Nearly-Optimal Mergesorts: Fast, Practical Sorting Methods That Optimally Adapt to Existing Runs

J. Ian Munro
University of Waterloo, Canada
imunro@uwaterloo.ca
https://orcid.org/0000-0002-7165-7988

Sebastian Wild
University of Waterloo, Canada
wild@uwaterloo.ca
https://orcid.org/0000-0002-6061-9177

Abstract

We present two stable mergesort variants, "peeksort" and "powersort", that exploit existing runs and find nearly-optimal merging orders with negligible overhead. Previous methods either require substantial effort for determining the merging order (Takaoka 2009; Barbay & Navarro 2013) or do not have an optimal worst-case guarantee (Peters 2002; Auger, Nicaud & Pivoteau 2015; Buss & Knop 2018). We demonstrate that our methods are competitive in terms of running time with state-of-the-art implementations of stable sorting methods.

2012 ACM Subject Classification Theory of computation → Sorting and searching

Keywords and phrases adaptive sorting, nearly-optimal binary search trees, Timsort

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.63

Related Version `arXiv: 1805.04154` (extended version with appendices)

Supplement Material `zenodo: 1241162` (code to reproduce running time study)

Funding This work was supported by the Natural Sciences and Engineering Research Council of Canada and the Canada Research Chairs Programme.

1 Introduction

Sorting is a fundamental building block for numerous tasks and ubiquitous in both the theory and practice of computing. While practical and theoretically (close-to) optimal comparison-based sorting methods are known, *instance-optimal sorting,* i.e., methods that *adapt* to the actual input and exploit specific structural properties if present, is still an area of active research. We survey some recent developments in Section 1.1.

Many different structural properties have been investigated in theory. Two of them have also found wide adoption in practice, e.g., in Oracle's Java runtime library: adapting to the presence of duplicate keys and using existing sorted segments, called *runs*. The former is achieved by a so-called fat-pivot partitioning variant of quicksort [8], which is also used in the OpenBSD implementation of `qsort` from the C standard library. It is an *unstable* sorting method, though, i.e., the relative order of elements with equal keys might be destroyed in the process. It is hence used in Java solely for primitive-type arrays.

© J. Ian Munro and Sebastian Wild;
licensed under Creative Commons License CC-BY

Making use of existing runs in the input is a well-known option in mergesort; e.g., Knuth [17] discusses a bottom-up mergesort variant that does this. He calls it "natural mergesort" and we will use this as an umbrella term for any mergesort variant that picks up existing runs in the input (instead of blindly starting with runs of size 1). The Java library uses *Timsort* [25, 15], a natural mergesort originally developed as Python's new library sort.

While fat-pivot quicksort provably adapts to the *entropy of the multiplicities* of keys [34] – it is optimal up to a factor of 1.088 on average with pseudomedian-of-9 ("ninther") pivots[1] – Timsort is much more heuristic in nature. It picks up existing runs and tries to perform merges in a favorable order (i.e., avoiding merges of runs with very different lengths) by distinguishing a handful of cases of how the lengths of the 4 most recent runs relate to each other. The rules are easy to implement and were empirically shown to be effective in most cases, but their interplay is quite intricate. Although announced as an $O(n \log n)$ worst-case method with its introduction in 2002 [24], a rigorous proof of this bound was only given in 2015 by Auger, Nicaud, and Pivoteau [2] and required a rather sophisticated amortization argument.

The core complication is that – unlike for standard mergesort variants – a given element might participate in more than a logarithmic number of merges. Indeed, Buss and Knop [9] have very recently shown that for some family of inputs, the average number of merges a single element participates in is at least $\left(\frac{3}{2} - o(1)\right) \cdot \lg n$. So in the worst case, Timsort does, e.g., asymptotically at least *1.5 times as many element moves as standard mergesort.* In terms of adapting to existing order, provable guarantees for Timsort had long remained elusive; an upper bound of $O(n + n \log r)$ was conjectured in [2] and [9]), but indeed only for this very conference, Auger, Jugé, Nicaud and Pivoteau [1] finally give a proof. The hidden constants in their bound are quite big (and far away from the coefficient $\frac{3}{2}$ of the above lower bound).

A further manifestation of the complexity of Timsort's rules was reported by de Gouw et al. [10]: The original rules to maintain the desired invariant for run lengths on the stack are insufficient in some cases. This (algorithmic!) bug had remained unnoticed until their attempt to formally verify the correctness of the Java implementation failed because of it. Gouw et al. proposed two options for correcting Timsort, one of which was applied for the Java library. But now, Auger et al. [1] demonstrated that *this correction is still insufficient:* as of this writing, the Java runtime library contains a flawed sorting method! All of this indicates that already the core algorithm in Timsort is (too?) complicated, and it raises the question whether Timsort's good properties cannot be achieved in a simpler way.

For its theoretical guarantees on adapting to existing runs this is certainly the case. Takaoka [29, 30] and Barbay and Navarro [5] independently described a sorting method that we call *Huffman-Merge* (see below why). It adapts even to the *entropy of the distribution of run lengths:* it sorts an input of r runs with respective lengths $L_1, \dots, L_r$ in time $O\big((\mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_r}{n}) + 1)n\big) \subseteq O(n + n \lg r)$, where $\mathcal{H}(p_1, \dots, p_r) = \sum_{i=1}^{r} p_i \lg(1/p_i)$ is the binary Shannon entropy.[2] Since $\mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_r}{n})n - O(n)$ comparisons are necessary for distinct keys, Huffman-Merge is optimal up to $O(n)$. The algorithm is also conceptually simple: find runs in a linear scan, determine an optimal merging order using a Huffman tree of

[1] The median of three elements is chosen as the pivot, each of which is a median of three other elements. This is a good approximation of the median of 9 elements and a recommended pivot selection rule [8].

[2] Note that $\mathcal{H}(L_1/n, \dots, L_r/n)$ can be significantly smaller than $\lg r$: Consider the run lengths $L_1 = n - \lceil n/\lg n \rceil$ and $L_2 = \cdots = L_r = 1$, i.e., $r = 1 + \lceil n/\lg n \rceil$. Then $\mathcal{H} \leq 2$, but $\lg r \sim \lg n$. (Indeed, $\mathcal{H} \to 1$ and the input can indeed be sorted with overall costs $2n$.)

the run lengths, and execute those merges bottom-up in the tree. However, straight-forward implementations add significant overhead in terms of time and space, which renders Huffman-Merge uncompetitive to (reasonable implementations of) elementary sorting methods.

Moreover, Huffman-Merge leads to an *unstable* sorting method since it merges non-adjacent runs. The main motivation for Timsort was to find a fast general-purpose sorting method that is *stable* [24], and the Java library even dictates the sorting method used for objects to be stable.[3] It is conceptually easy to modify the idea of Huffman-Merge to sort stably: replace the Huffman tree by an *optimal binary search tree* and otherwise proceed as before. Since we only have weights at the leaves of the tree, we can compute this tree in $O(n + r \log r)$ time using the Hu-Tucker- or Garsia-Wachs-algorithm. (Barbay and Navarro made this observation, as well; indeed they initially used the Hu-Tucker algorithm [4] and only switched to Huffman in the journal paper [5].) Since r can be $\Theta(n)$ and the algorithms are fairly sophisticated, this idea is not very appealing for use in practical sorting, though.

In this paper, we present two new stable, natural mergesort variants, "peeksort" and "powersort", that have the same optimal asymptotic running time $O\big((\mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_r}{n})+1)n\big)$ as Huffman-merge, but incur much less overhead. For that, we build upon classic algorithms for computing *nearly-optimal binary search trees* [21]; but the vital twist for practical methods is to neither explicitly store the full tree, nor the lengths of all runs at any point in time. In particular – much like Timsort – we only store a *logarithmic* number of runs at any point in time (in fact reducing their number from roughly $\log_\varphi \approx 1.44 \lg n$ in Timsort to $\lg n$), but – much *un*like Timsort – we retain the guarantee of an optimal merging order up to linear terms. Our methods require at most $n \lg n + O(n)$ comparison in the worst case and $\mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_r}{n})n + 3n$ for an input with runs of lengths $L_1, \dots, L_r$.

We demonstrate in a running-time study that our methods achieve guaranteed (leading-term) optimal adaptive sorting in practice with negligible overhead to compute the merging order: our methods are *not* slower than standard mergesort when no existing runs can be exploited, but outperform standard mergesort and quicksort when long runs are present in the input. Finally, we show that Timsort is slower than standard mergesort and our new methods on certain inputs that do have existing runs, but whose lengths pattern hits a weak point of Timsort's heuristic merging-order rule.

Outline: The rest of this paper is organized as follows. In the remainder of this section we survey related work. Section 2 contains notation and known results on optimal binary search trees that our work builds on. The new algorithms and their analytical guarantees are presented in Section 3. Section 4 reports on our running-time study, comparing the new methods to state-of-the-art sorting methods. Finally, Section 5 summarizes our findings. Details about the experimental setup and some proofs are found in the extended version of this article (`arXiv: 1805.04154`).

1.1 Adaptive Sorting

The idea to exploit existing "structure" in the input to speed up sorting dates (at least) back to methods from the 1970s [20] that sort faster when the number of inversions is small. A systematic treatment of this and many further measures of presortedness (e.g., the number of

[3] We remark that while stability is a much desired feature, practical, stable sorting methods do not try to *exploit* the presence of duplicate elements to speed up sorting, and we will focus on the performance for distinct keys in this article.

inversions, the number of runs, and the number of shuffled up-sequences), their relation and how to sort *adaptively* w.r.t. these measures are discussed by Estivill-Castro and Wood [12].

While the focus of earlier works is mostly on combinatorial properties of permutations, a recent trend is to consider more fine-grained statistical quantities. For example as mentioned above, Huffman-Merge adapts to the *entropy* of the vector of run lengths [29, 30, 5]. Similar measures are the entropy of the lengths of shuffled up-sequences [5] and the entropy of lengths of an LRM-partition [3], a novel measure that lies between runs and shuffled up-sequences.

For multiset sorting, the fine-grained measure, the *entropy of the multiplicities*, has been considered instead of the *number* of unique values already in early work in the field (e.g. [22, 26]). A more recent endeavor has been to find sorting methods that optimally adapt to *both presortedness and repeated values*. Barbay, Ochoa, and Satti refer to this as *synergistic sorting* [6] and present an algorithm based on quicksort that is optimal up to a constant factor. The method's practical performance is unclear.

We remark that (unstable) multiset sorting is the *only* problem from the above list for which a theoretically optimal algorithm has found wide-spread adoption in programming libraries: quicksort is known to almost optimally adapt to the entropy of multiplicities on average [32, 28, 34], when elements equal to the pivot are excluded from recursive calls (fat-pivot partitioning). Supposedly, sorting is so fast to start with that further improvements from exploiting specific input characteristics are only fruitful if they can be realized with minimal additional overhead. Indeed, for algorithms that adapt to the number of inversions, Elmasry and Hammad [11] found that the adaptive methods could only compete with good implementations of elementary sorting algorithms in terms of running time for inputs with extremely few inversions (fewer than 1.5%). Translating the theoretical guarantees of adaptive sorting into practical, efficient methods is an ongoing challenge.

1.2 Lower bound

How much does it help for sorting an array $A[1..n]$ to know that it contains r runs of respective sizes $L_1, \ldots, L_r$, i.e., to know the relative order of $A[1..L_1]$, $A[L_1+1..L_1+L_2]$ etc.? If we assume distinct elements, there are $\binom{n}{L_1,\ldots,L_r}$ permutations that are compatible with this setup, namely the number of ways to partition n keys into r subsets of given sizes. We thus need $\lg(n!) - \sum_{i=1}^{r} \lg(L_i!) = \mathcal{H}(\frac{L_1}{n}, \ldots, \frac{L_r}{n})n - O(n)$ comparisons to sort such an input. A formal argument for this lower bound is given by Barbay and Navarro [5] in the proof of their Theorem 2.

1.3 Results on Timsort and stack-based mergesort

Timsort's good performance in running-time studies, especially on partially sorted inputs, have lead to its adoption in several programming libraries, but until recently, no (nontrivial) worst-case adaptivity guarantee had been known. To make progress towards these, simplified variations of Timsort were considered [2, 9]. They maintain a stack of runs yet to be merged and proceed as follows: Find the next run in the input and push it onto the stack. Then consider the top k elements on the stack (for k a small constant like 3 or 4) and decide based on these if any pair of them is to be merged. If so, the two runs in the stack are replaced with the merged result and the rule is applied repeatedly until the stack satisfies some invariant. The invariant is chosen so as to keep the height of the stack small (logarithmic in n).

The simplest version, "α-stack sort" [2], merges the topmost two runs until the run lengths in the stack grow at least by a factor of α, (e.g., $\alpha = 2$). This method can lead to imbalanced merges (and hence runtime $\omega(n \log r)$ [9]; the authors of [2] also point this out

in their conclusion): if the topmost run is much longer than what is below it on the stack, merging the second and third runs (repeatedly until they are at least as big as the topmost run) is much better. This modification is called "α-merge sort". It achieves a worst-case guarantee of $O(n + n \log r)$, but the constant is provably not optimal [9].

Timsort is quite similar to α-merge sort for $\alpha = \varphi$ (the golden ratio) by forcing the run lengths to grow at least like Fibonacci numbers. The detailed rules for selecting merges are found in [1] or [9]. They imply a logarithmic stack height, but the actual bounds are much more involved than it appears at first sight [1]. As mentioned in the introduction, this has lead to the widespread deployment of faulty library code – twice! For inputs with specific run-lengths patterns, the implementations access stack cells beyond the (fixed) stack size.

Timsort was conjectured to always sort in $O(n + n \log r)$ time and in their contribution to these proceedings, Auger et al. [1] finally gave a proof for this bound. The constant of proportionality is not known exactly, but Buss and Knop [9] gave a family of inputs for which Timsort incurs asymptotically at least 1.5 times the required effort (in terms of merge costs, see Section 2.2). This proves that Timsort – like α-merge sort – is *not* optimally adaptive to the *number* of runs r (let alone the entropy of the run length distribution).

2 Preliminaries

Let $A[1..n]$ be an array of n elements that we sort using comparisons. By $\mathcal{H}$, we denote the binary Shannon entropy, i.e., for $p_1, \dots, p_m \in [0, 1]$ with $p_1 + \cdots + p_m = 1$ we have $\mathcal{H}(p_1, \dots, p_m) = \sum p_i \lg(1/p_i)$, where $\lg = \log_2$. We always let r denote the *number of runs* in the input and $L_1, \dots, L_r$ their respective lengths with $L_1 + \cdots + L_r = n$. In the literature, a *run* usually means a maximal (contiguous) weakly increasing (i.e., nondecreasing) region, but we adopt the convention from Timsort in this paper: a run is either a maximal *weakly increasing* region *or* a maximal *strictly decreasing* region. Decreasing runs are reversed upon detection; allowing only strictly decreasing runs makes their *stable* reversal trivial. (Note that the algorithms are not affected by these details of what constitutes a "run"; they only rely on a unique partition of the input into sorted segments that can be found by sequential scans.)

2.1 Nearly-Optimal Binary Search Trees

In the *optimal binary search tree problem,* we are given probabilities $\beta_1, \dots, \beta_m$ to access the m keys $K_1 < \cdots < K_m$ (internal nodes) and probabilities $\alpha_0, \dots, \alpha_m$ to access the gaps (leaves) between these keys (setting $K_0 = -\infty$ and $K_{m+1} = +\infty$) and we are interested in the binary search tree that minimizes the expected search cost C, i.e., the expected number of (ternary) comparisons when accesses follow the given distribution.[4] Nagaraj [23] surveys various versions of the problem. We confine ourselves to *approximation algorithms* here. Moreover, we only need the special case of *alphabetic trees* where all $\beta_j = 0$.

The following methods apply to the general problem, but we present them for the case of *nearly-optimal alphabetic trees.* So let $\alpha_0, \dots, \alpha_m$ with $\sum_{i=0}^{m} \alpha_i = 1$ be given. A greedy top-down approach produces provably good search trees if the details are done right [7, 19]: *choose the boundary closest to* $\frac{1}{2}$ *as the bisection at the root* (*"weight-balancing heuristic"*). Mehlhorn [21, §III.4.2] discusses two algorithms for nearly-optimal binary search trees that follow this scheme: "Method 1" is the straight-forward recursive application of the above rule, whereas "Method 2" (*"bisection heuristic"*) continues by strictly halving the *original* interval in the recursive calls; see Figure 1.

[4] We deviate from the literature convention and use m to denote the number of keys to avoid confusion with n, the length of the arrays to sort, in the rest of the paper.

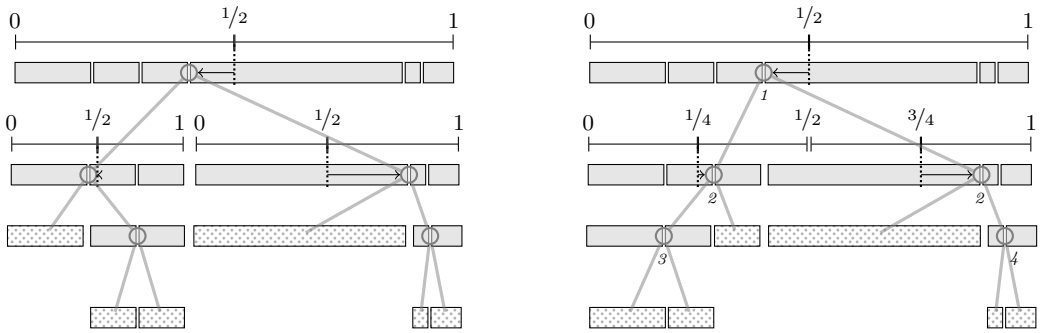

Figure 1 The two versions of weight-balancing for computing nearly-optimal alphabetic trees. The gap probabilities in the example are proportional to $5, 3, 3, 14, 1, 2$. **Left:** Mehlhorn's "Method 1" chooses the split closest to the midpoint of the subtree's actual weights ($1/2$ after renormalization). **Right:** "Method 2" continues to cut the original interval in half, irrespective of the total weight of the subtrees. The italic numbers are the powers of the nodes (see Definition 3 on page 8).

Method 1 was proposed in [31] and analyzed in [18, 7]; Method 2 is discussed in [19]. While Method 1 is arguably more natural, Method 2 has the advantage to yield splits that are predictable without going through all steps of the recursion. Both methods can be implemented to run in time $O(m)$ and yield very good trees. (Recall that in the case $\beta_j = 0$ the classic information-theoretic argument dictates $C \geq \mathcal{H}$; Bayer [7] gives lower bounds in the general case.)

▶ **Theorem 1** (Nearly-Optimal BSTs). *Let $\alpha_0, \beta_1, \alpha_1, \ldots, \beta_m, \alpha_m \in [0,1]$ with $\sum \alpha_i + \sum \beta_j = 1$ be given and let $\mathcal{H} = \sum_{i=0}^{m} \alpha_i \lg(1/\alpha_i) + \sum_{j=1}^{m} \beta_j \lg(1/\beta_j)$.*

- **(i)** *Method 1 yields a tree with search cost $C \leq \mathcal{H} + 2$. [7, Thm 4.8]*
- **(ii)** *If all $\beta_j = 0$, Method 1 yields a tree with search cost $C \leq \mathcal{H} + 2 - (m+3)\alpha_{\min}$, where $\alpha_{\min} = \min\{\alpha_0, \ldots, \alpha_m\}$. [16]*
- **(iii)** *Method 2 yields a tree with search cost $C \leq \mathcal{H} + 1 + \sum \alpha_i$. [19]*

2.2 Merge Costs

In this paper, we are primarily concerned with finding a good *order* of binary merges for the existing runs in the input. Following [2] and [9], we define the *merge cost* M for merging two runs of lengths m resp. n as $M = m + n$, i.e., the size of the result. This quantity has been studied earlier by Golin and Sedgewick [14] without giving it a name. Merge costs abstract away from key comparisons and element moves and simplify computations (see next subsection). Since any merge has to move most elements (except for rare lucky cases), and the average number of comparisons using standard merge routines is $m + n - \left(\frac{m}{n+1} + \frac{n}{m+1}\right)$, merge costs are a reasonable approximation, in particular when m and n are roughly equal. They always yield an upper bound for both the number of comparisons and moves.

2.3 Merge Trees

Let $L_1, \ldots, L_r$ with $\sum L_i = n$ be the lengths of the runs in the input. Any natural mergesort can be described as a rule to select some of the remaining runs, which are then merged and replaced by the merge result. If we always merge *two* runs this corresponds to a binary tree with the original runs at leaves $\boxed{1}, \ldots, \boxed{r}$. Internal nodes correspond to the result of merging their children. If we assign to internal node $\textcircled{j}$ the size M_j of the (intermediate) merge result it represents, then the overall merge cost is exactly $M = \sum_{\textcircled{j}} M_j$ (summing

over all internal nodes). Figure 1 shows two examples of merge trees; the merge costs are given by adding up all gray areas,[5] (ignoring the dotted leaves).

Let d_i be the *depth* of leaf $\boxed{i}$ (corresponding to the run of length L_i), where depth is the number of edges on the path to the root. Every element in the ith run is counted exactly d_i times in $\sum_{\textcircled{j}} M_j$, so we have $M = \sum_{i=1}^{r} d_i \cdot L_i$. Dividing by n yields $M/n = \sum_{i=1}^{r} d_i \cdot \alpha_i$ for $\alpha_i = L_i/n$, which happens to be the expected search time C in the merge tree if $\boxed{i}$ is requested with probability α_i for $i = 1, \dots, r$. So the minimal-cost merge tree for given run lengths $L_1, \dots, L_r$ is the optimal alphabetic tree for leaf probabilities $\frac{L_1}{n}, \dots, \frac{L_r}{n}$ and it holds $M \geq \mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_m}{n})n$. For distinct keys, the lower bound on comparisons (Section 1.2) coincides up to linear terms with this lower bound on the merge costs.

Combining this fact with the linear-time methods for nearly-optimal binary search trees from Section 2.1 immediately gives a stable sorting method that adapts optimally to existing runs up to an $O(n)$ term. We call such a method a *nearly-optimal (natural) mergesort*.

3 Nearly-Optimal Merging Orders

We now describe two new nearly-optimal mergesort variants that simulate nearly-optimal search-tree algorithms, but do so without ever storing the entire merge tree or run lengths.

3.1 Peeksort: A Simple Top-Down Method

The first method is similar to standard top-down mergesort in that it implicitly constructs a merge tree on the call stack. Instead of blindly cutting the input in half, however, we mimic Mehlhorn's Method 1. For that we need the *run boundary closest to the middle* of the input: this will become the root of the merge tree. Since we want to detect existing runs anyway at some point, we start by finding the run that contains the middle position. The end point closer to the middle determines the top-level cut and we recursively sort the parts left and right of it. A final merge completes the sort.

To avoid redundant scans, we pass on the information about detected runs. In the general case, we are sorting a range $A[\ell..r]$ whose prefix $A[\ell..e]$ and suffix $A[s..r]$ are (maximal) runs. Depending on whether the middle is contained in one of those runs, we have one of four different cases; apart from that the overall procedure (Algorithm 1) is quite straight-forward.

The following theorem shows that PeekSort is indeed a nearly-optimal mergesort. Unlike previous such methods, its code has very little overhead – in terms of both time and space, as well as in terms of conceptual complexity – over standard top-down mergesort, so it is a promising method for a practical nearly-optimal mergesort.

▶ **Theorem 2.** *The merge cost of* PeekSort *on an input consisting of r runs with respective lengths $L_1, \dots, L_r$ is at most $\mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_m}{n})n + 2n - (r+2)$, the number of comparisons is at most $\mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_m}{n}) + 3n - (2r+3)$. Both is optimal up to $O(n)$ terms (in the worst case).*

Proof. The recursive calls of Algorithm 1 produce the same tree as Mehlhorn's Method 1 with input $(\alpha_0, \dots, \alpha_m) = (\frac{L_1}{n}, \dots, \frac{L_r}{n})$ (i.e., $m = r - 1$) and $\beta_j = 0$. By Theorem 1–(ii), the search costs in this tree are $C \leq \mathcal{H} + 2 - (m+3)\alpha_{\min}$ with $\mathcal{H} = \mathcal{H}(\frac{L_1}{n}, \dots, \frac{L_m}{n})$. Since $L_j \geq 1$, we have $\alpha_{\min} \geq \frac{1}{n}$. As argued in Section 2.3, the overall merge costs are then given by $M = Cn \leq \mathcal{H}n + 2n - (r+2)$, which is within $O(n)$ of the lower bound for M.

[5] The left tree is obviously better here and this is a typical outcome. But there are also inputs where Method 2 yields a better tree than Method 1.

Algorithm 1 Peeksort: A simple top-down version of nearly-optimal natural mergesort. The initial call is PeekSort($A[1..n], 1, n$). Procedures ExtendRunLeft (-Right) scan left (right) starting at $A[m]$ as long as the run continues (and we did not cross the second parameter).

```
PeekSort(A[ℓ..r], e, s)
 1  if e == r or s == ℓ then return
 2  m = ℓ + ⌊(r−ℓ)/2⌋
 3  if m ≤ e                    // [ℓ ... m ... e | | s r]
 4      PeekSort(A[e+1..r], e+1, s)
 5      Merge(A[ℓ..e], A[e+1..r])
 6  else if m ≥ s               // [ℓ e | | s | m ... r]
 7      PeekSort(A[ℓ..s−1], e, s−1)
 8      Merge(A[ℓ..s−1], A[s..r])
 9  else   // Find existing run A[i..j] containing position m
10      i = ExtendRunLeft(A[m], ℓ);   j = ExtendRunRight(A[m], r)
11      if i == ℓ and j == r return
12      if m − i < j − m    // [ℓ e | | i | m ... j | | s r]
13          PeekSort(A[ℓ..i−1], e, i−1)
14          PeekSort(A[i..r], j, s)
15          Merge(A[ℓ..i−1], A[i..r])
16      else                // [ℓ e | | i | | m ... j | | s r]
17          PeekSort(A[ℓ..j], e, i)
18          PeekSort(A[j+1..r], j+1, s)
19          Merge(A[ℓ..j], A[j+1..r])
```

We can save at least one comparison per merge (namely for the last element). In total, we do exactly $r - 1$ merge operations. Apart from merging, we need a total of $n - 1$ additional comparisons to detect the existing runs in the input. Barbay and Navarro [5, Thm. 2] argued that $\mathcal{H}n - O(n)$ comparisons are necessary if the elements in the input are all distinct. ◀

3.2 Powersort: A Cache-Friendly Stack-Based Method

One little blemish remains in PeekSort: we have to use "random accesses" into the middle of the array to decide how to proceed. Even though we only use cache-friendly sequential scans, the I/O operations to load the middle run are effectively wasted, since it will often be merged only much later (after further recursive calls). Timsort proceeds in one left-to-right scan over the input and merges the top runs on the stack. This increases the likelihood to still have (parts of) the most recently detected run in cache when it is merged subsequently.

3.2.1 The power of top-down in a bottom-up method

In Method 2 for nearly-optimal search trees, we can unfold the recursion since all its potential cut points are predetermined: they have the form $c \cdot 2^{-\ell}$ for $\ell \in \mathbb{N}_{\geq 1}$ and odd c. The following definition characterizes the recursion depth ℓ, at which an internal node $\textcircled{j}$ is considered.

▶ **Definition 3** (Node Power). Let $\alpha_0, \ldots, \alpha_m, \sum \alpha_j = 1$ be leaf probabilities. For $1 \leq j \leq m$, let $\textcircled{j}$ be the internal node separating the $(j-1)$st and jth leaf. The *power* of (the split at)

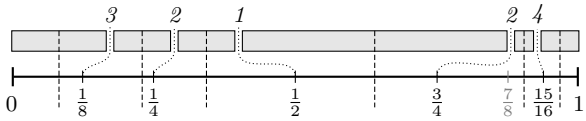

Figure 2 Illustration of node powers. The power of a run boundary is the smallest ℓ so that a multiple of $2^{-\ell}$ lies between the midpoints of the runs it separates.

node $\textcircled{j}$ is

$$P_j = \min\left\{\ell \in \mathbb{N} : \left\lfloor \frac{a}{2^{-\ell}} \right\rfloor < \left\lfloor \frac{b}{2^{-\ell}} \right\rfloor\right\}, \quad \text{where } a = \sum_{i=0}^{j-1} \alpha_i - \tfrac{1}{2}\alpha_{j-1}, \quad b = \sum_{i=0}^{j-1} \alpha_i + \tfrac{1}{2}\alpha_j.$$

(P_j is the index of the first bit where the (binary) fractional parts of a and b differ.)

Figure 2 illustrates Definition 3 for the nodes in our example (Figure 1). Intuitively, P_j is the "intended" depth of $\textcircled{j}$, but nodes occasionally end up higher in the tree if some leaf has a large weight relative to the current subtree, (see the rightmost branch in Figure 1). Mehlhorn's [19, 21] original implementation of Method 2, procedure *construct-tree*, does not single out the case that the next desired cut point lies *outside* the range of a subtree (cf. $\frac{7}{8}$ in Figure 2). This reduces the number of cases, but for our application, it is more convenient to explicitly check for this out-of-range case, and if it occurs to directly proceed to the next cut point. We refer to the modified algorithm as *Method 2′*. The extended version of this paper gives the details and shows that the changes do not affect the guarantee for Method 2 in Theorem 1. With this modification, the resulting tree is characterized by the node powers.

▶ **Lemma 4** (Path power monotonicity). *Consider the tree constructed by Method 2′. The powers of internal nodes along any root-to-leaf path are strictly increasing.*

The proof is given in the extended version.

3.2.2 Merging on-the-fly

Our second algorithm, "powersort", constructs the merge tree from left to right. Whenever the next internal node has a smaller power than the preceding one, we are following an edge from a left child up to its parent. That means that this subtree does not change anymore and we can execute any pending merges in it before continuing; the subtree then corresponds to the single resulting run. If we are instead following an edge down to a right child of a node, that subtree is still "open" and we only push the corresponding run onto the stack.

▶ Remark. The tree constructed by Method 2′ for leaf probabilities $\alpha_0, \ldots, \alpha_m$ is the (unique, min-oriented) Cartesian tree[6] for the sequence of node powers $P_1, \ldots, P_m$. It can be constructed iteratively (left to right) by the algorithm of Gabow, Bentley, and Tarjan [13], which effectively coincides with the procedure sketched above, except that we immediately collapse (merge) completed subtrees instead of keeping the entire tree.

As the runs on the stack correspond to a (right) path in the tree, the nodes have strictly increasing powers and we can bound the stack height by the maximal P_j. Since our leaf probabilities are $\alpha_j = \frac{L_j}{n} \geq \frac{1}{n}$, we have $P_j \leq \lfloor \lg n \rfloor + 1$. Algorithm 2 shows the detailed code.

[6] The Cartesian tree of an array of numbers is a binary tree. Its root corresponds to the global minimum in the array and its subtrees are the Cartesian trees of the numbers left resp. right of the minimum.

Algorithm 2 Powersort: A one-pass stack-based nearly-optimal natural mergesort. Procedure ExtendRunRight scans right as long as the run continues.

```
PowerSort(A[1..n])
 1  X = stack of runs                (capacity ⌊lg(n)⌋ + 1)
 2  P = stack of powers              (capacity ⌊lg(n)⌋ + 1)
 3  s1 = 1;  e1 = ExtendRunRight(A[1], n)    // A[s1..e1] is leftmost run
 4  while e1 < n
 5      s2 = e1 + 1;  e2 = ExtendRunRight(A[s2], n)    // A[s2..e2] next run
 6      p = NodePower(s1, e1, s2, e2, n)    // Pj for node (j) between A[s1..e1] and A[s2..e2]
 7      while P.top() > p      // previous merge deeper in tree than current
 8          P.pop()            // ⇝ merge and replace run A[s1..e1] by result
 9          (s1, e1) = Merge(X.pop(), A[s1..e1])
10      end while
11      X.push(A[s1, e1]);  P.push(p)
12      s1 = s2;  e1 = e2
13  end while // Now A[s1..e1] is the rightmost run
14  while ¬X.empty()
15      (s1, e1) = Merge(X.pop(), A[s1..e1])
16  end while

NodePower(s1, e1, s2, e2, n)
1  n1 = e1 − s1 + 1;  n2 = e2 − s2 + 1;  ℓ = 0
2  a = (s1 + n1/2 − 1)/n;  b = (s2 + n2/2 − 1)/n
3  while ⌊a · 2^ℓ⌋ == ⌊b · 2^ℓ⌋ do ℓ = ℓ + 1 end while
4  return ℓ
```

▶ **Theorem 5.** *The merge cost of* PowerSort *is at most* $\mathcal{H}(\frac{L_1}{n}, \ldots, \frac{L_m}{n}) + 2n$*, the number of comparisons is at most* $\mathcal{H}(\frac{L_1}{n}, \ldots, \frac{L_m}{n}) + 3n - r$*. Moreover, (apart from buffers for merging) only* $O(\log n)$ *words of extra space are required.*

Proof. The merge tree of PowerSort is exactly the search tree constructed by Method 2′ on leaf probabilities $(\alpha_0, \ldots, \alpha_m) = (\frac{L_1}{n}, \ldots, \frac{L_r}{n})$ and $\beta_j = 0$. By Theorem 1–(iii), the search costs are $C \le \mathcal{H} + 2$ with $\mathcal{H} = \mathcal{H}(\frac{L_1}{n}, \ldots, \frac{L_m}{n})$, so the overall merge costs are $M = Cn \le \mathcal{H}n + 2n$, which is within $O(n)$ of the lower bound for M. The comparisons follow as for Theorem 2. ◀

4 Running-Time Study

We conducted a running-time study comparing the two new nearly-optimal mergesorts with current state-of-the-art implementations and elementary mergesort variants. The code is available on github [33]. The main goal of this study is to show that

1. peeksort and powersort have very little overhead compared to standard (non-natural) mergesort variants (i.e., they are never (much) slower); and at the same time
2. peeksort and powersort outperform other mergesort variants on partially presorted inputs.

Timsort is arguably the most used adaptive sorting method; a secondary goal is hence to

3. investigate the typical merge costs of Timsort on different inputs, in particular in light of the recent negative theoretical results by Buss and Knop [9].

4.1 Setup

Oracle's Java runtime library includes a highly tuned Timsort implementation; to be able to directly compare with it, we chose to implement our algorithms in Java. We repeated all experiments with C++ ports of the algorithms and the results are qualitatively the same; we will comment on a few differences in the following subsections. The Timsort implementation is used for `Object[]`, i.e., arrays of *references* to objects. Since the location of objects on the heap is hard to control, this is likely to yield big variations in the number of cache misses. We chose to sort `int[]`s instead to obtain more reproducible results, and modified the library implementation of Timsort accordingly. This scenario makes key comparisons and element moves relatively cheap and thereby emphasizes the remaining overhead of the methods, which is in line with our primary goal 1) from above.

We compare our methods with standard top-down and bottom-up mergesort implementations. We use the code given by Sedgewick [27, Programs 8.3 and 8.5] as bases. In both cases, we add a check before calling merge to detect if the runs happen to already be in sorted order, and we use insertion sort for base cases of size $n \leq w = 24$. (For bottom-up mergesort, we start by sorting chunks of size $w = 24$.) Our implementations of peeksort and powersort are described in more detail in the extended version; apart from a mildly optimized version of the pseudocode, we added the same cutoff / minimal run length ($w = 24$) as above.

All our methods use Sedgewick's bitonic merge procedure [27, Program 8.2], whereas the Java library Timsort contains a custom merge method that tries to save key comparisons: when only elements from the same run are selected for the output repeatedly, Timsort enters a *"galloping mode"* and uses exponential search (instead of the conventional sequential search) to find the insertion position of the next element. Details are described by Peters [25]. Since saving comparisons is not of utmost importance in our scenario of sorting `int`s, we also added a version of Timsort, called "trotsort", that uses our plain merge method instead of galloping, but is otherwise identical to the library Timsort.

We use the following inputs types:

- *random permutations* are a case where no long runs are present to exploit;
- *"random-runs" inputs* are constructed from a random permutation by sorting segments of random lengths, where the lengths are chosen independently according to a geometric distribution with a given mean ℓ; since the geometric distribution has fairly large variance, these inputs tend to have runs whose sizes vary a lot;
- *"Timsort-drag" inputs* are special instances of random-runs inputs where the run lengths are chosen as $\mathcal{R}_{\text{tim}}$, the bad-case example for Timsort from [9, Thm. 3].

We remark that the above input types are chosen with our specific goals from above in mind; we do not attempt to model inputs from "real-world" applications in this study.

4.2 Overhead of Nearly-Optimal Merge Order

We first consider random permutations as inputs. Since random permutations contain (with high probability) no long runs that can be exploited, the adaptive methods will not find anything that would compensate for their additional efforts to identify runs. (This is confirmed by the fact that the total merge costs of all methods, including Timsort, are within 1.5% of each other in this scenario.) Figure 3 shows average running times for input sizes from 100 000 to 100 million ints. (Measurements for $n = 10\,000$ were too noisy to draw meaningful conclusions.)

Varying input sizes over several orders of magnitude, we consistently see the following picture: `Arrays.sort(int[])` (dual-pivot quicksort) is the fastest method. It is not a stable

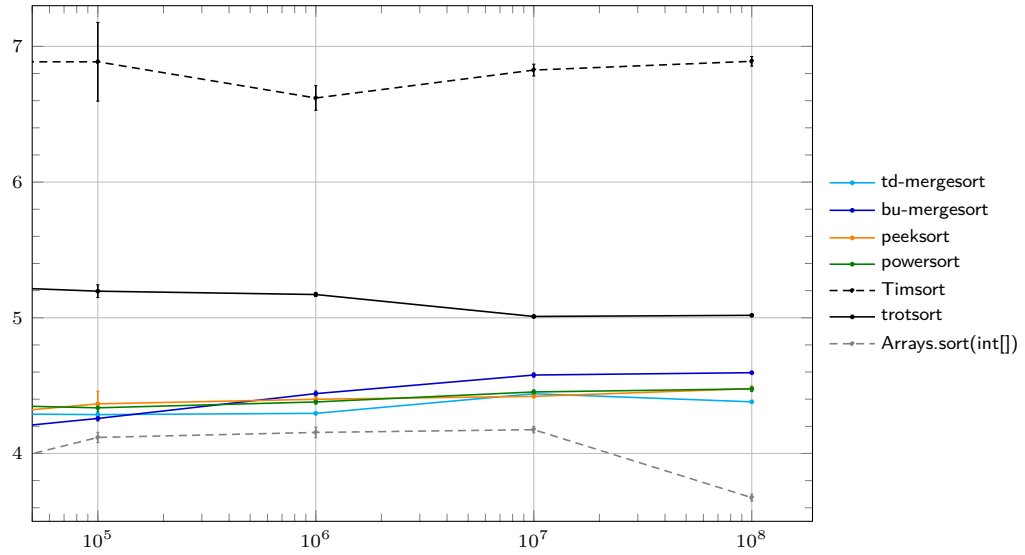

Figure 3 Normalized running times on random permutations. The logarithmic x-axis shows n, the y-axis shows the average of $t/(n \lg n)$ where t is the running time in ms. Error bars show one standard deviation (stdev). We plot averages over 1000 repetitions (200 resp. 20 for the largest n).

sort and thus merely serves as a baseline. Top-down and bottom-up mergesort, peeksort and powersort are 20–30% slower than dual-pivot quicksort. Comparing the four to each other, no large differences are visible; if anything, bottom-up mergesort was a bit slower than the others (for large inputs). Since the recursion cutoff resp. minimal run length $w = 24$ exceeded the length of *all* runs in all inputs, we effectively have equal-length runs. Merging left to right (as in bottom-up mergesort) then performs just fine, and top-down mergesort finds a close-to-optimal merging order in this case. That peek- and powersort perform essentially *as good as* elementary mergesorts on random permutations thus clearly indicates that their overhead for determining a nearly-optimal merging order is negligible.

The library Timsort performs surprisingly poorly on `int[]`s, probably due to the relatively cheap comparisons. Replacing the galloping merge with the ordinary merge alleviates this (see "trotsort"), but Timsort remains inferior on random permutations by a fair margin (10–20%). In C++, trotsort with straight insertion sort (instead of binary insertion sort) as base case was roughly as fast as the other mergesorts.

4.3 Practical speedups by adaptivity

After demonstrating that we do not lose much by using our adaptive methods when there is nothing to adapt to, we next investigate how much can be gained if there is. We consider random-runs inputs as described above. This input model instills a good dose of presortedness, but not in a way that gives any of the tested methods an obvious advantage or disadvantage over the others. We choose a representative size of $n = 10^7$ and an expected run length $\ell = 3\,000$, so that we expect roughly $\sqrt{n}$ runs of length $\sqrt{n}$.

If this was a random permutation, we would expect merge costs of roughly $n \lg(n/w) \approx 1.87 \cdot 10^8$ (indeed a bit above this). The right chart in Figure 4 shows that the adaptive methods can bring the merge cost down to a little over 60% of this number. Powersort achieved average merge costs of $1.14 \cdot 10^8 < n \lg r \approx 1.17 \cdot 10^8$, i.e., less than a method would that only adapts to the *number* of runs r.

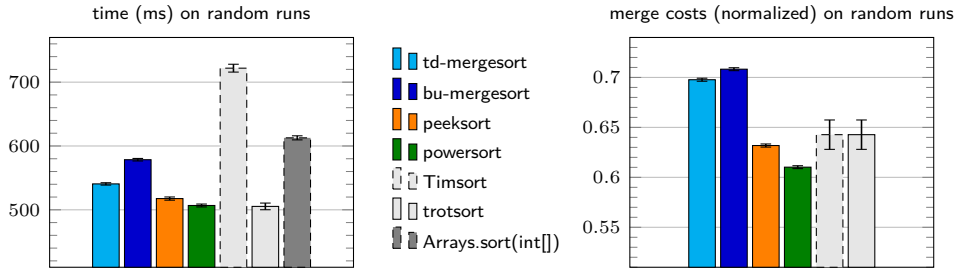

Figure 4 Average running times (left) and normalized merge cost (right) on random-runs inputs with $n = 10^7$ and $\ell = 3\,000 \approx \sqrt{n}$. Error bars show one stdev. Merge costs have been divided by $n \lg(n/w) \approx 1.87 \cdot 10^8$, which is the merge cost a (hypothetical) optimal mergesort that does not pick up existing runs, but starts with runs of length $w = 24$. Note that the simple sorted-check before a merge in the standard mergesorts already reduces the merge costs to roughly 70% of that number.

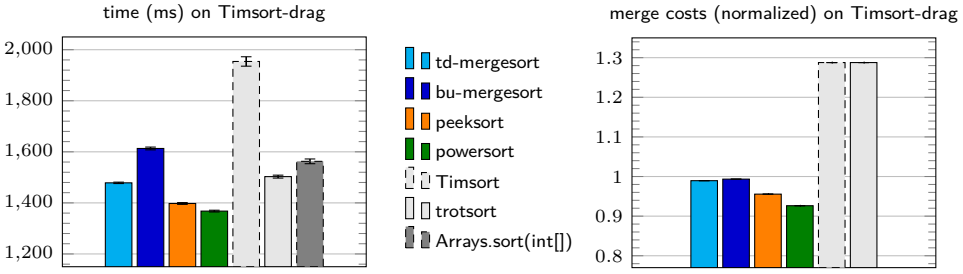

Figure 5 Average running times (left) and normalized merge cost (right) on "Timsort-drag" inputs with $n = 2^{24}$ and run lengths $\mathcal{R}_{\text{tim}}(2^{24}/32)$ multiplied by 32. Error bars show one standard deviation. Merge costs have been divided by $n \lg(n/w) \approx 3.26 \cdot 10^8$.

In terms of running time, powersort is again among the fastest stable methods, and indeed 20% *faster* than `Arrays.sort(int[])`. The best adaptive methods are also 10% faster than top-down mergesort. (Note that our standard mergesorts are "slightly adaptive" by skipping a merge of two runs that happened to already be in order, which can be checked with a single comparisons.) This supports the statement that significant speedups can be realized by adaptive sorting on inputs with existing order, and $\sqrt{n}$ runs suffice for that. If we increase ℓ to 100 000, so that we expect only roughly 100 long runs, the library quicksort becomes twice as slow as powersort and trotsort.

Timsort with galloping merges is again uncompetitive. Trotsort's running time is a bit anomalous. Even though it occasionally incurs 10% more merge costs on a given input than powersort, the Java running times were within 1% of each other. However, merge costs seem to more accurately predict running time in C++; there trotsort was 8% slower than powersort.

4.4 Non-optimality of Timsort

Finally, we consider "Timsort-drag" inputs, a sequence of run lengths $\mathcal{R}_{\text{tim}}(n)$ specifically crafted by Buss and Knop [9] to generate unbalanced merges (and hence large merge cost) in Timsort. Since actual Timsort implementations employ minimal run lengths of up to 32 elements we multiplied the run lengths by 32. Figure 5 shows running time and merge cost for all methods on a characteristic Timsort-drag input of length $2^{24} \approx 1.6 \cdot 10^7$.

In terms of merge costs, Timsort resp. trotsort now pays 30% more than even a simple non-adaptive mergesort, whereas peeksort and powersort obviously retain their proven nearly-optimal behavior. Also in terms of running time, trotsort is a bit slower than top-down mergesort, and 10% slower than powersort on these inputs. It is remarkable that the dramatically larger merge cost does not lead to a similarly drastic slow down in Java. In C++, however, trotsort was *40% slower* than powersort, in perfect accordance with the difference in merge costs.

It must thus be noted that Timsort's merging-order rules have weaknesses, and it is unclear if more dramatic examples are yet to be found.

5 Conclusion

In this paper, we have demonstrated that provably good merging orders for natural mergesort can be found with negligible overhead. The proposed algorithms, peeksort and powersort, offer more reliable performance than the widely used Timsort, and at the same time, are arguably simpler.

Powersort builds on a modified bisection heuristic for computing nearly-optimal binary search trees that might be of independent interest. It has the same quality guarantees as Mehlhorn's original formulation, but allows the tree to be built "bottom-up" as a Cartesian tree over a certain sequence, the "node powers". It is the only such method for nearly-optimal search trees to our knowledge.

Buss and Knop conclude with the question whether there exists a k-aware algorithm (a stack-based natural mergesort that only considers the top k runs in the stack) with merge cost $(1 + o_r(1))$ times the optimal merge cost [9, Question 37]. Powersort effectively answers this question in the affirmative with $k = 3$.[7]

5.1 Extensions and future work

Timsort's "galloping merge" procedure saves comparisons when we consistently consume elements from one run, but in "well-mixed" merges, it does not help (much). It would be interesting to compare this method with other comparison-efficient merge methods.

Another line of future research is to explore ways to profit from duplicate keys in the input. The ultimate goal would be a "synergistic" sorting method (in the terminology of [6]) that has practically no overhead for detecting existing runs and equals and yet exploits their *combined* presence optimally.

References

1 Nicolas Auger, Vincent Jugé, Cyril Nicaud, and Carine Pivoteau. On the worst-case complexity of TimSort. In Hannah Bast Yossi Azar and Grzegorz Herman, editors, *26th Annual European Symposium on Algorithms (ESA 2018)*, Leibniz International Proceedings in Informatics (LIPIcs), 2018. doi:10.4230/LIPIcs.ESA.2018.4.

2 Nicolas Auger, Cyril Nicaud, and Carine Pivoteau. Merge strategies: from merge sort to TimSort, 2015. URL: https://hal-upec-upem.archives-ouvertes.fr/hal-01212839.

[7] Strictly speaking, powersort needs a relaxation of the model of Buss and Knop. They require decisions to be made solely based on the *lengths* of runs, whereas node power takes the location of the runs within the array into account. Since the location of a run must be stored anyway, this appears reasonable to us.

3 Jérémy Barbay, Johannes Fischer, and Gonzalo Navarro. LRM-trees: Compressed indices, adaptive sorting, and compressed permutations. *Theoretical Computer Science*, 459:26–41, 2012. doi:10.1016/j.tcs.2012.08.010.

4 Jérémy Barbay and Gonzalo Navarro. Compressed representations of permutations, and applications. In Susanne Albers and Jean-Yves Marion, editors, *26th International Symposium on Theoretical Aspects of Computer Science (STACS 2009)*, pages 111–122, Freiburg, Germany, 2009. URL: https://hal.inria.fr/inria-00358018.

5 Jérémy Barbay and Gonzalo Navarro. On compressing permutations and adaptive sorting. *Theoretical Computer Science*, 513:109–123, 2013. doi:10.1016/j.tcs.2013.10.019.

6 Jérémy Barbay, Carlos Ochoa, and Srinivasa Rao Satti. Synergistic solutions on multisets. In Juha Kärkkäinen, Jakub Radoszewski, and Wojciech Rytter, editors, *28th Annual Symposium on Combinatorial Pattern Matching (CPM 2017)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 31:1–31:14, 2017. doi:10.4230/LIPIcs.CPM.2017.31.

7 Paul J. Bayer. *Improved Bounds on the Cost of Optimal and Balanced Binary Search Trees.* Master's thesis, Massachusetts Institute of Technology, 1975.

8 Jon L. Bentley and M. Douglas McIlroy. Engineering a sort function. *Software: Practice and Experience*, 23(11):1249–1265, 1993. doi:10.1002/spe.4380231105.

9 Sam Buss and Alexander Knop. Strategies for stable merge sorting, 2018. arXiv:1801.04641.

10 Stijn de Gouw, Frank S. de Boer, Richard Bubel, Reiner Hähnle, Jurriaan Rot, and Dominic Steinhöfel. Verifying OpenJDK's sort method for generic collections. *Journal of Automated Reasoning*, aug 2017. doi:10.1007/s10817-017-9426-4.

11 Amr Elmasry and Abdelrahman Hammad. Inversion-sensitive sorting algorithms in practice. *Journal of Experimental Algorithmics*, 13:11:1–11:18, 2009. doi:10.1145/1412228.1455267.

12 Vladmir Estivill-Castro and Derick Wood. A survey of adaptive sorting algorithms. *ACM Computing Surveys*, 24(4):441–476, dec 1992. doi:10.1145/146370.146381.

13 Harold N. Gabow, Jon Louis Bentley, and Robert E. Tarjan. Scaling and related techniques for geometry problems. In *16th annual ACM symposium on Theory of computing (STOC 1984)*, pages 135–143. ACM Press, 1984. doi:10.1145/800057.808675.

14 Mordecai J. Golin and Robert Sedgewick. Queue-mergesort. *Information Processing Letters*, 48(5):253–259, dec 1993. doi:10.1016/0020-0190(93)90088-q.

15 Chris Hegarty. Replace "modified mergesort" in java.util.Arrays.sort with timsort, 2009. URL: https://bugs.openjdk.java.net/browse/JDK-6804124.

16 Yasuichi Horibe. An improved bound for weight-balanced tree. *Information and Control*, 34(2):148–151, jun 1977. doi:10.1016/S0019-9958(77)80011-9.

17 Donald E. Knuth. *The Art Of Computer Programming: Searching and Sorting.* Addison Wesley, 2nd edition, 1998.

18 Kurt Mehlhorn. Nearly optimal binary search trees. *Acta Informatica*, 5(4), 1975. doi:10.1007/BF00264563.

19 Kurt Mehlhorn. A best possible bound for the weighted path length of binary search trees. *SIAM Journal on Computing*, 6(2):235–239, 1977. doi:10.1137/0206017.

20 Kurt Mehlhorn. Sorting presorted files. In *Theoretical Computer Science 4th GI Conference*, pages 199–212. Springer, 1979. doi:10.1007/3-540-09118-1_22.

21 Kurt Mehlhorn. *Data Structures and Algorithms 1: Sorting and Searching.* Springer, 1984.

22 Ian Munro and Philip M. Spira. Sorting and searching in multisets. *SIAM Journal on Computing*, 5(1):1–8, 1976. doi:10.1137/0205001.

23 S.V. Nagaraj. Optimal binary search trees. *Theoretical Computer Science*, 188(1-2):1–44, nov 1997. doi:10.1016/S0304-3975(96)00320-9.

24 Tim Peters. [Python-Dev] Sorting, 2002. URL: https://mail.python.org/pipermail/python-dev/2002-July/026837.html.

25 Tim Peters. Timsort, 2002. URL: http://hg.python.org/cpython/file/tip/Objects/listsort.txt.

26 Robert Sedgewick. Quicksort with equal keys. *SIAM Journal on Computing*, 6(2):240–267, 1977.

27 Robert Sedgewick. *Algorithms in Java*. Addison-Wesley, 2003.

28 Robert Sedgewick and Jon Bentley. Quicksort is optimal (talk slides), 2002. URL: http://www.cs.princeton.edu/~rs/talks/QuicksortIsOptimal.pdf.

29 Tadao Takaoka. A new measure of disorder in sorting – entropy. In *The Fourth Australasian Theory Symposium (CATS 1998)*, pages 77–86, 1998.

30 Tadao Takaoka. Partial solution and entropy. In *MFCS 2009*, pages 700–711, 2009. doi:10.1007/978-3-642-03816-7_59.

31 W.A. Walker and C.C. Gotlieb. A top-down algorithm for constructing nearly optimal lexicographic trees. In *Graph Theory and Computing*, pages 303–323. Elsevier, 1972. doi:10.1016/B978-1-4832-3187-7.50023-4.

32 Lutz M. Wegner. Quicksort for equal keys. *IEEE Transactions on Computers*, C-34(4):362–367, 1985. doi:10.1109/TC.1985.5009387.

33 Sebastian Wild. Accompanying code for running time study, 2018. doi:10.5281/zenodo.1241162.

34 Sebastian Wild. Quicksort is optimal for many equal keys. In *ANALCO 2018*, pages 8–22. SIAM, jan 2018. doi:10.1137/1.9781611975062.2.