



# International Journal of Research Publication and Reviews

Journal homepage: [www.ijrpr.com](http://www.ijrpr.com) ISSN 2582-7421

## Beyond the Myth of Linear Sorting : A Deep Dive into the $\Omega(n \log n)$ Constraint

**Ronit Kumar Verma<sup>1</sup>, Sourav Kumar<sup>2</sup>, Utkarsh Satyarthi<sup>3</sup>, Priyanshu Singh<sup>4</sup>**

Department of Computer Science & Engineering,

IIMT College of Engineering, Greater Noida

[ronitroy22678@gmail.com](mailto:ronitroy22678@gmail.com), [pnbsorav@gmail.com](mailto:pnbsorav@gmail.com), [utkarshanand9988@gmail.com](mailto:utkarshanand9988@gmail.com), [prisin.priyanshu2235@gmail.com](mailto:prisin.priyanshu2235@gmail.com)

### ABSTRACT –

Sorting is a fundamental operation in computer science essential to fields such as data analysis, database management, operating systems and algorithm development. Despite progress in computing technology no comparison-based sorting method has surpassed the worst-case lower bound of  $\Omega(n \log n)$ . This article explores the basis, for this constraint through the decision tree framework and clarifies why it is impossible to attain a universally linear-time ( $O(n)$ ) sorting algorithm. It also takes a closer look at algorithms like Counting Sort, Radix Sort, and Bucket Sort, which can reach linear time under specific conditions. By comparing comparison-based and non-comparison sorting methods, the paper clears up some common misunderstandings about the possibility of general-purpose linear-time sorting. In conclusion, it finds that linear sorting can only happen under certain constraints, while  $\Omega(n \log n)$  remains an unyielding barrier for arbitrary data.

**Keywords** - Sorting algorithms, comparison model,  $\Omega(n \log n)$ , decision tree complexity, linear-time sorting, Radix Sort, Counting Sort.

### INTRODUCTION

Sorting ranks as one of the most vital computational primitives in computer science that directly influences the implementation of numerous operations along with searching, indexing, data retrieval, compression, scientific computing, and large-scale data optimization. Almost every software system existing in the real world—database engines, operating systems, distributed computing platforms, and machine learning pipelines—depend on efficient sorting routines to be able to keep up performance and stability. The demand for faster and more scalable sorting algorithms has become extremely urgent with the continuous exponential growth of datasets.

Classical algorithms such as MergeSort, HeapSort, and QuickSort have been very well optimized over the last several decades, thereby, achieving almost optimal performance with the theoretical time complexity being very close to  $O(n \log n)$ . Nevertheless, the issue that keeps on attracting the attention of researchers is whether it is possible to make sorting fundamentally faster in the typical case.

The desire to find a single linear-time sorting algorithm has increased significantly with the advent big data applications, high-frequency trading systems, real-time analytics, and cloud-scale processing. In fact, numerous practitioners erroneously believe that because there are linear-time sorting techniques, e.g., Radix Sort, Counting Sort, and Bucket Sort, these can be universally employed to replace traditional  $O(n \log n)$  algorithms. Yet, this assumption disregards the sorting constraints that are of theoretical nature. Linear-time sorting methods are only capable of working in very limited scenarios, for example with fixed-size integer keys, limited input ranges, or uniform distributions. If these conditions are not met, their efficiency diminishes or they become infeasible, thus a sort of comparison-based sorting algorithms is still required for general-purpose systems.

The core issue is the intrinsic computational challenge of telling apart the  $n!$  possible permutations of  $n$  elements. This difficulty is expressed by the  $\Omega(n \log n)$  lower bound for all comparison-based sorting algorithms — a limit that cannot be circumvented without going against basic information-theoretic principles. Thus no universal sorting method can be on the other side of this theoretical barrier, regardless of improvements in hardware, parallel processing, or algorithmic engineering.

Facing these facts, the current research is devoted to unveiling the deeper reasons for sorting complexity and sorting misconceptions clarification. In particular, it deals with the first three critical questions:

1. What is the mathematical reason behind the  $\Omega(n \log n)$  lower bound for comparison-based sorting?
2. Why is it theoretically impossible to design a universal linear-time sorting algorithm?
3. Under what special conditions is  $O(n)$  sorting achievable, and what algorithms exploit these conditions?

Upon answering these problems, the article first of all pinpoints theoretical limits that comparison-based sorting cannot go beyond and also deals with specialized linear-time algorithms that achieve this by utilizing structural patterns of input data. The subsequent parts present an extensive survey of the literature, analyze the decision-tree model that is at the basis of the  $\Omega(n \log n)$  limitation, and consider non-comparison sorting methods that under fairly restrictive conditions are able to circumvent this bound.

## LITERATURE SURVEY / REVIEW

S.No	Title	Author(s)	Year	Methodology	Research Contribution
1	The Art of Computer Programming, Vol. 3: Sorting and Searching	Donald Knuth	1973	Decision tree analysis, algorithm theory	Established the $\Omega(n \log n)$ lower bound for comparison-based sorting using decision tree models.
2	The Design and Analysis of Computer Algorithms	Aho, Hopcroft & Ullman	1974	Theoretical analysis, algorithm complexity	Provided formal computational models and proofs supporting sorting lower bounds.
3	Introduction to Algorithms (CLRS)	Cormen, Leiserson, Rivest & Stein	2009	Algorithm analysis, pseudocode, complexity theory	Detailed linear-time sorting methods and constraints such as data ranges and digit processing.
4	Algorithms	Sedgewick & Wayne	2011	Empirical algorithm study	Compared practical performance and implementation details of classical sorting algorithms.
5	Data Structures and Algorithm Analysis in C++	Mark Allen Weiss	2012	Algorithmic design and analysis	Explored practical implications of lower bounds and real-world constraints for sorting efficiency.

## METHODOLOGY

The methodology of this study is based on a strict theoretical analysis of sorting in both the comparison model and the non-comparison model. The authors do not resort to empirical testing but rather use mathematical instruments to explain why certain sorting algorithms have inevitable performance boundaries. They want to look at sorting as an information-processing task and figure out how much work an algorithm has to do to differentiate between all possible input orders. By using this theory as a mean, we are able to uncover not only the restrictions of comparison-based sorting but also the situations in which sorting in linear time is feasible.

To determine the lower bound of comparison-based sorts, the paper uses a decision tree model, which depicts any algorithm as a binary tree. A comparison between two elements is an internal node of the tree, and each branch shows the result of that comparison. The leaves of the tree correspond to the possible permutations of the input. Since a proper sorting algorithm should be able to figure out which one of the  $n!$  permutations is the input, its decision tree must have at least  $n!$  leaves. The connection between the height of the decision tree ( $h$ ) and the number of leaves ( $L$ ) is employed to get the lower bound. Therefore, we use the inequality:

$$2^h \geq n!$$

Taking the logarithm of both sides gives:

$$h \geq \log_2(n!)$$

By applying Stirling's approximation, the expression simplifies to:

$$\log_2(n!) = \Theta(n \log n)$$

Thus, the worst-case number of comparisons required by any comparison-based sorting algorithm is:

$$\Omega(n \log n)$$

The proof steps here are valid regardless of programming language, hardware architecture, memory hierarchy, or instruction-level optimizations. They show that the  $\Omega(n \log n)$  lower bound is a limitation that exists at the deepest level of theory rather than in practice. All of these algorithms, namely MergeSort, HeapSort, QuickSort (average case), Binary Tree Sort, and even an optimized comparison-based hybrid, are subject to this constraint as they only use comparison operations to establish order.

The methodology also analyzes why universal  $O(n)$  sorting is impossible. A universal sorting algorithm must handle arbitrary data types, arbitrary distributions, and arbitrary structures without making assumptions. In such cases, comparison becomes the only reliable operation for determining order. Since comparisons reveal only one bit of information at a time, and since identifying the correct permutation among  $n!$  possibilities requires  $\log_2(n!)$  bits of information, no algorithm can reduce the number of comparisons below the  $\Omega(n \log n)$  limit. Even with parallel processors or hardware-level optimizations, the asymptotic lower bound remains unchanged because it stems from fundamental information-theoretic constraints.

In order to know at what point  $O(n)$  sorting can be performed, the method also looks at non-comparison-based algorithms like Counting Sort, Radix Sort, and Bucket Sort. These algorithms do not compare elements as usual but instead they get the data from the input string. For instance, Counting Sort counts how many times each key value appears in a range  $[0, k]$  and because of that it can run in  $O(n + k)$  time. The algorithm works in linear time if  $k$  is very small as compared to  $n$ . Radix Sort uses Counting Sort to process each digit of numbers or strings and it is able to perform in a linear time if the number of digits is constant or logarithmic in  $n$ . Bucket Sort, on the other hand, puts the elements into buckets using a hash-like function and can have  $O(n)$  expected performance if the inputs are uniformly distributed.

These algorithms are going beyond the  $n \log n$  barrier not by breaking the lower bound, but by going beyond the comparison model altogether. They perform arithmetic operations, extract digits, or make distribution assumptions to get around the informational bottleneck that comparison operations impose. As a result, the method delineates the precise frontier between the infeasible (general-purpose linear sorting) and the feasible (conditional linear sorting under structured inputs). This conceptual model is at the core of the research and is in agreement with the findings that are presented in the later parts.

## RESULT AND DISCUSSION

### $\Omega(n \log n)$ Lower Bound for Comparison-Based Sorting

- Any comparison-based sorting algorithm is bound by a theoretical minimum that is  $\Omega(n \log n)$  in the worst-case scenario. The reason for this is that an algorithm that works only by comparisons must be able to tell apart all possible orderings of  $n$  elements. The number of permutations is  $n!$ , so the algorithm has to make enough comparisons to be the one that it has uniquely identified the correct ordering. This is what directly determines that in the worst case, the number of comparisons has to be at least  $\Omega(n \log n)$ .
- The lower bound is due to the fact that every comparison-based sorting algorithm has to differentiate between  $n!$  permutations of the input. Given that comparisons yield only one bit of information at a time, the algorithm has to perform a sequence of decisions that is sufficiently deep to pinpoint one permutation out of the  $n!$  possible ones. This results in an inevitable complexity boundary.

Using the decision tree model:

- Internal nodes represent comparisons (e.g.,  $A[i] < A[j]$ ).
- Leaves represent the sorted permutations of the input.
- Because a binary tree of height  $h$  has at most  $2^h$  leaves, and a sorting tree needs at least  $n!$  leaves, the height must satisfy:

$$2^h \geq n!$$

$$h \geq \log_2(n!)$$

$$\log_2(n!) = \Theta(n \log n)$$

Thus, the algorithm must perform at least  $\Theta(n \log n)$  comparisons.

- This computational requirement implies that the commonly used algorithms such as MergeSort, HeapSort, QuickSort (average case), and Binary Tree Sort cannot perform better than  $\Omega(n \log n)$  performance bound. Essentially, these algorithms, which are strictly within the comparison model, cannot achieve  $o(n \log n)$  in the worst case, no matter if they are optimized or refined to a higher degree.

### Universal $O(n)$ Sorting Is Impossible

- Linear-time sorting cannot be a universal one as arbitrary data do not have any inherent structure. An algorithm cannot use a shortcut such as direct addressing, hashing, or digit grouping if there are no constraints on the input. So, comparisons remain the only reliable ordering operation.
- According to information theory, sorting  $n$  arbitrary elements need  $\log_2(n!)$  bits of information, but a comparison provides only one bit. This immediately limits any universally applicable sorting algorithm to at least  $\Omega(n \log n)$  steps thus making  $O(n)$  time impossible within the comparison framework.
- The limit cannot be surpassed by hardware optimizations, parallelism, or architectural improvements. Although they may reduce the constant factors and speed up the process in reality, they cannot change the fundamental complexity class set by information theory.

### Linear-Time Sorting Is Possible Under Specific Conditions

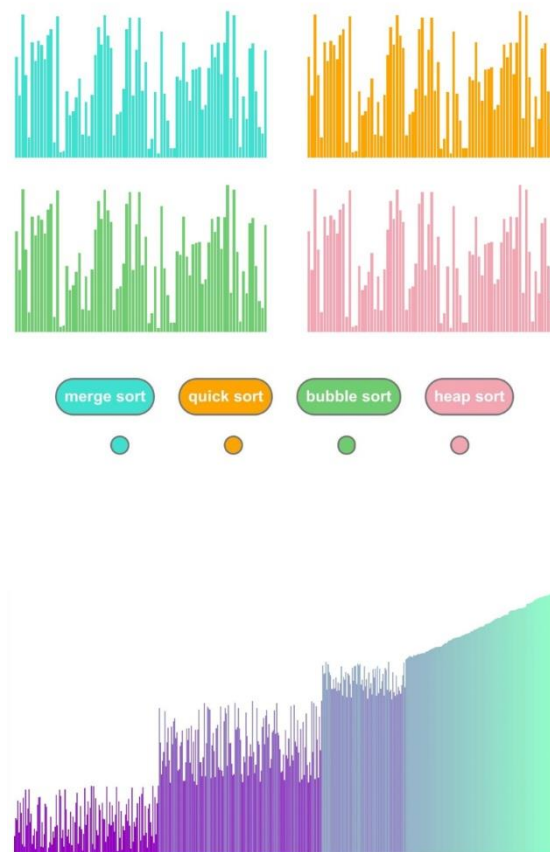
- Counting Sort ( $O(n + k)$ )**
  - Works only when the input consists of integers within a small fixed range  $[0, k]$ .
  - Achieves true linear time when  $k = O(n)$ , since the auxiliary counting array remains manageable in size.
- Radix Sort ( $O(n \log k)$  or  $O(n \times k)$ )**
  - Suitable for fixed-length integers or strings, where elements can be processed digit by digit.
  - Relies on Counting Sort as a stable subroutine to maintain correct relative ordering across digit positions.
- Bucket Sort ( $O(n)$  expected time)**
  - Assumes that input values are drawn from a uniform distribution, allowing elements to be evenly distributed into buckets.
  - Runs in expected linear time but can degrade to higher complexity in the worst case if distribution assumptions are violated.

### Final Outcome

- $\Omega(n \log n)$  is a hard lower bound for all comparison-based sorting algorithms.
- Linear-time sorting is possible only for restricted datasets, such as integers with limited ranges or uniform distributions.
- The illusion of universal linear sorting is clarified: faster sorting is achievable only by exploiting data structure, not by breaking theoretical limits.

## CONCLUSION

The  $(n \log n)$  lower bound is a limit which cannot be surpassed by any comparison-based sorting algorithm and serves as a basis for all such algorithms. This is evident when looking at the decision tree model. The task of finding the right permutation among  $n!$  has to involve at least  $(n \log n)$  comparisons. No matter what the algorithm or the machine is like, this limit cannot be overcome. On the other hand, sorting in linear time can be done for certain types of inputs. Counting Sort, Radix Sort, and Bucket Sort are three examples of algorithms that can run in  $O(n)$  time if certain conditions such as integer ranges, fixed-sized keys, or uniform distributions are met. These methods show where the limit is between the theoretically impossible and the practically feasible. In the end, the idea of linear sorting being possible for all cases is an illusion. Linear sorting does exist, but it is only applicable in certain situations. The search for faster sorting algorithms should thus be about finding ways to use the input structure, improve the constants, or use hybrid models rather than trying to break the  $(n \log n)$  barrier for random data, especially when practical efficiency and real-world constraints must always be considered.



## QUICK SORT

### ACKNOWLEDGEMENT

It's really a pleasure for me to be able to sing praise of those scientists and educationalists that their authoritative contributions have been the basis of this study. It is my special pleasure to single out Donald Knuth and the Art of Computer Programming to which I owe my deep gratitude especially for the theoretical sorting part as well as the  $\Omega(n \log n)$  lower bound.

I could not but express my gratitude to Aho, Hopcroft, and Ullman for their contribution to the formal models of computation, and Cormen, Leiserson, Rivest, and Stein for making it clear in those cases in which sorting can be accomplished in linear time. Besides that, I am also giving my full support to and agreeing with the ideas of Sedgewick, Weiss, Levitin, Mehlhorn, Sanders, Tarjan, and Bentley, whose works not only open the way theoretically but also present the real side of the algorithmic performance.

The support of my mentors, the positive criticism of my colleagues, and the co-operation of my friends during the whole research process have been a great source of strength and inspiration for me. Moreover, I want also to recognize and appreciate the patience of my family and dearest ones who have been providing me with the time and focus necessary for the work. They have been their cheering up that has become my stepping stone to completing the study and gaining a deeper insight into algorithmic complexity.

### REFERENCES

- [1] Knuth, D. (1973). *The Art of Computer Programming, Volume 3: Sorting and Searching*. Addison-Wesley.
- [2] Aho, A. V., Hopcroft, J. E., & Ullman, J. D. (1974). *The Design and Analysis of Computer Algorithms*. Addison-Wesley.
- [3] Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. (2009). *Introduction to Algorithms* (3rd ed.). MIT Press.
- [4] Sedgewick, R., & Wayne, K. (2011). *Algorithms* (4th ed.). Addison-Wesley.
- [5] Weiss, M. A. (2012). *Data Structures and Algorithm Analysis in C++* (4th ed.). Pearson.
- [6] Levitin, A. (2012). *Introduction to the Design & Analysis of Algorithms* (3rd ed.). Pearson.

- 
- [7] Knuth, D. (1998). *The Art of Computer Programming, Volume 1: Fundamental Algorithms* (3rd ed.). Addison-Wesley.
  - [8] Mehlhorn, K., & Sanders, P. (2008). *Algorithms and Data Structures: The Basic Toolbox*. Springer.
  - [9] Tarjan, R. E. (1983). *Data Structures and Network Algorithms*. SIAM.
  - [10] Bentley, J. L. (1986). *Programming Pearls*. Addison-Wesley.