

PV227 GPU programming

Marek Vinkler

Department of Computer Graphics and Design



GLSL

- officialy OpenGL Shading Language,
- part of OpenGL standard (from OpenGL 2.0),
- high-level procedural language (based on C and C++),
- independent on hardware,
- performance oriented (through custom compilers).



GLSL

- single set of instructions for all shaders (almost),
- native support for vectors and matrices,
- no pointers (hurray :D) and strings,
- strict with types,
- no length limit (language part).



GLSL

- part of the OpenGL driver – graphics driver,
- common front-end (should be), different (optimized) back-ends,
- **shaders** are combined into **programs**,
- linking resolves cross shader references,
- shaders are strings, not files (no **#include**).



Scalar data types

- **float**,
- **int**,
- **uint**,
- **bool**,
- declarations may appear anywhere.



Scalar data types

```
1 float f;  
2 float h = 2.4; // float constant in GLSL 3.3 and below  
3 f = 0.2f;  
4 float ff = 1.5e10;  
5 ff -= 1.E-3;  
6  
7 uint n = 5;  
8 n = 15u;  
9 int a = 0xA;  
10 a += 071;  
11  
12 bool skip = true;  
13 skip = skip && false;
```



Vector data types

- **vec2, vec3, vec4** – float,
- **ivec2, ivec3, ivec4** – int,
- **uvec2, uvec3, uvec4** – uint,
- **bvec2, bvec3, bvec4** – bool,
- two, three or four component vectors of scalar types.



Vector data types

- field selection or array access,
- **x, y, z, w** – for positions or directions,
- **r, g, b, a** – for colors,
- **s, t, p, q** – for texture coordinates,
- only for readability, all select certain vector coordinate (e.g. $v.x \equiv v.r \equiv v.s \equiv v[0]$).



Matrix data types

- only matrices of **floats**
- **mat2**, **mat3**, **mat4** – 2×2 , 3×3 , 4×4 matrices,
- **matmxn** – $m \times n$ (column \times row) matrix,
- column major order (first coordinate is column),
- as in OpenGL, unlike C/C++.

```
1 mat4 m;  
2 vec4 v = m[3]; // Fourth column  
3 float f = m[3][1]; // Second component (row) of the fourth  
   column vector
```



Sampler data types

- for texture access,
- variants for **floats**, **ints**, **unsigned ints** (no **bool**),
- **[i|u]sampler{1|2|3}D** – access one, two or three dimensional texture,
- **[i|u]samplerCube** – access cube-map texture,
- **[i|u]sampler2DRect** – access two-dimensional rectangle texture,
- **[i|u]sampler{1|2}DArray** – access one or two dimensional texture array,
- **[i|u]samplerBuffer** – access texture buffer,



Sampler data types

- **sampler{1|2|3}DShadow** – access one, two or three dimensional depth texture with comparison,
- **sampler{1|2|3}DShadow** – access one, two or three dimensional depth texture with comparison,
- **sampler2DRectShadow** – access two-dimensional rectangle depth texture with comparison,
- **sampler{1|2}DArrayShadow** – access one or two dimensional depth texture array with comparison.



Sampler data types

- application initializes the samplers,
- passed to shaders through **uniform** variables,
- samplers cannot be manipulated in shader,
- shadow textures and color samplers must not be mixed → undefined behaviour.

```
1 uniform sampler2D sampler;  
2 vec2 coord = vec2(0.f, 1.f);  
3 vec4 color = texture(sampler, coord);
```



Structures

- C++ style (name of structure → name of type),
- can be embedded and nested, contain arrays,
- bit-fields not allowed, no **union**, **enum**, **class**.

```
1 struct vertex
2 {
3     vec3 pos;
4     vec3 color;
5 };
6 vertex v;
```



Arrays

- available for any type,
- zero indexed,
- no pointers → always declared with [] and size,
- the array must be declared with same size in all shaders.



Declarations and scopes

- variable name format same as in C/C++ (case sensitive),
- names beginning with “gl_” or “__” are reserved,
- scoping similar to C++.

```
1 float f; // Declared from this point until the end of the block
2 for(int i = 0; i < 3; ++i) // i is declared only in this cycle
3     f *= f;
4
5 if(i == 1) // Invalid
6 {
7     ...
8 }
```

Initializers and constructors

- scalar variables may be initialized in declaration,
- constants must be initialized,
- **in** and **out** variables may not be initialized,
- **uniform** variables may be initialized.

```
1 int a = 0, b, c = 1;  
2 const float eps = 1e-3f;  
3 uniform float temp = 36.5f;
```



Initializers and constructors

- aggregate types are initialized/set with constructors,
- the number of components in vectors need not match.

```
1 vec2 v = vec2(0.f, 1.f);
2 v = vec2(1.f, 0.f);
3 vec3 v3 = vec3(v, 0.f);
4
5 v3 = vec3(1.f); // vec3(1.f, 1.f, 1.f);
6 v = vec2(v3); // vec2(1.f, 1.f);
7
8 float array[4] = float[4](0.f, 1.f, 2.f, 3.f);
9
10 struct person
11 {
12     struct attrib
13     {
14         vec3 color;
15         bool active;
16     };
17     vec3 pos;
18 } person1 = person(attrib(vec3(0.5f, 0.5f, 0.5f), true), v3);
```

Matrix constructors

- matrix components are read and written in column major order,
- matrices cannot be constructed from matrices.

```
1 mat2 matrix = mat2(1.f, 2.f, 3.f, 4.f); // 1.f, 3.f
2                                           // 2.f, 4.f);
3 mat2 identity = mat2(1.f); // Initializes diagonal
4                                           // mat2(1.f, 0.f, 0.f, 1.f);
5
6 vec2 v = vec2(1.0f);
7 mat2 identity2 = mat2(v);
```

Type matching and promotion

- strict matching (prevents ambiguity),
- assigned types, functions parameters must match exactly,
- scalar integers may be implicitly converted to scalar floats,
- may force the programmer to use explicit conversion.



Type conversions

- performed with constructors,
- no C-style typecast,
- no way to reinterpret a value,
- conversions to boolean → non-zero as **true**, zero as **false**,
- conversions from boolean → **true** as 1 (1.f), **false** as 0 (0.f).



GLSL qualifiers

- tell compiler where the value comes from,
- **in** – vertex attribute (vertex shader) vertex data (geometry shader) or interpolated value (fragment shader),
- **uniform** – constant variable in all shaders,
- **out** – varying variable passed from one shader to another, output to frame buffer,
- **const** – compile time constant variables,
- **in**, **uniform**, **out** are always global variables,
- qualifier are specified before variable type.



Uniform qualifier

- cannot be modified from shader,
- less frequently updated, max once per primitive,
- all data types supported,
- used for samplers,
- all shaders inside a program share uniform variables.



In qualifier (vertex shader)

- vertex attributes,
- can be changed as often as a single vertex,
- not all data types supported:
 - boolean scalars and vectors,
 - structures,
 - arrays.



Out qualifier (vertex shader/geometry shader)

- output to the geometry shader / rasterizer,
- interpolation qualifiers for computing fragments:
 - **smooth out** – perspective-correct interpolation,
 - **noperspective out** – interpolation without perspective correction,
 - **flat out** – no interpolation.
- floating point scalars, vectors, matrices and arrays,
- with **flat out**: [unsigned] integer scalars, vectors, arrays,
- no structures.



In qualifier (fragment shader)

- interpolated values from the rasterizer,
- must match the definition of **out** variables in vertex / geometry shader,
 - interpolation qualifier, type, size, name.



Out qualifier (fragment shader)

- passed to per fragment fixed-function stage,
- floating point/integer/unsigned integer scalars, vectors and arrays,
- no matrices or structures,
- can be preceded with `layout(location = x)`, where `x` is the number of the render target.



Constant qualifiers

- compile time constant,
- not visible outside the shader,
- individual structure items may not be constants,
- initializers may contain only literal values or other **const** variables.



No qualifiers

- can be both read and written,
- unqualified global variables,
 - shared between shader of the same type, not between shaders of different types,
 - not visible outside program,
 - lifetime limited to a single run of a shader (no “**static**”),
 - different variables for different processors → do NOT use.



Interface blocks

- common names for several variables,
- different meaning for each qualifier, same syntax,
- used for passing data between shaders, loading uniform variables.

```
1 storage_qualifier block_name
2 {
3     <members defition >
4 } [instance_name];
```

- block_name used from OpenGL,
- instance_name optional to create named instances inside GLSL,
- possible arrays of instances.



Inter shader communication

Name based matching:

```
1 // vertex shader
2 out vec4 color;
3
4 _____
5 // geometry shader
6 in vec4 color[];
7 out vec4 colorFromGeom;
8
9 _____
10 // fragment shader
11 in vec4 colorFromGeom;
```

- names in shaders must match,
- **in** and **out** cannot be named the same,
- cannot use the same shader for vertex → fragment and vertex → geometry → fragment.



Inter shader communication

Location based matching:

```
1 // vertex shader
2 layout (location = 0) out vec3 normalOut;
3 layout (location = 1) out vec4 colorOut;
4
5 -----
6 // geometry shader
7 layout (location = 0) in vec3 normalIn [];
8 layout (location = 1) in vec4 colorIn [];
9
10 layout (location = 0) out vec3 normalOut;
11 layout (location = 1) out vec4 colorOut;
12
13 -----
14 // fragment shader
15 layout (location = 0) in vec3 normalIn;
16 layout (location = 1) in vec4 colorIn;
```

- locations in shaders must match,
- location is per max vec4 item (not aggregate types),
- difficulty with assigning location numbers.



Inter shader communication

Interface based matching:

```
1 // vertex shader
2 out Data {
3     vec3 normal;
4     vec3 eye;
5     vec2 texCoord;
6 } DataOut;
7
8 // geometry shader
9 in Data {
10    vec3 normal;
11    vec3 eye;
12    vec2 texCoord;
13 } DataIn [];
14
15 out Data {
16    vec3 normal;
17    vec3 eye;
18    vec2 texCoord;
19 } DataOut;
```


Inter shader communication (cont.)

```
21 // fragment shader
22 in Data {
23     vec3 normal;
24     vec3 eye;
25     vec2 texCoord;
26 } DataIn;
27 ...
28 DataOut.normal = normalize(someVector);
```

- block names in shaders must match,
- data manipulation through instance name,
- same members in blocks.

Uniform interface blocks

- sharing uniforms between programs,
- setting multiple uniforms at once,
- named blocks of uniform variables (individual items are globally scoped),
- backed by buffers for data transfer,
- for setting transform matrices, common variables in shader families etc.



Uniform interface blocks

```
1 layout (xxx) uniform ColorBlock {
2     vec4 diffuse;
3     vec4 ambient;
4 };
5 ...
6 out vec4 outputF;
7
8 void main() {
9     outputF = diffuse + ambient;
10 }
```

- layout specifies storage (default is implementation dependent),
- **std140** – OpenGL specified layout, blocks can be shared between shaders,
- **shared** – implementation dependent layout, blocks can be shared between shaders,
- **packed** – unused variables are optimized-out, not shareable.



Uniform interface blocks

- uniform blocks are connected with buffers through binding points,
- block indices are assigned during program link,
- multiple blocks can be bound to the same binding point.

```
1 GLuint bindingPoint = 1, buffer, blockIndex;
2 float myFloats[8] = {1.0, 0.0, 0.0, 1.0, 0.4, 0.0, 0.0, 1.0};
3
4 // Assign the uniform block to the binding point
5 blockIndex = glGetUniformLocation(p, "ColorBlock");
6 glUniformBlockBinding(p, blockIndex, bindingPoint);
7
8 glGenBuffers(1, &buffer);
9 glBindBuffer(GL_UNIFORM_BUFFER, buffer);
10
11 // Assign the buffer to the binding point
12 glBufferData(GL_UNIFORM_BUFFER, sizeof(myFloats), myFloats,
13             GL_DYNAMIC_DRAW);
13 glBindBufferBase(GL_UNIFORM_BUFFER, bindingPoint, buffer);
```

Uniform interface blocks

- individual uniforms may be aligned in memory,
- to set them correctly we need to compute their offset,
- queried with **glGetActiveUniformBlockiv** and **glGetActiveUniformsiv**,
- set with **glBufferSubData**.

```
1 layout (std140) uniform ColorBlock2 {
2     vec3 diffuse;
3     vec3 ambient;
4 };
5
6 GLuint bindingPoint = 1, buffer, blockIndex;
7 float myFloats[3] = {0.4, 0.0, 0.0};
8
9 glGenBuffers(1, &buffer);
10 glBindBuffer(GL_UNIFORM_BUFFER, buffer);
11
12 glBufferSubData(GL_UNIFORM_BUFFER, 4*sizeof(float), sizeof(
    myFloats), myFloats); // Notice the offset
```

Program flow

- similar to C++,
- **main** is the entry point for a shader,
- global variables are initialized before **main** is executed,
- looping
 - **for**, **while**, **do-while**, **break**, **continue**,
- selection
 - **if**, **if-else**, **if-else if-else**, **?:** and **switch**,
- expressions must be booleans,
- partial evaluation of **&&** and **||**, **?:**,
- no **goto**,
- **discard** prevents fragment from updating frame buffer.



Functions

- support for C++ overload by parameter type,
- prototype declaration or definition before call to the function,
- exact matching of parameters, return values (**return**),
- no recursion,
- programs entry point is function `void main()`.



Calling conventions

- value-return,
- all input parameter values are copied to function before execution,
- all output parameter values are copied from the function after execution,
- parameter behaviour controlled by qualifiers **in** (default), **out** and **inout**,
- **in** parameters can be also **const** (not writeable inside function).



Functions continued

- arrays and structures are also copied by value,
- any return type (including structures).

```
1 void foo1(in vec3 normal, float eps, inout vec3 coord);
2 vec3 foo2(in vec3 normal, float eps, in vec3 coord);
3 void foo3(in vec3 normal, float eps, in vec3 coord, out vec3
  coordOut);
4
5 // Get coord
6 vec3 coord;
7 foo1(normal, eps, coord);
8 coord = foo2(normal, eps, coord);
9 foo3(normal, eps, coord, coord);
```



Swizzling

- used to select (rearrange) components of a vector,
- must use component names from the same set,
- must still be a valid type (no more than 4 components),
- R-values
 - any combination and repetition of components,
- L-values
 - no repetition of components.

```
1 vec4 pos = vec4(1.f, 2.f, 3.f, 4.f);
2 vec2 v1 = pos.xy;
3 vec3 v2 = pos.abb;
4 vec4 v3 = pos.xyrs; // Illegal different sets
5 vec4 o;
6 o.xw = v2; // (1.f, 2.f, 3.f, 4.f)
7 o.xx = vec2(0.f); // Illegal repetition
```

Operations on vectors and matrices

- mostly component-wise (independently for each component),
- vector sizes must match,
- vector * matrix and matrix * matrix are not component-wise,
- logical operations (!, &&, ||, ^^) only on scalar boolean,
- **not** component-wise logical not on boolean vectors.



Operations on vectors and matrices

- relational operators ($<$, $>$, $<=$, $>=$) on scalar floats and integers \rightarrow scalar boolean,
- build-in functions like **lessThanEqual** do component-wise relational operations on vectors,
- **==** and **!=** operate on all types except arrays \rightarrow scalar boolean,
- for component-wise comparison **equal** and **nonEqual** \rightarrow boolean vector,
- **any** and **all** turn boolean vector into boolean scalar,
- **=** and its variants (**+=**, **-=**, ***=**, **/=**) operates on all types except structures and arrays.



Preprocessor

- basically the same as in C,
- macros beginning with “GL_” or “_” are reserved,
- shaders should declare the GLSL version they are written for (**#version** number) as the first line of the code,
- usefull pragmas **optimize(on/off)** and **debug(on/off)**,
- language extensions can be accessed using **#extension**.



Build-in functions

- make shader programming easier,
- expose hardware functionality not writeable in the shader,
- provide optimized (possibly hardware accelerated) implementations of common functions,
- usually both scalar and vector variants,
- can be overridden by redeclaration,
- may be specific for a single shader type.



Shader specific functions

Geometry shader:

- `void EmitVertex(void);`,
 - use the current output state for a new vertex,
- `void EndPrimitive(void);`,
 - complete the current output primitive.



Keep up to date

- <http://www.opengl.org/sdk/docs/man/>
- <http://www.opengl.org/sdk/docs/manglsl/>
- <http://www.opengl.org/registry/>

