

pSTL-Bench: A Micro-Benchmark Suite for Assessing Scalability of C++ Parallel STL Implementations

Ruben Laso
ruben.laso@tuwien.ac.at
Faculty of Informatics
TU Wien
Vienna, Austria

Diego Krupitza
krupitza@par.tuwien.ac.at
Faculty of Informatics
TU Wien
Vienna, Austria

Sascha Hunold
sascha.hunold@tuwien.ac.at
Faculty of Informatics
TU Wien
Vienna, Austria

ABSTRACT

Since the advent of parallel algorithms in the C++17 Standard Template Library (STL), the STL has become a viable framework for creating performance-portable applications. Given multiple existing implementations of the parallel algorithms, a systematic, quantitative performance comparison is essential for choosing the appropriate implementation for a particular hardware configuration.

In this work, we introduce a specialized set of micro-benchmarks to assess the scalability of the parallel algorithms in the STL. By selecting different backends, our micro-benchmarks can be used on multi-core systems and GPUs.

Using the suite, in a case study on AMD and Intel CPUs and NVIDIA GPUs, we were able to identify substantial performance disparities among different implementations, including GCC+TBB, GCC+HPX, Intel’s compiler with TBB, or NVIDIA’s compiler with OpenMP and CUDA.

KEYWORDS

Performance Portability, C++, Standard Template Library, Threading Building Blocks, OpenMP, CUDA

1 INTRODUCTION

Writing efficient, parallel applications is notoriously hard, but writing performance-portable, efficient, parallel applications is harder. In the last decades, several types of parallel architectures (like GPUs or Xeon Phis) were introduced, which often required a complete rewrite of the parallelization approach to make applications efficient. To overcome the problem of having to combine several paradigms such as CUDA [16], OpenMP [6], or MPI [5] to write efficient programs, several frameworks, mainly using C++, such as Kokkos [23] or Raja [2], were proposed to allow scientists to write performance-portable applications. These frameworks allow for an efficient execution of parallel applications using different hardware architectures, i.e., the same program can run on one or more GPUs as well as on multi-core CPUs.

With the advent of C++17, parallel versions of the C++ Standard Template Library (STL) were standardized, which allows ISO C++ parallel programs to be performance portable [15]. Several works have already compared the resulting performance of various performance-portability layers [1, 10]. However, their focus lay on comparing full applications or mini-apps, where significant parts of a rewritten program may significantly influence the resulting performance.

In this work, we set out to devise a set of micro-benchmarks to assess the performance of the individual parallel STL algorithms found in C++ in a quantitative manner. Since different compiler

frameworks provide competing implementations of the STL, our goal is to capture the current state-of-the-art of the performance of parallel STL implementations. We compare several combinations of compiler frameworks, including GCC, Intel OneAPI compiler, and NVIDIA HPC SDK, and backends like Intel’s Threading Building Blocks (TBB), High-Performance ParalleX (HPX), and OpenMP.

In particular, we make the following contributions:

- (1) We introduce the benchmark suite pSTL-Bench, which is an extensible set of micro-benchmarks to assess the performance of parallel STL algorithms on different parallel architectures (multi-cores, GPUs).
- (2) Using the suite, we conduct a study over a selection of algorithms comparing the performance achieved on current multi-core architectures by different compiler frameworks and backends implementing the parallel STL.

The remainder of the paper is structured as follows. In Section 2, we give an overview of the field by summarizing the related work and current state of the art. Section 3 introduces the specifics of our proposed set of micro-benchmarks. In Section 4 we detail how the experiments were carried out before we show and analyze the experimental results in Section 5. Finally, we draw conclusions from the findings in Section 6 and outline future work.

2 RELATED WORK

Allowing for performance portability has always been a goal for programmers. This is especially true for developers on HPC systems, as novel HPC systems often provide new hardware architectures for which no efficient software solutions exist yet (cf. Jack Dongarra’s interview when receiving the ACM A.M. Turing Award [11]). The Message Passing Interface (MPI) is one of the standards that enables scientists to write efficient, parallel programs that are also portable across architectures. However, MPI has its limits, especially when it comes to efficiently programming multi-core systems or accelerators, such as GPUs or Xeon Phis.

For the reasons mentioned above, several parallel programming frameworks striving for performance portability were introduced. Note that the term “performance portability” may have completely different notions or definitions [18, 19]. Contrary to Pennycook et al.’s restrictive definition, in this work, we consider a program to be performance-portable if it can be executed efficiently on different architectures. We call a program efficient if it performs within a threshold of a specialized program, i.e., a scan algorithm implemented using a performance portability framework performs equally well on an NVIDIA GPU as the best-known CUDA implementation.

Our work focuses on micro-benchmarking algorithms from the ISO C++ standard library. There are several other feature-rich C++ frameworks that offer performance portability. For instance, Kokkos [23] and Raja [2] are one of the most prominent libraries that can be used to program portable HPC applications using different backends, such as SYCL [21], HPX [13], or OpenMP. Although Kokkos may use SYCL as a backend, SYCL itself is a C++ programming model that offers portability between heterogeneous compute resources. Similar to Kokkos, SYCL is a C++ abstraction layer that supports a range of processor architectures and accelerators [10]. For GPUs, the Thrust library [3] provides a high-level interface for STL to C++ programmers. Thrust is used as a backend in the C++ STL implementation of NVIDIA’s High-Performance Computing (HPC) Software Development Kit.

HPX [13] is another C++ abstraction layer for developing efficient parallel and concurrent applications. HPX can be used to write distributed applications using an Active Global Address Space (AGAS), which is an extension to PGAS for transparently moving global objects in between compute nodes [13]. SYCL, HPX, and Kokkos can also be used together to combine their strengths [7].

The benchmark suite SYCL-Bench [14] is a collection of different programs to analyze the performance of various SYCL implementations. Besides micro-benchmarks, SYCL-Bench also features real-world application benchmarks, such as 2DConvolution, as well as so-called runtime benchmarks, such as a DAG benchmark with many independent tasks. Pennycook et al. presented an analysis of the terminology used in SYCL and ISO C++ to shed light on the relationship between C++ concepts in both programming models [17].

This work was inspired by Lin et al. [15], who examined the ISO C++ capabilities for creating performance portable applications. In their work, they developed an ISO C++ parallel version of the BabelStream benchmark [8] and also adapted real-world applications, such as MiniBude, to utilize the parallel STL of ISO C++.

In the context of the present work, we will evaluate different implementations of the parallel STL for multi-core systems. That collection includes the HPX [13] library discussed above. We also analyze the performance of the GNU’s parallel STL, based on MC-STL [22] which uses OpenMP for parallelization and can be seen as the original parallel implementation of the STL. Currently, the GNU compiler and the Intel compiler use the Intel Threading Building Blocks (TBB) [20] for implementing the parallel C++ STL, whereas NVIDIA relies on several technologies, primarily OpenMP [6] and CUDA [16], for its parallel implementation of the STL.

Compared to the previous benchmarks, pSTL-Bench targets the scalability of the individual building blocks, in the same spirit as the OMPTB [12] or the EPCC OpenMP microbenchmark suite [4] for benchmarking OpenMP primitives.

3 MICRO-BENCHMARK SUITE PSTL-BENCH

We now present the micro-benchmark suite pSTL-Bench.

3.1 Idea and Goals

As mentioned previously, our primary goal for developing pSTL-Bench is to assess the scalability and the efficiency of the different implementations of the parallel C++ STL. Specifically, we aim to achieve several objectives. First, we would like to determine the

sweet-spot of problem size for each parallel STL algorithm, i.e., how large a problem has to be such that utilizing the parallel version is advantageous. Second, we aim to ascertain the maximum number of cores that can be effectively utilized by various parallel algorithms, many of which are memory-bound. Finally, we would like to examine the run-time achievable with each of the parallel STL implementations and backends. We intend to assess the run-time performance achievable with each of the parallel STL implementations and backends. To facilitate this comparison of performance across different standard library implementations, we have compiled these benchmarks into a suite named pSTL-Bench.¹

3.2 Description of Benchmark Kernels

For the sake of conciseness, we select and evaluate five algorithms of different computational structure in this paper. The `find` algorithm performs a linear search on the input array, while the `for_each` call performs a map operation on each element of the input array in parallel. The `reduce` call represents a parallel reduction operation, needed for map-reduce type programming. The `inclusive_scan` represents a typical, parallel prefix-sum operation, and `sort` was selected to include a parallel sorting function into the set of test algorithms.

A precise description of each of the algorithms is shown below:

- `find`: Given an array of n elements, $v = [1, 2, \dots, n]$, find a random element v_i , such that $v_i \in v$.
- `for_each`: Given an array of n elements, $v = [1, 2, \dots, n]$, compute for each element $\min\{\sin(v_i), \tan(v_i)\}$.
- `inclusive_scan`: Given an array of n elements, $v = [1, 2, \dots, n]$, compute the result of inclusive sum, r , such that $r_i = \sum_{j=1}^i v_j$.
- `reduce`: Given an array of n elements, $v = [1, 2, \dots, n]$, compute $\sum_{i=1}^n v_i$.
- `sort`: Sort the elements of v , where v is an array of n elements randomly shuffled, such that $v_i \in [1, n]$ and $v_i \neq v_j, \forall i \neq j$.

The full list of algorithms supported by pSTL-Bench can be found in Table 4 in Appendix A.

Since each implementation might have a slightly different interface, the calls to the algorithms are wrapped within lambda functions, as shown in Listings 1 and 2. These lambda functions are then called from a helper function that executes them repeatedly to measure their execution time with Google’s Benchmark library [9], see Listing 3. The wrapping function simply measures the time taken for the invocation of the function to be executed (function `f`).

This structure facilitates the incorporation of new benchmarks while the compiler can still inline the functions to avoid the potential overhead of calling the lambdas.

It should be mentioned that it is possible to manually select the input size and data types (`int`, `float`, `double`, etc.) used in the benchmarks, therefore, extending its customization and analysis capabilities.

¹Code can be provided on request.

Listing 1: Wrapper of the `for_each` benchmark using STL.

```

1 const auto for_each_std = [](auto && policy, const auto &
2   input_data, auto && f) {
3   std::for_each(policy, input_data.begin(),
4     input_data.end(), f);
};

```

Listing 2: Wrapper of the `for_each` benchmark using GNU's implementation.

```

1 const auto sort_gnu = []( [[maybe_unused]] auto &&
2   executionPolicy, auto & input_data) {
3   __gnu_parallel::sort(input_data.begin(),
4     input_data.end());
};

```

Listing 3: Example of the benchmark wrapper for `for_each`.

```

1 template<class Policy, class Function>
2 static void benchmark_for_each_wrapper(benchmark::State &
3   state, Function && f) {
4   constexpr auto execution_policy = Policy{};
5   const auto & size = state.range(0);
6   auto data = suite::generate_increment<Policy>(
7     execution_policy, size, 1);
8
9   for (auto _ : state) {
10    WRAP_TIMING(f(execution_policy, data, kernel));
11  }
12
13  state.SetBytesProcessed(state.iterations() * data.
14    size() * sizeof(int));
};

```

4 EXPERIMENTAL EVALUATION

For acquiring relevant results, experiments are carried out using different configurations in terms of hardware, software, number of threads, input sizes, and memory allocation.

4.1 Hardware and Software Setup

We conduct experiments on three different multi-core, shared-memory systems that comprise 32, 64, and 128 cores, which are called *Hydra*, *Nebula*, and *VSC-5*, respectively.² On these systems, we evaluate three different compilers: GCC, Intel oneAPI compiler, and NVIDIA HPC SDK. Also, several backends are tested, which are: Intel's Threading Building Blocks (TBB), High-Performance ParalleX (HPX), and OpenMP through the implementations of GNU and NVIDIA.

Furthermore, we also perform experiments on two GPU-based systems, *Tesla* and *Ampere*, using the NVIDIA HPC SDK with the CUDA compiler.

We provide an overview of the hardware and software details in Table 1.

4.2 Experimental Setup

In our experiments we test the parallel algorithms with different numbers of threads and problem sizes. Specifically, the thread counts range from 1 to the maximum available core count on each

²The names of the machines were intentionally anonymized for the dual-anonymous review process.

Listing 4: Function `allocate` for the custom parallel allocator.

```

1 pointer allocate(size_type cnt, void const * = nullptr) {
2   // allocate memory
3   auto p = static_cast<pointer*> (::operator new(cnt *
4     sizeof(T)));
5
6   // touch the first byte of every object
7   std::for_each(execution_policy, begin, begin + cnt,
8     [](T & val) { *reinterpret_cast<char*>(&val) =
9     0; });
10
11  // return the overall memory block
12  return p;
};

```

machine, so the values 1, 2, 4, ..., #cores are used. Concerning problem sizes, inputs vary from 2^3 to 2^{30} elements, using 32-bit integers, enabling the testing of a broad set of inputs that are accommodated in various cache levels or necessitate consistent DRAM traffic.

4.3 The Impact of Memory Allocation

Memory allocation might have a significant impact on performance in modern computer architectures. For that reason, a custom parallel allocator is used for the experiments. Through the first-touch policy of Linux, this allocator makes each thread touch the first byte of each object with the given parallel policy, as shown in Listing 4. This enforces the correct placement of threads and memory pages. The parallel allocator used in this work is an adapted version of the NUMA allocator that is part of HPX.³

Fig. 1 shows the speedup achieved when using the custom parallel allocator in *Hydra*. Using this allocator leads to an increase in performance in three of the five algorithms. For the `X::find` algorithm, the improvement goes from 11% to 39%, from -13% to 74% for the `X::inclusive_scan` and 56% to 194% for `X::reduce`. In the `X::for_each` and the `X::sort` algorithms, the differences are smaller, always below 2.5% and typically around 0.5%.

In only one experiment (`X::inclusive_scan`), the parallel allocator notably reduced performance, that is with the NVIDIA compiler and the OpenMP backend.

Considering the results of the parallel allocator, the rest of the benchmarks are executed using it except for those tests related to HPX, which is incompatible with it.

5 EXPERIMENTAL RESULTS


Table 2 summarises the results obtained, showing the speedup compared to the sequential implementation with GCC. In this table, we use a fixed baseline performance, GCC's sequential performance in this case, which may lead to speedups that are large than the total core count.

Table 3 shows the maximum number of threads for which parallel efficiency is above 70%. The purpose of this table is to analyze the number of threads that can be effectively utilized without excessively wasting resources. The threshold of 70% is somewhat arbitrarily defined, but it assists in estimating the effectiveness of the parallelization.

³For the NUMA allocator of HPX, see https://github.com/STELLAR-GROUP/hpx/blob/fe048ee6e01abedad0a60afd0c204116419871c3/libs/core/compute_local/include/hpx/compute_local/host/numa_allocator.hpp

Table 1: Summary of the hardware and software used in our study.

machine	<i>Hydra</i>	<i>Nebula</i>	<i>VSC-5</i>	<i>Tesla</i>	<i>Ampere</i>
CPU/GPU	Intel Xeon 6130F	AMD EPYC 7551	AMD EPYC 7713	NVIDIA Tesla P4	NVIDIA Ampere A2
core frequency	2.10 GHz	2.00 GHz	2.00 GHz	1.11 GHz	1.77 GHz
#sockets	2	2	2	—	—
#cores (node)	32	64	128	2560	1280
#threads (node)	32	64	256	2560	1280
compilers	g++ 12.1.0 nvc++ 22.11 icpx 2021.7.0	g++ 12.3.0 nvc++ 23.7	g++ 12.2.0 nvc++ 22.9 icpx 2022.2.1	g++ 10.2.1 nvc++ 23.5	g++ 10.2.1 nvc++ 23.5
libraries	TBB 2021.9.0 HPX 1.9.0 GOMP (from g++ 12.1.0) NVOMP 22.11	TBB 2021.10.0 HPX 1.9.1 GOMP (from g++ 12.3.0) NVOMP 22.11	TBB 2021.7.0 HPX 1.8.1 GOMP (from g++ 12.2.0) NVOMP 22.9	CUDA 11.8	CUDA 12.2



Figure 1: Speedup when using custom parallel allocator with 32 threads and a problem size of 2^{30} elements in *Hydra*. Higher is better.

All the data presented in this section are derived from the median of ten executions.

5.1 X::find

Figure 2 shows how the execution times grow with the size of the problem in *Hydra*. The behavior for other machines is shown in Appendix B. For small problem sizes, the sequential implementation offers better execution times, with a difference of orders of magnitude, highlighting the overhead of launching parallel sections. This difference is consistently reduced until problem sizes of around 2^{16} , where the performance of the parallelism starts to compensate for the overhead. It is interesting to note that the GNU parallel implementation falls back to the sequential implementation until the problem size reaches 2^9 . From problem sizes of 2^{20} , parallel implementations begin to be significantly faster than the sequential one. These behaviors, regarding the input size, can be extended to the rest of the benchmarks.

Related to the speedup, efficiency is an issue, where the maximum speedup achieved was 10x with the Intel compiler and using


Figure 2: Execution time scalability with the problem size in *Hydra*. All cores were used except for GCC’s sequential implementation. Lower is better.

TBB with 32 threads in *Hydra*. Also, using more than 32 threads results in little, if any, gains in performance.

5.2 X::for_each

Figure 3 shows the speedup obtained for a problem size of 2^{30} elements. Note that the Intel compiler and, particularly, the NVIDIA compiler achieve a better performance than GCC. Profiling of the code indicates that besides using its math library, the NVIDIA and Intel compilers make better use of vector instructions in comparison to GNU’s. Interestingly, the NVIDIA compiler uses 256-bit instructions, while the Intel compiler uses only 128-bit instructions.

Additionally, data in Table 3 shows an almost perfect scaling. This is reasonable given the *embarrassingly parallel* nature of the algorithm.

5.3 X::inclusive_scan

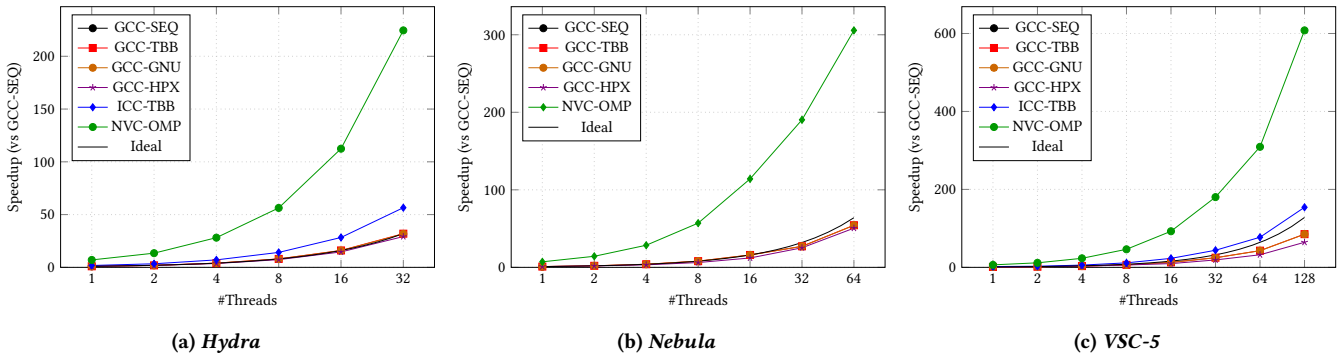
First, it should be noted that the GNU’s collection of parallel algorithms does not include the `inclusive_scan` function.

Table 2: Speedup against GCC’s sequential implementation for machines *Hydra*, *Nebula*, and *VSC-5*, with 32, 64, and 128 cores, respectively. Notation is *Hydra/Nebula/VSC-5*. Higher is better.

	X::find	X::for_each	X::inclusive_scan	X::reduce	X::sort
GCC-TBB	9.54 / 3.99 / 5.83	31.92 / 54.31 / 85.18	4.72 / 3.31 / 4.80	12.29 / 6.25 / 7.14	9.51 / 10.04 / 9.42
GCC-GNU	6.72 / 1.46 / 2.65	31.68 / 54.15 / 85.01	— / — / —	15.08 / 5.49 / 7.20	28.09 / 28.07 / 65.56
GCC-HPX	7.30 / 1.07 / 2.31	29.17 / 50.52 / 64.33	2.18 / 0.79 / 0.84	9.46 / 1.14 / 1.76	10.24 / 9.10 / 8.46
ICC-TBB	10.02 / — / 5.77	56.53 / — / 154.08	4.68 / — / 4.81	12.37 / — / 7.19	10.05 / — / 9.78
NVC-OMP	4.78 / 1.07 / 1.31	224.55 / 305.65 / 608.00	0.88 / 0.62 / 0.96	23.77 / 17.91 / 11.54	9.14 / 7.45 / 6.76

Table 3: Maximum number of threads for which parallel efficiency is above 70% (compared to the execution with 1 thread) for machines *Hydra*, *Nebula*, and *VSC-5*. Notation is *Hydra/Nebula/VSC-5*. Higher is better.

	X::find	X::for_each	X::inclusive_scan	X::reduce	X::sort
GCC-TBB	1 / 2 / 1	32 / 64 / 32	1 / 1 / 1	8 / 4 / 1	4 / 8 / 4
GCC-GNU	1 / 1 / 1	32 / 64 / 32	— / — / —	1 / 4 / 1	32 / 32 / 16
GCC-HPX	16 / 4 / 1	32 / 64 / 8	4 / 2 / 1	8 / 2 / 1	8 / 4 / 4
ICC-TBB	1 / — / 1	32 / — / 16	1 / — / 1	8 / — / 1	8 / — / 4
NVC-OMP	8 / 4 / 1	32 / 32 / 128	1 / 1 / 1	1 / 2 / 1	2 / 2 / 4

**Figure 3: Speedup against GCC’s sequential implementation for benchmark X::for_each. Higher is better.**

Based on the results shown in Figures 4, 5, and Table 2, it is evident that the NVIDIA implementation yields limited speedup for this operation, being even slower than the sequential version. This is because this implementation falls back to a sequential execution, which is negatively affected by the custom parallel allocator used.

The rest of the implementations produce slight improvements in performance when using parallelism. For example, a maximum speedup of 4.81 \times can be achieved using TBB as the backend. However, the efficiency of these implementations is far from ideal.

5.4 X::reduce

It should be noted that the GNU’s collection of algorithms does not include a reduction function so the accumulate function has been used instead.

As shown in Figure 6 the speedup behavior changes from system to system. In *Hydra* (Figure 6a), for example, a good speedup is achieved up to 16 threads, with a maximum speedup of 23.77 \times .

However, the performance achieved with the NVIDIA compiler with the OpenMP backend stands out, especially in *Nebula* and

VSC-5. In these systems, the rest of the compilers and backends present poor scalability, especially when more than 16 threads are used, see Figures 6b and 6c.

5.5 X::sort

Two different factors are relevant when analyzing the results for this benchmark, the problem size and the number of threads.

For small sizes (see Figure 7), below 2^{12} , the NVC-OMP shows a significantly higher overhead than the rest of compilers and backends. For bigger sizes, it is the GNU implementation that obtains better results.

Regarding the number of threads, when this is 8 or less, the NVIDIA compiler achieves a considerably better performance than the rest, see Figure 8. Nevertheless, with more threads, the implementation of GNU which uses OpenMP establishes itself as a clear winner.

It is interesting to note that GNU’s efficiency is above 70% up to 32 threads in *Hydra* and *Nebula*, but only up to 16 threads in

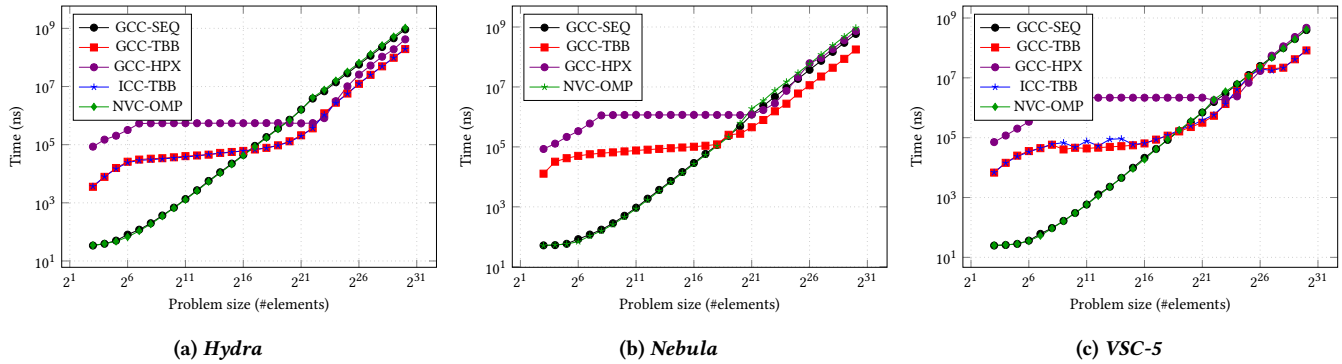


Figure 4: Execution time scalability with the problem size. All cores used except for GCC’s seq. implementation. Benchmark X::inclusive_scan. Lower is better.

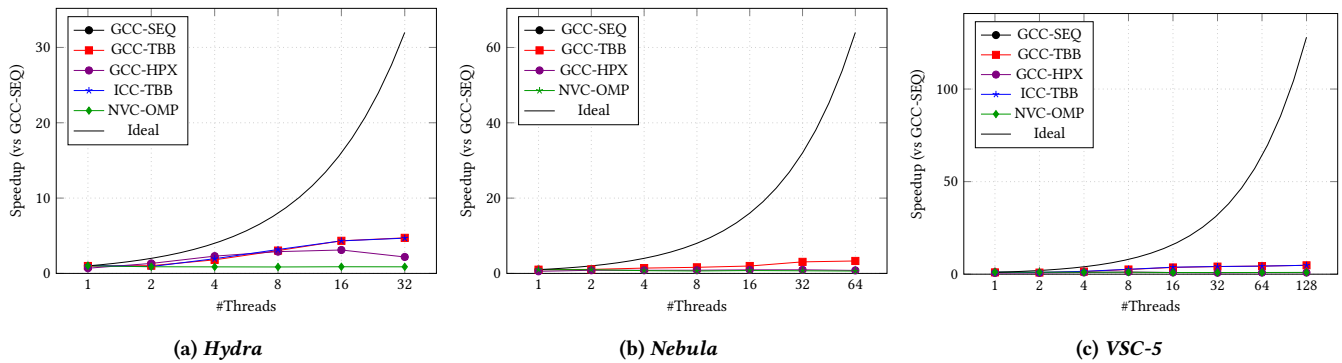


Figure 5: Speedup against GCC’s seq. implementation. Problem size is 2^{30} . Benchmark X::inclusive_scan. Higher is better.

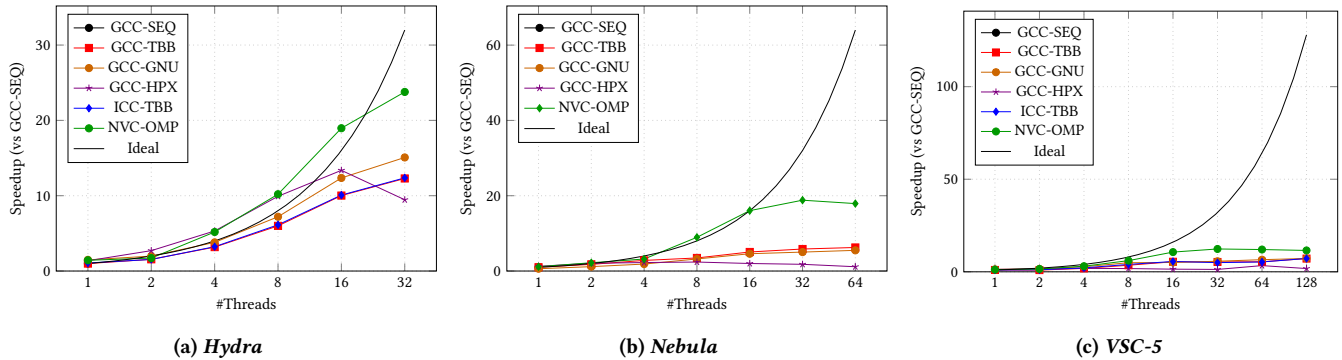


Figure 6: Speedup against GCC’s seq. implementation. Problem size is 2^{30} . Benchmark X::reduce. Higher is better.

VSC-5, see Table 3. The rest of the backends, cannot use more than 8 threads while keeping an efficiency of 70 % or better.

5.6 Performance on GPUs

Figure 9 includes the execution times of the benchmarks on *Hydra*, *Tesla* and *Ampere*. Results show that both the input size and the particular algorithm play a key role in terms of performance for the GPUs compared to the CPU.

First, a substantial penalty is paid in the form of overhead when using a small input size. Also, transferring a large amount of data seems to penalize the performance of the GPUs as well, see Figures 9c and 9d.

Regarding the algorithms, the X::find and X::reduce have poor performance on the GPUs. However, the X::for_each, the X::inclusive_scan, and the X::sort achieve significantly better execution times, especially for the sizes ranging from 2^{16} to 2^{27} elements.




Figure 7: Execution time scalability with the problem size. All cores used except for GCC’s seq. implementation. Benchmark X : : sort. Lower is better.

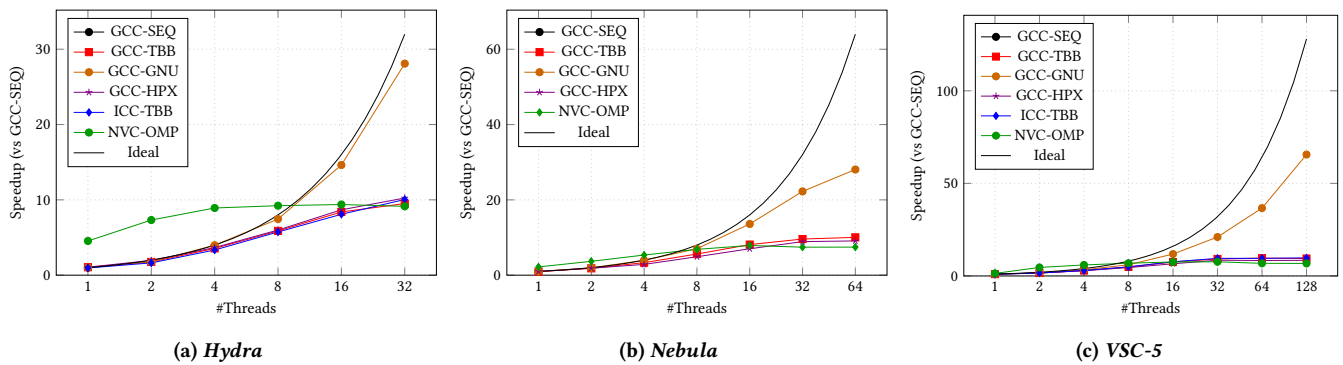


Figure 8: Speedup against GCC’s seq. implementation. Problem size is 2^{30} . Benchmark X : : sort. Higher is better.

6 CONCLUSIONS

In this work, we introduced a specialized set of benchmarks named pSTL-Bench to assess the performance of the parallel algorithms present in the C++ standard template library. With this suite, we aim to facilitate the users comparing compilers and backends to decide which fits best the particular characteristics of their system.

Additionally, using this set of benchmarks, we conducted a comprehensive analysis of performance on current multi-core architectures by benchmarking some of the most used algorithms in STL. This analysis took into account not only the compiler and backend being used but also the input size, the number of threads used, and the memory allocation.

First, we showed the potential performance improvements when using a custom parallel allocator in NUMA systems. Generally, execution times are improved, up to 194% in the best case, and minor losses in some scenarios.

Second, when using small input sizes it is possible to note the large overhead of launching parallel sections and the differences between backends, being the TBB typically lower than those based on OpenMP or HPX. It is interesting to note as well that the GNU implementation falls back to the sequential execution for inputs lower than 512 elements.

On the other hand, for large input sizes, the parallel implementations show better in the vast majority of scenarios. For example,

the speedup and efficiency in the X : : for_each stand out compared to the rest of the algorithms. Nevertheless, some algorithms like X : : find or X : : inclusive_scan show poor speedup and, consequently, a bad efficiency.

Third, it is interesting to note that there might be large differences in performance due to vectorization. In our experiments, we found that the NVIDIA compiler applies vector instruction more often than the rest of the compilers, particularly in the X : : for_each algorithm.

Finally, we evaluated the performance of the CUDA backend for the NVIDIA compiler with two different GPUs. We found high penalties for small inputs, but also for high inputs where the data transfers dominate the execution times. Nevertheless, for medium problem sizes and some particular algorithms, the GPUs are capable of outperforming the CPUs.

Regarding future work, we would like to expand our benchmark suite, to include more algorithms and also to support more compilers and backends. Similarly, it would be interesting to extend our experimental analysis on more architectures, particularly GPUs and, potentially, FPGAs.

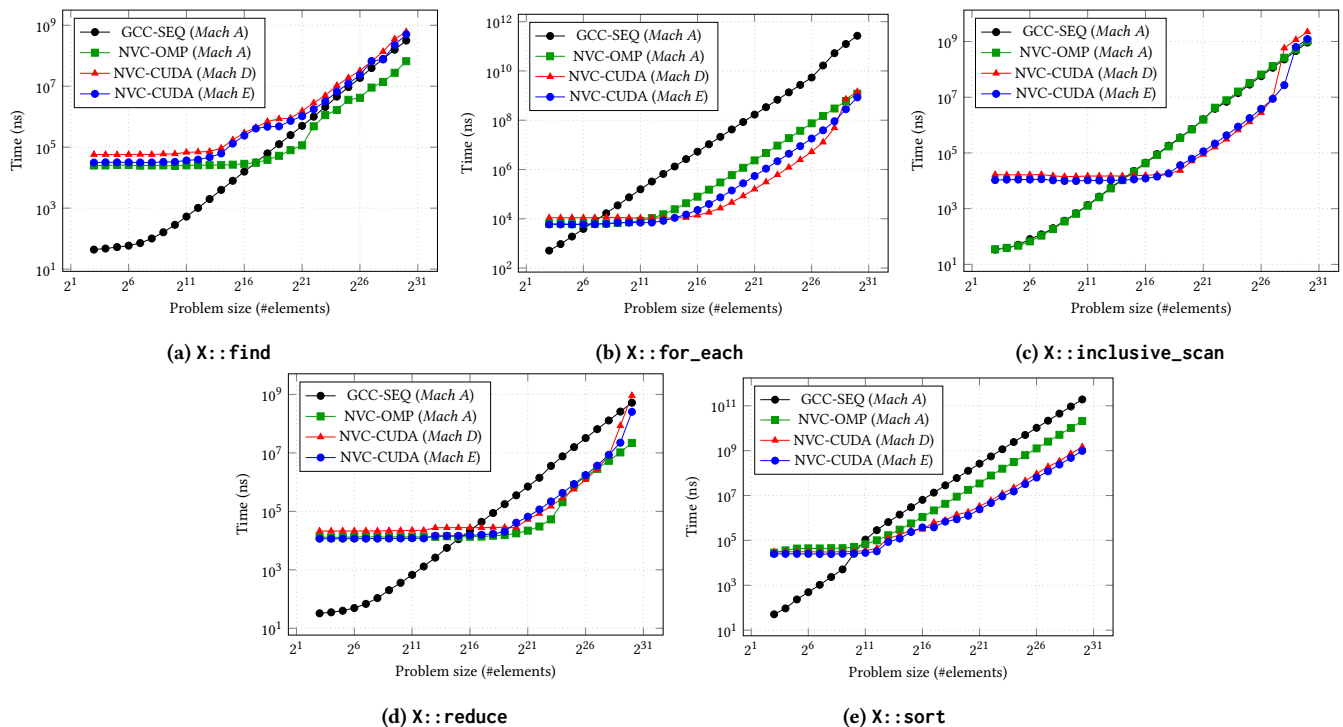


Figure 9: Execution time scalability with the problem size. All cores used except for GCC’s seq. implementation. Lower is better.

REFERENCES

- [1] Victor Artigues, Katharina Kormann, Markus Rampp, and Klaus Reuter. 2020. Evaluation of performance portability frameworks for the implementation of a particle-in-cell code. *Concurr. Comput. Pract. Exp.* 32, 11 (2020). <https://doi.org/10.1002/CPE.5640>
- [2] David Beckingsale, Thomas R. W. Scogland, Jason Burmark, Rich Hornung, Holger Jones, William Killian, Adam J. Kunen, Olga Pearce, Peter Robinson, and Brian S. Ryuji. 2019. RAJA: Portable Performance for Large-Scale Scientific Applications. In *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC, P3HPC@SC 2019, Denver, CO, USA, November 22, 2019*. IEEE, 71–81. <https://doi.org/10.1109/P3HPC49587.2019.00012>
- [3] Nathan Bell and Jared Hoberock. 2012. Chapter 26 - Thrust: A Productivity-Oriented Library for CUDA. In *GPU Computing Gems Jade Edition*, Wen mei W. Hwu (Ed.). Morgan Kaufmann, Boston, 359–371. <https://doi.org/10.1016/B978-0-12-385963-1.00026-5>
- [4] J. Mark Bull, Fiona Reid, and Nicola McDonnell. 2012. A Microbenchmark Suite for OpenMP Tasks. In *OpenMP in a Heterogeneous World - 8th International Workshop on OpenMP, IWOMP 2012, Rome, Italy, June 11-13, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7312)*, Barbara M. Chapman, Federico Massaioli, Matthias S. Müller, and Marco Rorro (Eds.). Springer, 271–274. https://doi.org/10.1007/978-3-642-30961-8_22
- [5] Lyndon Clarke, Ian Glendinning, and Rolf Hempel. 1994. The MPI Message Passing Interface Standard. In *Programming Environments for Massively Parallel Distributed Systems*, Karsten M. Decker and René M. Rehmman (Eds.). Birkhäuser Basel, Basel, 213–218.
- [6] Leonardo Dagum and Ramesh Menon. 1998. OpenMP: An industry standard API for shared-memory programming. *Computational Science & Engineering, IEEE* 5, 1 (1998), 46–55.
- [7] Gregor Daiß, Patrick Diehl, Hartmut Kaiser, and Dirk Pflüger. 2023. Stellar Mergers with HPX-Kokkos and SYCL: Methods of using an Asynchronous Many-Task Runtime System with SYCL. In *Proceedings of the 2023 International Workshop on OpenCL, IWOCCL 2023, Cambridge, United Kingdom, April 18-20, 2023*. ACM, 8:1–8:12. <https://doi.org/10.1145/3585341.3585354>
- [8] Tom Deakin, James Price, Matt Martineau, and Simon McIntosh-Smith. 2018. Evaluating attainable memory bandwidth of parallel programming models via BabelStream. *Int. J. Comput. Sci. Eng.* 17, 3 (2018), 247–262. <https://doi.org/10.1504/IJCSE.2018.095847>
- [9] Google. 2016. *Google Benchmark*. <https://github.com/google/benchmark>
- [10] Jeff R. Hammond, Michael Kinsner, and James C. Brodman. 2019. A comparative analysis of Kokkos and SYCL as heterogeneous, parallel programming models for C++ applications. In *Proceedings of the International Workshop on OpenCL, IWOCCL 2019, Boston, MA, USA, May 13-15, 2019*. ACM, 15:1–15:2. <https://doi.org/10.1145/3318170.3318193>
- [11] Leah Hoffmann. 2022. Learning new things and avoiding obstacles. *Commun. ACM* 65, 6 (2022). <https://doi.org/10.1145/3530690>
- [12] Sascha Hunold and Klaus Krafnitzner. 2022. A Quantitative Analysis of OpenMP Task Runtime Systems. In *Benchmarking, Measuring, and Optimizing - 14th BenchCouncil International Symposium, Bench 2022, Virtual Event, November 7-9, 2022, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 13852)*, Ana Gainaru, Ce Zhang, and Chunjie Luo (Eds.). Springer, 3–18. https://doi.org/10.1007/978-3-031-31180-2_1
- [13] Hartmut Kaiser, Patrick Diehl, Adrian S. Lemoine, Bryce Leibach, Parsa Amini, Agustín Berge, John Biddiscombe, Steven R. Brandt, Nikunj Gupta, Thomas Heller, Kevin A. Huck, Zahra Khatami, Alireza Kheirkhahan, Auriane Reverdell, Shahrzad Shirzad, Mikael Simberg, Bibek Wagle, Weile Wei, and Tianyi Zhang. 2020. HPX - The C++ Standard Library for Parallelism and Concurrency. *J. Open Source Softw.* 5, 53 (2020), 2352. <https://doi.org/10.21105/joss.02352>
- [14] Sohan Lal, Aksel Alpay, Philip Salzmann, Biagio Cosenza, Alexander Hirsch, Nicolai Stawinoga, Peter Thoman, Thomas Fahringer, and Vincent Heuveline. 2020. SYCL-Bench: A Versatile Cross-Platform Benchmark Suite for Heterogeneous Computing. In *Euro-Par 2020: Parallel Processing - 26th International Conference on Parallel and Distributed Computing, Warsaw, Poland, August 24-28, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12247)*, Maciej Malawski and Krzysztof Rzadca (Eds.). Springer, 629–644. https://doi.org/10.1007/978-3-030-57675-2_33
- [15] Wei-Chen Lin, Tom Deakin, and Simon McIntosh-Smith. 2022. Evaluating ISO C++ Parallel Algorithms on Heterogeneous HPC Systems. In *IEEE/ACM International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, PMBS@SC 2022, Dallas, TX, USA, November 13-18, 2022*. IEEE, 36–47. <https://doi.org/10.1109/PMBS56514.2022.00009>
- [16] John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. 2008. Scalable Parallel Programming with CUDA. *Queue* 6, 2 (March 2008), 40–53. <https://doi.org/10.1145/1365490.1365500>
- [17] Simon John Pennycook, Ben Ashbaugh, James C. Brodman, Michael Kinsner, Steffen Larsen, Gregory Lueck, Roland Schulz, and Michael Voss. 2023. Towards Alignment of Parallelism in SYCL and ISO C++. In *Proceedings of the 2023 International Workshop on OpenCL, IWOCCL 2023, Cambridge, United Kingdom, April*

- 18-20, 2023. ACM, 14:1–14:9. <https://doi.org/10.1145/3585341.3585371>
- [18] S. John Pennycook and Jason D. Sewall. 2021. Revisiting a Metric for Performance Portability. In *International Workshop on Performance, Portability and Productivity in HPC, P3HPC@SC 2021, St. Louis, MO, USA, November 14, 2021*. IEEE, 1–9. <https://doi.org/10.1109/P3HPC54578.2021.00004>
- [19] Simon J. Pennycook, Jason D. Sewall, and Victor W. Lee. 2019. Implications of a metric for performance portability. *Future Gener. Comput. Syst.* 92 (2019), 947–958. <https://doi.org/10.1016/j.future.2017.08.007>
- [20] James Reinders. 2007. *Intel Threading Building Blocks: Outfitting C++ for Multi-core Processor Parallelism*. O'Reilly. <http://www.oreilly.com/catalog/9780596514808/index.html>
- [21] Ruymán Reyes, Gordon Brown, Rod Burns, and Michael Wong. 2020. SYCL 2020: More than meets the eye. In *IWOCL '20: International Workshop on OpenCL, Virtual Event / Munich, Germany, April 27-29, 2020*, Simon McIntosh-Smith (Ed.). ACM, 4:1. <https://doi.org/10.1145/3388333.3388649>
- [22] Johannes Singler, Peter Sanders, and Felix Putze. 2007. MCSTL: The Multi-core Standard Template Library. In *Euro-Par 2007, Parallel Processing, 13th International Euro-Par Conference, Rennes, France, August 28-31, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4641)*, Anne-Marie Kermarrec, Luc Bougé, and Thierry Priol (Eds.). Springer, 682–694. [https://doi.org/10.1007/978-3-540-74466-5_7}{2}](https://doi.org/10.1007/978-3-540-74466-5_7}{2)
- [23] Christian R. Trott, Damien Lebrun-Grandié, Daniel Arndt, Jan Ciesko, Vinh Q. Dang, Nathan D. Ellingwood, Rahul Kumar Gayatri, Evan Harvey, Daisy S. Hollman, Dan Ibanez, Nevin Liber, Jonathan R. Madsen, Jeff Miles, David Poliakoff, Amy Powell, Sivasankaran Rajamanickam, Mikael Simberg, Dan Sunderland, Bruno Turcksin, and Jeremiah J. Wilke. 2022. Kokkos 3: Programming Model Extensions for the Exascale Era. *IEEE Trans. Parallel Distributed Syst.* 33, 4 (2022), 805–817. <https://doi.org/10.1109/TPDS.2021.3097283>

A SUPPORT OF PARALLEL IMPLEMENTATIONS

Table 4: List of algorithms with support for parallelization in STL. Those with benchmarks in pSTL-Bench are noted in gray.

adjacent_difference	adjacent_find	all_of
any_of	copy	copy_if
copy_n	count	count_if
destroy	destroy_n	equal
exclusive_scan	fill	fill_n
find	find_end	find_first_of
find_if	find_if_not	for_each
for_each_n	generate	generate_n
includes	inclusive_scan	inplace_merge
is_heap	is_heap_until	is_partitioned
is_sorted	is_sorted_until	lexicographical_compare
max_element	merge	min_element
minmax_element	mismatch	move
none_of	nth_element	partial_sort
partial_sort_copy	partition	partition_copy
reduce	remove	remove_copy
remove_copy_if	remove_if	replace
replace_copy	replace_copy_if	replace_if
reverse	reverse_copy	rotate
rotate_copy	search	search_n
set_difference	set_intersection	set_symmetric_difference
set_union	sort	stable_sort
stable_partition	swap_ranges	transform
transform_exclusive_scan	transform_inclusive_scan	transform_reduce
uninitialized_copy	uninitialized_copy_n	uninitialized_default_construct
uninitialized_default_construct_n	uninitialized_fill	uninitialized_fill_n
uninitialized_move	uninitialized_move_n	uninitialized_value_construct
uninitialized_value_construct_n	unique	unique_copy

B X::FIND




Figure 10: Results for benchmark X::find; Hydra.

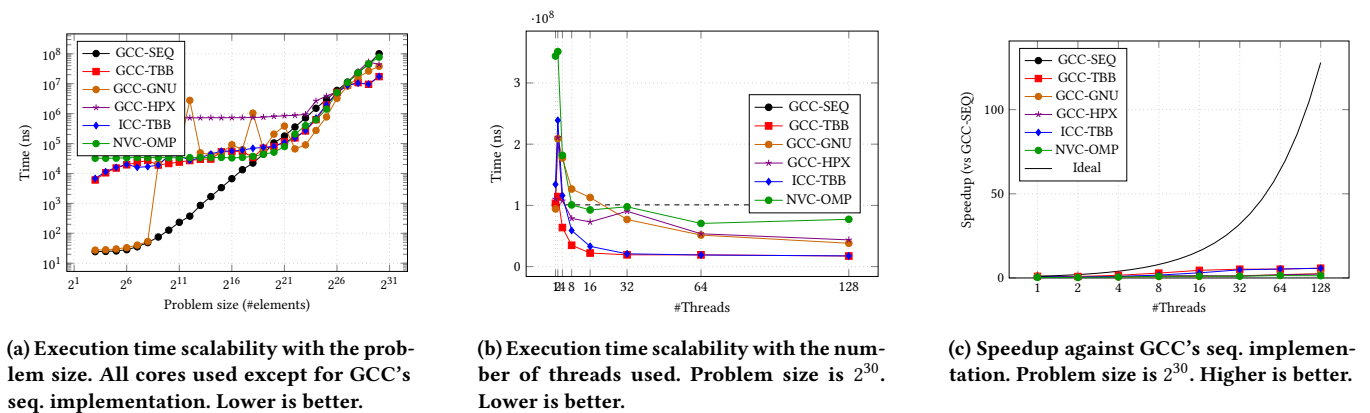

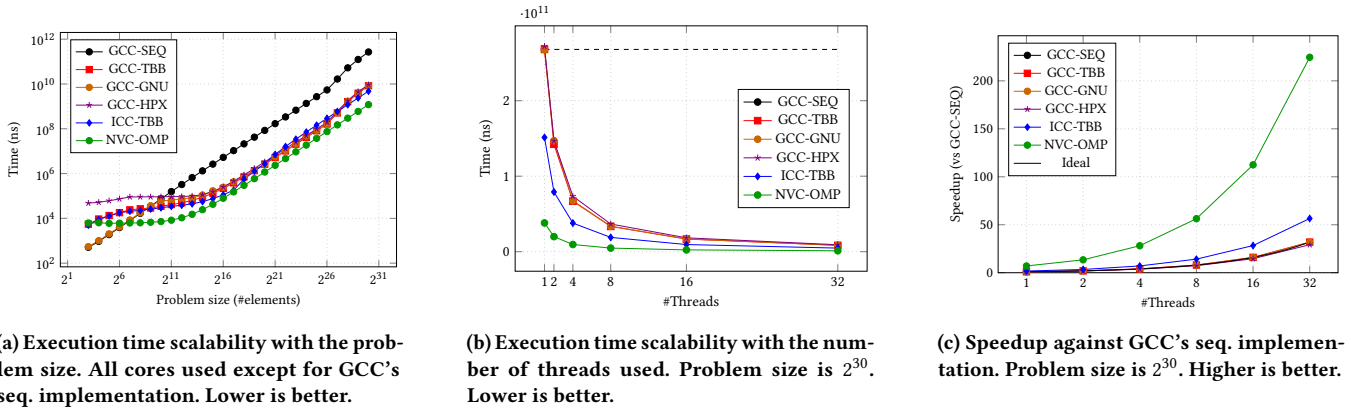


Figure 12: Results for benchmark X::find; VSC-5.

C X : : FOR_EACH

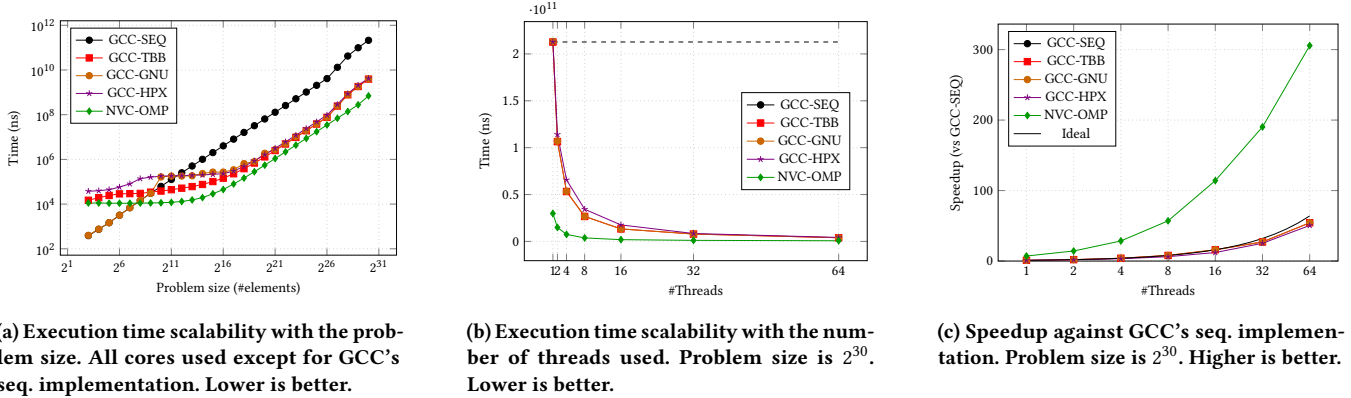


(a) Execution time scalability with the problem size. All cores used except for GCC's seq. implementation. Lower is better.

(b) Execution time scalability with the number of threads used. Problem size is 2^{30} . Lower is better.

(c) Speedup against GCC's seq. implementation. Problem size is 2^{30} . Higher is better.

Figure 13: Results for benchmark X : : for_each; Hydra.

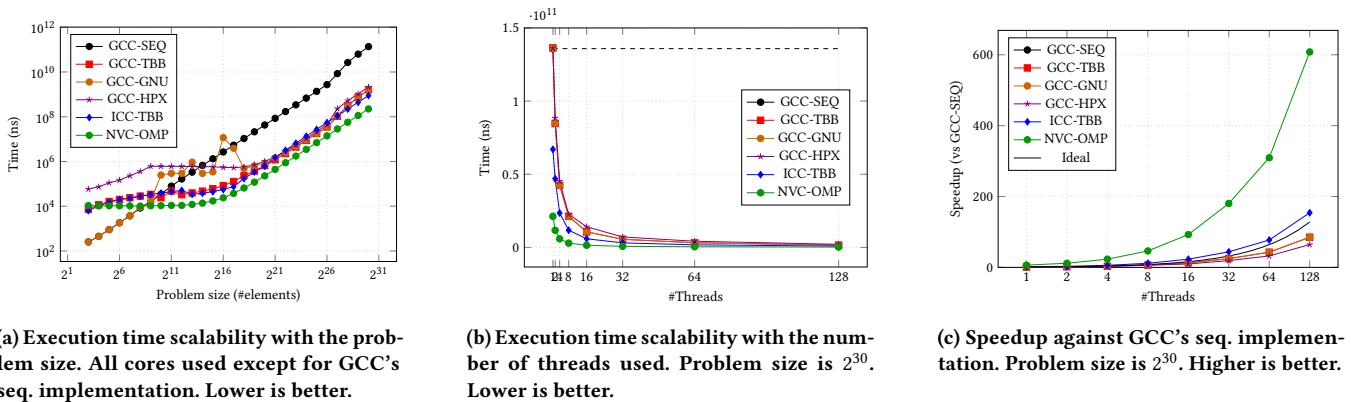


(a) Execution time scalability with the problem size. All cores used except for GCC's seq. implementation. Lower is better.

(b) Execution time scalability with the number of threads used. Problem size is 2^{30} . Lower is better.

(c) Speedup against GCC's seq. implementation. Problem size is 2^{30} . Higher is better.

Figure 14: Results for benchmark X : : for_each; Nebula.



(a) Execution time scalability with the problem size. All cores used except for GCC's seq. implementation. Lower is better.

(b) Execution time scalability with the number of threads used. Problem size is 2^{30} . Lower is better.

(c) Speedup against GCC's seq. implementation. Problem size is 2^{30} . Higher is better.

Figure 15: Results for benchmark X : : for_each; VSC-5.

D X::INCLUSIVE_SCAN

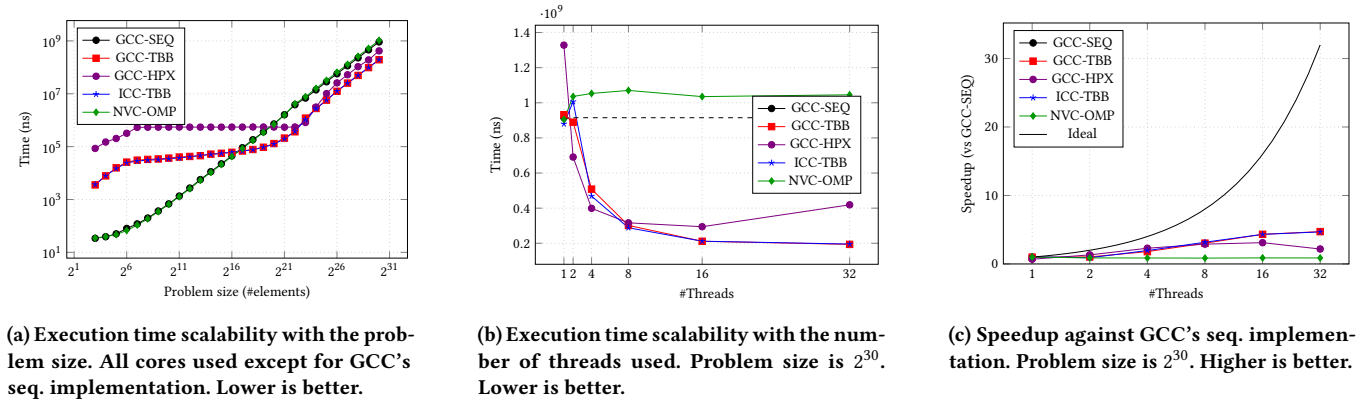


Figure 16: Results for benchmark X::inclusive_scan; Hydra.

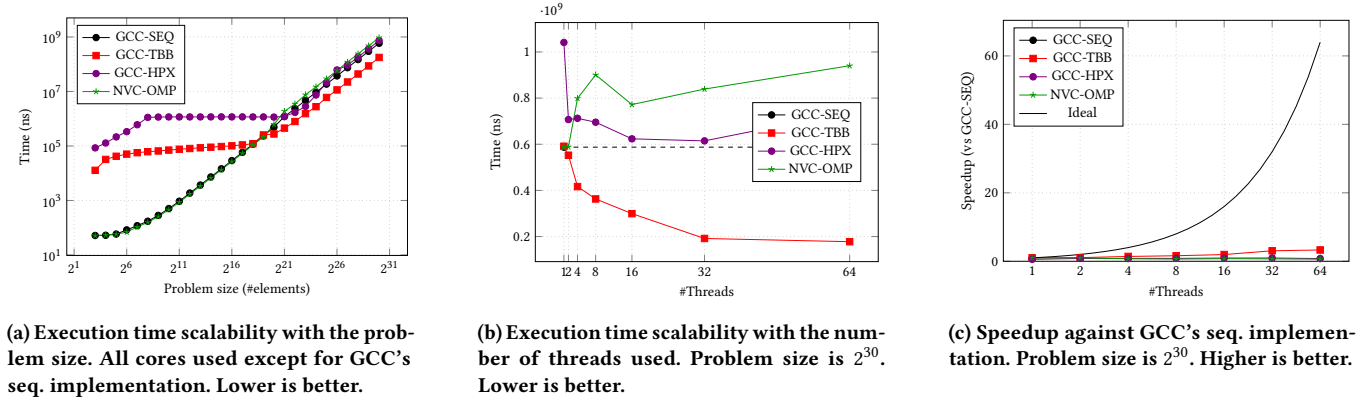


Figure 17: Results for benchmark X::inclusive_scan; Nebula.

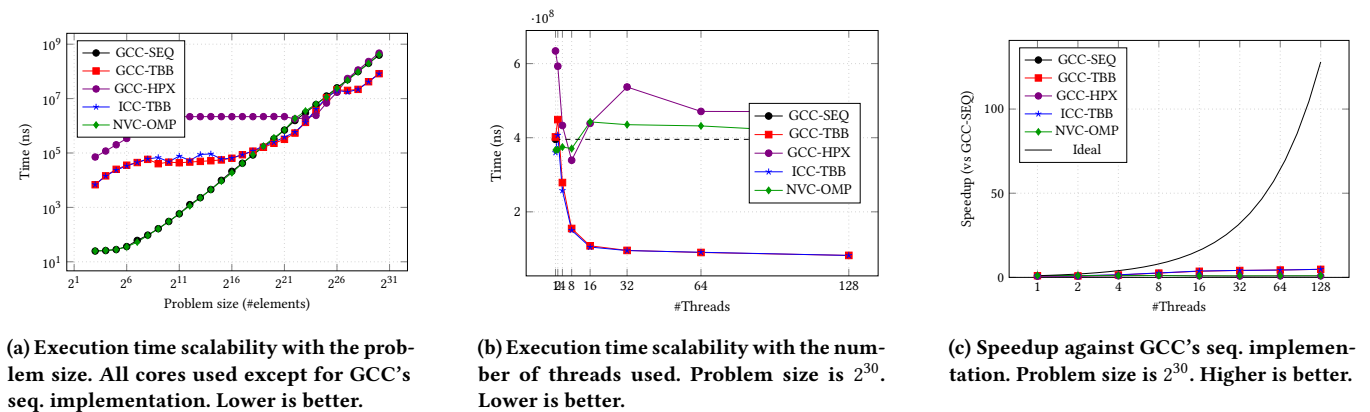
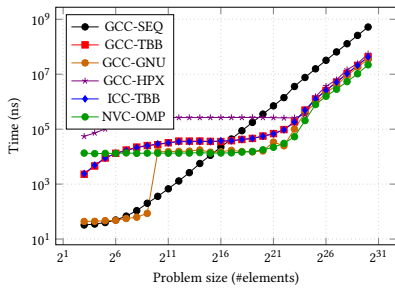
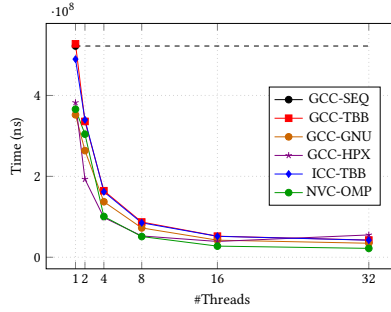


Figure 18: Results for benchmark X::inclusive_scan; VSC-5.


E X: :REDUCE



(a) Execution time scalability with the problem size. All cores used except for GCC's seq. implementation. Lower is better.

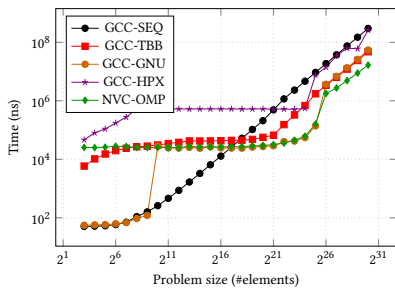


(b) Execution time scalability with the number of threads used. Problem size is 2^{30} . Lower is better.

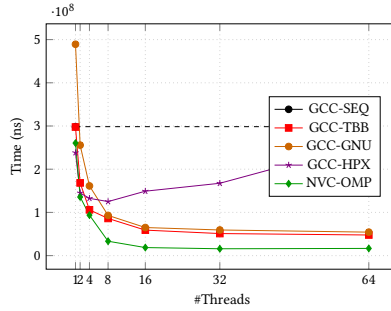


(c) Speedup against GCC's seq. implementation. Problem size is 2^{30} . Higher is better.


Figure 19: Results for benchmark X: :reduce; *Hydra*.



(a) Execution time scalability with the problem size. All cores used except for GCC's seq. implementation. Lower is better.

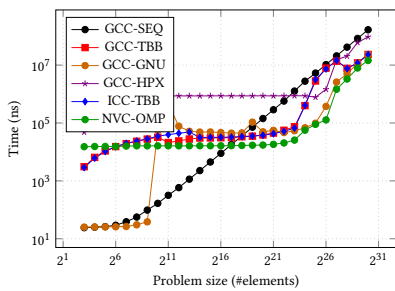


(b) Execution time scalability with the number of threads used. Problem size is 2^{30} . Lower is better.

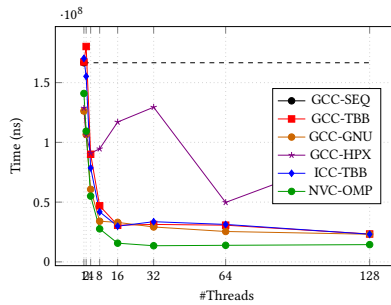


(c) Speedup against GCC's seq. implementation. Problem size is 2^{30} . Higher is better.


Figure 20: Results for benchmark X: :reduce; *Nebula*.



(a) Execution time scalability with the problem size. All cores used except for GCC's seq. implementation. Lower is better.



(b) Execution time scalability with the number of threads used. Problem size is 2^{30} . Lower is better.



(c) Speedup against GCC's seq. implementation. Problem size is 2^{30} . Higher is better.

Figure 21: Results for benchmark X: :reduce; *VSC-5*.

F X::SORT




Figure 22: Results for benchmark X::sort; Hydra.




Figure 23: Results for benchmark X::sort; Nebula.

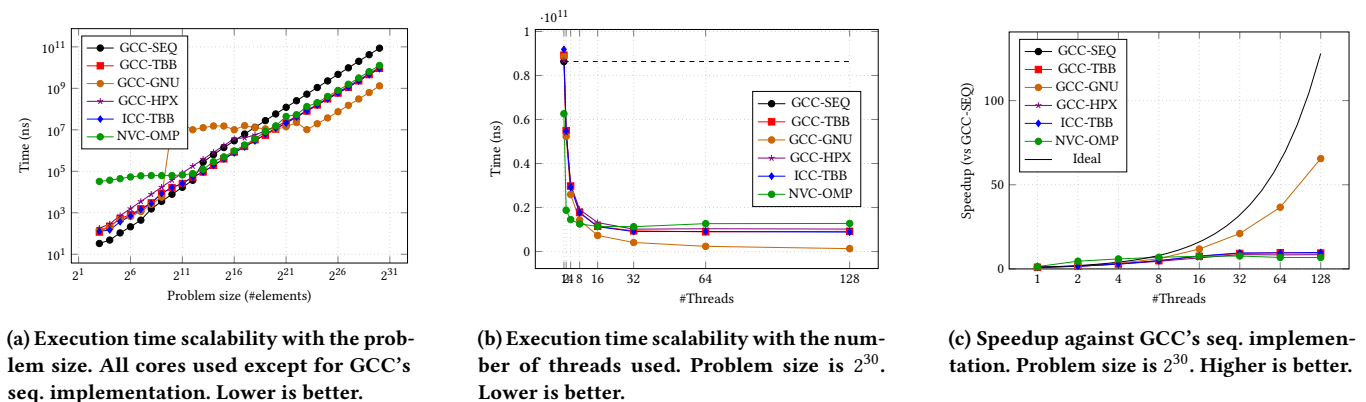


Figure 24: Results for benchmark X::sort; VSC-5.