



# Faster Set Intersection with SIMD Instructions by Reducing Branch Mispredictions



<u>Hiroshi Inoue</u><sup>†‡</sup>, Moriyoshi Ohara<sup>†</sup>, Kenjiro Taura<sup>‡</sup> <sup>†</sup> IBM Research – Tokyo <sup>‡</sup>University of Tokyo



#### What is Set Intersection?

- The operation to find common elements from two sets
- We think intersecting two sorted integer arrays (e.g. std::set\_intersection in STL of C++)





#### Does it matter?

 Heavily used in DBMS (join operator) and information retrieval systems (multiword AND query)





#### How can we implement this?



in each iteration

- 1. check the equality of two elements
- 2. advance a pointer by 1

#### Merge-based approach



## Existing intersection algorithms



- Many techniques have been proposed for intersecting two arrays of very different sizes (10x ~)
  - based on binary search (e.g. galloping)
  - based on additional data structures (e.g. skip list, hash etc)
- They focus on reducing the number of comparisons
- For arrays with similar sizes, the merge-based algorithm is faster than these advanced algorithms → our focus



#### Key observation



- The comparison to select an input array for the next block is hard to predict for branch prediction hardware
  - It will be taken in arbitrary order
- The comparison to check equality is much easier to predict
  - It is not taken frequently for many applications
- We reduce the hard-to-predict conditional branches



#### Our approach for reducing branch mispredictions



in each iteration:

- 1. to find any matching pairs in blocks of S elements, here S = 2
- 2. to advance a pointer by S
- In the second second
- Sincrease other (easy-to-predict) conditional branches by S times
- Based on a simple cost model, the block size of 3 is the best when misprediction penalty is 10~22 cycles



#### Pseudo code of our approach (with block size S = 2)

 $S^2$  easy-to-predict branches per S elements  $\rightarrow$  S times more

while (pA < pAend-1 / \_ < pBend-1) {</pre> A0 = \*pA; A1 = \*(A+1); B0 = \*pB; B1 = \*(pB+1);(A0 == B0) { \*pOut++ = A0; } if else if (A0 == B1) { \*pOut++ = A0; Bpos+=2; continue; } else if (A1 == B0) { \*pOut++ = A1; Apos+=2; continue; } if (A1 == B1) { \*pOut++ = A1; Apos+=2; Bpos+=2; } else if (A1 < B1) { Apos+=2; } increment a pointer by S else { Bpos+=**2**; } } only one while processing S elements  $\rightarrow$  reduced to 1/S



#### Determining the best block size

#### A simple cost model of branches for block size S

	execution per element	mispredicti on rate	total cost
if_equal branches	S	0%	S * cost <sub>exec</sub>
if_greater branches	1/S	50%	$(cost_{exec} + cost_{misp}^* 0.5) / S$

• Best block size is determined by  $r = cost_{misp} / cost_{exec}$ 



## Our approach for exploiting SIMD instructions

- Existing approach: <u>full</u> comparison by SIMD to <u>find matching pairs</u> [Lemire *et al.* 2015, Schlegel *et al.* 2011]
  - limited data parallelism
  - limited element size
- Our approach: <u>partial</u> comparison by SIMD to <u>filter out redundant comparisons</u>
  - We can enjoy higher data parallelism
  - We can support larger elements (e.g. 32-bit or 64-bit integers)
  - Optimized for the common case







### Partial comparison by SIMD

 We introduce partial comparison by SIMD before the scalar comparison to reduce redundant comparisons



We can skip the all-pairs comparison by scalar if the no matching pair found in the partial comparison by SIMD

#### **Performance Evaluations**

#### Systems

- 2.9-GHz Xeon E5-2690 (SandyBridge-EP) processors
  - using SSE instructions (128-bit SIMD)
  - Redhat Enterprise Linux 6.4, gcc-4.8.2
- -4.1-GHz POWER7+ processors
  - using VSX instructions (128-bit SIMD)
  - Redhat Enterprise Linux 6.4, gcc-4.8.3



#### Performance improvements by our scalar algorithm





#### Performance improvements with SIMD instructions





#### Numbers of branch mispredictions and instructions





#### Performance for arrays with different sizes



#### intersecting two random 32-bit integer arrays



#### Adaptive fallback to avoid pathological degradations



#### intersecting two 256k random 32-bit integers



#### Adaptive fallback to avoid pathological degradations



#### Our SIMD algorithm → Our non-SIMD algorithm → Naive algorithm













# Summary

- We proposed a new set intersection algorithm which is efficient on today's processors
  - by reducing branch mispredictions
  - by avoiding redundant comparisons using SIMD
- Our new algorithm accelerates set intersection for artificial dataset compared to STL
  - by up to 2x without SIMD
  - by up to 5x using SIMD
- It also achieves better performance in an emulated query serving system
  - by up to 2.3x with SIMD over STL
  - by up to 1.5x over existing SIMD algorithms [Lemire et al. '15]