

```
ics
& (depth < MAXDEPTH)
{
    inside / inside
    nt = nt / nc;
    pos2t = 1.0f - nnt;
    D, N );
    )
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0)
    Tr) R = (D * nnt - N * (a
    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, &align,
    e.x + radiance.y + radiance.z) > 0) && (depth <
    v = true;
    at brdfPdf = EvaluateDiffuse( L, N ) * Survive;
    at3 factor = diffuse * INVPI;
    at weight = Mis2( 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);
    sion = true;
```

/INFOMOV/

Optimization & Vectorization

J. Bikker - Sep-Nov 2016 - Lecture 8: "Data-Oriented Design"

Welcome!



```
...ics
& (depth < MAXDEPTH)
...
c = inside / inside;
nt = nt / nc;
cos2t = 1.0f - nnt;
D, N );
...
at a = nt - nc, b = nt;
at Tr = 1 - (R0 + (1 - R0) * c);
Tr) R = (D * nnt - N * (1 - D));
...
E * diffuse;
= true;
...
efl + refr)) && (depth < MAXDEPTH)
D, N );
-efl * E * diffuse;
= true;
...
MAXDEPTH)
survive = SurvivalProbability( diffuse );
estimation - doing it properly, clearly
if;
radiance = SampleLight( &rand, I, &t, &align,
e.x + radiance.y + radiance.z) && (rand.N < 1);
w = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Survive;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radiance
...
random walk - done properly, closely following
ive)
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf);
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;
```

2016, P2:
Avg. ?

a few remarks on
GRADING P1

2015:
Avg.7.9


2016:
Avg.9.0

2017:
Avg.7.5

Learn about caches

Performance feedback

Reward / distinction



Today's Agenda:

- OOP Performance Pitfalls
- DOD Concepts
- Practical DOD
- DOD or OO?

```
...ics
& (depth < MAXDEPTH)
{
    ...
    inside / 1.0;
    nt = nt / nc;
    ...
    os2t = 1.0f - nnt * nnt;
    ...
    D, N );
}

at a = nt - nc; b = nt;
at Tr = 1 - (R0 + (1 - R0) *
Tr) R = (D * nnt - N * (D0
E * diffuse;
= true;

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

MAXDEPTH)
survive = SurvivalProbability( diffuse );
estimation - doing it properly, classically
if;
radiance = SampleLight( &rand, I, &t, &align,
e.x + radiance.y + radiance.z) > 0) && (depth <
v = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Survive;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radiance
random walk - done properly, closely following
ive)

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



OOP

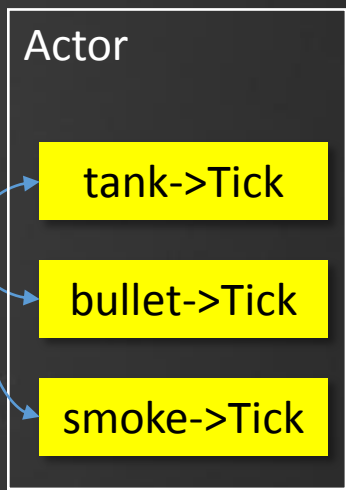
“Death by a Thousand Paper Cuts”

Object Oriented Programming:

- Objects
 - Data
 - Methods
 - Instances

```

...
    & (depth < MAXDEPTH)
...
    inside / inside
    nt = nt / nc;
    cos2t = 1.0f - nnt;
    D, N );
    )
...
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0)
    Tr) R = (D * nnt - N * (a
...
    E * diffuse;
    = true;
...
    refl + refr)) && (depth < MAXDEPTH)
...
    D, N );
    refl * E * diffuse;
    = true;
...
MAXDEPTH)
survive = SurvivalProbability( diffuse,
estimation - doing it properly, clean
if;
radiance = SampleLight(
e.x + radiance.y + radiance
...
v = true;
at brdfPdf = EvaluateDiffuse( L, N ) * radiance;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * radi
...
random walk - done properly, closely following
ive)
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, spot
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;
    
```



Cost of a virtual function call:

1. Virtual Function Table
2. No inlining

Calling such a function:

- cache miss** 1. Read pointer to VFT of base class
- cache miss** 2. Add function offset
- branch** 3. Read function address from VFT
- 4. Load address in PC (jump)

But, that isn't realistic, right?

It is, if we use OO for what it was designed for: operating on heterogeneous objects.



OOP

“Death by a Thousand Paper Cuts”

Characteristics of OO:

- Virtual calls
- Scattered individual objects

```


...
    & (depth < MAXDEPTH)
...
    c = inside / (1 + n);
    nt = nt / nc;
    cos2t = 1.0f - nt;
    D, N );
...
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0) * R);
    (Tr) R = (D * nnt - N * (a *
...
    E * diffuse;
    = true;
...
    refl + refr)) && (depth < MAXDEPTH)
...
    D, N );
    refl * E * diffuse;
    = true;
...
MAXDEPTH)
...
survive = SurvivalProbability( diffuse,
estimation - doing it properly, class
if;
radiance = SampleLight( @rand, I, R, Al, Aligned
e.x + radiance.y + radiance.z) > 0) && (survive
...
w = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Survive;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radiance
...
random walk - done properly, closely following
ive)
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf
urvive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;

```



OOP

“Death by a Thousand Paper Cuts”



Reading memory: 40 cycles @ 300Mhz



Reading memory: 600 cycles @ 3.2Ghz



OOP

“Death by a Thousand Paper Cuts”

Code performance is typically bound by memory access.

“The ideal data is in a format that we can use with the least amount of effort.”

➔ Effort = CPU-effort.

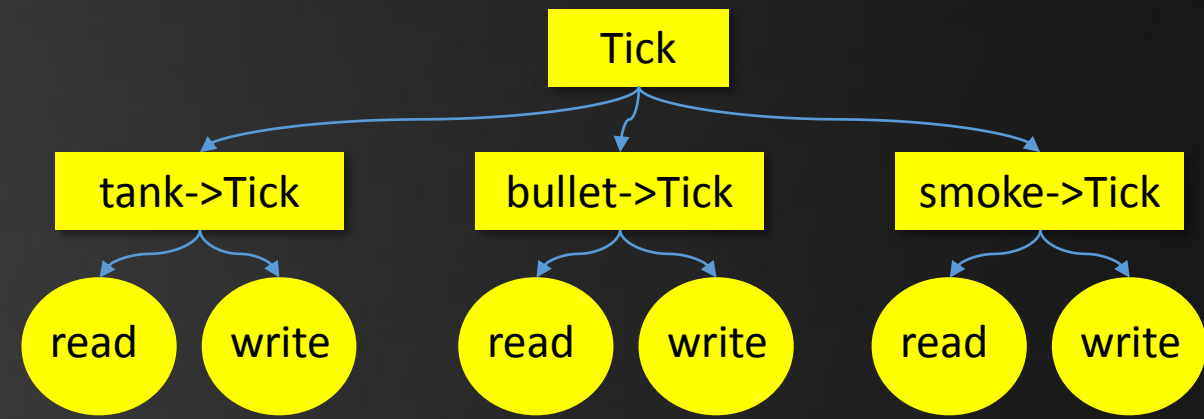
“You cannot be fast without knowing how data is touched.”



OOP

“Death by a Thousand Paper Cuts”

Parallel processing typically requires synchronization.




“You cannot multi-thread without knowing how data is touched.”



OOP

“Death by a Thousand Paper Cuts”

Parallel processing requires coherent program flow.




```

...
(depth + MAXDEPTH)
...
inside / ...
nt = nt / nc;
os2t = 1.0f - nnt;
D, N );
...
a = nt - nc; b = nt;
Tr = 1 - (R0 + ...);
Tr) R = (D * nnt;
...
E * diffuse;
= true;
...
refl + refr)) && (depth
D, N );
refl * E * diffuse;
= true;
...
MAXDEPTH)
survive = SurvivalProbability( diffuse;
estimation - doing it properly, clean
if;
radiance = SampleLight( @rand, I, R);
e.x + radiance.y + radiance.z) > 0) && (refl
...
v = true;
at brdfPdf = EvaluateDiffuse( L, N ) * P;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radiance
...
random walk - done properly, closely following
ive)
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;

```

“You cannot multi-thread without knowing how data is touched.”



OOP

“Death by a Thousand Paper Cuts”

```

class Bot : public Enemy
{
    ...
    vec3 m_position;
    ...
    float m_mod;
    ...
    float m_aimDirection;
    ...
    virtual void updateAim( vec3 target )
    {
        m_aimDirection = dot3( m_position, target ) * m_mod;
    }
}

```

cached but not used

cached but not used

branch

cache miss

cache miss

cache miss

cache miss

cache miss



OOP

“Death by a Thousand Paper Cuts”

```
void updateAims(
    float* aimDir,
    const AimingData* aim,
    vec3 target,
    uint count
```

only reads data that is actually needed to cache

```
for (uint i = 0; i < count; ++i)
{
    aimDir[i] = dot3(aim->positions[i],target) * aim->mod[i];
```

writes to linear array

actual functionality is unchanged



DOD

Data Oriented Design*

Origin: low-level game development.

Core idea: *focus software design on CPU- and cache-aware data layout.*

Take into account:

- Cache line size
- Data alignment
- Data size
- Access patterns
- Data transformations

Strive for a simple, linear access pattern as much as possible.

*: Nikos Drakos, “Data Oriented Design”, 2008. <http://www.dataorienteddesign.com/dodmain>



DOD

Bad Access Patterns: Linked List

The Perfect LinkedList™:

```

struct LLNode
{
    LLNode* next;
    int value;
};
    
```

```

nodes = new LLNode[...];
    
```

```

for( int i = 0; i < ...; i++ )
    nodes[i].next = &nodes[i + 1];
    
```



```

LLNode* NewNode( int value )
{
    LLNode* retval = pool;
    pool = pool->next;
    retval->value = value;
    return retval;
}
    
```

```

list = NewNode( -10000 );
list->next = NewNode( 10000 );
list->next->next = 0;
    
```



DOD

Bad Access Patterns: Linked List

KISS Array™:

```
data = new int[...];
memset( data, 0, ... * sizeof( int ) );
data[0] = -10000;
data[1] = 10000;
N = 2;
```

```
for( int i = 0; i < COUNT; i++ )
{
    int pos = 1, value = rand() & 8191;
    while (data[pos] < value) pos++;
    memcpy( data + pos + 1,
           data + pos,
           (N - pos + 1) * sizeof( int ) );
    data[pos] = value, N++;
}
```



DOD



```

...
    & (depth < MAXDEPTH)
...
    int i = 0; i < COUNT; i++ )
    {
        LLNode* node = NewNode( rand() & 8191);
        LLNode* iter = list;
        while (iter->next->value < node->value)
            iter = iter->next;
        node->next = iter->next;
        iter->next = node;
    }
...
    survive = SurvivalProbability( diffuse, ...
    estimation - doing it properly, clearly
    if;
    radiance = SampleLight( @rand, I, R, ...
    e.x + radiance.y + radiance.z) > 0) && (rand()
...
    v = true;
    brdfPdf = EvaluateDiffuse( L, N ) * Recursive
    factor = diffuse * INVPI;
    weight = Mis2( directPdf, brdfPdf );
    cosThetaOut = dot( N, L );
    E * ((weight * cosThetaOut) / directPdf) * (radiance
...
    random walk - done properly, closely following
    (vive)
...
    brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;

```

```

for( int i = 0; i < COUNT; i++ )
{
    int pos = 1, value = rand() & 8191;
    while (data[pos] < value) pos++;
    memcpy( data + pos + 1, data + pos,
            (N - pos + 1) * sizeof( int ) );
    data[pos] = value, N++;
}

```



DOD

Bad Access Patterns: Linked List*

Inserting elements in an array by shifting the remainder of the array is *significantly faster* than using an optimized linked list.

Why?

- Finding the location in the array: pure linear access
 - Shifting the remainder: pure linear access.
- ➔ Even though the amount of transferred memory is huge, this approach wins.



*: Also see: Nathan Reed, Data Oriented Hash Table, 2015.

<http://www.reedbeta.com/blog/data-oriented-hash-table>



DOD

Bad Access Patterns: Octree




DOD

Bad Access Patterns: Octree

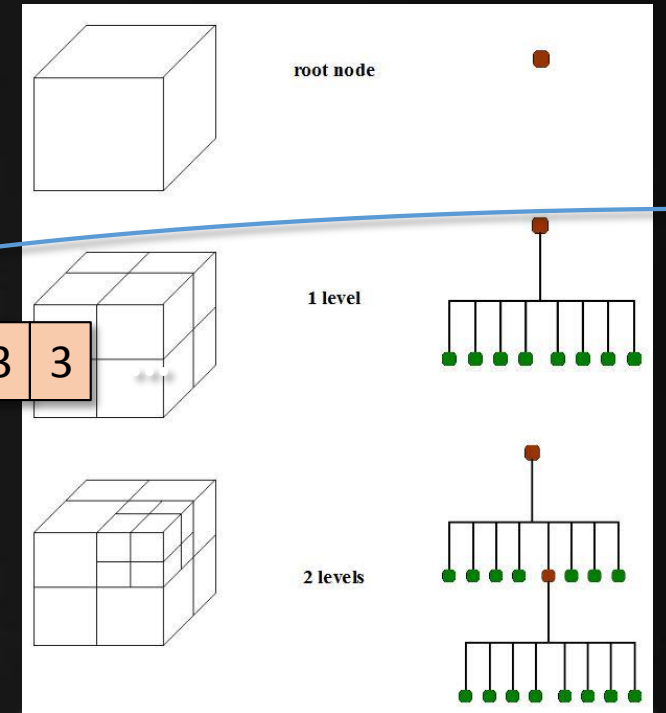
Query: find the color of a voxel visible through pixel (x,y).

Operation: ‘3DDDA’ (basically: Bresenham).

Data layout:



Color data: 32-bit (ARGB).



```

...
(depth < MAXDEPTH)
{
  ...
  inside / ...
  nt = nt / nc;
  pos2t = 1.0f - nnt;
  D, N );
}

at a = nt - nc; b = nt;
at Tr = 1 - (R0 + (1 - R0) * pos2t);
Tr) R = (D * nnt - N * (1 - pos2t));

E * diffuse;
= true;

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

MAXDEPTH)
survive = SurvivalProbability( diffuse );
estimation - doing it properly, close
if;
radiance = SampleLight( @rand, I, R, Align
e.x + radiance.y + radiance.z) > 0) && (rand
v = 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) * (radiance
random walk - done properly, closely following
ive)

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


```



DOD

Bad Access Patterns: Octree

Alternative layout (part 2):



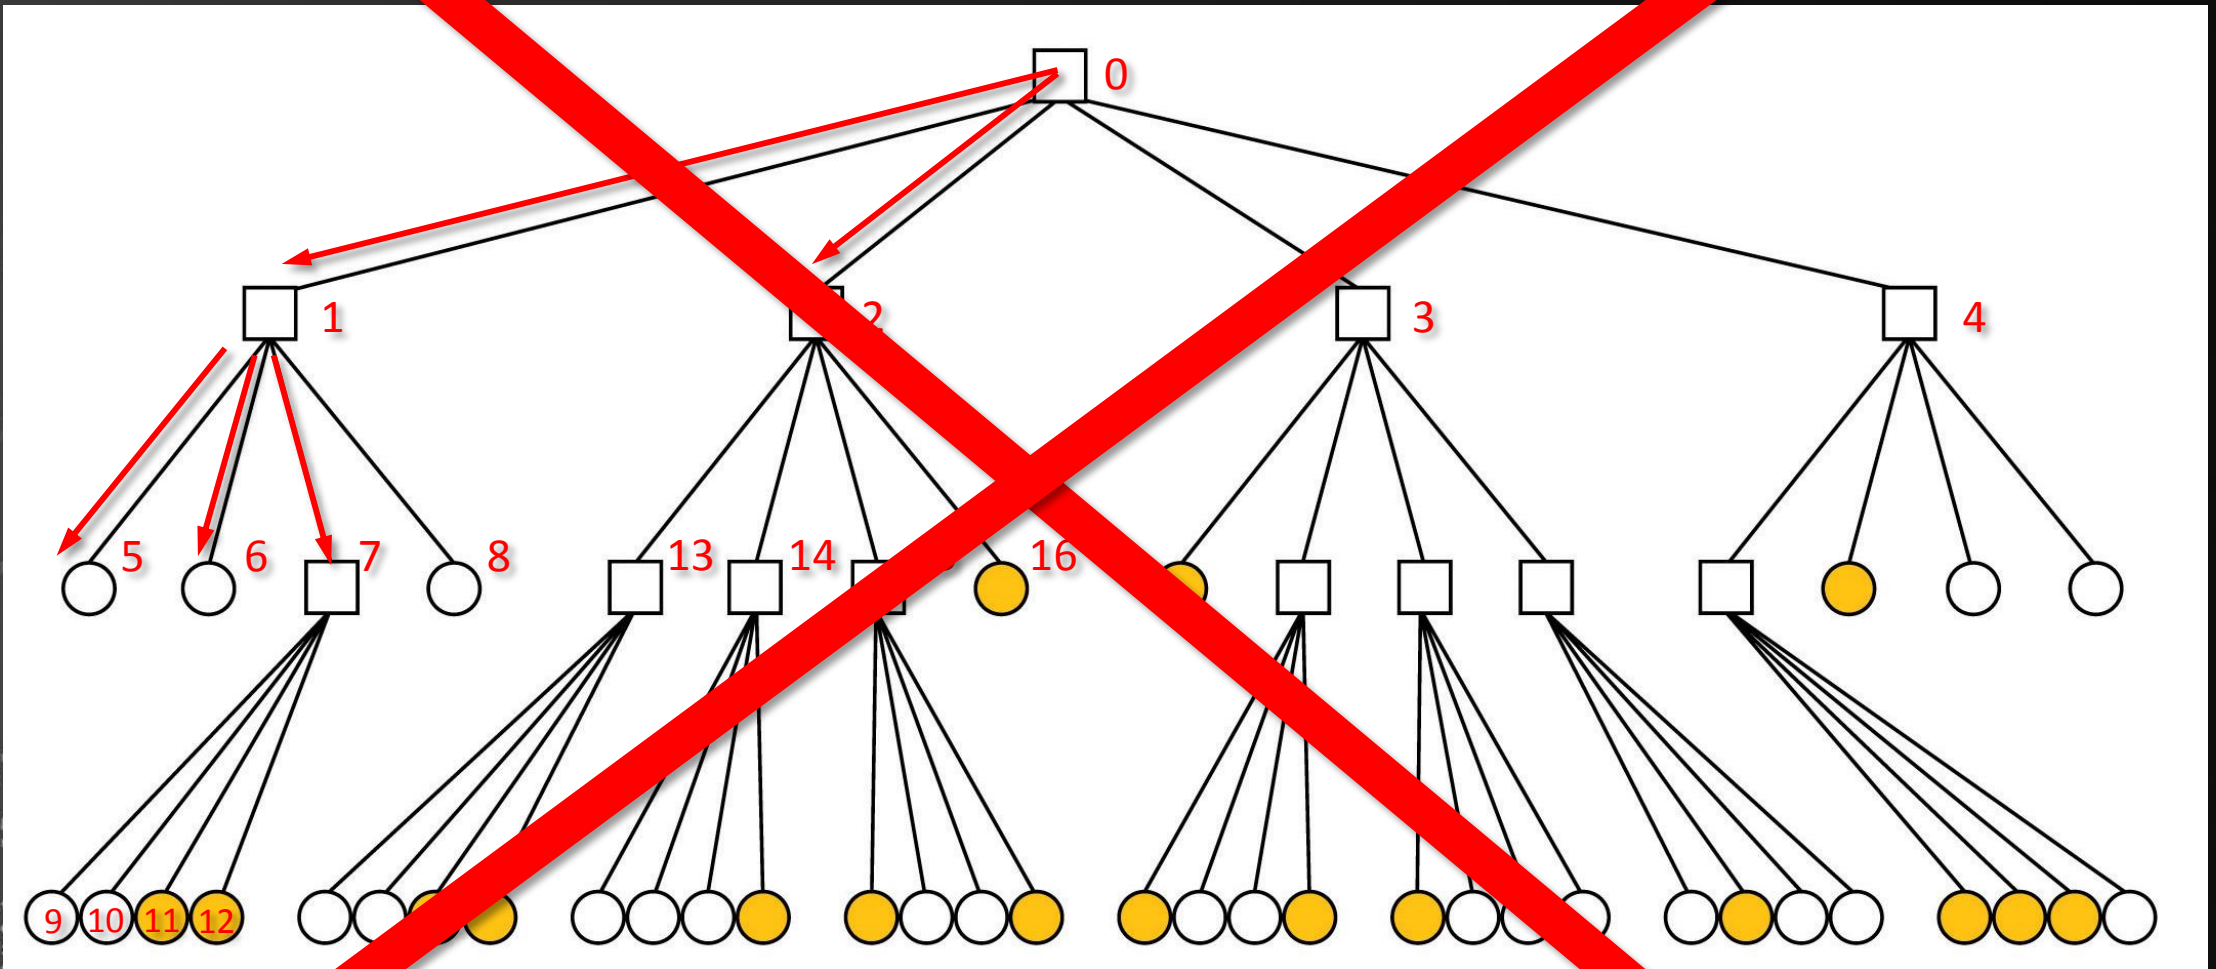
Trees are typically generated by a divide-and-conquer algorithm, in a depth-first fashion.

Compact storage:

```
struct OTNode
{
    int firstChild;
    // bit 31: empty
    // bit 30: leaf
};
```



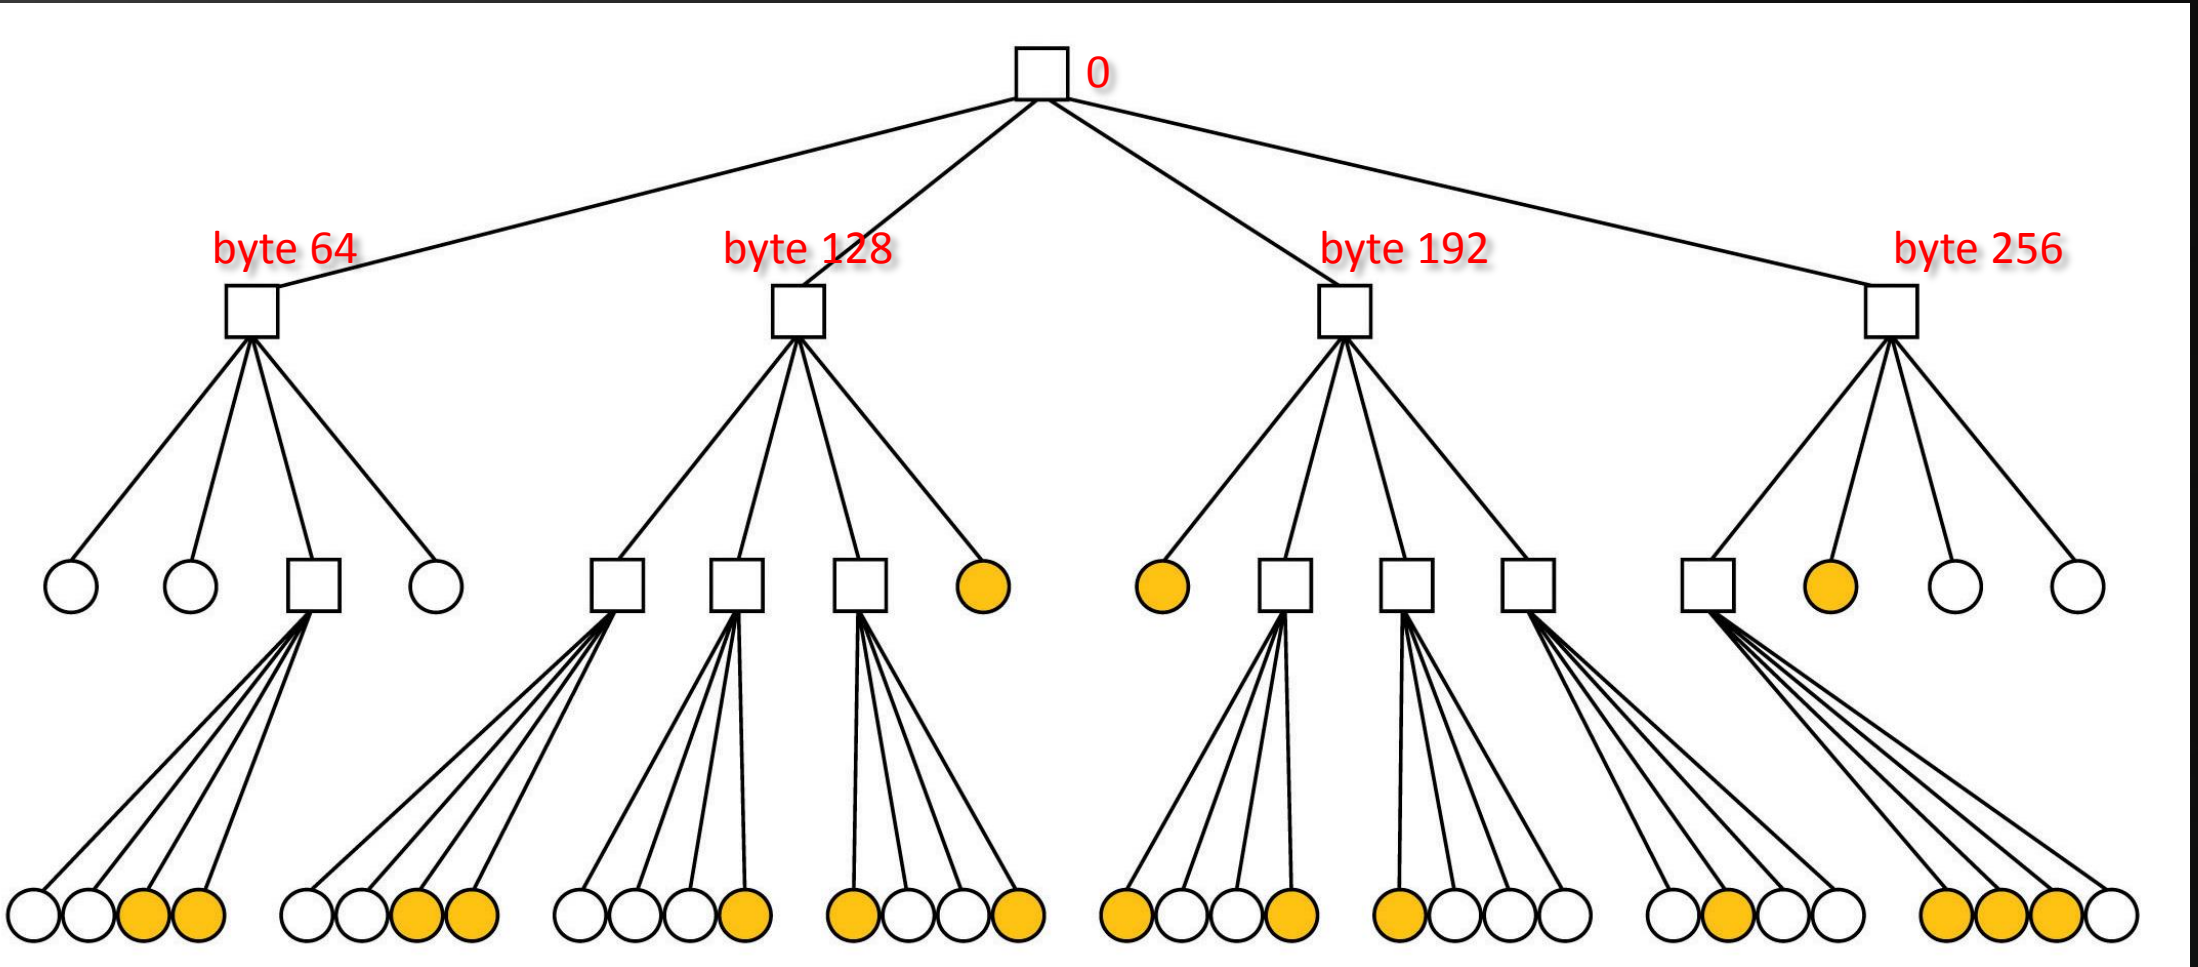
DOD



```
... (depth + MAXDEPTH) ...  
... inside ...  
... nt = nt / nc ...  
... os2t = 1.0f - nt ...  
... D, N ... ) ...  
... ) ...  
... at a = nt - nc, b ...  
... at Tr = 1 - (R0 + ...  
... Tr) R = (D * nnt ...  
... E * diffuse; ...  
... = true; ...  
... efl + refr) && (de ...  
... D, N ... ); ...  
... efl * E * diffuse; ...  
... = true; ...  
... MAXDEPTH) ...  
... survive = SurvivalP ...  
... estimation - doing ...  
... if; ...  
... radiance = SampleLi ...  
... e.x + radiance.y + ...  
... w = true; ...  
... at brdfPdf = Evalua ...  
... at3 factor = diffus ...  
... at weight = Mis2( d ...  
... at cosThetaOut = do ...  
... E * ((weight * cosThetaOut) / directPdf) ...  
... random walk - done properly, close to ...  
... (ive) ...  
... at3 brdf = SampleDiffuse( diffuse, N, r1, r2, R, sp ...  
... survive; ...  
... pdf; ...  
... n = E * brdf * (dot( N, R ) / pdf); ...  
... sion = true; ...
```



DOD



DOD

Bad Access Patterns: Octree

Alternative layout (part 2):

- Reorganize so that treelets are cacheline-aligned.

(you will waste some memory)

```

...
    & (depth < MAXDEPTH)
...
    t = inside / (1 + 3 * depth);
    nt = nt / nc;
    cos2t = 1.0f - nt;
    D, N );
)
...
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0) * a);
    Tr) R = (D * nnt - N * (1 - a));
...
    E * diffuse;
    = true;
...
    refl + refr)) && (depth < MAXDEPTH)
...
    D, N );
    -refl * E * diffuse;
    = true;
...
MAXDEPTH)
...
survive = SurvivalProbability( diffuse );
estimation - doing it properly, clearly
if;
radiance = SampleLight( &rand, I, &t, &light );
e.x + radiance.y + radiance.z) > 0) && (survive)
...
    v = true;
    at brdfPdf = EvaluateDiffuse( L, N ) * Survive;
    at3 factor = diffuse * INVPI;
    at weight = Mis2( directPdf, brdfPdf );
    at cosThetaOut = dot( N, L );
    E * ((weight * cosThetaOut) / directPdf) * (radiance
...
random walk - done properly, closely following
ive)
...
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;

```



DOD

Bad Access Patterns: Textures in a Ray Tracer

Typical process for tracing a ray:


- Traverse a tree (multiple kilobytes)
- Intersect triangles in the leaf nodes (quite a few bytes)
- If a hit is found, fetch texture.

This is almost always a cache miss.

```

...
    & (depth < MAXDEPTH)
...
    t = inside / (1.0f - outside);
    nt = nt / nc;
    cos2t = 1.0f - nt;
    D, N );
...
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0));
    Tr) R = (D * nnt - N * (a
...
    E * diffuse;
    = true;
...
    refl + refr)) && (depth < MAXDEPTH)
...
    D, N );
    refl * E * diffuse;
    = true;
...
MAXDEPTH)
...
    survive = SurvivalProbability( diffuse,
    estimation - doing it properly, close
    if;
    radiance = SampleLight( &rand, I, &t, &ll
    e.x + radiance.y + radiance.z) > 0) && (o
...
    w = true;
    at brdfPdf = EvaluateDiffuse( L, N ) * Pdf
    at3 factor = diffuse * INVPI;
    at weight = Mis2( directPdf, brdfPdf );
    at cosThetaOut = dot( N, L );
    E * ((weight * cosThetaOut) / directPdf)
...
    random walk - done properly, closely follow
    ve)
...
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;

```



DOD

Bad Access Patterns: Textures in a Ray Tracer



We suffer the cache miss *twice*:

- Once for the texture;
- Once for the normal map.

Note: both values are 32-bit.

```

...
    & (depth < MAXDEPTH)
...
    inside / inside
    nt = nt / nc;
    cos2t = 1.0f - nnt;
    D, N );
    )
...
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0));
    Tr) R = (D * nnt - N * (a
...
    E * diffuse;
    = true;
...
    refl + refr)) && (depth < MAXDEPTH)
...
    D, N );
    refl * E * diffuse;
    = true;
...
MAXDEPTH)
...
    survive = SurvivalProbability( diffuse,
    estimation - doing it properly, close
    if;
    radiance = SampleLight( &rand, I, M, All
    e.x + radiance.y + radiance.z) > 0) && (o
...
    w = true;
    at brdfPdf = EvaluateDiffuse( L, N ) * Pdf
    at3 factor = diffuse * INVPI;
    at weight = Mis2( directPdf, brdfPdf );
    at cosThetaOut = dot( N, L );
    E * ((weight * cosThetaOut) / directPdf)
...
    random walk - done properly, closely follow
    ve)
...
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, M
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;
    
```



DOD

Bad Access Patterns: Textures in a Ray Tracer

Interleaved texture / normal:

- One value now becomes 64-bit and contains the normal and the color.
- We still suffer a cache miss –
- But only once.

```

...
    & (depth < MAXDEPTH)
...
    t = inside / (1 + 0.5 * cos2t);
    nt = nt / nc; nnt = nnt * nt;
    cos2t = 1.0f - nnt;
    D, N );
    )
...
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0) * t);
    Tr) R = (D * nnt - N * (a *
...
    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) && (survive)
...
v = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Survive;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radiance
...
random walk - done properly, closely following
ive)
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;

```



DOD

Previously in INFOMOV

```

...ics
& (depth < MAXDEPTH)
...
c = inside / (1 + 1000 * ...
nt = nt / nc; dde = ...
ps2t = 1.0f - nnt * ...
D, N );
...
)
...
at a = nt - nc, b = nt - ...
at Tr = 1 - (R0 + (1 - R0) ...
Tr) R = (D * nnt - N * ...
...
E * diffuse;
= true;
...
efl + refr)) && (depth < MAXDEPTH)
...
D, N );
-efl * E * diffuse;
= true;
...
MAXDEPTH)
...
survive = SurvivalProbability( diffuse, ...
estimation - doing it properly, clearly ...
if;
...
radiance = SampleLight( &rand, I, &t, &align, ...
e.x + radiance.y + radiance.z) > 0) && (rand ...
...
v = true;
...
at brdfPdf = EvaluateDiffuse( L, N ) * ...
at3 factor = diffuse * INVPI;
...
at weight = Mis2( directPdf, brdfPdf );
...
at cosThetaOut = dot( N, L );
...
E * ((weight * cosThetaOut) / directPdf) * (radiance ...
...
random walk - done properly, closely following ...
ive)
...
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf ...
survive;
...
pdf;
...
n = E * brdf * (dot( N, R ) / pdf);
...
sion = true;

```

```

struct Particle
{
    float x, y, z;
    float vx, vy, vz;
    float mass;
};
// size: 28 bytes

```

Better:

```

struct Particle
{
    float x, y, z;
    float vx, vy, vz;
    float mass, dummy;
};
// size: 32 bytes

```



DOD

Previously in INFOMOV

```

struct Particle
{
    float x, y, z;
    int mass;
};
Particle particle[512];
    
```

AOS

```

union { float x[512]; __m128 x4[128]; };
union { float y[512]; __m128 y4[128]; };
union { float z[512]; __m128 z4[128]; };
union { int mass[512]; __m128i mass4[128]; };
    
```

SOA

structure of arrays




DOD

Previously in INFOMOV

```

...ics
& (depth < MAXDEPTH)
...
c = inside / 1.5;
nt = nt / nc;
cos2t = 1.0f - nt;
D, N );
)
...
at a = nt - nc, b = nt;
at Tr = 1 - (R0 + (1 - R0));
Tr) R = (D * nnt - N * (a
...
E * diffuse;
= true;
...
efl + refr)) && (depth < MAXDEPTH)
...
D, N );
-refl * E * diffuse;
= true;
...
MAXDEPTH)
...
survive = SurvivalProbability( diffuse,
estimation - doing it properly, class
if;
-radiance = SampleLight( @rand, I, @t, @lig
e.x + radiance.y + radiance.z) > 0) && (re
...
v = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Pearc
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radi
...
random walk - done properly, closely following
ive)
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, @R, @pdf
urvive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;

```



Method:

X = 1 1 0 0 0 1 0 1 1 0 1 1 0 1

Y = 1 0 1 1 0 1 1 0 1 0 1 1 1 0

M = 1101101000111001110011111001

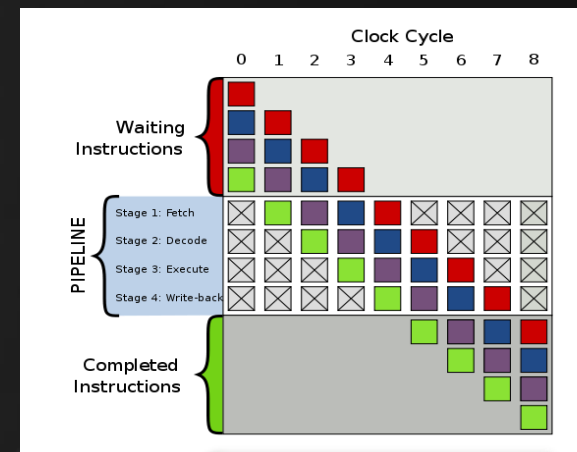




DOD

Algorithm Performance Factors

Estimating algorithm cost:

1. Algorithmic Complexity : $O(N)$, $O(N^2)$, $O(N \log N)$, ...
2. Cyclomatic Complexity* (or: Conditional Complexity)
3. Amdahl's Law / Work-Span Model
4. Cache Effectiveness



*: McCabe, A Complexity Measure, 1976.



Today's Agenda:

- OOP Performance Pitfalls
- DOD Concepts
- Practical DOD
- DOD or OO?



P2

Data-Oriented Optimization: P2

How many tanks fit in a grid cell?

How many tanks do we want to fit in a grid cell? 😊

How many cache lines is a grid cell?

How far away are the neighbors?

Can we fit more tanks in a grid cell without using more memory?

How can we efficiently point to a tank?

...

Profiler: does this stuff matter?

```

...
    & (depth < MAXDEPTH)
...
    inside / 1.0f;
    nt = nt / nc;
    cos2t = 1.0f - nt;
    D, N );
    )
...
    at a = nt - nc, b = nt;
    at Tr = 1 - (R0 + (1 - R0) *
    Tr) R = (D * nnt - N * (a
...
    E * diffuse;
    = true;
...
    refl + refr)) && (depth < MAXDEPTH)
...
    D, N );
    refl * E * diffuse;
    = true;
...
MAXDEPTH)
...
survive = SurvivalProbability( diffuse,
estimation - doing it properly, close
if;
radiance = SampleLight( &rand, I, &t, &light
e.x + radiance.y + radiance.z) > 0) && (survive
...
    v = true;
    at brdfPdf = EvaluateDiffuse( L, N ) * Pa
    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
ive)
...
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;

```



Today's Agenda:

- OOP Performance Pitfalls
- DOD Concepts
- Practical DOD
- DOD or OO?

```
...ics
& (depth < MAXDEPTH)
{
    ...
    inside / 1.0;
    nt = nt / nc;
    ...
    os2t = 1.0f - nnt * nnt;
    ...
    D, N );
}

...
at a = nt - nc; b = nt;
at Tr = 1 - (R0 + (1 - R0) *
Tr) R = (D * nnt - N * (D0
...
E * diffuse;
= true;
...
efl + refr)) && (depth < MAXDEPTH)
D, N );
-efl * E * diffuse;
= true;
...
MAXDEPTH)
survive = SurvivalProbability( diffuse );
estimation - doing it properly, classically
if;
radiance = SampleLight( &rand, I, &t, &align,
e.x + radiance.y + radiance.z) > 0) && (depth <
v = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Survive;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf) * (radiance
...
random walk - done properly, closely following
ive)
...
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf
urvive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;
```



