



Cg Toolkit

Cg 1.4.1
March 2006
Release Notes

Cg Toolkit Release Notes

The Cg Toolkit allows developers to write and run Cg programs using a wide variety of hardware 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, runtime libraries for CgFX, example applications, and extensive documentation.

This 1.4.1 release of Cg stabilizes several significant changes and features first introduced by Cg 1.4:

- ❑ Significantly improved code generation in most of the profiles
- ❑ Complete rewrite of CgFX, including an API integrated with the Cg runtime API, support for extensible, cross-platform graphics state management, and support for unsized arrays and interfaces.
- ❑ Substantial improvements in compiling many long shaders, especially those that make heavy use of interfaces.
- ❑ The Cg runtime is now substantially faster in many cases
- ❑ Cg 1.4.1 now ships with native implementations for Win32, Win64, Linux (32-bit and 64-bit), and MacOS 10.3 (Panther) and 10.4 (Tiger)

Note that some of this new functionality is currently only supported when using the CgGL OpenGL runtime library. The Direct3D-specific Cg runtime libraries currently do not support shared parameters, for example, nor does this release ship with a native D3D state manager for CgFX. This release also does not provide Direct3D profiles for Shader Model 3.0.

With minor exceptions, Cg 1.2 or Cg 1.3 applications should work with Cg 1.4.1 without the need to recompile the program. See “Compatibility Notes,” below, for more information.

Applications and effects that used previous versions of CgFX will require modification. Applications that use CgFX will have to be modified to use the new CgFX API. If the supplied OpenGL state manager for CgFX is used, the effect files themselves must be modified. The effects must be changed to use OpenGL state assignments, rather than the D3D state assignments supported in prior versions of CgFX. A subsequent release of Cg will provide D3D and cross-platform state managers for CgFX.

Cg is available for a wide variety of hardware and OS platforms (as mentioned above). Please visit the NVIDIA Cg website at developer.nvidia.com/Cg for complete availability and compatibility information.

Please report any bugs, issues, and feedback to NVIDIA by email at cgsupport@nvidia.com. We will expeditiously address any reported problems.

Compatibility Notes

Although the 1.4.1 release of Cg is generally compatible with previous releases, several improvements and other changes may affect existing applications. This section details these potential compatibility issues.

In the 1.2 and 1.3 releases of Cg, the Cg compiler would always generate generic attribute names (e.g., “ATTR0”) for varying input parameters, even if conventional semantic names (e.g., “POSITION”) were specified under OpenGL profiles. In this 1.4.1 release, the type of attribute now matches the type of semantic. As a result, applications that explicitly assumed and checked for generic attribute names may need to be modified to handle both generic and conventional attributes.

In this release, the Cg compiler uses an improved, more efficient constant register allocation scheme. As a result, uniform parameters in compiled programs may no longer reside in the same constant registers as they had in previous releases of Cg. Applications that (incorrectly) assume a particular constant register numbering or order should be modified to either remove this assumption, or to use register allocation semantics (e.g., “register(C0)”) to implement the layout assumed by the application.

CgFX has been redesigned and re-implemented. As noted above, existing applications that use CgFX must be modified if the new API is to be used. In addition, if you wish to use the supplied “native” OpenGL state manager, which assumes that OpenGL-style state names and values are used in effects, existing effects must be modified to use OpenGL-style states. See the “Introduction to CgFX” chapter in the Cg Users Manual for more details. If you have questions or problems related to porting existing CgFX applications, please contact NVIDIA developer relations.

Supported Profiles and Platforms

The Cg compiler currently supports the following hardware profiles:

OpenGL

- arbv1 (ARB_vertex_program 1.0)
- arbf1 (ARB_fragment_program 1.0)
- vp40 (ARB_vertex_program + NV_vertex_program2 option)
- vp30 (NV_vertex_program 2.0)
- fp40 (ARB_fragment_program + NV_fragment_program2 option)
- fp30 (NV_fragment_program 1.0)

- ❑ vp20 (NV_vertex_program 1.0)
- ❑ fp20 (NV_register_combiners and NV_texture_shader)

DirectX 8 & 9

- ❑ vs_1_1 (Vertex Shader 1.1)
- ❑ ps_1_1, ps_1_2 and ps_1_3 (Pixel Shader 1.1, 1.2, 1.3)

DirectX 9

- ❑ vs_2_0 and vs_2_x (Vertex Shader 2.0 and Extended VS 2.0)
- ❑ ps_2_0, and ps_2_x/ps_2_a (Pixel Shader PS 2.0 and Extended PS 2.0)

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 CgD3D8 runtime library for DirectX 8 based applications
- ❑ The CgD3D9 runtime library for DirectX 9 based applications
- ❑ CgFX has been incorporated directly into the Cg Runtime libraries

Improvements & Bug Fixes

Improvements

- ❑ CgFX has been completely rewritten
- ❑ Many changes to the infrastructure in both the compiler and runtime have been made in order to improve both performance and reliability

Improvement: CgFX rewrite

CgFX has been completely rewritten. Major goals of this rewrite:

- ❑ Create a more reliable base of software
- ❑ Improve efficiency of the CgFX Runtime
- ❑ Provide a means of supporting extensible, cross-graphics-API states and state managers
- ❑ Create a better mapping between CgFX files and OpenGL state, for those applications that focus exclusively on OpenGL
- ❑ Allow for compile-time execution of Cg expressions through a Cg virtual machine
- ❑ Portability across more platforms

CgFX is now more stable, faster, and feature-rich.

NOTE: This version of CgFX is not API-compatible, nor effect-file-compatible, with previous releases. NVIDIA is committed to supporting developers who have been using previous versions of CgFX. If you need help porting your application or effects to the new version of CgFX, please contact your NVIDIA Developer Support contacts.

Improvement: Cg Runtime efficiency

The Cg Runtime has been tuned to remove a significant amount of overhead from the runtime calls. In many cases, CPU overhead in the Cg runtime has been reduced on the order of 50%.

Improvement: Cg compile times

Many Cg programs should now compile faster than in previous releases. This is especially true of long shaders, and shaders that make heavy use of Cg's "interfaces" feature. Further improvements are still being worked on for a future Cg release.

Improvement: Windows Installer

The Cg 1.4.1 installer for Windows is thoroughly updated and now based on Inno Setup (rather than InstallShield). There is no longer a separate Cg 64-bit installer for Windows. The single Cg installer can optionally install 64-bit libraries for x86-64 systems. The new installer will automatically first remove a prior InstallShield-based installer.

The Cg 1.4.1 installer automatically makes the necessary updates to your `Path`, `CG_BIN_PATH`, `CG_INCLUDE_PATH`, and `CG_LIB_PATH` environment variables.

The Cg 1.4.1 installer has more control over the selected components for installation. For example, you can skip installation of documentation, examples, or 64-bit support. Everything is installed by default except the 64-bit libraries and executables.

A selectable install component (installed by default) provides Visual Studio integration for Cg for Visual C 6.0 and Visual Studio .NET 7.1.

Improvement: Documentation

- ❑ The Cg 1.4.1 Users Manual in PDF format is included in the install at `docs/CgUsersManual.pdf` or navigate the Windows Start menu to find *Start->All Programs->NVIDIA Corporation->Cg Toolkit->User's Manual*.
- ❑ Reference manual pages for the Cg runtime API, profiles, and Cg standard library are now included in HTML and PDF form. Find the HTML pages in the `docs/html` directory. Find the `CgReferenceManual.pdf` in the `docs` directory or navigate the Windows Start menu to find *Start->All Programs->NVIDIA Corporation->Cg Toolkit->Reference Manual* (select the PDF icon).
- ❑ The Microsoft HTML Help file is updated to include profile and standard library documentation. Navigate the Windows Start menu to find *Start->All Programs->NVIDIA Corporation->Cg Toolkit->Reference Manual* (select the HTML Help icon).

Removed features

- ❑ We have dropped support for the previous CgFX API
- ❑ We no longer provide debug `cgD3D8d.dll` and `cgD3D9d.dll` libraries.

Bug Fixes

- ❑ In CgFX, too many bugs have been fixed to list them here.
- ❑ All known memory leaks in CgFX and the Cg Runtime have been fixed
- ❑ The determinant standard library routine for a 4x4 matrix is correct now.
- ❑ Issues constant loading issues with the fp20 profile are resolved.

Known issues

Known runtime issues

- ❑ The DirectX 8 and DirectX 9 runtimes have not yet been updated to support the Cg 1.2 'interfaces' feature.
- ❑ There is currently no native CgFX support under D3D.
- ❑ The Cg runtime does not currently support created 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.
- ❑ The 'cg_explicit' and 'cgGL_explicit' libraries are not currently supported.
- ❑ Values set by `cgGLSetOptimalOptions(...)` can be un-set after a call to `cgDestroyContext()`. To work around this issue, `cgGLSetOptimalOptions()` should be called after each call to `cgDestroyContext()`, if more Cg contexts are going to be subsequently created.
- ❑ Error reporting should be improved in many cases.
- ❑ The only CgFX "state manager" provided in this release supports OpenGL-style state names and values, and only runs under OpenGL. State managers that support D3D-style state names and values, and cross-API state managers, will be provided in a subsequent release of Cg.
- ❑ More OpenGL state still needs to be exposed through CgFX state assignments. If you have specific feedback about missing state, please speak to your Developer Support contact.

Known compiler issues

- ❑ The 1.4.1 release of Cg is more accurate at detecting when program parameters are not used.
- ❑ Long shader programs that make heavy use of interfaces may still see very long compiler times.
- ❑ 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, and currently warnings are not always issued. Hence developers need to watch out for this case.
- ❑ Only one return statement is allowed per function. There is an error issued if there is any unreachable code. Return statements in if/for blocks are not supported.
- ❑ All matrices are assumed to be row-major. Currently, column-major matrices are not supported.
- ❑ 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 single induction variables are unrolled. Loops that require more than 1 induction variable will fail to compile.
- ❑ 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 cause the compiler to crash. Current hardware does not support this feature, but the compiler should not crash.
- ❑ Invalid Cg programs can, at times, generate invalid code, instead of a compiler error.

Known profile-specific issues

- ❑ The ps2* profiles do not yet support MRTs
- ❑ Because the underlying hardware support for the fp20 and ps_1_* profiles is quite limited and inflexible, it is not always possible to compile seemingly simple Cg programs under these profiles. For more details on these limitations, please see the NV_register_combiners, 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. (also true in Cg 1.2 and 1.3)



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 © 2004-2006 NVIDIA Corporation. All rights reserved.



NVIDIA.

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