




Cg Toolkit

Cg 2.1 beta
August 2008
Release Notes



Cg Toolkit Release Notes

The Cg Toolkit allows developers to write and run Cg programs using a wide variety of hardware and OS platforms and graphics APIs. Originally released in December 2002, the Toolkit now supports over 20 different DirectX and OpenGL profile targets. It provides a compiler for the Cg language, runtime libraries for use with the OpenGL and DirectX graphics APIs, support for the CgFX effect files, example applications, and extensive documentation.

The 2.1 beta release of Cg incorporates the following updates:

- ❑ New translation profiles for DirectX10 Shader Model 4
- ❑ Shader source virtual file system for compilation
- ❑ Callback for providing shader source included by the compiler
- ❑ Improved handling of generic attributes by the GLSL profile
- ❑ Performance improvements and bug fixes
- ❑ New examples including:
 - ❑ examples/D3D10/basic
 - ❑ 02_vertex_and_fragment_program
 - ❑ 03_uniform_parameter
 - ❑ 04_varying_parameter
 - ❑ 05_texture_sampling
 - ❑ 06_vertex_twisting
 - ❑ 07_two_texture_accesses
 - ❑ cgfx_buffer
 - ❑ cgfx_simple
 - ❑ cgfx_textureexamples
 - ❑ /OpenGL/basic (from *The Cg Tutorial*)
 - ❑ 23_bump_map_floor:
 - ❑ examples/OpenGL/advanced
 - ❑ generic_attribs:
 - ❑ include_string:

Visit the NVIDIA Cg website at developer.nvidia.com/page/cg_main.html for complete availability and compatibility information.

Bug reports, issues, and feedback can be sent to cgsupport@nvidia.com.

Supported OS/Hardware Platforms

Cg is available for these platforms:

- ❑ Windows 32
- ❑ Windows 64
- ❑ Linux x86
- ❑ Linux x86-64
- ❑ MacOS 10.4 (Tiger)
- ❑ MacOS 10.5 (Leopard)
- ❑ Solaris 10 x86

The Cg Runtime libraries include:

- ❑ The Cg core runtime library for managing parameters and loading programs
- ❑ The CgGL runtime library for OpenGL based applications
- ❑ The CgD3D10 runtime library for DirectX 10 based applications
- ❑ The CgD3D9 runtime library for DirectX 9 based applications
- ❑ The CgD3D8 runtime library for DirectX 8 based applications

Supported Profiles

The Cg compiler currently supports the following hardware profiles:

OpenGL

- ❑ **gpu_gp** NV_geomentry_program4
- ❑ **gpu_vp** NV_vertex_program4
- ❑ **gpu_fp** NV_fragment_program4
- ❑ **glslv** OpenGL Shading Language (GLSL) for OpenGL 2.0 vertex shader
- ❑ **glslf** OpenGL Shading Language (GLSL) for OpenGL 2.0 fragment shader
- ❑ **arbvp1** ARB_vertex_program 1.0
- ❑ **arbf1** ARB_fragment_program 1.0
- ❑ **vp40** ARB_vertex_program + NV_vertex_program2 option
- ❑ **fp40** ARB_fragment_program + NV_fragment_program2 option
- ❑ **vp30** NV_vertex_program 2.0
- ❑ **fp30** NV_fragment_program 1.0
- ❑ **vp20** NV_vertex_program 1.0
- ❑ **fp20** NV_register_combiners and NV_texture_shader

DirectX 10.0

- ❑ **vs_4_0** HLSL10 Vertex Shader
- ❑ **ps_4_0** HLSL10 Fragment Shader

DirectX 9.0c

- ❑ **hlslv** HLSL9 Vertex Shader
- ❑ **hlslf** HLSL9 Fragment Shader
- ❑ **vs_3_0** Vertex Shader 3.0
- ❑ **ps_3_0** Pixel Shader 3.0

DirectX 9

- ❑ **vs_2_x** Extended Vertex Shader 2.0
- ❑ **ps_2_x** Extended Pixel Shader 2.0
- ❑ **vs_2_0** Vertex Shader 2.0
- ❑ **ps_2_0** Pixel Shader 2.0

DirectX 8 & 9

- ❑ **vs_1_1** Vertex Shader 1.1
- ❑ **ps_1_3** Pixel Shader 1.3
- ❑ **ps_1_2** Pixel Shader 1.2
- ❑ **ps_1_1** Pixel Shader 1.1

Improvements & Bug Fixes

Improvements

- ❑ TEXCOORD8-15 are now available in SM 3.0 profiles.
- ❑ Added more 'clip' functions for improved compatibility with HLSL.
- ❑ cgGetProgramDomainProgram supports iterating over subprograms within a combined program object.
- ❑ cgGetParameterResourceName is particularly useful with the translation profiles.

Improvement: Documentation

- ❑ **Note:** The Cg Users Manual has **not** been updated for this release.
- ❑ New reference pages for the new Runtime API routines have been added.

Bug Fixes

- ❑ Various bugs have been fixed.

Compatibility Notes

There aren't any known compatibility issues with programs written against Cg 2.0. For programs written against Cg 1.5 or earlier, refer to the Compatibility Notes section of the release notes for Cg 2.0.

Known issues

Known runtime issues

- ❑ **cgGetParameterValues** incurs a penalty in both performance and memory footprint and its use is strongly discouraged. Use one of the **cgGetParameterValue** or **cgGetParameterDefaultValue** functions instead.
- ❑ **cgCopyProgram** and **cgCopyEffect** do not work.
- ❑ Loading precompiled code via **CG_OBJECT** in **cgCreateProgramFromFile** doesn't work for shaders which use semantic type modifiers.
- ❑ The DirectX 8 runtime does not support Cg interfaces.
- ❑ The Cg runtime does not support creating shared parameters containing varying members.
- ❑ Unsized arrays and interface parameters cannot currently be used on the right-hand side of state assignments. Doing so will trigger an error.
- ❑ Values set by **cgGLSetOptimalOptions(...)** can be un-set after a call to **cgDestroyContext()**. As a work around, call **cgGLSetOptimalOptions()** after each call to **cgDestroyContext()** when more Cg contexts are going to be created.

Known compiler issues

- ❑ Long shader programs that make heavy use of interfaces may still result in very long compilation time.
- ❑ Very little error checking is performed on the OpenGL state semantics string (**state.***); it is just copied to the output assembly. As a result, a typo in the string may compile correctly, and no error will be apparent until the application attempts to load the assembly shader.
- ❑ Error reporting: Some error and warning messages are not as clear as they could be. Some of the issues to be aware of are:
 - ❑ Reported line numbers do not match source code lines when standard library functions are being used
 - ❑ In some cases, errors are not reported in the order they appear in the program
 - ❑ Errors are not reported when constants are out of range for untyped constants.
- ❑ Side-effects in conditional expressions ('?:') and logical expressions ('&&' and '||') are always evaluated, regardless of the condition, as specified in the Cg language specification. Hence developers need to watch out for this case.
- ❑ At most one binding semantic per uniform variable is supported by the compiler. Multiple profile-specific binding semantics per uniform variable are not supported.

- ❑ Only loops with a single induction variable are unrolled. Loops that require more than 1 induction variable will fail to compile on older profiles that do not support loops.
- ❑ Local variable arrays which are written to in one block of code, and then read via a non-constant index in a different block will fail to compile on older hardware that does not support this feature. Current hardware supports this feature.
- ❑ Invalid Cg programs can, at times, generate invalid code, instead of a compiler error.

Known profile-specific issues

- ❑ The ps2* profiles do not support MRTs
- ❑ Because the underlying hardware support for the fp20 and ps_1_* profiles is quite limited and inflexible, it isn't always possible to compile even seemingly simple Cg programs under these profiles. For more details on these limitations, please see the NV_register_combiners and NV_texture_shader OpenGL extension specifications, or the DirectX PixelShader 1.* specifications.
- ❑ The FOG varying input semantic is not yet supported under the fp20 profile.

New API

Cg 2.1 adds new API for controlling include file processing by the compiler and for getting default parameter values. Here is the complete list of new routines:

cgSetCompilerIncludeString	cgGetParameterDefaultValuefc
cgSetCompilerIncludeFile	cgGetParameterDefaultValueir
cgSetCompilerIncludeCallback	cgGetParameterDefaultValueic
cgGetCompilerIncludeCallback	cgGetPassProgram
cgGetParameterDefaultValuedr	cgGetProgramDomainProgram
cgGetParameterDefaultValuedc	cgGetParameterResourceName
cgGetParameterDefaultValuefr	cgGLUnloadProgram

Release Types

Cg is released in two forms:

1. The Cg Toolkit provides a complete Cg Software Development Kit (SDK) including documentation, examples, standalone compiler, headers and libraries.
2. Cg Binary Distributions provide updated redistributable libraries that Cg-based applications can ship with.

SDK versions are released as platform specific installers containing the full toolkit (libraries, documentation, examples, etc.) They can be downloaded from

http://developer.nvidia.com/object/cg_toolkit.html

Binary distributions contain only the libraries, and all supported platforms are bundled in a single file. The libraries supplied in a binary distribution should be feature-for-feature and bug-for-bug compatible across all the platforms supported by a given distribution (meaning are all compiled from the same source code). Cross-platform software vendors are encouraged to redistribute Cg libraries from a single binary distribution to minimize platform variances in Cg.

Cg binary distributions can be found at

<http://developer.nvidia.com/object/cg-redistributable-binaries.html>

Distribution License

The docs directory contains a file `Cg_Redist_License.pdf` providing a non-exclusive, world-wide, royalty free licensee for redistributing Cg with your applications. See this license for details.

Release History

The following table summarizes release dates and library versions for Cg releases:

Cg Release Name	Release Date	Library Version
binary1	06/05/08	2.0.0016
SDK3	04/15/08	2.0.0015
SDK2	01/30/08	2.0.0012
SDK1	12/12/07	2.0.0010
beta	11/01/07	2.0.0008

The Cg library version is returned by `cgGetString(CG_VERSION)`

Change History

2.0.0016

- ❑ Fixed array uploading in GLSL profile
- ❑ Allow 'Texture' as a variable name
- ❑ Improved handling of complex CgFX files

2.0.0015

- ❑ Added support for samplerBUF to the runtime
- ❑ Fixed a bug in setting and getting string annotation values
- ❑ Fixed a crash at exit when running on 64-bit Windows
- ❑ Fixed a problem with `cgGLSetOptimalOptions` on the GPU Program4 profiles
- ❑ Fixed a problem with using 'discard' on MacOS

2.0.0012

- ❑ Added section on `cgc` and completed the `gp4` profile documentation
- ❑ The set and get parameter routines now silently ignore "extra" components. For example, if `cgSetParameter4f` is used to send 4 values to a `CG_FLOAT2` parameter, the first two values are used but the last two are ignored.
- ❑ The HLSL profiles can now run on SM2.0 and SM2.x hardware
- ❑ Fixed a multitexture sampler configuration bug with GLSL profiles in CgFX
- ❑ Fixed problems with using buffer semantics on pre-G80 class GPUs
- ❑ Fixed a bug with the path variable after installing Windows 64-bit dlls
- ❑ Added support for OSX x86_64



Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation.

Microsoft, Windows, the Windows logo, and DirectX are registered trademarks of Microsoft Corporation.

OpenGL is a trademark of SGI.

Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

Copyright © 2004-2008 NVIDIA Corporation. All rights reserved.



NVIDIA.

NVIDIA Corporation
2701 San Tomas Expressway
Santa Clara, CA 95050
www.nvidia.com