

# Orthogonal Range Searching in Moderate Dimensions: k-d Trees and Range Trees Strike Back<sup>\*†</sup>

Timothy M. Chan

Dept. of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, USA  
tmc@illinois.edu

---

## Abstract

---

We revisit the *orthogonal range searching* problem and the *exact  $\ell_\infty$  nearest neighbor searching* problem for a static set of  $n$  points when the dimension  $d$  is moderately large. We give the first data structure with near linear space that achieves truly sublinear query time when the dimension is any constant multiple of  $\log n$ . Specifically, the preprocessing time and space are  $O(n^{1+\delta})$  for any constant  $\delta > 0$ , and the expected query time is  $n^{1-1/O(c \log c)}$  for  $d = c \log n$ . The data structure is simple and is based on a new “augmented, randomized, lopsided” variant of k-d trees. It matches (in fact, slightly improves) the performance of previous combinatorial algorithms that work only in the case of offline queries [Impagliazzo, Lovett, Paturi, and Schneider (2014) and Chan (SODA’15)]. It leads to slightly faster combinatorial algorithms for *all-pairs shortest paths* in general real-weighted graphs and *rectangular Boolean matrix multiplication*.

In the offline case, we show that the problem can be reduced to the *Boolean orthogonal vectors* problem and thus admits an  $n^{2-1/O(\log c)}$ -time non-combinatorial algorithm [Abboud, Williams, and Yu (SODA’15)]. This reduction is also simple and is based on range trees.

Finally, we use a similar approach to obtain a small improvement to Indyk’s data structure [FOCS’98] for *approximate  $\ell_\infty$  nearest neighbor search* when  $d = c \log n$ .

**1998 ACM Subject Classification** F.2.2 Nonnumerical Algorithms and Problems

**Keywords and phrases** computational geometry, data structures, range searching, nearest neighbor searching

**Digital Object Identifier** 10.4230/LIPIcs.SoCG.2017.27

## 1 Introduction

In this paper, we revisit some classical problems in computational geometry:

- In *orthogonal range searching*, we want to preprocess  $n$  data points in  $\mathbb{R}^d$  so that we can detect if there is a data point inside any query axis-aligned box, or report or count all such points.
- In *dominance range searching*, we are interested in the special case when the query box is  $d$ -sided, of the form  $(-\infty, q_1] \times \cdots \times (-\infty, q_d]$ ; in other words, we want to detect if there is a data point  $(p_1, \dots, p_d)$  that is dominated by a query point  $(q_1, \dots, q_d)$ , in the sense that  $p_j \leq q_j$  for all  $j \in \{1, \dots, d\}$ , or report or count all such points.
- In  *$\ell_\infty$  nearest neighbor searching*, we want to preprocess  $n$  data points in  $\mathbb{R}^d$  so that we can find the nearest neighbor to the given query point under the  $\ell_\infty$  metric.

---

\* A full version of the paper is available at [http://tmc.web.engr.illinois.edu/high\\_ors3\\_17.pdf](http://tmc.web.engr.illinois.edu/high_ors3_17.pdf).

† This work was done while the author was at the Cheriton School of Computer Science, University of Waterloo.



© Timothy M. Chan;

licensed under Creative Commons License CC-BY

33rd International Symposium on Computational Geometry (SoCG 2017).

Editors: Boris Aronov and Matthew J. Katz; Article No. 27; pp. 27:1–27:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

All three problems are related. Orthogonal range searching in  $d$  dimensions reduces to dominance range searching in  $2d$  dimensions.<sup>1</sup> Furthermore, ignoring logarithmic factors,  $\ell_\infty$  nearest neighbor searching reduces to its decision problem (deciding whether the  $\ell_\infty$  nearest neighbor distance to a given query point is at most a given radius) by parametric search or randomized search [7], and the decision problem clearly reduces to orthogonal range searching.

The standard  $k$ - $d$  tree [22] has  $O(dn \log n)$  preprocessing time and  $O(dn)$  space, but the worst-case query time is  $O(dn^{1-1/d})$ . The standard *range tree* [22] requires  $O(n \log^d n)$  preprocessing time and space and  $O(\log^d n)$  query time, excluding an  $O(K)$  term for the reporting version of the problem with output size  $K$ . Much work in computational geometry has been devoted to small improvements of a few logarithmic factors. For example, the current best result for orthogonal range reporting has  $O(n \log^{d-3+\varepsilon} n)$  space and  $O(\log^{d-3} n / \log^{d-4} \log n + K)$  time [12]; there are also other small improvements for various offline versions of the problems [12, 13, 2].

In this paper, we are concerned with the setting *when the dimension is nonconstant*. Traditional approaches from computational geometry tend to suffer from exponential dependencies in  $d$  (the so-called “curse of dimensionality”). For example, the  $O(dn^{1-1/d})$  or  $O(\log^d n)$  query time bound for range trees or  $k$ - $d$  trees is sublinear only when  $d \ll \log n / \log \log n$ . By a more careful analysis [10], one can show that range trees still have sublinear query time when  $d \ll \alpha_0 \log n$  for a sufficiently small constant  $\alpha_0$ . The case when the dimension is close to logarithmic in  $n$  is interesting in view of known dimensionality reduction techniques [16] (although such techniques technically are not applicable to exact problems and, even with approximation, do not work well for  $\ell_\infty$ ). The case of polylogarithmic dimensions is also useful in certain *non-geometric* applications such as all-pairs shortest paths (as we explain later). From a theoretical perspective, it is important to understand when the time complexity transitions from sublinear to superlinear.

**Previous offline results.** We first consider the *offline* version of the problems where we want to answer a batch of  $n$  queries all given in advance. In high dimensions, it is possible to do better than  $O(dn^2)$ -time brute-force search, by a method of Matoušek [21] using fast (rectangular) matrix multiplication [20]; for example, we can get  $n^{2+o(1)}$  time for  $d \ll n^{0.15}$ . However, this approach inherently cannot give *subquadratic* bounds.

In 2014, a surprising discovery was made by Impagliazzo et al. [17]: range-tree-like divide-and-conquer can still work well even when the dimension goes a bit above logarithmic. Their algorithm can answer  $n$  offline dominance range queries (and thus orthogonal range queries and  $\ell_\infty$  nearest neighbor queries) in total time  $n^{2-1/O(c^{15} \log c)}$  (ignoring an  $O(K)$  term for reporting) in dimension  $d = c \log n$  for any possibly nonconstant  $c$  ranging from 1 to about  $\log^{1/15} n$  (ignoring  $\log \log n$  factors). Shortly after, by a more careful analysis of the same algorithm, Chan [8] refined the time bound to  $n^{2-1/O(c \log^2 c)}$ , which is subquadratic for  $c$  up to about  $\log n$ , i.e., dimension up to about  $\log^2 n$ .

At SODA’15, Abboud, Williams, and Yu [1] obtained an even better time bound for dominance range detection in the *Boolean* special case, where all coordinate values are 0’s and 1’s (in this case, the problem is better known as the *Boolean orthogonal vectors* problem<sup>2</sup>).

<sup>1</sup>  $(p_1, \dots, p_d)$  is inside the box  $[a_1, b_1] \times \dots \times [a_d, b_d]$  iff  $(-p_1, p_1, \dots, -p_d, p_d)$  is dominated by  $(-a_1, b_1, \dots, -a_d, b_d)$  in  $\mathbb{R}^{2d}$ .

<sup>2</sup> Two vectors  $(p_1, \dots, p_d), (q_1, \dots, q_d) \in \{0, 1\}^d$  are orthogonal iff  $\sum_{i=1}^d p_i q_i = 0$  iff  $(p_1, \dots, p_d)$  is dominated by  $(1 - q_1, \dots, 1 - q_d)$  (recalling that our definition of dominance uses non-strict inequality).

The total time for  $n$  offline Boolean dominance range detection queries is  $n^{2-1/O(\log c)}$ . The bound  $n^{2-1/O(\log c)}$  is a natural barrier, since a faster offline Boolean dominance algorithm would imply an algorithm for CNF-SAT with  $n$  variables and  $cn$  clauses that would beat the currently known  $2^{n(1-1/O(\log c))}$  time bound [1]; and an  $O(n^{2-\delta})$ -time algorithm for any  $c = \omega(1)$  would break the strong exponential-time hypothesis (SETH) [24]. Abboud et al.'s algorithm was based on the *polynomial method* pioneered by Williams [23] (see [4, 3] for other geometric applications). The algorithm was originally randomized but was subsequently derandomized by Chan and Williams [9] in SODA'16 (who also extended the result from detection to counting).

Abboud et al.'s approach has two main drawbacks, besides being applicable to the Boolean case only: 1. it is not “combinatorial” and relies on fast rectangular matrix multiplication, making the approach less likely to be practical, and 2. it only works in the offline setting.

Impagliazzo et al.'s range-tree method [17] is also inherently restricted to the offline setting – in their method, the choice of dividing hyperplanes crucially requires knowledge of all query points in advance. All this raises an intriguing open question: are there nontrivial results for online queries in  $d = c \log n$  dimensions?

**New online result.** In Section 2, we resolve this question by presenting a randomized data structure with  $O(n^{1+\delta})$  preprocessing time and space that can answer online dominance range queries (and thus orthogonal range queries and  $\ell_\infty$  nearest neighbor queries) in  $n^{1-1/O(c \log^2 c)}$  expected time for any  $d = c \log n \ll \log^2 n / \log \log n$  and for any constant  $\delta > 0$ . (We assume an oblivious adversary, i.e., that query points are independent of the random choices made by the preprocessing algorithm.) The total time for  $n$  queries is  $n^{2-1/O(c \log^2 c)}$ , matching the offline bound from Impagliazzo et al. [17] and Chan [8]. The method is purely combinatorial, i.e., does not rely on fast matrix multiplication.

More remarkable than the result perhaps is the simplicity of the solution: it is just a variant of k-d trees! More specifically, the dividing hyperplane is chosen in a “lopsided” manner, along a randomly chosen coordinate axis; each node is augmented with secondary structures for some lower-dimensional projections of the data points. The result is surprising, considering the longstanding popularity of k-d trees among practitioners. Our contribution lies in recognizing, and proving, that they can have good theoretical worst-case performance. (Simple algorithms with nonobvious analyses are arguably the best kind.)

In Appendix A.1, we also describe a small improvement of the query time to  $n^{1-1/O(c \log c)}$ . This involves an interesting application of so-called *covering designs* (from combinatorics), not often seen in computational geometry.

**Applications.** By combining with previous techniques [10, 8], our method leads to new results for two classical, non-geometric problems: *all-pairs shortest paths (APSP)* and *Boolean matrix multiplication (BMM)*.

- We obtain a new combinatorial algorithm for solving the APSP problem for arbitrary real-weighted graphs with  $n$  vertices (or equivalently the  $(\min, +)$  matrix multiplication problem for two  $n \times n$  real-valued matrices) in  $O((n^3 / \log^3 n) \text{poly}(\log \log n))$  time; see Appendix A.2. This is about a logarithmic factor faster than the best previous combinatorial algorithm [11, 15, 8], not relying on fast matrix multiplication à la Strassen. It also extends Chan's combinatorial algorithm for Boolean matrix multiplication from SODA'15 [8], which has a similar running time (although for Boolean matrix multiplication, Yu [26] has recently obtained a further logarithmic-factor improvement).

This extension is intriguing, as  $(\min, +)$  matrix multiplication over the reals appears tougher than other problems such as standard matrix multiplication over  $\mathbb{F}_2$ , for which the

well-known “four Russians” time bound of  $O(n^3/\log^2 n)$  [6] has still not been improved for combinatorial algorithms.

- We obtain a new combinatorial algorithm to multiply an  $n \times \log^2 n$  and a  $\log^2 n \times n$  Boolean matrix in  $O((n^2/\log n) \text{poly}(\log \log n))$  time, which is almost optimal in the standard word RAM model since the output requires  $\Omega(n^2/\log n)$  words; see the full paper. The previous combinatorial algorithm by Chan [8] can multiply an  $n \times \log^3 n$  and a  $\log^3 n \times n$  Boolean matrix in  $O(n^2 \text{poly}(\log \log n))$  time. The new result implies the old, but not vice versa.

**New offline result.** Returning to the offline dominance or orthogonal range searching problem, Abboud, Williams, and Yu’s non-combinatorial algorithm [1] has a better  $n^{2-1/O(\log c)}$  time bound but is only for the Boolean case, leading to researchers to ask whether the same result holds for the more general problem for real input. In one section of Chan and Williams’ paper [9], such a result was obtained but only for  $d \approx 2^{\Theta(\sqrt{\log n})}$ .

In Section 3, we resolve this question by giving a black-box reduction from the real case to the Boolean case, in particular, yielding  $n^{2-1/O(\log c)}$  time for any  $d = c \log n \ll 2^{\Theta(\sqrt{\log n})}$ .

This equivalence between general dominance searching and the Boolean orthogonal vectors problem is noteworthy, since the Boolean orthogonal vectors problem has been used recently as a basis for many conditional hardness results in P.

As one immediate application, we can now solve the *integer linear programming* problem on  $n$  variables and  $cn$  constraints in  $2^{(1-1/O(\log c))n}$  time, improving Impagliazzo et al.’s  $2^{(1-1/\text{poly}(c))n}$  algorithm [17].

Our new reduction is simple, this time, using a range-tree-like recursion.

**Approximate  $\ell_\infty$  nearest neighbor searching.** So far, our discussion has been focused on exact algorithms. We now turn to  $\ell_\infty$  nearest neighbor searching in the *approximate* setting. By known reductions (ignoring polylogarithmic factors) [16], it suffices to consider the fixed-radius decision problem: deciding whether the nearest neighbor distance is approximately less than a fixed value. Indyk [18] provided the best data structure for the problem, achieving  $O(\log_\rho \log d)$  approximation factor,  $O(dn^\rho \log n)$  preprocessing time,  $O(dn^\rho)$  space, and  $O(d \log n)$  query time for any  $\rho$  ranging from 1 to  $\log d$ . The data structure is actually based on traditional-style geometric divide-and-conquer. Andoni, Croitoru, and Pătraşcu [5] proved a nearly matching lower bound.

In Section 4, we improve the approximation factor of Indyk’s data structure to  $O(\log_\rho \log c)$  for dimension  $d = c \log n$ , for any  $\rho$  ranging from  $1 + \delta$  to  $\log c$  (as an unintended byproduct, we also improve Indyk’s query time to  $O(d)$ ). The improvement in the approximation factor is noticeable when the dimension is close to logarithmic. It does not contradict Andoni et al.’s lower bound [5], since their proof assumed  $d \gg \log^{1+\Omega(1)} n$ .

For example, by setting  $\rho \approx \log c$ , we get  $O(1)$  approximation factor,  $n^{O(\log c)}$  preprocessing time/space, and  $O(d)$  query time. By dividing into  $n^{1-\alpha}$  groups of size  $n^\alpha$ , we can lower the preprocessing time/space to  $n^{1-\alpha} \cdot (n^\alpha)^{O(\log(c/\alpha))}$  while increasing the query time to  $O(dn^{1-\alpha})$ . Setting  $\alpha \approx 1/\log c$ , we can thus answer  $n$  (online) queries with  $O(1)$  approximation factor in  $n^{2-1/O(\log c)}$  total time, which curiously matches our earlier result for exact  $\ell_\infty$  nearest neighbor search but by a purely combinatorial algorithm.

In the full paper, we also provide an alternative data structure with linear space but a larger  $O(c^{(1-\rho)/\rho^2})$  approximation factor, and  $O(dn^{\rho+\delta})$  query time for any  $\rho \in (\delta, 1 - \delta)$ .

The idea is to modify Indyk’s method to incorporate, once again, a range-tree-like recursion.

## 2 Online Dominance Range Searching

In this section, we study data structures for online orthogonal range searching in the reporting version (counting or detection can be dealt with similarly), using only combinatorial techniques without fast matrix multiplication. By doubling the dimension (footnote 1), it suffices to consider the dominance case.

Our data structure is an augmented, randomized lopsided variant of the k-d tree, where each node contains secondary structures for various lower-dimensional projections of the input.

**Data structure.** Let  $\delta \in (0, 1)$  and  $c \in [\delta C_0, (\delta/C_0) \log N / \log^2 \log N]$  be user-specified parameters, for a sufficiently large constant  $C_0$ , where  $N$  is a fixed upper bound on the size of the input point set. Let  $b \geq 2$  and  $\alpha \in (0, 1/2)$  be parameters to be chosen later.

Given a set  $P$  of  $n \leq N$  data points in  $d \leq c \log N$  dimensions, our data structure is simple and is constructed as follows:

0. If  $n \leq 1/\alpha$  or  $d = 0$ , then just store the given points.
1. Otherwise, let  $\mathcal{J}$  be the collection of all subsets of  $\{1, \dots, d\}$  of size  $\lfloor d/b \rfloor$ . Then  $|\mathcal{J}| = \binom{d}{\lfloor d/b \rfloor} = b^{O(d/b)}$ . For each  $J \in \mathcal{J}$ , recursively<sup>3</sup> construct a data structure for the projection  $P_J$  of  $P$  that keeps only the coordinate positions in  $J$ .
2. Pick a random  $i^* \in \{1, \dots, d\}$ . Let  $\mu(i^*)$  be the  $\lceil (1 - \alpha)n \rceil$ -th smallest  $i^*$ -th coordinate value in  $P$ ; let  $p(i^*)$  be the corresponding point in  $P$ . Store  $n$ ,  $i^*$ , and  $p(i^*)$ . Recursively construct data structures for
  - the subset  $P_L$  of all points in  $P$  with  $i^*$ -th coordinate less than  $\mu(i^*)$ , and
  - the subset  $P_R$  of all points in  $P$  with  $i^*$ -th coordinate greater than  $\mu(i^*)$ .

**Analysis.** The preprocessing time and space satisfy the recurrence

$$T_d(n) \leq T_d(\lfloor \alpha n \rfloor) + T_d(\lfloor (1 - \alpha)n \rfloor) + b^{O(d/b)} T_{\lfloor d/b \rfloor}(n) + O(n),$$

with  $T_d(n) = O(n)$  for the base case  $n \leq 1/\alpha$  or  $d = 0$ . This solves to

$$\begin{aligned} T_d(N) &\leq b^{O(d/b + d/b^2 + \dots)} N (\log_{1/(1-\alpha)} N)^{O(\log_b d)} \\ &= b^{O(d/b)} N ((1/\alpha) \log N)^{O(\log_b d)} \\ &= N^{1+O((c/b) \log b)} 2^{O(\log((1/\alpha) \log N) \log_b d)} \leq N^{1+O(\delta)} 2^{O(\log^2((1/\alpha) \log N))} \end{aligned}$$

by setting  $b := (c/\delta) \log(c/\delta)$ .

**Query algorithm.** Given the preprocessed set  $P$  and a query point  $q = (q_1, \dots, q_d)$ , our query algorithm proceeds as follows.

0. If  $n \leq 1/\alpha$  or  $d = 0$ , then answer the query directly by brute-force search.
1. Otherwise, let  $J_q = \{i \in \{1, \dots, d\} : q_i \neq \infty\}$ . If  $|J_q| \leq d/b$ , then recursively answer the query for  $P_{J_q}$  and the projection of  $q$  with respect to  $J_q$ .
2. Else,
  - if  $q_{i^*} \leq \mu(i^*)$ , then recursively answer the query for  $P_L$  and  $q$ ;
  - if  $q_{i^*} > \mu(i^*)$ , then recursively answer the query for  $P_R$  and  $q$ , and recursively answer the query for  $P_L$  and  $q' = (q_1, \dots, q_{i^*-1}, \infty, q_{i^*+1}, \dots, q_d)$ ;
  - in addition, if  $q$  dominates  $p(i^*)$ , then output  $p(i^*)$ .

<sup>3</sup> There are other options beside recursion here; for example, we could just use a range tree for  $P_J$ .

**Analysis.** We assume that the query point  $q$  is independent of the random choices made during the preprocessing of  $P$ . Let  $L_q = \{i \in \{1, \dots, d\} : \mu(i) < q_i \neq \infty\}$ . Let  $j = |J_q|$  and  $\ell = |L_q|$ .

Suppose that  $j > d/b$ . The probability that we make a recursive call for  $P_R$  is equal to  $\Pr[(i^* \in L_q) \vee (i^* \notin J_q)] = \ell/d + (1 - j/d)$ . We always make a recursive call for  $P_L$ , either for  $q$  or a point  $q'$  with  $j - 1$  non- $\infty$  values; the probability of the latter is equal to  $\Pr[i^* \in L_q] = \ell/d$ .

Hence, the expected number of leaves in the recursion satisfies the following recurrence:

$$Q_{d,j}(n) \leq \begin{cases} Q_{\lfloor d/b \rfloor, j}(n) & \text{if } j \leq d/b \\ \max_{\ell \leq j} \left[ \left( \frac{\ell}{d} + 1 - \frac{j}{d} \right) Q_{d,j}(\lfloor \alpha n \rfloor) + \left( \frac{\ell}{d} \right) Q_{d,j-1}(\lfloor (1-\alpha)n \rfloor) \right. \\ \quad \left. + \left( 1 - \frac{\ell}{d} \right) Q_{d,j}(\lfloor (1-\alpha)n \rfloor) \right] & \text{if } j > d/b, \end{cases} \quad (1)$$

with  $Q_{d,j}(n) = 1$  for the base case  $n \leq 1/\alpha$  or  $d = 0$ .

This recurrence looks complicated. Following [8], one way to solve it is by ‘‘guessing’’. We guess that

$$Q_{d,j}(n) \leq (1 + \gamma)^j n^{1-\varepsilon}$$

for some choice of parameters  $\gamma, \varepsilon \in (0, 1/2)$  to be specified later. We verify the guess by induction.

The base case  $n \leq 1/\alpha$  or  $d = 0$  is trivial. Assume that the guess is true for lexicographically smaller tuples  $(d, j, n)$ . For  $j \leq d/b$ , the induction trivially goes through. So assume  $j > d/b$ . Let  $\ell$  be the index that attains the maximum in (1). Then

$$\begin{aligned} Q_{d,j}(n) &\leq \left( \frac{\ell}{d} + 1 - \frac{j}{d} \right) (1 + \gamma)^j (\alpha n)^{1-\varepsilon} + \left( \frac{\ell}{d} \right) (1 + \gamma)^{j-1} ((1-\alpha)n)^{1-\varepsilon} + \\ &\quad \left( 1 - \frac{\ell}{d} \right) (1 + \gamma)^j ((1-\alpha)n)^{1-\varepsilon} \\ &= \left[ \left( \frac{\ell}{d} + 1 - \frac{j}{d} \right) \alpha^{1-\varepsilon} + \left( \frac{\ell}{d} \cdot \frac{1}{1+\gamma} + 1 - \frac{\ell}{d} \right) (1-\alpha)^{1-\varepsilon} \right] (1 + \gamma)^j n^{1-\varepsilon} \\ &\leq \left[ \left( 1 - \frac{j-\ell}{d} \right) \alpha^{1-\varepsilon} + \left( 1 - \frac{\gamma\ell}{2d} \right) (1-\alpha)^{1-\varepsilon} \right] (1 + \gamma)^j n^{1-\varepsilon} \\ &\leq (1 + \gamma)^j n^{1-\varepsilon}. \end{aligned}$$

For the last inequality, we need to upper-bound the following expression by 1:

$$\left( 1 - \frac{j-\ell}{d} \right) \alpha^{1-\varepsilon} + \left( 1 - \frac{\gamma\ell}{2d} \right) (1-\alpha)^{1-\varepsilon}. \quad (2)$$

■ **Case I:**  $j - \ell > d/(2b)$ . Then (2) is at most

$$\begin{aligned} \left( 1 - \frac{1}{2b} \right) \alpha^{1-\varepsilon} + (1-\alpha)^{1-\varepsilon} &\leq \left( 1 - \frac{1}{2b} \right) \alpha e^{\varepsilon \ln(1/\alpha)} + 1 - (1-\varepsilon)\alpha \\ &\leq \left( 1 - \frac{1}{2b} \right) \alpha (1 + 2\varepsilon \log(1/\alpha)) + 1 - (1-\varepsilon)\alpha \\ &\leq 1 - \frac{\alpha}{2b} + 3\alpha\varepsilon \log(1/\alpha), \end{aligned}$$

which is indeed at most 1 by setting  $\varepsilon := 1/(6b \log(1/\alpha))$ .

- **Case II:**  $\ell > d/(2b)$ . Then (2) is at most

$$\begin{aligned} \alpha^{1-\varepsilon} + 1 - \frac{\gamma}{4b} &\leq \alpha e^{\varepsilon \ln(1/\alpha)} + 1 - \frac{\gamma}{4b} \\ &\leq \alpha(1 + 2\varepsilon \log(1/\alpha)) + 1 - \frac{\gamma}{4b} \\ &\leq 2\alpha + 1 - \frac{\gamma}{4b}, \end{aligned}$$

which is indeed at most 1 by setting  $\gamma := 8b\alpha$ .

We can set  $\alpha := 1/b^4$ , for example. Then  $\gamma = O(1/b^3)$ . We conclude that

$$Q_d(N) \leq (1 + \gamma)^d N^{1-\varepsilon} \leq e^{\gamma d} N^{1-\varepsilon} \leq N^{1-\varepsilon+O(c\gamma)} \leq N^{1-1/O(b \log b)}.$$

Now,  $Q_d(N)$  only counts the number of leaves in the recursion. The recursion has depth  $O(\log_{1/(1-\alpha)} N + \log d)$ . Each internal node of the recursion has cost  $O(d)$ , and each leaf has cost  $O(d/\alpha)$ , excluding the cost of outputting points (which occurs during the base case  $d = 0$ ). Thus, the actual expected query time can be bounded by  $Q_d(N)(bd \log N)^{O(1)}$ , which is  $N^{1-1/O(b \log b)}$  for  $b \ll \log N / \log^2 \log N$ . As  $b = (c/\delta) \log(c/\delta)$ , the bound is  $N^{1-1/O((c/\delta) \log^2(c/\delta))}$ .

Slight improvement of one  $\log(c/\delta)$  factor in the exponent is possible, by an interesting application of *covering designs*. The details are explained in Appendix A.1. Thus:

► **Theorem 1.** *Let  $\delta > 0$  be any fixed constant and  $c \in [C_1, (1/C_1) \log N / \log^2 \log N]$  for a sufficiently large constant  $C_1$ . Given  $N$  points in  $d = c \log N$  dimensions, we can construct a data structure in  $O(N^{1+\delta})$  preprocessing time and space, so that for any query point, we can answer a dominance range reporting query in  $N^{1-1/O(c \log c)} + O(K)$  expected time where  $K$  is the number of reported points. For dominance range counting, we get the same time bound but without the  $K$  term.*

We mention one application to online  $(\min, +)$  matrix-vector multiplication. The corollary below follows immediately from a simple reduction [10] to  $d$  instances of  $d$ -dimensional dominance range reporting with disjoint output.<sup>4</sup>

► **Corollary 2.** *Let  $\delta > 0$  be any fixed constant and  $d = (1/C_1) \log^2 N / \log^2 \log N$  for a sufficiently large constant  $C_1$ . We can preprocess an  $N \times d$  real-valued matrix  $A$  in  $O(N^{1+\delta})$  time, so that given a query real-valued  $d$ -dimensional vector  $x$ , we can compute the  $(\min, +)$ -product of  $A$  and  $x$  in  $O(N)$  expected time.*

Applying the above corollary  $N/d$  times yields:

► **Corollary 3.** *Let  $\delta > 0$  be any fixed constant. We can preprocess an  $N \times N$  real-valued matrix  $A$  in  $O(N^{2+\delta})$  time, so that given a query  $N$ -dimensional real-valued vector  $x$ , we can compute the  $(\min, +)$ -product of  $A$  and  $x$  in  $O((N^2 / \log^2 N) \log^2 \log N)$  expected time.*

A similar result was obtained by Williams [25] for online *Boolean* matrix-vector multiplication. Recently Larsen and Williams [19] have found a faster algorithm, in the Boolean case, but it is not combinatorial, requires amortization, and does not deal with the rectangular matrix case in Corollary 2.

In Appendix A.2, we further show how to reduce the  $O(K)$  term in Theorem 1 by about a logarithmic factor in the offline case, by modifying the algorithm to incorporate

<sup>4</sup> For any  $j_0 \in \{1, \dots, d\}$ , the key observation is that  $\min_{j=1}^d (a_{ij} + x_j) = a_{ij_0} + x_{j_0}$  iff  $(a_{ij_0} - a_{i1}, \dots, a_{ij_0} - a_{id})$  is dominated by  $(x_1 - x_{j_0}, \dots, x_d - x_{j_0})$  in  $\mathbb{R}^d$ .

bit-packing tricks. This has applications to speeding up combinatorial algorithms for  $(\min,+)$  matrix-matrix multiplication and all-pairs shortest paths.

In the full paper, we note that the method can be simplified in the Boolean case – the data structure becomes just an augmented, randomized variant of the *trie*. This has an application to combinatorial algorithms for Boolean matrix multiplication.

### 3 Offline Dominance Range Searching

In this section, we study the offline orthogonal range searching problem in the counting version (which includes the detection version), allowing the use of fast matrix multiplication. By doubling the dimension (footnote 1), it suffices to consider the dominance case: given  $n$  data/query points in  $\mathbb{R}^d$ , we want to count the number of data points dominated by each query point. We describe a black-box reduction of the real case to the Boolean case.

We use a recursion similar to a degree- $s$  range tree (which bears some resemblance to a low-dimensional algorithm from [13]).

**Algorithm.** Let  $\delta \in (0, 1)$  and  $s$  be parameters to be set later. Let  $[s]$  denote  $\{0, 1, \dots, s-1\}$ .

Given a set  $P$  of  $n \leq N$  data/query points in  $\mathbb{R}^j \times [s]^{d-j}$ , with  $d \leq c \log N$ , our algorithm is simple and proceeds as follows:

0. If  $j = 0$ , then all points are in  $[s]^d$  and we solve the problem directly by mapping each point  $(p_1, \dots, p_d)$  to a binary string  $1^{p_1}0^{s-p_1} \dots 1^{p_d}0^{s-p_d} \in \{0, 1\}^{ds}$  and running a known Boolean offline dominance algorithm in  $ds$  dimensions.
1. Otherwise, for each  $i \in [s]$ , recursively solve the problem for the subset  $P_i$  of all points in  $P$  with ranks from  $i(n/s) + 1$  to  $(i+1)(n/s)$  in the  $j$ -th coordinate.
2. “Round” the  $j$ -th coordinate values of all data points in  $P_i$  to  $i+1$  and all query points in  $P_i$  to  $i$ , and recursively solve the problem for  $P$  after rounding (which now lies in  $\mathbb{R}^{j-1} \times [s]^{d-j+1}$ ); add the results to the existing counts of all the query points.

**Analysis.** Suppose that the Boolean problem for  $n$  points in  $d \leq c \log n$  dimensions can be solved in  $d^C n^{2-f(c)}$  time for some absolute constant  $C \geq 1$  and some function  $f(c) \in [0, 1/4]$ . The following recurrence bounds the total cost of the leaves of the recursion in our algorithm (assuming that  $n$  is a power of  $s$ , for simplicity):

$$T_{d,j}(n) = sT_{d,j}(n/s) + T_{d,j-1}(n).$$

For the base cases,  $T_{d,j}(1) = 1$ ; and if  $n > \sqrt{N}$ , then  $T_{d,0}(n) \leq (ds)^C n^{2-f(2cs)}$  (since the Boolean subproblems have dimension  $ds \leq cs \log N \leq 2cs \log n$ ). On the other hand, if  $n \leq \sqrt{N}$ , we can use brute force to get  $T_{d,0}(n) \leq dn^2 \leq dn^{3/2}N^{1/4}$ . In any case,  $T_{d,0}(n) \leq (ds)^C n^{3/2}N^{1/2-f(2cs)} = An^{3/2}$  where we let  $A := (ds)^C N^{1/2-f(2cs)}$ .

One way<sup>5</sup> to solve this recurrence is again by “guessing”. We guess that

$$T_{d,j}(n) \leq (1 + \gamma)^j An^{3/2}$$

for some choice of parameter  $\gamma \in (0, 1)$  to be determined later. We verify the guess by induction.

<sup>5</sup> Since this particular recurrence is simple enough, an alternative, more direct way is to expand  $T_{d,d}(N)$  into a sum  $\sum_{i \geq 0} \binom{d+i}{i} s^i T_{d,0}(N/s^i) \leq \sum_{i \geq 0} O\left(\frac{d+i}{i\sqrt{s}}\right)^i \cdot AN^{3/2}$ , and observe that the maximum term occurs when  $i$  is near  $d/\sqrt{s}$ ...



The base cases are trivial. Assume that the guess is true for lexicographically smaller  $(j, n)$ . Then

$$\begin{aligned} T_{d,j}(n) &\leq (1 + \gamma)^j A s(n/s)^{3/2} + (1 + \gamma)^{j-1} A n^{3/2} \\ &= \left[ \frac{1}{\sqrt{s}} + \frac{1}{1 + \gamma} \right] (1 + \gamma)^j A n^{3/2} \leq (1 + \gamma)^j A n^{3/2}, \end{aligned}$$

provided that

$$\frac{1}{\sqrt{s}} + \frac{1}{1 + \gamma} \leq 1,$$

which is true by setting  $\gamma := 2/\sqrt{s}$ .

We can set  $s := c^4$ , for example. Then  $\gamma = O(1/c^2)$ . We conclude that

$$\begin{aligned} T_{d,d}(N) &\leq (1 + \gamma)^d A N^{3/2} \leq e^{\gamma d} (ds)^{O(1)} N^{2-f(2cs)} \\ &\leq (ds)^{O(1)} N^{2-f(2cs)+O(\gamma c)} \\ &= d^{O(1)} N^{2-f(2c^5)+O(1/c)}. \end{aligned}$$

Now,  $T_{d,d}(N)$  excludes the cost at internal nodes of the recursion. Since the recursion has depth at most  $\log_s N + d$ , the actual running time can be bounded by  $T_{d,d}(n)(d \log N)^{O(1)}$ .

Aboud, Williams, and Yu's algorithm [1] for the Boolean case, as derandomized by Chan and Williams [9], achieves  $f(c) = 1/O(\log c)$ , yielding an overall time bound of  $N^{2-1/O(\log c)}(d \log N)^{O(1)}$ , which is  $N^{2-1/O(\log c)}$  for  $\log c \ll \sqrt{\log N}$ .

► **Theorem 4.** *Let  $c \in [1, 2^{(1/C_1)\sqrt{\log N}}]$  for a sufficiently large constant  $C_1$ . Given  $N$  points in  $d = c \log N$  dimensions, we can answer  $N$  offline dominance range counting queries in  $N^{2-1/O(\log c)}$  time.*

We remark that if the Boolean problem could be solved in truly subquadratic time  $d^{O(1)}N^{2-\varepsilon}$ , then the above analysis (with  $s := (c \log N)^2$ , say) would imply that the general problem could be solved in truly subquadratic time with the *same*  $\varepsilon$ , up to  $(d \log N)^{O(1)}$  factors.

## 4 Approximate $\ell_\infty$ Nearest Neighbor Searching

In this section, we study (online, combinatorial) data structures for  $t$ -approximate  $\ell_\infty$  nearest neighbor search. By known reductions [16, 18], it suffices to solve the fixed-radius approximate decision problem, say, for radius  $r = 1/2$ : given a query point  $q$ , we want to find a data point of distance at most  $t/2$  from  $q$ , under the promise that the nearest neighbor distance is at most  $1/2$ .

Our solution closely follows Indyk's divide-and-conquer method [18], with a simple modification that incorporates a range-tree-like recursion.

**Data structure.** Let  $\delta \in (0, 1)$ ,  $\rho > 1$ , and  $c \geq 4$  be user-specified parameters. Let  $s$  and  $k$  be parameters to be chosen later.

Given a set  $P$  of  $n \leq N$  data points in  $d \leq c \log N$  dimensions, our data structure is constructed as follows:

0. If  $n \leq s$  or  $d = 0$ , then just store the points in  $P$ .

Otherwise, compute and store the median first coordinate  $\mu$  in  $P$ . Let  $P_{>i}$  (resp.  $P_{<i}$ ) denote the subset of all points in  $P$  with first coordinate greater than (resp. less than)  $\mu + i$ . Let  $\alpha_i := |P_{>i}|/n$  and  $\beta_i := |P_{<-i}|/n$ . Note that the  $\alpha_i$ 's and  $\beta_i$ 's are decreasing sequences with  $\alpha_0 = \beta_0 = 1/2$ .

## 27:10 Orthogonal Range Searching in Moderate Dimensions

1. If  $\alpha_k > 1/s$  and  $\alpha_{i+1} > \alpha_i^\rho$  for some  $i \in \{0, 1, \dots, k-1\}$ , then set  $type = (1, i)$  and recursively construct a data structure for  $P_{>i}$  and for  $P_{<i+1}$ .
2. Else if  $\beta_k > 1/s$  and  $\beta_{i+1} > \beta_i^\rho$  for some  $i \in \{0, 1, \dots, k-1\}$ , then set  $type = (2, i)$  and recursively construct a data structure for  $P_{<-i}$  and for  $P_{>-(i+1)}$ .
3. Else if  $\alpha_k, \beta_k \leq 1/s$ , then set  $type = 3$  and recursively construct a data structure for
  - the set  $P_{>k} \cup P_{<-k}$  and
  - the  $(d-1)$ -dimensional projection of  $P - (P_{>k+1} \cup P_{<-(k+1)})$  that drops the first coordinate (this recursion in  $d-1$  dimensions is where our algorithm differs from Indyk's).

We set  $k := \lceil \log_\rho \log s \rceil$ . Then one of the tests in steps 1–3 must be true. To see this, suppose that  $\alpha_k > 1/s$  (the scenario  $\beta_k > 1/s$  is symmetric), and suppose that  $i$  does not exist in step 1. Then  $\alpha_k \leq (1/2)^{\rho^k} \leq 1/s$ , a contradiction.

**Analysis.** The space usage is proportional to the number of points stored at the leaves in the recursion, which satisfies the following recurrence (by using the top expression with  $(\alpha, \alpha') = (\alpha_i, \alpha_{i+1})$  for step 1 or  $(\alpha, \alpha') = (\beta_i, \beta_{i+1})$  for step 2, or the bottom expression for step 3):

$$S_d(n) \leq \max \begin{cases} \max_{\alpha, \alpha': \alpha' > \alpha^\rho, 1/s < \alpha' \leq \alpha \leq 1/2} [S_d(\alpha n) + S_d((1-\alpha')n)] \\ S_d(2n/s) + S_{d-1}(n), \end{cases} \quad (3)$$

with  $S_d(n) = n$  for the base case  $n \leq s$  or  $d = 0$ .

We guess that

$$S_d(n) \leq (1 + \gamma)^d n^\rho$$

for some choice of parameter  $\gamma \in (0, 1)$ . We verify the guess by induction.

The base case is trivial. Assume that the guess is true for lexicographically smaller  $(d, n)$ .

- **Case I:** the maximum in (3) is attained by the top expression and by  $\alpha, \alpha'$ . Then

$$\begin{aligned} S_d(n) &\leq (1 + \gamma)^d [( \alpha n )^\rho + ((1 - \alpha')n)^\rho] \\ &\leq [\alpha^\rho + 1 - \alpha'] (1 + \gamma)^d n^\rho \\ &\leq (1 + \gamma)^d n^\rho. \end{aligned}$$

- **Case II:** the maximum in (3) is attained by the bottom expression. Then

$$\begin{aligned} S_d(n) &\leq (1 + \gamma)^d (2n/s)^\rho + (1 + \gamma)^{d-1} n^\rho \\ &\leq \left[ \left( \frac{2}{s} \right)^\rho + \frac{1}{1 + \gamma} \right] (1 + \gamma)^d n^\rho \\ &\leq (1 + \gamma)^d n^\rho \end{aligned}$$

by setting  $s := 2(2/\gamma)^{1/\rho}$ .

Set  $\gamma := \delta/c$ . Then  $s = O((c/\delta)^{1/\rho})$  and  $k = \log_\rho \log(c/\delta) + O(1)$ . We conclude that

$$S_d(N) \leq e^{\gamma d} N^\rho \leq N^{\rho + O(\gamma c)} = N^{\rho + O(\delta)}.$$

For the preprocessing time, observe that the depth of the recursion is  $h := O(\log_{s/(s-1)} N + d)$  (since at each recursive step, the size of the subsets drops by a factor of  $1 - 1/s$  or the dimension decreases by 1). Now,  $h = O(s \log N + d) \leq O((c/\delta) \log N + d) = O((c/\delta) \log N)$ . Hence, the preprocessing time can be bounded by  $O(S_d(N)h) = O((c/\delta) N^{\rho + \delta} \log N)$ .

**Query algorithm.** Given the preprocessed set  $P$  and a query point  $q = (q_1, \dots, q_d)$ , our query algorithm proceeds as follows:

0. If  $n \leq s$  or  $d = 0$ , then answer the query directly by brute-force search.
1. If  $type = (1, i)$ : if  $q_1 > i + 1/2$ , then recursively answer the query in  $P_{>i}$ , else recursively answer the query in  $P_{<i+1}$ .
2. If  $type = (2, i)$ : proceed symmetrically.
3. If  $type = 3$ :
  - if  $q_1 > k + 1/2$  or  $q_1 < -(k + 1/2)$ , then recursively answer the query in  $P_{>k} \cup P_{<-k}$ ;
  - else recursively answer the query in  $P - (P_{>k+1} \cup P_{<-(k+1)})$ , after dropping the first coordinate of  $q$ .

Note that in the last subcase of step 3, any returned point has distance at most  $2k + 3/2$  from  $q$  in terms of the first coordinate. By induction, the approximation factor  $t$  is at most  $4k + 3 = O(\log_\rho \log(c/\delta))$ .

**Analysis.** The query time is clearly bounded by the depth  $h$ , which is  $O((c/\delta) \log N)$ .

► **Theorem 5.** *Let  $\delta > 0$  be any fixed constant. Let  $\rho > 1$  and  $c \geq \Omega(1)$ . Given  $N$  points in  $d = c \log N$  dimensions, we can construct a data structure in  $O(dN^{\rho+\delta})$  time and  $O(dN + N^{\rho+\delta})$  space, so that we can handle the fixed-radius decision version of approximate  $\ell_\infty$  nearest neighbor queries in  $O(d)$  time with approximation factor  $O(\log_\rho \log c)$ .*

---

## References

- 1 Amir Abboud, Ryan Williams, and Huacheng Yu. More applications of the polynomial method to algorithm design. In *Proc. 26th ACM–SIAM Symp. Discrete Algorithms (SODA)*, pages 218–230, 2015.
- 2 Peyman Afshani, Timothy M. Chan, and Konstantinos Tsakalidis. Deterministic rectangle enclosure and offline dominance reporting on the RAM. In *Proc. 41st Int'l Colloq. Automata, Languages, and Programming (ICALP), Part I*, pages 77–88, 2014.
- 3 Josh Alman, Timothy M. Chan, and Ryan Williams. Polynomial representation of threshold functions with applications. In *Proc. 57th IEEE Symp. Found. Comput. Sci. (FOCS)*, pages 467–476, 2016.
- 4 Josh Alman and Ryan Williams. Probabilistic polynomials and Hamming nearest neighbors. In *Proc. 56th IEEE Symp. Found. Comput. Sci. (FOCS)*, pages 136–150, 2015.
- 5 Alexandr Andoni, Dorian Croitoru, and Mihai M. Pătraşcu. Hardness of nearest neighbor under  $L_\infty$ . In *Proc. 49th IEEE Symp. Found. Comput. Sci. (FOCS)*, pages 424–433, 2008.
- 6 V. Z. Arlazarov, E. A. Dinic, M. A. Kronrod, and I. A. Faradzhev. On economical construction of the transitive closure of a directed graph. *Soviet Mathematics Doklady*, 11:1209–1210, 1970.
- 7 Timothy M. Chan. Geometric applications of a randomized optimization technique. *Discrete Comput. Geom.*, 22(4):547–567, 1999.
- 8 Timothy M. Chan. Speeding up the Four Russians algorithm by about one more logarithmic factor. In *Proc. 26th ACM–SIAM Symp. Discrete Algorithms (SODA)*, pages 212–217, 2015.
- 9 Timothy M. Chan and Ryan Williams. Deterministic APSP, orthogonal vectors, and more: Quickly derandomizing Razborov–Smolensky. In *Proc. 27th ACM–SIAM Symp. Discrete Algorithms (SODA)*, pages 1246–1255, 2016.
- 10 T. M. Chan. All-pairs shortest paths with real weights in  $O(n^3/\log n)$  time. *Algorithmica*, 50:236–243, 2008.

- 11 T. M. Chan. More algorithms for all-pairs shortest paths in weighted graphs. *SIAM J. Comput.*, 39:2075–2089, 2010.
- 12 T. M. Chan, K. G. Larsen, and M. Pătraşcu. Orthogonal range searching on the RAM, revisited. In *Proc. 27th ACM Symp. Comput. Geom. (SoCG)*, pages 1–10, 2011.
- 13 T. M. Chan and M. Pătraşcu. Counting inversions, offline orthogonal range counting, and related problems. In *Proc. 21st ACM–SIAM Symp. Discrete Algorithms (SODA)*, pages 161–173, 2010.
- 14 Daniel M. Gordon, Oren Patashnik, Greg Kuperberg, and Joel Spencer. Asymptotically optimal covering designs. *J. Combinatorial Theory, Series A*, 75(2):270–280, 1996.
- 15 Y. Han and T. Takaoka. An  $O(n^3 \log \log n / \log^2 n)$  time algorithm for all pairs shortest paths. In *Proc. 13th Scand. Symp. and Workshops on Algorithm Theory (SWAT)*, pages 131–141, 2012.
- 16 Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani. Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory Comput.*, 8(1):321–350, 2012.
- 17 R. Impagliazzo, S. Lovett, R. Paturi, and S. Schneider. 0-1 integer linear programming with a linear number of constraints, 2014.
- 18 Piotr Indyk. On approximate nearest neighbors under  $l_\infty$  norm. *J. Comput. Sys. Sci.*, 63(4):627–638, 2001.
- 19 Kasper Green Larsen and Ryan Williams. Faster online matrix-vector multiplication. In *Proc. 28th ACM–SIAM Symp. Discrete Algorithms (SODA)*, pages 2182–2189, 2017.
- 20 François Le Gall. Faster algorithms for rectangular matrix multiplication. In *Proc. 53rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 514–523, 2012.
- 21 Jirí Matoušek. Computing dominances in  $E^n$ . *Inform. Process. Lett.*, 38(5):277–278, 1991.
- 22 F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, 1985.
- 23 R. Williams. Faster all-pairs shortest paths via circuit complexity. In *Proc. 46th ACM Symp. Theory Comput. (STOC)*, pages 664–673, 2014.
- 24 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005.
- 25 Ryan Williams. Matrix-vector multiplication in sub-quadratic time (some preprocessing required). In *Proc. 18th ACM–SIAM Symp. Discrete Algorithms (SODA)*, pages 995–1001, 2007. URL: <http://dl.acm.org/citation.cfm?id=1283383.1283490>.
- 26 Huacheng Yu. An improved combinatorial algorithm for Boolean matrix multiplication. In *Proc. 42nd Int'l Colloq. Automata, Languages, and Programming (ICALP), Part I*, pages 1094–1105, 2015.

## **A** Online Dominance Range Searching (Continued)

### A.1 Slightly Improved Version

We now describe a small improvement to the data structure in Section 2. The idea is to replace  $\mathcal{J}$  with a collection of slightly larger subsets, but with fewer subsets, so that any set  $J_q$  of size  $t := \lfloor d/b \rfloor$  is covered by some subset in  $J \in \mathcal{J}$ . Such a collection is called a *covering design* (e.g., see [14]), which can be constructed easily by random sampling, as explained in see part (i) of the lemma below. In our application, we also need a good time bound for finding such a  $J \in \mathcal{J}$  for a given query set  $J_q$ ; this is addressed in part (ii) of the lemma. (Proofs are deferred to the full paper.)

► **Lemma 6** (Covering designs). *Given numbers  $v \geq k \geq t$  and  $N$ , and given a size- $v$  ground set  $V$ ,*

- (i) *we can construct a collection  $\mathcal{J}$  of at most  $\binom{v}{t} / \binom{k}{t} \ln N$  size- $k$  subsets of  $V$  in  $O(v|\mathcal{J}|)$  time, so that given any query size- $t$  subset  $J_q \subset V$ , we can find a subset  $J \in \mathcal{J}$  containing  $J_q$  in  $O(v|\mathcal{J}|)$  time with success probability at least  $1 - 1/N$ ;*
- (ii) *alternatively, with a larger collection  $\mathcal{J}$  of at most  $\binom{v}{t} / \binom{k}{t}^2 \ln^2(vN)$  subsets, we can reduce the query time to  $O(v^3 \log^2(vN))$ .*

We now modify the data structure in Section 2 as follows. In step 1, we change  $\mathcal{J}$  to a collection of size- $\lfloor d/2 \rfloor$  subsets of  $\{1, \dots, d\}$  obtained from Lemma 6(ii) with  $(v, k, t) = (d, \lfloor d/2 \rfloor, \lfloor d/b \rfloor)$ . Then  $|\mathcal{J}| \leq \left( \binom{d}{\lfloor d/b \rfloor} / \binom{\lfloor d/2 \rfloor}{\lfloor d/b \rfloor} \right)^2 \ln^2(dN) \leq 2^{O(d/b)} \log^2 N$ . The recurrence for the preprocessing time and space then improves to

$$T_d(n) \leq T_d(\lfloor \alpha n \rfloor) + T_d(\lfloor (1 - \alpha)n \rfloor) + (2^{O(d/b)} \log^2 N) T_{\lfloor d/b \rfloor}(n) + O(n),$$

which solves to  $T_d(N) \leq 2^{O(d/b+d/b^2+\dots)} N (\log_{1/(1-\alpha)} N)^{O(\log_b d)} \leq N^{1+O(\delta)} 2^{O(\log^2((1/\alpha) \log N))}$ , this time by setting  $b := c/\delta$  (instead of  $b := (c/\delta) \log(c/\delta)$ ).

In the query algorithm, we modify step 1 by finding a set  $J \in \mathcal{J}$  containing  $J_q$  by Lemma 6(ii) and recursively querying  $P_J$  (instead of  $P_{J_q}$ ). If no such  $J$  exists, we can afford to switch to brute-force search, since this happens with probability less than  $1/N$ . The analysis of the recurrence for  $Q_d(N)$  remains the same. Each internal node of the recursion now has cost  $O(d^3 \log^2 N)$  by Lemma 6(ii); the extra factor will not affect the final bound. The overall query time is still  $N^{1-1/O(b \log b)}$ , which is now  $N^{1-1/O((c/\delta) \log(c/\delta))}$ .

## A.2 Offline Packed-Output Version, with Application to APSP

In this subsection, we discuss how to refine the algorithm in Section 2, so that the output can be reported in roughly  $O(K/\log n)$  time instead of  $O(K)$  in the offline setting. The approach is to combine the algorithm with bit-packing tricks.

We assume a  $w$ -bit word RAM model which allows for certain exotic word operations. In the case of  $w := \delta_0 \log N$  for a sufficiently small constant  $\delta_0 > 0$ , exotic operations can be simulated in constant time by table lookup; the precomputation of the tables requires only  $N^{O(\delta_0)}$  time.

We begin with techniques to represent and manipulate sparse sets of integers in the word RAM model. Let  $z$  be a parameter to be set later. In what follows, an interval  $[a, b]$  refers to the integer set  $\{a, a + 1, \dots, b - 1\}$ . A *block* refers to an interval of the form  $[kz, (k + 1)z)$ . Given a set  $S$  of integers over an interval  $I$  of length  $n$ , we define its *compressed representation* to be a doubly linked list of *mini-sets*, where for each of the  $O(\lceil n/z \rceil)$  blocks  $B$  intersecting  $I$  (in sorted order), we store the *mini-set*  $\{j \bmod z : j \in S \cap B\}$ , which consists of small  $(\log z)$ -bit numbers and can be packed in  $O((|S \cap B|/w) \log z + 1)$  words. The total number of words in the compressed representation is  $O((|S|/w) \log z + n/z + 1)$ . Proofs of the following facts can be found in the full paper.

► **Lemma 7** (Bit-packing tricks).

- (i) *Given compressed representations of two sets  $S_1$  and  $S_2$  over two disjoint intervals, we can compute the compressed representation of  $S_1 \cup S_2$  in  $O(1)$  time.*
- (ii) *Given compressed representations of  $S_0, \dots, S_{m-1} \subset [0, n)$ , we can compute the compressed representations of  $T_0, \dots, T_{n-1} \subset [0, m)$  with  $T_j = \{i : j \in S_i\}$  (called the transposition of  $S_0, \dots, S_{m-1}$ ), in  $O((K/w) \log^2 z + mn/z + m + n + z)$  time, where  $K = \sum_{i=0}^{m-1} |S_i|$ .*

(iii) Given compressed representations of  $S_0, \dots, S_{m-1} \subset [0, n)$  and a bijective function  $\pi : [0, n) \rightarrow [0, n)$  which is evaluable in constant time, we can compute compressed representations of  $\pi(S_1), \dots, \pi(S_m)$  in  $O((K/w) \log^2 z + mn/z + m + n + z)$  time, where  $K = \sum_{i=0}^{m-1} |S_i|$ .

► **Theorem 8.** Assume  $z \leq N^{o(1)}$ . Let  $\delta > 0$  be any fixed constant and  $c \in [C_1, (1/C_1) \log N / \log^2 \log N]$  for a sufficiently large constant  $C_1$ . Given a set  $P$  of  $N$  points in  $d = c \log N$  dimensions, we can construct a data structure in  $O(N^{1+\delta})$  preprocessing time and space, so that we can answer  $N$  offline dominance range reporting queries (with a compressed output representation) in  $N^{2-1/O(c \log c)} + O(((K/w) \log^2 z + N^2/z) \log d)$  time where  $K$  is the total number of reported points over the  $N$  queries.

**Proof.** We adapt the preprocessing and query algorithm in Section 2, with the improvement from Appendix A.1. A *numbering* of a set  $S$  of  $n$  elements refers to a bijection from  $S$  to  $n$  consecutive integers. For each point set  $P$  generated by the preprocessing algorithm, we define a numbering  $\phi_P$  of  $P$  simply by recursively “concatenating” the numberings  $\phi_{P_L}$  and  $\phi_{P_R}$  and appending  $p(i^*)$ . The output to each query for  $P$  will be a compressed representation of the subset of dominated points after applying  $\phi_P$ .

In step 2 of the query algorithm, we can union the output for  $P_L$  and for  $P_R$  in  $O(1)$  time by Lemma 7(i). In step 1 of the query algorithm, we need additional work since the output is with respect to a different numbering  $\phi_{P_J}$ , for some set  $J \in \mathcal{J}$ . For each  $J \in \mathcal{J}$ , we can change the compressed representation to follow the numbering  $\phi_P$  by invoking Lemma 7(iii), after collecting all query points  $Q(P_J)$  that are passed to  $P_J$  (since queries are offline). To account for the cost of this invocation to Lemma 7(iii), we charge (a)  $(1/w) \log^2 z$  units to each output feature, (b)  $1/z$  units to each point pair in  $P_J \times Q(P_J)$ , (c) 1 unit to each point in  $P_J$ , and (d) 1 unit to each point in  $Q(P_J)$ , and (e)  $z$  units to the point set  $P_J$  itself.

Each output feature or each point pair is charged  $O(\log d)$  times, since  $d$  decreases to  $\lfloor d/2 \rfloor$  with each charge. Thus, the total cost for (a) and (b) is  $O((K/w) \log^2 z \log d + (N^2/z) \log d)$ . The total cost of (c) is  $N^{1+o(1)}$  by the analysis of our original preprocessing algorithm; similarly, the total cost of (e) is  $zN^{1+o(1)}$ . The total cost of (d) is  $N^{2-1/O(c \log c)}$  by the analysis of our original query algorithm.

We can make the final compressed representations to be with respect to any user-specified numbering of  $P$ , by one last invocation to Lemma 7(iii). The algorithm can be derandomized, as noted in the full paper. ◀

One may wonder whether the previous range-tree-like offline algorithm by Impagliazzo et al. [17, 8] could also be adapted; the problem there is that  $d$  is only decremented rather than halved, which makes the cost of re-numbering too large.

The main application is to  $(\min, +)$  matrix multiplication and all-pairs shortest paths (APSP). The corollary below follows immediately from a simple reduction [10] (see footnote 4) to  $d$  instances of  $d$ -dimensional offline dominance range reporting where the total output size  $K$  is  $O(n^2)$ . Here, we set  $w := \delta_0 \log N$  and  $z := \text{poly}(\log N)$ .

► **Corollary 9.** Let  $d = (1/C_1) \log^2 N / \log^2 \log N$  for a sufficiently large constant  $C_1$ . Given an  $N \times d$  and a  $d \times N$  real-valued matrix, we can compute their  $(\min, +)$ -product (with a compressed output representation) in  $O((N^2 / \log N) \log^3 \log N)$  expected time.

The corollary below follows from applying Corollary 9  $q/d$  times, in conjunction with a subroutine by Chan [11, Corollary 2.5]. (The result improves [11, Corollary 2.6].)

► **Corollary 10.** Let  $q = \log^3 N / \log^5 \log N$ . Given an  $N \times q$  and a  $q \times N$  real-valued matrix, we can compute their  $(\min, +)$ -product in  $O(N^2)$  time.

Applying Corollary 10  $N/q$  times (and using a standard reduction from APSP to  $(\min,+)$ -multiplication), we obtain:

► **Corollary 11.** *Given two  $N \times N$  real-valued matrices, we can compute their  $(\min,+)$ -product by a combinatorial algorithm in  $O((N^3/\log^3 N) \log^5 \log N)$  time. Consequently, we obtain a combinatorial algorithm for APSP for arbitrary  $N$ -vertex real-weighted graphs with the same time bound.*

Note that Williams' algorithm [23] is faster (achieving  $N^3/2^{\Omega(\sqrt{\log N})}$  time), but is non-combinatorial and gives a worse time bound ( $O(N^2 \log^{O(1)} N)$ ) for the rectangular matrix case in Corollary 10.