

point when the program is executed. Lastly, **args** is a null-terminated list of null-terminated strings that is passed as an argument to the compiler.

Loading a Program

After you compile a program, you need to pass the resulting object code to the 3D API that you're using. For this, you need to invoke the Cg runtime's API-specific functions.

The Direct3D-specific functions require the Direct3D device structure in order to make the necessary Direct3D calls. The application passes it to the runtime using the following call:

```
cgD3D9SetDevice(Device);
```

You must do this every time a new Direct3D device is created, typically only at the beginning of the application.

You can then load a Cg program in this way for the Direct3D 9 Cg runtime:

```
cgD3D9LoadProgram(program, CG_FALSE, 0);
```

or this way for the Direct3D 8 Cg runtime:

```
cgD3D8LoadProgram(program, CG_FALSE, 0, 0, vertexDeclaration);
```

The parameter **vertexDeclaration** is the Direct3D 8 vertex declaration array that describes where to find the necessary vertex attributes in the vertex streams. (See [“Expanded Interface Program Execution”](#) on page 103 for the details on the arguments to **cgD3D8LoadProgram()** and **cgD3D9LoadProgram()**).

In OpenGL, the equivalent call is

```
cgGLLoadProgram(program);
```

Modifying Program Parameters

The runtime gives you the option of modifying the values of your program parameters. The first step is to get a handle to the parameter:

```
Cgparameter myParameter = cgGetNamedParameter(
    program, "myParameter");
```

The variable **myParameter** is the name of the parameter as it appears in the program source code.

The second step is to set the parameter value. The function used depends on the parameter type.

Here is an example in OpenGL:

```
cgGLSetParameter4fv(myParameter, value);
```

Here is the same example in Direct3D:

```
cgD3D9SetUniform(myParameter, value);
```

Numeric parameters may also be set using core Cg runtime calls, such as:

```
cgSetParameterValuefr(myParameter, 4, value);
```

These function calls assign the four floating-point values contained in the array **value** to the parameter **myParameter**, which is assumed to be of type **float4**.

In both APIs, there are variants of these calls to set matrices, arrays, textures, and texture states. The core Cg runtime provides variants of these calls to set the value of numeric parameters, including scalars, vectors, arrays, and structures. The graphics API-specific runtimes must be used to set API-specific values, such as sampler handles.

Executing a Program

Before you can execute a program in OpenGL, you must enable its corresponding profile:

```
cgGLEnableProfile(CG_PROFILE_ARBVP1);
```

In Direct3D, nothing explicitly needs to be done to enable a specific profile.

Next, you bind the program to the current state. This means that in subsequent drawing calls the program is executed for every vertex in the case of a vertex program and for every fragment in the case of a fragment program.

Here's how to bind a program in OpenGL:

```
cgGLBindProgram(program);
```

Here's how to bind a program in Direct3D:

```
cgD3D9BindProgram(program);
```

You can only bind one vertex and one fragment program at a time for a particular profile. Therefore, the same vertex program is executed until another vertex program is bound. Similarly, the same fragment program is executed as long as no other fragment program is bound.

In OpenGL, you disable profiles by the following call:

```
cgGLDisableProfile(CG_PROFILE_ARBVP1);
```

Disabling a profile also disables the execution of the corresponding vertex or fragment program.

Releasing Resources

When your application is ready to close, it is good programming practice to free resources that you've acquired.

Because the Direct3D runtime keeps an internal reference to the Direct3D device, you must tell it to release this reference when you are done using the runtime. This is done with the following call:

```
cgD3D9SetDevice(0);
```

To free resources allocated for a program, call this function:

```
cgDestroyProgram(program);
```

To free resources allocated for a context, use this function:

```
cgDestroyContext(context);
```

Note that destroying a context destroys all the programs it contains as well.

Core Cg Runtime

The core Cg runtime provides all the functions necessary to manage Cg programs from within the application. It makes no assumption about which 3D API the applications uses, so that any application could easily ignore the API-specific Cg runtime libraries and content itself with the core Cg runtime.

The core Cg runtime is built around three main concepts: context, program, and parameter, which are represented by the **CGcontext**, **CGprogram**, and **CGparameter** object types. Those concepts are hierarchically related one to each other: a program has several parameters, a context contains several programs and shared parameters, and the application can define several contexts.

The next sections describe these three basic object types and the runtime entry points that operate on them. The three object types have some points in common:

- ❑ The use of **CGbool**, which is an integer type equal to either **CG_TRUE** or **CG_FALSE**
- ❑ The use of **CGenum**, which is an enumerate type used to specify various enumerate values that are not necessarily related
- ❑ The convention that functions that return a value of type **CGcontext**, **CGprogram**, **CGparameter**, or **const char*** indicate failure by returning zero

Core Cg Context

The Cg runtime provides functions for creating, destroying, and querying contexts.

Context Creation and Destruction

Programs can only be created as part of a context that acts as a program container. A context is created by calling **cgCreateContext()**:

```
CGcontext cgCreateContext();
```

A context is destroyed by **cgDestroyContext()**:

```
void cgDestroyContext(CGcontext context);
```

cgDestroyContext() deletes all data associated with the context, including all programs it contains. **cgDestroyContext()** should be called before destroying any associated OpenGL context or Direct3D device.

Context Query

To check whether a context handle references a valid context or not, use **cgIsContext()**:

```
CGbool cgIsContext(CGcontext context);
```

Core Cg Program

There are Cg functions for creating, destroying, iterating over, and querying programs.

Program Creation and Destruction

A program is created by calling either **cgCreateProgram()**:

```
CGprogram cgCreateProgram(CGcontext    context,
                           CGLenum      programType,
                           const char*  program,
                           CGprofile    profile,
                           const char*  entry,
                           const char** args);
```

or **cgCreateProgramFromFile()**:

```
CGprogram cgCreateProgramFromFile(CGcontext context,
                                   CGLenum      programType,
                                   const char*  program,
                                   CGprofile    profile,
                                   const char*  entry,
                                   const char** args);
```

These functions create a program object, add it to the specified context and compile the associated source code. For both of them,

- ❑ **context** is a valid context handle.
- ❑ **profile** is an enumerant specifying the profile to which the program must be compiled.
- ❑ **entry** is the name of the function that must be considered as the main entry point by the compiler. If the value is zero, the name **main** is used.
- ❑ **args** is a pointer to a null-terminated array of null-terminated strings that are passed as arguments to the compiler. The pointer may itself be null.

The only difference between the two functions is how **program** is interpreted. For **cgCreateProgramFromFile()**, **program** is a string containing the name of a file containing source code; for **cgCreateProgram()**, **program** directly contains source code. If the enumerant **programType** is equal to **CG_SOURCE**, the source code is Cg source code; if it is equal to **CG_OBJECT**, the source code is precompiled object code and does not require any further compilation.

The **CGprogram** handle returned by **cgCreateProgramFromFile()** is valid if it is different from zero, which means that the program has been successfully created and compiled. The program is destroyed by passing its handle to **cgDestroyProgram()**:

```
void cgDestroyProgram(CGprogram program);
```

The Cg runtime allows for either automatic or manual compilation of programs. Compilation of a program is required before the program may be used when drawing. As such, program compilation is necessary sometime after the program is first created, or whenever it enters an uncompiled state. A program may enter an uncompiled state for a variety of reasons, including

- ❑ Changing variability of parameters
Parameters may be changed from uniform variability to literal variability (compile time constant). See the **cgSetParameterVariability** manual page for more information.
- ❑ Changing value of literal parameters
Changing the value of a literal parameter will require recompilation since the value is used at compile time. See the **cgSetParameter** and **cgSetMatrixParameter** manual pages for more information.
- ❑ Resizing unsized arrays
Changing the length of a parameter array may require recompilation depending on the capabilities of the program profile. See the

cgSetArraySize and **cgSetMultiDimArraySize** manual pages for more information.

- ❑ **Connecting structures to interface parameters**
Structure parameters can be connected to interface program parameters to control the behavior of the program. Changing these connections requires recompilation on all current profiles. See the **cgConnectParameter** manual page and the Interfaces section of this document for more details.

When a program enters an uncompiled state, it is automatically unloaded and unbound. In order to be used again, the program must be recompiled (either automatically or manually—see the following), and then reloaded and rebound.

Compilation can be performed manually by the application via

```
cgCompileProgram(CGprogram program);
```

or automatically by the runtime.

Compilation behavior is controlled via

```
void cgSetAutoCompile(CGcontext ctx, CGenum flag);
```

Here, *flag* may be one of the following enumerants:

- ❑ **CG_COMPILE_MANUAL**
In this mode, the application is responsible for manually compiling a program. The application may check to see if a program requires recompilation with the entry point **cgIsProgramCompiled**. The program may then be compiled via **cgCompileProgram()**. This mode provides the application with the most control over how and when program recompilation occurs.
- ❑ **CG_COMPILE_IMMEDIATE**
In this mode, the Cg runtime will force compilation automatically and immediately when a program enters an uncompiled state, or when the program is first created. This is the default mode.
- ❑ **CG_COMPILE_LAZY**
This mode is similar to **CG_COMPILE_IMMEDIATE**, but will delay program compilation until the program object code is needed. The advantage of this method is the reduction of extraneous recompilations. The disadvantage is that compile time errors will not be encountered when the program enters an uncompiled state, but will instead be encountered at some later time (most likely when the program is loaded or bound).

A call to **cgIsProgramCompiled()** determines whether a program needs to be recompiled:

```
CGbool cgIsProgramCompiled(CGprogram program);
```

To recompile a program, use **cgCompileProgram()**:

```
cgCompileProgram(CGprogram program);
```

Program Iteration

The programs within a context are sequentially ordered and can be iterated over by using **cgGetFirstProgram()** and **cgGetNextProgram()**:

```
CGprogram cgGetFirstProgram(CGcontext context);  
CGprogram cgGetNextProgram(CGprogram program);
```

The first program of the sequence is retrieved by **cgGetFirstProgram()**. If the context is invalid or does not contain any program, the function returns zero. Given a program, **cgGetNextProgram()** returns the program immediately next in the sequence, or zero if there is none. Here is how those two functions would typically be used given a valid context named **context**:

```
CGprogram program = cgGetFirstProgram(context);  
while (program != 0) {  
    /* Here is the code that handles the program */  
    program = cgGetNextProgram(program);  
}
```

Nothing is guaranteed regarding the order of the programs in the sequence or how **cgGetFirstProgram()** and **cgGetNextProgram()** behave when programs are created or destroyed during iteration.

Program Query

Program queries encompass validity, compilation results, and attributes.

Program Validity

Use **cgIsProgram()** to check whether a program handle references a valid program:

```
CGbool cgIsProgram(CGprogram program);
```

Compilation Result

You can query the result of the compilation resulting from the last call to **cgCreateProgram()** for a given context by using **cgGetLastListing()**:

```
const char* cgGetLastListing(CGcontext context);
```

If no call to **cgCreateProgram()** has been made for the context, **cgGetLastListing()** returns zero. Otherwise, it returns a string containing the output you would typically get from the command-line version of the compiler.

Program Attributes

To retrieve the context the program belongs to, use

cgGetProgramContext():

```
CGcontext cgGetProgramContext(CGprogram program);
```

Retrieving the profile the program has been compiled to is done with

cgGetProgramProfile():

```
CGprofile cgGetProgramProfile(CGprogram program);
```

The function pair **cgGetProfile()** and **cgGetProfileString()** allows you to find the correspondence between a profile enumerator and its corresponding string:

```
CGprofile cgGetProfile(const char* profileString);  
const char* cgGetProfileString(CGprofile profile);
```

If the string passed to **cgGetProfile()** does not correspond to any profile, **CG_PROFILE_UNKNOWN** is returned.

The function **cgGetProgramString()** retrieves various strings related to the program depending on the value of the enumerator **stringType**:

```
const char* cgGetProgramString(CGprogram program,  
                                CGenum stringType);
```

The variable **stringType** can have any of these values:

- ❑ **CG_PROGRAM_SOURCE**: The original Cg source program is returned.
- ❑ **CG_PROGRAM_ENTRY**: The main entry point of the Cg source program is returned.
- ❑ **CG_PROGRAM_PROFILE**: The profile string is returned.
- ❑ **CG_COMPILED_PROGRAM**: The resulting compiled program is returned.

Core Cg Parameters

Cg parameters fall into three broad categories: program parameters, effect parameters, and shared parameters.

Program parameters are associated with Cg programs. A parameter that is declared as part of the program's entry point belongs to the program's

namespace. A parameter that is declared globally in the file scope of the Cg program belongs to the program's global namespace.

Effect parameters are associated with Cg Effects. See the Introduction to CgFX chapter for more information on managing effect parameters.

Shared parameters are associated with Cg contexts. See [“Shared Parameters” on page 59](#), for more details.

Cg functions exist for retrieving, creating, and querying program parameters.

Program Parameter Retrieval

Parameters associated with Cg programs may be retrieved iteratively or directly.

Iteration

A program has a sequence of parameters that can be iterated over by using **cgGetFirstParameter()** and **cgGetNextParameter()**:

```
CGparameter cgGetFirstParameter(CGprogram program,
                                CGenum namespace);
CGparameter cgGetNextParameter(CGparameter parameter);
```

A call to **cgGetFirstParameter()** returns the first parameter of the sequence. If the program is invalid or does not contain any parameter, the call returns zero. Given a parameter, **cgGetNextParameter()** returns the parameter immediately next in the sequence or zero if there is none. The **namespace** argument of **cgGetFirstParameter()** specifies the name space of the parameters returned by this function and subsequent calls to **cgGetNextParameter()**. Every parameter belongs to a particular name space that defines its scope. When **CG_GLOBAL** is specified, the program's global parameters (i.e., those parameters that are in the file scope of the program's entry point), are iterated over. When **CG_PROGRAM** is specified, the parameters specified in the program's entry point declaration are iterated over.

Here is how those two functions would typically be used given a valid program called **program**:

```
CGparameter parameter = cgGetFirstParameter(program,
                                             CG_PROGRAM);
while (parameter != 0) {
    /* Here is the code that handles the parameter */
    parameter = cgGetNextParameter(parameter);
}
```

These functions don't provide access to the fields of a structure parameter (type **CG_STRUCT**) or the elements of an **array** parameter (type **CG_ARRAY**). In other words, if a **struct** or **array** parameter is declared, these entry points return will return a handle to the **struct** or **array** itself.

One way to access the fields of a structure is to use **cgGetFirstStructParameter()** along with **cgGetNextParameter()**:

```
CGparameter cgGetFirstStructParameter(CGparameter parameter);
```

If **parameter** is not of type **CG_STRUCT**, **cgGetFirstStructParameter()** returns zero.

Similarly, to get access to the elements of an array, you can use **cgGetArrayDimension()**, **cgGetArraySize()**, **cgGetArrayParameter()**, and **cgGetNextParameter()**:

```
int cgGetArrayDimension(CGparameter parameter);  
int cgGetArraySize(CGparameter parameter, int dimension);  
CGparameter cgGetArrayParameter(CGparameter parameter,  
int index);
```

These three functions return 0 if **parameter** is not of type **CG_ARRAY**. Function **cgGetArrayDimension()** gives the dimension of the array. It returns 1 for **float4 array[10]**, 2 for **float4 array[10][100]**, and so on. Next, **cgGetArraySize()** gives the size of every dimension. For example, for **float4 array[10][100]**, **cgGetArraySize(array, 0)** returns 10 and **cgGetArraySize(array, 1)** returns 100. An array, **anArray**, has **cgGetArraySize(anArray, 0)** elements. If its dimension is greater than one, those elements are themselves arrays.

Here is how these iteration functions could be used given a valid program named **program**:

```
void IterateProgramParameters(CGprogram program) {  
    RecurseProgramParameters(cgGetFirstParameter(program,  
                            CG_PROGRAM));  
}  
  
void RecurseProgramParameters(CGparameter parameter) {  
    if (parameter == 0)  
        return;  
    do {  
        switch(cgGetParameterType(parameter)) {  
            case CG_STRUCT:  
                RecurseProgramParameters(  
                    cgGetFirstStructParameter(parameter));  
                break;  
        }  
    }
```

```

    case CG_ARRAY:
        int arraySize = cgGetArraySize(parameter, 0);
        for (int i = 0; i < arraySize; ++i)
            RecurseProgramParameters(
                cgGetArrayParameter(parameter, i));
        break;
    default:
        /* Here is the code that handles the parameter */
        break;
}
} while((parameter = cgGetNextParameter(parameter)) != 0);
}

```

In practice, it is usually simpler to iterate over all of the “leaf” parameters (that is, non-aggregate parameters) directly using

cgGetNextLeafParameter():

```

CGparameter cgGetFirstLeafParameter(CGprogram program,
                                     CGenum namespace);
CGparameter cgGetNextLeafParameter(CGparameter parameter);

```

These functions iterate through all the simple parameters, including structure fields and array elements that serve as inputs to the program. Nothing is guaranteed regarding the order of the parameters in the sequence.

Direct Retrieval

Any parameter of a program can also be retrieved directly by using its name with **cgGetNamedParameter()**:

```

CGparameter cgGetNamedProgramParameter(CGprogram program,
                                       CGenum namespace,
                                       const char* name);

```

Here, *namespace* may be either **CG_GLOBAL** or **CG_PROGRAM**, as above. If the program has no parameter corresponding to *name*, **cgGetNamedParameter()** returns zero.

The Cg syntax is used to retrieve structure fields or array elements. Let’s take the following code snippet as an example:

```

struct FooStruct {
    float4 A;
    float4 B;
};
struct BarStruct {
    FooStruct Foo[2];
};

```

```
void main(BarStruct Bar[3]) {
    // ...
}
```

The following are valid names for retrieving the corresponding parameter:

```
"Bar"
"Bar[1]"
"Bar[1].Foo"
"Bar[1].Foo[0]"
"Bar[1].Foo[0].B"
```

Parameter Values

The core Cg runtime provides a number of entry points for setting and retrieving parameter values. In addition, the graphics-API-specific Cg runtimes provide additional entry points for managing parameter values.

When managing numeric parameters, choosing which set of entry points to use is largely a matter of programmer preference. In some circumstances, it may be slightly more efficient to use the core Cg runtime entry points. However, parameters that hold graphics-API-specific quantities, such as sampler handles, must be set using the API-specific entry points. The API-specific entry points must be used because the core Cg runtime, which is graphics-API-agnostic, provides no such entry points.

The most often-used parameter value routines are used to set and get a parameter's current values. A parameter's current value is initialized to any default value assigned in the Cg source, or 0 otherwise. The current value of a numeric parameter can be queried using the family of entry points:

```
int cgGetParameterValue{i,f,d}{r,c}(CGparameter param,
                                     int nvals, type *v);
```

The given parameter must be a scalar, vector, matrix, or an (possibly-multidimensional) array of scalars, vectors, or matrices. There are versions of each function to retrieve the values into an int, float, or double buffer; these are signified by the *i*, *f*, and *d* in the entry point name, respectively. Similarly, there are versions of each function that retrieve any matrices in the given parameter in row-major or column-major order. These are specified using *r* or *c*, respectively. At most, *nvals* values will be copied into the given array, *v*. The total number of values copied into *v* is returned.

For example, `cgGetParameterValueic()` retrieves the values of the given parameter into the supplied array of integer data, and copies matrix data in column-major order. The total number of values associated with a given

parameter, and hence the required length of the given array, can be computed using the core Cg runtime:

```
int nrows = cgGetParameterRows(param);
int ncols = cgGetParameterColumns(param);
int asize = cgGetArrayTotalSize(param);
int ntotal = nrows*ncols;
if (asize > 0) ntotal *= asize;
```

A similar family of entry points exist for setting a parameter's values:

```
void cgSetParameterValue{i,f,d}{r,c}(CGparameter param,
                                     int nvals, type *v);
```

The entry points in this family are identical to those of the **cgGetParameterValue** family. The total number of values in a parameter may be computed as above. If **nvals** is less than the total size of the parameter, an error is generated.

The core Cg runtime also allows the application to query a parameter's default values:

```
const double* cgGetParameterValues(CGparameter parameter,
                                    CGenum valueType,
                                    int* numberOfValuesReturned);
```

This entry point retrieves the parameter's default value if **valueType** is equal to **CG_DEFAULT**. The components of the value are returned in row-major order as a pointer to an array containing type **double** elements. The number of components available in the array is returned in **numberOfValuesReturned**. Function **cgGetParameterValues()** can also be used to retrieve a parameter's constant values, but this functionality is rarely used; see the corresponding manual page for more details.

Shared Parameters

The core Cg runtime supports the creation of instances of any type of concrete parameter (e.g., built-in types, user-defined structures) within a Cg context. A parameter instance may be connected to any number of compatible parameters, including any program or effect parameter within the context.

When an instance is connected to another parameter, the second parameter will inherit its values from the instance. Furthermore, if the variability of the second parameter has not been explicitly set by a call to **cgSetParameterVariability()**, its variability will also be inherited from the instance.

The ability to create and easily manage shared, context-global parameters provides a powerful means for creating parameter trees, and for sharing data and user-defined objects between multiple Cg programs or effects.

Shared Parameter Creation

Shared parameters are associated with a **CGcontext**. They may be created with the following entry points:

```
CGparameter cgCreateParameter(CGcontext ctx, CGtype type);  
CGparameter cgCreateParameterArray(CGtype type, int length);  
CGparameter cgCreateParameterMultiDimArray(CGtype type,  
int dim, int *lengths);
```

Only parameters of concrete types may be created. In particular, parameters of abstract interface types may not be created. By default, a created parameter has uniform variability and undefined values.

Shared Parameter Deletion

Shared parameters may be deleted using

```
Void cgDeleteParameter(CGparameter param);
```

When a shared parameter is deleted, all parameters connected to it are disconnected, and vice-versa.

Connecting Parameters

Once created, a shared parameter may be connected to any number of program, effect, or shared parameters using

```
void cgConnectParamteer(CGparameter source, CGparameter sink);
```

where **source** is the shared parameter, and **sink** is the target parameter that will inherit the shared parameter's values.

Once a parameter has had a **source** connected to it, its value should no longer be set directly. Instead, its value can be set indirectly by setting the value of the associated **sink**.

A parameter that has been connected to a shared **source** parameter may be disconnected using

```
Void cgDisconnectParameter(param);
```

Shared Parameters and Interfaces

Using Cg, it is possible to create families of code "modules" that share a common interface, each member of which has a different implementation. This ability makes it easy for applications to construct material trees on the

fly, to change the number or type of texture maps applied to an object at application runtime, and so on.

Specifying which particular implementation of an interface to use is accomplished through “connecting” parameters. In particular, a shared instance of a **struct** that implements the interface is created by the application. This shared instance is then connected to the interface parameter. The act of connecting the parameters causes the interface parameter to inherit the shared parameter’s implementation of the interface. This process can be thought of as implementing compile-time polymorphism.

It is legal to connect a shared parameter of a user-defined structure type to an interface parameter, as long as the structure type implements that interface type. At runtime, the entry point’s **cgIsParentType**, coupled with **cgGetParameterNamedType**, can be used to determine type parenthood.

When a structure parameter is connected to an interface parameter, copies of any child (that is, member) variables associated with the **source** structure parameter are automatically created as children of the **sink** parameter. Under most circumstances, these member variable copies can be ignored by the application, since their values and variability are automatically set by the Cg runtime. However, in some situations it may be useful to query a “sink-side” member parameter for its underlying resource, for example.

A shared instance of a structure whose type is defined in one Cg program or effect may be connected to parameters of other programs or effects, provided that the entities involved define the **source** structure types and destination interface types equivalently. See [“Parameter Type Equivalency” on page 65](#) or more details. If the types are not equivalent, **cgConnectParameter()** generates a runtime error.

The following example illustrates structure-to-interface connection by creating three programs, all of which define a type named **Foo**, with one program’s definition differing from the others:

```
interface MyInterface {
    float Val(float x);
};
struct MyStruct : MyInterface {
    float Scale;
    float Val(float x) { return(Scale * x);
};
float4 main(MyInterface foo) : COLOR {
    return(foo.Val(.2).xxxx);
}
```

Listing 1: Cg Program 1

```

interface MyInterface {
    float Val(float x);
};
struct MyStruct : MyInterface {
    float Scale;
    float Val(float x) { return(Scale * x);
};
float4 main(MyInterface foo) : COLOR {
    return(foo.Val(.3).xxxx);
}

```

Listing 2: Cg Program 2

```

interface MyInterface {
    half Val(half x);
};
struct MyStruct : MyInterface {
    float Scale;
    half Val(half x) { return(Scale * x);
};
float4 main(MyInterface foo) : COLOR {
    return(foo.Val(.5).xxxx);
}

```

Listing 3: Cg Program 3

Notice that both Cg Program 1 and Cg Program 2 define the **Val()** method of the **MyInterface** and **MyStruct** types using the **float** type, whereas Cg Program 3 does so using the **half** type. As a result, the **MyInterface** and **MyStruct** types defined in Cg Program Three are not equivalent to types in the other two programs, even though the types have the same names.

The following C program creates all three of the above Cg programs and connects shared parameter instances to their input parameters:

```

static CGprogram CreateProgram(const char *program_str) {
    return cgCreateProgram(Context, CG_SOURCE,
                          program_str, CG_PROFILE_ARBFP1,
                          "main", NULL);
}
int main(int argc, char *argv[]) {
    CGContext Context;
    CGprogram Program1, Program2, Program3;
    CGparameter ms1, ms3;
    // Disable automatic compilation, since the
    // programs cannot be compiled until concrete structs
    // are connected to each program's interface parameters.
}

```

```

Context = cgCreateContext();
cgSetAutoCompile(Context, CG_COMPILE_MANUAL);
// Create the programs
Program1 = CreateProgram(Program1String);
Program2 = CreateProgram(Program2String);
Program3 = CreateProgram(Program3String);
// Create two shared parameters,
// one of the MyStruct type from Program1, and
// one of the MyStruct type from Program3.
ms1 = cgCreateParameter(cgGetNamedUserType(Program1,
                                             "MyStruct"));
ms3 = cgCreateParameter(cgGetNamedUserType(Program3,
                                             "MyStruct"));
/* Connect the same shared parameter to Program1 and
   Program2 */
cgConnectParameter(Foo1, cgGetNamedParameter(Program1,
                                             "foo"));
cgConnectParameter(Foo1, cgGetNamedParameter(Program2,
                                             "foo"));
// The following would generate an error because the type
// of the Foo1 parameter is not equivalent to type
// "MyStruct" from Program3.
// cgConnectParameter(ms1,
//                    cgGetNamedParameter(Program3, "foo"));
cgConnectParameter(ms3, cgGetNamedParameter(Program3,
                                             "foo"));
// Now we can compile all three programs.
cgCompileProgram(Program1);
cgCompileProgram(Program2);
cgCompileProgram(Program3);
// ... and so on ...
}

```

Parameter Properties

Parameter properties encompass validity, references, size, and other attributes.

Parameter Type

The Cg language defines a number of built-in parameter types, such as **float4**, **int3x3**, and so on. In addition, user-defined types may be specified in a program when declaring structure and interface types. For example, if the following Cg code is included in the source to a **CGprogram** created via **cgCreateProgram()**, the types **MyInterface** and **MyStruct** will be added to the resulting **CGprogram**.

```
interface MyInterface {
    float SomeMethod(float x);
};
struct MyStruct : MyInterface {
    float Scale;
    SomeMethod(float x) {
        return(Scale * x);
    }
};
```

In order to obtain the unique enumerant associated with a parameter's type, the following entry point should be used

```
CGtype cgGetParameterNamedType(CGparameter param);
```

The **CGtype** associated with a named user-defined type in a program can be retrieved using

```
CGtype cgGetNamedUserType(CGhandle handle, const char *name);
```

Here, **handle** can be either a **CGprogram** or a **CGeffect**.

The **struct** types can implement a given interface. In such a case, the indicated interface is known as a *parent* type of the **struct** type. In the example above, **MyStruct** has a single parent type, **MyInterface**. The parent types of a given named type may be obtained with the following entry points:

```
int cgGetNumParentTypes(CGtype type);  
CGtype cgGetParentType(CGtype type, int index);
```

Note that the Cg language specification currently makes it impossible for a **struct** type to have more than a single parent type.

All of the user-defined types associated with a program may be obtained with the following entry points:

```
int cgGetNumUserTypes(CGprogram program);  
CGtype cgGetType(CGprogram program, int index);
```

Note that the runtime treats interface program parameters as if they were structure parameters with no concrete data or function members.

In older applications that use the Cg runtime, you may encounter the deprecated entry point:

```
CGtype cgGetParameterType(CGparameter parameter);
```

This entry point differs from **cgGetNamedUserType()** in that it always returns **CG_STRUCT** for any **struct** parameter, rather than returning the enumerant associated with the user-defined type of the **struct**.

The name associated with a given type enumerant can be queried using

```
const char* cgGetTypeString(CGtype type);
```

If the string passed to **cgGetType()** does not correspond to any type, **CG_UNKNOWN_TYPE** is returned.

Function **cgGetParameterBaseType()** returns the basic type of vector matrix and matrix parameters. For example, given a **float4x4** parameter, **cgGetParameterBaseType()** returns the **CG_FLOAT** type. Similarly, given a multidimensional array of **float4x4s**, it also returns **CG_FLOAT**.

It is also possible to determine the general class of the type of a parameter:

```
CGparameterclass cgGetParameterClass(CGparameter param);
```

It returns one of the following enumerated values:

CG_PARAMETERCLASS_UNKNOWN	CG_PARAMETERCLASS_SCALAR
CG_PARAMETERCLASS_VECTOR	CG_PARAMETERCLASS_OBJECT
CG_PARAMETERCLASS_MATRIX	CG_PARAMETERCLASS_STRUCT
CG_PARAMETERCLASS_ARRAY	

Parameter Type Equivalency

If a program containing a user-defined type is created in a context that already contains another program or effect that defines a user type with the same name, the two type definitions are compared. If both type definitions are found to be equivalent, the **CGtype** enumerant associated with the user type in the new program will be identical to that of the identical user type in the existing program or effect. If the types are not equivalent, the new type will be assigned a unique **CGtype**. In this way, type equivalency of

parameters shared between multiple programs and effects can be assured simply by comparing **CGtype** enumerants.

In order for two types to be considered equivalent, they must meet the following requirements:

- ❑ The type names must match.
Both types must have the exact same name.
- ❑ The parent types, if any, must match.
If the type is a structure, both must either not implement an interface, or both implement interfaces that are type-equivalent.
- ❑ The member variables and methods must match.
They must both have the exact same member variables and methods. The order and name of the variables must match exactly, and the order and name of the methods must match. The signature of the methods, including argument and return types, must be identical.

Type equivalency is useful when using shared parameters instances with multiple programs by connecting them with **cgConnectParameter()**.

Parameter Validity

The function **cgIsParameter()** allows you to check whether a parameter handle references a valid parameter or not:

```
CGbool cgIsParameter(CGparameter parameter);
```

A parameter handle becomes invalid when the program or the context of the program it corresponds to is destroyed.

Parameter References

A parameter that is referenced by the original Cg source code may be optimized out of the compiled program by the compiler, in which case the application can simply ignore it and not set its value. Calling **cgIsParameterReferenced()** allows you to check whether a parameter is potentially used by the final compiled program:

```
CGbool cgIsParameterReferenced(CGparameter parameter);
```

Note that the value returned by this entry point is conservative, but not always exact, particularly if the program has not yet been compiled. Also, note that no error is generated if you set the value of a parameter that is not referenced.

Parameter Size

A number of core Cg runtime entry points are provided for querying and setting parameter size and length.

The number of rows or columns associated with a parameter can be retrieved using

```
int cgGetParameterRows(CGparameter param);  
int cgGetParameterColumns(CGparameter param);
```

A scalar parameter is considered to have a single row and a single column, while a vector parameter has a single row and columns equal to the length of the vector. If *param* is a matrix parameter, the values returned correspond to those of the matrix. If *param* is an array, the number of rows or columns associated with each element of the array is returned. If *param* is not a numeric type, 0 is returned by either entry point.

The dimensionality of an array is queried using

```
int cgGetArrayDimension(CGparameter param);
```

Dimensions are enumerated starting at 0 (zero). The length of a particular dimension of an array can be retrieved by calling

```
int cgGetArraySize(CGparameter param, int dimension);
```

The total number of elements in an array may be queried using

```
int cgGetArrayTotalSize(CGparameter param);
```

Here, *param* may be an array of any dimension; the returned value is the total number of elements across all dimensions of the array.

The type of each element of an array can be queried using

```
Cgtype cgGetArrayType(CGparameter param);
```

For example, if a parameter were declared

```
float4 array[2][3];
```

cgGetArrayType() would return **CG_FLOAT4**. If it were declared

```
mystruct array[3];
```

cgGetArrayType() would return the enumerant corresponding to the user-defined **mystruct** type.

Unsigned Array Length

Unsigned arrays can be assigned concrete sizes via the runtime. Under many profiles, setting the size of unsigned arrays associated with a Cg program is required before the program can be compiled.

The length of one-dimensional unsized arrays can be set using

```
void cgSetArraySize(CGparameter param, int size);
```

The size of multidimensional arrays may be set using

```
void cgSetMultiDimArraySize(CGparameter param, int *sizes);
```

Note that arrays with completely determined lengths may not have their size changed using either entry point. Only unsized arrays may be modified using these entry points.

Parameter Attributes

A parameter's general class can be queried using

```
CGparameterclass cgGetParameterClass(CGparameter param);
```

The returned **CGparameterclass** value enumerates the high-level parameter classes:

- ❑ **CG_PARAMETERCLASS_SCALAR**
A scalar type, such as **CG_INT** or **CG_FLOAT**
- ❑ **CG_PARAMETERCLASS_VECTOR**
A vector type, such as **CG_INT1** or **CG_FLOAT4**
- ❑ **CG_PARAMETERCLASS_MATRIX**
A matrix type, such as **CG_INT1X2** or **CG_FLOAT4X4**
- ❑ **CG_PARAMETERCLASS_STRUCT**
A **struct** or interface
- ❑ **CG_PARAMETERCLASS_SAMPLER**
A sampler type, such as **sampler1D** or **samplerCUBE**
- ❑ **CG_PARAMETERCLASS_OBJECT**
A texture, string, or program

The program that the parameter corresponds to is found using

```
cgGetParameterProgram():
```

```
CGprogram cgGetParameterProgram(CGparameter parameter);
```

To determine whether the parameter is varying, uniform, or constant, **cgGetParameterVariability()** is used:

```
CGenum cgGetParameterVariability(CGparameter parameter);
```

The call returns **CG_VARYING** if the parameter is a varying parameter, **CG_UNIFORM** if the parameter is a uniform parameter, or **CG_CONSTANT** if the parameter is a constant parameter. A *constant parameter* is a parameter whose value never changes for the life of a compiled program, so that changing its

value requires recompiling the program. For some profiles, the compiler has to add some that correspond to literal constant values in the code.

A parameter's variability can also be modified via the core Cg runtime using

```
void cgSetParameterVariability(CGparameter parameter,  
                               CGenum vary);
```

Here, ***vary*** may be one of:

❑ **CG_UNIFORM**

The parameter is set to uniform variability.

❑ **CG_LITERAL**

The parameter is marked as a literal, whose value can be assumed to be a compile-time constant compilation. This feature can be used to “bake” parameter values into the compiled Cg program, which often produces much more efficient compiled code.

❑ **CG_DEFAULT**

The parameter reverts to its default variability as specified in the program text, or is made to inherit its variability from any source it has been connected to.

Note that parameters may not currently be set to **CG_VARYING** variability.

To obtain the parameter direction, use **cgGetParameterDirection()**:

```
CGenum cgGetParameterDirection(CGparameter parameter);
```

It returns **CG_IN** if the parameter is an input parameter, **CG_OUT** if the parameter is an output parameter, or **CG_INOUT** if the parameter is both an input and an output parameter.

The entry point **cgGetParameterType()** retrieves the parameter name:

```
const char* cgGetParameterName(CGparameter parameter);
```

Use **cgGetParameterSemantic()** to retrieve the parameter semantic string:

```
const char* cgGetParameterSemantic(CGparameter parameter);
```

If the parameter does not have any semantic, an empty string is returned.

There is a one-to-one correspondence between a set of predefined semantics (**POSITION**, **COLOR**, and so on) and hardware resources (registers, texture units, and so on). In the Cg runtime, a hardware resource is represented by the type **CGresource** and **cgGetParameterResource()** retrieves the resource assigned to a parameter:

```
CGresource cgGetParameterResource(CGparameter parameter);
```

If the parameter does not have any associated resource, **cgGetParameterResource()** returns **CG_UNDEFINED**.

The two functions **cgGetResource()** and **cgGetResourceString()** allow you to determine the correspondence between a resource enumerant and its corresponding string:

```
Cgresource cgGetResource(const char* resourceString);  
const char* cgGetResourceString(Cgresource resource);
```

If the string passed to **cgGetResource()** does not correspond to any resource, **CG_UNDEFINED** is returned.

Using **cgGetParameterBaseResource()** allows you to retrieve the base resource for a parameter in a Cg program:

```
Cgresource cgGetParameterBaseResource(  
    Cgparameter parameter);
```

The base resource is the first resource in a set of sequential resources. For example, if a given parameter has a resource equal to **CG_TEXCOORD7**, its base resource is **CG_TEXCOORD0**. Only parameters with resources whose name ends with a number have a base resource. All other parameters return **CG_UNDEFINED** when **cgGetParameterBaseResource()** is called.

Function **cgGetParameterResourceIndex()** retrieves the numerical portion of the resource:

```
unsigned long cgGetParameterResourceIndex(  
    Cgparameter parameter);
```

For example, if the resource for a given parameter is **CG_TEXCOORD7**, **cgGetParameterResourceIndex()** returns 7.

The **cgGetParameterValues()** function retrieves the default or constant value of a uniform parameter:

```
const double* cgGetParameterValues(Cgparameter parameter,  
    Cgenum valueType, int* numberOfValuesReturned);
```

It retrieves the default value if *valueType* is equal to **CG_DEFAULT** and the constant value if *valueType* is equal to **CG_CONSTANT**. The components of the value are returned in row-major order as a pointer to an array containing type **double** elements. After **cgGetParameterValues()** is called, the number of components available in the array is pointed to by *numberOfValuesReturned*.

Core Cg Error Reporting

An error code is associated with each type of runtime error that can be generated. The runtime caches both the most recently generated error, as well as the error that was first generated since the error code was last checked by the application. Applications can query the cached error codes, as well as the error message corresponding to either, using

```
CGerror error = cgGetError();
CGerror error = cgGetFirstError();
const char* errorString = cgGetErrorString(error);
```

An error code of 0 indicates no error. When either error-fetching entry point is called, its cached error value is reset to 0.

More comprehensive error checking and handling can be achieved using Cg's error handler callback mechanism. Each time an error occurs, the core Cg runtime calls an error handler callback function, optionally provided by the application. The application registers the error handler using

```
typedef void (*CGErrorHandlerFunc)(CGcontext ctx, CGerror err,
                                   void *appdata);
void cgSetErrorHandler(CGErrorHandlerFunc func, void *data);
```

When an error occurs, the Cg runtime calls the specified function, passing the **CGcontext** in which the error occurred, the code associated with the triggering error, and a copy of the data pointer registered by the application. A typical implementation of the error handler might look like this:

```
void HandleCgError(CGcontext ctx, CGerror err, void *appdata)
{
    fprintf(stderr, "Cg error: %s\n", cgGetErrorString(err));
    const char *listing = cgGetLastListing(ctx);
    if (listing != NULL)
        fprintf(stderr, "    last listing: %s\n", listing);
}
```

Here is a list of some of the **CGerror** codes specific to the core Cg runtime:

- ❑ **CG_NO_ERROR**: Returned when no error has occurred.
- ❑ **CG_COMPILER_ERROR**: Returned when the compiler generated an error. A call to **cgGetLastListing()** should be made to get more details on the actual compiler error.
- ❑ **CG_INVALID_PARAMETER_ERROR**: Returned when the parameter used is invalid.
- ❑ **CG_INVALID_PROFILE_ERROR**: Returned when the profile is not supported.

- ❑ **CG_INVALID_VALUE_TYPE_ERROR:** Returned when an unknown value type is assigned to a parameter.
- ❑ **CG_NOT_MATRIX_PARAM_ERROR:** Returned when the parameter is not of a matrix type.
- ❑ **CG_INVALID_ENUMERANT_ERROR:** Returned when the enumerant parameter has an invalid value.
- ❑ **CG_NOT_4x4_MATRIX_ERROR:** Returned when the parameter must be a 4x4 matrix type.
- ❑ **CG_FILE_READ_ERROR:** Returned when the file cannot be read.
- ❑ **CG_FILE_WRITE_ERROR:** Returned when the file cannot be written.
- ❑ **CG_MEMORY_ALLOC_ERROR:** Returned when a memory allocation fails.
- ❑ **CG_INVALID_CONTEXT_HANDLE_ERROR:** Returned when an invalid context handle is used.
- ❑ **CG_INVALID_PROGRAM_HANDLE_ERROR:** Returned when an invalid program handle is used.
- ❑ **CG_INVALID_PARAM_HANDLE_ERROR:** Returned when an invalid parameter handle is used.
- ❑ **CG_UNKNOWN_PROFILE_ERROR:** Returned when the specified profile is unknown.
- ❑ **CG_VAR_ARG_ERROR:** Returned when the variable arguments are specified incorrectly.
- ❑ **CG_INVALID_DIMENSION_ERROR:** Returned when the dimension value is invalid.
- ❑ **CG_ARRAY_PARAM_ERROR:** Returned when the parameter must be an array.
- ❑ **CG_OUT_OF_ARRAY_BOUNDS_ERROR:** Returned when the index into an array is out of bounds.

API-Specific Cg Runtimes

Each API-specific Cg runtime provides an additional set of functions on top of the core Cg runtime to ease the integration of Cg to an application based on this API. They essentially interface between the core runtime data structures and the API data structures to provide the following facilities:

- ❑ Setting the parameter values: A distinction is made between texture, matrix, array, vector and scalar values as those various types are handled differently by each API and have different data structures.
- ❑ Executing the program: Program execution is divided into program loading (passing the result of the Cg compiler to the API) and program binding (setting the program as the one to execute for any subsequent draw calls). This is because those two operations are usually done at a different time: A program is loaded each time it is recompiled and it is bound each time it needs to be executed for a particular draw call.

Parameter Shadowing

When the value of a uniform parameter is set by some function of the OpenGL Cg runtime, it is actually stored internally (or *shadowed*) by either the Cg or the OpenGL runtime so that it does not need to be reset every time the program is about to be executed. This behavior is referred to as *parameter shadowing*.

If the Direct3D Cg runtime expanded interface (described in “[Direct3D Expanded Interface](#)” on page 98) is used, parameter shadowing can be turned on or off on a per-program basis. When parameter shadowing is turned off for a given program and the value of any of its uniform parameters is set by some function of the Direct3D Cg runtime, it is immediately downloaded to the GPU constant memory (the memory containing the values of all the uniform parameters). When parameter shadowing is turned on, the value is shadowed instead and no Direct3D call is made at the time it is set; only when the program is bound are all of its parameters actually downloaded to the constant memory. This means that a parameter value set after binding the program is not used during the execution of the program until the next time the program is bound. Parameter shadowing applies to all parameter settings including texture state stage and texture mode.

Disabling parameter shadowing allows the runtime to consume less memory, but forces the application to do the work of making sure that the constant memory contains all the right values every time it activates a program.

OpenGL Cg Runtime

This section discusses setting parameters and program execution for the OpenGL Cg runtime.

Note: Before any OpenGL Cg runtime functions can be executed, an OpenGL context must be created with either **wglCreateContext()** or **glXCreateContext()**.

Setting Parameters in OpenGL

In accordance with the OpenGL convention, many of the functions described below come in two versions: a version operating on **float** values, marked with an **f**, and a version operating on **double** values, marked with a **d**.

Setting Uniform Scalar and Uniform Vector Parameters

To set the values of scalar parameters or vector parameters, use the **cgGLSetParameter** functions:

```
void cgGLSetParameter1f(CGparameter parameter, float x);
void cgGLSetParameter1fv(CGparameter parameter,
                          const float* array);
void cgGLSetParameter1d(CGparameter parameter, double x);
void cgGLSetParameter1dv(CGparameter parameter,
                          const double* array);

void cgGLSetParameter2f(CGparameter parameter, float x,
                        float y);
void cgGLSetParameter2fv(CGparameter parameter,
                          const float* array);
void cgGLSetParameter2d(CGparameter parameter, double x,
                        double y);
void cgGLSetParameter2dv(CGparameter parameter,
                          const double* array);

void cgGLSetParameter3f(CGparameter parameter, float x,
                        float y, float z);
void cgGLSetParameter3fv(CGparameter parameter,
                          const float* array);
void cgGLSetParameter3d(CGparameter parameter, double x,
                        double y, double z);
void cgGLSetParameter3dv(CGparameter parameter,
                          const double* array);

void cgGLSetParameter4f(CGparameter parameter, float x,
                        float y, float z, float w);
void cgGLSetParameter4fv(CGparameter parameter,
                          const float* array);
```

```

void cgGLSetParameter4d(CGparameter parameter, double x,
                        double y, double z, double w);
void cgGLSetParameter4dv(CGparameter parameter,
                        const double* array);

```

The digit in the name of those functions indicates how many scalar values are set by the function. The **v** suffix is for functions that operate on an array of values as opposed to individual arguments.

If more values are set than the parameter requires, the extra values are ignored. If less values are set than the parameter requires, the last value is smeared. The **cgGLSetParameter** functions may be called for either uniform or varying parameters. When called for a varying parameter, the appropriate immediate mode OpenGL entry point is called.

The corresponding parameter value retrieval functions are as follows:

```

cgGLGetParameter1f(CGparameter parameter, float* array);
cgGLGetParameter1d(CGparameter parameter, double* array);
cgGLGetParameter2f(CGparameter parameter, float* array);
cgGLGetParameter2d(CGparameter parameter, double* array);
cgGLGetParameter3f(CGparameter parameter, float* array);
cgGLGetParameter3d(CGparameter parameter, double* array);
cgGLGetParameter4f(CGparameter parameter, double* array);
cgGLGetParameter4d(CGparameter parameter, type* array);

```

Setting Uniform Matrix Parameters

The **cgGLSetMatrixParameter** functions are used to set any matrix:

```

void cgGLSetMatrixParameterfr(CGparameter parameter,
                              const float* matrix);
void cgGLSetMatrixParameterfc(CGparameter parameter,
                              const float* matrix);
void cgGLSetMatrixParameterdr(CGparameter parameter,
                              const double* matrix);
void cgGLSetMatrixParameterdc(CGparameter parameter,
                              const double* matrix);

```

The matrix is passed as an array of floating point values whose size matches the number of coefficients of the matrix. The **r** suffix is for functions that assume the matrix is laid out in row order, and the **c** suffix is for functions that assume the matrix is laid out in column order.

The corresponding parameter value retrieval functions are

```

void cgGLGetMatrixParameterfr(CGparameter parameter,
                              float* matrix);
void cgGLGetMatrixParameterfc(CGparameter parameter,
                              float* matrix);

```

```

void cgGLGetMatrixParameterdr(CGparameter parameter,
                             double* matrix);
void cgGLGetMatrixParameterdc(CGparameter parameter,
                             double* matrix);

```

Use **cgGLSetStateMatrixParameter()** to set a OpenGL 4x4 state matrix:

```

void cgGLSetStateMatrixParameter(CGparameter parameter,
                                GLenum stateMatrixType, GLenum transform);

```

The variable ***stateMatrixType*** is an enumerate type specifying the state matrix to be used to set the parameter:

- ☐ **CG_GL_MODELVIEW_MATRIX** for the current model-view matrix
- ☐ **CG_GL_PROJECTION_MATRIX** for the current projection matrix
- ☐ **CG_GL_TEXTURE_MATRIX** for the current texture matrix
- ☐ **CG_GL_MODELVIEW_PROJECTION_MATRIX** for the concatenated model-view and projection matrices

The variable ***transform*** is an enumerate type specifying a transformation applied to the state matrix before it is used to set the parameter value:

- ☐ **CG_GL_MATRIX_IDENTITY** for applying no transformation at all
- ☐ **CG_GL_MATRIX_TRANSPOSE** for transposing the matrix
- ☐ **CG_GL_MATRIX_INVERSE** for inverting the matrix
- ☐ **CG_GL_MATRIX_INVERSE_TRANSPOSE** for inverting and transposing the matrix

Setting Uniform Arrays of Scalar, Vector, and Matrix Parameters

To set the values of arrays of uniform scalar or vector parameters, use the **cgGLSetParameterArray** functions:

```

void cgGLSetParameterArray1f(CGparameter parameter,
                             long startIndex, long numberOfElements,
                             const float* array);
void cgGLSetParameterArray1d(CGparameter parameter,
                             long startIndex, long numberOfElements,
                             const double* array);
void cgGLSetParameterArray2f(CGparameter parameter,
                             long startIndex, long numberOfElements,
                             const float* array);
void cgGLSetParameterArray2d(CGparameter parameter,
                             long startIndex, long numberOfElements,
                             const double* array);

```

```

void cgGLSetParameterArray3f(CGparameter parameter,
    long startIndex, long numberOfElements,
    const float* array);
void cgGLSetParameterArray3d(CGparameter parameter,
    long startIndex, long numberOfElements,
    const double* array);
void cgGLSetParameterArray4f(CGparameter parameter,
    long startIndex, long numberOfElements,
    const float* array);
void cgGLSetParameterArray4d(CGparameter parameter,
    long startIndex, long numberOfElements,
    const double* array);

```

The digit in the name of those functions indicates the type of the parameter array elements: **1** for arrays of **float1**, **2** for arrays of **float2**, and so on. The variables **startIndex** and **numberOfElements** specify which elements of the **array** parameter are set: They are the **numberOfElements** elements of the indices that range from **startIndex** to **startIndex+numberOfElements-1**. Passing a value of 0 for **numberOfElements** tells the functions to set all the values starting at index **startIndex** up to the last valid index of the array, namely **cgGLGetArraySize(parameter, 0)-1**. This is equivalent to setting **numberOfElements** to **cgGLGetArraySize(parameter, 0) - startIndex**. The parameter **array** is an array of scalar values. It must have **numberOfElements** for the **cgGLSetParameterArray1** functions, **2*numberOfElements** for the **cgGLSetParameterArray2** functions, and so on.

The corresponding parameter value retrieval functions are as follows:

```

void cgGLGetParameterArray1f(CGparameter parameter,
    long startIndex, long numberOfElements, float* array);
void cgGLGetParameterArray1d(CGparameter parameter,
    long startIndex, long numberOfElements, double* array);
void cgGLGetParameterArray2f(CGparameter parameter,
    long startIndex, long numberOfElements, float* array);
void cgGLGetParameterArray2d(CGparameter parameter,
    long startIndex, long numberOfElements, double* array);
void cgGLGetParameterArray3f(CGparameter parameter,
    long startIndex, long numberOfElements, float* array);
void cgGLGetParameterArray3d(CGparameter parameter,
    long startIndex, long numberOfElements, double* array);
void cgGLGetParameterArray4f(CGparameter parameter,
    long startIndex, long numberOfElements, float* array);
void cgGLGetParameterArray4d(CGparameter parameter,
    long startIndex, long numberOfElements, double* array);

```

Similar functions exist to set the values of arrays of uniform matrix parameters:

```
void cgGLSetMatrixParameterArrayfr(CGparameter parameter,
    long startIndex, long numberOfElements,
    const float* array);
void cgGLSetMatrixParameterArrayfc(CGparameter parameter,
    long startIndex, long numberOfElements,
    const float* array);
void cgGLSetMatrixParameterArraydc(CGparameter parameter,
    long startIndex, long numberOfElements,
    const double* array);
void cgGLSetMatrixParameterArraydc(CGparameter parameter,
    long startIndex, long numberOfElements,
    const double* array);
```

and to query those values:

```
void cgGLGetMatrixParameterArrayfr(CGparameter parameter,
    long startIndex, long numberOfElements, float* array);
void cgGLGetMatrixParameterArrayfc(CGparameter parameter,
    long startIndex, long numberOfElements, float* array);
void cgGLGetMatrixParameterArraydc(CGparameter parameter,
    long startIndex, long numberOfElements, double* array);
void cgGLGetMatrixParameterArraydc(CGparameter parameter,
    long startIndex, long numberOfElements, double* array);
```

The **c** and **r** suffixes have the same meaning as they do for the **cgGLSetMatrixParameter** functions.

Setting Varying Parameters

The values of fragment program varying parameters are set as the result of the interpolation across the triangles performed by the GPU, so only the values of vertex program varying parameters are set by the application.

Setting a vertex varying parameter requires two steps.

The first step consists in passing a pointer to an array containing the values for each vertex. This is done using **cgGLSetParameterPointer()**:

```
void cgGLSetParameterPointer(CGparameter parameter,
    GLint size, GLenum type, GLsizei stride,
    GLvoid* array);
```

The variable **size** indicates the number of values per vertex that are stored in **array**. It is equal to 1, 2, 3, or 4. If fewer values are set than the parameter requires, the non-specified values default to 0 for **x**, **y**, and **z**, and 1 for **w**.

The enumerate type **type** specifies the data type of the values stored in **array**: **GL_SHORT**, **GL_INT**, **GL_FLOAT**, or **GL_DOUBLE**.

The parameter **stride** is the byte offset between any two consecutive vertices. Passing a value of zero for **stride** is equivalent to passing a byte offset equal to **size** multiplied by the size of **type** in bytes; in other words, it means that there is no gap between two consecutive vertex values. Note that the minimum size for **array** is implicitly defined by the biggest vertex index specified in the triangles drawn.

The second step consists in enabling the varying parameter for a specific drawing call:

```
void cgGLEnableClientState(CGparameter parameter);
```

The equivalent disabling function is

```
void cgGLDisableClientState(CGparameter parameter);
```

Another way to set the vertex varying parameter is to use the **cgGLSetParameter** functions. When a **cgGLSetParameter** function is called for a varying parameter, the appropriate immediate-mode OpenGL entry point is called. The **cgGLGetParameter** functions do not apply to varying parameters.

Setting Sampler Parameters

Setting a sampler parameter requires two steps. First, an OpenGL texture object handle must be assigned to the sampler parameter. Next, the texture unit associated with the sampler must be enabled prior to drawing. The first step must be done explicitly by the application. The second step may also be performed explicitly by the application, or the OpenGL Cg runtime can be instructed to automatically manage texture units itself.

The first step consists in assigning an OpenGL texture object to the sampler parameter using

```
void cgGLSetTextureParameter(CGparameter parameter,  
                             GLuint textureName);
```

where **textureName** is the OpenGL texture name. Note that when your application makes OpenGL calls to initialize the texture environment for a given **sampler**, it is important to remember to set the active texture unit to that associated with the **sampler** before doing so. The **sampler**'s texture unit can be retrieved by calling **cgGLGetTextureEnum()**; see the following discussion.

The second step consists of enabling the texture unit associated with the **sampler** parameter for a specific drawing call. It is strongly recommended

that applications allow the Cg OpenGL runtime library to perform this second step itself. This is accomplished by calling:

```
void cgGLSetManageTextureParameters(CGcontext context,  
                                     CGbool enable);
```

with **enable** set to a non-zero value after the Cg context has been created. When automatic texture parameter management is in effect, the Cg OpenGL runtime will automatically enable all appropriate texture units when a CGprogram is bound.

If, despite the above, you wish to manage texture parameters yourself, you can use the helper function

```
void cgGLEnableTextureParameter(CGparameter parameter);
```

which must be called after **cgGLSetTextureParameter()** and before the actual drawing call.

The equivalent disabling function is:

```
void cgGLDisableTextureParameter(CGparameter parameter);
```

You can retrieve the texture object assigned to a sampler parameter using

```
GLuint cgGLGetTextureParameter(CGparameter parameter);
```

You can retrieve the OpenGL enumerator for the texture unit associated with a sampler parameter using

```
GLenum cgGLGetTextureEnum(CGparameter parameter);
```

The returned enumerator has the form **GL_TEXTURE#_ARB** where # is the texture unit index.

OpenGL Profile Support

A convenient function is provided that gives the best available profile for vertex or fragment programs depending on the available OpenGL extensions.

```
CGprofile cgGLGetLatestProfile(CGGLenum profileType);
```

Parameter **profileType** is equal to **CG_GL_VERTEX** or **CG_GL_FRAGMENT**. Function **cgGLGetLatestProfile()** may be used in conjunction with **cgCreateProgram()** or **cgCreateProgramFromFile()** to ensure that the best available vertex and fragment profiles are used for compilation. This allows you to make your application future-ready, because the Cg programs are automatically compiled for the best profiles that are available at runtime, even if these profiles did not exist at the time the application was written. Another function that allows you optimal compilation is **cgGLSetOptimalOptions()**. It sets implicit compiler arguments that are

appended to the argument list passed to **cgCreateProgram()** or **cgCreateProgramFromFile()**.

```
void cgGLSetOptimalOptions(CGprofile profile);
```

OpenGL Program Execution

All programs must be loaded before they can be bound. To load a program use **cgGLLoadProgram()**:

```
void cgGLLoadProgram(CGprogram program);
```

Binding a program only works if its profile is enabled. This is done by calling **cgGLEnableProfile()** with the program profile:

```
void cgGLEnableProfile(CGprofile profile);
```

The binding itself is done using **cgGLBindProgram()**:

```
void cgGLBindProgram(CGprogram program);
```

Only one vertex program and one fragment program can be bound at any given time, so binding a program implicitly unbinds any other program of that type.

Profiles are disabled using **cgGLDisableProfile()**:

```
void cgGLDisableProfile(CGprofile profile);
```

Some profiles may not be supported on some systems. For example, a given profile is not supported if the OpenGL extensions it requires are not available. You can check if a profile is supported by using

```
cgGLIsProfileSupported();
```

```
CGbool cgGLIsProfileSupported(CGprofile profile);
```

It returns **CG_TRUE** if *profile* is supported and **CG_FALSE** otherwise.

OpenGL Program Examples

This section presents code that illustrates how to use functions from the OpenGL Cg interface to make Cg programs work with OpenGL. The vertex and fragment programs below are used in [“OpenGL Application” on page 82](#).

OpenGL Vertex Program

The following Cg code is assumed to be in a file called **VertexProgram.cg**.

```
void VertexProgram(
    in float4 position    : POSITION,
    in float4 color       : COLOR0,
    in float4 texCoord    : TEXCOORD0,
```

```

    out float4 position0 : POSITION,
    out float4 color0     : COLOR0,
    out float4 texCoord0  : TEXCOORD0,
    const uniform float4x4 ModelViewMatrix )
{
    position0 = mul(position, ModelViewMatrix);
    color0 = color;
    texCoord0 = texCoord;
}

```

OpenGL Fragment Program

The following Cg code is assumed to be in a file called **FragmentProgram.cg**.

```

void FragmentProgram(
    in float4 color      : COLOR0,
    in float4 texCoord   : TEXCOORD0,
    out float4 color0    : COLOR0,
    const uniform sampler2D BaseTexture,
    const uniform float4 SomeColor)
{
    color0 = color * tex2D(BaseTexture, texCoord) + SomeColor;
}

```

OpenGL Application

This C code links the previous vertex and fragment programs to the application.

```

#include <cg/cg.h>
#include <cg/cgGL.h>

float* vertexPositions; // Initialized somewhere else
float* vertexColors;    // Initialized somewhere else
float* vertexTexCoords; // Initialized somewhere else
GLuint texture;         // Initialized somewhere else
float constantColor[];  // Initialized somewhere else
CGcontext context;
CGprogram vertexProgram, fragmentProgram;
CGprofile vertexProfile, fragmentProfile;
CGparameter position, color, texCoord, baseTexture, someColor,
               modelViewMatrix;

// Called at initialization
void CgGLInit()
{
    // Create context
    context = cgCreateContext();
}

```

```

// Initialize profiles and compiler options
vertexProfile = cgGLGetLatestProfile(CG_GL_VERTEX);
cgGLSetOptimalOptions(vertexProfile);
fragmentProfile = cgGLGetLatestProfile(CG_GL_FRAGMENT);
cgGLSetOptimalOptions(fragmentProfile);

// Create the vertex program
vertexProgram = cgCreateProgramFromFile(
    context, CG_SOURCE, "VertexProgram.cg",
    vertexProfile, "VertexProgram", 0);

// Load the program
cgGLLoadProgram(vertexProgram);

// Create the fragment program
fragmentProgram = cgCreateProgramFromFile(
    context, CG_SOURCE, "FragmentProgram.cg",
    fragmentProfile, "FragmentProgram", 0);

// Load the program
cgGLLoadProgram(fragmentProgram);

// Grab some parameters.
position = cgGetNamedParameter(vertexProgram, "position");
color = cgGetNamedParameter(vertexProgram, "color");
texCoord = cgGetNamedParameter(vertexProgram, "texCoord");
modelViewMatrix = cgGetNamedParameter(vertexProgram,
    "ModelViewMatrix");
baseTexture = cgGetNamedParameter(fragmentProgram,
    "BaseTexture");
someColor = cgGetNamedParameter(fragmentProgram,
    "SomeColor");

// Set parameters that don't change:
// They can be set only once because of parameter shadowing.
cgGLSetTextureParameter(baseTexture, texture);
cgGLSetParameter4fv(someColor, constantColor);
}

// Called to render the scene
void Display()
{
    // Set the varying parameters
    cgGLEnableClientState(position);

```

```

    cgGLSetParameterPointer(position, 3, GL_FLOAT, 0,
                           vertexPositions);
    cgGLEnableClientState(color);
    cgGLSetParameterPointer(color, 1, GL_FLOAT, 0,
                           vertexColors);
    cgGLEnableClientState(texCoord);
    cgGLSetParameterPointer(texCoord, 2, GL_FLOAT, 0,
                           vertexTexCoords);

    // Set the uniform parameters that change every frame
    cgGLSetStateMatrixParameter(modelViewMatrix,
                               CG_GL_MODELVIEW_PROJECTION_MATRIX,
                               CG_GL_MATRIX_IDENTITY);

    // Enable the profiles
    cgGLEnableProfile(vertexProfile);
    cgGLEnableProfile(fragmentProfile);

    // Bind the programs
    cgGLBindProgram(vertexProgram);
    cgGLBindProgram(fragmentProgram);

    // Enable texture
    cgGLEnableTextureParameter(baseTexture);

    // Draw scene
    // ...

    // Disable texture
    cgGLDisableTextureParameter(baseTexture);

    // Disable the profiles
    cgGLDisableProfile(vertexProfile);
    cgGLDisableProfile(fragmentProfile);

    // Set the varying parameters
    cgGLDisableClientState(position);
    cgGLDisableClientState(color);
    cgGLDisableClientState(texCoord);
}

// Called before application shuts down
void CgShutdown()
{
    // This frees any runtime resource.

```

```
cgDestroyContext(context);
}
```

OpenGL Error Reporting

Here is the list of the **Cgerror** errors specific to the OpenGL Cg runtime:

- ❑ **CG_PROGRAM_LOAD_ERROR**: Returned when the program could not be loaded.
- ❑ **CG_PROGRAM_BIND_ERROR**: Returned when the program could not be bound.
- ❑ **CG_PROGRAM_NOT_LOADED_ERROR**: Returned when the program must be loaded before the operation may be used.
- ❑ **CG_UNSUPPORTED_GL_EXTENSION_ERROR**: Returned when an unsupported Open GL extension is required to perform the operation.

Any OpenGL Cg runtime function can generate an OpenGL error in addition to the Cg-specific error. These errors are checked in Cg, as in any OpenGL application, by using **glGetError()**.

Direct3D Cg Runtime

The Direct3D Cg runtime is composed of two interfaces:

- ❑ *Minimal interface*: This interface makes no Direct3D calls itself and should be used when you prefer to keep the Direct3D code in the application itself.
- ❑ *Expanded interface*: This interface makes the Direct3D calls necessary to provide enhanced program and parameter management and should be used when you prefer to let the Cg runtime manage the Direct3D shaders.

Direct3D Minimal Interface

The minimal interface simply supplies convenient functions to convert some information provided by the core runtime to information specific to Direct3D.

Vertex Declaration

In Direct3D, you have to supply a vertex declaration that establishes a mapping between the vertex shader input registers and the data provided by the application as data streams. In Direct3D 9, this vertex declaration is bound to the current state the same way the vertex shader is (see the

Direct3D 9 documentation on

IDirect3DDevice9::CreateVertexDeclaration() and **IDirect3DDevice9::SetVertexDeclaration()** for a detailed explanation). In Direct3D 8, the vertex declaration is required at the time you create the vertex shader (for more information, see the Direct3D 8 documentation on **IDirect3DDevice8::CreateVertexShader()**).

A data stream is basically an array of data structures. Each of those structures is of a particular type called the *vertex format* of the stream. Here is an example of a vertex declaration for Direct3D 9:

```
const D3DVERTEXELEMENT9 declaration[] = {
    { 0, 0 * sizeof(float),
      D3DDECLTYPE_FLOAT3, D3DDECLMETHOD_DEFAULT,
      D3DDECLUSAGE_POSITION, 0 }, // Position
    { 0, 3 * sizeof(float),
      D3DDECLTYPE_FLOAT3, D3DDECLMETHOD_DEFAULT,
      D3DDECLUSAGE_NORMAL, 0 }, // Normal
    { 0, 8 * sizeof(float),
      D3DDECLTYPE_FLOAT2, D3DDECLMETHOD_DEFAULT,
      D3DDECLUSAGE_TEXCOORD, 0 }, // Base texture
    { 1, 0 * sizeof(float),
      D3DDECLTYPE_FLOAT3, D3DDECLMETHOD_DEFAULT,
      D3DDECLUSAGE_TEXCOORD, 1 }, // Tangent
    D3DD3CL_END()
};
```

Here is an example of a vertex declaration for Direct3D 8:

```
const DWORD declaration[] = {
    D3DVSD_STREAM(0),
    D3DVSD_REG(D3DVSDE_POSITION, D3DVSDT_FLOAT3), // Position
    D3DVSD_REG(D3DVSDE_NORMAL, D3DVSDT_FLOAT3),   // Normal
    D3DVSD_SKIP(2),                               // Skip the diffuse and specular color
    D3DVSD_REG(D3DVSDE_TEXCOORD0,
               D3DVSDT_FLOAT2), // Base texture
    D3DVSD_STREAM(1), // Tangent basis stream
    D3DVSD_REG(D3DVSDE_TEXCOORD1, D3DVSDT_FLOAT3), // Tangent
    D3DVSD_END()
};
```

Both declarations tell the Direct3D runtime to find (1) the positions of the vertices in stream **0** as the first three floating point values of the vertex format, (2) the normals as the next three floating point values following the three floating point values in stream **0**, and (3) the texture coordinates as the two floating point values located at an offset equal to twice the size of a **DWORD** from the end of the normal data in stream **0**. The tangents are

provided in stream **1** as a second texture coordinate set that is found as the first three floating point values of the vertex format.

To get a vertex declaration from a Cg vertex program for the Direct3D 9 Cg runtime use **cgD3D9GetVertexDeclaration()**:

```
Cgbool cgD3D9GetVertexDeclaration(CGprogram program,
                                   D3DVERTEXELEMENT9 declaration[MAXD3DDECLLENGTH]);
```

MAXD3DDECLLENGTH is a Direct3D 9 constant that gives the maximum length of a Direct3D 9 declaration. If no declaration can be derived from the program, **cgD3D9GetVertexDeclaration()** fails and returns **CG_FALSE**.

To get a vertex declaration from a Cg vertex program for the Direct3D 8 Cg runtime use **cgD3D8GetVertexDeclaration()**:

```
Cgbool cgD3D8GetVertexDeclaration(CGprogram program,
                                   DWORD declaration[MAX_FVF_DECL_SIZE]);
```

MAX_FVF_DECL_SIZE is a Direct3D constant that gives the maximum length of a Direct3D declaration. If no declaration can be derived from the program, **cgD3D8GetVertexDeclaration()** fails and returns **CG_FALSE**.

The declaration returned by **cgD3D9GetVertexDeclaration()** or **cgD3D8GetVertexDeclaration()** is for a single stream, so that for the following program:

```
void main(in float4 position : POSITION,
          in float4 color    : COLOR0,
          in float4 texCoord : TEXCOORD0,
          out float4 hpos    : POSITION)
{ }
```

it is equivalent to:

```
const D3DVERTEXELEMENT9 declaration[] = {
    { 0, 0 * sizeof(float),
      D3DDECLTYPE_FLOAT4, D3DDECLMETHOD_DEFAULT,
      D3DDECLUSAGE_POSITION, 0 },
    { 0, 4 * sizeof(float),
      D3DDECLTYPE_FLOAT4, D3DDECLMETHOD_DEFAULT,
      D3DDECLUSAGE_COLOR, 0 },
    { 0, 8 * sizeof(float),
      D3DDECLTYPE_FLOAT4, D3DDECLMETHOD_DEFAULT,
      D3DDECLUSAGE_TEXCOORD, 0 },
    D3DD3CL_END()
};
```

for the Direct3D 9 Cg runtime, and it is equivalent to:

```
const DWORD declaration[] = {
```

for the Direct3D 8 Cg runtime. This is vertex shader

Usually though, you want to apply a vertex program to geometric data that come in multiple streams or with specific vertex formats. In this case, the vertex declaration is based on the vertex

```

        D3DDECLTYPE_FLOAT2, D3DDECLMETHOD_DEFAULT,
        D3DDECLUSAGE_TEXCOORD, 0 },
    D3DD3CL_END()
};

```

and the following Direct3D 8 vertex declaration is valid:

```

DWORD declaration[] = {
    D3DVSD_STREAM(0),
    D3DVSD_REG(D3DVSDE_POSITION, D3DVSDT_FLOAT3),
    D3DVSD_REG(D3DVSDE_DIFFUSE, D3DVSDT_D3DCOLOR),
    D3DVSD_STREAM(1),
    D3DVSD_SKIP(4),
    D3DVSD_REG(D3DVSDE_TEXCOORD0, D3DVSDT_FLOAT2),
    D3DVSD_END()
};

```

This is true because **D3DDECLUSAGE_POSITION** and **D3DVSDE_POSITION** match the hardware register associated with the predefined semantic **POSITION**, **D3DDECLUSAGE_DIFFUSE** and **D3DVSDE_DIFFUSE** match the register associated with **COLOR0**, and **D3DDECLUSAGE_TEXCOORD0** and **D3DVSDE_TEXCOORD0** match the register associated with **TEXCOORD0**.

The above declarations can also be written the following way using **cgD3D9ResourceToDeclUsage()** or **cgD3D8ResourceToInputRegister()**:

```

const D3DVERTEXELEMENT9 declaration[] = {
    { 0, 0 * sizeof(float),
      D3DDECLTYPE_FLOAT3, D3DDECLMETHOD_DEFAULT,
      cgD3D9ResourceToDeclUsage(CG_POSITION), 0 },
    { 0, 3 * sizeof(float),
      D3DDECLTYPE_D3DCOLOR, D3DDECLMETHOD_DEFAULT,
      cgD3D9ResourceToDeclUsage(CG_COLOR0), 0 },
    { 1, 4 * sizeof(float),
      D3DDECLTYPE_FLOAT2, D3DDECLMETHOD_DEFAULT,
      cgD3D9ResourceToDeclUsage(CG_TEXCOORD0), 0 },
    D3DD3CL_END()
};

```

```

DWORD declaration[] = {
    D3DVSD_STREAM(0),
    D3DVSD_REG(cgD3D8ResourceToInputRegister(CG_POSITION),
               D3DVSDT_FLOAT3),
    D3DVSD_REG(cgD3D8ResourceToInputRegister(CG_COLOR0),
               D3DVSDT_D3DCOLOR),
    D3DVSD_STREAM(1),
    D3DVSD_SKIP(4),
    D3DVSD_REG(cgD3D8ResourceToInputRegister(CG_TEXCOORD0),
               D3DVSDT_FLOAT2),
    D3DVSD_END()
};

```

```

D3DVSD_END()
};

```

If it is possible to do so, the functions **cgD3D9ResourceToDeclUsage()** and **cgD3D8ResourceToInputRegister()** convert a **CGresource** enumerated type into a Direct3D vertex shader input register:

```

BYTE  cgD3D9ResourceToDeclUsage(CGresource resource);
DWORD cgD3D8ResourceToInputRegister(CGresource resource);

```

If the resource is not a vertex shader input resource, the call to **cgD3D9ResourceToDeclUsage()** returns **CGD3D9_INVALID_REG** and the call to **cgD3D8ResourceToInputRegister()** returns **CGD3D8_INVALID_REG**.

To write the vertex declarations described above based on the program parameters, which eliminates the reference to any semantic, use **cgD3D9ResourceToDeclUsage()** or **cgD3D8ResourceToInputRegister()**:

```

CGparameter position =
    cgGetNamedParameter(program, "position");
CGparameter color =
    cgGetNamedParameter(program, "color");
CGparameter texCoord =
    cgGetNamedParameter(program, "texCoord");

const D3DVERTEXELEMENT9 declaration[] = {
    { 0, 0 * sizeof(float),
      D3DDECLTYPE_FLOAT3, D3DDECLMETHOD_DEFAULT,
      cgD3D9ResourceToDeclUsage(
          cgGetParameterResource(position)),
      cgGetParameterResourceIndex(position) },
    { 0, 3 * sizeof(float),
      D3DDECLTYPE_D3DCOLOR, D3DDECLMETHOD_DEFAULT,
      cgD3D9ResourceToDeclUsage(cgGetParameterResource(color)),
      cgGetParameterResourceIndex(color) },
    { 1, 4 * sizeof(float),
      D3DDECLTYPE_FLOAT2, D3DDECLMETHOD_DEFAULT,
      cgD3D9ResourceToDeclUsage(
          cgGetParameterResource(texCoord)),
      cgGetParameterResourceIndex(texCoord) },
    D3DDCL_END()
};

DWORD declaration[] = {
    D3DVSD_STREAM(0),
    D3DVSD_REG(cgD3D8ResourceToInputRegister(
        cgGetParameterResource(position)), D3DVSDT_FLOAT3),
    D3DVSD_REG(cgD3D8ResourceToInputRegister(

```

```

        cgGetParameterResource(color)), D3DVSDT_D3DCOLOR),
    D3DVSD_STREAM(1),
    D3DVSD_SKIP(4),
    D3DVSD_REG(cgD3D8ResourceToInputRegister(
        cgGetParameterResource(texCoord)), D3DVSDT_FLOAT2),
    D3DVSD_END()
};

```

The size specified as the second argument of the **D3DVSD_REG()** macro call of a Direct3D 8 declaration does not need to match the size of the corresponding parameter for the vertex declaration to be valid. Those sizes are specified to describe how the data is laid out in the streams, not to perform any type checking with the shader code. The data referred to by a **D3DVSD_REG()** macro call is expanded to the four floating point values of the corresponding hardware register, and the missing values are set to 0 for **x**, **y**, and **z**, and to 1 for **w**.

Minimal Interface Type Retrieval

Use **cgD3D9TypeToSize()** to retrieve the size of a **CGtype** enumerated type in terms of floating-point numbers:

```
DWORD cgD3D9TypeToSize(CGtype type);
```

More precisely, it is the number of floating-point values required to store a parameter of type **type**. This function does not apply to some types, like the **sampler** types, in which case it returns zero. It is useful because applications can determine how many floating-point values they have to provide to set the value of a given parameter.

Minimal Interface Program Examples

In this section we provide some code samples that illustrate how and when to use functions from the minimal interface to make Cg programs work with Direct3D. To enhance clarity, the examples do very little error checking, but a production application should check the return values of all Cg functions. The vertex and fragment programs below are referenced in [“Direct3D 9 Application” on page 92](#) and [“Direct3D 8 Application” on page 95](#).

Vertex Program

The following Cg code is assumed to be in a file called **VertexProgram.cg**.

```

void VertexProgram(
    in float4 position : POSITION,
    in float4 color    : COLOR0,
    in float4 texCoord : TEXCOORD0,
    out float4 position0 : POSITION,

```

```

    out float4 color0      : COLOR0,
    out float4 texCoord0   : TEXCOORD0,
    const uniform float4x4 ModelViewMatrix)
{
    position0 = mul(position, ModelViewMatrix);
    color0 = color;
    texCoord0 = texCoord;
}

```

Fragment Program

The following Cg code is assumed to be in a file called **FragmentProgram.cg**.

```

void FragmentProgram(
    in float4 color      : COLOR0,
    in float4 texCoord   : TEXCOORD0,
    out float4 color0    : COLOR0,
    const uniform sampler2D BaseTexture,
    const uniform float4 SomeColor)
{
    color0 = color * tex2D(BaseTexture, texCoord) + SomeColor;
}

```

Direct3D 9 Application

The following C code links the previous vertex and fragment programs to the Direct3D 9 application.

```

#include <cg/cg.h>
#include <cg/cgD3D9.h>

IDirect3DDevice9* device; // Initialized somewhere else
IDirect3DTexture9* texture; // Initialized somewhere else
D3DXMATRIX matrix; // Initialized somewhere else
D3DXCOLOR constantColor; // Initialized somewhere else
CGcontext context;
CGprogram vertexProgram, fragmentProgram;
IDirect3DVertexDeclaration9* vertexDeclaration;
IDirect3DVertexShader9* vertexShader;
IDirect3DPixelShader9* pixelShader;
CGparameter baseTexture, someColor, modelViewMatrix;

// Called at application startup
void OnStartup()
{
    // Create context
    context = cgCreateContext();
}

```

```

// Called whenever the Direct3D device needs to be created
void OnCreateDevice()
{
    // Create the vertex shader
    vertexProgram = cgCreateProgramFromFile(context, CG_SOURCE,
        "VertexProgram.cg", CG_PROFILE_VS_2_0, "VertexProgram", 0);
    CComPtr<ID3DXBuffer> byteCode;
    const char* progSrc = cgGetProgramString(vertexProgram,
        CG_COMPILED_PROGRAM);
    D3DXAssembleShader(progSrc, strlen(progSrc), 0, 0, 0,
        &byteCode, 0);
    // If your program uses explicit binding semantics (like
    // this one), you can create a vertex declaration
    // using those semantics.
    const D3DVERTEXELEMENT9 declaration[] = {
        { 0, 0 * sizeof(float),
          D3DDECLTYPE_FLOAT3, D3DDECLMETHOD_DEFAULT,
          D3DDECLUSAGE_POSITION, 0 },
        { 0, 3 * sizeof(float),
          D3DDECLTYPE_D3DCOLOR, D3DDECLMETHOD_DEFAULT,
          D3DDECLUSAGE_COLOR, 0 },
        { 0, 4 * sizeof(float),
          D3DDECLTYPE_FLOAT2, D3DDECLMETHOD_DEFAULT,
          D3DDECLUSAGE_TEXCOORD, 0 },
        D3DD3CL_END()
    };
    // Make sure the resulting declaration is compatible with
    // the shader. This is really just a sanity check.
    assert(cgD3D9ValidateVertexDeclaration(vertexProgram,
        declaration));

    device->CreateVertexDeclaration(
        declaration, &vertexDeclaration);
    device->CreateVertexShader(
        byteCode->GetBufferPointer(), &vertexShader);

    // Create the pixel shader.
    fragmentProgram = cgCreateProgramFromFile(context,
        CG_SOURCE, "FragmentProgram.cg",
        CG_PROFILE_PS_2_0, "FragmentProgram", 0);
    {
        CComPtr<ID3DXBuffer> byteCode;
        const char* progSrc = cgGetProgramString(fragmentProgram,
            CG_COMPILED_PROGRAM);
        D3DXAssembleShader(progSrc, strlen(progSrc), 0, 0, 0,

```

```

        &byteCode, 0);
device->CreatePixelShader(byteCode->GetBufferPointer(),
                        &pixelShader)
}

// Grab some parameters.
modelViewMatrix = cgGetNamedParameter(vertexProgram,
                                      "ModelViewMatrix");
baseTexture = cgGetNamedParameter(fragmentProgram,
                                   "BaseTexture");
someColor = cgGetNamedParameter(fragmentProgram,
                                "SomeColor");

// Sanity check that parameters have the expected size
assert(cgD3D9TypeToSize(cgGetParameterType(
                        modelViewMatrix)) == 16);
assert(cgD3D9TypeToSize(cgGetParameterType(someColor))
      == 4);
}

// Called to render the scene
void OnRender()
{
    // Get the Direct3D resource locations for parameters
    // This can be done earlier and saved
    DWORD modelViewMatrixRegister =
        cgGetParameterResourceIndex(modelViewMatrix);
    DWORD baseTextureUnit =
        cgGetParameterResourceIndex(baseTexture);
    DWORD someColorRegister =
        cgGetParameterResourceIndex(someColor);

    // Set the Direct3D state.
    device->SetVertexShaderConstantF(modelViewMatrixRegister,
                                    &matrix, 4);
    device->SetPixelShaderConstantF(someColorRegister,
                                    &constantColor, 1);
    device->SetVertexDeclaration(vertexDeclaration);
    device->SetTexture(baseTextureUnit, texture);
    device->SetVertexShader(vertexShader);
    device->SetPixelShader(pixelShader);

    // Draw scene.
    // ...
}

```

```

// Called before the device changes or is destroyed
void OnDestroyDevice() {
    vertexShader->Release();
    pixelShader->Release();
    vertexDeclaration->Release();
}

// Called before application shuts down
void OnShutdown() {
    // This frees any core runtime resources.
    // The minimal interface has no dynamic storage to free.
    cgDestroyContext(context);
}

```

Direct3D 8 Application

The following C code links the previous vertex and fragment programs to the Direct3D 8 application.

```

#include <cg/cg.h>
#include <cg/cgD3D8.h>

IDirect3DDevice8* device; // Initialized somewhere else
IDirect3DTexture8* texture; // Initialized somewhere else
D3DXMATRIX matrix; // Initialized somewhere else
D3DXCOLOR constantColor; // Initialized somewhere else
CGcontext context;
CGprogram vertexProgram, fragmentProgram;
DWORD vertexShader, pixelShader;
CGparameter baseTexture, someColor, modelViewMatrix;

// Called at application startup
void OnStartup()
{
    // Create context
    context = cgCreateContext();
}

// Called whenever the Direct3D device needs to be created
void OnCreateDevice()
{
    // Create the vertex shader
    vertexProgram = cgCreateProgramFromFile(context, CG_SOURCE,
        "VertexProgram.cg", CG_PROFILE_VS_1_1, "VertexProgram", 0);
    CComPtr<ID3DXBuffer> byteCode;
    const char* progSrc = cgGetProgramString(vertexProgram,

```

```

        CG_COMPILED_PROGRAM);
// Normally, you also grab the constants and prepend them
// to your vertex declaration. Not shown here for brevity.
D3DXAssembleShader(progSrc, strlen(progSrc), 0, 0, 0,
    &byteCode, 0);
// If your program uses explicit binding semantics (like
// this one), you can create a vertex declaration
// using those semantics.
DWORD declaration[] = {
    D3DVSD_STREAM(0),
    D3DVSD_REG(D3DVSDE_POSITION, D3DVSDT_FLOAT3),
    D3DVSD_REG(D3DVSDE_DIFFUSE, D3DVSDT_D3DCOLOR),
    D3DVSD_REG(D3DVSDE_TEXCOORD0, D3DVSDT_FLOAT2),
    D3DVSD_END()
}
// Make sure the resulting declaration is compatible with
// the shader. This is really just a sanity check.
assert(cgD3D8ValidateVertexDeclaration(vertexProgram,
    declaration));
// Create the shader handle using the declaration.
device->CreateVertexShader(declaration,
    byteCode->GetBufferPointer(), &vertexShader, 0);

// Create the pixel shader.
fragmentProgram = cgCreateProgramFromFile(context,
    CG_SOURCE, "FragmentProgram.cg",
    CG_PROFILE_PS_1_1, "FragmentProgram", 0);
{
    CComPtr<ID3DXBuffer> byteCode;
    const char* progSrc = cgGetProgramString(fragmentProgram,
        CG_COMPILED_PROGRAM);
    D3DXAssembleShader(progSrc, strlen(progSrc), 0, 0, 0,
        &byteCode, 0);
    device->CreatePixelShader(byteCode->GetBufferPointer(),
        &pixelShader);
}

// Grab some parameters.
modelViewMatrix = cgGetNamedParameter(vertexProgram,
    "ModelViewMatrix");
baseTexture = cgGetNamedParameter(fragmentProgram,
    "BaseTexture");
someColor = cgGetNamedParameter(fragmentProgram,
    "SomeColor");

```

```

// Sanity check that parameters have the expected size
assert(cgD3D8TypeToSize(cgGetParameterType(
                                modelViewMatrix)) == 16);
assert(cgD3D8TypeToSize(cgGetParameterType(someColor))
        == 4);
}

// Called to render the scene
void OnRender()
{
    // Get the Direct3D resource locations for parameters
    // This can be done earlier and saved
    DWORD modelViewMatrixRegister =
        cgGetParameterResourceIndex(modelViewMatrix);
    DWORD baseTextureUnit =
        cgGetParameterResourceIndex(baseTexture);
    DWORD someColorRegister =
        cgGetParameterResourceIndex(someColor);

    // Set the Direct3D state.
    device->SetVertexShaderConstant(modelViewMatrixRegister,
                                    &matrix, 4);
    device->SetPixelShaderConstant(someColorRegister,
                                   &constantColor, 1);
    device->SetTexture(baseTextureUnit, texture);
    device->SetVertexShader(vertexShader);
    device->SetPixelShader(pixelShader);

    // Draw scene.
    // ...
}

// Called before the device changes or is destroyed
void OnDestroyDevice() {
    device->DeleteVertexShader(vertexShader);
    device->DeletePixelShader(pixelShader);
}

// Called before application shuts down
void OnShutdown() {
    // This frees any core runtime resources.
    // The minimal interface has no dynamic storage to free.
    cgDestroyContext(context);
}

```

Direct3D Expanded Interface

If you use the expanded interface for a program, in order to avoid any unfortunate inconsistencies it is advisable to stick with the expanded interface for all shader-related operations that can be performed through its functions, such as shader setting, shader activation, and parameter setting—including setting texture stage states.

Setting the Direct3D Device

The expanded interface encapsulates more functionality than the minimal interface to ease program and parameter management. It does this by making the appropriate Direct3D calls at the appropriate times. Because some of these calls require the Direct3D device, it must be communicated to the Cg runtime:

```
HRESULT cgD3D9SetDevice(IDirect3DDevice9* device);
```

You can get the Direct3D device currently associated with the runtime using **cgD3D9GetDevice()**:

```
IDirect3DDevice9* cgD3D9GetDevice();
```

When **cgD3D9SetDevice()** is called with zero as an input, all Direct3D resources used by the expanded interface are released. Since a Direct3D device is destroyed only when all references to it are removed, the application should call **cgD3D9SetDevice()** with zero as an input when it is done with a Direct3D device so that it gets destroyed when the application shuts down. Otherwise, Direct3D does not shut down properly and reports memory leaks to the debug console.

Note that calling **cgD3D9SetDevice()** with zero as an input does not affect the Cg core runtime resources in any way: all the related core runtime handles (of type **CGprogram**, **CGparameter**, and so on) remain valid.

If you call **cgD3D9SetDevice()** a second time with a different device, all programs managed by the old device are rebuilt using the new device.

Responding to Lost Direct3D Devices

The expanded interface may hold references to Direct3D resources that need to be recreated in response to a lost device. In particular, certain sampler parameters might need to be released before a Direct3D device can be reset from a lost state. The expanded interface is holding a reference to a texture that needs to be reset in response to a lost device if both of the following are true for a texture:

- ❑ It was created in the **D3DPPOOL_DEFAULT** pool.

- ❑ It was bound to a sampler parameter (using **cgD3D9SetTexture()**) of a program for which parameter shadowing is enabled.

In this case, the parameter must be set to zero (using **cgD3D9SetTexture()**) to remove the expanded interface's reference to that texture so it can be destroyed and the Direct3D device can be reset from a lost state. Later, after resetting the Direct3D device and recreating the texture, it needs to be re-bound to the sampler parameter. For example,

```
IDirect3DDevice9* device; // Initialized elsewhere
IDirect3DTexture9* myDefaultPoolTexture;
CGprogram program;

void OneTimeLoadScene()
{
    // Load the program with cgD3D9LoadProgram and
    // enable parameter shadowing
    /* ... */
    cgD3D9LoadProgram(program, TRUE, 0, 0, 0);
    /* ... */
    // Bind sampler parameter
    GCparameter parameter;
    parameter = cgGetParameterByName(program, "MySampler");
    cgD3D9SetTexture(parameter, myDefaultPoolTexture);
}

void OnLostDevice()
{
    // First release all necessary resources
    PrepareForReset();
    // Next actually reset the Direct3D device
    device->Reset( /* ... */ );
    // Finally recreate all those resource
    OnReset();
}

void PrepareForReset()
{
    /* ... */
    // Release expanded interface reference
    cgD3D9SetTexture(mySampler, 0);
    // Release local reference
    // and any other references to the texture
    myDefaultPoolTexture->Release();
    /* ... */
}
```

```

void OnReset()
{
    // Recreate myDefaultPoolTexture in D3DP00L_DEFAULT
    /* ... */
    // Since the texture was just recreated,
    // it must be re-bound to the parameter
    GCparameter parameter;
    parameter = cgGetParameterByName(prog, "MySampler");
    cgD3D9SetTexture(mySampler, myDefaultPoolTexture);
    /* ... */
}

```

See the Direct3D documentation for a full explanation of lost devices and how to properly handle them.

Setting Expanded Interface Parameters

This section discusses setting the various types of parameters of the expanded interface, including uniform scalar, uniform vector, uniform matrix, uniform arrays of the three previous types, and sampler.

Setting Uniform Scalar, Vector, and Matrix Parameters

The function **cgD3D9SetUniform()** sets floating-point parameters like **float3** and **float4x3**:

```

HRESULT cgD3D9SetUniform(CGparameter parameter,
                        const void* value);

```

The amount of data required depends on the type of *parameter*, but is always specified as an array of one or more floating point values. The type is **void*** so a user-defined structure that is compatible can be passed in without type casting. Here is some code illustrating the use of **cgD3D9SetUniform()** for setting a *vectorParam* of type **float3**, *matrixParam* of type **float2x3**, and *arrayParam* of type **float2x2[3]**:

```

D3DXVECTOR3 vectorData(1,2,3);
float matrixData[2][3] = {{1, 2, 3}, {4, 5, 6}};
float arrayData[3][2][2] =
    {{{1, 2}, {3, 4}}, {{5, 6}, {7, 8}}, {{9, 10}, {11, 12}}};
cgD3D9SetUniform(vectorParam, &vectorData);
cgD3D9SetUniform(matrixParam, matrixData);
cgD3D9SetUniform(arrayParam, arrayData);

```

As mentioned previously, **cgD3D9TypeToSize()** can be used to determine how many values are required for setting a parameter of a particular type.

For convenience, there is also a function to set a parameter from a 4x4 matrix of type **D3DMATRIX**:

```
HRESULT cgD3D9SetUniformMatrix(CGparameter parameter,
                                const D3DMATRIX* matrix);
```

The upper-left portion of the matrix is extracted to fit the size of the input parameter, so that you could set *matrixParam* this way as well:

```
D3DMATRIX matrix(
    1, 1, 1, 0,
    1, 1, 1, 0,
    0, 0, 0, 0,
    0, 0, 0, 0,
);
cgD3D9SetUniformMatrix(matrixParam, &matrix);
```

In the example above, every element of *matrixParam* is set to 1.

Setting Uniform Arrays of Scalar, Vector, and Matrix Parameters

To set an *array* parameter, use **cgD3D9SetUniformArray()**:

```
HRESULT cgD3D9SetUniformArray(CGparameter parameter,
                                DWORD startIndex, DWORD numberOfElements,
                                const void* array);
```

The parameters *startIndex* and *numberOfElements* specify which elements of the array parameter are set: Those are the *numberOfElements* elements of indices ranging from *startIndex* to *startIndex + numberOfElements - 1*. It is assumed that *array* contains enough values to set all those elements. As with **cgD3D9SetUniform()**, **cgD3D9TypeToSize()** can be used to determine how many values are required, and the type is **void*** so a compatible user-defined structure can be passed in without type casting.

There is a convenience function equivalent to **cgD3D9SetUniformMatrix()**:

```
HRESULT cgD3D9SetUniformMatrixArray(CGparameter parameter,
                                       DWORD startIndex, DWORD numberOfElements,
                                       const D3DMATRIX* matrices);
```

The parameters *startIndex* and *numberOfElements* have the same meanings as for **cgD3D9SetUniformMatrix()**.

The upper-left portion of each matrix of the array *matrices* is extracted to fit the size of the element of the array parameter *parameter*. Array *matrices* is assumed to have *numberOfElements* elements.

Setting Sampler Parameters

You assign a Direct3D texture to a sampler parameter using

```
HRESULT cgD3D9SetTexture(CGparameter parameter,  
                          IDirect3DBaseTexture9* texture);
```

To set the sampler state in the Direct3D 9 Cg runtime, use

```
HRESULT cgD3D9SetSamplerState(CGparameter parameter,  
                               D3DSAMPLERSTATETYPE type, DWORD value);
```

Parameter **type** is any of the **D3DSAMPLERSTATETYPE** enumerants and parameter **value** is a value appropriate for the corresponding **type**. Here is an example of how to use this function:

```
cgD3D9SetSamplerState(parameter, D3DSAMP_MAGFILTER,  
                       D3DTEXF_LINEAR);
```

To set the texture stage state in the Direct3D 8 Cg runtime, use:

```
HRESULT cgD3D8SetTextureStageState(CGparameter parameter,  
                                    D3DTEXTURESTAGESTATETYPE type, DWORD value);
```

Parameter **type** must be one of the following values:

D3DTSS_ADDRESSU	D3DTSS_ADDRESSV
D3DTSS_ADDRESSW	D3DTSS_BORDERCOLOR
D3DTSS_MAGFILTER	D3DTSS_MINFILTER
D3DTSS_MIPFILTER	D3DTSS_MIPMAPLODBIAS
D3DTSS_MAXMIPLEVEL	D3DTSS_MAXANISOTROPY

Parameter **value** is a value appropriate for the corresponding **type**. Here is an example of how to use this function:

```
cgD3D8SetTextureStageState(parameter, D3DTSS_MAGFILTER,  
                             D3DTEXF_LINEAR);
```

The texture wrap mode is set using

```
HRESULT cgD3D9SetTextureWrapMode(CGparameter parameter,  
                                   DWORD value);
```

The input **value** is either zero or a combination of **D3DWRAP_U**, **D3DWRAP_V**, and **D3DWRAP_W**. Here is an example of how to use this function:

```
cgD3D9SetTextureWrapMode(parameter, D3DWRAP_U | D3DWRAP_V);
```

Parameter Shadowing

Parameter shadowing can be enabled or disabled on a per-program basis:

- ❑ When loading the program (see [“Expanded Interface Program Execution” on page 103](#))

- At any time using

```
HRESULT cgD3D9EnableParameterShadowing(
    CGprogram program, CGbool enable);
```

for which **enable** should be set to **CG_TRUE** to enable parameter shadowing and to **CG_FALSE** to disable it.

To know if parameter shadowing is enabled for a given program, use:

```
CGbool cgD3D9IsParameterShadowingEnabled(CGprogram program);
```

This function returns **CG_TRUE** if parameter shadowing is enabled for **program**.

Expanded Interface Program Execution

To load a program in Direct3D 9 use **cgD3D9LoadProgram()**:

```
HRESULT cgD3D9LoadProgram(CGprogram program,
    CG_BOOL parameterShadowingEnabled,
    DWORD assembleFlags);
```

This function assembles the result of the compilation of **program** using **D3DXAssembleShader()** with **assembleFlags** as the **D3DXASM** flags.

Depending on the program's profile, it then either uses

IDirect3DDevice9::CreateVertexShader() to create a Direct3D 9 vertex shader, or uses **IDirect3DDevice9::CreatePixelShader()** to create a Direct3D 9 pixel shader.

Here is a typical use of the function:

```
HRESULT hresult = cgD3D9LoadProgram(vertexProgram, TRUE,
    D3DXASM_DEBUG);
HRESULT hresult = cgD3D9LoadProgram(fragmentProgram, TRUE, 0);
```

To load a program in Direct3D 8 use **cgD3D8LoadProgram()**:

```
HRESULT cgD3D8LoadProgram(CGprogram program,
    BOOL parameterShadowingEnabled, DWORD assembleFlags,
    DWORD vertexShaderUsage, const DWORD* declaration);
```

This function assembles the result of the compilation of **program** using **D3DXAssembleShader()** with **assembleFlags** as the **D3DXASM** flags.

Depending on the program's profile, it then either uses

IDirect3DDevice8::CreateVertexShader() to create a Direct3D vertex shader with **declaration** as the vertex declaration and **vertexShaderUsage** as the usage control, or uses **IDirect3DDevice8::CreatePixelShader()** to create a Direct3D pixel shader.

The value of ***parameterShadowingEnabled*** should be set to **TRUE** to enable parameter shadowing for the program. This behavior can be changed after the program is created by calling **cgD3D8EnableParameterShadowing()**. Here is a typical use of the function:

```
HRESULT hresult = cgD3D8LoadProgram(vertexProgram, TRUE,
                                     D3DXASM_DEBUG, D3DUSAGE_SOFTWAREVERTEXPROCESSING,
                                     declaration);
HRESULT hresult = cgD3D8LoadProgram(fragmentProgram, TRUE,
                                     0, 0, 0);
```

If you want to apply the same vertex program to several sets of geometric data, each having a different layout, you need to load the program with different vertex declarations in Direct3D 8. To do so, you need to make a duplicate of the program, using **cgCopyProgram()**, for each of these declarations. Here is a code sample illustrating this operation:

```
CGprogram program1, program2;
program1 = cgCreateProgramFromFile(context, CG_SOURCE,
                                   "VertexProgram.cg", CG_PROFILE_VS_1_1, 0, 0);
const DWORD declaration1 =
    cgD3D8GetVertexDeclaration(program1);
cgD3D8LoadProgram(program1, TRUE, 0, 0, declaration1);
program2 = cgCopyProgram(program1);
const DWORD declaration2[] = {
    //... Custom declaration ...
};
if (cgD3D8ValidateVertexDeclaration(program2, declaration2))
    cgD3D8LoadProgram(program2, TRUE, 0, 0, declaration2);
```

Only the loading functions differ between Direct3D 9 and Direct3D 8; the unloading and binding functions are the same.

To release the Direct3D resources allocated by **cgD3D9LoadProgram()**, such as the Direct3D shader object and any shadowed parameter, use

```
HRESULT cgD3D9UnloadProgram(CGprogram program);
```

Note that **cgD3D9UnloadProgram()** does not free any core runtime resources, such as ***program*** and any of its parameter handles. On the other hand, destroying a program with **cgDestroyProgram()** or **cgDestroyContext()** releases any Direct3D resources by indirectly calling **cgD3D9UnloadProgram()**.

Function **cgD3D9IsProgramLoaded()** returns **CG_TRUE** if a ***program*** is loaded:

```
CGbool cgD3D9IsProgramLoaded(CGprogram program);
```

All programs must be loaded before they can be bound. Binding a program is done by calling **cgD3D9BindProgram()**:

```
HRESULT cgD3D9BindProgram(CGprogram program);
```

This function basically activates the Direct3D shader corresponding to **program** by calling **IDirect3DDevice9::SetVertexShader()** or **IDirect3DDevice9::SetPixelShader()** depending on the program's profile. If parameter shadowing is enabled for **program**, it also sets all the shadowed parameters and their associated Direct3D states (such as texture stage states for the **sampler** parameters). No value or state tracking is performed by the runtime so that this setting is done regardless of what the current values of these parameters or of their states are. If a shadowed parameter has not been set by the time **cgD3D9BindProgram()** is called, no Direct3D call of any sort is issued for this parameter.

Only one vertex program and one fragment program can be bound at any given time, so binding a program of a given type implicitly unbinds any other program of the same type.

Expanded Interface Profile Support

Two convenient functions are provided that give the highest vertex and pixel shader versions supported by the device:

```
CGprofile cgD3D9GetLatestVertexProfile();  
CGprofile cgD3D9GetLatestPixelProfile();
```

This allows you to make your application future-ready, because the Cg programs are automatically compiled for the best profiles that are available at runtime, even if these profiles did not exist at the time the application was written. Another function that allows you optimal compilation is **cgD3D9GetOptimalOptions()**. It returns a string representing the optimal set of compiler options for a given profile:

```
char const* cgD3D9GetOptimalOptions(CGprofile profile);
```

This string is meant to be used as part of the argument parameter to **cgCreateProgram()**. It does not need to be destroyed by the application. However, its content could change if **cgD3D9GetOptimalOptions()** is called again for the same profile but for a different Direct3D device.

Expanded Interface Program Examples

In this section we provide programs that illustrates how and when to use functions from the expanded interface to make Cg programs work with Direct3D. For the sake of clarity, the examples do very little error checking, but a production application should check the return values of all Cg

functions. The vertex and fragment programs that follow are referenced in [“Expanded Interface Direct3D 9 Application” on page 106](#) and [“Expanded Interface Direct3D 8 Application” on page 109](#).

Expanded Interface Vertex Program

The following Cg code is assumed to be in a file called **VertexProgram.cg**.

```
void VertexProgram(
    in float4 position : POSITION,
    in float4 color     : COLOR0,
    in float4 texCoord  : TEXCOORD0,
    out float4 position0 : POSITION,
    out float4 color0    : COLOR0,
    out float4 texCoord0 : TEXCOORD0,
    const uniform float4x4 ModelViewMatrix)
{
    position0 = mul(position, ModelViewMatrix);
    color0 = color;
    texCoord0 = texCoord; }

```

Expanded Interface Fragment Program

The following Cg code is assumed to be in a file called **FragmentProgram.cg**.

```
void FragmentProgram(
    in float4 color     : COLOR0,
    in float4 texCoord  : TEXCOORD0,
    out float4 color0   : COLOR0,
    const uniform sampler2D BaseTexture,
    const uniform float4 SomeColor)
{
    color0 = color * tex2D(BaseTexture, texCoord) + SomeColor;
}

```

Expanded Interface Direct3D 9 Application

The following C code links the previous vertex and fragment programs to the Direct3D 9 application.

```
#include <cg/cg.h>
#include <cg/cgD3D9.h>

IDirect3DDevice9* device; // Initialized somewhere else
IDirect3DTexture9* texture; // Initialized somewhere else
D3DXCOLOR constantColor; // Initialized somewhere else
CGcontext context;
IDirect3DVertexDeclaration9* vertexDeclaration;
CGprogram vertexProgram, fragmentProgram;
CGparameter baseTexture, someColor, modelViewMatrix;

```

```

// Called at application startup
void OnStartup()
{
    // Create context
    context = cgCreateContext();
}

// Called whenever the Direct3D device needs to be created
void OnCreateDevice()
{
    // Pass the Direct3D device to the expanded interface.
    cgD3D9SetDevice(device);

    // Determine the best profiles to use
    CGprofile vertexProfile = cgD3D9GetLatestVertexProfile();
    CGprofile pixelProfile = cgD3D9GetLatestPixelProfile();

    // Grab the optimal options for each profile.
    const char* vertexOptions[] = {
        cgD3D9GetOptimalOptions(vertexProfile), 0 };
    const char* pixelOptions[] = {
        cgD3D9GetOptimalOptions(pixelProfile), 0 };

    // Create the vertex shader.
    vertexProgram = cgCreateProgramFromFile(
        context, CG_SOURCE, "VertexProgram.cg",
        vertexProfile, "VertexProgram", vertexOptions);
    // If your program uses explicit binding semantics, you
    // can create a vertex declaration using those semantics.
    const D3DVERTEXELEMENT9 declaration[] = {
        { 0, 0 * sizeof(float),
          D3DDECLTYPE_FLOAT3, D3DDECLMETHOD_DEFAULT,
          D3DDECLUSAGE_POSITION, 0 },
        { 0, 3 * sizeof(float),
          D3DDECLTYPE_D3DCOLOR, D3DDECLMETHOD_DEFAULT,
          D3DDECLUSAGE_COLOR, 0 },
        { 0, 4 * sizeof(float),
          D3DDECLTYPE_FLOAT2, D3DDECLMETHOD_DEFAULT,
          D3DDECLUSAGE_TEXCOORD, 0 },
        D3DD3CL_END()
    };

    // Ensure the resulting declaration is compatible with the
    // shader. This is really just a sanity check.

```

```

assert(cgD3D9ValidateVertexDeclaration(vertexProgram,
                                       declaration));
device->CreateVertexDeclaration(
    declaration, &vertexDeclaration);
// Load the program with the expanded interface.
// Parameter shadowing is enabled (second parameter = TRUE).
cgD3D9LoadProgram(vertexProgram, TRUE, 0);

// Create the pixel shader.
fragmentProgram = cgCreateProgramFromFile(
    context, CG_SOURCE, "FragmentProgram.cg",
    pixelProfile, "FragmentProgram", pixelOptions);

// Load the program with the expanded interface. Parameter
// shadowing is enabled (second parameter = TRUE). Ignore
// vertex shader specific flags, such as declaration usage.
cgD3D9LoadProgram(fragmentProgram, TRUE, 0);

// Grab some parameters.
modelViewMatrix = cgGetNamedParameter(vertexProgram,
                                       "ModelViewMatrix");
baseTexture = cgGetNamedParameter(fragmentProgram,
                                   "BaseTexture");
someColor = cgGetNamedParameter(fragmentProgram,
                                 "SomeColor");

// Sanity check that parameters have the expected size
assert(cgD3D9TypeToSize(cgGetParameterType(
    modelViewMatrix)) == 16);
assert(cgD3D9TypeToSize(cgGetParameterType(someColor))
    == 4);

// Set parameters that don't change. They can be set
// only once since parameter shadowing is enabled
cgD3D9SetTexture(baseTexture, texture);
cgD3D9SetUniform(someColor, &constantColor);
}

// Called to render the scene
void OnRender()
{
    // Load model-view matrix.
    D3DXMATRIX modelViewMatrix;
    // ...

```

```

// Set the parameters that change every frame
// This must be done before binding the programs
cgD3D9SetUniformMatrix(modelViewMatrix, &modelViewMatrix);

// Set the vertex declaration
device->SetVertexDeclaration(vertexDeclaration);

// Bind the programs. This downloads any parameter values
// that have been previously set.
cgD3D9BindProgram(vertexProgram);
cgD3D9BindProgram(fragmentProgram);

// Draw scene.
// ...
}

// Called before the device changes or is destroyed
void OnDestroyDevice()
{
    // Calling this function tells the expanded interface to
    // release its internal reference to the Direct3D device
    // and free its Direct3D resources.
    cgD3D9SetDevice(0);
}

// Called before application shuts down
void OnShutdown()
{
    // This frees any core runtime resource.
    cgDestroyContext(context);
}

```

Expanded Interface Direct3D 8 Application

The following C code links the previous vertex and fragment programs to the Direct3D 8 application.

```

#include <cg/cg.h>
#include <cg/cgD3D8.h>

IDirect3DDevice8* device; // Initialized somewhere else
IDirect3DTexture8* texture; // Initialized somewhere else
D3DXCOLOR constantColor; // Initialized somewhere else
CGcontext context;
CGprogram vertexProgram, fragmentProgram;
CGparameter baseTexture, someColor, modelViewMatrix;

```

```

// Called at application startup
void OnStartup()
{
    // Create context
    context = cgCreateContext();
}

// Called whenever the Direct3D device needs to be created
void OnCreateDevice()
{
    // Pass the Direct3D device to the expanded interface.
    cgD3D8SetDevice(device);

    // Determine the best profiles to use
    CGprofile vertexProfile = cgD3D8GetLatestVertexProfile();
    CGprofile pixelProfile = cgD3D8GetLatestPixelProfile();

    // Grab the optimal options for each profile.
    const char* vertexOptions[] = {
        cgD3D8GetOptimalOptions(vertexProfile), 0 };
    const char* pixelOptions[] = {
        cgD3D8GetOptimalOptions(pixelProfile), 0 };

    // Create the vertex shader.
    vertexProgram = cgCreateProgramFromFile(
        context, CG_SOURCE, "VertexProgram.cg",
        vertexProfile, "VertexProgram", vertexOptions);
    // If your program uses explicit binding semantics (like
    // this one), you can create a vertex declaration
    // using those semantics.
    DWORD declaration[] = {
        D3DVSD_STREAM(0),
        D3DVSD_REG(D3DVSDE_POSITION, D3DVSDT_FLOAT3),
        D3DVSD_REG(D3DVSDE_DIFFUSE, D3DVSDT_D3DCOLOR),
        D3DVSD_REG(D3DVSDE_TEXCOORD0, D3DVSDT_FLOAT2),
        D3DVSD_END()
    }

    // Ensure the resulting declaration is compatible with the
    // shader. This is really just a sanity check.
    assert(cgD3D8ValidateVertexDeclaration(vertexProgram,
        declaration));

    // Load the program with the expanded interface.
    // Parameter shadowing is enabled (second parameter = TRUE).

```

```

cgD3D8LoadProgram(vertexProgram, TRUE, 0, 0, declaration);

// Create the pixel shader.
fragmentProgram = cgCreateProgramFromFile(
    context, CG_SOURCE, "FragmentProgram.cg",
    pixelProfile, "FragmentProgram", pixelOptions);

// Load the program with the expanded interface.
// Parameter shadowing is enabled (second parameter = TRUE).
// Ignore vertex shader specific flags, like declaration and
// usage.
cgD3D8LoadProgram(fragmentProgram, TRUE, 0, 0, 0);

// Grab some parameters.
modelViewMatrix = cgGetNamedParameter(vertexProgram,
                                       "ModelViewMatrix");
baseTexture = cgGetNamedParameter(fragmentProgram,
                                   "BaseTexture");
someColor = cgGetNamedParameter(fragmentProgram,
                                 "SomeColor");

// Sanity check that parameters have the expected size
assert(cgD3D8TypeToSize(cgGetParameterType(
    modelViewMatrix)) == 16);
assert(cgD3D8TypeToSize(cgGetParameterType(someColor))
    == 4);

// Set parameters that don't change. They can be set
// only once since parameter shadowing is enabled
cgD3D8SetTexture(baseTexture, texture);
cgD3D8SetUniform(someColor, &constantColor);
}

// Called to render the scene
void OnRender()
{
    // Load model-view matrix.
    D3DXMATRIX modelViewMatrix;
    // ...

    // Set the parameters that change every frame
    // This must be done before binding the programs
    cgD3D8SetUniformMatrix(modelViewMatrix, &modelViewMatrix);

    // Bind the programs. This downloads any parameter values

```

```

    // that have been previously set.
    cgD3D8BindProgram(vertexProgram);
    cgD3D8BindProgram(fragmentProgram);

    // Draw scene.
    // ...
}

// Called before the device changes or is destroyed
void OnDestroyDevice()
{
    // Calling this function tells the expanded interface to
    // release its internal reference to the Direct3D device
    // and free its Direct3D resources.
    cgD3D8SetDevice(0);
}

// Called before application shuts down
void OnShutdown()
{
    // This frees any core runtime resource.
    cgDestroyContext(context);
}

```

Direct3D Debugging Mode

In addition to the error reporting mechanisms described in [“Direct3D Error Reporting” on page 114](#), a debug version of the Direct3D 9 or Direct3D 8 Cg runtime DLL is provided to assist you with the development of applications using the Direct3D 9 or Direct3D 8 Cg runtime. This version does not have debug symbols, but when used in place of the regular version, it uses the Win32 function **OutputDebugString()** to output many helpful messages and traces to the debug output console. Examples of information the debug DLL outputs are the following:

- ❑ Any Direct3D or Cg core runtime errors
- ❑ Debugging information about parameters that are managed by the expanded interface
- ❑ Potential performance warnings

Here is a sample trace:

```

cgD3D(TRACE): Creating vertex shader for program 3
cgD3D(TRACE): Discovering parameters for vertex program 3
cgD3D(TRACE): Discovered uniform parameter 'ModelViewProj'
of type float4x4

```


3. Turn on and turn off tracing of portions of your code using **cgD3D9EnableDebugTracing()**:

```
void cgD3D9EnableDebugTracing(CGbool enable);
```

Here is how you would enable debug tracing for part of the application code:

```
cgD3D9EnableDebugTracing(CG_TRUE);
// ...
// Application code that is traced
// ...
cgD3D9EnableDebugTracing(CG_FALSE);
```

Note that each debug trace output sets an error equal to **cgD3D9DebugTrace**. So, if an error callback has been registered with the core runtime using **cgSetErrorCallback()**, each debug trace output triggers a call to this error callback (see [“Using Error Callbacks” on page 116](#)).

Direct3D Error Reporting

Error reporting in Cg includes defined error types, functions that allow testing for errors, and support for error callbacks.

Direct3D Error Types

The Direct3D runtime generates errors of type **CGerror**, reported by the Cg core runtime and of type **HRESULT**, reported by the Direct3D runtime. In addition, it returns the errors listed in the next two groups that are specific to the Direct3D Cg runtime.

❑ **CGerror**

- ↳ **cgD3D9Failed**: Set when a Direct3D runtime function makes a Direct3D call that returns an error.
- ↳ **cgD3D9DebugTrace**: Set when a debug message is output to the debug console when using the debug DLL (see [“Direct3D Debugging Mode” on page 112](#)).

❑ **HRESULT**

- ↳ **CGD3D9ERR_INVALIDPARAM**: Returned when a parameter value cannot be set.
- ↳ **CGD3D9ERR_INVALIDPROFILE**: Returned when a program with an unexpected profile is passed to a function.
- ↳ **CGD3D9ERR_INVALIDSAMPLERSTATE**: Returned when a parameter of type **D3DTEXTURESTAGESTATETYPE**, which is not a valid **sampler** state, is passed to a **sampler** state function.

- ✚ **CGD3D9ERR_INVALIDVEREXDECL:** Returned when a program is loaded with the expanded interface, but the given declaration is incompatible.
- ✚ **CGD3D9ERR_NODEVICE:** Returned when a required Direct3D device is 0. This typically occurs when an expanded interface function is called and a Direct3D device has not been set with **cgD3D9SetDevice()**.
- ✚ **CGD3D9ERR_NOTMATRIX:** Returned when a parameter that is not a matrix type is passed to a function that expects one.
- ✚ **CGD3D9ERR_NOTLOADED:** Returned when a parameter has not been loaded with the expanded interface by **cgD3D9LoadProgram()**.
- ✚ **CGD3D9ERR_NOTSAMPLER:** Returned when a parameter that is not a **sampler** parameter is passed to a function that expects one.
- ✚ **CGD3D9ERR_NOTUNIFORM:** Returned when a parameter that is not uniform is passed to a function that expects one.
- ✚ **CGD3D9ERR_NULLVALUE:** Returned when a value of zero is passed to a function that requires a non-zero value.
- ✚ **CGD3D9ERR_OUTOFRANGE:** Returned when an array range specified to a function is out of range.
- ✚ **CGD3D9_INVALID_REG:** Returned when a register number is requested for an invalid parameter type. This error is specific to the minimal interface functions and does not trigger an error callback.

Testing for Errors

When a Direct3D runtime function is called that returns an error of type **HRESULT**, the proper method of testing for success or failure is to use the Win32 macros **FAILED()** and **SUCCEEDED()**. Simply testing the error against zero or **D3D_OK** is not sufficient, because there could be more than one success value.

As an added convenience, and for uniformity with the core runtime, the Direct3D runtime also supplies **cgD3D9GetLastError()**, which is analogous to **cgGetLastError()** but returns the last Direct3D runtime error of type **HRESULT** for which the **FAILED()** macro returns **TRUE**:

```
HRESULT cgD3D9GetLastError();
```

The last error is always cleared immediately after the call.

The function **cgD3D9TranslateHRESULT()** converts an error of type **HRESULT** into a string:

```
const char* cgD3D9TranslateHRESULT(HRESULT hr);
```

This function should be called instead of **DXGetErrorDescription9()** because it also translates errors that the Cg Direct3D runtime generates.

Using Error Callbacks

Here is an example of a possible error callback that sorts out debug trace errors from core runtime errors and from Direct3D runtime errors:

```
void MyErrorCallback() {
    CGError error = cgGetError();
    if (error == cgD3D9DebugTrace) {
        // This is a debug trace output.
        // A breakpoint could be set here to step from one
        // debug output to the other.
        return;
    }
    char buffer[1024];
    if (error == cgD3D9Failed)
        sprintf(buffer, "A Direct3D error occurred: %s'\n",
                cgD3D9TranslateHRESULT(cgD3D9GetLastError()));
    else
        sprintf(buffer, "A Cg error occurred: '%s'\n",
                cgD3D9TranslateCGError(error));
    OutputDebugString(buffer);
}
cgSetErrorCallback(MyErrorCallback);
```

Introduction to CgFX

CgFX Overview

CgFX is an extended file format for Cg. In addition to Cg programs, CgFX files can also represent both fixed-function graphics state and meta-information about shader parameters. The CgFX API makes it possible to load CgFX effects files, traverse the data in them, set the associated graphics state, and so on. This chapter introduces this new API and the ideas behind it and is intended to make it easy to get started using CgFX.

This chapter assumes that the OpenGL state manager, implemented as part of the CgGL runtime, is being used. Because CgFX allows for extensible, custom state managers, alternate state managers that accept different state syntax may also be available. For example, a Direct3D state manager might accept Direct3D-style state names, while a Direct3D Under OpenGL state manager might accept Direct3D-style state names, but allow for rendering using OpenGL.

Key Concepts

Effect

An *effect* file contains a collection of shader source code, parameters, and rendering techniques. An effect encapsulates one or more different methods to render a particular visual effect. For example, the effect might provide one approach intended for use on fixed-function hardware, and a different approach on more modern, programmable hardware.

Technique

Each effect contains one or more *techniques*. A technique is intended to encapsulate the information needed to produce a visual effect—graphics state, shaders, and at least one rendering pass.

Pass

Each technique contains one or more rendering *passes*. Passes store graphics state, possibly including fixed-function state settings and vertex and

fragment shaders. The passes are generally processed in order: CgFX sets the graphics state for a pass, the application draws the scene geometry, the state for the next pass is set, geometry is drawn again, and so on.

State assignment

Passes hold *state assignments* that describe the graphics state for the pass.

Annotation

Annotations make it possible to associate meta-data with parameters, techniques, passes, and so on. For example, a parameter like ***LightIntensity*** might have annotations indicating the minimum and maximum valid values for the parameter.

Effect parameter

Parameters declared in the global scope of the effect file are *effect parameters*. Effect parameter values may be set and queried using the Cg runtime API. Effect parameters may be referenced on the right-hand side of state assignments and also as global parameters within Cg functions and programs defined within the effect.

Getting Started

We expect that the reader is generally familiar with the Cg runtime. See [“Introduction to the Cg Runtime Library” on page 43](#) for more details.

Consider the following effect:

```
float3 DiffuseColor<
string type = "color";
float3 minValue = float3(0,0,0);
float3 maxValue = float3(10,10,10);
> = { 1, 1, 1 };

technique FixedFunctionLighting {
    pass {
        LightingEnable = true;
        LightEnable[0] = true;
        LightPosition[0] = float4(-10, 10, 10, 1);
        LightAmbient[0] = float4(.1, .1, .1, .1);
        LightDiffuse[0] = (float4(2*DiffuseColor, 1));
        LightSpecular[0] = float4(1,1,1,1);

        MaterialShininess = 10.f;
        MaterialAmbient = float4(1,1,1,1);
```

```

        MaterialDiffuse = float4(.5, .5, .5, 1);
        MaterialSpecular = float4(.5, .5, .5, 1);
    }
}

```

The effect defines a single effect parameter, **DiffuseColor**, with three associated annotations: a string named **type** and two **float3s** named **minValue** and **maxValue**. These annotations exist purely for the use of the application using the effect file; the Cg runtime does not interpret the annotation names or values in any way. The effect parameter is initialized to the value [1,1,1].

The effect also defines a single technique, named **FixedFunctionLighting**, which in turn contains a single rendering pass. The rendering pass sets the appropriate OpenGL state to perform per-vertex lighting using the built-in fixed-function material model of OpenGL. The complete set of supported OpenGL states is listed in the section “[OpenGL State](#)” on page 129.

Note that the **LightDiffuse[0]** state value, corresponding to the fixed-function light's diffuse color, is set with an expression involving the **DiffuseColor** effect parameter. If the value of this parameter is changed by the application and the pass's state is later set, the parameter's new value is used in the expression that sets the light's diffuse color.

Note also that this expression is parenthesized. In general, CgFX requires that most expressions, like this one, involving effect parameters be in parenthesis. This is necessary so that CgFX can distinguish between effect parameters and built-in enumerant values representing constants.

The code below demonstrates how to create an effect given the name of an effect file. After creating a Cg context, **cgGLRegisterStates()** sets up the state assignments that support the standard OpenGL state manager. Most applications will want to do this immediately after creating the **CGcontext**. Next, the effect is created and associated with the given context.

```

CGcontext context = cgCreateContext();
cgGLRegisterStates(context);
CGeffect effect = cgCreateEffectFromFile(context,
"simple.cgfx", NULL);
    if (!effect) {
        fprintf(stderr, "Unable to create effect!\n");
        const char *listing = cgGetLastListing(context);
        if (listing)
            fprintf(stderr, "%s\n", listing);
        exit(1);
    }

```

Technique Validation

Before using any of the techniques in an effect, it's important to validate the techniques. Validation fails, for instance, if a technique includes a "compile" state assignment that references a profile that isn't supported on the current graphics hardware. Similarly, validation fails if the technique includes a state assignment that uses an unsupported OpenGL extension. Effects are commonly written such that the application can iterate over the given techniques in order and then choose the first technique that passes validation to apply the effect. For this reason, techniques are usually given in order of decreasing quality.

The code below iterates through the techniques in a **CGeffect** in turn, attempting to validate each of them and printing an error for the ones that fail.

```
CGtechnique technique = cgGetFirstTechnique(effect);
while (technique) {
    if (cgValidateTechnique(technique) == CG_FALSE)
        fprintf(stderr,
            "Technique %s did not validate. Skipping.\n",
            cgGetTechniqueName(technique));
    technique = cgGetNextTechnique(technique);
}
```

The function **cgIsTechniqueValidated()** can be used to check if the given technique has been validated.

Note that any Cg programs referenced in a technique are not compiled until the technique is validated. This makes it possible to modify the uncompiled program by connecting concrete shared **structs** to interface effect parameters, marking uniforms as literals, changing the program's profile, and so on.

Passes and Pass State

The heart of CgFX is applying the state defined in the passes in a technique. The loop below demonstrates the standard approach for looping over a technique's passes and applying their states in turn.

```
CGpass pass = cgGetFirstPass(technique);
while (pass) {
    cgSetPassState(pass);
    drawGeom();
    cgResetPassState(pass);
    pass = cgGetNextPass(pass);
}
```

Each of the state assignments in a pass translates directly to an OpenGL API call. For example, **LightingEnable = true;** translates to the call **glEnable(GL_LIGHTING)**, and **LightPosition[0] = float4(-10, 10, 10, 1)** translates to the call **glLightfv(GL_LIGHT0, GL_POSITION, v)** where **v** is an array of four **GLfloat** values.

Before or after the call to **cgSetPassState()**, the application is of course free to set other OpenGL state as desired. However, any state set before the call to **cgSetPassState()** may be overridden by the pass.

Note that if the technique containing the indicated pass has not been validated, calling **cgSetStatePass()** triggers an attempted validation of the technique. If validation fails, a runtime error results.

After the geometry has been drawn, **cgResetPassState()** resets the state that was set by the pass to the default values as specified by OpenGL. Note that it does *not* reset state to its values before **cgSetPassState()**—an application that desires this behavior should either push and pop OpenGL state, or should manually examine the state assignments in the pass in order to determine what state was changed, so that it can set it back to the desired values. (The routines to manually traverse the state in a pass are explained in “OpenGL State” on page 129.)

Effect Parameters

Handles to effect parameters can be retrieved using **cgGetNamedEffectParameter()**. Given such a handle, the name of the parameter can be found with **cgGetParameterName()**, its value can be set using the Cg runtime value-setting entry points, and so on.

```
CGparameter c = cgGetNamedEffectParameter(effect, "Color");
cgSetParameter3fv(c, Color);
CGparameter mvp = cgGetNamedEffectParameter(effect,
                                              "ModelViewProjection");
cgGLSetStateMatrixParameter(mvp,
                             CG_GL_MODELVIEW_PROJECTION_MATRIX,
                             CG_GL_MATRIX_IDENTITY);
```

Vertex and Fragment Programs

With the OpenGL state manager, vertex and fragment programs are defined via assignments to the **VertexProgram** and **FragmentProgram** states, respectively. Three different classes of expressions can be given on the right-hand side of these state assignments:

- ❑ Compile statements

- ❑ In-line assembly
- ❑ NULL

These three possibilities are demonstrated in the effect file below:

```
float4 main(uniform float foo, float4 uv : TEXCOORD0) : COLOR
{
    return (foo > 0) ? uv : 2 * uv;
}

technique SimpleFrag {
    pass {
        VertexProgram = NULL;
        FragmentProgram = compile arbfp1 main(-2.f);
    }
}

technique AsmFrag {
    pass {
        FragmentProgram = asm {
            !!FP1.0
            TEX o[COLR], {0}.x, TEX6, 2D;
            END
        };
    }
}
```

The most common of these three options for specifying programs is using `compile` statements. The first argument following the **compile** keyword is the name of the profile to which the program is to be compiled (for example, **fp30**, **fp40**, **arbfp1**, or **vp20**). The next argument gives the name of the function in the effect file that serves as the program entry point, followed by a list of expressions (for example, **-2.f**). These expressions have a one-to-one correspondence with the uniform parameters of the program being compiled—there must be exactly one for each uniform program parameter, no more, and no less.

In the example above, the expression “**-2.f**” sets the value for the **foo** parameter to **main()**. Because it is a literal value, CgFX is able to compile the program to a particularly efficient version that just includes returning the **uv** value.

It is also possible to include references to effect parameters in the expression used in the `compile` statement; for example:

```
float4 main(uniform float foo, float4 uv : TEXCOORD0) : COLOR
{
    return (foo > 0) ? uv : 2 * uv;
}
```

```

}

float bar;

technique NewSimpleFrag {
    pass {
        VertexProgram = NULL;
        FragmentProgram = compile arbfp1 main(2 * bar);
    }
}

```

Here, the value “**2 * bar**” is associated with the **foo** parameter of **main()**. When the value of **bar** is changed by the application, the value of **foo** in **main()** is set appropriately.

The second class of program state assignment types is assembly code. In-line assembly is indicated using the **asm** keyword, with the assembly language code between braces, as in the example above. CgFX depends on having the appropriate header at the start of the assembly — **!!FP1.0** for **fp30**, **!!ARBvp1.0** for **arbvp1**, and so on—to determine the profile for which the code is given.

Finally, vertex or fragment programs may be assigned the value **NULL** in the state assignment. This signifies that no such program should be used in this pass.

Textures and Samplers

CgFX also makes it possible to define state related to textures in the effect file. The effect file below shows an example. The full set of supported OpenGL texture state is listed in “[OpenGL State](#)” on page 129.

```

sampler2D samp = sampler_state {
    generateMipMap = true;
    minFilter = LinearMipMapLinear;
    magFilter = Linear;
};

float4 texsimple(uniform sampler2D sampler,
               float2 uv : TEXCOORD0) : COLOR {
    return tex2D(sampler, uv);
}

technique TextureSimple {
    pass {
        FragmentProgram = compile arbfp1 texsimple(samp);
    }
}

```

```

    }
}

```

Given this effect file, the application must take an extra step or two when setting up the texture in OpenGL. First, the application must indicate which texture handle should be used for the **sampler2D** in the effect file. Secondly, the application must use the Cg runtime to set the texture state given in the **sampler_state** block at the appropriate time.

Under OpenGL, the easiest way to achieve these goals is to call **cgGLSetupSampler(param, textureID)**. This entry point binds the given texture, associates the texture handle with the given parameter, and initializes the sampler state by calling **cgSetSamplerState()**.

Alternately, an application can perform these steps itself. The code below shows this in practice:

```

CGparameter p = cgGetNamedEffectParameter(effect, "samp");

GLuint handle;
glGenTextures(1, &handle);
glBindTexture(GL_TEXTURE_2D, handle);

cgGLSetTextureParameter(p, handle);
cgSetSamplerState(p);
...
glTexImage2D(GL_TEXTURE_2D, 0, GL_RGBA, RES, RES, 0, GL_RGBA,
             GL_FLOAT, data);

```

Note the calls to **cgGLSetTextureParameter()** and **cgSetSamplerState()**. The first call is the usual runtime call that needs to be made to tell the runtime which OpenGL texture object is associated with a given parameter.

The **cgSetSamplerState()** call ends up making the **glTexParameter** calls that set up the texture state defined in the **sampler_state** block. It expects that the appropriate texture object has been bound with **glBindTexture** first.

After the sampler has been initialized in either of these manners, there are two possibilities for how the texture parameters are managed. By far the easiest method is to enable texture management in the context:

```

cgGLSetManageTextureParameters(context, CG_TRUE);

```

If this is done, then when the **CGprogram** is bound by a call to **cgSetPassState()**, the texture parameters used are associated with the appropriate hardware texture units automatically.

Alternatively, the mapping of texture parameters to hardware units can be handled explicitly by the application, using the routine

cgGLEnableTextureParameter():

```
CGparameter progParam = cgGetNamedParameter(prog, "sampler");
cgGLEnableTextureParameter(progParam);
```

However, note that it is *not* possible to call **cgGLEnableTextureParameter()** with a handle to an effect's sampler parameter; the handle must be to an actual *program* parameter.

In general, the first approach is to be preferred for its simplicity.

Interfaces and Unsized Arrays

CgFX also supports Cg's interfaces and unsized arrays features. Given an effect file with Cg programs that use these features, the compile statement can be used in two different ways to resolve the interfaces and unsized arrays so that the program can be compiled. The abstract types may be resolved using Cg code itself, or they may be resolved using the Cg runtime.

Consider the following example: a **Light** interface has been defined with **SpotLight** implementing the interface. The **main()** program takes an unsized array of **Light** interface objects, loops over them, and returns the sum of the values returned by their respective **value()** methods.

```
interface Light {
    float4 value();
};

struct SpotLight : Light {
    float4 value() { return float4(1,2,3,4); }
};

float4 main(uniform Light l[]) : COLOR {
    float4 v = float4(0,0,0,0);
    for (int i = 0; i < l.length; ++i)
        v += l[i].value();
    return v;
}
```

Recall that all uniform parameters to the program must have expressions in the parenthesized list in the compile statement, and therefore one expression is necessary here for the **l** parameter.

Resolution using Cg

The first way that **main()** can be compiled is to provide the name of an effect parameter that resolves both the actual size of the array as well as the concrete type that implements the **Light** interface:

```
SpotLight spots[4];

technique {
    pass {
        FragmentProgram = compile arbfp1 main(spots);
    }
}
```

Resolution using the Cg runtime

Alternatively, the application can leave the resolution of the concrete types and array size until later so that they may be set via Cg runtime calls from the application, as one typically does for Cg programs that are not CgFX.

For this case, the expression passed to the compile statement should just be an unsized array of the abstract interface type:

```
Light lights[];

technique {
    pass {
        FragmentProgram = compile arbfp1 main(lights);
    }
}
```

The application must then create a shared array of concrete **light** instances. To do so, the application proceeds as it would when operating on a **CGprogram**—by retrieving the **CGtype** corresponding to each type of concrete instance to be created, and calling **cgCreateParameter()** or **cgCreateParameterArray()** to create the shared parameter of the given type. Lastly, the shared parameter is connected to the **effect** parameter. This process is illustrated below:

```
CGtype spotType = cgGetNamedUserType(effect, "SpotLight");
CGparameter spots = cgCreateParameterArray(context,
                                           spotType, 4);
CGparameter lights = cgGetNamedEffectParameter(effect,
                                                "lights");
cgConnectParameter(spots, lights);
```

Note that **cgGetNamedUserType()** in this case is passed a **Cgeffect** handle, rather than a **CGprogram** handle.

Later, when the associated technique is validated, any programs that make use of the abstract effect parameters are compiled.

Note that abstract parameters may *not* be used on the right-hand side of any state assignments other than **compile** state assignments. Doing so results in an error at effect creation time.

Evaluating Cg Programs using the Virtual Machine

There are many situations where it is useful to execute Cg programs on the CPU using the Cg runtime Virtual Machine (VM). Although running Cg programs on the CPU doesn't offer the same performance as execution on the GPU, it is sometimes useful, as in tabularizing complex functions into texture maps.

Programs that are to run on the VM are declared as follows:

```
float foo = 4.f;
float4 func(float2 p : POSITION, float2 delta : PSIZE) : COLOR
{
    return foo * p.xyxy;
}
```

The **POSITION** semantic denotes the parameter or parameters that are initialized with the coordinates of each point at which the function is evaluated. The value passed varies from zero to one in each of the dimensions over which the function is being evaluated. The **PSIZE** semantic denotes the parameter that is initialized with the spacing between samples at which the function is being evaluated. Lastly, the **COLOR** semantic denotes which parameter (or function return value) holds the computed value. Thus, the function above could have been written as a **void** function but with an **out float4 ret : COLOR** parameter and an assignment to **ret**, instead of using a **return** statement.

Given an effect file with such a program, a **Cgprogram** handle to it can be retrieved by creating a program using the **CG_PROFILE_GENERIC** profile:

```
Cgprogram tp = cgCreateProgramFromEffect(effect,
                                         CG_PROFILE_GENERIC,
                                         "func", NULL);
```

Given such a program handle, **cgEvaluateProgram** evaluates the program over the same one-, two-, or three-dimensional domain:

```
cgEvaluateProgram(Cgprogram prog, float *obuf, int ncomp,
                 int nx, int ny, int nz);
```

Where **prog** is the **Cgprogram** handle retrieved using **cgCreateProgramFromEffect()**, **obuf** is the buffer to which output values

are to be written, ***ncomp*** is the number of components per pixel in the output buffer (1, 2, 3, or 4), and ***nx***, ***ny***, and ***nz*** indicate the number of positions at which the function should be evaluated in each of the x, y, and z dimensions.

The total size of the buffer should be equal to the product of the number of positions in each of the dimensions and the number of components in the buffer, as in the example below:

```
#define RES 256
#define NCOMPS 4
float *buf = new float[NCOMPS*RES*RES];
cgEvaluateProgram(tp, buf, NCOMPS, RES, RES, 1);
// do something with buf
delete[] buf;
```

It is a error to pass a **CGprogram** that doesn't have the **CG_PROFILE_GENERIC** profile to **cgEvalauteProgram()**.

Annotations

Using annotations, it is possible to attach additional information to parameters, techniques, programs, and passes in the effect file for use by the application. An annotation is a list of variables and values denoted by angle brackets immediately following a declaration, as in the effect below:

```
float3 LightDir < string UItype = "direction"; >;

technique fancyHalo <
    bool optional = true;
> {
    pass < string geometry = "character";
        string destination = "texture"; > {
        ...
    }
}
```

CgFX does not interpret the meaning of annotations in any way; annotations exist solely for the convenience of the application. The example above shows a few common uses for annotations: the annotation of **LightDir** indicates what sort of user interface widget would be appropriate to provide the user for setting that parameter. The technique's annotation might indicate that applying the technique was optional when rendering the scene. In the example above, the pass annotations indicates to the application which part of the scene geometry to draw when rendering that pass, as well as where to store the image from rendering the pass.

Given a handle to a technique, pass, or parameter, there are API entry points for iterating through the annotations in turn:

```
CGannotation cgGetFirstTechniqueAnnotation(CGtechnique);
CGannotation cgGetFirstPassAnnotation(CGpass);
CGannotation cgGetFirstParameterAnnotation(CGparameter);
CGannotation cgGetFirstProgramAnnotation(CGprogram);
CGannotation cgGetNextAnnotation(CGannotation);
```

In addition, there are entry points for retrieving annotations by name:

```
CGannotation cgGetNamedTechniqueAnnotation(CGtechnique,
                                             const char *);
CGannotation cgGetNamedPassAnnotation(CGpass, const char *);
CGannotation cgGetNamedParameterAnnotation(CGparameter,
                                             const char *);
CGannotation cgGetNamedProgramAnnotation(CGprogram,
                                           const char *);
```

Given an annotation handle, its values may be retrieved through the use of one of the **cgGet*AnnotationValues()** entry points:

```
const float *cgGetFloatAnnotationValues(CGannotation,
                                          int *nvalues);
const int *cgGetIntAnnotationValues(CGannotation,
                                     int *nvalues);
const char *cgGetStringAnnotationValue(CGannotation);
const int *cgGetBooleanAnnotationValues(CGannotation,
                                         int *nvalues);
```

OpenGL State

When **cgGLRegisterStates()** is called, the CgFX OpenGL runtime initializes state assignments that correspond to almost all appropriate or useful OpenGL API calls. The set of states and state callbacks that are registered by this call compose the CgFX OpenGL state manager.

There is a one-to-one mapping between the state assignments that are provided by the OpenGL state manager and the corresponding OpenGL calls. Given an OpenGL call of interest, it is intended to be simple to determine which state assignment it corresponds to, and vice versa. For example, the state assignment **ClearColor = float4(0,1,0,1)** leads to the call **glClearColor(0,1,0,1)** when the state assignment is executed during a call to **cgSetPassState()**.

For calls that take enumerated values (for example, **GL_DEST_COLOR** for **glBlendFunc()**), corresponding enumerants are defined by the CgFX

OpenGL state manager, again with a straightforward mapping:

GL_DEST_COLOR corresponds to **DestColor**, and so forth. When an OpenGL call takes multiple parameters or multiple enumerants, a corresponding vector type is used; for example, a call to **glBlendFunc(GL_ZERO, GL_DST_ALPHA)** corresponds to the CgFX state assignment **BlendFunc = int2(Zero, DstAlpha)**.

When a state assignment depends on the presence of an OpenGL extension (for example, **BlendFuncSeparate** requires either **EXT_blend_func_separate** or the presence of OpenGL 1.4), it is possible to successfully load an effect file that uses that extension in one of its techniques, even if the OpenGL context doesn't support that extension. However, validation of any technique that uses such an unsupported extension in of its passes will fail.

The following table lists the names of the states supported by the CgFX OpenGL state manager, their types, and valid enumerants. The “Requires” column in the tables below indicates what OpenGL version or extension is required for each state assignment.

Table 6. CgFX OpenGL State Manager States

State Name	Type	Valid Enumerants	Requires
AlphaFunc	float2 (enum, reference_ value)	Never, Less, LEqual, Equal, Greater, NotEqual, GEqual, Always	OpenGL 1.0
BlendFunc	int2 (src_ factor, dst_factor)	Zero, One, DestColor, OneMinusDestColor, SrcAlpha, OneMinusSrcAlpha, DstAlpha, OneMinusDstAlpha, SrcAlphaSaturate, SrcColor, OneMinusSrcColor, ConstantColor, OneMinusConstantColor, ConstantAlpha, OneMinusConstantAlpha	1.0; 1.4 or NV_blend_square for SrcColor or OneMinusSrcColor for src_factor , and DstColor or OneMinusDstColor for dst_factor

Table 6. CgFX OpenGL State Manager States (continued)

State Name	Type	Valid Enumerants	Requires
BlendFuncSeparate	int4 (rgb_src, rgb_dst, a_src, a_dst)	Zero, One, DestColor, OneMinusDestColor, SrcAlpha, OneMinusSrcAlpha, DstAlpha, OneMinusDstAlpha, SrcAlphaSaturate, SrcColor, OneMinusSrcColor, ConstantColor, OneMinusConstantColor, ConstantAlpha, OneMinusConstantAlpha	OpenGL 1.4 or EXT_blend_func_separate; 1.4 or NV_blend_square for SrcColor or OneMinusSrcColor for rgb_src, and DestColor or OneMinusDstColor for rgb_dst
BlendEquation	int	FuncAdd, FuncSubtract, Min, Max, LogicOp	1.4 or ARB_imaging; or EXT_blend_subtract for FuncSubtract or FuncReverseSubtract; or EXT_blend_minmax for Min or Max; or EXT_blend_logic_op for LogicOp
BlendEquationSeparate	int2 (rgb, alpha)	FuncAdd, FuncSubtract, Min, Max, LogicOp	EXT_blend_equation_ separate; or 1.4, ARB_imaging, or EXT_blend_subtract for FuncSubtract or FuncReverseSubtract; or 1.4, ARB_imaging, or EXT_blend_minmax for Min or Max; or EXT_blend_logic_op for LogicOp
BlendColor	float4		1.4, ARB_imaging, or EXT_blend_color
ClearColor	float4		1.0
ClearStencil	int		1.0
ClearDepth	float		1.0

Table 6. CgFX OpenGL State Manager States (continued)

State Name	Type	Valid Enumerants	Requires
ClipPlane[<i>ndx</i>]	float4		OpenGL 1.0; <i>ndx</i> must be greater than or equal to zero and less than the value of GL_MAX_CLIP_PLANES
ColorMask	bool4		1.0
ColorMatrix	float4x4		ARB_imaging
ColorMaterial	int2	Front, Back, FrontAndBack, Emission, Ambient, Diffuse, Specular, AmbientAndDiffuse	1.0
CullFace	int	Front, Back, FrontAndBack	1.0
DepthBounds	float2		EXT_depth_bounds_test
DepthFunc	int	Never, Less, LEqual, Equal, Greater, NotEqual, GEqual, Always	1.0
DepthMask	bool		1.0
DepthRange	float2		1.0
FogMode	int	Linear, Exp, Exp2	1.0
FogDensity	float		1.0
FogStart	float		1.0
FogEnd	float		1.0
FogColor	float4		1.0
FragmentEnvParameter [<i>ndx</i>]	float4		ARB_fragment_program ; <i>ndx</i> must be greater than or equal to zero and less than the value of GL_MAX_PROGRAM_ENV_PARAMETERS_ARB for the GL_FRAGMENT_PROGRAM_ARB target to glGetProgramivARB

Table 6. CgFX OpenGL State Manager States (continued)

State Name	Type	Valid Enumerants	Requires
FragmentLocalParameter [<i>ndx</i>]	float4		ARB_fragment_program ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_PROGRAM_LOCAL_ PARAMETERS_ARB for the GL_FRAGMENT_PROGRAM_ARB target to glGetProgramivARB
FogCoordSrc	int	FragmentDepth, FogCoord	OpenGL 1.4 or EXT_fog_coord
FogDistanceMode	int	EyeRadial, EyePlane, EyePlaneAbsolute	NV_fog_distance
FragmentProgram	compile statement		ARB_fragment_program or NV_fragment_program
FrontFace	int	CW, CCW	1.0
LightModelAmbient	float4		1.0
LightAmbient [<i>ndx</i>]	float4		1.0; <i>ndx</i> must be greater or equal to 0 and less than the value of GL_MAX_LIGHTS
LightConstantAttenuation [<i>ndx</i>]	float		Same as LightAmbient
LightDiffuse [<i>ndx</i>]	float4		Same as LightAmbient
LightLinearAttenuation [<i>ndx</i>]	float		Same as LightAmbient
LightPosition [<i>ndx</i>]	float4		Same as LightAmbient
LightQuadraticAttenuation [<i>ndx</i>]	float		Same as LightAmbient
LightSpecular [<i>ndx</i>]	float4		Same as LightAmbient
LightSpotCutoff [<i>ndx</i>]	float		Same as LightAmbient
LightSpotDirection [<i>ndx</i>]	float3		Same as LightAmbient

Table 6. CgFX OpenGL State Manager States (continued)

State Name	Type	Valid Enumerants	Requires
LightSpotExponent [<i>ndx</i>]	float		Same as LightAmbient
LightModelColorControl	int	SingleColor , SeparateSpecular	OpenGL 1.2 or EXT_separate_specular_color
LineStipple	int2		1.0
LineWidth	float		1.0
LogicOp	int	Clear , And , AndReverse , Copy , AndInverted , Noop , Xor , Or , Nor , Equiv , Invert , OrReverse , CopyInverted , Nand , Set	1.0
MaterialAmbient	float4		1.0
MaterialDiffuse	float4		1.0
MaterialEmission	float4		1.0
MaterialShininess	float		1.0
MaterialSpecular	float4		1.0
ModelViewMatrix	float4x4		1.0
PointDistanceAttenuation	float3		1.4, ARB_point_parameters , or EXT_point_parameters
PointFadeThresholdSize	float		1.4, ARB_point_parameters , or EXT_point_parameters
PointSize	float		1.0
PointSizeMin	float		1.4, ARB_point_parameters , or EXT_point_parameters

Table 6. CgFX OpenGL State Manager States (continued)

State Name	Type	Valid Enumerants	Requires
PointSizeMax	float		OpenGL 1.4, ARB_point_parameters, or EXT_point_parameters
PointSpriteCoordOrigin	int	LowerLeft, UpperLeft	2.0
PointSpriteCoordReplace [ndx]	bool		2.0, ARB_point_sprite, or NV_point_sprite; <i>ndx</i> must be greater than or equal to zero and less than the value of GL_MAX_TEXTURE_COORDS
PointSpriteRMode	int	Zero, R, S	NV_point_sprite
PolygonMode	int2	Front, Back, FrontAndBack, Point, Line, Fill	1.0
PolygonOffset	float2		1.1
ProjectionMatrix	float4x4		1.0
Scissor	int4		1.0
ShadeModel	int	Flat, Smooth	1.0
StencilFunc	int3	Never, Less, LEqual, Equal, Greater, NotEqual, GEqual, Always	1.0
StencilMask	int		1.0
StencilOp	int3	Keep, Zero, Replace, Incr, Decr, Invert, IncrWrap, DecrWrap	1.0
StencilFuncSeparate	int4	Front, Back, FrontAndBack, Never, Less, LEqual, Equal, Greater, NotEqual, GEqual, Always	2.0 or EXT_stencil_two_side

StencilMaskSeparate	int2	Front, Back, FrontAndBack	OpenGL 2.0 or EXT_stencil_two_side
StencilOpSeparate	int4	Keep, Zero, Replace, Incr, Decr, Invert, IncrWrap, DecrWrap	2.0 or EXT_stencil_two_side
TexGenSMode[<i>ndx</i>]	int	ObjectLinear, EyeLinear, SphereMap, ReflectionMap, NormalMap	1.0; or 1.3, ARB_texture_cube_map , EXT_texture_cube_map , or NV_texgen_reflection for ReflectionMap , or NormalMap ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_COORDS
TexGenTMode[<i>ndx</i>]	int		Same as 0 7.98 408 442.0233 -j.7 Tw(

Table 6. CgFX OpenGL State Manager States (continued)

State Name	Type	Valid Enumerants	Requires
TexGenQEyePlane [<i>ndx</i>]	float4		Same as TexGenSEyePlane
TexGenSObjectPlane [<i>ndx</i>]	float4		Same as TexGenSEyePlane
TexGenTObjectPlane [<i>ndx</i>]	float4		Same as TexGenSEyePlane
TexGenRObjectPlane [<i>ndx</i>]	float4		Same as TexGenSEyePlane
TexGenQObjectPlane [<i>ndx</i>]	float4		Same as TexGenSEyePlane
Texture1D [<i>ndx</i>]	sampler1D		OpenGL 1.0; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_UNITS
Texture2D [<i>ndx</i>]	sampler2D		Same as Texture1D
Texture3D [<i>ndx</i>]	sampler3D		1.2 or EXT_texture3D ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_UNITS
TextureRectangle [<i>ndx</i>]	samplerRECT		ARB_texture_rectangle , EXT_texture_rectangle (Apple), or NV_texture_rectangle ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_UNITS

Table 6. CgFX OpenGL State Manager States (continued)

State Name	Type	Valid Enumerants	Requires
TextureCubeMap [<i>ndx</i>]	samplerCUBE		1.3, ARB_texture_cube_map , or EXT_texture_cube_map ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_ UNITS
TextureEnvColor [<i>ndx</i>]	float4		OpenGL 1.0; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_UNITS
TextureEnvMode [<i>ndx</i>]	int	Modulate , Decal , Blend , Replace , Add	1.0; 1.3, ARB_texture_env_add , or EXT_texture_env_add for Add; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_UNITS
VertexEnvParameter [<i>ndx</i>]	float4		ARB_vertex_program ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_PROGRAM_LOCAL_ PARAMETERS_ARB for the GL_VERTEX_PROGRAM_ARB target to glGetProgramivARB
VertexLocalParameter [<i>ndx</i>]	float4		ARB_vertex_program ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_PROGRAM_LOCAL_ PARAMETERS_ARB for the GL_VERTEX_PROGRAM_ARB target to glGetProgramivARB
VertexProgram	compile statement		ARB_vertex_program or NV_vertex_program

Similarly, there is a simple algorithm for determining the relationship between enumerants for **glEnable()** and for **glDisable()** and each of the states in the table below: for example, the state assignment **BlendEnable = false** corresponds to a call to **glDisable(GL_BLEND)**.

Table 7. Enable/Disable States

Enable/Disable State Name	Type	Requires
AlphaTestEnable	bool	OpenGL 1.0
AutoNormalEnable	bool	1.0
BlendEnable	bool	1.0
ClipPlaneEnable[<i>ndx</i>]	bool	1.0; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_CLIP_PLANES
ColorLogicOpEnable	bool	1.2
CullFaceEnable	bool	1.0
DepthBoundsEnable	bool	EXT_depth_bounds
DepthClampEnable	bool	NV_depth_clamp
DepthTestEnable	bool	1.0
DitherEnable	bool	1.0
FogEnable	bool	1.0
LightEnable[<i>ndx</i>]	bool	1.0; <i>ndx</i> must be greater or equal to 0 and less than the value of GL_MAX_LIGHTS
LightingEnable	bool	1.0
LightModelLocalViewerEnable	bool	1.0
LightModelTwoSideEnable	bool	1.0
LineSmoothEnable	bool	1.0
LineStippleEnable	bool	1.0
LogicOpEnable	bool	1.0
MultisampleEnable	bool	1.3 or ARB_multisample
NormalizeEnable	bool	1.0
PointSmoothEnable	bool	1.0

Table 7. Enable/Disable States (continued)

Enable/Disable State Name	Type	Requires
PointSpriteEnable	bool	2.0, ARB_point_sprite , or NV_point_sprite
PolygonOffsetFillEnable	bool	OpenGL 1.1
PolygonOffsetLineEnable	bool	1.1
PolygonOffsetPointEnable	bool	1.1
PolygonSmoothEnable	bool	1.0
PolygonStippleEnable	bool	1.0
RescaleNormalEnable	bool	1.2 or EXT_rescale_normal
SampleAlphaToCoverageEnable	bool	1.3 or ARB_multisample
SampleAlphaToOneEnable	bool	1.3 or ARB_multisample
SampleCoverageEnable	bool	1.3 or ARB_multisample
ScissorTestEnable	bool	1.0
StencilTestEnable	bool	1.0
TexGenSEnable [<i>ndx</i>]	bool	1.0; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_COORDS
TexGenTEnable [<i>ndx</i>]	bool	Same as TexGenSEnable
TexGenREnable [<i>ndx</i>]	bool	Same as TexGenSEnable
TexGenQEnable [<i>ndx</i>]	bool	Same as TexGenSEnable
Texture1DEnable [<i>ndx</i>]	bool	1.0; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_UNITS
Texture2DEnable [<i>ndx</i>]	bool	same as Texture1DEnable
Texture3DEnable [<i>ndx</i>]	bool	1.2 or EXT_texture3D ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_UNITS

Table 7. Enable/Disable States (continued)

Enable/Disable State Name	Type	Requires
TextureRectangleEnable [<i>ndx</i>]	bool	ARB_texture_rectangle , EXT_texture_rectangle (Apple), or NV_texture_rectangle ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_UNITS
TextureCubeMapEnable [<i>ndx</i>]	bool	OpenGL 1.3, ARB_texture_cube_map , or EXT_texture_cube_map ; <i>ndx</i> must be greater or equal to zero and less than the value of GL_MAX_TEXTURE_IMAGE_UNITS

OpenGL Sampler State

The following table lists the state assignments available in **sampler_state** blocks when using the CgFX OpenGL state manager. Any state values given are set when the **cgSetSamplerState()** routine is called with the **CGparameter** handle for a particular sample.

Note that some of these states are defined in OpenGL extensions—for example, **MirrorClampToBorder** is defined in the **EXT_texture_mirror_clamp** extension. Any state used that is based on an extension not supported by the current OpenGL context is ignored by the CgFX runtime.

Table 8. sampler_state State Assignments

Name	Type	Valid Values	Requires
WrapS , WrapT , WrapR	int	Repeat , Clamp , ClampToEdge , ClampToBorder , MirroredRepeat , MirrorClamp , MirrorClampToEdge , MirrorClampToBorder	OpenGL 1.2 or EXT_texture3D for WrapR ; 1.2 or EXT_texture_edge_clamp for ClampToEdge ; 1.3 or ARB_texture_border_clamp for ClampToBorder ; 1.4, ARB_texture_mirrored_repeat , or IBM_texture_mirrored_repeat for MirroredRepeat ; EXT_texture_mirror_clamp or ATI_texture_mirror_once for MirrorClamp or MirrorClampToEdge ; EXT_texture_mirror_clamp for MirrorClampToBorder

Table 8. sampler_state State Assignments (continued)

Name	Type	Valid Values	Requires
BorderColor	float4		OpenGL 1.0
CompareMode	int	None, CompareRToTexture	1.4 or ARB_shadow
CompareFunc	int	Never, Less, LEqual, Equal, Greater, NotEqual, GEqual, Always	1.4 or ARB_shadow ; 1.5 or EXT_shadow_funcs for Never, Less, Equal, Greater, NotEqual , or Always
DepthMode	int	Alpha, Intensity, Luminance	1.4 or ARB_depth_texture
GenerateMipMap	bool		1.4 or SGIS_generate_mipmap
LODBias	float		1.4
MinFilter	int	Nearest, Linear, LinearMipMapNearest, NearestMipMapNearest, NearestMipMapLinear, LinearMipMapLinear	1.0
MagFilter	int	Nearest, Linear	1.0
MaxMipLevel	float		1.2 or EXT_texture_lod
MaxAnisotropy	float		EXT_texture_filter_anisotropic
MinMipLevel	float		1.2 or EXT_texture_lod
Texture	texture	(Reference to texture parameter)	

OpenGL State Not Specifiable with State Assignments

By design, state assignments are limited to OpenGL state related to rendering geometric primitives. OpenGL state that is not assignable using the built-in OpenGL state manager includes the following:

- ☐ Pixel path state (such as pixel transfer and convolution state)
- ☐ Per-vertex attributes (such as **glColor** or **glNormal**)
- ☐ Client-side state such as vertex arrays and pixel store modes

- ❑ Vertex and pixel buffer object state
- ❑ Miscellaneous state for evaluators, feedback, selection, or occlusion queries
- ❑ Texture environment **GL_COMBINE** state
Although related to rendering, it is complex and redundant with fragment color operations better specified with Cg fragment programs.

Future enhancements may allow assignments for currently unassignable OpenGL state.

A Brief Tutorial

This section walks you through the sample Cg Microsoft Visual Studio workspace we have provided, along with a simple Cg program that you can use for experimentation.

Loading the Workspace

When you load the **Cg_Simple** file, your workspace should look like the image in Fig. 3.




Fig. 3. The **Cg_Simple** Workspace

As usual, click the **FileView** tab to view the various files in the project. What's different in this case, though, is that in addition to the usual Source Files and Header Files folders, there is also a Cg Programs folder.

This Cg Programs folder should contain one Cg program, **simple.cg**, which is what you can use for experimentation. Double-click **simple.cg** to open it for editing. While you are editing **simple.cg**, you can press **Control+F7** at any time to compile it. Because of the way the project is set up, any errors in your code will be shown just as when you compile a normal C or C++ program.

You can also double-click on an error, which takes you to the location in the source code that caused the error.

Understanding simple.cg

The **Cg_Simple** application runs the shader defined in **simple.cg** on a torus. The provided version of **simple.cg** calculates diffuse and specular lighting for each vertex. A screenshot of the shader is shown in [Fig. 4](#).




Fig. 4. The simple.cg Shader

Program Listing for simple.cg

The following is the program listing for **simple.cg**:

```
// Define inputs from application.
struct appin
{
    float4 Position      : POSITION;
    float4 Normal        : NORMAL;
};

// Define outputs from vertex shader.
struct vertout
{
    float4 HPosition     : POSITION;
    float4 Color         : COLOR;
};

vertout main(appin IN,
             uniform float4x4 ModelViewProj,
             uniform float4x4 ModelViewIT,
             uniform float4 LightVec)
{
    vertout OUT;

    // Transform vertex position into homogenous clip-space.
    OUT.HPosition = mul(ModelViewProj, IN.Position);

    // Transform normal from model-space to view-space.
    float3 normalVec = normalize(mul(ModelViewIT,
                                     IN.Normal).xyz);

    // Store normalized light vector.
    float3 lightVec = normalize(LightVec.xyz);

    // Calculate half angle vector.
    float3 eyeVec = float3(0.0, 0.0, 1.0);
    float3 halfVec = normalize(lightVec + eyeVec);

    // Calculate diffuse component.
    float diffuse = dot(normalVec, lightVec);

    // Calculate specular component.
    float specular = dot(normalVec, halfVec);

    // Use the lit function to compute lighting vector from
```

```
// diffuse and specular values.
float4 lighting = lit(diffuse, specular, 32);

// Blue diffuse material
float3 diffuseMaterial = float3(0.0, 0.0, 1.0);

// White specular material
float3 specularMaterial = float3(1.0, 1.0, 1.0);

// Combine diffuse and specular contributions and
// output final vertex color.
OUT.Color.rgb = lighting.y * diffuseMaterial +
                lighting.z * specularMaterial;
OUT.Color.a = 1.0;

return OUT;
}
```

Definitions for Structures with Varying Data

The first thing to notice is the definitions of structures with binding semantics for varying data.

Let's take a look at the **appin** structure:

```
// define inputs from application
struct appin
{
    float4 Position      : POSITION;
    float4 Normal        : NORMAL;
};
```

This structure contains only two members: **Position** and **Normal**. Because this data varies per-vertex, the binding semantics **POSITION** and **NORMAL** tell the compiler that the position information is associated with the predefined attribute **POSITION** and that the normal information is associated with the predefined attribute **NORMAL**.

The other structure that is defined in **simple.cg** is **vertout**, which connects the vertex to the fragment:

```
// define outputs from vertex shader
struct vertout
{
    float4 HPosition     : POSITION;
    float4 Color          : COLOR;
};
```

The **vertout** structure also contains only two members: **Hposition**, the vertex position in homogeneous coordinates, and **Color**, the vertex color. Again, binding semantics are used to specify register locations for the variables. In this case, the homogeneous position information resides in the hardware register corresponding to **POSITION** and that the color information resides in the hardware register corresponding to **COLOR**.

Passing Arguments

Now let's take a look at the body of the program, section by section, starting with the declaration of **main()**:

```
vertout main(appin IN,
             uniform float4x4 ModelViewProj,
             uniform float4x4 ModelViewIT,
             uniform float4 LightVec)
```

As required for a vertex program, **main()** takes an application-to-vertex structure as input and returns a vertex-to-fragment structure. In this case, we are using the two structure types we have already defined: **appin** and **vertout**. Notice that **main()** takes in three uniform parameters: two matrices and one vector. All three parameters are passed to **simple.cg** by the application, using the run-time library.

The first matrix, **ModelViewProj**, is the concatenation of the modelview and projection matrices. Together, these matrices transform points from model space to clip space. The second matrix, **ModelViewIT**, is the inverse transpose of the modelview matrix. The third parameter, **LightVec**, is a vector that specifies the location of the light source.

Basic Transformations

Now we start the body of the vertex program:

```
vertout OUT;

OUT.HPosition = mul(ModelViewProj, IN.Position);
```

A vertex program is responsible for calculating the homogenous clip-space position of the vertex (given the vertex's model-space coordinates). Therefore, the vertex's model-space position (given by **IN.Position**) needs to be transformed by the concatenation of the modelview and projection matrices (called **ModelViewProj** in this example). The transformed position is assigned directly to **OUT.HPosition**. Note that you are not responsible for

the perspective division when using vertex programs. The hardware automatically performs the division after executing the vertex program.

Since we want to do our lighting in eye space, we have to transform the model space normal **IN.Normal** to eye space:

```
// transform normal from model-space to view-space
float3 normalVec = normalize(mul(ModelViewIT,
                               IN.Normal).xyz);
```

Remember that when transforming normals, we need to multiply by the inverse transpose of the modelview matrix. Then we normalize the eye space normal vector and store it as **normalVec**.

Prepare for Lighting

The subsequent steps prepare for lighting:

```
// store normalized light vector
float3 lightVec = normalize(LightVec.xyz);

// calculate half angle vector
float3 eyeVec = float3(0.0, 0.0, 1.0);
float3 halfVec = normalize(lightVec + eyeVec);
```

At this point we have to ensure that all our vectors are normalized. We start by normalizing **LightVec**¹. Then, in preparation for specular lighting, we have to define the “half-angle” vector **halfVec**, which is the vector halfway between the light and the eye vectors (that is, **(lightVec+eyeVec)/2**). We normalize **halfVec**, so we don’t need to bother with the division by two, because it cancels out after normalization anyway. In this example, we assume that the eye is at **(0,0,1)**, but an application would typically pass the eye position also as a uniform parameter, since it would be unchanged from vertex to vertex. We use Cg’s inline vector construction capability to build a 3-component **float** vector that contains the eye position, and then we assign this value to **eyeVec**.

1. Because **LightVec** is uniform, it is more efficient to normalize it once in the application rather than on a per-vertex basis. It is done here for illustrative purposes.

Calculating the Vertex Color

Now we have to calculate the vertex color to output.

Calculating the Diffuse and Specular Lighting Contributions

In this example, we're going to calculate just a simple combination of diffuse and specular lighting:

```
// calculate diffuse component
float diffuse = dot(normalVec, lightVec);

// calculate specular component
float specular = dot(normalVec, halfVec);

// Use the lit function to compute lighting vector from
// diffuse and specular values
float4 lighting = lit(diffuse, specular, 32);
```

Here we use the Cg Standard Library to perform dot products (using **dot()**). We also make use of the Standard Library's **lit()** function to calculate a Blinn-style lighting vector based on the previously computed dot products. The returned vector holds the diffuse lighting contribution in the y-coordinate, and the specular lighting contribution in the z-coordinate.

Remember to take advantage of the Standard Library to help speed up your development cycle.

Modulating the Diffuse and Specular Lighting Contributions

Once the diffuse and specular lighting contributions **lighting.y** and **lighting.z** have been calculated, we need to modulate them with the object's material properties:

```
// blue diffuse material
float3 diffuseMaterial = float3(0.0, 0.0, 1.0);

// white specular material
float3 specularMaterial = float3(1.0, 1.0, 1.0);

// combine diffuse and specular contributions and
// output final vertex color
OUT.Color.rgb = lighting.y * diffuseMaterial +
                lighting.z * specularMaterial;
OUT.Color.a = 1.0;
return OUT;
```

We define the object's diffuse material color as blue. We modulate the lighting contributions with the material properties to get the final vertex color, and we assign it to the output structure's color field, **OUT.Color**. Finally, we set the alpha channel of the final color to 1.0, so that our object will be opaque, and return the computed position and color values stored in the **OUT** structure.

Further Experimentation

Use `simple.cg` as a framework to try more advanced experiments, perhaps by adding more parameters to the program or by performing more complex calculations in the vertex program. Have fun experimenting!



Advanced Profile Sample Shaders

This chapter provides a set of advanced profile sample shaders written in Cg. Each shader comes with an accompanying snapshot, description, and source code.

Examples shown are

- ❑ Improved Skinning
- ❑ Improved Water
- ❑ Melting Paint
- ❑ MultiPaint
- ❑ Ray-Traced Refraction
- ❑ Skin
- ❑ Thin Film Effect
- ❑ Car Paint 9

Improved Skinning

Description

This shader takes in a set of all the transformation matrices that can affect a particular bone. Each bone also sends in a list of matrices that affect it. There is then a simple loop that for each vertex goes through each bone that affects that vertex and transforms it. This allows just one Cg program to do the entire skinning for vertices affected by any number of bones. This is a simple

Vertex Shader Source Code for Improved Skinning

```

struct inputs
{
    float4 position      : POSITION;
    float4 weights       : BLENDWEIGHT;
    float4 normal        : NORMAL;
    float4 matrixIndices : TESSFACTOR;
    float4 numBones      : SPECULAR;
};

struct outputs
{
    float4 hPosition     : POSITION;
    float4 color         : COLOR0;
};

outputs main(inputs IN,
              uniform float4x4 modelViewProj,
              uniform float3x4 boneMatrices[30],
              uniform float4 color,
              uniform float4 lightPos)
{
    outputs OUT;

    float4 index = IN.matrixIndices;
    float4 weight = IN.weights;

    float4 position;
    float3 normal;

    for (float i = 0; i < IN.numBones.x; i += 1) {
        // transform the offset by bone i
        position = position + weight.x *
            float4(mul(boneMatrices[index.x], IN.position).xyz,
                  1.0);

        // transform normal by bone i
        normal = normal + weight.x *
            mul((float3x3)boneMatrices[index.x],
              IN.normal.xyz).xyz;

        // shift over the index/weight variables; this moves
        // the index and weight for the current bone into
        // the .x component of the index and weight variables
    }
}

```

```
        index = index.yzwx;  
        weight = weight.yzwx;  
    }  
  
    normal = normalize(normal);  
  
    OUT.hPosition = mul(modelViewProj, position);  
    OUT.color = dot(normal, lightPos.xyz) * color;  
  
    return OUT;  
}
```

Improved Water

Description

This demo gives the appearance that the viewer is surrounded by a large grid of vertices (because of the free rotation), but switching to wireframe or increasing the frustum angle makes it apparent that the vertices are a static mesh with the height, normal, and texture coordinates being calculated on-the-fly based on the direction and height of the viewer. This technique allows for very GPU-friendly water animations because the static mesh can be precomputed. The vertices are displaced using sine waves, and in this example a loop is used to sum five sine waves to achieve realistic effects.



Fig. 6. Example of Improved Water

Vertex Shader Source Code for Improved Water

```

struct app2vert
{
    float4 Position    : POSITION;
};

struct vert2frag
{
    float4 HPosition   : POSITION;
    float4 TexCoord0   : TEXCOORD0;
    float4 TexCoord1   : TEXCOORD1;
    float4 Color0      : COLOR0;
    float4 Color1      : COLOR1;
};

void calcWave(out float disp, out float2 normal,
              float dampening, float3 viewPosition,
              float waveTime, float height,
              float frequency, float2 waveDirection)
{
    float distance1 = dot(viewPosition.xy, waveDirection);
    distance1 = frequency * distance1 + waveTime;

    disp = height * sin(distance1) / dampening;
    normal = -cos(distance1) * height * frequency *
              (waveDirection.xy) / (.4*dampening);
}

vert2frag main(
    app2vert IN,
    uniform float4x4 ModelViewProj,
    uniform float4x4 ModelView,
    uniform float4x4 ModelViewIT,
    uniform float4x4 TextureMat,
    uniform float   Time,
    uniform float4   Wave1,
    uniform float4   Wave1Origin,
    uniform float4   Wave2,
    uniform float4   Wave2Origin,
    const uniform float4   WaveData[5])
{
    vert2frag OUT;

```

```

float4 position = float4(IN.Position.x, 0,
                        IN.Position.y,1);
float4 normal = float4(0,1,0,0);
float dampening = 1 + dot(position.xyz, position.xyz)/1000;
float i, disp;
float2 norm;

for (i = 0; i < 5; i = i + 1)
{
    float waveTime = Time.x * WaveData[i].z;
    float frequency = WaveData[i].z;
    float height = WaveData[i].w;
    float2 waveDir = WaveData[i].xy;

    calcWave(disp, norm, dampening, IN.Position.xyz,
            waveTime, height, frequency, waveDir);
    position.y = position.y + disp;
    normal.xz = normal.xz + norm;
}

OUT.HPosition = mul(ModelViewProj, position);

// transfrom normal into eye-space
normal = mul(ModelViewIT, normal);
normal.xyz = normalize(normal.xyz);

// get a vector from the vertex to the eye
float3 eyeToVert = mul(ModelView, position).xyz;
eyeToVert = normalize(eyeToVert);

// calculate the reflected vector for cubemap look-up
float4 reflected = mul(TextureMat,
                        reflect(eyeToVert, normal.xyz).xyzz);

// output two reflection vectors for the two
// environment cubemaps
OUT.TexCoord0 = reflected;
OUT.TexCoord1 = reflected;

// Calculate a fresnel term (note that f0 = 0)
float fres = 1+dot(eyeToVert,normal.xyz);
fres = pow(fres, 5);

// set the two color coefficients (the magic constants
// are arbitrary), these two color coefficients are used

```

```

// to calculate the contribution from each of the two
// environment cubemaps (one bright, one dark)
OUT.Color0 = (fres*1.4 + min(reflected.y,0)).xxxx +
    float4(.2, .3, .3, 0);
OUT.Color1 = (fres*1.26).xxxx;

return OUT;
}

```

Pixel Shader Source Code for Improved Water

```

float4 main(in float3 color0      : COLOR0,
            in float3 color1      : COLOR1,
            in float3 reflectVec   : TEXCOORD0,
            in float3 reflectVecDark : TEXCOORD1,
            uniform samplerCUBE environmentMaps[2]
            ) : COLOR
{
    float3 reflectColor = texCUBE(environmentMaps[0],
        reflectVec).rgb;
    float3 reflectColorDark = texCUBE(environmentMaps[1],
        reflectVecDark).rgb;

    float3 color = (reflectColor * color0) +
        (reflectColorDark * color1);
    return float4(color, 1.0);
}

```

Melting Paint

Description

This shader uses an environment map with procedurally modified texture lookups to create a melting effect on the surface texture (the NVIDIA logo in this example). The reflection vector is shifted using a noise function, giving the appearance of a bumpy surface. The surface texture's texture coordinates are shifted in a time-dependent manner, also based on a noise texture.



Fig. 7. Example of Melting Paint

Vertex Shader Source Code for Melting Paint

```
// define inputs from application
struct app2vert
{
    float4 Position      : POSITION;
    float4 Normal        : NORMAL;
```

```

    float4 Color0      : COLOR0;
    float4 TexCoord0   : TEXCOORD0;
};

struct vert2frag
{
    float4 HPosition    : POSITION;
    float3 OPosition    : TEXCOORD2;
    float3 EPosition    : TEXCOORD3;
    float3 Normal       : TEXCOORD1;
    float3 TexCoord0    : TEXCOORD0;
    float4 Color0       : COLOR0;

    float3 LightPos     : TEXCOORD4;
    float3 ViewerPos    : TEXCOORD5;
};

vert2frag main(app2vert In,
               uniform float4x4 ModelViewProj,
               uniform float4x4 ModelView,
               uniform float4x4 ModelViewI,
               uniform float4 ViewerPos,
               uniform float4 LightPos)
{
    vert2frag Out;

    // Vertex positions:
    // In clip space
    Out.HPosition = mul(ModelViewProj, In.Position);
    // In object space
    Out.OPosition = In.Position.xyz;
    // In eye space
    Out.EPosition = mul(ModelView, In.Position).xyz;

    Out.Normal = normalize(In.Normal.xyz);
    // Copy the texture coordinates
    Out.TexCoord0 = In.TexCoord0.xyz;
    // Generate a white color
    Out.Color0 = LightPos;

    Out.LightPos = mul(ModelViewI, LightPos).xyz;
    Out.ViewerPos = mul(ModelViewI, float4(0,0,0,1)).xyz;

    return Out;
}

```

Pixel Shader Source Code for Melting Paint

```

struct vert2frag
{
    float4 HPosition : POSITION;
    float3 OPosition : TEXCOORD2;
    float3 EPosition : TEXCOORD3;
    float3 Normal     : TEXCOORD1;
    float3 TexCoord0  : TEXCOORD0;
    float4 Color0     : COLOR0;

    float3 LightPos   : TEXCOORD4;
    float3 ViewerPos  : TEXCOORD5;
};

void calcLighting(out float diffuse, out float specular,
    float3 normal, float3 fragPos, float3 lightPos,
    float3 eyePos, float specularExp)
{
    float3 light = lightPos - fragPos;
    float len = length(light);
    light = light / len;

    float3 eye = normalize(eyePos - fragPos);
    float3 halfVec = normalize(eyePos + light);

    float attenuation = 1. / (.3 * len);

    float4 lighting = lit(dot(light, normal),
        dot(halfVec, normal), specularExp);
    diffuse = lighting.y * attenuation;
    specular = lighting.z * attenuation;
}

float4 main(vert2frag IN,
    uniform float4  LightPos,
    uniform sampler3D noise_map,
    uniform sampler2D nv_map,
    uniform samplerCUBE cube_map,
    uniform float4  interpolate
    ) : COLOR
{
    float diffuse, specular;

    float3 biVariate = float3(IN.OPosition.x-IN.OPosition.z,

```

```

    IN.OPosition.y+IN.OPosition.z, 0);
float3 uniVariate = float3(IN.OPosition.x+IN.OPosition.z,
    0, 0);

float3 normal = normalize(IN.Normal);
float3 noiseTex = float3((IN.OPosition.x+IN.OPosition.z)*6,
    IN.OPosition.y/2, 0);
float3 noiseSum = tex3D(noise_map, biVariate/3).rgb/12 +
    tex3D(noise_map, noiseTex).rgb/18 +
    tex3D(noise_map, biVariate*6).rgb/18;
normal = normalize(normal + noiseSum);

calcLighting(diffuse, specular, normal, IN.OPosition,
    IN.LightPos, IN.ViewerPos, 32);

float3 nvShift = tex3D(noise_map, uniVariate/3).rgb / 2 +
    tex3D(noise_map, uniVariate).rgb / 4 +
    tex3D(noise_map, biVariate*3).rgb / 16;
nvShift.x = nvShift.x*nvShift.x * interpolate.x * 3;
nvShift.y = 0;

biVariate = float3(IN.OPosition.x - IN.OPosition.z,
    IN.OPosition.y, 0);
float2 texCoord = biVariate.xy/4 + float2(1.1, .5) +
    nvShift.yx + float2(0, interpolate.x/8);
float3 nvDecal =
    tex2D(nv_map, float2(1-texCoord.x, texCoord.y)).rgb *
    (1-interpolate.x * .7).xxx;

float3 eye = IN.ViewerPos - IN.OPosition;
float3 lightMetal = texCUBE(cube_map,
    reflect(normal, eye)).rgb;
float3 darkMetal = (diffuse * float3(.5, .25, 0) +
    specular * float3(.7, .4, 0));

float3 finalColor = lerp(lightMetal, darkMetal, nvDecal.x);
return float4(finalColor, 1);
}

```

MultiPaint

Description

MultiPaint presents a single-pass solution to a common production problem: mixing multiple kinds of materials on a single polygonal surface. MultiPaint provides a simple BRDF (bidirectional reflectance distribution function) that is still complex enough to represent many common metallic and dielectric surfaces, and controls all key factors of the variable BRDF through texturing. This permits you to create multiple materials without switching shaders, splitting your model, or resorting to multiple passes.

Uses for MultiPaint might include complex armor built of inlaid metals, woods, and stones—all modeled on a single, simple poly mesh; buildings composed of multiple types of stone, glass, and metal, expressed as simple cubes; cloth with inlaid metallic threads; or as in this demo, metal partially covered with peeling paint.

Using multiple BRDFs is common in the offline world, but rarely optimized; instead, two different shaders may be evaluated and their results blended using a mask texture or chained through **if** statements. For maximum real-time performance, MultiPaint instead integrates all of the key parts of the BRDFs as multiple painted textures so that only one pass through the shader is required to create the mixed appearance. This permits a single-pass shader containing diffuse, specular, and environmental lighting effects in a compact, fast-executing package.



Fig. 8. Example of MultiPaint

Vertex Shader Source Code for MultiPaint

```

// define inputs from vertex buffer
struct appin
{
    float4 Position      : POSITION;
    float4 UV            : TEXCOORD0;
    float4 Tangent       : TEXCOORD1;
    float4 Binormal      : TEXCOORD2;
    float4 Normal        : TEXCOORD3;
};

// output -- same struct is the input to "cg_multipaint.cg"
struct MultiPaintV2F {
    float4 HPosition     : POSITION; // position (clip space)
    float4 TexCoords     : TEXCOORD0; // base ST coordinates
    float3 OPosition     : TEXCOORD1; // position (obj space)
    float3 Normal        : TEXCOORD2; // normal (eye space)
    float3 VPosition     : TEXCOORD3; // view pos (obj space)
    float3 T             : TEXCOORD4; // tangent (obj space)
    float3 B             : TEXCOORD5; // binormal (obj space)
    float3 N             : TEXCOORD6; // normal (obj space)
    float4 LightVec0     : TEXCOORD7; // light dir (obj space)
};

MultiPaintV2F main(appin IN,
                    uniform float4x4 ModelViewProj,
                    uniform float4x4 ModelViewIT,
                    uniform float4x4 ModelViewI,
                    uniform float4 TexRepeats,
                    uniform float4 LightVec) // (eye space)
{
    MultiPaintV2F OUT;

    OUT.HPosition = mul(ModelViewProj, IN.Position);

    // pass through object-space position
    OUT.OPosition = IN.Position.xyz;

    // transform normal to eye space
    OUT.Normal = normalize(mul(ModelViewIT, IN.Normal).xyz);

    OUT.TexCoords = IN.UV * TexRepeats;

    // pass through object-space normal, tangent, binormal.

```

```

    OUT.N = normalize(IN.Normal.xyz);
    OUT.T = IN.Tangent.xyz;
    OUT.B = IN.Binormal.xyz;

    // transform view pos (origin) to obj space
    OUT.VPosition = mul(ModelViewI, float4(0,0,0,1)).xyz;

    // transform light vector to obj space
    OUT.LightVec0 = mul(ModelViewI, LightVec);

    return OUT;
}

```

Pixel Shader Source Code for MultiPaint

```

#define WHITE half4(1.0h,1.0h,1.0h,1.0h)

// input -- same struct is output from "cg_multipaintVP.cg"
struct MultiPaintV2F {
    float4 HPosition      : POSITION;    // position (clip space)
    float4 TexCoords      : TEXCOORD0; // base ST coordinates
    float3 OPosition      : TEXCOORD1; // position (obj space)
    float3 Normal         : TEXCOORD2; // normal (eye space)
    float3 VPosition      : TEXCOORD3; // view pos (obj space)
    float3 T              : TEXCOORD4; // tangent (obj space)
    float3 B              : TEXCOORD5; // binormal (obj space)
    float3 N              : TEXCOORD6; // normal (obj space)
    float4 LightVec0      : TEXCOORD7; // light dir (obj space)
};

// channels in our material map:
#define SPEC_STR x
#define METALNESS y
#define NORM_SPEC_EXPON z

// subfields in "SpecData"
#define MINPOWER x
#define MAXPOWER y
#define MAXSPEC z

// subfields in "Refldata"
#define FRESNEL_MIN x
#define FRESNEL_MAX y
#define FRESNEL_EXPON z
#define REFL_STRENGTH w

```

```
// subfields in "BumpData"
#define BUMP_SCALE x

half4 main(MultiPaintV2F IN,
    uniform sampler2D ColorMap,    // color
    uniform sampler2D MaterialMap, // see above
    uniform sampler2D NormalMap,   // tangent-space normals
    uniform samplerCUBE EnvMap,    // environment skybox
    uniform float4 SpecData,       // see above
    uniform float4 ReflData,       // see above
    uniform float4 BumpData        // see above
) : COLOR
{
```

```
half4 paintShine = fakeFresnel * reflColor;  
half4 metalShine = surfCol * reflColor;  
half4 shineCol = ReflData.REFL_STRENGTH *  
                lerp(paintShine, metalShine,  
                    material.METALNESS);  
  
half4 finalColor = specResult + diffResult + shineCol;  
finalColor.w = 1.0h;  
  
return finalColor;  
}
```

Ray-Traced Refraction

Description

This shader presents a method for adding high-quality details to small objects using a single-bounce, ray-traced pass. In this example, the polygonal surface is sampled and a refraction vector is calculated. This vector is then intersected with a plane that is defined as being perpendicular to the object's x-axis. The intersection point is calculated and used as texture indices for a painted iris.

The demo permits varying the index of refraction, the depth and density of the lens. Note that the choice of geometry is arbitrary — this sample is a sphere, but any polygonal model can be used.




Fig. 9. Example of Ray-Traced Refraction

Vertex Shader Source Code for Ray-Traced Refraction

```

struct appin
{
    float4 Position    : POSITION;
    float4 Normal      : NORMAL;
};

// output -- same struct is the input to fragment shader
struct EyeV2F {
    float4 HPosition   : POSITION; // clip space pos
    float3 OPosition   : TEXCOORD0; // Obj-coords location
    float3 VPosition   : TEXCOORD1; // eye pos (obj space)
    float3 N           : TEXCOORD2; // normal (obj space)
    float4 LightVec0    : TEXCOORD3; // light dir (obj sp)
};

EyeV2F main(appin IN,
            uniform float4x4 ModelViewProj,
            uniform float4x4 ModelViewI,
            uniform float4 LightVec) // in EYE coords
{
    EyeV2F OUT;

    // calculate clip space position for rasterizer use
    OUT.HPosition = mul(ModelViewProj, IN.Position);

    // pass through object space position
    OUT.OPosition = IN.Position.xyz;

    // object-space normal
    OUT.N = normalize(IN.Normal.xyz);

    // transform view pos and light vec to obj space
    OUT.VPosition = mul(ModelViewI, float4(0,0,0,1)).xyz;
    OUT.LightVec0 = normalize(mul(ModelViewI, LightVec));

    return OUT;
}

```

Pixel Shader Source Code for Ray-Traced Refraction

```
// Assume ray direction is normalized.
// Vector "planeEq" is encoded half3(A,B,C,D) where
// (Ax+By+Cz+D)=0 and half3(A,B,C) has been normalized.
// Returns distance along to to intersection; distance is
// negative if no intersection.
half intersect_plane(half3 rayOrigin,half3 rayDir,
                    half4 planeEq) {
    half3 planeN = planeEq.xyz;
    half denominator = dot(planeN, rayDir);
    half result = -1.0h;

    // d==0 -> parallel || d>0 -> faces away
    if (denominator < 0.0h) {
        half top = dot(planeN,rayOrigin) + planeEq.w;
        result = -top/denominator;
    }
    return result;
}

// subfields in "BallData"
#define RADIUS x
#define IRIS_DEPTH y
#define ETA z
#define LENS_DENSITY w

// subfields in "SpecData"
#define PHONG x
#define GLOSS1 y
#define GLOSS2 z
#define DROP w

struct EyeV2F {
    float4 HPosition : POSITION;
    float3 OPosition : TEXCOORD0;
    float3 VPosition : TEXCOORD1;
    float3 N : TEXCOORD2;
    float4 LightVec0 : TEXCOORD3;
};

half4 main(EyeV2F IN,
    uniform sampler2D ColorMap, // color
    // components: {radius,irisDepth,eta,lensDensity)
    uniform float4 BallData,
```

```

// components: {phongExp,gloss1,gloss2,drop)
uniform float4 GlossData,
uniform float3 AmbiColor,
uniform float3 DiffColor,
uniform float3 SpecColor,
uniform float3 LensColor,
uniform float3 BgColor) : COLOR
{
    const half3 baseTex = half3(1.0h,1.0h,1.0h);
    const half GRADE = 0.05h;
    const half3 yAxis = half3(0.0h,1.0h,0.0h);
    const half3 xAxis = half3(1.0h,0.0h,0.0h);
    const half3 ballCtr = half3(0.0h,0.0h,0.0h);

    // (actually constants - could be done in VP or on CPU)
    half irisSize = BallData.RADIUS *
        sqrt(1.0h-BallData.IRIS_DEPTH * BallData.IRIS_DEPTH);
    half irisScale = 0.3333h / max(0.01h, irisSize);
    half irisDist = BallData.RADIUS * BallData.IRIS_DEPTH;
    half3 pupilCenter = ballCtr + half3(irisDist,0.0h,0.0h);
    // if x axis, returns simple -irisDist
    half D = -dot(pupilCenter, xAxis);
    half slice = IN.OPosition.x - irisDist;
    half4 planeEquation = half4(xAxis, D);

    // view vector TO surface
    half3 Vn = normalize(IN.OPosition - IN.VPosition);
    half3 Nf = normalize(IN.N);
    half3 Ln = IN.LightVec0.xyz;
    half3 DiffLight = DiffColor * saturate(dot(Nf, -Ln));
    half3 missColor = AmbiColor + baseTex * DiffLight;
    half3 DiffPupil = AmbiColor + saturate(dot(xAxis, -Ln));

    half3 halfAng = normalize(-Ln - Vn);
    half ndh = abs(dot(Nf,halfAng));
    half spec1 = pow(ndh, GlossData.PHONG);
    half s2 = smoothstep(GlossData.GLOSS1, GlossData.GLOSS2,
        spec1);
    spec1 = lerp(GlossData.DROP, spec1, s2);
    half3 SpecularLight = SpecColor * spec1;

    half3 hitColor = missColor;

    if (slice >= 0.0h) {
        half gradedEta = BallData.ETA;
    }
}

```

```

gradedEta = 1.0h/gradedEta;
half3 faceColor = BgColor;

half3 refVector = refract(Vn, Nf, gradedEta);
if (dot(refVector, refVector) > 0) {
    // now let's intersect with the iris plane
    half irist = intersect_plane(IN.0Position, refVector,
        planeEquation);
    half fadeT = irist * BallData.LENS_DENSITY;
    fadeT = fadeT * fadeT;
    faceColor = DiffPupil.xxx;
    if (irist > 0) {
        half3 irisPoint = IN.0Position + irist*refVector;
        half3 irisST = (irisScale*irisPoint) +
            half3(0.0h, 0.5h, 0.5h);
        faceColor = tex2D(ColorMap, irisST.yz).rgb;
    }
    faceColor = lerp(faceColor, LensColor, fadeT);
    hitColor = lerp(missColor, faceColor,
        smoothstep(0.0h, GRADE, slice));
}

hitColor = hitColor + SpecularLight;
return half4(hitColor, 1.0h);
}

```

Skin

Description

This effect demonstrates some techniques for rendering skin ranging from simple Blinn-Phong Bump-Mapping to more complex Subsurface Scattering lighting models. It also illustrates the use of “Rim” lighting and simple translucency for capturing some of the more subtle properties of skin resulting from complex, non-local lighting interactions. Finally, it shows how the various techniques can be combined to produce compelling, stylized skin.



Fig. 10. Example of Skin

Pixel Shader Source Code for Skin

```
struct fragin
{
    float2 texcoords          : TEXCOORD0;
```

```

float4 shadowcoords      : TEXCOORD1;
float4 tangentToEyeMat0  : TEXCOORD4;
float3 tangentToEyeMat1  : TEXCOORD5;
float3 tangentToEyeMat2  : TEXCOORD6;
float3 eyeSpacePosition  : TEXCOORD7;
};

float3 hgphase( float3 v1, float3 v2, float3 g )
{
    float costheta;
    float3 g2;
    float3 gtemp;

    costheta = dot( -v1, v2 );
    g2 = g*g;
    gtemp = 1.0.xxx + g2 - 2.0*g*costheta;
    gtemp = pow( gtemp, 1.5.xxx );
    gtemp = (1.0.xxx - g2) / gtemp;
    return gtemp;
}

// Computes the single-scattering approximation to
// scattering from a one-dimensional volumetric surface.
float3 singleScatter( float3 wi, float3 wo, float3 n,
                     float3 g, float3 albedo,
                     float thickness )
{
    float win = abs(dot(wi,n));
    float won = abs(dot(wo,n));
    float eterm;
    float3 result;

    eterm = 1.0 - exp( (-((1./win)+(1./won))*thickness) );
    result = eterm * (albedo * hgphase( wo, wi, g ) /
                     (win + won));

    return result;
}

// i is the incident ray
// n is the surface normal
// eta is the ratio of indices of refraction
// r is the reflected ray
// t is the transmitted ray

```

```

float fresnel( float3 i, float3 n, float eta,
               out float3 r, out float3 t )
{
    float result;
    float c1;
    float cs2;
    float tflag;

    // Refraction vector courtesy Paul Heckbert.
    c1 = dot(-i,n);
    cs2 = 1.0-eta*eta*(1.0-c1*c1);
    tflag = (float) (cs2 >= 0.0);
    t = tflag * ((eta*c1-sqrt(cs2))*n) + eta*i);
    // t is already unit length or (0,0,0)

    // Compute Fresnel terms
    // (From Global Illumination Compendium.)
    float ndott;
    float cosr_div_cosi;
    float cosi_div_cosr;
    float fs;
    float fp;
    float kr;

    ndott = dot(-n,t);
    cosr_div_cosi = ndott / c1;
    cosi_div_cosr = c1 / ndott;
    fs = (cosr_div_cosi - eta) / (cosr_div_cosi + eta);
    fs = fs * fs;
    fp = (cosi_div_cosr - eta) / (cosi_div_cosr + eta);
    fp = fp * fp;
    kr = 0.5 * (fs+fp);
    result = tflag*kr + (1.-tflag);
    r = reflect( i, n );

    return result;
}

float4 main( fragin In,
             uniform sampler2D tex0,
             uniform sampler2D tex1,
             uniform sampler2D tex2,
             uniform sampler2D tex3,
             uniform float3 eyeSpaceLightPosition,
             uniform float thickness,

```

```

uniform float4 ambient ) : COLOR
{
    float bscale = In.tangentToEyeMat0.w;

    float eta = (1.0/1.4);

    // ratio of indices of refraction (air/skin)
    float m = 34.; // specular exponent
    float4 lightColor = { 1, 1, 1, 1 }; // light color
    float4 sheenColor = { 1, 1, 1, 1 }; // sheen color
    float4 skinColor = tex2D( tex1, In.texcoords );
    float3 g = { 0.8, 0.3, 0.0 };
    float3 albedo = { 0.8, 0.5, 0.4 };

    // oiliness mask
    float4 oiliness = 0.9 * tex2D( tex2, In.texcoords);

    // Get eye-space eye vector.
    float3 v = normalize( -In.eyeSpacePosition );

    // Get eye-space light and halfangle vectors.
    float3 l = normalize( eyeSpaceLightPosition -
                          In.eyeSpacePosition );
    float3 h = normalize( v + l );

    // Get tangent-space normal vector from normal map.
    float3 tangentSpaceNormal = tex2D(tex0, In.texcoords).rgb;
    float3 bumpscale = { bscale, bscale, 1.0 };
    tangentSpaceNormal = tangentSpaceNormal * bumpscale;

    // Transform it into eye-space.
    float3 n;
    n[0] = dot( In.tangentToEyeMat0.xyz, tangentSpaceNormal );
    n[1] = dot( In.tangentToEyeMat1, tangentSpaceNormal );
    n[2] = dot( In.tangentToEyeMat2, tangentSpaceNormal );
    n = normalize( n );

    // Compute the lighting equation.
    float ndotl = max( dot(n,l), 0 ); // clamp 0 to 1
    float ndoth = max( dot(n,h), 0 ); // clamp 0 to 1
    float flag = (float)(ndotl > 0);

    // Compute oil, sheen, subsurf scattering contributions.
    float4 oil;
    float4 sheen;

```

```

float4 subsurf;
float  Kr, Kr2;
float  Kt, Kt2;
float3 T, T2;
float3 R, R2;

// Compute fresnel at sheen layer, ramp it up a bit.
Kr = fresnel( -v, n, eta, R, T );
Kr = smoothstep( 0.0, 0.5, Kr );
Kt = 1.0 - Kr;

// Compute the refracted light ray and the refraction
// coefficient.
Kr2 = fresnel( -l, n, eta, R2, T2 );
Kr2 = smoothstep( 0.0, 0.5, Kr2 );
Kt2 = 1.0 - Kr2;

// For oil contribution, modulate the oiliness mask by a
// specular term.
oil = 0.5 * oiliness * pow( ndoth, m );

// For sheen contribution, modulate Fresnel term by
// sheen color times specular. Modulate by additional
// diffuse term to soften it a bit.
sheen = 2.5*Kr*sheenColor*(ndotl*(0.2 + pow( ndoth, m)));

// Compute single scattering approximation to subsurface
// scattering. Here we compute 3 scattering terms
// simultaneously and the results end up in the x,y,z
// components of a float3. Using 3 terms approximates
// distribution of multiply-scattered light. For
// details see: Matt Pharr's SIGGRAPH 2001 RenderMan
// course notes "Layered Media for Surface Shaders".
float3 temp = singleScatter( T2, T, n, g, albedo,
                           thickness );
subsurf = 2.5 * skinColor * ndotl * Kt * Kt2 *
          (temp.x+temp.y+temp.z);

// Add contributions from oil, sheen, and subsurface
// scattering and modulate by light color and result
// of a shadow map lookup.
return lightColor*tex2Dproj( tex3, In.shadowcoords ).r *
       (oil + sheen + subsurf);
}

```

Thin Film Effect

Description

This demo shows a thin film interference effect. Specular and diffuse lighting are computed per-vertex in a Cg program, along with a view depth parameter, which is computed using the view vector, surface normal, and the depth of the thin film on the surface of the object. The view depth is then perturbed in an ad-hoc manner per-fragment by the underlying decal texture, and is then used to lookup into a 1D texture containing the precomputed destructive interference for red / green / blue wavelengths given a particular view depth. This interference value is then used to modulate the specular lighting component of the standard lighting equation.




Fig. 11. Example of Thin Film Effect

Vertex Shader Source Code for Thin Film Effect

```
// define inputs from application
struct a2v
{
    float4 Position : POSITION;
```

```

    float3 Normal    : NORMAL;
};

// define outputs from vertex shader
struct v2f
{
    float4 HPOS      : POSITION;
    float4 diffCol    : COLOR0;
    float4 specCol    : COLOR1;
    float2 filmDepth  : TEXCOORD0;
};

v2f main(a2v IN,
        uniform float4x4 WorldViewProj,
        uniform float4x4 WorldViewIT,
        uniform float4x4 WorldView,
        uniform float4 LightVector,
        uniform float4 FilmDepth,
        uniform float4 EyeVector)
{
    v2f OUT;

    //transform position to clip space
    OUT.HPOS = mul(WorldViewProj, IN.Position);

    float4 tempnorm = float4(IN.Normal, 0.0);

    // transform normal from model-space to view-space
    float3 normalVec = mul(WorldViewIT, tempnorm).xyz;
    normalVec = normalize(normalVec);

    // compute the eye->vertex vector
    float3 eyeVec = EyeVector.xyz;

    // compute the view depth for the thin film
    float viewdepth = (1.0 / dot(normalVec, eyeVec)) *
        FilmDepth.x;

    OUT.filmDepth = viewdepth.xx;

    // store normalized light vector
    float3 lightVec = normalize((float3)LightVector);

    // calculate half angle vector
    float3 halfAngleVec = normalize(lightVec + eyeVec);

```

```

    // calculate diffuse component
    float diffuse = dot(normalVec, lightVec);

    // calculate specular component
    float specular = dot(normalVec, halfAngleVec);

    // use the lit instruction to calculate lighting,
    // automatically clamp
    float4 lighting = lit(diffuse, specular, 32);

    // output final lighting results
    OUT.diffCol = (float4)lighting.y;
    OUT.specCol = (float4)lighting.z;

    return OUT;
}

```

Pixel Shader Source Code for Thin Film Effect

```

struct v2f
{
    float3 diffCol    : COLOR0;
    float3 specCol    : COLOR1;
    float2 filmDepth : TEXCOORD0;
};

void main( v2f IN,
           out float4 color : COLOR,
           uniform sampler2D fringeMap,
           uniform sampler2D diffMap)
{
    // diffuse material color
    float3 diffCol = float3(0.3, 0.3, 0.5);

    // lookup fringe value based on view depth
    float3 fringeCol = (float3)tex2D(fringeMap, IN.filmDepth);

    // modulate specular lighting by fringe color,
    // combine with regular lighting
    color.rgb = fringeCol*IN.specCol + IN.diffCol*diffCol;
    color.a = 1.0;
}

```

Car Paint 9

Description

This car paint shader uses gonioreflectometric paint samples measured by Cornell University. The samples were converted into a 2D texture map which is indexed using **NdotL** and **NdotH** as the **(s, t)** coordinate pair, and which provides the diffuse component of our lighting equation. The specular term is calculated using the Blinn model, and also includes a term which simulates the clear coat's metallic flecks.

The fleck normal mipmap chain has randomly generated vectors which reside within a positive Z cone in tangent space. The cone is reduced gradually at every level such that in the distance the flecks are pointing mostly up. The flecks' specular power and their contribution are reduced by distance, to give it a grainier appearance up close and a more uniform appearance from afar. Next, the view vector is reflected off a wavy normal map—which represents the object's natural undulations—to index into the environment map. The shininess of the clear coat itself is calculated by scaling the Fresnel term by the luminance of the environment map. (The luminance transfer function selects only the perceptually bright areas of the environment map in order not to reflect the darker areas of the scene.) Finally, the shader lerp's between the diffuse paint color and the reflection based on the Fresnel term, and adds the specular highlights.




Fig. 12. Example of Car Paint 9

Vertex Shader Source Code for Car Paint 9

```

// This shader is based on the Time Machine temporal rust
// shader. Car paint data was measured by Cornell
// University from samples provided by Ford Motor Company.

struct a2v {
    float4 OPosition : POSITION;
    float3 ONormal    : NORMAL;
    float2 uv         : TEXCOORD0;
    float3 Tangent    : TEXCOORD1;
    float3 Binormal   : TEXCOORD2;
    float3 Normal     : TEXCOORD3;
};

struct VS_OUTPUT {
    float4 HPosition : POSITION; // coord position in window
    float2 uv        : TEXCOORD0; // wavy/fleckmap coords
    float3 light     : TEXCOORD1; // light pos (tangent space)
    float4 halfangle : TEXCOORD2; // Blinn halfangle
    float3 reflection: TEXCOORD3; // Refl vector (per-vertex)
    float4 view      : TEXCOORD4; // view (tangent space)
    float3 tangent   : TEXCOORD5; // view-tangent matrix
    float3 binormal  : TEXCOORD6; // ...
    float3 normal    : TEXCOORD7; // ...
    float fresn      : COLOR0;
};

VS_OUTPUT main( a2v vert,
    // TRANSFORMATIONS
    uniform float4x4 ModelView,
    uniform float4x4 ModelViewIT,
    uniform float4x4 ModelViewProj,
    uniform float3  LightVector,    // Obj space
    uniform float3  EyePosition )   // Obj space
{
    VS_OUTPUT O;

    // Generate homogeneous POSITION
    O.HPosition = mul(ModelViewProj, vert.OPosition);

    // Generate BASIS matrix
    float3x3 ModelTangent = { normalize(vert.Tangent),
                             normalize(vert.Binormal),
                             normalize(vert.Normal) };

```

```

// FRESNEL          = { OFFSET, SCALE, POWER, UNUSED };
float4 Fresnel      = { 0.1f, 4.2f, 4.4f, 0.0f };

float3x3 ViewTangent = mul(ModelTangent,
                           (float3x3)ModelViewIT);

// Generate VIEW SPACE vectors
float3 viewN = normalize(mul((float3x3)ModelView,
                           vert.ONormal));
float4 viewP = mul(ModelView, vert.OPosition);
viewP.w = 1-saturate(sqrt(dot(viewP.xyz,
                             viewP.xyz))*0.01);
float3 viewV = -viewP.xyz;

// Generate OBJECT SPACE vectors
float3 objV = normalize(EyePosition-vert.OPosition.xyz);
float3 objL = normalize(LightVector);
float3 objH = normalize(objL + objV);

// Generate TANGENT SPACE vectors
float3 tanL = mul(ModelTangent, objL);
float3 tanV = mul(ModelTangent, objV);
float3 tanH = mul(ModelTangent, objH);

// Generate REFLECTION vector for per-vertex
// reflection look-up
float3 reflection = reflect(-viewV, viewN);

// Generate FRESNEL term
float ndv = saturate(dot(viewN, viewV));
float FresnelApprox = (pow((1-ndv),Fresnel.z)*Fresnel.y +
                      Fresnel.x);

// Fill OUTPUT parameters
0.uv.xy      = vert.uv;           // TEXCOORD0.xy
0.light      = tanL;             // Tangent space LIGHT
// Tangent space HALF-ANGLE
0.halfangle  = float4(tanH.x, tanH.y,
                    tanH.z, 1-exp(-viewP.w));
0.reflection = reflection;        // View space REFLECTION
// Tangent space VIEW + distance attenuation
0.view      = float4(tanV.x, tanV.y,
                    tanV.z, viewP.w);

```

```

// VIEWTANGENT
0.tangent    = normalize(ViewTangent[0]); // column 0
0.binormal   = normalize(ViewTangent[1]); // column 1
0.normal     = normalize(ViewTangent[2]); // column 2
0.fresn      = FresnelApprox;

return 0;
}

```

Pixel Shader Source Code for Car Paint 9

```

// This shader is based on the Time Machine temporal rust
// shader. Car paint data was measured by Cornell
// University from samples provided by Ford Motor Company.
//

struct VS_OUTPUT {
    float4 HPosition : POSITION; // coord position in window
    float2 uv        : TEXCOORD0; // wavy/fleckmap coords
    float3 light     : TEXCOORD1; // light pos (tangent space)
    float4 halfangle : TEXCOORD2; // Blinn halfangle
    float3 reflection: TEXCOORD3; // Refl vector (per-vertex)
    float4 view      : TEXCOORD4; // view (tangent space)
    float3 tangent   : TEXCOORD5; // view-tangent matrix
    float3 binormal  : TEXCOORD6; // ...
    float3 normal    : TEXCOORD7; // ...
    float fresn      : COLOR0;
};

// PIXEL SHADER
float4 main( VS_OUTPUT vert,
             uniform sampler2D WavyMap          : register(s0),
             uniform samplerCUBE EnvironmentMap : register(s1),
             uniform sampler2D PaintMap         : register(s2),
             uniform sampler2D FleckMap         : register(s3),
             uniform float Ambient ) : COLOR
{
    // NEWPAINTSPEC      = { UNUSED, SPEC POWER, GLOSSINESS,
    //                       FLECK SPEC POWER }
    float4 NewPaintSpec = { 0.0f, 64.0f, 3.8f, 8.0f };
    float3 ClearCoat    = { 0.299f, 0.587f, 0.114f };
    float3 FleckColor   = { 0.9, 1.05, 1.0 };
    float3 WavyScale    = { 0.2, -0.2, 1.0 };
}

```

```

// Tangent space LIGHT vector
float3 L = normalize(vert.light);

// Tangent space HALF-ANGLE vector
float3 H = normalize(vert.halfangle.xyz);

// Tangent space VIEW vector
float3 V = normalize(vert.view.xyz);
float v_dist = vert.view.w;

// Tangent space WAVY_NORMAL
float3 wavyN = (float3)tex2D(WavyMap, vert.uv)*2-1;
wavyN = normalize(wavyN*WavyScale);

// PAINT
// A normal map map could be loaded here instead if
// we wanted more detail. In this case we have a
// uniform tangent space normal (0,0,1)
float n_d_l = L.z;
float n_d_h = H.z;
float3 paint_color = (float3)tex2D(PaintMap,
                                float2(n_d_l, n_d_h));

// SPECULAR POWER - use a saturated diffuse term
// to clamp the backlighting
n_d_h = saturate(n_d_l*4)*pow(n_d_h, NewPaintSpec.y);

// REFLECTION ENVIRONMENT
// Reflect view vector about wavy normal and bring
// to view space
float3 R = reflect(-V, wavyN);
R = R.x*vert.tangent + R.y*vert.binormal +
    R.z*vert.normal;
float3 reflect_color = (float3)texCUBE(EnvironmentMap, R);

// FLECKS
// Load random 3-vector flecks from fleck_map
// Reduce tiling artifacts by sampling at
// different frequencies
float3 fleckN = (float3)tex2D(FleckMap, vert.uv*37)*2-1;
float3 fleckN = ((float3)tex2D(FleckMap, vert.uv*23)*2-1)/2 +
    fleckN/2;
float fleck_n_d_h = saturate(dot(fleckN, H));
float3 fleck_color = FleckColor * pow(fleck_n_d_h,

```

```

        lerp(NewPaintSpec.y, NewPaintSpec.w, v_dist));
// Control the ambient fleckiness and also
// attenuate with distance
fleck_color = fleck_color*Ambient*vert.halfangle.w;

// DIFFUSE
float k_d = saturate(n_d_l*1.2);
float3 paintResult = lerp(Ambient*paint_color,
                        paint_color, k_d);

// FRESNEL
float Fresnel = saturate(dot(ClearCoat, reflect_color));
Fresnel = pow(Fresnel, NewPaintSpec.z);

// This helps make the clear coat less omnipresent --
// only the really (perceptually) bright areas reflect
// the most.
Fresnel = saturate(vert.fresn*Fresnel);
// Show more of the specular reflection environment
// when in fresnel zones
// diffuse * (1-fresnel) + environment * (fresnel)
paintResult = lerp(paintResult, reflect_color, Fresnel);

// SPECULAR
// diffuse + specular + flecks
paintResult = paintResult + n_d_h + fleck_color;

// OUTPUT
return paintResult.xyz;
}

```



Basic Profile Sample Shaders

This chapter provides a set of basic profile sample shaders written in Cg. Each shader comes with an accompanying snapshot, description, and source code.

Examples shown are:

- ❑ Anisotropic Lighting
- ❑ Bump Dot3x2 Diffuse and Specular
- ❑ Bump-Reflection Mapping
- ❑ Fresnel
- ❑ Grass
- ❑ Refraction
- ❑ Shadow Mapping
- ❑ Shadow Volume Extrusion
- ❑ Sine Wave Demo
- ❑ Matrix Palette Skinning

Anisotropic Lighting

Description

The anisotropic lighting effect (Fig. 13.) shows the vertex program's half-angle vector calculation. It uses **HdotN** and **LdotN** per-vertex to look up into a 2D texture to achieve interesting lighting effects.




Fig. 13. Example of Anisotropic Lighting

Vertex Shader Source Code for Anisotropic Lighting

```

struct appdata {
    float3 Position : POSITION;
    float3 Normal : NORMAL;
};

struct vpconn {
    float4 Hposition : POSITION;
    float4 TexCoord0 : TEXCOORD0;
};

vpconn main(appdata IN,
    uniform float4x4 WorldViewProj,
    uniform float3x3 WorldIT,
    uniform float3x4 World,
    uniform float3 LightVec,
    uniform float3 EyePos)
{
    vpconn OUT;

    float3 worldNormal = normalize(mul(WorldIT, IN.Normal));

    //build float4
    float4 tempPos;
    tempPos.xyz = IN.Position.xyz;
    tempPos.w = 1.0;

    //compute world space position
    float3 worldSpacePos = mul(World, tempPos);
    //vector from vertex to eye, normalized
    float3 vertToEye = normalize(EyePos - worldSpacePos);

    //h = normalize(l + e)
    float3 halfAngle = normalize(vertToEye + LightVec);

    OUT.TexCoord0.x = max(dot(LightVec, worldNormal), 0.0);
    OUT.TexCoord0.y = max(dot(halfAngle, worldNormal), 0.0);
    // transform into homogeneous-clip space
    OUT.Hposition = mul(WorldViewProj, tempPos);

    return OUT;
}

```

Bump Dot3x2 Diffuse and Specular

Description

The bump dot3x2 diffuse and specular effect mixes bump mapping with diffuse and specular lighting based on the **texm3x2tex** DirectX 8 pixel shader instruction (**DOT_PRODUCT_TEXTURE_2D** in OpenGL). This instruction computes the dot product of the normal and the light vector, corresponding to the diffuse light component, and the dot product of the normal and the half angle vector, corresponding to the specular light component. This results into two scalar values that are used as texture coordinates to look up a 2D illumination texture containing the diffuse color and the specular term in its alpha component. Since the normal fetched from the normal map is in tangent space, both the light vector and the half angle vector are transformed to this space by the vertex shader (Fig. 14.).



Fig. 14. Example of Bump Dot3x2 Diffuse and Specular

Vertex Shader Source Code for Bump Dot3x2

```

struct a2v {
    float4 Position : POSITION; //in object space
    float3 Normal : NORMAL; //in object space
    float2 TexCoord : TEXCOORD0;
    float3 T : TEXCOORD1; //in object space
    float3 B : TEXCOORD2; //in object space
    float3 N : TEXCOORD3; //in object space
};

struct v2f {
    float4 Position : POSITION; //in projection space
    float4 Normal : COLOR0; //in tangent space
    float4 LightVectorUnsigned : COLOR1; //in tangent space
    float3 TexCoord0 : TEXCOORD0;
    float3 TexCoord1 : TEXCOORD1;
    float4 LightVector : TEXCOORD2; //in tangent space
    float4 HalfAngleVector : TEXCOORD3; //in tangent space
};

v2f main(a2v IN,
    uniform float4x4 WorldViewProj,
    uniform float4 LightVector, //in object space
    uniform float4 EyePosition //in object space
    )
{
    v2f OUT;

    // pass texture coordinates for
    // fetching the diffuse map
    OUT.TexCoord0.xy = IN.TexCoord.xy;

    // pass texture coordinates for
    // fetching the normal map
    OUT.TexCoord1.xy = IN.TexCoord.xy;

    // compute the 3x3 transform from
    // tangent space to object space
    float3x3 objToTangentSpace;
    objToTangentSpace[0] = IN.T;
    objToTangentSpace[1] = IN.B;
    objToTangentSpace[2] = IN.N;

    // transform normal from

```

```

// object space to tangent space
OUT.Normal.xyz = 0.5 * mul(objToTangentSpace, IN.Normal) +
    0.5;

// transform light vector from
// object space to tangent space
float3 lightVectorInTangentSpace =
    mul(objToTangentSpace, LightVector.xyz);
OUT.LightVector.xyz = lightVectorInTangentSpace;
OUT.LightVectorUnsigned.xyz = 0.5 *
    lightVectorInTangentSpace + 0.5;

// compute view vector
float3 viewVector =
    normalize(EyePosition.xyz - IN.Position.xyz);

// compute half angle vector
float3 halfAngleVector =
    normalize(LightVector.xyz + viewVector);

// transform half-angle vector from
// object space to tangent space
OUT.HalfAngleVector.xyz =
    mul(objToTangentSpace, halfAngleVector);

// transform position to projection space
OUT.Position = mul(WorldViewProj, IN.Position);

return OUT;
}

```

Pixel Shader Source Code for Bump Dot3x2

```

struct v2f {
    float4 Position : POSITION; //in projection space
    float4 Normal : COLOR0; //in tangent space
    float4 LightVectorUnsigned : COLOR1; //in tangent space
    float3 TexCoord0 : TEXCOORD0;
    float3 TexCoord1 : TEXCOORD1;
    float4 LightVector : TEXCOORD2; //in tangent space
    float4 HalfAngleVector : TEXCOORD3; //in tangent space
};

float4 main(v2f IN,
    uniform sampler2D DiffuseMap,

```

```

        uniform sampler2D NormalMap,
        uniform sampler2D IlluminationMap,
        uniform float Ambient) : COLOR
{
    // fetch base color
    float4 color = tex2D(DiffuseMap, IN.TexCoord0.xy);

    // fetch bump normal and expand it to [-1,1]
    float4 bumpNormal = 2 *
        (tex2D(NormalMap, IN.TexCoord1.xy) - 0.5);

    // compute the dot product between
    //   the bump normal and the light vector,
    // compute the dot product between
    //   the bump normal and the half angle vector,
    // fetch the illumination map using
    //   the result of the two previous dot products
    //   as texture coordinates

    // returns the diffuse color in the
    //   color components and the specular color in the
    //   alpha component

    float2 illumCoord =
        float2(dot(IN.LightVector.xyz, bumpNormal.xyz),
              dot(IN.HalfAngleVector.xyz, bumpNormal.xyz));
    float4 illumination = tex2D(IlluminationMap, illumCoord);

    // expand iterated normal to [-1,1]
    float4 normal = 2 * (IN.Normal - 0.5);

    // compute self-shadowing term
    float shadow = saturate(4 * dot(normal.xyz,
        IN.LightVectorUnsigned.xyz));

    // compute final color
    return (Ambient * color + shadow)
        * (illumination * color + illumination.wwww);
}

```

Bump-Reflection Mapping

Description

This effect mixes bump mapping and reflection mapping based on the **texm3x3vspec** DirectX 8 pixel shader instruction (**DOT_PRODUCT_REFLECT_CUBE_MAP** in OpenGL). This instruction computes three dot products to transform the normal fetched from the normal map into the environment cube space, reflects the transformed normal with respect to the eye vector and fetches a cube map to get the final color. The vertex shader is responsible for computing the transform matrix and the eye vector (Fig. 15.).



Fig. 15. Example of Bump-Reflection Mapping

Vertex Shader Source Code for Bump-Reflection Mapping

```

struct a2v {
    float4 Position : POSITION; // in object space
    float2 TexCoord : TEXCOORD0;
    float3 T : TEXCOORD1; // in object space
    float3 B : TEXCOORD2; // in object space
    float3 N : TEXCOORD3; // in object space
};

struct v2f {
    float4 Position : POSITION; // in projection space
    float4 TexCoord : TEXCOORD0;

    // first row of the 3x3 transform
    // from tangent to cube space
    float4 TangentToCubeSpace0 : TEXCOORD1;

    // second row of the 3x3 transform
    // from tangent to cube space
    float4 TangentToCubeSpace1 : TEXCOORD2;

    // third row of the 3x3 transform
    // from tangent to cube space
    float4 TangentToCubeSpace2 : TEXCOORD3;
};

v2f main(a2v IN,
    uniform float4x4 WorldViewProj,
    uniform float3x4 ObjToCubeSpace,
    uniform float3 EyePosition, // in cube space
    uniform float BumpScale)
{
    v2f OUT;

    // pass texture coordinates for
    // fetching the normal map
    OUT.TexCoord.xy = IN.TexCoord.xy;

    // compute 3x3 transform from tangent to object space
    float3x3 objToTangentSpace;

    // first rows are the tangent and binormal
    // scaled by the bump scale
    objToTangentSpace[0] = BumpScale * IN.T;

```

```

objToTangentSpace[1] = BumpScale * IN.B;
objToTangentSpace[2] = IN.N;
// compute the 3x3 transform from
//   tangent space to cube space:
// TangentToCubeSpace
//   = object2cube * tangent2object
//   = object2cube * transpose(objToTangentSpace)
// (since the inverse of a rotation is its transpose)
//
// So a row of TangentToCubeSpace is the transform by
//   objToTangentSpace of the corresponding row of
//   ObjToCubeSpace

OUT.TangentToCubeSpace0.xyz =
    mul(objToTangentSpace, ObjToCubeSpace[0].xyz);
OUT.TangentToCubeSpace1.xyz =
    mul(objToTangentSpace, ObjToCubeSpace[1].xyz);
OUT.TangentToCubeSpace2.xyz =
    mul(objToTangentSpace, ObjToCubeSpace[2].xyz);

// compute the eye vector
//   (going from eye to shaded point) in cube space
float3 eyeVector = mul(ObjToCubeSpace, IN.Position) -
    EyePosition;

OUT.TangentToCubeSpace0.w = eyeVector.x;
OUT.TangentToCubeSpace1.w = eyeVector.y;
OUT.TangentToCubeSpace2.w = eyeVector.z;

// transform position to projection space
OUT.Position = mul(WorldViewProj, IN.Position);

return OUT;
}

```

Pixel Shader Source Code for Bump and Reflection Mapping

```

struct v2f {
    float4 Position : POSITION; //in projection space
    float4 TexCoord : TEXCOORD0;

    // first row of the 3x3 transform
    //   from tangent to cube space
    float4 TangentToCubeSpace0 : TEXCOORD1;

    // second row of the 3x3 transform
    //   from tangent to cube space
    float4 TangentToCubeSpace1 : TEXCOORD2;

    // third row of the 3x3 transform
    //   from tangent to cube space
    float4 TangentToCubeSpace2 : TEXCOORD3;
};

float4 main(v2f IN,
            uniform sampler2D NormalMap,
            uniform samplerCUBE EnvironmentMap,
            uniform float3 EyeVector) : COLOR
{
    // fetch the bump normal from the normal map
    float4 normal = tex2D(NormalMap, IN.TexCoord.xy);

    // transform the bump normal into cube space
    //   then use the transformed normal and eye vector
    //   to compute the reflection vector that is
    //   used to fetch the cube map
    return texCUBE_reflect_eye_dp3x3(EnvironmentMap,
                                       IN.TangentToCubeSpace2.xyz,
                                       IN.TangentToCubeSpace0,
                                       IN.TangentToCubeSpace1,
                                       normal,
                                       EyeVector);
}

```

Fresnel

Description

This effect computes a reflection vector to lookup into an environment map for reflections, and modulates this by a Fresnel term. The result is reflections only at grazing angles (Fig. 16.).



Fig. 16. Example of Fresnel

Vertex Shader Source Code for Fresnel

```
struct app2vert
{
    float4 Position      : POSITION;
    float4 Normal        : NORMAL;
    float4 TexCoord0     : TEXCOORD0;
};
```

```

struct vert2frag
{
    float4 HPosition      : POSITION;
    float4 Color0         : COLOR0;
    float4 TexCoord0      : TEXCOORD0;
};

vert2frag main(app2vert IN,
               uniform float4x4 ModelViewProj,
               uniform float4x4 ModelView,
               uniform float4x4 ModelViewIT)
{
    vert2frag OUT;

#ifdef PROFILE_ARBVP1
    ModelViewProj = glstate.matrix.mvp;
    ModelView = glstate.matrix.modelview[0];
    ModelViewIT = glstate.matrix.invtrans.modelview[0];
#endif

    OUT.HPosition = mul(ModelViewProj, IN.Position);

    float3 normal = normalize(mul(ModelViewIT,
                                  IN.Normal).xyz);
    float3 eyeToVert = normalize(mul(ModelView,
                                      IN.Position).xyz);

    // reflect the eye vector across the normal vector
    // for reflection
    OUT.TexCoord0 = float4(reflect(eyeToVert, normal), 1.0);

    float f0 = .1;

    // compute the fresnel term
    float oneMCosAngle = 1+dot(eyeToVert,normal);
    oneMCosAngle = pow(oneMCosAngle, 5);
    OUT.Color0 = lerp(oneMCosAngle, 1, f0).xxxx;

    return OUT;
}

```

Grass

Description

This effect shows procedural animation of geometry using a Sine function, along with calculation of a normal for the procedurally deformed geometry (Fig. 17.).




Fig. 17. Example of Grass

Vertex Shader Source Code for Grass

```
struct app2vert {  
    float4 Position : POSITION;  
    float4 Normal : NORMAL;
```

```

    float4 TexCoord0 : TEXCOORD0;
    float4 Color0 : COLOR0;
};

struct vertout {
    float4 Hposition : POSITION;
    float4 Color0 : COLOR0;
    float4 TexCoord0 : TEXCOORD0;
};

vertout main(app2vert IN,
             uniform float4x4 ModelViewProj,
             uniform float4x4 ModelView,
             uniform float4x4 ModelViewIT,
             uniform float4 Constants)
{
    vertout OUT;

    // we need to figure OUT what the position is
    float4 position = IN.Position;
    position.z = 0;
    position.y = 0;

    // add IN the actual base location of
    // the straw (stored IN Color0.xz)
    position.x = position.x + IN.Color0.x;
    position.z = position.z + IN.Color0.z;

    // figure OUT where the wind is coming from
    float4 origin = float4(20,0,20,0);
    float4 dir = position - origin;

    // find the intensity of the wind
    float inten = sin(Constants.x + .2*length(dir)) *
        IN.Position.y;
    dir = normalize(dir);

    // we need to do some Bezier curve stuff here.
    float4 ctrl1 = float4(0,0,0,0);
    float4 ctrl2 = float4(0,IN.Color0.y/2,0,0);
    float4 ctrl3 = float4(dir.x*inten, IN.Color0.y,
                          dir.z*inten, 0);
    // do the Bezier linear interpolation steps
    float t = IN.Color0.w;

```

```

float4 temp = lerp(ctrl1, ctrl2, t);
float4 temp2 = lerp(ctrl2, ctrl3, t);
float4 result = lerp(temp, temp2, t);

// add IN the height and wind displacement components
position = position + result;
position.w = 1;

// transform for sending to the reg. combiners
OUT.Hposition = mul(ModelViewProj, position);

// calculate the texture coordinate
//   from the position passed IN
OUT.TexCoord0 = float4((IN.Position.x + .05)*10,t,1,1);

// find the normal
// we need one more point to do a partial
temp = lerp(ctrl1, ctrl2, t+0.05);
temp2 = lerp(ctrl2, ctrl3, t+0.05);
float4 newResult = lerp(temp, temp2, t+0.05);

// do a crossproduct with a vector that
//   is horizontal across the screen
float normal = cross((result - newResult).xyz,
                    float3(1,0,0));
normal = normalize(normal);

// calculate diffuse lighting off the normal
//   that was just calculated
float3 lightPos = float3(0,5,15);
float3 lightVec = normalize(lightPos - position);
float diffuseInten = dot(lightVec, normal);

// Set up the final color
// The first term is a semi random term based
//   on the total height of this straw
// The second term is the diffuse lighting component
OUT.Color0 = normalize(ctrl3) * diffuseInten *
    IN.Position.z;

return OUT;
}

```

Refraction

Description

This effect performs custom texture coordinate generation to compute a refracted vector per-vertex that is then used to look up in a cube map. Fresnel is also calculated to blend between reflection and refraction ([Fig. 18.](#)).




Fig. 18. Example of Refraction

Vertex Shader Source Code for Refraction

```

struct inputs
{
    float4 Position      : POSITION;
    float4 Normal        : NORMAL;
};

struct outputs
{
    float4 hPosition     : POSITION;
    float4 fresnelTerm   : COLOR0;
    float4 refractVec     : TEXCOORD0;
    float4 reflectVec    : TEXCOORD1;
};

// fresnel approximation
fixed fast_fresnel(float3 I, float3 N,
                  float3 fresnelValues)
{
    fixed power = fresnelValues.x;
    fixed scale = fresnelValues.y;
    fixed bias = fresnelValues.z;

    return bias + pow(1.0 - dot(I, N), power) * scale;
}

outputs main(inputs IN,
             uniform float4x4 ModelViewProj,
             uniform float4x4 ModelView,
             uniform float4x4 ModelViewIT,
             uniform float theta)
{
    outputs OUT;

    OUT.hPosition = mul(ModelViewProj, IN.Position);

    // convert the position and normal into
    // appropriate spaces
    float3 eyeToVert = mul(ModelView, IN.Position).xyz;
    eyeToVert = normalize(eyeToVert);
    float3 normal = mul(ModelViewIT, IN.Normal).xyz;
    normal = normalize(normal);

    OUT.refractVec.xyz = refract(eyeToVert, normal, theta);
}

```

```

    OUT.refractVec.w = 1;

    OUT.reflectVec.xyz = reflect(eyeToVert, normal);
    OUT.reflectVec.w = 1;

    // calculate the fresnel reflection
    OUT.fresnelTerm = fast_fresnel(-eyeToVert, normal,
                                   float3(5.0, 1.0, 0.0));
    return OUT;
}

```

Pixel Shader Source Code for Refraction

```

float4 main(in float3 refractVec    : TEXCOORD0,
            in float3 reflectVec    : TEXCOORD1,
            in float3 fresnelTerm    : COLOR0,

            uniform samplerCUBE environmentMaps[2],
            uniform float enableRefraction,
            uniform float enableFresnel) : COLOR
{
    float3 refractColor = texCUBE(environmentMaps[0],
                                   refractVec).rgb;
    float3 reflectColor = texCUBE(environmentMaps[1],
                                   reflectVec).rgb;

    float3 reflectRefract = lerp(refractColor, reflectColor,
                                   fresnelTerm);

    float3 finalColor = enableRefraction ?
        (enableFresnel ? reflectRefract : refractColor) :
        (enableFresnel ? reflectColor : fresnelTerm);

    return float4(finalColor, 1.0);
}

```

Shadow Mapping

Description

This effect shows generating texture coordinates for shadow mapping, along with using the shadow map in the lighting equation per pixel ([Fig. 19](#)).



Fig. 19. Example of Shadow Mapping

Vertex Shader Source Code for Shadow Mapping

```

struct appdata {
    float3 Position : POSITION;
    float3 Normal : NORMAL;
};

struct vpconn {
    float4 Hposition : POSITION;
    float4 TexCoord0 : TEXCOORD0;
    float4 TexCoord1 : TEXCOORD1;
    float4 Color0 : COLOR0;
};

vpconn main(appdata IN,
            uniform float4x4 WorldViewProj,
            uniform float4x4 TexTransform,
            uniform float3x3 WorldIT,
            uniform float3 LightVec)
{
    vpconn OUT;

    float3 worldNormal = normalize(mul(WorldIT, IN.Normal));

    float ldotn = max(dot(LightVec, worldNormal), 0.0);

    OUT.Color0.xyz = ldotn.xxx;

    float4 tempPos;
    tempPos.xyz = IN.Position.xyz;
    tempPos.w = 1.0;

    OUT.TexCoord0 = mul(TexTransform, tempPos);
    OUT.TexCoord1 = mul(TexTransform, tempPos);

    OUT.Hposition = mul(WorldViewProj, tempPos);

    return OUT;
}

```

Pixel Shader Source Code for Shadow Mapping

```
struct v2f_simple {
    float4 Hposition : POSITION;
    float4 TexCoord0 : TEXCOORD0;
    float4 TexCoord1 : TEXCOORD1;
    float4 Color0 : COLOR0;
};

float4 main(v2f_simple IN,
            uniform sampler2D ShadowMap,
            uniform sampler2D SpotLight) : COLOR
{
    float4 shadow    = tex2D(ShadowMap, IN.TexCoord0.xy);
    float4 spotlight = tex2D(SpotLight, IN.TexCoord1.xy);
    float4 lighting  = IN.Color0;

    return shadow * spotlight * lighting;
}
```

Shadow Volume Extrusion

Description

This effect uses vertex programs to generate shadow volumes by extruding geometry along the light vector (Fig. 20.).



Fig. 20. Example of Shadow Volume Extrusion

Vertex Shader Source Code for Shadow Volume Extrusion

```

struct appdata
{
    float4 Position : POSITION;
    float3 Normal : NORMAL;
    float4 DiffuseColor : COLOR0;
    float2 TexCoord0 : TEXCOORD0;
};

struct vpconn {
    float4 Hposition : POSITION;
    float4 Color0 : COLOR0;
    float2 TexCoord0 : TEXCOORD0;
};

vpconn main(appdata IN,
            uniform float4x4 WorldViewProj,
            uniform float4 LightPos, // (in object space)
            uniform float4 Fatness,
            uniform float4 ShadowExtrudeDist,
            uniform float4 Factors
        )
{
    vpconn OUT;

    // Create normalized vector from vertex to light
    float4 light_to_vert = normalize(IN.Position - LightPos);

    // N dot L to decide if point should be moved away
    //   from the light to extrude the volume
    float ndotl = dot(-light_to_vert.xyz, IN.Normal.xyz);

    // Inset the position along
    // the normal vector direction
    // This moves the shadow volume points
    // inside the model slightly to minimize
    // popping of shadowed areas as
    // each facet comes in and out of shadow.
    // The Fatness value should be negative
    float4 inset_pos = (IN.Normal * Fatness.xyz +
        IN.Position.xyz).xyzz;
    inset_pos.w = IN.Position.w;

    // scale the vector from light to vertex

```

```
float4 extrusion_vec = light_to_vert * ShadowExtrudeDist;

// if ndotl < 0 then the vertex faces
// away from the light, so move it.
// It will be moved along the direction from
// light to vertex to extrude the shadow volume.
float away = (float)(ndotl < 0);

// Move the back-facing shadow volume points
float4 new_position = extrusion_vec * away + inset_pos;

// Transform position to hclip space;
OUT.Hposition = mul(WorldViewProj, new_position);

// Set the color to blue for when the shadow volume
// is rendered in color for illustrative purposes
float4 color = float4(0, 0, Factors.x, 0);

OUT.Color0 = color;
OUT.TexCoord0.xy = IN.TexCoord0;
return OUT;
}
```

Sine Wave Demo

Description

This effect modifies the vertex positions using a sine function based on the current time. It demonstrates use of the built-in **sin()** function. It also computes a normal based on the perturbed mesh, and uses this to compute a reflection vector to look up in a cube map (Fig. 21.).




Fig. 21. Example of Sine Wave

Vertex Shader Source Code for Sine Wave

```

struct appdata {
    float4 TexCoord0 : TEXCOORD0;
};

struct vpconn {
    float4 HPOS : POSITION;
    float4 COL0 : COLOR0;
    float4 TEX0 : TEXCOORD0;
};

vpconn main(appdata IN,
            uniform float4x4 WorldViewProj,
            uniform float3x4 WorldView,
            uniform float3x3 WorldViewIT,
            uniform float3    WavesX,
            uniform float3    WavesY,
            uniform float3    WavesH,
            uniform float3    Time
        )
{
    vpconn OUT;

    float3 angle = WavesX * IN.TexCoord0.x +
                  WavesY * IN.TexCoord0.y;
    angle = angle + Time;

    float3 sine, cosine;
    sincos(angle, sine, cosine);

    // position is: (u, sum(hi * sin(anglei)), v, 1)
    float4 position;
    position.xz = IN.TexCoord0.xy;
    position.y = dot(WavesH, sine);
    position.w = 1.0f;

    OUT.HPOS = mul(WorldViewProj, position);

    // normal is (t h WaveX cos(angle),
    // -1,
    // t h WaveY cos(angle))
    float3 normal;
    normal.x = dot(WavesH * WavesX, cosine);
    normal.y = -1.0f;

```

```
normal.z = dot(WavesH * WavesY, cosine);

// transform normal into eye-space
normal = mul(WorldViewIT, normal);
normal = normalize(normal);

// Transform vertex to eye-space and
//   compute the vector from the eye to the vertex.
// Because the eye is at 0, no subtraction is
//   necessary. Because the reflection of this vector
//   looks into a cube-map normalization is also
//   unnecessary!
float3 eyeVector = mul(WorldView, position);
OUT.TEX0.xyz = reflect(eyeVector, normal);

return OUT;
}
```

Matrix Palette Skinning

Description

This effect performs matrix palette skinning using two bones per vertex. All the bones for the mesh are set in the constant memory, and each vertex includes two indices that indicate which bones influence this vertex. The final skinned positions are computed using these bones, along with the weights supplied per vertex. Tangent-space bases are skinned in a similar fashion and then used to transform the light vector into tangent space for per-pixel bump mapping (Fig. 22.).



Fig. 22. Example of Matrix Palette Skinning

Vertex Shader Source Code for Matrix Palette Skinning

```

struct appdata {
    float3 Position : POSITION;
    float2  Weights : BLENDWEIGHT0;
    float2  Indices : BLENDINDICES;
    float3 Normal : NORMAL;
    float2  TexCoord0 : TEXCOORD0;
    float3 S : TEXCOORD1;
    float3 T : TEXCOORD2;
    float3 SxT : TEXCOORD3;
};

struct vpconn {
    float4 Hposition : POSITION;
    float4 TexCoord0 : TEXCOORD0;
    float4 TexCoord1 : TEXCOORD1;
    float4 Color0 : COLOR0;
};

vpconn main(appdata IN,
            uniform float4x4 WorldViewProj,
            uniform float3x4 Bones[26],
            uniform float3 LightVec)
{
    vpconn OUT;

    float4 tempPos;
    tempPos.xyz = IN.Position.xyz;
    tempPos.w = 1.0;

    // grab first bone matrix
    float i = IN.Indices.x;

    //transform position
    float3 pos0 = mul(Bones[i], tempPos);

    //create 3x3 version of bone matrix
    float3x3 m;
    m._m00_m01_m02 = Bones[i]._m00_m01_m02;
    m._m10_m11_m12 = Bones[i]._m10_m11_m12;
    m._m20_m21_m22 = Bones[i]._m20_m21_m22;

    // transform S, T, SxT
    float3 s0 = mul(m, IN.S);

```

```

float3 t0    = mul(m, IN.T);
float3 sxt0 = mul(m, IN.SxT);

// next bone
i = IN.Indices.y;

// create 3x3 version of bone
m._m00_m01_m02 = Bones[i]._m00_m01_m02;
m._m10_m11_m12 = Bones[i]._m10_m11_m12;
m._m20_m21_m22 = Bones[i]._m20_m21_m22;

float3 pos1 = mul(Bones[i], tempPos);

// transform S, T, SxT
float3 s1    = mul(m, IN.S);
float3 t1    = mul(m, IN.T);
float3 sxt1 = mul(m, IN.SxT);

// final blending

// blend s, t, sxt
float3 finals    = s0 * IN.Weights.x + s1 * IN.Weights.y;
float3 finalT    = t0 * IN.Weights.x + t1 * IN.Weights.y;
float3 finalSxT  = sxt0 * IN.Weights.x + sxt1 * IN.Weights.y;

// blend between the two positions
float3 finalPos = pos0 * IN.Weights.x + pos1 * IN.Weights.y;

float3x3 worldToTangentSpace;

worldToTangentSpace._m00_m01_m02 = finals;
worldToTangentSpace._m10_m11_m12 = finalT;
worldToTangentSpace._m20_m21_m22 = finalSxT;

float3 tangentLight =
normalize(mul(worldToTangentSpace, LightVec));

// scale and bias, add bit of ambient
tangentLight = ((tangentLight + 1.0) * 0.5) + 0.2;

// create float4 with 1.0 alpha
float4 templight;
templight.xyz = tangentLight.xyz;
templight.w = 1.0;
OUT.Color0 = templight;

```

```
// pass through texcoords
OUT.TexCoord0.xy = IN.TexCoord0.xy;
OUT.TexCoord1.xy = IN.TexCoord0.xy;

float4 tempPos2;
tempPos2.xyz = finalPos.xyz;
tempPos2.w = 1.0;

OUT.Hposition = mul(WorldViewProj, tempPos2);

return OUT;
}
```



Appendix A

Cg Language Specification

Language Overview

The Cg language is primarily modeled on ANSI C, but adopts some ideas from modern languages such as C++ and Java, and from earlier shading languages such as RenderMan and the Stanford shading language. The language also introduces a few new ideas. In particular, it includes features designed to represent data flow in stream-processing architectures such as GPUs. Profiles, which are specified at compile time, may subset certain features of the language, including the ability to implement loops and the precision at which certain computations are performed.

Silent Incompatibilities

Most of the changes from ANSI C are either omissions or additions, but there are a few potentially *silent incompatibilities*. These are changes within Cg that could cause a program that compiles without errors to behave in a manner different from C:

- ❑ The type promotion rules for constants are different when the constant is not explicitly typed using a type cast or type suffix. In general, a binary operation between a constant that is not explicitly typed and a variable is performed at the variable's precision, rather than at the constant's default precision.
- ❑ Declarations of **struct** perform an automatic **typedef** (as in C++) and thus could override a previously declared type.
- ❑ Arrays are first-class types that are distinct from pointers. As a result, array assignments semantically perform a copy operation for the entire array.

Similar Operations That Must be Expressed Differently

There are several changes that force the same operation to be expressed differently in Cg than in C:

- ❑ A Boolean type, **bool**, is introduced, with corresponding implications for operators and control constructs.
- ❑ Arrays are first-class types because Cg does not support pointers.
- ❑ Functions pass values by value/result, and thus use an **out** or **inout** modifier in the formal parameter list to return a parameter. By default, formal parameters are **in**, but it is acceptable to specify this explicitly. Parameters can also be specified as **in out**, which is semantically the same as **inout**.

Differences from ANSI C

Cg was developed based on the ANSI-C language with the following major additions, deletions, and changes. (This is a summary—more detail is provided later in this document):

- ❑ Language profiles (described in [“Profiles” on page 225](#)) may subset language capabilities in a variety of ways. In particular, language profiles may restrict the use of **for** and **while** loops. For example, some profiles may only support loops that can be fully unrolled at compile time.
- ❑ A *binding semantic* may be associated with a structure tag, a variable, or a structure element to denote that object’s mapping to a specific hardware or API resource. See [“Binding Semantics” on page 242](#).
- ❑ Reserved keywords **goto**, **break**, and **continue** are not supported.
- ❑ Reserved keywords **switch**, **case**, and **default** are not supported. Labels are not supported either.
- ❑ Pointers and pointer-related capabilities (such as the **&** and **->** operators) are not supported.
- ❑ Arrays are supported, but with some limitations on size and dimensionality. Restrictions on the use of computed subscripts are also permitted. Arrays may be designated as **packed**. The operations allowed on packed arrays may be different from those allowed on unpacked arrays. Predefined **packed** types are provided for vectors and matrices. It is *strongly* recommended these predefined types be used.

- ❑ *Unsize*d arrays can be created by declaring an array's dimension as `[]`. The array's actual dimension can be set at runtime before a final compilation step.
- ❑ There is a built-in swizzle operator: `.xyzw` or `.rgba` for vectors. This operator allows the components of a vector to be rearranged and also replicated. It also allows the creation of a vector from a scalar.
- ❑ For an lvalue, the swizzle operator allows components of a vector or matrix to be selectively written.
- ❑ There is a similar built-in swizzle operator for matrices:

`._m<row><col>[_m<row><col>][...]`

This operator allows access to individual matrix components and allows the creation of a vector from elements of a matrix. For compatibility with DirectX 8 notation, there is a second form of matrix swizzle, which is described later.

- ❑ Numeric data types are different. Cg's primary numeric data types are **float**, **half**, and **fixed**. Fragment profiles are required to support all three data types, but may choose to implement **half** and **fixed** at **float** precision. Vertex profiles are required to support **half** and **float**, but may choose to implement **half** at **float** precision. Vertex profiles may omit support for **fixed** operations, but must still support definition of **fixed** variables. Cg allows profiles to omit run-time support for **int**. Cg allows profiles to treat **double** as **float**.
- ❑ Many operators support per-element vector operations.
- ❑ The `?:`, `||`, `&&`, `!`, and comparison operators can be used with **bool** four-vectors to perform four conditional operations simultaneously. The side effects of all operands to the `?:`, `||`, and `&&` operators are always executed.
- ❑ Non-static global variables and parameters to top-level functions—such as `main()`—may be designated as **uniform**. A **uniform** variable may be read and written within a program, just like any other variable. However, the **uniform** modifier indicates that the initial value of the variable or parameter is expected to be constant across a large number of invocations of the program.
- ❑ A new set of **sampler*** types represents handles to texture objects.
- ❑ Functions may have default values for their parameters, as in C++. These defaults are expressed using assignment syntax.
- ❑ Function overloading is supported.

- ❑ There is no **enum** or **union**.
- ❑ Bit-field declarations in structures are not allowed.
- ❑ There are no bit-field declarations in structures.
- ❑ Variables may be defined anywhere before they are used, rather than just at the beginning of a scope as in C. (That is, we adopt the C++ rules that govern where variable declarations are allowed.) Variables may not be redeclared within the same scope.
- ❑ Vector constructors, such as the form **float4(1,2,3,4)**, may be used anywhere in an expression.
- ❑ A **struct** definition automatically performs a corresponding **typedef**, as in C++.
- ❑ An **interface** can be specified to define a set of methods that comprises an abstract interface.
- ❑ A **struct** type can be declared as implementing an **interface** by adding a colon ":" and the name of the interface after the name of the **struct**.
- ❑ Methods can be defined in the body of a **struct** definition.
- ❑ C++-style `//` comments are allowed in addition to C-style `/*...*/` comments.

Detailed Language Specification

Definitions

The following definitions are based on the ANSI C standard:

- ❑ *Object*
An object is a region of data storage in the execution environment, the contents of which can represent values. When referenced, an object may be interpreted as having a particular type.
- ❑ *Declaration*
A declaration specifies the interpretation and attributes of a set of identifiers.
- ❑ *Definition*
A declaration that also causes storage to be reserved for an object or code that will be generated for a function named by an identifier is a definition.

Profiles

Compilation of a Cg program, a top-level function, always occurs in the context of a compilation profile. The profile specifies whether certain optional language features are supported. These optional language features include certain control constructs and standard library functions. The compilation profile also defines the precision of the **float**, **half**, and **fixed** data types, and specifies whether the **fixed** and **sampler*** data types are fully or only partially supported. The choice of a compilation profile is made externally to the language, by using a compiler command-line switch, for example.

The profile restrictions are only applied to the top-level function that is being compiled and to any variables or functions that it references, either directly or indirectly. If a function is present in the source code, but not called directly or indirectly by the top-level function, it is free to use capabilities that are not supported by the current profile.

The intent of these rules is to allow a single Cg source file to contain many different top-level functions that are targeted at different profiles. The core Cg language specification is sufficiently complete to allow all of these functions to be parsed. The restrictions provided by a compilation profile are only needed for code generation, and are therefore only applied to those functions for which code is being generated. This specification uses the word *program* to refer to the top-level function, any functions the top-level function calls, and any global variables or **typedef** definitions it references.

Each profile must have a separate specification that describes its characteristics and limitations.

This core Cg specification requires certain minimum capabilities for all profiles. In some cases, the core specification distinguishes between vertex-program and fragment-program profiles, with different minimum capabilities for each.

The Uniform Modifier

Non-static global variables and parameters passed to functions, such as **main()**, can be declared with an optional qualifier **uniform**. To specify a **uniform** variable, use this syntax:

uniform <type> <variable>

For example,

```
uniform float4 myVector;
```

or

```
float4 foo(uniform float4 uv);
```

If the **uniform** qualifier is specified for a function that is not top level, it is meaningless and is ignored. The intent of this rule is to allow a function to serve either as a top-level function or as one that is not.

Note that **uniform** variables may be read and written just like non-**uniform** variables. The **uniform** qualifier simply provides information about how the initial value of the variable is to be specified and stored, through a mechanism external to the language.

Typically, the initial value of a **uniform** variable or parameter is stored in a different class of hardware register. Furthermore, the external mechanism for specifying the initial value of **uniform** variables or parameters may be different than that used for specifying the initial value of non-**uniform** variables or parameters. Parameters qualified as **uniform** are normally treated as persistent state, while non-**uniform** parameters are treated as streaming data, with a new value specified for each stream record (such as within a vertex array).

Function Declarations

Functions are declared essentially as in C. A function that does not return a value must be declared with a **void** return type. A function that takes no parameters may be declared in one of two ways:

- ❑ As in C, using the **void** keyword: **functionName(void)**
- ❑ With no parameters at all: **functionName()**

Functions may be declared as **static**. If so, they may not be compiled as a program and are not visible from other compilation units.

Overloading of Functions by Profile

Cg supports overloading of functions by compilation profile. This capability allows a function to be implemented differently for different profiles. It is also useful because different profiles may support different subsets of the language capabilities, and because the most efficient implementation of a function may be different for different profiles.

The profile name must immediately precede the type name in the function declaration. For example, to define two different versions of the function **myfunc()** for the **profileA** and **profileB** profiles:

```
profileA float myfunc(float x) { /* ... */ };
profileB float myfunc(float x) { /* ... */ };
```

If a type is defined (using a **typedef**) that has the same name as a profile, the identifier is treated as a type name and is not available for profile overloading at any subsequent point in the file.

If a function definition does not include a profile, the function is referred to as an *open profile* function. Open-profile functions apply to all profiles.

Several wildcard profile names are defined. The name **vs** matches any vertex profile, while the name **ps** matches any fragment or pixel profile.

The names **ps_1** and **ps_2** match any DirectX 8 pixel shader 1.x profile or DirectX 9 pixel shader 2.x profile, respectively. Similarly, the names **vs_1** and **vs_2** match any DirectX vertex shader 1.x or 2.x, respectively. Additional valid wildcard profile names may be defined by individual profiles.

In general, the most specific version of a function is used. More details are provided in “[Function Overloading](#)” on page 240, but roughly speaking, the search order is the following:

1. Version of the function with the exact profile overload
2. Version of the function with the most specific wildcard profile overload (such as **vs** or **ps_1**)
3. Version of the function with no profile overload

This search process allows generic versions of a function to be defined that can be overridden as needed for particular hardware.

Syntax for Parameters in Function Definitions

Functions are declared in a manner similar to C, but the parameters in function definitions may include a binding semantic (see “[Binding Semantics](#)” on page 242) and a default value.

Each parameter in a function definition takes the following form:

[uniform] <type> identifier [: <binding_semantic>] [= <default>]

where

- ❑ **<type>** may include the qualifiers **in**, **out**, **inout**, and **const**, as discussed in “[Type Qualifiers](#)” on page 233.

- **<default>** is an expression that resolves to a constant at compile time.

Default values are only permitted for **uniform** parameters, and for **in** parameters to functions that are not top-level.

Function Calls

A function call returns an rvalue. Therefore, if a function returns an array, the array may be read but not written. For example, the following is allowed:

```
y = myfunc(x)[2];
```

But, this is not: `myfunc(x)[2] = y;`

For multiple function calls within an expression, the calls can occur in *any* order—it is undefined.

Method Calls

Structures may have methods declared and defined in their structure definitions. For example,

```
struct Foo {
    float value;
```

Structure methods are called using the “.” notation: given an object **f** of type **Foo**, the **valueTimesTwo()** method is called by **f.valueTimesTwo()**.

Interfaces

Interfaces may be declared in order to define a set of methods that a structure must provide in order to implement that interface.

Programs and functions can take interfaces as parameters, where the specific structure types being passed to them may be resolved at runtime. Depending on hardware limitations, some profiles may require that the concrete types associated with a particular usage of interfaces be resolved by the runtime before the program can execute.

Interfaces are specified with the **interface** keyword:

A structure indicates that it implements a particular interface with a colon and the name of the interface:

```
struct PointLight : Light {
    float3 illuminate(float3 position) { ... }
};
```

A structure may only implement a single interface and inheritance between structures is not supported.

Types

Cg's types are as follows:

- ❑ The **int** type is preferably 32-bit two's complement. Profiles may optionally treat **int** as **float**.
- ❑ The **float** type is as close as possible to the IEEE single precision (32-bit) floating point. Profiles must support the **float** data type.
- ❑ The **half** type is lower-precision IEEE-like floating point. Profiles must support the **half** type, but may choose to implement it with the same precision as the **float** type.
- ❑ The **fixed** type is a signed type with a range of at least [-2,2) and with at least 10 bits of fractional precision. Overflow operations on the data type clamp rather than wrap. Fragment profiles must support the **fixed** type, but may implement it with the same precision as the **half** or **float** types. Vertex profiles are required to provide partial support (see [“Partial Support of Types” on page 231](#)) for the **fixed** type. Vertex profiles have the option to provide full support for the **fixed** type or to implement the **fixed** type with the same precision as the **half** or **float** types.
- ❑ The **bool** type represents Boolean values. Objects of **bool** type are either true or false.
- ❑ The **cint** type is 32-bit two's complement. This type is meaningful only at compile time; it is not possible to declare objects of type **cint**.
- ❑ The **cfloat** type is IEEE single-precision (32-bit) floating point. This type is meaningful only at compile time; it is not possible to declare objects of type **cfloat**.
- ❑ The **void** type may not be used in any expression. It may only be used as the return type of functions that do not return a value.

- ❑ The **sampler*** types are handles to texture objects. Formal parameters of a program or function may be of type **sampler***. No other definition of **sampler*** variables is permitted. A **sampler*** variable may only be used by passing it to another function as an **in** parameter. Assignment to **sampler*** variables is not permitted, and **sampler*** expressions are not permitted.

The following **sampler*** types are always defined: **sampler**, **sampler1D**, **sampler2D**, **sampler3D**, **samplerCUBE**, and **samplerRECT**. The base **sampler** type may be used in any context in which a more specific sampler type is valid. However, a **sampler** variable must be used in a consistent way throughout the program. For example, it cannot be used in place of both a **sampler1D** and a **sampler2D** in the same program.

Fragment profiles are required to fully support the **sampler**, **sampler1D**, **sampler2D**, **sampler3D**, and **samplerCUBE** data types. Fragment profiles are required to provide partial support (see [“Partial Support of Types” on page 231](#)) for the **samplerRECT** data type and may optionally provide full support for this data type.

Vertex profiles are required to provide partial support for the six sampler data types and may optionally provide full support for these data types.

- ❑ An **array** type is a collection of one or more elements of the same type. An **array** variable has a single index.
- ❑ Some array types may be optionally designated as *packed*, using the **packed** type modifier. The storage format of a packed type may be different from the storage format of the corresponding unpacked type. The storage format of packed types is implementation dependent, but must be consistent for any particular combination of compiler and profile. The operations supported on a packed type in a particular profile may be different than the operations supported on the corresponding unpacked type in that same profile. Profiles may define a maximum allowable size for packed arrays, but must support at least size 4 for packed vector (one-dimensional array) types, and 4x4 for packed matrix (two-dimensional array) types.
- ❑ When declaring an array of arrays in a single declaration, the **packed** modifier only refers to the outermost array. However, it is possible to declare a packed array of packed arrays by declaring the first level of array in a **typedef** using the **packed** keyword and then declaring a packed array of this type in a second statement. It is not possible to have a packed array of unpacked arrays.

- ❑ For any supported numeric data type **TYPE**, implementations must support the following packed array types, which are called *vector types*. Type identifiers must be predefined for these types in the global scope:

```
typedef packed TYPE TYPE1[1];
typedef packed TYPE TYPE2[2];
typedef packed TYPE TYPE3[3];
typedef packed TYPE TYPE4[4];
```

For example, implementations must predefine the type identifiers **float1**, **float2**, **float3**, **float4**, and so on for any other supported numeric type.

- ❑ For any supported numeric data type **TYPE**, implementations must support the following packed array types, which are called *matrix types*. Implementations must also predefine type identifiers (in the global scope) to represent these types:

```
packed TYPE1 TYPE1x1[1];      packed TYPE1 TYPE3x1[3];
packed TYPE2 TYPE1x2[1];      packed TYPE2 TYPE3x2[3];
packed TYPE3 TYPE1x3[1];      packed TYPE3 TYPE3x3[3];
packed TYPE4 TYPE1x4[1];      packed TYPE4 TYPE3x4[3];
packed TYPE1 TYPE2x1[2];      packed TYPE1 TYPE4x1[4];
packed TYPE2 TYPE2x2[2];      packed TYPE2 TYPE4x2[4];
packed TYPE3 TYPE2x3[2];      packed TYPE3 TYPE4x3[4];
packed TYPE4 TYPE2x4[2];      packed TYPE4 TYPE4x4[4];
```

For example, implementations must predefine the type identifiers **float2x1**, **float3x3**, **float4x4**, and so on. A **typedef** follows the usual matrix-naming convention of **TYPE_rows_X_columns**. If we declare **float4x4 a**, then **a[3]** is equivalent to **a._m30_m31_m32_m33**.

Both expressions extract the third row of the matrix.

- ❑ Implementations are required to support indexing of vectors and matrices with constant indices.
- ❑ A **struct** type is a collection of one or more members of possibly different types.
- ❑ An **interface** type defines a collection of methods that comprises an abstract interface.

Partial Support of Types

This specification mandates *partial support* for some types. Partial support for a type requires the following:

- ❑ Definitions and declarations using the type are supported.

- ❑ Assignment and copy of objects of that type are supported (including implicit copies when passing function parameters).
- ❑ Top-level function parameters may be defined using that type.

If a type is partially supported, variables may be defined using that type but no useful operations can be performed on them. Partial support for types makes it easier to share data structures in code that is targeted at different profiles.

Type Categories

- ❑ The *integral* type category includes types **cint** and **int**.
- ❑ The *floating* type category includes types **cfloat**, **float**, **half**, and **fixed**. (Note that floating really means floating or fixed/fractional.)
- ❑ The *numeric* type category includes integral and floating types.
- ❑ The *compile time* type category includes types **cfloat** and **cint**. These types are used by the compiler for constant type conversions.
- ❑ The *concrete* type category includes all types that are not included in the compile-time type category.
- ❑ The *scalar* type category includes all types in the numeric category, the **bool** type, and all types in the compile-time category. In this specification, a reference to a **<category>** type (such as a reference to a numeric type) means one of the types included in the category (such as **float**, **half**, or **fixed**).

Constants

A constant may be explicitly typed or implicitly typed. *Explicit typing* of a constant is performed, as in C, by suffixing the constant with a single character indicating the type of the constant:

- ❑ **f** for **float**
- ❑ **d** for **double**
- ❑ **h** for **half**
- ❑ **x** for **fixed**

Any constant that is not explicitly typed is *implicitly typed*. If the constant includes a decimal point, it is implicitly typed as **cfloat**. If it does not include a decimal point, it is implicitly typed as **cint**.

By default, constants are base 10. For compatibility with C, integer hexadecimal constants may be specified by prefixing the constant with **0x**, and integer octal constants may be specified by prefixing the constant with **0**.

Compile-time constant folding is preferably performed at the same precision that would be used if the operation were performed at run time. Some compilation profiles may allow some precision flexibility for the hardware; in such cases the compiler should ideally perform the constant folding at the highest hardware precision allowed for that data type in that profile.

If constant folding cannot be performed at run-time precision, it may optionally be performed using the precision indicated below for each of the numeric data types:

- ❑ **float**: s23e8 (**fp32**) IEEE single-precision floating point
- ❑ **half**: s10e5 (**fp16**) floating point with IEEE semantics
- ❑ **fixed**: s1.10 fixed point, clamping to [-2, 2)
- ❑ **double**: s52e11 (**fp64**) IEEE double-precision floating point
- ❑ **int**: signed 32-bit integer

Type Qualifiers

The type of an object may be qualified with one or more qualifiers. Qualifiers apply only to objects. Qualifiers are removed from the value of an object when used in an expression. The qualifiers are

❑ **const**

The value of a **const** qualified object cannot be changed after its initial assignment. The definition of a **const** qualified object that is not a parameter must contain an initializer. Named compile-time values are inherently qualified as **const**, but an explicit qualification is also allowed.

The value of a **static const** cannot be changed after compilation, and thus its value may be used in constant folding during compilation. A **uniform const**, on the other hand, is only **const** for a given execution of the program; its value may be changed via the runtime between executions.

❑ **in** and **out**

Formal parameters may be qualified as **in**, **out**, or both (by using **in out** or **inout**). By default, formal parameters are **in** qualified. An **in** qualified parameter is equivalent to a call-by-value parameter. An **out** qualified parameter is equivalent to a call-by-result parameter, and an

inout qualified parameter is equivalent to a value/result parameter. An **out** qualified parameter cannot be **const** qualified, nor may it have a default value.

Type Conversions

Some type conversions are allowed implicitly, while others require an cast. Some implicit conversions may cause a warning, which can be suppressed by using an explicit cast. Explicit casts are indicated using C-style syntax: casting **variable** to the **float4** type can be achieved using

(float4)variable.

❑ Scalar conversions

Implicit conversion of any scalar numeric type to any other scalar numeric type is allowed. A warning may be issued if the conversion is implicit and a loss of precision is possible. Implicit conversion of any scalar object type to any compatible scalar object type is allowed. Conversions between incompatible scalar object types or between object and numeric types are not allowed, even with an explicit cast. A **sampler** is compatible with **sampler1D**, **sampler2D**, **sampler3D**, **samplerCube**, and **samplerRECT**. No other object types are compatible—**sampler1D** is not comparable with **sampler2D**, even though both are compatible with **sampler**.

Scalar types may be implicitly converted to vectors and matrices of compatible type. The scalar is replicated to all elements of the vector or matrix. Scalar types may also be explicitly cast to structure types if the scalar type can be legally cast to every member of the structure.

❑ Vector conversions

Vectors may be converted to scalar types (the first element of the vector is selected). A warning is issued if this is done implicitly. A vector may also be implicitly converted to another vector of the same size and compatible element type.

A vector may be converted to a smaller compatible vector or a matrix of the same total size, but a warning is issued if an explicit cast is not used.

❑ Matrix conversions

Matrices may be converted to a scalar type—element (0,0) is selected. As with vectors, this causes a warning if it is done implicitly. A matrix may also be converted implicitly to a matrix of the same size and shape and compatible element type.

A matrix may be converted to a smaller matrix type (the upper-left submatrix is selected) or to a vector of the same total size, but a warning is issued if an explicit cast is not used.

❑ Structure conversions

A structure may be explicitly cast to the type of its first member or to another structure type with the same number of members, if each member of the **struct** can be converted to the corresponding member of the new **struct**. No implicit conversions of **struct** types are allowed.

❑ Array conversions

No conversions of array types are allowed.

[Table 9.](#) summarizes the type conversions discussed here. The table entries have the following meanings, but please pay attention to the footnotes:

❑ *Allowed*

- ❑ *Numeric* type when applied to expressions of numeric or compile-time type
- ❑ *Numeric vector* type when applied to another vector type of the same number of elements
- ❑ *Numeric matrix* type when applied to another matrix type of the same number of rows and columns

Type Equivalency

Type T1 is equivalent to type T2 if any of the following are true:

- ❑ T2 is equivalent to T1.
- ❑ T1 and T2 are the same scalar, vector, or structure type.
A packed array type is *not* equivalent to the same size unpacked array.
- ❑ T1 is a **typedef** name of T2.
- ❑ T1 and T2 are arrays of equivalent types with the same number of elements.
- ❑ The unqualified types of T1 and T2 are equivalent, and both types have the same qualifications.
- ❑ T1 and T2 are functions with equivalent return types, the same number of parameters, and all corresponding parameters are pair-wise equivalent.

Type-Promotion Rules

The **cfloat** and **cint** types behave like **float** and **int** types except for the usual arithmetic conversion behavior and function-overloading rules (see [“Function Overloading”](#) on page 240).

The *usual arithmetic conversions* for binary operators are defined as follows:

1. If either operand is **double**, the other is converted to **double**.
2. Otherwise, if either operand is **float**, the other operand is converted to **float**.
3. Otherwise, if either operand is **half**, the other operand is converted to **half**.
4. Otherwise, if either operand is **fixed**, the other operand is converted to **fixed**.

5. Otherwise, if either operand is **cfloat**, the other operand is converted to **cfloat**.
6. Otherwise, if either operand is **int**, the other operand is converted to **int**.
7. Otherwise, both operands have type **cint**.

Note that conversions happen prior to performing the operation.

Assignment

Assignment of an expression to an object or compile-time typed value converts the expression to the type of the object or value. The resulting value is then assigned to the object or value.

The value of the assignment

Arrays and Subscripting

Arrays are declared as in C, except that they may optionally be declared to be **packed**, as described under “Types” on page 229. Arrays in Cg are first-class types, so array parameters to functions and programs must be declared using array syntax, rather than pointer syntax. Likewise, assignment of an **array**-typed object implies an array copy rather than a pointer copy.

Arrays with size **[1]** may be declared but are considered a different type from the corresponding non-array type.

Because the language does not currently support pointers, the storage order of arrays is only visible when an application passes parameters to a vertex or fragment program. Therefore, the compiler is currently free to allocate temporary variables as it sees fit.

The declaration and use of arrays of arrays is in the same style as in C. That is, if the 2D array **A** is declared as

```
float A[4][4];
```

then, the following statements are true:

- ❑ The array is indexed as **A[*row*][*column*]**.
- ❑ The array can be built with a constructor using

```
A = { {A[0][0], A[0][1], A[0][2], A[0][3]},
      {A[1][0], A[1][1], A[1][2], A[1][3]},
      {A[2][0], A[2][1], A[2][2], A[2][3]},
      {A[3][0], A[3][1], A[3][2], A[3][3]} };
```

- ❑ **A[0]** is equivalent to **{A[0][0], A[0][1], A[0][2], A[0][3]}**.

Support must be provided for any **struct** containing arrays.

Minimum Array Requirements

Profiles are required to provide partial support for certain kinds of arrays. This partial support is designed to support vectors and matrices in all profiles. For vertex profiles, it is additionally designed to support arrays of light state (indexed by light number) passed as uniform parameters, and arrays of skinning matrices passed as uniform parameters.

Profiles must support subscripting, copying, and swizzling of vectors and matrices. However, subscripting with run-time computed indices is not required to be supported.

Vertex profiles must support the following operations for any non-packed array that is a uniform parameter to the program, or is an element of a

structure that is a uniform parameter to the program. This requirement also applies when the array is indirectly a uniform program parameter (that is, it and or the structure containing it has been passed via a chain of **in** function parameters). There are two operations that must be supported:

- ❑ Rvalue subscripting by a run-time computed value or a compile-time value
- ❑ Passing the entire array as a parameter to a function, where the corresponding formal function parameter is declared as **in**

The following operations are explicitly not required to be supported:

- ❑ Lvalue subscripting
- ❑ Copying
- ❑ Other operators, including multiply, add, compare, and so on

Note that when the array is rvalue subscripted, the result is an expression, and this expression is no longer considered to be a **uniform** program parameter. Therefore, if this expression is an array, its subsequent use must conform to the standard rules for array usage.

These rules are not limited to arrays of numeric types, and thus imply support for arrays of **struct**, arrays of matrices, and arrays of vectors when the array is a **uniform** program parameter. Maximum array sizes may be limited by the number of available registers or other resource limits, and compilers are permitted to issue error messages in these cases. However, profiles must support sizes of at least **float arr[8]**, **float4 arr[8]**, and **float4x4 arr[4][4]**.

Fragment profiles are not required to support any operations on arbitrarily sized arrays; only support for vectors and matrices is required.

Unsize Array

An *unsize array* may be declared by declaring an array with no length specified between the brackets: **float a[]**. The actual length of the array may then be set by the runtime before program execution. In program code, the length of any array can be queried using the syntax **a.length**, where **length** acts like an undeclared structure parameter that holds the actual length of the array at runtime.

Function Overloading

Multiple functions may be defined with the same name, as long as the definitions can be distinguished by unqualified parameter types and do not have an open-profile conflict (see [“Overloading of Functions by Profile” on page 226](#)).

Function-matching rules:

1. Add all visible functions with a matching name in the calling scope to the set of function candidates.
2. Eliminate functions whose profile conflicts with the current compilation profile.
3. Eliminate functions with the wrong number of formal parameters. If a candidate function has excess formal parameters, and each of the excess parameters has a default value, do not eliminate the function.
4. If the set is empty, fail.
5. For each actual parameter expression in sequence, perform the following:
 - a. If the type of the actual parameter matches the unqualified type of the corresponding formal parameter in any function in the set, remove all functions whose corresponding parameter does not match exactly.
 - b. If there is a defined promotion for the type of the actual parameter to the unqualified type of the formal parameter of any function, remove all functions for which this is not true from the set.
 - c. If there is a valid implicit cast that converts the type of the actual parameter to the unqualified type of the formal parameter of any function, remove all functions without this cast.
 - d. Fail.
6. Choose a function based on profile:
 - a. If there is at least one function with a profile that exactly matches the compilation profile, discard all functions that don't exactly match.
 - b. Otherwise, if there is at least one function with a wildcard profile that matches the compilation profile, determine the “most specific” matching wildcard profile in the candidate set. Discard all functions except those with this most specific wildcard profile. How “specific” a given wildcard profile name is relative to a particular profile is determined by the profile specification.

7. If the number of functions remaining in the set is not one, then fail.

Global Variables

Global variables are declared and used as in C. Uniform non-static variables may have a semantic associated with them. Uniform non-static variables may have their value set through the run-time API.

Use of Uninitialized Variables

It is incorrect for a program to use an uninitialized variable. However, the compiler is not obligated to detect such errors, even if it would be possible to do so by compile-time data-flow analysis. The value obtained from reading an uninitialized variable is undefined. This same rule applies to the implicit use of a variable that occurs when it is returned by a top-level function. In particular, if a top-level function returns a **struct**, and some element of that **struct** is never written, then the value of that element is undefined.

Note: Variables are not defined as being initialized to zero because this would result in a performance penalty in cases where the compiler is unable to determine if a variable is properly initialized by the programmer.

Preprocessor

Cg profiles must support the full ANSI C standard preprocessor capabilities: **#if**, **#define**, and so on. However, Cg profiles are not required to support macro-like **#define** or the use of **#include** directives.

Overview of Binding Semantics

In stream-processing architectures, data packets flow between different programmable units. On a GPU, for example, packets of vertex data flow from the application to the vertex program.

Because packets are produced by one program (the application, in this case), and consumed by another (the vertex program), there must be some method for defining the interface between the two. The approach used in Cg is to associate a binding semantic with each element of the packet. This is a *bind by name* approach. For example, an output with the binding semantic **F00** is fed to an input with the binding semantic **F00**. Profiles may allow the user to define arbitrary identifiers in this “semantic namespace,” or they may restrict

the allowed identifiers to a predefined set. Often, these predefined names correspond to the names of hardware registers or API resources.

In some cases, predefined names may control non-programmable parts of the hardware. For example, vertex programs normally compute a position that is fed to the rasterizer, and this position is stored in an output with the binding semantic **POSITION**.

For any profile, there are two namespaces for predefined binding semantics—the namespace used for **in** variables and the namespace used for **out** variables. The primary implication of having two namespaces is that the binding semantic cannot be used to implicitly specify whether a variable is **in** or **out**.

Binding Semantics

A binding semantic may be associated with an input to a top-level function in one of three ways:

- ❑ The binding semantic is specified in the formal parameter declaration for the function. The syntax for formal parameters to a function is

```
[const] [in | out | inout]
<type> <identifier> [ : <binding-semantic> ][= <initializer>]
```

- ❑ If the formal parameter is a **struct**, the binding semantic may be specified with an element of the **struct** when the **struct** is defined:

```
struct <struct-tag> {
  <type> <identifier>[ : <binding-semantic>];
  /*...*/ };
```

- ❑ If the input to the function is implicit (a non-static global variable that is read by the function), the binding semantic may be specified when the non-static global variable is declared:

```
<type> <identifier>[ : <binding-semantic> ][= <initializer>]
```

If the non-static global variable is a **struct**, the binding semantic may be specified when the **struct** is defined, as described in the second bullet above.

- ❑ A binding semantic may be associated with the output of a top-level function in a similar manner:

```
<type> <identifier> ( <parameter-list> ) [ : <binding-semantic> ]
{ <body> }
```

Another method available for specifying a semantic for an output value is to return a **struct** and to specify the binding semantic(s) with

elements of the **struct** when the **struct** is defined. In addition, if the output is a formal parameter, the binding semantic may be specified using the same approach used to specify binding semantics for inputs.

Aliasing of Semantics

Semantics must honor a copy-on-input and copy-on-output model. Thus, if the same input binding semantic is used for two different variables, those variables are initialized with the same value, but the variables are not aliased thereafter. Output aliasing is illegal, but implementations are not required to detect it. If the compiler does not issue an error on a program that aliases output binding semantics, the results are undefined.

Restrictions on Semantics Within a Structure

For a particular profile, it is illegal to mix input binding semantics and output binding semantics within a particular **struct**. That is, for a particular top-level function, a **struct** must be either input-only or output-only. Likewise, a **struct** must consist exclusively of uniform inputs or exclusively of non-uniform inputs. It is illegal to use binding semantics to mix the two within a single **struct**.

Additional Details for Binding Semantics

The following rules are somewhat redundant, but provide extra clarity:

- ❑ Semantics names are case-insensitive.
- ❑ Semantics attached to parameters to non-main functions are ignored.
- ❑ Input semantics may be aliased by multiple variables.
- ❑ Output semantics may not be aliased.

How Programs Receive and Return Data

A program is just a non-static function that has been designated as the main entry point at compilation time. The varying inputs to the program come from this top-level function's varying **in** parameters. The uniform inputs to the program come from the top-level function's uniform **in** parameters and from any non-static global variables that are referenced by the top-level function or by any functions that it calls. The output of the program comes from the return value of the function (which is always implicitly varying), and from any **out** parameters, which must also be varying.

Parameters to a program of type **sampler*** are implicitly **const**.

Statements

Statements are expressed just as in C, unless an exception is stated elsewhere in this document. Additionally,

- ❑ The **if**, **while**, and **for** statements require **bool** expressions in the appropriate places.
- ❑ Assignment is performed using **=**. The assignment operator returns a value, just as in C, so assignments may be chained.
- ❑ The new **discard** statement terminates execution of the program for the current data element—such as the current vertex or current fragment—and suppresses its output. Vertex profiles may choose to omit support for **discard**.

Minimum Requirements for **if**, **while**, and **for** Statements

The minimum requirements are as follows:

- ❑ All profiles should support **if**, but such support is not strictly required for older hardware.
- ❑ All profiles should support **for** and **while** loops if the number of loop iterations can be determined at compile time.

“Can be determined at compile time” is defined as follows:

The loop-iteration expressions can be evaluated at compile time by use of intra-procedural constant propagation and folding, where the variables through which constant values are propagated do not appear as lvalues within any kind of control statement (**if**, **for**, or **while**) or **?:** construct.

Profiles may choose to support more general constant propagation techniques, but such support is not required.

- ❑ Profiles may optionally support fully general **for** and **while** loops.

New Vector Operators

These new operators are defined for vector types:

- ❑ Vector construction operator: **<TypeID>(...)**
This operator builds a vector from multiple scalars or shorter vectors:

```
float4(scalar, scalar, scalar, scalar)
float4(float3, scalar)
```
- ❑ Matrix construction operator: **<TypeID>(...)**

This operator builds a matrix from multiple rows. Each row may be specified either as multiple scalars or as any combination of scalars and vectors with the appropriate size.

```
float3x3(1, 2, 3, 4, 5, 6, 7, 8, 9)
float3x3(float3, float3, float3)
float3x3(1, float2, float3, float3, 1, 1, 1)
```

□ Swizzle operator: (.)

```
a = b.xyz; // A swizzle operator example
```

- ↳ At least one swizzle character must follow the operator.
- ↳ There are two sets of swizzle characters and they may not be mixed. Set one is **xyzw** = **0123**, and set two is **rgba** = **0123**.
- ↳ The vector swizzle operator may only be applied to vectors or to scalars.
- ↳ Applying the vector swizzle operator to a scalar gives the same result as applying the operator to a vector of length one. Thus, **myscalar.xxx** and **float3(myscalar,myscalar,myscalar)** yield the same value.
- ↳ If only one swizzle character is specified, the result is a scalar, not a vector of length one. Therefore, the expression **b.y** returns a scalar.
- ↳ Care is required when swizzling a constant scalar because of ambiguity in the use of the decimal point character. For example, to create a three-vector from a scalar, use one of the following:
(1).xxx or **1..xxx** or **1.0.xxx** or **1.0f.xxx**
- ↳ The size of the returned vector is determined by the number of swizzle characters. Therefore, the size of the result may be larger or smaller than the size of the original vector.
For example, **float2(0,1).xxyy** and **float4(0,0,1,1)** yield the same result.

□ Matrix swizzle operator:

For any matrix type of the form **<type><rows>x<columns>**, the notation **<matrixObject>._m<row><col>[_m<row><col>][...]**

can be used to access individual matrix elements (in the case of only one **<row><col>** pair) or to construct vectors from elements of a matrix (in the case of more than one **<row><col>** pair). The row and column numbers are zero-based.

For example,

```
float4x4 myMatrix;
float    myFloatScalar;
float4    myFloatVec4;

// Set myFloatScalar to myMatrix[3][2].
myFloatScalar = myMatrix._m32;

// Assign the main diagonal of myMatrix to myFloatVec4.
myFloatVec4 = myMatrix._m00_m11_m22_m33;
```

- ✎ For compatibility with the **D3DMatrix** data type, Cg also allows one-based swizzles, using a form with the **m** omitted after the **_** symbol:

<matrixObject>._<row><col>[_<row><col>][...]

In this form, the indexes for **<row>** and **<col>** are one-based, rather than the C standard zero-based. So, the two forms are functionally equivalent:

```
float4x4 myMatrix;
float4    myVec;

// These two statements are functionally equivalent:
myVec = myMatrix._m00_m23_m11_m31;
myVec = myMatrix._11_34_22_42;
```

Because of the confusion that can be caused by the one-based indexing, use of the latter notation is strongly discouraged.

- ✎ The matrix swizzles may only be applied to matrices. When multiple components are extracted from a matrix using a swizzle, the result is an appropriately sized vector. When a swizzle is used to extract a single component from a matrix, the result is a scalar.
- The write-mask operator: **(.)**
It can only be applied to an lvalue that is a vector. It allows assignment to particular elements of a vector or matrix, leaving other elements unchanged. The only restriction is that a component cannot be repeated.

Arithmetic Precision and Range

Some hardware may not conform exactly to IEEE arithmetic rules. Fixed-point data types do not have IEEE-defined rules.

Optimizations are allowed to produce slightly different results than unoptimized code. Constant folding must be done with approximately the

correct precision and range, but is not required to produce bit-exact results. It is recommended that compilers provide an option either to forbid these optimizations or to guarantee that they are made in bit-exact fashion.

Operator Precedence

Cg uses the same operator precedence as C for operators that are common between the two languages.

The swizzle and write-mask operators (`.`) have the same precedence as the structure member operator (`.`) and the array index operator (`[]`).

Operator Enhancements

The standard C arithmetic operators (`+`, `-`, `*`, `/`, `%`, **unary-**) are extended to support vectors and matrices. Sizes of vectors and matrices must be appropriately matched, according to standard mathematical rules. Scalar-to-vector promotion (see [“Smearing of Scalars to Vectors” on page 237](#)) allows relaxation of these rules.

Table 10. Expanded Operators

Operator	Description
M [<i>n</i>][<i>m</i>]	Matrix with <i>n</i> rows and <i>m</i> columns
V [<i>n</i>]	Vector with <i>n</i> elements
- V [<i>n</i>] -> V [<i>n</i>]	Unary vector negate
- M [<i>n</i>] -> M [<i>n</i>]	Unary matrix negate
V [<i>n</i>] * V [<i>n</i>] -> V [<i>n</i>]	Componentwise *
V [<i>n</i>] / V [<i>n</i>] -> V [<i>n</i>]	Componentwise /
V [<i>n</i>] % V [<i>n</i>] -> V [<i>n</i>]	Componentwise %
V [<i>n</i>] + V [<i>n</i>] -> V [<i>n</i>]	Componentwise +
V [<i>n</i>] - V [<i>n</i>] -> V [<i>n</i>]	Componentwise -
M [<i>n</i>][<i>m</i>] * M [<i>n</i>][<i>m</i>] -> M [<i>n</i>][<i>m</i>]	Componentwise *
M [<i>n</i>][<i>m</i>] / M [<i>n</i>][<i>m</i>] -> M [<i>n</i>][<i>m</i>]	Componentwise /
M [<i>n</i>][<i>m</i>] % M [<i>n</i>][<i>m</i>] -> M [<i>n</i>][<i>m</i>]	Componentwise %
M [<i>n</i>][<i>m</i>] + M [<i>n</i>][<i>m</i>] -> M [<i>n</i>][<i>m</i>]	Componentwise +
M [<i>n</i>][<i>m</i>] - M [<i>n</i>][<i>m</i>] -> M [<i>n</i>][<i>m</i>]	Componentwise -

Operators

Boolean

&& || !

Boolean operators may be applied to **bool** packed **bool** vectors, in which case they are applied in elementwise fashion to produce a result vector of the same size. Each operand must be a **bool** vector of the same size.

Both sides of **&&** and **||** are always evaluated; there is no short-circuiting as there is in C.

Comparisons

< > <= >= != ==

Comparison operators may be applied to numeric vectors. Both operands must be vectors of the same size. The comparison operation is performed in elementwise fashion to produce a **bool** vector of the same size.

Comparison operators may also be applied to **bool** vectors. For the purpose of relational comparisons, **true** is treated as one and **false** is treated as zero. The comparison operation is performed in elementwise fashion to produce a **bool** vector of the same size.

Comparison operators may also be applied to numeric or **bool** scalars.

Arithmetic

+ - * / % ++ -- unary- unary+

The arithmetic operator **%** is the remainder operator, as in C. It may only be applied to two operands of **cint** or **int** type.

When **/** or **%** is used with **cint** or **int** operands, C rules for integer **/** and **%** apply.

The C operators that combine assignment with arithmetic operations (such as **+=**) are also supported when the corresponding arithmetic operator is supported by Cg.

Conditional Operator

?:

If the first operand is of type **bool**, one of the following statements must hold for the second and third operands:

- Both operands have compatible structure types.

- ❑ Both operands are scalars with numeric or **bool** type.
- ❑ Both operands are vectors with numeric or **bool** type, where the two vectors are of the same size, which is less than or equal to four.

If the first operand is a packed vector of **bool**, then the conditional selection is performed on an elementwise basis. Both the second and third operands must be numeric vectors of the same size as the first operand.

Unlike C, side effects in the expressions in the second and third operands are always executed, regardless of the condition.

Miscellaneous Operators

(typecast) ,

Cg supports C's typecast and comma operators.

Reserved Words

The following are the reserved words in Cg:

asm*	asm_fragment	auto
bool	break	case
catch	char	class
column major	compile	const
const_cast	continue	decl*
default	delete	discard
do	double	dword*
dynamic_cast	else	emit
enum	explicit	extern
false	fixed	float*
for	friend	get
goto	half	if
in	inline	inout
int	interface	long
matrix*	mutable	namespace
new	operator	out
packed	pass*	pixelfragment*
pixelshader*	private	protected
public	register	reinterpret_cast
return	row major	sampler
sampler_state	sampler1D	sampler2D
sampler3D	samplerCUBE	shared
short	signed	sizeof
static	static_cast	string*
struct	switch	technique*
template	texture*	texture1D

texture2D	texture3D	textureCUBE
textureRECT	this	throw
true	try	typedef
typeid	typename	uniform
union	unsigned	using
vector*	vertexfragment*	vertexshader*
virtual	void	volatile
while	__identifier (two underscores before identifier)	

Cg Standard Library Functions

Cg provides a set of built-in functions and predefined structures with binding semantics to simplify GPU programming. These functions are discussed in [“Cg Standard Library Functions” on page 33](#).

Vertex Program Profiles

A few features of the Cg language that are specific to vertex program profiles are required to be implemented in the same manner for all vertex program profiles.

Mandatory Computation of Position Output

Vertex program profiles may (and typically do) require that the program compute a position output. This homogeneous clip-space position is used by the hardware rasterizer and must be stored in a program output with an output binding semantic of **POSITION** (or **HPOS** for backward compatibility).

Position Invariance

In many graphics APIs, the user can choose between two different approaches to specifying per-vertex computations: use a built-in configurable *fixed function* pipeline or specify a user-written vertex program. If the user wishes to mix these two approaches, it is sometimes desirable to guarantee that the position computed by the first approach is bit-identical to the position computed by the second approach. This *position invariance* is particularly important for multipass rendering.

Support for position invariance is optional in Cg vertex profiles, but for those vertex profiles that support it, the following rules apply:

- ❑ Position invariance with respect to the fixed function pipeline is guaranteed if two conditions are met:

↳ The vertex program is compiled using a compiler option indicating position invariance (**-posinv**, for example).

↳ The vertex program computes position as follows:

OUT_POSITION = mul(MVP, IN_POSITION)

where

OUT_POSITION is a variable (or structure element) of type **float4** with an output binding semantic of **POSITION** or **HPOS**.

IN_POSITION is a variable (or structure element) of type **float4** with an input binding semantic of **POSITION**.

MVP is a uniform variable (or structure element) of type **float4x4** with an input binding semantic that causes it to track the fixed-function modelview-projection matrix. (The name of this binding semantic is currently profile-specific—for OpenGL profiles, the semantic **_GL_MVP** is recommended).

- ❑ If the first condition is met but not the second, the compiler is encouraged to issue a warning.
- ❑ Implementations may choose to recognize more general versions of the second condition (such as the variables being copy propagated from the original inputs and outputs), but this additional generality is not required.

Binding Semantics for Outputs

As shown in [Table 11.](#), there are two output binding semantics for vertex program profiles:

Table 11. Vertex Output Binding Semantics

Name	Meaning	Type	Default Value
POSITION	Homogeneous clip-space position; fed to rasterizer.	float4	Undefined
PSIZE	Point size	float	Undefined

Profiles may define additional output binding semantics with specific behaviors, and these definitions are expected to be consistent across commonly used profiles.

Fragment Program Profiles

A few features of the Cg language that are specific to fragment program profiles are required to be implemented in the same manner for all fragment program profiles.

Binding Semantics for Outputs

As shown in [Table 12.](#), there are three output binding semantics for fragment program profiles. Profiles may define additional output binding semantics with specific behaviors, and these definitions are expected to be consistent across commonly used profiles.

Table 12. Fragment Output Binding Semantics

Name	Meaning	Type	Default Value
COLOR	RGBA output color	float4	Undefined
COLOR0	Same as COLOR	—	—
DEPTH	Fragment depth value (in range [0,1])	float	Interpolated depth from rasterizer (in range [0,1])

If a program desires an output color alpha of 1.0, it should explicitly write a value of 1.0 to the **w** component of the **COLOR** output. The language does *not* define a default value for this output.

Note: If the target hardware uses a default value for this output, the compiler may choose to optimize away an explicit write specified by the user if it matches the default hardware value. Such defaults are not exposed in the language.

In contrast, the language does define a default value for the **DEPTH** output. This default value is the interpolated depth obtained from the rasterizer. Semantically, this default value is copied to the output at the beginning of the execution of the fragment program.

Note: Although the **DEPTH** output is assigned a default value, as with all outputs its value cannot be read in a Cg program.

As discussed earlier, when a binding semantic is applied to an output, the type of the output variable is not required to match the type of the binding semantic. For example, the following is legal, although not recommended:

```
struct myfragoutput {  
    float2 mycolor : COLOR;  
}
```

In such cases, the variable is implicitly copied (with a **typecast**) to the semantic upon program completion. If the variable's vector size is shorter than the semantic's vector size, the larger-numbered components of the semantic receive their default values, if applicable, and otherwise are undefined. In the case above, the **R** and **G** components of the output color are obtained from *mycolor*, while the **B** and **A** components of the color are undefined.

Appendix B Language Profiles

This appendix describes the language capabilities that are available in each of the following profiles supported by the Cg compiler:

- ❑ OpenGL ARB Vertex Program Profile (arbvp1)
- ❑ OpenGL ARB Fragment Program Profile (arbfvp1)
- ❑ OpenGL NV_vertex_program 3.0 Profile (vp40)
- ❑ OpenGL NV_fragment_program 2.0 Profile (fp40)
- ❑ OpenGL NV_vertex_program 2.0 Profile (vp30)
- ❑ OpenGL NV_fragment_program Profile (fp30)
- ❑ OpenGL NV_vertex_program 1.0 Profile (vp20)
- ❑ OpenGL NV_texture_shader and NV_register_combiners Profile (fp20)
- ❑ DirectX Vertex Shader 2.x Profiles (vs_2_*)
- ❑ DirectX Pixel Shader 2.x Profiles (ps_2_*)
- ❑ DirectX Vertex Shader 1.1 Profile (vs_1_1)
- ❑ DirectX Pixel Shader 1.x Profiles (ps_1_*)

In each case, the capabilities are a subset of the full capabilities described by the Cg language specification in “Cg Language Specification” on page 221.

OpenGL ARB Vertex Program Profile (**arbvp1**)

The OpenGL ARB Vertex Program Profile is used to compile Cg source code to vertex programs compatible with version 1.0 of the **GL_ARB_vertex_program** extension.

- ❑ **Profile name:** **arbvp1**
- ❑ **How to invoke:** Use the compiler option **-profile arbvp1**.

This section describes the capabilities and restrictions of Cg when using the **arbvp1** profile.

Overview

- ❑ The **arbvp1** profile is similar to the **vp20** profile except for the format of its output and its capability of accessing OpenGL state easily.
- ❑ **ARB_vertex_program** has the same capabilities as **NV_vertex_program** and DirectX 8 vertex shaders, so the limitations that this profile places on the Cg source code written by the programmer is the same as the **NV_vertex_program**¹ profile.

Accessing OpenGL State

The **arbvp1** profile allows Cg programs to refer to the OpenGL state directly, unlike the **vp20** profile. However, if you want to write Cg programs that are compatible with **vp20**, **vp30**, and **dx8vs** profiles, you should use the alternate mechanism of setting uniform variables with the necessary state using the Cg run time. The compiler relies on the feature of ARB vertex assembly programs that enables parts of the OpenGL state to be written automatically to program parameter registers as the state changes. The OpenGL driver handles this state-tracking feature.

A special variable semantic called **state** can be used to refer to every part of the OpenGL state that ARB vertex programs can reference. Following this paragraph are three lists of the **state** fields that can be accessed. The array indexes are shown as **0**, but an array can be accessed using any positive integer that is less than the limit of the array. For example, the diffuse component of the second light would be accessed by using the semantic

1. See “OpenGL NV_vertex_program 1.0 Profile (vp20)” on page 279 for a full explanation of the data types, statements, and operators supported by this profile.

`state.light[1].diffuse`, assuming that `GL_MAX_LIGHTS` is at least 2, as shown in the following code:

```
void main( uniform float4 lightColor : state.light[1].diffuse,
          ... )
```

The **state** semantics of type **float4x4** that can be accessed are in [Table 13](#).

Table 13. **float4x4 state** Semantics

<code>state.matrix.modelview[0]</code>	<code>state.matrix.projection</code>
<code>state.matrix.mvp</code>	<code>state.matrix.texture[0]</code>
<code>state.matrix.palette[0]</code>	<code>state.matrix.program[0]</code>
<code>state.matrix.inverse.modelview[0]</code>	<code>state.matrix.inverse.projection</code>
<code>state.matrix.inverse.mvp</code>	<code>state.matrix.inverse.texture[0]</code>
<code>state.matrix.inverse.palette[0]</code>	<code>state.matrix.inverse.program[0]</code>
<code>state.matrix.transpose.modelview[0]</code>	<code>state.matrix.transpose.projection</code>
<code>state.matrix.transpose.mvp</code>	<code>state.matrix.transpose.texture[0]</code>
<code>state.matrix.transpose.palette[0]</code>	<code>state.matrix.transpose.program[0]</code>
<code>state.matrix.invtrans.modelview[0]</code>	<code>state.matrix.invtrans.projection</code>
<code>state.matrix.invtrans.mvp</code>	<code>state.matrix.invtrans.texture[0]</code>
<code>state.matrix.invtrans.palette[0]</code>	<code>state.matrix.invtrans.program[0]</code>

Accessible **state** semantics of type **float4** are listed in [Table 14](#).

Table 14. **float4 state** Semantics

<code>state.material.ambient</code>	<code>state.material.diffuse</code>
<code>state.material.specular</code>	<code>state.material.emission</code>
<code>state.material.shininess</code>	<code>state.material.front.ambient</code>
<code>state.material.front.diffuse</code>	<code>state.material.front.specular</code>
<code>state.material.front.emission</code>	<code>state.material.front.shininess</code>
<code>state.material.back.ambient</code>	<code>state.material.back.diffuse</code>
<code>state.material.back.specular</code>	<code>state.material.back.emission</code>

Table 14. **float4 state** Semantics (continued)

<code>state.material.back.shininess</code>	<code>state.light[0].ambient</code>
<code>state.light[0].diffuse</code>	<code>state.light[0].specular</code>
<code>state.light[0].position</code>	<code>state.light[0].attenuation</code>
<code>state.light[0].spot.direction</code>	<code>state.light[0].half</code>
<code>state.lightmodel.ambient</code>	<code>state.lightmodel.scenecolor</code>
<code>state.lightmodel.front.scenecolor</code>	<code>state.lightmodel.back.scenecolor</code>
<code>state.lightprod[0].ambient</code>	<code>state.lightprod[0].diffuse</code>
<code>state.lightprod[0].specular</code>	<code>state.lightprod[0].front.ambient</code>
<code>state.lightprod[0].front.diffuse</code>	<code>state.lightprod[0].front.specular</code>
<code>state.lightprod[0].back.ambient</code>	<code>state.lightprod[0].back.diffuse</code>
<code>state.lightprod[0].back.specular</code>	<code>state.texgen[0].eye.s</code>
<code>state.texgen[0].eye.t</code>	<code>state.texgen[0].eye.r</code>
<code>state.texgen[0].eye.q</code>	<code>state.texgen[0].object.s</code>
<code>state.texgen[0].object.t</code>	<code>state.texgen[0].object.r</code>
<code>state.texgen[0].object.q</code>	<code>state.fog.color</code>
<code>state.fog.params</code>	<code>state.clip[0].plane</code>

The **state** semantics of type **float** that can be accessed are listed in [Table 15](#).

Table 15. **float state** Semantics

<code>state.point.size</code>	<code>state.point.attenuation</code>
-------------------------------	--------------------------------------

Position Invariance

- ❑ The **arbvp1** profile supports position invariance, as described in the core language specification.
- ❑ The modelview-projection matrix is not specified using a binding semantic of **_GL_MVP**.

Data Types

This profile implements data types as follows:

- ❑ **float** data type is implemented as defined in the **ARB_vertex_program** specification.
- ❑ **half** data type is implemented as float.
- ❑ **fixed** or **sampler*** data types are not supported, but the profile does provide the minimal partial support that is required for these data types by the core language specification—that is, it is legal to declare variables using these types as long as no operations are performed on the variables.

Compatibility with the **vp20** Vertex Program Profile

Programs that work with the **vp20** profile are compatible with the **arbvp1** profile as long as they use the Cg run time to manage all uniform parameters, including OpenGL state. That is, **arbvp1** and **vp20** profiles can be used interchangeably without changing the Cg source code or the application program except for specifying a different profile. However, if any of the **glProgramParameterxxNV()** routines are used the application program needs to be changed to use the corresponding ARB functions.

Since there is no ARB function corresponding to **glTrackMatrixNV()**, an application using **glTrackMatrixNV()** and the **arbvp1** profile needs to be modified. One solution is to change the Cg source code to refer to the matrix using the **state** structure so that the matrix is automatically tracked by the OpenGL driver as part of its **GL_ARB_vertex** support. Another solution is for the application to use the Cg run-time routine **cgGLSetStateMatrixParameter()** to load the appropriate matrix or matrices when necessary.

Another potential incompatibility between the **arbvp1** and **vp20** profiles is the way that input varying semantics are handled. In the **vp20** profile, semantic names such as **POSITION** and **ATTR0** are aliases of each other the same way **NV_vertex_program** aliases Vertex and Attribute 0 (see [Table 30, “vp20 Varying Input Binding Semantics,”](#) on page 281). In the **arbvp1** profile, the semantic names are not aliased because **ARB_vertex_program** allows the conventional attributes (such as vertex position) to be separate from the generic attributes (such as Attribute 0). For this reason it is important to follow the conventions given in [Table 17, “arbvp1 Varying Input Binding Semantics,”](#) on page 261 so that **arbvp1** programs work for all implementations of **ARB_vertex_program**. The **arbvp1** conventions are compatible with the **vp20** and **vp30** profiles.

Loading Constants

Applications that do not use the Cg run time are no longer required to load constant values into program parameters registers as indicated by the **#const** expressions in the Cg compiler output. The compiler produces output that causes the OpenGL driver to load them. However, uniform variables that have a default definition still require constant values to be loaded into the appropriate program parameter registers, as ARB vertex programs do not support this feature. Application programs either have to use the Cg run time, parse, and handle the **#default** commands, or have to avoid initializing uniform variables in the Cg source code.

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **arbvp1** profile are summarized in [Table 16](#).

Table 16. **arbvp1** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(c0)–register(c255) c0–c255	Local parameter with index <i>n</i> , <i>n</i> = [0..255]. The aliases c0–c255 (lowercase) are also accepted. If used with a variable that requires more than one constant register (for example, a matrix), the semantic specifies the first local parameter that is used.

Binding Semantics for Varying Input/Output Data

The valid binding semantics for uniform parameters in the **arbvp1** profile are summarized in [Table 17](#).

The set of binding semantics for varying input data to **arbvp1** consists of **POSITION**, **BLENDWEIGHT**, **NORMAL**, **COLOR0**, **COLOR1**, **TESSFACTOR**, **PSIZE**, **BLENDINDICES**, and **TEXCOORD0–TEXCOORD7**. One can also use **TANGENT** and **BINORMAL** instead of **TEXCOORD6** and **TEXCOORD7**. Additionally, a set of generic binding semantics of **ATTR0–ATTR15** can be used. In OpenGL implementations, conventional and generic vertex attributes may or may not be aliases for each other; see the **ARB_vertex_program** specification for more

details. The mapping of these semantics to corresponding setting command is listed in the table.

Table 17. **arbvp1** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION	Input Vertex, through Vertex command
BLENDWEIGHT	Input vertex weight through WeightARB , VertexWeightEXT command
NORMAL	Input normal through Normal command
COLOR0 , DIFFUSE	Input primary color through Color command
COLOR1 , SPECULAR	Input secondary color through SecondaryColorEXT command
FOGCOORD	Input fog coordinate through FogCoordEXT command
TEXCOORD0 - TEXCOORD7	Input texture coordinates (texcoord0 - texcoord7) through MultiTexCoord command
ATTR0 - ATTR15	Generic Attribute 0-15 through VertexAttrib command
PSIZE , ATTR6	Generic Attribute 6

The valid binding semantics for varying output parameters in the **arbvp1** profile are found in [Table 18](#). These binding semantics map to **ARB_vertex_program** output registers. The two sets act as aliases to each other.

Table 18. **arbvp1** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION , HPOS	Output position
PSIZE , PSIZ	Output point size
FOG , FOGC	Output fog coordinate
COLOR0 , COL0	Output primary color
COLOR1 , COL1	Output secondary color
BCOL0	Output backface primary color

Table 18. **arbvp1** Varying Output Binding Semantics (continued)

Binding Semantics Name	Corresponding Data
BCOL1	Output backface secondary color
TEXCOORD0-TEXCOORD7, TEX0-TEX7	Output texture coordinates

Note: The application must call **glEnable(GL_COLOR_SUM_ARB)** in order to enable **COLOR1** output when using the **arbvp1** profile.

The profile also allows **WPOS** to be present as binding semantics on a member of a structure of a varying output data structure, provided the member with this binding semantics is not referenced. This allows Cg programs to have the same structure specify the varying output of an **arbvp1** profile program and the varying input of an **fp30** profile program.

Options

The **arbvp1** profile supports the following profile-specific options:

NumTemps=<n>	(where $1 \leq n \leq 32$; default 32)
MaxAddressRegs=<n>	(where $1 \leq n \leq 8$; default 1)
MaxInstructions=<n>	(where $16 \leq n \leq 4096$; default 1024)
MaxLocalParams=<n>	(where $16 \leq n \leq 256$; default 96)

OpenGL ARB Fragment Program Profile (**arbfp1**)

The OpenGL ARB Fragment Program Profile is used to compile Cg source code to fragment programs compatible with version 1.0 of the **GL_ARB_fragment_program** OpenGL extension.²

- ❑ **Profile name:** **arbfp1**
- ❑ **How to invoke:** Use the compiler option **-profile arbfp1**.

The **arbfp1** profile limits Cg to match the capabilities of OpenGL ARB fragment programs. This section describes the capabilities and restrictions of Cg when using the **arbfp1** profile.

Accessing OpenGL State

The **arbfp1** profile supports access to OpenGL state with the same set of **state** semantics provided by the **arbvp1** profile. See [“Accessing OpenGL State” on page 256](#) for more information about this feature.

MRT Support

This profile supports multiple render targets (MRTs). When MRTs are used, up to three additional four-component outputs may be written in addition to the **COLOR** and **DEPTH** outputs supported in other profiles. These new outputs are available via the output semantics **COLOR1** through **COLOR3**.

The use of MRTs is an optional feature of the **ARB_fragment_program** and the DirectX PixelShader 2 specifications; consequently, not all hardware that supports these profiles supports MRTs. The **MaxDrawBuffers** profile option may be used to explicitly set the number of draw buffers (that is, render targets) available on the target hardware. If the input program requires more than the specified number of draw buffers, compilation fails.

If the **MaxDrawBuffers** profile option is not specified, the stand-alone Cg compiler, **cgc**, assumes that the target hardware supports MRTs to whatever extent required by the input program.

When compiling programs using the Cg runtime, be sure to call **cgGLSetOptimalOptions()** under OpenGL, or call **cgD3D9GetOptimalOptions()** under DirectX3D. These functions allow you to

2. To understand the capabilities of OpenGL ARB fragment programs and the code produced by the compiler, refer to the ARB fragment program extension in the OpenGL Extensions documentation.

automatically determine the value for the **MaxDrawBuffers** profile option that is appropriate for the graphics hardware on the target machine.

Resource Limits

The **ARB_fragment_profile** specifications allows an OpenGL implementation to place limits on the numbers and types of resources that a fragment program may use. If these resource limits must be exceeded to compile a Cg program, the compilation will fail. Resources that may be limited include the number of instructions, the number of registers, and the number of dependent texture reads.

The **arbfp1** profile supports a number of options that allow these limits to be specified on the compiler command line; see [“Options” on page 262](#) for details. These limits may also be values appropriate for the host computer's GPU, which are set using the **cgGLSetOptimalOptions()** Cg runtime call.

Language Constructs and Support

Data Types

This profile implements data types as follows:

- ❑ **float** data type is implemented as IEEE 32-bit single precision.
- ❑ **half**, **fixed**, and **double** data types are treated as float.
- ❑ **int** data type is supported using floating point operations.
- ❑ **sampler*** types are supported to specify sampler objects used for texture fetches.

Statements and Operators

With the ARB fragment program profiles **while**, **do**, and **for** statements are allowed only if the loops they define can be unrolled because there is no dynamic branching in ARB fragment program 1.

Comparison operators are allowed (**>**, **<**, **>=**, **<=**, **==**, **!=**) and Boolean operators (**|**, **&&**, **?:**) are allowed. However, the logic operators (**&**, **|**, **^**, **~**) are not.

Using Arrays and Structures

Variable indexing of arrays is not allowed. Array and structure data is not packed.

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **arbfp1** profile are found in [Table 19](#).

Table 19. **arbfp1** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(s0)–register(s15) TEXUNIT0-TEXUNIT15	Texunit image unit <i>N</i> , where <i>N</i> is in range [0..15] May only be used with uniform inputs with sampler* types.
register(c0)-register(c31) C0-C31	Local Parameter <i>N</i> , where <i>N</i> is in range [0..31] May only be used with uniform inputs.

Binding Semantics for Varying Input/Output Data

The valid binding semantics for varying input parameters in the **arbfp1** profile are summarized in [Table 20](#).

Table 20. **arbfp1** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data (type)
COLOR0	Input color 0 (float4)
COLOR1	Input color 1 (float4)
TEXCOORD0-TEXCOORD7	Input texture coordinates (float4)

The valid binding semantics for varying output parameters in the **arbfp1** profile are summarized in [Table 21](#).

Table 21. **arbfp1** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
COLOR, COLOR0	Output color (float4)
DEPTH	Output depth (float)

Options

The ARB fragment program profile allows the following profile specific options:

NumTemps=<<i>n</i>>	(where $0 \leq n \leq 32$; default 32)
NumInstructionSlots=<<i>n</i>>	(where $n \geq 0$; default 1024)
NumMathInstructionSlots=<<i>n</i>>	(where $n \geq 0$; default 1024)
NoDependentReadLimit=<<i>b</i>>	(where $b = 0$ or 1 ; default 1)
NumTexInstructionSlots=<<i>n</i>>	(where $n \geq 0$; default 1024)
MaxTexIndirections=<<i>n</i>>	(where $n \geq 1$; default infinite)
NumDrawBuffers=<<i>n</i>>	(where $1 \leq n \leq 4$; default 1)

OpenGL NV_vertex_program 3.0 Profile (**vp40**)

The **vp40** profile is an extended version of the **arbvp1** profile. It has all of the capabilities of **arbvp1** and the added capability described in this section.

Vertex Texturing

The **vp40** profile supports accessing texture maps in programs. Textures are available via the usual **sampler*** types and the **tex*()** standard library calls.

OpenGL NV_fragment_program 2.0 Profile (**fp40**)

The **fp40** profile is an extended version of the **arbpfp1** profile. It has all of the capabilities of **arbpfp1** as well as the added capabilities described in this section.

Branching

The branching support in **fp40** allows some **if** statements and looping constructs to be implemented with branching. In profiles such as **fp30**, conditional execution of code was always implemented with predicated instructions, and loops were always unrolled.

In the GeForce 6800 GPU, there is a cost associated with executing a branch in the fragment shading engine. As such, it is possible that the cost of the branch will out-weigh the savings from skipping over a block of conditionally executed code or of executing an unrolled loop. (Please refer to the NVIDIA developer Web site for more information about the performance of this and other NVIDIA GPUs.) The **fp40** profile, therefore, provides two options to control whether the compiler should emit branches or conditionally executed code for the **if** statements and loops within Cg shaders. The options are described in [Table 22](#).

Setting both **-ifcvt** and **-unroll** to **all** yields behavior similar to the **fp30** profile, for which branch instructions are not available. Using **-ifcvt=none** places the burden on the Cg fragment program author to use **if** statements where they want true branches and to use conditional expressions otherwise.

FACE Semantic

The **FACE** semantic can be applied to a varying parameter to a program. The value of such a parameter has a value less than zero if the fragment

OpenGL NV_vertex_program 2.0 Profile (**vp30**)

The **vp30** Vertex Program profile is used to compile Cg source code to vertex programs for use by the **NV_vertex_program2** OpenGL extension.

- ❑ **Profile name:** **vp30**
- ❑ **How to invoke:** Use the compiler option **-profile vp30**.

The **vp30** profile limits Cg to match the capabilities of the **NV_vertex_program2** extension. This section describes the capabilities and restrictions of Cg when using the **vp30** profile.

Position Invariance

Under **vp30**, unlike other profiles, the following points can be made:

- ❑ The **-posinv** option won't cause an **OPTION** driver directive to be added to the assembly code header (see the OpenGL specification for more details on this directive).
- ❑ The instructions for transforming the position using the modelview-projection matrix are emitted.

They are true because the final assembly code itself guarantees that the position calculation is invariant compared to the fixed pipeline calculation.

Language Constructs

Data Types

This profile implements data types as follows:

- ❑ **float** data type is implemented as IEEE 32-bit single precision.
- ❑ **half** data type is implemented as **float**.
- ❑ **int** data type is supported using floating point operations, which adds extra instructions for proper truncation for divides, modulus, and casts from floating point types.
- ❑ **fixed** or **sampler*** data types are not supported, but the profile does provide the minimal partial support that is required for these data types by the core language specification—that is, it is legal to declare variables using these types, as long as no operations are performed on the variables.

Statements and Operators

This profile is a superset of the **vp20** profile. Any program that compiles for the **vp20** profile should also compile for the **vp30** profile, although the converse is not true.

The additional capabilities of the **vp30** profile, beyond those of **vp20** are

- ❑ **for**, **while**, and **do** loops are supported without requiring loop unrolling
- ❑ Full support for **if/else** allowing non-constant conditional expressions

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **vp30** profile are summarized in [Table 23](#).

Table 23. **vp30** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(c0)–register(c255) c0–c255	Constant register [0..255]. The aliases c0–c255 (lowercase) are also accepted. If used with a variable that requires more than one constant register (for example, a matrix), the semantic specifies the first register that is used.

Binding Semantics for Varying Input/Output Data

The valid binding semantics for varying input parameters in the **vp30** profile are summarized in [Table 24](#).

One can also use **TANGENT** and **BINORMAL** instead of **TEXCOORD6** and **TEXCOORD7**. These binding semantics map to **NV_vertex_program2** input attribute parameters. The two sets act as aliases to each other.

Table 24. **vp30** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION, ATTR0	Input Vertex, Generic Attribute 0
BLENDWEIGHT, ATTR1	Input vertex weight, Generic Attribute 1
NORMAL, ATTR2	Input normal, Generic Attribute 2
COLOR0, DIFFUSE, ATTR3	Input primary color, Generic Attribute 3
COLOR1, SPECULAR, ATTR4	Input secondary color, Generic Attribute 4
TESSFACTOR, FOGCOORD, ATTR5	Input fog coordinate, Generic Attribute 5
PSIZE, ATTR6	Input point size, Generic Attribute 6
BLENDINDICES, ATTR7	Generic Attribute 7
TEXCOORD0-TEXCOORD7, ATTR8-ATTR15	Input texture coordinates (texcoord0-texcoord7), Generic Attributes 8–15
TANGENT, ATTR14	Generic Attribute 14
BINORMAL, ATTR15	Generic Attribute 15

The valid binding semantics for varying output parameters in the **vp30** profile are summarized in [Table 25](#).

These binding semantics map to **NV_vertex_program2** output registers. The two sets act as aliases to each other.

Table 25. **vp30** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION, HPOS	Output position
PSIZE, PSIZ	Output point size

Table 25. **vp30** Varying Output Binding Semantics (continued)

Binding Semantics Name	Corresponding Data
FOG, FOGC	Output fog coordinate
COLOR0, COL0	Output primary color
COLOR1, COL1	Output secondary color
BCOL0	Output backface primary color
BCOL1	Output backface secondary color
TEXCOORD0 - TEXCOORD7, TEX0 - TEX7	Output texture coordinates
CLP0 - CL5	Output Clip distances

The profile allows **WPOS** to be present as binding semantics on a member of a structure of a varying output data structure, provided the member with this binding semantics is not referenced. This allows Cg programs to have same structure specify the varying output of a **vp30** profile program and the varying input of an **fp30** profile program.

OpenGL NV_fragment_program Profile (**fp30**)

The **fp30** Fragment Program Profile is used to compile Cg source code to fragment programs for use by the **NV_fragment_program** OpenGL extension.

- ❑ **Profile name:** **fp30**
- ❑ **How to invoke:** Use the compiler option **-profile fp30**.

This section describes the capabilities and restrictions of Cg when using the **fp30** profile.

Language Constructs and Support

Data Types

- ❑ **fixed** type (s1.10 fixed point) is supported
- ❑ **half** type (s10e5 floating-point) is supported

It is recommended that you use **fixed**, **half**, and **float** in that order for maximum performance. Reversing this order provides maximum precision. You are encouraged to use the fastest type that meets your needs for precision.

Statements and Operators

- ❑ Full support for **if/else**
- ❑ No **for** and **while** loops, unless they can be unrolled by the compiler
- ❑ Support for flexible texture mapping
- ❑ Support for screen-space derivative functions
- ❑ No support for variable indexing of arrays

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **fp30** profile are summarized in [Table 26](#).

Table 26. **fp30** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(s0)-register(s15) TEXUNIT0-TEXUNIT15	Texunit N , where N is in the range [0..15]. May be used only with uniform inputs with sampler* types.
register(c0)-register(c31) C0-C31	Constant register N , where N is in range [0..15] May only be used with uniform inputs.

Binding Semantics for Varying Input/Output Data

The valid binding semantics for varying input parameters in the **fp30** profile are summarized in [Table 27](#).

These binding semantics map to **NV_fragment_program** input registers. The two sets act as aliases to each other. The profile also allows **POSITION**, **FOG**, **PSIZE**, **HPOS**, **FOGC**, **PSIZ**, **BCOL0**, **BCOL1**, and **CLP0-CLP5** to be present as binding semantics on a member of a structure of a varying input data structure, provided the member with this binding semantics is not referenced. This allows Cg programs to have the same structure specify the varying output of a **vp30** profile program and the varying input of an **fp30** profile program.

Table 27. **fp30** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data (type)
COLOR0 , COL0	Input color0 (float4)
COLOR1 , COL1	Input color1 (float4)
TEXCOORD0-TEXCOORD7 , TEX0-TEX7	Input texture coordinates (float4)
WPOS	Window Position Coordinates (float4)

The valid binding semantics for varying output parameters in the **fp30** profile are summarized in [Table 28](#).

Table 28. **fp30** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
COLOR , COLOR0 , COL	Output color (float4)
DEPTH , DEPR	Output depth (float)

Pack and Unpack Functions

The **fp30** profile provides a number of functions for packing multiple floating point values into a single 32-bit result. Corresponding unpacking functions are also provided. These functions map directly to the packing and unpacking instructions defined by the **NV_fragment_program** OpenGL extension.

pack_2half()

```
float pack_2half(float2 a);
float pack_2half(half2 a);
```

Converts the components of **a** into a pair of 16-bit floating point values. The two converted components are then packed into a single 32-bit result. This operation can be reversed using the **unpack_2half()** function.

```
// C Pseudocode
result = (((half)a.y) << 16) | (half)a.x;
```

unpack_2half()

```
half2 unpack_2half(float a);
```

Unpacks a 32-bit value into two 16-bit floating point values.

```
// C Pseudocode
result.x = (a >> 0) & 0xFF;
result.y = (a >> 16) & 0xFF;
```

pack_2ushort()

```
float pack_2ushort(float2 a);
float pack_2ushort(half2 a);
```

Converts the components of **a** into a pair of 16-bit unsigned integers. The two converted components are then packed into a single 32-bit return value. This operation can be reversed using the **unpack_2ushort()** function.

```
// C Pseudocode
ushort.x = round(65535.0 * clamp(a.x, 0.0, 1.0));
ushort.y = round(65535.0 * clamp(a.y, 0.0, 1.0));
result = (ushort.y << 16) | ushort.y;
```

unpack_2ushort()

```
float2 unpack_2ushort(float a);
```

Unpacks two 16-bit unsigned integer values from **a** and scales the results into individual floating point values between 0.0 and 1.0.

```
// C Pseudocode
result.x = ((x >> 0) & 0xFFFF) / 65535.0;
result.y = ((x >> 16) & 0xFFFF) / 65535.0;
```

pack_4byte()

```
float pack_4byte(float4 a);
float pack_4byte(half4 a);
```

Converts the four components of **a** into 8-bit signed integers. The signed integers are such that a representation with all bits set to 0 corresponds to the value $-(128/127)$, and a representation with all bits set to 1 corresponds to $+(127/127)$. The four signed integers are then packed into a single 32-bit result. This operation may be reversed using the **unpack_4byte()** function.

```
// C Pseudocode
ub.x = round(127 * clamp(a.x, -128/127, 127/127) + 128);
ub.y = round(127 * clamp(a.y, -128/127, 127/127) + 128);
ub.z = round(127 * clamp(a.z, -128/127, 127/127) + 128);
ub.w = round(127 * clamp(a.w, -128/127, 127/127) + 128);
result = (ub.w << 24) | (ub.z << 16) | (ub.y << 8) | ub.x;
```

unpack_4byte()

```
half4 unpack_4byte(float a);
```

Unpacks four 8-bit integers from **a** and scales the results into individual 16-bit floating point values between $-(128/127)$ and $+(127/127)$.

```
// C Pseudocode
result.x = (((a >> 0) & 0xFF) - 128) / 127.0;
result.y = (((a >> 8) & 0xFF) - 128) / 127.0;
result.z = (((a >> 16) & 0xFF) - 128) / 127.0;
result.w = (((a >> 24) & 0xFF) - 128) / 127.0;
```

pack_4byte()

```
float pack_4byte(float4 a);  
float pack_4byte(half4 a);
```

Converts the four components of **a** into 8-bit unsigned integers. The unsigned integers are such that a representation with all bits set to 0 corresponds to 0.0, and a representation with all bits set to 1 corresponds to 1.0. The four unsigned integers are then packed into a single 32-bit result. This operation can be reversed using the **unpack_4byte()** function.

```
// C Psuedocode
ub.x = round(255.0 * clamp(a.x, 0.0, 1.0));
ub.y = round(255.0 * clamp(a.y, 0.0, 1.0));
ub.z = round(255.0 * clamp(a.z, 0.0, 1.0));
ub.w = round(255.0 * clamp(a.w, 0.0, 1.0));
result = (ub.w << 24) | (ub.z << 16) | (ub.y << 8) | ub.x;
```

unpack_4ubyte()

```
half4 unpack_4ubyte(float a);
```

Unpacks the four 8-bit integers in **a** and scales the results into individual 16-bit floating point values between 0.0 and 1.0.

```
// C Pseudocode
result.x = ((a >> 0) & 0xFF) / 255.0;
result.y = ((a >> 8) & 0xFF) / 255.0;
result.z = ((a >> 16) & 0xFF) / 255.0;
result.w = ((a >> 24) & 0xFF) / 255.0;
```

OpenGL NV_vertex_program 1.0 Profile (**vp20**)

The **vp20** Vertex Program profile is used to compile Cg source code to vertex programs for use by the **NV_vertex_program** OpenGL extension³.

- ❑ **Profile name:** **vp20**
- ❑ **How to invoke:** Use the compiler option **-profile vp20**.

This section describes the capabilities and restrictions of Cg when using the **vp20** profile.

Overview

The **vp20** profile limits Cg to match the capabilities of the **NV_vertex_program** extension. **NV_vertex_program** has the same capabilities as DirectX 8 vertex shaders, so the limitations that this profile places on the Cg source code written by the programmer is the same as the DirectX VS 1.1 shader profile⁴.

Aside from the syntax of the compiler output, the only difference between the **vp20** Vertex Shader profile and the DirectX VS 1.1 profile is that the **vp20** profile supports two additional outputs: **BCOL0** (for back-facing primary color) and **BCOL1** (for back-facing secondary color).

Position Invariance

- ❑ The **vp20** profile supports position invariance, as described in the core language specification.
- ❑ The modelview-projection matrix must be specified using a binding semantic of **_GL_MVP**.

Data Types

This profile implements data types as follows:

- ❑ **float** data types are implemented as IEEE 32-bit single precision.
- ❑ **half** and **double** data types are implemented as **float**.

3. To understand the **NV_vertex_program** and the code produced by the compiler using the **vp20** profile, see the **GL_NV_vertex_program** extension documentation.

4. See “[OpenGL NV_vertex_program 1.0 Profile \(vp20\)](#)” on page 279 for a full explanation of the data types, statements, and operators supported by this profile.

- ❑ **int** data type is supported using floating point operations, which add extra instructions for proper truncation for divides, modulus, and casts from floating point types.
- ❑ **fixed** or **sampler*** data types are not supported, but the profile does provide the minimal partial support that is required for these data types by the core language specification—that is, it is legal to declare variables using these types, as long as no operations are performed on the variables.

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **vp20** profile are summarized in [Table 29](#).

Table 29. **vp20** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(c0)–register(c95) C0–C95	Constant register [0..95]. The aliases c0–c95 (lowercase) are also accepted. If used with a variable that requires more than one constant register (for example, a matrix), the semantic specifies the first register that is used.

Binding Semantics for Varying Input/Output Data

The valid binding semantics for varying input parameters in the **vp20** profile are summarized in [Table 30](#).

One can also use **TANGENT** and **BINORMAL** instead of **TEXCOORD6** and **TEXCOORD7**. A second set of binding semantics, **ATTR0-ATTR15**, can also be used. The two sets act as aliases to each other.

Table 30. **vp20** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION, ATTR0	Input Vertex, Generic Attribute 0
BLENDWEIGHT, ATTR1	Input vertex weight, Generic Attribute 1
NORMAL, ATTR2	Input normal, Generic Attribute 2
COLOR0, DIFFUSE, ATTR3	Input primary color, Generic Attribute 3
COLOR1, SPECULAR, ATTR4	Input secondary color, Generic Attribute 4
TESSFACTOR, FOGCOORD, ATTR5	Input fog coordinate, Generic Attribute 5
PSIZE, ATTR6	Input point size, Generic Attribute 6
BLENDINDICES, ATTR7	Generic Attribute 7
TEXCOORD0-TEXCOORD7, ATTR8-ATTR15	Input texture coordinates (texcoord0-texcoord7), Generic Attributes 8-15
TANGENT, ATTR14	Generic Attribute 14
BINORMAL, ATTR15	Generic Attribute 15

The valid binding semantics for varying output parameters in the **vp20** profile are summarized in [Table 31](#).

These binding semantics map to **NV_vertex_program** output registers. The two sets act as aliases to each other.

Table 31. **vp20** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION, HPOS	Output position
PSIZE, PSIZ	Output point size
FOG, FOGC	Output fog coordinate

Table 31. **vp20** Varying Output Binding Semantics (continued)

Binding Semantics Name	Corresponding Data
COLOR0 , COL0	Output primary color
COLOR1 , COL1	Output secondary color
BCOL0	Output backface primary color
BCOL1	Output backface secondary color
TEXCOORD0 - TEXCOORD3 , TEX0 - TEX3	Output texture coordinates

The profile also allows **WPOS** to be present as binding semantics on a member of a structure of a varying output data structure, provided the member with this binding semantics is not referenced. This allows Cg programs to have the same structure specify the varying output of a **vp20** profile program and the varying input of an **fp30** profile program.

OpenGL NV_texture_shader and NV_register_combiners Profile (**fp20**)

The OpenGL NV_texture_shader and NV_register_combiners profile is used to compile Cg source code to the **nvparse** text format for the NV_texture_shader and NV_register_combiners family of OpenGL extensions⁵.

- ❑ **Profile name:** **fp20**
- ❑ **How to invoke:** Use the compiler option **-profile fp20**.

This document describes the capabilities and restrictions of Cg when using the **fp20** profile.

Overview

Operations in the **fp20** profile can be categorized as texture shader operations and arithmetic operations. Texture shader operations are operations which generate texture shader instructions, arithmetic operations are operations which generate register combiners instructions.

The underlying instruction set and machine architecture limit programmability in this profile compared to what is allowed by Cg constructs. Thus, this profile places additional restrictions on what can and cannot be done in a Cg program.

Restrictions

A Cg program in one of these profiles is limited to generating a maximum of four texture shader instructions and eight register combiner instructions. Since these numbers are quite small, users need to be very aware of this limitation while writing Cg code for these profiles.

The **fp20** profile also restricts when a texture shader operation or arithmetic operation can occur in the program. A texture shader operation may not have any dependency on the output of an arithmetic operation unless

- ❑ the arithmetic operation is a valid input modifier for the texture shader operation

5. For more details about the underlying instruction sets, their capabilities, and their limitations, please refer to the NV_texture_shader and NV_register_combiners extensions in the OpenGL Extensions documentation.

- the arithmetic operation is part of a complex texture shader operation (which are summarized in the section [“Auxiliary Texture Functions” on page 290](#))

Modifiers

There are certain simple arithmetic operations that can be applied to inputs of texture shader operations and to inputs and outputs of arithmetic operations without generating a register combiner instruction. These operations are referred to as input modifiers and output modifiers.

Instead of generating a register combiners instruction, the arithmetic operation modifies the assembly instruction or source registers to which it is applied. For example, the following Cg expression

$$z = (x - 0.5 + y) / 2$$

could generate the following register combiner instruction (assuming **x** is in **tex0**, **y** is in **tex1**, and **z** is in **col0**)

```
rgb
{
    discard = half_bias(tex0.rgb);
    discard = tex1.rgb;
    col0 = sum();
    scale_by_one_half();
}
alpha
{
    discard = half_bias(tex0.a);
    discard = tex1.a;
    col0 = sum();
    scale_by_one_half();
}
```

How different NV_texture_shader and NV_register_combiners instruction set modifiers are expressed in Cg programs are summarized in [Table 32](#). For more details on the context in which each modifier is allowed and ways in which modifiers may be combined refer to the NV_texture_shader and NV_register_combiners documentation.

Table 32. NV_texture_shader and NV_register_combiners Instruction Set Modifiers

Instruction/Register Modifier	Cg Expression
scale_by_two()	$2 * x$
scale_by_four()	$4 * x$
scale_by_one_half()	$x / 2$
bias_by_negative_one_half()	$x - 0.5$
bias_by_negative_one_half_scale_by_two()	$2 * (x - 0.5)$
unsigned(reg)	saturate(x) (i.e. $\min(1, \max(0, x))$)
unsigned_invert(reg)	$1 - \text{saturate}(x)$
half_bias(reg)	$x - 0.5$
- reg	$-x$
expand(reg)	$2 * (x - 0.5)$

Language Constructs and Support

Data Types

In the **fp20** profile, operations occur on signed clamped floating-point values in the range -1 to 1. These profiles allow all data types to be used, but all operations are carried out in the above range. Refer to the NV_texture_shader and NV_register_combiners documentation for more details.

Statements and Operators

The **fp20** profile supports all of the Cg language constructs, with the following exceptions:

- ❑ Arbitrary swizzles are not supported (though arbitrary write masks are). Only the following swizzles are allowed
 - `.x/.r .y/.g .z/.b .w/.a`
 - `.xy/.rg .xyz/.rgb .xyzw/.rgba`
 - `.xxx/.rrr .yyy/.ggg .zzz/.bbb .www/.aaa`
 - `.xxxx/.rrrr .yyyy/.gggg .zzzz/.bbbb .www/.aaaa`

- ❑ Matrix swizzles are not supported.
- ❑ Boolean operators other than `<`, `<=`, `>` and `>=` are not supported. Furthermore, `<`, `<=`, `>` and `>=` are only supported as the condition in the `?:` operator.
- ❑ Bitwise integer operators are not supported.
- ❑ `/` is not supported unless the divisor is a non-zero constant or it is used to compute the depth output.
- ❑ `%` is not supported.
- ❑ Ternary `?:` is supported if the boolean test expression is a compile-time boolean constant, a uniform scalar boolean or a scalar comparison to a constant value in the range `[-0.5, 1.0]` (for example, `a > 0.5 ? b : c`).
- ❑ **do**, **for**, and **while** loops are supported only when they can be completely unrolled.
- ❑ arrays, vectors, and matrices may be indexed only by compile-time constant values or index variables in loops that can be completely unrolled.
- ❑ The **discard** statement is not supported. The similar but less general **clip()** function is supported.
- ❑ The use of an **allocation-rule-identifier** for an input or output **struct** is optional.

Standard Library Functions

Because the **fp20** profile has limited capabilities, not all of the Cg standard library functions are supported.

The Cg standard library functions that are supported by this profile are presented in [Table 33](#). See the standard library documentation for descriptions of these functions.

Table 33. Supported Standard Library Functions

dot(floatN, floatN)
lerp(floatN, floatN, floatN)
lerp(floatN, floatN, float)
tex1D(sampler1D, float)
tex1D(sampler1D, float2)

Table 33. Supported Standard Library Functions (continued)

<code>tex1Dproj(sampler1D, float2)</code>
<code>tex1Dproj(sampler1D, float3)</code>
<code>tex2D(sampler2D, float2)</code>
<code>tex2D(sampler2D, float3)</code>
<code>tex2Dproj(sampler2D, float3)</code>
<code>tex2Dproj(sampler2D, float4)</code>
<code>texRECT(samplerRECT, float2)</code>
<code>texRECT(samplerRECT, float3)</code>
<code>texRECTproj(samplerRECT, float3)</code>
<code>texRECTproj(samplerRECT, float4)</code>
<code>tex3D(sampler3D, float3)</code>
<code>tex3Dproj(sampler3D, float4)</code>
<code>texCUBE(samplerCUBE, float3)</code>
<code>texCUBEproj(samplerCUBE, float4)</code>

Note: The nonprojective texture lookup functions are actually done as projective lookups on the underlying hardware. Because of this, the **w** component of the texture coordinates passed to these functions from the application or vertex program must contain the value 1.

Texture coordinate parameters for projective texture lookup functions must have swizzles that match the swizzle done by the generated texture shader instruction. While this may seem burdensome, it is intended to allow **fp20** profile programs to behave correctly under other pixel shader profiles.

The swizzles required on the texture coordinate parameter to the projective texture lookup functions are listed in [Table 34](#).

Table 34. Required Projective Texture Lookup Swizzles

Texture Lookup Function	Texture Coordinate Swizzle
tex1Dproj	.xw/.ra
tex2Dproj	.xyw/.rga
texRECTproj	.xyw/.rga
tex3Dproj	.xyzw/.rgba
texCUBEproj	.xyzw/.rgba

Bindings

Manual Assignment of Bindings

The Cg compiler can determine bindings between texture units and uniform sampler parameters/texture coordinate inputs automatically. This automatic assignment is based on the context in which uniform sampler parameters and texture coordinate inputs are used together.

To specify bindings between texture units and uniform parameters/texture coordinates to match their application, all sampler uniform parameters and texture coordinate inputs that are used in the program must have matching binding semantics—for example, **TEXUNIT<n>** may only be used with **TEXCOORD<n>**. Partially specified binding semantics may not work in all cases. Fundamentally, this restriction is due to the close coupling between texture samplers and texture coordinates in the NV_texture_shader extension.

Binding Semantics for Uniform Data

If a binding semantic for a uniform parameter is not specified, then the compiler will allocate one automatically. Scalar uniform parameters may be allocated to either the **xyz** or the **w** portion of a constant register depending on how they are used within the Cg program. When using the output of the compiler without the Cg runtime, you must set all values of a scalar uniform to the desired scalar value, not just the **x** component.

The valid binding semantics for uniform parameters in the **fp20** profile are summarized in [Table 35](#).

Table 35. **fp20** Uniform Binding Semantics

Binding Semantics Name	Corresponding Data
register(s0)–register(s3) TEXTUNIT0–TEXTUNIT3	Texture unit N , where N is in range [0..3]. May be used only with uniform inputs with sampler* types.

The **ps_1_X** profiles allow the programmer to decide which constant register a uniform variable will reside in by specifying the **C<n>/register(c<n>)** binding semantic. This is not allowed in the **fp20** profile since the **NV_register_combiners** extension does not have a single bank of constant registers. While the **NV_register_combiners** extension does describe constant registers, these constant registers are per-combiner stage and specifying bindings to them in the program would overly constrain the compiler.

Binding Semantics for Varying Input/Output Data

The varying input binding semantics in the **fp20** profile are the same as the varying output binding semantics of the **vp20** profile.

Varying input binding semantics in the **fp20** profile consist of **COLOR0**, **COLOR1**, **TEXCOORD0**, **TEXCOORD1**, **TEXCOORD2** and **TEXCOORD3**. These map to output registers in vertex shaders.

The valid binding semantics for varying input parameters in the **fp20** profile are summarized in [Table 36](#).

Table 36. **fp20** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data
COLOR , COLOR0 COL , COL0	Input color value v0
COLOR1 COL1	Input color value v1
TEXCOORD0–TEXCOORD3 TEX0–TEX3	Input texture coordinates t0–t3
FOGP FOG	Input fog color and factor

Additionally, the **fp20** profile allows **POSITION**, **PSIZE**, **TEXCOORD4**, **TEXCOORD5**, **TEXCOORD6**, and **TEXCOORD7** to be specified on varying inputs, provided these inputs are not referenced. This allows Cg programs to have the same structure specify the varying output of a **vp20** profile program and the varying input of a **fp20** profile program.

The valid binding semantics for varying output parameters in the **fp20** profile are summarized in [Table 37](#).

Table 37. **fp20** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
COLOR , COLOR0 COL , COL0	Output color (float4)
DEPR DEPTH	Output depth (float)

The output depth value is special in that it may only be assigned a value of the form

```
...
float4 t = <texture shader operation>;
float z = dot(texCoord<n>, t.xyz);
float w = dot(texCoord<n+1>, t.xyz);
depth = z / w;
...
```

Auxiliary Texture Functions

Because the capabilities of the texture shader instructions are limited in `NV_texture_shader`, a set of auxiliary functions are provided in these profiles that express the functionality of the more complex texture shader instructions. These functions are merely provided as a convenience for writing **fp20** Cg programs. The same result can be achieved by writing the expanded form of each function directly. Using the expanded form has the additional advantage of being supported on other profiles.

These functions are summarized in [Table 38](#).

Table 38. **fp20** Auxiliary Texture Functions

Texture Function
Description
<pre> offsettex2D(uniform sampler2D tex, float2 st, float4 prevlookup, uniform float4 m) offsettexRECT(uniform samplerRECT tex, float2 st, float4 prevlookup, uniform float4 m) </pre> Performs the following: <pre> float2 newst = st + m.xy * prevlookup.xx + m.zw * prevlookup.yy; return tex2D/RECT(tex, newst); </pre> where st are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, and m is the offset texture matrix. This function can be used to generate the offset_2d or offset_rectangle NV_texture_shader instructions.
<pre> offsettex2DScaleBias(uniform sampler2D tex, float2 st, float4 prevlookup, uniform float4 m, uniform float scale, uniform float bias) offsettexRECTScaleBias(uniform samplerRECT tex, float2 st, float4 prevlookup, uniform float4 m, uniform float scale, uniform float bias) </pre> Performs the following <pre> float2 newst = st + m.xy * prevlookup.xx + m.zw * prevlookup.yy; float4 result = tex2D/RECT(tex, newst); return result * saturate(prevlookup.z * scale + bias); </pre> where st are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, m is the offset texture matrix, scale is the offset texture scale, and bias is the offset texture bias. This function can be used to generate the offset_2d_scale or offset_rectangle_scale NV_texture_shader instructions.

Table 38. **fp20** Auxiliary Texture Functions (continued)

Texture Function
Description
tex1D_dp3(sampler1D tex, float3 str, float4 prevlookup) Performs the following <pre>return tex1D(tex, dot(str, prevlookup.xyz));</pre> where str are texture coordinates associated with sampler tex , and prevlookup is the result of a previous texture operation. This function can be used to generate the dot_product_1d NV_texture_shader instruction.
tex2D_dp3x2(uniform sampler2D tex, float3 str, float4 intermediate_coord, float4 prevlookup) texRECT_dp3x2(uniform samplerRECT tex, float3 str, float4 intermediate_coord, float4 prevlookup) Performs the following <pre>float2 newst = float2(dot(intermediate_coord.xyz, prevlookup.xyz), dot(str, prevlookup.xyz));</pre> <pre>return tex2D/RECT(tex, newst);</pre> where str are texture coordinates associated with sampler tex , prevlookup is the result of a previous texture operation, and intermediate_coord are texture coordinates associated with the previous texture unit. This function can be used to generate the dot_product_2d or dot_product_rectangle NV_texture_shader instruction combinations.
tex3D_dp3x3(sampler3D tex, float3 str, float4 intermediate_coord1, float4 intermediate_coord2, float4 prevlookup) texCUBE_dp3x3(samplerCUBE tex, float3 str, float4 intermediate_coord1, float4 intermediate_coord2, float4 prevlookup)

Table 38. **fp20** Auxiliary Texture Functions (continued)

Texture Function
Description
<pre>texCUBE_reflect_eye_dp3x3(uniform samplerCUBE tex, float3 str, float4 intermediate_coord1, float4 intermediate_coord2, float4 prevlookup, uniform float3 eye)</pre>
Performs the following <pre>float3 N = float3(dot(intermediate_coord1.xyz, prevlookup.xyz), dot(intermediate_coord2.xyz, prevlookup.xyz), dot(coords.xyz, prevlookup.xyz)); return texCUBE(tex, 2 * dot(N, E) / dot(N, N) * N - E);</pre> where str are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, intermediate_coord1 are texture coordinates associated with the n-2 texture unit, intermediate_coord2 are texture coordinates associated with the n-1 texture unit, and eye is the eye-ray vector. This function can be used generate the dot_product_reflect_cube_map_const_eye NV_texture_shader instruction combination.
<pre>tex_dp3x2_depth(float3 str, float4 intermediate_coord, float4 prevlookup)</pre>
Performs the following <pre>float z = dot(intermediate_coord.xyz, prevlookup.xyz); float w = dot(str, prevlookup.xyz); return z / w;</pre> where str are texture coordinates associated with the nth texture unit, intermediate_coord are texture coordinates associated with the n-1 texture unit, and prevlookup is the result of a previous texture operation. This function can be used in conjunction with the DEPTH varying out semantic to generate the dot_product_depth_replace NV_texture_shader instruction combination.

Examples

The following examples show how a developer can use Cg to achieve NV_texture_shader and NV_register_combiners functionality.

Example 1

```
struct VertexOut {
    float4 color      : COLOR0;
    float4 texCoord0  : TEXCOORD0;
    float4 texCoord1  : TEXCOORD1;
};

float4 main(VertexOut IN,
            uniform sampler2D diffuseMap,
            uniform sampler2D normalMap) : COLOR
{
    float4 diffuseTexColor = tex2D(diffuseMap, IN.texCoord0.xy);
    float4 normal = 2 * (tex2D(normalMap, IN.texCoord1.xy)-0.5);
    float3 light_vector = 2 * (IN.color.rgb - 0.5);
    float4 dot_result = saturate(
        dot(light_vector, normal.xyz).xxxx);
    return dot_result * diffuseTexColor;
}
```

Example 2

```
struct VertexOut {
    float4 texCoord0 : TEXCOORD0;
    float4 texCoord1 : TEXCOORD1;
    float4 texCoord2 : TEXCOORD2;
    float4 texCoord3 : TEXCOORD3;
};

float4 main(VertexOut IN,
            uniform sampler2D normalMap,
            uniform sampler2D intensityMap,
            uniform sampler2D colorMap) : COLOR
{
    float4 normal = 2 * (tex2D(normalMap, IN.texCoord0.xy)-0.5);
    float2 intensCoord = float2(
        dot(IN.texCoord1.xyz, normal.xyz),
        dot(IN.texCoord2.xyz, normal.xyz));
    float4 intensity = tex2D(intensityMap, intensCoord);
    float4 color = tex2D(colorMap, IN.texCoord3.xy);
    return color * intensity;
}
```

DirectX Vertex Shader 2.x Profiles (**vs_2_***)

The DirectX Vertex Shader 2.0 profiles are used to compile Cg source code to DirectX 9 VS 2.0 vertex shaders⁶ and DirectX 9 VS 2.0 Extended vertex shaders.

- ❑ **Profile names**
 - vs_2_0** (for DirectX 9 VS 2.0 vertex shaders)
 - vs_2_x** (for DirectX 9 VS 2.0 extended vertex shaders)
- ❑ **How to invoke:** Use the compiler options
 - profile vs_2_0**
 - profile vs_2_x**

This section describes how using the **vs_2_0** and **vs_2_x** profiles affects the Cg source code that the developer writes.

Overview

The **vs_2_0** profile limits Cg to match the capabilities of DirectX VS 2.0 vertex shaders. The **vs_2_x** profile is the same as the **vs_2_0** profile but allows extended features such as dynamic flow control (branching).

Memory

DirectX 9 vertex shaders have a limited amount of memory for instructions and data.

Program Instruction Limit

DirectX 9 vertex shaders are limited to 256 instructions. If the compiler needs to produce more than 256 instructions to compile a program, it reports an error.

Vector Register Limit

Likewise, there are limited numbers of registers to hold program parameters and temporary results. Specifically, there are 256 read-only vector registers and 12–32 read/write vector registers. If the compiler needs more registers to compile a program than are available, it generates an error.

6. To understand the DirectX VS 2.0 Vertex Shaders and the code the compiler produces, see the Vertex Shader Reference in the DirectX 9 SDK documentation.

Statements and Operators

If the **vs_2_0** profile is used, then **if**, **while**, **do**, and **for** statements are allowed only if the loops they define can be unrolled because there is no dynamic branching in unextended VS 2.0 shaders.

If the **vs_2_x** profile is used, then **if**, **while**, and **do** statements are fully supported as long as the **DynamicFlowControlDepth** option is not 0.

Comparison operators are allowed (**>**, **<**, **>=**, **<=**, **==**, **!=**) and Boolean operators (**||**, **&&**, **?:**) are allowed. However, the logic operators (**&**, **|**, **^**, **~**) are not.

Data Types

The profiles implement data types as follows:

- ❑ **float** data types are implemented as IEEE 32-bit single precision.
- ❑ **half** and **double** data types are treated as **float**.
- ❑ **int** data type is supported using floating point operations, which adds extra instructions for proper truncation for divides, modulus and casts from floating point types.
- ❑ **fixed** or **sampler*** data types are not supported, but the profiles do provide the minimal partial support that is required for these data types by the core language specification—that is, it is legal to declare variables using these types, as long as no operations are performed on the variables.

Using Arrays

Variable indexing of arrays is allowed as long as the array is a uniform constant. For compatibility reasons arrays indexed with variable expressions need not be declared **const** just uniform. However, writing to an array that is later indexed with a variable expression yields unpredictable results.

Array data is not packed because vertex program indexing does not permit it. Each element of the array takes a single 4-float program parameter register. For example, **float arr[10]**, **float2 arr[10]**, **float3 arr[10]**, and **float4 arr[10]** all consume 10 program parameter registers.

It is more efficient to access an array of vectors than an array of matrices. Accessing a matrix requires a floor calculation, followed by a multiply by a constant to compute the register index. Because vectors (and scalars) take one register, neither the floor nor the multiply is needed. It is faster to do

matrix skinning using arrays of vectors with a premultiplied index than using arrays of matrices.

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **vs_2_0** and **vs_2_X** profiles are summarized in [Table 39](#).

Table 39. **vs_2_*** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(c0)–register(c255) C0–C255	Constant register [0..95]. The aliases c0-c95 (lowercase) are also accepted. If used with a variable that requires more than one constant register (for example, a matrix), the semantic specifies the first register that is used.

Binding Semantics for Varying Input/Output Data

Only the binding semantic names need be given for these profiles. The vertex parameter input registers are allocated dynamically. All the semantic names, except **POSITION**, can have a number from 0 to 15 after them.

Table 40. **vs_2_*** Varying Input Binding Semantics

POSITION	PSIZE
BLENDWEIGHT	BLENDINDICES
NORMAL	TEXCOORD
COLOR	TANGENT
TESSFACTOR	BINORMAL

The valid binding semantics for varying output parameters in the **vs_2_0** and **vs_2_X** profiles are summarized in [Table 41](#).

These map to output registers in DirectX 9 vertex shaders.

Table 41. **vs_2_*** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION	Output position: oPos
PSIZE	Output point size: oPts
FOG	Output fog value: oFog
COLOR0 - COLOR1	Output color values: oD0 , oD1
TEXCOORD0 - TEXCOORD7	Output texture coordinates: oT0 - oT7

Options

The **vs_2_x** profile allows the following profile specific options:

DynamicFlowControlDepth=<n> (where **n** = 0 or 24; default 24)
NumTemps=<n> (where 12 ≤ **n** ≤ 32; default 16)
Predication (default true)

DirectX Pixel Shader 2.x Profiles (**ps_2_***)

The DirectX Pixel Shader 2.0 Profiles are used to compile Cg source code to DirectX 9 PS 2.0 pixel shaders⁷ and DirectX 9 PS 2.0 extended pixel shaders.

- ❑ **Profile names**
 - ps_2_0** (for DirectX 9 PS 2.0 pixel shaders)
 - ps_2_x** (for DirectX 9 PS 2.0 extended pixel shaders)
- ❑ **How to invoke:** Use the compiler options
 - profile ps_2_0**
 - profile ps_2_x**

The **ps_2_0** profile limits Cg to match the capabilities of DirectX PS 2.0 pixel shaders. The **ps_2_x** profile is the same as the **ps_2_0** profile but allows extended features such as arbitrary swizzles, larger limit on number of instructions, no limit on texture instructions, no limit on texture dependent reads, and support for predication.

This section describes the capabilities and restrictions of Cg when using these profiles.

Memory

Program Instruction Limit

DirectX 9 Pixel shaders have a limit on the number of instructions in a pixel shader.

- ❑ PS 2.0 (**ps_2_0**) pixel shaders are limited to 32 texture instructions and 64 arithmetic instructions.
- ❑ Extended PS 2 (**ps_2_x**) shaders have a limit of maximum number of total instructions between 96 to 1024 instructions.

There is no separate texture instruction limit on extended pixel shaders.

If the compiler needs to produce more than the maximum allowed number of instructions to compile a program, it reports an error.

Vector Register Limit

Likewise, there are limited numbers of registers to hold program parameters and temporary results. Specifically, there are 32 read-only vector registers

7. To understand the capabilities of DirectX PS 2.0 Pixel Shaders and the code produced by the compiler, refer to the Pixel Shader Reference in the DirectX 9 SDK documentation.

and 12-32 read/write vector registers. If the compiler needs more registers to compile a program than are available, it generates an error.

Language Constructs and Support

Data Types

This profile implements data types as follows:

- ❑ **float** data type is implemented as IEEE 32-bit single precision.
- ❑ **half**, **fixed**, and **double** data types are treated as float.
half data types can be used to specify partial precision hint for pixel shader instructions.
- ❑ **int** data type is supported using floating point operations.
- ❑ **sampler*** types are supported to specify sampler objects used for texture fetches.

Statements and Operators

With the **ps_2_0** profiles **while**, **do**, and **for** statements are allowed only if the loops they define can be unrolled because there is no dynamic branching in PS 2.0 shaders. In current Cg implementation, extended **ps_2_x** shaders also have the same limitation.

Comparison operators are allowed (**>**, **<**, **>=**, **<=**, **==**, **!=**) and Boolean operators (**||**, **&&**, **?:**) are allowed. However, the logic operators (**&**, **|**, **^**, **~**) are not.

Using Arrays and Structures

Variable indexing of arrays is not allowed. Array and structure data is not packed.

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **ps_2_0** and **ps_2_x** profiles are summarized in [Table 42](#).

Table 42. **ps_2_*** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(s0)–register(s15) TEXUNIT0-TEXUNIT15	Texunit unit <i>N</i> , where <i>N</i> is in range [0..15] May only be used with uniform inputs with sampler* types.
register(c0)-register(c31) C0-C31	Constant register <i>N</i> , where <i>N</i> is in range [0..31] May only be used with uniform inputs.

Binding Semantics for Varying Input/Output Data

The valid binding semantics for varying input parameters in the **ps_2_0** and **ps_2_x** profiles are summarized in [Table 43](#).

Table 43. **ps_2_*** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data (type)
COLOR0	Input color 0 (float4)
COLOR1	Input color 1 (float4)
TEXCOORD0-TEXCOORD7	Input texture coordinates (float4)

The valid binding semantics for varying output parameters in the **ps_2_0** and **ps_2_x** profiles are summarized in [Table 44](#).

Table 44. **ps_2_*** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
COLOR, COLOR0	Output color (float4)
DEPTH	Output depth (float)

Options

The **ps_2_x** profile allows the following profile specific options:

NumTemps=<n>	(where $0 \leq n \leq 32$; default 32)
NumInstructionSlots=<n>	(where $n \geq 0$; default 1024)
Predication=	(where $b = 0$ or 1 ; default 1)
ArbitrarySwizzle=	(where $b = 0$ or 1 ; default 1)
GradientInstructions=	(where $b = 0$ or 1 ; default 1)
NoDependentReadLimit=	(where $b = 0$ or 1 ; default 1)
NoTexInstructionLimit=	(where $b = 0$ or 1 ; default 1)

Limitations in this Implementation

Currently, this profile implementation has the following limitations:

- ❑ Dynamic flow control is not supported in extended pixel shaders.
- ❑ Multiple color outputs are not supported in pixel shaders. Only **color0** is supported.

DirectX Vertex Shader 1.1 Profile (**vs_1_1**)

The DirectX Vertex Shader 1.1 profile is used to compile Cg source code to DirectX 8.1 Vertex Shaders and DirectX 9 VS 1.1 shaders⁸.

- ❑ **Profile name:** **vs_1_1**
- ❑ **How to invoke:** Use the compiler option **-profile vs_1_1**.

The **vs_1_1** profile limits Cg to match the capabilities of DirectX Vertex Shaders.

This section describes how using the **vs_1_1** profile affects the Cg source code that the developer writes.

Memory Restrictions

DirectX 8 vertex shaders have a limited amount of memory for instructions and data.

Program Instruction Limits

The DirectX 8 vertex shaders are limited to 128 instructions. If the compiler needs to produce more than 128 instructions to compile a program, it reports an error.

Vector Register Limits

Likewise, there are limited numbers of registers to hold program parameters and temporary results. Specifically, there are 96 read-only vector registers and 12 read/write vector registers. If the compiler needs more registers to compile a program than are available, it generates an error.

Language Constructs and Support

Data Types

This profile implements data types as follows:

- ❑ **float** data types are implemented as IEEE 32-bit single precision.
- ❑ **half** and **double** data types are treated as **float**.

8. To understand the DirectX VS 1.1 Vertex Shaders and the code the compiler produces, see the Vertex Shader Reference in the DirectX 8.1 SDK documentation.

- ❑ **int** data type is supported using floating point operations, which adds extra instructions for proper truncation for divides, modulus and casts from floating point types.
- ❑ **fixed** or **sampler*** data types are not supported, but the profile does provide the minimal partial support that is required for these data types by the core language specification—that is, it is legal to declare variables using these types, as long as no operations are performed on the variables.

Statements and Operators

The **if**, **while**, **do**, and **for** statements are allowed only if the loops they define can be unrolled, because there is no branching in VS 1.1 shaders.

There are no subroutine calls either, so all functions are inlined. Comparison operators are allowed (**>**, **<**, **>=**, **<=**, **==**, **!=**) and Boolean operators (**||**, **&&**, **?:**) are allowed. However, the logic operators (**&**, **|**, **^**, **~**) are not allowed.

Using Arrays

Variable indexing of arrays is allowed as long as the array is a uniform constant. For compatibility reasons arrays indexed with variable expressions need not be declared **const** just uniform. However, writing to an array that is later indexed with a variable expression yields unpredictable results.

Array data is not packed because vertex program indexing does not permit it. Each element of the array takes a single 4-float program parameter register. For example, **float arr[10]**, **float2 arr[10]**, **float3 arr[10]**, and **float4 arr[10]** all consume ten program parameter registers.

It is more efficient to access an array of vectors than an array of matrices. Accessing a matrix requires a floor calculation, followed by a multiply by a constant to compute the register index. Because vectors (and scalars) take one register, neither the floor nor the multiply is needed. It is faster to do matrix skinning using arrays of vectors with a premultiplied index than using arrays of matrices.

Constants

Literal constants can be used with this profile, but it is not possible to store them in the program itself. Instead the compiler will issue, as comments, a list of program parameter registers and the constants that need to be loaded into them. The Cg run-time system will handle loading the constants, as directed by the compiler.

Note: If the Cg run-time system is not used, it is the responsibility of the programmer to make sure that the constants are loaded properly.

Bindings

Binding Semantics for Uniform Data

The valid binding semantics for uniform parameters in the **vs_1_1** profile are summarized in [Table 45](#).

Table 45. **vs_1_1** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(c0)–register(c95) C0–C95	Constant register [0..95]. The aliases c0–c95 (lowercase) are also accepted. If used with a variable that requires more than one constant register (for example, a matrix), the semantic specifies the first register that is used.

Binding Semantics for Varying Input/Output Data

The valid binding semantics for uniform parameters in the **vs_1_1** profile are summarized in [Table 46](#). These map to the input registers in DirectX 8.1 vertex shaders.

Table 46. **vs_1_1** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION	Vertex shader input register: v0
BLENDWEIGHT	Vertex shader input register: v1
BLENDINDICES	Vertex shader input register: v2
NORMAL	Vertex shader input register: v3
PSIZE	Vertex shader input register: v4
COLOR0, DIFFUSE	Vertex shader input register: v5

Table 46. **vs_1_1** Varying Input Binding Semantics (continued)

Binding Semantics Name	Corresponding Data
COLOR1, SPECULAR	Vertex shader input register: v6
TEXCOORD0-TEXCOORD7	Vertex shader input register: v7-v14
TANGENTⁱ	Vertex shader input register: v14
BINORMAL	Vertex shader input register: v15

i. **TANGENT** is an alias for **TEXCOORD7**.

The valid binding semantics for varying output parameters in the **vs_1_X** profile. These map to output registers in DirectX 8.1 vertex shaders are summarized in [Table 47](#).

Table 47. **vs_1_1** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
POSITION	Output position: oPos
PSIZE	Output point size: oPts
FOG	Output fog value: oFog
COLOR0-COLOR1	Output color values: oD0, oD1
TEXCOORD0-TEXCOORD7	Output texture coordinates: oT0-oT7

Options

When using the **vs_1_1** profile under DirectX 9 it is necessary to tell the compiler to produce **dcl** statements to declare varying inputs. The option **-profileopts dcls** causes **dcl** statements to be added to the compiler output.

DirectX Pixel Shader 1.x Profiles (**ps_1_***)

The DirectX pixel shader 1_X profiles are used to compile Cg source code to DirectX PS 1.1, PS 1.2, or PS 1.3 pixel shader assembly.

- ❑ **Profile names**
 - ps_1_1** (for DirectX PS 1.1 pixel shaders)
 - ps_1_2** (for DirectX PS 1.2 pixel shaders)
 - ps_1_3** (for DirectX PS 1.3 pixel shaders)
- ❑ **How to invoke:** Use the compiler options
 - profile ps_1_1**
 - profile ps_1_2**
 - profile ps_1_3**

The deprecated profile **dx8ps** is also available and is synonymous with **ps_1_1**.

This document describes the capabilities and restrictions of Cg when using the DirectX pixel shader 1_X profiles.

Overview

DirectX PS 1.4 is not currently supported by any Cg profile; all statements about **ps_1_X** in the remainder of this document refer only to **ps_1_1**, **ps_1_2** and **ps_1_3**.

The underlying instruction set and machine architecture limit programmability in these profiles compared to what is allowed by Cg constructs⁹. Thus, these profiles place additional restrictions on what can and cannot be done in a Cg program.

The main differences between these profiles from the Cg perspective is that additional texture addressing operations are exposed in **ps_1_2** and **ps_1_3** and the depth value output is made available (in a limited form) in **ps_1_3**.

Operations in the DirectX pixel shader 1_X profiles can be categorized as texture addressing operations and arithmetic operations. Texture addressing operations are operations which generate texture addressing instructions, arithmetic operations are operations which generate arithmetic instructions. A Cg program in one of these profiles is limited to generating a maximum of four texture addressing instructions and eight arithmetic instructions. Since

9. For more details about the underlying instruction sets, their capabilities, and their limitations, refer to the MSDN documentation of DirectX pixel shaders 1.1, 1.2 and 1.3.

these numbers are quite small, users need to be very aware of this limitation while writing Cg code for these profiles.

There are certain simple arithmetic operations that can be applied to inputs of texture addressing operations and to inputs and outputs of arithmetic operations without generating an arithmetic instruction. From here on, these operations are referred to as input modifiers and output modifiers.

The **ps_1_x** profiles also restrict when a texture addressing operation or arithmetic operation can occur in the program. A texture addressing operation may not have any dependency on the output of an arithmetic operation unless

- ❑ The arithmetic operation is a valid input modifier for the texture addressing operation.
- ❑ The arithmetic operation is part of a complex texture addressing operation (which are summarized in the section on [Auxiliary Texture Functions](#)).

Modifiers

Input and output modifiers may be used to perform simple arithmetic operations without generating an arithmetic instruction. Instead, the arithmetic operation modifies the assembly instruction or source registers to which it is applied. For example, the following Cg expression:

$$z = (x - 0.5 + y) / 2$$

could generate the following pixel shader instruction (assuming **x** is in **t0**, **y** is in **t1**, and **z** is in **r0**):

```
add_d2 r0, t0_bias, t1
```

How different DirectX pixel shader 1_X instruction set modifiers are expressed in Cg programs are summarized in [Table 48](#). For more details on the context in which each modifier is allowed and ways in which modifiers may be combined refer to the DirectX pixel shader 1_X documentation.

Table 48. **ps_1_x** Instruction Set Modifiers

Instruction/Register Modifier	Cg Expression
instr_x2	2*x
instr_x4	4*x
instr_d2	x/2

Table 48. **ps_1_x** Instruction Set Modifiers (continued)

Instruction/Register Modifier	Cg Expression
instr_sat	saturate(x) (i.e. $\min(1, \max(0, x))$)
reg_bias	x-0.5
1-reg	1-x
-reg	-x
reg_bx2	2*(x-0.5)

Language Constructs and Support

Data Types

In the **ps_1_X** profiles, operations occur on signed clamped floating point values in the range **MaxPixelShaderValue** to **MaxPixelShaderValue**, where **MaxPixelShaderValue** is determined by the DirectX implementation. These profiles allow all data types to be used, but all operations are carried out in the above range. Refer to the DirectX pixel shader 1_X documentation for more details.

Statements and Operators

The DirectX pixel shader 1_X profiles support all of the Cg language constructs, with the following exceptions:

- ❑ Arbitrary swizzles are not supported (though arbitrary write masks are). Only the following swizzles are allowed
`.x/.r .y/.g .z/.b .w/.a`
`.xy/.rg .xyz/.rgb .xyzw/.rgba`
`.xxx/.rrr .yyy/.ggg .zzz/.bbb .www/.aaa`
`.xxxx/.rrrr .yyyy/.gggg .zzzz/.bbbb .www/.aaaa`
- ❑ Matrix swizzles are not supported.
- ❑ Boolean operators other than **<**, **<=**, **>** and **>=** are not supported. Furthermore, **<**, **<=**, **>** and **>=** are only supported as the condition in the **?:** operator.
- ❑ Bitwise integer operators are not supported.
- ❑ **/** is not supported unless the divisor is a non-zero constant or it is used to compute the depth output in **ps_1_3**.

- ❑ `%` is not supported.
- ❑ Ternary `?:` is supported if the boolean test expression is a compile-time boolean constant, a uniform scalar boolean or a scalar comparison to a constant value in the range `[-0.5, 1.0]` (for example, `a > 0.5 ? b : c`).
- ❑ **do**, **for**, and **while** loops are supported only when they can be completely unrolled.
- ❑ arrays, vectors, and matrices may be indexed only by compile-time constant values or index variables in loops that can be completely unrolled.
- ❑ The **discard** statement is not supported. The similar but less general **clip()** function is supported.
- ❑ The use of an **allocation-rule-identifier** for an input or output **struct** is optional.

Standard Library Functions

Because the DirectX pixel shader 1_X profiles have limited capabilities, not all of the Cg standard library functions are supported. [Table 49](#) presents the Cg standard library functions that are supported by these profiles. See the standard library documentation for descriptions of these functions.

Table 49. Supported Standard Library Functions

<code>dot(floatN, floatN)</code>
<code>lerp(floatN, floatN, floatN)</code>
<code>lerp(floatN, floatN, float)</code>
<code>tex1D(sampler1D, float)</code>
<code>tex1D(sampler1D, float2)</code>
<code>tex1Dproj(sampler1D, float2)</code>
<code>tex1Dproj(sampler1D, float3)</code>
<code>tex2D(sampler2D, float2)</code>
<code>tex2D(sampler2D, float3)</code>
<code>tex2Dproj(sampler2D, float3)</code>
<code>tex2Dproj(sampler2D, float4)</code>
<code>tex3D(sampler3D, float3)</code>

Table 49. Supported Standard Library Functions (continued)

tex3Dproj(sampler3D, float4)
texCUBE(samplerCUBE, float3)
texCUBEproj(samplerCUBE, float4)

Note: The non-projective texture lookup functions are actually done as projective lookups on the underlying hardware. Because of this, the **w** component of the texture coordinates passed to these functions from the application or vertex program must contain the value 1.

Texture coordinate parameters for projective texture lookup functions must have swizzles that match the swizzle done by the generated texture addressing instruction. While this may seem burdensome, it is intended to allow **ps_1_X** profile programs to behave correctly under other pixel shader profiles.

The swizzles required on the texture coordinate parameter to the projective texture lookup functions are listed in [Table 50](#).

Table 50. Required Projective Texture Lookup Swizzles

Texture Lookup Function	Texture Coordinate Swizzle
tex1Dproj	.xw/.ra
tex2Dproj	.xyw/.rga
texRECTproj	.xyw/.rga
tex3Dproj	.xyzw/.rgba
texCUBEproj	.xyzw/.rgba

Bindings

Manual Assignment of Bindings

The Cg compiler can determine bindings between texture units and uniform sampler parameters/texture coordinate inputs automatically. This automatic assignment is based on the context in which uniform sampler parameters and texture coordinate inputs are used together.

To specify bindings between texture units and uniform parameters/texture coordinates to match their application, all sampler uniform parameters and texture coordinate inputs that are used in the program must have matching binding semantics—that is, **TEXUNIT<n>** may only be used with **TEXCOORD<n>**.

Partially specified binding semantics may not work in all cases. Fundamentally, this restriction is due to the close coupling between texture samplers and texture coordinates in DirectX pixel shaders 1_X.

Binding Semantics for Uniform Data

If a binding semantic for a uniform parameter is not specified then the compiler will allocate one automatically. Scalar uniform parameters may be allocated to either the **xyz** or the **w** portion of a constant register depending on how they are used within the Cg program. When using the output of the compiler without the Cg runtime, you must set all values of a scalar uniform to the desired scalar value, not just the **x** component.

The valid binding semantics for uniform parameters in the **ps_1_X** profiles are summarized in [Table 51](#).

Table 51. **ps_1_x** Uniform Input Binding Semantics

Binding Semantics Name	Corresponding Data
register(s0)–register(s3) TEXUNIT0–TEXUNIT3	Texture unit N , where N is in range [0..3]. May be used only with uniform inputs with sampler* types.
register(c0)–register(c7) C0–C7	Constant register [0..7]

Binding Semantics for Varying Input/Output Data

The varying input binding semantics in the **ps_1_X** profiles are the same as the varying output binding semantics of the **vs_1_1** profile.

Varying input binding semantics in the **ps_1_X** profiles consist of **COLOR0**, **COLOR1**, **TEXCOORD0**, **TEXCOORD1**, **TEXCOORD2** and **TEXCOORD3**. These map to output registers in DirectX vertex shaders.

The valid binding semantics for varying input parameters in the **ps_1_X** profiles are summarized in [Table 52](#).

Table 52. **ps_1_x** Varying Input Binding Semantics

Binding Semantics Name	Corresponding Data
COLOR, COLOR0 COL, COL0	Input color value v0
COLOR1 COL1	Input color value v1
TEXCOORD0–TEXCOORD3 TEX0–TEX3	Input texture coordinates t0–t3

Additionally, the **ps_1_X** profiles allow **POSITION**, **FOG**, **PSIZE**, **TEXCOORD4**, **TEXCOORD5**, **TEXCOORD6**, and **TEXCOORD7** to be specified on varying inputs, provided these inputs are not referenced. This allows Cg programs to have the same structure specify the varying output of a **vs_1_1** profile program and the varying input of a **ps_1_X** profile program.

The valid binding semantics for varying output parameters in the **ps_1_X** profile are summarized in [Table 53](#).

Table 53. **ps_1_x** Varying Output Binding Semantics

Binding Semantics Name	Corresponding Data
COLOR, COLOR0 COL, COL0	Output color (float4)
DEPTH DEPR	Output depth (float)

The output depth value is special in that it may only be assigned a value in the **ps_1_3** profile, and must be of the form

```
...
float4 t = <texture addressing operation>;
float z = dot(texCoord<n>, t.xyz);
float w = dot(texCoord<n+1>, t.xyz);
depth = z / w;
...
```

Auxiliary Texture Functions

Because the capabilities of the texture addressing instructions are limited in DirectX pixel shader 1_X, a set of auxiliary functions is provided in these profiles that express the functionality of the more complex texture addressing instructions. These functions are provided merely as a convenience for writing **ps_1_X** Cg programs. The same result can be achieved by writing the expanded form of each function directly. The expanded form has the added advantage of being supported on other profiles.

These functions are summarized in [Table 54](#).

Table 54. **ps_1_x** Auxiliary Texture Functions

Texture Function
Description
offsettex2D(uniform sampler2D tex, float2 st, float4 prevlookup, uniform float4 m)
Performs the following: <pre>float2 newst = st + m.xy * prevlookup.xx + m.zw * prevlookup.yy; return tex2D(tex, newst);</pre> where st are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, and m is the 2-D bump environment mapping matrix. This function can generate the texbem instruction in all ps_1_X profiles.
offsettex2DScaleBias(uniform sampler2D tex, float2 st, float4 prevlookup, uniform float4 m, uniform float scale, uniform float bias)
Performs the following: <pre>float2 newst = st + m.xy * prevlookup.xx + m.zw * prevlookup.yy; float4 result = tex2D(tex, newst); return result * saturate(prevlookup.z * scale + bias);</pre> where st are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, m is the 2-D bump environment mapping matrix, scale is the 2-D bump environment mapping scale factor, and bias is the 2-D bump environment mapping offset. This function can generate the texbem1 instruction in all ps_1_X profiles.

Table 54. **ps_1_x** Auxiliary Texture Functions (continued)

Texture Function
Description
tex1D_dp3(sampler1D tex, float3 str, float4 prevlookup) Performs the following: <pre>return tex1D(tex, dot(str, prevlookup.xyz));</pre> where str are texture coordinates associated with sampler tex , and prevlookup is the result of a previous texture operation. This function can be used to generate the texdp3tex instruction in the ps_1_2 and ps_1_3 profiles.
tex2D_dp3x2(uniform sampler2D tex, float3 str, float4 intermediate_coord, float4 prevlookup) Performs the following: <pre>float2 newst = float2(dot(intermediate_coord.xyz, prevlookup.xyz), dot(str, prevlookup.xyz));</pre> <pre>return tex2D(tex, newst);</pre> where str are texture coordinates associated with sampler tex , prevlookup is the result of a previous texture operation, and intermediate_coord are texture coordinates associated with the previous texture unit. This function can be used to generate the texm3x2pad/texm3x2tex instruction combination in all ps_1_X profiles.
tex3D_dp3x3(sampler3D tex, float3 str, float4 intermediate_coord1, float4 intermediate_coord2, float4 prevlookup) texCUBE_dp3x3(samplerCUBE tex, float3 str, float4 intermediate_coord1, float4 intermediate_coord2, float4 prevlookup) Performs the following: <pre>float3 newst = float3(dot(intermediate_coord1.xyz, prevlookup.xyz), dot(intermediate_coord2.xyz, prevlookup.xyz), dot(str, prevlookup.xyz));</pre> <pre>return tex3D/CUBE(tex, newst);</pre> where

Table 54. **ps_1_x** Auxiliary Texture Functions (continued)

Texture Function
Description
str are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, intermediate_coord1 are texture coordinates associated with the n-2 texture unit, and intermediate_coord2 are texture coordinates associated with the n-1 texture unit. This function can be used to generate the texm3x3pad/texm3x3pad/texm3x3tex instruction combination in all ps_1_X profiles.
<pre>texCUBE_reflect_dp3x3(uniform samplerCUBE tex, float4 strq, float4 intermediate_coord1, float4 intermediate_coord2, float4 prevlookup)</pre> Performs the following: <pre>float3 E = float3(intermediate_coord2.w, intermediate_coord1.w, strq.w); float3 N = float3(dot(intermediate_coord1.xyz, prevlookup.xyz), dot(intermediate_coord2.xyz, prevlookup.xyz), dot(strq.xyz, prevlookup.xyz)); return texCUBE(tex, 2 * dot(N, E) / dot(N, N) * N - E);</pre> where strq are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, intermediate_coord1 are texture coordinates associated with the n-2 texture unit, and intermediate_coord2 are texture coordinates associated with the n-1 texture unit. This function can be used to generate the texm3x3pad/texm3x3pad/texm3x3vspec instruction combination in all ps_1_X profiles.

Table 54. **ps_1_x** Auxiliary Texture Functions (continued)

Texture Function
Description
texCUBE_reflect_eye_dp3x3 (uniform samplerCUBE tex, float3 str, float4 intermediate_coord1, float4 intermediate_coord2, float4 prevlookup, uniform float3 eye)
Performs the following: <pre>float3 N = float3(dot(intermediate_coord1.xyz, prevlookup.xyz), dot(intermediate_coord2.xyz, prevlookup.xyz), dot(coords.xyz, prevlookup.xyz)); return texCUBE(tex, 2 * dot(N, E) / dot(N, N) * N - E);</pre> where strq are texture coordinates associated with sampler tex, prevlookup is the result of a previous texture operation, intermediate_coord1 are texture coordinates associated with the n-2 texture unit, intermediate_coord2 are texture coordinates associated with the n-1 texture unit, and eye is the eye-ray vector. This function can be used to generate the texm3x3pad/texm3x3pad/texm3x3spec instruction combination in all ps_1_X profiles.
tex_dp3x2_depth (float3 str, float4 intermediate_coord, float4 prevlookup)
Performs the following: <pre>float z = dot(intermediate_coord.xyz, prevlookup.xyz); float w = dot(str, prevlookup.xyz); return z / w;</pre> where str are texture coordinates associated with the nth texture unit, intermediate_coord are texture coordinates associated with the n-1 texture unit, and prevlookup is the result of a previous texture operation. This function can be used with the DEPTH varying out semantic to generate the texm3x2pad/texm3x2depth instruction combination in ps_1_3.

Examples

The following examples illustrate how a developer can use Cg to achieve DirectX pixel shader 1_X functionality.

Example 1

```
struct VertexOut {
    float4 color      : COLOR0;
    float4 texCoord0  : TEXCOORD0;
    float4 texCoord1  : TEXCOORD1;
};

float4 main(VertexOut IN,
            uniform sampler2D diffuseMap,
            uniform sampler2D normalMap) : COLOR
{
    float4 diffuseTexColor = tex2D(diffuseMap, IN.texCoord0.xy);
    float4 normal = 2 * (tex2D(normalMap, IN.texCoord1.xy)-0.5);
    float3 light_vector = 2 * (IN.color.rgb - 0.5);
    float4 dot_result = saturate(dot(light_vector,
                                    normal.xyz).xxxx);
    return dot_result * diffuseTexColor;
}
```

Example 2

```
struct VertexOut {
    float4 texCoord0 : TEXCOORD0;
    float4 texCoord1 : TEXCOORD1;
    float4 texCoord2 : TEXCOORD2;
    float4 texCoord3 : TEXCOORD3;
};

float4 main(VertexOut IN,
            uniform sampler2D normalMap,
            uniform sampler2D intensityMap,
            uniform sampler2D colorMap) : COLOR
{
    float4 normal = 2 * (tex2D(normalMap, IN.texCoord0.xy)-0.5);
    float2 intensCoord = float2(
        dot(IN.texCoord1.xyz, normal.xyz),
        dot(IN.texCoord2.xyz, normal.xyz));
    float4 intensity = tex2D(intensityMap, intensCoord);
    float4 color = tex2D(colorMap, IN.texCoord3.xy);
    return color * intensity;
}
```


Appendix C

Nine Steps to High-Performance Cg

Writing Cg code that compiles to efficient programs requires techniques and approaches that are different from efficient programming in C, C++, or Java. While some of the basic lessons are the same (such as using efficient underlying algorithms), the hardware programming model of modern GPUs is substantially different from that of modern CPUs. This can lead to pitfalls—where you may be disappointed by your shader’s performance—as well as to opportunities—where you can push the GPU to its limits through careful programming.

The Cg language shields you from the majority of the low-level details of GPU hardware, enabling you to think about your shaders at a higher level than the low-level GPU instruction sets. However, just as an understanding of modern computer architecture (such as cache and memory hierarchy issues) is important for writing fast C and C++ code, understanding a bit about the GPU can help you write better Cg code. This appendix focuses on techniques for maximizing performance from vertex and fragment programs written in Cg and running on the NVIDIA GeForce FX architecture (specifically the **vp30**, **fp30**, **arbfp1**, **ps_2_0**, **ps_2_x**, **vs_2_0**, and **vs_2_x** profiles), although many of the principles are more broadly applicable.

1. Program for Vectorization

The GPU can generally perform four arithmetic operations as quickly as it can perform a single operation. Therefore, if you have two vectors of four floating point values,

```
float4 a, b;
```

you can add the two vectors together

```
float4 c = a+b;
```

with no more computational expense than adding together two of their elements

```
float d = a.x + b.x;
```

This has two implications for efficient programming. First, you should try to write code that naturally maps to these vector operations. If you want to add two **float4** variables together, it may be substantially less efficient to write it this way:

```
float4 c = float4(a.x + b.x, a.x + b.y, a.z + b.z,
                  a.w + b.w);
```

than to write it this way:

```
float4 c = a+b;
```

The compiler does its best to find vectorization in your programs, but the more vectorized your original code is, the better starting place it has to work from.

A more specific example comes from a common computation done for tangent-space bump mapping. Given a texture map that encodes a bump map by storing the offset along the tangent direction in **x**, the offset along the binormal in **y**, and the offset along the normal in **z**, the bump-mapped normal is computed by scaling the tangent, binormal, and normal appropriately. In C or C++, the natural way to write this computation is as shown:

```
// Tangent, binormal, normal. Passed in from vertex program.
Float3 T, B, N;
Float3 Nbump;    // Bump-mapped normal
Float3 bump = tex2D(bumpSampler, uv);
Nbump.x = bump.x * T.x + bump.y * B.x + bump.z * N.x;
Nbump.y = bump.x * T.y + bump.y * B.y + bump.z * N.y;
Nbump.z = bump.x * T.z + bump.y * B.z + bump.z * N.z;
```

However, here we have written a series of computations that add and multiply single pairs of floating point values at a time. After a little algebra, we can rewrite this as three multiplies of a **float3** and a **float** and two **float3** additions—which runs several times faster than the original!

```
Nbump = bump.x * T + bump.y * B + bump.z * N;
```

2. Use Swizzles to Make the Most of Vectorization

The GPU can swizzle the values in vectors with no performance penalty (recall that a swizzle can be used to rearrange the elements of a vector).

Given a vector:

```
float3 a = float3(0, 1, 2);
```

swizzles construct new vectors:

```
a.xxx = float3(0, 0, 0);
a.yzz = float3(1, 2, 2);
a.zy  = float2(2, 1);
```

and so forth. By swizzling your data carefully, you can still take advantage of vectorization, even when you don't want to use the same component of both vectors on both sides of your computation. For example, consider the computation of the cross product. Given two three-dimensional vectors, the cross product returns a new vector that is perpendicular to the given vectors. It is computed by

```
float3 a, b;
float3 c = float3(a.y*b.z - a.z*b.y, a.z*b.x - a.x*b.z,
                  a.x*b.y - a.y*b.x);
```

Here we've again got a lot of arithmetic operations, each using a single pair of **float** values. Some cleverness lets us turn this into a vectorized operation. Below is the implementation of the **cross()** function from the Cg Standard Library, requiring just two vector multiply operations and one vector subtraction operation:

```
float3 cross(float3 a, float3 b) {
    return a.yzx * b.zxy - a.zxy * b.yzx;
}
```

Confirm for yourself that this computes the same value as the first section of code for the cross product; note that it exposes much more vectorized computation for the GPU to efficiently process.

3. Use the Cg Standard Library

The functions in the Cg Standard Library have been carefully written for both efficiency and correctness. By using Standard Library functions when appropriate, you can automatically take advantage of the work that went into making sure they compile to fast code on GPUs while you concentrate on the hard problems you're solving in your own shaders.

Particularly fast Standard Library functions include **dot()**, which computes the dot product of two vectors, **abs()**, which computes the absolute value of a variable, **saturate()**, which clamps a value to be between zero and one, and **min()** and **max()**, which return the minimum and maximum of a pair of values. You won't be able to write more efficient implementations of these functions than the Standard Library provides because many of them compile directly to GPU assembly language instructions. Writing a dot product function of your own,

```
float mydot(float3 a, float3 b) {  
    return a.x*b.x + a.y*b.y + a.z*b.z;  
}
```

compiles to a handful of instructions, while the built-in **dot()** function compiles to a single specialized dot product instruction. There's no other way to get to this instruction other than by using the Standard Library.

Two functions deserve particular attention. The **abs()** function usually has no cost in either vertex or fragment programs because the GPU can evaluate the function while executing other instructions. Similarly, the **saturate()** function usually has no cost in fragment programs. Do not hesitate to use these functions when appropriate.

4. Use Texture Maps to Encode Complex Functions

For profiles that support texture maps, filtered texture map lookups are extraordinarily efficient. If you have a complex function that takes more than a handful of arithmetic operations to evaluate, you might want to encode the function in a texture map. Say that you have written a function **f(x, y)** that is a bottleneck in your shader. Assume for now that it is always called with values of **x** and **y** between zero and one, and that the value that **f(x, y)** computes is always between zero and one. If the function is reasonably smooth and you don't need to compute it at extremely high precision, you

can precompute the function in your application and store it in a texture map, replacing calls like

```
float val = f(x, y);
```

with code like

```
float val = tex2D(fSampler, float2(x, y)).x;
```

This method can also be applied to one- and three-dimensional functions, using 1D and 3D texture maps.

More generally, the values you pass to the function may not be in the range **[0, 1]**, and the values your function returns may not be in the range **[0, 1]**. In this case, the following two utility functions can serve as a base:

remapTo01() remaps the range **[low, high]** into **[0, 1]**, **remapFrom01()** does the opposite.

```
float4 remapTo01(float4 v, float4 low, float4 high) {
    return saturate((v - low)/(high-low));
}

float4 remapFrom01(float4 v, float4 low, float4 high) {
    return lerp(low, high, v);
}
```

Don't forget vectorization here as well. If two **float**-valued functions have the same domain and range, you can pack them into two texture components of the same texture. Only one texture lookup is needed to load them both, and vectorized versions of the **remap*()** can be used to do the remapping more efficiently as well.

5. Use Data Types with Minimum Sufficient Precision

For profiles that support multiple precisions, a general rule of thumb is that if you can do a computation with **fixed** precision variables, the computation is faster than if you use **half**; and if you use **half**, the computation is faster than if you use **float**. Although sometimes you need the range and extra precision that **half** and **float** offer, you should avoid using them unless necessary.

6. Use the Right Standard Library Routines for Shading Computations

If you're implementing a shading model (such as Lambertian, Blinn, or Phong), you'll generally be performing some dot product routines, clamping negative results to zero, and raising some of the values to a power, to compute a specular exponent. There are a few tricks that can speed up this process:

- ❑ Be sure to use the **dot()** function when computing dot products.
- ❑ If you need to clamp the result of a dot product computation to the range **[0, 1]** in a fragment program, use the **saturate()** function instead of **max()**. This is often written as **max(0, dot(N, L))**, but as long as the **N** and **L** vectors are normalized, this can be written equivalently as **saturate(dot(N, L))** because the dot product of two normalized vectors is never greater than one. Given that **saturate()** is free in fragment programs (see ["3. Use the Cg Standard Library" on page 324](#)), this compiles to more efficient code.
- ❑ Use the **lit()** Standard Library function, if appropriate. The **lit()** function implements a diffuse-glossy Blinn shading model. It takes three parameters:
 - ↳ The dot product of the normalized surface normal and the light vector
 - ↳ The dot product of a half-angle vector and the normal
 - ↳ The specular exponent
 It returns a 4-vector, where
 - ↳ The **x** and **w** components are always one.
 - ↳ The **y** component is equal to the diffuse dot product or to zero if the product is less than zero.
 - ↳ The **z** component is equal to the specular dot product raised to the given exponent or to zero if the diffuse dot product was less than zero.

All this is done substantially more efficiently than if the corresponding operations were written out in Cg code.

7. Take Advantage of the Different Levels of Computation Frequency

Always keep in mind the fact that fragment programs generally are executed many more times than vertex programs. Therefore, move computation from fragment programs into vertex programs whenever possible. Recall that varying outputs from vertex programs are automatically linearly interpolated before being passed to the fragment program.

There are three main cases where you can move computation from a fragment program into a vertex program:

- ❑ The result is constant over all fragments
If the vertex shader computes a value that is the same for all vertices, so that all fragments receive the same value after interpolation, any computation that the fragment shaders do that is based solely on such values can be moved to the vertex shader (as long as it doesn't require texture map lookups or other fragment-only operations).
- ❑ The result is linear across a triangle.
If the fragment shader is computing a value that varies linearly over the face of the triangle (for example, the distance from the fragment to a light source, to be used for attenuation), the value can be computed in the vertex shader at each vertex, passed to the fragment shader, and automatically interpolated by the GPU along the way.
- ❑ The result is nearly linear across a triangle.
When a value computed by a fragment shader varies slowly over triangles, it may be an acceptable approximation to compute its value at each vertex and use its linearly interpolated value in the fragment shader. For example, the usual Gouraud shading algorithm takes advantage of this situation to compute lighting per-vertex, rather than per-pixel.

In a similar manner, it may be advantageous to move any vertex shader computation that is solely dependent on the values of uniform parameters to the CPU and then to pass the result of the computation into the vertex shader with different uniform parameters. For example, if the vertex shader is passed a **float3** vector giving the direction of a distant light source, the vector should be normalized on the CPU and passed to the vertex shader. This avoids the need to repeatedly and unnecessarily recompute **normalize(lightvector)** in the vertex shader.

8. Avoid Matrix Transposes Just for Multiplication

Computing the transpose of a matrix can often be avoided. If you would like to multiply transposed **float3x3** matrix *m* by a **float3** *v*,

```
mul(v, m);
```

is equivalent to and more efficient than

```
mul(transpose(m), v);
```

9. Minimize Conditional Code in Fragment Programs

GPUs don't currently support branching in fragment programs; a program with a large amount of code that is conditionally executed—for example in an **if/else** expression—tends to run at the same speed as if all of it were executed. Therefore, if you have a large amount of conditional code *and* it is possible to evaluate the condition on the CPU, it may be advantageous to have multiple versions of the shader source code and to bind the one with the appropriate code path at run-time.

An example of this situation would be a fragment shader that supported a generic light source model for shading. Depending on how its parameters were set, it might implement a point light, a spotlight, or a light source that projected a texture map to determine the light distribution. Rather than having a series of **if/else** tests to determine which light model to use, having a separate version of the shader for each light type is generally more efficient.

Appendix D

Cg Compiler Options

This appendix describes the command-line options for the Cg compiler. What follows are the command-line options for the Cg compiler, **cgc.exe**:

- ❑ **-profile *prof***
Compile for the ***prof*** profile.
- ❑ **-profileopts *profopts***
Specify a comma-separated list of profile-specific options. See the profile specification for valid options.
- ❑ **-entry *fname***
Specify the main function name as ***fname***.
- ❑ **-o *fname***
Write the output to file ***fname***.
- ❑ **-Dmacro[=*value*]**
Define a ***macro***, with optional ***value***.
- ❑ **-I*pathname***
Specify path to an include directory.
- ❑ **-l *filename***
Write compiler messages to ***filename*** rather than to standard output.
- ❑ **-strict**
Enforce strict type checking.
- ❑ **-nofx**
Do not treat CgFX keywords as reserved words.
- ❑ **-quiet**
Suppress printing the header to **stdout**.
- ❑ **-nocode**
Compile, but do not generate any code.
- ❑ **-nostdlib**
Do not include the **stdlib.h** header file before compilation.

- ❑ **-longprogs**
Allow code generation that is longer than a profile's limit.
- ❑ **-debug**
Activate the **debug()** function.
- ❑ **-v**
Print the compiler's version to **stdout**.
- ❑ **-h**
Print a short help message.
- ❑ **-maxunrollcount *N***
Set the maximum loop unroll count to ***N***. Loops with greater than ***N*** iterations are not unrolled. Defaults to 256.
- ❑ **-posinv**
Generate a position-invariant vertex program if position invariance is supported by the current profile.

Index

A

- abs() for performance 324
- animation of geometry 202
- anisotropic lighting
 - sample shader 190
 - vertex shader code example 191
- Annotation 118
- ANSI C
 - differences from Cg 222
 - relation to Cg 221
- arbf1 profile 263
- arbv1 profile 256
- arithmetic operators 20, 248
- arithmetic precision 246
- arithmetic range 246
- array type, specification 230
- arrays
 - declaration and use of 238
 - support of 14

B

- binding semantics 242
 - defined 6
 - overview 241
- Blinn-Phong Bump-Mapping 175
- bool data type 11
- bool type, specification 229
- boolean operators 21, 248
- built-in functions 33
- bump dot3x2 diffuse and specular
 - pixel shader code example 194
 - sample shader 192
 - vertex shader code example 193
- bump-reflection mapping
 - pixel shader code example 199
 - sample shader 196
 - vertex shader code example 197

C

- C preprocessor

- supporting 241
- C++, relation to Cg 221
- Car Paint 9
 - pixel shader code example 186
 - vertex shader code example 184
- cfloat type, specification 229
- Cg
 - brief tutorial 145
 - defined 1
 - language, introduction 1
 - necessity for xiv
 - standard library functions 33
- Cg compiler
 - cgc.exe 329
 - command-line options 329
- Cg runtime
 - API specific 72
 - benefits 44
 - compiling 46
 - context creation 46
 - Direct3D 85
 - cgD3D9GetLastError() 115
 - CLError 114
 - debugging mode 112
 - error callbacks 116
 - error testing 115
 - error types 114
 - Direct3D
 - cgD3D9EnableDebugTracing() 114
 - Direct3D
 - cgD3D9TranslateHRESULT() 116
 - Direct3D expanded interface 98
 - cgD3D8LoadProgram() 103
 - cgD3D8SetSamplerState() 102
 - cgD3D9BindProgram() 105
 - cgD3D9EnableParameterShadowing() 103
 - cgD3D9GetDevice() 98
 - cgD3D9GetLatestPixelProfile() 105
 - cgD3D9GetLatestVertexProfile() 105

- cgD3D9GetOptimalOptions() 105
- cgD3D9IsParameterShadowingEnabled() 103
- cgD3D9IsProgramLoaded() 104
- cgD3D9LoadProgram() 103
- cgD3D9SetDevice() 98
- cgD3D9SetSamplerState() 102
- cgD3D9SetTexture() 102
- cgD3D9SetTextureWrapMode() 102
- cgD3D9SetUniform() 100
- cgD3D9SetUniformArray() 101
- cgD3D9SetUniformMatrix() 101
- cgD3D9SetUniformMatrixArray() 101
- cgD3D9UnloadProgram() 104
- Direct3D 8 application 109
- Direct3D 9 application 106
- Direct3D device 98
- fragment program 106
- lost devices 98
- parameters 100
 - array 101
 - sampler 102
 - uniform 100
- profile support 105
- program execution 103
- vertex program 106
- Direct3D HRESULT 114
- Direct3D minimal interface 85
 - cgD3D8ResourceToDeclUsage() 90
 - cgD3D8ValidateVertexDeclaration() 88
 - cgD3D9ResourceToDeclUsage() 90
 - cgD3D9ValidateVertexDeclaration() 88
 - Direct3D 8 application 95
 - Direct3D 9 application 92
 - fragment program 92
 - type retrieval 91
 - vertex declaration 85
 - vertex declaration for Direct3D 8 86
 - vertex declaration for Direct3D 9 86
 - vertex program 91
- header files 46
- loading 47
- modifying parameters 47
- OpenGL 73
 - error reporting 85
 - OpenGL application 82
 - OpenGL parameter setting 74
 - parameter shadowing 73
 - program execution 48
 - releasing resources 49
- Cg Runtime Library
 - overview 45
- Cg standard library 33
- Cg_Simple file 145
- cgc.exe, Cg compiler 329
- cgD3D9EnableParameterShadowing() 103
- CLError
 - Direct3D 114
 - OpenGL 85
- cint type, specification 229
- command-line options, Cg compiler 329
- comparison operators 248
 - introduction 21
- compilation profiles, use of 225
- compiler options
 - command-line 329
 - debug 330
 - Dmacro 329
 - entry 329
 - h 330
 - Ipathname 329
 - I filename 329
 - longprogs 330
 - maxunrollcount 330
 - nocode 329
 - nofx 329
 - nostdlib 329
 - o 329
 - profile 329
 - profileopts 329
 - quiet 329
 - strict 329
 - v 330
- compile-time type category 232
- computation frequency for performance 327
- concrete type category 232
- conditional code in fragment programs and performance 328
- conditional operator 248

- conditional operators 22
- constants, typing of 232
- construction operator, described 244
- context
 - core Cg 50
- control constructs used 19
- core Cg context 50
- Core Cg error reporting 71
- Core Cg parameter 54
- Core Cg program 50
- core Cg runtime 49

D

- data types
 - bool 11
 - fixed 11
 - float 11
 - half 11
 - int 11
 - sampler 11
 - supported 11
- data types for performance 325
- debugging function 41
- declaration, Cg definition 224
- definition, as used in Cg 224
- derivative functions 41
- Direct3D Cg runtime 85
 - cgD3D9EnableDebugTracing() 114
 - cgD3D9GetLastError() 115
 - cgD3D9TranslateHRESULT() 116
 - CGError 114
 - debugging mode 112
 - error callbacks 116
 - error testing 115
 - error types 114
 - expanded interface 98
 - cgD3D8LoadProgram() 103
 - cgD3D8SetSamplerState() 102
 - cgD3D9BindProgram() 105
 - cgD3D9EnableParameterShadowing() 103
 - cgD3D9GetDevice() 98
 - cgD3D9GetLatestPixelProfile() 105
 - cgD3D9GetLatestVertexProfile() 105
 - cgD3D9GetOptimalOptions() 105

- cgD3D9IsParameterShadowingEnabled() 103
- cgD3D9IsProgramLoaded() 104
- cgD3D9LoadProgram() 103
- cgD3D9SetDevice() 98
- cgD3D9SetSamplerState() 102
- cgD3D9SetTexture() 102
- cgD3D9SetTextureWrapMode() 102
- cgD3D9SetUniform() 100
- cgD3D9SetUniformArray() 101
- cgD3D9SetUniformMatrix() 101
- cgD3D9SetUniformMatrixArray() 101
- cgD3D9UnloadProgram() 104
- Direct3D 8 application 109
- Direct3D 9 application 106
- Direct3D device 98
- fragment program 106
- lost devices 98
- parameters 100
 - array 101
 - sampler 102
 - uniform 100
- profile support 105
- program execution 103
- vertex program 106
- HRESULT 114
- minimal interface 85
 - cgD3D8ResourceToDeclUsage() 90
 - cgD3D8ValidateVertexDeclaration() 88
 - cgD3D9ResourceToDeclUsage() 90
 - cgD3D9ValidateVertexDeclaration() 88
 - Direct3D 8 application 95
 - Direct3D 9 application 92
 - fragment program 92
 - type retrieval 91
 - vertex declaration 85
 - vertex declaration for Direct3D 8 86
 - vertex declaration for Direct3D 9 86
 - vertex program 91
- Direct3D debug DLL, using 113
- DirectX pixel shader 1.x profiles 308
- DirectX pixel shader 2.x profile 300
- DirectX vertex shader 1.1 profile 304

DirectX vertex shader 2.x profile 296
 dot() for performance 324
 dx8ps profile, deprecated 308

E

effect 117
 Effect parameter 118
 effect parameters 121
 evaluating Cg programs 127
 explicit casts
 compile-time 235
 numeric 236
 numeric matrix 236
 numeric vector 236

F

fixed data type 11
 fixed type, specification 229
 float data type 11
 float type, specification 229
 floating type category 232
 for statements 244
 fp20 profile 283
 fp30 profile 274
 fragment profiles
 texture lookups 23
 fragment program 121
 predefined output structures 42
 varying output 9
 fragment program profiles 252
 OpenGL ARB 263
 OpenGL NV_fragment_program 274
 fragment program, defined 3
 fresnel 200
 sample shader 200
 vertex shader code example 200
 function
 calls 228
 multiplying 20
 open profile 227
 function definitions
 introduction 19
 function overloading 240
 introduction 19
 functions

debugging 41
 declaring 226
 derivative 41
 geometric 38
 mathematical 33
 overloading by profile 226
 standard library 33
 texture map 38

G

geometric functions 38
 GL_ARB_vertex 256
 global variables 241
 graphics hardware, evolution of xiii
 grass
 sample shader 202
 vertex shader code example 202

H

half data type 11
 half type, specification 229

I

if statements 244
 inputs
 uniform 5
 varying 5, 6
 int data type 11
 int type, specification 229
 integral type category 232
 interfaces 125

J

Java, relation to Cg 221

L

language profiles
 concept of 3

M

mathematical functions 33
 matrices, multiplying 20
 matrices, support of 12
 matrix palette skinning 217

- sample shader 217
- vertex shader code example 218
- matrix transposes and performance 328
- melting paint
 - pixel shader code example 163
 - sample shader 161
 - vertex shader code example 161
- min() for performance 324
- miscellaneous operators 249
- modifiable function parameters, passing 19
- multipaint
 - pixel shader code example 167
 - sample shader 165
 - vertex shader code example 166

N

- namespaces 237
- numeric type category 232

O

- object, Cg definition 224
- open profile functions 227
- OpenGL Cg runtime 73
 - error reporting 85
 - OpenGL application 82
 - parameter setting 74
- OpenGL CLError 85
- OpenGL profiles
 - ARB fragment program 263
 - ARB vertex program 256
 - NV_fragment_program 274
 - NV_register_combiners 283
 - NV_texture_shader 283
 - NV_vertex_program 279
 - NV_vertex_program 2.0 270
- operations
 - expressed differently from C 222
- operator
 - enhancements 247
 - precedence 247
- operators
 - arithmetic 20
 - boolean 21
 - conditional 22
 - introduction 18

- swizzle 22
- write-mask 22

P

- packed, type modifier 230
- parameter shadowing 73
- parameters
 - modifiable function, passing 19
- parameters in function definitions, syntax 227
- pass 117, 120
- pass state 120
- performance techniques
 - abs() 324
 - avoiding matrix transposes 328
 - computation frequency 327
 - conditional code in fragment
 - programs 328
 - data types 325
 - dot() 324
 - min() 324
 - saturate() 324
 - shading computations 326
 - swizzle 323
 - texture maps 324
 - vectorization 321
- pixel program, defined 3
- pixel shader, defined 3
- position invariance 250
- profile
 - arbfp1 263
 - arbvp1 256
 - fp20 283
 - fp30 274
 - ps_1_1, ps_1_2, ps_1_3 308
 - ps_2_0, ps_2_x 300
 - vp20 279
 - vp30 270
 - vs_1_1 304
 - vs_2_0, vs_2_x 296
- profile, defined 3
- program
 - declaring 5
 - kinds of inputs 5
- program profiles
 - fragment 252

- vertex 250
- programming model, GPU 2
- ps_1_x profile 308
- ps_2_0 profile 300
- ps_2_x profile 300

R

- ray-traced refraction
 - pixel shader code example 172
 - sample shader 170
 - vertex shader code example 171
- recursion, function 19
- reflection vector 200
- refraction
 - pixel shader code example 207
 - sample shader 205
 - vertex shader code example 206
- release notes xvi
- Renderman, relation to Cg 221
- reserved words 249
- runtime
 - core Cg 49

S

- sampler data type 11
- sampler type, specification 230
- samplers 123
- saturate() for performance 324
- scalar type category 232
- semantics
 - aliasing 243
 - restrictions 243
- shader sample
 - anisotropic lighting 190
 - bump dot 3x2 diffuse and specular 192
 - bump-reflection mapping 196
 - fresnel 200
 - grass 202
 - improved skinning 154
 - improved water 157
 - matrix palette skinning 217
 - melting paint 161
 - multi-paint 165
 - ray-traced refraction 170
 - refraction 205

- shadow mapping 208
- shadow volume extrusion 211
- sine wave demo 214
- skin 175
- shader, simple.cg example 146
- shaders
 - advanced profile samples 153
 - basic profile samples 189
- shading computations for performance 326
- shadow mapping 208
 - pixel shader code example 210
 - sample shader 208
 - vertex shader code example 209
- shadow volume extrusion
 - sample shader 211
 - vertex shader code example 212
- shadow volumes 211
- silent incompatibilities with C 221
- simple.cg
 - basic transformations 149
 - passing arguments 149
- Sine function 202, 214
- sine wave demo
 - sample shader 214
 - vertex shader code example 215
- sinh(x) 37
- skin
 - pixel shader code example 175
 - sample shader 175
- skinning, improved
 - sample shader 154
 - vertex shader code example 155
- smearing, scalar to vector 237
- Stanford shading language, relation to Cg 221
- State assignment 118
- statements
 - introduction 18
- statements, in Cg 244
- structures
 - introduction 13
- swizzle
 - for performance 323
- swizzle operator 22
- swizzle operator, described 245

T

- technique 117
- technique validation 120
- texture lookups 23
- texture map functions 38
- texture maps for performance 324
- textures 123
- thin film effect
 - pixel shader code example 182
 - vertex shader code example 180
- tutorial 145
- type conversions 12, 234
 - array 235
 - matrix 234
 - scalar 234
 - structure 235
 - vector 234
- type equivalency 236
- type promotion 236
 - assignment 237
 - smearing 237
- type qualifiers 233
 - const 233
 - in 233
 - out 233
- types
 - general discussion 229
 - partial support 231

U

- uniform inputs 5
- uniform modifier, use of 225
- uninitialized variables, use of 241
- unsized arrays 125

V

- variables
 - global 241
 - uninitialized, use of 241
- varying inputs 5, 6
- vector data types 12
- vector operators, new 244
- vectorization
 - for performance 321
- vectors, constructing 21

- vertex color 149
- vertex position 149
- vertex program 121
 - varying output 7
- vertex program profiles 250
- vertex programs, defined 3
- virtual machine 127
- void type, specification 229
- vp20 profile 279
- vp30 profile 270
- vs_1_1 profile 304
- vs_2_0 profile 296
- vs_2_x profile 296

W

- water, improved
 - pixel shader code example 160
 - sample shader 157
 - vertex shader code example 158
- web site, NVIDIA xvi
- while statements 244
- workspace, loading 145
- write-mask operator 22
 - described 246

