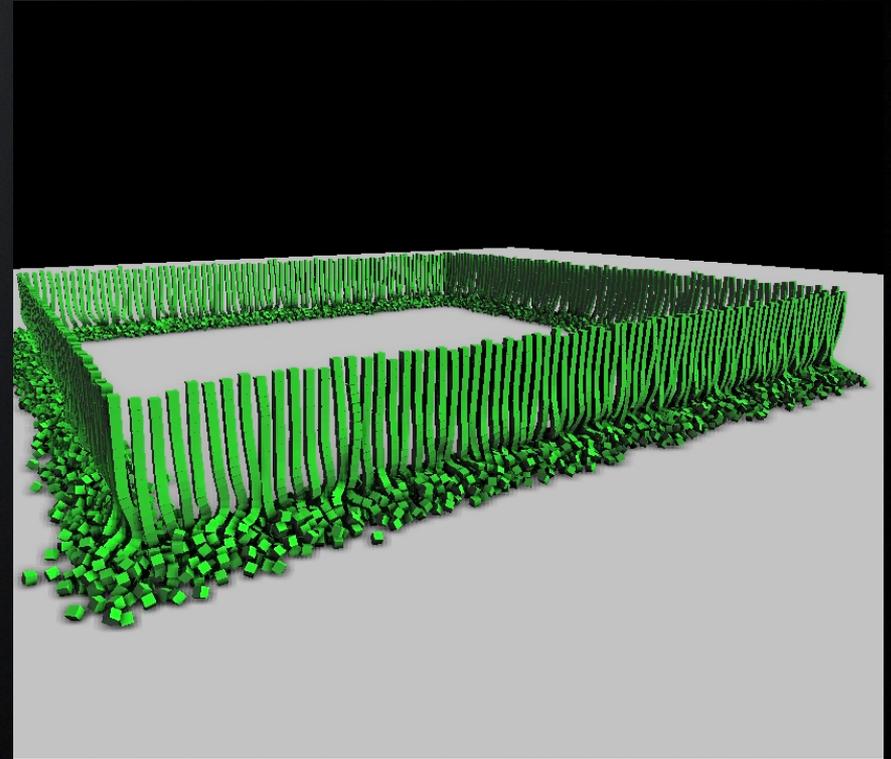
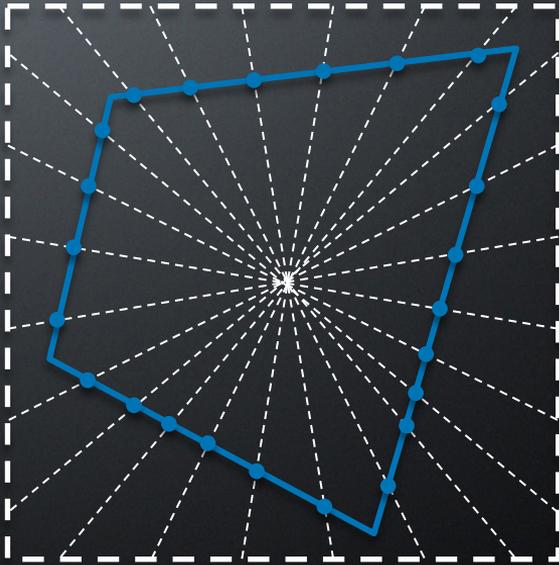


- Variety of computation
  - Many shape combinations -> divergence
    - Box-box, box-sphere, box-convex, ...
  - Sequential algorithm
    - GJK





- Unified shape representation
- Parallelize per pair, parallelize in a pair





- Distance Query

```
Parallel For All Surface Points
```

```
  d = Calculate distance
```

```
  if( d < surface height )
```

```
    Append Contact Point
```

- Contact point reduction

```
p0 = Parallel Reduce(Direction0)
```

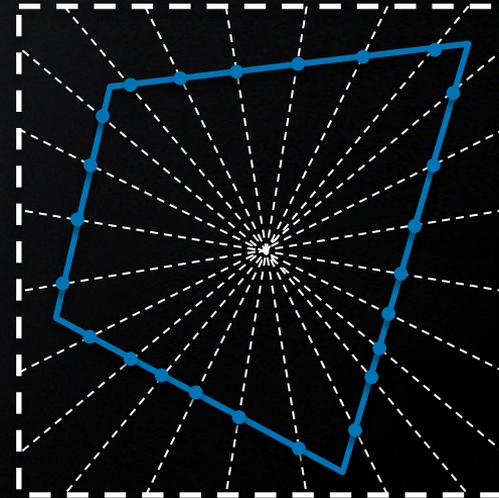
```
p1 = Parallel Reduce(Direction1)
```

```
p2 = Parallel Reduce(Direction2)
```

```
p3 = Parallel Reduce(Direction3)
```

- Export

```
Append contact(p0, p1, p2, p3)
```

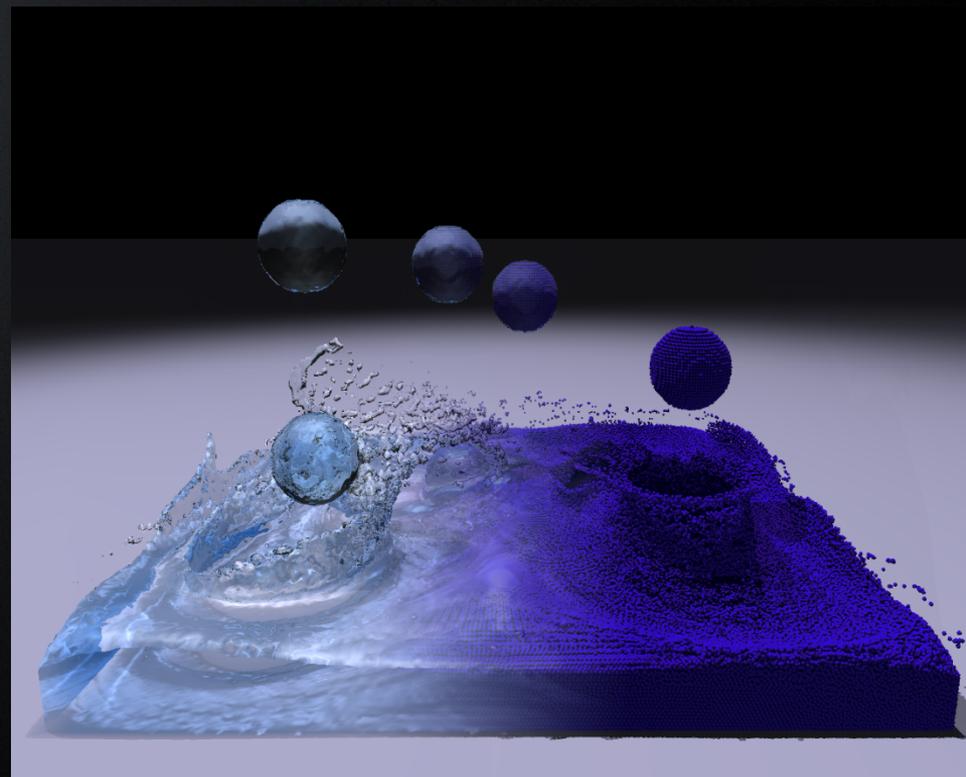




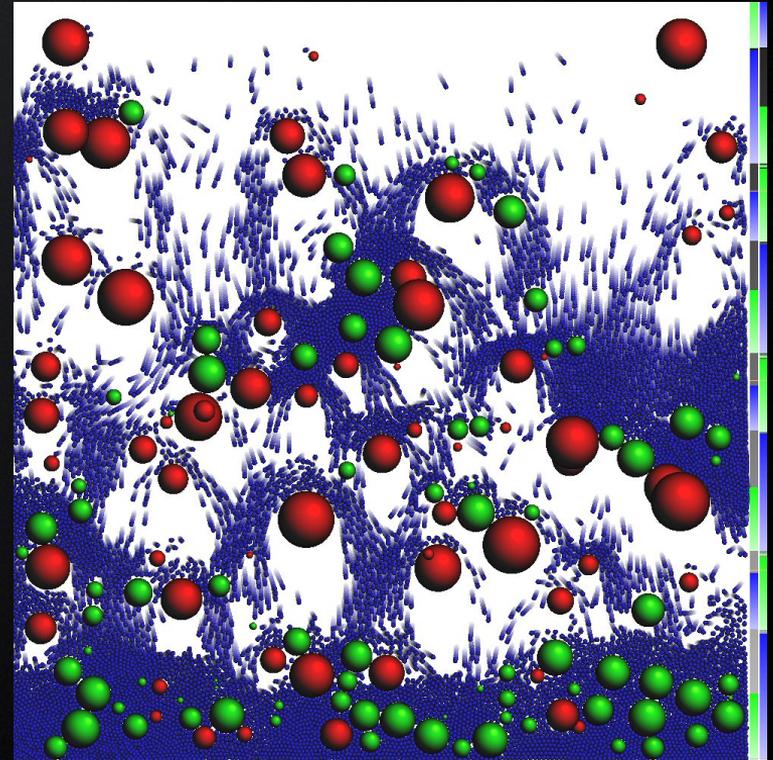
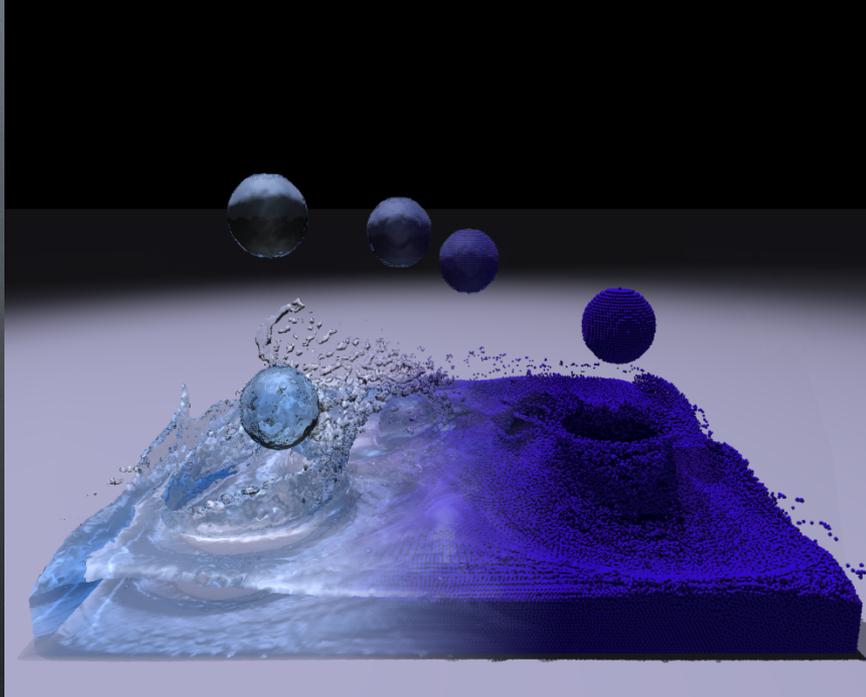
# ***HETEROGENEOUS COLLISION DETECTION***



- Large number of particles
- Particles with identical size
  - Work granularity is almost the same
  - Good for the wide SIMD architecture

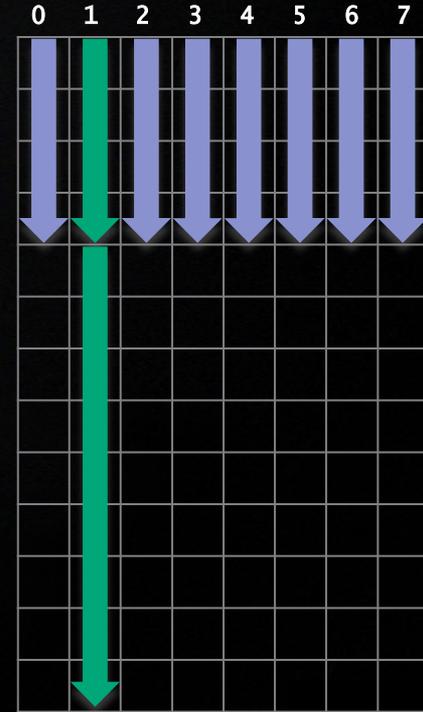
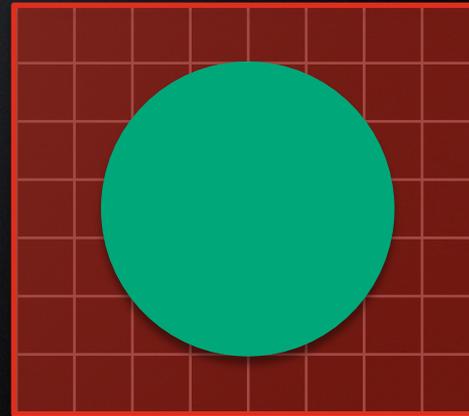
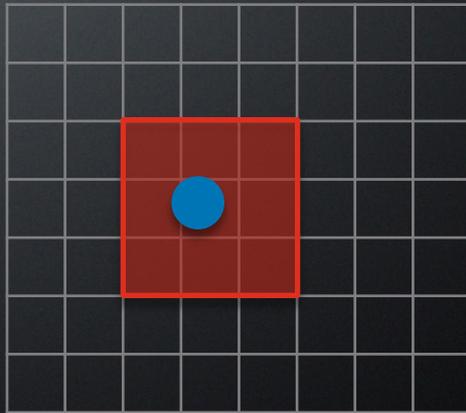


Harada et al. 2007



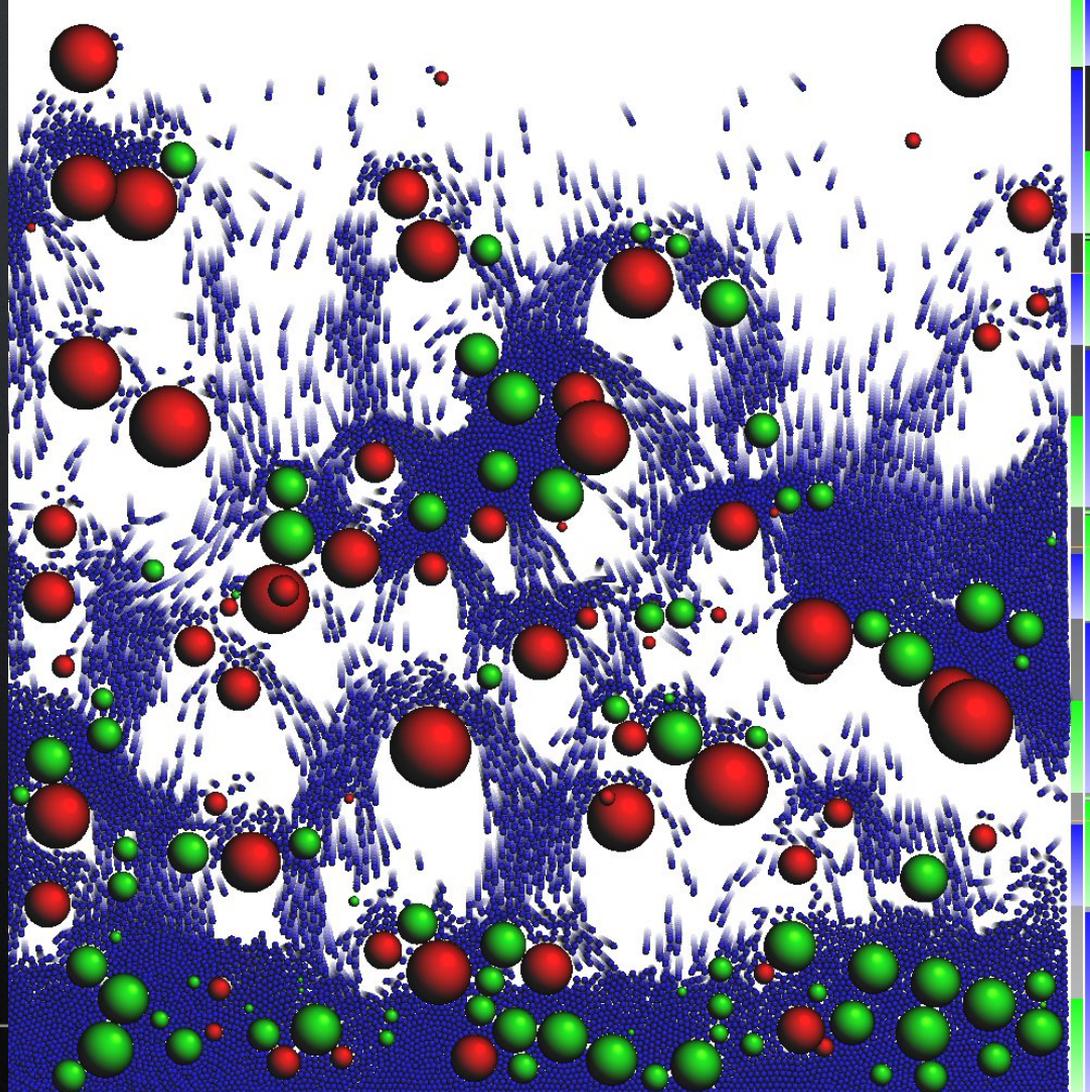


- Not only small particles
- Difficulty for GPUs
  - Large particles interact with small particles
  - Large-large collision



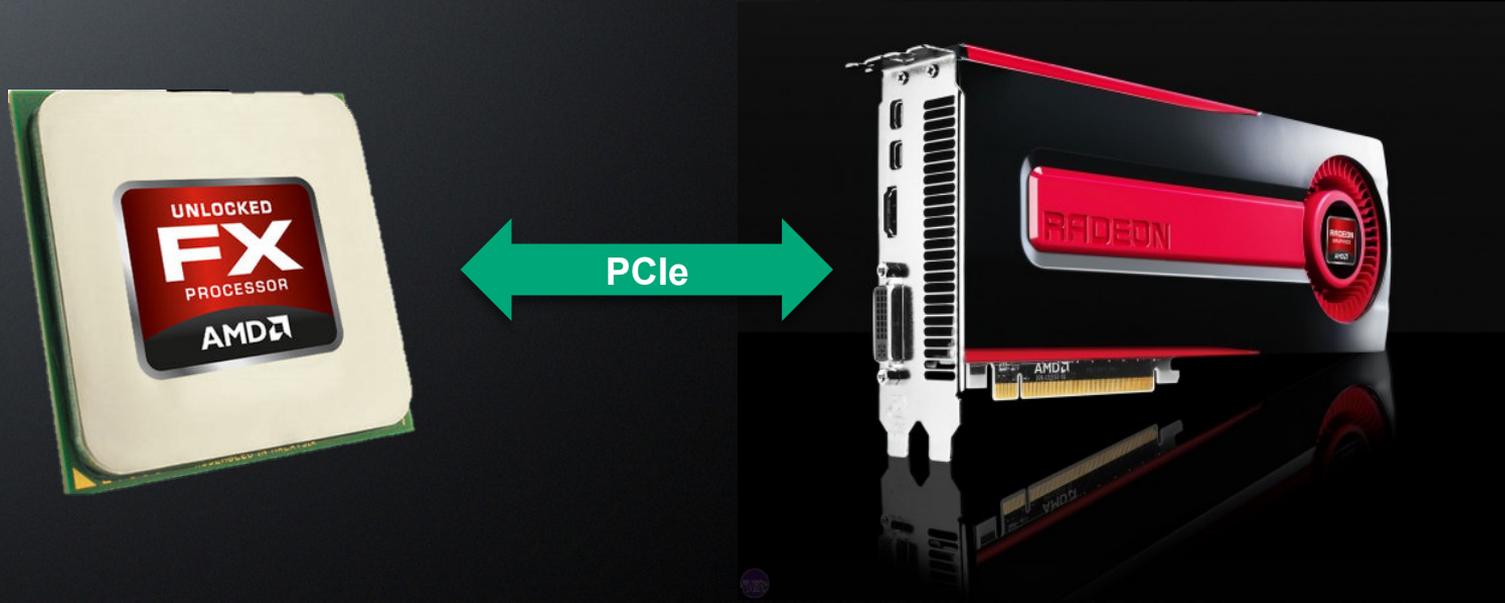
# CHALLENGE

- Non uniform work granularity
  - Small-small(SS) collision
    - Uniform, GPU
  - Large-large(LL) collision
    - Non Uniform, CPU
  - Large-small(LS) collision
    - Non Uniform, CPU
- Can use a CPU and a discrete GPU
  - PCIe data transfer bottleneck



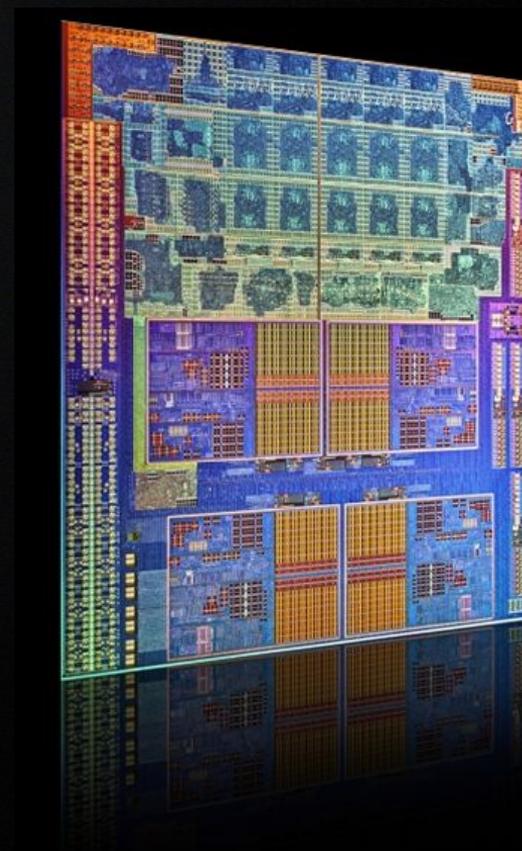


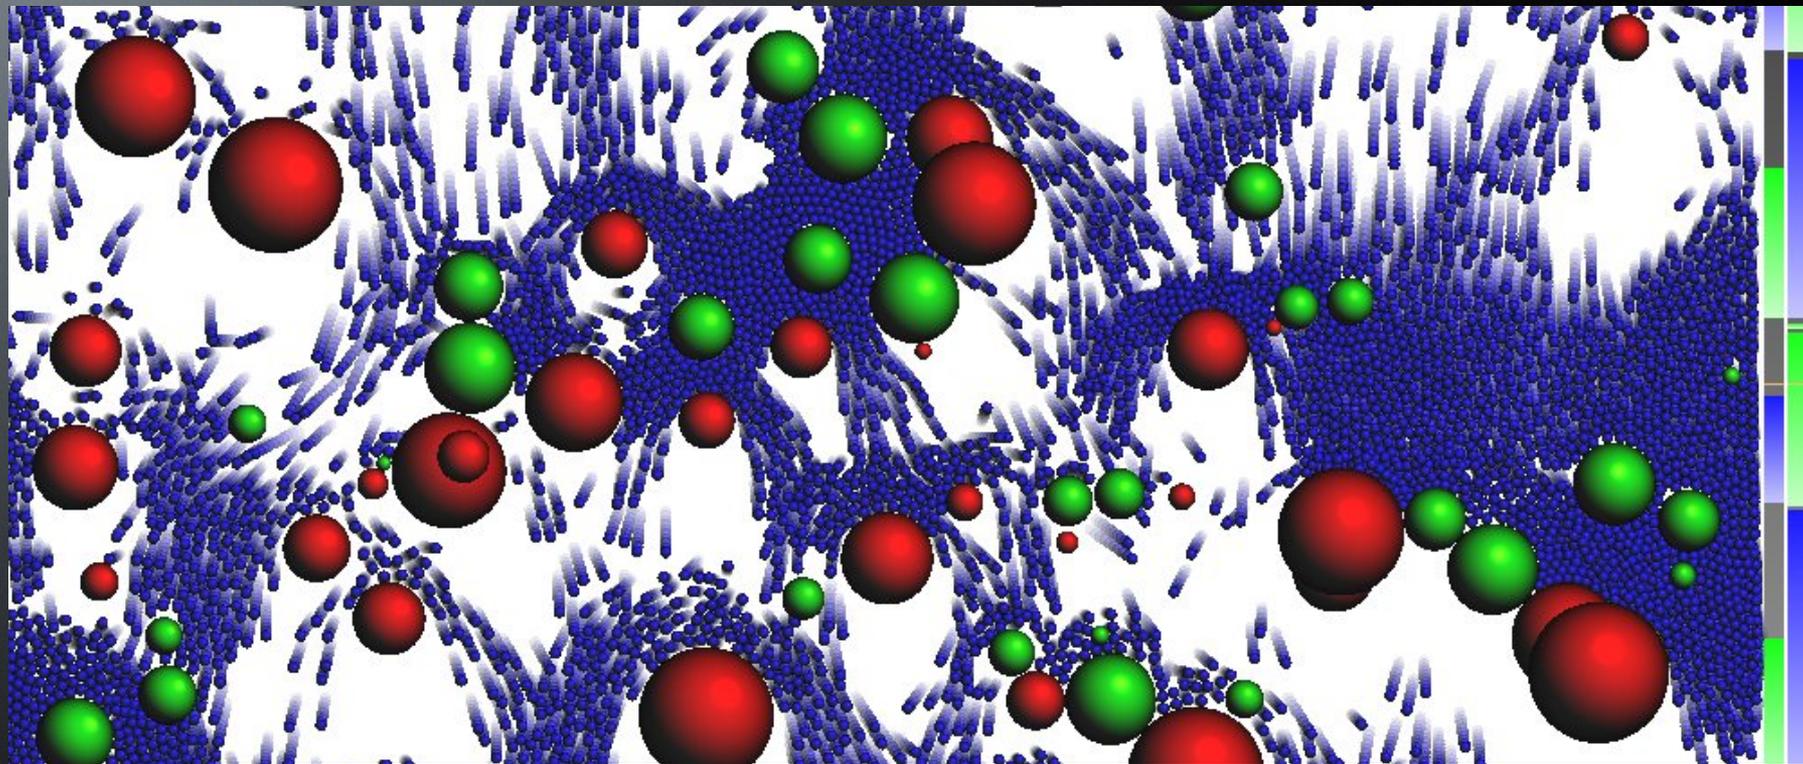
- If we can use CPU and GPU, problem is solved
- PCI Express data transfer kills the performance





- CPU + GPU on the same die
- AMD Fusion (APU), Intel Sandybridge
- No PCI Express bus
- Memory sharing
- Able to dispatch a job to the GPU without data transfer
  - Serial work: dispatch to the CPU
  - Parallel work: dispatch to the GPU
- Can use this for a non-uniform simulation





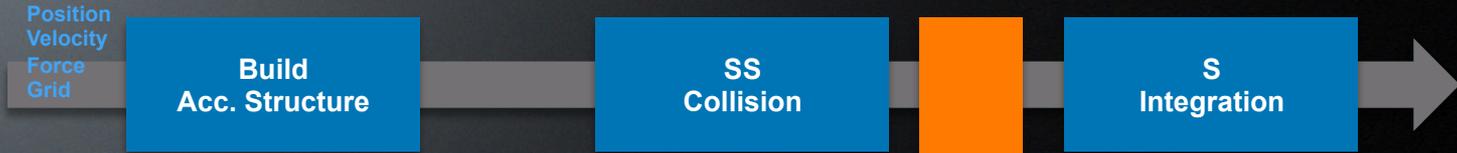


# ***METHOD***

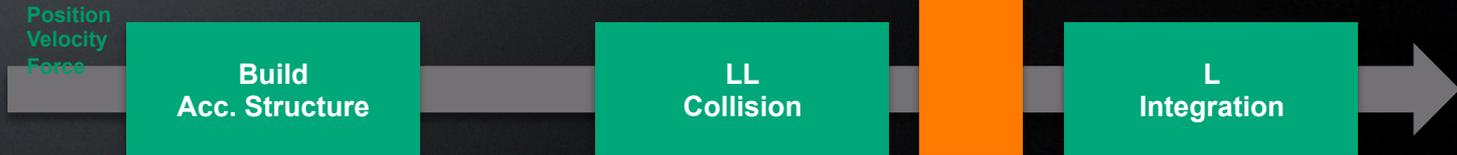
# TWO SIMULATIONS



- Small particles

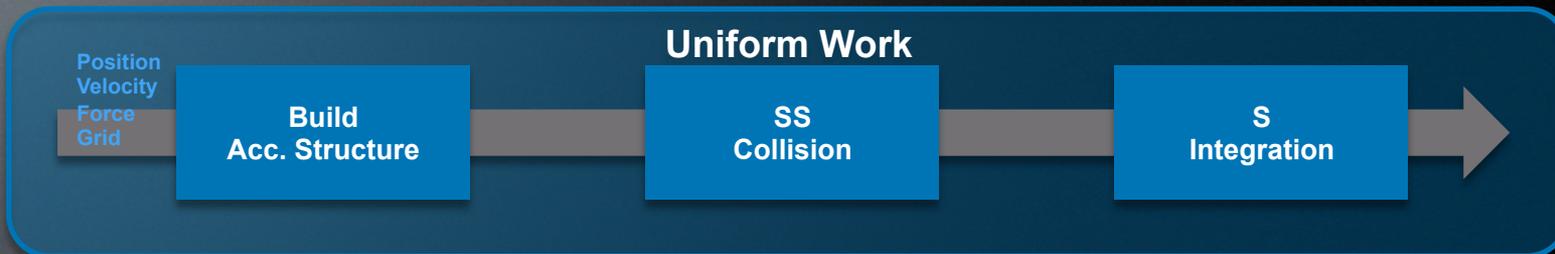


- Large particles





- Small particles



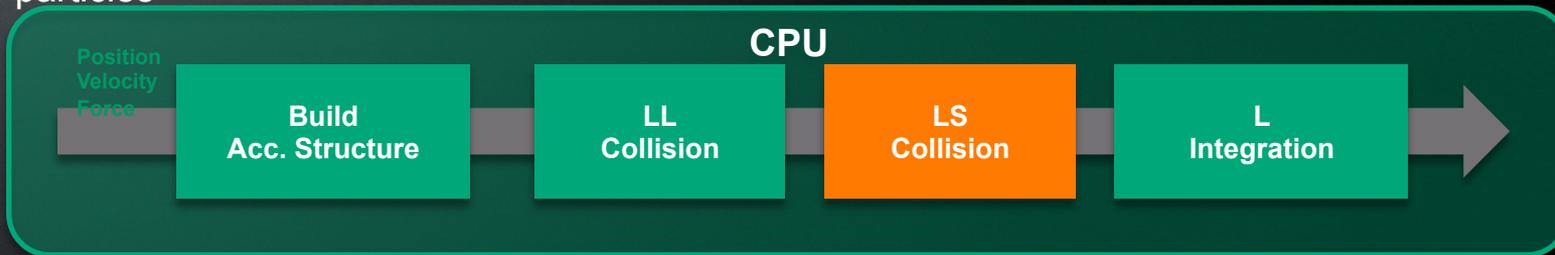
- Large particles



- Small particles

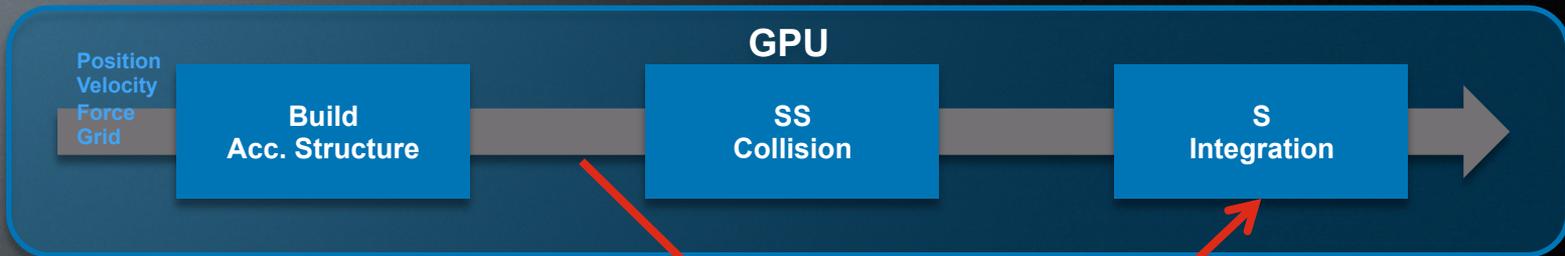


- Large particles

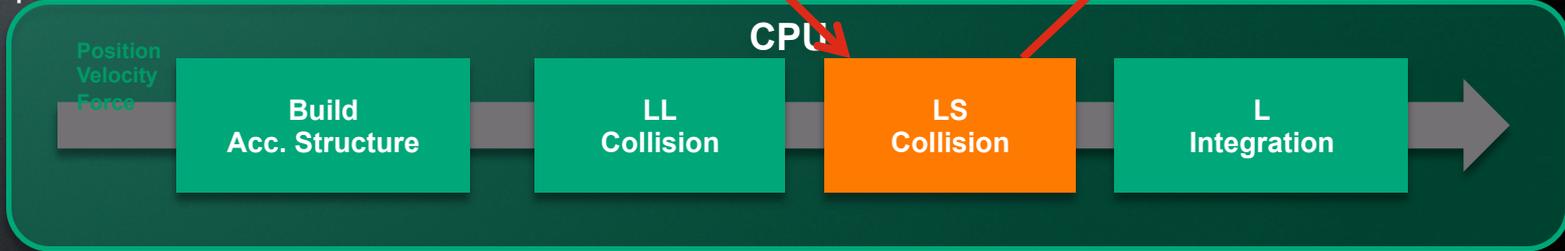




- Small particles



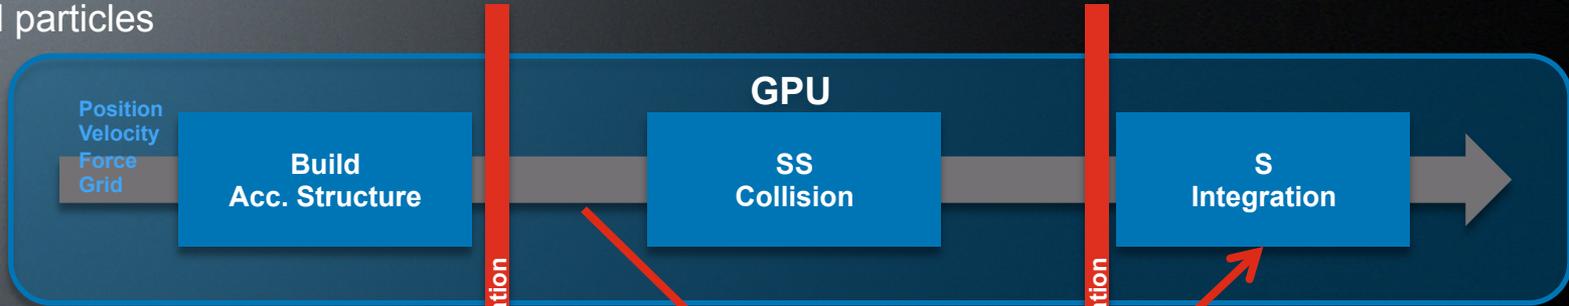
- Large particles



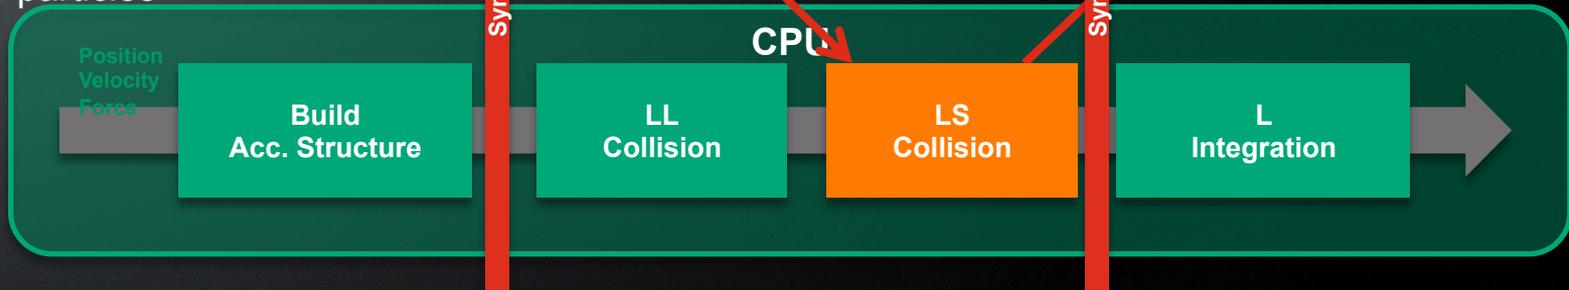
- Grid, small particle data has to be shared with the CPU for LS collision
  - Allocated as zero copy buffer



- Small particles

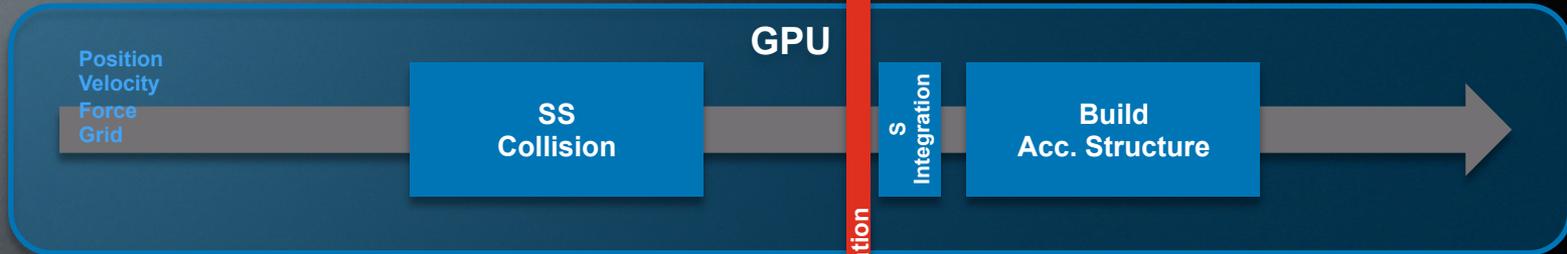


- Large particles



- Grid, small particle data has to be shared with the CPU for LS collision
  - Allocated as zero copy buffer

- Small particles



- Large particles

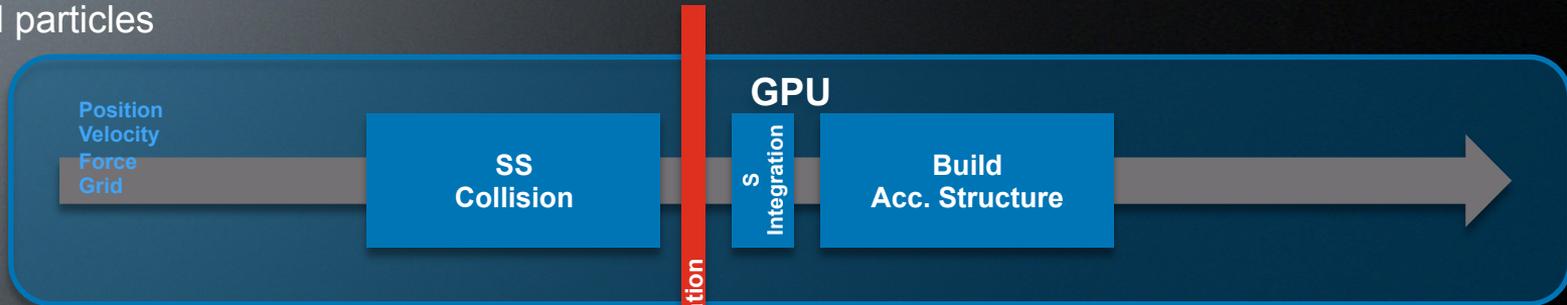


- Grid construction can be moved at the end of the pipeline

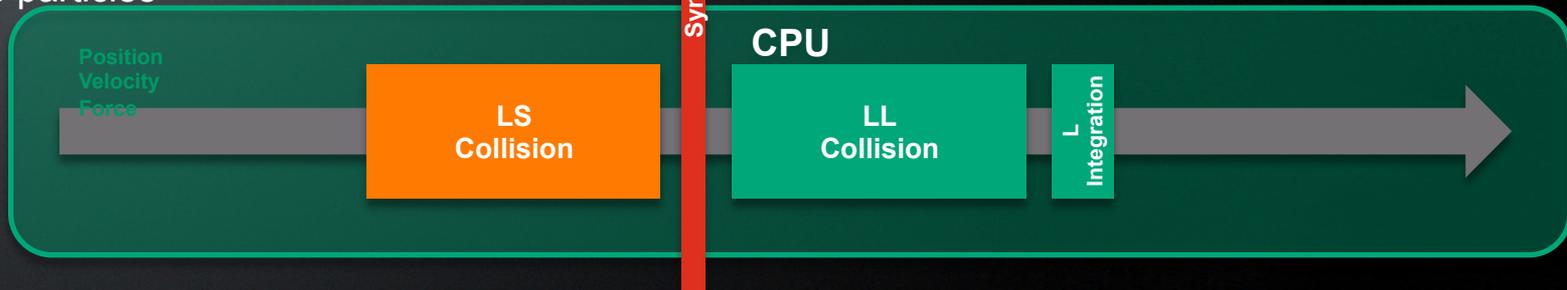
- Unbalanced workload



- Small particles



- Large particles

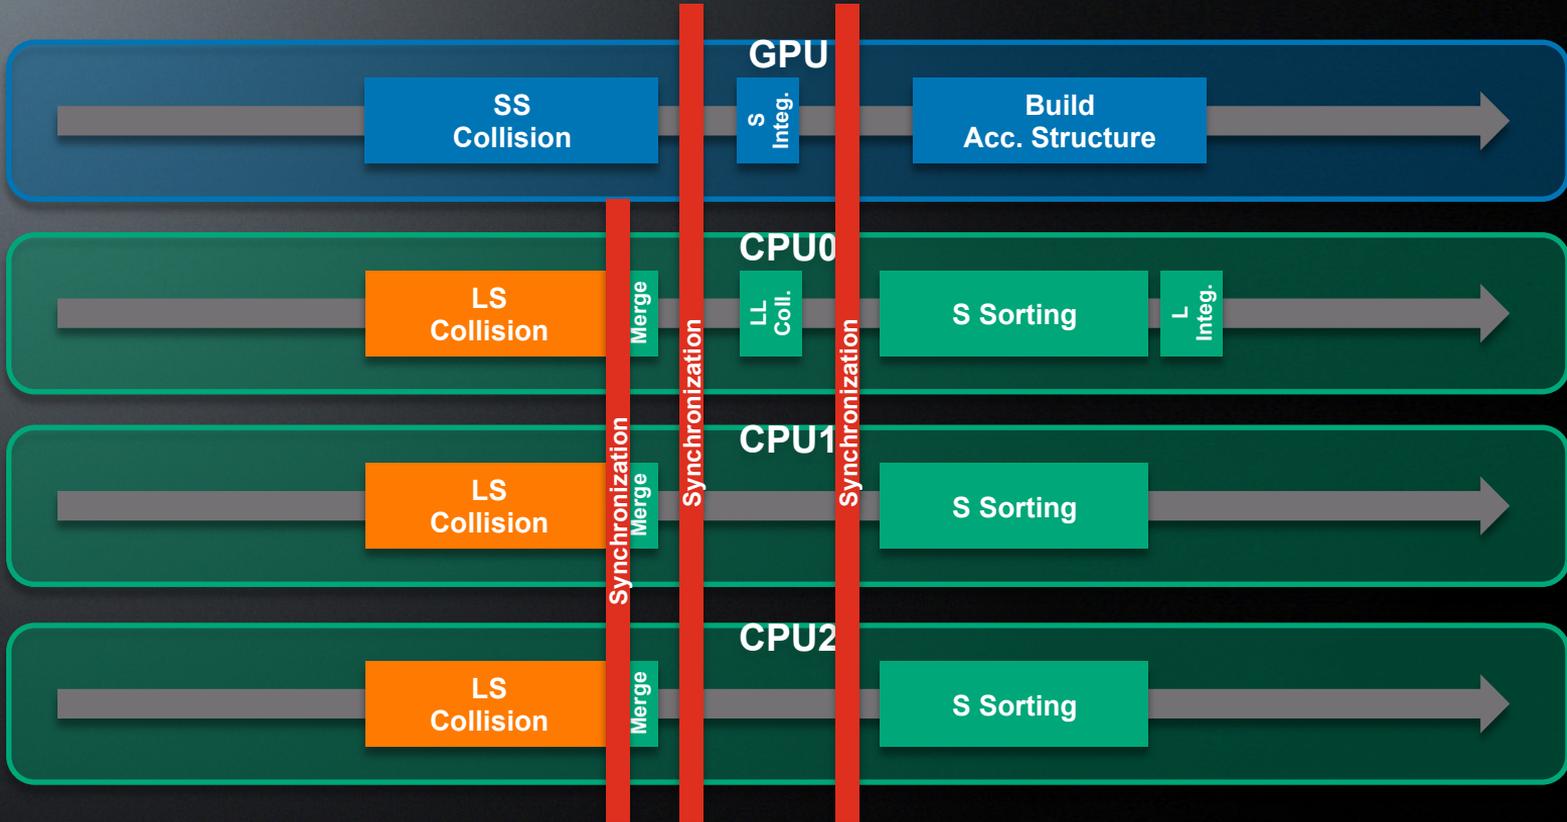


- To get better load balancing

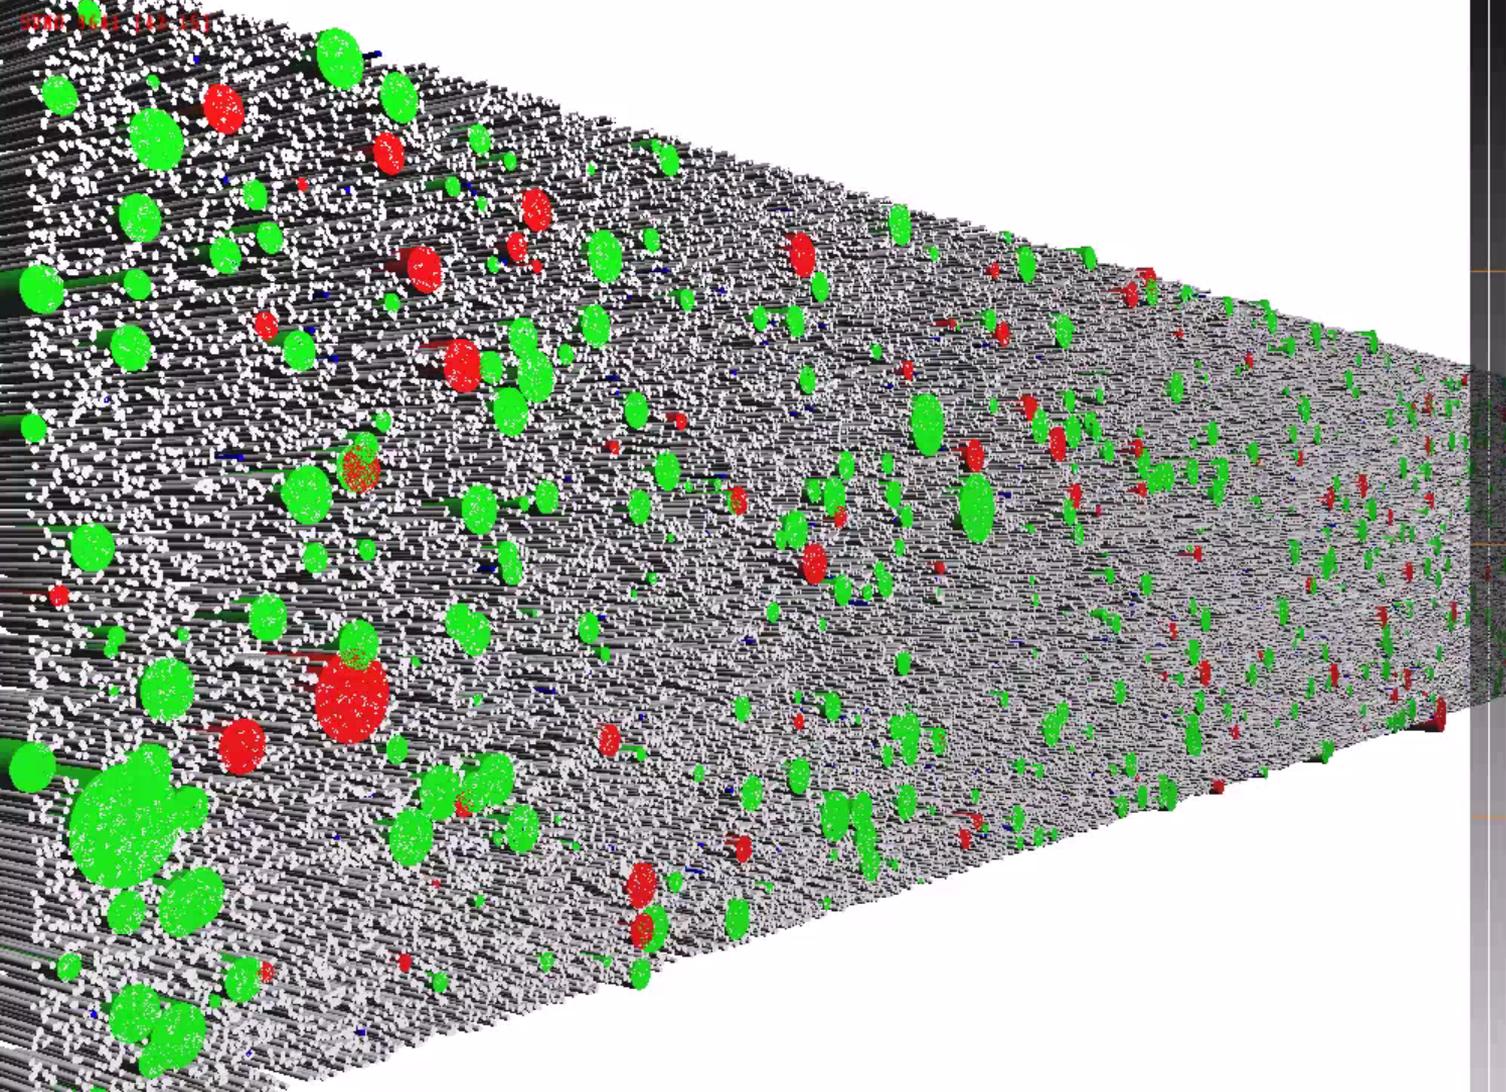
- The sync is for passing the force buffer filled by the CPU to the GPU
- Move the LL collision after the sync



# ***MULTI THREADING (4 THREADS)***



Harada, Heterogeneous Particle-Based Simulation, SIG ASIA Talk(2011)



graphics 2012

May 13-18



EUROPEAN ASSOCIATION FOR COMPUTER GRAPHICS

CPU Work

GPU Work



***QUESTIONS?***

# PROBLEM SOLVED??

- Uniform grid can solve a lot of problems
- However, uniform grid is not perfect solution
- More sophisticated collision detection is necessary
  
- Rigid body simulation
  - 2 step collision detection
  - Broad-phase collision detection
    - Grid, Sweep & prune, Tree
  - Narrow-phase collision detection
    - SAT, GJK, approximated solutions