

---

# Uniforms

by Ian Romanick

This work is licensed under the Creative Commons Attribution Non-commercial Share Alike (by-nc-sa) License. To view a copy of this license, (a) visit <http://creativecommons.org/licenses/by-nc-sa/3.0/>; or, (b) send a letter to Creative Commons, 171 2nd Street, Suite 300, San Francisco, California, 94105, USA.

## 1 Introduction

While attributes change per vertex, uniforms are constant across a group of vertexes. Uniforms can be used for things like transformation matrices, lights, and other per-object parameters. Figure 1 uses a uniform called `model_view_projection` to transform the attribute `gl_Vertex`.

---

**Figure 1** Simple vertex shader

```
1 uniform mat4 model_view_projection;
2 void main(void)
3 {
4     gl_Position = model_view_projection * gl_Vertex;
5 }
```

---

### 1.1 Uniform Storage Limits

Uniforms are shared by all program units. If there is a uniform `light_position` in the vertex shader and a uniform of the same name in the fragment shader, they will have the same value. By implication, they must also have the same type. A program will fail to link if uniforms with the same name in different compilation units have different types.

Each program unit has a limited amount of storage for uniforms. Classically, uniforms have been implemented as a small, read-only register array on the GPU. GPUs designed in this way have very limited uniform storage, which may be as small as 512 floating-point values. GPUs a generation or two later are capable of storing uniforms in much larger off-chip buffers. GPUs designed in this way can typically support 4,096 floating-point values.

The actual number of uniforms available for each program unit is queryable. Limits are queried by calling `glGetIntegerv` with one of the query enumerants listed in Table 1. Each OpenGL version specifies a minimum value for each of these queries<sup>1</sup>. Note that OpenGL 2.x implementations are only required to support 64 floating-point components for fragment shader uniforms.

---

**Table 1** Uniform Limits

Program Unit	Query Enumerant	OpenGL 2.x Minimum	OpenGL 3.x Minimum
Vertex	<code>GL_MAX_VERTEX_UNIFORM_COMPONENTS</code>	512	1024
Fragment	<code>GL_MAX_FRAGMENT_UNIFORM_COMPONENTS</code>	64	1024

---

OpenGL 3.0 introduced a new mechanism, called *uniform buffer objects*, that allows shaders to access even larger amounts of uniform storage. Uniform buffer objects will be covered in a later chapter.

Each GLSL type has a specific number of components that it occupies. These values are listed in Table 2. Many GPUs only allow indexed array access to arrays that are stored internally as `vec4`. As a result, each element in an array can occupy as many as four components. In practice this can depend on the GPU and on the manner in which the array is accessed in all shaders. For example, `data_blob` in Figure 2 would *likely* only use 16 components. If that vertex shader were linked with the fragment shader in Figure 3, `data_blob` would almost certainly require 32 components. This means

---

<sup>1</sup> Minimum requirements for implementation limits queried by `GL_MAX_` enumerants are often called "minimum maximums".

that a particular uniform used in a shader can use a different number of components in each program where it is linked.

**Table 2** Uniform Components

GLSL Type	Components	Components in Array
float, int, or bool	1	4
vec2, ivec2, or bvec2	2	4
vec3, ivec3, or bvec3	4	4
vec4, ivec4, or bvec4	4	4
mat2	4	4
mat3	12	12
mat4	16	16

**Figure 2** Vertex shader with a uniform array

```

1 uniform vec2 data_blob[8];
2 void main(void)
3 {
4     if (gl_Vertex.x < 6.0) {
5         gl_Position = vec4(data_blob[4], data_blob[1]);
6     } else {
7         gl_Position = vec4(data_blob[2], data_blob[7]);
8     }
9 }
```

**Figure 3** Fragment shader with a uniform array

```

1 uniform vec2 data_blob[8];
2 void main(void)
3 {
4     gl_Color = vec4(data_blob[int(7.0 * abs(sin(gl_FragPos.x)))]],
5                   data_blob[int(7.0 * abs(sin(gl_FragPos.y)))]);
6 }
```

Notice also that a `vec4` can occupy four components, and a `mat3` can occupy 12 components. This is again due to the `vec4`-centric nature of many GPUs. In some cases the compiler may place a non-array scalar and a non-array `vec3` in the single `vec4`.

Implementations are required to implement certain procedures for packing data and eliminating unused uniforms. As with vertex shader inputs (see Chapter 3), the requirements for detecting and eliminating unused uniforms only require that the compiler perform simple static analysis of the shader. This means that a uniform may be eliminated if it does not appear as an operand in the shader text. This means that `should_eliminate` in Figure 4 may not be eliminated! Components of uniforms that do not appear in the shader text also will be eliminated. In Figure 5 the `z` component of `offset` will be eliminated.

Compilers are not required to adhere to the rules of uniform packing for uniform arrays that are accessed with non-constant indexes. Even if the `z` of the 5th element of some `vec4` array is never accessed, the compiler is not required to eliminate it. At the same time, the compiler is also not required to not eliminate it.

## 1.2 Setting Uniform Values

Setting the value of a uniform is a two step process. After a program has been linked, uniforms are assigned numeric locations. The application must query the location associated with a particular name. The location is queried by calling `glGetUniformLocation`. The `program` is the name of the program object, and `name` is the name of the uniform whose location is to be queried. The value returned is the position of the uniform in the program. The value -1 will be returned in the case of an error.

**Figure 4** Shader with unexecuted use of a uniform

```

1 uniform vec4 should_eliminate;
2 uniform mat4 mvp;
3 void main(void)
4 {
5     if (false) {
6         gl_Position = mvp * should_eliminate;
7     } else {
8         gl_Position = mvp * gl_Vertex;
9     }
10 }

```

**Figure 5** Shader with unused uniform component

```

1 uniform vec4 offset;
2 uniform mat4 mvp;
3 void main(void)
4 {
5     gl_Position = mvp * gl_Vertex + vec4(offset.xyw, 0.0);
6 }

```

```
GLint glGetUniformLocation(GLuint program, const GLchar *name);
```

After determining the location, the value for the location can be set. There are a variety of functions for setting uniform values based on the uniform's type. For now only the three most common variants will be discussed. `glUniform4f` and `glUniform4fv` can be used to set the value of a `vec4` varying, and `glUniformMatrix4fv` can be used to set the value of a `mat4` varying.

```

void glUniform4f(GLint location,
                GLfloat v0, GLfloat v1, GLfloat v2, GLfloat v3);

void glUniform4fv(GLint location, GLsizei count, const GLfloat *value);

void glUniformMatrix4fv(GLint location, GLsizei count,
                       GLboolean transpose, const GLfloat *value);

```

For all three functions, `location` is the uniform location returned by `glGetUniformLocation`. Notice that there is no explicit program object parameter. Instead the program most recently bound with `glUseProgram` is used. For `glUniform4f`, the remaining parameters are the values to set for the components of the `vec4` uniform.

For the vectored functions (i.e., the ones with `v` at the end of the name), the `count` parameter is the number of uniforms to set. This is used when setting values for a uniform array. For non-arrays, 1 should be used. The `value` is a pointer to the data to copy to the uniform. It is important that this point to the correct number of elements.

For `glUniformMatrix4fv` the `transpose` parameter specifies whether the supplied data is in column-major or row-major ordering. If `GL_FALSE` is passed, the first four elements of `value` make up the first column of the matrix. If `GL_TRUE` is passed, the first four elements of `value` make up the first row of the matrix.

Figure 6 shows how the `model_view_projection` uniform of Figure 1 could be set.

### 1.3 Pre-initialized Uniforms

In addition to being set through the `glUniform` family of functions, uniforms can be initialized in the shader itself<sup>2</sup>. This is common practice when a uniform has a typical value that only needs to be changed in infrequent circumstances.

<sup>2</sup> GLSL 1.20 or later is required.

**Figure 6** Set a matrix uniform

---

```

1 float.mvp[16];
2 GLint.mvp_location =
3     glGetUniformLocation(program, "model_view_projection");
4 if (mvp_location < 0)
5     /* ... error path ... */
6
7 calculate.mvp_matrix(mvp);           // Matrix in column-major order.
8
9 glUseProgram(program);
10 glUniformMatrix4fv(mvp_location, 1, GL_FALSE,.mvp);

```

---

The fragment shader in Figure 7 initializes *color* to the vector { 0.0, 1.0, 0.0, 1.0 }. When *prog1* in Figure 8 is used, *color* will have the value initialized in the shader. However, when *prog2* is used, *color* will have the value { 1.0, 0.0, 0.0, 1.0 }.

**Figure 7** Initializing a uniform

---

```

1 uniform vec4 color = vec4(0.0, 1.0, 0.0, 1.0);
2 void main(void)
3 {
4     gl_FragColor = color * gl_Color;
5 }

```

---

**Figure 8** Overriding a pre-initialized uniform

---

```

1 GLuint.vs = glCreateShader(GL_VERTEX_SHADER);
2 glShaderSource(vs, 1, (const GLchar **) &code_from_figure_7, NULL);
3 glCompileShader(vs);
4
5 prog1 = glCreateProgram();
6 glAttachShader(prog1, vs);
7 glAttachShader(prog1, fsl);       // fsl is initialized elsewhere
8 glLinkProgram(prog1);
9
10 prog2 = glCreateProgram();
11 glAttachShader(prog2, vs);
12 glAttachShader(prog2, fsl);       // fsl is initialized elsewhere
13 glLinkProgram(prog2);
14
15 GLint.color_location = glGetUniformLocation(prog2, "color");
16 glUseProgram(prog2);
17 glUniform4f(color_location, 1.0, 0.0, 0.0, 1.0);

```

---

## 1.4 Complex Datatypes

As seen in the preceding sections, uniforms can be any of the scalar, vector, and matrix built-in types. Uniforms can also be structures, arrays of built-in types, and arrays of structures.