



Bitonic Sorting

Intel® OpenCL SDK Sample Documentation

Document Number: 325262-002US



Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to <http://www.intel.com/design/literature.htm>.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Go to: http://www.intel.com/products/processor_number/.

This document contains information on products in the design phase of development.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Intel, Intel logo, Intel Core, VTune, Xeon are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

Microsoft product screen shot(s) reprinted with permission from Microsoft Corporation.

Copyright © 2010-2011 Intel Corporation. All rights reserved.



Optimization Notice

Intel compilers, associated libraries and associated development tools may include or utilize options that optimize for instruction sets that are available in both Intel and non-Intel microprocessors (for example SIMD instruction sets), but do not optimize equally for non-Intel microprocessors. In addition, certain compiler options for Intel compilers, including some that are not specific to Intel micro-architecture, are reserved for Intel microprocessors. For a detailed description of Intel compiler options, including the instruction sets and specific microprocessors they implicate, please refer to the "Intel Compiler User and Reference Guides" under "Compiler Options." Many library routines that are part of Intel compiler products are more highly optimized for Intel microprocessors than for other microprocessors. While the compilers and libraries in Intel compiler products offer optimizations for both Intel and Intel-compatible microprocessors, depending on the options you select, your code and other factors, you likely will get extra performance on Intel microprocessors.

Intel compilers, associated libraries and associated development tools may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Supplemental Streaming SIMD Extensions 3 (SSSE3) instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors.

While Intel believes our compilers and libraries are excellent choices to assist in obtaining the best performance on Intel and non-Intel microprocessors, Intel recommends that you evaluate other compilers and libraries to determine which best meet your requirements. We hope to win your business by striving to offer the best performance of any compiler or library; please let us know if you find we do not.

Notice revision #20110307



Contents

About Bitonic Sorting.....	5
Path	5
Introduction	5
Motivation.....	5
Algorithm	6
OpenCL™ Implementation.....	6
Code Highlights	6
Limitations	6
Understanding OpenCL Performance Characteristics.....	6
Benefits of Using Vector Data Types.....	6
Work-group Size Considerations.....	7
Reference (Native) Implementation	7
Project Structure	7
Controlling the Sample	7
References.....	8



About Bitonic Sorting

Bitonic Sorting sample illustrates implementing calculation kernels using OpenCL™ C99 and parallelizing kernels by running several work-groups in parallel.

Path

Location	Executable
<INSTALL_DIR>samples\BitonicSort	Win32\Release\BitonicSort.exe – 32-bit executable x64\Release\BitonicSort.exe – 64-bit executable Win32\Debug\BitonicSort.exe – 32-bit debug executable x64\Debug\BitonicSort.exe – 64- bit debug executable

Introduction

This sample demonstrates how to sort arbitrary input array of integer values with OpenCL using Single Instruction Multiple Data (SIMD) bitonic sorting networks. This implementation is very general, so it permits you to add *<key/value>* sorting with relatively low effort.

Motivation

Sorting algorithms are among most widely used building blocks. Bitonic sorting algorithm implemented in this sample is based on properties of bitonic sequence and principles of so-called sorting networks. It enables efficient SIMD-style parallelism through OpenCL vector data types.



Algorithm

For an array of length $2N \times 4$, this algorithm completes N stages of sorting. The first stage has one pass. As a result, the kernel forms bitonic sequences of size 4 using SIMD sorting network inside each item of the input array.

For each successive stage, the number of passes is incremented by one and the sequence size is doubled by merging two neighboring items.

For general reference on bitonic sorting networks, see [1]. For reference on sorting networks using SIMD data types, see [2].

OpenCL™ Implementation

Code Highlights

Bitonic sort OpenCL™ kernel of `BitonicSort.cl` file performs the specified stage of each pass. Every input array item or item pair (depending on the pass number) corresponds to a unique global ID that the kernel uses for their identification. The full sorting sequence consists of repetitive kernel calls performed in `ExecuteSortKernel()` function of `BitonicSort.cpp` file.

Limitations

For the sake of simplicity, the current version of the sample requires input array of size of 4×2^N 32-bit integer items, where N is a positive integer.

Understanding OpenCL Performance Characteristics

Benefits of Using Vector Data Types

This sample implements the bitonic sort algorithm using vector data types. Explicit usage of these types, such as `int4` or `float4`, enables the following optimizations:

- You can work with quads instead of single integers. This removes unnecessary branches, saves memory bandwidth, and optimizes CPU cache usage.



- You can use sorting network inside a single vector item during the last pass on every stage. This permits merging two last passes together to save an extra kernel invocation per stage, thus decreasing execution overhead.

Beside the maximum possible 4x speedup brought by SIMD register usage, these optimizations bring additional 25% speedup to the explicitly vectorized version. As a result, you get approximately 5x speedup in total.

Work-group Size Considerations

Valid work-group sizes on Intel platforms range from 1 to 1024 elements. To achieve peak performance, use work-groups of 64-128 elements.

Reference (Native) Implementation

Reference implementation is done in `ExecuteSortReference()` routine of `BitonicSort.cpp` file. This is single-threaded code that performs exactly the same bitonic sort sequence as OpenCL code, but uses pure scalar C nested loop.

Project Structure

This sample project has the following structure:

- `BitonicSort.cpp` - the host code, with OpenCL initialization and processing functions
- `BitonicSort.cl` – OpenCL sorting kernel source code
- `BitonicSort.vcproj` – Microsoft* Visual Studio* 2008 project file containing all the required dependencies.

Controlling the Sample

The sample executable is a console application. To set the sorting direction and input array size, use command line arguments. If the command line is empty, the sample uses the default values:

- Sorting direction is ascending.
 - Input array size is $1048576(2^{20})$ items.
- h** command line argument prints help information;



- s <arraySize> command line argument setups input/output array size;
- d command line argument sets descending sorting direction instead of default ascending.

References

- [1] H. W. Lang. Bitonic Sort. <http://www.iti.fh-flensburg.de/lang/algorithmen/sortieren/bitonic/bitonicen.htm>.
- [2] Jatin Chhugani, Anthony D. Nguyen, Victor W. Lee, William Macy, Mostafa Hagog, Yen-Kuang Chen, Akram Baransi, Sanjeev Kumar, Pradeep Dubey: Efficient implementation of sorting on multi-core SIMD CPU architecture. PVLDB 1(2): 1313-1324 (2008)
<http://portal.acm.org/citation.cfm?id=1454159.1454171&coll=GUIDE&dl=GUIDE&CFID=105910684&CFTOKEN=82233064>