

Paralelní a distribuované výpočty (B4B36PDV)

Jakub Mareček, Michal Jakob

jakub.marecek@fel.cvut.cz

Artificial Intelligence Center
Department of Computer Science
Faculty of Electrical Engineering
Czech Technical University in Prague

Use Cases

Via Theoretical Computer Science

A large part of theoretical computer science is devoted to the study of the limitations of parallelizability of algorithms, starting with the very clear question:

does every problem with a polynomial-time sequential algorithm also have an efficient parallel algorithm?

In particular, an algorithm solving the problem in time $O(\log^c n)$ using $O(n^k)$ parallel processors for some constants c and k .

In short, the answer is no.

Use Cases

Via Theoretical Computer Science

Let us consider the **P-Complete** problems i.e., decision problems that are in P and where every problem in P can be reduced to the P-Complete problem in logarithmic space (L). The prime examples of P-Complete problems are

- circuit evaluation: given a Boolean circuit and an input, is the output of the circuit 0?
- linear programming: given a linear function subject to linear inequality constraints, is the minimum greater or equal to 0?
- many graph problems, such as Lexicographically First Depth-first Search Ordering (LFDFS): given an undirected graph with fixed ordered adjacency lists, and two vertices u and v , is vertex u visited before vertex v in the depth-first search of the graph?
- many compression algorithms: given strings s and t , will compressing s with LZ78 add t to the dictionary?
- many problems related to Markov decision processes: is the minimum expected cost over all policies equal to 0 (for both finite, and long-run versions)?
- many tests of local optimality in combinatorial optimization: in an instance of Traveling Salesman, and a sequence of tours, is this a 2-Opt sequence?

Use Cases

Via Theoretical Computer Science

In contrast, Nick's class NC^c are decision problems solvable by a uniform family of Boolean circuits with polynomial size, depth $O(\log^c(n))$, and fan-in 2.

More usefully, a problem in NC with input of length n can be solved in time $O(\log^c n)$ using $O(n^k)$ parallel processors for some constants c and k . (Notice that for a constant k , $O(n^k)$ is a polynomial.

Thus, NC can be thought of as problems that can be efficiently solved on a parallel computer, and hence "easier" than P-Complete problems. For example:

- integer arithmetics (addition, multiplication and division),
- matrix arithmetics (multiplication, determinant, inverse, rank), or
- some graph problems (shortest path, maximal matching with some restrictions on the weights)

Use Cases

Via Theoretical Computer Science

Let us consider the problem of sorting an array of n elements.
Is it P-Complete?

Use Cases

Via Theoretical Computer Science

Let us consider the problem of sorting an array of n elements.
Is it P-Complete?

Algorithm	$p(n)$ Processors	Time
Sequential algorithms	1	$O(n \log n)$
Parallel divide and conquer	$O(1)$	$O(n \log n)$
	$O(\log n)$	$O\left(\frac{n(\log n)}{p(n)}\right)$
Parallel Ranking	$\omega(\log n)$	$O(n)$
	$O(n^2)$	$O(\log n)$

Use Cases

Via Theoretical Computer Science

Let us consider the problem of sorting an array of n elements.
Is it P-Complete?

Algorithm	$p(n)$ Processors	Time
Sequential algorithms	1	$O(n \log n)$
Parallel divide and conquer	$O(1)$	$O(n \log n)$
	$O(\log n)$	$O\left(\frac{n(\log n)}{p(n)}\right)$
	$\omega(\log n)$	$O(n)$
Parallel Ranking	$O(n^2)$	$O(\log n)$

As the table suggests, sorting is Nick's class: for n items, we can consider all pairs of items in parallel, compare them to obtain a binary value, and then for each item, obtain its rank in the sorted order by adding the binary values. This is known as the parallel ranking. There are many other parallel algorithms based on picking minimum and maximum from a small set, as well as algorithms based on hashing.

Use Cases

What to do in a particular use case?

In the second lecture, we have seen that within shared-memory parallel programming, we have broadly four options:

- **Confinement:** Do not share memory between threads.
- **Immutability:** Do not share any mutable data between threads.
- **Thread-safe code:** Use data types with additional guarantees for storing any mutable data shared between threads, or even better, use implementations of algorithms that are already parallelized and handle the concurrency issues for you.
- **Synchronization:** Use synchronization primitives to prevent accessing the variable at the same time.

Thread-safe Code in C++20

Use STL!

We have seen that the header `execution` defines objects `std::execution::par` and `std::execution::par_unseq`, which can be passed as the first argument of any standard algorithm:

```
1  #include <algorithm>
2  #include <chrono>
3  #include <execution>
4  #include <iostream>
5  #include <random>
6  #include <vector>
7  using namespace std::chrono;
8
9  int main() {
10     const int N = 1000000;
11     std::vector<int> v(N);
12     std::mt19937 rng;
13     rng.seed(std::random_device{}());
14     std::uniform_int_distribution<int> dist(0, 255);
15     std::generate(begin(v), end(v), [&]() { return dist(rng);
16     ↪ });
17     auto start = high_resolution_clock::now();
18     std::sort(std::execution::par, begin(v), end(v));
19     auto finish = high_resolution_clock::now();
20     auto duration = duration_cast<milliseconds>(finish - start);
21     std::cout << "\nElapsed time = " << duration.count() << "
22     ↪ ms\n";
23     return 0;
24 }
```

Thread-safe Code in C++20

Use STL!

What does the code do?

```
1 #include <algorithm>
2 #include <chrono>
3 #include <execution>
4 #include <iostream>
5 #include <random>
6 #include <vector>
7 using namespace std::chrono;
8
9 int main() {
10     const int N = 1000000;
11     std::vector<int> v(N);
12     std::mt19937 rng;
13     rng.seed(std::random_device()());
14     std::uniform_int_distribution<int> dist(0, 255);
15     std::generate(begin(v), end(v), [&]() { return dist(rng);
    ↪ });
16     auto start = high_resolution_clock::now();
17     std::sort(std::execution::par, begin(v), end(v));
18     auto finish = high_resolution_clock::now();
19     auto duration = duration_cast<milliseconds>(finish - start);
20     std::cout << "\nElapsed time = " << duration.count() << "
    ↪ ms\n";
21     return 0;
22 }
```

Thread-safe Code in C++20

Use STL!

What does the code do?

One may imagine:

- templates
- iterators
- execution strategies.

Elegant code.

```
1  template<class ForwardIt>
2  void quicksort(ForwardIt first, ForwardIt last) {
3      if (first == last) return;
4      std::size_t distance = std::distance(first, last);
5      auto pivot = *std::next(first, distance / 2);
6      ForwardIt middle1; ForwardIt middle2;
7      if (distance < threshold) {
8          middle1 = std::partition(std::execution::seq, first,
9              ↪ last, [pivot](const auto &em) { return em < pivot;
10             ↪ });
11         middle2 = std::partition(std::execution::seq, middle1,
12             ↪ last, [pivot](const auto &em) { return !(pivot <
13             ↪ em); });
14     } else {
15         middle1 = std::partition(std::execution::par, first,
16             ↪ last, [pivot](const auto &em) { return em < pivot;
17             ↪ });
18         middle2 = std::partition(std::execution::par, middle1,
19             ↪ last, [pivot](const auto &em) { return !(pivot <
20             ↪ em); });
21     }
22     quicksort(first, middle1);
23     quicksort(middle2, last);
24 }
```

Thread-safe Code in C++20

Use STL!

Reality can be
much more messy:

```
1  template <class _ExecutionPolicy, typename
   ↪  _RandomAccessIterator, typename _Compare, typename
   ↪  _LeafSort>
2  void
3  __parallel_stable_sort(_ExecutionPolicy&&,
   ↪  _RandomAccessIterator __xs, _RandomAccessIterator __xe,
   ↪  _Compare __comp,
4
   ↪  _LeafSort __leaf_sort, std::size_t
   ↪  __nsort = 0)
5  {
6      tbb::this_task_arena::isolate([=, &__nsort]() {
7          //sorting based on task tree and parallel merge
8          typedef typename
   ↪  std::iterator_traits<_RandomAccessIterator>::value_ty
   ↪  _ValueType;
```

Thread-safe Code in C++20

Use STL!

Reality can be
much more messy:

```
9         typedef typename
           ↪ std::iterator_traits<_RandomAccessIterator>::difference_type
           ↪ _DifferenceType;
10        const _DifferenceType __n = __xe - __xs;
11        if (__nsort == __n)
12            __nsort = 0; // 'partial_sort' becomes 'sort'
13
14        const _DifferenceType __sort_cut_off =
           ↪ _PSTL_STABLE_SORT_CUT_OFF;
15        if (__n > __sort_cut_off)
16        {
17            __buffer<_ValueType> __buf(__n);
18
           ↪ __root_task<__stable_sort_func<_RandomAccessIterator,
           ↪ _ValueType*, _Compare, _LeafSort>> __root{
19                __xs, __xe, __buf.get(), true, __comp,
           ↪ __leaf_sort, __nsort, __xs,
           ↪ __buf.get()};
20            __task::spawn_root_and_wait(__root);
21            return;
22        }
23        //serial sort
24        __leaf_sort(__xs, __xe, __comp);
25    });
26 }
```

Thread-safe Code in C++20

Use STL!

In the previous slides, we have seen the implementation in Intel Thread Building Blocks (TBB) backend of the GCC. This uses:

- many megabytes of a library (TBB)
- “sorting based on task tree and parallel merge”,

while making use of several non-trivial tricks, including

- `tbb::task_scheduler_init`,
- `std::thread::hardware_concurrency()`,
- `std::hardware_constructive_interference_size`.

(Contrast this with the serial version of GCC **sort**, which uses a multi-way mergesort, and GCC **stable_sort**, which uses a quicksort.) We wish to make use of the STL, rather than redevelop it, in the first instance.

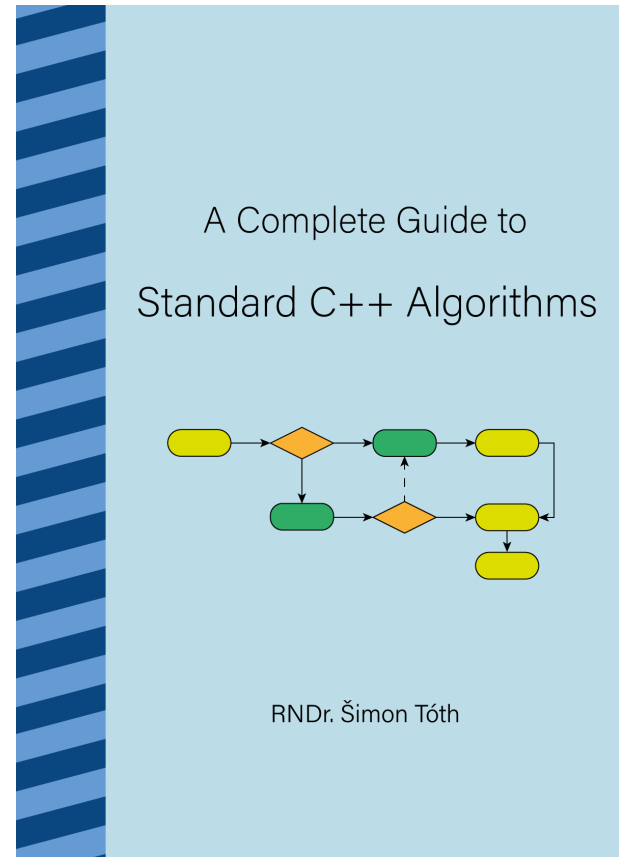
Thread-safe Code in C++20

Use STL!

Even making full use of the STL is quite non-trivial.

In the lecture notes, we present an overview of the sorting-related routines in verbatim from the fantastic book "A Complete Guide to Standard C++ Algorithms" of Simon Toth, in compliance with the license.

<https://github.com/HappyCerberus/book-cpp-algorithms>



Thread-safe Code in C++20

Use STL!

What is wrong with comparators:

```
1 struct Point {
2
3 int x;
4 int y;
5
6 // pre-C++20 lexicographical less-than
7 friend bool operator<(const Point& left, const Point& right) {
8     if (left.x != right.x)
9         return left.x < right.x;
10    return left.y < right.y;
11 }
12
13 // default C++20 spaceship version of lexicographical
14   ↪ comparison
15 friend auto operator<=>(const Point&, const Point&) = default;
16
17 // manual version of lexicographical comparison using operator
18   ↪ <=>
19 friend auto operator<=>(const Point& left, const Point& right)
20   ↪ {
21     if (left.x != right.x)
22         return left.x <=> right.x;
23     return left.y <=> right.y;
24 }
```


Thread-safe Code in C++20

Use STL!

What is wrong with projections:

```
1 struct Account {
2     double value() { return value_; }
3     double value_;
4 };
5
6 std::vector<Account> accounts{{0.1}, {0.3}, {0.01}, {0.05}};
7 std::ranges::sort(accounts, std::greater<>{},
8     ↪ &Account::value);
9 // accounts = { {0.3}, {0.1}, {0.05}, {0.01} }
```

Thread-safe Code in C++20

If you know your problem and STL well, you may benefit from reformulating the problem, e.g., to partial sort:

Example of using `std::partial_sort` to sort the first three elements of a range.

```
1  std::vector<int> data{9, 8, 7, 6, 5, 4, 3, 2, 1};
2  std::partial_sort(data.begin(), data.begin()+3, data.end());
3  // data == {1, 2, 3, -unspecified order-}
4
5  std::ranges::partial_sort(data, data.begin()+3,
6  ↪  std::greater<>());
7  // data == {9, 8, 7, -unspecified order-}
```

Synchronization

Now, let us move to our own parallel sorting algorithms.

We will see:

- Bubble sort with OpenMP
- Quick sort variants with task construct in OpenMP
- Merge sort variants with task construct in OpenMP
- Many variants, including intrinsics.

Synchronization

Bubble Sort

Bubble sort is essentially using loops:

```
1  bool compare_swap(std::vector<int>& vector_to_sort, const int&
   ↪ val1, const int& val2) {
2      if (vector_to_sort[val1] > vector_to_sort[val2]) {
3          std::iter_swap(vector_to_sort.begin() + val1,
   ↪ vector_to_sort.begin() + val2);
4          return true;
5      }
6      return false;
7  }
8
9  void bubble(std::vector<int>& vector_to_sort, int from, int
   ↪ to) {
10     bool change = true;
11     while (change) {
12         change = false;
13         for (int i = from + 1; i < to; i++) {
14             change |= compare_swap(vector_to_sort, i - 1, i);
15         }
16     }
17 }
```

Synchronization

Bubble Sort

Bubble sort is essentially using loops, which are easy to parallize:

```
1 void parallel_bubble (std::vector<int>& vector_to_sort,  
  ↪ unsigned int from,unsigned int to) {  
2     while (change) {  
3         change = false;  
4 #pragma omp parallel for num_threads(thread_count)  
  ↪ schedule(static) shared(vector_to_sort)  
  ↪ reduction(|:change)  
5         for (int i = from + 1; i < to; i += 2) {  
6             change |= compare_swap(vector_to_sort, i - 1, i);  
7         }  
8  
9 #pragma omp parallel for num_threads(thread_count)  
  ↪ schedule(static) shared(vector_to_sort)  
  ↪ reduction(|:change)  
10        for (int i = from + 2; i < to; i += 2) {  
11            change |= compare_swap(vector_to_sort, i - 1, i);  
12        }  
13    }  
14 }
```

Synchronization

Quick Sort

Quick sort may benefit from, OpenMP construct **task**

```
1 void qs(std::vector<int> &vector_to_sort, int from, int to) {
2     if (to - from <= base_size) {
3         std::sort(vector_to_sort.begin() + from,
4                 ↪ vector_to_sort.begin() + to);
5         return;
6     }
7     // cf. the pivot (vector_to_sort[from])
8     int part2_start = partition(vector_to_sort, from, to,
9                               ↪ vector_to_sort[from]);
10
11     if (part2_start - from > 1) {
12 #pragma omp task shared(vector_to_sort) firstprivate(from,
13 ↪ part2_start)
14     {
15         qs(vector_to_sort, from, part2_start);
16     }
17 }
18 if (to - part2_start > 1) {
19     qs(vector_to_sort, part2_start, to);
20 }
21 }
```

Synchronization

Quick Sort

One can improve upon this:

- using three-way sort
- using task mergeable
(As suggested by Intel.)

```
1  template<class RanIt, class _Pred>
2  void qsort3w(RanIt _First, RanIt _Last, _Pred compare) {
3      if (_First >= _Last) return;
4
5      std::size_t _Size = 0L;
6      g_depth++;
7      if ((_Size = std::distance(_First, _Last)) > 0) {
8          RanIt _LeftIt = _First, _RightIt = _Last;
9          bool is_swapped_left = false, is_swapped_right =
10             ↪ false;
11             typename std::iterator_traits<RanIt>::value_type
12             ↪ _Pivot = *_First;
13
14             RanIt _FwdIt = _First + 1;
15             while (_FwdIt <= _RightIt) {
16                 if (compare(*_FwdIt, _Pivot)) {
17                     is_swapped_left = true;
18                     std::iter_swap(_LeftIt, _FwdIt);
19                     _LeftIt++;
20                     _FwdIt++;
21                 } else if (compare(_Pivot, *_FwdIt)) {
22                     is_swapped_right = true;
23                     std::iter_swap(_RightIt, _FwdIt);
24                     _RightIt--;
25                 } else _FwdIt++;
26             }
27         }
```

Synchronization

Quick Sort

One can improve upon this:

- using three-way sort
- using task mergeable

```
26         if (_Size >= cutoff) {
27 #pragma omp taskgroup
28         {
29 #pragma omp task untied mergeable
30             if ((std::distance(_First, _LeftIt) > 0) &&
31                 ↪ (is_swapped_left))
32                 qsort3w(_First, _LeftIt - 1, compare);
33 #pragma omp task untied mergeable
34             if ((std::distance(_RightIt, _Last) > 0) &&
35                 ↪ (is_swapped_right))
```


Synchronization

Quick Sort

One can improve upon this:

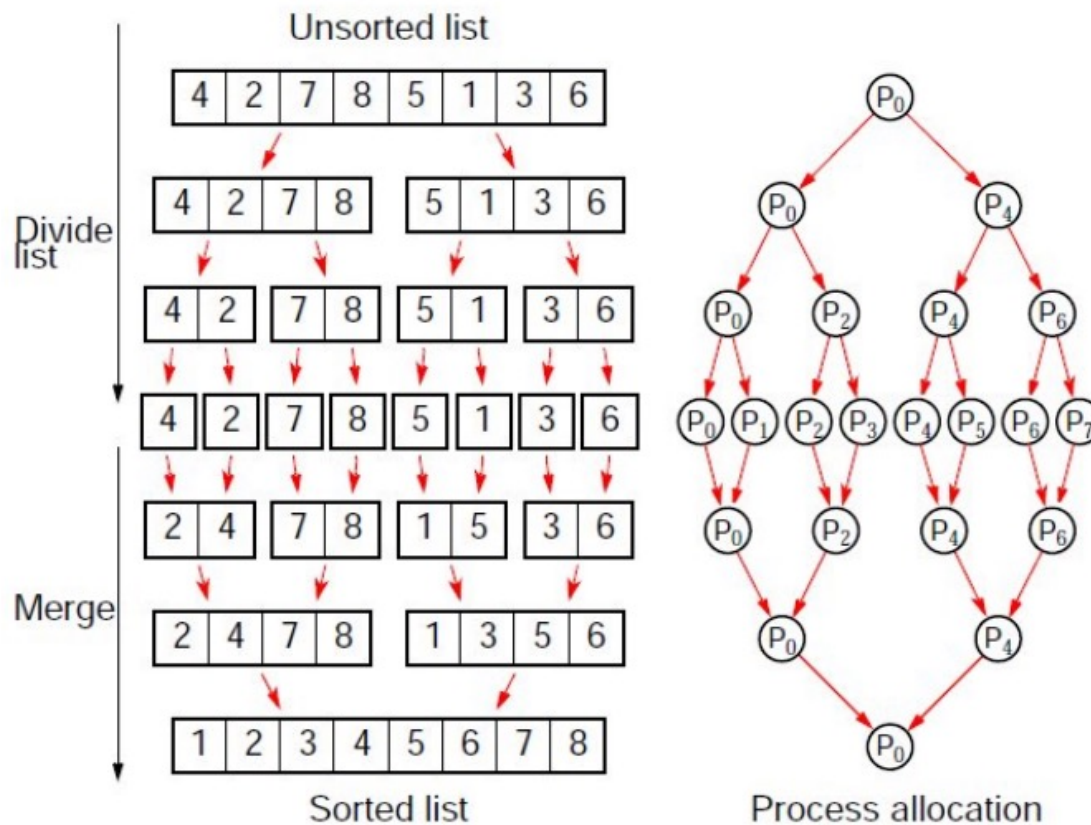
- using three-way sort
- using task mergeable

```
35             qsort3w(_RightIt + 1, _Last, compare);
36         }
37     } else {
38 #pragma omp task untied mergeable
39     {
40         if ((std::distance(_First, _LeftIt) > 0) &&
41             ↪ is_swapped_left)
42             qsort3w(_First, _LeftIt - 1, compare);
43         if ((std::distance(_RightIt, _Last) > 0) &&
44             ↪ is_swapped_right)
45             qsort3w(_RightIt + 1, _Last, compare);
46     }
47 }
48 }
```

Synchronization

Merge Sort

Similarly, one can use task to parallelise merge sort:



Synchronization

Merge Sort

Similarly, one can use task to parallelise merge sort:

```
1 void ms_parallel(std::vector<int>& vector_to_sort, int from,
  ↪ int to) {
2     if (to - from <= 1) {
3         return;
4     }
5     int middle = (to - from)/2 + from;
6
7     ms_serial(vector_to_sort, from, middle);
8     ms_serial(vector_to_sort, middle, to);
9     std::inplace_merge(vector_to_sort.begin()+from,
  ↪ vector_to_sort.begin()+middle,
  ↪ vector_to_sort.begin()+to);
10 }
11
12 void ms(std::vector<int>& vector_to_sort, int from, int to) {
13     if (to - from <= base_size) {
14         ms_serial(vector_to_sort,from,to);
15         return;
16     }
17     int middle = (to - from)/2 + from;
18
```

Synchronization

Merge Sort

Similarly, one can use task to parallelise merge sort:

```
19 #pragma omp task shared(vector_to_sort)
   ↪ firstprivate(from,middle)
20     ms(vector_to_sort, from, middle);
21
22     ms(vector_to_sort, middle, to);
23
24 #pragma omp taskwait
25     std::inplace_merge(vector_to_sort.begin()+from,
   ↪ vector_to_sort.begin()+middle,
   ↪ vector_to_sort.begin()+to);
26 }
```










Synchronization

Merge Sort

RESEARCH-ARTICLE

Efficient implementation of sorting on multi-core SIMD CPU architecture



Authors:  [Jatin Chhugani](#),  [Anthony D. Nguyen](#),  [Victor W. Lee](#),  [William Macy](#),  [Mostafa Hagog](#),
 [Yen-Kuang Chen](#),  [Akram Baransi](#),  [Sanjeev Kumar](#),  [Pradeep Dubey](#) [Authors Info & Affiliations](#)

Publication: Proceedings of the VLDB Endowment • August 2008 • <https://doi.org/10.14778/1454159.1454171>



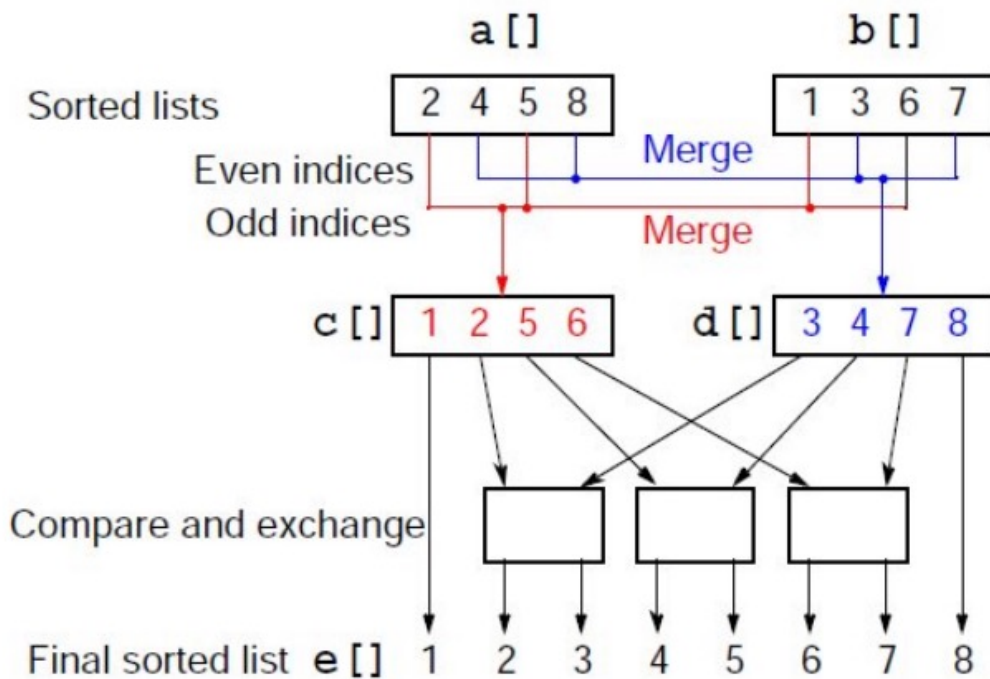


Figure 8: Parallel performance of the scalar and SIMD implementations.

Synchronization

Going beyond Merge Sort



Odd-Even Merge Sort:



Synchronization

Going beyond Merge Sort

Odd-Even Merge Sort:



Synchronization

Going beyond Merge Sort

Odd-Even Merge Sort:


```
1 void odd-even-merge (std::vector<int>& vector_to_sort, int
  ↪ from, int to, int step) {
2     auto new_step = step * 2;
3     if (new_step < to - from) {
4         odd-even-merge(vector_to_sort,from,to,new_step);
5         odd-even-merge(vector_to_sort,from+step,to,new_step);
6         for (int i=from+step; i<to-step; i += new_step) {
7             compare_and_swap(vector_to_sort,i,i+step);
8         }
9     } else {
10        compare_and_swap(vector_to_sort,from,from+step);
11    }
12 }
```

One can go even further with bitonic sort.

Synchronization

Bitonic Sort

- Efficient implementations can vectorize the compare and swap
- This is the most efficient approach on GPGPUs
- One can also experiment with intrinsics:
<https://software.intel.com/sites/landingpage/IntrinsicsGuide/>



intel® Intrinsics Guide

Technologies

- MMX
- SSE
- SSE2
- SSE3
- SSSE3
- SSE4.1
- SSE4.2
- AVX
- AVX2
- FMA
- AVX-512
- KNC
- AMX
- SVML
- Other

The Intel Intrinsics Guide is an interactive reference tool for Intel intrinsic instructions, which are C style functions that provide access to many Intel instructions - including Intel® SSE, AVX, AVX-512, and more - without the need to write assembly code. ✕


```
void __mm_2intersect_epi32 (__m128i a, __m128i b, __mmask8* k1, __mmask8* k2)                vp2intersectd
void __mm256_2intersect_epi32 (__m256i a, __m256i b, __mmask8* k1, __mmask8* k2)          vp2intersectd
void __mm512_2intersect_epi32 (__m512i a, __m512i b, __mmask16* k1, __mmask16* k2)       vp2intersectd
void __mm_2intersect_epi64 (__m128i a, __m128i b, __mmask8* k1, __mmask8* k2)          vp2intersectq
void __mm256_2intersect_epi64 (__m256i a, __m256i b, __mmask8* k1, __mmask8* k2)       vp2intersectq
void __mm512_2intersect_epi64 (__m512i a, __m512i b, __mmask8* k1, __mmask8* k2)       vp2intersectq
__m512i __mm512_4dpwssd_epi32 (__m512i src, __m512i a0, __m512i a1, __m512i a2, __m512i a3,
__m128i * b)                                    vp4dpwssd
__m512i __mm512_mask_4dpwssd_epi32 (__m512i src, __mmask16 k, __m512i a0, __m512i a1, __m512i a2,
__m512i a3, __m128i * b)                        vp4dpwssd
__m512i __mm512_maskz_4dpwssd_epi32 (__mmask16 k, __m512i src, __m512i a0, __m512i a1, __m512i a2,
__m512i a3, __m128i * b)                        vp4dpwssd
```

An Aside

Intrinsics

- Modern processors implement vector instructions (SIMD)
- On Intel and AMD, there are Streaming SIMD Extensions (SSE) and Advanced Vector Extensions (AVX) incl. 512-bit vectors.
- In SSE, `__m128d` stores 2 doubles, 4 ints, 16 chars, but in the reverse order `float[4] {0f,1f,2f,3f}`
- You can compile for this with `-march=native-mavx`
- You can run pairwise sorting using minima and maxima:


3	2	1	0
---	---	---	---



An Aside

Intrinsics

```
15     __m256i v1;  
16     __m256i v2;  
17     __m256i r1,r2;  
18  
19     for (int i=0; i<SIZE; i += 8) {  
20         v1 = _mm256_loadu_si256((__m256i *) &vec1[i]);  
21         v2 = _mm256_loadu_si256((__m256i *) &vec2[i]);  
22         r1 = _mm256_min_epi32(v1, v2);  
23         r2 = _mm256_max_epi32(v1, v2);  
24         _mm256_storeu_si256((__m256i *) &vec1[i], r1);  
25         _mm256_storeu_si256((__m256i *) &vec2[i], r2);  
26     }
```



An Aside

Intrinsics

- We can also pad, shift, truncate.
(Illustrations by Brano Bosansky.)

Pad:

				x3	x2	x1	x0
0	0	0	0	2	5	6	9

Shift to the right:

					x3	x2	x1
0	0	0	0	0	2	5	6

Truncate:

	x3	x2	x1
0	2	5	6

Compare:

x3	x2	x1	x0
2	5	6	9

An Aside

Intrinsics

- We can also pad, shift, truncate.
(Illustrations by Brano Bosansky.)

```
7  __m128i mask_llhllhh =
   ↪  _mm_set_epi32(0xffffffff,0,0xffffffff,0);
8  __m128i mask_hhllhll =
   ↪  _mm_set_epi32(0,0xffffffff,0,0xffffffff);
9  __m128i v1;
10 __m128i v2;
11 __m128i r1,r2;
12 for (int i=0; i<SIZE; i += 4) {
13 v1 = _mm_loadu_si128((__m128i *) &vec1[i]);
14 v2 = _mm_alignr_epi8(_mm_setzero_si128(), v1 ,1*4);
15 r1 = _mm_min_epi32(v1, v2);
16 r1 = _mm_and_si128(r1,mask_hhllhll);
17 v2 = _mm_alignr_epi8(v1, _mm_setzero_si128(),3*4);
18 r2 = _mm_max_epi32(v1, v2);
19 r2 = _mm_and_si128(r2,mask_llhllhh);
20 r1 = _mm_or_si128(r1,r2);
21 _mm_storeu_si128((__m128i *) &vec1[i], r1);
22 }
```

An Aside

Intrinsics

- We can also pad, shift, truncate.
(Illustrations by Brano Bosansky.)
- See <https://xhad1234.github.io/Parallel-Sort-Merge-Join-in-Peloton/> for a comprehensive illustration.

The Upshot

Intrinsics need not win

ARTICLE

Samplesort: A Sampling Approach to Minimal Storage Tree Sorting



Authors: [W. D. Frazer](#), [A. C. McKellar](#) [Authors Info & Affiliations](#)

Publication: Journal of the ACM • July 1970 • <https://doi.org/10.1145/321592.321600>



The Intel Intrinsic Guide is an interactive reference tool for Intel intrinsic instructions, which are C style functions that provide access to many Intel instructions - including Intel® SSE, AVX, AVX-512, and more - without the need to write assembly code. ✕

Technologies

- MMX
- SSE
- SSE2
- SSE3
- SSSE3
- SSE4.1
- SSE4.2
- AVX
- AVX2
- FMA
- AVX-512
- KNC

<code>void _mm_2intersect_epi32 (__m128i a, __m128i b, __mmask8* k1, __mmask8* k2)</code>	<code>vp2intersectd</code>
<code>void _mm256_2intersect_epi32 (__m256i a, __m256i b, __mmask8* k1, __mmask8* k2)</code>	<code>vp2intersectd</code>
<code>void _mm512_2intersect_epi32 (__m512i a, __m512i b, __mmask16* k1, __mmask16* k2)</code>	<code>vp2intersectd</code>
<code>void _mm_2intersect_epi64 (__m128i a, __m128i b, __mmask8* k1, __mmask8* k2)</code>	<code>vp2intersectq</code>
<code>void _mm256_2intersect_epi64 (__m256i a, __m256i b, __mmask8* k1, __mmask8* k2)</code>	<code>vp2intersectq</code>
<code>void _mm512_2intersect_epi64 (__m512i a, __m512i b, __mmask8* k1, __mmask8* k2)</code>	<code>vp2intersectq</code>
<code>__m512i _mm512_4dpwssd_epi32 (__m512i src, __m512i a0, __m512i a1, __m512i a2, __m512i a3, __m128i * b)</code>	<code>vp4dpwssd</code>
<code>__m512i _mm512_mask_4dpwssd_epi32 (__m512i src, __mmask16 k, __m512i a0, __m512i a1, __m512i a2, __m512i a3, __m128i * b)</code>	<code>vp4dpwssd</code>

The Upshot

<https://arxiv.org/pdf/2009.13569.pdf>

Engineering In-place (Shared-memory) Sorting Algorithms

MICHAEL AXTMANN, Karlsruhe Institute of Technology

SASCHA WITT, Karlsruhe Institute of Technology

DANIEL FERIZOVIC, Karlsruhe Institute of Technology

PETER SANDERS, Karlsruhe Institute of Technology

Type	Distribution	IPS ⁴ _o	PBBS	PS ⁴ _o	MCSTLmwm	MCSTLbq	TBB	RegionSort	PBBR	RADULS2	ASPaS
double	Sorted	1.42	10.96	2.02	15.47	13.36	1.06				42.23
double	ReverseSorted	1.06	1.34	1.98	1.76	11.00	3.01				5.34
double	Zero	1.54	12.83	1.80	14.55	166.67	1.06				41.78
double	Exponential	1.00	1.82	1.97	2.60	3.20	10.77				4.97
double	Zipf	1.00	1.96	2.12	2.79	3.55	11.56				5.33
double	RootDup	1.00	1.54	2.22	2.52	3.88	5.54				6.28
double	TwoDup	1.00	1.93	1.88	2.45	2.99	5.52				4.44
double	EightDup	1.00	1.82	2.01	2.48	3.19	10.37				5.02
double	AlmostSorted	1.00	1.73	2.40	5.12	2.18	3.54				6.37
double	Uniform	1.00	2.00	1.85	2.53	2.99	9.16				4.39
Total		1.00	1.82	2.06	2.83	3.10	7.46				5.21
Rank		1	2	3	4	5	7				6

Table 4. Average slowdowns of parallel algorithms for different data types and input distributions. The slowdowns average over the machines and input sizes with at least $2^{21}t$ bytes.

Conclusions

- First use case in parallelization
- Highlights importance of “thinking out of the box”