

FESIA: A Fast and SIMD-Efficient Set Intersection Approach on Modern CPUs

Jiyuan Zhang (jiyuanz@andrew.cmu.edu), Daniele G. Spampinato, Franz Franchetti

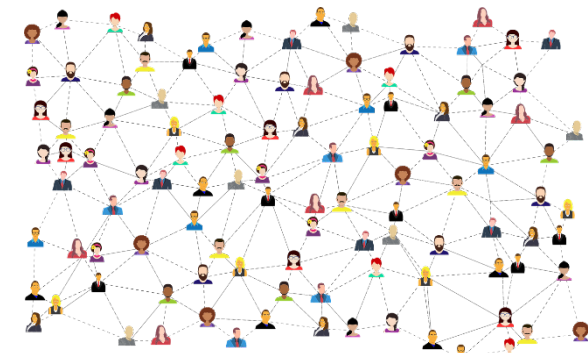
CMU ECE

Background

Set intersection is a functional primitive used in:



Data query processing



Social network analytics

Data parallelism is a common feature in the architectures of modern processors. And it has the trend of becoming wider and wider.



Approaches for set intersection

Example: hash, tree-based data structures for set intersections

✓ Good runtime complexity

✗ Hard to leverage vectorization features on modern processors, therefore cannot achieve speedups in practice

Vectorizations for set intersection

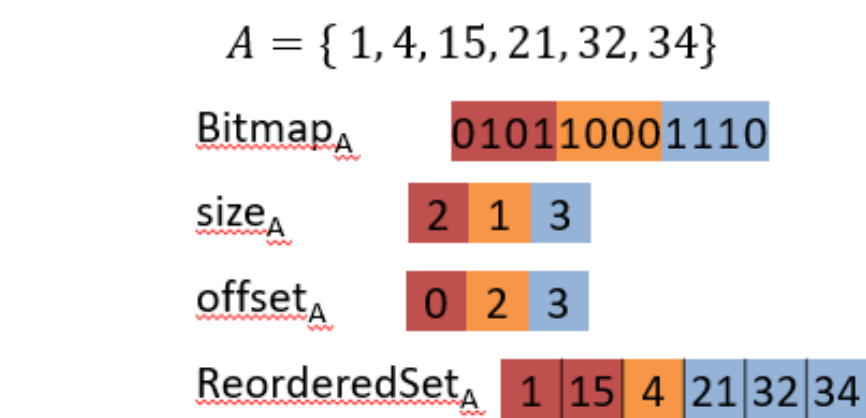
Merge-based intersection has irregular patterns of computations. It is non-trivial on how to leverage SIMD acceleration.

```

1 int scalar_merge_intersection(int L1[],
2                             int n1, int L2[], int n2) {
3     int i = 0, j = 0, r = 0;
4     while (i < n1 && j < n2) {
5         if (L1[i] < L2[j]) {
6             i++;
7         } else if (L1[i] > L2[j]) {
8             j++;
9         } else {
10            i++; j++; r++;
11        }
12    }
13    return r;
14 }
    
```

Listing 1. A code example of scalar merge-based set intersection

FESIA: A Fast and SIMD-Efficient Set Intersection Approach



Algorithm 1: The intersection algorithm in FESIA

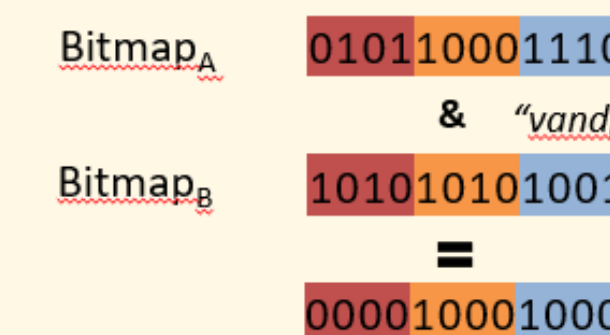
Input: List L_A , List L_B , Bitmap_A, Bitmap_B, ReorderedSet_A and ReorderedSet_B

Output: the intersection size r

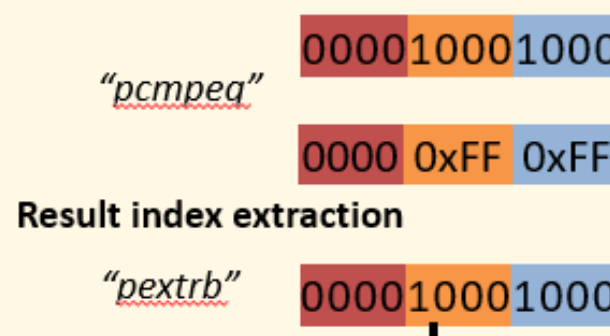
```

1 r = 0
2 N = m/s // the number of segments
3 for i in 0..N-1 do
4     if BitmapA & BitmapB != 0 then
5         r = r + Intersect(ReorderedSetA,
6                           ReorderedSetB)
7 return r
    
```

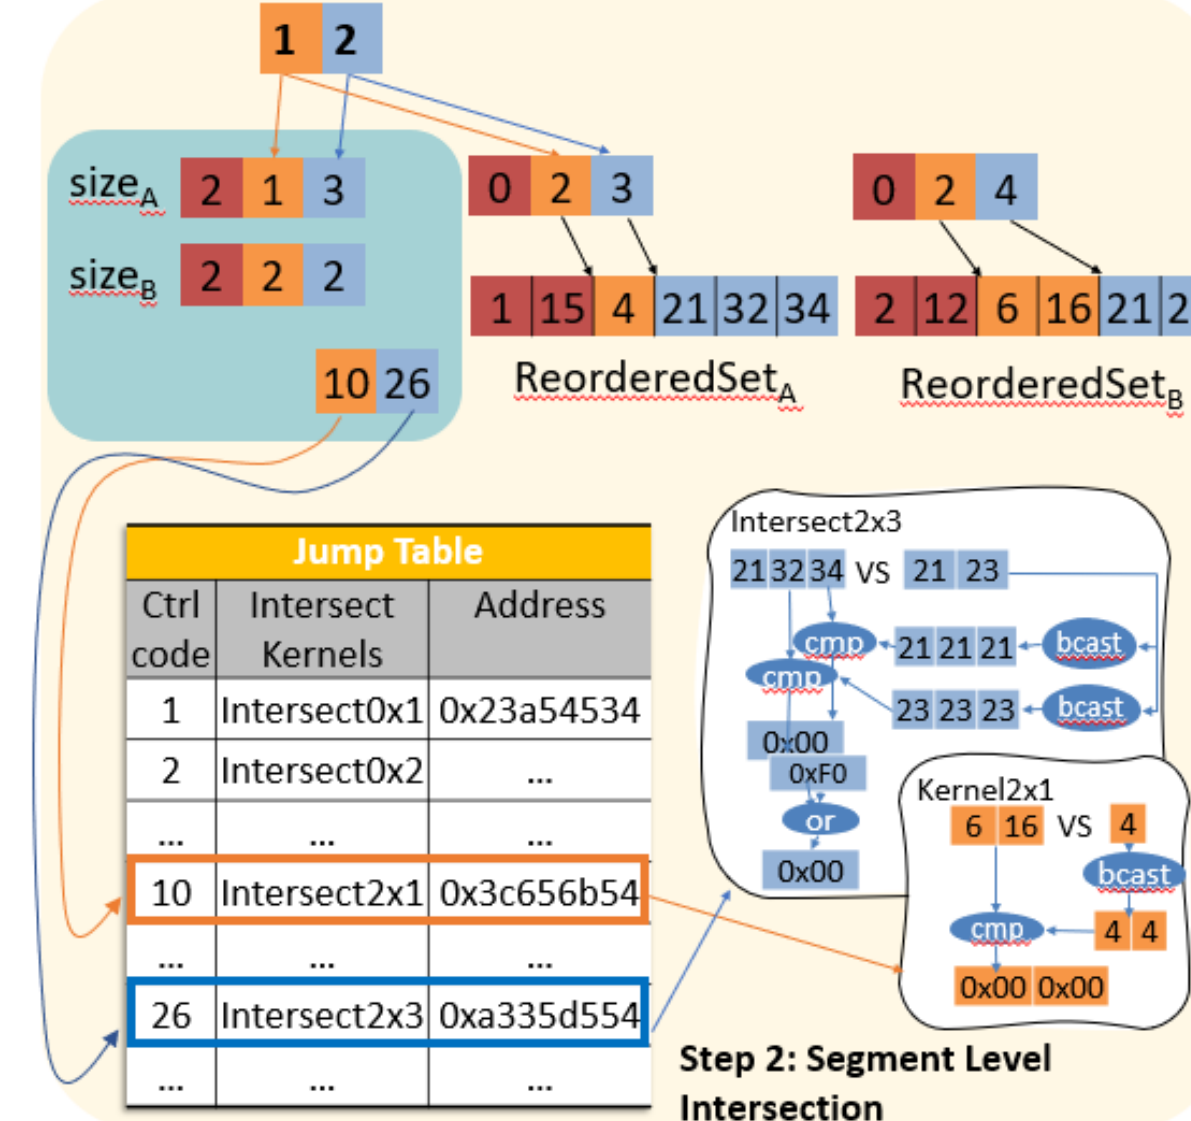
Segment comparison



Segment transformation



Step 1: Bitmap Level Intersection



A segmented-bitmap data structure, and a two-step intersection approach: (1) the bitmaps are used to filter out unmatched elements, and (2) a segment-by-segment comparison is conducted to compute the final set intersection using specialized SIMD kernels.

Step 1: Bitmap-level intersection

- Bitwise-AND on bitmaps e.g., *vands*
- Segment transformation
 - e.g. *pcmpeqw*, *pcmpeqq*, with $s =$ segment size
 - the output grouped by the segment size, e.g., *0xFFFF* ($s=16$)
- Non-zero segment index extraction e.g., *pextrb*

Step 2: Segment-level intersection

Specialized 2x4 code

```

vb = _mm_load_si128(B); v4 = _mm_set1_epi32(A[3]);
v1 = _mm_set1_epi32(A[0]); c4 = _mm_cmp_epi32(v4, vb);
c1 = _mm_cmp_epi32(v1, vb);
v2 = _mm_set1_epi32(A[1]); cmp =
c2 = _mm_or_si128(_mm_or_si128(c1,
c2), _mm_or_si128(c3, c4));
mask = _mm_movemask_ps(cmp);
v3 = _mm_set1_epi32(A[2]); count += _mm_popcnt_32(mask);
c3 = _mm_cmp_epi32(v3, vb); return count;
    
```

Function: Intersect4by5(A, B)

```

1 va = _mm_load_si128(A)
2 v1 = _mm_set1_epi32(B1)
3 c1 = _mm_cmp_epi32(v1, va)
4 v2 = _mm_set1_epi32(B2)
5 c2 = _mm_cmp_epi32(v2, va)
6 v3 = _mm_set1_epi32(B3)
7 c3 = _mm_cmp_epi32(v3, va)
8 v4 = _mm_set1_epi32(B4)
9 c4 = _mm_cmp_epi32(v4, va)
10 /* v5 and c5 are omitted */
11 t1 = _mm_or_si128(c1, c2)
12 t2 = _mm_or_si128(c3, c4)
13 t3 = _mm_or_si128(t1, t2)
14 cmp = _mm_or_si128(t3, c5)
15 mask = _mm_movemask_ps(cmp)
16 count = _mm_popcnt_32(mask)
17 return count
18
    
```

Specialized Intersect Kernels

The summary of FESIA VS. state-of-the-art intersection approaches

Methods	Complexity	Small Intersect	SIMD	Multicore	$n_1 \ll n_2$	k -way intersection	Portable
FESIA	$n/\sqrt{w} + r$	✓	✓	✓	$\min(n_1, n_2)$	$kn/\sqrt{w} + r$	✓
BMiss ^[1]	$n_1 + n_2$	✓	✓		$n_1 + n_2$	$n_1 \dots + n_k$	✓
Galloping ^[2]	$n_1 \log n_2$	✓	✓		$n_1 \log n_2$	$n_1 (\log n_2 + \dots + \log n_k)$	
Hiera ^[3]	$n_1 + n_2$	✓	✓		$n_1 + n_2$	$n_1 \dots + n_k$	
Fast ^[4]	$n/\sqrt{w} + r$	✓			$n/\sqrt{w} + r$	$n/\sqrt{w} + kr$	

w indicates the SIMD width, and n and r indicate the size of the input set and the intersection size

```

1 /* the dispatch control code */
2 int ctrl = ((Sa << 3) | Sb);
3 /* jump to a specialized intersection kernel */
4 switch (ctrl & 0x7F) {
5     /* Each specialized kernel is a macro */
6     case 0: break; //kernel0x0
7     case 1: break; //kernel0x1
8     case 2: break; //kernel0x2
9     /* case 3 to 62 are omitted ... */
10    case 63: Kernel17x7; break;
11    default: GeneralIntersection(); break;
12 }
    
```

Listing 2. The jump table for specialized kernels

Runtime dispatch

Small-size intersect kernel speedups

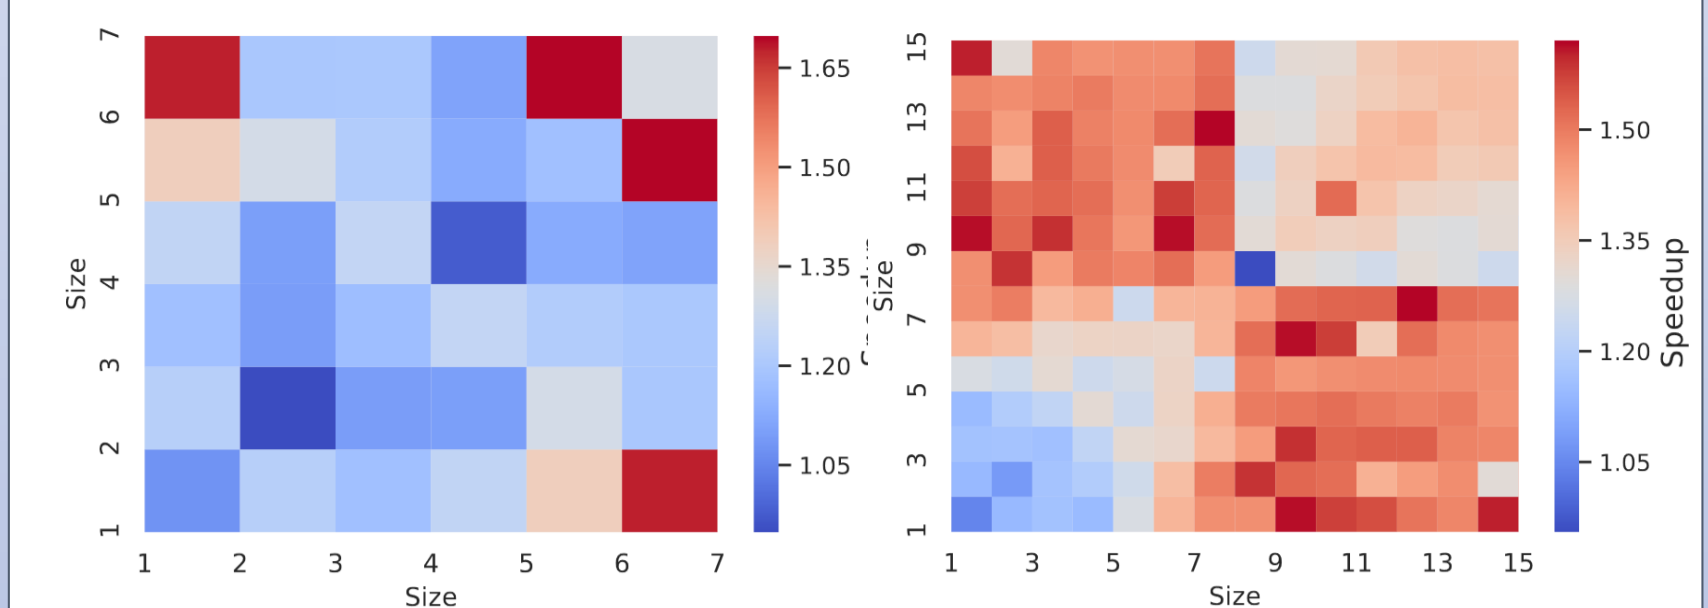
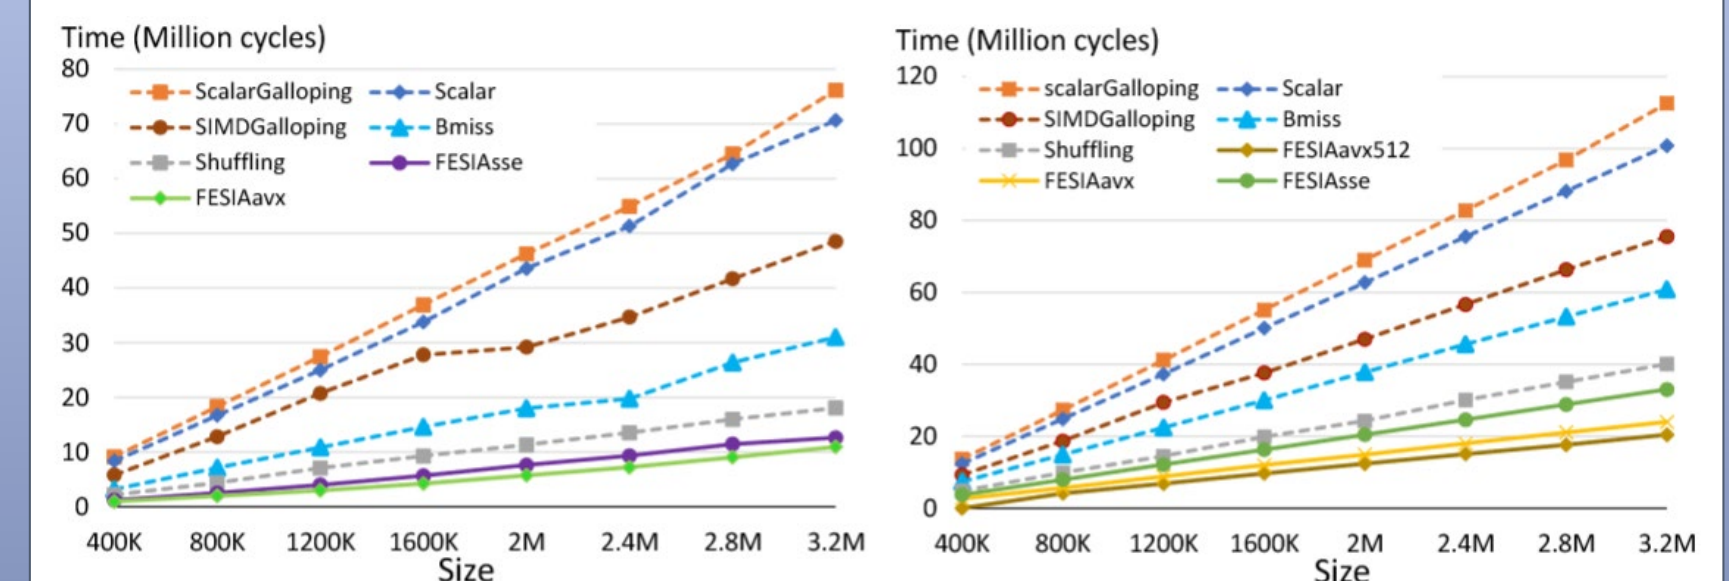


Fig. 4. Speedups of SSE kernels

Fig. 5. Speedups of AVX kernels

Performance on varying input size

Results are shown in two Intel architectures: Haswell and Skylake. The relative performance of these methods remains consistent as we increase the input size. FESIA outperforms other scalar and SIMD set intersection methods.



Real-world datasets

We study the performance on two real-world tasks: (1) a database query task, and (2) a triangle counting task in graph analytics

