

Graph Sketching Against Adaptive Adversaries Applied to the Minimum Degree Algorithm

Matthew Fahrbach* Georgia Tech matthew.fahrbach@gatech.edu	Gary L. Miller† CMU glmiller@cs.cmu.edu	Richard Peng‡ Georgia Tech rpeng@cc.gatech.edu
Saurabh Sawlani‡ Georgia Tech sawlani@gatech.edu	Junxing Wang CMU junxingw@cs.cmu.edu	Shen Chen Xu† Facebook§ shenchex@cs.cmu.edu

April 13, 2018

Abstract

Motivated by the study of matrix elimination orderings in combinatorial scientific computing, we utilize graph sketching and local sampling to give a data structure that provides access to approximate fill degrees of a matrix undergoing elimination in $O(\text{polylog}(n))$ time per elimination and query. We then study the problem of using this data structure in the minimum degree algorithm, which is a widely-used heuristic for producing elimination orderings for sparse matrices by repeatedly eliminating the vertex with (approximate) minimum fill degree. This leads to a nearly-linear time algorithm for generating approximate greedy minimum degree orderings. Despite extensive studies of algorithms for elimination orderings in combinatorial scientific computing, our result is the first rigorous incorporation of randomized tools in this setting, as well as the first nearly-linear time algorithm for producing elimination orderings with provable approximation guarantees.

While our sketching data structure readily works in the oblivious adversary model, by repeatedly querying and greedily updating itself, it enters the adaptive adversarial model where the underlying sketches become prone to failure due to dependency issues with their internal randomness. We show how to use an additional sampling procedure to circumvent this problem and to create an independent access sequence. Our technique for decorrelating the interleaved queries and updates to this randomized data structure may be of independent interest.

*Supported in part by a National Science Foundation Graduate Research Fellowship under grant DGE-1650044.

†This material is based on work supported by the National Science Foundation under Grant No. 1637523.

‡This material is based on work supported by the National Science Foundation under Grant No. 1637566.

§Part of this work was done while at CMU.

1 Introduction

Randomization has played an increasingly fundamental role in the design of modern data structures. The current best algorithms for fully-dynamic graph connectivity [KKM13, NSW17, NS17, Wu17], shortest paths [HKN14, ACK17], graph spanners [BKS12], maximal matchings [BGS15, Sol16], and the dimensionality-reductions of large matrices [Woo14, CMP16, KLM⁺17, KPPS17] all critically rely on randomization. An increasing majority of these data structures operate under the *oblivious adversary model*, which assumes that updates are generated independently of the internal randomness used in the data structure. In contrast, many applications of data structures are *adaptive*—meaning that subsequent updates may depend on the output of previous queries. A classical example of this paradigm is the combination of greedy algorithms with data structures, including Dijkstra’s algorithm for computing shortest paths and Kruskal’s algorithm for finding minimum spanning trees. The limitations imposed by adaptive adversaries are beginning to receive attention in the dynamic connectivity [NS17, NSW17] and spanner [BK16] literature, but even for these problems there remains a substantial gap between algorithms that work in the adaptive adversary model and those that work only against oblivious adversaries [BKS12, KKM13].

Motivated by a practically important example of adaptive invocations to data structures for greedy algorithms, we study the minimum degree algorithm for sparse matrix factorization and linear system solving [DRSL16]. This heuristic for precomputing an efficient pivot ordering is ubiquitous in numerical linear algebra libraries that handle large sparse matrices [Mat17], and relies on a graph-theoretic interpretation of Gaussian elimination. In particular, the variables and nonzeros in a linear system correspond to vertices and edges in a graph, respectively. When the variable associated with vertex u is eliminated, a clique is induced on the neighborhood of u , and then u is deleted from the graph. This heuristic repeatedly eliminates the vertex of minimum degree in this graph, which corresponds to the variable with the fewest nonzeros in its row and column.

Computing elimination orderings that minimize the number of additional nonzeros, known as *fill*, has been shown to be computationally hard [Yan81, NSS00], even in parameterized settings [KST99, FV13, WAPL14, BCK⁺16, CS17]. However, the practical performance of direct methods has greatly benefited from more efficient algorithms for analyzing elimination orderings [ADD04, DGLN04]. Tools such as elimination trees [Liu90, GNP94] can implicitly represent fill in time that is nearly-linear in the number of *original* nonzeros, which allows for efficient prediction and reorganization of future computation and, more importantly, memory bandwidth. In contrast to the abundance of algorithms built on examining elimination orderings via implicit representation [HP07, NS12], surprisingly little attention has been given to producing elimination orderings implicitly. In the survey by Heggernes et al. [HEKP01], the authors give an $O(n^2m)$ algorithm for computing a minimum degree ordering, which is more than the cost of Gaussian elimination itself and significantly more than the nearly-linear time algorithms for analyzing such orderings [GNP94].

Main Results. We begin our study by combining implicit representations of fill with graph sketching. The nonzero entries of a partially eliminated matrix can be represented as the set of vertices reachable within two hops in a graph that undergoes edge contractions [GNP94]. This allows us to incorporate ℓ_0 -sketches [Coh97], which were originally developed to estimate the cardinality of reachable sets of vertices in directed graphs. By augmenting ℓ_0 -sketches with suitable data structures, we obtain the following result for dynamically maintaining fill structure.

Theorem 1.1. *Against an oblivious adversary, we can maintain $(1 \pm \epsilon)$ -approximations to the degrees of the graph representation of a matrix undergoing elimination in $O(\log^3 n \epsilon^{-2})$ per operation.*

We also give an exact version of this data structure for cases where the minimum degree is always small (e.g., empirical performance of Gaussian elimination on grid graphs [BMMR97]). Ignoring issues of potential dependent randomness, the approximation guarantees of this data structure provide us with an ordering that we call an *approximate greedy minimum degree ordering*, where at each step a vertex whose degree is close to the minimum is pivoted. It is unclear if such an ordering approximates a true minimum degree ordering, but such guarantees are more quantifiable than previous heuristics for approximating minimum degree orderings [ADD96, HEKP01].

However, using this randomized data structure in a greedy manner exposes the severe limitations of data structures that only work in the oblivious adversary model. The updates (i.e. the vertices we eliminate) depend on the output to previous minimum-degree queries, and hence its own internal randomness. The main result in this paper is an algorithm that uses dynamic sketching, as well as an additional routine for estimating degrees via local sampling, to generate an approximate greedy minimum degree sequence in nearly-linear time against adaptive adversaries.

Theorem 1.2. *Given an $n \times n$ matrix A with nonzero graph structure G containing m nonzeros, we can produce a $(1 + \epsilon)$ -approximate greedy minimum degree ordering in $O(m \log^5 n \epsilon^{-2})$ time.*

Techniques. Several components of our algorithm are highly tailored to the minimum degree algorithm. For example, our dynamic sketches and local degree estimation routine depend on the implicit representation of intermediate states of Gaussian elimination [GNP94]. That said, our underlying randomized techniques (e.g., ℓ_0 -sketches [Coh97] and wedge sampling [KP17, ELRS17]) are new additions to combinatorial scientific computing.

The primary focus of this paper is modifying the guarantees in the oblivious adversary model from Theorem 1.1 to work within a greedy loop (i.e. an adaptive adversary) to give Theorem 1.2. However, we do not accomplish this by making the queries deterministic or worst-case as in [BK16, NS17, NSW17]. Instead, we use an external randomized routine for estimating fill degrees to create a fixed sequence of updates. The randomness within the sketching data structure then becomes independent to the update sequence, but its internal state is still highly useful for determining which vertices could have approximate minimum degree. We then efficiently construct the update sequence using recent developments for randomized graph algorithms that use exponential random variables [MPX13, MPVX15]. Our use of sketching can also be viewed as a pseudodeterministic algorithm whose goal is to efficiently recover a particular sequence of vertices [GG11, GGR13]. We believe that both of these views are valuable to the study of randomness and for better understanding the relationship between oblivious and adaptive adversaries.

Organization. In Section 2 we formalize the implicit representation of fill and variants of minimum degree orderings. In Section 3 we give an overview of our results, along with a brief description of the algorithms and techniques we employ. The use of sketching and sampling to obtain our exact and approximate algorithms are given in Section 4 and Section 5, respectively. We also detail our derandomization routine in Section 5, which is crucial for using our randomized data structure against an adaptive adversary. In Section 6 we demonstrate how to estimate fill degrees via local sampling, and in Section 7 we show how to maintain sketches as vertices are pivoted. Lastly, in Section 8 we discuss hardness results for computing the minimum degree of a vertex in a partially eliminated system and also for producing a minimum degree ordering.

2 Preliminaries

We assume that function arguments are pointers to objects instead of the objects themselves, and thus passing an object of size $O(n)$ does not cost $O(n)$ time and space. This is essentially the “pass by reference” construct in high-level programming languages.

2.1 Gaussian Elimination and Fill Graphs

Gaussian elimination is the process of repeatedly eliminating variables from a system of linear equations, while maintaining an equivalent system on the remaining variables. Algebraically, this involves taking an equation involving a target variable and subtracting (a scaled version of) this equation from all others involving the target variable. We assume throughout the paper that the systems are symmetric positive definite (SPD) and thus the diagonal will remain positive, allowing for any pivot order. This further implies that we can apply elimination operations to columns in order to isolate the target variable, resulting in the Schur complement.

A particularly interesting fact about Gaussian elimination is that the numerical Schur complement is unique irrespective of the pivoting order. Under the now standard assumption that nonzero elements do not cancel each other out [GL89], this commutative property also holds for the combinatorial nonzero structure. By interpreting the nonzero structure of a symmetric matrix A as an adjacency matrix for a graph G , we can define the change to the nonzero structure of A as a graph-theoretic operation on G analogous to the Schur complement.

Our notation extends that of Gilbert, Ng, and Peyton [GNP94], who worked with known elimination orderings and treated the entire fill pattern (i.e. additional nonzeros entries) statically. Because we work with partially eliminated states, we will need to distinguish between the *eliminated* and *remaining* vertices in G . We implicitly address this by letting x and y denote eliminated vertices and by letting u , v , and w denote remaining vertices. The following definition of a fill graph allows us to determine the nonzero structure on the remaining variables of a partially eliminated system.

Definition 2.1. The *fill graph* $G^+ = (V^+, E^+)$ is a graph on the remaining vertices such that the edge $(u, v) \in E^+$ if u and v are connected by a (possibly empty) path of eliminated vertices.

This characterization of fill means that we can readily compute the *fill degree* of a vertex v , denoted by $\deg^+(v) = |N^+(v)|$, in a partially eliminated state without explicitly constructing the matrix. We can also iteratively form G^+ from the original graph G by repeatedly removing an eliminated vertex x along with its incident edges, and then adding edges between all of the neighbors of x to form a clique. This operation gives the nonzero structure of the Schur complement.

Lemma 2.2. For any graph $G = (V, E)$ and vertex $v \in V$, given an elimination ordering S we can compute $\deg^+(v)$ at the step when v is eliminated in $O(m)$ time.

Proof. Mark all the vertices appearing in S before v as eliminated, and mark the rest as remaining. Run a breadth-first search from v that terminates at remaining vertices (not including v). Let T be the set of vertices where the search terminated. By the definition of G^+ we have $\deg^+(v) = |T|$. \square

This kind of path finding among eliminated vertices adds an additional layer of complexity to our data structures. To overcome this, we contract eliminated vertices into their connected components (with respect to their induced subgraph in G), which leads to the component graph.

Definition 2.3. We use $G^\circ = (V_{\text{comp}}^\circ, V_{\text{rem}}^\circ, E^\circ)$ to denote the *component graph*. The set of vertices in V_{comp}° is formed by contracting edges between eliminated vertices, and the set of vertices that have not been eliminated is V_{rem}° . The set of edges E° is implicitly given by the contractions.

Note that G° is quasi-bipartite, as the contraction rule implies there are no edges between vertices in V_{comp}° . It will be useful to refer to two different kinds of neighborhoods in a component graph. For any vertex v in G° , let $N_{\text{rem}}^\circ(v)$ be the set of neighbors of v are in V_{rem}° , and let $N_{\text{comp}}^\circ(v)$ denote the neighbors of v that are in V_{comp}° . Analogously, we use the notation $\text{deg}_{\text{rem}}^\circ(v) = |N_{\text{rem}}^\circ(v)|$ and $\text{deg}_{\text{comp}}^\circ(v) = |N_{\text{comp}}^\circ(v)|$.

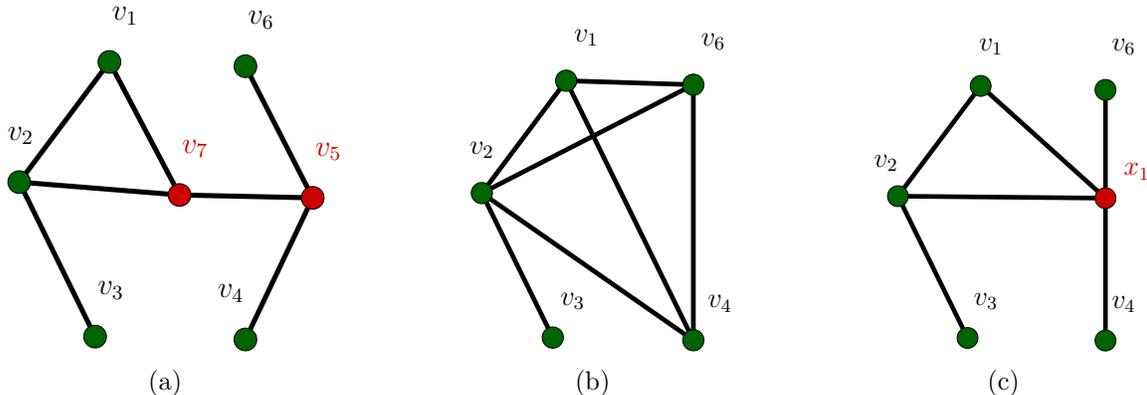


Figure 1: The (a) original graph, (b) fill graph, and (c) component graph after pivoting v_5 and v_7 .

For example, let us consider Figure 1. The original graph G has seven vertices v_1, v_2, \dots, v_7 , and the algorithm decides to pivot v_5 and v_7 marked in red. Eliminating these vertices induces a clique on v_1, v_2, v_4, v_6 in the fill graph because each pair of vertices is connected by a path through the eliminated vertices v_5 and v_7 . Our algorithms implicitly maintain the fill graph by maintaining the component graph, where v_5 and v_7 merge to form the connected component x_1 with edges incident to all remaining neighbors of v_5 and v_7 in the original graph. Note that an upper bound for the number of edges in a component graph is the number of edges in the original graph. We repeatedly exploit this property when proving the time and space bounds of our algorithms.

2.2 Minimum Degree Orderings

The minimum degree algorithm is a greedy heuristic for reducing the cost of solving sparse linear systems that repeatedly eliminates the variable involved in the fewest number of equations [GL89]. Although there are many situations where this is suboptimal, it is remarkably effective and widely used in practice. For example, the approximate minimum degree algorithm (AMD) [ADD96] is a heuristic for generating minimum degree orderings that plays an integral role in the sparse linear algebra packages in MATLAB [Mat17], Mathematica [Wol18], and Julia [BKSE12].

For any elimination ordering (u_1, u_2, \dots, u_n) , we let G_i be the graph with vertices u_1, u_2, \dots, u_i marked as eliminated and $u_{i+1}, u_{i+2}, \dots, u_n$ marked as remaining. We denote the corresponding sequence of fill graphs by $(G_0^+, G_1^+, \dots, G_n^+)$, where $G_0^+ = G$ and G_n^+ is the empty graph. Throughout the paper, we frequently use the notation $[n] = \{1, 2, \dots, n\}$ when iterating over sets.

Definition 2.4. A *minimum degree ordering* is an elimination ordering such that for all $i \in [n]$, the vertex u_i has minimum fill degree in G_{i-1}^+ . Concretely, this means that

$$\deg_{i-1}^+(u_i) = \min_{v \in V_{i-1}^+} \deg_{i-1}^+(v).$$

The data structures we use for finding the vertices with minimum fill degree are randomized, so we need to be careful to not introduce dependencies between different steps of the algorithm when several vertices are of minimum degree. To avoid this problem, we simply require that the lexicographically-least vertex be eliminated in the event of a tie.

Our notion for approximating a minimum degree ordering is based on finding a vertex at each step whose degree is close to the minimum in G_t^+ . Note that this is the goal of the AMD algorithm.

Definition 2.5. A $(1 + \epsilon)$ -*approximate greedy minimum degree ordering* is an elimination ordering such that at each step $i \in [n]$, we have

$$\deg_{i-1}^+(u_i) \leq (1 + \epsilon) \min_{v \in V_{i-1}^+} \deg_{i-1}^+(v).$$

This decision process has no lookahead, and thus does not in any way approximate the minimum possible total fill incurred during Gaussian elimination, which is known to be NP-complete [Yan81].

2.3 Related Works

Gaussian Elimination and Fill. The study of pivoting orderings is a fundamental question in combinatorial scientific computing. Work by George [Geo73] led to the study of nested dissection algorithms, which utilize separators to give provably smaller fill bounds for planar [RTL76, LRT79] and separable graphs [GT87, AY10]. A side effect of this work is the better characterization of fill via component graphs [Liu85], which is used to compute the total fill-in of a specific elimination ordering [GNP94]. This characterization is also used to construct elimination trees, which are ubiquitous in scientific computing to preallocate memory and optimize cache behaviors [Liu90].

Finding Low-Fill Orderings. The goal of an elimination ordering is to minimize the total fill. Unfortunately, this problem is NP-complete [Yan81, BS90]. Algorithms that approximate the minimum fill-in within polynomial factors have been studied [NSS00], as well as algorithms [KST99, FV13] and hardness results [WAPL14, BCK⁺16, CS17] for parameterized variants. Partially due to the high overhead of the previous algorithms, the minimum degree heuristic remains as one of the most widely-used methods for generating low-fill orderings [GL89].

Somewhat surprisingly, we were not able to find prior works that compute the minimum degree ordering in time faster than $O(n^3)$ or works that utilize the implicit representation of fill provided by elimination trees.¹ On the other hand, there are various heuristics for finding minimum degree-like orderings, including multiple minimum degree (MMD) [Liu85] and the approximate minimum degree algorithm (AMD) [ADD96]. While both of these methods run extremely well in practice, they have theoretically tight performances of $\Theta(n^2m)$ for MMD and $\Theta(nm)$ for AMD [HEKP01]. Furthermore, AMD is not always guaranteed to produce a vertex of approximate minimum degree.

¹ We use speculative language here due to the vastness of the literature on variants of minimum degree algorithms.

Removing Dependencies in Randomized Algorithms. Our size estimators are dynamic—the choice of the pivot, which directly affects the subsequent fill graph, is a result of the randomness used to generate the pivot in the previous step—and prone to dependency problems. Independence between the access sequence and internal randomness is a common requirement in recent works on data structures for maintaining spanning trees, spanners, and matchings [BGS15, KKM13, Sol16]. Often these algorithms only have guarantees in the oblivious adversary model, which states that the adversary can choose the graph and the sequence of updates, but it cannot choose updates adaptively in response to the randomly-guided choices of the algorithm.

Recent works in randomized dimensionality-reduction have approached this issue of dependency by injecting additional randomness to preserve independence [LS15]. Quantifying the amount of randomness that is “lost” over the course of an algorithm has recently been characterized using mutual information [KNP⁺17], but their results do not allow for us to consider n adversarial vertex pivots. Our analysis also has tenuous connections to recent works utilizing matrix martingales to analyze repeated introductions of randomness into graph sparsification algorithms [KS16, KPPS17].

3 Overview

We discuss the main components of our algorithms in three parts. In Section 3.1 we explore how dynamic graph sketching can be used to construct a randomized data structure that maintains approximate degrees under vertex eliminations. In Section 3.2 we demonstrate how data structures that work against oblivious adversaries can fail against adaptive adversaries. We also describe our approach to circumvent this problem for approximate minimum degree sequences. In Section 3.3 we discuss a local degree estimation routine (the new key primitive) in the context of estimating the number of nonzero columns of a matrix via sampling. Finally, in Section 3.4 we explain the implications of our results to the study of algorithms for computing elimination orderings.

3.1 Dynamically Sketching Fill Graphs

The core problem of estimating fill degrees can be viewed as estimating the cardinality of sets undergoing unions and deletion of elements. Cardinality estimation algorithms in the streaming algorithm literature often trade accuracy for space [FM85, CM05], but our degree-approximation data structures use sketching to trade space for accuracy and more efficient update operations.

We first explain the connection between computing fill degrees and estimating the size of reachable sets. Assume for simplicity that no edges exist between the remaining vertices in the component graph G° . Split each remaining vertex u into two vertices u_1 and u_2 , and replace every edge (u, x) to a component vertex x by the directed edges (u_1, x) and (x, u_2) . The fill degree of u is the number of remaining vertices v_2 reachable from u_1 (not including u_1). Cohen [Coh97] developed a nearly-linear time size-estimation framework for reachability problems using sketching and ℓ_0 -estimators. Adapting this framework to our setting for fill graphs leads to the following kind of ℓ_0 -sketch data structure. We refer to the set $N(u) \cup \{u\}$ as the *1-neighborhood* of u , and we call its cardinality $\deg(u) + 1$ the *1-degree* of u .

Definition 3.1. A *1-neighborhood ℓ_0 -sketch* of a graph G is constructed as follows:

1. Each vertex $u \in V$ independently generates a random key $R(u)$ uniformly from $[0, 1)$.

2. Then each vertex determines which of its neighbors (including itself) has the smallest key. We denote this by the function

$$\text{MINIMIZER}(u) \stackrel{\text{def}}{=} \arg \min_{v \in N(u) \cup \{u\}} R(v).$$

To give some intuition for how sketching is used to estimate cardinality, observe that choosing keys independently and uniformly at random essentially assigns a random vertex $N(u) \cup \{u\}$ to be $\text{MINIMIZER}(u)$. Therefore, the key value $R(\text{MINIMIZER}(u))$ is correlated with $\deg(u) + 1$. This correlation is the cornerstone of sketching. If we construct $k = \Omega(\log n \epsilon^{-2})$ independent sketches, then by concentration we can use an order statistic of $R_i(\text{MINIMIZER}(u))$ over all k sketches to give an ϵ -approximation of $\deg(u) + 1$ with high probability. We give the full details in Appendix A.

To maintain sketches of the fill graph as it undergoes vertex eliminations, we first need to implicitly maintain the component graph G° (Lemma 6.2). We demonstrate how to efficiently propagate key values in a sketch as vertices are pivoted in Section 7. For now, it is sufficient to know that each vertex in a sketch has an associated min-heap that it uses to report and update its minimizer. Because eliminating vertices leads to edge contractions in the component graph, there is an additional layer of intricacies that we need to resolve using amortized analysis.

Suppose v is the vertex eliminated as we go from G_t° to G_{t+1}° . The sketch propagates this information to relevant vertices in the graph using a two-level notification mechanism. The neighbors of v are informed first, and then they notify their neighbors about the change, all the while updating the key values in their heaps. We outline the subroutine $\text{PIVOTVERTEX}(v)$ that accomplishes this:

1. Update the min-heaps of every remaining neighbor of v .
2. For each component neighbor x of v , if the minimum in its heap changes, then propagate this change to the remaining neighbors of x and merge x with v .

While it is simple enough to see that this algorithm correctly maintains key values, bounding its running time is nontrivial and requires a careful amortized analysis to show that the bottleneck operation is the merging of component vertices.

We can merge two min-heaps in $O(\log^2 n)$ time, so merging at most n heaps takes $O(n \log^2 n)$ time in total. To bound the cost of heap updates due to merges, we define the potential of the component graph as

$$\Phi(G_t^\circ) \stackrel{\text{def}}{=} \sum_{u \in V_{\text{comp},t}^\circ} D(u) \log(D(u)),$$

where $D(u)$ is the sum of the original degrees of vertices merged into u . Using the fact that a merge operation only informs neighbors of at most one of the two merged vertices, we are able to show that the number of updates to produce G_t° is of the order of $\Phi(G_t^\circ)$. It follows that the total number of updates is at most $O(m \log n)$, which gives us a total update cost of $O(m \log^2 n)$.

3.2 Correlation and Decorrelation

We now discuss how we use the randomized sketching data structure within a greedy algorithm. We start with a simple concrete example to illustrate a problem that an adaptive adversary can cause. Consider a data structure that uses sketching to estimate the cardinality of a subset $S \subseteq [n]$

under the insertion and deletion of elements. This data structure randomly generates a subset of keys $T \subseteq [n]$ such that $|T| = \Theta(\log n \epsilon^{-2})$, and it returns as its estimate the scaled intersection

$$n \cdot \frac{|S \cap T|}{|T|},$$

which is guaranteed to be within an ϵn -additive error of the true value $|S|$ by Chernoff bounds, assuming that T is generated independently of S . Clearly this cardinality-estimation algorithm works in the oblivious adversary model.

However, an adaptive adversary can use answers to previous queries to infer the set of secret keys T in $O(n)$ updates and queries. Consider the following scheme in Figure 2 that returns $S = T$.

1. Initialize $S = [n]$.
2. For each $i = 1$ to n :
 - (a) Delete i from S . If the estimated size of S changed, reinsert i into S .
3. Return S .

Figure 2: An adaptive routine that amplifies the error of a cardinality-estimation scheme that uses a fixed underlying sketch.

While the updates performed by a greedy algorithm are less extreme than this, in the setting where we maintain the cardinality of the smallest of k dynamic sets, having access to elements in the minimizer does allow for this kind of sketch deduction. Any accounting of correlation (in the standard sense) also allows for worst-case kinds of adaptive behavior, similar to the scheme above.

To remove potential correlation, we use an external routine that is analogous to the local degree-estimation algorithm used in the approximate minimum degree algorithm, which runs in time close to the degree it estimates. In this simplified example, suppose for each cardinality query that the data structure first regenerates T . Then the probability that i belongs to S is $\Theta(\log n \epsilon^2/n)$. Stepping through all $i \in [n]$, it follows that the expected number of deletions is $\Theta(\log n \epsilon^{-2})$, and hence S remains close to size n with high probability.

Reinjecting randomness is a standard method for decorrelating a data structure across steps. However, if we extend this example to the setting where we maintain the cardinality of k sets (similar to our minimum degree algorithm), then the previous idea requires that we reestimate the size of every set to determine the one with minimum cardinality. As a result, this approach is prohibitively expensive. However, these kinds of cardinality estimations are actually local—meaning that it is sufficient to instead work with a small and accurate subset of candidates sets. If we compute the set with minimum cardinality among the candidates using an external estimation scheme, then this decision is independent of the random choice of T in the sketching data structure, which then allows us to use the sketching data structure to generate the list candidates.

Our algorithm for generating an approximate greedy minimum degree ordering relies on a similar external routine called `ESTIMATEFILL1DEGREE(u, ϵ)`, which locally estimates the fill 1-degree of u at any step of the algorithm in time proportional to $\deg(u)$ in the original graph. We further

describe this estimator in Section 3.3 and present the full sampling algorithm in Section 6. In Section 5 we show that to generate an approximate greedy minimum degree sequence, it is instead sufficient to pivot the vertex

$$\arg \min_{u \in V^+} \left(1 - \frac{\epsilon \cdot \text{Exp}(1)}{\log n} \right) \cdot \text{ESTIMATEFILL1DEGREE} \left(u, \frac{\epsilon}{\log n} \right)$$

at each step, which we call the ϵ -decayed minimum over all external estimates.

Analogous to the discussion about the set cardinality estimation above, evaluating the degrees of every remaining vertex using ESTIMATEFILL1DEGREE at each step is expensive and leads to a total cost of $\Omega(nm)$. However, we can reincorporate the sketching data structure and use the following observations about the perturbation coefficient involving the exponential random variable $\text{Exp}(1)$ to sample a small number of candidate vertices that contains the ϵ -decayed minimum.

- For a set of vertices whose degrees are within $1 \pm \epsilon / \log n$ of each other, it suffices to randomly select and consider $O(1)$ of them by generating the highest order statistics of exponential random variables in decreasing order.
- By the memoryless property of the exponential distribution, if we call ESTIMATEFILL1DEGREE, then with constant probability it will be for the vertex we pivot. Therefore, we can charge the cost of these evaluations to the original edge count and retain a nearly-linear running time.

Invoking ESTIMATEFILL1DEGREE only on the candidate vertices allows us to efficiently find the ϵ -decayed minimizer in each step, which leads to the nearly-linear runtime as stated in Theorem 1.2. The key idea is that any dependence on the ℓ_0 -sketches stops after the candidates are generated, since their degrees only depend on the randomness of an external cardinality-estimation routine.

3.3 Local Estimation of Fill Degrees

A critical part of the approximate min-degree algorithm is the ESTIMATEFILL1DEGREE function, which estimates the fill 1-degree of a vertex $u \in V^+$ using fresh randomness and $O(\deg(u) \log^2 n \epsilon^{-2})$ oracle queries to the component graph G° . At the beginning of Section 6 we show how to construct a $(0, 1)$ -matrix A where each row corresponds to a remaining neighborhood of a component neighbor of u . The number of nonzero columns in A is equal to $\deg^+(u)$. Using only the following matrix operations (which correspond to component graph oracle queries), we analyze the more general problem of counting the number of nonzero columns in a matrix. We note that this technique is closely related to recent results in wedge sampling for triangle counting [KP17, ELRS17].

- $\text{ROWSIZE}(A, i)$: Returns the number of nonzero elements in row i of A .
- $\text{SAMPLEFROMROW}(A, i)$: Returns a column index j uniformly at random from the nonzero entries of row i of A .
- $\text{QUERYVALUE}(A, i, j)$: Returns the value of $A(i, j)$.

Lemma 6.1. *There is a routine ESTIMATENONZEROCOLUMNS using the three operations above that takes as input (implicit) access to a matrix A and an error ϵ , and returns an ϵ -approximation to the number of nonzero columns in A with high probability. The expected total number of operations used is $O(r \log^2 n \epsilon^{-2})$, where r is the number of rows and n is the number of columns in A .*

We now give an overview of how ESTIMATENONZEROCOLUMNS works. Let B be the normalized version of A where every nonzero entry is divided by its column sum. The sum of the nonzero entries in B is the number of nonzero columns in A , denoted by $\text{NONZEROCOLUMNS}(A)$. If we uniformly sample a nonzero entry of B , then the mean of this distribution is $\text{NONZEROCOLUMNS}(A)/\text{nnz}(A)$. Because random variables sampled from this distribution take their value in $[0, 1]$, we can estimate their mean using an ESTIMATEMEAN subroutine (Lemma 6.4), which does the following:

1. Set a threshold σ depending on the accuracy of the desired estimate.
2. Sample k independent random variables from the distribution until their sum first exceeds σ .
3. Return σ/k .

Using the matrix operations above, we can easily sample indices (i, j) of nonzero entries in B , but evaluating $B(i, j)$ requires that we know the j -th column sum of A . Therefore, to compute this column sum we estimate the mean of a Bernoulli distribution on the j -th column of A defined by selecting an entry from $A[:, j]$ uniformly at random. This distribution has mean $\text{nnz}(A[:, j])/r$, and it is amenable to sampling using the provided operations.

While the previous estimator works satisfactorily, we show how to combine these distributions and use a hitting time argument to reduce the sample complexity by a factor of $O(\epsilon^{-2})$. Specifically, for a fixed column, we consider a random variable that has a limited number of attempts to find a nonzero entry by uniformly sampling rows. By optimizing the number of attempts, we can reduce our error overhead in the runtime at the expense of a $1/\text{poly}(n)$ perturbation to the approximation.

3.4 Significance to Combinatorial Scientific Computing

Despite the unlikelihood of theoretical gains for solving linear systems by improved direct methods for sparse Gaussian elimination, we believe our study could influence combinatorial scientific computing in several ways. First, we provide evidence in Section 8 for the nonexistence of nearly-linear time algorithms for finding exact minimum degree orderings by proving conditional hardness results. Our reduction uses the observation that determining if a graph can be covered by a particular union of cliques (or equivalently, that the fill graph is a clique after eliminating certain vertices) is equivalent to the orthogonal vectors problem [Wil05]. Assuming the strong exponential time hypothesis, this leads to a conditional hardness of $\Omega(m^{4/3-\theta})$ for computing a minimum degree ordering. However, we believe that this result is suboptimal and that a more careful construction could lead to $\Omega(nm^{1-\theta})$ -hardness.

On the other hand, advances in minimum degree algorithms cannot be justified in practice solely by worst-case asymptotic arguments. In general, nested dissection orderings are asymptotically superior in quality to minimum degree orderings [HR98]. Furthermore, methods based on Krylov spaces, multiscale analysis, and iterative methods [Gut07, GGLN13] are becoming increasingly popular as they continue to improve state-of-the-art solvers for large sparse systems. Such advancements are also starting to be reflected in theoretical works. As a result, from both a theoretical and practical perspective, we believe that the most interesting question related to minimum degree algorithms is whether or not such sequences lead to computational gains for problems of moderate size.

In our approximate minimum degree algorithm, the $O(\log^5 n)$ term and convenient choice of constants preclude it from readily impacting elimination orderings in practice. However, the under-

lying sketching technique is quite flexible. For example, consider modifying the dynamic ℓ_0 -sketches such that:

1. Each vertex maintains k random numbers in the range $[0, 1)$.
2. Each component vertex maintains the smallest k numbers of its remaining neighbors.
3. Each remaining vertex then maintains the smallest k numbers among its component neighbors.

If we repeatedly eliminate the vertex whose median is largest, this routine is similar to using k copies of the previous type of the sketch. Letting $k = \Omega(\log n \epsilon^{-2})$, we can analyze this variant against an oblivious adversary using slight modifications to our original sketching algorithm [Mas00]. Although our analysis demonstrates that new tools are necessary for studying its behavior within a greedy algorithm, we experimentally observed desirable behavior for such sketches. Therefore, we plan to continue studying this kind of adaptive graph sketching both theoretically and experimentally.

4 Sketching Algorithms for Computing Degrees

Let us recall a few relevant definitions from Section 2 for convenience. For a given vertex elimination sequence (u_1, u_2, \dots, u_n) , let G_t^+ denote the fill graph obtained by pivoting vertices u_1, u_2, \dots, u_t , and let δ_t denote the minimum degree in G_t^+ . An ℓ_0 -sketch data structure consists of the following:

1. Each vertex u independently generates a key $R(u)$ from $[0, 1)$ uniformly at random.
2. Then each vertex u determines which neighbor (including itself) has the smallest key value. We denote this neighbor by $\text{MINIMIZER}(u)$.

In this section we show that if an ℓ_0 -sketch can efficiently be maintained for a dynamic graph, then we can use the same set of sketches at each step to determine the vertex with minimum fill degree and eliminate it. We explore the dynamic ℓ_0 -sketch data structure for efficiently propagating key values under pivots in detail in Section 7 (and for now we interface it via Theorem 4.4). This technique leads to improved algorithms for computing the minimum degree ordering of a graph, which we analyze in three different settings.

First, we consider the case where the minimum degree at each step is bounded. In this case we choose a fixed number of ℓ_0 -sketches and keep track of every minimizer of a vertex over all of the sketch copies. Note that we can always use n as an upper bound on the minimum fill degree.

Theorem 4.1. *There is an algorithm DELTACAPPEDMINDEGREE that, when given a graph with a lexicographically-first min-degree ordering whose minimum degree is always bounded by Δ , outputs this ordering with high probability in expected time $O(m\Delta \log^3 n)$ and uses space $O(m\Delta \log n)$.*

Next, we relax the bound on the minimum degrees over all steps of the algorithm and allow the time and space complexity to be output sensitive by adaptively increasing the number of ℓ_0 -sketches as the algorithm progresses.

Theorem 4.2. *There is an algorithm OUTPUTSENSITIVEMINDEGREE that, when given a graph with a lexicographically-first min-degree sequence $(\delta_1, \delta_2, \dots, \delta_n)$, outputs this ordering with high probability in expected time $O(m \cdot \max_{t \in [n]} \delta_t \cdot \log^3 n)$ and uses space $O(m \cdot \max_{t \in [n]} \delta_t \cdot \log n)$.*

Lastly, we modify the algorithm to compute an approximate minimum degree vertex at each step. By maintaining $\Theta(\log n \epsilon^{-2})$ copies of the ℓ_0 -sketch data structure, we are able to accurately approximate the 1-degree of a vertex using the $(1 - 1/e)$ -th² order statistic of the key values of its minimizers. We abstract this idea using the following approximate degree data structure, which when given an elimination ordering directly leads to a nearly-linear time algorithm.

Theorem 4.3. *There is a data structure `ApproxDegreeDS` that supports the following methods:*

- `APPROXDEGREEEDS_PIVOT(u)`, which pivots a remaining vertex u .
- `APPROXDEGREEEDS_REPORT()`, which provides balanced binary search tree (BST) containers V_1, V_2, \dots, V_B such that all the vertices in the bucket V_i have 1-degree in the range

$$\left[(1 + \epsilon)^{i-2}, (1 + \epsilon)^{i+2} \right].$$

The memory usage of this data structure is $O(m \log n \epsilon^{-2})$. Furthermore, if the pivots are picked independently from the randomness used in this data structure (i.e., we work under the oblivious adversary model) then:

- The total cost of all the calls to `APPROXDEGREEEDS_PIVOT` is bounded by $O(m \log^3 n \epsilon^{-2})$.
- The cost of each call to `APPROXDEGREEEDS_REPORT` is bounded by $O(\log^2 n \epsilon^{-1})$.

4.1 Computing the Exact Minimum Degree Ordering

We first consider the case where the minimum degree in each of the fill graphs G_t^+ is at most Δ . In this case, we maintain $k = O(\Delta \log n)$ copies of the ℓ_0 -sketch data structure. By a coupon collector argument, any vertex with degree at most Δ contains all of its neighbors in its list of minimizers with high probability. This implies that for each $t \in [n]$, we can obtain the exact minimum degree in G_t^+ with high probability. Figure 3 briefly describes the data structures we will maintain for this version of the algorithm.

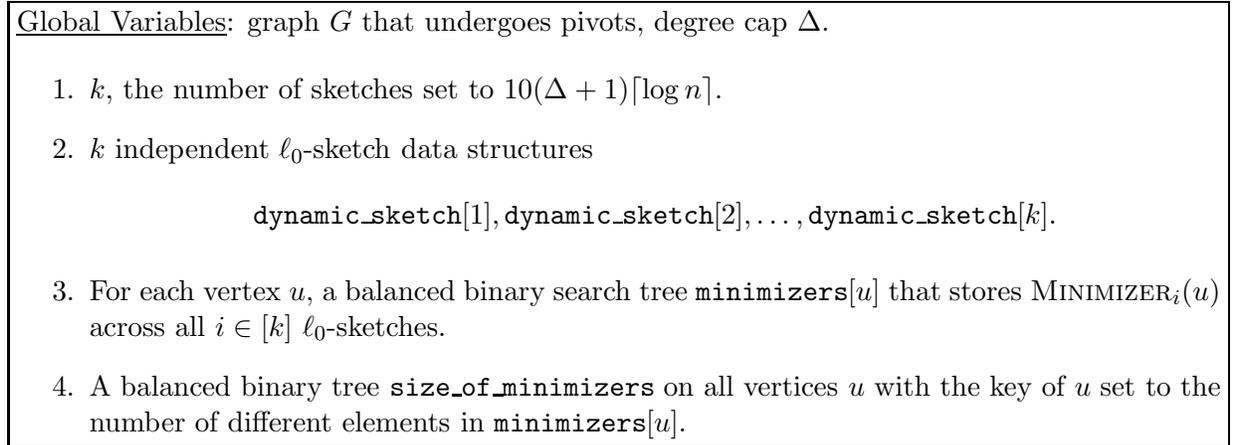


Figure 3: Global variables for the Δ -capped min-degree algorithm `DELTACAPPEDMINDEGREE`.

²Note that we use e to refer to the base of the natural logarithm.

Note that if we can efficiently maintain the data structures in Figure 3, then querying the minimum element in `size_of_minimizers` returns the (lexicographically-least) vertex with minimum degree. Theorem 4.4 demonstrates that we can maintain the ℓ_0 -sketch data structures efficiently.

Theorem 4.4. *Given i.i.d. random variables $R(v)$ associated with each vertex $v \in V_t^+$, there is a data structure `DynamicSketch` that, for each vertex u , maintains the vertex with minimum $R(v)$ among itself and its neighbors in G_t^+ . This data structure supports the following methods:*

- `QUERYMIN(u)`, which returns `MINIMIZER(u)` for a remaining vertex u in $O(1)$ time.
- `PIVOTVERTEX(u)`, which pivots a remaining vertex u and returns the list of all remaining vertices v whose value of `MINIMIZER(v)` changed immediately after this pivot.

The memory usage of this data structure is $O(m)$. Moreover, for any choice of key values $R(v)$:

- The total cost of all the pivots is $O(m \log^2 n)$.
- The total size of all lists returned by `PIVOTVERTEX` over all steps is $O(m \log n)$.

This theorem relies on intermediate data structures described in Section 7, so we defer the proof until the end of that section. Note that this `DynamicSketch` data structure will be essential to all three min-degree algorithms.

Now consider a sketch of G^+ and a vertex u with degree $\deg^+(u) \leq \Delta$. By symmetry of the $R(v)$ values, each vertex in $N^+(u) \cup \{u\}$ is the minimizer of u with probability $1/(\deg^+(u) + 1)$. Therefore, if we maintain $O(\Delta \log n)$ independent ℓ_0 -sketches, we can ensure that we have an accurate estimation of the minimum fill degree with high probability. The pseudocode for this routine is given in Figure 4. We formalize the probability guarantees in Lemma 4.5 and Lemma 4.6, which are essentially a restatement of [Coh97, Theorem 2.1].

<p>DELTA_CAPPED_MIN_DEGREE(G, Δ) <u>Input:</u> graph $G = (V, E)$, threshold Δ. <u>Output:</u> exact lexicographically-first min-degree ordering (u_1, u_2, \dots, u_n).</p> <ol style="list-style-type: none"> 1. For each step $t = 1$ to n: <ol style="list-style-type: none"> (a) Set $u_t \leftarrow \text{min}(\text{size_of_minimizers})$. (b) DELTA_CAPPED_MIN_DEGREE_PIVOT(u_t). 2. Return (u_1, u_2, \dots, u_n).
--

<p>DELTA_CAPPED_MIN_DEGREE_PIVOT(u) <u>Input:</u> vertex to be pivoted u. <u>Output:</u> updated global state.</p> <ol style="list-style-type: none"> 1. For each sketch $i = 1$ to k: <ol style="list-style-type: none"> (a) $(v_1, v_2, \dots, v_\ell) \leftarrow \text{dynamic_sketch}[i].\text{PIVOT_VERTEX}(u)$, the set of vertices in the i-th sketch whose minimizers changed after pivoting out u. (b) For each $j = 1$ to ℓ: <ol style="list-style-type: none"> i. Update the values corresponding to sketch i in $\text{minimizers}[v_j]$. ii. Update the entry for v_j in $\text{size_of_minimizers}$ with the size of $\text{minimizers}[v_j]$.

Figure 4: Pseudocode for the exact Δ -capped min-degree algorithm, which utilizes the global data structures for DELTA_CAPPED_MIN_DEGREE defined in Figure 3.

Lemma 4.5. *With high probability, for all remaining vertices u such that $\text{deg}^+(u) \leq 2\Delta$ we have*

$$\text{size_of_minimizers}[u] = \text{deg}^+(u) + 1.$$

Proof. The only way we can have $\text{size_of_minimizers}[u] < \text{deg}^+(u) + 1$ is if at least one neighbor of u or u itself is not present in $\text{minimizers}[u]$. Let v be an arbitrary vertex in $N^+(u) \cup \{u\}$. The probability of v not being the minimizer in any of the $k = 10(\Delta + 1)\lceil \log n \rceil$ sketches is

$$\begin{aligned} \Pr [\text{MINIMIZER}_i(u) \neq v \text{ for all } i \in [k]] &= \left(1 - \frac{1}{\text{deg}^+(u) + 1}\right)^k \\ &\leq \left(1 - \frac{1}{2\Delta + 1}\right)^{10(\Delta+1)\log n} \\ &\leq \exp\left(-\frac{10(\Delta + 1)\log n}{2\Delta + 1}\right) \\ &\leq \frac{1}{n^5}. \end{aligned}$$

We can upper bound the probability that there exists a vertex $v \in N^+(u) \cup \{u\}$ not in `minimizers`[u] using a union bound. It follows that

$$\begin{aligned} \Pr [\text{size_of_minimizers}[u] < \deg^+(u) + 1] &\leq |N^+(u) \cup \{u\}| \cdot \Pr [\text{MINIMIZER}_i(u) \neq v \text{ for all } i \in [k]] \\ &\leq \frac{1}{n^4}. \end{aligned}$$

Using a second union bound for the event that there exists a vertex $u \in V^+$ such that $\deg^+(u) \leq 2\Delta$ and `size_of_minimizers`[u] $< \deg^+(u) + 1$ completes the proof. \square

Lemma 4.6. *With high probability, for all remaining vertices u with $\deg^+(u) > 2\Delta$ we have*

$$\text{size_of_minimizers}[u] > \Delta + 1.$$

Proof. We first upper bound the probability of the event `size_of_minimizers`[u] $\leq \Delta + 1$. Let S be any subset of $N^+(u) \cup \{u\}$ of size $\deg^+(u) - \Delta > 0$. Using the assumption that $\deg^+(u) \geq 2\Delta + 1$,

$$\begin{aligned} \Pr [S \cap \text{minimizers}[u] = \emptyset] &= \left(1 - \frac{\deg^+(u) - \Delta}{\deg^+(u) + 1}\right)^k \\ &\leq \exp\left(-\frac{10(\deg^+(u) - \Delta)(\Delta + 1) \log n}{\deg^+(u) + 1}\right) \\ &\leq \exp\left(-\frac{5(\deg^+(u) + 1)(\Delta + 1) \log n}{\deg^+(u) + 1}\right) \\ &= \frac{1}{n^{5(\Delta+1)}}. \end{aligned}$$

Next, sum over all choices of the set S and use a union bound. Note that this over counts events, but this suffices for an upper bound. It follows that

$$\begin{aligned} \Pr [\text{size_of_minimizers}[u] \leq \Delta + 1] &\leq \binom{\deg^+(u) + 1}{\deg^+(u) - \Delta} \cdot \Pr [S \cap \text{minimizers}[u] = \emptyset] \\ &\leq \binom{\deg^+(u) + 1}{\Delta + 1} \cdot \frac{1}{n^{5(\Delta+1)}} \\ &\leq \frac{n^{\Delta+1}}{n^{5(\Delta+1)}} \\ &\leq \frac{1}{n^4}. \end{aligned}$$

Using a second union bound for the event that there exists a vertex $u \in V^+$ such that $\deg^+(u) > 2\Delta$ and `size_of_minimizers`[u] $\leq \Delta + 1$ completes the proof. \square

Proof of Theorem 4.1. The algorithm correctly pivots the minimum degree vertex by Lemma 4.5 and Lemma 4.6. For the space complexity, each of the k ℓ_0 -sketch data structures uses $O(m)$ memory by Theorem 4.4, and for each vertex there is a corresponding balanced binary search tree `minimizers` which uses $O(k)$ space. We also have `size_of_minimizers`, which uses $O(n)$ space. Therefore, since we assume $m \geq n$, the total space is $O(km + nk + n) = O(m\Delta \log n)$.

For the running time, Theorem 4.4 gives a cost of $O(m \log^2 n)$ across all PIVOTVERTEX calls per sketch, and thus a total cost of $O(m \Delta \log^3 n)$ over all sketches. Theorem 4.4 also states the sum of ℓ (the length of the update lists) across all steps is at most $O(m \log n)$. Each of these updates leads to two BST updates, so the total overhead is $O(m \log^2 n)$, which is an equal order term. \square

4.2 Modifying the Algorithm to be Output Sensitive by Adaptive Sketching

If we do away with the condition that minimum fill degrees are bounded above by Δ , then the number of copies of the ℓ_0 -sketch data structure needed depends only on the values of the minimum fill degree at each step. Therefore, we can modify DELTACAPPEDMINDEGREE to potentially be more efficient by adaptively maintaining the required number of sketches.

To accurately estimate degrees in G_t^+ we need $\Omega(\delta_t \log n)$ copies of the ℓ_0 -sketch data structure, but we do not know the values of δ_t a priori. To rectify this, consider the following scheme that adaptively keeps a sufficient number of copies of the ℓ_0 -sketch data structures. First, initialize the value $c = \delta_0$ (the minimum degree in G). Then for each step $t = 1$ to n update c according to:

1. Let $\delta_t(c)$ be the candidate minimum degree in G_t^+ using $k = 10c \lceil \log n \rceil$ sketches.
2. If $\delta_t(c) > c/2$, then set $c \leftarrow 2c$ and repeat.

The core idea of the routine above is that if the candidate minimum degree is at most $c/2$, then with high probability the true minimum degree is at most c . It follows that using $O(c \log n)$ sketching data structures guarantees the minimum degree estimate is correct with high probability.

Proof of Theorem 4.2. The proof is analogous to that of Theorem 4.1. The upper bound for the minimum degrees is now $\Delta = 2 \cdot \max_{t \in [n]} \delta_t$, and so the time and space complexities follow. \square

4.3 Computing an Approximate Minimum Degree

To avoid bounding the minimum fill degree over all steps and to make the running time independent of the output, we modify the previous algorithms to obtain an approximate min-degree vertex at each step. We reduce the number of ℓ_0 -sketches and use the reciprocal of the $(1 - 1/e)$ -th order statistic to approximate the cardinality `size_of_minimizers`[u] (and hence the 1-degree of u) to obtain a nearly-linear time approximation algorithm.

There is, however, a subtle issue with the randomness involved with this algorithm. A necessary condition for the algorithm to succeed as intended is that the sketches at each step are independent of the past decisions of the algorithm. Therefore, we must remove all dependencies between previous and current queries. In Section 3.2 we demonstrate how correlations between steps can amplify. To avoid this problem, we must decorrelate the current state of the sketches from earlier pivoting updates to the data structures. We carefully address this issue in Section 5. Instead of simply selecting a vertex with an approximate min-degree, this algorithm instead requires access to all vertices whose estimated degree is within a certain range of values. Therefore, this approximation algorithm uses a bucketing data structure, as opposed to the previous two versions that output the vertex to be pivoted. Figure 5 describes the global data structures for this version of the algorithm.

Global Variables: graph G , error tolerance $\epsilon > 0$.

1. k , the number of sketches set to $50 \lceil \log n \epsilon^{-2} \rceil$.

2. k independent ℓ_0 -sketch data structures

`dynamic_sketch[1], dynamic_sketch[2], ..., dynamic_sketch[k]`.

3. For each vertex u , a balanced binary search tree `minimizers[u]` that stores `MINIMIZERi(u)` across all $i \in [k]$ ℓ_0 -sketches, and maintains the element in `minimizers[u]` with rank

$$\left\lfloor k \left(1 - \frac{1}{e}\right) \right\rfloor.$$

4. A balanced binary tree `quantile` over all vertices u whose key is the $\lfloor k(1 - 1/e) \rfloor$ -ranked element in `minimizers[u]`.

Figure 5: Global variables and data structures for APPROXDEGREEDS, which returns (implicit) partitions of vertices into buckets with ϵ -approximate degrees.

To successfully use fewer sketches, for a given vertex u we estimate the cardinality of the set of its minimizers via its order statistics instead of using the exact cardinality as we did before with the binary search tree `size_of_minimizers[u]`. Exploiting correlations in the order statistics of sketches is often the underlying idea behind efficient cardinality estimation. In particular, we make use of the following lemma, which is essentially a restatement of [Coh97, Propositions 7.1 and 7.2].

Lemma 4.7. *Suppose that we have k copies of the ℓ_0 -sketch data structure, for $k = 50 \lceil \log n \epsilon^{-2} \rceil$. Let u be any vertex such that $\deg(u) + 1 > 2\epsilon^{-1}$, and let $Q(u)$ denote the $\lfloor k(1 - 1/e) \rfloor$ -ranked key value in the list `minimizers[u]`. Then, with high probability, we have*

$$\frac{1 - \epsilon}{\deg(u) + 1} \leq Q(u) \leq \frac{1 + \epsilon}{\deg(u) + 1}.$$

In [Coh97] they assume that the random keys $R(v)$ are drawn from the exponential distribution (and hence the minimum key value is also), whereas we assume that $R(v)$ is drawn independently from the uniform distribution. When $\deg(u)$ is large enough though, the minimum of $\deg(u)$ random variables from either distribution is almost identically distributed. For completeness, we prove Lemma 4.7 when the keys $R(v)$ are drawn from the uniform distribution in Appendix A.

This idea leads to the following subroutine for providing implicit access to all vertices with approximately the same degree. This is critical for our nearly-linear time algorithm, and we explain its intricacies in Section 5. The pseudocode for this subroutine is given in Figure 6.

APPROXDEGREEDS_PIVOT(u)
Input: vertex to be pivoted, u .
Output: updated global state.

1. For each sketch $i = 1$ to k :
 - (a) $(v_1, v_2, \dots, v_\ell) \leftarrow \text{dynamic_sketch}[i].\text{PIVOTVERTEX}(u)$, the set of vertices in the i -th sketch whose minimizers changed after we pivot out u .
 - (b) For each $j = 1$ to ℓ :
 - i. Update the values corresponding to sketch i in $\text{minimizers}[v_j]$, which in turn updates its $\lfloor k(1 - 1/e) \rfloor$ -ranked quantile.
 - ii. Update the entry for v_j in quantile with the new value of the $\lfloor k(1 - 1/e) \rfloor$ -ranked quantile of $\text{minimizers}[v_j]$.

APPROXDEGREEDS_REPORT()
Output: approximate bucketing of the vertices by their fill 1-degrees.

1. For each $i = 0$ to $B = O(\log n\epsilon^{-1})$:
 - (a) Set V_i to be the split binary tree in quantile that contains all nodes with $\lfloor k(1 - 1/e) \rfloor$ -ranked quantiles in the range
$$\left[(1 + \epsilon)^{-(i+1)}, (1 + \epsilon)^{-i} \right].$$
2. Return (V_1, V_2, \dots, V_B) .

Figure 6: Pseudocode for the data structure that returns pointers to binary trees containing partitions of the remaining vertices into sets with ϵ -approximate degrees.

Observe that because 1-degrees are bounded by n , whenever we call **APPROXDEGREEDS_REPORT** we have $B = O(\log n\epsilon^{-1})$ with high probability by Lemma 4.7. Therefore, this data structure can simply return pointers to the first element in each of the partitions V_1, V_2, \dots, V_B .

Proof of Theorem 4.3. By construction, all vertices in V_i have their $\lfloor k(1 - 1/e) \rfloor$ -ranked quantile in the range

$$\left[(1 + \epsilon)^{-(i+1)}, (1 + \epsilon)^{-i} \right].$$

By Lemma 4.7, the 1-degree of any vertex in bucket V_i lies in the range

$$\left[(1 - \epsilon)(1 + \epsilon)^i, (1 + \epsilon)^{i+2} \right]$$

with high probability, which is within the claimed range for $\epsilon \leq 1/2$.

The proof of time and space complexities is similar to that of Theorem 4.1. Letting the number of sketches $k = O(\log n \epsilon^{-2})$ instead of $O(\Delta \log n)$ proves the space bound. One of the main differences in this data structure is that we need to store information about the $\lfloor k(1 - 1/e) \rfloor$ -ranked quantiles. These queries can be supported in $O(\log n)$ time by augmenting a balanced binary search tree with information about sizes of the subtrees in standard ways (e.g., [CLRS09, Chapter 14]). It follows that the total cost of all calls to APPROXDEGREE_DS_PIVOT is $O(m \log^3 n \epsilon^{-2})$. To analyze each call to APPROXDEGREE_DS_REPORT, we use standard splitting operations for binary search trees (e.g., treaps [SA96]), which allows us to construct each bucket in $O(\log n)$ time. \square

Note that there will be overlaps between the 1-degree intervals, so determining which bucket contains a given vertex is ambiguous if its order statistic is near the boundary of an interval.

An immediate corollary of Theorem 4.3 is that we can provide access to approximate min-degree vertices for a fixed sequence of updates by always returning an entry from the first nonempty bucket.

Corollary 4.8. *For a fixed elimination ordering (u_1, u_2, \dots, u_n) , we can find $(1 + \epsilon)$ -approximate minimum degree vertices in each of the intermediate states in $O(m \log^3 n \epsilon^{-2})$ time.*

It is also possible to adaptively choose the number of sketches for the $(1 + \epsilon)$ -approximate minimum degree algorithm by using a subroutine that is similar to the one in Section 4.2.

5 Generating Decorrelated Sequences

In this section we present a nearly-linear $(1 + \epsilon)$ -approximate marginal min-degree algorithm. This algorithm relies on degree approximation via sketching, as described in Theorem 4.3. In particular, it uses the randomized data structure APPROXDEGREE_DS, which provides access to buckets of vertices where the i -th bucket contains vertices with fill 1-degree in the range $[(1 + \epsilon)^{i-2}, (1 + \epsilon)^{i+2}]$.

Theorem 5.1. *There is an algorithm APPROXMINDEGREE_SEQUENCE that produces a $(1 + \epsilon)$ -approximate marginal min-degree ordering in expected $O(m \log^5 n \epsilon^{-2})$ time with high probability.*

At each step of this algorithm, reporting any member of the first nonempty bucket gives an approximate minimum degree vertex to pivot. However, such a choice must not have any dependence on the randomness used to get to this step, and more importantly, it should not affect pivoting decisions in future steps. To address this issue, we introduce an additional layer of randomization that decorrelates the ℓ_0 -sketches and the choice of vertices to pivot. Most of this section focuses on our technique for efficiently decorrelating such sequences.

The pseudocode for APPROXMINDEGREE_SEQUENCE is given in Figure 7. This algorithm makes use of the following global data structures and subroutines.

- APPROXDEGREE_DS: Returns buckets of vertices with approximately equal 1-degrees (Section 4.3).
- EXPDECAYEDCANDIDATES: Takes a sequence of values that are within $1 \pm \epsilon$ of each other, randomly perturbs the elements, and returns the new (ϵ -decayed) sequence (Section 5.2).
- ESTIMATEFILL1DEGREE: Gives an ϵ -approximation to the 1-degree of any vertex (Section 6).

We give the formal statement for ESTIMATEFILL1DEGREE in the following result.

Theorem 5.2. *There is a data structure that maintains a component graph G° under (adversarial) vertex pivots in a total of $O(m \log^2 n)$ time and supports the operation $\text{ESTIMATEFILL1DEGREE}(u, \epsilon)$, which given a vertex u and error threshold $\epsilon > 0$, returns with high probability a ϵ -approximation to the fill 1-degree of u by making $O(\deg(u) \log^2 n \epsilon^{-2})$ oracle queries to G° .*

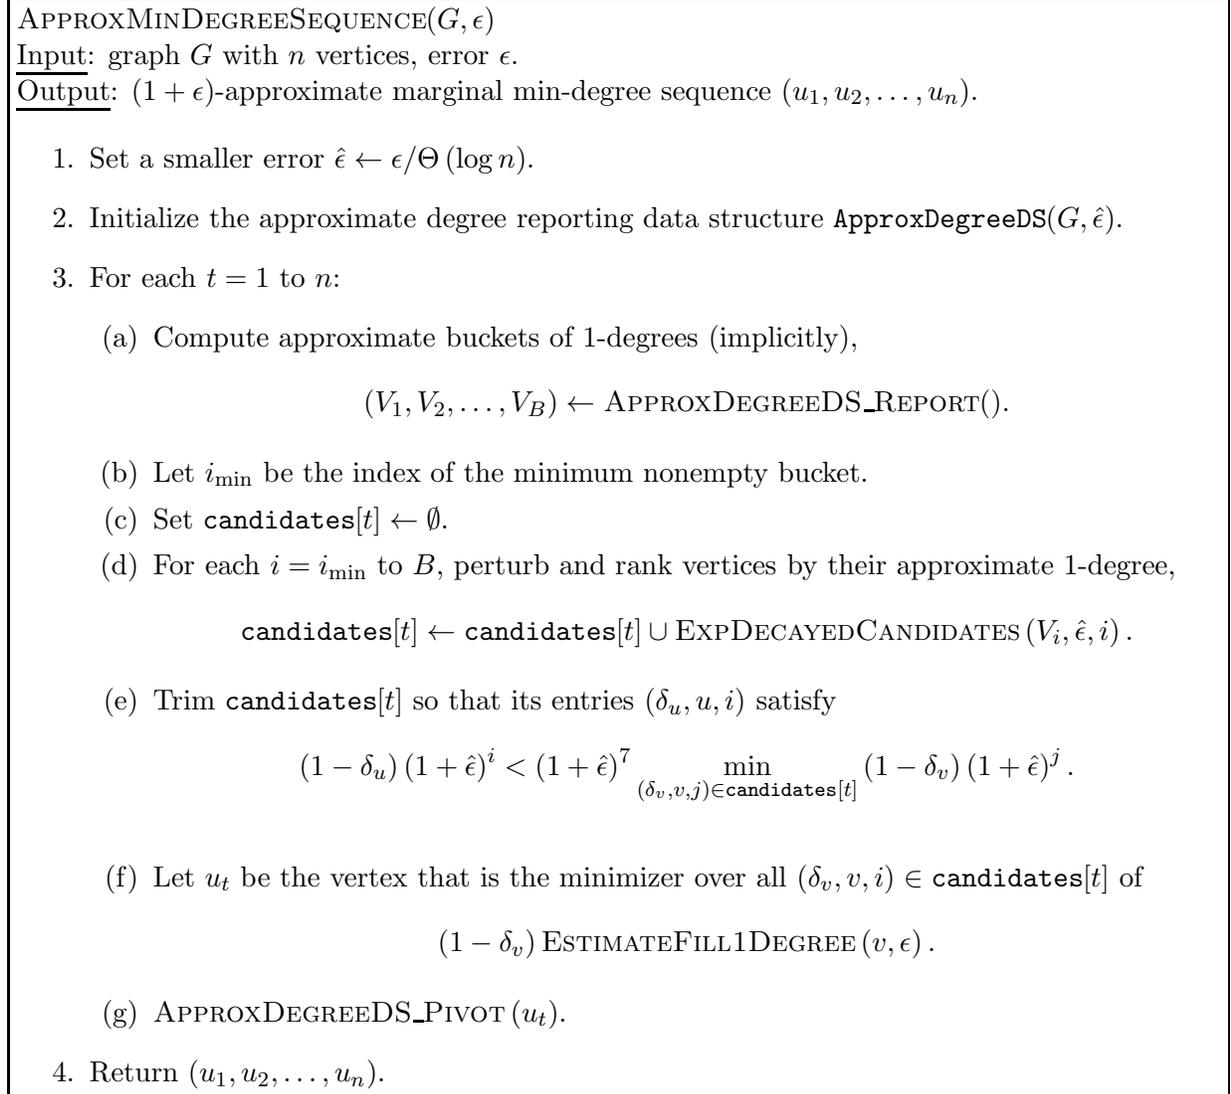


Figure 7: Pseudocode for the $(1 + \epsilon)$ -approximate marginal minimum degree ordering algorithm.

The most important part of this algorithm is arguably the use of exponential random variables to construct a list of candidates that is completely uncorrelated with the randomness used to generate the ℓ_0 -sketches and the choice of previous vertex pivots. The next subsection summarizes some desirable properties of exponential distributions that we exploit for efficient perturbations.

5.1 Exponential Random Variables

The exponential distribution is a continuous analog of the geometric distribution that describes the time between events in a Poisson point process. We utilize well-known facts about its order statistics, which have also appeared in the study of fault tolerance and distributed graph decompositions [MPX13]. For a rate parameter λ , the exponential distribution $\text{Exp}(\lambda)$ is defined by the probability density function (PDF)

$$f_{\text{Exp}(\lambda)}(x) = \begin{cases} \lambda \exp(-\lambda x) & \text{if } x \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

We will also make use of its cumulative density function (CDF)

$$F_{\text{Exp}(\lambda)}(x) = \begin{cases} 1 - \exp(-\lambda x) & \text{if } x \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

A crucial property of the exponential distribution is that it is memoryless. This means that for any rate $\lambda > 0$ and $s, t \geq 0$, an exponentially distributed random variable X satisfies the relation

$$\Pr[X > s + t \mid X > s] = \Pr[X > t].$$

A substantial portion of our analysis relies on the *order statistics* of exponential random variables. Given n random variables X_1, X_2, \dots, X_n , the i -th order statistic is the value of the i -th minimum random variable. A useful fact about i.i.d. exponential random variables is that the difference between consecutive order statistics also follows an exponential distribution. The algorithmic consequences of this property are that we can sample the smallest (or largest) k of n exponential random variables in increasing (or decreasing) order without ever generating all n random variables.

Lemma 5.3 ([Fel71]). *Let $X_{(i)}^n$ denote the i -th order statistic of n i.i.d. random variables drawn from the distribution $\text{Exp}(\lambda)$. Then, the n variables $X_{(1)}^n, X_{(2)}^n - X_{(1)}^n, \dots, X_{(n)}^n - X_{(n-1)}^n$ are independent, and the density of $X_{(k+1)}^n - X_{(k)}^n$ is given by the distribution $\text{Exp}((n-k)\lambda)$.*

One approach to prove Lemma 5.3 uses the i.i.d. assumption to show that the CDF of $X_{(1)}^n$ is

$$\begin{aligned} F_{X_{(1)}^n}(x) &= 1 - (1 - F_{\text{Exp}(\lambda)}(x))^n \\ &= 1 - \exp(-n\lambda x). \end{aligned}$$

This proves that $X_{(1)}^n$ follows an exponential distribution with rate $n\lambda$. Conditioning on $X_{(1)}^n$, we see that $X_{(2)}^n - X_{(1)}^n$ follows an exponential distribution equal to $X_{(1)}^{n-1}$ by the memoryless property. Therefore, one can repeat this argument to get the density of $X_{(k+1)}^n - X_{(k)}^n$ for all k up to $n-1$.

5.2 Implicitly Sampling ϵ -Decayed Minimums

The key idea in this section is the notion of ϵ -decay, which we use to slightly perturb approximate 1-degree sequences. It is motivated by the need to decorrelate the list of vertices grouped approximately by their 1-degree from previous sources of randomness in the algorithm. In the following definition, n is the number of vertices in the original graph before pivoting and $c_1 > 1$ is a constant.

Definition 5.4. Given a sequence $(x_1, x_2, \dots, x_k) \in \mathbb{R}^k$, we construct the corresponding ϵ -decayed sequence (y_1, y_2, \dots, y_k) by independently sampling the exponential random variables

$$\delta_i \sim \hat{\epsilon} \cdot \text{Exp}(1),$$

where $\hat{\epsilon} = \epsilon / (c_1 \log n)$ as in line 1 in APPROXMINDEGREESEQUENCE, and letting

$$y_i \leftarrow (1 - \delta_i) x_i.$$

We say that the ϵ -decayed minimum of (x_1, x_2, \dots, x_k) is the value $\min(y_1, y_2, \dots, y_k)$.

Definition 5.5. Given an error parameter $\epsilon > 0$ and an ϵ -approximate 1-degree estimation routine ESTIMATE1DEGREE(G, u), an ϵ -decayed minimum degree ordering is a sequence such that:

1. The vertex u_t corresponds to the ϵ -decayed minimum of ESTIMATE1DEGREE(G_{t-1}^+, v) over all remaining vertices $v \in V_{t-1}^+$.
2. The fill graph G_t^+ is obtained after eliminating u_t from G_{t-1}^+ .

Observe that the randomness of this perturbed degree estimator is regenerated at each step and thus removes any previous dependence. Next, we show that this adjustment is a well-behaved approximation, and then we show how to efficiently sample an ϵ -decayed minimum degree.

Lemma 5.6. Let Y be an ϵ -decayed minimum of (x_1, x_2, \dots, x_k) . With high probability, we have

$$Y \geq (1 - \epsilon) \min(x_1, x_2, \dots, x_k).$$

Proof. We bound the probability of the complementary event

$$\Pr[Y < (1 - \epsilon) \min(x_1, x_2, \dots, x_k)].$$

Observe that we can upper bound this probability by the probability that some x_i decreases to less than $1 - \epsilon$ times its original value. Recall that we set $\hat{\epsilon} = \epsilon / (c_1 \log n)$ for some constant $c_1 > 1$. Consider k i.i.d. exponential random variables $X_1, X_2, \dots, X_k \sim \text{Exp}(1)$, and let

$$\delta_i = \hat{\epsilon} \cdot X_i,$$

as in the definition of an ϵ -decayed minimum. Using the CDF of the exponential distribution, for each $i \in [k]$ we have

$$\begin{aligned} \Pr[\delta_i > \epsilon] &= \Pr\left[\frac{\epsilon}{c_1 \log n} \cdot X_i > \epsilon\right] \\ &= \Pr[X_i > c_1 \log n] \\ &= \exp(-c_1 \log n). \end{aligned}$$

It follows by a union bound that

$$\begin{aligned} \Pr\left[\max_{i \in [k]} \delta_i > \epsilon\right] &\leq \sum_{i=1}^k \Pr[\delta_i > \epsilon] \\ &= n^{1-c_1}, \end{aligned}$$

which completes the proof since $c_1 > 1$. □

By the previous lemma, to produce a $(1 + \epsilon)$ -approximate marginal minimum degree ordering, it suffices to compute an ϵ -decayed minimum degree ordering. Specifically, at each step we only need to find the ϵ -decayed minimum among the approximate fill 1-degrees of the remaining vertices. It turns out, however, that computing the approximate 1-degree for each remaining vertex in every iteration is expensive, so we avoid this problem by using `EXPDECAYEDCANDIDATES` on each bucket of vertices to carefully select a representative subset of candidates, and then we pivot out the minimizer over all buckets. The pseudocode for this subroutine is given in Figure 8, where we again let $\hat{\epsilon} = \epsilon / (c_1 \log n)$ for some constant $c_1 > 1$. Next, we show that this sampling technique is equivalent to finding the ϵ -decayed minimum over all remaining vertices with high probability.

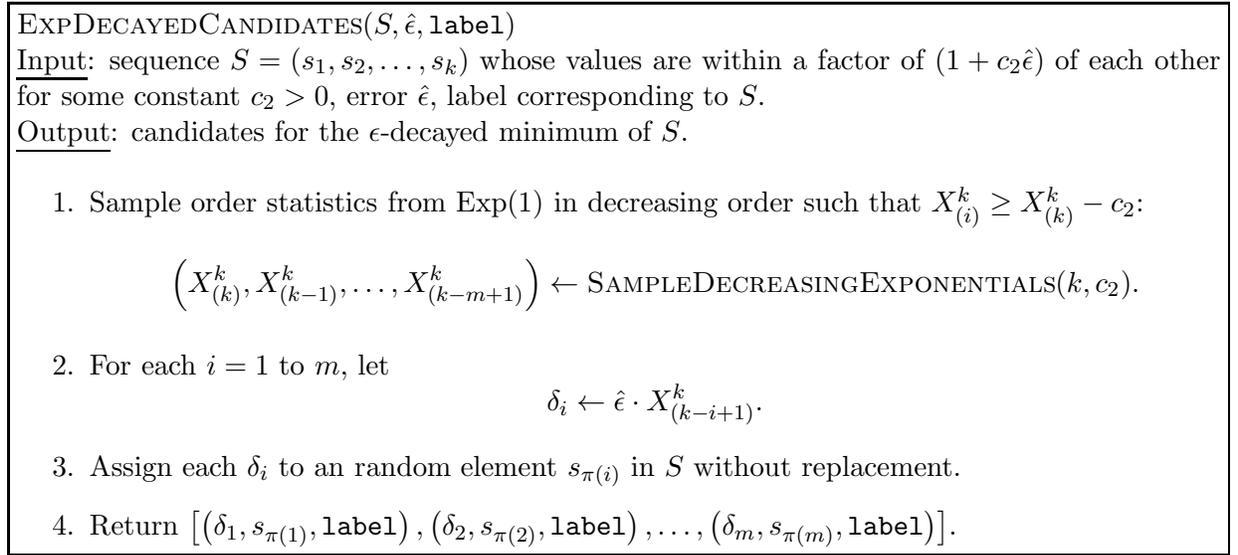


Figure 8: Pseudocode for generating an expected constant-size list of candidates for the ϵ -decayed minimum of a sequence of values that are within $(1 + c_2 \hat{\epsilon})$ of each other.

Note that the input sequence to `EXPDECAYEDCANDIDATES` requires that all its elements are within a factor of $(1 + c_2 \hat{\epsilon})$ of each other. We achieve this easily using the vertex buckets returned by `APPROXDEGREEDS_REPORT` in Section 4.3 when $\hat{\epsilon}$ is the error tolerance. The next lemma shows that the approximate vertex 1-degrees in any such bucket satisfy the required input condition.

Lemma 5.7. *For any bucket V_i of vertices returned by `APPROXDEGREEDS_REPORT`, there exists a constant $c_2 > 0$ such that all of the approximate 1-degrees are within a factor of $(1 + c_2 \hat{\epsilon})$ of each other. Alternatively, all of the approximate 1-degrees are within a factor of $(1 + \hat{\epsilon})^7$ of each other.*

Proof. The bucket V_i has vertices with approximate 1-degrees in the range

$$\left[(1 + \hat{\epsilon})^{i-2}, (1 + \hat{\epsilon})^{i+2} \right].$$

by Theorem 4.3. We have oracle access to the component graph G_t° by Theorem 5.2 and therefore can invoke `ESTIMATEFILL1DEGREE` on it. Instead of treating calls to `ESTIMATEFILL1DEGREE`

and the values of δ_i used to generate ϵ -decayed minimums as random variables, we view them as fixed values by removing the bad cases with high probability. That is, we define

$$\widetilde{\text{deg}}_t^+(u) + 1 \stackrel{\text{def}}{=} \text{ESTIMATEFILL1DEGREE}(u, \hat{\epsilon}).$$

Every call to ESTIMATEFILL1DEGREE is correct with high probability by Theorem 5.2, so we have

$$(1 - \hat{\epsilon}) (\text{deg}_t^+(u) + 1) \leq \widetilde{\text{deg}}_t^+(u) + 1 \leq (1 + \hat{\epsilon}) (\text{deg}_t^+(u) + 1).$$

This implies that all the approximate 1-degrees in bucket V_i are in the range

$$\left[(1 + \hat{\epsilon})^{i-4}, (1 + \hat{\epsilon})^{i+3} \right],$$

for $\hat{\epsilon}$ sufficiently small. Therefore, all of these values are within a factor

$$\frac{(1 + \hat{\epsilon})^{i+3}}{(1 + \hat{\epsilon})^{i-4}} = (1 + \hat{\epsilon})^7 \leq (1 + c_2 \hat{\epsilon})$$

of each other for some $c_2 > 0$, which completes the proof. \square

The most important part of EXPDECAYEDCANDIDATES is generating the order statistics efficiently. In Figure 9 we show how to iteratively sample the variables $X_{(i)}^k$ in decreasing order using Lemma 5.3. This technique is critically important for us because we only consider order statistics satisfying the condition $X_{(k)}^k - c_2$, which is at most a constant number of variables in expectation.

SAMPLEDECREASINGEXPONENTIALS(k, c_2)
Input: integer $k \geq 0$, real-valued threshold $c_2 > 0$.
Output: order statistics $X_{(k)}^k, X_{(k-1)}^k, \dots, X_{(k-m+1)}^k$ from Exp(1) such that $X_{(k-m+1)}^k \geq X_{(k)}^k - c_2$ and $X_{(k-m)}^k < X_{(k)}^k - c_2$.

1. Sample $X_{(k)}^k$ using its CDF
$$\Pr \left[X_{(k)}^k \leq x \right] = (1 - e^{-x})^k.$$
2. For each $i = 1$ to $k - 1$:
 - (a) Sample the difference $Y \sim \text{Exp}(i)$ and let
$$X_{(k-i)}^k \leftarrow X_{(k-i+1)}^k - Y.$$
 - (b) If $X_{(k-i)}^k < X_{(k)}^k - c_2$, let $m \leftarrow i$ and exit the loop.
3. Return $\left(X_{(k)}^k, X_{(k-1)}^k, \dots, X_{(k-m+1)}^k \right)$.

Figure 9: Pseudocode for iteratively generating order statistics of exponential random variables in decreasing order within a threshold c_2 of the maximum value $X_{(k)}^k$.

To show that our algorithm is correct, we must prove that (1) the algorithm selects a bounded number of candidates in expectation at each step, and (2) the true ϵ -decayed minimum belongs to the candidate list. We analyze both of these conditions in the following lemma.

Lemma 5.8. *If x_1, x_2, \dots, x_k are within a factor of $(1 + c_2\hat{\epsilon})$ of each other, then the ϵ -decayed minimum is among the candidates returned by $\text{EXPDECAYEDCANDIDATES}((x_1, x_2, \dots, x_k), \hat{\epsilon}, \cdot)$. Furthermore, the expected number of candidates returned is bounded by the constant e^{c_2} .*

Proof. Let $X_1, X_2, \dots, X_k \sim \text{Exp}(1)$ be i.i.d. and let the order statistic $X_{(k)}^k = \max\{X_1, X_2, \dots, X_k\}$. We first verify the correctness of $\text{SAMPLEDECREASINGEXPONENTIALS}$. Using the CDF of the exponential distribution, it follows that the CDF of $X_{(k)}^k$ is

$$\begin{aligned} F_{X_{(k)}^k}(x) &= \Pr[\max\{X_1, X_2, \dots, X_n\} \leq x] \\ &= \prod_{i=1}^k \Pr[X_i \leq x] \\ &= (1 - e^{-x})^k. \end{aligned}$$

The memoryless property implies that we can generate $X_{(k-1)}^k, X_{(k-2)}^k, \dots, X_{(1)}^k$ iteratively by sampling their differences from an exponential distribution whose rate is given by Lemma 5.3.

Now we show that the ϵ -decayed minimum is among the candidates. We claim that it suffices to sample every $X_{(i)}^k$ such that $X_{(i)}^k \geq X_{(k)}^k - c_2$, or equivalently every δ_i such that $\delta_i \geq \delta_k - c_2\hat{\epsilon}$. To see this, suppose for contradiction that the ϵ -decayed minimum $x_{\pi(j)}$ is not included. Then we have $\delta_j < \delta_k - c_2\hat{\epsilon}$, so it follows that

$$(1 - \delta_j) x_{\pi(j)} > (1 - \delta_k + c_2\hat{\epsilon}) \frac{x_{\pi(k)}}{1 + c_2\hat{\epsilon}} \geq (1 - \delta_k) x_{\pi(k)},$$

which is a contradiction. Therefore, our candidate list contains the ϵ -decayed minimum.

Lastly, to count the expected number of candidates, we count the number of values δ_j generated. Let Z_i be the indicator variable for the event $X_i \geq X_{(k)}^k - c_2$. Then $Z = \sum_{i=1}^k Z_i$ indicates the size of the candidate list. Using the memoryless property of exponential random variables, we have

$$\begin{aligned} \mathbb{E}[Z] &= \sum_{i=1}^k \mathbb{E}[Z_i] \\ &= \sum_{i=1}^k \Pr[X_{(i)}^k \geq X_{(k)}^k - c_2] \\ &= 1 + \sum_{i=1}^{k-1} \Pr[X_{(i)}^k \geq X_{(k)}^k - c_2] \\ &= 1 + \sum_{i=1}^{k-1} \Pr[X_{(i)}^i \leq c_2] \\ &= 1 + \sum_{i=1}^{k-1} (1 - e^{-c_2})^i \\ &\leq e^{c_2}, \end{aligned}$$

where the final equality considers the geometric series. Therefore, at most a constant number of exponential random variables are generated as we sample backwards from the maximum. \square

We cannot simply work with the first nonempty bucket because the randomness introduces a $1 \pm \epsilon$ perturbation. Furthermore, the bucket containing the vertex with minimum degree is dependent on the randomness of the sketches (as discussed in Theorem 4.3). To bypass this problem we inject additional, uncorrelated randomness into the algorithm at each step to find $O(1)$ candidates for each of the $O(\log n \hat{\epsilon}^{-1})$ buckets, which increases the number of global candidates to $O(\log n \hat{\epsilon}^{-1})$. Then in the penultimate step of each iteration, before we compute the approximate 1-degrees of candidate vertices (which is somewhat expensive), we carefully filter the global list so that the global ϵ -decayed minimum remains in the list with high probability.

Lemma 5.9. *Let (δ_u, u, i) be the entry over all $(\delta_v, v, j) \in \text{candidates}[t]$ that minimizes*

$$(1 - \delta_v) \text{ESTIMATEFILL1DEGREE}(v, \epsilon).$$

Then, with high probability, we have

$$(1 - \delta_u) (1 + \hat{\epsilon})^i \leq (1 + \hat{\epsilon})^7 \min_{(\delta_v, v, j) \in \text{candidates}[t]} (1 - \delta_v) (1 + \hat{\epsilon})^j.$$

Proof. Let (δ_v, v, j) be an arbitrary entry in $\text{candidates}[t]$. By assumption, we have

$$(1 - \delta_u) \text{ESTIMATEFILL1DEGREE}(u, \epsilon) \leq (1 - \delta_v) \text{ESTIMATEFILL1DEGREE}(v, \epsilon).$$

Using inequalities in Lemma 5.7, it follows that

$$\text{ESTIMATEFILL1DEGREE}(u, \epsilon) \geq (1 + \hat{\epsilon})^{i-4}$$

and

$$\text{ESTIMATEFILL1DEGREE}(v, \epsilon) \leq (1 + \hat{\epsilon})^{j+3}$$

with high probability. Substituting these into the previous inequality gives us the result. \square

5.3 Analysis of the Approximation Algorithm

Now that we have all of the building blocks from the previous subsection, we prove the correctness of the $(1 + \epsilon)$ -approximate marginal minimum degree algorithm and bound its running time.

Lemma 5.10. *For any graph G and any error ϵ , the output of $\text{APPROXMINDEGREESEQUENCE}(G, \epsilon)$ is a $(1 + \epsilon)$ -approximate marginal minimum degree sequence with high probability.*

Proof. We prove by induction that for some constant $c > 1$, after t steps of the algorithm, the output is a $(1 + \epsilon)$ -approximate marginal min-degree ordering with probability at least $1 - tn^{-c}$. The base case when $t = 0$ follows trivially because nothing has happened. For the inductive hypothesis, assume that after t steps the sequence (u_1, u_2, \dots, u_t) is a $(1 + \epsilon)$ -approximate min-degree ordering and let the graph state be G_t^+ , where the eliminated vertices u_1, u_2, \dots, u_t have been pivoted.

By Lemma 5.7, all values in a given bucket are within a factor of $1 + c_2 \hat{\epsilon}$ of each other. We use the guarantees of Lemma 5.8 to compute the ϵ -decayed minimum candidate of each bucket. It follows from Lemma 5.9 that after we trim the candidate list, one of the remaining candidates is

the original minimizer of $(1 - \delta_v)\text{ESTIMATEFILL1DEGREE}(v, \epsilon)$ with high probability. Therefore, u_t is the ϵ -decayed minimum over all values of $\text{ESTIMATEFILL1DEGREE}(v, \epsilon)$ with high probability. Lastly, invoking the bound on distortions incurred by ϵ -decay in Lemma 5.6 and accounting for the error of $\text{ESTIMATEFILL1DEGREE}$, the 1-degree of u_t is within $1 + \epsilon$ of the minimum 1-degree in the fill graph G_t^+ with high probability. Taking a union bound over all the high probability claims, we have a failure probability of at most n^{-c} . Thus, the inductive hypothesis also holds for $t + 1$. \square

We now analyze the cost of the algorithm. To do this, we first show that if a vertex is close to the global ϵ -decayed minimum, then there is a reasonable chance that it actually is the minimizer. In other words, if the algorithm queries the approximate degree of a vertex, then it is likely that this vertex belongs to the ϵ -decayed approximate degree sequence. This explains the trimming condition in $\text{APPROXMINDEGREESEQUENCE}$.

Lemma 5.11. *For any constant $c_3 \geq 1$, choice of error $\hat{\epsilon}$, sequence of values (x_1, x_2, \dots, x_k) , and index $i \in [k]$, we have*

$$\begin{aligned} \Pr [i \text{ corresponds to the } \epsilon\text{-decayed minimum of } (x_1, x_2, \dots, x_k)] \\ \geq \exp(-2c_3) \Pr \left[(1 - \delta_i) x_i < (1 + \hat{\epsilon})^{c_3} \min_{j \in [k]} (1 - \delta_j) x_j \right]. \end{aligned}$$

Proof. Without loss of generality, suppose we generate δ_i last. Let

$$m = \min_{j \in [k] \setminus \{i\}} (1 - \delta_j) x_j$$

be the previous ϵ -decayed minimum. If $m \geq x_i$ then both probabilities in the claim are equal to 1 and the result holds trivially. Otherwise, consider the probability that i corresponds to the minimizer conditioned on the event

$$(1 - \delta_i) x_i < (1 + \hat{\epsilon})^{c_3} m.$$

Equivalently, assume that $\delta_i > \gamma$, for some γ such that $(1 - \gamma)x_i = (1 + \hat{\epsilon})^{c_3} m$. By the memoryless property of the exponential distribution, we have

$$\begin{aligned} \Pr [\delta_i > \gamma + 2c_3\hat{\epsilon} \mid \delta_i > \gamma] &= \Pr [\delta_i > 2c_3\hat{\epsilon}] \\ &= \exp(-2c_3). \end{aligned}$$

Therefore, with probability at least $\exp(-2c_3)$, it follows that

$$\begin{aligned} (1 - \delta_i) x_i &< (1 - \gamma - 2c_3\hat{\epsilon}) x_i \\ &\leq (1 - 2c_3\hat{\epsilon}) (1 - \gamma) x_i \\ &= (1 - 2c_3\hat{\epsilon}) (1 + \hat{\epsilon})^{c_3} m \\ &\leq m. \end{aligned}$$

This means that if the decayed value of x_i is within a threshold of the previous minimum m , then x_i itself will decay below m with at least constant probability. \square

Making the substitution $c_3 = 7$, as in line 3e of $\text{APPROXMINDEGREESEQUENCE}$ in Figure 7, gives the following corollary, which allows us to prove our main result.

Corollary 5.12. *If a vertex v is in $\mathbf{candidates}[t]$ after line 3e of APPROXMINDEGREESEQUENCE, then with probability at least $\exp(-14)$, v is the ϵ -decayed minimum.*

Proof of Theorem 5.1. The correctness follows from Lemma 5.10. We can maintain access to all of the buckets across the sequence of pivots in a total time of

$$O(m \log^3 n \hat{\epsilon}^{-2}) = O(m \log^5 n \epsilon^{-2})$$

by Theorem 4.3, so all that remains is bounding the total cost of the calls to ESTIMATEFILL1DEGREE.

By Theorem 5.2, the cost of maintaining the component graph under pivots is $O(m \log^2 n)$, a lower order term. To analyze the aggregate cost of the calls to ESTIMATEFILL1DEGREE, we utilize Corollary 5.12, which states that any vertex in $\mathbf{candidates}[t]$ (after trimming) is the one we pivot with constant probability. Specifically, we prove by induction on the number of remaining vertices that for some constant c_4 the expected cost of calling ESTIMATEFILL1DEGREE is bounded by

$$c_4 \left(\sum_{u \in V_t^+} \deg_{\text{rem},t}^\circ(u) \right) \log^2 n \hat{\epsilon}^{-2}.$$

The base case $t = n$ follows trivially since no vertices remain. Now suppose that the claim is true for $t + 1$ vertices. By the induction hypothesis, the expected cost of the future steps is bounded by

$$\begin{aligned} & \sum_{u \in V_t^+} \Pr[u \text{ is the } \epsilon\text{-decayed minimum}] \cdot c_4 \left(-\deg_{\text{rem},t}^\circ(u) + \sum_{v \in V_t^+} \deg_{\text{rem},t}^\circ(v) \right) \log^2 n \hat{\epsilon}^{-2} \\ &= c_4 \left(\sum_{u \in V_t^+} \deg_{\text{rem},t}^\circ(u) \right) \log^2 n \hat{\epsilon}^{-2} - c_4 \sum_{u \in V_t^+} \Pr[u \text{ is the } \epsilon\text{-decayed minimum}] \cdot \deg_{\text{rem},t}^\circ(u) \log^2 n \hat{\epsilon}^{-2}. \end{aligned}$$

Now we consider the cost of evaluating ESTIMATEFILL1DEGREE($u, \hat{\epsilon}$) at time t if $u \in \mathbf{candidates}[t]$. By Corollary 5.12 and expanding the conditional probability, we have

$$\Pr[u \text{ is the } \epsilon\text{-decayed minimum}] \geq \exp(-14) \Pr[u \in \mathbf{candidates}[t]].$$

Therefore, using Theorem 5.2, the expected cost of these calls is

$$\begin{aligned} & \sum_{u \in V_t^+} \Pr[u \in \mathbf{candidates}[t]] \cdot c_3 \deg_{\text{rem},t}^\circ(u) \log^2 n \hat{\epsilon}^{-2} \\ & \leq c_3 \exp(14) \sum_{u \in V_t^+} \Pr[u \text{ is the } \epsilon\text{-decayed minimum}] \cdot \deg_{\text{rem},t}^\circ(u) \log^2 n \hat{\epsilon}^{-2}. \end{aligned}$$

It follows that the inductive hypothesis holds for all t by letting $c_4 = c_3 \exp(14)$. The initial sum of the remaining degrees is $O(m)$, so the total expected cost of calling ESTIMATEFILL1DEGREE is

$$O(m \log^2 n \hat{\epsilon}^{-2}) = O(m \log^4 n \epsilon^{-2}),$$

which completes the proof. \square

6 Estimating the Fill 1-Degree of a Vertex

This section discusses routines for approximating the fill 1-degree of a vertex in a partially eliminated graph. We also show how to maintain the partially eliminated graph throughout the course of the algorithm, which allows us to prove Theorem 5.2. The partially eliminated graph we use for degree estimation is the component graph G° , where connected components of the eliminated vertices are contracted into single vertices called *component* vertices. See Section 2.1 for a detailed explanation.

Our goal is to efficiently approximate the fill 1-degree of a given remaining vertex u . By the definition of fill 1-degree and the neighborhoods of component graphs, it follows that

$$\deg^+(u) + 1 = \left| \{u\} \cup N_{\text{rem}}^\circ(u) \cup \bigcup_{x \in N_{\text{comp}}^\circ(u)} N_{\text{rem}}^\circ(x) \right|.$$

In other words, the fill 1-neighborhood of u is set of remaining 1-neighbors of u in the original graph in addition to the remaining neighbors of each component neighbor of u .

This union-of-sets structure has a natural $(0,1)$ -matrix interpretation, where columns correspond to remaining vertices and rows correspond to neighboring component neighborhoods of u (along with an additional row for the 1-neighborhood of u). For each row i , set the entry $A(i, j) = 1$ if vertex j is in the i -th neighborhood set and let $A(i, j) = 0$ otherwise. The problem can then be viewed as querying for the number of nonzero columns of A . Specifically, we show how one can accurately estimate fill 1-degrees using the following matrix queries:

- **ROWSIZE**(A, i): Return the number of nonzero elements in row i of A .
- **SAMPLEFROMROW**(A, i): Returns a column index j uniformly at random from the nonzero entries of row i of A .
- **QUERYVALUE**(A, i, j): Returns the value of $A(i, j)$.

The main result in this section is the follow matrix sampler.

Lemma 6.1. *There is a routine ESTIMATENONZEROCOLUMNS using the three operations above that takes as input (implicit) access to a matrix A and an error ϵ , and returns an ϵ -approximation to the number of nonzero columns in A with high probability. The expected total number of operations used is $O(r \log^2 n \epsilon^{-2})$, where r is the number of rows and n is the number of columns in A .*

Before analyzing this matrix-based estimator, we verify that Lemma 6.1 can be used in the graph-theoretic setting to prove Theorem 5.2. We use the following tools for querying degrees and sampling neighbors in a component graph as it undergoes pivots.

Lemma 6.2. *We can maintain a component graph under vertex pivots in a total time of $O(m \log^2 n)$. Additionally, this component graph data structure grants $O(\log n)$ time oracle access for:*

- *Querying the state of a vertex.*
- *Querying the component or remaining neighborhood (and hence degree) of a vertex.*
- *Uniformly sampling a remaining neighbor of a component or remaining vertex.*
- *Uniformly sampling a random component vertex.*

We defer the proof of Lemma 6.2 to Section 7.

Assuming the correctness of Lemma 6.1 and Lemma 6.2, we can easily prove Theorem 5.2, which allows us to efficiently estimate the fill 1-degrees of vertices throughout the algorithm.

Proof of Theorem 5.2. We can implicitly construct A and simulate the matrix operations as follows. Storing the adjacency list of the component graph using binary search trees clearly implies that we can implement ROWSIZE and SAMPLEFROMROW. Moreover, in this setting QUERYVALUE corresponds to querying connectivity, which again is simple because we use binary search trees. Substituting in the runtime bounds gives the desired result. \square

The rest of the section is outlined as followed. We prove a weaker but relevant version of the matrix estimator (Lemma 6.1) in Section 6.1. This algorithm relies on a subroutine to estimate the mean of a distribution, which we discuss in Section 6.2. Then by more carefully analyzing the previous two algorithms, we prove the original estimation result in Section 6.3.

6.1 Approximating the Number of Nonzero Columns using Mean Estimators

We begin by defining an estimator for counting the number of nonzero columns of A . Let

$$\text{COLUMNSUM}(A, j) \stackrel{\text{def}}{=} \sum_{i=1}^r A(i, j),$$

and consider the normalized matrix B such that

$$B(i, j) = \begin{cases} 0 & \text{if } A(i, j) = 0 \\ \text{COLUMNSUM}(A, j)^{-1} & \text{if } A(i, j) = 1. \end{cases}$$

If we know the number of nonzeros in B and can uniformly sample nonzero entries of B , then we can use this distribution on B as an unbiased estimator for the number of nonzero columns in A . We explicitly capture this idea with the following lemma.

Lemma 6.3. *If X is a uniformly random nonzero entry of B , then*

$$\mathbb{E}[X] = \frac{\text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)}.$$

Assuming that we can uniformly sample indices (i, j) of nonzero entries of B , Lemma 6.3 implies that it is sufficient to estimate column sums. We show how to do this by estimating the mean of an appropriately chosen Bernoulli distribution on the column. All of our estimators in this section use a general-purpose ESTIMATEMEAN algorithm (given in Figure 10), for any distribution over $[0, 1]$. We present its accuracy guarantees and sample complexity next, and defer the proof to Section 6.4.

<p>ESTIMATEMEAN(D, σ)</p> <p><u>Input</u>: access to a distribution D over $[0, 1]$, cutoff threshold $\sigma > 0$.</p> <p><u>Output</u>: estimation of the mean of D.</p> <ol style="list-style-type: none"> 1. Initialize <code>counter</code> $\leftarrow 0$ and <code>sum</code> $\leftarrow 0$. 2. While <code>sum</code> $< \sigma$: <ol style="list-style-type: none"> (a) Generate $X \sim D$. (b) <code>sum</code> \leftarrow <code>sum</code> $+ X$. (c) <code>counter</code> \leftarrow <code>counter</code> $+ 1$. 3. Return $\sigma/\text{counter}$.

Figure 10: Pseudocode for an algorithm that estimates the mean of a distribution D on $[0, 1]$.

Lemma 6.4. *Let D be any distribution over $[0, 1]$ with (an unknown) mean μ . For any cutoff $\sigma > 0$ and error $\epsilon > 0$, with probability at least $1 - \exp(-\epsilon^2\sigma/5)$, the algorithm ESTIMATEMEAN(D, σ):*

- Generates $O(\sigma/\mu)$ samples from the distribution.
- Produces an estimated mean $\bar{\mu}$ such that $(1 - \epsilon)\mu \leq \bar{\mu} \leq (1 + \epsilon)\mu$.

An immediate corollary of Lemma 6.4 is a routine APPROXCOLUMNSUM for estimating column sums of A , where the running time depends on the column sum itself. We give the pseudocode for this estimator in Figure 11, and then we prove its correctness and running time in Lemma 6.5.

<p>APPROXCOLUMNSUM(A, j, ϵ, δ)</p> <p><u>Input</u>: matrix A with r rows, column index j, error $\epsilon > 0$, failure probability $\delta > 0$.</p> <p>Implicit access to the number of rows r and the number of remaining vertices n.</p> <p><u>Output</u>: estimation for COLUMNSUM(A, j).</p> <ol style="list-style-type: none"> 1. Let $D_{\text{col}}(j)$ denote the distribution for the random variable that: <ol style="list-style-type: none"> (a) Chooses a row $i \in [r]$ uniformly at random. (b) Returns the value of $A(i, j)$. 2. Set $\sigma \leftarrow 5\epsilon^{-2} \log(1/\delta)$. 3. Return $r \cdot \text{ESTIMATEMEAN}(D_{\text{col}}(j), \sigma)$.
--

Figure 11: Pseudocode for approximating the column sum of a matrix.

Lemma 6.5. For any $(0, 1)$ -matrix $A \in \mathbb{R}^{r \times n}$, column index j , error $\epsilon > 0$, and failure rate $\delta > 0$, invoking APPROXCOLUMNSSUM returns an ϵ -approximation to $\text{COLUMNSSUM}(A, j)$ with probability at least $1 - \delta$ while making

$$O\left(\frac{r \log(1/\delta)}{\text{COLUMNSSUM}(A, j)\epsilon^2}\right)$$

oracle calls to the matrix A in expectation.

Proof. Observe that $D_{\text{col}}(j)$ is a Bernoulli distribution with mean $\text{COLUMNSSUM}(A, j)/r$. The success probability follows directly from our choice of σ and the success probability of ESTIMATEMEAN for Bernoulli distributions (Lemma 6.4). Moreover, the total number of matrix-entry queries is

$$O\left(\frac{\sigma}{\text{COLUMNSSUM}(A, j)/r}\right) = O\left(\frac{r \log(1/\delta)}{\text{COLUMNSSUM}(A, j)\epsilon^2}\right),$$

as desired. □

We can simulate sampling from the nonzero entries of the reweighted matrix B by using the inverse of APPROXCOLUMNSSUM to estimate $B(i, j)$. Moreover, by sampling enough entries of B so that their sum is $O(\log n \epsilon^{-2})$, we can accurately estimate $\text{NONZEROCOLUMNS}(A)$ by Lemma 6.3. We give the slow version of this estimator in Figure 12, and we prove its running time in Lemma 6.1.

ESTIMATENONZEROCOLUMNS_SLOW(A, ϵ)

Input: oracle access to the matrix A with r rows and n columns, error threshold $\epsilon > 0$.

Output: estimation for the number of nonzero columns in A .

1. Compute $\text{nnz}(A)$, the total number of nonzeros in A .
2. Let D_{global} denote the distribution for the random variable that:
 - (a) Chooses a uniformly random nonzero index (i, j) in A (by first picking a row with probability proportional to its number of nonzeros and then picking a random nonzero entry from that row).
 - (b) Returns

$$\frac{1}{\text{APPROXCOLUMNSSUM}(A, j, \epsilon, n^{-10})},$$
 where the value of $\text{APPROXCOLUMNSSUM}(A, j, \epsilon, n^{-10})$ is generated once per each column and reused on subsequent calls (via storage in a binary search tree).
3. Set $\sigma \leftarrow 50\epsilon^{-2} \log(n)$
4. Return $\text{nnz}(A) \cdot \text{ESTIMATEMEAN}(D_{\text{global}}, \sigma)$.

Figure 12: Pseudocode for (slowly) estimating the number of nonzero columns of a $(0, 1)$ -matrix.

Next, we prove the correctness of ESTIMATENONZEROCOLUMNS_SLOW, and then we bound the expected number of times it generates samples from $D_{\text{global}}(j)$.

Lemma 6.6. *With high probability, the estimation $\text{ESTIMATENONZEROCOLUMNS_SLOW}(A, \epsilon)$ is within a factor of $1 \pm \epsilon$ of the number of nonzero columns of A .*

Proof. To start, we extract all of the randomness out of $\text{ESTIMATENONZEROCOLUMNS_SLOW}$ by considering running all calls to $\text{APPROXCOLUMNSUM}(A, j, \epsilon, n^{-10})$ beforehand. By Lemma 6.5, with high probability, for each column j we have

$$(1 - \epsilon) \text{COLUMNSUM}(A, j) \leq \text{APPROXCOLUMNSUM}(A, j, \epsilon, n^{-10}) \leq (1 + \epsilon) \text{COLUMNSUM}(A, j).$$

Therefore, by Lemma 6.3 it follows that

$$(1 - 2\epsilon) \frac{\text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)} \leq \mu(D_{\text{global}}) \leq (1 + 2\epsilon) \frac{\text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)},$$

for sufficiently small ϵ . Incorporating the accuracy guarantee from Lemma 6.4 gives

$$(1 - 4\epsilon) \frac{\text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)} \leq \text{ESTIMATEMEAN}(D_{\text{global}}, \sigma) \leq (1 + 4\epsilon) \frac{\text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)}.$$

The desired bound follows by quartering ϵ and multiplying by $\text{nnz}(A)$. \square

Initial Proof of Lemma 6.1 (using $O(r \log^2 n \epsilon^{-4})$ operations). The correctness is a consequence of Lemma 6.6, so we must bound the total number of queries to entries of A . Using Lemma 6.4 and Lemma 6.3, the expected number of queries made to D_{global} is

$$O\left(\frac{\text{nnz}(A) \log n}{\text{NONZEROCOLUMNS}(A) \epsilon^2}\right).$$

Therefore, it suffices to bound the expected cost of each sample generated from D_{global} .

Applying Lemma 6.5 to each column j , the expected number of queries to A made by $\text{COLUMNSUM}(A, j)$ is

$$O\left(\frac{r \log n}{\text{COLUMNSUM}(A, j) \epsilon^2}\right).$$

Summing over all $\text{COLUMNSUM}(A, j)$ nonzero entries in column j gives

$$O\left(\frac{r \log n}{\epsilon^2}\right)$$

queries to A per nonzero column. It follows that the expected number of queries to A per sample generated from D_{global} is

$$O\left(\frac{\text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)} \cdot \frac{r \log n}{\epsilon^2}\right).$$

Multiplying this by the expected number of queries to D_{global} gives the overall result. \square

We note that this sample complexity bound also holds with high probability (instead of only in expectation) by invoking Chernoff bounds. This is because the cost of each query to D_{global} is bounded by $O(r \log n \epsilon^{-2})$ and the overall cost bound is larger by a factor of at least $\Omega(\log n)$.

6.2 Estimating the Mean of a Distribution

We now provide the details for the mean estimation algorithm, which proves the correctness of the column sum estimator. We analyze the following scheme:

1. Generate an infinite stream of i.i.d. samples X_1, X_2, \dots from any distribution D over $[0, 1]$.
2. Let $\mathbf{counter} = \min \{t \geq 0 : \sum_{i=1}^t X_i \geq \sigma\}$.
3. Output $\sigma/\mathbf{counter}$.

This process generates more samples than ESTIMATEMEAN in Figure 10, but the extra evaluations happen after the subroutine terminates and thus does not affect the outcome.

Let μ be the (hidden) mean of the distribution D . For any error $\epsilon > 0$, define the two cutoffs

$$L(D, \epsilon) \stackrel{\text{def}}{=} \frac{\sigma}{(1 + \epsilon)\mu}$$

and

$$R(D, \epsilon) \stackrel{\text{def}}{=} \frac{\sigma}{(1 - \epsilon)\mu}.$$

For convenience we write $L = L(D, \epsilon)$ and $R = R(D, \epsilon)$. We claim that if $L \leq \mathbf{counter} \leq R$, then the estimation we output will be sufficiently accurate. Therefore, we first bound the probabilities of the complementary events $\mathbf{counter} < L$ and $\mathbf{counter} > R$.

Lemma 6.7. *Let D be any distribution over $[0, 1]$. For any sequence X_1, X_2, X_3, \dots of i.i.d. random variables generated from D and any choice of $\sigma \geq 0$, we have*

$$\Pr \left[\sum_{i=1}^L X_i \geq \sigma \right] \leq \exp \left(-\frac{\epsilon^2 \sigma}{4} \right).$$

Proof. By the linearity of expectation we have

$$\mathbb{E} \left[\sum_{i=1}^L X_i \right] = L\mu.$$

Since X_1, X_2, \dots, X_L are independent random variables with values in $[0, 1]$, Chernoff bounds give

$$\begin{aligned} \Pr \left[\sum_{i=1}^L X_i \geq (1 + \epsilon) L\mu \right] &\leq \exp \left(-\frac{\epsilon^2 L\mu}{3} \right) \\ &\leq \exp \left(-\frac{\epsilon^2 \sigma}{4} \right) \end{aligned}$$

by letting $\sigma = (1 + \epsilon)L\mu$ and considering ϵ sufficiently small. This completes the proof. \square

Lemma 6.8. *Let D be any distribution over $[0, 1]$. For any sequence X_1, X_2, X_3, \dots of i.i.d. random variables generated from D and any choice of $\sigma \geq 0$, we have*

$$\Pr \left[\sum_{i=1}^R X_i \leq \sigma \right] \leq \exp \left(-\frac{\epsilon^2 \sigma}{4} \right).$$

Proof. Consider the proof of Lemma 6.7, and use a Chernoff bound for the lower tail instead. \square

Proof of Lemma 6.4. We will show that the estimator behaves as intended when $L \leq \text{counter} \leq R$. Considering the complementary events, it is easy to see that

$$\Pr[\text{counter} \leq L] \leq \Pr\left[\sum_{i=1}^L X_i \geq \sigma\right]$$

and

$$\Pr[\text{counter} \geq R] \leq \Pr\left[\sum_{i=1}^R X_i \leq \sigma\right].$$

Therefore, by Lemma 6.7 and Lemma 6.8 it follows that

$$\begin{aligned} \Pr[L \leq \text{counter} \leq R] &\geq 1 - 2 \exp\left(-\frac{\epsilon^2 \sigma}{4}\right) \\ &\geq 1 - \exp\left(-\frac{\epsilon^2 \sigma}{5}\right). \end{aligned}$$

Assume that $L \leq \text{counter} \leq R$ and recall the definitions of L and R . It follows that the number of samples generated is $O(\sigma/\mu)$. To prove that $\sigma/\text{counter}$ is an accurate estimate, observe that

$$\frac{\sigma}{(1+\epsilon)\mu} \leq \text{counter} \leq \frac{\sigma}{(1-\epsilon)\mu},$$

and therefore

$$(1-\epsilon)\mu \leq \frac{\sigma}{\text{counter}} \leq (1+\epsilon)\mu.$$

This completes the proof. \square

6.3 Improving the Error Bounds Through a More Holistic Analysis

We now give a better running time bound by combining the analyses of the two previous estimators in a more global setting. Pseudocode for this final estimation routine is given in Figure 13.

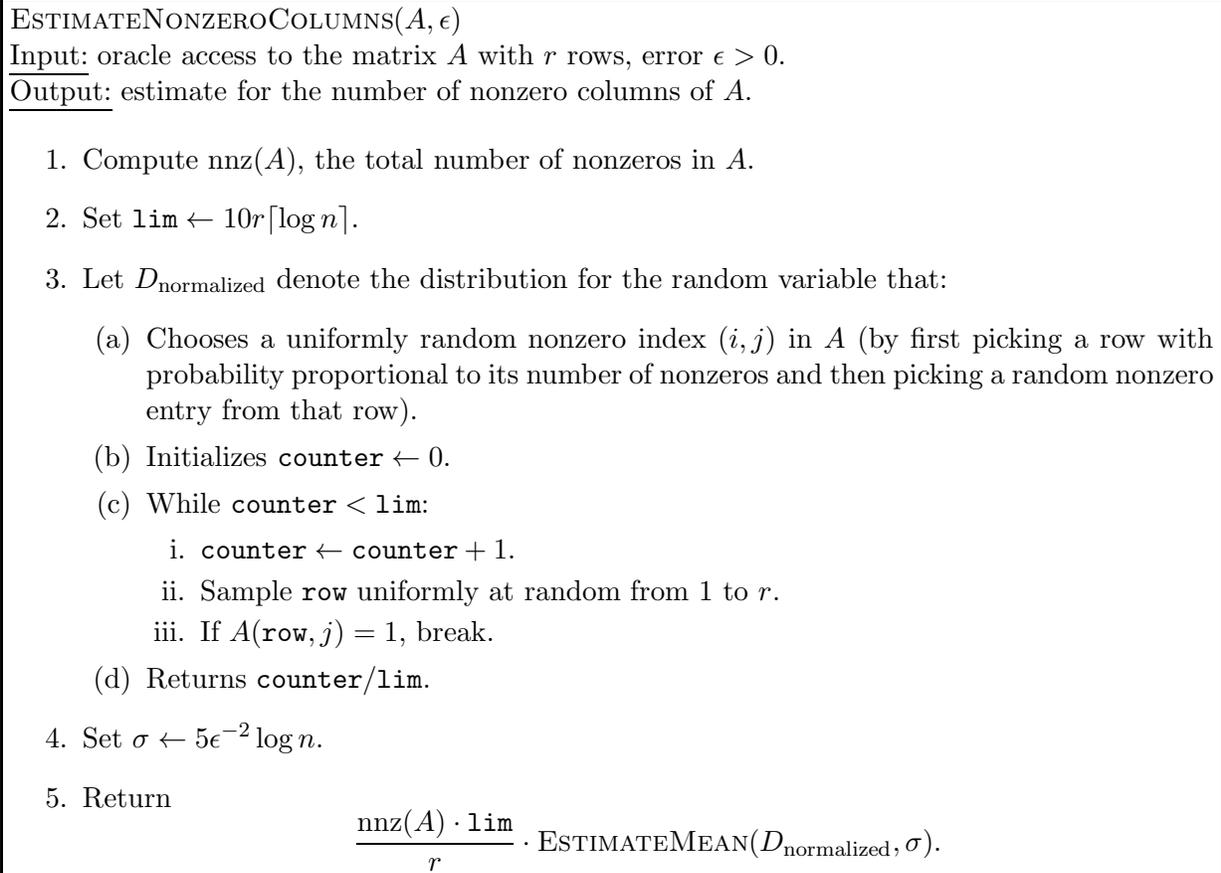


Figure 13: Pseudocode for a faster estimation of the number of nonzero columns of a matrix.

The essence of this algorithm can be better understood by analyzing a simpler version of $D_{\text{normalized}}$. The difference between these two distributions is that we artificially force $D_{\text{normalized}}$ to be a distribution over $[0, 1]$ (first by truncating and then normalizing) so that we can use the ESTIMATEMEAN algorithm. We show that our threshold of $O(r \log n)$ for truncating the number of sampled rows can perturb the expected value by at most a factor of $1/\text{poly}(n)$.

Definition 6.9. We define the simpler variant D_{simple} of $D_{\text{normalized}}$ as follows:

1. Sample a nonzero index (i, j) from A uniformly at random.
2. Return the minimum of $10r \lceil \log n \rceil$ and the number of random rows row until $A(\text{row}, j) = 1$.

Now we analyze the expected value of D_{simple} and relate it to that of $D_{\text{normalized}}$ to prove the correctness of ESTIMATENONZEROCOLUMNS and bound the time needed to sample from $D_{\text{normalized}}$.

Lemma 6.10. *If X is a random variable drawn from $D_{\text{normalized}}$, then*

$$\left(1 - \frac{1}{n}\right) \frac{r \cdot \text{NONZEROCOLUMNS}(A)}{\text{lim} \cdot \text{nnz}(A)} \leq \mathbb{E}[X] \leq \frac{r \cdot \text{NONZEROCOLUMNS}(A)}{\text{lim} \cdot \text{nnz}(A)},$$

and the expected cost of each sample generated from $D_{\text{normalized}}$ is

$$O\left(\frac{r \cdot \text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)}\right).$$

Proof. For each column j , denote by $\text{nnz}(A[:, j])$ the number of nonzero entries in the column, and let the probability of picking a nonzero entry from this column be

$$p_j \stackrel{\text{def}}{=} \frac{\text{nnz}(A[:, j])}{r}.$$

When considering columns selected by D_{combined} and D_{simple} , the column necessarily has a nonzero entry so $p_j > 0$.

Next, define the random variable H_j to be the number of times it takes to independently sample a row row between 1 and r uniformly at random such that $A(\text{row}, j) = 1$. For all integers $k \geq 1$,

$$\Pr[H_j = k] = (1 - p_j)^{k-1} p_j.$$

It follows that the expected value of H_j is

$$\begin{aligned} \mathbb{E}[H_j] &= \sum_{k=1}^{\infty} k \Pr[H_j = k] \\ &= p_j \sum_{k=1}^{\infty} k (1 - p_j)^{k-1} \\ &= \frac{p_j}{(1 - (1 - p_j))^2} \\ &= \frac{r}{\text{nnz}(A[:, j])}, \end{aligned}$$

where the second to last line uses the Maclaurin series

$$\sum_{k=1}^{\infty} kx^{k-1} = \frac{1}{(1-x)^2},$$

for all $|x| < 1$. Note that we can apply this formula because $p_j > 0$.

To account for truncation, consider the random variable $Z_j = \min(H_j, \text{lim})$ and observe that

$$\begin{aligned} \mathbb{E}[Z_j] &= \sum_{k=1}^{\text{lim}} k \Pr[H_j = k] + \sum_{k=1}^{\infty} (\text{lim} + k - k) \Pr[H_j = \text{lim} + k] \\ &= \mathbb{E}[H_j] - (1 - p_j)^{\text{lim}} p_j \sum_{k=1}^{\infty} k (1 - p_j)^{k-1} \\ &= \mathbb{E}[H_j] - \frac{(1 - p_j)^{\text{lim}}}{p_j}, \end{aligned}$$

where we again use the Maclaurin series above. Recalling that $\text{nnz}(A[:, j]) \geq 1$ and $r \leq n$, it follows from the definition of p_j that

$$\begin{aligned} \frac{(1 - p_j)^{\text{lim}}}{p_j} &\leq n(1 - p_j)^{\text{lim}} \\ &\leq n \exp\left(-\frac{\text{nnz}(A[:, j])}{r} \cdot 10r \log n\right) \\ &= n \left(\frac{1}{n}\right)^{10 \cdot \text{nnz}(A[:, j])} \\ &\leq \frac{1}{n^9}. \end{aligned}$$

Therefore, by truncating the number of row samples, we deviate from $\mathbb{E}[H_j]$ by at most $1/\text{poly}(n)$. Letting Z be a random variable drawn from D_{simple} , we have the conditional expectation bound

$$\frac{r}{\text{nnz}(A[:, j])} - \frac{1}{n^9} \leq \mathbb{E}[Z \mid Z \text{ chooses column } j] \leq \frac{r}{\text{nnz}(A[:, j])}.$$

If we now consider the probability that Z selects column j , it follows that

$$\frac{r \cdot \text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)} \left(1 - \frac{1}{n^7}\right) \leq \mathbb{E}[Z] \leq \frac{r \cdot \text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)},$$

since $\text{nnz}(A) = O(n^2)$. To prove the claim for the second distribution $D_{\text{normalized}}$, we use the same argument and the additional fact that $\text{lim} = O(n^2)$. The expected running time per sample follows from the expected value of D_{simple} . \square

Proof of Lemma 6.1. Since $D_{\text{normalized}}$ is a distribution over $[0, 1]$, we can use ESTIMATEMEAN to approximate the expected value within a factor of ϵ with high probability by Lemma 6.4. The correctness of ESTIMATENONZEROCOLUMNS follows from our choice of σ and Lemma 6.10.

It remains to bound the running time of the estimator. By Lemma 6.4 and Lemma 6.10, the expected number of times we sample $D_{\text{normalized}}$ is

$$O\left(\frac{\text{nnz}(A) \cdot \text{lim} \cdot \log n \epsilon^{-2}}{r \cdot \text{NONZEROCOLUMNS}(A)}\right).$$

The expected cost per sample generated from $D_{\text{normalized}}$ is

$$O\left(\frac{r \cdot \text{NONZEROCOLUMNS}(A)}{\text{nnz}(A)}\right)$$

by Lemma 6.10. Multiplying these expectations gives that the total expected running time is

$$O(\text{lim} \cdot \log n \epsilon^{-2}) = O(r \log^2 n \epsilon^{-2}).$$

Furthermore, since the cost per sample from $D_{\text{normalized}}$ is bounded by $\text{lim} = O(r \log n)$, it follows that the running time is concentrated around this value with high probability. \square

7 Maintaining Graphs Under Pivots

In this section we first show how to efficiently maintain the component graph under pivots in such a way that supports component and remaining neighborhood queries. This proves Lemma 6.2, which in turns completes the proof of Theorem 5.2. Then we spend the rest of the section demonstrating how to maintain a dynamic 1-neighborhood sketch (described in Definition 3.1) of a fill graph as it undergoes vertex eliminations to prove Theorem 4.4.

For convenience, we restate our claim about supporting dynamic component graphs G° .

Lemma 6.2. *We can maintain a component graph under vertex pivots in a total time of $O(m \log^2 n)$. Additionally, this component graph data structure grants $O(\log n)$ time oracle access for:*

- *Querying the state of a vertex.*
- *Querying the component or remaining neighborhood (and hence degree) of a vertex.*
- *Uniformly sampling a remaining neighbor of a component or remaining vertex.*
- *Uniformly sampling a random component vertex.*

Proof. We maintain the adjacency list of G° explicitly, where each node stores its state as well as its neighbors in balanced binary search trees. When we pivot a remaining vertex v , we examine all all of its neighbors that are component vertices (i.e. $N_{\text{comp}}^\circ(u)$) and merge the neighborhood lists of these vertices. By always inserting elements from the smaller list into the larger, we can guarantee that each element is inserted at most $O(\log n)$ times. The total number of elements in the neighborhood lists is m , so since each insertion costs $O(\log n)$ time, it follows that the total cost across all m edges is $O(m \log^2 n)$.

When a vertex is pivoted, we also need to move it from the remaining vertex list to the component vertex list for each of its neighbors. This can be done by iterating over all the edges of the vertex once. The cost is $O(m)$ because each vertex is pivoted at most once and prior to pivoting no edges are added to the graph. By maintaining all lists using balanced binary search trees, we can make all updates and sample a random remaining (or component) neighbor in $O(\log n)$ time. A global list that tracks all remaining and component vertices also allows for $O(\log n)$ time uniform sampling. \square

We now focus on proving Theorem 4.4. We maintain a 1-neighborhood ℓ_0 -sketch data structure for a fill graph as it undergoes pivots (starting with the original graph) similarly to how we maintain the adjacency list of the component graph. Because the minimum key value $R(v)$ in the 1-neighborhood of a vertex u continually changes, we show how to track the minimizer of a vertex via an eager-propagation routine. This protocol informs the neighbors of a pivoted vertex about its minimum key, which ultimately propagates minimum key values throughout the graph as needed.

In Figure 14, we give a brief description about the data structures we use to achieve this. When we refer to maintaining sets of vertices, the underlying data structure is a balanced binary search tree. Recall that min heaps (e.g. binary heaps) of size $O(n)$ support the methods MIN, INSERT, and DELETE which require $O(1)$, $O(\log n)$, and $O(\log n)$ time respectively. Additionally, we use a standard subroutine HEAPMERGE to merge two heaps in $O(\log^2 n)$ time.

1. A set V_{rem}° containing the remaining vertices.
2. A set V_{comp}° containing the component vertices.
3. For each $x \in V_{\text{comp}}^\circ$, a corresponding min heap

remaining[x]

that contains the key values $R(v)$ of its remaining neighbors $v \in N_{\text{rem}}^\circ(x)$.

4. For each $u \in V_{\text{rem}}^\circ$, a corresponding min heap

fill[u]

that contains the union of **remaining**[x].MIN() for each component vertex $x \in N_{\text{comp}}^\circ(u)$, as well as the key values of the vertices in $N_{\text{rem}}^\circ(u)$.

Figure 14: Data structures needed to maintain G° and an ℓ_0 -sketch of G^+ under vertex pivots.

As a brief intuition behind the algorithm, consider the case where no vertex is deleted, but we merge neighborhoods of vertices. In this case, as the neighborhood of a particular vertex grows, the expected number of times the minimum R value in this neighborhood changes is $O(\log n)$. To see this, consider the worst case where each time, the neighborhood of v increases by 1. Then the expected number of changes in the minimum is

$$1/2 + 1/3 + \dots + 1/(n-1) \leq O(\log n).$$

The major difficulty dealing with this is that deletions reduce degrees. In particular, it is possible for the min at some vertex to change $\Omega(n)$ times due to repeated deletions. As a result, we can only bound the total, or average number of propagations. This leads to a much more involved amortized analysis, where we also use backwards analysis to explicitly bound the probability of each informing operation.

Given a component graph G_t° and a (remaining) vertex u to be pivoted, we use the routine PIVOTVERTEX to produce a new graph G_{t+1}° . In terms of the structure of the graph, our routine does the same thing as the traditional quotient graph model for symmetric factorization [GL81].

Therefore we turn our attention to the problem of maintaining the minimum R values of the neighborhoods. For a subset of vertices $V' \subseteq V_{\text{comp}}^\circ$, let $R_{\min}(V')$ denote the minimum R value among all its vertices. Specifically, we want to maintain the values $R_{\min}(N_{\text{rem}}^\circ(w))$ for every $w \in V_{\text{comp}}^\circ$ and $R_{\min}(N^+(v))$ for every $v \in V_{\text{rem}}^\circ$. This update procedure is basically a notification mechanism. When the status of a vertex changes, we update the data structures of its neighbors correspondingly. The **fill**[u] heap will then give $R_{\min}(N^+(u))$ and be used to estimate the fill-degree of each remaining vertex as described in Section 4.

Suppose a remaining vertex v is pivoted. Then, for a component vertex w , the content of **remaining**[w] changes only if v is its neighbor. Pseudocode of this update (PIVOTVERTEX) is given in Figure 15. In particular, since v is no longer a remaining vertex, its entry needs to be

removed from `remaining[z]`. Since v is now a component vertex, we need to construct `remaining[v]`, and update the `fill` heaps of its remaining neighbors appropriately. Furthermore, if $R(v)$ was the minimum element in `remaining[w]`, this is no longer the case and the other remaining neighbors of w need to be notified of this (so they can update their `fill` heaps). This is done via the call to `INFORMREMAINING` in Line 5b of the algorithm. The last step consists of melding the (now component) vertex v with its existing component neighbors via calls to `MELD`. The pseudocode for this routine is in Figure 16. Note that, at all times, we make a note of any remaining vertex whose `MINIMIZER` is updated due to the pivoting.

<p><code>PIVOTVERTEX(v)</code> <u>Input:</u> (implicitly as a global variable) a component graph $G_t^\circ = \langle V_{\text{rem}}^\circ, V_{\text{comp}}^\circ, E^\circ \rangle$ along with associated data structures. A vertex $v \in V_{\text{rem}}^\circ$ to be pivoted, <u>Output:</u> A list of vertices in V_{rem}° whose <code>MINIMIZERS</code> have changed.</p> <ol style="list-style-type: none"> 1. Initialize <code>changed_list</code> $\leftarrow \emptyset$. 2. Create an empty min-heap <code>remaining[v]</code> 3. For each vertex $y \in N_{\text{rem}}^\circ(v)$ in lexicographical order <ol style="list-style-type: none"> (a) <code>fill[y].DELETE(R(v))</code> (b) <code>remaining[v].INSERT(R(y))</code> (c) If $R(v)$ was the old minimum in <code>fill[y]</code>: <code>changed_list</code> \leftarrow <code>changed_list</code> \cup $\{y\}$ 4. For each vertex $y \in N_{\text{rem}}^\circ(v)$ in lexicographical order <ol style="list-style-type: none"> (a) <code>fill[y].INSERT(remaining[v].MIN())</code> (if not already present) (b) If the minimum in <code>fill[y]</code> changes: <code>changed_list</code> \leftarrow <code>changed_list</code> \cup $\{y\}$ 5. For each vertex $w \in N_{\text{comp}}^\circ(v)$ in lexicographical order <ol style="list-style-type: none"> (a) <code>remaining[w].DELETE(R(v))</code>. (b) If $R(v)$ was the old minimum in <code>remaining[w]</code>: <code>changed_list</code> \leftarrow <code>changed_list</code> \cup <code>INFORMREMAINING(w, R(v), remaining[w].MIN())</code> (c) <code>MELD(v, w)</code>; 6. Update $V_{\text{comp}}^\circ, V_{\text{rem}}^\circ$ and E° to form G_{t+1}°; 7. Return <code>changed_list</code>.
--

Figure 15: Pseudocode for pivoting a vertex

<p>MELD(v, w)</p> <p><u>Input:</u> (implicitly as a global variable) A graph state $G = \langle V_{\text{rem}}^{\circ}, V_{\text{comp}}^{\circ}, E \rangle$ along with associated data structures.</p> <p>Two component vertices v (the pivoted vertex) and w to be melded.</p> <p><u>Output:</u> A list of vertices in $N_{\text{rem}}^{\circ}(v) \cup N_{\text{rem}}^{\circ}(w)$ whose MINIMIZERS have changed.</p> <ol style="list-style-type: none"> 1. Initialize <code>changed_list</code> $\leftarrow \emptyset$. 2. If <code>remaining[v].MIN()</code> < <code>remaining[w].MIN()</code> <ol style="list-style-type: none"> (a) <code>changed_list</code> \leftarrow INFORMREMAINING($w, \text{remaining}[w].\text{MIN}(), \text{remaining}[v].\text{MIN}()$); 3. Else If <code>remaining[w].MIN()</code> < <code>remaining[v].MIN()</code> <ol style="list-style-type: none"> (a) <code>changed_list</code> \leftarrow INFORMREMAINING($v, \text{remaining}[v].\text{MIN}(), \text{remaining}[w].\text{MIN}()$); 4. <code>remaining[v]</code> \leftarrow HEAPMERGE(<code>remaining[v]</code>, <code>remaining[w]</code>) 5. Return <code>changed_list</code>.
--

Figure 16: Pseudocode for melding two component vertices, and informing their neighbors of any changes in the minimizers of N_{rem}° .

For every component vertex w such that $R(v)$ is the minimum value in `remaining(w)`, the routine INFORMREMAINING (Pseudocode in Figure 17) is responsible for updating the contents in the `fill` heaps of remaining vertices adjacent to w . This routine is also required when we merge two component vertices in the algorithm MELD, since there are now more entries in the `fill` heaps of adjacent remaining vertices.

<p>INFORMREMAINING(w, R_{old}, R_{new})</p> <p><u>Input:</u> (implicitly as a global variable) a component graph $G^\circ = \langle V_{rem}^\circ, V_{comp}^\circ, E^\circ \rangle$ along with associated data structures;</p> <p>a vertex $w \in V_{comp}^\circ$ that's causing updates;</p> <p>old and new values for $R_{min}(N_{rem}^\circ(w))$: R_{old} and R_{new}.</p> <p><u>Output:</u> A list of vertices $v \in N_{rem}^\circ(w)$ whose MINIMIZERS have changed.</p> <ol style="list-style-type: none"> 1. Initialize <code>changed_list</code> $\leftarrow \emptyset$. 2. For each $v \in N_{rem}^\circ(w)$ <ol style="list-style-type: none"> (a) Delete the entry R_{old} from <code>fill[v]</code> if it exists (b) Add the entry R_{new} to <code>fill[v]</code> (c) If <code>fill[v].MIN()</code> changed, <code>changed_list</code> \leftarrow <code>changed_list</code> \cup $\{v\}$. 3. Return <code>changed_list</code>.

Figure 17: Pseudocode for propagating to remaining vertex neighbors

We break down the cost of calls to INFORMREMAINING into two parts: when it is invoked by PIVOTVERTEX, and when it is invoked by MELD. The first type of calls happens only when a remaining vertex v is pivoted, and v is the minimum entry of the *remaining* heap of a component vertex. The following lemma gives an upper bound on the expected cost of such calls to INFORMREMAINING by arguing that this event happens with low probability.

Lemma 7.1. *The expected total number of updates to remaining vertices made by INFORMREMAINING when invoked from PIVOTVERTEX (Line 5b) over any sequence of n pivots that are independent of the R values is $O(m)$.*

Proof. Let G° be the component graph at a certain instant in the algorithm. Let $v \in V_{rem}^\circ$ be the vertex to be pivoted, and let $w \in N_{comp}^\circ(v)$ be a neighboring component vertex. We only invoke INFORMREMAINING if $R(v)$ is the minimum value in `remaining[w]`, which occurs with probability $1/|N_{rem}^\circ(w)|$ and would cost $O(|N_{rem}^\circ(w)|)$ updates. Therefore the expected number of updates is only $O(1)$ for each edge between a remaining vertex and a component vertex. When a remaining vertex v is pivoted, its degree is the same as in the original graph. Therefore the number of edges between v and a component vertex is bounded by the degree of v and hence the total expected number of updates is $O(\sum_{v \in V} \deg(v)) = O(m)$. \square

The calls to MELD are the primary bottlenecks in the running time, but will be handled similarly. Its pseudocode is given in Figure 16.

We will show that the expected number of vertices updated by INFORMREMAINING that result from any fixed sequence of calls to MELD is bounded by $O(m \log n)$. We first analyze the number of updates during a single meld in the following lemma.

Lemma 7.2. *Let u and v be two component vertices in a graph stage G° . Then the expected number of updates to vertices by INFORMREMAINING when melding u and v is at most:*

$$\frac{2 |N_{\text{rem}}^\circ(u)| \cdot |N_{\text{rem}}^\circ(v)|}{|N_{\text{rem}}^\circ(u)| + |N_{\text{rem}}^\circ(v)|},$$

Proof. Let us define:

$$\begin{aligned} n_{\text{common}} &= |N_{\text{rem}}^\circ(u) \cap N_{\text{rem}}^\circ(v)|, \\ n_u &= |N_{\text{rem}}^\circ(u) \setminus N_{\text{rem}}^\circ(v)|, \\ n_v &= |N_{\text{rem}}^\circ(v) \setminus N_{\text{rem}}^\circ(u)|. \end{aligned}$$

If the minimum R value is generated by a vertex from $N_{\text{rem}}^\circ(u) \cap N_{\text{rem}}^\circ(v)$, then no cost is incurred. If it is generated by a vertex from $N_{\text{rem}}^\circ(u) \setminus N_{\text{rem}}^\circ(v)$, we need to update the every vertex in $N_{\text{rem}}^\circ(v)$ (line 2a). This happens with probability

$$\begin{aligned} \frac{n_u}{n_{\text{common}} + n_u + n_v} &\leq \frac{n_u + n_{\text{common}}}{2n_{\text{common}} + n_u + n_v} \\ &= \frac{|N_{\text{rem}}^\circ(u)|}{|N_{\text{rem}}^\circ(u)| + |N_{\text{rem}}^\circ(v)|}. \end{aligned}$$

Therefore the expected number of updates is bounded by:

$$\frac{|N_{\text{rem}}^\circ(u)| \cdot |N_{\text{rem}}^\circ(v)|}{|N_{\text{rem}}^\circ(u)| + |N_{\text{rem}}^\circ(v)|},$$

and we get the other term (for updating u 's neighborhood) similarly. \square

This allows us to carry out an amortized analysis for the number of updates to remaining vertices. We will define the potential function of an intermediate component graph during elimination in terms of the degrees of component vertices *in the original graph G* , in which adjacent component vertices are *not* contracted. Let u° denote the set of vertices in $V(G)$ which have been melded into u in G° .

$$\Phi(G^\circ) \stackrel{\text{def}}{=} \sum_{u \in V_{\text{comp}}^\circ} D(u) \log(D(u)),$$

where $D(u)$ for a vertex $u \in V_{\text{comp}}^\circ$ is defined to be

$$D(u) = \sum_{v \in u^\circ} \deg_G(v).$$

This function starts out at 0, and can be at most $m \log n$.

Lemma 7.3. *When melding two neighboring component vertices in a graph G_t° to create G_{t+1}° , the expected number of vertex updates by INFORMREMAINING is at most*

$$2 (\Phi(G_{t+1}^\circ) - \Phi(G_t^\circ)).$$

Proof. When melding two component vertices u and v in G_t° to form G_{t+1}° , the change in potential is given by

$$\Phi(G_{t+1}^\circ) - \Phi(G_t^\circ) = (D(u) + D(v)) \log(D(u) + D(v)) - D(u) \log D(u) - D(v) \log D(v).$$

On the other hand, by Lemma 7.2 the expected number of remaining vertices updated is

$$\frac{2|N_{\text{rem}}^\circ(u)| \cdot |N_{\text{rem}}^\circ(v)|}{|N_{\text{rem}}^\circ(u)| + |N_{\text{rem}}^\circ(v)|} \leq \frac{2D(u)D(v)}{D(u) + D(v)}.$$

To see that the above statement is true, observe that $N_{\text{rem}}^\circ(u) \leq D(u)$, $N_{\text{rem}}^\circ(v) \leq D(v)$, and that both the LHS and RHS can be viewed as two resistors in parallel. Now it suffices to show the following algebraic identity:

$$2x \log x + 2y \log y + \frac{2xy}{x+y} \leq 2(x+y) \log(x+y),$$

and let $x = D(u)$ and $y = D(v)$. By symmetry, we can assume $x \leq y$ without loss of generality. Then we get

$$\begin{aligned} \frac{xy}{x+y} &\leq \frac{xy}{y} \\ &= y \cdot \frac{x}{y} \\ &\leq y \cdot \log\left(1 + \frac{x}{y}\right), \end{aligned}$$

where the last inequality follows from $\log(1+z) \geq z$ when $z \leq 1$. Plugging this in then gives:

$$\begin{aligned} 2x \log x + 2y \log y + \frac{2xy}{x+y} &\leq 2x \log x + 2y \left(\log y + \log\left(1 + \frac{x}{y}\right) \right) \\ &= 2x \log x + 2y \log(x+y) \\ &\leq 2(x+y) \log(x+y). \quad \square \end{aligned}$$

Lemma 7.4. *Over any fixed sequence of calls to MELD, the expected number of updates to the fill heaps in remaining vertices (lines 2a and 3a) is bounded by $O(m \log n)$.*

Proof. By Lemma 7.3, the number of updates is within a constant of the potential increase. Since our potential function Φ is bounded between 0 and $O(m \log n)$, and at no point can it decrease. Hence, the total number of updates is also bounded by $O(m \log n)$. \square

Combining the above lemmas gives our main theorem from Section 4 on maintaining one copy of the 1-neighborhood sketch.

Proof. (of Theorem 4.4) Given any graph G and a fixed sequence of vertices for pivoting, we use the PIVOTVERTEX routine to produce the sequence of graph states

$$G = G_0^\circ, G_1^\circ, G_2^\circ, \dots, G_n^\circ = \emptyset.$$

Recall that the goal is to maintain $R_{\min}(N_{\text{rem}}^\circ(w))$ for all $w \in V_{\text{comp}}^\circ$ and $R_{\min}(N^+(v))$ for all $v \in V_{\text{rem}}^\circ$. This is achieved by maintaining the two min-heaps, **remaining** and **fill**.

When pivoting a remaining vertex v , PIVOTVERTEX first makes a constant number of updates to v 's remaining neighbors, which are bounded above by the original degree of v . Next, it removes v from the **remaining** heaps among v 's component vertex neighbors (line 5a), which are again at most as many as the original degree of v . Therefore the total cost of this part of the algorithm is $O(m \log n)$. The major chunk of the running time cost is incurred by updates to the **fill** heaps in INFORMREMAINING. By Lemma 7.1 and Lemma 7.4, the number of such updates is bounded by $O(m \log n)$. As each update is a $O(\log n)$ operation on a heap, the total running time is $O(m \log^2 n)$. The final step of a meld consists of merging two **remaining** heaps. Since two min-heaps can be merged in time $O(\log^2 n)$, and the number of merges for a pivoted vertex can be bounded by its original degree, the cost of this step can be bounded by $O(m \log^2 n)$. \square

8 SETH-Hardness for Computing Minimum Degree Orderings

Our hardness results for computing the minimum fill degree and the min-degree ordering are based on the *strong exponential time hypothesis* (SETH), which states that for all $\theta > 0$ there exists a k such that solving k -SAT requires $\Omega(2^{(1-\theta)n})$ time. Many hardness results based on SETH, including ours, go through the ORTHOGONALVECTORS problem and make use of the following result.

Theorem 8.1 ([Wil05]). *Assuming SETH, for any $\theta > 0$, there does not exist an $O(n^{2-\theta})$ time algorithm that takes n binary vectors with $\Theta(\log^2 n)$ bits and decides if there is an orthogonal pair.*

We remark that ORTHOGONALVECTORS is often stated as deciding if there exists a pair of orthogonal vectors from two different sets [Wil15], but we can reduce the problem to a single set by appending $[1; 0]$ to all vectors in the first set and $[0; 1]$ to all vectors in the second set.

The first hardness observation for computing the minimum fill degree in a partially eliminated graph is a direct reduction to ORTHOGONALVECTORS. To show this, we construct a bipartite graph that demonstrates how ORTHOGONALVECTORS can be interpreted as deciding if a union of cliques covers the edges of a clique on the remaining vertices of a partially eliminated graph.

Lemma 8.2. *Assuming SETH, for any $\theta > 0$, there does not exist an $O(m^{2-\theta})$ time algorithm that takes as input G with a set of eliminated vertices and computes the minimum fill degree in G^+ .*

Proof. Consider an ORTHOGONALVECTORS instance with n vectors $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n \in \{0, 1\}^d$, and construct a bipartite graph $G = (V_{\text{vec}}, V_{\text{dim}}, E)$ such that each vertex in V_{vec} corresponds to a vector \mathbf{a}_i and each vertex in V_{dim} uniquely corresponds to a dimension $1 \leq j \leq d$. For the edges, we connect vertices $i \in V_{\text{vec}}$ with $j \in V_{\text{dim}}$ if and only if $\mathbf{a}_i(j) = 1$.

Consider the graph state with all of V_{dim} eliminated and all of V_{vec} remaining. We claim that there exists a pair of orthogonal vectors among $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$ if and only if there exists a remaining vertex $v \in V(G^+)$ with $\deg(v) < n - 1$. Let $u, v \in V_{\text{vec}}$ be any two different vertices, and let \mathbf{a}_u and \mathbf{a}_v be their corresponding vectors. The vertices u and v are adjacent in G^+ if and only if there exists an index $1 \leq j \leq d$ such that $\mathbf{a}_u(j) = \mathbf{a}_v(j) = 1$.

Suppose there exists an $O(m^{2-\theta})$ time algorithm for implicitly finding the minimum fill degree in a partially eliminated graph, for some $\theta > 0$. Then for $d = \Theta(\log^2 n)$ we can use this algorithm to compute the vertex with minimum fill degree in the bipartite graph described above in time

$$O\left(m^{2-\theta}\right) = O\left((n \log^2 n)^{2-\theta}\right) = O\left(n^{2-\theta/2}\right),$$

which contradicts SETH by Theorem 8.1. \square

Building on the previous observation, we now show that an exact linear-time algorithm for computing min-degree elimination orderings is unlikely. In particular, our main hardness result is:

Theorem 8.3. *Assuming SETH, for any $\theta > 0$, there does not exist an $O(m^{4/3-\theta})$ time algorithm for producing a min-degree elimination ordering.*

The main idea behind our construction is to modify the bipartite graph in the proof of Lemma 8.2 in such a way that a minimum degree ordering has the effect of eliminating the d vertices in V_{dim} before any vertex in V_{vec} . This allows us to use MINDEGREEORDERING to efficiently solve any instance of ORTHOGONALVECTORS. A limitation of the initial construction is that vertices in V_{dim} can have degree as large as n , so requiring that they be eliminated first is difficult to guarantee. To overcome this problem, we create a degree hierarchy by splitting each vertex in V_{dim} into $\Theta(n)$ vertices with degree $O(\sqrt{n})$. We call this construction a *covering set system* because it maintains all two-step connections between vertices in V_{vec} .

Lemma 8.4. *Given any positive integer n , we can construct in $O(n^{3/2})$ time a covering set system of the integers $[n] = \{1, 2, \dots, n\}$. This system is collection of subsets $I_1, I_2, \dots, I_k \subseteq [n]$ such that:*

- *The number of subsets $k = O(n)$.*
- *The cardinality $|I_j| \leq 10\sqrt{n}$, for all $1 \leq j \leq k$.*
- *For each $(i_1, i_2) \in [n]^2$ there exists a subset I_j such that $i_1, i_2 \in I_j$.*

We also pad each vertex in V_{vec} with $\Omega(\sqrt{n})$ edges to ensure that it is eliminated after the vertices introduced by the covering set systems. We formally describe this construction in Figure 18.

1. Create one vertex per input vector $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$, and let these vertices be V_{vec} .
2. For each dimension $j = 1$ to d :
 - (a) Construct a covering set system for $[n]$.
 - (b) Create a vertex in V_{dim} for each subset in this covering set system.
 - (c) For each vector \mathbf{a}_i such that $\mathbf{a}_i(j) = 1$, add an edge between its vertex in V_{vec} and every vertex corresponding to a subset in this covering system that contains i .
3. Introduce $20\sqrt{n}$ extra vertices called V_{pad} :
 - (a) Connect all pairs of vertices in V_{pad} .
 - (b) Connect each vertex in V_{pad} with every vertex in V_{vec} .

Figure 18: Construction for reducing ORTHOGONALVECTORS to MINDEGREEORDERING.

Lemma 8.5. *For any ORTHOGONALVECTORS instance with n vectors of dimension d , let G be the graph produced by the construction in Figure 18. We have $|V| = O(nd)$ and $|E| = O(n^{3/2}d)$.*

Proof. The number of vertices in G is

$$|V| = 20\sqrt{n} + n + d \cdot O(n) = O(nd).$$

Similarly, an upper bound on the number of edges in G is

$$|E| = \binom{20\sqrt{n}}{2} + 20\sqrt{n} \cdot n + d \cdot 10\sqrt{n} \cdot O(n) = O\left(n^{3/2}d\right),$$

where the terms on the left-hand side of the final equality correspond to edges contained in V_{pad} , the edges between V_{pad} and V_{vec} , and edges between V_{vec} and V_{dim} , respectively. \square

Lemma 8.6. *Consider a graph G constructed from an ORTHOGONALVECTORS instance as described in Figure 18. For any min-degree ordering of G , the vertices in V_{dim} are the first to be eliminated. Furthermore, the fill degree of the next vertex to be eliminated is $\min_{v \in V_{\text{vec}}} \deg^+(v)$.*

Proof. Let the graph be $G = (V, E)$, such that V is partitioned into

$$V = V_{\text{vec}} \cup V_{\text{dim}} \cup V_{\text{pad}},$$

as described in Figure 18. Initially, for every vertex $v_{\text{pad}} \in V_{\text{pad}}$ we have

$$\deg(v_{\text{pad}}) = (20\sqrt{n} - 1) + n.$$

For every vertex $v_{\text{vec}} \in V_{\text{vec}}$ we have

$$\deg(v_{\text{vec}}) = 20\sqrt{n} + |E(v_{\text{vec}}, V_{\text{dim}})| \geq 20\sqrt{n},$$

and for every vertex $v_{\text{dim}} \in V_{\text{dim}}$ we have

$$\deg(v_{\text{dim}}) \leq 10\sqrt{n}.$$

Pivoting out a vertex in V_{dim} does not increase the fill degree of any other vertex in V_{dim} since no two vertices in V_{dim} are adjacent. As these vertices are pivoted, we still maintain

$$\deg^+(v) \geq 20\sqrt{n},$$

for all $v \in V_{\text{vec}}$. Therefore, the first vertices to be pivoted must be all $v \in V_{\text{dim}}$. After all the vertices in V_{dim} have been pivoted, the next vertex has fill degree $\min_{v \in V_{\text{vec}}} \deg^+(v)$, because either a vertex in V_{vec} will be eliminated or all remaining vertices have fill degree $20\sqrt{n} + n - 1$. \square

Proof of Theorem 8.3. Suppose for some $\theta > 0$ there exists an $O(m^{4/3-\theta})$ time algorithm for MIN-DEGREEORDERING. Construct the graph $G = (V, E)$ with covering sets as described in Figure 18. For $d = \Theta(\log^2 n)$, it follows from Lemma 8.5 that $|V| = O(n \log^2 n)$ and $|E| = O(n^{3/2} \log^2 n)$. Therefore, by the assumption, we can obtain a min-degree ordering of G in time

$$O\left(m^{4/3-\theta}\right) = O\left(\left(n^{3/2} \log^2 n\right)^{4/3-\theta}\right) = O\left(n^{2-\theta}\right).$$

By Lemma 8.6, after the first $|V_{\text{dim}}|$ vertices have been pivoted, the fill graph G^+ is essentially identical to the partially eliminated state of the bipartite graph in the proof of Lemma 8.2. We can then compute the fill degree of the next vertex to be eliminated in $O(m) = O(n^{2-\theta})$ time by Lemma 2.2. Checking whether the fill degree of this vertex is $20\sqrt{n} + n - 1$ allows us to solve ORTHOGONALVECTORS in $O(n^{2-\theta})$ time, which contradicts SETH. \square

All that remains is to efficiently construct the covering set system defined in Lemma 8.4. We can interpret this construction as a way to cover all the edges of K_n using $O(n)$ $K_{10\sqrt{n}}$ subgraphs. We note that our construction is closely related to Steiner systems obtained via finite affine planes as well as existence results for covering problem with fixed-size subgraphs [CCLW13, CY98].

Proof of Lemma 8.4. We use a simple property of finite fields. Let $p = \text{NEXTPRIME}(\sqrt{n})$, which we can compute in $O(n)$ since $p < 4\sqrt{n}$ by Bertrand's postulate. Clearly $[n] \subseteq [p^2]$, so it suffices to find a covering for $[p^2]$. Map the elements of $[p^2]$ to the coordinates of a $p \times p$ array in the canonical way so that $1 \mapsto (0, 0), 2 \mapsto (0, 1), \dots, p^2 \mapsto (p-1, p-1)$. For all $(a, b) \in \{0, 1, \dots, p-1\}^2$, define

$$D(a, b) = \left\{ (x, y) \in \{0, 1, \dots, p-1\}^2 : y \equiv ax + b \pmod{p} \right\}$$

to be the diagonal subsets of the array, and define

$$R(a) = \left\{ (x, y) \in \{0, 1, \dots, p-1\}^2 : x \equiv a \pmod{p} \right\}$$

to be the row subsets of the array. Let the collection of these subsets be

$$S = \{D(a, b) : a, b \in \{0, 1, \dots, p-1\}\} \cup \{R(a) : a \in \{0, 1, \dots, p-1\}\}.$$

The construction clearly satisfies the first two conditions. Consider any $(a, b) \in [p^2]^2$ and their coordinates in the array (x_1, y_1) and (x_2, y_2) . If $x_1 = x_2$, then $(x_1, y_1), (x_2, y_2) \in R(x_1)$. Otherwise, it follows that (x_1, y_1) and (x_2, y_2) are solutions to the line

$$y \equiv \frac{y_1 - y_2}{x_1 - x_2} (x - x_1) + y_1 \pmod{p},$$

so the third condition is satisfied. □

Acknowledgements

We thank John Gilbert and Gramoz Goranci for many helpful discussions regarding various topics in this paper. We also would like to acknowledge Animesh Fatehpuria for independently obtaining the construction of covering set systems in Lemma 8.4.

References

- [ACK17] Ittai Abraham, Shiri Chechik, and Sebastian Krininger. Fully dynamic all-pairs shortest paths with worst-case update-time revisited. In *Symposium on Discrete Algorithms (SODA)*, pages 440–452, 2017.
- [ADD96] Patrick R. Amestoy, Timothy A. Davis, and Iain S. Duff. An approximate minimum degree ordering algorithm. *SIAM J. Matrix Anal. Appl.*, 17(4):886–905, October 1996.
- [ADD04] Patrick R. Amestoy, Timothy A. Davis, and Iain S. Duff. Algorithm 837: AMD, an approximate minimum degree ordering algorithm. *ACM Trans. Math. Softw.*, 30(3):381–388, September 2004.

- [AY10] Noga Alon and Raphael Yuster. Solving linear systems through nested dissection. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 225–234. IEEE, 2010.
- [BCK⁺16] Ivan Bliznets, Marek Cygan, Pawel Komosa, Lukáš Mach, and Michał Pilipczuk. Lower bounds for the parameterized complexity of minimum fill-in and other completion problems. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1132–1151. Society for Industrial and Applied Mathematics, 2016.
- [BGS15] Surender Baswana, Manoj Gupta, and Sandeep Sen. Fully dynamic maximal matching in $O(\log n)$ update time. *SIAM Journal on Computing*, 44(1):88–113, 2015. Announced at FOCS’11.
- [BK16] Greg Bodwin and Sebastian Krinninger. Fully dynamic spanners with worst-case update time. In *24th Annual European Symposium on Algorithms, ESA 2016, August 22–24, 2016, Aarhus, Denmark*, pages 17:1–17:18, 2016.
- [BKS12] Surender Baswana, Sumeet Khurana, and Soumojit Sarkar. Fully dynamic randomized algorithms for graph spanners. *ACM Transactions on Algorithms (TALG)*, 8(4):35, 2012.
- [BKSE12] Jeff Bezanson, Stefan Karpinski, Viral B Shah, and Alan Edelman. Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*, 2012.
- [BMMR97] Claudson Bornstein, Bruce Maggs, Gary Miller, and R Ravi. Parallelizing elimination orders with linear fill. In *Foundations of Computer Science, 1997. Proceedings., 38th Annual Symposium on*, pages 274–283. IEEE, 1997.
- [BS90] Piotr Berman and Georg Schnitger. On the performance of the minimum degree ordering for gaussian elimination. *SIAM J. Matrix Anal. Appl.*, 11(1):83–88, 1990.
- [CCLW13] Yeow Meng Chee, Charles J Colbourn, Alan CH Ling, and Richard M Wilson. Covering and packing for pairs. *Journal of Combinatorial Theory, Series A*, 120(7):1440–1449, 2013.
- [CLRS09] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.
- [CM05] Graham Cormode and S. Muthukrishnan. An improved data stream summary: The count-min sketch and its applications. *J. Algorithms*, 55(1):58–75, April 2005.
- [CMP16] Michael B. Cohen, Cameron Musco, and Jakub W. Pachocki. Online row sampling. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, (APPROX/RANDOM)*, pages 7:1–7:18, 2016.
- [Coh97] Edith Cohen. Size-estimation framework with applications to transitive closure and reachability. *J. Comput. Syst. Sci.*, 55(3):441–453, December 1997.

- [CS17] Yixin Cao and R. B. Sandeep. Minimum fill-in: Inapproximability and almost tight lower bounds. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 875–880, 2017.
- [CY98] Yair Caro and Raphael Yuster. Covering graphs: The covering problem solved. *Journal of Combinatorial Theory, Series A*, 83(2):273–282, 1998.
- [DGLN04] Timothy A. Davis, John R. Gilbert, Stefan I. Larimore, and Esmond G. Ng. A column approximate minimum degree ordering algorithm. *ACM Trans. Math. Softw.*, 30(3):353–376, 2004.
- [DRSL16] Timothy A Davis, Sivasankaran Rajamanickam, and Wissam M Sid-Lakhdar. A survey of direct methods for sparse linear systems. *Acta Numerica*, 25:383–566, 2016.
- [ELRS17] Talya Eden, Amit Levi, Dana Ron, and C Seshadhri. Approximately counting triangles in sublinear time. *SIAM Journal on Computing*, 46(5):1603–1646, 2017.
- [Fel71] William Feller. *An Introduction to Probability Theory and its Applications. Vol. II*. Second edition. John Wiley & Sons Inc., New York, 1971.
- [FM85] Philippe Flajolet and G. Nigel Martin. Probabilistic counting algorithms for data base applications. *J. Comput. Syst. Sci.*, 31(2):182–209, September 1985.
- [FV13] Fedor V. Fomin and Yngve Villanger. Subexponential parameterized algorithm for minimum fill-in. *SIAM Journal on Computing*, 42(6):2197–2216, 2013.
- [Geo73] Alan George. Nested dissection of a regular finite element mesh. *SIAM Journal on Numerical Analysis*, 10(2):345–363, 1973.
- [GG11] Eran Gat and Shafi Goldwasser. Probabilistic search algorithms with unique answers and their cryptographic applications. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:136, 2011.
- [GGLN13] André Gaul, Martin H Gutknecht, Jorg Liesen, and Reinhard Nabben. A framework for deflated and augmented krylov subspace methods. *SIAM Journal on Matrix Analysis and Applications*, 34(2):495–518, 2013.
- [GGR13] Oded Goldreich, Shafi Goldwasser, and Dana Ron. On the possibilities and limitations of pseudodeterministic algorithms. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, pages 127–138. ACM, 2013.
- [GL81] Alan George and Joseph W. Liu. *Computer Solution of Large Sparse Positive Definite*. Prentice Hall Professional Technical Reference, 1981.
- [GL89] A. George and W. H. Liu. The evolution of the minimum degree ordering algorithm. *SIAM Rev.*, 31(1):1–19, March 1989.
- [GNP94] John R Gilbert, Esmond G Ng, and Barry W Peyton. An efficient algorithm to compute row and column counts for sparse Cholesky factorization. *SIAM Journal on Matrix Analysis and Applications*, 15(4):1075–1091, 1994.

- [GT87] J. R. Gilbert and R. E. Tarjan. The analysis of a nested dissection algorithm. *Numerische Mathematik*, 50(4):377–404, 1987.
- [Gut07] Martin H Gutknecht. A brief introduction to Krylov space methods for solving linear systems. In *Frontiers of Computational Science*, pages 53–62. Springer, 2007.
- [HEKP01] Pinar Heggernes, S. C. Eisestad, Gary Kumfert, and Alex Pothen. The computational complexity of the minimum degree algorithm. Technical report, Institute for Computer Applications in Science and Engineering, 2001.
- [HKN14] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. In *Symposium on Foundations of Computer Science (FOCS)*, pages 146–155, 2014.
- [HP07] Bruce Hendrickson and Alex Pothen. Combinatorial scientific computing: The enabling power of discrete algorithms in computational science. In *Proceedings of the 7th International Conference on High Performance Computing for Computational Science, VECPAR’06*, pages 260–280, 2007.
- [HR98] Bruce Hendrickson and Edward Rothberg. Improving the run time and quality of nested dissection ordering. *SIAM Journal on Scientific Computing*, 20(2):468–489, 1998.
- [KKM13] Bruce M Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1131–1142. Society for Industrial and Applied Mathematics, 2013.
- [KLM⁺17] Michael Kapralov, Yin Tat Lee, C. N. Musco, C. P. Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. *SIAM Journal on Computing*, 46(1):456–477, 2017.
- [KNP⁺17] Michael Kapralov, Jelani Nelson, Jakub Pachocki, Zhengyu Wang, David P Woodruff, and Mobin Yahyazadeh. Optimal lower bounds for universal relation, and for samplers and finding duplicates in streams. *arXiv preprint arXiv:1704.00633*, 2017.
- [KP17] John Kallaugher and Erie Price. A hybrid sampling scheme for triangle counting. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1778–1797. SIAM, 2017.
- [KPPS17] Rasmus Kyng, Jakub Pachocki, Richard Peng, and Sushant Sachdeva. A framework for analyzing resparsification algorithms. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2032–2043. SIAM, 2017.
- [KS16] Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians-fast, sparse, and simple. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 573–582. IEEE, 2016.
- [KST99] Haim Kaplan, Ron Shamir, and Robert E Tarjan. Tractability of parameterized completion problems on chordal, strongly chordal, and proper interval graphs. *SIAM Journal on Computing*, 28(5):1906–1922, 1999.

- [Liu85] Joseph W. H. Liu. Modification of the minimum-degree algorithm by multiple elimination. *ACM Transactions on Mathematical Software (TOMS)*, 11(2):141–153, 1985.
- [Liu90] Joseph WH Liu. The role of elimination trees in sparse factorization. *SIAM Journal on Matrix Analysis and Applications*, 11(1):134–172, 1990.
- [LRT79] R. J. Lipton, D. J. Rose, and R. E. Tarjan. Generalized nested dissection. *SIAM J. on Numerical Analysis*, 16:346–358, 1979.
- [LS15] Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, FOCS '15, pages 230–249, 2015.
- [Mas00] Pascal Massart. About the constants in Talagrand’s concentration inequalities for empirical processes. *Annals of Probability*, pages 863–884, 2000.
- [Mat17] MATLAB optimization toolbox, 2017. The MathWorks, Natick, MA, USA.
- [MPVX15] Gary L Miller, Richard Peng, Adrian Vladu, and Shen Chen Xu. Improved parallel algorithms for spanners and hopsets. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 192–201. ACM, 2015.
- [MPX13] Gary L. Miller, Richard Peng, and Shen Chen Xu. Parallel graph decompositions using random shifts. In *25th ACM Symposium on Parallelism in Algorithms and Architectures, (SPAA)*, pages 196–203, 2013.
- [NS12] Uwe Naumann and Olaf Schenk. *Combinatorial Scientific Computing*. Chapman & Hall/CRC, 1st edition, 2012.
- [NS17] Danupon Nanongkai and Thatchaphol Saranurak. Dynamic spanning forest with worst-case update time: adaptive, las vegas, and $o(n^{1/2} - \epsilon)$ -time. In *Symposium on Theory of Computing (STOC)*, pages 1122–1129, 2017.
- [NSS00] Assaf Natanzon, Ron Shamir, and Roded Sharan. A polynomial approximation algorithm for the minimum fill-in problem. *SIAM Journal on Computing*, 30(4):1067–1079, 2000.
- [NSW17] Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen. Dynamic minimum spanning forest with subpolynomial worst-case update time. In *Symposium on Foundations of Computer Science (FOCS)*, pages 950–961, 2017.
- [RTL76] Donald J. Rose, Robert Endre Tarjan, and George S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM J. Comput.*, 5(2):266–283, 1976.
- [SA96] Raimund Seidel and Cecilia R Aragon. Randomized search trees. *Algorithmica*, 16(4):464–497, 1996.
- [Sol16] Shay Solomon. Fully dynamic maximal matching in constant update time. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 325–334. IEEE, 2016.

- [WAPL14] Yu Wu, Per Austrin, Toniann Pitassi, and David Liu. Inapproximability of treewidth and related problems. *Journal of Artificial Intelligence Research*, 49:569–600, 2014.
- [Wil05] Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science*, 348(2-3):357–365, 2005.
- [Wil15] Virginia V. Williams. Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis. In *Proceedings of the International Symposium on Parameterized and Exact Computation*, pages 16–28, 2015.
- [Wol18] Wolfram Research, Inc. Mathematica, Version 11.3, 2018. Champaign, IL.
- [Woo14] David P Woodruff. Sketching as a tool for numerical linear algebra. *Theoretical Computer Science*, 10(1-2):1–157, 2014.
- [Wul17] Christian Wulff-Nilsen. Fully-dynamic minimum spanning forest with improved worst-case update time. In *Symposium on Theory of Computing (STOC)*, pages 1130–1143, 2017.
- [Yan81] Mihalis Yannakakis. Computing the minimum fill-in is np-complete. *SIAM Journal on Algebraic Discrete Methods*, 2(1):77–79, 1981.

A Guarantees for Selection-Based Estimators

In this section we prove Lemma 4.7, which states that the reciprocal of the $\lfloor k(1 - 1/e) \rfloor$ -quantile in $\text{minimizers}[u]$ can be used to accurately approximate $\deg(u)$. Our proofs follow the same outline as in [Coh97, Section 7], but we consider keys $R(u)$ drawn uniformly from $[0, 1)$ instead of the exponential distribution. We restate the lemma for convenience.

Lemma 4.7. *Suppose that we have k copies of the ℓ_0 -sketch data structure, for $k = 50 \lceil \log n \epsilon^{-2} \rceil$. Let u be any vertex such that $\deg(u) + 1 > 2\epsilon^{-1}$, and let $Q(u)$ denote the $\lfloor k(1 - 1/e) \rfloor$ -ranked key value in the list $\text{minimizers}[u]$. Then, with high probability, we have*

$$\frac{1 - \epsilon}{\deg(u) + 1} \leq Q(u) \leq \frac{1 + \epsilon}{\deg(u) + 1}.$$

We start by stating Hoeffding’s tail inequality for sums of independent Bernoulli random variables, and then we give a useful numerical bound that relates the approximation error and degree of a vertex to the probability that the quantile $Q(u)$ variable deviates from its expected value.

Lemma A.1 (Hoeffding’s inequality). *Let X_1, X_2, \dots, X_n be i.i.d. Bernoulli random variables such that $\Pr[X_i = 1] = p$ and $\Pr[X_i = 0] = 1 - p$. Then for any $\delta > 0$ we have the inequalities*

$$\Pr \left[\sum_{i=1}^n X_i \leq (p - \delta)n \right] \leq \exp(-2\delta^2 n),$$

$$\Pr \left[\sum_{i=1}^n X_i \geq (p + \delta)n \right] \leq \exp(-2\delta^2 n).$$

Lemma A.2. For any $|\epsilon| < 0.1$ and $d \geq 1$, we have

$$\exp\left(-1 + \epsilon - \frac{1}{d+1}\right) \leq \left(1 - \frac{1-\epsilon}{d+1}\right)^{d+1} \leq \exp(-1 + \epsilon).$$

Proof. The Maclaurin series for $\log(1-x)$ is

$$\log(1-x) = -x - \frac{1}{2}x^2 - \frac{1}{3}x^3 - \frac{1}{4}x^4 \dots,$$

for $-1 \leq x < 1$. Whenever $|x| \leq 0.1$, we have the inequality

$$\left|\frac{1}{3}x + \frac{1}{4}x^2 + \frac{1}{5}x^3 + \dots\right| \leq \frac{0.1}{3} + \frac{0.01}{4} + \frac{0.001}{5} + \dots \leq \frac{1}{2}.$$

It follows that

$$-x - x^2 \leq \log(1-x) \leq -x.$$

Applying this inequality when $d+1 \geq 10$ and $x = (1-\epsilon)/(d+1) < 0.1$ gives

$$-\frac{1-\epsilon}{d+1} - \frac{1}{(d+1)^2} \leq \log\left(1 - \frac{1-\epsilon}{d+1}\right) \leq -\frac{1-\epsilon}{d+1}.$$

The result for $d+1 \geq 10$ follows by multiplying the inequalities by $d+1$ and then exponentiating. Checking the remaining cases numerically completes the proof. \square

For convenience, we split the proof of Lemma 4.7 into two parts—one for the upper tail inequality and one for the lower tail inequality.

Lemma A.3. Assuming the hypothesis in Lemma 4.7, we have

$$\Pr\left[Q(u) \geq \frac{1+\epsilon}{\deg(u)+1}\right] \leq \frac{1}{n^4}.$$

Proof. For each sketch $i \in [k]$, we have

$$\begin{aligned} \Pr\left[R(\text{MINIMIZER}_i(u)) \geq \frac{1+\epsilon}{\deg(u)+1}\right] &= \prod_{v \in N(u) \cup \{u\}} \Pr\left[R_i(v) \geq \frac{1+\epsilon}{\deg(u)+1}\right] \\ &= \left(1 - \frac{1+\epsilon}{\deg(u)+1}\right)^{\deg(u)+1}. \end{aligned}$$

Letting I_i be the indicator variable for the event $R(\text{MINIMIZER}_i(u)) \geq (1+\epsilon)/(\deg(u)+1)$, it follows that

$$\mathbb{E}[I_i] = \left(1 - \frac{1+\epsilon}{\deg(u)+1}\right)^{\deg(u)+1}$$

and

$$\Pr\left[Q(u) \geq \frac{1+\epsilon}{\deg(u)+1}\right] = \Pr\left[\sum_{i=1}^k I_i \geq \lceil k/e \rceil\right].$$

Since $\mathbb{E}[I_i] \leq \exp(-(1+\epsilon))$, we let $\delta = 1/e - \mathbb{E}[I_i] > 0$ and use Hoeffding's inequality to show that

$$\begin{aligned} \Pr \left[\sum_{i=1}^k I_i \geq k/e \right] &\leq \exp(-2k\delta^2) \\ &\leq \exp\left(-100 \log n (\delta/\epsilon)^2\right), \end{aligned}$$

where the last inequality uses the fact that $k = 50 \lceil \log n \epsilon^{-2} \rceil$. For any $\epsilon < 1$, we have

$$\frac{\delta}{\epsilon} \geq \frac{1}{\epsilon} \left(\frac{1}{e} - \frac{1}{e^{1+\epsilon}} \right) \geq \frac{1}{5}.$$

Therefore, it follows that

$$\Pr \left[Q(u) \geq \frac{1+\epsilon}{\deg(u)+1} \right] \leq \Pr \left[\sum_{i=1}^k I_i \geq k/e \right] \leq \frac{1}{n^4},$$

as desired. □

Lemma A.4. *Assuming the hypothesis in Lemma 4.7, we have*

$$\Pr \left[Q(u) \leq \frac{1-\epsilon}{\deg(u)+1} \right] \leq \frac{1}{n^4}.$$

Proof. For each sketch $i \in [k]$, we have

$$\begin{aligned} \Pr \left[R(\text{MINIMIZER}_i(u)) \geq \frac{1-\epsilon}{\deg(u)+1} \right] &= \prod_{v \in N(u) \cup \{u\}} \Pr \left[R_i(v) \geq \frac{1-\epsilon}{\deg(u)+1} \right] \\ &= \left(1 - \frac{1-\epsilon}{\deg(u)+1} \right)^{\deg(u)+1}. \end{aligned}$$

Letting J_i be the indicator variable for the event $R(\text{MINIMIZER}_i(u)) \geq (1-\epsilon)/(\deg(u)+1)$, it follows that

$$\mathbb{E}[J_i] = \left(1 - \frac{1-\epsilon}{\deg(u)+1} \right)^{\deg(u)+1}$$

and

$$\Pr \left[Q(u) < \frac{1-\epsilon}{\deg(u)+1} \right] = \Pr \left[\sum_{i=1}^k J_i \leq \lceil k/e \rceil \right].$$

Let $\delta = \mathbb{E}[J_i] - 1/e$. Using Lemma A.2 and the assumption that $\deg(u)+1 > 2\epsilon^{-1}$, observe that

$$\delta \geq \exp\left(-1 + \epsilon - \frac{1}{\deg(u)+1}\right) - 1/e > 0.$$

Therefore, by Hoeffding's inequality we have

$$\begin{aligned} \Pr \left[\sum_{i=1}^k J_i \leq k/e \right] &\leq \exp(-2k\delta^2) \\ &\leq \exp\left(-100 \log n (\delta/\epsilon)^2\right). \end{aligned}$$

For any $\epsilon < 1$, using the lower bound for δ and the assumption that $\deg(u) + 1 > 2\epsilon^{-1}$ gives

$$\frac{\delta}{\epsilon} \geq \frac{1}{e \cdot \epsilon} \left(\exp\left(\epsilon - \frac{1}{\deg(u) + 1}\right) - 1 \right) \geq \frac{e^{\epsilon/2} - 1}{e \cdot \epsilon} \geq \frac{1}{2e}.$$

Therefore, it follows that

$$\Pr \left[Q(u) \leq \frac{1 - \epsilon}{\deg(u) + 1} \right] \leq \Pr \left[\sum_{i=1}^k J_i \leq \lceil k/e \rceil \right] \leq \frac{1}{n^4},$$

which completes the proof. □