



SIMD- and Cache-Friendly Algorithm for Sorting an Array of Structures



<u>Hiroshi Inoue</u>^{†‡} and Kenjiro Taura[‡] † IBM Research – Tokyo ‡ University of Tokyo

© 2015 IBM Corporation



Our Goal

Develop a fast algorithm for sorting an array of structures



Our approach:

- To use SIMD-based multiway mergesort as the basis
- To exploit SIMD instructions to accelerate mergesort kernel
- To avoid random memory accesses that cause excessive cache misses

Outline

✓Goal

- Background: SIMD multiway mergesort for integers
- Performance problems in existing approaches
- Our approach
- Performance results



Background: Mergesort with SIMD for sorting integers [Inoue *et al.* PACT 2007, Chhugani *et al.* PVLDB 2008 etc.]

- Merge operation for 32-bit or 64-bit integers can be efficiently implemented with SIMD by:
 - merging multiple values in vector registers using SIMD min and max instructions (i.e. without conditional branches)
 - integrating the in-register vector merge into usual comparison-based merge operation
- SIMD min and max instructions can accelerate sorting by
 - parallelizing comparisons
 - avoiding unpredictable conditional branches



Background: SIMD-based merge for values in two vector registers



SIMD- and Cache-Friendly Algorithm for Sorting an Array of Structures



























Background: Multiway mergesort



Multiway (*k*-way) mergesort, here k = 4



© 2015 IBM Corporation

Background: Multiway merge with SIMD



(8-way merge as an example. we used 32-way merge in actual implementation)

Outline

✓Goal

Background: SIMD multiway mergesort for integers

Performance problems in existing approaches

- Our approach
- Performance results

Existing approach for sorting structures with SIMD

- For sorting structures with SIMD mergesort, a frequently used approach (*key-index approach*) is
 - 1. pack key and index for each record into an integer.
 - 2. sort the key-index pairs with SIMD, and then
 - 3. rearrange the records based on the sorted key-

index pairs

8 Costly due to random accesses for main memory

Performance problems in existing approaches

- Existing approaches for sorting structures with SIMD instructions have performance problems
 - Key-index approach: SIMD friendly but not cache friendly due to random memory accesses
 - Direct approach: cache friendly but not SIMD friendly because the keys are not stored contiguously in memory

Our approach takes the benefits of both approaches



Our approach: Rearranging at each multiway merge step



SIMD- and Cache-Friendly Algorithm for Sorting an Array of Structures



Our approach: Rearranging at each multiway merge step



SIMD- and Cache-Friendly Algorithm for Sorting an Array of Structures

© 2015 IBM Corporation



Our approach: Multiway Merge with SIMD





Our approach: Overall sorting scheme



(read the paper on vectorized combsort)



Performance Evaluations

System

- 2.9-GHz Xeon E5-2690 (SandyBridge-EP) processors

- 2 sockets x 8 cores = 16 cores
- using SSE instructions (128-bit SIMD)
- 96 GB of system memory
- Redhat Enterprise Linux 6.4, gcc-4.8.2

We compared the performance of following algorithms

- Multiway mergesort (k = 32 way)

- Our approach, key-index approach, direct approach
- Radix sort, STL std::stable_sort, STL (unstable) std::sort



Performance for Sorting 512M 16-byte records with and without SIMD instructions



Performance Scalability with multiple cores



using 512M 16-byte records



Performance with various sizes of a record



Summary

- We developed a new algorithm for sorting an array of structures, that enables
 - efficient use of SIMD instructions
 - efficient use of cache memory (no random accesses)
- Our new algorithm outperformed the key-index approaches in multiway mergesort especially when each record is smaller than a cache line
- It outperformed the radix sort for records larger than 16 byte and showed better scalability with multiple cores
- Read the paper for more detail: efficient 4-wide SIMD exploitation, results with 10-byte ASCII key, sorting for variable-length strings

IBM

backup



Optimization techniques: partial key comparison

- To use 4-wide SIMD (32 bit x 4) instead of 2-wide SIMD (64 bit x 2)
 - Increasing data parallelism and giving the higher performance
 - Sorting based on <u>a partial key</u>
 - We confirm that the sorted order is correct using the entire key when rearranging records using scalar comparison



Sorting for variable-length strings

256M 12-byte to 20-byte records (16 bytes on average)

64M 12-byte to 84-byte records (48 bytes on average)



Performance with various numbers of records





Effect of number of ways in multi-way merge

