

# Applying Sorting Networks to Synthesize Optimized Sorting Libraries <sup>\*</sup>

Michael Codish<sup>1</sup>, Luís Cruz-Filipe<sup>2</sup>,  
Markus Nebel<sup>2</sup>, and Peter Schneider-Kamp<sup>2</sup>

<sup>1</sup> Department of Computer Science, Ben-Gurion University of the Negev, Israel

<sup>2</sup> Dept. Mathematics and Computer Science, Univ. of Southern Denmark, Denmark

**Abstract.** This paper shows an application of the theory of sorting networks to facilitate the synthesis of optimized general purpose sorting libraries. Standard sorting libraries are often based on combinations of the classic Quicksort algorithm with insertion sort applied as the base case for small fixed numbers of inputs. Unrolling the code for the base case by ignoring loop conditions eliminates branching and results in code which is equivalent to a sorting network. This enables the application of further program transformations based on sorting network optimizations, and eventually the synthesis of code from sorting networks. We show that if considering the number of comparisons and swaps then theory predicts no real advantage of this approach. However, significant speed-ups are obtained when taking advantage of instruction level parallelism and non-branching conditional assignment instructions, both of which are common in modern CPU architectures. We provide empirical evidence that using code synthesized from efficient sorting networks as the base case for Quicksort libraries results in significant real-world speed-ups.

## 1 Introduction

General-purpose sorting algorithms are based on comparing, and possibly exchanging, pairs of inputs. If the order of these comparisons is predetermined by the number of inputs to sort and does not depend on their concrete values, then the algorithm is said to be data-oblivious. Such algorithms are well suited for e.g. parallel sorting or secure multi-party computations.

Sorting functions in state-of-the-art programming language libraries (such as e.g. the GNU C Library) are typically based on a variant of Quicksort, where the base cases of the recursion apply insertion sort. In other words, once the subsequence to sort considered by Quicksort falls under a certain limit  $M$ , this subsequence is sorted using insertion sort. The reasons for using such base cases is that, both theoretically and empirically, insertion sort is faster than Quicksort for sorting up to  $M$  elements when  $M$  is sufficiently small. Typical values of  $M$  for resorting to a base case are 4 or 8 (e.g. the GNU C library uses 4).

Generalizing this construction, we can take any sorting algorithm based on the divide-and-conquer approach (e.g. Quicksort, merge sort) and use another sorting method once the number of elements to sort in one partition does not

---

<sup>\*</sup> Supported by the Israel Science Foundation, grant 182/13 and by the Danish Council for Independent Research, Natural Sciences.

exceed a pre-defined limit  $M$ . The guiding idea here is that, by supplying optimized code for sorting up to  $M$  inputs, the overall performance of the sorting algorithm can be improved.

One obvious way to supply optimized code for sorting up to  $M$  inputs is to provide a unique optimized implementation of sorting  $m$  elements for each  $m \leq M$ . This approach leads directly to the following problem: *For a given fixed number  $M$ , how can we obtain an efficient way to sort  $M$  elements on a modern CPU?* Similar questions have been asked since the 50's, though obviously with a different notion of what constitutes a modern CPU. Sorting networks are a classical model of comparison-based sorting that provide a framework for addressing such questions. In a sorting network,  $n$  inputs are fed into networks of  $n$  channels, which are connected pairwise by comparators. Each comparator takes the two inputs from its two channels, compares them, and outputs them sorted back to the same two channels. Consecutive comparators can be viewed as a "parallel layer" if no two touch the same channel.

Sorting networks are data-oblivious algorithms, since the sequence of comparisons performed is independent of the actual input. For this reason, they are typically viewed as hardware-oriented algorithms where data-obliviousness is a requirement and a fixed number of inputs is given.

In recent years there has been significant progress in finding sorting networks that are optimal for sorting particular numbers of inputs. There are two main notions of optimality in common use: *size* optimality, where we want to minimize the number of comparators used in the network; and *depth* optimality, where we want to minimize the number of execution steps, taking into account that some comparators can be executed in parallel. New results for size optimality are reported in [2] and for depth optimality in [1, 5].

In this paper we examine how the theory of sorting networks can improve the performance of general-purpose software sorting algorithms. We show that replacing the insertion sort base case of a Quicksort implementation as found in standard C libraries by optimized code synthesized from logical descriptions of sorting networks leads to significant improvements in execution times.

The idea of using sorting networks to guide the synthesis of optimized code for base cases of sorting algorithms may seem rather obvious, and, indeed, has been the subject of related work. A straightforward attempt [10] of utilizing sorting networks for base cases has not resulted in significant improvements, though.

In this paper we show that this is not unexpected, providing theoretical and empirical insight into the reasons for these rather discouraging results. In a nutshell, we provide an average case analysis of the complexity w.r.t. measures such as number of comparisons and number of swaps. From the complexity point of view, code synthesized from sorting networks can be expected to perform slightly worse than unrolled insertion sort. Fortunately, for small numbers (asymptotic) complexity arguments are not always a good predictor of real-world performance.

Instead, the approach in [7] matches the advantages of sorting networks with the vectorization instruction sets available in some modern CPU architectures. They do obtain significant speedups by implementing parallel comparators as

vector operations, but require a complex heuristic algorithm to generate sequences of bit shuffling code that needs to be executed between comparators.

Our approach combines the best of both previous attempts by providing a straightforward implementation of sorting networks that still takes advantage of the features of modern CPU architectures. We obtain speedups comparable to [7], while our requirements to the instruction set are satisfied by virtually all modern CPUs including those without vector operations. The success of our approach is based on two observations about sorting networks: **(1)** Sorting networks are data-oblivious and the order of comparisons is fully determined at compile time, i.e., they are free of any control-flow branching. Comparators can also be implemented without branching, and on modern CPU architectures even efficiently so. **(2)** Sorting networks are inherently parallel, i.e., comparators at the same level can be performed in parallel. Conveniently, this maps directly to implicit *instruction level parallelism* (ILP) common in modern CPU architectures. This feature allows to execute several instructions in parallel on a single thread of a single core as long as they are working on disjoint sets of registers.

Avoiding branching and exploiting ILP are tasks also performed through program transformations by the optimization stages of modern C compilers, e.g., by unrolling loops and reordering instructions to minimize data-dependence between neighbouring instructions. They are though both restricted by the data-dependencies of the algorithms being compiled and, consequently, of only limited use for data-dependent sorting algorithms like insertion sort.


Throughout the paper, for empirical evaluations we run all code on an Intel Core i7 measuring runtime in CPU cycles using the time stamp counter register using the RDTSC instruction. As a compiler for all benchmarks we used LLVM 6.1.0 with clang-602.0.49 as frontend on Max OS X 10.10.2. We also tried GCC 4.8.2 on Ubuntu with Linux kernel 3.13.0-36, yielding comparable results.

The rest of the paper is organized as follows. In the following section, we provide necessary background information and formal definitions for both sorting algorithms and hardware features. In Section 3 we theoretically compare Quicksort and the best known sorting networks w.r.t. numbers of comparisons and swaps. We aggressively unroll insertion sort until we obtain a sorting network in Section 4. Then we show how to efficiently implement the comparators that constitute sorting networks in Section 5. We empirically evaluate our contribution as a base case of Quicksort in Section 6 before concluding and giving an outlook on future work in Section 7.

## 2 Background

### 2.1 Quicksort with Insertion Sort for Base Case

For decades Quicksort has been used in practice due to its efficiency in the average case. Since its first publication by Hoare [8], many modifications were suggested to improve it further. Examples are the clever choice of the pivot or the use of a different sorting algorithm like, e.g., insertion sort, for small subproblem sizes. Most such suggestions have in common that the empirically observed efficiency can be explained on theoretical grounds by analyzing the expected number of comparisons, swaps, and partitioning stages (see [13] for details).



**Fig. 1.** Comparison of different sorting algorithms for small numbers of inputs


Figure 1 clarifies the motivation behind using insertion sort for the base case of Quicksort. The figure compares the common spectrum of data-dependent sorting algorithms for small numbers of inputs. The figure plots the number of inputs to sort on the  $x$ -axis with respect to the number of cycles required to sort on the  $y$ -axis (averaged over 100 million random executions). The upper curve in the figure is derived from the standard Quicksort implementation in the C library (which is at some disadvantage as it requires a general compare function as an argument). The rest of the curves derive from applying standard sorting algorithms as detailed by Sedgewick [14] (the code was taken directly from the web page <http://algs4.cs.princeton.edu/home/> affiliated with the book). Insertion sort is the clear winner.

## 2.2 Sorting Networks

A *comparator network* on  $n$  channels is a finite sequence  $C = c_1, \dots, c_k$  of *comparators*, where each comparator  $c_\ell$  is a pair  $(i_\ell, j_\ell)$  with  $1 \leq i_\ell < j_\ell \leq n$ . The *size* of  $C$  is the number  $k$  of comparators it contains. Let  $D$  be any totally ordered domain. Given an input  $\mathbf{x} \in D^n$ , the *output* of  $C$  on  $\mathbf{x}$  is the sequence  $C(\mathbf{x}) = \mathbf{x}^n$ , where  $\mathbf{x}^\ell$  is defined inductively as follows:  $\mathbf{x}^0 = \mathbf{x}$ , and  $\mathbf{x}^\ell$  is obtained from  $\mathbf{x}^{\ell-1}$  by swapping the elements in positions  $i_\ell$  and  $j_\ell$ , in case  $x_{i_\ell} < x_{j_\ell}$ . The set of outputs of  $C$  is  $\text{outputs}(C) = \{C(\mathbf{x}) \mid \mathbf{x} \in D^n\}$ , and  $C$  is a *sorting network* if this set only contains sorted sequences.

Comparators may act in parallel. A comparator network  $C$  has *depth*  $d$  if  $C$  is the concatenation of  $L_1, \dots, L_d$ , where each  $L_i$  is a *layer*: a comparator network with the property that no two of its comparators act on a common channel.

Figure 2 depicts a sorting network on 5 channels in the graphical notation we will use throughout this paper. Comparators are depicted as vertical lines, and layers are separated by a dashed line. The numbers illustrate how the input  $10101 \in \{0, 1\}^5$  propagates through the network. This network has 6 layers and optimal size (9 comparators).



**Fig. 2.** A sorting network on 5 channels operating on the input 10101.

Given  $n$  inputs, finding the minimal size  $s_n$  and depth  $t_n$  of a sorting network is an extremely hard problem. The table below details the best currently known bounds. The values for  $n \leq 8$  are already listed in [9]; the values of  $t_9$  and  $t_{10}$  were proven exact by Parberry [11], those of  $t_{11}$ – $t_{16}$  by Bundala and Závodný [1], and  $t_{17}$  was recently computed by Ehlers and Müller [5] using results from [3, 4]. Regarding optimal size, there was no progress for fifty years, until the computer-assisted proof that the known upper bounds for  $s_9$  and  $s_{10}$  were tight [2].

$n$	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
$s_n$	0	1	3	5	9	12	16	19	25	29	35	39	45	51	56	60	73
$t_n$	0	1	3	3	5	5	6	6	7	7	8	8	9	9	9	9	10


Oblivious versions of classic sorting algorithms can also be implemented as sorting networks, as described in [9]. Figure 3 (a) shows an oblivious version of insertion-sort. Here vertical dashed lines highlight the 4 iterations of “insertion” required to sort 5 elements. Figure 3 (b) shows the same network where comparators are arranged according to their parallel layers. Bubble-sort can also be implemented as a sorting network as illustrated in Figure 3 (c) where vertical dashed lines illustrate the 4 iterations of the classic bubble-sort algorithm. Ordered according to layers, this network is also identical to the one in Figure 3 (b).

### 2.3 Modern CPU Architectures

Modern CPU architectures allow multiple instructions to be performed in parallel on a single thread. This phenomenon is called instruction-level parallelism (ILP) and is built on three modern micro-architectural techniques<sup>3</sup>:

- superscalar instruction pipelines, i.e., pipelines with the ability to hold and execute multiple instructions at the same time

<sup>3</sup> For details on these features of modern microarchitectures see e.g. [15, 6].



**Fig. 3.** Insertion sort on 5 inputs viewed as a sorting network, with dashed lines separating iterations (a), and the corresponding parallelized version (b). Bubble-sort (c) also reduces to the same sorting network when parallelized.

- dynamic out-of-order execution, i.e., dynamic reordering of instructions respecting data dependencies
- redundant execution units, i.e., multiple ALUs per core

Together, these features allow to execute instructions in an order such that data dependencies are minimized and multiple redundant execution units can be used at the same time. This is often termed *implicit* ILP, in contrast to the explicit ILP found in vector operations.

**Example.** Consider the C expression  $(x+y)*(z+u)$ . Assume the variables  $x$ ,  $y$ ,  $z$ , and  $u$  are loaded in registers  $eax$ ,  $ebx$ ,  $ecx$ , and  $edx$ . Then the evaluation of the above expression is compiled to three machine instructions: `ADD eax, ebx`; `ADD ecx, edx`; `MUL eax, ecx` with the result being in  $ecx$ . Here, the first two instructions are data-independent and can be executed in parallel, while the first one depends on the results of those and is executed in another CPU cycle.

Conditional branching instructions are the most expensive instructions on pipelined CPUs as they require flushing and refilling the pipeline. In order to minimize their cost, modern CPU architectures employ dynamic branch prediction. By keeping the pipeline filled with the instructions of the predicted branch, the cost of branching is severely alleviated. Unfortunately, branch prediction cannot be perfect, and when the wrong branch is predicted, the pipeline needs to be flushed and refilled – an operation taking many CPU cycles.

In order to avoid branching instructions for “small” decisions, e.g., for deciding whether to assign a value or not, modern CPU architectures feature conditional instructions. Depending on the flags set by e.g. a comparison, an assignment of a value of a register will either be performed or the instruction will be ignored. In both cases, the pipeline is filled with the following instructions and the cost of the operation is smaller than a possible branch prediction failure.

**Example.** Consider the C statement `if (x == 42) x = 23`; and assume that the variable  $x$  is loaded in  $eax$ . Without conditional move instructions, this is compiled to code with a conditional branching instruction, e.g. `CMP eax, 42`; `JNZ after`; `MOV eax, 23` where `after` is the address of the instruction following the `MOV` instruction. Alternatively, using conditional instruction we obtain `CMP eax, 42`; `CMOVZ eax, 23`. This code saves one machine code instruction, but most importantly it avoids the huge performance impact of a mispredicted branch.

### 3 Quicksort with Sorting Networks for Base Case

A general theme in this paper is to derive, from sorting networks, optimized code to sort small numbers of inputs and then to apply this code as the base case in a Quicksort algorithm. In this section we compare precise average case results for the number of comparisons and swaps performed by a classic Quicksort algorithm and by a modification that uses sorting networks on subproblems of size at most 14. We choose 14 for this analysis, as it is the largest value  $n$  for which we conveniently could measure the number of comparisons and swaps for all  $n!$  permutations. We used the smallest (w.r.t. size) sorting networks known (for up to 10 they are optimal) in order to obtain the most favorable comparison numbers for sorting networks. To this end we assume the algorithm to act on random permutations of size  $n$  each being the input with equal probability.

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14
comparisons	0	1	3	5	9	12	16	19	25	29	35	39	45	51
swaps	0.0	0.5	1.5	2.7	4.8	6.6	8.6	10.6	13.0	11.1	19.4	22.4	20.0	26.5

**Table 1.** The average number of comparisons and swaps when executing optimal sorting networks with at most  $M = 14$  inputs.

Let  $C_n$  (resp.  $S_n$ ) denote the expected number of comparisons (resp. swaps) performed by classic Quicksort on (random) inputs of size  $n$ . Let furthermore  $\hat{C}_n$  and  $\hat{S}_n$  denote the corresponding quantities for Quicksort using sorting networks for inputs smaller than 15. It is standard to set up recurrence relations for those quantities which typically obey a pattern such as:

$$T_n(a, b) = \begin{cases} a \cdot n + b + \frac{1}{n} \sum_{1 \leq j \leq n} T_{j-1}(a, b) + T_{n-j}(a, b) & \text{if } n > M, \\ g(n) & \text{otherwise.} \end{cases}$$

Here  $a$  and  $b$  have to be chosen properly to reflect the parameter's (comparisons, swaps) behavior,  $M$  determines the maximum subproblem size for which a different algorithm (insertion sort, sorting networks) is used and  $g$  accounts for the costs of that algorithm. In order to analyze classic Quicksort as proposed by Hoare, we have to choose  $a = 1$ ,  $b = -1$  (resp.  $a = \frac{1}{6}$ ,  $b = \frac{2}{3}$ ) for comparisons (resp. swaps), together with  $M = 0$  and  $g(0) = 0$ . For the analysis of our proposed modification that uses sorting networks for subproblems of small sizes we set  $M = 14$  together with the values for  $g$  as given in Table 1. Using standard algebraic manipulations it is possible to solve this recurrence explicitly to obtain a formula for  $T_n(a, b)$  in terms of  $n$ ,  $M$ ,  $a$  and  $b$ . Defining  $t_n = a \cdot n + b$  and  $\nabla t_n = t_n - t_{n-1}$  one finds (see [12] for details) that:

$$T_n(a, b) = 2(n+1) \sum_{M+2 \leq k \leq n} \frac{\nabla t_k}{k+1} + \frac{n+1}{M+2} (t_{M+1} + T_{M+1}(a, b)) - t_n \text{ for } n > M.$$

Computing closed form expressions for  $\sum_{M+2 \leq k \leq n} \frac{\nabla t_k}{k+1}$  for the different choices of  $t_n$ , we finally get

$$\begin{aligned} C_n &= 2n \ln(n+1) - 2.84557n + o(n) \\ \hat{C}_n &= 2n \ln(n+1) - 2.44869n + o(n) \\ S_n &= \frac{1}{3}n \ln(n+1) + 0.359072n + o(n) \\ \hat{S}_n &= \frac{1}{3}n \ln(n+1) + 0.524887n + o(n) \end{aligned}$$

It is obvious that when increasing  $n$ , both parameters get worse by our modification of classic Quicksort. Even for small  $n$  and optimal size sorting networks, there is no advantage w.r.t. the numbers of comparisons or swaps. In conclusion, we cannot hope to get a faster sorting algorithm by switching to sorting networks for small subproblems – at least not on grounds of our theoretical investigations. And by transitivity, replacing insertion sort by sorting networks in the base case should result in an even worse behavior w.r.t. both parameters.

## 4 Unrolling the Base Case

In this section we show how to unroll an implementation of insertion sort, step by step, until finally obtaining code which is equivalent to a sorting network. We take the basic insertion sort code from Sedgewick [14] and for illustration assume the fixed number of inputs is  $n = 5$ . We experimented also with optimized variants, e.g., making use of sentinels to avoid the  $j > 0$  check, but did not find any of them to be faster for small inputs given a modern C compiler.

```
#define SWAP(x,y) {int tmp = a[x]; a[x] = a[y]; a[y] = tmp;}
static inline void sort5(int *a, int n) {
    n=5
    for (int i = 1; i < n; i++)
        for (int j = i; j > 0 && a[j] < a[j-1]; j--)
            SWAP(j-1, j)
}
```

Applying partial evaluation and (outer) loop unrolling results in:

```
static inline void sort5_unrolled(int *a) {
    for (int j = 1; j > 0 && a[j] < a[j-1]; j--)
        SWAP(j-1, j)
    for (int j = 2; j > 0 && a[j] < a[j-1]; j--)
        SWAP(j-1, j)
    for (int j = 3; j > 0 && a[j] < a[j-1]; j--)
        SWAP(j-1, j)
    for (int j = 4; j > 0 && a[j] < a[j-1]; j--)
        SWAP(j-1, j)
}
```

The condition in the inner loop is data-dependent, hence no sound and complete program transformation can be applied to unroll them. To address this, we move the data-dependent part of the loop condition to the statement in the body of the loop while always iterating the variable  $j$  down to 1.

```
static inline void sort5_oblivious(int *a) {
    for (int j = 1; j > 0; j--)
        if (a[j] < a[j-1]) SWAP(j-1, j)
    for (int j = 2; j > 0; j--)
        if (a[j] < a[j-1]) SWAP(j-1, j)
    for (int j = 3; j > 0; j--)
        if (a[j] < a[j-1]) SWAP(j-1, j)
    for (int j = 4; j > 0; j--)
        if (a[j] < a[j-1]) SWAP(j-1, j)
}
```


Now we can now apply (inner) loop unrolling and obtain:

```
static inline void sort5_oblivious_unrolled(int *a) {
    if (a[1] < a[0]) SWAP(0, 1)
    if (a[2] < a[1]) SWAP(1, 2)
    if (a[1] < a[0]) SWAP(0, 1)
    if (a[3] < a[2]) SWAP(2, 3)
    if (a[2] < a[1]) SWAP(1, 2)
    if (a[1] < a[0]) SWAP(0, 1)
    if (a[4] < a[3]) SWAP(3, 4)
    if (a[3] < a[2]) SWAP(2, 3)
    if (a[2] < a[1]) SWAP(1, 2)
    if (a[1] < a[0]) SWAP(0, 1)
}
```

All the statements in the body of `sort5_oblivious_unrolled` are now conditional swaps. For readability, we move the condition into the macro. **COMPs** on the same line indicate that they originate from the same iteration of insertion sort:

```
#define COMP(x,y) { if (a[y] < a[x]) SWAP(x,y) }
static inline void sort5_fig3a(int *a) {
```





**Fig. 4.** Comparison of insertion sort with (unrolled) comparator based code for small numbers of inputs

```

COMP(0, 1)
COMP(1, 2) COMP(0, 1)
COMP(2, 3) COMP(1, 2) COMP(0, 1)
COMP(3, 4) COMP(2, 3) COMP(1, 2) COMP(0, 1)
}

```

This sequence is equivalent to the sorting network in Figure 3 (a). Thus, we can apply the reordering of comparators that resulted in Figure 3 (b) to obtain the following implementation, where we reduce the number of layers from 10 to 7 (here COMPs on the same line indicate a layer in the sorting network):

```

#define COMP(x,y) { if (a[y] < a[x]) SWAP(x,y) }
static inline void sort5_fig3b(int *a) {
    COMP(0, 1)
    COMP(1, 2)
    COMP(0, 1) COMP(2, 3)
    COMP(1, 2) COMP(3, 4)
    COMP(0, 1) COMP(2, 3)
    COMP(1, 2)
    COMP(0, 1)
}

```

Figure 4 presents a comparison of a standard insertion sort (code from [14]) with the several optimized versions. The figure plots the number of inputs to sort on the  $x$ -axis with respect to the number of cycles required to sort on the  $y$ -axis (averaged over 100 million random executions). The curve labeled “insertion sort” portrays the same data as the corresponding curve in Figure 1. The curve labeled “unrolled insertion sort” corresponds to the unrolled version of insertion sort (in the style of function `sort5_unrolled`). The other three curves correspond to code derived from different types of sorting networks: “insertion sorting network” (in the style of Figure 3 (a) and function `sort5_fig3a`), “compressed insertion sorting network” (in the style of Figure 3 (b) and function `sort5_fig3b`), and “optimal sorting network” corresponding to the use of a best (smallest) known sorting network.

The figure clarifies that standard sorting network optimizations such as reordering of independent comparators [9] gives a slight performance boost. But

there is another clear take home messages: even going beyond standard program transformations by breaking data-dependence and obtaining a sequence of conditional swaps (i.e., a sorting network), we do not manage to make any significant improvements of the performance of sorting implementations for small numbers of inputs. Furthermore, even when using size optimal sorting networks, we see no real benefit over compiler-optimized insertion sort. This is inline with the theoretical results on average case complexity obtained in the previous section.

## 5 Implementing Sorting Networks Efficiently

The results in the previous two sections explained the rather discouraging results obtained when attempting to use sorting networks as the base case of a divide-and-conquer sorting algorithm: they are simply not faster than e.g. using insertion sort – at least when implemented naively. In this section we will show how we can exploit two main properties of sorting networks together with features of modern CPU architectures in order to obtain speed-ups of more than a factor 3 compared to unrolled insertion sort.

First, note that sorting networks are data-oblivious and the order of comparisons is fully determined at compile time, i.e., the implementation is free of any control-flow branching. Unfortunately, the naive implementation of each comparator involves branching when deciding whether to perform a swap. Which path is taken depends entirely on the specific inputs to be sorted, such that branch prediction necessarily does not perform very well.

Luckily, we can also implement a comparator without branching. To this end we use a conditional assignment (defined by the macro `COND` below), which can be compiled to a conditional move (`CMOV`) instruction available on modern CPU architectures. This approach proved to be very fruitful. For illustration, from the optimal size sorting network for 5 inputs, portrayed as Figure 2, we synthesize the following C function `sort5_best` where each row in the code corresponds to a layer in the sorting network:

```
#define COND(c,x,y) { x = (c) ? y : x; }
#define COMP(x,y) { int ax = a[x]; COND(a[y] < ax,a[x],a[y]); \
                  COND(a[y] < ax,a[y],ax ); }

static inline void sort5_best(int *a) {
    COMP(0, 1) COMP(3, 4)
    COMP(2, 4)
    COMP(2, 3) COMP(1, 4)
    COMP(0, 3)
    COMP(0, 2) COMP(1, 3)
    COMP(1, 2)
}
```

The comparator macro that compares and conditionally swaps the values at indices `x` and `y` works as follows:

1. Keep a copy of the value at index `x`.
2. Compare (once) the value at index `y` with the stored value from `x`.
3. If the value was greater, we copy the value at index `y` to index `x`. Otherwise, we do nothing in this step.
4. If the value was greater, we write the old copied value from `x` to index `y`. Otherwise, we do nothing in this step.

Correctness follows directly by case analysis: If the value at index `y` was not greater than the value at index `x`, the two conditional assignments do not change anything and all we did was an unnecessary copy of the value at index `x`. If the value at index `y` was greater than the value at index `x`, we essentially perform a classic swap using `ax` as the temporary variable.

Given a sufficient optimization level (`-O2` and above), the above code is compiled by the LLVM (or GNU) C compiler to use two conditional move (`CMOV`) instructions, resulting in a totally branching free code for `sort5_best`. As can be expected, the other two instructions are a move (`MOV`) and a compare (`CMP`) instruction. In other words, each comparator is implemented by exactly four non-branching machine code instructions.

Alternatively, we could implement the comparator applying the folklore idea of swapping values using XORs to eliminate one conditional assignment<sup>4</sup>:

```
#define COND(c,x,y) { x = (c) ? y : x; }
#define COMP(x,y) { int ax = a[x]; COND(a[y] < ax, a[x], a[y]); \
                  a[y] ^= ax ^ a[x]; }
```

This alternative comparator performs a conditional swap as follows:

1. Keep a copy of the value at index `x`.
2. If the value at index `y` is greater than the value at index `x`, copy the value at index `y` to index `x`.
3. Bitwise XOR the value at index `y` with the copied old and the new value at index `x`.

Step 3 works because if the condition holds, then `ax` and the value at index `x` cancel out leaving the value at `y` unchanged, while the value at `y` and `ax` cancel out otherwise, effectively assigning the original value from index `x` to index `y`.

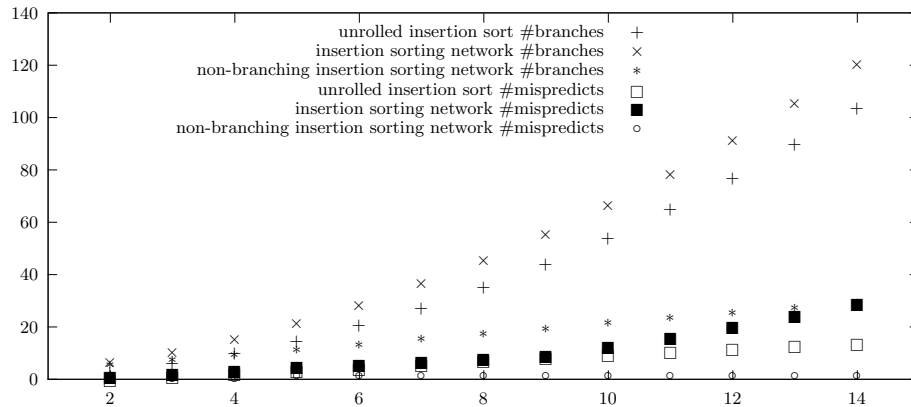
We also implemented this variant and observed that it compiles down to five instructions (`MOV`, `CMP`, `CMOV`, and two `XORs`). We benchmarked the two variants and observed that they are indistinguishable in practice with differences well within the margin of measurement error. Thus, we decided to continue with this version, as the `XOR` instructions are more “basic” and can be expected to behave better w.r.t. e.g. instruction level parallelism.

A third approach would be to define branching free minimum and maximum operations<sup>5</sup> and use them to assign the minimum to the upper channel and the maximum to the lower channel of the comparator. We tested this approach, but found that it did not compile to branching-free code. Even if it was, the number of instructions involved would be rather large, alleviating any chance of competing with the two previous variants.


The reader might wonder as to whether using a different `SWAP` macro could similarly speed up the working of standard insertion sort. The answer here is a clear no, as the standard swapping operation is implemented by only three operations. Tricks like using `XORs` only increase the number of instructions to be executed while not reducing the number of branching in the code. We, of

<sup>4</sup> See <https://graphics.stanford.edu/~seander/bithacks.html#SwappingValuesXOR>

<sup>5</sup> See <https://graphics.stanford.edu/~seander/bithacks.html#IntegerMinOrMax>



**Fig. 5.** Comparing the number of branches, encountered and mispredicted, in optimized sorting algorithms for small numbers of inputs

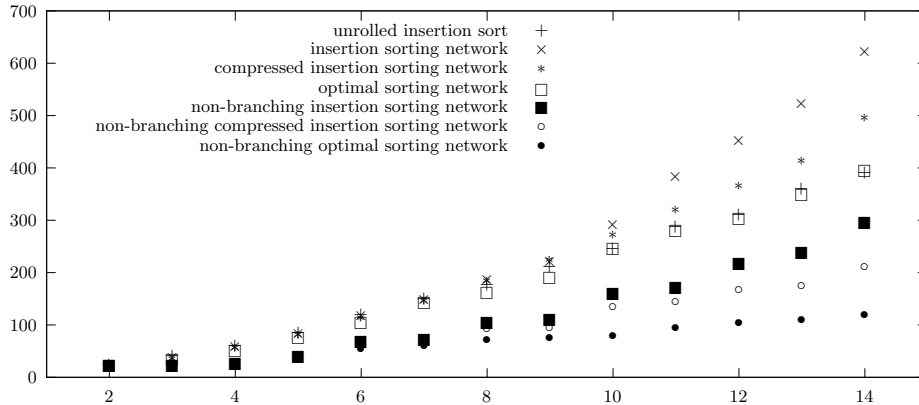


**Fig. 6.** ILP on comparator networks of length 1000 with differing levels of parallelism

course, implemented and benchmarked several alternative **SWAP** macros, finding only detrimental effects on measured performance.

Figure 5 compares three sorting algorithms for small numbers of inputs: (1) the unrolled insertion sort (also plotted in Figure 4); (2) code derived from a standard insertion sorting network (also plotted in Figure 4); (3) the same insertion sorting network but with a non-branching version of the **COMP** macro. We compare the number of branches encountered and mispredicted (averaged over 100 million random executions). The figure illustrates that the number of branches encountered (and mispredicted) is larger for both unrolled insertion sort and a naive implementation of sorting networks. In contrast, for the branching free implementation we see a nearly constant level of branches encountered and mispredicted. In fact these branches originate from the surrounding test code (filling an array with random numbers, computing random numbers, checking that the result is actually sorted).

Our second observation is that sorting networks are inherently parallel, i.e., comparators at the same level can be performed in parallel. In the remainder of this section we show that this parallelism can be mapped directly to instruction level parallelism (ILP) commonly encountered in modern CPU architectures.



**Fig. 7.** Comparing sorting networks for small numbers of inputs: non branching sorting networks are fastest

The ability to make use of ILP, i.e., to execute several instructions in parallel on a single thread of a single core, has further performance potential.


In order to demonstrate this potential, we constructed artificial test cases with varying levels of data dependency. Given a natural number  $m$  defining the length of subsequences of parallel comparators in a comparator network of size 1000. We would expect that as  $m$  grows, we would see more use of ILP.

In Figure 6 the values for  $m$  are represented on the  $x$ -axis while the  $y$ -axis (as usual) is the averaged number of CPU cycles. And indeed, we see significant performance gains when going from  $m = 1$  to  $m = 2$  and  $m = 3$ . After this, the performance stays unchanged. This has to be understood on the background of each comparator being compiled to 5 assembler instructions when using optimization level `-O3`. Then we obtain slightly below 2 CPU cycles per comparator.

Combining the gains from ILP with the absence of branching, we obtain large speed-ups for small inputs when comparing to both insertion sort and naive implementations of sorting networks. In Figure 7 we show the magnitude of the improvements obtained. Once again we plot the number of inputs to sort on the  $x$ -axis with respect to the number of cycles required to sort on the  $y$ -axis (averaged over 100 million random executions). We consider the unrolled insertion sort, the three sorting networks from Figure 4 (insertion sorting network, compressed insertion sorting network, and optimal sorting network), and we consider these same three sorting networks using non-branching comparators (non-branching insertion sorting network, non-branching compressed insertion sorting network, and non-branching optimal sorting network). The figure illustrates that using the best known (optimal) sorting networks, in their non-branching forms, results in a more than 3 times speed-up.

## 6 Quicksort with Sorting Network Base Case

We now demonstrate that optimizing the code in the base case of a Quicksort algorithm translates to real-world savings when applying the sorting function.



**Fig. 8.** Quicksort: comparing insertion sort at the base case with non-branching optimal sorting networks at the base case. Plotting base case size (x axis) and number of cycles (averaged over one million random runs)

To this end, we used as base cases (1) the empirically best variant of insertion sort unrolled by applying program transformations to the algorithm from Sedgewick; and (2) the fastest non-branching code derived from optimal (size) sorting networks.

In Figure 8 we describe the experiment when sorting lists of 10,000 elements where the y-axis illustrates the number of cycles (averaged over one million random runs). The x-axis details the limit at which Quicksort resorts to a base case. For example, the value 8 indicates that we use a base case whenever the algorithm is required to sort a sequence of length at most 8 elements. The value 2 corresponds to the case where the base case has no impact. To quantifying the impact of the choice of base case we compare to the case for value 2 (on the x-axis). For insertion sort we see a 2 – 12% reduction in runtime depending on the limit, and for non-branching sorting networks we achieve instead 7 – 23% reduction in runtime.

## 7 Conclusion

In this paper we showed both theoretically and empirically that using code derived naively from sorting networks when sorting small numbers of inputs is not advantageous compared to the use of standard data-dependent sorting algorithms like e.g. insertion sort. Furthermore, we showed that program transformations are of only limited utility for improving insertion sort on small numbers of inputs.

We show how to synthesize simple yet efficient implementations of sorting networks and gave insight into the microarchitectural features that enable this implementation. We demonstrated the significant speed ups obtained and provided evidence that they translate to the use of efficient sorting networks as a base case in divide-and-conquer sorting algorithms such as e.g. Quicksort.

Regarding future work, we observe that using different types of sorting networks has a measurable impact on the efficiency of the synthesized C code. While previous research on finding optimal sorting networks has focused on optimal depth or optimal size, we plan to identify the criteria that will lead to optimality in our setting. What are the criteria that determine the real-world efficiency of the synthesized code, and how do we find the optimal sorting networks to meet this criteria? We also plan to make an empirical comparison to the vector operation-based approach of [7] and to benchmark our approach as base case for other sorting algorithms such as merge sort.

## References

1. Daniel Bundala and Jakub Závodný. Optimal sorting networks. In *LATA 2014*, volume 8370 of *LNCS*, pages 236–247. Springer, 2014.
2. Michael Codish, Luís Cruz-Filipe, Michael Frank, and Peter Schneider-Kamp. Twenty-five comparators is optimal when sorting nine inputs (and twenty-nine for ten). In *ICTAI 2014*, pages 186–193. IEEE, December 2014.
3. Michael Codish, Luís Cruz-Filipe, and Peter Schneider-Kamp. The quest for optimal sorting networks: Efficient generation of two-layer prefixes. In F. Winkler, V. Negru, T. Ida, T. Jebelan, D. Petcu, S.M. Watt, and D. Zaharie, editors, *SYNASC 2014*, pages 359–366. IEEE, 2015.
4. Michael Codish, Luís Cruz-Filipe, and Peter Schneider-Kamp. Sorting networks: the end game. In A.-H. Dediu, E. Formenti, C. Martín-Vide, and B. Truthe, editors, *LATA 2015*, volume 8977 of *LNCS*, pages 664–675. Springer, 2015.
5. Thorsten Ehlers and Mike Müller. New bounds on optimal sorting networks. *CoRR*, abs/1501.06946, 2015.
6. Joseph A. Fisher, Paolo Faraboschi, and Cliff Young. *Embedded Computing: A VLIW Approach to Architecture, Compilers, and Tools*. Morgan Kaufman, 2005.
7. Timothy Furtak, José Nelson Amaral, and Robert Niewiadomski. Using simd registers and instructions to enable instruction-level parallelism in sorting algorithms. In *SPAA '07*, pages 348–357. ACM, 2007.
8. C.A.R. Hoare. Quicksort. *Comput. J.*, 5(1):10–15, 1962.
9. Donald E. Knuth. *The Art of Computer Programming, Volume III: Sorting and Searching*. Addison-Wesley, 1973.
10. Blanca Lopez and Nareli Cruz-Cortes. On the usage of sorting networks to big data. In Hamid R. Arabnia, Mary Q. Yang, George Jandieri, James J. (Jong Hyuk) Park, Ashu M. G. Solo, and Fernando G. Tinetti (Editors), editors, *Advances in Big Data Analytics: The 2014 WorldComp International Conference Proceedings*. Mercury Learning and Information, 2014.
11. Ian Parberry. A computer-assisted optimal depth lower bound for nine-input sorting networks. *Mathematical Systems Theory*, 24(2):101–116, 1991.
12. Robert Sedgewick. The analysis of quicksort programs. *Acta Inf.*, 7:327–355, 1977.
13. Robert Sedgewick and Philippe Flajolet. *An introduction to the analysis of algorithms*. Addison-Wesley-Longman, 1996.
14. Robert Sedgewick and Kevin Wayne. *Algorithms, 4th Edition*. Addison-Wesley, 2011.
15. Jurij Silc, Borut Robic, and Theo Ungerer. *Processor Architecture: From Dataflow to Superscalar and Beyond*. Springer, 1999.