

# /INFOMOV/

## Optimization & Vectorization

J. Bikker - Sep-Nov 2019 - Lecture 7: "Data-Oriented Design"

# Welcome!

```
rics  
 3 (depth < MAXDEPTH)  
  n = inside ? 1 : 1.2f;  
  nt = nt / nc, ddn = ddn / nc;  
  pos2t = 1.0f - nnt - nnr;  
  D, N );  
)  
  
at a = nt - nc, b = nt + nc;  
at Tr = 1 - (R0 + (1 - R0) *  
Tr) R = (D * nnt - N * (ddn -  
E * diffuse;  
= true;  
  
-  
refl + refr)) && (depth < MAXDEPTH);  
D, N );  
refl * E * diffuse;  
= true;  
  
MAXDEPTH)  
  
survive = SurvivalProbability( diffuse );  
estimation - doing it properly, closer  
if;  
radiance = SampleLight( &rand, I, &L, &lightDir );  
e.x + radiance.y + radiance.z ) > 0) && (dot( N,  
v = true;  
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;  
at3 factor = diffuse * INVPi;  
at weight = Mis2( directPpdf, brdfPpdf );  
at cosThetaOut = dot( N, L );  
E * ((weight * cosThetaOut) / directPpdf) * (radiance  
random walk - done properly, closely following Smiley  
ive);  
;  
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );  
urvive;  
pdf;  
n = E * brdf * (dot( N, R ) / pdf);  
ision = true;
```



# Today's Agenda:

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



# Fact Checking

“Floating point code is (typically) undeterministic”

```

rics
  & (depth < MAXDEPTH)
  c = inside ? 1 : 1.2f;
  nt = nt / nc; ddn = ddn / nc;
  pos2t = 1.0f - nnt * nnt;
  D, N );
}

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

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

MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
if;
radiance = SampleLight( &rand, I, &L, &lighting,
e.x + radiance.y + radiance.z ) > 0) && (dot( N,
v = true;
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPpdf, brdfPpdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPpdf) * (radiance -
random walk - done properly, closely following Sampling
alive);
;

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

```

<b>float</b> v0 = 1; <b>float</b> v1 = 1; <b>float</b> v2 = 1; <b>float</b> v3 = 1; <b>float</b> v4 = 1; <b>float</b> v5 = 1; <b>float</b> v6 = 1; <b>float</b> v7 = 1; <b>for</b> ( <b>int</b> i = 0; i < 2000000; i++) { v0 *= 1.00001f; v1 *= 1.00001f; v2 *= 1.00001f; v3 *= 1.00001f; v4 *= 1.00001f; v5 *= 1.00001f; v6 *= 1.00001f; v7 *= 1.00001f;	<b>fld1</b> <b>fld</b> st(0) <b>fld</b> st(1) <b>fld</b> st(2) <b>fld</b> st(3) <b>fld</b> st(4) <b>fld</b> st(5) <b>fld</b> st(6) <b>fmul</b> st(7),st ; <b>fxch</b> st(7) ; <b>fstp</b> [v0] <b>fxch</b> st(5) ; <b>fmul</b> st,st(6) <b>fxch</b> st(4) ; <b>fmul</b> st,st(6) <b>fxch</b> st(3) ; <b>fmul</b> st,st(6) <b>fxch</b> st(2) ; <b>fmul</b> st,st(6) <b>fxch</b> st(1) ; <b>fmul</b> st,st(6) <b>fxch</b> st(5) ; <b>fmul</b> st,st(6) <b>fld</b> [v7] <b>fmul</b> st,st(7) <b>fstp</b> [v7]
---	---



# Fact Checking

“Doubles are slower than floats (4x)”

This statement is **mostly true**. The real story, CPU (win32, x64):

- A **float** takes 32-bit in memory, but gets promoted to 80 bits in an FPU register.
- A **double** takes 64-bit in memory, but gets promoted to 80 bits in an FPU register.
- A **long double** takes 64-bit in memory, but gets promoted to 80 bits in an FPU register.

Calculation time on 80-bit FPU registers does not depend on the source of the data.  
HOWEVER: the fp registers are rarely used anymore...

The real story, GPU (Nvidia, AMD): <https://www.geeks3d.com/20140305/amd-radeon-and-nvidia-geforce-fp32-fp64-gflops-table-computing>

- Titan V: FP64 = 1/2 \* FP32 (6900 vs 13800 GFLOPS)
- Titan X Pascal: FP64 = 1/32 \* FP32 (350 vs 11300 GFLOPS) (same for all 10xx)
- Radeon RX Vega 64: FP64 = 1/16 \* FP32 (790 vs 12700 GFLOPS)
- Radeon HD 7990: FP64 = 1/4 \* FP32 (1946 vs 7782)

FP16 (GPU only): <https://www.anandtech.com/show/10325/the-nvidia-geforce-gtx-1080-and-1070-founders-edition-review/5>

- GTX 1080Ti: FP16 = 1/64 \* FP32 (ouch)
- Radeon RX Vega 64: FP16 = 2 \* FP32 (!)



# Today's Agenda:

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



# OOP

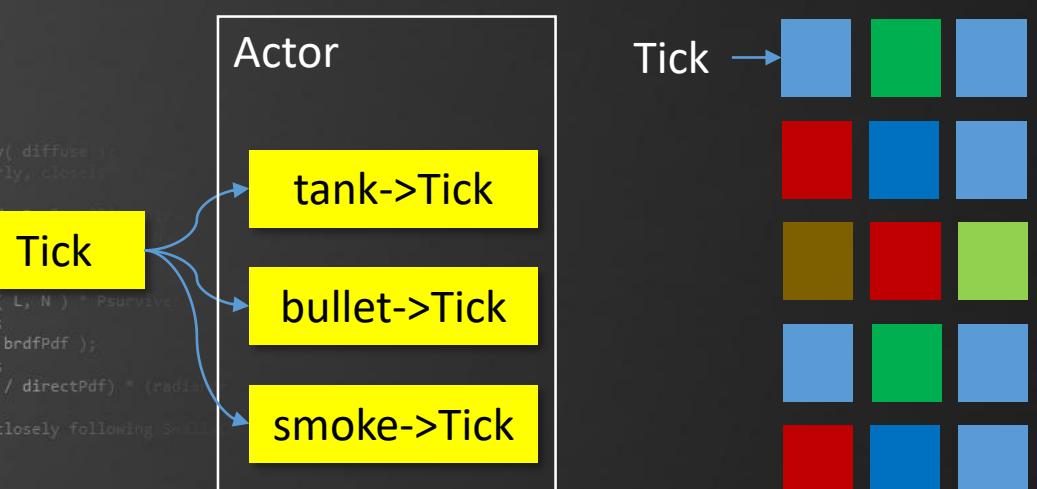
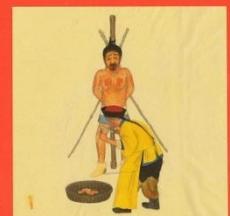
## “Death by a Thousand Cuts”

Object Oriented Programming:

- Objects
  - Data
  - Methods
  - Instances



D E A T H   B Y  
A T H O U S A N D  
C U T S



Timothy Brook  
Jérôme Bourgon  
Gregory Blue

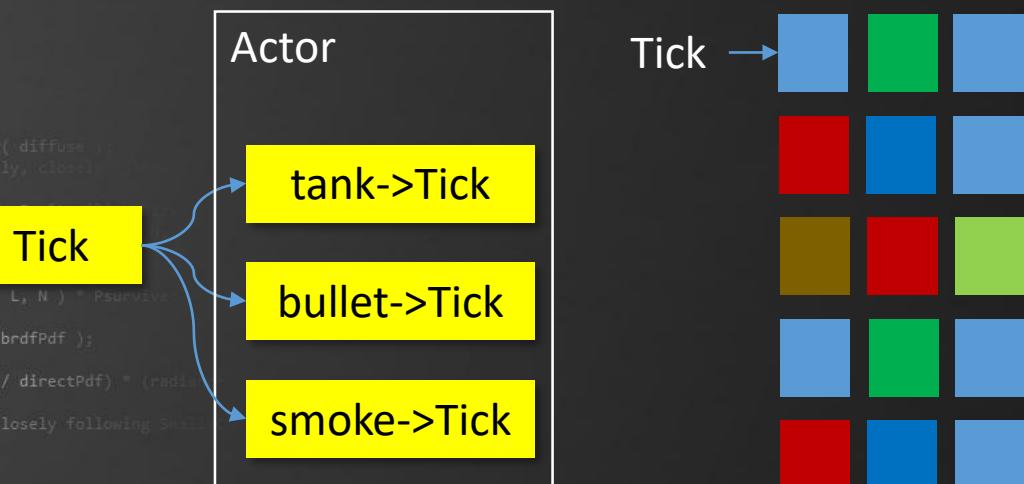


# OOP

## “Death by a Thousand Cuts”

### Object Oriented Programming:

- Objects
  - Data
  - Methods
  - Instances



### Cost of a virtual function call:

1. Virtual Function Table
2. No inlining

...

### Calling such a function:

1. Read pointer to VFT of base class
2. Add function offset
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.

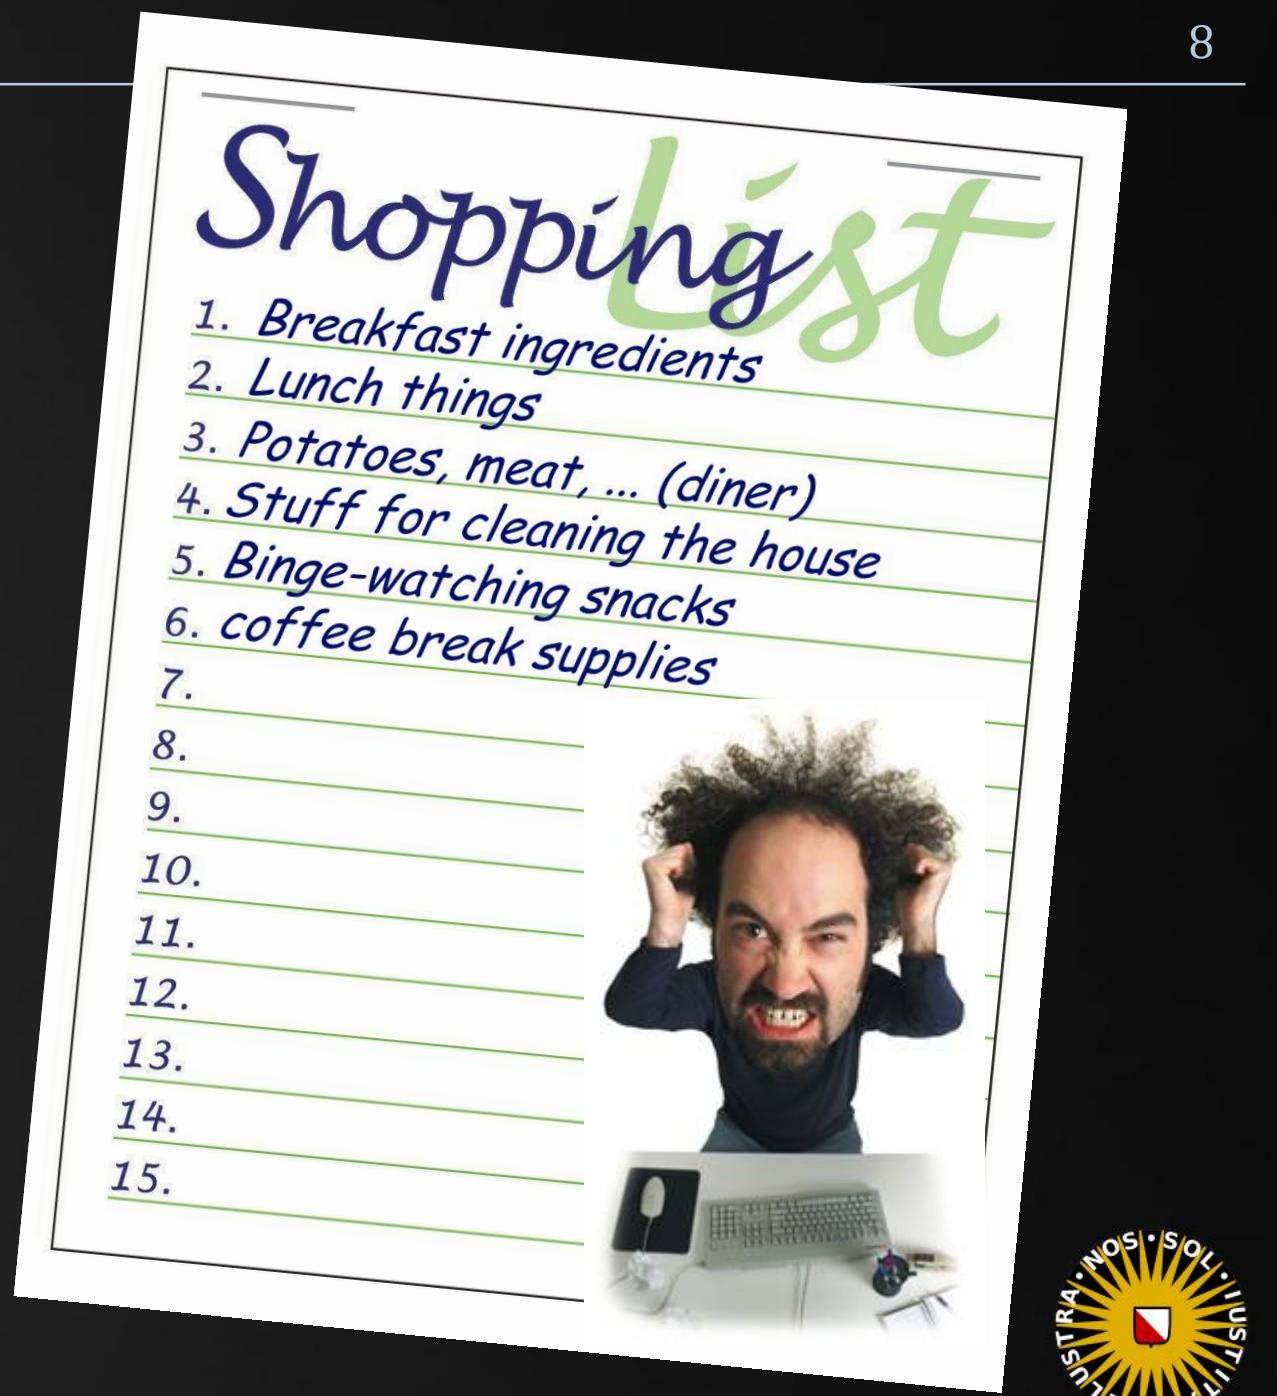
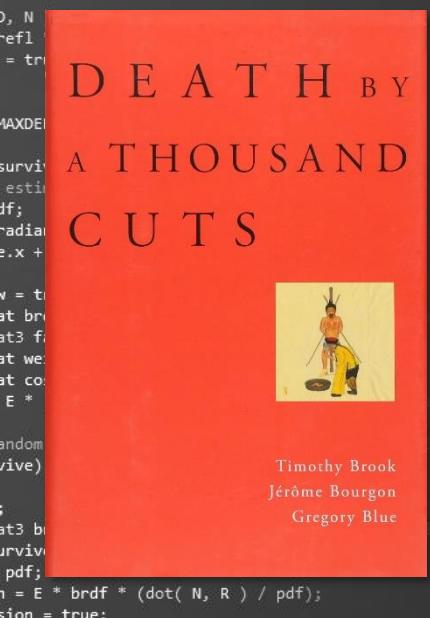


# OOP

## "Death by a Thousand Cuts"

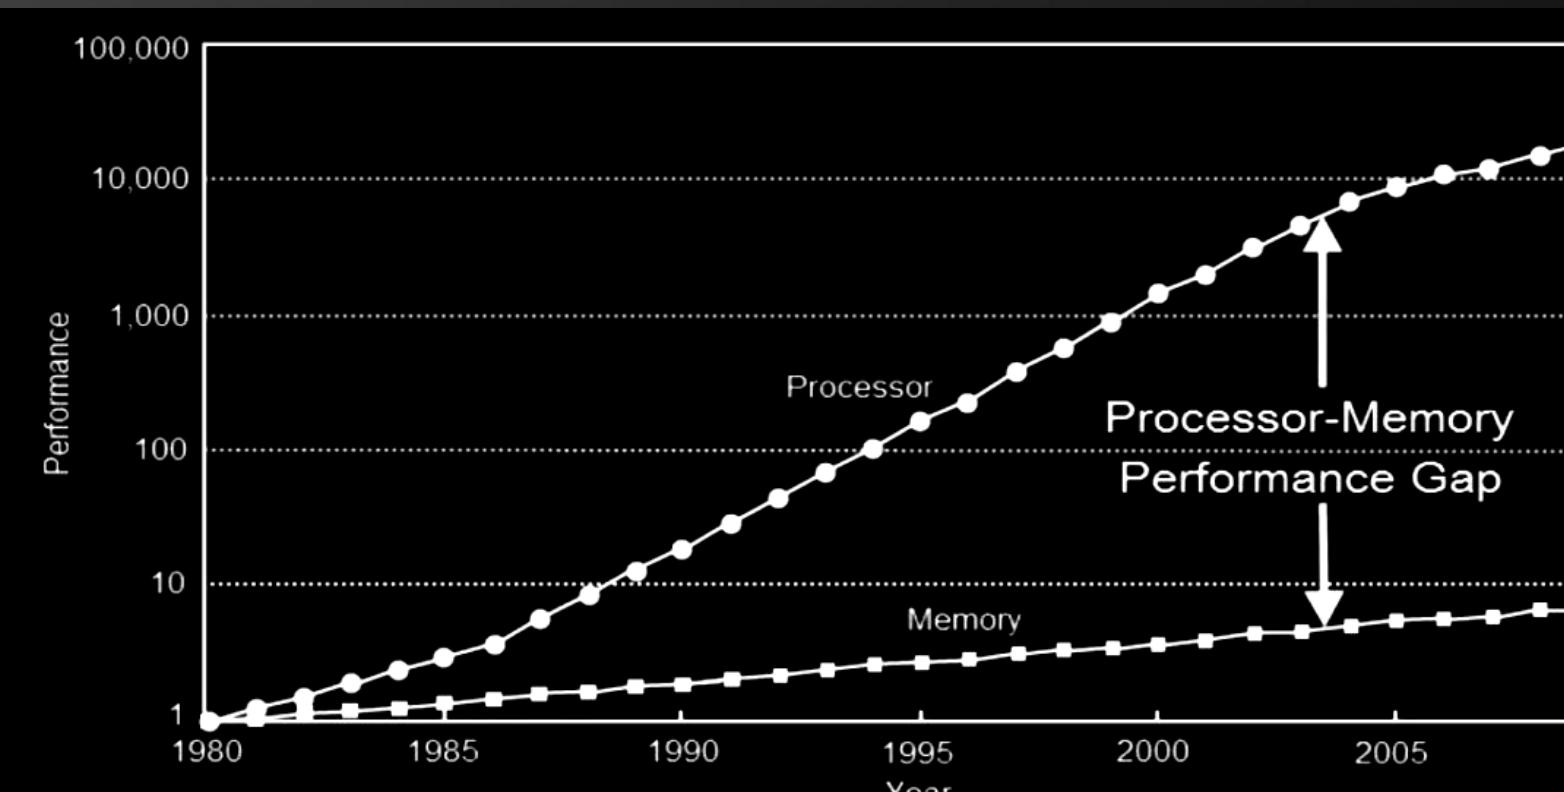
Characteristics of OOP:

- Virtual calls
- Scattered individual objects



## OOP

“Death by a Thousand Cuts”



*The problem is growing with time.*



# OOP

## “Death by a Thousand Cuts”

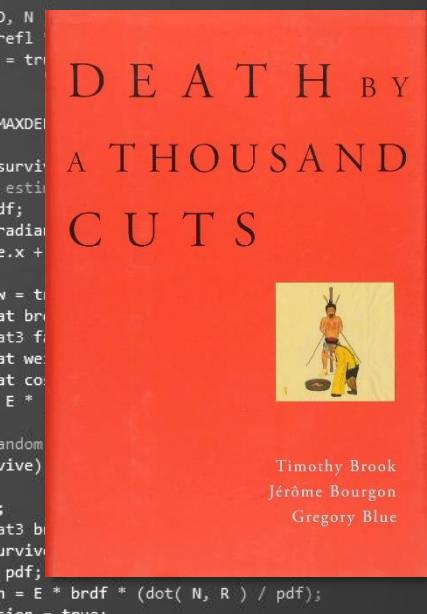
Dealing with “bandwidth starvation”:

Caching

Continuous memory access (full cache lines)

Large array continuous memory access

(caches ‘read ahead’)



# OOP

## "Death by a Thousand Cuts"

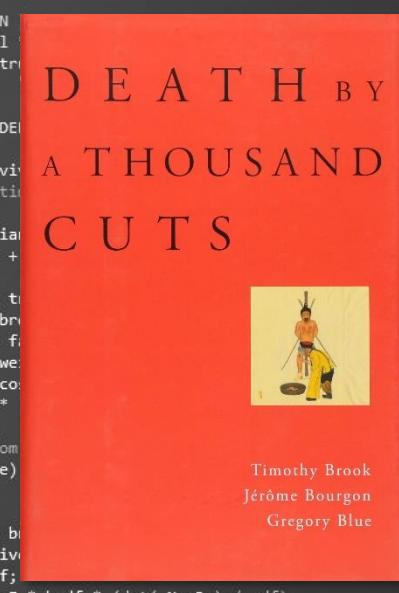
Code performance is typically bound by memory access.

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

→ Effort = CPU-effort.

"Most programs are made faster if we improve their memory access patterns."  
(this will be more true every year)

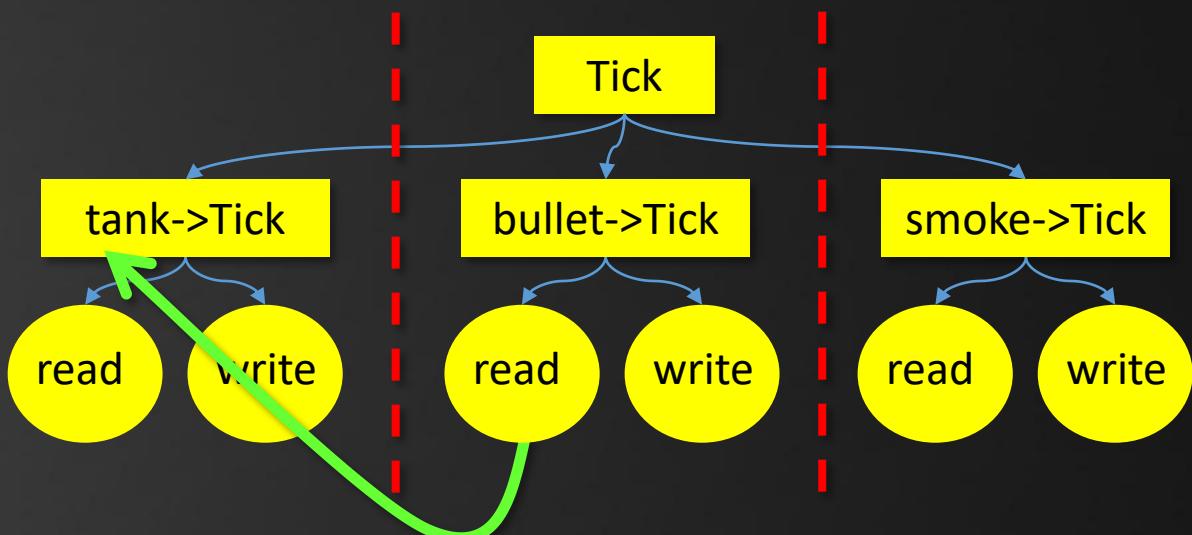
*"You cannot be fast without knowing how data is touched."*



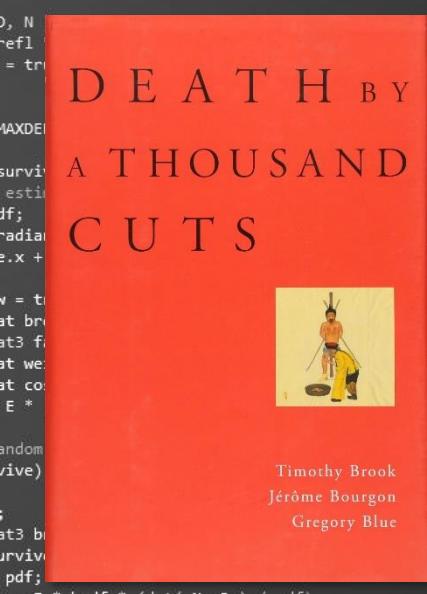
## OOP

**"Death by a Thousand Cuts"**

Parallel processing typically requires synchronization.



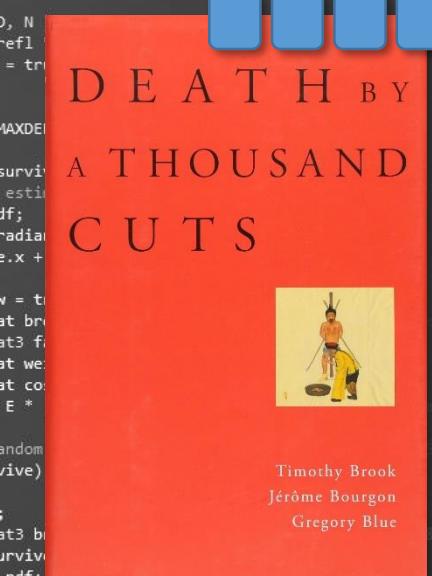
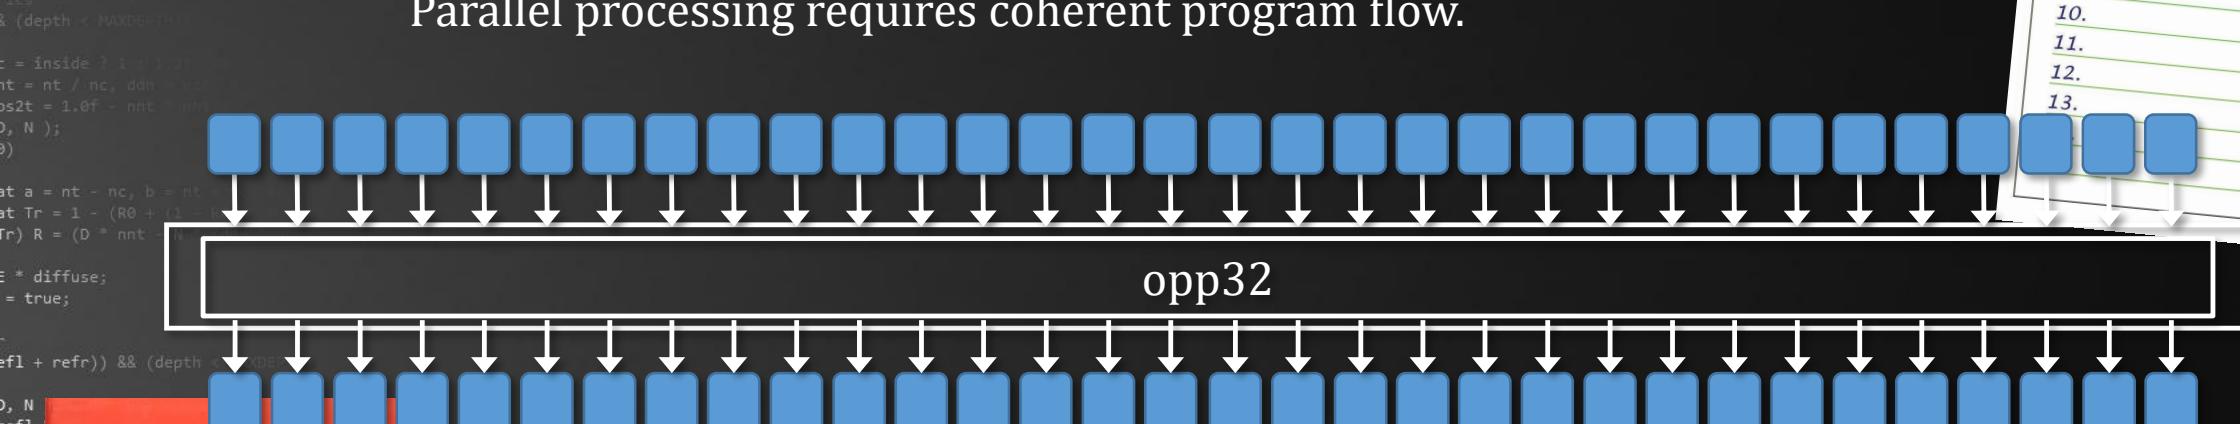
*"You cannot multi-thread without knowing how data is touched."*



## OOP

**"Death by a Thousand Cuts"**

Parallel processing requires coherent program flow.



*"You cannot multi-thread without knowing how data is touched."*



# OOP

## "Death by a Thousand Cuts"

```

rics
& (depth < MAXDEPTH)
nt = inside ? 1 : 1.0f;
nt = nt / nc; ddn = ddn / nc;
pos2t = 1.0f - nnt * nnt;
D, N );
)

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

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

MAXDEPTH)
survive = SurvivalProbability(
estimation - doing it properly;
if;
radiance = SampleLight( &rand
e.x + radiance.y + radiance.z ) > 0) && (depth <
true;
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;
at3 factor = diffuse * INVPi;
at weight = Mis2( directPpdf, brdfPpdf );
at cosThetaOut = dot( N, L );
E * (weight * cosThetaOut) / directPpdf ) * (radiance
random walk - done properly, closely following Smiley's
alive);

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

```

**branch**

**cache miss**

**cache miss**

**cached but not used**

**cached but not used**

**cache miss**

**cache miss**

**cache miss**



## OOP

## “Death by a Thousand Cuts”

```

rics
& (depth < MAXDEPTH)
    = inside ? 1 : 1.2f;
nt = nt / nc; ddn = ddn / nc;
pos2t = 1.0f - nnt - nnt;
D, N );
)
at a = nt - nc, b = nt + nc;
at Tr = 1 - (R0 + (1 - R0) *
Tr) R = (D * nnt + N * (ddn -
E * diffuse;
= true;

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

MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
if;
radiance = SampleLight( &rand, I, &L, &light );
e.x + radiance.y + radiance.z ) > 0) && (doe
e = true;
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;
at t3 factor = diffuse * INVPI;
at weight = Mis2( directPpdf, brdfPpdf );
at cosThetaOut = dot( N, L );
E * (weight * cosThetaOut) / directPpdf ) * (radiance
random walk - done properly, closely following Smiley
ive);

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

```

**void updateAims(**

**float\* aimDir,**

**const AimingData\* aim,**

**vec3 target,**

**uint count**

**)**

**{**

**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**

**only reads data that is actually needed to cache**

**reads from linear array**

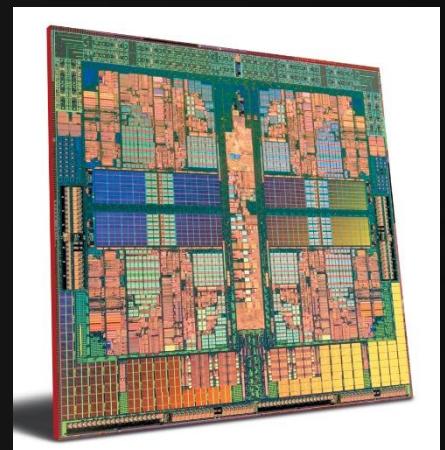
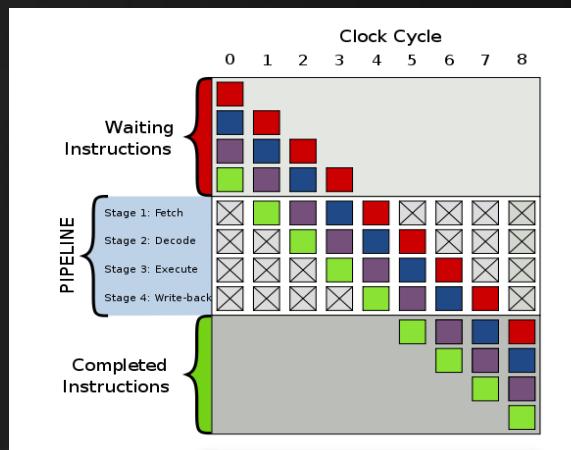
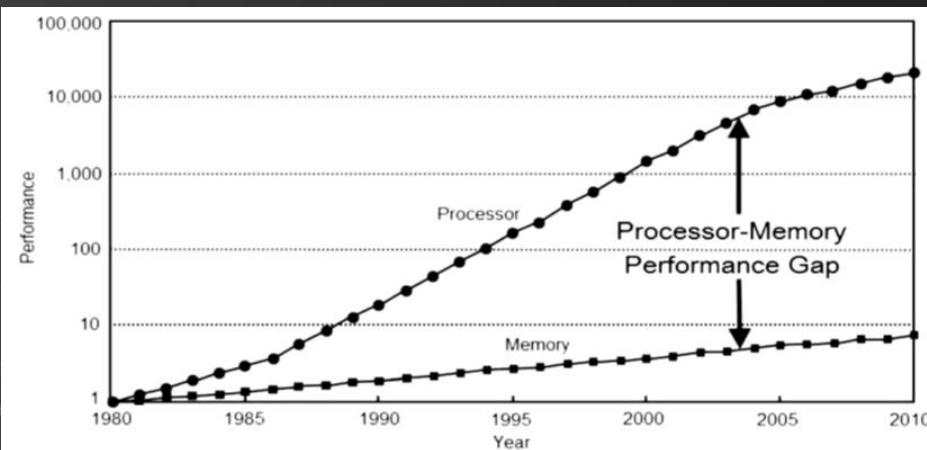


# OOP

## 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
- DOD or OO?



# 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

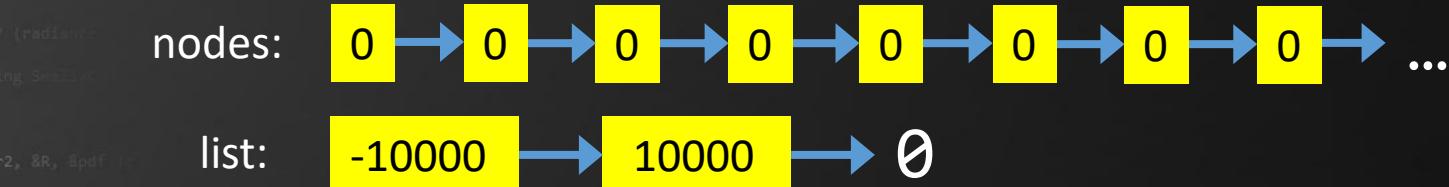
## Bad Access Patterns: Linked List

The Perfect LinkedList™:

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

LLNode* nodes = new LLNode[...];
LLNode* pool = nodes;

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



# DOD

## Bad Access Patterns: Linked List

The Perfect LinkedList™, experiment:

*Insert 25000 random values in the list so that we obtain a sorted sequence.*

```
for( 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;
}
```

```
rics
  & (depth < MAXDEPTH)
  c = inside ? 1 : 1.2f;
  nt = nt / nc, ddn = ddc;
  pos2t = 1.0f - nnt * nnt;
  D, N );
}
at a = nt - nc, b = nt + nc;
at Tr = 1 - (R0 + (1 - R0) *
Tr) R = (D * nnt - N * (ddn
E * diffuse;
= true;

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

MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
if;
radiance = SampleLight( &rand, I, &L, &lighting );
e.x + radiance.y + radiance.z ) > 0) && (dot( N,
e.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 Smiley
ive);

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



# 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



```

rics
    & (depth < MAXDEPTH)
    for( 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;
    }
    MAXDEPTH)
    survive = SurvivalProbability( diffuse );
    estimation - doing it properly, closely
    if;
    radiance = SampleLight( &rand, I, &L, &light,
        e.x + radiance.y + radiance.z ) > 0) && (dot( N,
        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 Smiley
    /alive)
    ;
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf
    /survive;
    /pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    ision = 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.reedbetta.com/blog/data-oriented-hash-table>

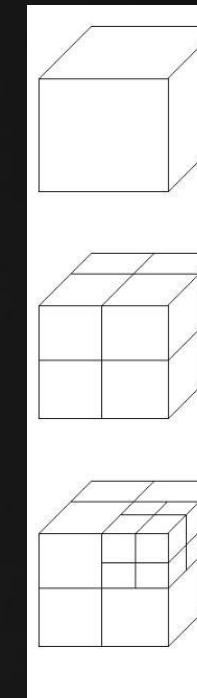
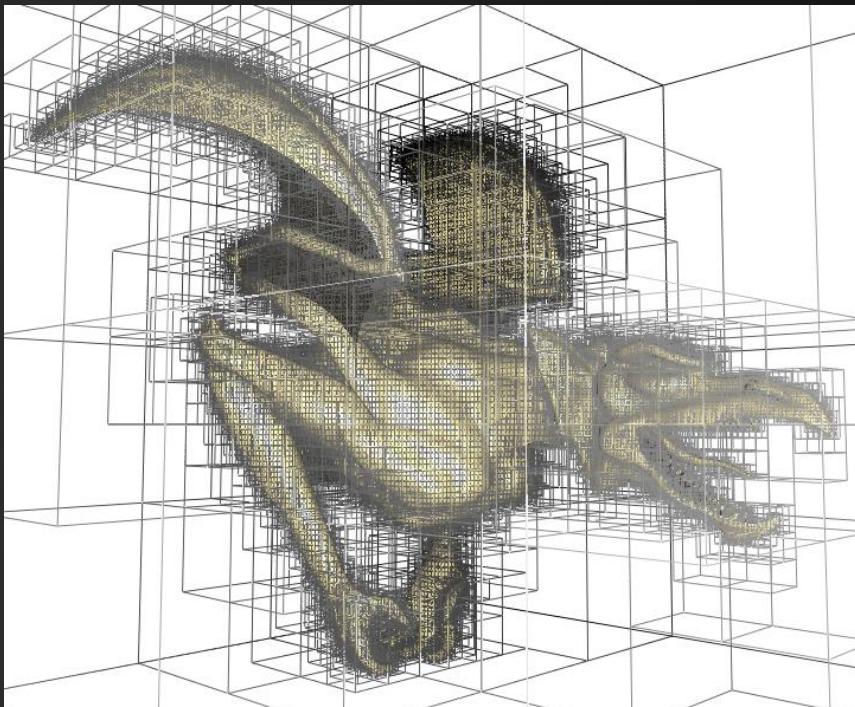


## DOD

## Bad Access Patterns: Octree



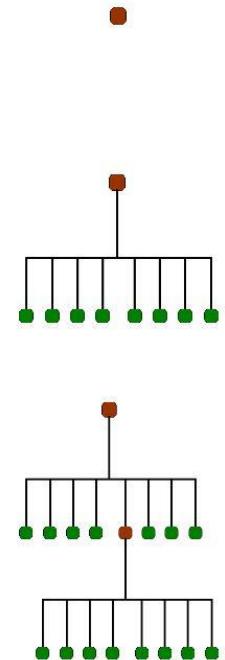
```
rics  
3 (depth  
nt = inside  
nt = nt / 2  
pos2t = 1.0  
(D, N );  
(D)  
at a = nt;  
at Tr = 1 -  
(Tr) R = (D  
E * diffuse  
= true;  
  
-  
refl + refl  
(D, N );  
refl * E *  
= true;  
  
MAXDEPTH)  
survive = $  
estimation  
df;  
radiance =  
e.x + radi  
  
v = true;  
at brdfPdf = evaluateDiffuse(  
at3 factor = diffuse * INVPI;  
at weight = Mis2( directPdf, brdfPdf );  
at cosThetaOut = dot( N, L );  
E * ((weight * cosThetaOut) / directPdf);  
random walk - done properly, closely following S  
survive)  
  
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf;  
survive;  
pdf;  
n = E * brdf * (dot( N, R ) / pdf);  
survive = true;
```



Root

Level 1

Level 2



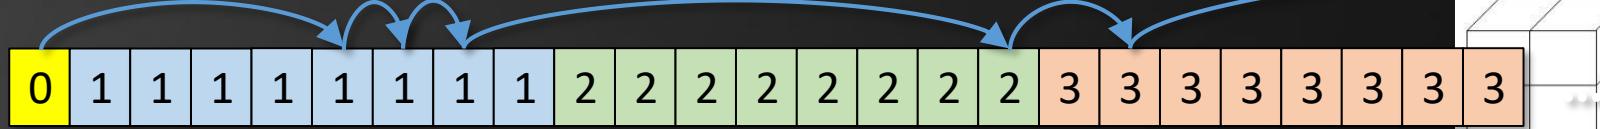
# DOD

## Bad Access Patterns: Octree

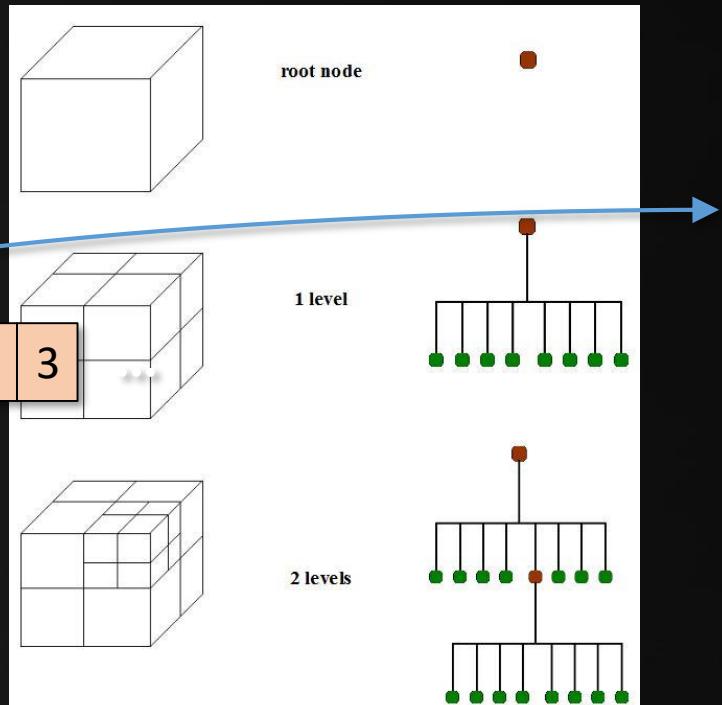
Query: find the color of a voxel visible through pixel (x,y).

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

Data layout:



Color data: 32-bit (ARGB).



# DOD

## Bad Access Patterns: Octree

Alternative layout:

1. Tree 1: occlusion (1 bit per voxel);
2. Tree 2: color information (32 bits per voxel).

Use tree 1 to find the voxel you are looking for.

Lookup the correct voxel (incurring a single cache miss) in tree 2.

Caching in tree 1:

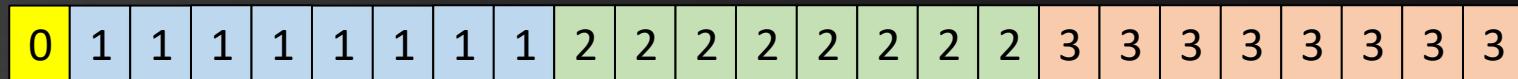
- A cache line holds  $64 \times 8 = 512$  voxels
- Accessing the root gets several levels in L1 cache



# DOD

## Bad Access Patterns: Octree

Alternative layout (part 2):



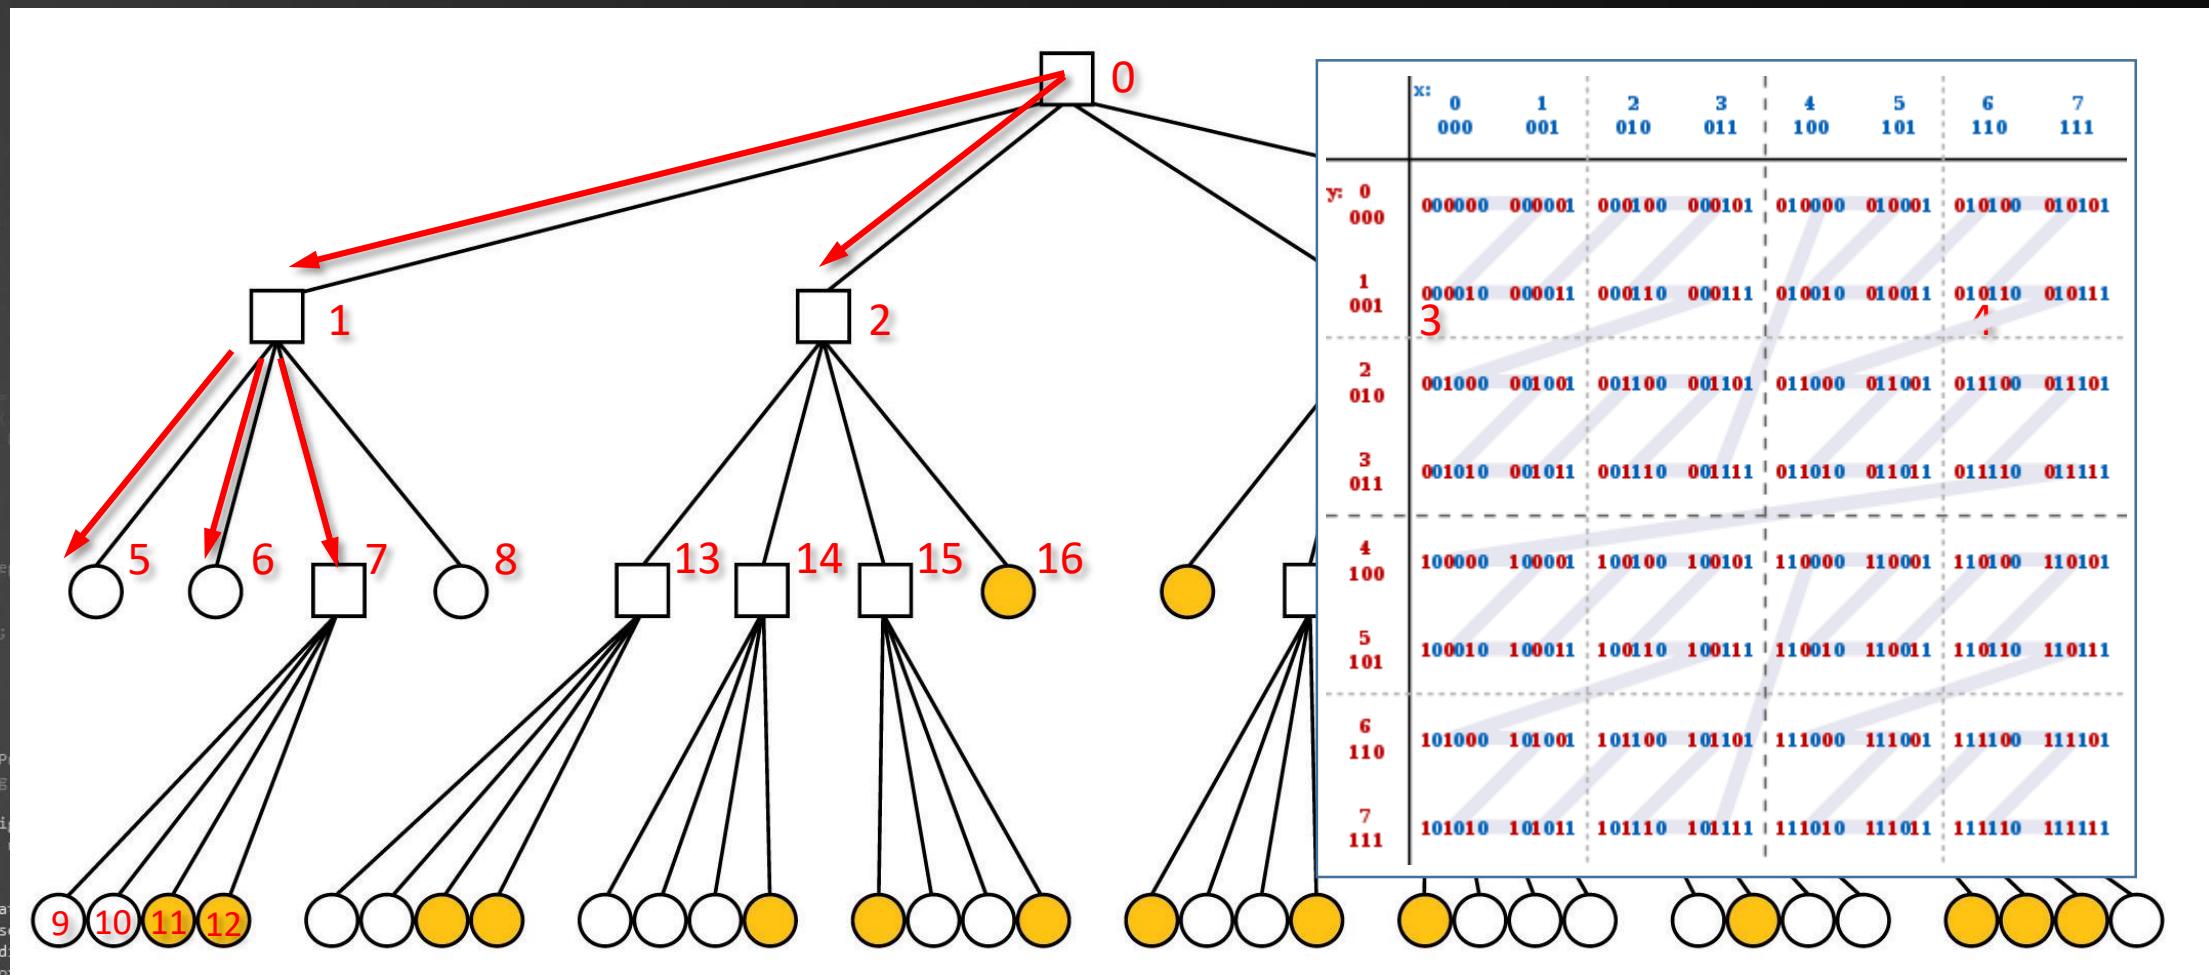
Trees are typically generated by a divide-and-conquer algorithm, in a depth-first fashion.

Compact storage:

```
struct OTNode
{
    int firstChild;
    // bit 31 set: empty
```



## DOD



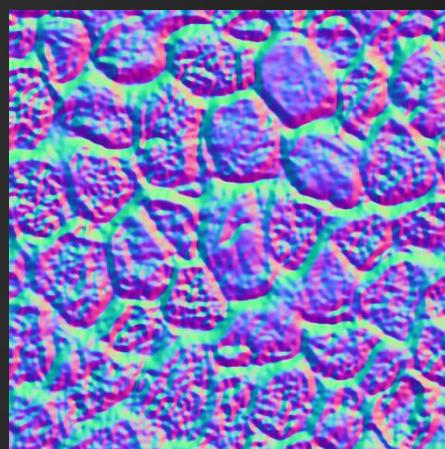
# DOD

## 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.*



# 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.

```
rics
& (depth < MAXDEPTH)
    c = inside ? 1 : 1.2f;
    nt = nt / nc; ddn = ddn / nc;
    pos2t = 1.0f - nnt * nnt;
    D, N );
}
}

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

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

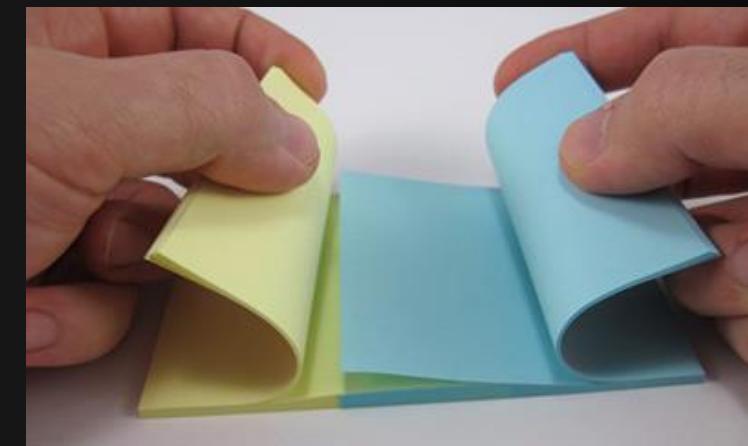
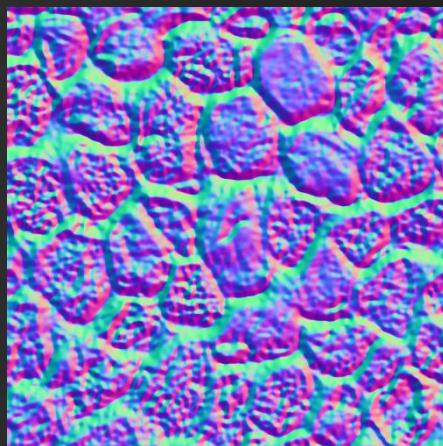
D, N );
refl * E * diffuse;
= true;

MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
df;
radiance = SampleLight( &rand, I, &L, &L1,
e.x + radiance.y + radiance.z) > 0) && (d
e = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Pst
at3 factor = diffuse * INVPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPdf)

random walk - done properly, closely follow
survive)

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

```

rics
& (depth < MAXDEPTH)
    c = inside ? 1 : 1.2f;
    nt = nc / nc, ddn = nc;
    pos2t = 1.0f - nnt * nnt;
    D, N );
}
at a = nt - nc, b = nt + nc;
at Tr = 1 - (R0 + (1 - R0) *
Tr) R = (D * nnt + N * (ddn
E * diffuse;
= true;

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

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

MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
if;
radiance = SampleLight( &rand, I, &L, &light,
e.x + radiance.y + radiance.z) > 0) && (dot( N,
e.y + radiance.z) > 0) && (dot( L,
e.z + radiance.x) > 0) && (dot( N,
e.x + radiance.y) > 0) && (dot( L,
e.y + radiance.z) > 0) && (dot( N,
e.z + radiance.x) > 0);
v = true;
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;
at3 factor = diffuse * INVPI;
at weight = Mis2( directPpdf, brdfPpdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPpdf) * (radiance
random walk - done properly, closely following Smiley
ive);

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

```



# DOD

## Previously in INFOMOV

```
rics
  & (depth < MAXDEPTH)
  n = nt / nc, ddn = ddc;
  os2t = 1.0f - nnt * nnt;
  D, N );
}

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

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

```
MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
if;
radiance = SampleLight( &rand, I, &L, &lightDir,
e.x + radiance.y + radiance.z) > 0) && (dot( N,
e, v = true;
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;
at t3 factor = diffuse * INVPI;
at weight = Mis2( directPpdf, brdfPpdf );
at cosThetaOut = dot( N, L );
E * (weight * cosThetaOut) / directPpdf) * (radiance
random walk - done properly, closely following Smiley
ive);

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

Better:

```
struct Particle
{
    float x, y, z;
    float vx, vy, vz;
    float mass;
};

// size: 28 bytes
```

```
struct Particle
{
    float x, y, z;
    float vx, vy, vz;
    float mass, dummy;
};

// size: 32 bytes
```



# DOD

Previously in INFOMOV

```
rics
  & (depth < MAXDEPTH)
  c = inside ? 1 : 1.2f;
  nt = nc / c, ddn = ddc;
  pos2t = 1.0f - nnt * nnt;
  D, N );
  )
at a = nt - nc, b = nt + nc;
at Tr = 1 - (R0 + (1 - R0) *
Tr) R = (D * nnt - N * (ddn *
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, &L, &lighting);
e.x + radiance.y + radiance.z) > 0) && (dot( N,
  v = true;
  at brdfPpdf = EvaluateDiffuse( L, N ) * Psumvive;
  at3 factor = diffuse * INVPi;
  at weight = Mis2( directPpdf, brdfPpdf );
  at cosThetaOut = dot( N, L );
  E * (weight * cosThetaOut) / directPpdf) * (radiance -
random walk - done properly, closely following Smiley's
alive);

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

# AOS

# SOA

structure  
of  
arrays



# DOD

## Previously in INFOMOV

```

rics
  & (depth < MAXDEPTH)
  n = nt / nc, ddn = ddc * n;
  pos2t = 1.0f - nnt * nnt;
  D, N );
}
}

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

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

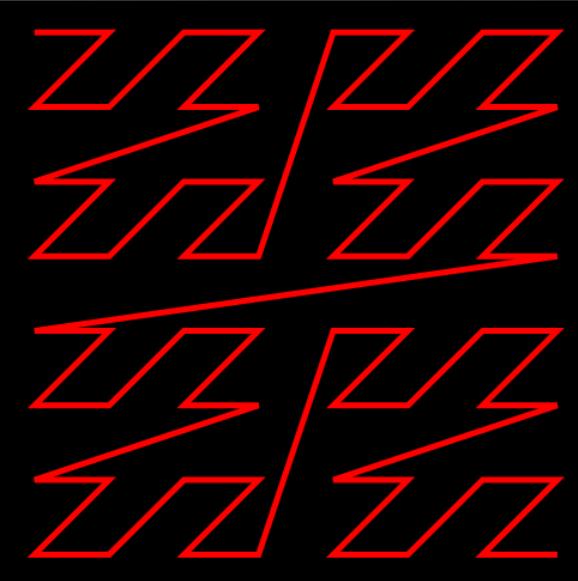
D, N );
refl * E * diffuse;
= true;

MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
if;
radiance = SampleLight( &rand, I, &L, &light
e.x + radiance.y + radiance.z ) > 0) && (do
e = true;
at brdfPdf = EvaluateDiffuse( L, N ) * Psurvive;
at3 factor = diffuse * INMPI;
at weight = Mis2( directPdf, brdfPdf );
at cosThetaOut = dot( N, L );
E * (weight * cosThetaOut) / directPdf) * (radiance
random walk - done properly, closely following Smith's
survive);

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



# Today's Agenda:

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



## 2B|~2B

OO = Evil, DO = Good?

10% of your code runs 90% of the time. DO is good for this 10%.

For all other code, please:

- Use STL
- Apply OO
- Program in C#
- Use event handling
- Check return values
- Focus on productivity



2B|~2B



<https://www.youtube.com/watch?v=rX0ItVEVjHc>



2B|~2B

[Next](#) [Up](#) [Previous](#) [Contents](#)Next: [Contents](#) [Contents](#)**Online release of Data-Oriented Design :**

This is the free, online, reduced version. Some inessential chapters are excluded from this version, but in the spirit of this being an education resource, the essentials are present for anyone wanting to learn about data-oriented design. Expect some odd formatting and some broken images and listings as this is auto generated and the Latex to html converters available are not perfect. If the source code listing is broken, you should be able to find the referenced source on [github](#). If you like what you read here, consider purchasing the real paper book from [here](#), as not only will it look a lot better, but it will help keep this version online for those who cannot afford to buy it. Please send any feedback to [support@dataorienteddesign.com](mailto:support@dataorienteddesign.com)

# Data-Oriented Design

**Richard Fabian**

- [Contents](#)
- [Data-Oriented Design](#)
  - [It's all about the data](#)
  - [Data is not the problem domain](#)
  - [Data and statistics](#)
  - [Data can change](#)
  - [How is data formed?](#)
  - [The framework](#)
  - [Conclusions and takeaways](#)

<http://www.dataorienteddesign.com/dodbook/>

2B|~2B

# Data Oriented Design Resources

A curated list of awesome data oriented design resources.

Feel free to contribute by sending PR!

- [Data Oriented Design](#)
  - [Presentations](#)
  - [Blog Posts](#)
  - [Videos](#)
  - [Other](#)
  - [Code Examples](#)

## Presentations

- [A Step Towards Data Orientation \(2010\)](#) - Johan Torp
- [Introduction To Data Oriented Design \(2010\)](#) - DICE
- [Memory Optimization \(2003\)](#) - Christer Ericson
- [Practical Examples In Data Oriented Design \(2013\)](#) - Niklas Frykholm
- [Three Big Lies \(2008\)](#) - Mike Acton
- [Typical C++ Bullshit \(2008\)](#) - Mike Acton
- [Data-Oriented Design and C++ \(2014\)](#) - Mike Acton
- [Entity Component Systems & Data Oriented Design \(2018\)](#) - Aras Pranckevičius

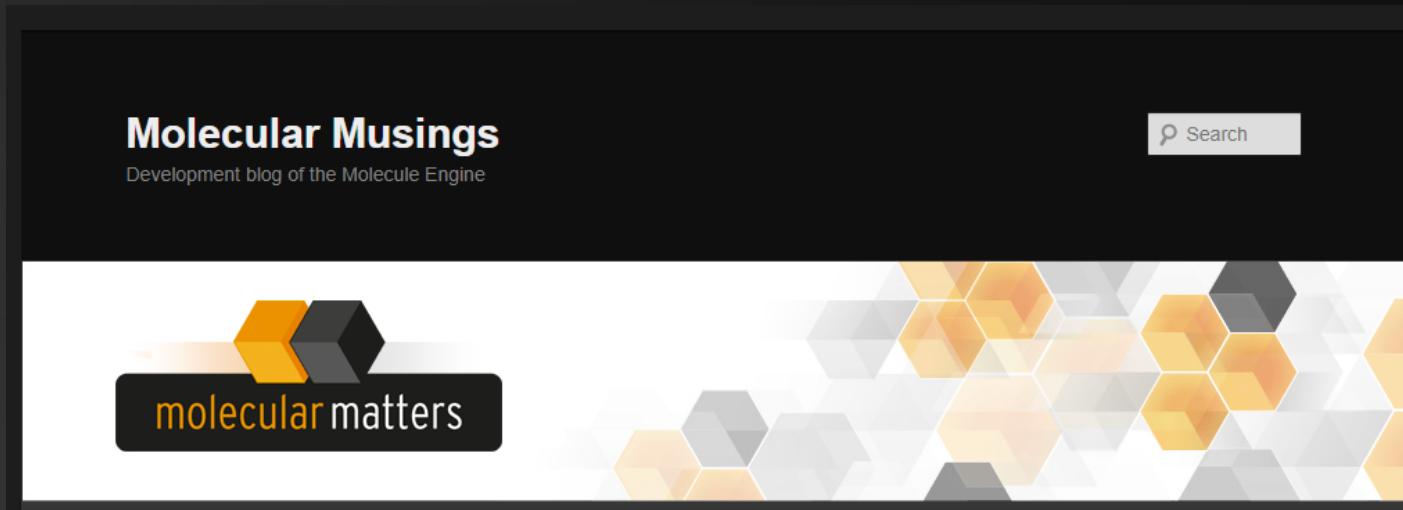
## Blog Posts

<https://github.com/dbartolini/data-oriented-design>

- [Adventures in data-oriented design – Part 1: Mesh data \(2011\)](#) - Stefan Reinalter



2B|~2B



<https://blog.molecular-matters.com/2011/11/03/adventures-in-data-oriented-design-part-1-mesh-data-3/>  
<https://blog.molecular-matters.com/2013/02/22/adventures-in-data-oriented-design-part-2-hierarchical-data/>  
<https://blog.molecular-matters.com/2013/05/02/adventures-in-data-oriented-design-part-3a-ownership/>  
<https://blog.molecular-matters.com/2013/05/17/adventures-in-data-oriented-design-part-3b-internal-references/>

## Adventures in data-oriented design – Part 1: Mesh data

Let's face it, performance on modern processors (be it PCs, consoles or mobiles) is mostly governed by memory access patterns. Still, data-oriented design is considered something new and novel, and only slowly creeps into programmers' brains, and this really needs to change. Having co-workers fix your code and improving its performance really is no excuse for writing crappy code (from a performance point-of-view) in the first place.

This post is the first in an ongoing series about how certain things in the Molecule Engine are done in a data-oriented fashion, while still making use of OOP concepts. A common misconception about data-oriented design is that it is “C-like”, and “not OOP”, and therefore less maintainable – but that need not be the case. The concrete example we are going to look at today is how to organize your mesh data, but let's start with the pre-requisites first.



# Today's Agenda:

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



# /INFOMOV/

## END of “Data-Oriented Design”

next lecture: “GPGPU (1)”

```
rics
3 (depth < MAXDEPTH)
    t = inside ? 1 : 1.2f;
    nt = nt / nc, ddn = ddn / nc;
    os2t = 1.0f - nnt * nnt;
    D, N );
}
}

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

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

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely
if;
radiance = SampleLight( &rand, I, &L, &lighting,
e.x + radiance.y + radiance.z) > 0) && (dot( N,
e, L ) > 0);
v = true;
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;
at t3 factor = diffuse * INVPi;
at weight = Mis2( directPpdf, brdfPpdf );
at cosThetaOut = dot( N, L );
E * ((weight * cosThetaOut) / directPpdf) * (radiance -
random walk - done properly, closely following Smiley
alive);

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

