

1 What's New

Installer

The AMD APP SDK can be installed on Linux with root as well as non-root permissions.

New Samples

- **HeatPDE:** This sample demonstrates simulating heat field using shared virtual memory feature of OpenCL. The host and the device interactively access the same virtual memory and use platform atomics for ensuring memory consistency. The sample also graphically shows the simulation of heat field using OpenGL APIs.
- **AdvancedConvolution:** This sample demonstrates the optimal implementation of separable and non-separable filters. The sample implements 3X3 and 5X5 Sobel, Box, and Gaussian Filters.

Updated Samples

- **SVMAtomicsBinaryTreeInsert:** The sample has been enhanced to show visualization of the nodes inserted. It shows nodes in different colors depending on whether it contains child nodes inserted by host only (blue), device only (Red), both host and device (Green) or a leaf node (Gray).
- **SimpleConvolution:** The sample has been extended to include separable filter in addition to the existing non-separable filter.
- **NBody:** The sample now prints performance metric in GFLOPS.

Restructured OpenCL Samples Directory

OpenCL samples have now been restructured such that directory in which they are located is based on the OpenCL version features that they use.

The samples that use 2.0 features are located in the `\\samples\opencl\cl\2.0` directory.

The samples that use 1.x features are located in the `\\samples\opencl\cl\1.x` directory and the `\\samples\opencl\cpp_cl\1.x` directory.

APARAPI Samples Excluded

The Aparapi samples are not included in the latest APP SDK. For reference on usage of Aparapi, developers can see earlier versions of the AMD APP SDK (now AMD APP SDK) or visit <https://code.google.com/p/aparapi/>.

2 Important Notes

- For a successful build and correct operation of individually downloaded samples, update to the AMD Catalyst 14.501 driver or newer.
- AMD APP SDK samples are located in the `$(AMDAPPSDKSamplesInstallPath)\samples` folder. See the *AMD APP SDK Installation Notes* and the *AMD APP SDK Getting Started Guide* for more information. These documents also provide information about using OpenCL, BOLT, C++ AMP, and OpenCV samples.
- The SDK samples provided with this release of the AMD APP SDK are not necessarily tuned for optimal performance. AMD is improving the samples continually; check <http://developer.amd.com/AMDAPPSDK> for new and updated samples.
- OpenCL 2.0 runtime support is limited to 64-bit applications running on 64-bit Windows and Linux operating systems only.
- 32-bit AMD APP SDK installation on 64-bit Linux is not supported.
- The minimum CMake version required is 2.8.0.
- When building on Linux systems, the GL samples require the development files for OpenGL and for the OpenGL utility library. If these are not already installed on your system, install them on your system in order for the sample to build. If required, create a symlink `libGLU.so` to `libGLU.so.1`.
- On Windows, the SDK samples have Microsoft® Visual Studio® 2010 and 2012 projects. To generate Microsoft® Visual Studio® 2013 project files, use the included CMake files.
- All the samples in `$(AMDAPPSDKSamplesInstallPath)/samples/opengl/benchmark`, as well as GL-based, DirectX, and MultiGPU samples do not work on CPU-only machines.
- Bolt and OpenCV-CL libraries make use of OpenCL kernels that require a workgroup size of 256. Hence AMD APP SDK Bolt and OpenCV-CL samples can be run only on devices that support a workgroup size of 256 or above.
- If GL samples fail on Linux, ensure that the `libGL.so.1` file is linked to `fglrx-libGL.so.1.2`, which for Ubuntu is found in `/usr/lib/fglrx/` and `/usr/lib32/fglrx/`, or for RHEL is found in `/usr/lib64/fglrx/`. The sample fails when linked to `/usr/lib/x86_64-linux-gnu/mesa/libGL.so.1`, `/usr/lib/i386-linux-gnu/mesa/libGL.so.1` in Ubuntu or `/usr/lib/libGL.so.1`, `/usr/lib64/libGL.so.1` in RHEL. For more details, see: <http://phoronix.com/forums/showthread.php?7351-Does-fglrx-s-libGL-so-1-2-have-wrong-soname>
<https://bugs.launchpad.net/ubuntu/+source/mesa/+bug/943162>

3 Resolved Issues and Improvements

- Memory leaks in several OpenCL samples are fixed.
- RegionGrowingSegmentation: Updated the algorithm to reduce the % of mismatch in cpu and gpu.
- Boundary value checks and memory size limit checks in several samples: MonteCarloAsian, BufferBandwidth.
- BasicDebug: Fixed the out-of-boundary write into local memory.
- PrefixSum: Fixed the incorrect output generated when input length is 2.

- FineGrainSVM: Fixed issue of soft hang at input size > 16000.
- LDSBandwidth: Modified kernels to accommodate 128 work-group size.

4 Known Issues

- For correct operation of the samples when using RHEL 5.5 or RHEL 5.8, it may be necessary to build the samples using that OS.
- SimpleDX9, SimpleDX10, and SimpleDX11 samples: These samples do not operate correctly when using the MinGW compiler.
- Running samples can result in an error if the `TEMP` environment variable contains multi-byte characters.
- The OpenCL NBody may report differences when run with the `--verify` option and a higher number of iterations due to precision differences between the OpenCL and C functions.
- The FluidSimulation2D and NBody samples use the `glut` library; thus, they may crash if the "X" button is used to close the application window. This is a known issue in glut64 for all Windows 64-bit operating systems.
- For MinGW64 issues, please read KnowledgeBase article KB125:
<http://developer.amd.com/resources/documentation-articles/knowledge-base/>
- Increased error tolerance due to reduced accuracy when using native versions of SIN and COS functions on the AMD Radeon™ HD 6970 and AMD Radeon™ HD 6950 relative to other AMD GPUs.
- Bolt samples, C++AMP samples, as well as SimpleDX10 and FluidSimulation2D cannot be built using the Intel C Compiler.
- Samples that include `cl.hpp` do not build on MinGW. These samples are in the `cpp_cl` folder.
- On a machine with [Intel/AMD CPU + integrated Intel/AMD GPU] + discrete AMD GPU, C++ AMP samples do not run on the discrete GPU when the Power-Express(PX) mode is enabled. This limitation holds true for any DirectX samples as well. On a desktop, to enable running on discrete GPU, connect 2 monitors to both iGPU and dGPU ports. This way DirectX will detect both the cards in power-express mode. In case of a mobile device, such as a laptop, disable iGPU using the BIOS settings. This way DirectX will be use the dGPU.
- The verification of C++AMP samples may fail on Windows 8 64-bit systems when run on the Windows emulator. If the executable is set to "High Performance mode," the verification should pass.
- Bolt C++ AMP samples using sort APIs fail to build in Visual Studio 2012. The compilation error is caused by a known issue in the Bolt 1.2/1.3 library wherein the compilation of particular modules fails with Visual Studio 2012. However, the sample can be successfully built by using Visual Studio 2013.
- An intermittent CPU reference verification failure occurs in the GEMM C++AMP sample on Tahiti GPU, and the FFT C++ AMP on Firepro W8100.
- When invoking `clEnqueueSVMMap` on Win 8.1 64-bit systems for large SVM buffers, an `CL_OUT_OF_RESOURCES` error occurs intermittently.
- When the MatrixMullImage sample is run on a CPU OpenCL device on Windows 8.1, a CPU verification failure occurs intermittently.

- The following samples have memory leaks: TransferOverlapCPP, SVMAtomicsBinaryTreeInsert, HeatPDE, HDR ToneMapping, GaussianNoiseGL, FluidSimulation2D, FineGrainSVMCAS, FineGrainSVM, CplusplusWrapper, SimpleConvolution, SimpleGL, SimpleMultiDevice, UnsharpMask, and URNGNoiseGL.
- CodeXL 1.8 displays API errors when samples are run in the Application Trace Mode. Ignore these messages.

Contact

Advanced Micro Devices, Inc.
One AMD Place
P.O. Box 3453
Sunnyvale, CA, 94088-3453
Phone: +1.408.749.4000

For AMD APP SDK:
URL: developer.amd.com/appsdk
Developing: developer.amd.com/



The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. The information contained herein may be of a preliminary or advance nature and is subject to change without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Copyright and Trademarks

© 2015 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ATI, the ATI logo, Radeon, FireStream, and combinations thereof are trademarks of Advanced Micro Devices, Inc. OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos. Other names are for informational purposes only and may be trademarks of their respective owners.