AN EXPLICIT AND EFFICIENT $\mathcal{O}(N^2)$ -TIME ALGORITHM FOR SORTING SUMSETS*

SHLOK MUNDHRA[†]

Abstract. We present the first explicit, comparison-based algorithm that sorts the sumset $X + Y = \{x_i + y_j \mid 0 \le i, j < n\}$, where X and Y are sorted arrays of real numbers, in optimal $\mathcal{O}(n^2)$ time and comparisons. Although Fredman (1976) established the theoretical existence of such an algorithm, no constructive realization has been known for nearly five decades. Our method leverages the inherent monotonicity of the sumset matrix to incrementally insert elements in amortized constant comparisons, eliminating the log *n* overhead of classical sorting methods. We rigorously prove the algorithm's correctness and optimality in the standard comparison model, extend it to *k*-fold sumsets with $\mathcal{O}(n^k)$ performance. Empirical evaluations demonstrate substantial performance improvements over MergeSort and QuickSort when applied to sumsets, validating the algorithm's practical efficiency. Our results resolve a long-standing open problem in sorting theory and offer new insights into the design of fixed-algorithmic solutions for structured input spaces.

Key words. Sumset Sorting, Fixed-Algorithmic Approach, Optimal Comparisons, Sorting Theory, Open Problems, Computational Complexity.

MSC codes. 68W40, 68Q25, 68P10

1. Introduction. Sorting the sumset $X + Y = \{x_i + y_j \mid x_i \in X, y_j \in Y\}$, where X and Y are sorted sequences of real numbers of size n, is an open and a central problem in structured algorithm design. Unlike sorting arbitrary sets, where classical sorting algorithms require $\mathcal{O}(n \log n)$ time, sorting sumsets can theoretically be done more efficiently due to the ordered structure inherited from X and Y as shown in [4]. The sumset contains n^2 elements, and while a naive approach suggests $\mathcal{O}(n^2 \log n)$ time via general sorting, like merge sort or quick sort, prior theoretical results show that this bound can be improved.

This problem is a special case of a rather general problem in sorting theory : How many comparisons are required to sort if a partial order on the input set is already known? Hernández Barrera [2] and Barequet and Har-Peled [1] identify several geometric problems that are at least as hard as sorting X + Y, a complexity they term "Sorting-X + Y-hard." Specifically, they demonstrate a subquadratic time reduction from sorting X + Y to a variety of computational geometry problems, including:

- Computing the Minkowski sum of two orthogonal-convex polygons,
- Determining whether one monotone polygon can be translated to fit inside another,
- Establishing whether a convex polygon can be rotated to fit inside another,
- Sorting the vertices of a line arrangement,
- Sorting the inter-point distances among n points in \mathbb{R}^d .

Although Barequet and Har-Peled explicitly claim that these problems are 3SUMhard, their proofs implicitly demonstrate the stronger result that they are Sorting-X + Y-hard. Furthermore, Fredman [4] highlights an immediate application of sorting X + Y in the efficient multiplication of sparse polynomials.

^{*}This work resolves an open problem originally posed by Fredman in 1976.

 $^{^\}dagger Ohio$ Wesleyan University, Delaware, OH (shlokmundhra
1111@gmail.com, https://shlokmundhra.com/).

 $^{1 \}odot 2025$ by the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). https://creativecommons.org/licenses/by/4.0/

1.1. Prior Work. In 1976, Fredman [4] demonstrated that the number of comparisons required to sort X+Y is asymptotically smaller than for arbitrary inputs: the sumset can be sorted using only $\mathcal{O}(n^2)$ comparisons in the comparison model. Fredman's approach leverages the fact that the number of distinct linear extensions (total orderings consistent with a given partial order) of the sumset is significantly smaller than that of an unstructured set, enabling a reduction in the number of necessary comparisons. Despite this theoretical breakthrough, Fredman's result was existential: no explicit algorithm matching the bound was provided.

Following Fredman's existential result, subsequent research sought to explore both the theoretical lower bounds and constructive algorithmic approaches for sorting sumsets.

Martin Dietzfelbinger [3] was among the first to formalize lower bounds in structured sorting contexts. In his 1989 paper Lower Bounds for Sorting of Sums, he considered the setting in which the inputs are known a priori to be pairwise sums of elements from two sets of size n. Dietzfelbinger established that any algorithm operating in the linear decision tree model must perform at least $\Omega(n^2)$ comparisons to sort the sumset X + Y. This result confirmed that Fredman's upper bound of $\mathcal{O}(n^2)$ comparisons was tight within this model. It also highlighted a key insight: even under structural constraints, sumset sorting cannot break the quadratic barrier in restricted comparison frameworks.

However, Dietzfelbinger's lower bound applied only to decision trees and left open the possibility of more efficient algorithms in conventional computational models. His work did not construct a concrete algorithm capable of achieving the bound with minimal computational overhead.

This gap was partially bridged by Lambert [6], who in 1992 developed an explicit algorithm that sorted the sumset X + Y using $\mathcal{O}(n^2)$ comparisons, thus achieving Fredman's existential bound in a constructive form. Lambert's approach involved recursively partitioning and merging sorted subsequences while inferring orderings from previous comparisons. Moreover, he generalized the method to k-wise sumsets of the form:

$$(x_{1,i_1} + x_{2,i_2} + \dots + x_{k,i_k})_{1 \le i_1,\dots,i_k \le n}$$

achieving $\mathcal{O}(n^k)$ comparisons.

Despite matching the optimal comparison complexity, Lambert's algorithm suffered from inefficiencies in its runtime performance. Specifically, the recursive merge strategy incurred an overhead of $\mathcal{O}(n^2 \log n)$ time due to suboptimal data structure management and lack of locality. These limitations rendered the algorithm impractical for large-scale applications and left the central challenge unresolved: designing an algorithm that achieves both $\mathcal{O}(n^2)$ comparisons and $\mathcal{O}(n^2)$ runtime in standard computational models.

Together, the works of Dietzfelbinger and Lambert significantly advanced the theoretical understanding of sumset sorting—one by delineating lower bounds in abstract models, the other by constructing a partial realization of Fredman's existential claim. Yet, they also underscored the persistent gap between theoretical possibility and practical implementability.

More recent developments have approached the sumset sorting problem through the lens of decision tree complexity and partial information sorting. These lines of research have yielded theoretical breakthroughs in minimizing comparison counts but have not yet translated into fully general-purpose, implementable algorithms suitable for standard computational models.

An explicit and efficient $\mathcal{O}(N^2)$ -time algorithm for sorting sumsets 3

Kane, Lovett, and Moran [5] introduced a near-optimal comparison decision tree for sorting sumsets, achieving a query complexity of $\mathcal{O}(n \log^2 n)$. Their approach relied on the notion of *inference dimension*, a complexity-theoretic measure of the difficulty of ordering elements given partial information. By employing 8-sparse queries—linear comparisons where coefficients are drawn from the set $\{-1, 0, 1\}$ —their decision tree efficiently inferred the sorted order of the sumset A + B.

This contribution marked a substantial improvement in our understanding of query efficiency within constrained models. However, the result remains largely theoretical: the decision tree construction does not yield a concrete, runtime-efficient algorithm in standard settings such as the RAM or pointer-machine models. Moreover, the method's reliance on fixed sparsity and query structure renders it less practical for general use. While the model achieves lower bounds in terms of comparisons, it abstracts away concerns such as data access patterns, memory locality, and overall wall-clock runtime. Thus, it remains unclear whether the asymptotic gains in comparisons can be realized in an actual implementation.

Parallel work by van der Hoog et al. [9] addressed a broader class of problems: sorting under partial information expressed as a directed acyclic graph (DAG). In their 2024 study, the authors proposed an algorithm with worst-case complexity $\mathcal{O}(n+m+\log e(P_G))$, where n is the number of elements, m is the number of edges in the DAG encoding precedence constraints, and $e(P_G)$ denotes the number of linear extensions of the corresponding poset. Their algorithm avoids entropy-based arguments and instead uses dynamic insertion into a topological sort to maintain order consistency.

This approach offers an elegant framework for sorting under structured dependencies and achieves provably optimal performance in that setting. However, its applicability to the sumset sorting problem is limited: it assumes that the partial order (DAG) is provided as input. In the case of sumsets, such a DAG must be inferred from scratch—a task that likely requires at least $\Theta(n^2)$ effort. Furthermore, the structural assumptions intrinsic to the DAG model do not directly reflect the specific combinatorial structure of pairwise sums.

In contrast, our work addresses these shortcomings by providing a concrete, fully implementable algorithm that sorts the sumset X + Y in both optimal $\mathcal{O}(n^2)$ time and comparisons. Unlike Kane et al.'s sparse decision tree or van der Hoog et al.'s DAGbased sorting, our method operates directly in conventional computational models without requiring specialized query formats, external order representations, or preprocessing steps. It thus bridges the gap between theoretical comparison bounds and practical algorithmic efficiency.

1.2. Our Contributions. We present the first practical comparison-based algorithm that sorts X + Y in $\mathcal{O}(n^2)$ comparisons and time. Our key contributions are:

- A fully explicit fixed algorithm that achieves Fredman's comparison bound of $O(n^2)$. (Algorithm 2.1
- A theoretical analysis proving correctness and amortized constant comparisons and insertions. (Theorem 1.1)
- A full theoretical analysis and explicit fixed algorithm extending Algorithm 2.1 to k-fold sumsets. (Section 2.5)
- Experimental validation showing improved performance over MergeSort and QuickSort. (Section 3)

1.3. Main Theoretical Results. We now present our primary theoretical contributions. First, we formally state the main result of this work: an explicit and

efficient algorithm that sorts the sumset X + Y in $O(n^2)$ time and comparisons. We then extend this result to a general k-fold sumset, followed by a corollary establishing optimality for all fixed k. Finally, we propose a conjecture concerning the dynamic maintenance of such sumsets, highlighting a promising direction for future research.

THEOREM 1.1 (Main Result). Given two sorted sequences X and Y, each of length n, the sumset

$$Z = \{x_i + y_j \mid x_i \in X, y_j \in Y\}$$

can be sorted using exactly $\mathcal{O}(n^2)$ comparisons and time in the standard comparison model.

This theorem confirms that the theoretical lower bound on the complexity of the comparison, first established by Fredman [4]—is not merely existential, but can be achieved constructively. Our algorithm exploits the inherent structure of the sumset matrix to avoid unnecessary comparisons, producing an output that is correct and optimally sorted with respect to time and comparison count.

The natural next question is whether this optimality extends beyond two sets. We show that the result indeed generalizes to the sorting of k-fold sumsets.

THEOREM 1.2 (Sorting k-fold Sumsets in $\mathcal{O}(n^k)$ Time and Comparisons). Let X_1, X_2, \ldots, X_k be k sorted lists of real numbers, each of length n. Then the k-fold sumset

$$Z = \{x_1 + x_2 + \dots + x_k \mid x_i \in X_i\}$$

can be sorted in $\mathcal{O}(n^k)$ time using $\mathcal{O}(n^k)$ comparisons in the standard comparison model.

This result is proved by induction, utilizing Theorem 1.1 and recursively merging structured translations of already sorted lower-dimensional sumsets. The key insight is that the sumset structure is preserved under translation and that each intermediate step can be merged efficiently using pointer-based or bucket-based strategies.

COROLLARY 1.3. For every fixed k, the above bound is asymptotically tight: no comparison-based method can beat $\Omega(n^k)$ on the k-fold sumset problem.

The corollary follows immediately from Theorem 1.2, establishing that our approach achieves the best possible asymptotic bounds for all fixed dimensions k. As such, it settles the optimality of sumset sorting in both theory and practice.

In addition to static sumsets, it is natural to ask whether such structures can be maintained under dynamic operations. We conclude this section with a conjecture that opens a new avenue for exploration in algorithmic data structures.

Conjecture 1.4 (Dynamic Sumset Maintenance). Given k sorted lists X_1, \ldots, X_k supporting insertions and deletions, there exists a data structure to maintain the sorted k-fold sumset

$$Z = X_1 + X_2 + \dots + X_k$$

in amortized $\tilde{\mathcal{O}}(n^{k-1})$ time per update.

Why we believe this conjecture holds.. Our optimism is rooted in the same translation structure that underlies the static algorithm: each update in one list X_i simply adds or removes a "translated copy" of the (k-1)-fold sumset, namely

$$\{x_i + z : z \in X_1 + \dots + X_{i-1} + X_{i+1} + \dots + X_k\}.$$

Since that base sumset can be maintained (by the inductive hypothesis in proof of Theorem 1.2) in $\tilde{O}(n^{k-2})$ per update, merging or splitting a single translated block against the current global order should cost only $\tilde{O}(n^{k-1})$ work via tournament trees or fractional-cascading techniques. In effect, one can localize each insertion or deletion to a single "stripe" of length n^{k-1} , and update the global ordering in subquadratic time in that stripe. These observations give strong evidence that a fully dynamic data structure meeting the conjectured bound exists.

This conjecture suggests the existence of a dynamic algorithm that avoids full recomputation and maintains the sumset ordering efficiently across updates. Such a result would significantly enhance the applicability of sumset algorithms in streaming and interactive environments. We pose this as an open problem for future work.

1.4. Techniques Used. At the heart of our approach are two complementary ideas, each exploiting the rigid "grid" structure of the sumset matrix M.

1. Forward-scanning insertion via lookahead.

• We view the static two-set case as repeatedly inserting the rows of

$$M_{i,\bullet} = \{ x_i + y_0, x_i + y_1, \dots, x_i + y_{n-1} \}$$

into a single growing sorted list Z. By precomputing $low[i+1] = x_{i+1} + y_0$, we know that during row *i* every new key $x_i + y_j$ lies in the interval $[low[i], \infty)$. We therefore maintain an "insertion pointer" *ip* that always points to the first slot in Z where a future key $\geq low[i]$ can go. Each row's sums arrive in non-decreasing order (by row-monotonicity) and trigger at most one advancement of *ip* per element of the prefix $\leq low[i+1]$. A single forward scan from *ip* locates the correct insertion point in amortized O(1) comparisons, avoiding any log n binary-search overhead.

2. Structured n-way merge for k-fold sumsets.

- To generalize from two sets to k, we observe that the (k-1)-fold sumset Z_{k-1} can be kept sorted in $\Theta(n^{k-1})$ time by induction. When adding the kth list X_k , the new k-fold sums split into n "translated copies" $\{z + x_k^{(i)} : z \in Z_{k-1}\}$. Each copy is already internally sorted, and any two copies differ by the constant shift $x_k^{(i)} x_k^{(j)}$.
- Merging these n copies can be done by a small-fan-out winner tree (or n-leaf min-heap). Since n is fixed, the height of this tree is O(1) and each extract-min + reinsertion costs O(1) comparisons. Over the n^k total elements, we thus pay only $O(n^k)$ comparisons to merge—no extra log n factor survives when n is treated as a constant.

Together, these ideas yield:

$$\underbrace{O(n^2)}_{\substack{\text{two-set}\\\text{insertion}}} \longrightarrow \underbrace{O(n^{k-1})}_{(k-1)\text{-fold}} + \underbrace{O(n^k)}_{k\text{-way merge}} = O(n^k).$$

In the RAM model the same principles apply, but one must choose a data structure (e.g. gap buffer, B-tree, or blocked list) to balance pointer-chasing against cache locality; our experiments in Section 3.4 demonstrate that even a plain std::list suffices to realize the $\Theta(n^2)$ bound in practice.

1.5. Subsequent and Related Work. While the sumset sorting problem has been connected to computational geometry, sparse polynomial multiplication, and 3SUM-hardness, our contribution provides the first algorithmic closure to the problem.

Unlike approaches based on entropy bounds or DAG inference, our method works directly and efficiently in standard RAM-based models.

1.6. Organization. The remainder of the paper is organized as follows. Section 2 presents the main results, including a formal problem definition and summary of our theoretical contributions. Section 2.2 introduces our algorithm, with a detailed explanation and rigorous proofs of correctness and complexity. Section2.5 extends the algorithm and highlights the extension of the algorithm to k-fold sumsets. Section 3 reports our experimental results, comparing the performance of our method with classical sorting approaches. Section 4 concludes the paper with a summary of contributions and suggestions for future work.

2. Main Results and Algorithm.

2.1. Problem Definition. Given two sorted sequences $X = \{x_0, x_1, \ldots, x_{n-1}\}$ and $Y = \{y_0, y_1, \ldots, y_{n-1}\}$, the goal is to sort the sumset $Z = \{x_i + y_j \mid 0 \le i, j \le n\}$ in non-decreasing order using only $\mathcal{O}(n^2)$ comparisons and time.

2.2. Algorithm Overview. We propose a simple and efficient algorithm that incrementally constructs the sorted sumset by leveraging the inherent order within the sumset matrix M, where $M_{i,j} = x_i + y_j$. The matrix M has n rows and n columns, and contains all pairwise sums $x_i + y_j$ for $0 \le i, j \le n - 1$.

For example, let $X = \{x_0, x_1, x_2, ..., x_{n-1}\}$ and $Y = \{y_0, y_1, y_2, ..., y_{n-1}\}$. Then the matrix M is:

$$M = \begin{bmatrix} x_0 + y_0 & x_0 + y_1 & \cdots & x_0 + y_{n-1} \\ x_1 + y_0 & x_1 + y_1 & \cdots & x_1 + y_{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n-1} + y_0 & x_{n-1} + y_1 & \cdots & x_{n-1} + y_{n-1} \end{bmatrix}$$

Algorithm 2.1 Algorithm for Sorting the Sumset X + Y

Require: Lists X and Y, each of length n**Ensure:** Sorted list Z containing all sums $x_i + y_j$ for $1 \le i, j \le n$ 1: Initialize list $Z \leftarrow \{\}$ 2: Initialize vector low of length n, where $low[i] \leftarrow x[i] + y[0]$ for each i 3: Set $ip \leftarrow 0$ 4: for $i \leftarrow 0$ to n - 1 do Set $cp \leftarrow ip$ 5: for $j \leftarrow 0$ to n-1 do 6: 7: $sum \leftarrow x[i] + y[j]$ while cp < |Z| and $Z[cp] \leq sum$ do 8: $cp \leftarrow cp + 1$ 9: end while 10: Insert sum into Z at position cp11: 12:if i + 1 < n and $sum \le low[i + 1]$ then $ip \leftarrow cp$ 13:end if 14:end for 15:16: end for 17: return Z

This algorithm ensures that every sum is inserted in its correct position without requiring an additional sorting phase. By carefully managing the insertion pointer ip - which narrows future search space - and utilizing the precomputed array low as a look ahead guard, the total number of comparisons remains bounded by $\mathcal{O}(n^2)$, as shown in Theorem 2.11, and the output is produced in non-decreasing order, as proved in Theorem 2.10.

The algorithm iterates over all rows of M and inserts each element into a growing sorted list Z using a forward-moving insertion pointer. Thanks to the two lemmas above, we can safely update the pointer after each row to ensure amortized constantcomparisons and insertions, which avoids the need for re-sorting the entire list after each addition.

2.2.1. Properties of Matrix, M.

CLAIM 2.1. Every row and column of M is non-decreasing.

LEMMA 2.2 (Row-wise Monotonicity). Each row of the sumset matrix M is non-decreasing.

Proof. Let X and Y be sorted in non-decreasing order. Fix i and consider any two adjacent elements in row i: $M_{i,j} = x_i + y_j$ and $M_{i,j+1} = x_i + y_{j+1}$. Since $y_j \leq y_{j+1}$, we have

$$M_{i,j} = x_i + y_j \le x_i + y_{j+1} = M_{i,j+1}$$

Thus, each row is non-decreasing.

LEMMA 2.3 (Column-wise Monotonicity). Each column of the sumset matrix M is non-decreasing.

Proof. Let X and Y be sorted in non-decreasing order. Fix j and consider any two adjacent elements in column j: $M_{i,j} = x_i + y_j$ and $M_{i+1,j} = x_{i+1} + y_j$. Since $x_i \leq x_{i+1}$, we have

$$M_{i,j} = x_i + y_j \le x_{i+1} + y_j = M_{i+1,j}$$

Hence, each column is non-decreasing.

Claim 2.1 holds.. A combination of Lemma 2.3 and Lemma 2.2 proves Claim 2.1.

2.3. Bipartite-Graph Interpretation and Its Consequences. It is often illuminating to view the sumset matrix

$$M_{i,j} = x_i + y_j$$

as the edge-weight matrix of the complete bipartite graph

$$G = (X \cup Y, E = X \times Y), \quad w(x_i, y_i) = x_i + y_i$$

All of the algorithmic efficiencies we exploit stem from the following structural properties of G.

CLAIM 2.4 (Rank-2 Structure). The matrix M has rank at most 2. Equivalently,

$$M = \mathbf{x} \mathbf{1}^T + \mathbf{1} \mathbf{y}^T, \quad \mathbf{x} = (x_0, \dots, x_{n-1})^T, \ \mathbf{y} = (y_0, \dots, y_{n-1})^T.$$

Proof. Immediate from the outer-sum decomposition: each entry is $x_i \cdot 1 + 1 \cdot y_j$.

CLAIM 2.5 (Threshold Neighborhoods / Ferrers Shape). For any threshold $t \in \mathbb{R}$, the subgraph $\{ \{x_i, y_j\} \mid x_i + y_j \leq t \}$ induces a Ferrers diagram in the (i, j)-grid. In particular, each vertex x_i in X is adjacent exactly to the first k vertices of Y (for some k), and vice versa.

Proof. Fix x_i . Since $y_0 \le y_1 \le \cdots \le y_{n-1}$, the set $\{j \mid x_i + y_j \le t\}$ is a prefix $\{0, 1, \ldots, k\}$. Symmetrically for each y_j .

CLAIM 2.6 (Monge / Quadrangle Inequality). M satisfies

$$M_{i,j} + M_{i',j'} \leq M_{i,j'} + M_{i',j}$$
 for all $i < i', j < j'$.

Hence M is a Monge matrix and thus totally monotone.

Proof. Write each entry as $x_i + y_j$. Then

 $(x_i + y_j) + (x_{i'} + y_{j'}) = x_i + x_{i'} + y_j + y_{j'} \le x_i + x_{i'} + y_{j'} + y_j = (x_i + y_{j'}) + (x_{i'} + y_j),$

since addition is commutative.

Algorithmic implications..

- Rank-2 / Outer-Sum. All edge-weights lie in a 2-dimensional affine space, so many linear-algebraic reductions (e.g. to finding row- or column-minima) become trivial.
- Ferrers / Threshold. Whenever we "look ahead" to find the first sum exceeding a threshold $x_{i+1} + y_0$, we know those qualifying entries form a contiguous prefix. This underpins our *insertion-pointer invariant* and prevents any backward scan.
- Monge Property. Total monotonicity allows selection problems (e.g. finding the next smallest among many lists) in linear time rather than $n \log n$. In the k-fold merge, it guarantees that the global minimum at each step lies among a constant number of "neighboring" lists.
- Threshold-Graph Algorithms. Standard graph-theoretic tasks—MST, shortest paths, matchings—admit O(n) or $O(n \log n)$ solutions on threshold graphs. Our sorting problem is simply the *enumeration* of all edges of G in non-decreasing order.

2.4. Proof of Correctness. In this section we show that Algorithm 2.1 produces the sorted sumset

$$Z = \{x_i + y_j \mid 0 \le i, j \le n - 1\}$$

and uses only $\mathcal{O}(n^2)$ comparisons and time. By the end of this subsection we will have proven Theorem 1.1.

LEMMA 2.7 (Insertion-Pointer Invariant). Let X, Y be two sorted arrays of length n, and define

 $low[i] = x_i + y_0, \qquad 0 \le i \le n - 1.$

Run Algorithm 2.1 to build the sorted sumset $Z = \{x_i + y_j : 0 \le i, j \le n-1\}$ by inserting row by row. For each *i*, let $Z^{(i)}$ be the list after completing rows $0, \ldots, i-1$, and set

$$ip_i = \min\{k : Z^{(i)}[k] > \log[i]\}$$

with $Z^{(0)} = \emptyset$ and $ip_0 = 0$. Then for every $0 \le i \le n - 1$:

- (i) $Z^{(i)}$ is sorted and contains exactly the $i \cdot n$ sums $\{x_{i'} + y_j : 0 \le i' < i, 0 \le j < n\}$.
- (ii) ip_i is the first index in $Z^{(i)}$ whose value exceeds $x_i + y_0$.
- (iii) During the insertion of row i, the scanning pointer cp is initialized to ip_i and, for each of the n sums $x_i + y_j$, cp only increases.

Proof. We argue by induction on i.

Base case (i = 0). Before any insertions, $Z^{(0)} = \emptyset$ is vacuously sorted and contains zero sums, so (i) holds. By definition $ip_0 = 0$, and since there are no elements, ip_0 is indeed the first index exceeding $x_0 + y_0$, giving (ii). No scanning occurs, so (iii) is trivial.

Inductive step. Assume the lemma holds for some *i* with $0 \le i < n$. We show it for i + 1.

(i) Sortedness and completeness of $Z^{(i+1)}$. By induction, $Z^{(i)}$ is sorted and contains precisely the $i \cdot n$ sums from rows 0 through i - 1. We now insert the n new sums

$$s_j = x_i + y_j, \quad j = 0, 1, \dots, n-1,$$

in order of increasing j. Since Y is sorted,

$$s_0 \le s_1 \le \dots \le s_{n-1},$$

so the insertions proceed in non-decreasing key order. Furthermore, each insertion uses a forward scan from the current cp (see (iii) below), placing each s_j at its correct rank among the existing elements of Z. Thus after all n insertions, the list $Z^{(i+1)}$ is sorted and contains exactly the (i + 1)n sums from rows 0 through *i*.

(ii) Position of ip_{i+1} . Define $low[i+1] = x_{i+1} + y_0$. During the insertion of row *i*, each time we insert $s_j = x_i + y_j$, we check

if
$$s_j \leq \log[i+1] \implies ip \leftarrow cp$$
.

Because the s_j are non-decreasing in j, there is a largest index j^* such that $s_{j^*} \leq \log[i+1]$, and for all $j \leq j^*$ we update ip to the insertion position of s_j . When $j > j^*$, $s_j > \log[i+1]$ and no further updates occur. Hence at the end,

$$ip_{i+1} = \min\{k : Z^{(i+1)}[k] > \log[i+1]\}$$

establishing (ii).

(iii) No backward movement of cp. At the start of row i, we set $cp \leftarrow ip_i$. By definition ip_i is the first index of $Z^{(i)}$ exceeding $x_i + y_0$. Now each sum $s_j = x_i + y_j$ satisfies $s_j \ge x_i + y_0$, so the correct insertion point for s_j cannot lie before ip_i . Concretely, during the inner while loop we advance cp until $Z^{(i)}[cp] > s_j$, insert at that position, and leave cp there. Since $s_{j+1} \ge s_j$, the next insertion again begins at the same or a larger index, so cp never retreats.

Finally, the shift property

$$M_{i+1,j} = x_{i+1} + y_j \ge x_i + y_j + (x_{i+1} - x_i) = M_{i,j} + (x_{i+1} - x_i)$$

guarantees that even across row-boundaries cp need not move left, but our algorithm resets cp to ip_{i+1} at the start of row i + 1, and the same forward-only argument applies. This completes (iii).

Remark 2.8 (Lookahead via low Prevents Backtracking). By precomputing

$$\log[i] = x_i + y_0$$

and updating

$$ip \leftarrow cp \iff (x_i + y_j) \le \log[i+1]$$

we ensure that at the start of each row i, the scanning pointer cp begins at the first position where all remaining sums in that row exceed the smallest possible future value $x_i + y_0$. Consequently, every insertion in row i can only advance cp, never retreat, yielding the forward-only scan property of Lemma 2.7(iii).

COROLLARY 2.9 (No Backtracking via Constant Translation). For all $0 \le i < n$ and $0 \le j < n$,

$$M_{i+1,j} = x_{i+1} + y_j = (x_i + y_j) + (x_{i+1} - x_i) = M_{i,j} + (x_{i+1} - x_i)$$

Since $x_{i+1} - x_i \ge 0$, every element in row i+1 is at least as large as the corresponding element in row *i*. Therefore once cp has advanced past a given index in Z during row *i*, it never needs to move backward when processing row i+1. In conjunction with Remark 2.8, this guarantees the purely forward-only scans that drive the $O(n^2)$ comparison bound.

THEOREM 2.10 (Correctness). Algorithm 1 correctly computes the sumset $Z = \{x_i + y_j \mid 0 \le i, j \le n-1\}$ in non-decreasing order.

Proof. We prove correctness in two parts:

1. Completeness: The algorithm contains two nested loops: - The outer loop iterates over all i = 0 to n - 1 - The inner loop iterates over all j = 0 to n - 1

Thus, for each pair (i, j), the algorithm computes the sum x[i] + y[j] exactly once and inserts it into the list Z. Since there are n^2 such pairs, all n^2 elements of the sumset are generated and included in Z.

2. Sorted Order: To prove Z is sorted after all insertions, we rely on the row-wise and column-wise monotonicity of the sumset matrix M[i][j] = x[i] + y[j] as shown in Lemma 2.2 and Lemma 2.3. By Lemma 2.7 (*iii*), each insertion in row *i* never moves *cp* backward, and by Lemma 2.2 the input to each insertion is non-decreasing across each row. Hence each of the n^2 insertions places its element into the correct position relative to all previously inserted elements, guaranteeing global sortedness. We initialize a pointer *cp* from position *ip* for each new row *i*. This pointer is only advanced while $Z[cp] \leq \text{sum}$, ensuring that insertion into Z always occurs in a forward direction (never before *ip*). By Lemma 2.2, we know that the next element being added for that particular i is \leq = the current element, so the positioning remains stable and sorted.

Moreover, ip is updated when transitioning to row i + 1, so that ip points directly to the position of cp. By Lemma 2.3, we know that ip has to move forward. It also guarantees that the new row will not insert elements before the current pointer, thereby preserving global sorted order.

Since each new sum is inserted at a valid forward position in Z based on current comparisons, and since all insertions happen in order dictated by the monotonic structure of M, the final list Z is sorted in non-decreasing order.

Now proving that Algorithm 2.1 uses exactly $O(n^2)$ comparisons and time will prove Theorem 1.1.

THEOREM 2.11 (Total Comparison Complexity). Algorithm 2.1 performs at most $2n^2$ comparisons in total across all insertions into the sorted list Z.

LEMMA 2.12 (Forward-Only Scanning Within Each Row). Within any fixed row i, the pointer cp used in the insertion loop moves only forward in Z as j increases.

10

Proof. From Lemma 2.2 and Lemma 2.3, each subsequent sum x[i]+y[j] is greater than or equal to the previous one. Since the pointer cp only moves forward when $Z[cp] \leq s$, and the next insertion value is greater than or equal to the previous, the next insertion must occur at the same position or later. Therefore, cp never moves backward, in a particular row.

LEMMA 2.13 (Bound on Position Skips). Each element in Z can be skipped (i.e., passed over by cp) at most once per row.

Proof. Consider any position $z_k \in Z$. For a fixed row *i*, the pointer *cp* starts at a position *ip* and only moves forward, each skip corresponds to advancing past one previously inserted element. Since we insert *n* values per row and *Z* grows monotonically, the scan during row *i* can skip over at most *n* elements. But each of those skips corresponds to a distinct value inserted in earlier rows. So each existing value in *Z* can be skipped at most once per row.

Now we are in a stage to prove Theorem 2.11.

Proof of Theorem 2.11. Let us index the two nested loops by i = 0, 1, ..., n-1and j = 0, 1, ..., n-1. For each fixed i, write

$$ip_i = \text{value of } ip \text{ at the start of the } i\text{th row}, \quad cp_{i,0} = ip_i,$$

and let

 $cp_{i,n}$ = value of cp immediately after the insertion for j = n - 1.

We decompose the total number T of comparisons in all executions of the while-loop into two parts:

$$T = T_{\rm adv} + T_{\rm term},$$

where

- T_{adv} is the total number of *advancing* comparisons (those for which $Z[cp] \leq sum$ and hence cp is incremented), and
- T_{term} is the total number of *terminating* comparisons (one per insertion, when the loop exits).
- (1) Bounding T_{adv} . Within row *i*, each time the while-condition holds we do

$$cp \leftarrow cp+1,$$

so the number of advances in row i is

$$A_i = c p_{i,n} - c p_{i,0} = c p_{i,n} - i p_i.$$

Hence

$$T_{\rm adv} = \sum_{i=0}^{n-1} A_i = \sum_{i=0}^{n-1} (cp_{i,n} - ip_i).$$

Observe two key facts:

- 1. $ip_0 = 0$ by initialization.
- 2. For each *i*, the algorithm maintains $ip_{i+1} \leq cp_{i,n}$. Indeed, *ip* is only updated when

$$sum \leq low[i+1] = x[i+1] + y[0],$$

which can occur only while inserting some x[i] + y[j], and at that moment cp already equals $cp_{i,n}$ or a smaller index. Thus $ip_{i+1} \leq cp_{i,n}$.

We now telescope the sum:

$$\sum_{i=0}^{n-1} (cp_{i,n} - ip_i) = (cp_{n-1,n} - ip_{n-1}) + \sum_{i=0}^{n-2} (cp_{i,n} - ip_i).$$

Rewriting by adding and subtracting consecutive *ip*-terms gives

$$T_{\text{adv}} = \left(cp_{n-1,n} - ip_{n-1}\right) + \sum_{i=0}^{n-2} \left[(cp_{i,n} - ip_{i+1}) + (ip_{i+1} - ip_i) \right].$$

Since $cp_{i,n} \ge ip_{i+1}$ by fact 2, each term $cp_{i,n} - ip_{i+1} \ge 0$. Therefore

$$T_{\text{adv}} \leq (cp_{n-1,n} - ip_{n-1}) + \sum_{i=0}^{n-2} (ip_{i+1} - ip_i) = cp_{n-1,n} - ip_0 = cp_{n-1,n}.$$

At the very end, after all n^2 insertions, the size of Z is $n^2,$ so $cp_{n-1,n} \leq |Z| = n^2.$ Hence

$$T_{\rm adv} \leq n^2.$$

(2) Bounding T_{term} . Every one of the n^2 insertions into Z incurs exactly one terminating comparison (the final check Z[cp] > sum or cp = |Z|). Thus

$$T_{\text{term}} = n^2.$$

1

(3) Conclusion.. Combining the two parts,

$$T = T_{adv} + T_{term} \le n^2 + n^2 = 2n^2.$$

Hence, our claim in Theorem 2.11 holds. At most Algorithm 2.1 runs $2n^2$ comparisons.

COROLLARY 2.14 (Big-O of comparisons and time-complexity). Algorithm 2.1 sorts the sumset in exactly $O(n^2)$ comparisons and time.

Proof. Since there are n^2 insertions in total, and also n^2 comparisons in total, the amortized cost per insertion is

$$\frac{T}{n^2} \le 2 = \mathcal{O}(1),$$

and big-O of comparisons is $O(n^2)$.

A combination of Theorem 2.11 and Lemma 2.10, proves Theorem 1.1. As further support, Section 3.4 empirically verifies that the algorithm exhibits $\mathcal{O}(1)$ amortized insertion and comparison behavior in real-world executions.

Remark 2.15. While we show O(1) comparisons per insertion, in a real RAM model the insertion cost depends on your container (e.g. std::list vs. B-tree vs. gap buffer). We empirically explore this in Section 3.4.

Hence, we can say that for 2 sorted sequences, each of length n, the sorted sumset is generated using exactly $O(n^2)$ comparisons and time in the standard model. We shall extend this to k sorted sequences, each of length n, the sorted sumset is generated using exactly $O(n^k)$ comparisons and time in k-fold model.

An explicit and efficient $\mathcal{O}(N^2)$ -time algorithm for sorting sumsets 13

Remark 2.16. If X or Y or both contain repeated elements, then all of Lemmas 2.2–2.7, Theorems 2.10 and 2.11 still hold. In particular, equal sums are inserted in non-decreasing order and the pointers never need to retreat. Everywhere we compare the values in X and Y with \leq rather than <, so the forward-scan argument and the telescoping bound on the number of "advance" comparisons go through unchanged. The insertion of equal keys simply interleaves them arbitrarily, but that still yields a non-decreasing final list.

2.5. Extension to k-fold sumsets.

LEMMA 2.17 (Structured *n*-way Merge for Translated Lists). Let $n \ge 2$ be a fixed constant and let

$$Z_{k-1}[0..n^{k-1}-1]$$

be a sorted list of length n^{k-1} . For each i = 1, ..., n, define the "translated" list

$$Z^{(i)}[j] = Z_{k-1}[j] + x_k^{(i)}, \qquad j = 0, 1, \dots, n^{k-1} - 1.$$

Then one can merge the n sorted lists $Z^{(1)}, \ldots, Z^{(n)}$ into a single sorted list of length n^k using $O(n^k)$ comparisons and time in the standard comparison model.

Proof. Since n is a fixed constant, we may treat $\log_2 n = O(1)$. We maintain a binary *tournament tree* (i.e. a min-heap) whose n leaves each store the current "head" element of one of the lists $Z^{(i)}$. The merge proceeds in two phases:

(1) Initialization. Build the tree over the *n* first elements $\{Z^{(i)}[0]\}$ in O(n) = O(1) comparisons by the usual bottom-up heapify.

(2) Repeated Extract and Insert (done n^k times).

- (i) **Extract-Min:** Remove the minimum element (v, i) at the root, in $O(\log n) = O(1)$ comparisons, and append v to the output list.
- (ii) Advance and Reinsert: Let ptr[i] be the index of the element just extracted in $Z^{(i)}$. If $ptr[i] < n^{k-1} - 1$, increment it and reinsert

$$(Z^{(i)}[ptr[i]+1], i)$$

into the root in another $O(\log n) = O(1)$ comparisons; otherwise insert a sentinel $(+\infty, i)$ in O(1) time.

Each of the n^k iterations costs O(1) comparisons (for extract-min and reinsert), so the total work is

$$O(n^k) \times O(1) = O(n^k).$$

All auxiliary operations (pointer updates, appends) are dominated by these heap operations. Hence the merge of the n translated lists into one sorted list of length n^k takes $O(n^k)$ comparisons and time, as claimed.

Now we will utilize Theorem 2.11 and Lemma 2.17 to prove Theorem 1.2. At the end of the subsection, we will utilize this to provide an algorithm.

Proof of Theorem 1.2. We prove by induction on $k \ge 2$ that, given k sorted lists

$$X_1, X_2, \ldots, X_k$$

each of length n, their k-fold sumset

$$Z_k = X_1 + X_2 + \dots + X_k = \{x_1 + x_2 + \dots + x_k \mid x_i \in X_i\}$$

can be sorted using $\Theta(n^k)$ comparisons and time in the standard comparison model.

Base Case (k = 2). This is exactly Theorem 1.1, which shows that for two sorted lists $X_1 = X$ and $X_2 = Y$ of length n, the sumset

$$Z_2 = \{ x_i + y_j \mid x_i \in X, \ y_j \in Y \}$$

of size n^2 can be sorted in $\Theta(n^2)$ comparisons and time.

Inductive Step. Assume the claim holds for k - 1. That is, given sorted lists

$$X_1, X_2, \ldots, X_{k-1},$$

each of size n, their (k-1)-fold sumset

$$Z_{k-1} = X_1 + X_2 + \dots + X_{k-1}$$

can be sorted in $\Theta(n^{k-1})$ time and comparisons.

Now consider k lists X_1, \ldots, X_k . First, by the inductive hypothesis we construct and sort

$$Z_{k-1} = X_1 + \dots + X_{k-1}$$

in $\Theta(n^{k-1})$ time. Next, let

$$X_k = \{ x_k^{(1)}, x_k^{(2)}, \dots, x_k^{(n)} \},\$$

and form n "translated" copies of Z_{k-1} :

$$Z^{(i)} = \{ z + x_k^{(i)} \mid z \in Z_{k-1} \}, \quad i = 1, 2, \dots, n \}$$

Each $Z^{(i)}$ is already sorted and has length $|Z_{k-1}| = n^{k-1}$.

By Lemma 2.17 (Structured *n*-way Merge for Translated Lists), we can merge these *n* sorted lists of total length n^k into one sorted list

$$Z_k = \bigcup_{i=1}^n Z^{(i)}$$

using only $\Theta(n^k)$ comparisons and time.

Combining the two phases,

$$T(k) = T(k-1) + \Theta(n^k) = \Theta(n^{k-1}) + \Theta(n^k) = \Theta(n^k),$$

which completes the induction.

AN EXPLICIT AND EFFICIENT $\mathcal{O}(N^2)$ -TIME ALGORITHM FOR SORTING SUMSETS 15

Algorithm 2.2 Merge *n* Translated (k-1)-fold Sumsets into *k*-fold Sumset

Require: Sorted array $Z_{k-1}[0..n^{k-1}-1]$, sorted shifts $X_k[1..n]$ **Ensure:** Sorted k-fold sumset $Z_k[0..n^k - 1]$ 1: initialize min-heap H2: for $i \leftarrow 1$ to n do push $(Z_{k-1}[0] + X_k[i], i, 0)$ into H 3: 4: end for 5: $t \leftarrow 0$ 6: while $t < n^k$ do 7: pop (v, i, j) from H $Z_k[t] \leftarrow v$ 8: if $j + 1 < n^{k-1}$ then 9: push $(Z_{k-1}[j+1] + X_k[i], i, j+1)$ into H 10: else 11: push $(+\infty, i, j+1)$ into H 12:end if 13: $t \leftarrow t + 1$ 14: 15: end while 16: return Z_k

Remark 2.18. Each extract-min + re-insert on size-n heap costs O(logn) comparisons, and we do that once for each of the n^k output elements. In addition, a O(n)-time heapify upfront yields a total time (and total number of comparisons) is $O(n^k logn) + O(n)$. Since n is treated as a fixed constant, by definition log(n) = 1, this shows Algorithm 2.2 runs in $O(n^k)$ time and complexity.

Note 2.19. It is important to note that Lemma 2.17 and Theorem 1.2 relies on the assumption that n is a fixed constant.

2.6. Computational Models and Limits. Our algorithm and analysis sit at the intersection of two standard computational models. In this subsection we spell out the subtleties and inherent limitations that accompany each.

1. Comparison Model vs. RAM Model..

- Comparison Model. Here we count only the number of key-comparisons between sums. Theorem 2.11 establishes a tight $\Theta(n^2)$ bound on comparisons (amortized O(1) per insertion). This bound is information-theoretically optimal, matching Fredman's lower bound in the decision-tree model [[4]].
- *RAM Model.* In practice, each insertion involves pointer or index manipulations that incur real clock-time costs:
 - std::list: true O(1) pointer-updates but poor spatial locality (cache misses on random jumps).
 - std::vector: O(n) element shifts per insertion but excellent locality (sequential memory).
 - Hybrid structures: B-trees or skip-lists offer $O(\log n)$ search and O(1) insertion with tunable node-fanout to trade off between pointer-chasing and block locality.

In Section 3.4 we empirically compare these containers under realistic CPU cache hierarchies. While the *comparison* count remains $O(n^2)$, the *wall-clock* time can differ by constant factors of $2-3\times$ depending on locality.

2. Structural Assumptions and Lower Bounds. Our $O(n^2)$ algorithm critically relies on the exact translation property of sumsets (Corollary 2.9): each row of the matrix is a constant shift of the previous. If one only assumes the weaker Monge or "pseudoline" monotonicity (i.e. each row and column is sorted but without a constant offset), then the best known algorithms require $\Omega(n^2 \log n)$ comparisons:

- Pseudoline arrangements. Steiger et al. [7] show that sorting the intersection points of n pseudolines requires $\Omega(n^2 \log n)$ comparisons, even though the matrix is totally monotone.
- 3SUM-hardness. Barequet and Har-Peled [1] reduce many geometric problems to sorting a general monotone matrix, inheriting an $\Omega(n^2)$ decision-tree lower bound, but with an extra log n factor in the absence of exact translation.

Thus our result exploits a strictly stronger combinatorial structure than mere Mongetype monotonicity.

3. Beyond Static Sorting.. Finally, we comment on dynamic and parallel extensions:

- Dynamic updates. Conjecture 1.4 asks whether one can maintain the sorted sumset under insertions/deletions in $\tilde{O}(n^{k-1})$ time per update. Known dynamic order-maintenance data structures (e.g. balanced BSTs) pay $\Omega(\log n)$ per operation and currently no subquadratic dynamic sumset algorithm is known.
- Parallel and external memory. In PRAM or cache-oblivious models, one must balance parallel merge overheads or block transfers. While our comparison count remains $O(n^2)$, achieving matching work-depth or I/O bounds (e.g. $O(\frac{n^2}{B} \log_{M/B} n)$ in external memory) is an open direction.

In summary, our algorithm attains the optimal $\Theta(n^2)$ comparison bound in the classical decision-tree model by leveraging exact translations, but practical performance and extensions to weaker structures or dynamic settings remain constrained by well-known lower bounds and memory-hierarchy costs.

2.7. Example. To illustrate Algorithm 2.1 in action, we walk through an example using the input sets:

- $X = \{2, 4, 6\}$
- $Y = \{1, 3, 5\}$

2.7.1. Initialization.

• Compute the *low* vector:

low = {
$$x[0] + y[0], x[1] + y[0], x[2] + y[0]$$
} = {3, 5, 7}

- Set the insertion pointer ip = 0.
- Initialize $Z = \{\}$ (empty vector).

A visual representation of the processing, is given in section 2.7.3 for better clarity.

2.7.2. Processing Steps. The algorithm iterates over i and j, computing sums and inserting them into Z. The insertion position is determined using cp, which starts at ip and moves forward.

Processing i = 0:

- Set cp = ip = 0.
- For j = 0:

$$S = x[0] + y[0] = 2 + 1 = 3$$

- Insert 3 at position 0. - Update: $Z = \{3\}, cp = 0, ip = 0.$

• For j = 1:

$$S = x[0] + y[1] = 2 + 3 = 5$$

- Insert 5 at position 1. - Update: $Z = \{3, 5\}, cp = 1, ip = 1.$

• For j = 2:

$$S = x[0] + y[2] = 2 + 5 = 7$$

- Insert 7 at position 2. - Update: $Z = \{3, 5, 7\}, cp = 2, ip = 1$. Processing i = 1:

- Set cp = ip = 2.
- For j = 0:

$$S = x[1] + y[0] = 4 + 1 = 5$$

- Insert 5 at position 2. - Update: $Z = \{3, 5, 5, 7\}, cp = 2, ip = 2$. • For j = 1:

$$S = x[1] + y[1] = 4 + 3 = 7$$

- Insert 7 at position 4. - Update: $Z = \{3, 5, 5, 7, 7\}, cp = 4, ip = 4.$ • For j = 2:

$$S = x[1] + y[2] = 4 + 5 = 9$$

- Insert 9 at position 5. - Update: $Z=\{3,5,5,7,7,9\},\, cp=5,\, ip=4.$ Processing $i=2{:}$

• Set cp = ip = 4.

• For j = 0:

$$S = x[2] + y[0] = 6 + 1 = 7$$

Insert 7 at position 5. - Update: Z = {3,5,5,7,7,7,9}, cp = 5, ip = 4.
For j = 1:

$$S = x[2] + y[1] = 6 + 3 = 9$$

- Insert 9 at position 7. - Update: $Z = \{3, 5, 5, 7, 7, 7, 9, 9\}, cp = 7, ip = 4.$ • For j = 2:

$$S = x[2] + y[2] = 6 + 5 = 11$$

- Insert 11 at position 8. - Update: $Z = \{3, 5, 5, 7, 7, 7, 9, 9, 11\}, cp = 8, ip = 4.$

2.7.3. Visual Example. To make the roles of the insertion pointer ip and scanning pointer cp crystal-clear, we show three snapshots of the list Z as it grows. Blue arrows mark ip; red arrows mark cp.

After row i = 0 (inserted 3,5,7):

3	₿ <i>ip</i>	♥ cp

During row i = 1, after inserting 5,7:



Final Z after row i = 2:



Fig. 1: Three snapshots of the list Z. Blue arrow = insertion-pointer ip. Red arrow = scanning-pointer cp.

Explanation..

- After row i = 0. We have inserted $\{2+1, 2+3, 2+5\} = \{3, 5, 7\}$. Here low[1] = 4 + 1 = 5, so ip moves to the first element > 5, namely position 1; the scan pointer cp ended at position 2.
- During row i = 1. We insert 4 + 1 = 5 at index 2, then 4 + 3 = 7 at index 4. *ip* remains at 2, and *cp* advances to 4.
- Final (row i = 2). After inserting {6+1, 6+3, 6+5}, we obtain {3, 5, 5, 7, 7, 7, 9, 9, 11}. *ip* stays at 4, *cp* ends at 8.

2.7.4. Final Output. After all iterations, the sorted sumset is:

$$Z = \{3, 5, 5, 7, 7, 7, 9, 9, 11\}$$

This walk through demonstrates the execution of the algorithm, showing how the insertion pointer ip and scanning pointer cp optimize the search for the correct insertion position.

3. Experimental Results. To validate our theoretical findings and gauge practical performance, we implemented the proposed sumset-sorting algorithm in C++ and compared it against classical full-sort methods (QuickSort and MergeSort) on the generated sumsets. All experiments were conducted on a machine with an Intel i7@3.6 GHz CPU and 16 GB RAM, compiling with -O3 under GCC.

3.1. Experimental Setup. We generated two arrays X and Y of length n, each filled with independent uniform integers in [0, 10000]. After sorting each input array, we formed the sumset X + Y of size n^2 . We measured:

- **Proposed Algorithm:** Our pointer-based insertion approach, implemented using a std::list to maintain the growing sorted output dynamically.
- QuickSort / MergeSort: Standard sorting (QuickSort and MergeSort) on the full n^2 element vector.

Each configuration was run ten times for

 $n \in \{100, 200, 500, 1000, 2000, 5000, 10000\},\$

and we report the average wall-clock time in milliseconds, measured with C++17's std::chrono::high_resolution_clock for all three methods. Comparison counts for the proposed algorithm were recorded via manual counters inside the insertion routine; for QuickSort and MergeSort, we likewise instrumented their comparison functions to obtain the exact number of comparisons at runtime. This rigorous setup allows us to observe both the true asymptotic behavior and the practical performance impacts of cache locality and pointer-chasing.

Before dwelving into the experimental results, we shall address practical trade-offs of using std::list instead of std::vector or other data structures like B-trees.

Practical Data-Structure Trade-off. In our benchmarks we used std::list to achieve true $\mathcal{O}(1)$ insertion per element. While a contiguous container such as std::vector offers better cache locality, each search adds $\mathcal{O}(n)$ insertion into a vector would actually slow down the overall routine on large sumsets. In preliminary tests, switching the insertion structure from std::list to std::vector increased wall-clock time by over 30% for $n \geq 2000$. Hence, although linked lists suffer pointer-chasing overhead, they remain the fastest choice for our amortized-constant-time insertion pattern. We leave a more detailed study of hybrid or gap-buffer approaches to future work. Beyond the linked-list vs. vector dichotomy, there exist intermediate structures that may offer even better overall performance. For example:

- Skip lists maintain multiple forward-pointers per node, providing expected $\mathcal{O}(\log n)$ search and $\mathcal{O}(1)$ insertion, while still using pointer-based storage. In practice, a skip list can reduce the number of cache misses compared to a simple singly-linked list.
- Balanced search trees (e.g. red-black trees or B-trees) support $\mathcal{O}(\log n)$ search and insertion with good node-occupancy and cache-aware fan-out. A B-tree tuned for large nodes can amortize pointer-chasing across many elements per cache line.
- Gap buffers or rope-like arrays reserve small "gaps" within a dynamic array to allow fast insertions without a full shift of all tail elements, trading a slight increase in memory usage for $\mathcal{O}(1)$ amortized insertion near the current gap.

In future implementations, one could benchmark these alternatives under our sumset workload. In particular, a hybrid approach—using a small vector chunk per list-node—may combine the low latency of array accesses with the amortized insertion guarantee of a list. We anticipate that such cache-blocked or B-tree-based structures would further narrow the gap between our theoretical $\mathcal{O}(n^2)$ bound and optimal wall-clock performance, especially on modern CPUs with deep memory hierarchies.

QuickSort Proposed MergeSort n100 1.1 9 7 20052088 215001411,4981.000600 23.59288 2,000 $4.06 \cdot 10^{5}$ 387 2,5225,0002,23717,597 $1.06 \cdot 10^{6}$ 10,000 $2.86\cdot 10^6$ 9,145 73,627

Table 1: Average Number of Data Comparisons for Sorting X + Y

3.2. Comparison Count. Table 1 shows that the proposed algorithm performs $\mathcal{O}(n^2)$ comparisons exactly, while QuickSort and MergeSort incur the additional $\log(n^2)$ factor, matching their $\mathcal{O}(n^2 \log n)$ behavior. It is also easy to notice, that QuickSort is slower then both our Proposed Algorithm and MergeSort.

3.3. Execution Time. Figure 2 and Figure 3 shows that for large n, the proposed algorithm outperforms both QuickSort and MergeSort. Figure 3 highlighting the consistent advantage of our method over MergeSort. Finally, the log–log plot in Figure 4 confirms the asymptotic slopes: the proposed algorithm scales as $\mathcal{O}(n^2)$ (slope 2), whereas the full sorts exhibit the additional logarithmic factor.



Fig. 2: Execution time for all three algorithms on the sumset.



Fig. 4: Log–log plot of execution time vs. n for all three methods.



Fig. 3: Zoom-in comparing our method against MergeSort only.

3.4. Empirical Validation of $\mathcal{O}(1)$ **Comparisons and Insertions.** To empirically validate the claim that each insertion into the sorted output list incurs only $\mathcal{O}(1)$ amortized work, we measured the quantity T/n^2 , where T is the total runtime in milliseconds and n^2 is the number of elements in the sumset.

For each value of $n \in \{100, 200, 500, 1000, 2000, 5000, 10000\}$, we executed the algorithm ten times on randomly generated sorted arrays X and Y, and recorded the mean and standard deviation of T/n^2 .

Table 2: Mean and standard deviation of T/n^2 over 10 runs.

n	Mean T/n^2 (ms)	Std. Dev. (ms)
100	$4.564e{-5}$	$1.433e{-5}$
200	$3.334e{-5}$	4.582e - 6
500	$1.943e{-5}$	2.411e - 6
1000	$1.745e{-5}$	$2.944e{-7}$
2000	$1.748e{-5}$	$2.801e{-7}$
5000	$1.756e{-5}$	$2.431e{-7}$
10000	$1.835e{-5}$	2.342e - 7

As shown in Table 2, the values of T/n^2 remain remarkably stable, and the standard deviation shrinks as n increases. This indicates a tight concentration of insertion times around a constant mean, supporting the $\mathcal{O}(1)$ amortized insertion claim.



Fig. 5: Stability of T/n^2 over 10 trials. The trend appears nearly flat, supporting the $\mathcal{O}(1)$ insertion claim.

3.5. Discussion. These experiments corroborate our theoretical analysis: by attaining exactly $\mathcal{O}(n^2)$ comparisons and leveraging the structure of the sumset, our algorithm achieves superior practical performance on large inputs.

The pointer-based insertion method performs consistently across different input sizes, and empirical measurements of T/n^2 demonstrate near-constant values. As shown in Figure 5, the error bars around each measurement are narrow, and the standard deviation shrinks as n increases, indicating tighter concentration around the expected $\mathcal{O}(1)$ insertion and comparison cost.

This validates the amortized constant-time insertion behavior, confirming that the theoretical efficiency of our algorithm extends to real-world implementations.

Future work will explore alternative data structures (e.g. contiguous buffers, B-trees, or gap buffers) to reduce pointer-chasing overhead and improve cache locality. Additional directions include parallel and external-memory variants, as well as applications to structured domains such as geometry, sparse data joins, and layout optimization.

4. Conclusion and Future Work. We have presented the first explicit, implementable algorithm that sorts the sumset

$$X + Y = \{ x_i + y_j \mid x_i \in X, \ y_j \in Y \}$$

in optimal $O(n^2)$ time and comparisons. By exploiting the row-wise and column-wise monotonicity of the sumset matrix, our forward-scanning insertion strategy achieves amortized constant-time insertion per element, matching Fredman's existential bound with a concrete procedure. We proved correctness and tight comparison complexity in the standard comparison model, and demonstrated via extensive C++ benchmarks that our algorithm outperforms classical $O(n^2 \log n)$ methods (Merge Sort and Quick Sort) on large inputs.

Moreover, we showed that the same ideas extend naturally to the k-fold sumset

$$X_1 + X_2 + \dots + X_k = \{x_1 + x_2 + \dots + x_k \mid x_i \in X_i\}$$

of k sorted lists of length n, yielding an $O(n^k)$ -time and comparison-optimal algorithm by induction and a structured n-way merge of translated partial sumsets.

AN EXPLICIT AND EFFICIENT $\mathcal{O}(N^2)$ -TIME ALGORITHM FOR SORTING SUMSETS 23

This work closes a long-standing gap between theory and practice in structured sorting, resolving an open problem that has stood for nearly fifty years. Our algorithms not only attain the information-theoretic lower bound on comparisons, but also exhibit strong real-world performance, making them directly applicable to tasks in computational geometry, sparse polynomial multiplication, VLSI design, and other areas where multi-way combinations arise. We achieve $O(n^2)$ RAM time in our prototype with a linked-list; designing a cache-friendly structure to provably attain the same bound (or beat it in practice) is an interesting open problem.

Open Problems and Future Directions. Despite its optimality, our approach suggests several avenues for further research:

- Cache-Optimized Data Structures. Replacing the std::list in our prototype with cache-aware or hardware-friendly structures—such as blocked linked lists, B-trees, or gap buffers—may yield additional speedups by reducing pointer-chasing overhead.
- **Parallel and External Memory Algorithms.** The current algorithm is inherently sequential. Designing parallel variants for multi-core or GPU architectures, or adapting it to external-memory models, could extend its scalability to even larger and higher-dimensional sumsets.
- Hybrid Decision-Tree Techniques. Kane, Lovett, and Moran's decision-tree framework achieves sub-quadratic query complexity under sparsity constraints. Investigating hybrid algorithms that combine fixed-structure insertion with sparse decision-tree inference may further reduce comparisons in practice.
- Dynamic and Streaming Sumsets. Maintaining a sorted k-fold sumset under insertions and deletions to each X_i remains open. Data structures supporting updates in $O(n^k)$ time would enable real-time applications and streaming scenarios.
- Empirical Evaluation on Real-World Data. Beyond synthetic benchmarks, applying our algorithms to domain-specific workloads—such as high dimensional distance computations, database joins, or signal processing pipelines—will validate their utility and uncover practical refinements.

We believe these directions will deepen our understanding of structured sorting and broaden the impact of optimal comparison-based algorithms in both theory and practice.

Acknowledgments. We thank Dr. C. Jackson and Dr. S. McCulloh for their insightful discussions and valuable feedback throughout the development of this work. Their guidance helped shape both the theoretical and empirical aspects of this paper.

REFERENCES

- G. BAREQUET AND S. HAR-PELED, Polygon containment and translational in-hausdorffdistance between segment sets are 3sum-hard, International Journal of Computational Geometry & Applications, 11 (2001), pp. 465–474, https://doi.org/10.1142/ S0218195901000596, https://doi.org/10.1142/S0218195901000596, https://arxiv.org/abs/ https://doi.org/10.1142/S0218195901000596.
- [2] A. H. BARRERA*, Finding an o(n² log n) algorithm is sometimes hard, McGill-Queen's University Press, Montreal, 1996, pp. 289–294, https://doi.org/doi:10.1515/9780773591134-051, https://doi.org/10.1515/9780773591134-051.
- M. DIETZFELBINGER, Lower bounds for sorting of sums, Theoretical Computer Science, 66 (1989), pp. 137–155, https://doi.org/https://doi.org/10.1016/0304-3975(89)90132-1, https://www. sciencedirect.com/science/article/pii/0304397589901321.

- M. L. FREDMAN, How good is the information theory bound in sorting?, Theoretical Computer Science, 1 (1976), pp. 355–361, https://doi.org/https://doi.org/ 10.1016/0304-3975(76)90078-5, https://www.sciencedirect.com/science/article/pii/ 0304397576900785.
- D. M. KANE, S. LOVETT, AND S. MORAN, Near-optimal linear decision trees for k-sum and related problems, J. ACM, 66 (2019), https://doi.org/10.1145/3285953, https://doi.org/10. 1145/3285953.
- J.-L. LAMBERT, Sorting the sums (xi + yj) in o(n2) comparisons, Theoretical Computer Science, 103 (1992), pp. 137–141, https://doi.org/https://doi.org/10.1016/0304-3975(92)90089-X, https://www.sciencedirect.com/science/article/pii/030439759290089X.
- [7] W. STEIGER AND I. STREINU, A pseudo-algorithmic separation of lines from pseudo-lines, Information processing letters, 53 (1995), p. 295.
- [8] C. D. TOTH, J. O'ROURKE, AND J. E. GOODMAN, Handbook of discrete and computational geometry, CRC press, 2017.
- I. VAN DER HOOG, E. ROTENBERG, AND D. RUTSCHMANN, Simpler Optimal Sorting from a Directed Acyclic Graph, pp. 350–355, https://doi.org/10.1137/1.9781611978315.26, https:// epubs.siam.org/doi/abs/10.1137/1.9781611978315.26, https://arxiv.org/abs/https://epubs. siam.org/doi/pdf/10.1137/1.9781611978315.26.