

INFOGR – Computer Graphics

J. Bikker - April-July 2016 - Lecture 7: “Accelerate”

Welcome!



Today's Agenda:

- High-speed Ray Tracing
- Acceleration Structures
- The Bounding Volume Hierarchy
- BVH Construction
- BVH Traversal
- Optimizing Construction
- High-speed Traversal



High-speed Ray Tracing

Ray Tracing – Needful things

Whitted-style ray tracing:

1 primary ray per pixel

1 shadow ray per pixel per light

Optional: rays for reflections & refraction

Estimate:

- 10 rays per pixel
- 1M pixels (~1280x800)
- 30 fps

→ 300Mrays/s

How does one intersect 300Mrays/s on a 3Ghz CPU?

Easy: use no more than 10 cycles per ray.



High-speed Ray Tracing

Actually...

- We have 8 cores (so 80 cycles)
- Executing AVX code (so 640 cycles)
- Plus 20% gains from hyperthreading (768 cycles).

But really...

Assuming we get a linear increase in performance for the number of cores and AVX, how do we intersect thousands of triangles in 768 cycles?



High-speed Ray Tracing

Optimization

1. Measure: performance & scalability
2. High level optimizations: improve algorithmic complexity
3. Low level optimization: instruction level & thread-level parallelism, caching
4. GPGPU

More in the master course Optimization & Vectorization.



Today's Agenda:

- High-speed Ray Tracing
- Acceleration Structures
- The Bounding Volume Hierarchy
- BVH Construction
- BVH Traversal
- Optimizing Construction
- High-speed Traversal



High-speed Ray Tracing

Optimization: reduce algorithmic complexity

Complexity:

number of ray/primitive intersections

= pixels * paths per pixel * average path length * primitives


= 1M * 1 * 2 * 1M

```

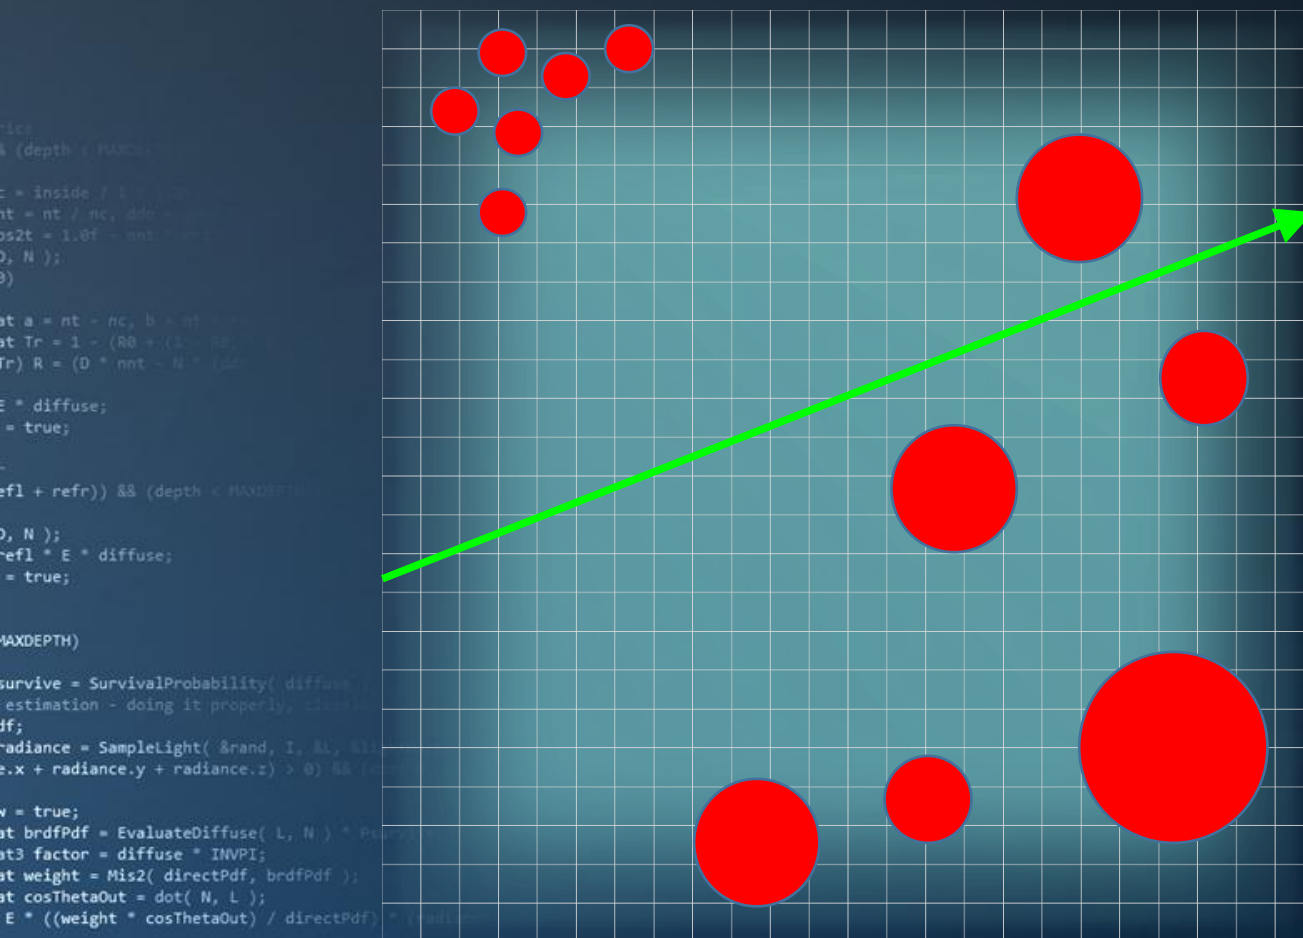
...
    & (depth < MAXDEPTH)
{
    // Inside / Outside test
    nt = nt / nc; add = add / nc;
    pos2t = 1.0f - nnt; // weight
    D, N );
}

// Russian roulette
// at a = nt - nc, b = nt - nc;
// at Tr = 1 - (R0 + (1 - R0) * a);
// Tr) R = (D * nnt - N * (add
//
// E * diffuse;
// = true;
//
//
// refl + refr)) && (depth < MAXDEPTH)
//
// D, N );
// refl * E * diffuse;
// = true;
//
// MAXDEPTH)
//
// survive = SurvivalProbability( diffuse );
// estimation - doing it properly, close to
// if;
// radiance = SampleLight( &rand, I, &t, &light);
// e.x + radiance.y + radiance.z) > 0) && (cosTheta > 0)
//
// v = true;
// at brdfPdf = EvaluateDiffuse( L, N ) * Psurvive;
// at3 factor = diffuse * INVPI;
// at weight = Mix2( directPdf, brdfPdf );
// at cosThetaOut = dot( N, L );
// E * ((weight * cosThetaOut) / directPdf) * (radiance
//
// random walk - done properly, closely following well known
// survive)
//
//
// at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf);
// survive;
// pdf;
// n = E * brdf * (dot( N, R ) / pdf);
// sion = true;

```



Acceleration Structures



Option 1:

Use a grid.

- Each grid cell has a list of primitives that overlap it.
- The ray traverses the grid, and intersects only primitives in the grid cells it visits.

Problems:

- Many primitives will be checked more than once.
- It costs to traverse the grid.
- How do we choose grid resolution?
- What if scene detail is not uniform?




Acceleration Structures

```

    // Ray-sphere intersection
    float t = inside / 1.0f;
    float nt = nt / nc;
    float cos2t = 1.0f - nnt * nnt;
    if (cos2t < 0) return 0;
    float t = t * cos2t;
    float D, N;
    // Ray-sphere intersection
    float a = nt - nc;
    float Tr = 1 - (R0 + (1 - R0) * t);
    float R = (D * nnt - N * (1 - R0));
    // Ray-sphere intersection
    float E * diffuse;
    // Ray-sphere intersection
    float refl + refr;
    // Ray-sphere intersection
    float D, N;
    float refl * E * diffuse;
    // Ray-sphere intersection
    float MAXDEPTH;
    // Ray-sphere intersection
    float survive = SurvivalProbability( diffuse );
    // Ray-sphere intersection
    float estimation - doing it properly, closely following
    // Ray-sphere intersection
    float radiance = SampleLight( &rand, I, &t, &R, &pdf );
    // Ray-sphere intersection
    float e.x + radiance.y + radiance.z > 0;
    // Ray-sphere intersection
    float v = true;
    // Ray-sphere intersection
    float brdfPdf = EvaluateDiffuse( L, N );
    // Ray-sphere intersection
    float factor = diffuse * INVPI;
    // Ray-sphere intersection
    float weight = Mis2( directPdf, brdfPdf );
    // Ray-sphere intersection
    float cosThetaOut = dot( N, L );
    // Ray-sphere intersection
    float E * ((weight * cosThetaOut) / directPdf) * (radiance);
    // Ray-sphere intersection
    float random walk - done properly, closely following
    // Ray-sphere intersection
    float survive;
    // Ray-sphere intersection
    float pdf;
    // Ray-sphere intersection
    float n = E * brdf * (dot( N, R ) / pdf);
    // Ray-sphere intersection
    float ion = true;

```



Option 2:

Use a nested grid.


- We use fewer cells. Each grid cell that overlaps multiple primitives has a smaller grid in it.
- The ray rapidly traverses empty space, and checks the nested grids when needed.

Problems:

- How do we choose grid resolutions?
- Is this the optimal way to traverse space?



Acceleration Structures



Option 3:

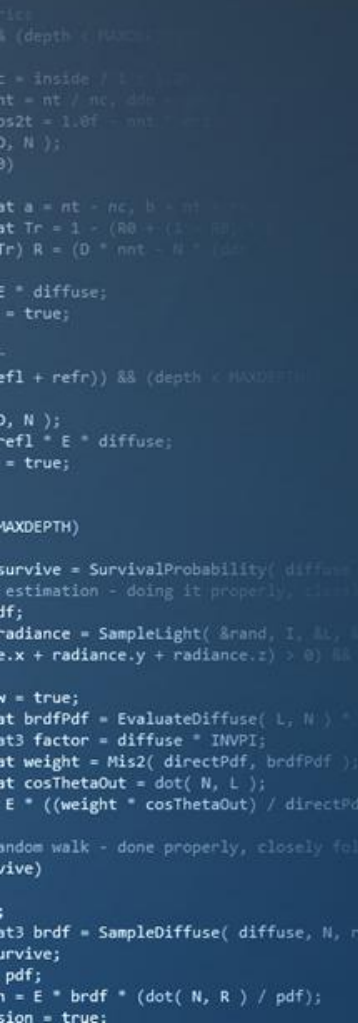
Use an octree.

- We start with a bounding box of the scene;
- The box is recursively subdivided in 8 equal boxes as long as it contains more than X primitives.

Problems:

- What if all the detail is exactly in the centre of the scene?
- Splitting in 8 boxes: is that the optimal subdivision?





Use an kD -tree.

- ## Problems:

- Primitives may end up in multiple leaf nodes.
- How hard is it to build such a tree?



```

    rics
    & (depth < MAXD);

    t = inside / 1.5;
    nt = nt / nc; dde = dde / nc;
    pos2t = 1.0f - nnt * dde;
    D, N );
    )

    at a = nt - nc, b = nt - nc;
    at Tr = 1 - (R0 + (1 - R0) * t);
    Tr) R = (D * nnt - N * (d

    E * diffuse;
    = true;

    -
    refl + refr)) && (depth < MAXDEPTH)

    D, N );
    refl * E * diffuse;
    = true;

    MAXDEPTH)

    survive = SurvivalProbability( diffuse );
    estimation - doing it properly, closely following
    if;
    radiance = SampleLight( &rand, I, &t, &light;
    e.x + radiance.y + radiance.z) > 0) && (exit &
    w = true;
    at brdfPdf = EvaluateDiffuse( L, N ) * Psurvive;
    at3 factor = diffuse * INVPI;
    at weight = Mis2( directPdf, brdfPdf );
    at cosThetaOut = dot( N, L );
    E * ((weight * cosThetaOut) / directPdf) * (radiant

    random walk - done properly, closely following well
    rive)

    ;
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;

```

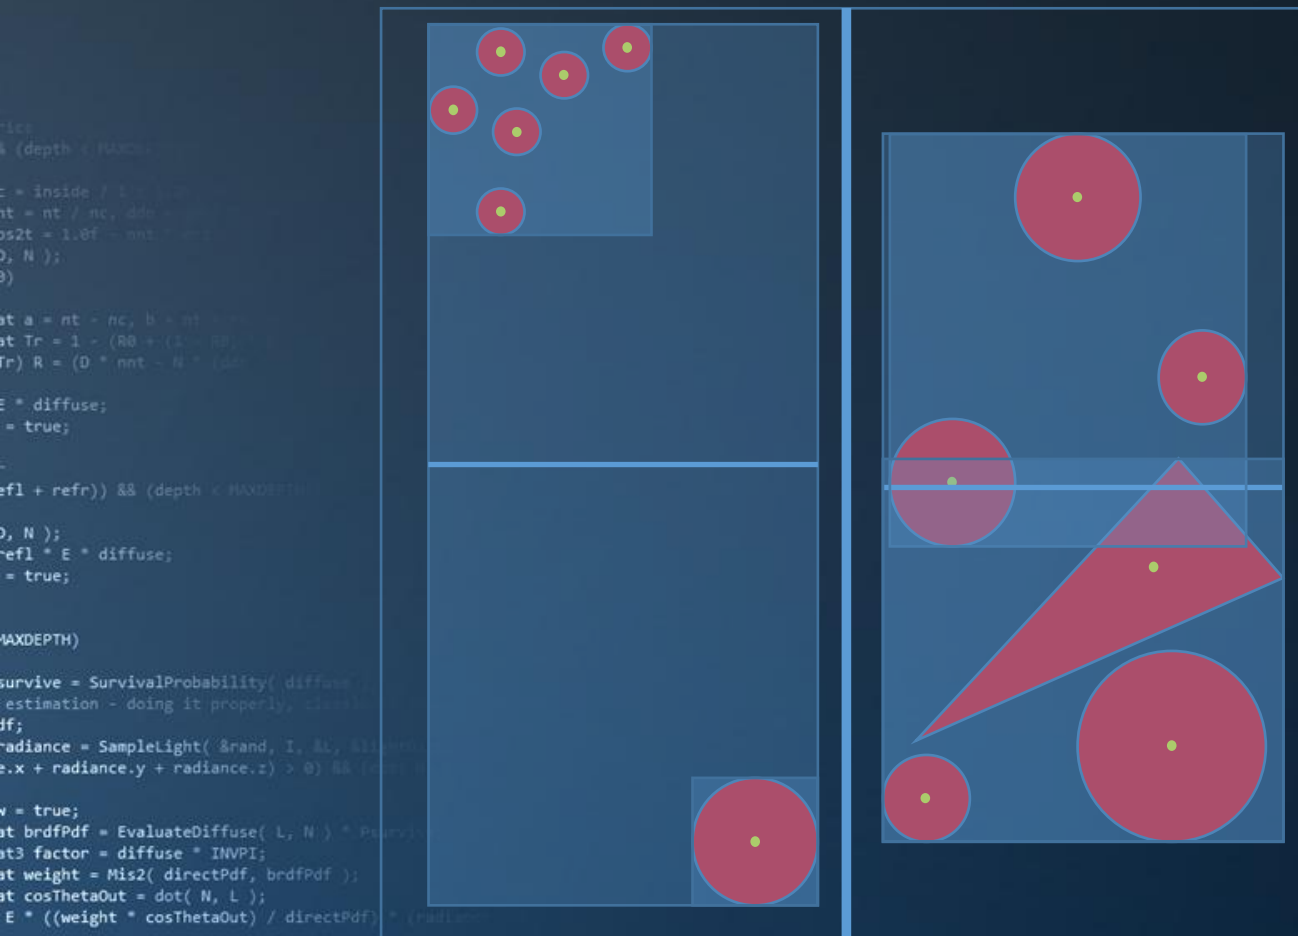


Today's Agenda:

- High-speed Ray Tracing
- Acceleration Structures
- The Bounding Volume Hierarchy
- BVH Construction
- BVH Traversal
- Optimizing Construction
- High-speed Traversal




BVH



Option 5:

Use a bounding volume hierarchy.



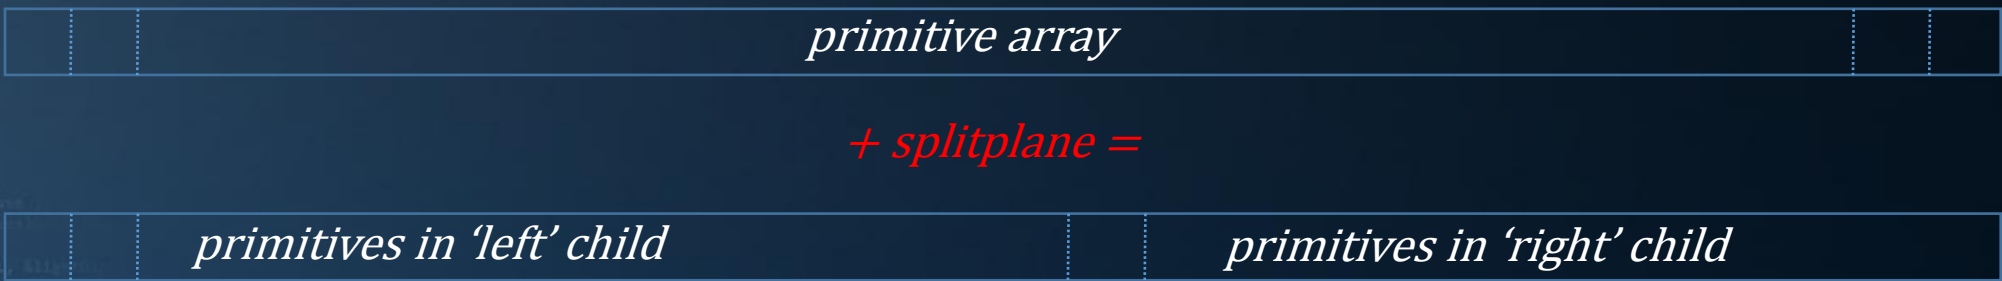
BVH

The Bounding Volume Hierarchy

BSPs, grids, octrees and kD-trees are examples of *spatial subdivisions*.

The BVH is of a different category: it is an *object partitioning scheme*:

Rather than recursively splitting space, it splits collections of objects.

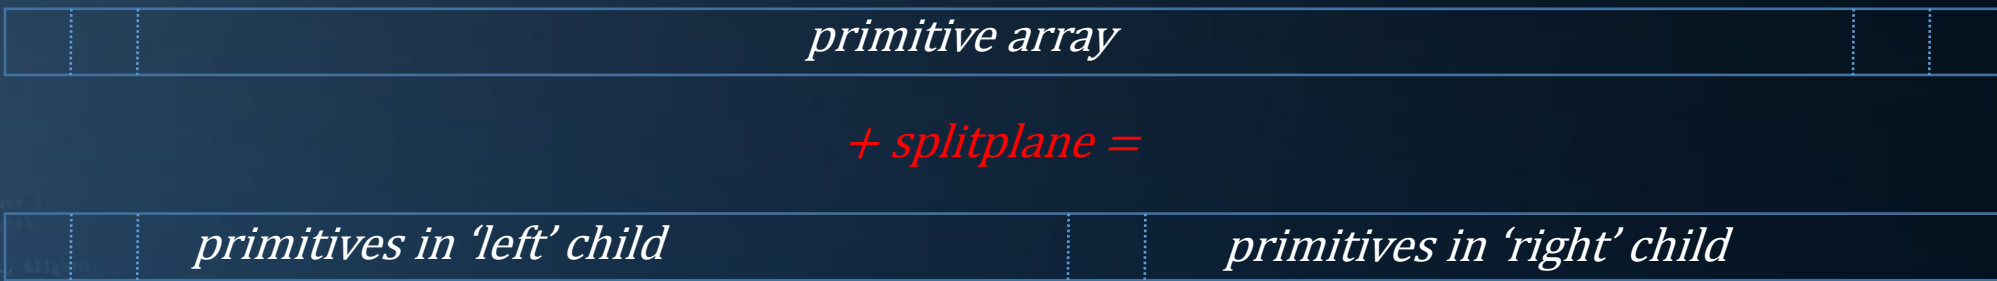


BVH

The Bounding Volume Hierarchy

Sorting an array of elements based on a value:
BVH is very similar to *QuickSort*.

In the BVH construction algorithm, the split plane position is the pivot.



Acceleration Structures

Bounding Volume Hierarchy: data structure

```
struct BVHNode
```

```
{
    BVHNode* left;           // 4 or 8 bytes
    BVHNode* right;          // 4 or 8 bytes
    aabb bounds;             // 2 * 3 * 4 = 24 bytes
    bool isLeaf;              // ?
    vector<Primitive*> primitives; // ?
};
```



Acceleration Structures

Bounding Volume Hierarchy: construction

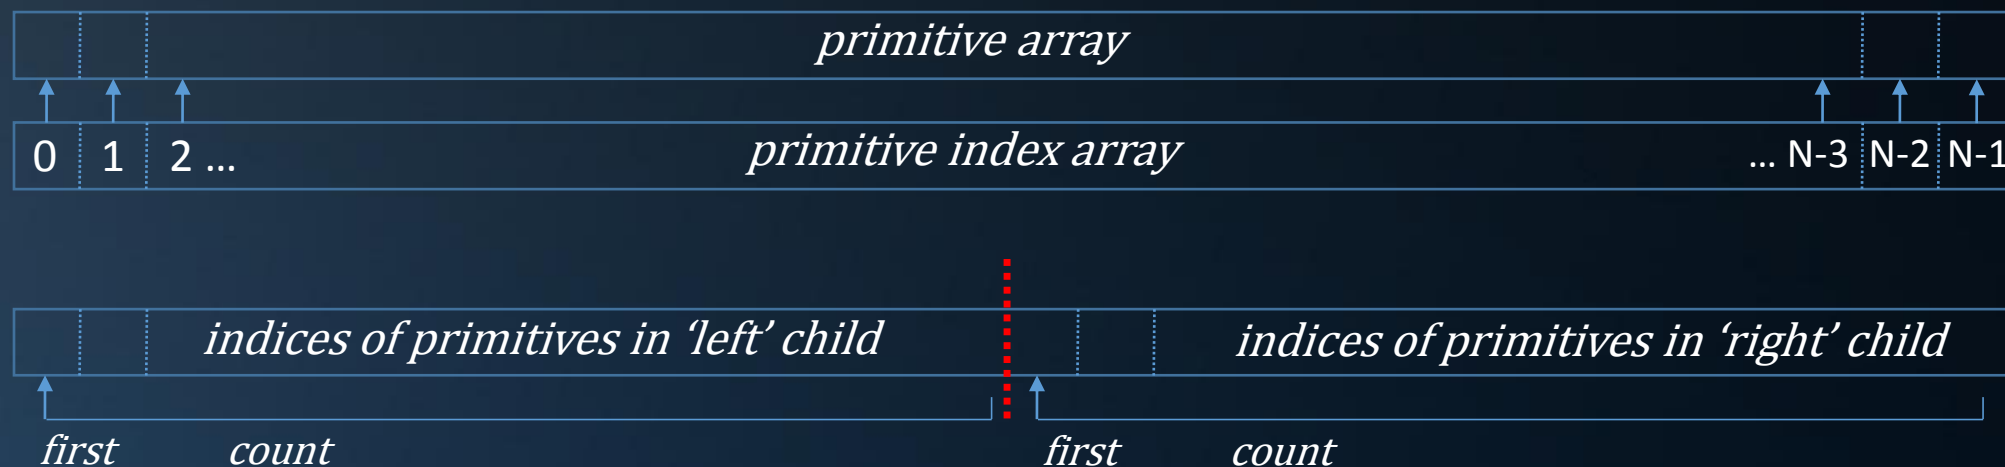
```
void ConstructBVH( Primitive* primitives )
{
    BVHNode* root = new BVHNode();
    root->primitives = primitives;
    root->bounds = CalculateBounds( primitives );
    root->isLeaf = true;
    root->Subdivide();
}
```

```
void BVHNode::Subdivide()
{
    if (primitives.size() < 3) return;
    this->left = new BVHNode(), this->right = new BVHNode();
    ...split 'bounds' in two halves, assign primitives to each half...
    this->left->Subdivide();
    this->right->Subdivide();
    this->isLeaf = false;
}
```



Acceleration Structures

Bounding Volume Hierarchy: construction



Construction consequences:

- Construction happens *in place*:
primitive array is constant,
index array is changed
- Very similar to Quicksort
(split plane = pivot)

Data consequences:

- 'Primitive list' for node becomes
offset + count
- No pointers!
- No pointers?
(what about left / right?)



Acceleration Structures

Bounding Volume Hierarchy: data structure

```
struct BVHNode
{
    BVHNode* left;
    BVHNode* right;
    aabb bounds;
    bool isLeaf;
    vector<Primitive*> primitives;
};
```

```
struct BVHNode
{
    uint left;           // 4 bytes
    uint right;          // 4 bytes
    aabb bounds;         // 24 bytes
    bool isLeaf;         // 4 bytes
    uint first;          // 4 bytes
    uint count;          // 4 bytes
    // -----
    // 44 bytes
};
```



Acceleration Structures

Bounding Volume Hierarchy: data structure

```

// ...
// (depth < MAXDEPTH)
// ...
// inside / ...
// nt = nt / nc; add ...
// ps2t = 1.0f - nnt; ...
// D, N );
// ...
// a = nt - nc; b = nt - nc;
// at Tr = 1 - (R0 + (1 - R0) * ...
// Tr) R = (D * nnt - N * ( ...
// ...
// E * diffuse;
// = true;
// ...
// refl + refr)) && (depth < MAXDEPTH)
// ...
// D, N );
// refl * E * diffuse;
// = true;
// ...
// MAXDEPTH)
// survive = SurvivalProbability( diffuse, ...
// estimation - doing it properly, ...
// if;
// radiance = SampleLight( &rand, I, &t, &light ...
// e.x + radiance.y + radiance.z) > 0) && ( ...
// ...
// v = true;
// at brdfPdf = EvaluateDiffuse( L, N ) * Psurvive;
// at3 factor = diffuse * INVPI;
// at weight = Mix2( directPdf, brdfPdf );
// at cosThetaOut = dot( N, L );
// E * ((weight * cosThetaOut) / directPdf) * (radiance ...
// ...
// random walk - done properly, closely following ...
// survive)
// ...
// at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
// survive;
// pdf;
// n = E * brdf * (dot( N, R ) / pdf);
// ion = true;

```

```

struct BVHNode
{
    union                // 4 bytes
    {
        uint left;
        uint first;
    };
    aabb bounds;        // 24 bytes
    uint count;          // 4 bytes
};                       // -----
                        // 32 bytes

```

```

struct BVHNode
{
    uint left;           // 4 bytes
    uint right;       // 4 bytes
    aabb bounds;        // 24 bytes
    bool isLeaf;        // 4 bytes
    uint first;          // 4 bytes
    uint count;          // 4 bytes
};                       // -----
                        // 44 bytes

```



Acceleration Structures

Bounding Volume Hierarchy: data structure

```

struct BVHNode
{
    float3 bmin;           // bounds: minima
    uint leftFirst;        // or a union
    float3 bmax;           // bounds: maxima
    uint count;            // leaf if 0
};                          // -----
                          // 32 bytes

```




Today's Agenda:

- High-speed Ray Tracing
- Acceleration Structures
- The Bounding Volume Hierarchy
- BVH Construction
- BVH Traversal
- Optimizing Construction
- High-speed Traversal



BVH Traversal



BVH Traversal Algorithm:

Starting with the root:

If the node is a leaf node:

Intersect triangles.

Else:

If the ray intersects the left child AABB:

Traverse left child

If the ray intersects the right child AABB:

Traverse right child

How efficient is this?

In this case: we check every AABB, but we only try to intersect one red sphere.
(total: 8 tests)



BVH Traversal

BVH Efficiency

The number of nodes in a BVH is at most $2N - 1$.
Example:

16	1
8 + 8	2
(4 + 4) + (4 + 4)	4
((2+2) + (2+2)) + ((2+2) + (2+2))	8
1+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1	16

	31

- In this case, we get from the root to a leaf in 5 steps, or: $\log_2 N + 1$.
- For 1024 primitives, we get to a leaf in 11 steps.
- For 1M primitives, we get to a leaf in 21 steps.



Today's Agenda:

- High-speed Ray Tracing
- Acceleration Structures
- The Bounding Volume Hierarchy
- BVH Construction
- BVH Traversal
- Optimizing Construction
- High-speed Traversal



```

    if (depth < MAXDEPTH) {
        nc = inside / 1.0f * RREFR;
        nt = nt / nc; ddn = dot(N, N);
        cos2t = 1.0f - nt * nt; r = sqrt(cos2t);
        D, N );
    }
}

at a = nt + nc, b = nt * r, c = nt * r;
at Tr = 1 - (RB + (1 - RB) * r);
Fr) R = (D * nnt - N * (ddn * Tr + Fr));

E * diffuse;
= true;

-
efl + refr)) && (depth < MAXDEPTH)) {
    D, N );
    refl * E * diffuse;
    = true;

MAXDEPTH)


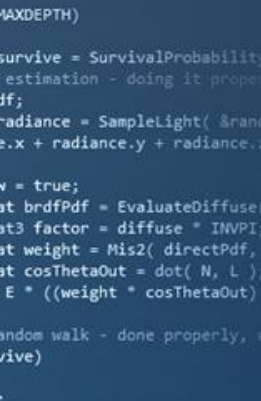
survive = SurvivalProbability(
estimation - doing it properly
diff;
radiance = SampleLight( $randc
e.x + radiance.y + radiance.z;

w = true;
at brdfPdf = EvaluateDiffuse(
at3 factor = diffuse * INVPI;
at3 weight = Mis2( directPdf,
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut)

random walk - done properly,
(survive)

;
at3 brdf = SampleDiffuse( diffuse,
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;

```



What is a good BVH?

- ➔ One that minimizes the number of ray/primitive intersections, and the number of ray/AABB intersections.



Optimizing Construction

BVH Quality

A good BVH minimizes the number of intersections.

Concrete:

$$Q_{bvh} = \sum_1^N P_{AABB} (C_{AABB} + N_{tri} C_{tri})$$

Where:

N is the number of BVH nodes;

P_{AABB} is the probability of a ray hitting the AABB;

C_{AABB} is the cost of a ray intersecting the AABB;

N_{tri} is the number of triangles in the node;

C_{tri} is the cost of intersecting a triangle.


Probability of hitting an AABB
with an arbitrary ray:

*Proportional to the surface
area of the AABB.*




Binned BVH Construction

Surface Area Heuristic (*Or: what is the best way to slice a bunny?*)



Binned BVH Construction



Cost:

$$N_{\text{left}} * A_{\text{left}} + N_{\text{right}} * A_{\text{right}}$$

Select the split with the lowest cost.



Optimizing Construction

Surface Area Heuristic

We construct a BVH by minimizing the cost after each split, i.e. we use the split plane position and orientation that minimizes the cost function:

$$C_{split} = N_{left} A_{left} + N_{right} A_{right}$$

The split is not made at all if the best option is more expensive than *not* splitting, i.e.

$$C_{nosplit} = N A$$

This provides a natural termination criterion for BVH construction.



Efficiency of the Surface Area Heuristic

A BVH constructed with the Surface Area Heuristic is typically *twice as efficient* as a tree constructed with naïve midpoint subdivision.

```

    int nc = inside / l * n; // number of cells
    nt = nt / nc; ddx = 0.0f;
    cos2t = 1.0f - nnt * nnt;
    D, N );
    )

    at a = nt - nc, b = nt + nc;
    at Tr = 1 - (RB + ((1-RB) * ddx));
    Tr) R = (D * nnt - N * (ddx > 0))

    E * diffuse;
    = true;

    .
    refl + refr)) && (depth < MAXDEPTH))

    D, N );
    refl * E * diffuse;
    = true;

    MAXDEPTH)

survive = SurvivalProbability( diffuse, // survival probability
estimation - doing it properly, closely following Monte Carlo);
if;

radiance = SampleLight( &rand, I, &l, &align, &pdf, &N );
e.x + radiance.y + radiance.z) > 0) && (cosThetaOut > 0)

w = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Psurvive;
at3 factor = diffuse * INVPI;
at weight = Mix2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radiance

random walk - done properly, closely following Monte Carlo
(survive)

;
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
ision = true;

```



Today's Agenda:

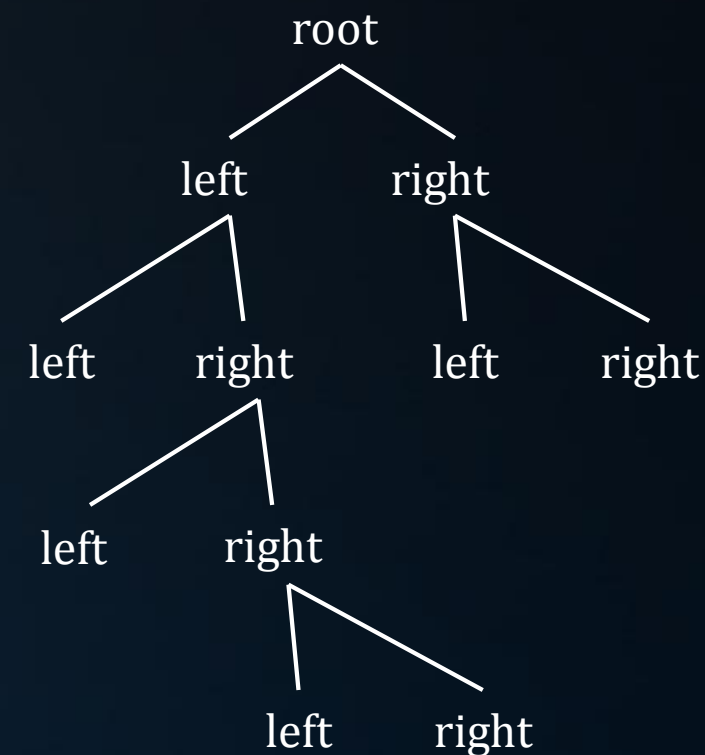

- High-speed Ray Tracing
- Acceleration Structures
- The Bounding Volume Hierarchy
- BVH Construction
- BVH Traversal
- Optimizing Construction
- High-speed Traversal



```

100
    k (depth < MAXDEPTH)
    {
        // Inside of the sphere
        Vec inside = Vec(0,0,0);
        Vec nt = inside / nc, ddx = Vec(0,0,0);
        Vec nnt = nt * nt, ddx2t = 1.0f - nnt * ddx * ddx;
        Vec D, N;
        Vec R;
        Vec a = nt - nc, b = nt + nc;
        Vec Tr = 1 - (RB + (1 - RB) * ddx2t);
        Vec Rr = (D * nnt - N * (ddx * ddx2t));
        Vec E = diffuse;
        Vec refl = true;
        Vec refl + refr)) && (depth < MAXDEPTH)
        {
            Vec D, N;
            Vec refl = E * diffuse;
            Vec refl = true;
        }
    }
    Vec MAXDEPTH)
    {
        Vec survive = SurvivalProbability( diffuse );
        Vec estimation - doing it properly, please
        Vec df;
        Vec radiance = SampleLight( &rand, 1, &g, Vec(0,0,0) );
        Vec r.x + radiance.y + radiance.z) > 0)
        {
            Vec w = true;
            Vec brdfPdf = EvaluateDiffuse( diffuse );
            Vec t3 factor = diffuse * IN;
            Vec weight = Mis2( dir, brdfPdf );
            Vec cosThetaOut = dot( N, dir );
            Vec E = ((weight * cosThetaOut) / d );
            Vec random walk - done properly, only following one
            Vec survive)

```



Ray Packet Traversal


We can exploit this by explicitly traversing *ray packets*.

pdf;

BONUS, NOT EXAM MATERIAL



Fast Traversal



Packet Traversal Algorithm:

Starting with the root:

If the node is a leaf node:

Intersect triangles.

Else:

If **any** ray intersects the left child AABB:

Traverse left child

If **any** ray intersects the right child AABB:

Traverse right child

**BONUS, NOT
EXAM MATERIAL**



Fast Traversal

Ray Packet Traversal

Quickly determining if *any* ray intersects a node:

- Test the first one.

If it intersects, we’re done. Else:

- Test if the AABB is outside the frustum encapsulating the packet.

If it misses, we’re done. Else:

- Brute force test all rays. The first one that hits the AABB will be the ray we check first while processing the child nodes.

**BONUS, NOT
EXAM MATERIAL**



Fast Traversal


Ray Packet Traversal Efficiency

Using the packet traversal approach, we can very efficiently traverse large packets of rays that travel roughly in the same direction. For primary rays, this can be 32x faster than single ray traversal.

Note that this requires the rays in the packet to traverse a similar set of BVH nodes. The ray packet must be *coherent* (as opposed to *divergent*). Ray coherence can be expressed as the extend to which rays in a packet travel the same nodes, or:

$$\text{coherence} = \frac{\text{\#rays in packet}}{\text{average \#rays intersecting a node}}$$

Combined with an efficient BVH, we now have the performance needed for real-time ray tracing.



**BONUS, NOT
EXAM MATERIAL**



Today's Agenda:

- High-speed Ray Tracing
- Acceleration Structures
- The Bounding Volume Hierarchy
- BVH Construction
- BVH Traversal
- Optimizing Construction
- High-speed Traversal



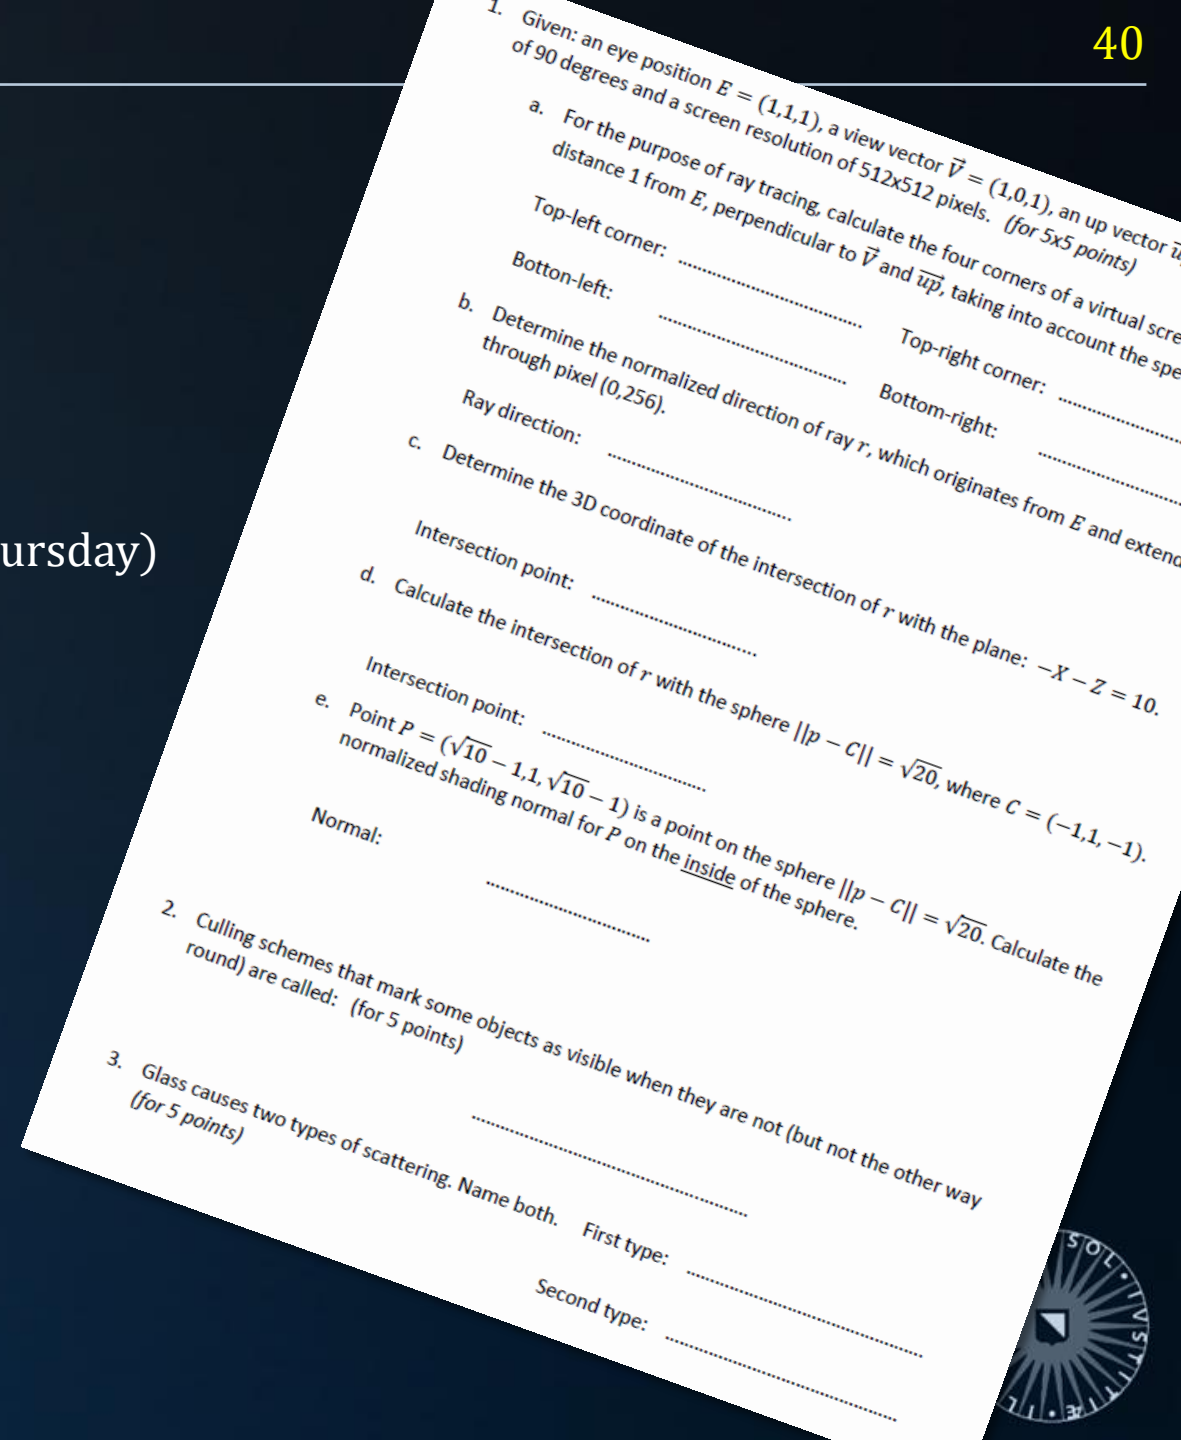
Mid-term Exam

What to study for the exam?

1. Slides (mind cursive terminology!)
2. Example exam (now online, discuss next Thursday)
3. Tutorial sheets

Expectations:

- Fluency with vectors, including dot product, cross product, normalization and all combinations thereof.
- Good understanding of the ray tracing algorithm and light transport.
- Knowledge of terminology used in the lectures.



INFOGR – Computer Graphics

J. Bikker - April-July 2016 - Lecture 7: “Accelerate”

END of “Accelerate”

next up: “Mid-term Exam”

