

Reconstructing Random Graphs from Distance Queries

Michael Krivelevich 

School of Mathematical Sciences, Tel Aviv University, Israel

Maksim Zhukovskii 

School of Computer Science, The University of Sheffield, UK

Abstract

We estimate the minimum number of distance queries that is sufficient to reconstruct the binomial random graph $G(n, p)$ with constant diameter with high probability. We get a tight (up to a constant factor) answer for all $p > n^{-1+o(1)}$ outside “threshold windows” around $n^{-k/(k+1)+o(1)}$, $k \in \mathbb{Z}_{>0}$: with high probability the query complexity equals $\Theta(n^{4-d}p^{2-d})$, where d is the diameter of the random graph. This demonstrates the following non-monotone behaviour: the query complexity jumps down at moments when the diameter gets larger; yet, between these moments the query complexity grows. We also show that there exists a non-adaptive algorithm that reconstructs the random graph with $O(n^{4-d}p^{2-d} \ln n)$ distance queries with high probability, and this is best possible.

2012 ACM Subject Classification Mathematics of computing \rightarrow Random graphs; Mathematics of computing \rightarrow Graph algorithms

Keywords and phrases random graphs, graph reconstruction, distance queries, query complexity

Digital Object Identifier 10.4230/LIPIcs.ESA.2025.30

Related Version *Full Version*: <https://arxiv.org/pdf/2404.18318>

Funding *Michael Krivelevich*: Research supported in part by NSF-BSF grant 2023688.

1 Introduction

Reconstruction of graphs was thoroughly studied in many different contexts and has various applications, e.g., in DNA sequencing [2, 13, 34], phylogenetics [1, 23, 29], and recovering neural networks [39]. A vast amount of literature is devoted to the average-case problem, i.e. to the reconstruction of random graphs [8, 12, 19, 21, 28, 31, 33, 35, 41]. In particular, the conjecture of Kelly and Ulam [24, 42] which is considered as one of the main reconstruction challenges is known to be true in binomial random graphs whp¹ [8, 28, 35].

Let G be a connected graph on $[n] := \{1, \dots, n\}$, and let $\mathcal{Q} \subset \binom{[n]}{2}$ be a set of pairs of vertices. In the usual way, we denote by $d_G(x, y)$ the graph distance between x and y in G . Let us say that \mathcal{Q} *reconstructs* G , if G is the only graph on $[n]$ with distances $d_G(x, y)$ between $\{x, y\} \in \mathcal{Q}$, i.e. for every graph G' on $[n]$ such that $d_{G'}(x, y) = d_G(x, y)$ for all $\{x, y\} \in \mathcal{Q}$, we have $G' = G$. We will also call the graph \mathcal{Q} -*reconstructible* in this case.

In this paper, we study the model of reconstruction with an access to distance query oracle introduced in [5], although its variants were considered long before that for modelling reconstruction of a phylogenetic tree [18, 25, 43]. For every queried pair of vertices, the oracle answers the distance between these vertices in the hidden unknown graph G . For a queried pair of vertices $\{x, y\}$, we denote the oracle’s answer to this query by $d(x, y)$, omitting the unknown graph in the subscript. An adaptive algorithm, at every step, selects the next query (a pair of vertices) based on the responses from the oracle to earlier queries.

¹ With high probability, that is, with probability tending to 1 as $n \rightarrow \infty$.



Let $q \in \mathbb{N}$. If there exists an adaptive algorithm A such that, for the hidden input graph G on $[n]$, it queries at most q pairs of vertices comprising a set $\mathcal{Q} \subset \binom{[n]}{2}$ that reconstructs G , then we call G q -reconstructible by A . Let us call the minimum q such that G is q -reconstructible by some algorithm the (distance) query complexity of G . It is easy to see that, in the worst case, the query complexity equals $\binom{n}{2}$: all pairs should be queried in order to reconstruct the graph. For example, this is the case when G is a clique. Reyzin and Srivastava [37] showed that, even for some trees, $\Omega(n^2)$ pairs are required to query. On the other hand, Mathieu and Zhou [31] presented an algorithm such that whp a uniformly random d -regular graph on $[n]$ is $(\lfloor n \ln^2 n \rfloor + dn/2)$ -reconstructible by this algorithm, and wondered whether methods similar to those developed in their paper are applicable to sparse binomial random graphs $G(n, p)$. In this paper, in contrast, we study the distance query complexity of relatively dense random graphs $G(n, p)$ assuming $p > n^{-1+\varepsilon}$ for some $\varepsilon > 0$, i.e., when the random graph has bounded diameter whp. Although our main motivation lies in attempting to achieve tight bounds for the average-case complexity in a situation where the general setup is hard to analyse, additional interest in this problem is sparked by a surprising phenomenon, which will be discussed in the next two paragraphs.

One can expect that dense graphs are typically harder to recover through distance queries, as the influence of a single edge to be recovered on distances between vertices might be less pronounced and thus harder to detect. In order to get some intuition of how the query complexity of the random graph evolves, let us first assume that p is significantly above the threshold for the property of having diameter at most 2, namely, $p^2 n - 2 \ln n - \ln \ln n \rightarrow \infty$ as $n \rightarrow \infty$. In this case, whp every pair of non-adjacent vertices has at least two common neighbours. Indeed, probability that a fixed pair of vertices has at most one common neighbour equals

$$\mathbb{P}(\text{Bin}(n-2, p^2) \leq 1) = (1-p^2)^{n-2} + (n-2)p^2(1-p^2)^{n-3} = O(np^2 e^{-np^2}) = o(n^{-2}).$$

Therefore, by the union bound, every non-adjacent pair of vertices of $G(n, p)$ has at least two common neighbours whp. Then whp the query complexity equals $\binom{n}{2}$. Indeed, for any set \mathcal{Q} of pairs of vertices of size less than $\binom{n}{2}$, changing any single adjacency relation of a pair from $\binom{[n]}{2} \setminus \mathcal{Q}$ does not influence the distance between x and y for any $\{x, y\} \in \mathcal{Q}$. Therefore, any set \mathcal{Q} that does not contain all pairs does not reconstruct $G(n, p)$ whp.

On the other hand, it seems likely that for p close to the connectivity threshold $\ln n/n$, the query complexity might be much less – quasilinear, as in the case of random regular graphs. Note that, in general the query complexity does not decrease as the graph becomes sparser since, in particular, as we mention above, for certain trees, the query complexity equals $\Omega(n^2)$. Nevertheless, we show that, on average, as density decreases, the query complexity jumps down at specific moments around hitting times of diameter's increments. However, between these moments the query complexity grows quite rapidly. In particular, there exist $n^{-1+\varepsilon} \ll p_1 \ll p_2 \ll n^{-1/2}$ such that whp the diameter of $G(n, p_1)$ is larger than the diameter of $G(n, p_2)$ while the query complexity of the former is also larger.

Related work

Kannan, Mathieu, and Zhou [22, 30] presented a reconstruction algorithm that uses $\tilde{O}(n^{3/2})$ distance queries for bounded degree graphs. They also proved an information-theoretic lower bound $\Omega(n \log n / \log \log n)$ for trees with maximum degree 3 and asked whether $\tilde{O}(n)$ is achievable for all bounded degree graphs. Tight results for certain families of bounded degree graphs – trees, chordal graphs, and graphs with bounded treelength – were recently obtained

by Bastide and Groenland in [4]. In particular, they proved that for every $\Delta \geq 3$, any randomised algorithm requires $\Omega(n \log n)$ queries to reconstruct n -vertex trees of maximum degree Δ for a certain sequence of sizes of trees, which is tight since there exists a deterministic algorithm to reconstruct such trees using $O(n \log n)$ queries. In [31], Mathieu and Zhou answered the question from [22] for random d -regular graphs with constant d : whp $n \ln^2 n$ distance queries is enough to reconstruct a uniformly random d -regular graph.

A modified version of the problem where the hidden graph is a tree and it is only allowed to query pairs of leaves (that are known) comes from biology: it serves as a model of reconstructing evolutionary (phylogenetic) trees and was introduced by Waterman, Smith, Singh, and Beyer [43]. This problem is well studied: lower bounds on query complexity for deterministic and randomised algorithms were obtained in [25] and [4] respectively, and upper bounds were investigated in [11, 23, 25, 27].

Other types of query oracles explored in the literature include *all-shortest-path*, *all-distances*, and *shortest-path* queries. The first two oracles answer all shortest paths [5, 38] and all distances [5, 14] from a queried vertex respectively. The shortest-path oracle answers a shortest path between a queried pair of vertices [22, 37]. Erlebach, Hall, and Mihal'ák [15] studied reconstruction of $G(n, p)$ from all-shortest-path queries for $p = \text{const}$ and $p = n^{-1+\varepsilon}$ for specific $\varepsilon > 0$.

Mathieu and Zhou [31] also showed that their construction of the set of distance queries can be used to get an upper bound for the metric dimension of random d -regular graphs. Let us recall that the *metric dimension* $\beta(G)$ of a graph G is the minimum cardinality of $S \subset V(G)$ such that all vertices of G have different vectors of distances to the vertices from S . Such *resolving sets* S or their variants were used, in particular, to label canonically vertices of random graphs [3, 26]. The metric dimension of $G(n, p)$ was studied by Bollobás, Mitsche, and Pralat in [10]. In particular, they proved that the metric dimension of $G(n, p)$ undergoes the following “zigzag” behaviour: if $p = n^{-\alpha}$, then $\log_n \beta(G(n, p)) = 1 - (1 - \alpha) \lfloor 1/(1 - \alpha) \rfloor$ whp; see also [36] for related results. Although in our proof of the upper bound for the distance query complexity of $G(n, p)$ we use a resolving set whose size is far from being optimal, as we show below the “jumps” of the metric dimension and of the query complexity are synchronous – both happen around $n^{-k/(k+1)}$, $k \in \mathbb{Z}_{>0}$.

Our contribution

As follows from the next result, the query complexity drops for the first time when p passes a threshold $n^{-1/2+o(1)}$ on its way down. More generally, we prove that, for every integer constant $k \geq 1$, as soon as p passes a threshold $n^{-(k+1)/(k+2)+o(1)}$, the query complexity drops by a factor of $\Theta(np)$.

► **Theorem 1.** *Let $k \geq 1$ be a fixed integer, $\varepsilon > 0$, $n^{-1+\varepsilon} < p < n^{-k/(k+1)-\varepsilon}$, and let $G_n \sim G(n, p)$. Then there exist $C = C(k, \varepsilon) > 0$ and an adaptive algorithm **A** such that, for $q = \lfloor C/(n^{k-2}p^k) \rfloor$, whp G_n is q -reconstructible by **A**.*

We prove Theorem 1 in Section 3.

► **Theorem 2.** *Let $k \geq 1$ be a fixed integer, $\varepsilon > 0$, and let $n^{k+1}p^{k+2} - 2 \ln n \rightarrow \infty$ as $n \rightarrow \infty$. Then whp $G_n \sim G(n, p)$ has query complexity at least $\frac{1}{2(k+1+\varepsilon)n^{k-2}p^k}$.*

Theorem 2 is proven in Section 5. Let us emphasise that its proof is based on an explicit argument which is more efficient than the usual information-theoretic approach: we show a lower bound on the query complexity of a deterministic graph in terms of its diameter and maximum degree, see Claim 14. Unlike the information-theoretic approach, our approach


30:4 Reconstructing Random Graphs from Distance Queries

allows to bound the query complexity of any given graph. So, it actually demonstrates that, when the number of queries is less than the bound, even knowing the input graph G does not help choosing the optimal querying strategy – no matter which pairs of vertices are queried, there is always a graph $G' \neq G$ with exactly the same answers.

Note that, for any $w_n \rightarrow \infty$ as $n \rightarrow \infty$ and any

$$\left(\frac{2 \ln n + w_n}{n^{k+1}} \right)^{1/(k+2)} < p < n^{-k/(k+1)-\varepsilon}, \quad (1)$$

the bounds in Theorem 1 and Theorem 2 differ by a constant factor, see Figure 1. So, roughly speaking (ignoring the constant factor) the query complexity in the interval (1) increases with p .



■ **Figure 1** Upper and lower bounds for query complexity of $G(n, p)$ for p satisfying (1).

We also prove tight (up to a constant factor) bounds for “non-adaptive reconstructibility”. It turns out that if the set of queried pairs \mathcal{Q} is fixed, then both upper and lower bounds on the minimum $|\mathcal{Q}|$ such that $G(n, p)$ is \mathcal{Q} -reconstructible increase by a factor of $\log n$ when p satisfies (1).

► **Theorem 3.** *Let $k \geq 1$ be a fixed integer, $\varepsilon > 0$, and let $p > n^{-1+\varepsilon}$. Then there exists a set $\mathcal{Q} = \mathcal{Q}(n) \subset \binom{[n]}{2}$ of size at most $\frac{3}{n^{k-2}p^k} e^{n^k p^{k+1}} \ln n$ such that whp $G_n \sim G(n, p)$ is \mathcal{Q} -reconstructible.*

Theorem 3 is proven in Section 4.

► **Theorem 4.** *Let $k \geq 1$ be a fixed integer, $\varepsilon > 0$, and let $w_n \rightarrow \infty$ as $n \rightarrow \infty$. If*

$$\left(\frac{2 \ln n + w_n}{n^{k+1}} \right)^{1/(k+2)} < p < \left(\left(\frac{1}{k+1} - \varepsilon \right) \frac{\ln n}{n^k} \right)^{1/(k+1)},$$

then there is no set $\mathcal{Q} = \mathcal{Q}(n)$ of size at most $\frac{1}{8(k+1)^2 n^{k-2} p^k} e^{n^k p^{k+1}} \ln n$ such that whp $G_n \sim G(n, p)$ is \mathcal{Q} -reconstructible.

Theorem 4 is proven in Section 6.

Proofs outline

For simplicity of presentation, let us assume that $p = n^{-\alpha}$, where $\frac{k}{k+1} < \alpha < \frac{k+1}{k+2}$ for some $k \in \mathbb{Z}_{>0}$. The crucial observation that allows to prove tight (up to a constant) upper bounds both for adaptive and non-adaptive query complexities is the following: if $\{v_1, x\}$ and $\{v_2, x\}$ are queried and $d(v_2, x) \geq d(v_1, x) + 2$, then v_1 and v_2 are not adjacent.

Although whp the diameter of $G_n \sim G(n, p)$ equals $k + 2$, for a sufficiently large fixed set $X \subset [n]$, every vertex $v \in [n]$ has a vertex from X in its k -neighbourhood. Indeed, if X has size $\Theta\left(\frac{\ln n}{n^{k-1}p^k}\right)$, then the expected number of k -paths from a fixed vertex $v \in [n]$ to X is $\Theta(\ln n)$. Thus, using standard concentration inequalities, it is possible to overcome the union bound over the choice of $v \in [n]$ for $|X| = C \frac{\ln n}{n^{k-1}p^k}$ and C sufficiently large. Actually, the following stronger property can be proved: whp for every $v_1, v_2 \notin X$, there exists $x \in X$ such that $d_{G_n}(v_1, x) \leq k$ while $d_{G_n \setminus \{v_1, v_2\}}(v_2, x) \geq k + 2$, where $G_n \setminus \{v_1, v_2\}$ is a spanning subgraph of G_n that is obtained by excluding the edge $\{v_1, v_2\}$. Thus, as soon as all pairs of vertices with at least one vertex in X are queried, all the remaining adjacencies are reconstructible whp. Indeed, for non-adjacent v_1, v_2 , there exists x such that $d(v_1, x) \leq k$ and $d(v_2, x) \geq k + 2$, thus we may reconstruct all non-adjacencies. All the other pairs of vertices must be adjacent since otherwise if we could exclude, say, the edge $\{v_1, v_2\}$, then there is an $x \in X$ such that $d(v_2, x) \leq k + 1$ while, in the corrupted graph, the distance between v_2 and x would be at least $k + 2$. This proves Theorem 3, see details in Section 4. A counterpart of this very simple algorithm was employed by Mathieu and Zhou in [31], who showed that it allows typically to reconstruct random regular graphs with $(1 + o(1))n \ln^2 n$ queries, which is tight, up to a $(\log n)^{1+o(1)}$ -factor.

The proof of Theorem 1 in Section 3 is based on similar ideas but also utilises the possibility of updating queries adaptively: indeed, instead of querying all pairs $\{x \in X, v \in [n]\}$ simultaneously, we may split the set X into $\nu = \Theta(\log n)$ equal parts X_1, \dots, X_ν and then, for every $i = 1, \dots, \nu$, query a pair $\{x \in X_i, v \in [n]\}$ only if v has not yet been excluded from the consideration. A vertex v is excluded from the consideration at time i , if it has sufficiently many (at least M for some large enough constant M) vertices in the intersection of its k -neighbourhood with $X_1 \cup \dots \cup X_i$. This allows to query the desired number of pairs whp. In order to complete the proof, it remains to observe that whp there are no v_1, v_2 that have M different vertices $x \in X$ so that $d_{G_n \setminus \{v_1, v_2\}}(v_1, x) \leq k$ and $d_{G_n \setminus \{v_1, v_2\}}(v_2, x) \leq k + 1$. This would imply the existence of a good $x \in X$ as in the proof of Theorem 3 for any pair $\{v_1, v_2\}$.

Proofs of both lower bounds in Theorem 2 and Theorem 4 use the notion of an *undetectable pair* in a graph G having diameter $d \geq 3$ with a set of queried pairs $\mathcal{Q} \subset \binom{V(G)}{2}$: a pair of vertices $\{u_1, u_2\}$ is undetectable if $\{u_1, u_2\} \notin \mathcal{Q} \cup E(G)$ and there is no “path” $u_1 w_1 \dots w_\ell u_2$ in G , where $\ell \leq d - 2$, such that ℓ of its “edges” are actual edges of G , and the one remaining pair of consecutive vertices belongs to \mathcal{Q} . It is easy to see that if G has at least one undetectable pair, then it is not reconstructible from \mathcal{Q} since the addition of the edge $\{u_1, u_2\}$ does not change answers to queries. When \mathcal{Q} is small, then a simple counting proof shows that there are undetectable pairs in G_n whp. This proves Theorem 2, see details in Section 5.

In order to improve the lower bound by a factor of $\log n$ for the “non-adaptive reconstructibility” (Theorem 4), we first refine the notion of undetectability, and then introduce a randomised algorithm that finds an undetectable pair (in the refined sense) in G_n whp. Let us call a pair of vertices $\{u_1, u_2\}$ ℓ -undetectable in G , if $\{u_1, u_2\} \notin \mathcal{Q} \cup E(G)$ and there is no “path” $u_1 w_1 \dots w_\ell u_2$ in G such that ℓ of its “edges” are actual edges of G , the one remaining pair of consecutive vertices belongs to \mathcal{Q} , and the distance between the two vertices in this pair is at least $\ell + 2$. Though the refined notion is weaker, it still allows to prove non-reconstructibility: if G has diameter d and has a pair which is ℓ -undetectable for all $\ell \leq d - 2$, then G is not \mathcal{Q} -reconstructible. Notice that, to decide whether a given pair $\{u_1, u_2\}$ is undetectable, it is sufficient to know only the edges induced by the union of k -neighbourhoods of u_1 and u_2 . In Section 6, we show that the following algorithm finds an undetectable pair whp: at every step, choose a pair of vertices $\{u_1, u_2\}$ uniformly at random from the set of *unconsidered* vertices, consider all vertices in k -neighbourhoods of u_1, u_2 , and check whether $\{u_1, u_2\}$ is ℓ -undetectable for all $\ell \leq k$.

Structure of the paper

The rest of the paper is organised as follows. In Section 2 we recall properties of random graphs that we will use in our proofs: namely, asymptotic behaviour of sizes of balls of small radii and the diameter. Proofs of upper bounds – Theorem 1 and Theorem 3 – appear in Sections 3 and 4 respectively. In Sections 5 and 6 we prove the lower bounds, Theorems 2 and 4 respectively. Finally, in Section 7 we raise several natural questions that we leave unanswered.

2 Properties of random graphs

In this section, we describe certain properties of $G_n \sim G(n, p)$ that we will use in the proofs. We also refer the reader to monographs [9, 16, 20] for a comprehensive exposition of properties of random graphs and of probabilistic tools – especially, to [20, Chapter 2] containing well-known standard concentration inequalities (in particular, Chernoff bounds for binomial distribution [20, Theorem 2.1], Hoeffding’s exponential tail bound for the hypergeometric distribution [20, Theorem 2.10], and Harris’ lemma [20, Theorem 2.12]) that we will use in the proofs and will not state them explicitly.

Let us state a technical claim about the sizes of k -neighbourhoods in G_n that we will use frequently both for lower and upper bounds. For $v \in [n]$ and $r \in \mathbb{Z}_{\geq 0}$, everywhere in the paper we denote by $N_r(v)$ the (random) r -ball around v in G_n . We will also denote $N_r(U) := \cup_{v \in U} N_r(v)$ for $U \subset [n]$.

▷ **Claim 5.** Let k be a positive integer and let $\frac{\log^2 n}{n} \ll p \ll (n^{k-1} \log n)^{-1/k}$. Fix $v \in [n]$. With probability $1 - o(n^{-2})$, in $G(n, p)$ we have $|N_r(v)| = (np)^r(1 + o(\log^{-1} n))$, for all $r \in [k]$.

Proof. From the Chernoff bound, it follows that $|N_1(v)| = np(1 + o(1/\log n))$ with probability $1 - o(n^{-2})$. Let us prove by induction over r that, for every, $r \leq k$, $|N_r(v)| = (np)^r(1 + o(1/\log n))$ with probability $1 - o(n^{-2})$. Assume that for an $r \in [k-1]$, this probability bound is already obtained. Let us expose the r -ball around v and denote the set of vertices at distance exactly r by N'_r . Due to the induction assumption, we may suppose that $|N'_r| = (np)^r(1 + o(1/\log n)) = o(n/\log n)$. Thus, the number of vertices in $[n] \setminus N_r(v)$ having neighbours in N'_r (denoted by $|N'_{r+1}|$) is distributed as $\text{Bin}(n(1 - o(1/\log n)), 1 - (1 - p)^{|N'_r|})$. Clearly, $N_{r+1}(v) = N_r(v) \sqcup N'_{r+1}$. Since $\mathbb{E}|N'_{r+1}| = n^{r+1}p^{r+1}(1 + o(1/\log n)) \gg \log^4 n$, we get that

$$\mathbb{P}\left(\left||N'_{r+1}| - \mathbb{E}|N'_{r+1}|\right| > \mathbb{E}|N'_{r+1}|/(\ln n)^{3/2}\right) = \exp\left[-\Omega(\mathbb{E}|N'_{r+1}|/\ln^3 n)\right] = o(n^{-2})$$

by the Chernoff bound, as needed. \triangleleft

Let us recall the result of Bollobás [7] about the asymptotic distribution of the diameter of G_n , see also [9, Chapter 10.2].

► **Theorem 6.** Let $c \in \mathbb{R}$ be a constant and $p = \left(\frac{2 \ln n + c}{n^k}\right)^{1/(k+1)}$. Then the diameter of G_n converges in distribution to a random variable taking values $k+1$ and $k+2$ with probabilities $e^{-e^{-c}/2}$ and $1 - e^{-e^{-c}/2}$ respectively.

Since the property of having diameter at most d is monotone, we immediately get the following corollary.

► **Corollary 7.** If $n^k p^{k+1} - 2 \ln n \rightarrow \infty$ as $n \rightarrow \infty$, then whp the diameter of G_n is at most $k+1$. If $n^k p^{k+1} - 2 \ln n \rightarrow -\infty$ as $n \rightarrow \infty$, then whp the diameter of G_n is at least $k+2$.

3 Proof of Theorem 1

Let us recall that $k \geq 1$ is a fixed integer, $\varepsilon > 0$, and $n^{-1+\varepsilon} < p < n^{-k/(k+1)-\varepsilon}$. We let $G_n \sim G(n, p)$. Let $M \in \mathbb{Z}$ be large enough and let us choose a large enough real $C \gg M$. Let us also set $\nu := \lceil \log_2 n \rceil$ and fix a set $X \subset [n]$ of size $\nu \lfloor C/(2n^{k-1}p^k) \rfloor$. We will need the following auxiliary statements.

► **Lemma 8.** *Let $i = i(n) \in [\nu]$ and let $X' \subset X$ be a fixed set of size $i \lfloor C/(n^{k-1}p^k) \rfloor$. Then with probability $1 - o(1/\log n)$, at least $(n - |X|)(1 - 2^{-i})$ vertices of $[n] \setminus X$ have at least M vertices at distance at most k in X' .*

For an edge e of a graph G , we denote by $G \setminus e$ the spanning subgraph of G obtained by deleting e . If $e \notin E(G)$, then we let $G \setminus e = G$.

► **Lemma 9.** *Whp there are no $v_1, v_2 \in [n]$ and $u_1, \dots, u_M \in X$ such that, for all $j \in [M]$,*

$$d_{G_n \setminus \{v_1, v_2\}}(v_1, u_j) \leq k \quad \text{and} \quad d_{G_n \setminus \{v_1, v_2\}}(v_2, u_j) \leq k + 1. \quad (2)$$

We prove Lemma 8 and Lemma 9 in Sections 3.1 and 3.2, respectively.

Let us consider a balanced partition of X into ν disjoint sets X_1, \dots, X_ν and assume that G_n is a graph satisfying conclusions of Lemma 8 and Lemma 9 deterministically, i.e.:

- 1 for every $i \in [\nu]$, the number of vertices $v \in [n] \setminus X$ that have at least M vertices $u \in X_1 \cup \dots \cup X_i$ such that $d(v, u) \leq k$ is at least $(n - |X|)(1 - 2^{-i})$;
- 2 there are no $v_1, v_2 \in [n]$ and $u_1, \dots, u_M \in X$ such that (2) holds for all $j \in [M]$.

Let us prove that these conclusions imply the statement of Theorem 1.

Let us first query all pairs from $\binom{X}{2}$ in G_n . Then, we query all pairs from $X_1 \times ([n] \setminus X)$ and find all vertices $v \notin X$ such that there are at least M vertices $u \in X_1$ satisfying $d(v, u) \leq k$. Let $V_1 \subset [n] \setminus X$ be the set of all such vertices v . If $V_1 = [n] \setminus X$, then we set $\hat{i} := 1$. We then proceed by induction: assume that at step $i \in \{1, \dots, \nu - 1\}$ disjoint sets $V_1, \dots, V_i \subset [n] \setminus X$ are chosen and $V_1 \cup \dots \cup V_i \neq [n] \setminus X$. We then query all pairs from $X_{i+1} \times ([n] \setminus (X \cup V_1 \cup \dots \cup V_i))$ and find all vertices $v \notin X \cup V_1 \cup \dots \cup V_i$ such that there are at least M vertices $u \in X_1 \cup \dots \cup X_i \cup X_{i+1}$ satisfying $d(v, u) \leq k$. We let $V_{i+1} \subset [n] \setminus (X \cup V_1 \cup \dots \cup V_i)$ be the set of all such v . If $V_1 \cup \dots \cup V_{i+1} = [n] \setminus X$, then we set $\hat{i} := i + 1$. From 1, we get that, for every i , $|V_1 \cup \dots \cup V_i| \geq (1 - 2^{-i})(n - |X|)$. Thus, the inductive procedure stops at some moment $\hat{i} \leq \nu$, i.e. $V_1 \cup \dots \cup V_{\hat{i}} = [n] \setminus X$, and we query

$$\binom{|X|}{2} + \sum_{i=1}^{\hat{i}} |X_i| \cdot |[n] \setminus (X \cup V_1 \cup \dots \cup V_{i-1})| \leq \left\lfloor \frac{C}{2n^{k-1}p^k} \right\rfloor \sum_{i=1}^{\hat{i}} n \cdot 2^{1-i} < 2n \left\lfloor \frac{C}{2n^{k-1}p^k} \right\rfloor \leq q$$

pairs.

Let us now show that G_n is reconstructible from the set of queried pairs. For every $i \leq \hat{i}$, we have to recover all adjacencies between $v_1 \in V_i$ and $v_2 \in [n]$. Without loss of generality we may assume that v_2 is either from X or from V_j for some $j > i$. Since all pairs from $V_i \times (X_1 \cup \dots \cup X_i)$ are queried, we may consider only $v_2 \in (X_{i+1} \cup \dots \cup X_\nu) \cup (V_{i+1} \cup \dots \cup V_{\hat{i}})$. It is crucial that for all such v_2 , the pairs $\{v_2, u\}$ for $u \in X_1 \cup \dots \cup X_i$ are queried, so that we got both $d(v_1, u)$ and $d(v_2, u)$ from the oracle.

Let us first assume that v_1 and v_2 are non-adjacent in G_n . Let $u_1, \dots, u_M \in X_1 \cup \dots \cup X_i$ be at distance at most k from v_1 . From 2, it follows that there exists $u \in \{u_1, \dots, u_M\}$ such that $d(v_1, u) \leq k$ while $d(v_2, u) \geq k + 2$. It may only happen when v_1, v_2 are not adjacent. Thus, we may “reconstruct non-adjacencies” for all pairs of non-adjacent vertices in G_n .

Now, assume that G' is obtained from G_n by removing some edges that belong to $\cup_{i \leq i} V_i \times ([n] \setminus (X_1 \cup \dots \cup X_i))$ and that for G' we get exactly the same answers as for G_n for all queried pairs. Let $\{v_1, v_2\} \in G_n \setminus G'$, where $v_1 \in V_i$ and $v_2 \in (X_{i+1} \cup \dots \cup X_\nu) \cup (V_{i+1} \cup \dots \cup V_i)$. As above, we let $u_1, \dots, u_M \in X_1 \cup \dots \cup X_i$ be at distance at most k from v_1 in G_n . From **2**, we get that there exists $u \in \{u_1, \dots, u_M\}$ such that $d_{G'}(v_2, u) \geq k + 2$. This inequality holds true since distances in G' cannot be less than in G_n . But then $d_{G_n}(v_2, u) \geq k + 2$ as $d_{G_n}(v_2, u) = d_{G'}(v_2, u) = d(v_2, u)$ by the assumption since the pair $\{v_2, u\}$ was queried. On the other hand, $v_1 \sim v_2$ in G_n and $d(v_1, u) = d_{G_n}(v_1, u) \leq k$. So $d(v_2, u) \leq d(v_1, u) + 1 \leq k + 1$ – a contradiction. Thus, edges are also reconstructible, Theorem 1 follows.

3.1 Proof of Lemma 8

Fix $i \in [\nu]$ and $X' \subset X$ of size $i \lfloor C/(2n^{k-1}p^k) \rfloor$. For $v \in [n] \setminus X$, we let \mathcal{B}_v be the event saying that at least M vertices in X' are at distance at most k from v . We denote $\xi := \sum_{v \in [n] \setminus X} \mathbb{1}_{\mathcal{B}_v}$ the number of such vertices. We have to prove that $\mathbb{P}(\xi < (n - |X|)(1 - 2^{-i})) = o(1/\log n)$.

Fix a vertex $v \in [n] \setminus X$. For any $N \in [n]$, subject to the event $\{|N_k(v)| = N\}$, the set $N_k(v)$ is a uniformly random N -subset of $[n]$ that contains v . We derive the following inequality from Hoeffding's exponential tail bound for the hypergeometric distribution (see [20, Theorem 2.10]): for every $N = (np)^k(1 + o(1))$,

$$\mathbb{P}\left(|N_k(v) \cap X'| < \frac{1}{4}iC \mid |N_k(v)| = N\right) \leq \exp\left[-\frac{iC}{16 - o(1)}\right]. \quad (3)$$

First of all, (3) implies that $\mathbb{P}(|N_k(v) \cap X'| < M \mid |N_k(v)| = N) = o(1/(n \log n))$ uniformly over $N = (np)^k(1 + o(1))$ if $i \geq 17 \ln n/C$. By the union bound and Claim 5, we get that $\xi = n - |X|$ for such i with probability $1 - o(1/\log n)$. We then further assume that $i < 17 \ln n/C$ and that C is so large that $e^{-2i} > n^{-2/(k+1)}$.

For $v \in [n] \setminus X'$, denote by $N_r^{-X'}(v)$ the r -ball around v in $G_n[[n] \setminus X']$. Due to Claim 5, with probability $1 - o(1/n)$, $|N_{k-1}^{-X'}(v)| = (np)^{k-1}(1 + o(1))$ for all $v \in [n] \setminus X'$. We then expose only edges induced by $[n] \setminus X'$ and assume that the latter event holds deterministically.

For $v \in [n] \setminus X$, we let \mathcal{C}_v be the event saying that less than M vertices in X' have neighbours in $N_{k-1}^{-X'}(v)$. Let $\eta = \sum_{v \in [n] \setminus X} \mathbb{1}_{\mathcal{C}_v}$. Note that $\eta \geq n - |X| - \xi$. Therefore, it suffices to prove that $\mathbb{P}(\eta \geq 2^{-i}(n - |X|)) = o(1/\log n)$. A vertex from X' has a neighbour in $N_{k-1}^{-X'}(v)$ for a fixed $v \in [n] \setminus X$ with probability

$$1 - (1 - p)^{(np)^{k-1}(1+o(1))} = 1 - e^{-n^{k-1}p^k(1+o(1))} \sim n^{k-1}p^k(1 + o(1))$$

since $p < n^{-k/(k+1)-\varepsilon}$. Therefore, by the Chernoff bound,

$$\mathbb{P}(\mathcal{C}_v) \leq \mathbb{P}\left(\text{Bin}(|X'|, n^{k-1}p^k(1 + o(1))) < \frac{1}{4}iC\right) \leq \exp\left[-\frac{iC}{16 - o(1)}\right].$$

We get that $\mathbb{E}\eta \leq e^{-i}(n - |X|)$ for $C > 16$ and for large n . Since the relation $u \in N_{k-1}^{-X'}(v)$ on pairs (u, v) of vertices from $[n] \setminus X'$ is symmetric, every vertex from $[n] \setminus X'$ belongs to at most $(np)^{k-1}(1 + o(1))$ balls $N_{k-1}^{-X'}(v)$, $v \in [n] \setminus X$. We shall use the following bounded difference inequality applied to η considered as a function of edges between X' and $[n] \setminus X'$:

► **Theorem 10** ([32]). *Let $\xi_1, \dots, \xi_N \sim \text{Bern}(p)$ be independent, and let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ satisfy*

$$|f(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_N) - f(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_N)| \leq c_i$$

for some $c_1, \dots, c_N > 0$ and every $i \in [N]$, $x_1, \dots, x_N, x'_i \in \mathbb{R}$. Then, for every $t > 0$,

$$\mathbb{P}(|f(\xi_1, \dots, \xi_N) - \mathbb{E}f(\xi_1, \dots, \xi_N)| \geq t) \leq \exp\left[-\frac{t^2}{2p(1-p) \sum_{i=1}^N c_i^2 + \frac{2}{3}t \max_i c_i}\right].$$

Assuming that \hat{G}_n is a graph on $[n]$ such that $\hat{G}_n[[n] \setminus X'] = G_n[[n] \setminus X']$ contains exactly the exposed edges, we get that, for any pair $e \in X' \times ([n] \setminus X')$, values $\eta(\hat{G}_n)$ and $\eta(\hat{G}_n \triangle e)$ differ by at most $(np)^{k-1}(1 + o(1))$ since e may change only the values of $\mathbb{1}_{C_v}$ such that the vertex of e from $[n] \setminus X'$ belongs to $N_{k-1}^{-X'}(v)$. Therefore, we may set all c_i equal to $(np)^{k-1}(1 + o(1))$. Then Theorem 10 implies that, for some constant $c > 0$,

$$\begin{aligned} \mathbb{P}(\eta \geq 2^{-i}(n - |X|)) &\leq \mathbb{P}\left(\eta \geq \frac{e}{2} \cdot e^{-i}(n - |X|)\right) \\ &\leq \exp \left[-c \frac{e^{-2i}(n - |X|)^2}{p(n - |X'|)|X'|((np)^{2k-2} + e^{-i}(n - |X|)(np)^{k-1})} \right] \\ &\leq \exp \left[-\frac{2c}{C+1} \cdot \frac{e^{-2i}n}{i(np)^{k-1}} \right] \leq \exp \left[-n^{\varepsilon(k-1)+o(1)} \right] = o\left(\frac{1}{\log n}\right), \end{aligned}$$

where the penultimate inequality is due to $e^{-2i} > n^{-2/(k+1)}$, completing the proof.

3.2 Proof of Lemma 9

Assume that $M \gg M' \gg 1/\varepsilon$.

▷ **Claim 11.** Fix a vertex $v \in [n]$. With probability $1 - o(n^{-2})$, $|N_{k-1}(v) \cap X| \leq M'$.

Proof. Due to Claim 5, we have that $|N_{k-1}(v)| = (np)^{k-1}(1 + o(1))$ with probability $1 - o(n^{-2})$. For simplicity of presentation, let us assume that $v \notin X$ (the case $v \in X$ does not differ much). Subject to $|N_{k-1}(v)| = N + 1$, the random variable $|N_{k-1}(v) \cap X|$ has a hypergeometric distribution, implying that

$$\begin{aligned} \mathbb{P}(|N_{k-1}(v) \cap X| > M') &\leq (1 + o(1)) \max_{N=(np)^{k-1}(1+o(1))} \frac{\binom{|X|}{M'} \binom{n-|X|-1}{N-M'}}{\binom{n-1}{N}} \\ &\leq (1 + o(1)) \left(\frac{|X| - M'}{n} \right)^{M'} = o(n^{-2}). \quad \triangleleft \end{aligned}$$

Let us define the following family of graphs \mathcal{H} on $[n]$: fix a vertex v_1 and draw from this vertex $\lfloor M/M' \rfloor$ disjoint paths of length at most k to distinct vertices from X and avoiding the edge $\{v_1, v_2\}$; fix another vertex v_2 and draw from this vertex a path of length at most $k+1$ to each of the paths starting at v_1 and avoiding the edge $\{v_1, v_2\}$: these $\lfloor M/M' \rfloor$ paths from v_2 are not necessarily disjoint.

▷ **Claim 12.** Whp G_n does not contain any subgraph from \mathcal{H} .

Proof. First of all note that any $H \in \mathcal{H}$ has at least $\lfloor M/M' \rfloor$ and at most $2k\lfloor M/M' \rfloor + 2$ vertices. Moreover, for any x in this range, the number of graphs in \mathcal{H} with exactly x vertices is at most $c_x n^{x-\lfloor M/M' \rfloor} |X|^{\lfloor M/M' \rfloor}$ for a certain constant c_x – there are at most $|X|^{\lfloor M/M' \rfloor}$ choices for the $\lfloor M/M' \rfloor$ ends in X of disjoint paths emanating from v_1 and at most $n^{x-\lfloor M/M' \rfloor}$ ways to choose the remaining $x - \lfloor M/M' \rfloor$ vertices. Let us bound from below the number of edges in such a graph and apply the union bound. Let us fix vertices v_1, v_2 and $u_1, \dots, u_{\lfloor M/M' \rfloor}$ that play the role of ends of paths from v_1 in X . Note that a graph from \mathcal{H} can be obtained from this set of vertices by a sequential addition of $2\lfloor M/M' \rfloor$ paths joining two vertices that are presented in the graph at the previous step. First $\lfloor M/M' \rfloor$ paths are of length at most k , and the last $\lfloor M/M' \rfloor$ paths are of length at most $k+1$. Assuming that at the j -th step the number of added vertices equals x_j , we get that the constructed graph H has density

$$\rho(H) = \frac{|E(H)|}{|V(H)|} = \frac{\sum_{j=1}^{2\lfloor M/M' \rfloor} (x_j + 1)}{2 + \lfloor M/M' \rfloor + \sum_{j=1}^{2\lfloor M/M' \rfloor} x_j} = \frac{x - 2 + \lfloor M/M' \rfloor}{x} \geq \frac{(2k+1)\lfloor M/M' \rfloor}{2k\lfloor M/M' \rfloor + 2}.$$

Finally, assuming $\lfloor M/M' \rfloor > \frac{3}{\varepsilon(k+1)}$, the expected number of subgraphs $H \in \mathcal{H}$ in G_n is at most

$$\sum_x c_x n^{x - \lfloor M/M' \rfloor} |X|^{\lfloor M/M' \rfloor} p^{x \frac{(2k+1)\lfloor M/M' \rfloor}{2k\lfloor M/M' \rfloor + 2}} = O\left(\left(\nu n^{k+2/\lfloor M/M' \rfloor} p^{k+1}\right)^{\lfloor M/M' \rfloor}\right) = O(n^{-1}). \triangleleft$$

Let us now fix v_1, v_2 and assume that the graph $G' := G_n \setminus \{v_1, v_2\}$ satisfies the following properties deterministically:

- for all $v \in [n]$, $|N_{k-1}(v) \cap X| \leq M'$;
- $u_1, \dots, u_M \in N_k(v_1)$;
- G' does not contain a graph from family \mathcal{H} .

It is then sufficient to prove that there exists $j \in [M]$ such that $u_j \notin N_{k+1}(v_2)$.

Take a shortest path P_1 (of length at most k) between v_1 and u_1 . We then proceed by induction. Assume that at step $j < \lfloor M/M' \rfloor$, we found j disjoint shortest paths P_1, \dots, P_j from v_1 to u_1, \dots, u_j respectively. Let w_1, \dots, w_j be the first vertices (after v_1) on P_1, \dots, P_j respectively. We have that there exists $j' \in [M] \setminus [j]$ such that $u_{j'} \notin N_{k-1}(w_1) \cup \dots \cup N_{k-1}(w_j)$. Without loss of generality let $j' = j+1$. Observe that the shortest path P_{j+1} between v_1 and u_{j+1} cannot intersect any of P_1, \dots, P_j outside of v_1 . Eventually we get $\lfloor M/M' \rfloor$ disjoint paths from v_1 to $u_1, \dots, u_{\lfloor M/M' \rfloor}$. If, for some $j \in [\lfloor M/M' \rfloor]$, $u_j \notin N_{k+1}(v_2)$, we are done. Otherwise, G_n contains a graph from \mathcal{H} – a contradiction.

4 Proof of Theorem 3

Let $\lambda = n^k p^{k+1}$. Without loss of generality we may assume that

$$p \leq \left(\frac{(k + o(1)) \ln n}{(k+1)n^k} \right)^{1/(k+1)} \quad (4)$$

so that $\frac{3e^\lambda \ln n}{n^{k-2} p^k} \leq \binom{n}{2}$. Let X be a fixed set of vertices of size $\lfloor 3e^\lambda \ln n / (n^{k-1} p^k) \rfloor \leq \frac{n}{2}$. Let us query all the pairs from $[X \times ([n] \setminus X)] \cup \binom{X}{2}$. Thus, we query at most $\frac{3}{n^{k-2} p^k} e^\lambda \ln n$ pairs. As in the previous section, we let $G_n \sim G(n, p)$, and, for $v \in [n]$, we let $N_k(v)$ denote the set of vertices in G_n at distance at most k from v .

We further show that whp in G_n for every $v_1, v_2 \in [n] \setminus X$, there is an $x \in X$ such that x is at distance at most k from v_1 and at least $k+2$ from v_2 in the graph obtained from G_n by removing the edge $\{v_1, v_2\}$. It would immediately imply that whp G_n is q -reconstructible. Indeed, we need only to reconstruct adjacency relations between pairs of vertices $v_1, v_2 \notin X$. If v_1, v_2 are non-adjacent, then we queried $\{v_1, x\}, \{v_2, x\}$ such that $d(v_1, x) \leq k, d(v_2, x) \geq k+2$. This may only happen when v_1, v_2 are non-adjacent. Then we can reconstruct all non-adjacencies. It remains to prove that we cannot delete any subset of edges from $G_n[[n] \setminus X]$ without changing the answers to the queries. If $v_1 \sim v_2$, then we queried $\{v_1, x\}, \{v_2, x\}$ such that $d(v_1, x) \leq k, d(v_2, x) \leq d(v_1, x) + 1$, and any shortest path from x to v_2 follows the edge $\{v_1, v_2\}$. Thus, the edge $\{v_1, v_2\}$ must be in the graph.

Therefore, the following lemma immediately implies Theorem 3.

► **Lemma 13.** *Whp for every $v_1, v_2 \in [n] \setminus X$, there is an $x \in X$ such that*

$$d_{G_n \setminus \{v_1, v_2\}}(v_1, x) \leq k \quad \text{and} \quad d_{G_n \setminus \{v_1, v_2\}}(v_2, x) \geq k+2.$$

Proof. Fix vertices $v_1, v_2 \notin X$ and remove the edge $\{v_1, v_2\}$ if it is in G_n . Let N' be the (closed) $(k-1)$ -neighbourhood of v_1 in $G_n \setminus [X \cup \{v_2\}]$, and let N''_i be the (closed) i -neighbourhood of v_2 in $G_n \setminus [N' \cup X]$ for $i \in \{k-1, k\}$. Note that, as soon as N' is

exposed, N''_{k-1} and N''_k are defined solely by adjacencies that are entirely outside N' and their cardinalities do not depend on edges induced by N' – only on the size of N' . Therefore, due to Claim 5, we may assume that $|N'| = (np)^{k-1}(1+o(1/\log n))$, $|N''_i| = (np)^i(1+o(1/\log n))$, $i \in \{k-1, k\}$. By the Chernoff bound and due to (4), the number of neighbours of vertices from N''_{k-1} in N' is $o((np)^{k-1}/\log n)$ with probability $1-o(n^{-2})$. Note that the above makes sense only when $k \geq 2$. In the case $k = 1$, we do not need to apply the Chernoff bound since $N' = \{v_1\}$, $N''_{k-1} = \{v_2\}$, and there is no edge $\{v_1, v_2\}$ in the considered graph.

We then exclude from N' the neighbours of N''_{k-1} and get that the refined set N'_0 still has the size $(np)^{k-1}(1+o(1/\log n))$. If we can prove that with probability $1-o(n^{-2})$, there exists a vertex $x \in X$ that has a neighbour in N'_0 and does not have neighbours in $N''_k \sqcup (N' \setminus N'_0)$, we immediately get the statement of Lemma 13: if P is a path of length at most $k+1$ from v_2 to x in $G_n \setminus \{v_1, v_2\}$, then by construction the neighbour of x on P is in N'_0 . Then the farthest from x vertex in $P \cap N'_0$ has a neighbour in $P \cap N''_{k-1}$ – a contradiction. The probability bound on the latter event is also immediate since

$$\begin{aligned} & \mathbb{P}\left(\text{Bin}\left(|X|, \left(1 - (1-p)^{|N'_0|}\right) (1-p)^{|N''_k|+|N' \setminus N'_0|}\right) = 0\right) \\ &= \left(1 - \left(1 - (1-p)^{(np)^{k-1}(1+o(1/\log n))}\right) (1-p)^{(np)^k(1+o(1/\log n))}\right)^{|X|} \\ &\stackrel{(4)}{=} \exp\left[-(1+o(1))|X|n^{k-1}p^k e^{-\lambda}\right] = o(n^{-2}). \quad \blacktriangleleft \end{aligned}$$

5 Proof of Theorem 2

Without loss of generality, we assume $p = o(1)$. Let us show that Theorem 2 follows from

▷ Claim 14. Let $d \geq 3$, $\varepsilon > 0$. There exists $C = C(d, \varepsilon) > 0$ such that, for every graph G on $[n]$ with diameter at most d and maximum degree $\Delta > C$, its query complexity q satisfies:

$$q(d-1+\varepsilon)\Delta^{d-2} \geq \binom{n}{2} - |E(G)|. \quad (5)$$

Indeed, whp G_n has $(1+o(1))\binom{n}{2}p$ edges, diameter at most $k+2$ (by Corollary 7), and maximum degree $np(1+o(1))$ (by Claim 5, see also [6]), implying that the query complexity of $G(n, p)$ is at least $\frac{1}{2(k+1+\varepsilon)n^{k-2}p^k}$ whp, as needed.

Proof of Claim 14. Let q be an integer that does not satisfy (5). Let $\mathcal{Q} \subset \binom{[n]}{2}$ have size q . Let us call a pair of vertices $\{u_1, u_2\}$ *undetectable* in G if:

- u_1, u_2 are not adjacent in G ,
- the pair $\{u_1, u_2\}$ was not queried,
- there is no “path” $u_1 w_1 \dots w_\ell u_2$, $\ell \leq d-2$, such that ℓ of its “edges” are actual edges of G , and the one remaining pair of consecutive vertices belongs to \mathcal{Q} .

Assuming that the pair $\{u_1, u_2\}$ is undetectable in G , observe that the addition of the edge $\{u_1, u_2\}$ agrees with the answers to all the queries. Indeed, if it affects some queried pair $\{w, w'\}$, then there should be a path $w \dots u_1 u_2 \dots w'$ of length at most $d-1$ in $G \cup \{u_1, u_2\}$ – a contradiction.

It remains to prove that G has an undetectable pair. For an appropriate choice of C , the number of pairs that do not satisfy the third property of an undetectable pair is at most

$$\sum_{\ell=1}^{d-2} (\ell+1)q\Delta^\ell < q((d-1)\Delta^{d-2} + (d-2)^2\Delta^{d-3}) < q(d-1+\varepsilon)\Delta^{d-2} - q$$

On the other hand, the number of pairs that satisfy the first two properties is at least $\binom{n}{2} - |E(G)| - q > q(d-1+\varepsilon)\Delta^{d-2} - q$, implying the existence of an undetectable pair. ◁

6

 Proof of Theorem 4

Let $G_n \sim G(n, p)$, $N := \left\lfloor \frac{e^\lambda \ln n}{8(k+1)^2 n^{k-2} p^k} \right\rfloor$, where $\lambda = n^k p^{k+1}$. We assume that a set $\mathcal{Q} \subset \binom{[n]}{2}$ of N pairs of vertices from $[n]$ is fixed in advance to be queried. We call a pair $\{u_1, u_2\}$ *small* if

- the number of $\{v_1, v_2\} \in \mathcal{Q}$ such that $v_1 \in N_{k'}(u_1)$ and $v_2 \in N_{k-k'}(u_2)$ for some $k' \in \{0, 1, \dots, k\}$ is at most $M := 7(k+1)(np)^k N/n^2$;
- there are no $\{v_1, v_2\} \in \mathcal{Q}$ such that $v_1 \in N_{k'}(u_1)$ and $v_2 \in N_{k-1-k'}(u_2)$ for some $k' \in \{0, 1, \dots, k-1\}$;
- there are no $\{v_1, v_2\} \in \mathcal{Q}$ such that, for some $k' \in \{0, 1, \dots, k\}$, $v_1 \in N_{k'}(u_1)$, $v_2 \in N_{k-k'}(u_2)$, and either $v_2 \in N_k(u_1)$ or $v_1 \in N_k(u_2)$.

▷ **Claim 15.** Whp at least $\frac{1}{2} \binom{n}{2}$ pairs of vertices are small.

Proof. Due to Claim 5, whp, for every $k' \in \{0, 1, \dots, k\}$, for every $\{v_1, v_2\} \in \mathcal{Q}$, there are at most $(1+o(1))(np)^k$ ordered pairs (u_1, u_2) such that $v_1 \in N_{k'}(u_1)$ and $v_2 \in N_{k-k'}(u_2)$. Then whp the following holds: if the number of pairs $\{u_1, u_2\}$ satisfying the first condition in the definition of a small pair is less than $\frac{5}{6} \binom{n}{2}$, then $\frac{1}{6} n(n-1)M < (1+o(1))(k+1)(np)^k N$. The later inequality is false for large enough n . Similarly, whp, for every $\{v_1, v_2\} \in \mathcal{Q}$, there are at most $(1+o(1))k(np)^{k-1}$ ordered pairs (u_1, u_2) such that $v_1 \in N_{k'}(u_1)$ and $v_2 \in N_{k-1-k'}(u_2)$ for some $k' \in \{0, 1, \dots, k-1\}$. Then whp the following holds: if the number of pairs $\{u_1, u_2\}$ satisfying the second condition in the definition of a small pair is less than $\frac{5}{6} \binom{n}{2}$, then $\frac{1}{6} n(n-1) < (1+o(1))k(np)^{k-1} N$. The later inequality is false for large enough n as well.

Let us finally show that whp the number of pairs $\{u_1, u_2\}$ satisfying the third condition in the definition of a small pair is at least $\frac{5}{6} \binom{n}{2}$. We will do it in two steps: first, we will show that whp there are only $o(n^2)$ pairs $\{u_1, u_2\}$ that have in their “neighbourhood” $N_{k'}(u_1) \times N_{k-k'}(u_2)$ queried pairs $\{v_1, v_2\}$ with $d(v_1, v_2) \leq k$; second, we will observe that whp as soon as a queried pair $\{v_1, v_2\}$ has $d(v_1, v_2) > k$, then the intersection of the k' -ball around v_1 and the k -ball around v_2 has size around $(np)^{k+k'}/n$ – that will be sufficient to finish the proof.

Fix a queried pair $\{v_1, v_2\}$. Since the expected number of paths of length at most k between v_1, v_2 is at most $n^{k-1}p^k(1+o(1))$, we get that the distance between v_1, v_2 is at most k with probability at most $n^{k-1}p^k(1+o(1))$ by Markov’s inequality. Therefore, the expected number of pairs in \mathcal{Q} that are at distance at most k is at most $n^{k-1}p^k(1+o(1)) \cdot N$. By Markov’s inequality, we get that the number of such pairs in \mathcal{Q} is $O_p(n^{k-1}p^k N)$. As we noted above, whp, for every pair $\{v_1, v_2\} \in \mathcal{Q}$, there are at most $(1+o(1))(k+1)(np)^k$ pairs (u_1, u_2) such that $v_1 \in N_{k'}(u_1)$, $v_2 \in N_{k-k'}(u_2)$ for some $k' \in \{0, 1, \dots, k\}$. Let us fix a sequence w_n growing to infinity with n sufficiently slowly. We get that whp, if the number of pairs $\{u_1, u_2\}$ that have $v_1 \in N_{k'}(u_1)$ and $v_2 \in N_{k-k'}(u_2)$ for some $k' \in \{0, 1, \dots, k\}$ such that $\{v_1, v_2\} \in \mathcal{Q}$ and $d(v_1, v_2) \leq k$ is at least $\frac{1}{12} \binom{n}{2}$, then $\frac{1}{12} n(n-1) < w_n (np)^k n^{k-1} p^k N$. This is false for all large enough n . This proves that whp the number of pairs satisfying the last property in the definition of small pairs with the restriction that $d(v_1, v_2) \leq k$ (note that it always holds for $k' \in \{0, k\}$) is at least $\frac{11}{12} \binom{n}{2}$. It then remains to prove that whp the number of pairs $\{u_1, u_2\}$ with no $\{v_1, v_2\} \in \mathcal{Q}$ such that $d(v_1, v_2) > k$ and, for some $k' \in [k-1]$, $v_1 \in N_{k'}(u_1)$, $v_2 \in N_{k-k'}(u_2)$, and either $v_2 \in N_k(u_1)$ or $v_1 \in N_k(u_2)$, is at least $\frac{11}{12} \binom{n}{2}$ as well.

As promised, we first show that whp, for any $\{v_1, v_2\}$ such that $d(v_1, v_2) > k$ and any $k' \in [k]$, the intersection of the k' -ball around v_1 and the k -ball around v_2 has size at most $(1+o(1))(np)^{k+k'}/n + \ln^2 n$. Fix v_2 and expose the k -ball around it. Due to Claim 5, with probability $1 - o(n^{-2})$, this ball $N_k(v_2)$ has size $(1+o(1))(np)^k$. Let us now observe that we may assume that all vertices from $S := N_k(v_2) \setminus N_{k-1}(v_2)$ have degrees at most $\ln^2 n$

in $N_k(v_2)$. Indeed, the complementary event implies that there exists a vertex $w \in S$ such that the total number of $(k+1)$ -paths and k -paths between u in v_2 is at least $\ln^2 n$, which is false with probability $1 - o(n^{-2})$: due to [40, Theorem 5], there exists $C > 0$ such that, with probability $1 - o(n^{-2})$, the number of paths of length at most $k+1$ between any two vertices in $G(n, p)$ is at most $C \ln n$.

Then, fix $v_1 \notin N_k(v_2)$. For $i \in \{0, 1, \dots, k' - 1\}$, let N'_i be the i -neighbourhood of v_1 in $[n] \setminus N_k(v_2)$, and let N_i^\cap be the set of neighbours of N'_i in $N_k(v_2)$. Note that, for every vertex $w \in N_{k'}(v_1) \cap N_k(v_2)$, a shortest path from w to v_1 leaves $N_k(v_2)$ at some vertex $w' \in S$. Thus, the set $N_{k'}(v_1) \cap N_k(v_2)$ coincides with $\cup_{i=0}^{k'-1} [N_{k'-1-i}(N_i^\cap) \cap N_k(v_2)]$. Then

$$|N_k(v_2) \cap N_{k'}(v_1)| \leq \sum_{i=0}^{k'-2} |N_i^\cap| \ln^2 n (np)^{k'-2-i} + |N_{k'-1}^\cap| \quad (6)$$

with probability $1 - o(n^{-2})$. Moreover, due to Claim 5, with probability $1 - o(n^{-2})$, for every $i \in \{0, 1, \dots, k' - 1\}$, $|N'_i| \leq (np)^i(1 + o(1))$.

Expose all N'_i , $i \in \{0, 1, \dots, k' - 1\}$, and assume that $|N'_i| \leq (np)^i(1 + o(1))$ and $|N_k(v_2)| \leq (1 + o(1))(np)^k$ hold deterministically. Since $|N_i^\cap|$ is the number of vertices in $N_k(v_2)$ that have at least one neighbour in N'_i , we get, by the Chernoff bound,

$$\begin{aligned} \mathbb{P}(|N_i^\cap| > (np)^{k+i} p(1 + (\ln \ln n)^{-1}) + \ln^2 n) &\leq \\ \mathbb{P}\left(\text{Bin}\left((np)^k(1 + o(1)), 1 - (1 - p)^{(np)^i(1+o(1))}\right) > (np)^{k+i} p \left(1 + \frac{1}{\ln \ln n}\right) + \ln^2 n\right) &= \\ o(n^{-2}). \end{aligned}$$

We finally get from (6) that with probability $1 - o(n^{-2})$,

$$\begin{aligned} |N_{k'}(v_1) \cap N_k(v_2)| &\leq \sum_{i=0}^{k'-2} \ln^2 n (np)^{k'-2-i} ((np)^{k+i} p(1 + o(1)) + \ln^2 n) + \\ &\quad (np)^{k+k'-1} p(1 + o(1)) + \ln^2 n = (1 + o(1))(np)^{k+k'-1} p + \ln^2 n, \end{aligned}$$

as needed. From this and since whp, for every vertex v_2 and every $k' \in [k-1]$, $|N_{k-k'}(v_2)| = (1 + o(1))(np)^{k-k'}$ due to Claim 5, we get that whp, for every $\{v_1, v_2\} \in \mathcal{Q}$ satisfying $d(v_1, v_2) > k$, there at most $k \left(\frac{(np)^{2k}}{n} + (np)^{k-1} \ln^2 n \right)$ pairs (u_1, u_2) such that,

$$\text{for some } k' \in [k-1], \quad v_1 \in N_{k'}(u_1), \quad v_2 \in N_{k-k'}(u_2), \quad \text{and} \quad v_2 \in N_k(u_1). \quad (7)$$

So, the event that the number of pairs $\{u_1, u_2\}$ that have a pair $\{v_1, v_2\}$ satisfying $d(v_1, v_2) > k$ and (7) is at least $\frac{1}{24} \binom{n}{2}$ implies

$$\frac{1}{24} n(n-1) < k \left(\frac{(np)^{2k}}{n} + (np)^{k-1} \ln^2 n \right) N,$$

which is false for all large n . Due to symmetry, we get that whp the number of pairs $\{u_1, u_2\}$ that have a pair $\{v_1, v_2\}$ satisfying $d(v_1, v_2) > k$, $v_1 \in N_{k'}(u_1)$, $v_2 \in N_{k-k'}(u_2)$, and $v_1 \in N_k(u_2)$ is less than $\frac{1}{24} \binom{n}{2}$ as well, completing the proof. \triangleleft

Let $\ell \in [k]$. We call a pair of vertices (u_1, u_2) ℓ -undetectable, if

- u_1, u_2 are not adjacent in G_n and $\{u_1, u_2\} \notin \mathcal{Q}$,
- there is no “path” $u_1 w_1 \dots w_\ell u_2$, such that ℓ of the pairs of consecutive vertices $\{u_1, w_1\}$, $\{w_1, w_2\}$, \dots , $\{w_\ell, u_2\}$ are actual edges of G_n , the one remaining pair of consecutive vertices is from \mathcal{Q} , and the distance between the vertices in this pair is at least $\ell + 2$ in G_n .

Let us recall that whp G_n has diameter $k + 2$ due to Corollary 7. Assuming that the pair $\{u_1, u_2\}$ is ℓ -undetectable for all $\ell \in [k]$, observe that whp the addition of the edge $\{u_1, u_2\}$ agrees with the answers to all the queries. Indeed, if it affects some queried pair $\{w, w'\}$, then there should be a path $w \dots u_1 u_2 \dots w'$ of length $\ell + 1 \leq k + 1$ in $G_n \cup \{u_1, u_2\}$ while in G_n the distance between w, w' should be at least $\ell + 2$. This contradicts the fact that $\{u_1, u_2\}$ is ℓ -undetectable since if we consider the sequence $u_2 \dots w' w \dots u_1$ obtained from the path $w \dots u_1 u_2 \dots w'$ by replacing the prefix $w \dots u_1$ and the suffix $u_2 \dots w'$, we get that the pair $\{w, w'\}$ is queried, $d(w, w') \geq \ell + 2$, and all the other consecutive pairs of vertices are edges of G_n . Thus, whp the event that there exists a pair $\{u_1, u_2\}$ which is ℓ -undetectable for all $\ell \in [k]$ implies that G_n is not \mathcal{Q} -reconstructible. It then remains to prove that whp G_n has a pair $\{u_1, u_2\}$ which is ℓ -undetectable for all $\ell \in [k]$.

Let us now consider the following iterative procedure. At every step $i \geq 1$, we are given with a set of “considered” vertices $X_i \subset [n]$ (initially $X_1 = \emptyset$), then we sample a pair $\{u_1, u_2\}$ uniformly at random from $\binom{[n] \setminus X_i}{2}$ and expose $G_n[N_k(u_1)]$ and $G_n[N_k(u_2)]$. If the pair $\{u_1, u_2\}$ is not small or u_1, u_2 are adjacent (note that the exposed edges are enough to get these decisions), then we just skip this step and switch to the step $i + 1$ with $X_{i+1} := X_i \cup N_k(u_1) \cup N_k(u_2)$.

Now, assume that at the current step i , we observe that the pair $\{u_1, u_2\}$ is small and u_1, u_2 are not adjacent. Note that the pair $\{u_1, u_2\}$ is ℓ -undetectable for every $\ell \in [k - 1]$ since it is small (so, due to the first bullet point in the definition of a small pair, it cannot have any “path” of length at most k from the definition of an undetectable pair). The crucial observation is that the edges between $N_k(u_1) \setminus N_k(u_2)$ and $N_k(u_2) \setminus N_k(u_1)$ have not been exposed yet unless they are entirely inside X_i . We then expose these missing edges and show that they define an event that has sufficiently large probability and implies that the pair $\{u_1, u_2\}$ is k -undetectable (see Claim 16 below). After this step, we switch to the step $i + 1$ with $X_{i+1} := X_i \cup N_k(u_1) \cup N_k(u_2)$ as before.

We perform $\tau = \lfloor n^{1/(k+1)-o(1)} \rfloor$ steps, where τ is chosen so that $\tau(np)^k = o(n/\log n)$. Due to Claim 5, whp (in the measure of the union of the balls exposed before this step) $|X_\tau| = o(n/\log n)$, implying that whp (in the same measure), for every $i \leq \tau$, the probability that the pair $\{u_1, u_2\}$ we sample at step i is small (in the measure of the random pair $\{u_1, u_2\}$) is at least $1/2 - o(1)$ by Claim 15.

▷ **Claim 16.** There exists an event \mathcal{B}_{u_1, u_2} that depends only on edges between $N_k(u_1) \setminus N_k(u_2)$ and $N_k(u_2) \setminus N_k(u_1)$, implies that the pair $\{u_1, u_2\}$ is k -undetectable, and satisfies

$$\mathbb{P}(\mathcal{B}_{u_1, u_2} \mid G_n[X_i], \{u_1, u_2\} \notin E(G_n), \{u_1, u_2\} \text{ is small}) \geq n^{-7/(8(k+1))+o(1)}.$$

Proof. Fix $k' \leq k$, $v_1 \in N_{k'}(u_1)$, $v_2 \in N_{k-k'}(u_2)$ such that the pair $\{v_1, v_2\}$ was queried. Note that since $\{u_1, u_2\}$ is small, it may only happen when $v_1 \notin N_k(u_2)$, $v_2 \notin N_k(u_1)$.

Without loss of generality, assume that $k' \leq k - k'$. Note that, if there exists an edge in G_n between $A(v_1) := N_{k-k'}(v_1) \setminus (X_i \cup N_k(u_2))$ and $A(v_2) := N_{k'}(v_2) \setminus N_k(u_1)$, then the distance between v_1, v_2 is at most $k + 1$. Due to Claim 5, whp

$$\begin{aligned} |N_{k-k'}(v_1)| &= (1 + o(1/\log n))(np)^{k-k'}, & |N_k(u_1)| &= (1 + o(1/\log n))(np)^k, \\ |N_{k'}(v_2)| &= (1 + o(1/\log n))(np)^{k'}, & |N_k(u_2)| &= (1 + o(1/\log n))(np)^k. \end{aligned}$$

Since $|X_i| = o(n/\log n)$, due to Hoeffding’s exponential tail bound for the hypergeometric distribution, whp $|A(v_1)| = (1 + o(\log^{-1} n))(np)^{k-k'}$ and $|A(v_2)| = (1 + o(\log^{-1} n))(np)^{k'}$. Then

$$\mathbb{P}(\text{there is an edge between } A(v_1) \text{ and } A(v_2)) = 1 - (1 - p)^{(1+o(1/\log n))(np)^k} = 1 - e^{-\lambda+o(1)}.$$

Since $\{u_1, u_2\}$ is small, the number of queried $\{v_1, v_2\}$ in $\bigcup_{k'=0}^k N_{k'}(u_1) \times N_{k-k'}(u_2)$ is at most M . Now, we define the desired event \mathcal{B}_{u_1, u_2} : for every k' and every queried pair $\{v_1 \in N_{k'}(u_1), v_2 \in N_{k-k'}(u_2)\}$, there exists an edge between the respective $A(v_1)$ and $A(v_2)$. Then, due to the Harris inequality [17],

$$\mathbb{P}(\mathcal{B}_{u_1, u_2}) \geq (1 - e^{-\lambda + o(1)})^M \sim \exp[-(7/8 + o(1)) \ln n / (k+1)]. \quad \triangleleft$$

Whp in at least $(1/2 - o(1))\tau$ steps we get a small non-adjacent pair. Due to Claim 16, we then get that the probability that, for at least one of these small pairs $\{u_1, u_2\}$, the respective event \mathcal{B}_{u_1, u_2} happens is at least

$$1 - \left(1 - n^{-7/(8(k+1)) + o(1)}\right)^{(1/2 - o(1))\tau} = 1 - \exp\left[-\frac{1}{2} n^{\frac{1+o(1)}{8(k+1)}}\right] = 1 - o(1),$$

completing the proof.

7 Open questions

Although we get tight bounds on distance query complexity for the binomial random graph $G(n, p)$ when the edge probability p satisfies (1), as p gets closer to a hitting time of a diameter increment, the difference between the bounds increases. In particular, there is a polynomial gap between the bounds even for the “non-adaptive reconstructibility” when $p = \Theta((n^{-k} \ln n)^{1/(k+1)})$ satisfies

$$(1 - o(1)) \left(\frac{\ln n}{(k+1)n^k}\right)^{1/(k+1)} < p < \left(\frac{2 \ln n + O(1)}{n^k}\right)^{1/(k+1)}.$$

It would be interesting to get at least the right power of n in the query complexity in this case.

The further natural step is to generalise our results to growing diameter, i.e. to $p = n^{-1+o(1)}$. Although Theorems 2, 3, and 4 are generalised directly to $p \gg \ln^2 n / n$, it is unclear whether it is possible to extend Theorem 1 without any significant modification of the proof method. The main complication is in the application of the bounded difference inequality (Theorem 10) since we use the fact that $e^{-2^i} > n^{-2/(k+1)}$ for all $i < \ln n / C$ and some sufficiently large constant C . If k is growing, then there is no such C .

As we mention in Introduction, Mathieu and Zhou [31] proved that, for constant d , the query complexity of the random d -regular graph on $[n]$ is at most $n \ln^2 n$ whp. On the other hand, it is easy to prove the information-theoretic lower bound $\Omega(n \log n / \log \log n)$. It seems likely that the query complexity of the random d -regular graph actually equals $n(\log n)^{1-o(1)}$ whp, and we are asking whether the lower bound is tight.

References

- 1 R. Afshar, M. T. Goodrich, P. Matias, and M. C. Osegueda. Reconstructing Biological and Digital Phylogenetic Trees in Parallel. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms (ESA 2020)*, volume 173 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:24, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2020.3.
- 2 R. Arratia, D. Martin, G. Reinert, and M. S. Waterman. Poisson process approximation for sequence repeats, and sequencing by hybridization. *Journal of Computational Biology*, 3(3):425–463, 1996. doi:10.1089/CMB.1996.3.425.

- 3 L. Babai, P. Erdős, and S. M. Selkow. Random graph isomorphism. *SIAM Journal on Computing*, 9(3):628–635, 1980. doi:10.1137/0209047.
- 4 P. Bastide and C. Groenland. Optimal distance query reconstruction for graphs without long induced cycles. *arXiv preprint arXiv:2306.05979*, 2023.
- 5 Z. Beerliova, F. Eberhard, T. Erlebach, A. Hall, M. Hoffmann, M. Mihal’ák, and L. S. Ram. Network discovery and verification. *IEEE Journal on selected areas in communications*, 24(12):2168–2181, 2006. doi:10.1109/JSAC.2006.884015.
- 6 B. Bollobás. The distribution of the maximum degree of a random graph. *Discrete Mathematics*, 32:201–203, 1980. doi:10.1016/0012-365X(80)90054-0.
- 7 B. Bollobás. The diameter of random graphs. *Transactions of the American Mathematical Society*, 267:41–52, 1981.
- 8 B. Bollobás. Almost every graph has reconstruction number three. *Journal of Graph Theory*, 14(1):1–4, 1990. doi:10.1002/JGT.3190140102.
- 9 B. Bollobás. *Random Graphs*. Cambridge University Press, 2 edition, 2001.
- 10 B. Bollobás, D. Mitsche, and P. Prałat. Metric dimension for random graphs. *Electronic Journal of Combinatorics*, 20(4):P1, 2013.
- 11 R. Brodal, G. S. Fagerberg, C. N. S. Pedersen, and A. Östlin. The complexity of constructing evolutionary trees using experiments. *BRICS Report Series*, 8(1), 2001.
- 12 Y. Demidovich, Y. Panichkin, and M. Zhukovskii. Reconstruction of graph colourings. *arXiv preprint*, 2023. arXiv:2308.01671.
- 13 M. Dyer, A. Frieze, and S. Suen. The probability of unique solutions of sequencing by hybridization. *Journal of Computational Biology*, 1(2):105–110, 1994. doi:10.1089/CMB.1994.1.105.
- 14 T. Erlebach, A. Hall, M. Hoffmann, and M. Mihal’ák. Network discovery and verification with distance queries. *Algorithms and Complexity*, pages 69–80, 2006.
- 15 T. Erlebach, A. Hall, and M. Mihal’ák. Approximate discovery of random graphs. In *Stochastic Algorithms: Foundations and Applications, SAGA 2007*, volume 4665 of *Lecture Notes in Computer Science*. Springer, 2007. doi:10.1007/978-3-540-74871-7_8.
- 16 A. Frieze and M. Karoński. *Introduction to Random Graphs*. Cambridge University Press, 2015.
- 17 T. E. Harris. A lower bound for the critical probability in a certain percolation process. *Mathematical Proceedings of the Cambridge Philosophical Society*, 56:13–20, 1960.
- 18 J. J. Hein. An optimal algorithm to reconstruct trees from additive distance data. *Bulletin of mathematical biology*, 51(5):597–603, 1989.
- 19 H. Huang and K. Tikhomirov. Shotgun assembly of unlabeled Erdős-Rényi graphs. *Probab. Theory Relat. Fields*, 192:575–624, 2025. doi:10.1007/s00440-024-01347-4.
- 20 S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs*. J. Wiley & Sons, 2000.
- 21 T. Johnston, G. Kronenberg, A. Roberts, and A. Scott. Shotgun assembly of random graphs. *Probability Theory and Related Fields*, 2025. doi:10.1007/s00440-025-01380-x.
- 22 S. Kannan, C. Mathieu, and H. Zhou. Graph reconstruction and verification. *ACM Transactions on Algorithms*, 14(4):1–30, 2018. doi:10.1145/3199606.
- 23 S. K. Kannan, E. L. Lawler, and T. J. Warnow. Determining the evolutionary tree using experiments. *Journal of Algorithms*, 21(1):26–50, 1996. doi:10.1006/JAGM.1996.0035.
- 24 P. J. Kelly. A congruence theorem for trees. *Pacific Journal of Mathematics*, 7:961–968, 1957.
- 25 V. King, L. Zhang, and Y. Zhou. On the complexity of distance-based evolutionary tree reconstruction. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’03)*, pages 444–453, 2003.
- 26 L. Kučera. Canonical labeling of regular graphs in linear average time. In *Proceedings of 28th Annual Symposium on Foundations of Computer Science (FOCS’87)*, pages 271–279, 1987.
- 27 A. Lingas, H. Olsson, and A. Östlin. Efficient merging, construction, and maintenance of evolutionary trees. In *Automata, Languages and Programming*, pages 544–553, 1999.

- 28 N. Linial and J. Mosheiff. On the rigidity of sparse random graphs. *Journal of Graph Theory*, 85(2):466–480, 2017. doi:10.1002/JGT.22073.
- 29 G. D. Marmerola, M. A. Oikawa, Z. Dias, S. Goldenstein, and A. Rocha. On the reconstruction of text phylogeny trees: evaluation and analysis of textual relationships. *PloS one*, 11(12):e0167822, 2016.
- 30 C. Mathieu and H. Zhou. Graph reconstruction via distance oracles. In *International Colloquium on Automata, Languages and Programming (ICALP’13)*, pages 733–744, 2013.
- 31 C. Mathieu and H. Zhou. A simple algorithm for graph reconstruction. In *29th Annual European Symposium on Algorithms (ESA’21)*, pages 68:1–68:18, 2021. doi:10.4230/LIPIcs.ESA.2021.68.
- 32 C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics: Invited Papers at the Twelfth British Combinatorial Conference*, London Mathematical Society Lecture Note Series, pages 148–188. Cambridge University Press, 1989.
- 33 E. Mossel and N. Ross. Shotgun assembly of labelled graphs. *IEEE Transactions on Network Science and Engineering*, 6(2):145–157, 2019. doi:10.1109/TNSE.2017.2776913.
- 34 A. S. Motahari, G. Bresler, and D. N. Tse. Information theory of dna shotgun sequencing. *IEEE Transactions on Information Theory*, 59(10):6273–6289, 2013. doi:10.1109/TIT.2013.2270273.
- 35 V. Müller. Probabilistic reconstruction from subgraphs. *Commentationes Mathematicae Universitatis Carolinae*, 17(4):709–719, 1976.
- 36 G. Odor. *The role of adaptivity in source identification with time queries*. PhD thesis, EPFL, 2022.
- 37 L. Reyzin and N. Srivastava. Learning and verifying graphs using queries with a focus on edge counting. In *Algorithmic Learning Theory*, pages 285–297. Springer, 2007. doi:10.1007/978-3-540-75225-7_24.
- 38 S. Sen and V. N. Muralidhara. The covert set-cover problem with application to network discovery. In *International Conference and Workshops on Algorithms and Computation (WALCOM’10)*, pages 228–239, 2010.
- 39 D. Soudry, S. Keshri, P. Stinson, M.-h. Oh, G. Iyengar, and L. Paninski. Efficient “shotgun” inference of neural connectivity from highly sub-sampled activity data. *PLoS Computational Biology*, 11(12):e1004657, 2015.
- 40 J. Spencer. Counting extensions. *Journal of Combinatorial Theory, Series A*, 55:247–255, 1990. doi:10.1016/0097-3165(90)90070-D.
- 41 H. Spinoza and D. B. West. Reconstruction from the deck of k -vertex induced subgraphs. *Journal of Graph Theory*, 90:497–522, 2019. doi:10.1002/JGT.22409.
- 42 S. M. Ulam. *A collection of mathematical problems*. Number 8 in Interscience Tracts in Pure and Applied Mathematics. Interscience Publishers, New York–London, 1960.
- 43 M. S. Waterman, T. F. Smith, M. Singh, and W. A. Beyer. Additive evolutionary trees. *Journal of Theoretical Biology*, 64(2):199–213, 1977.