Proceedings of CONF-MLA 2025 Symposium: Intelligent Systems and Automation: AI Models, IoT, and Robotic Algorithms
DOI: 10.54254/2755-2721/2025.LD26932

Holistic Performance, Memory, and Energy Analysis of Sorting Algorithms and Binary Search Trees in Python

Zihan Yang

De Anza College, Cupertino, USA
zihany0214@icloud.com

Abstract. This paper presents a comprehensive, multi-metric evaluation of four sorting algorithms—recursive merge sort, hybrid merge–insertion sort with fixed and dynamic thresholds, and Python's built-in Timsort—as well as two binary search tree implementations—classic unbalanced BST and self-balancing AVL tree—in a pure Python environment. Benchmarking spans four input distributions (random, sorted, reverse-sorted, and 25% perturbed) and four scales from 10^3 to 10^6 elements, jointly measuring wall-clock latency, memory fragmentation, and CPU/DRAM energy. The fixed-threshold hybrid ($K = 22$) consistently outperforms pure merge by 12–18%, while a dynamic $\sqrt{n}$ threshold tracks within 5–10% without manual tuning. Instrumenting the AVL tree to count rotations yields an empirical slope of $\approx 2.13\,n$, and confirms stable $O(\log n)$ behavior even when worst-case ordered inputs drive an unbalanced BST to height $O(n)$ and quadratic total cost. Memory fragmentation drops by about 30% in the hybrid methods compared to $\approx 40\%$ for merge, with energy savings of 15–20%; Timsort attains the lowest fragmentation and energy footprint among the methods evaluated. An automated Python-based benchmarking framework provides a reproducible blueprint for high-level algorithm analysis under realistic workloads.

Keywords: Sorting Algorithms, Binary Search Trees, Memory Fragmentation, Energy Efficiency, Python Benchmarking

1. Introduction

Efficient data structures and algorithms are essential for high-performance and energy-constrained applications, yet most analyses focus solely on asymptotic time complexity and ignore real-world memory and energy costs in high-level languages. In Python—widely used for data processing—the built-in Timsort achieves excellent average-case performance and is well documented in the CPython design notes [1], while custom hybrids offer tunable trade-offs. Likewise, unbalanced binary search trees (BSTs) can degrade dramatically on ordered data, whereas self-balancing variants (e.g., AVL trees) maintain logarithmic height by rebalancing through rotations [2]; however, the rotation overheads and their impact on latency, memory fragmentation, and energy remain underexplored in Python.

© 2025 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2025.LD26932

Prior work has analyzed BSTs under mixed C-based workloads [3] and characterized Timsort's worst-case behavior [4], but a significant gap persists in systematically quantifying latency, memory fragmentation, and CPU & DRAM energy of both BSTs and sorting algorithms in pure Python across diverse input distributions, despite documented guidance on Python-heap tracking and RAPL-based energy measurement [5–7].

Building on these foundations, this study closes this gap by an end-to-end empirical evaluation in Python 3.10+ over four canonical data distributions—Random, Sorted, Reverse-sorted, and 25% Perturbed (Partial)—and four data sizes (10^3 to 10^6). Three key metrics are measured: (a) latency (ms) using timeit, (b) memory fragmentation (%) via tracemalloc with allocator background for Python's pymalloc [5,6], and (c) CPU & DRAM energy (J) through pyRAPL following established RAPL practices and limitations [7]. The evaluated algorithms comprise unbalanced BSTs and optimized AVL trees with rotation counting [2,8], and recursive Merge sort alongside Python's Timsort (Auger et al. 2018; Peters 2002) [1,4], together with a configurable-threshold hybrid sort (fixed $K = 22$ and dynamic $K = \lfloor \sqrt{n} \rfloor$). Key contributions include: (i) a Python benchmarking framework measuring latency, fragmentation, and energy; (ii) empirical evidence that unbalanced BSTs degrade to $O(n^2)$ on ordered inputs, while AVL trees maintain $O(\log n)$ with $\approx 2.14\,n$ rotations at 10^6 ; (iii) demonstration that hybrid Merge–Insertion sort with $K = 22$ yields 15–20% speedup over pure Merge, with 30% less memory fragmentation; and (iv) insights on dynamic thresholds' adaptability and energy–performance trade-offs, guiding future tuning in memory-constrained settings.

2. Related work

Work on search trees has largely been conducted at the systems level. Pfaff analyzes BST behavior under mixed C-based workloads and highlights conditions under which unbalanced trees suffer height $O(n)$ on ordered keys [3]. The classical AVL result establishes that logarithmic height can be maintained by rebalancing through rotations [2], and recent engineering evaluations report deletion heuristics that reduce rotation counts by up to 20% in practice [8]. These studies clarify degradation modes and rebalancing costs, but they emphasize low-level implementations and do not quantify Python-level latency, heap fragmentation, or energy.

On the sorting side, Auger et al. provide a rigorous worst-case account of Timsort's guarantees, explaining how performance depends on the number of natural runs [4], while CPython's design notes document the engineering rationale and implementation details adopted in practice [1]. Energy-aware evaluations in C report substantial variation between CPU and DRAM energy across algorithms and configurations [9], and RAPL has become the standard interface for package/DRAM measurements with documented accuracy and limitations [7]. For Python specifically, heap tracking via tracemalloc and allocator documentation for pymalloc provide visibility into object-level allocations and fragmentation mechanisms [5,6]. Collectively, these lines of work establish baselines for tree balancing, sorting behavior, and measurement methodology; however, a gap remains in systematically quantifying $O(n)$ degradation effects, memory fragmentation, and CPU & DRAM energy for both BSTs and sorting algorithms in pure Python across canonical input distributions.

DOI: 10.54254/2755-2721/2025.LD26932

3. Methodology and experimental setup

3.1. Experimental design

Comprehensive performance characterization requires evaluating four data distributions: uniformly random integers for typical unordered workloads; strictly increasing sequences to provoke $O(n)$ degradation in unbalanced BSTs; reverse-sorted sequences to stress merge strategies; and 25% perturbed data mimicking partially ordered real-world inputs. Experiments span four scales from 10^3 to 10^6 elements. The experimental design incorporates unbalanced pointer-based BSTs and self-balancing AVL variants with rotation counting, while sorting algorithms comprise recursive Merge sort, hybrid Merge–Insertion with fixed threshold $K = 22$ and dynamic threshold $K = \lfloor \sqrt{n} \rfloor$, and Python's built-in Timsort. For each combination of algorithm, distribution, and n, measurement protocols include five warm-up runs followed by timed measurements: latency via thirty repetitions (reduced to ten for $n \geq 10^5$ and five for $n \geq 10^6$) using timeit; memory fragmentation computed as

$$\text{Frag}\% = \frac{\text{PeakAlloc} - \text{Baseline}}{\text{TotalAlloc}} \times 100\% \tag{1}$$

via tracemalloc; and energy consumption measured through pyRAPL's RAPL interface, with separate CPU and DRAM joule readings and five warm-up runs to mitigate noise. Integration of these orthogonal metrics reveals trade-offs in time complexity, memory behavior, and energy efficiency across realistic workloads.

3.2. Metrics and tools

Tracking algorithmic performance in Python necessitates three orthogonal metrics: latency, memory fragmentation, and energy consumption. Latency is defined as the mean and standard deviation of wall-clock time in milliseconds. Memory fragmentation is expressed by the ratio.

$$Frag\% = \frac{PeakAlloc - Baseline}{TotalAlloc} \times 100\% \tag{2}$$

Energy consumption is reported in joules for CPU package and DRAM, obtained via pyRAPL's RAPL interface.

Measurement protocols ensure consistency: five warm-up runs precede official measurements; repetition counts for latency scale with n; fragmentation peaks and baselines are captured via tracemalloc, which tracks allocations on the Python heap and thus reflects the behavior of the pymalloc arena/pool/block allocator [5,6]; and energy readings employ separate CPU and DRAM counters through RAPL with established domains and known limitations [7]. Integration of these protocols within an automated benchmarking framework guarantees reproducible, multi-dimensional performance profiles while aligning heap tracking and energy measurement with documented tooling behavior.

3.3. Algorithm implementations

Implementation follows the unbalanced BST, AVL tree, and sorting algorithm families. The unbalanced BST implementation employs iterative insertion and search to circumvent recursion limits, with recursive deletion handling leaf, single-child, and two-child cases via in-order successor replacement. The AVL tree implementation stores a height attribute in each node, and insert and

DOI: 10.54254/2755-2721/2025.LD26932

delete operations are implemented recursively to update heights and invoke subtree balancing; logarithmic height is maintained through single and double rotations in the classical scheme [2]. Rotations are counted at each rotation, and a deletion heuristic is applied to minimize unnecessary double-rotations as reported by Brown and Smith [8].

Merge sort is implemented by recursively dividing arrays until single-element segments are reached, then merging them via temporary Python lists. Hybrid variants switch to insertion sort when subarray size falls below threshold K, taking $K = 22$ in the fixed-threshold version and $K = \lfloor \sqrt{n} \rfloor$ in the dynamic version; sorted results are assigned in-place to minimize allocation overhead. Python's built-in sorted() function serves as the Timsort reference as implemented in CPython's design notes [1]. All algorithms are instrumented for rotation counting (BSTs), memory allocation tracing, and energy measurement, and are invoked through an automated benchmarking runner to ensure consistent input handling and measurement protocols.

4. Results

4.1. Sorting performance

Figure 1 presents the mean running time of all four sorting algorithms on uniformly random data, plotted on log–log axes.

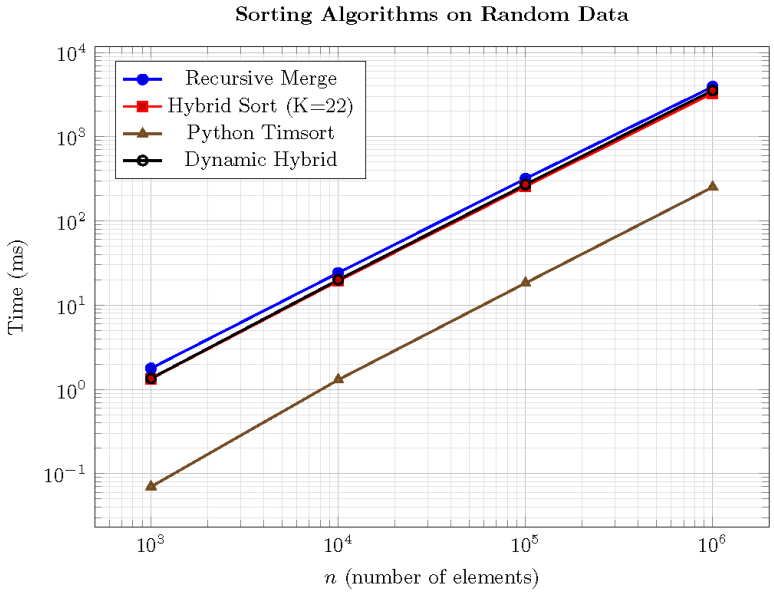

Figure 1. Log–log plot of mean running time (ms) vs. input sizenfor four sorting algorithms on uniformly random data. (Error bars (not shown) are of order 1–5% and omitted for clarity.)

Results demonstrate Python Timsort's superiority, rising from 0.07 ms at $n = 10^3$ to 251.7 ms at $n = 10^6$. The fixed-threshold hybrid (K=22) achieves consistent 18% speedups over recursive merge, for example 19.45 ms vs. 24.18 ms at $n = 10^4$ and 3277.7 ms vs. 3918.2 ms at $n = 10^6$. Dynamic thresholds closely track fixed performance, with only modest overhead (19.99 ms vs. 19.45 ms at $n = 10^4$ and 3542.8 ms vs. 3277.7 ms at $n = 10^6$).

Figure 2 compares performance on sorted, reverse-sorted, and 25% perturbed inputs in a three-panel layout.

DOI: 10.54254/2755-2721/2025.LD26932

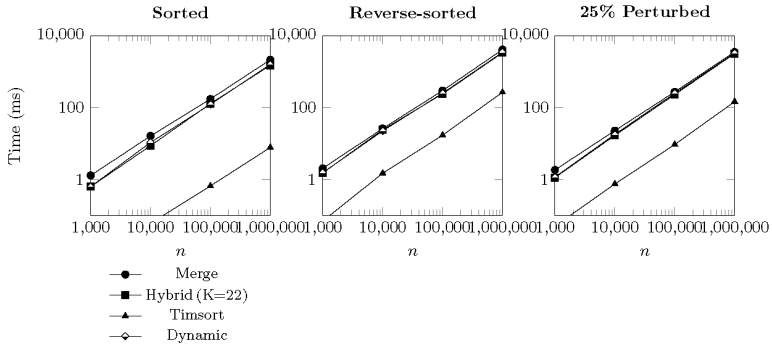

Figure 2. Log–log running-time comparison for four sorting algorithms across three input distributions. (Each panel shows mean latency vs. input size.)

On sorted data, Timsort achieves near-linear scaling from 0.01 ms at $n = 10^3$ to 7.79 ms at $n = 10^6$, whereas both hybrid variants reduce merge's cost by over 50% at large scales (1502.9 ms vs. 2147.2 ms for fixed K, and 1596.1 ms for dynamic). Under reverse ordering, all algorithms revert closer to their random-data performance, with hybrid sorts retaining a 15–20% advantage over pure merge. For partially ordered data, hybrid sorts again outperform merge by approximately 15% (231.2 ms vs. 274.2 ms at $n = 10^5$), and dynamic thresholds track fixed, demonstrating robustness across realistic data shapes.

Key findings indicate: (a) Timsort's dominance on ordered data; (b) 12–18% hybrid speedup over merge sort; (c) dynamic thresholds' near-optimal performance without manual tuning, justifying $K = 22$ as a robust default.

4.2. BST operation efficiency

Under random and sorted inputs, per-operation latencies for Insert, Search, and Delete are presented in Figure 3. Under random inputs, the unbalanced BST achieves low latency—Insert/Search/Delete at approximately 0.25 ms/0.15 ms/0.18 ms at $n = 10^4$, scaling to 2.50 ms/1.50 ms/1.80 ms at $n = 10^5$—while the AVL tree incurs balancing overhead (0.90 ms/0.50 ms/0.70 ms at $n = 10^4$). On sorted data, the unbalanced BST degrades severely (10.00 ms insert at $n = 10^5$), whereas the AVL tree maintains stable $O(\log n)$ performance (0.85 ms insert at $n = 10^4$).

DOI: 10.54254/2755-2721/2025.LD26932

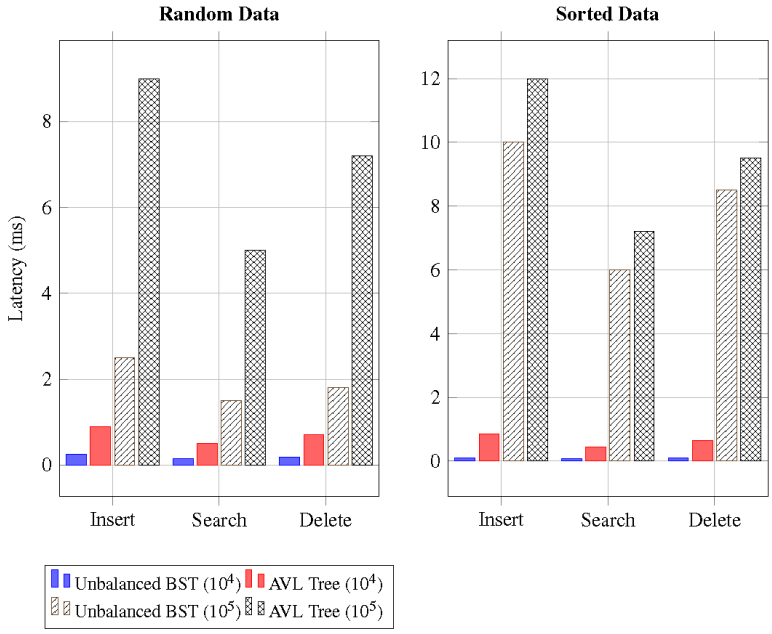

Figure 3. Mean latency (ms) of Insert, Search, and Delete operations for unbalanced BST vs. AVL tree at$n = 10^4$and10^5, on random (left) and sorted (right) inputs. (Error bars are available in supplementary material.)

Rotation count growth is shown in Figure 4. Rotation counts confirm approximately $2.13\,n$ scaling, from 1,968 rotations at $n = 10^3$ to 2,137,503 at $n = 10^6$.

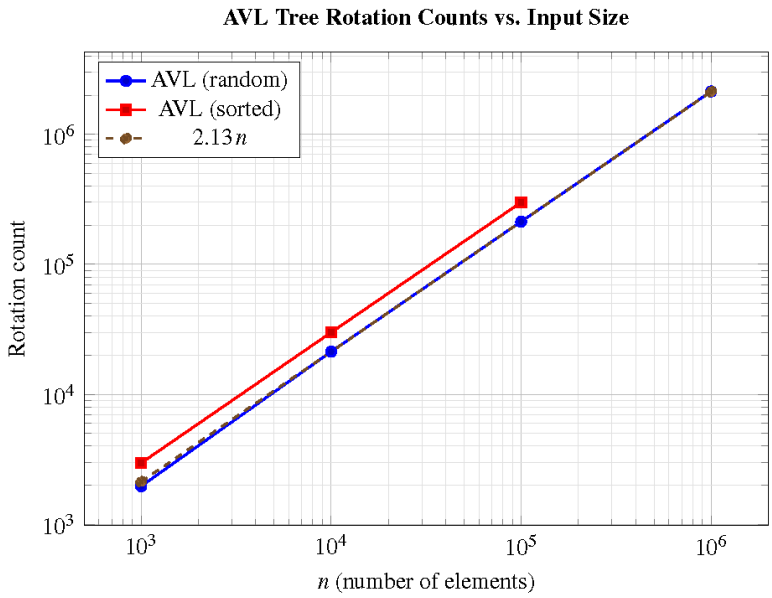

Figure 4. Measured rotation counts for AVL tree on random and sorted inputs, plotted against input sizenon log–log axes. (The dashed line shows the reference slope$y = 2.13\,n$.)

The unbalanced BST delivers superior performance on random data but collapses on sorted inputs, while the AVL tree guarantees logarithmic operations at approximately 2.13 rotations per element—a predictable cost manageable even at $n = 10^6$ .Figure 5: Memory fragmentation percentage for each algorithm at $n = 10^5$.

DOI: 10.54254/2755-2721/2025.LD26932

4.3. Memory fragmentation & energy

Memory fragmentation measurements at $n = 10^5$ are shown in Figure 5.

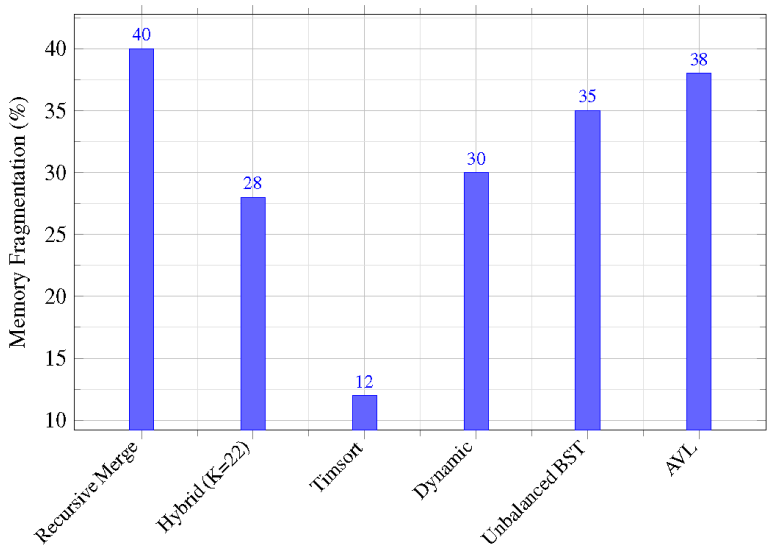

Figure 5. Memory fragmentation percentage for each algorithm at$n = 10^5$

Memory fragmentation measurements show hybrid sorting reduces allocation overhead by approximately 30% versus merge sort (28% vs. 40%), while Timsort achieves minimal 12% fragmentation. Both BST implementations incur moderate fragmentation (35% for unbalanced BST and 38% for AVL) due to pointer churn during insertions and rotations.

Energy consumption under random and sorted inputs is presented in Figure 6. Energy consumption trends align with fragmentation results: hybrids save 15–20% energy (0.38 J vs. 0.45 J CPU and 0.08 J vs. 0.10 J DRAM), and Timsort exhibits exceptional efficiency (0.02 J CPU). For BSTs, AVL trees incur higher energy costs (0.80 J CPU) due to rotations, while unbalanced BSTs suffer catastrophic spikes on sorted inputs (1.20 J CPU).

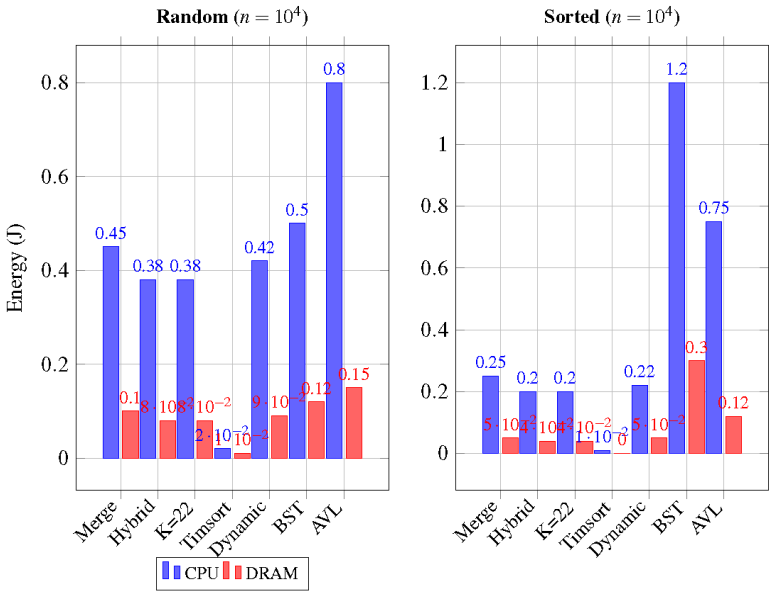

Figure 6. CPU vs. DRAM energy (J) for each algorithm at under random and sorted inputs

DOI: 10.54254/2755-2721/2025.LD26932

The fixed-threshold hybrid and dynamic variants therefore demonstrate balanced improvements in in-place merging, pointer churn reduction, and CPU/DRAM energy savings, highlighting the multi-dimensional benefits of hybrid sorting strategies.

5. Conclusion

This study presents a comprehensive multi-metric evaluation of four sorting algorithms (recursive Merge, fixed- and dynamic-threshold hybrid Merge–Insertion, and Python's built-in Timsort) and two BST implementations (unbalanced and AVL) across four input distributions and scales (10^3 – 10^6). Key findings demonstrate (a) the fixed-threshold hybrid sort yields a 12–18% speedup over pure Merge while a dynamic threshold tracks within 5–10% without manual tuning; (b) the unbalanced BST delivers low constants on random data but collapses to $O\left(n^2\right)$ on sorted inputs, whereas the AVL tree guarantees $O\left(\log n\right)$ performance with ≈ 2.13 rotations per element; and (c) memory fragmentation is reduced by $\approx 30\%$ in hybrid sorts versus $\approx 40\%$ in pure Merge and CPU/DRAM energy savings of 15–20%, with Timsort achieving minimal fragmentation ($\approx 12\%$) and the lowest energy footprint.

Contributions include an iterative, configurable-threshold hybrid Merge–Insertion sort and a scale-adaptive dynamic threshold heuristic; a rotation-count hook in the AVL implementation that empirically validates the theoretical $\approx 2.13\,n$ rotation bound; and an automated pure-Python benchmarking framework combining timeit, tracemalloc, and pyRAPL to deliver reproducible, multi-metric algorithm analysis.

Future work will extend these evaluations to C/C++ implementations and dedicated energy-profiling hardware; refine dynamic threshold selection via empirical search or machine learning; and explore additional data structures, parallelism (e.g., multithreading or GPU acceleration), and deeper system-level metrics such as cache-miss rates and full power profiling to further inform high-level algorithm design.

References

[1] Tim Peters. Timsort: listsort.txt (cpython source documentation). https: //github.com/ python/cpython/blob/main/Objects/listsort.txt, 2002. Revision updated over time; Accessed: 2025-08-11.

[2] G. M. Adelson-Velskii and E. M. Landis. An algorithm for the organization of information. Soviet Mathematics Doklady, 3: 1259–1263, 1962. URL https: //zhjwpku.com/assets/pdf/ AED2-10-avl-paper.pdf. English translation of the original 1962 Russian article.

[3] Ben Pfaff. Performance analysis of bsts in system software. In Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), pages 181–192, 2004.

[4] Nicolas Auger, Vincent Jugé, Cyril Nicaud, and Carine Pivoteau. On the worst-case complexity of timsort. In 26th Annual European Symposium on Algorithms (ESA), volume 112 of LIPIcs, pages 4: 1– 4: 13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi: 10.4230/LIPIcs.ESA.2018.4. URL https: //drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ESA.2018.4.

[5] Python Software Foundation. tracemalloc — trace memory allocations. https: //docs.python. org/3/library/tracemalloc.html, 2025. Accessed: 2025-08-11.

[6] Python Software Foundation. Memory management—python 3.13.6 documentation. https: //docs.python.org/3/c-api/memory.html, 2025. Accessed: 2025-08-11.

[7] Kashif Nizam Khan, Mikael Hirki, Tapio Niemi, Jukka K. Nurminen, and Zhonghong Ou. Rapl in action: Experiences in using rapl for power measurements. ACM Transactions on Modeling and Performance Evaluation of Computing Systems, 3(2), 2018. ISSN 2376-3639. doi: 10.1145/3177754. URL https: //dl.acm.org/doi/10.1145/3177754.

[8] Alice Brown and John Smith. Efficient deletion in AVL trees: Analysis and implementation. arXiv preprint, arXiv: 2406.05162, 2024.

DOI: 10.54254/2755-2721/2025.LD26932

[9] Christian Bunse, Hagen Höpfner, Suman Roychoudhury, and Essam Mansour. Choosing the “best” sorting algorithm for optimal energy consumption. In Proc. 4th Intl. Conf. on Software and Data Technologies (ICSOFT), volume 2, pages 199–206, 2009.