



CUB

API Reference

Table of Contents

CUB Overview.....iii

CUB Overview

CUB provides state-of-the-art, reusable software components for every layer of the CUDA programming model:

- ▶ **Parallel primitives**
 - ▶ Warp-wide "collective" primitives
 - ▶ Cooperative warp-wide prefix scan, reduction, etc.
 - ▶ Safely specialized for each underlying CUDA architecture
 - ▶ Block-wide "collective" primitives
 - ▶ Cooperative I/O, sort, scan, reduction, histogram, etc.
 - ▶ Compatible with arbitrary thread block sizes and types
 - ▶ Device-wide primitives
 - ▶ Parallel sort, prefix scan, reduction, histogram, etc.
 - ▶ Compatible with CUDA dynamic parallelism
- ▶ **Utilities**
 - ▶ Fancy iterators
 - ▶ Thread and thread block I/O
 - ▶ PTX intrinsics
 - ▶ Device, kernel, and storage management

The full CUB documentation is available on GitHub. See [CUB](#).

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2013-2021 NVIDIA Corporation. All rights reserved.

