

```
ics
& (depth < MAXDEPTH)
= inside ? 1.0f : 0.0f;
nt = nt / nc; ddn = abs(
os2t = 1.0f - nnt * nnt;
D, N );
)
at a = nt - nc, b = nt + nc;
at Tr = 1 - (R0 + (1 - R0) *
Fr) R = (D * nnt - N * (ddn
E * diffuse;
= true;
efl + refr)) && (depth < MAXDEPTH)
D, N );
refl * E * diffuse;
= true;
MAXDEPTH)
survive = SurvivalProbability( diffuse, r1, r2, &R, $pdf );
estimation - doing it properly, closely following
if;
radiance = SampleLight( &rand, I, &L, &light,
e.x + radiance.y + radiance.z) > 0) && (abs(
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;
```

/INFOMOV/

Optimization & Vectorization

J. Bikker - Sep-Nov 2018 - Lecture 5: "SIMD (1)"

Welcome!




Meanwhile, on ars technica

```

ics
& (depth < MAXDEPTH)
    c = inside ? 1.0f : 0.0f;
    nt = nt / nc, dd = dd / dc;
    cos2t = 1.0f - nnt * nnt;
    D, N );
    )
    at a = nt - nc, b = nt + nc;
    at Tr = 1 - (R0 + (1 - R0) * c);
    Tr) R = (D * nnt - N * (dd
    )
    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, &L, Alignment
    e.x + radiance.y + radiance.z) > 0) && (depth <
    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) * (radiance
    random walk - done properly, closely following
    ve)
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, Spdf
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;
    
```

TECH— HBM3: Cheaper, up to 64GB on-package, and terabytes-per-second bandwidth



Plus, Samsung unveils GDDR6 and "low cost" HBM technologies.

MARK WALTON (UK) - 8/23/2016, 3:02 PM

46



Despite first- and second-generation High Bandwidth Memory having made few appearances in shipping products, Samsung and Hynix are already working on a followup: HBM3. Teased at the Hot Chips symposium in Cupertino, California, HBM3 will offer improved density, bandwidth, and power efficiency. Perhaps most importantly though, given the high cost of HBM1 and HBM2, HBM3 will be cheaper to produce.

With conventional memory setups, RAM chips are placed next to each other on a circuit board, usually as close as possible to the logic device (CPU or GPU) that needs access to the RAM. HBM, however, stacks a bunch of RAM dies (dice?) on top of each other, connecting them directly with through-silicon vias (TSVs). These stacks of RAM are then placed on the logic chip package, which reduces the surface area of the device (AMD's Fury Nano is a prime example), and potentially provides a massive boost in bandwidth.

HBM Gen3

- Plan to extend capacity / bandwidth 2X or more with similar power budget

	HBM	Gen2	Gen3
Schedule	Year	2016	2019/2020
	Density	8Gb	Higher (>2X)
	Stack height	2,4,8	Higher (>8H)
	# of p-Channel	16 p-CH	Similar or more

Meanwhile, the job market

```

ics
& (depth < MAXDEPTH)
    c = inside ? 1.0f : 0.0f;
    nt = nt / nc, ddn = ddn * c;
    r2s2t = 1.0f - nnt * ddn;
    D, N );
    )
    at a = nt - nc, b = nt + nc;
    at Tr = 1 - (R0 + (1 - R0) * ddn);
    R = (D * nnt - N * (ddn * r2s2t +
    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, &L, &light,
    e.x + radiance.y + radiance.z) > 0) && (depth <
    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) * (radiance
    random walk - done properly, closely following
    ve)
    ,
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf;
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    ion = true;
    
```

[Schedule a meeting](#) or [Call +31 854865760](#)

- Industries
- Our Services and Products
- Technologies and Expertise
- Our Partners and Activities
- Our Company
- Blog
- Contact us

JOBS

StreamHPC > > About Us > > Jobs

We only have jobs for people who get bored easily.

In return we offer **a solution to boredom**, as performance engineering is hard.

As the market demand for affordable high performance software grows, StreamHPC continuously looks for people to join the team. With upcoming products and new markets like OpenCL on low-power ARM processors, we expect continuous growth for the years to come.

To apply send your motivation and a recent version of your CV to jobs@streamhpc.com. If you have LinkedIn, you can easily build a CV with [LinkedIn Resume Builder](#) and send us the link.

- **OpenCL/CUDA expert** (Also accepting freelancers)
- **Sales support**

The procedure for the technical roles is as follows:

- You send a CV and tell us why you are the perfect candidate.
- After that you are invited for a longer online test. You show your skills on C/C++ and algorithms. You will receive a PDF with useful feedback.
- If you selected GPGPU or mentioned it, we send you a GPU assignment. You need to pick out the right optimisations, code it and explain your decisions. (Hopefully under 30 minutes)
- If all goes well, you'll have a videochat with Vincent (CEO) on personal and practical matters. You can also ask us anything, to find out if we fit you. (Around 1 hour)

```
ics
& (depth < MAXDEPTH)
{
    if ( ! inside )
    {
        nt = nt / nc; ddn = ddn * ddn;
        cos2t = 1.0f - nnt * ddn;
        D, N );
    }
}

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

E * diffuse;
= true;

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

MAXDEPTH)

survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely following
if;
radiance = SampleLight( &rand, I, &L, &light);
e.x + radiance.y + radiance.z) > 0) && (depth <
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) * (radiance

random walk - done properly, closely following
vive)
;
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;
```

Today's Agenda:

- Introduction
- Intel: SSE
- Streams
- Vectorization



Introduction

Consistent Approach

(0.) Determine optimization requirements

1. Profile: determine hotspots
2. Analyze hotspots: determine scalability
3. Apply high level optimizations to hotspots
4. Profile again.
5. Parallelize / vectorize / use GPGPU
6. Profile again.
7. Apply low level optimizations to hotspots
8. Repeat steps 7 and 8 until time runs out
9. Report.

Rules of Engagement

1. Avoid Costly Operations
2. Precalculate
3. Pick the Right Data Type
4. Avoid Conditional Branches
5. Early Out
6. Use the Power of Two
7. Do Things Simultaneously



Introduction

S.I.M.D.

Single Instruction Multiple Data:

Applying the same instruction to several input elements.

In other words: if we are going to apply the same sequence of instructions to a large input set, this allows us to do this in parallel (and thus: faster).

SIMD is also known as *instruction level parallelism*.

Examples:

```
union { uint a4; unsigned char a[4]; };
do
{
    GetFourRandomValues( a );
}
while (a4 != 0);
```

```
unsigned char a[4] = { 1, 2, 3, 4 };
unsigned char b[4] = { 5, 5, 5, 5 };
unsigned char c[4];
*(uint*)c = *(uint*)a + *(uint*)b;
// c is now { 6, 7, 8, 9 }.
```



Introduction

uint = unsigned char[4]

Pinging google.com yields: 74.125.136.101

Each value is an unsigned 8-bit value (0..255).

Combing them in one 32-bit integer:

$$101 + 256 * 136 + 256 * 256 * 125 + 256 * 256 * 256 * 74 = 1249740901.$$

Browse to: <http://1249740901> (works!)

Evil use of this:

We can specify a user name when visiting a website, but any username will be accepted by google. Like this:

<http://infomov@google.com>

Or:

<http://www.ing.nl@1249740901>

Replace the IP address used here by your own site which contains a copy of the ing.nl site to obtain passwords, and send the link to a ‘friend’.



Introduction

Example: color scaling

Assume we represent colors as 32-bit ARGB values using unsigned ints:



To scale this color by a specified percentage, we use the following code:

```
uint ScaleColor( uint c, float x ) // x = 0..1
{
    uint red = (c >> 16) & 255;
    uint green = (c >> 8) & 255;
    uint blue = c & 255;
    red = red * x, green = green * x, blue = blue * x;
    return (red << 16) + (green << 8) + blue;
}
```



Introduction



Example: color scaling


```
uint ScaleColor( uint c, float x ) // x = 0..1
{
    uint red = (c >> 16) & 255, green = (c >> 8) & 255, blue = c & 255;
    red = red * x, green = green * x, blue = blue * x;
    return (red << 16) + (green << 8) + blue;
}
```

Improved:

```
uint ScaleColor( uint c, uint x ) // x = 0..255
{
    uint red = (c >> 16) & 255, green = (c >> 8) & 255, blue = c & 255;
    red = (red * x) >> 8;
    green = (green * x) >> 8;
    blue = (blue * x) >> 8;
    return (red << 16) + (green << 8) + blue;
}
```



Introduction



Example: color scaling

```
uint ScaleColor( uint c, uint x ) // x = 0..255
{
    uint red = (c >> 16) & 255, green = (c >> 8) & 255, blue = c & 255;
    red = (red * x) >> 8, green = (green * x) >> 8, blue = (blue * x) >> 8;
    return (red << 16) + (green << 8) + blue;
}
```

7 shifts, 3 ands, 3 muls, 2 adds


Improved:

```
uint ScaleColor( const uint c, const uint x ) // x = 0..255
{
    uint redblue = c & 0x00FF00FF;
    uint green    = c & 0x0000FF00;
    redblue = ((redblue * x) >> 8) & 0x00FF00FF;
    green = ((green * x) >> 8) & 0x0000FF00;
    return redblue + green;
}
```

2 shifts, 4 ands, 2 muls, 1 add



Introduction



Example: color scaling

```
uint ScaleColor( uint c, uint x ) // x = 0..255
{
    uint red = (c >> 16) & 255, green = (c >> 8) & 255, blue = c & 255;
    red = (red * x) >> 8, green = (green * x) >> 8, blue = (blue * x) >> 8;
    return (red << 16) + (green << 8) + blue;
}
```

*7 shifts, 3 ands, 3 muls, 2 adds
(15 ops)*

Further improved:

```
uint ScaleColor( const uint c, const uint x ) // x = 0..255
{
    uint redblue = c & 0x00FF00FF;
    uint green    = c & 0x0000FF00;
    redblue = (redblue * x) & 0xFF00FF00;
    green = (green * x) & 0x00FF0000;
    return (redblue + green) >> 8;
}
```

*1 shift, 4 ands, 2 muls, 1 add
(8 ops)*



Introduction

Other Examples

Rapid string comparison:

```

char a[] = "optimization skills rule";
char b[] = "optimization is so nice!";
bool equal = true;
int l = strlen( a );
for ( int i = 0; i < l; i++ )
{
    if (a[i] != b[i])
    {
        equal = false;
        break;
    }
}

```

Likewise, we can copy byte arrays faster.

```

char a[] = "optimization skills rule";
char b[] = "optimization is so nice!";
bool equal = true;
int q = strlen( a ) / 4;
for ( int i = 0; i < q; i++ )
{
    if (((int*)a)[i] != ((int*)b)[i])
    {
        equal = false;
        break;
    }
}

```



Introduction

SIMD using 32-bit values - Limitations

Mapping four chars to an int value has a number of limitations:

$$\{ 100, 100, 100, 100 \} + \{ 1, 1, 1, 200 \} = \{ 101, 101, 102, 44 \}$$

$$\{ 100, 100, 100, 100 \} * \{ 2, 2, 2, 2 \} = \{ \dots \}$$

$$\{ 100, 100, 100, 200 \} * 2 = \{ 200, 200, 201, 144 \}$$

In general:

- Streams are not separated (prone to overflow into next stream);
- Limited to small unsigned integer values;
- Hard to do multiplication / division.

```

ics
& (depth < MAXDEPTH)
{
    int inside = 1;
    int nt = nt / nc;
    double ns2t = 1.0f * nnt * nnt;
    double D, N );
}

at a = nt - nc; b = nt + nc;
at Tr = 1 - (R0 + (1 - R0) * ns2t);
Tr) R = (D * nnt - N * (D + ns2t));

E * diffuse;
= true;

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

MAXDEPTH)
survive = SurvivalProbability( diffuse, ns2t );
estimation - doing it properly, closely following
if;
radiance = SampleLight( &rand, I, &L, &light );
e.x + radiance.y + radiance.z) > 0) && (depth <
{
    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;
}
    
```



Introduction

SIMD using 32-bit values - Limitations

Ideally, we would like to see:

- Isolated streams
- Support for more data types (char, short, uint, int, float, double)
- An easy to use approach

Meet SSE!

```

ics
& (depth < MAXDEPTH)
{
    if (inside ? 1 : 0)
    {
        nt = nt / nc; ddn = ddn * 2;
        ps2t = 1.0f - nnt * nnt;
        D, N );
    }
}

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

E * diffuse;
= true;

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

MAXDEPTH)
{
    survive = SurvivalProbability( diffuse );
    estimation - doing it properly, closely following
    if;
    radiance = SampleLight( &rand, I, &L, &light);
    e.x + radiance.y + radiance.z) > 0) && (survive)
    {
        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) * (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;

```



```
ics
& (depth < MAXDEPTH)
{
    if ( ! inside )
        return 0;
    double nt = nt / nc, ddn = ddn / nc;
    double r0s2t = 1.0f - nnt * nnt;
    double D, N );
    double a = nt - nc, b = nt + nc;
    double Tr = 1 - (R0 + (1 - R0) * r0s2t);
    double R = (D * nnt - N * (ddn > 0 ? 1 : -1));
    double E * diffuse;
    double refl + refr)) && (depth < MAXDEPTH)
    double D, N );
    double refl * E * diffuse;
    double = true;
    double MAXDEPTH)
    double survive = SurvivalProbability( diffuse );
    double estimation - doing it properly, closely following
    double if;
    double radiance = SampleLight( &rand, I, &L, &light );
    double e.x + radiance.y + radiance.z) && (abs(e.x + radiance.y + radiance.z) > 0) && (abs(e.x + radiance.y + radiance.z) > 0)
    double w = true;
    double brdfPdf = EvaluateDiffuse( L, N ) * Psurvive;
    double at3 factor = diffuse * INVPI;
    double at weight = Mis2( directPdf, brdfPdf );
    double at cosThetaOut = dot( N, L );
    double E * ((weight * cosThetaOut) / directPdf) * (radiance
    double random walk - done properly, closely following
    double survive)
    double ;
    double at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
    double survive;
    double pdf;
    double n = E * brdf * (dot( N, R ) / pdf);
    double sion = true;

```

Today's Agenda:

- Introduction
- Intel: SSE
- Streams
- Vectorization



SSE

A Brief History of SIMD

Early use of SIMD was in vector supercomputers such as the CDC Star-100 and TI ASC (image).



Intel’s MMX extension to the x86 instruction set (1996) was the first use of SIMD in commodity hardware, followed by Motorola’s AltiVec (1998), and Intel’s SSE (P3, 1999).

SSE:

- 70 assembler instructions
- Operates on 128-bit registers
- Operates on vectors of 4 floats.



SSE

SIMD Basics

C++ supports a 128-bit vector data type: `__m128`
Henceforth, we will pronounce to this as ‘quadfloat’. ☺

`__m128` literally is a small array of floats:

```
union { __m128 a4; float a[4]; };
```

Alternatively, you can use the integer variety `__m128i`:

```
union { __m128i a4; int a[4]; };
```



SSE

SIMD Basics

We operate on SSE data using *intrinsics*: in the case of SSE, these are keywords that translate to a single assembler instruction.

Examples:

```

__m128 a4 = _mm_set_ps( 1, 0, 3.141592f, 9.5f );
__m128 b4 = _mm_setzero_ps();
__m128 c4 = _mm_add_ps( a4, b4 ); // not: __m128 = a4 + b4;
__m128 d4 = _mm_sub_ps( b4, a4 );

```

Here, ‘_ps’ stands for *packed scalar*.



SSE

SIMD Basics

Other instructions:

```


__m128 c4 = _mm_div_ps( a4, b4 ); // component-wise division
__m128 d4 = _mm_sqrt_ps( a4 ); // four square roots
__m128 d4 = _mm_rcp_ps( a4 ); // four reciprocals
__m128 d4 = _mm_rsqrt_ps( a4 ); // four reciprocal square roots (!)

__m128 d4 = _mm_max_ps( a4, b4 );
__m128 d4 = _mm_min_ps( a4, b4 );
    
```

Keep the assembler-like syntax in mind:

```

__m128 d4 = dx4 * dx4 + dy4 * dy4;
__m128 d4 = _mm_add_ps(
    _mm_mul_ps( dx4, dx4 ),
    _mm_mul_ps( dy4, dy4 )
);
    
```



CODING TIME




SSE

SIMD Basics

In short:

- Four times the work at the price of a single scalar operation (if you can feed the data fast enough)
- Potentially even better performance for min, max, sqrt, rsqrt
- Requires four independent streams.

And, with AVX we get `__m256...`



```
...ics
& (depth < MAXDEPTH)
{
    if ( ! inside )
        return 0;
    int nt = nt / nc, ndn = ndn / nc;
    double r0s2t = 1.0f + nnt * nnt;
    double D, N );
    double R = (D * nnt - N * (Dn));
    double E * diffuse;
    bool = true;
    ...
    refl + refr)) && (depth < MAXDEPTH)
    double D, N );
    refl * E * diffuse;
    = true;
    ...
    MAXDEPTH)
    survive = SurvivalProbability( diffuse );
    estimation - doing it properly, closely following
    if;
    radiance = SampleLight( &rand, I, &L, &light);
    e.x + radiance.y + radiance.z) > 0) && (abs(
    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) * (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;
}
```

Today's Agenda:

- Introduction
- Intel: SSE
- Streams
- Vectorization



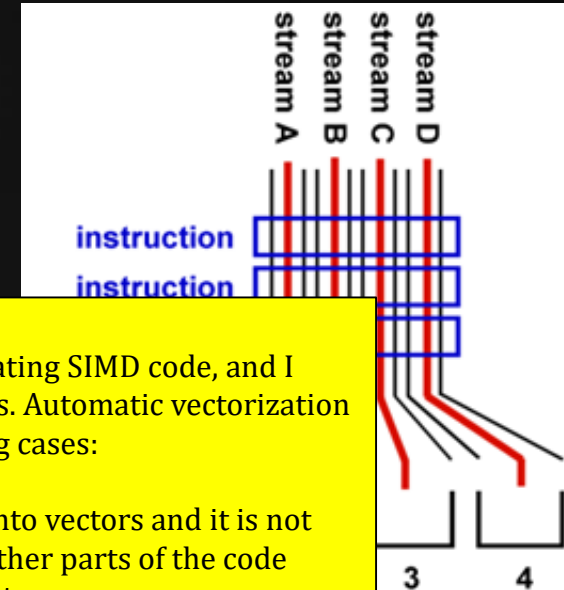
Streams

SIMD According To Visual Studio

```
vec3 A( 1, 0, 0 );
vec3 B( 0, 1, 0 );
vec3 C = (A + B) * 0.1f;
vec3 D = normalize( C );
```

The compiler will notice that we are operating on vectors, and it will use SSE instructions. This results in a modest speedup. Note that operating on vectors

To get maximum throughput, we want for each processor running in parallel.



Agner Fog:

“Automatic vectorization is the easiest way of generating SIMD code, and I would recommend to use this method when it works. Automatic vectorization may fail or produce suboptimal code in the following cases:

- when the algorithm is too complex.
- when data have to be re-arranged in order to fit into vectors and it is not obvious to the compiler how to do this or when other parts of the code needs to be changed to handle the re-arranged data.
- when it is not known to the compiler which data sets are bigger or smaller than the vector size.
- when it is not known to the compiler whether the size of a data set is a multiple of the vector size or not.
- when the algorithm involves calls to functions that are defined elsewhere or cannot be inlined and which are not readily available in vector versions.
- when the algorithm involves many branches that are not easily vectorized.
- when floating point operations have to be reordered or transformed and it is not known to the compiler whether these transformations are permissible with respect to precision, overflow, etc.
- when functions are implemented with lookup tables.



Streams


SIMD According To Visual Studio

```

float Ax = 1, Ay = 0, Az = 0;
float BX = 0, By = 1, Bz = 0;
float Cx = (Ax + Bx) * 0.1f;
float Cy = (Ay + By) * 0.1f;
float Cz = (Az + Bz) * 0.1f;
float l = sqrtf( Cx * Cx + Cy * Cy + Cz * Cz);
float Dx = Cx / l;
float Dy = Cy / l;
float Dz = Cz / l;
    
```

```

float Ax = 1, Ay = 0, Az = 0;
float BX = 0, By = 1, Bz = 0;
float Cx = (Ax + Bx) * 0.1f;
float Cy = (Ay + By) * 0.1f;
float Cz = (Az + Bz) * 0.1f;
float l = sqrtf( Cx * Cx + Cy * Cy + Cz * Cz);
float Dx = Cx / l;
float Dy = Cy / l;
float Dz = Cz / l;
    
```



Streams

SIMD According To Visual Studio

```


...ics
& (depth < MAXDEPTH)
...
= inside ? 1.0f : 0.0f;
nt = nt / nc; ddn = ddn * ddn;
ps2t = 1.0f - nnt * nnt;
D, N );
)
...
at a = nt - nc; b = nt + nc;
at Tr = 1 - (R0 + (1 - R0) * ps2t);
Tr) R = (D * nnt - N * (ddn * nnt + ps2t));
...
E * diffuse;
= true;
...
efl + refr)) && (depth < MAXDEPTH)
D, N );
refl * E * diffuse;
= true;
...
MAXDEPTH)
...
survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely following
if;
radiance = SampleLight( &rand, I, &L, &light );
e.x + radiance.y + radiance.z) > 0) && (rand < survive);
...
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) * (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;

```

```

float Ax[4] = {...}, Ay[4] = {...}, Az[4] = {...};
float Bx[4] = {...}, By[4] = {...}, Bz[4] = {...};
float Cx[4] = ...;
float Cy[4] = ...;
float Cz[4] = ...;
float l[4] = ...;
float Dx[4] = ...;
float Dy[4] = ...;
float Dz[4] = ...;

```




Streams

SIMD According To Visual Studio

```

__m128 Ax4 = {...}, Ay4 = {...}, Az4 = {...};
__m128 Bx4 = {...}, By4 = {...}, Bz4 = {...};
__m128 Cx4 = ...;
__m128 Cy4 = ...;
__m128 Cz4 = ...;
__m128 l4 = ...;
__m128 Dx4 = ...;
__m128 Dy4 = ...;
__m128 Dz4 = ...;

```



Streams

SIMD According To Visual Studio

```


...ics
& (depth < MAXDEPTH)
...
= inside ? 1.0f : 0.0f;
nt = nt / nc; add = add / nc;
ps2t = 1.0f - nnt * nnt;
D, N );
...
at a = nt - nc, b = nt - nc;
at Tr = 1 - (R0 + (1 - R0) * r);
Tr) R = (D * nnt - N * (add
...
E * diffuse;
= true;
...
efl + refr)) && (depth < MAXDEPTH)
D, N );
refl * E * diffuse;
= true;
...
MAXDEPTH)
survive = SurvivalProbability( diffuse );
estimation - doing it properly, closely following
if;
radiance = SampleLight( &rand, I, &L, &light, &N, &D, &N );
e.x + radiance.y + radiance.z) > 0) && (add
...
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;

```

```

__m128 Ax4 = {...}, Ay4 = {...}, Az4 = {...};
__m128 Bx4 = {...}, By4 = {...}, Bz4 = {...};
__m128 X4 = _mm_set1_ps( 0.1f );
__m128 Cx4 = _mm_mul_ps( _mm_add_ps( Ax4, Bx4 ), X4 );
__m128 Cy4 = _mm_mul_ps( _mm_add_ps( Ay4, By4 ), X4 );
__m128 Cz4 = _mm_mul_ps( _mm_add_ps( Az4, Bz4 ), X4 );
__m128 l4 = ...;
__m128 Dx4 = ...;
__m128 Dy4 = ...;
__m128 Dz4 = ...;

```



Streams

SIMD According To Visual Studio

```


...
    & (depth < MAXDEPTH)
...
    inside ? 1 : 0;
    nt = nt / nc;
    ps2t = 1.0f - nnt;
    D, N );
    )
...
    at a = nt - nc, b = nt - nc;
    at Tr = 1 - (R0 + (1 - R0)
    Tr) R = (D * nnt - N * (D0
...
    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) && (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 3e31
    (vive)
...
    at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf
    survive;
    pdf;
    n = E * brdf * (dot( N, R ) / pdf);
    sion = true;

```

```

__m128 Ax4 = _mm_set_ps( Ax[0], Ax[1], Ax[2], Ax[3] );
__m128 Ay4 = _mm_set_ps( Ay[0], Ay[1], Ay[2], Ay[3] );
__m128 Az4 = _mm_set_ps( Az[0], Az[1], Az[2], Az[3] );
__m128 Bx4 = {...}, By4 = {...}, Bz4 = {...};
__m128 X4 = _mm_set1_ps( 0.1f );
__m128 Cx4 = _mm_mul_ps( _mm_add_ps( Ax4, Bx4 ), X4 );
__m128 Cy4 = _mm_mul_ps( _mm_add_ps( Ay4, By4 ), X4 );
__m128 Cz4 = _mm_mul_ps( _mm_add_ps( Az4, Bz4 ), X4 );
__m128 l4 = ...;
__m128 Dx4 = ...;
__m128 Dy4 = ...;
__m128 Dz4 = ...;

```



Streams


SIMD Friendly Data Layout

Consider the following data structure:

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

AoS

SoA



```
ics
& (depth < MAXDEPTH)
= inside ? 1 : 0;
nt = nt / nc; ddn = sqrt(
ps2t = 1.0f - nnt * nnt;
D, N );
)
at a = nt - nc, b = nt + nc;
at Tr = 1 - (R0 + (1 - R0)
Tr) R = (D * nnt - N * (dd
E * diffuse;
= true;
efl + refr)) && (depth < MAXDEPTH)
D, N );
refl * E * diffuse;
= true;
MAXDEPTH)
survive = SurvivalProbability( diffuse
estimation - doing it
if;
radiance = SampleLight( R, r1, r2, &R, &pdf );
e.x + radiance.y + radiance.z;
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 S
ive)
;
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf );
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
sion = true;
```

```
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]; };
```



Streams

SIMD Data Naming Conventions

```

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


Notice that SoA is breaking our OO...

Consider adding the struct name to the variables:

```
float particle_x[512];
```

Or put an amount of particles in a struct.

Also note the convention of adding ‘4’ to any SSE variable.



CODING TIME



Today's Agenda:

- Introduction
- Intel: SSE
- Streams
- Vectorization



Vectorization

Converting your Code

1. Locate a significant bottleneck in your code
(converting is going to be labor-intensive, be sure it's worth it)
2. Keep a copy of the original code (use #ifdef)
(you may want to compile on some other platform later)
3. Prepare the scalar code
(add a 'for(int stream = 0; stream < 4; stream++)' loop)
4. Reorganize the data
(make sure you don't have to convert all the time)
5. Union with floats
6. Convert one line at a time, verifying functionality as you go
7. Check MSDN for exotic SSE instructions
(some odd instructions exist that may help your problem)



/INFOMOV/

END of “SIMD (1)”

next lecture: “SIMD (2)”

```
ics
& (depth < MAXDEPTH)
{
    if (inside ? 1 : 0)
    {
        nt = nt / nc; ddn = ddn * n;
        cos2t = 1.0f - nnt * nnt;
        D, N );
    }

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

    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, &L, &light);
e.x + radiance.y + radiance.z) > 0) && (cosThetaOut > 0)
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) * (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;
```

