

/INFOMOV/

Optimization & Vectorization

J. Bikker - Sep-Nov 2019 - Lecture 9: “GPGPU (1)”

Welcome!



Global Agenda:

1. GPGPU (1) : Introduction, architecture, concepts
2. GPGPU (2) : Practical Code using GPGPU
3. GPGPU (3) : Parallel Algorithms, Optimizing for GPU

Today's Agenda:

- Introduction to GPGPU
- Example: Voronoi Noise
- GPGPU Programming Model
- OpenCL Template



```
rics  
    & (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 );  
if (e.x + radiance.y + radiance.z) > 0) && (dot( N, L ) > 0)  
    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;  
E * brdf * (dot( N, R ) / pdf);  
ision = true;
```

“If you were plowing a field, which would you rather use? Two strong oxen, or 1024 chickens?”

- Seymour Cray



Introduction

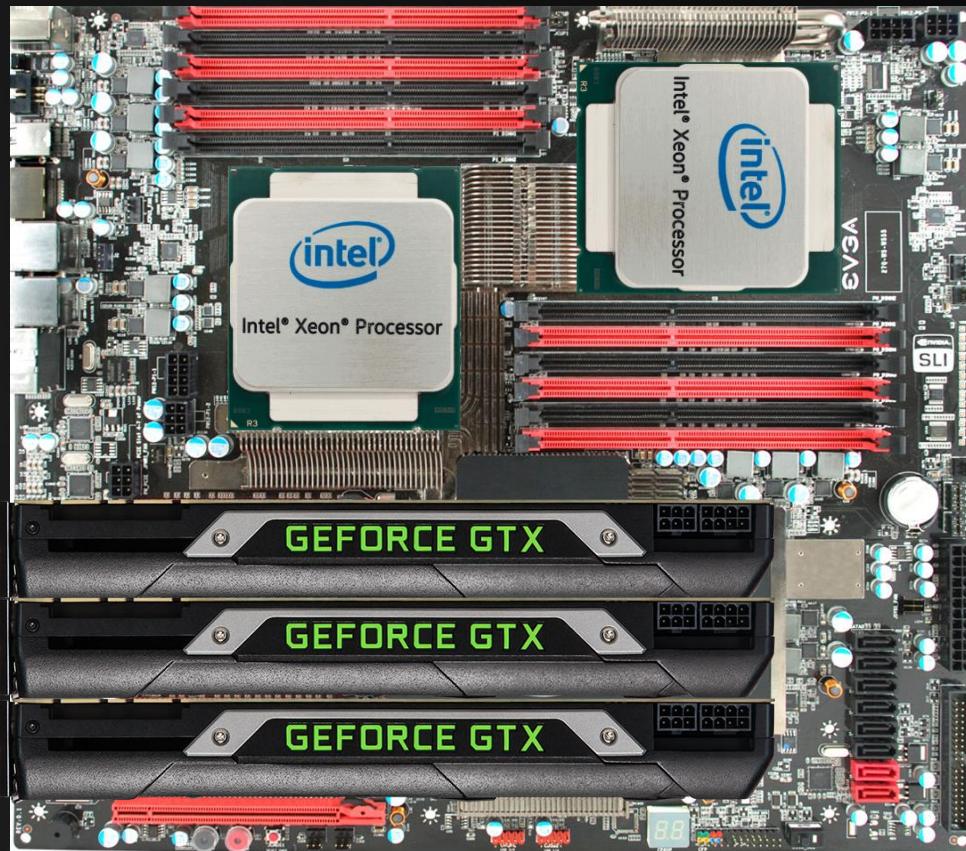
Heterogeneous Processing

The average computer contains:

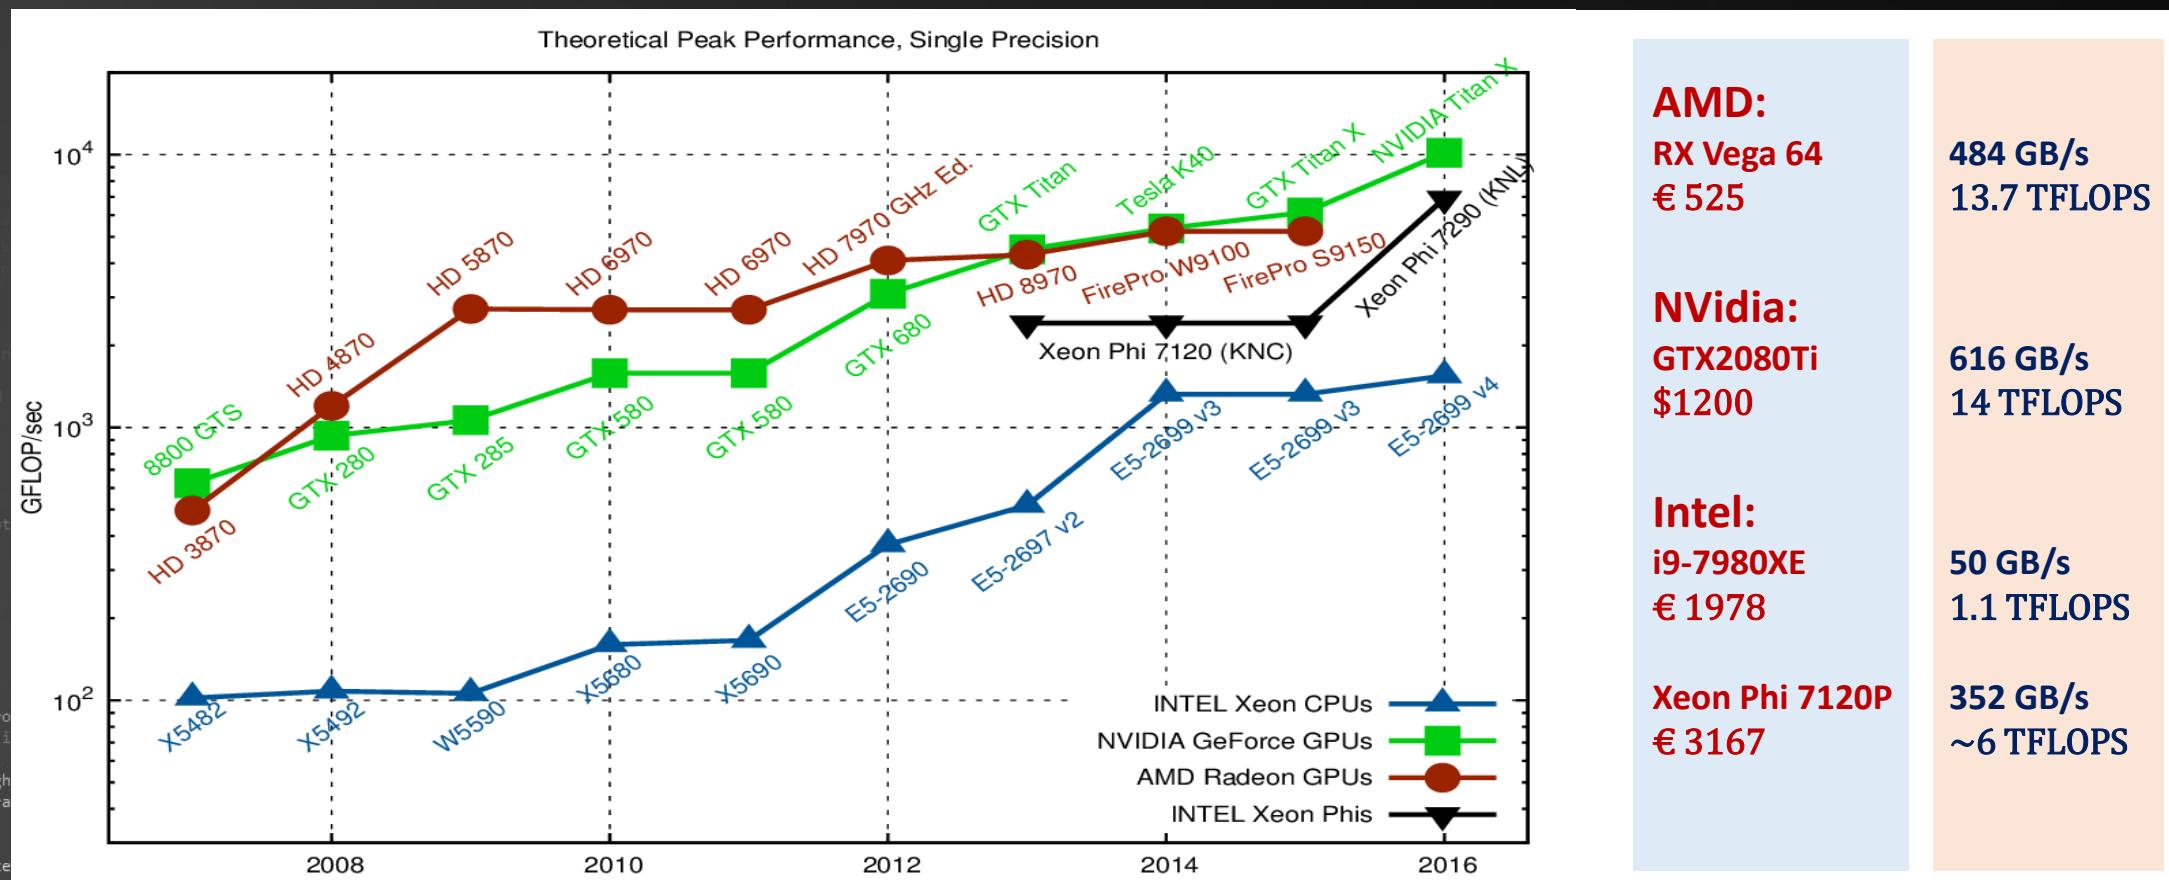
- 1 or more CPUs;
- 1 or more GPUs.

We have been optimizing CPU code.
A vast source of compute power has remained unused:

The Graphics Processing Unit.



Introduction



Introduction

A Brief History of GPGPU

```
rics
  & (depth < MAXDEPTH)
  n = 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;

at 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, &light,
e.x + radiance.y + radiance.z) > 0) && (dot( N,
e = 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;
R = E * brdf * (dot( N, R ) / pdf);
ision = true;
```



Introduction

A Brief History of GPGPU

```
rics
  & (depth < MAXDEPTH)
  c = inside ? 1 : 1.2f;
  nt = nc / ncc; 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;

-
at 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);
e.x + radiance.y + radiance.z ) > 0) && (dot( N,
e.y + radiance.z ) > 0);
e.x + radiance.y + radiance.z ) > 0) && (dot( N,
e.y + radiance.z ) > 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;
```



Introduction

NVidia NV-1
(Diamond Edge 3D)
1995



A Brief History of GPGPU

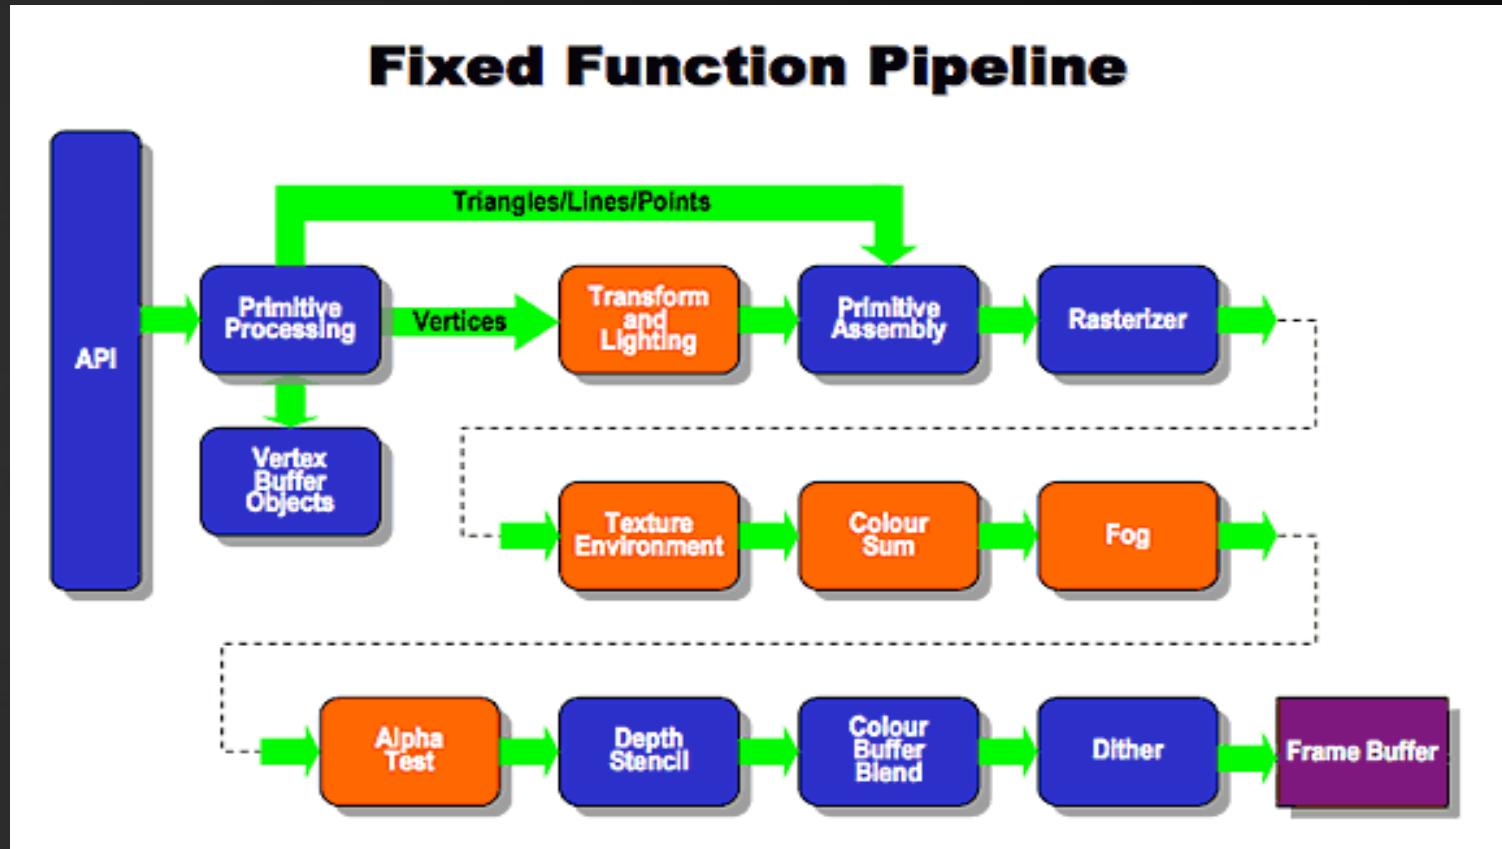


3Dfx -
Diamond Monster 3D
1996



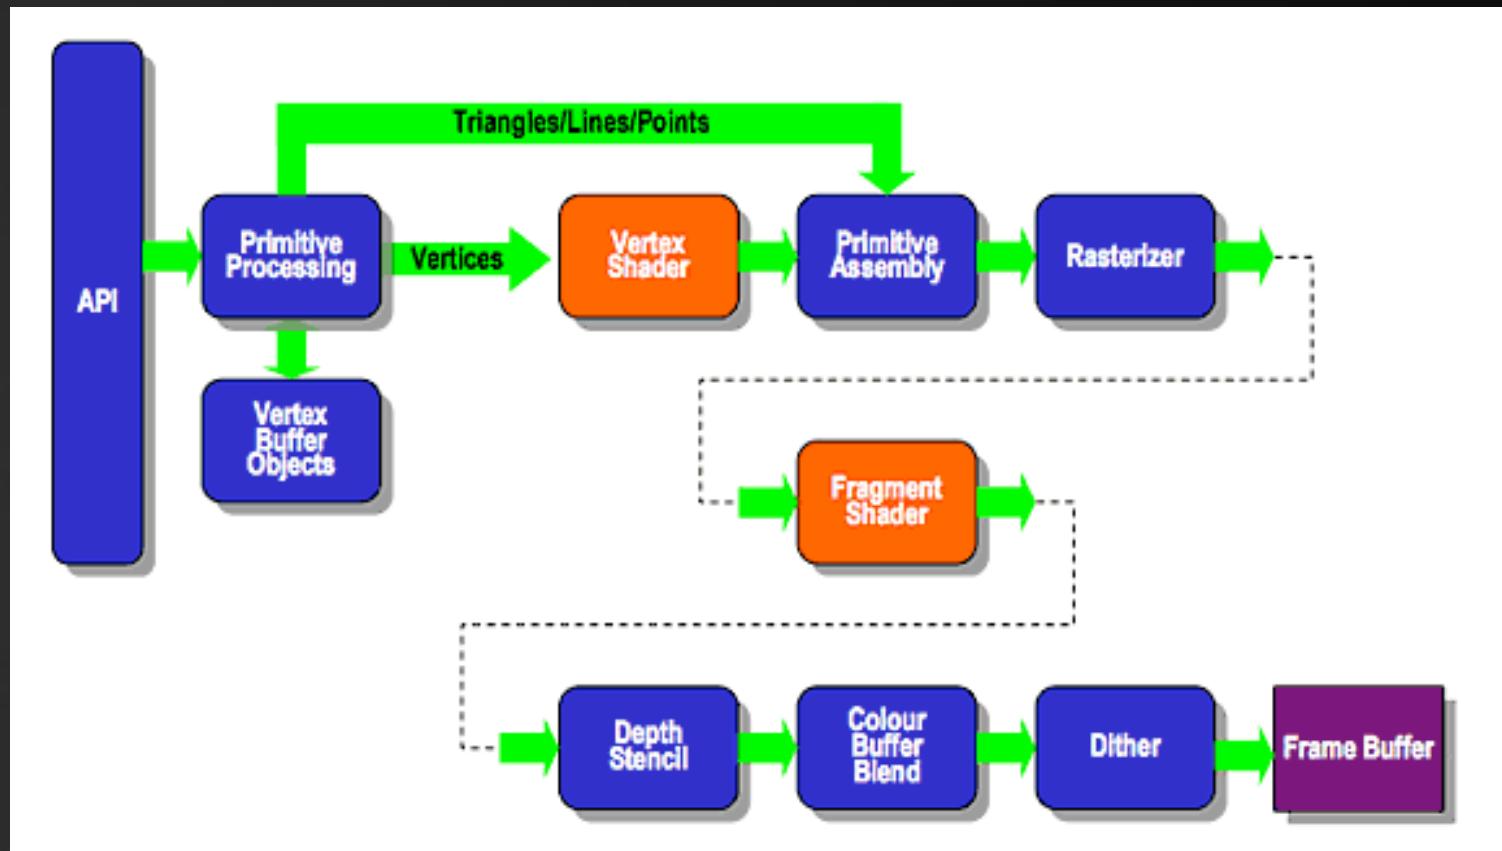
Introduction

A Brief History of GPGPU



Introduction

A Brief History of GPGPU



Introduction

A Brief History of GPGPU

```

rics
3 (depth < MAXDEPTH)
at = inside ? 1 : 1/2;
nt = nt / nc, ddn = ddc;
os2t = 1.0f - nnt * nnt;
D, N );
)
at a = nt - nc, b = nt
at Tr = 1 - (R0 + (1 - Tr) R = (D * nnt - N *
E * diffuse;
= true;
.
.
.
at refl + refr)) && (depth
D, N );
at refl * E * diffuse;
= true;
MAXDEPTH)
survive = SurvivalProbab
estimation - doing it p
if;
radiance = SampleLight(
e.x + radiance.y + radia
.
.
.
= true;
at brdfPpdf = EvaluateDiff
at3 factor = diffuse * I
at weight = Mis2( direct
at cosThetaOut = dot( N,
E * ((weight * cosTheta
random walk - done proper
ive)

;
at3 brdf = SampleDiffuse( diffuse, N, m1, m2, &R, &pdf );
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
ision = true;

```

Ok
load"cas:GUSANO"
Skip :RANA
Found:GUSANO
Ok
list 1-100
65 OPEN "GRP:" FOR OUTPUT AS#1
70 STOP ON:ON STOP GOSUB 4000
71 COLOR15,15,15:SCREEN2
72 FORF=1TO8
73 READA:S\$\$=S\$\$+CHR\$(A)
74 NEXTF
75 SPRITE\$(1)=S\$
76 COLOR15,4,13:S\$\$=""
80 SOUND 7,255:SOUND8,15:RR=25
90 SCREEN2,2:PRESET(130,5):PRINT #1,"
MAX. :";PEEK(-1200)
100 DIMX3(500),Y3(500):PSET(35,5),4:P
RINT#1,"PUNTOS: 0":PSET(200,5),4:PRIN
T#1,"@":PSET(212,5),4:PRINT#1,"@":XB=
200:Q=1:RR=25
Ok
color auto goto list run

POTATO SALAD

Some good cooks sprinkle grated pimiento cheese on this

4 cups diced cooked potatoes
1 cup sliced celery
3 hard-cooked eggs, cut up
½ cup finely cut onion or sliced green onions
¼ cup sliced radishes
1 cup mayonnaise
1 tablespoon vinegar
1 teaspoon prepared mustard
1½ to 2 teaspoons salt
⅛ teaspoon pepper
Lettuce

Mix all the ingredients in a bowl. Cover and refrigerate several hours so flavors can blend. Serve on crisp lettuce. Makes 6 servings.



Introduction

A Brief History of GPGPU

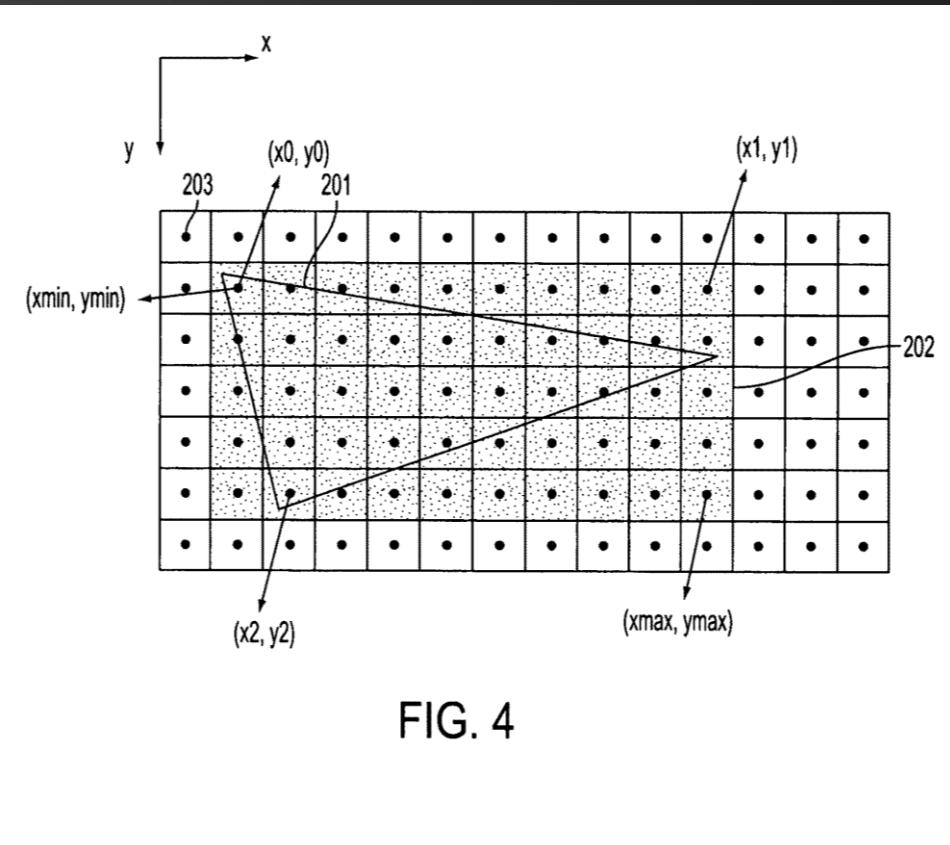


FIG. 4

GPU - conveyor belt:

input = vertices + connectivity

step 1: transform

step 2: rasterize

step 3: shade

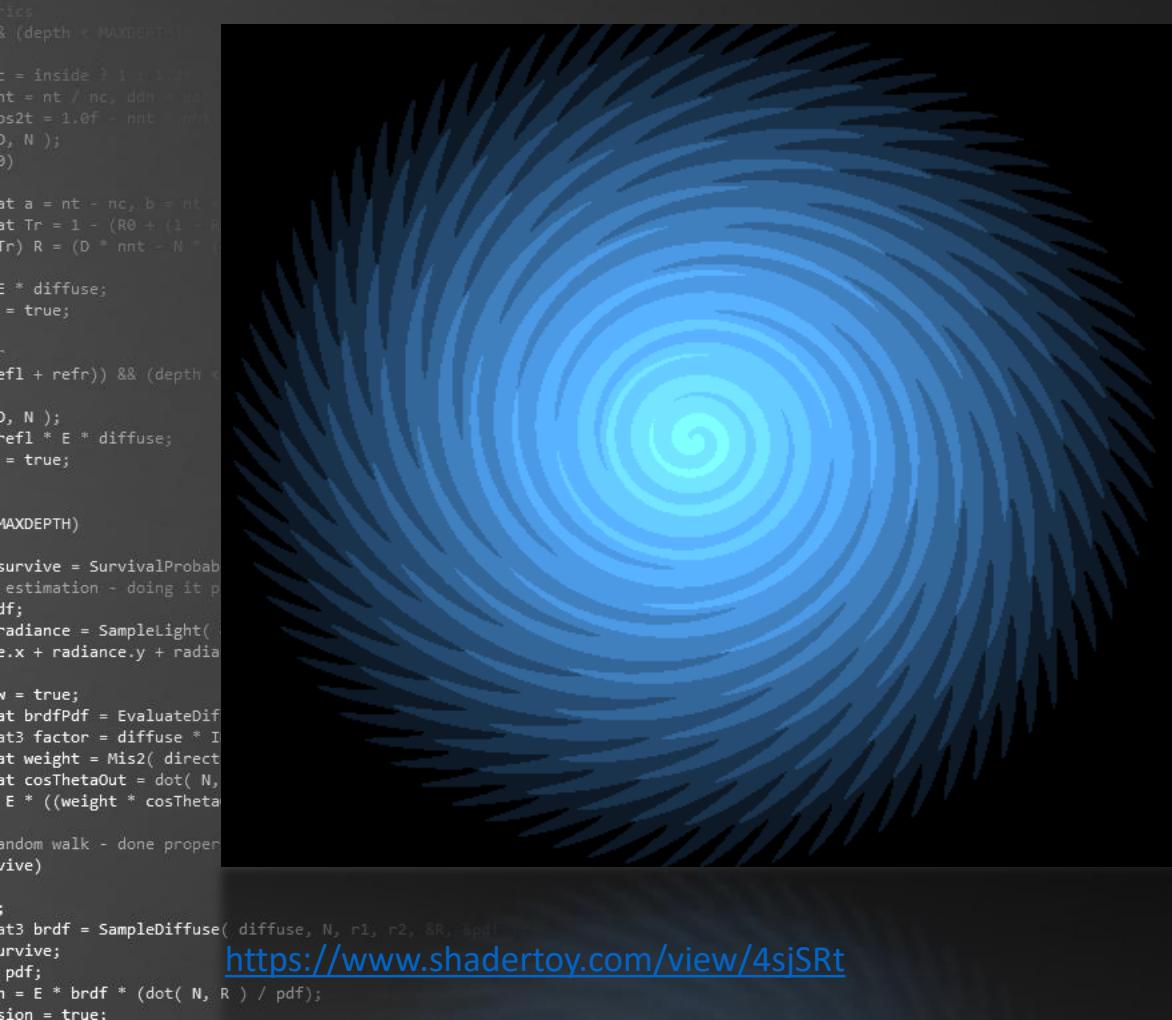
step 4: z-test

output = pixels



Introduction

A Brief History of GPGPU



<https://www.shadertoy.com/view/4sjSRt>

```
void main(void)
{
    float t = iGlobalTime;
    vec2 uv = gl_FragCoord.xy / iResolution.y;
    float r = length(uv), a = atan(uv.y,uv.x);
    float i = floor(r*10);
    a *= floor(pow(128,i/10));
    a += 20.*sin(0.5*t)+123.34*i-100.*  

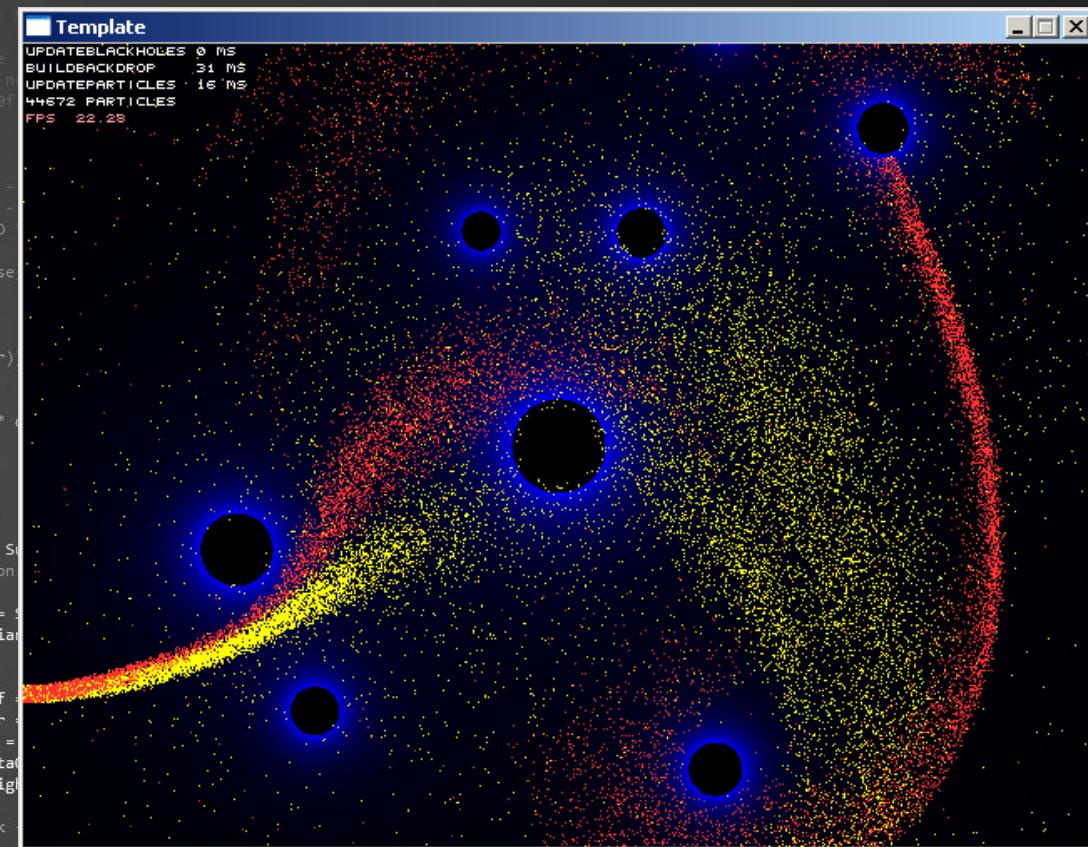
        (r*i/10)*cos(0.5*t);
    r += (0.5+0.5*cos(a)) / 10;
    r = floor(N*r)/10;
    gl_FragColor = (1-r)*vec4(0.5,1,1.5,1);
}
```

GLSL ES code



Introduction

A Brief History of GPGPU

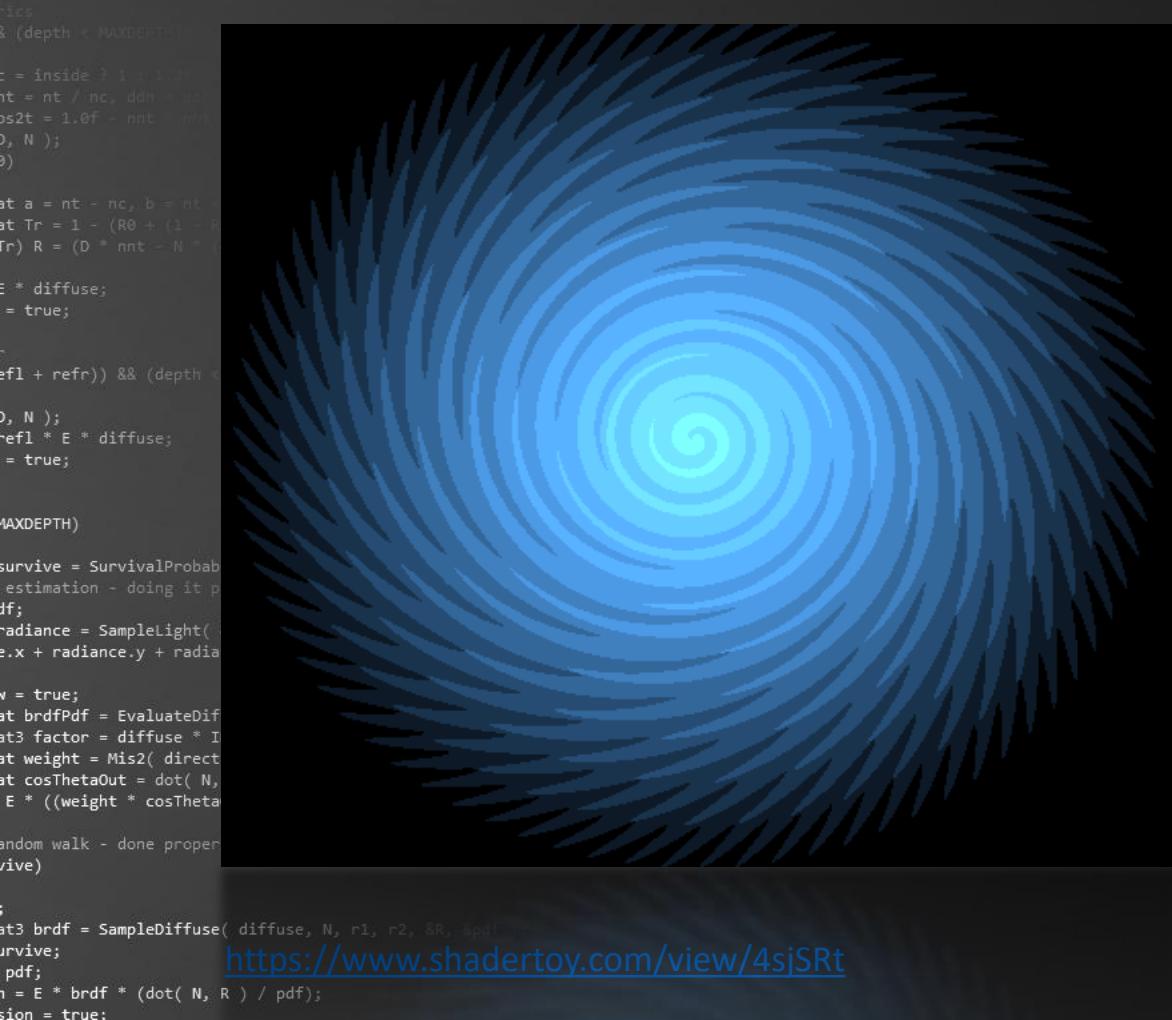


```
void Game::BuildBackdrop()
{
    Pixel* dst = m_Surface->GetBuffer();
    float fy = 0;
    for ( unsigned int y = 0; y < SCRHEIGHT; y++ )
    {
        float fx = 0;
        for ( unsigned int x = 0; x < SCRWIDTH; x++ )
        {
            float g = 0;
            for ( unsigned int i = 0; i < HOLES; i++ )
            {
                float dx = m_Hole[i]->x - fx, dy =
                float squareddist = ( dx * dx + dy *
                g += (250.0f * m_Hole[i]->g) / squ
            }
            if ( g > 1 ) g = 0;
            *dst++ = (int)(g * 255.0f);
        }
    }
}
```



Introduction

A Brief History of GPGPU



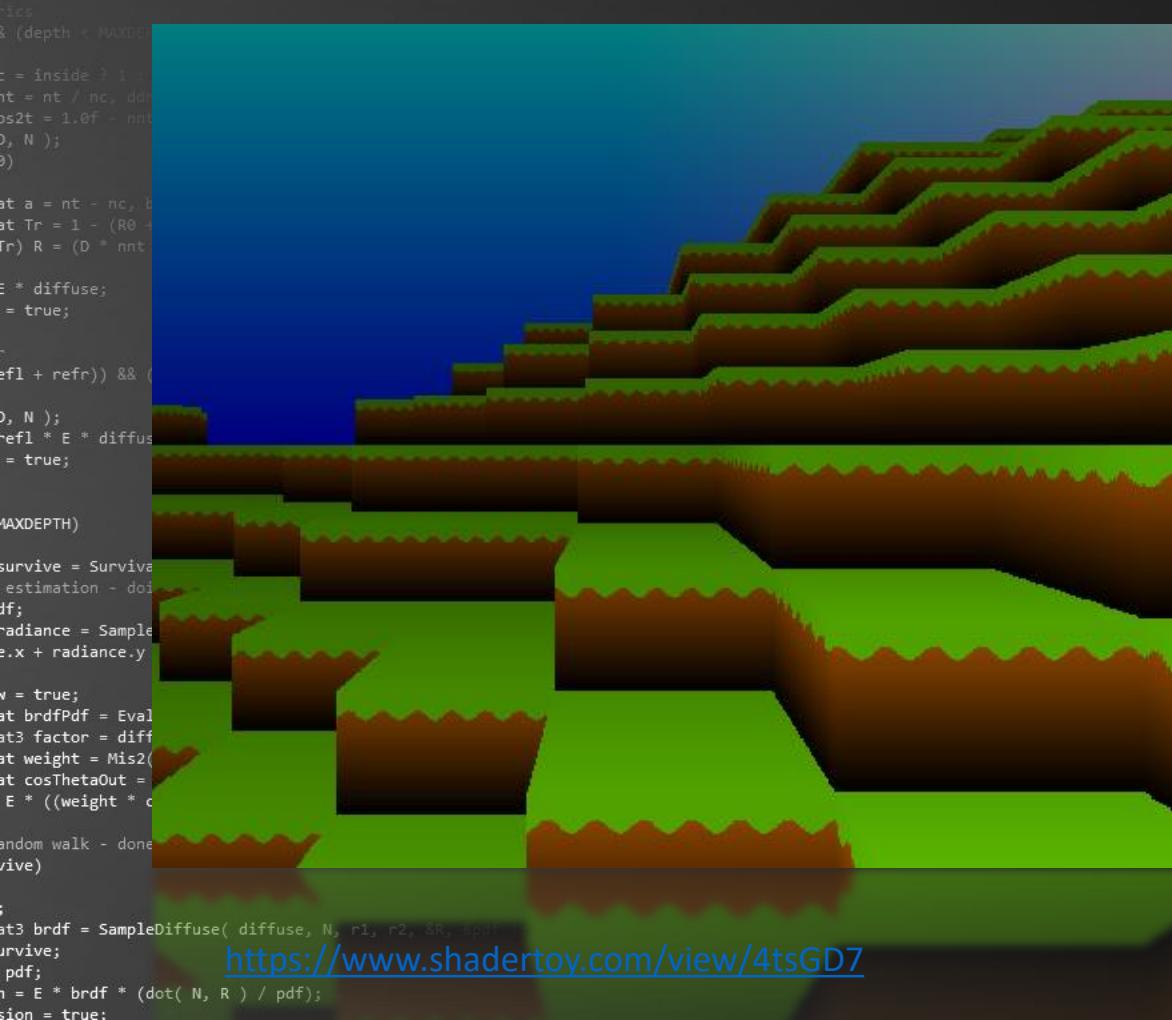
```
void main(void)
{
    float t = iGlobalTime;
    vec2 uv = gl_FragCoord.xy / iResolution.y;
    float r = length(uv), a = atan(uv.y,uv.x);
    float i = floor(r*10);
    a *= floor(pow(128,i/10));
    a += 20.*sin(0.5*t)+123.34*i-100. *
        (r*i/10)*cos(0.5*t);
    r += (0.5+0.5*cos(a)) / 10;
    r = floor(N*r)/10;
    gl_FragColor = (1-r)*vec4(0.5,1,1.5,1);
}
```

GLSL ES code



Introduction

A Brief History of GPGPU

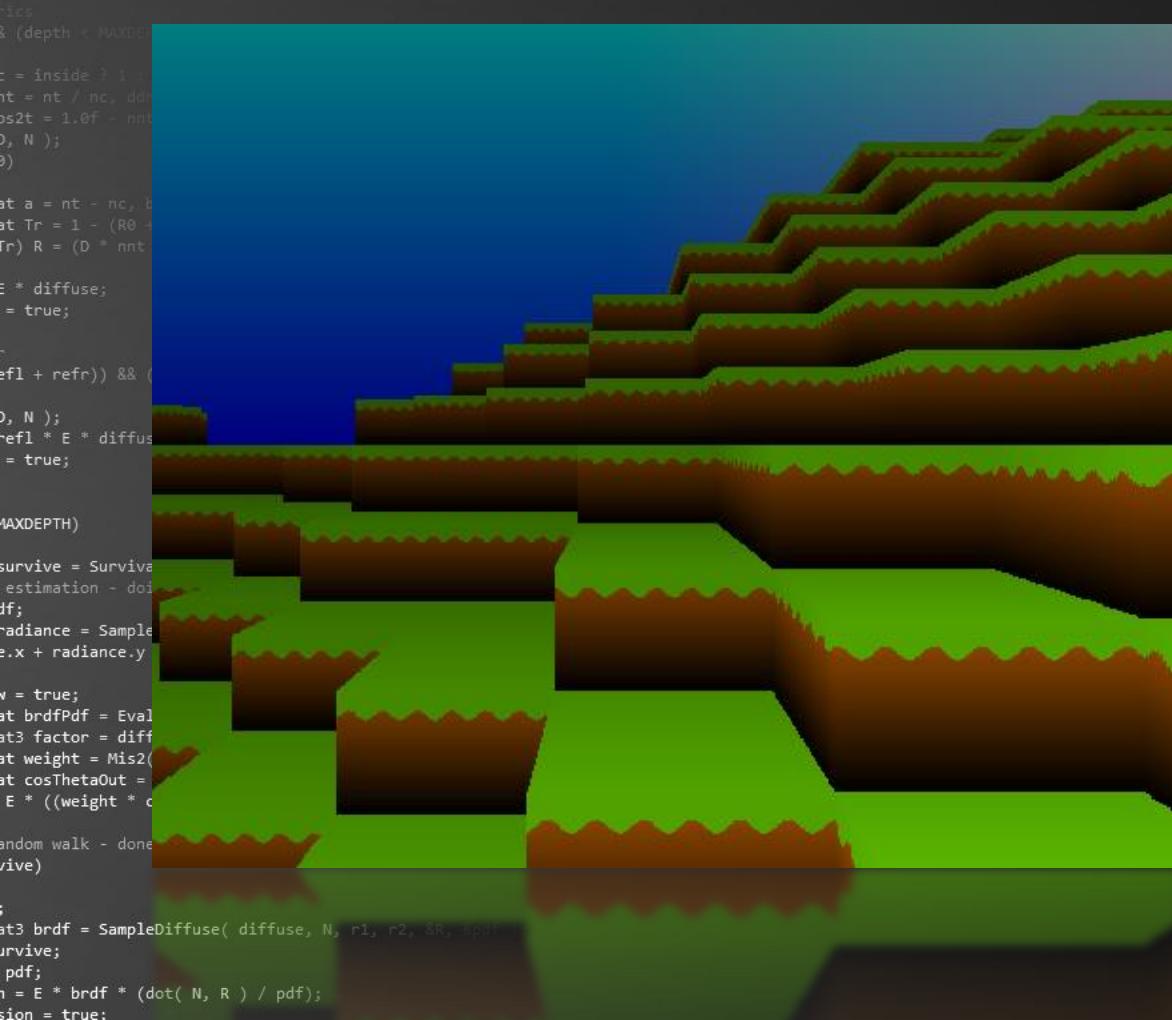


```
void mainImage( out vec4 z, in vec2 w ) {
    vec3 d = vec3(w,1)/iResolution-.5, p, c, f;
    vec3 g = d, o, y = vec3( 1,2,0 );
    o.y = 3. * cos((o.x=.3)*(o.z = iDate.w));
    for( float i=.0; i<9.; i+=.01 ) {
        f = fract(c = o += d*i*.01),
        p = floor( c )*.3;
        if( cos(p.z) + sin(p.x) > ++p.y ) {
            g = (f.y - .04*cos((c.x+c.z)*40.)>.8?y:
                f.y * y.yxz) / i;
            break;
        } } z.xyz = g; }
```

GLSL ES code

Introduction

A Brief History of GPGPU



GPUs perform well because they have a constrained execution model, based on massive parallelism.

CPU: Designed to run one thread as fast as possible.

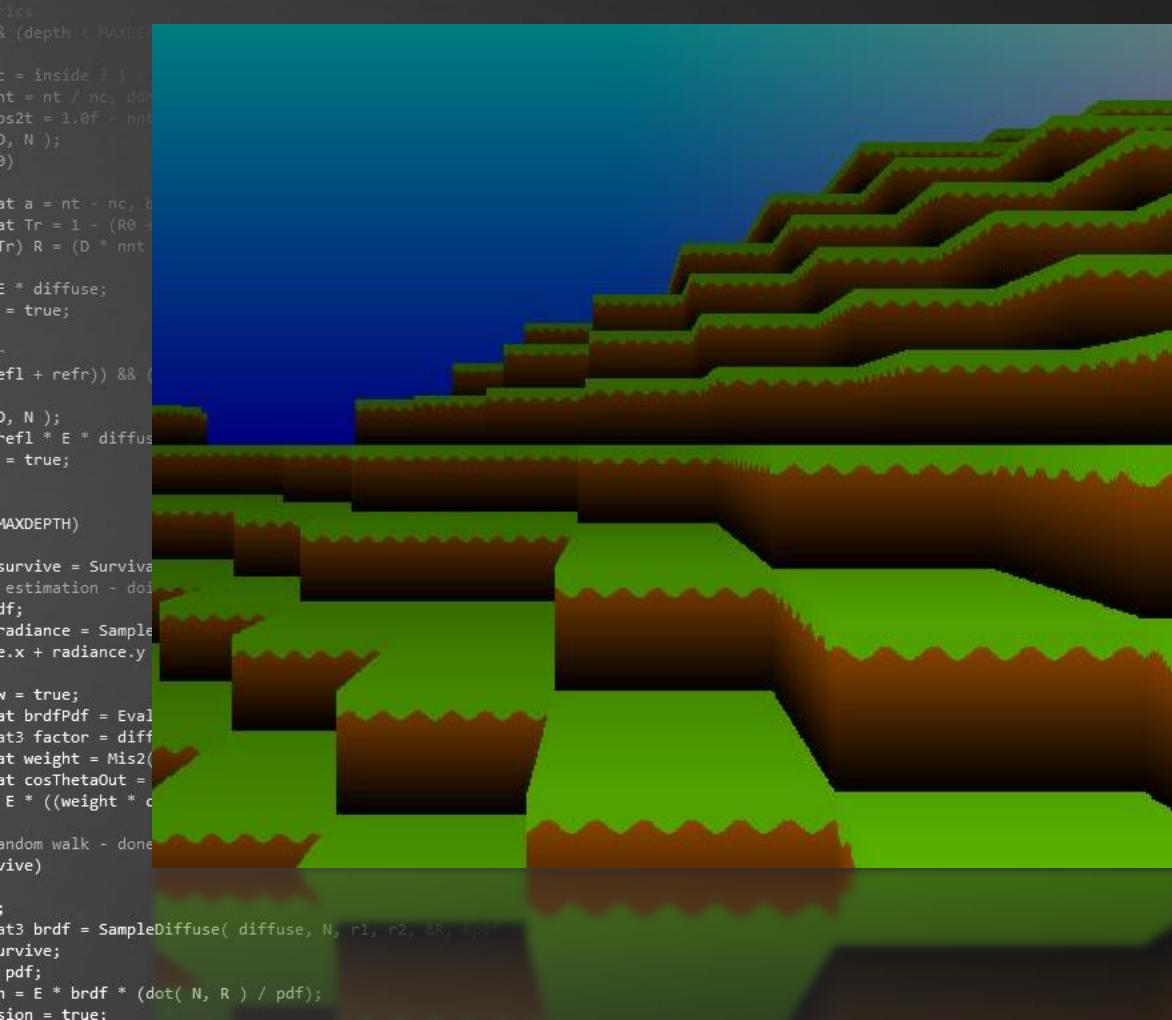
- Use caches to minimize memory latency
- Use pipelines and branch prediction
- Multi-core processing: *task parallelism*

Tricks:

- SIMD
- “Hyperthreading”

Introduction

A Brief History of GPGPU



GPUs perform well because they have a constrained execution model, based on massive parallelism.

GPU: Designed to combat latency using many threads.

- Hide latency by computation
- Maximize parallelism
- Streaming processing → Data parallelism → SIMD

Tricks:

- Use typical GPU hardware (filtering etc.)
- Cache anyway



Introduction

GPU Architecture

```

rics
& (depth < MAXDEPTH)
    n = inside ? 1 : 1.2f;
    nt = nc / nc, ddn = ddc;
    pos2t = 1.0f - nnt * nnt;
    R, 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)

```

CPU

- Multiple tasks = multiple threads
- Tasks run different instructions
- 10s of complex threads execute on a few cores
- Thread execution managed explicitly

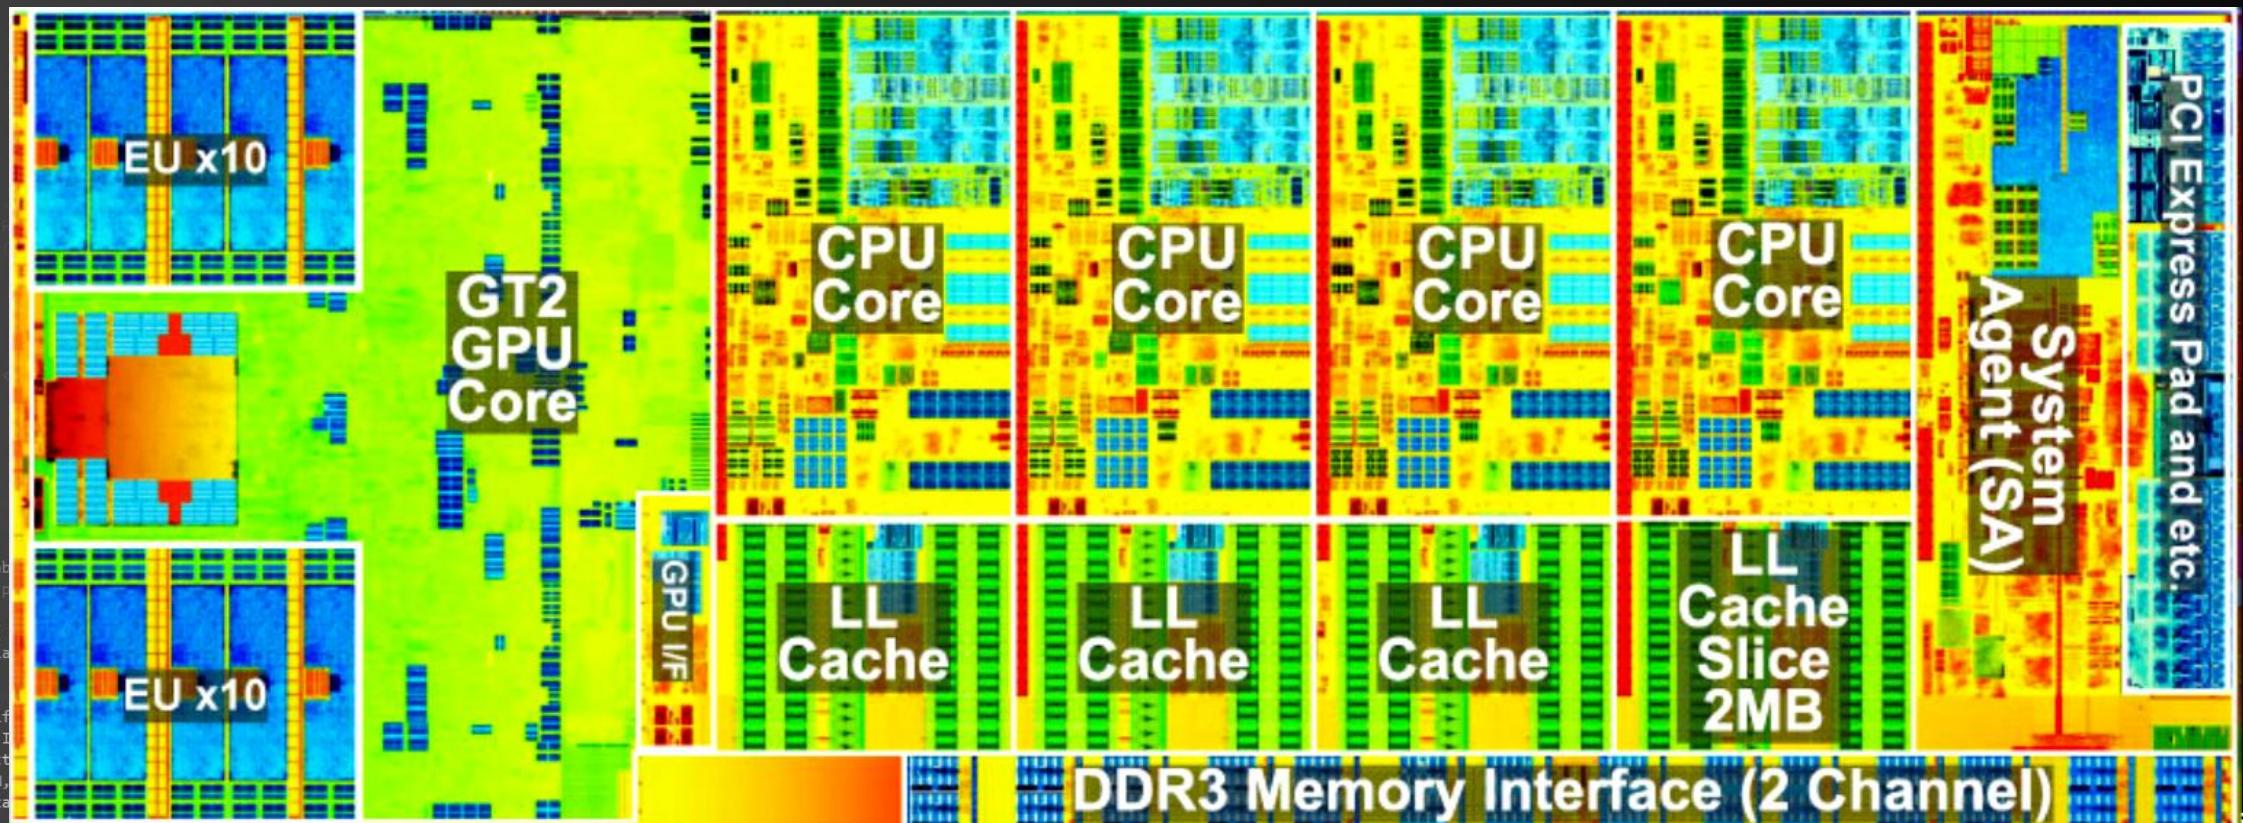
GPU

- SIMD: same instructions on multiple data
- 10.000s of light-weight threads on 100s of cores
- Threads are managed and scheduled by hardware



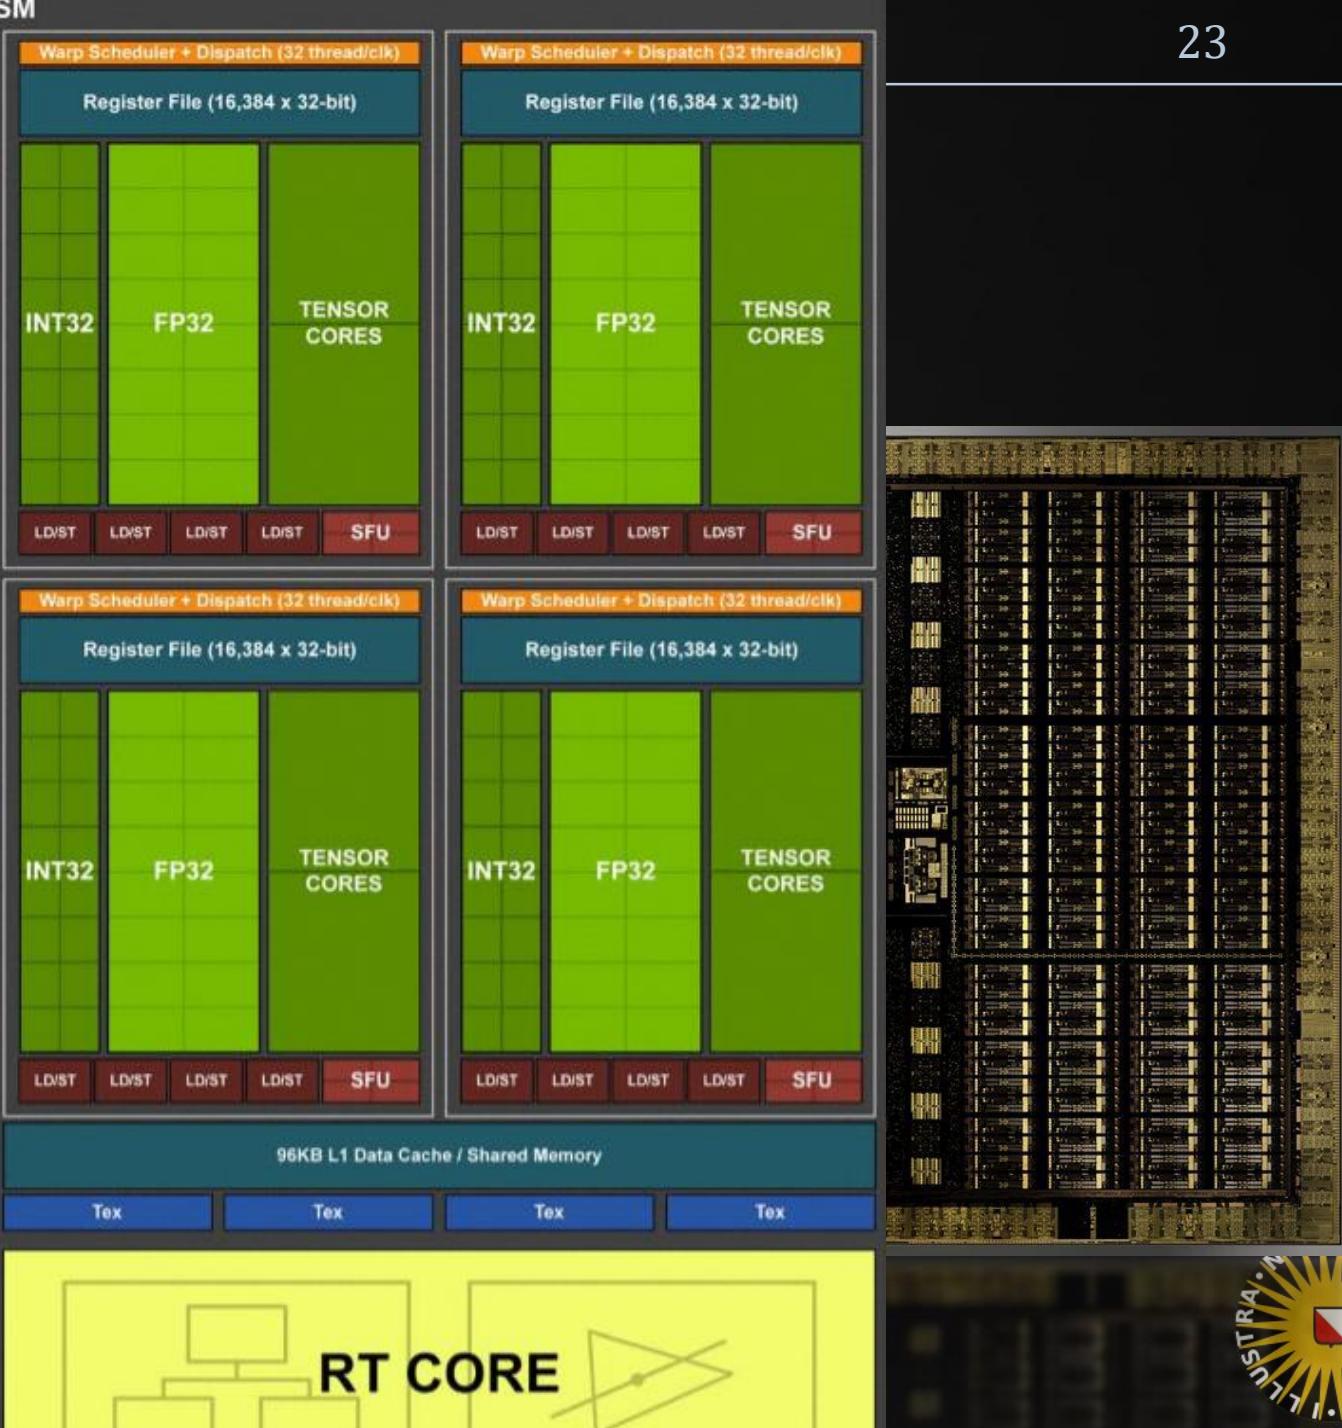
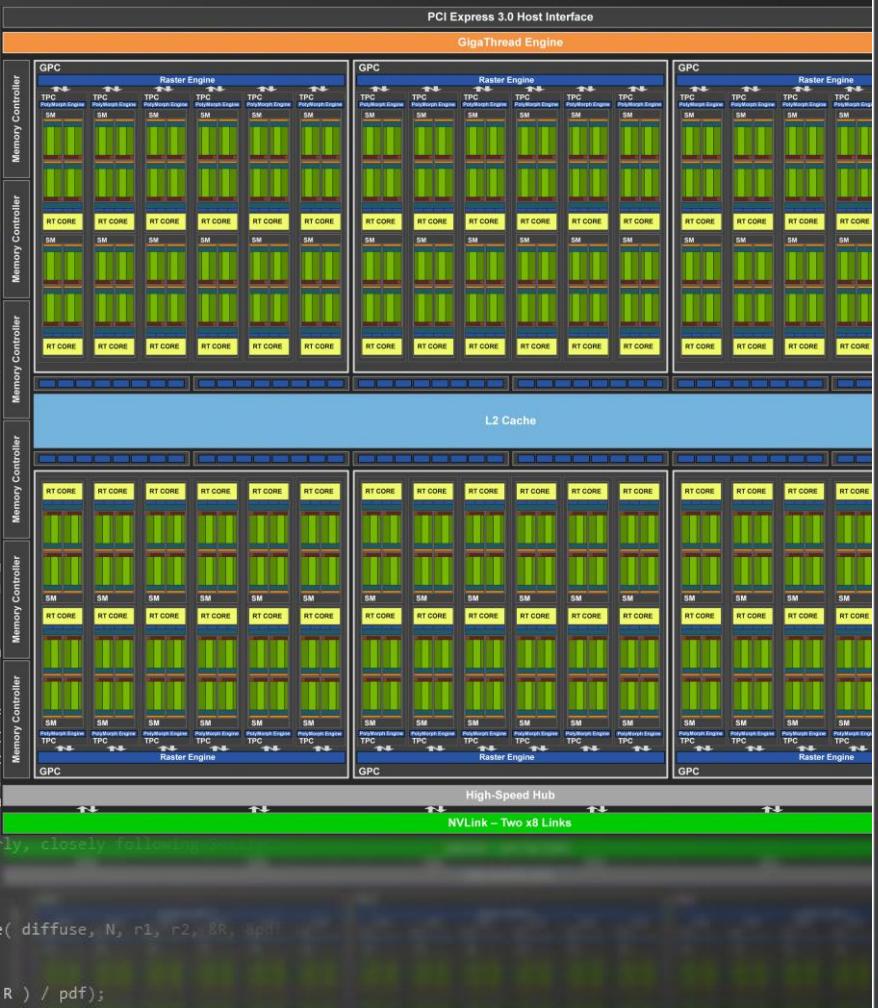
Introduction

CPU Architecture...



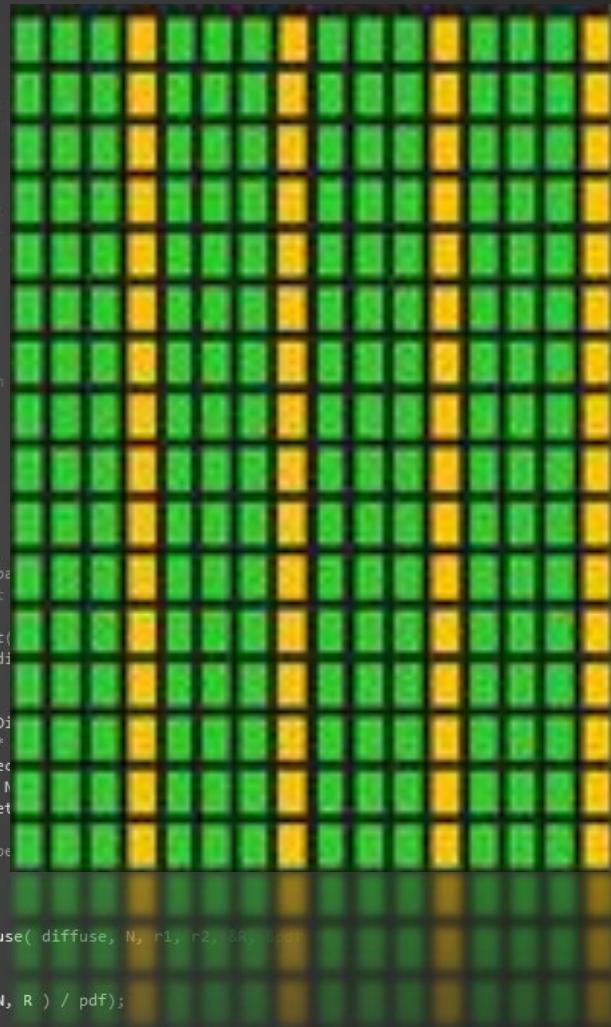
Introduction

versus GPU Architecture:



Introduction

GPU Architecture

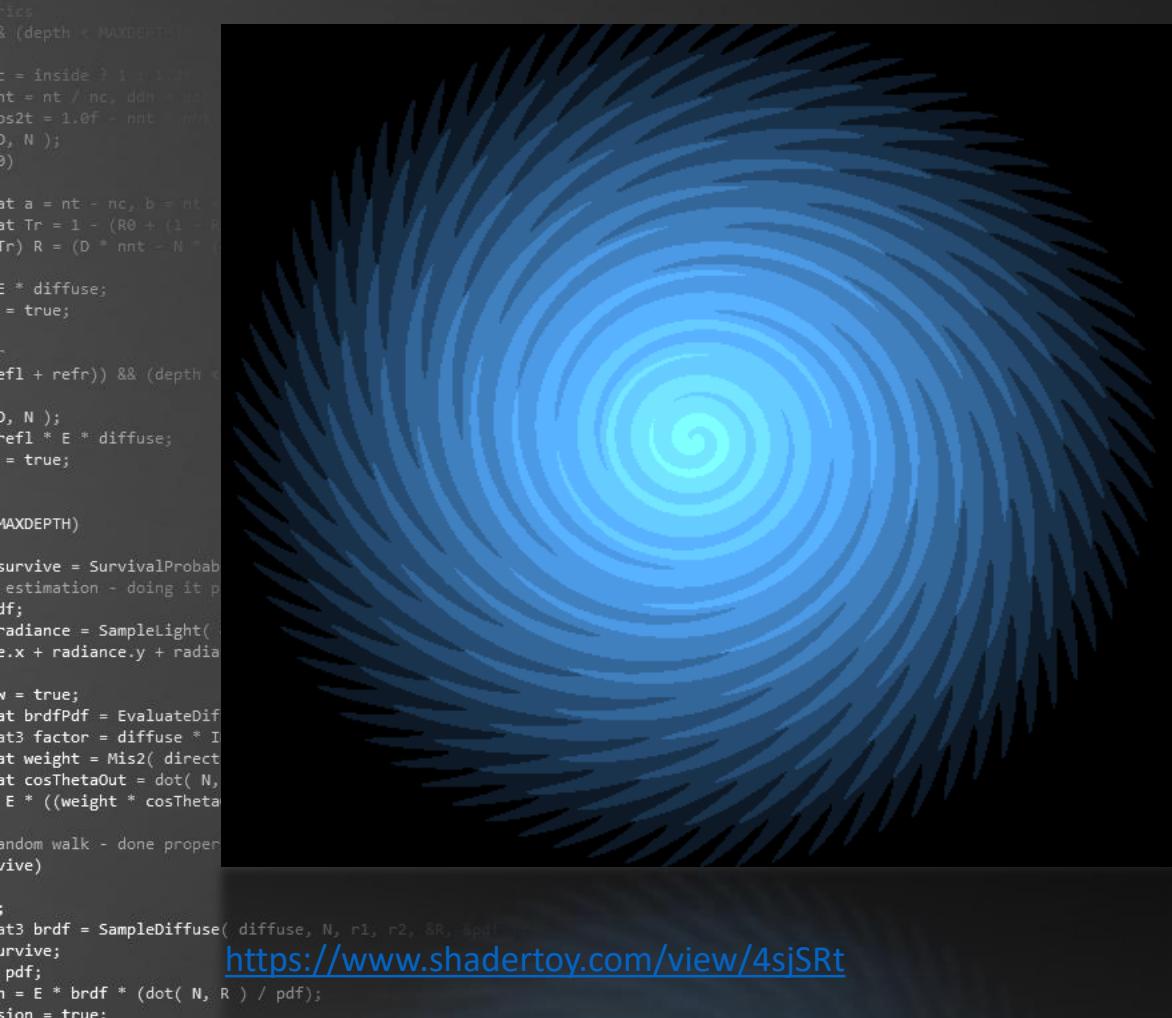


SIMT Thread execution:

- Group 32 threads (vertices, pixels, primitives) into *warps*
- Each warp executes the same instruction
- In case of latency, switch to different warp (thus: switch out 32 threads for 32 different threads)
- Flow control: ...

Introduction

GPGPU Programming



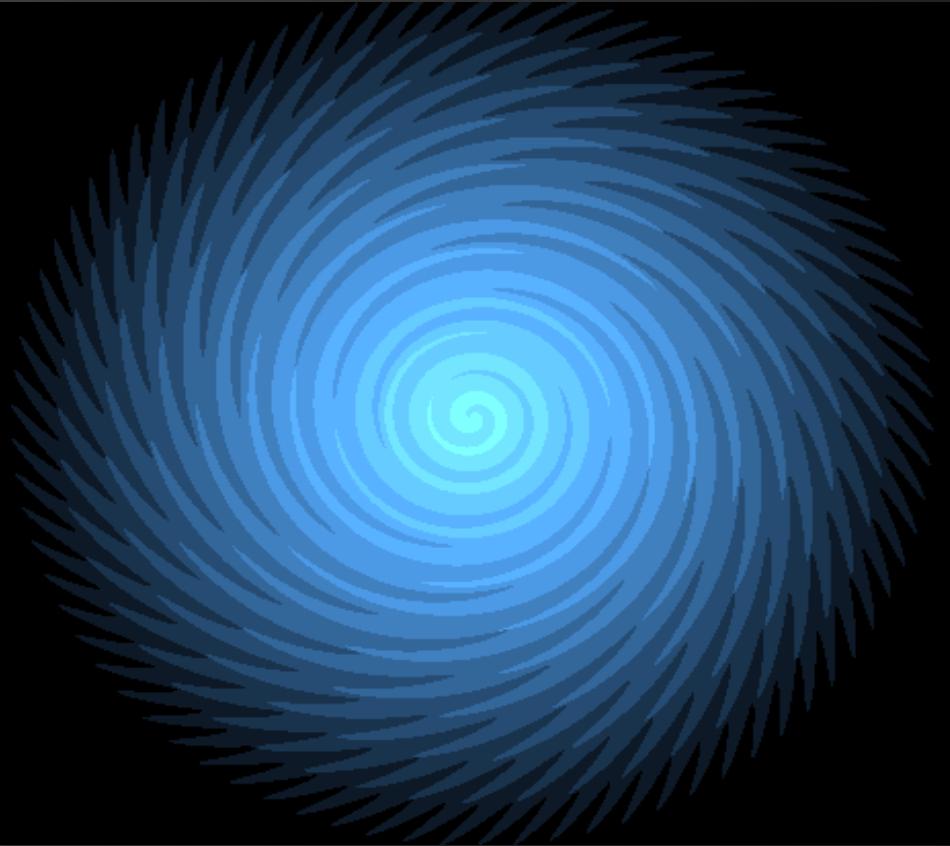
```
void main(void)
{
    float t = iGlobalTime;
    vec2 uv = gl_FragCoord.xy / iResolution.y;
    float r = length(uv), a = atan(uv.y,uv.x);
    float i = floor(r*10);
    a *= floor(pow(128,i/10));
    a += 20.*sin(0.5*t)+123.34*i-100. *
        (r*i/10)*cos(0.5*t);
    r += (0.5+0.5*cos(a)) / 10;
    r = floor(N*r)/10;
    gl_FragColor = (1-r)*vec4(0.5,1,1.5,1);
}
```



Introduction

GPGPU Programming

```
rics  
    & (depth < MAXDEPTH)  
    c = inside ? 1 : 1.2f;  
    nt = nc; ddn = ddc;  
    os2t = 1.0f - nnt * nnt;  
    D, N );  
)  
  
at a = nt - nc, b = nt  
at Tr = 1 - (R0 + (1 - R  
Tr) R = (D * nnt - N *  
E * diffuse;  
= true;  
  
-  
refl + refr)) && (depth <  
D, N );  
refl * E * diffuse;  
= true;  
  
MAXDEPTH)  
  
survive = SurvivalProbab  
estimation - doing it p  
if;  
radiance = SampleLight(  
e.x + radiance.y + radi  
  
v = true;  
at brdfPdf = EvaluateIf  
at3 factor = diffuse * I  
at weight = Mis2( direct  
at cosThetaOut = dot( N,  
E * ((weight * cosTheta  
  
random walk - done proper  
alive)  
  
;  
at3 brdf = SampleDiffuse( diffuse, N, r1, r2, &R, &pdf  
survive;  
pdf;  
n = E * brdf * (dot( N, R ) / pdf);  
ision = true;
```



Easy to port to GPU:

- Image postprocessing
- Particle effects
- Ray tracing
- ...

Today's Agenda:

- Introduction to GPGPU
- Example: Voronoi Noise
- GPGPU Programming Model
- OpenCL Template



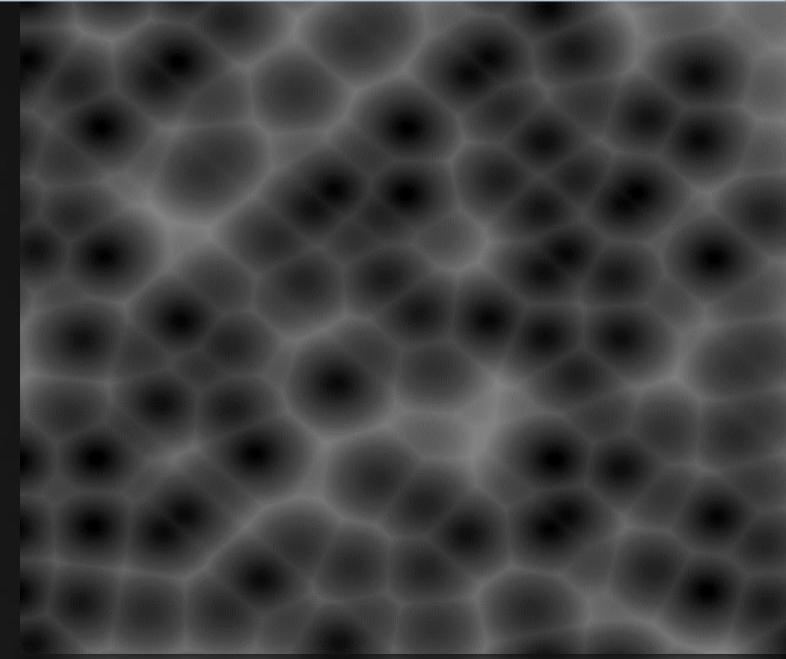
Example

Voronoi Noise / Worley Noise*

Given a random set of uniformly distributed points, and a position x in \mathbb{R}^2 , $\mathbf{F}_1(x) = \text{distance of } x \text{ to closest point.}$

For Worley noise, we use a Poisson distribution for the points. In a lattice, we can generate this as follows:

1. The expected number of points in a region is constant (Poisson);
2. The probability of each point count in a region is computed using the discrete Poisson distribution function;
3. The point count and coordinates of each point can be determined using a random seed based on the coordinates of the region in the lattice (so: *on the fly*)



*A Cellular Texture Basis Function, Worley, 1996





Mountains

shadertoy.com/view/4s1GD4

gmail news misc tech games coding research tools ompf2 huis Netflix glTF-Tutorials/gltfT... Ray Tracing Denoisi...

Shadertoy Search... Browse New Sign In

Mountains ★★

Views: 33228, Tags: raymarching, terrain, landscape, mountains, vr, soundcloud

Created by **Dave_Hoskins** in 2013-06-06

A Shader version of my terrain renderer:-
<http://www.youtube.com/watch?v=qzkBnCBpQAM>

USE MOUSE > TO SHIFT TIME

Video of my OpenGL version that uses streaming texture normals for speed...
<http://www.youtube.com/watch?v=qzkBnCBpQAM>

Comments (43)

Sign in to post a comment.

Image

Shader Inputs

```
1 // Mountains. By David Hoskins - 2013
2 // License Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License.
3
4 // https://www.shadertoy.com/view/4s1GD4
5 // A ray-marched version of my terrain renderer which uses
6 // streaming texture normals for speed-
7 // http://www.youtube.com/watch?v=qzkBnCBpQAM
8
9 // It uses binary subdivision to accurately find the height map.
10 // Lots of thanks to Inigo and his noise functions!
11
12 // Video of my OpenGL version that
13 // http://www.youtube.com/watch?v=qzkBnCBpQAM
14
15 // Stereo version code thanks to Croqueteer :)
16 // #define STEREO
17
18 float treeLine = 0.0;
19 float treeCol = 0.0;
20
21
22 vec3 sunLight = normalize( vec3( -0.4, 0.4, 0.48 ) );
23 vec3 sunColour = vec3(1.0, .9, .83);
24 float specular = 0.0;
25 vec3 cameraPos;
26 float ambient;
27 vec2 add = vec2(1.0, 0.0);
28 #define HASHSCALE1 .1031
29 #define HASHSCALE3 vec3(.1031, .1030, .0973)
30 #define HASHSCALE4 vec4(1031, .1030, .0973, .1099)
31
32 // This perturbs the fractal positions for each iteration down...
33 // Helps make nice twisted landscapes...
34 const mat2 rotate2D = mat2(1.3623, 1.7531, -1.7131, 1.4623);
35
36 // Alternative rotation-
37 // const mat2 rotate2D = mat2(1.2323, 1.999231, -1.999231, 1.22);
38
39
40 // 1 out, 2 in...
41 float Hash12(vec2 p)
42 +
```

Compiled in 0.2 secs 7149 chars

Audio not loaded

Example

Voronoi Noise / Worley Noise*

```

rics
  & (depth < MAXDEPTH)
  c = inside ? 1 : 1.2f;
  nt = nc / nc; ddn = nc;
  pos2t = 1.0f - nt * 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;
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* https://www.shadertoy.com/view/4djGRh
survive;
pdf;
n = E * brdf * (dot( N, R ) / pdf);
ision = true;

```

Characteristics of this code:

- Pixels are independent, and can be calculated in arbitrary order;
- No access to data (other than function arguments and local variables);
- Very compute-intensive;
- Very little input data required.



Example

Voronoi Noise / Worley Noise*

Timing of the Voronoi code in C++:

~250ms per image (1280 x 720 pixels), ~65 with multiple threads.

Executing the same code in OpenCL (GPU: GTX1060, mobile):

~1.2ms (faster).



```

rics
& (depth < MAXDEPTH)
    = 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, &light,
e.x + radiance.y + radiance.z) > 0) && (dot( N,
e.x + radiance.y + radiance.z) > 0);
e.z = radiance.z;
e = true;
at brdfPpdf = EvaluateDiffuse( L, N ) * Psurvive;
at t3 factor = diffuse * INMPI;
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;

```

Example

Voronoi Noise / Worley Noise

GPGPU allows for efficient execution of tasks that expose a lot of potential parallelism.

- Tasks must be independent;
- Tasks must come in great numbers;
- Tasks must require little data from CPU.

Notice that these requirements are met for rasterization:

- For thousands of pixels,
- fetch a pixel from a texture,
- apply illumination from a few light sources,
- and draw the pixel to the screen.



Today's Agenda:

- Introduction to GPGPU
- Example: Voronoi Noise
- GPGPU Programming Model
- OpenCL Template



Programming Model

GPU Architecture

A typical GPU:

- Has a small number of ‘shading multiprocessors’ (comparable to CPU cores);
- Each core runs a small number of ‘warps’ (comparable to hyperthreading);
- Each warp consists of 32 ‘threads’ that run in lockstep (comparable to SIMD).



Core 0



Core 1



Programming Model

GPU Architecture

Multiple warps on a core:

The core will switch between warps whenever there is a stall in the warp (e.g., the warp is waiting for memory). Latencies are thus hidden by having many tasks.

This is only possible if you feed the GPU enough tasks: $\text{cores} \times \text{warps} \times 32$.



Core 0



Core 1



Programming Model

GPU Architecture

Threads in a warp running in lockstep:

At each cycle, all ‘threads’ in a warp must execute the same instruction. Conditional code is handled by temporarily disabling threads for which the condition is not true. If-then-else is handled by sequentially executing the ‘if’ and ‘else’ branches. Conditional code thus reduces the number of active threads (occupancy). Note the similarity to SIMD code!



Core 0



Core 1



Programming Model

SIMT

The GPU execution model is referred to as SIMT: Single Instruction, Multiple Threads.

A GPU is therefore a very wide vector processor.

Converting code to GPGPU is similar to vectorizing code on the CPU.



Core 0

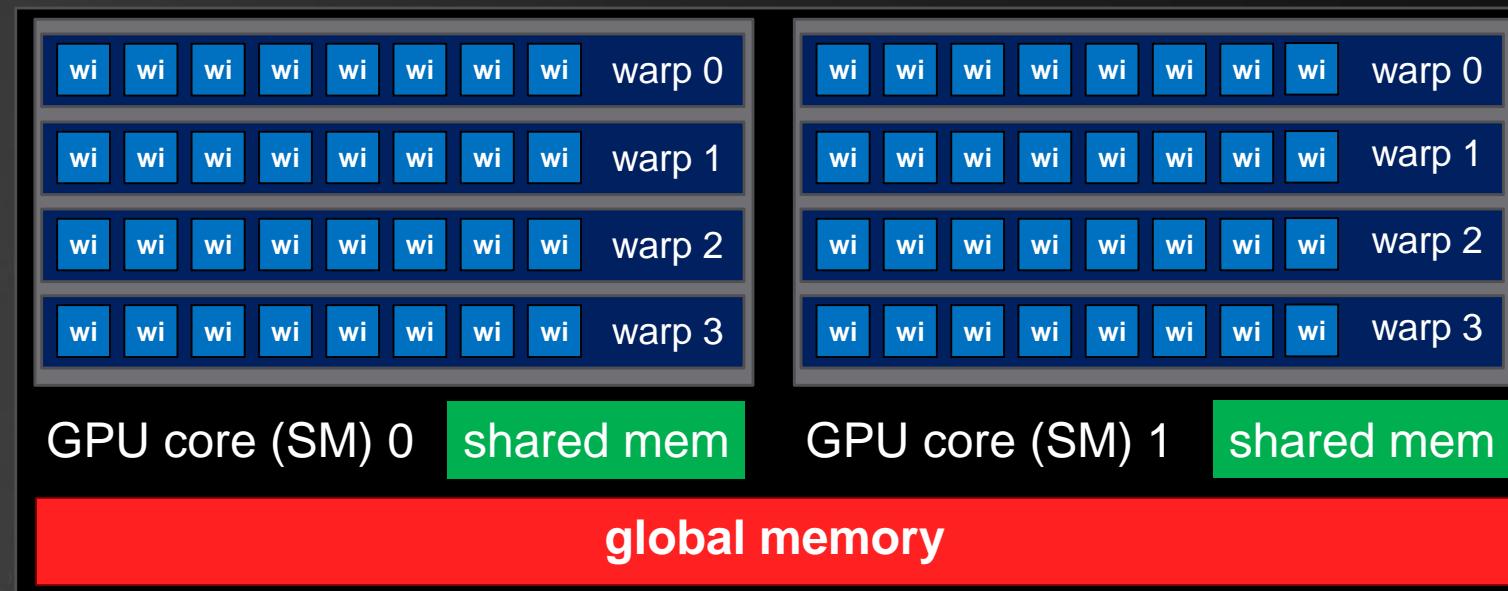


Core 1



Programming Model

GPU Memory Model

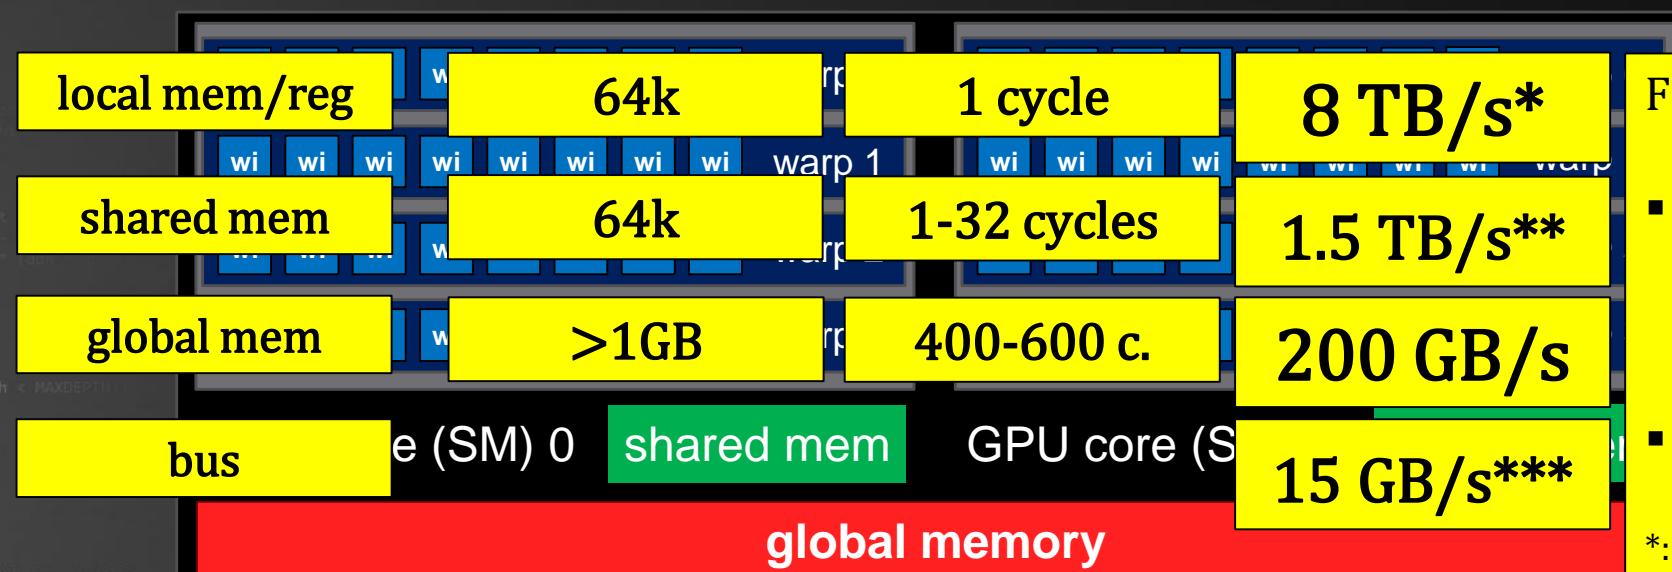


- Each SM has a large number of registers, which is shared between the warps.
- Each SM has shared memory, comparable to L1 cache on a CPU.
- The GPU has global memory, comparable to CPU RAM.
- The GPU communicates with the ‘host’ over a bus.



Programming Model

GPU Memory Model



For reference, Core i7-3960X:

- RAM bandwidth for quad-channel DDR3-1866 memory: **18.1GB/s**
- L2 bandwidth: **46.8GB/s***

*: Molka et al., Main Memory and Cache Performance of Intel Sandy Bridge and AMD Bulldozer. 2014.

* Values for NVidia G80 (Tesla)

** Fermi uses L1 cache

*** PCIe 3.0



Programming Model

GPU Memory Model

There appear to be many similarities between a CPU and a GPU:

- Cores, with hyperthreading
- A memory hierarchy
- SIMD

However, there are fundamental differences in each of these.

- One GPU core will execute up to 64 warps (instead of 2 on the CPU);
- The memory hierarchy is explicit on the GPU, rather than implicit on the CPU;
- GPU SIMD on the other hand is implicit (SIMT model).



Programming Model

GPGPU Programming Model

A number of APIs is available to run general purpose GPU code:

Pixel shaders:

- Executed as part of the rendering pipeline
- The number of tasks is equal to the number of pixels

Compute shaders:

- Executed as part of the rendering pipeline
- More control over the number of tasks

OpenCL / CUDA:

- Executed independent of rendering pipeline
- Full control over memory hierarchy and division of tasks over hardware

Graphics-centric work:

Shading, postprocessing (using a full-screen quad)

Graphics-centric work:

Preparing data, output to textures / vertex buffers / ...

General Purpose



Programming Model

GPGPU Programming Model

APIs like CUDA and OpenCL may look like C, but are in fact heavily influenced by the underlying hardware model.

```
__kernel void task( write_only image2d_t outimg, __global uint* logBuffer )  
{  
    float t = 1;  
    int column = get_global_id( 0 );  
    int line = get_global_id( 1 );  
    float c = Cells( (float2)((float)column / 500, (float)line / 500), t );  
    write_imagedf( outimg, (int2)(column, line), c );  
}
```

- Kernel: one task (of which we need thousands to run efficiently);
- `get_global(0,1)`: identifies a single task from a 2D array of tasks.

Many threads will execute the same kernel. We can not execute different code in parallel.



Programming Model

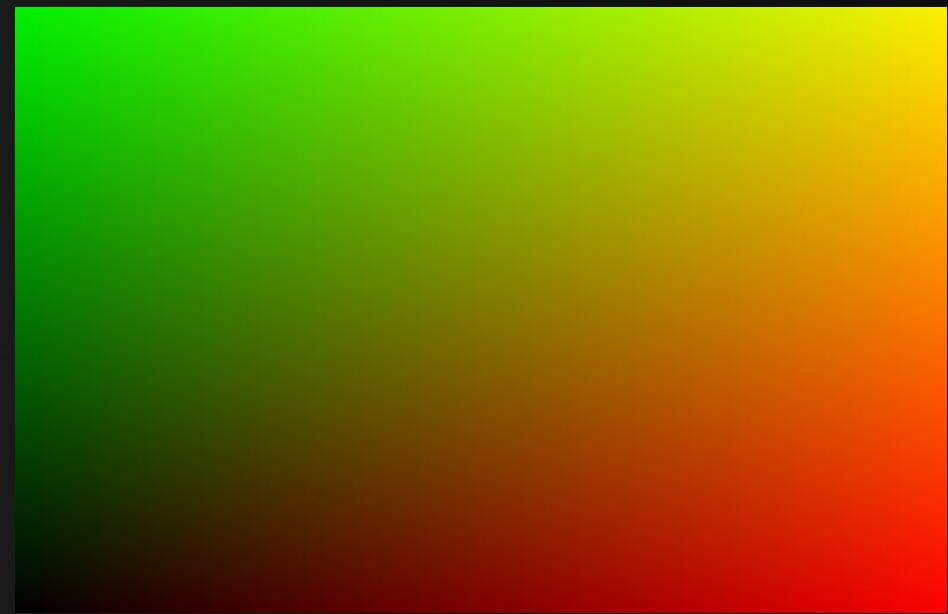
GPGPU Programming Model

Kernels are invoked from the host:

```
size_t workSize[2] = { SCRWIDTH, SCRHEIGHT };
void Kernel::Run( cl_mem* buffers, int count )
{
    ...
    clEnqueueNDRangeKernel( queue, kernel, 2, 0, workSize, NULL, 0, 0 );
    ...
}
```

Device code:

```
__kernel void main( write_only image2d_t outimg )
{
    int column = get_global_id( 0 );
    int line = get_global_id( 1 );
    float red = column / 800.;
    float green = line / 480.;
    float4 color = { red, green, 0, 1 };
    write_imagef( outimg, (int2)(column, line), color );
}
```



Programming Model

GPGPU Programming Model

Kernels are invoked from the host:

```

rics
& (depth < MAXDEPTH)
c = 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, &light
e.x + radiance.y + radiance.z ) > 0) && (dot( N,
E = 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.z
random walk - done properly, closely following Sm
alive);
at3 brdf = SampleDiffuse( diffuse, N, r1, r2 } R, Pdf;
survive;
pdf;
E * brdf * (dot( N, R ) / pdf);
ision = true;

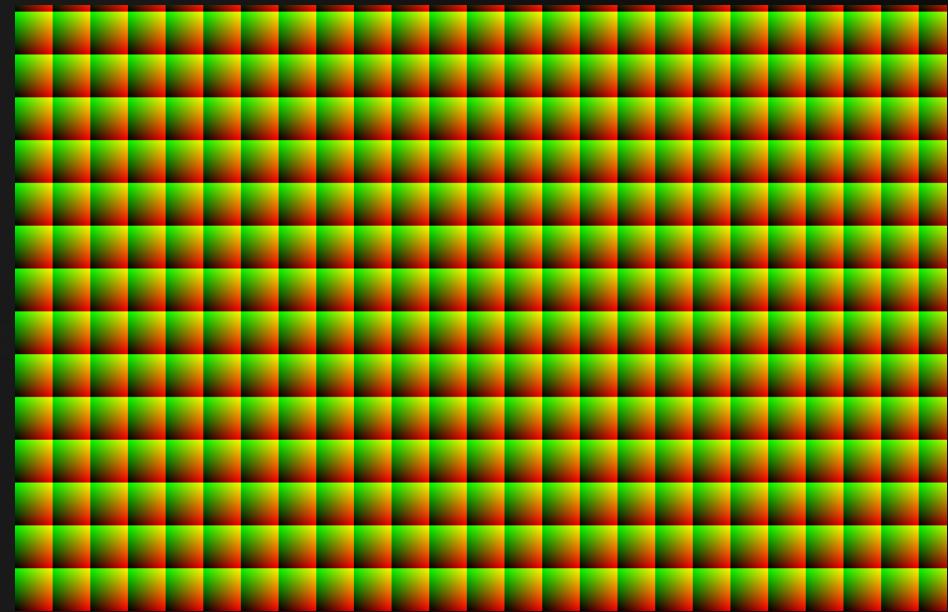
```

Device code:

```

__kernel void main( write_only image2d_t outimg )
{
    int column = get_global_id( 0 );
    int line = get_global_id( 1 );
    float red = get_local_id( 0 ) / 32.;
    float green = get_local_id( 1 ) / 32.;
    float4 color = { red, green, 0, 1 };
    write_imagef( outimg, (int2)(column, line), color );
}

```



Today's Agenda:

- Introduction to GPGPU
- Example: Voronoi Noise
- GPGPU Programming Model
- OpenCL Template



Template

OCL_Lab: The Familiar Template

The OpenCL template is a basic experimentation framework for OpenCL.
Game::Tick implements the following functionality:

1. Set arguments for the OpenCL kernel;
2. Execute the OpenCL kernel (which stores output in an OpenGL texture);
3. Draw a full-screen quad using a shader.

You can find the OpenCL code in program.cl;
The shader is defined in vignette.frag.

```
rics
3 & (depth < MAXDEPTH)
    = 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, &lightDir,
e.x + radiance.y + radiance.z) > 0) && (dot( N,
e.x + e.y + e.z ) > 0);
e.w = 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;
```



/INFOMOV/

END of “GPGPU (1)”

next lecture: “GPGPU (2)”

```
rics
3 & (depth < MAXDEPTH)
    n = inside ? 1 : 1.0f;
    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, &lightDir,
e.x + radiance.y + radiance.z) > 0) && (dot( N,
e, radiance ) > 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
ive);
;

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

