

Introduction to Programmable GPUs

CPSC 314

News


- Homework: no new homework this week (focus on quiz prep)
- Quiz 2
 - this Friday
 - topics:
 - everything after transformations up until last Friday's lecture
 - questions on rendering pipeline as a whole
- Office hours (Wolfgang) Thursday, Friday 11:30-12:30

Real Time Graphics




GPUs vs CPUs

- 800 GFLOPS vs 80 GFLOPS
- 86.4 GB/s vs 8.4 GB/s




GPUs vs CPUs

- 800 GFLOPS vs 80 GFLOPS
- 86.4 GB/s vs 8.4 GB/s




Programmable Pipeline

- so far:
 - have discussed rendering pipeline as specific set of stages with fixed functionality



Programmable Pipeline

- now: programmable rendering pipeline!



Vertex Shader

- performs all per-vertex computation (transform & lighting):


- model and view transform
- perspective transform
- texture coordinate transform
- per-vertex lighting

Vertex Shader

- input:
 - vertex position and normal (sometimes tangent)
 - (multi-)texture coordinate(s)
 - modelview, projection, and texture matrix
 - vertex material or color
 - light sources – color, position, direction etc.
- output:
 - 2D vertex position
 - transformed texture coordinates
 - vertex color

Vertex Shader - Applications


- deformable surfaces: skinning
- different parts have different rigid transformations
- vertex positions are blended
- used in facial animations – many transformations!



Fragment Shader

- performs all per-fragment computation:
 - texture mapping
 - fog
- input (interpolated over primitives by rasterizer):
 - texture coordinates
 - color
- output:
 - fragment color

Fragment Shader - Applications



Vertex & Fragment Shader

- massively parallel computing by parallelization
- same shader is applied to all data (vertices or fragments) – SIMD (single instruction multiple data)
- parallel programming issues:
 - main advantage: high performance
 - main disadvantage: no access to neighboring vertices/fragments

Vertex Shader - Instructions

- Arithmetic Operations on 4-vectors:
 - ADD, MUL, MAD, MIN, MAX, DP3, DP4
- Operations on Scalars
 - RCP (1/x), RSQ (1/sqrt(x)), EXP, LOG
- Specialty Instructions
 - DST (distance: computes length of vector)
 - LIT (quadratic falloff term for lighting)
- Later generation:
 - Loops and conditional jumps

Vertex Shader - Example

- morph between cube and sphere & lighting
- vertex attributes: $v[0..N]$, matrices $c[1..N]$, registers R

```

#blend normal and position
# m = cv + (1-c)v2, # c = (v1.v1) / (v1.v1 + v2.v2)
MOV R3, v[2]
MOV R4, v[2]
ADD R8, v[1], -R3
ADD R9, v[1], -R3
MAD R8, v[15], x, R8, R3
MAD R9, v[15], x, R9, R5

# transform normal to eye space
DP3 R9, v, R8, c[12]
DP3 R9, v, R8, c[13]
DP3 R9, v, R8, c[14]

# transform position and output
DP4 c[RP00], v, R8, c[4]
DP4 c[RP08], v, R8, c[5]
DP4 c[RP16], v, R8, c[6]
DP4 c[RP24], v, R8, c[7]

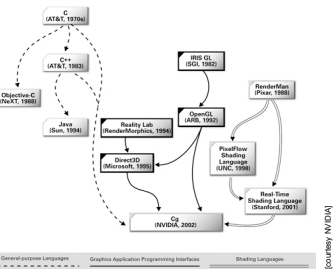
# normalize normal
DP3 R9, v, R8, R9
RSQ R9, v, R9, w
MUL R9, R9, w

# apply lighting and output color
DP3 R0, x, R9, c[20]
DP3 R0, y, R9, c[22]
MOV R0, w, c[21]
LIT R1, R0
DP3 c[CO0], c[21], R1
    
```

Shading languages

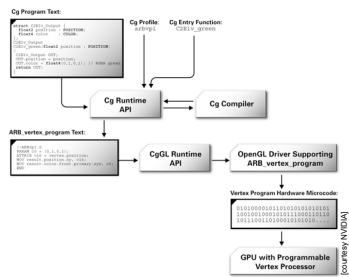
- Cg (C for Graphics – NVIDIA)
- GLSL (GL Shading Language – OpenGL)
- HLSL (High Level Shading Language – MS Direct3D)

Cg History



Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

Cg – How does it work?



Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

Cg – Integration into OpenGL



```
void initShader(void) {
    // get fragment shader profile
    _cgFragmentProfile =
        cgGLGetLatestProfile(CG_GL_FRAGMENT);
    // init Cg context
    _cgContext = cgCreateContext();
    // load shader from file
    _cgProgram =
        cgCreateProgramFromFile(_cgContext,
            CG_SOURCE,
            "MyShader.cg",
            cgFragmentProfile,
            NULL, NULL);
    // upload shader on GPU
    cgGLLoadProgram(_cgProgram);
    // get handles to shader parameters
    _cgTexture =
        cgGLSetNamedParameter(_cgProgram, "texture");
    _cgParameter =
        cgGLNamedParameter(_cgProgram, "parameter");
}
```

```
void displayLoop(void) {
    // setup transformation
    ...
    // enable shader and set parameters
    cgGLEnableProfile(_cgFragmentProfile);
    cgGLBindProgram(_cgProgram);
    // set Cg texture
    cgGLSetTextureParameter(_cgTexture, _texturalID);
    cgGLSetTextureParameter(_cgTexture);
    // set gamma
    cgGLSetParameter1f(_cgParameter, _parameter);
    // draw geometry
    ...
    // disable Cg texture and profile
    cgGLDisableTextureParameter(_cgTexture);
    cgGLDisableProfile(_cgFragmentProfile);
    // swap buffers
    ...
}
```

Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

Cg Example – Fragment Shader



- Fragment Shader: gamma mapping



```
void main(float4 texCoord : TEXCOORD,
         uniform samplerRECT texture,
         uniform float gamma,
         out float4 color : COLOR)
{
    // perform texture look up
    float3 textureColor = f4texRECT(texture, texCoord.xy).rgb;
    // set out color
    color.rgb = pow(textureColor, gamma);
}
```

Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

Cg Example – Vertex Shader



- Vertex Shader: animated teapot



```
void main()
{
    // user parameters
    uniform float4x4 objectMatrix;
    uniform float4x4 objectMatrixInv;
    uniform float4x4 modelViewMatrix;
    uniform float4x4 modelViewMatrixInv;
    uniform float4x4 projectionMatrix;
    uniform float deformation;
    uniform float3 lightPosition;
    uniform float3 lightAmbient;
    uniform float3 lightDiffuse;
    uniform float3 lightSpecular;
    uniform float3 materialEmission;
    uniform float3 materialAmbient;
    uniform float3 materialDiffuse;
    uniform float3 materialSpecular;
    uniform float3 materialShininess;

    // object coordinate system matrix
    // object coordinate system matrix inverse transpose
    // modelview matrix
    // modelview matrix inverse transpose
    // projection matrix
    // deformation parameter
    // light position
    // light ambient parameter
    // light diffuse parameter
    // light specular parameter
    // light attenuation parameter - constant, linear, quadratic
    // material emission parameter
    // material ambient parameter
    // material diffuse parameter
    // material specular parameter
    // material shininess parameter

    // output
    out float4 outPosition : POSITION; // position in object coordinates
    out float4 outNormal : NORMAL; // normal

    // compute the diffuse term
    float4 positionObject = mul(ObjectMatrix, POSITION);
    float4 positionWorld = mul(modelViewMatrix, positionObject);
    float4 viewerPositionWorld = float4(0.0, 0.0, 0.0, 1.0);
    // assume viewer position is in origin
    // apply deformation
    positionObject.xyz = positionObject.xyz +
        deformation * normalize(normalObject.xyz);
    float4 positionWorld = mul(modelViewMatrix, positionObject);
    outPosition = mul(projectionMatrix, positionWorld);

    // two vectors
    float3 P = positionWorld.xyz;
    float3 N = normalize(normalWorld.xyz);

    // compute the ambient term
    float3 ambient = materialAmbient * lightAmbient;

    // compute the diffuse term
    float3 L = normalize(lightPositionWorld.xyz - P);
    float diffuseFactor = max(dot(N, L), 0);
    float3 diffuse = materialDiffuse * lightDiffuse * diffuseFactor;

    // compute the specular term
    float3 H = normalize(viewerPositionWorld.xyz -
        positionWorld.xyz);
    float specularFactor = 1;
    float shininess = materialShininess;
    if (diffuseFactor <= 0) specularFactor = 0;
    float3 specular =
        materialSpecular *
        lightSpecular *
        specularFactor;

    // attenuation factor
    float distanceLightVertex = length(P - lightPositionWorld.xyz);
    float attenuationFactor = 1 / (1 + lightAttenuation.x * distanceLightVertex + lightAttenuation.y * distanceLightVertex + lightAttenuation.z * distanceLightVertex);

    // set output color
    outColor.rgb = materialEmission + ambient + diffuse + specular;
    outColor.w = 1;
}
```

Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

Cg Example – Vertex Shader



```
// transform position from object space to clip space
float4 positionObject = mul(ObjectMatrix, POSITION);
float4 positionWorld = mul(modelViewMatrix, positionObject);
float4 viewerPositionWorld = float4(0.0, 0.0, 0.0, 1.0);
// assume viewer position is in origin
// apply deformation
positionObject.xyz = positionObject.xyz +
    deformation * normalize(normalObject.xyz);
float4 positionWorld = mul(modelViewMatrix, positionObject);
outPosition = mul(projectionMatrix, positionWorld);

// two vectors
float3 P = positionWorld.xyz;
float3 N = normalize(normalWorld.xyz);

// compute the ambient term
float3 ambient = materialAmbient * lightAmbient;

// compute the diffuse term
float3 L = normalize(lightPositionWorld.xyz - P);
float diffuseFactor = max(dot(N, L), 0);
float3 diffuse = materialDiffuse * lightDiffuse * diffuseFactor;

// compute the specular term
float3 H = normalize(viewerPositionWorld.xyz -
    positionWorld.xyz);
float specularFactor = 1;
float shininess = materialShininess;
if (diffuseFactor <= 0) specularFactor = 0;
float3 specular =
    materialSpecular *
    lightSpecular *
    specularFactor;

// attenuation factor
float distanceLightVertex = length(P - lightPositionWorld.xyz);
float attenuationFactor = 1 / (1 + lightAttenuation.x * distanceLightVertex + lightAttenuation.y * distanceLightVertex + lightAttenuation.z * distanceLightVertex);

// set output color
outColor.rgb = materialEmission + ambient + diffuse + specular;
outColor.w = 1;
}
```

Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

Cg Example – Phong Shading



vertex shader

```
void main()
{
    // user parameters
    ...
    // output
    out float4 outTexCoord0 : TEXCOORD0; // world normal
    out float4 outTexCoord1 : TEXCOORD1; // world position
    out float4 outTexCoord2 : TEXCOORD2; // world light position
    out float4 outPosition : POSITION; // position in clip space
    ...
    // transform position from object space to clip space
    ...
    // compute the ambient term
    ...
    // compute the diffuse term
    ...
    // compute the specular term
    ...
    // set world normal as out texture coordinate0
    outTexCoord0 = normalWorld;
    // set world position as out texture coordinate1
    outTexCoord1 = positionWorld;
    // set output color
    outTexCoord2 = mul(modelViewMatrix, float4(lightPosition, 1));
}
```



Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

Cg Example – Phong Shading



fragment shader

```
void main(float4 normal : TEXCOORD0, // normal
         float4 position : TEXCOORD1, // position
         out float4 outColor : COLOR)
{
    // compute the ambient term
    ...
    // compute the diffuse term
    ...
    // compute the specular term
    ...
    // attenuation factor
    ...
    // set output color
    outColor.rgb = materialEmission + ambient + attenuationFactor * (diffuse + specular);
}
```



Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

GPGPU

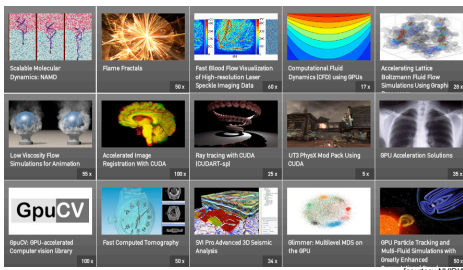


- general purpose computation on the GPU
- in the past: access via shading languages and rendering pipeline
- now: access via cuda interface in C environment



Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09

GPGPU Applications



Introduction to GPU Programming | CS314 Gordon Wetzstein, 09/03/09