

On Approximating the Number of k -Cliques in Sublinear Time

Talya Eden*
Tel Aviv University
Israel
talyaa01@gmail.com

Dana Ron†
Tel Aviv University
Israel
danaron@tau.ac.il

C. Seshadhri
University of California at Santa Cruz
USA
sesh@ucsc.edu

ABSTRACT

We study the problem of approximating the number of k -cliques in a graph when given query access to the graph. We consider the standard query model for general graphs via (1) degree queries, (2) neighbor queries and (3) pair queries. Let n denote the number of vertices in the graph, m the number of edges, and C_k the number of k -cliques. We design an algorithm that outputs a $(1 + \epsilon)$ -approximation (with high probability) for C_k , whose expected query complexity and running time are $O\left(\frac{n}{C_k^{1/k}} + \frac{m^{k/2}}{C_k}\right) \text{poly}(\log n, 1/\epsilon, k)$. Hence, the complexity of the algorithm is sublinear in the size of the graph for $C_k = \omega(m^{k/2-1})$. Furthermore, the query complexity of our algorithm is essentially optimal (up to the dependence on $\log n$, $1/\epsilon$ and k).

The previous results in this vein are by Feige (SICOMP 06) and by Goldreich and Ron (RSA 08) for edge counting ($k = 2$) and by Eden et al. (FOCS 2015) for triangle counting ($k = 3$). Our result matches the complexities of these results.

The previous result by Eden et al. hinges on a certain amortization technique that works for triangle counting, and does not generalize to all k . We obtain a general algorithm that works for any $k \geq 3$ by designing a procedure that samples each k -clique incident to a given set S of vertices with approximately equal probability. The primary difficulty is in finding cliques incident to purely high-degree vertices, since random sampling within neighbors has a low success probability. This is achieved by an algorithm that samples uniform random high degree vertices and a careful tradeoff between estimating cliques incident purely to high-degree vertices and those that include a low-degree vertex.

CCS CONCEPTS

• **Theory of computation** → **Streaming, sublinear and near linear time algorithms**; *Graph algorithms analysis*; *Approximation algorithms analysis*;

*This research was partially supported by a grant from the Blavatnik fund. The author is grateful to the Azrieli Foundation for the award of an Azrieli Fellowship.

†This research was partially supported by the Israel Science Foundation grant No. 671/13 and by a grant from the Blavatnik fund.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

STOC'18, June 25–29, 2018, Los Angeles, CA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5559-9/18/06...\$15.00

<https://doi.org/10.1145/3188745.3188810>

KEYWORDS

Sublinear algorithms, Approximation algorithms, Counting cliques.

ACM Reference Format:

Talya Eden, Dana Ron, and C. Seshadhri. 2018. On Approximating the Number of k -Cliques in Sublinear Time. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC'18)*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3188745.3188810>

1 INTRODUCTION

Counting the number of k -cliques in a graph is a classic problem in theoretical computer science, and the special case of $k = 3$ (triangle counting) is itself an important problem. In practice, clique counting has received much attention due to its significance for analyzing real-world graphs [4–6, 12, 16, 23, 25, 26, 31, 32, 39, 44, 47, 52]. From a theoretical standpoint, the best exact algorithms use matrix multiplication [22, 40]. Better bounds for sparse graphs can be obtained by combinatorial methods [9, 55].

There is a long line of algorithms for approximately counting the number of cliques (especially triangles) in various computational models, including distributed and streaming settings [2, 3, 9, 11, 34–36, 45, 46, 48–51, 53, 54]. All these algorithms begin by reading their entire input graph, and hence must run in at least linear time. Recently, Eden et al. [17] gave the first sublinear-time algorithm for triangle counting. The query model used is the standard model for sublinear algorithms on general graphs (refer to Chapter 10 of Goldreich’s book [27]). We assume that the vertex set is $V = [n]$. Algorithms can make the following queries. (1) Degree queries: given $v \in V$, get the degree $d(v)$. (2) Neighbor queries, given $v \in V$ and $i \leq d(v)$ get the i^{th} neighbor of v . (3) Pair queries: given vertices u, v , determine if (u, v) is an edge.

We show that there is a sublinear-time algorithm for approximating the number of k -cliques in this model, for any given k , subsuming the result of [17] for $k = 3$.

1.1 Results

Let $G = (V, E)$ be a graph over n vertices and m edges, where we view edges as ordered pairs, so that m equals the sum of the degrees of all vertices. Let C_k denote that number of cliques of size k in G .

Our main theorem follows.

THEOREM 1. *There exists an algorithm that, given n , k , an approximation parameter $0 < \epsilon < 1$, and query access to a graph G , outputs an estimate \hat{C}_k , such that with high constant probability (over the randomness of the algorithm),*

$$(1 - \epsilon) \cdot C_k \leq \hat{C}_k \leq (1 + \epsilon) \cdot C_k.$$

The expected query complexity of the algorithm is

$$O\left(\frac{n}{C_k^{1/k}} + \min\left\{\frac{m^{k/2}}{C_k}, m\right\}\right) \cdot \text{poly}(\log n, 1/\epsilon, k),$$

and the expected running time is $O\left(\frac{n}{C_k^{1/k}} + \frac{m^{k/2}}{C_k}\right) \cdot \text{poly}(\log n, 1/\epsilon, k)$.

Below we state a nearly matching lower bound. The lower bound was first given in an earlier version of this paper [18], and a simplified proof is given in [20].

THEOREM 2 ([18, 20]). *The query complexity of any multiplicative-approximation algorithm for C_k is*

$$\Omega\left(\frac{n}{C_k^{1/k}} + \min\left\{\frac{m^{k/2}}{C_k \cdot (c \cdot k)^k}, m\right\}\right),$$

for a (small) constant c .

1.2 Main Ideas and Techniques

In all that follows, a random vertex refers to a vertex selected uniformly at random. Our starting point for approximating the number of k -cliques is similar to that of Eden et al. [17] (henceforth ELRS) for approximating the number of triangles (i.e., $k = 3$). However, we diverge rather quickly, as the heart of our algorithm, a process for sampling k -cliques, is conceptually and technically different. For the sake of simplicity, assume ϵ is a constant. Also, it is convenient to assume that constant factor estimates of C_k and m are known.¹

The starting point that is shared with ELRS: vertex sampling and clique assignment. Our algorithm starts by uniformly and independently selecting a (multi-)set S of vertices of size roughly $n/C_k^{1/k}$. Denoting the number of k -cliques incident to a vertex u by $c_k(u)$, a natural estimate for C_k is $\frac{n}{|S|} \sum_{u \in S} c_k(u)$. Unfortunately, for a random u , $c_k(u)$ can have extremely large variance. ELRS reduce the variance by considering triangles only from endpoints u where $c_3(u)$ is not too large. In our setting, we formalize this by defining “sociable” vertices. A vertex is sociable if it participates in a number of k -cliques that is above a certain threshold τ or if its degree is above another threshold τ' . A k -clique with vertices $\{u_1, \dots, u_k\}$ is *assigned* to the smallest degree vertex u_j that is not sociable. We set the parameters τ, τ' to ensure that the number of k -cliques that are not assigned to any vertex is at most $\epsilon \cdot C_k$. Let $\alpha(u)$ be the number of k -cliques assigned to u . We can prove that if $|S|$ is roughly $n/C_k^{1/k}$, then $\frac{n}{|S|} \sum_{u \in S} \alpha(u)$ is a $(1 + \epsilon)$ -approximation of C_k with high probability. The problem of approximating C_k is now reduced to approximating $\alpha(S) = \sum_{u \in S} \alpha(u)$. From this point on the ELRS approach does not generalize (as we explain towards the end of this section). We continue by describing our approach.

Approximating $\alpha(S)$: Sampling k -cliques incident to S almost uniformly. Let $\mathcal{A}(S)$ denote the (multi-)set of k -cliques assigned to vertices in S . We estimate $\alpha(S) = |\mathcal{A}(S)|$ by sampling k -cliques incident to S , and checking whether they belong to $\mathcal{A}(S)$. The key is to sample each k -clique incident to S (and in particular in $\mathcal{A}(S)$) with (approximately) the same probability. Specifically, as explained

next, we will do so with probability proportional to $\frac{1}{m(S) \cdot m^{k/2-1}}$, where $m(S) = |E(S)|$ and $E(S)$ denotes the (multi-)set of (ordered) edges incident to S .

Each k -clique containing a vertex $u \in S$ is associated with one of the edges incident to u in the clique. For reasons that will become clear shortly, the other endpoint of this edge is selected to be the lowest degree vertex among the vertices in the clique (excluding u , and breaking ties arbitrarily but consistently). The procedure for sampling k -cliques starts by sampling an edge (u, v) uniformly in $E(S)$. It then attempts to extend this edge to a k -clique, where the other $k - 2$ vertices have degree higher than v . This is done in one of two different ways, depending on the degree of v , where in either way, the algorithm may fail to output any k -clique. For the sake of simplicity, in what follows we refer to a vertex as a *low-degree vertex* if its degree is at most \sqrt{m} , and as a *high-degree vertex* otherwise. (The technical part of the paper uses a slightly different definition.)

The (easy) “low case”: If v is a low-degree vertex, then we sample $k - 2$ random neighbors of v and check whether we obtained a k -clique in which v is the lowest degree vertex other than u . In order to ensure that all k -cliques incident to S are output with the same probability, we apply rejection sampling and keep each sampled neighbor w of v with probability $d(v)/\sqrt{m}$ (and conditioned on $d(v) \leq d(w)$). Hence, each such clique (i.e., in which the lowest degree vertex other than $u \in S$ is a low-degree vertex) is output with equal probability $\frac{(k-2)!}{m(S) \cdot \sqrt{m}^{k-2}}$. (The $(k - 2)!$ factor is due to the fact that we may obtain the $k - 2$ vertices in the clique other than u and v in any order.)

The (hard) “high case”: The challenging case is when v is a high-degree vertex. Rejection sampling, as in the low case, is too expensive now. However, observe that we are interested only in sampling neighbors of v with degree higher than v , and that the number of vertices with high degree is at most \sqrt{m} . Therefore, if we had a way to efficiently sample each high-degree vertex with probability (approximately) $1/\sqrt{m}$, we would obtain the same probability over cliques incident to S as in the low case. We next explain how this can be done.

Consider selecting a random multi-set T of roughly $t = \frac{n}{\log n / \sqrt{m}}$ vertices. The setting of t is such that with high probability, for every high-degree vertex w , the number of neighbors that w has in T is close to its expected value, that is, $d(w) \cdot \frac{t}{n}$. This implies that if we select an edge (x, y) uniformly at random in $E(T)$ (the (ordered) edges incident to T , whose number is $m(T)$), then the probability that $y = w$ for a fixed high-degree vertex w , is approximately $\frac{d(w) \cdot (t/n)}{m(T)}$.

Assume that $m(T)$ is not much larger than its expected value, $m \cdot \frac{t}{n}$ (which can be ensured with high probability). Let $p(w) = \frac{m(T)}{d(w) \cdot (t/n) \cdot \sqrt{m}}$, so that under this assumption on $m(T)$, and the fact that w is a high-degree vertex, $p \in (0, 1]$. If we now keep w with probability $p(w)$, then we have a subroutine that samples each high-degree vertex with approximately equal probability $1/\sqrt{m}$. This in turn implies that we can select any fixed subset of $k - 2$ high-degree vertices with probability very close to $\frac{(k-2)!}{\sqrt{m}^{k-2}}$. We then

¹The assumption on m can be removed by applying [24], and the assumption on C_k can be removed by a geometric search (for details see Subsection 3.7).

check whether whether the chosen vertices form a clique together with u and v (and that the clique is associated with (u, v)).

We now have a procedure that outputs each clique incident to S with (roughly) the same probability, $\frac{(k-2)!}{m(S) \cdot \sqrt{m}^{k-2}}$. In the next paragraph, we discuss a procedure for deciding whether a k -clique belongs to $\mathcal{A}(S)$. Given this decision procedure, we can estimate $|\mathcal{A}(S)|$ by performing $\frac{m(S) \cdot (2\sqrt{m})^{k-2}}{(k-2)! \cdot |\mathcal{A}(S)|} = O\left(\frac{m^{k/2}}{C_k}\right)$ calls to the k -clique sampling procedure. (We assume that $|\mathcal{A}(S)|$ is close to its expected value, $C_k \cdot \frac{S}{n}$ and that $m(S)$ is not much larger than its expected value, $m \cdot \frac{S}{n}$).

Deciding whether a k -clique is in $\mathcal{A}(S)$. Deciding whether a k -clique incident to a vertex $u \in S$ should be assigned to u , requires to determine which of the clique vertices are sociable. Recall that a vertex is considered sociable if the number of k -cliques it participates in is more than τ or if its degree is above τ' (for appropriate settings of τ and τ'). The second condition can be easily verified by a single degree query. As for the first condition, given a vertex u , we verify whether the number of k -cliques that it participates in is more than τ by running the k -clique sampling procedure with $S = \{u\}$ for a sufficient number of times (roughly $\frac{d(u) \cdot (2\sqrt{m})^{k-2}}{(k-2)! \cdot \tau} \leq \frac{\tau' \cdot (2\sqrt{m})^{k-2}}{(k-2)! \cdot \tau}$ where $\frac{\tau'}{\tau} = O\left(\frac{m}{C_k}\right)$). The procedure may err on “almost sociable” vertices, but the analysis can be modified to deal with this. While the procedure might have high query complexity, it is only invoked when the clique-sampling procedure returns a clique. The frequency of the latter can be bounded appropriately to get the final query complexity.

Why the ELRS algorithm does not generalize. The success of ELRS hinges on the following bound: $\sum_{(u,v) \in E} \min(d(u), d(v)) = O(m^{3/2})$, discovered in [9] in the context of triangle enumeration. The short answer to why the ELRS algorithm does not generalize is that the above bound does not have analogues for $k > 3$. Indeed, this is why the simple algorithm of [9] for triangle enumeration does not work for cliques of larger size.²

Let us revisit the triangle estimator of ELRS. Recall that in this context, $\alpha(u)$ denotes the number of triangles assigned to a vertex u , and that the aim is to estimate the average value of $\alpha(u)$ for $u \in S$. ELRS first “transfer” the assignment of triangles from vertices to edges. Letting $\alpha(u, v)$ denote the number of triangles assigned to the edge (u, v) we have that $\alpha(S) = \sum_{(u,v) \in E(S)} \alpha(u, v)$. For a random edge (u, v) , a triangle can be detected by sampling a random neighbor of the lower degree vertex among u and v and performing a pair query with the other endpoint. Since the probability of finding a triangle decreases as $\min(d(u), d(v))$ increases, ELRS select $\left\lceil \frac{\min(d(u), d(v))}{\sqrt{m}} \right\rceil$ random neighbors. Hence, the expected number of queries performed per edge is $\frac{1}{m} \sum_{(u,v) \in E} \left\lceil \frac{\min(d(u), d(v))}{\sqrt{m}} \right\rceil$, which by the aforementioned bound from [9], is $O(1)$. ELRS prove that in order to estimate $\alpha(S)$, it suffices to sample $O(m^{3/2}/C_3)$ random edges from $E(S)$, and therefore they get the desired bound.

A generalization of ELRS to k -cliques would require the following bound: $\sum_{(u,v) \in E} \min(d(u), d(v))^{k-2} = O(m^{k/2})$. This bound is

false for $k > 4$. To exemplify this, consider a graph over the vertex set $\{1, \dots, n\}$, with the following edges. There is an edge $(1, 2)$, and both vertices 1 and 2 have an edge to all other vertices. The left-hand-side of the bound is $\Theta(n^{k-2})$, while the right-hand-side is $O(n^{k/2})$. Thus, the ELRS analysis depends on a seeming singularity for $k = 3, 4$, and does not generalize to all k .

1.3 Related Work

A significant portion of the work on clique counting focuses on triangle counting. Because our focus is on general k , we avoid a detailed discussion of results for triangle counting. We point the interested reader to [17].

Nešetřil and Poljak give the first non-trivial algorithm for k -clique counting by reducing to matrix multiplication [40]. Specifically, their algorithm runs in time $O(n^{\omega \lfloor k/3 \rfloor + k \pmod{3}})$, where ω is the matrix multiplication exponent. Eisenbrand and Grandoni refine this bound for certain values of k by careful reductions to rectangular matrix multiplication [22]. They also give better dependencies on m for sparse graphs. The general dependence of the form $n^{\omega k/3}$ is believed to be optimal. Recent work by Abboud et al. builds on this conjecture to prove hardness for various parsing algorithms [1]. More relevant to our work, Chiba and Nishizeki give an algorithm for k -clique enumeration, based on the arboricity of the graph, from which the $O(n + m^{k/2})$ bound for general graphs follows immediately. The use of degree/degeneracy orientations have appeared in recent practical works on clique counting [25, 33]. In the current work we design various primitives to sample random k -cliques, by either extending smaller cliques or by sampling high degree vertices. The idea of extending smaller cliques to large ones using degree orientations is an important feature of previous practical approaches [25, 33]. It would be of interest to see if the new techniques given by our result could be used for practical algorithms.

In the context of sublinear algorithms, our work follows a line of results on sublinear estimation of subgraph counts. Our analysis builds on several techniques developed in these results. The starting point is the average degree estimation results of Feige [24] and Goldreich and Ron [28]. Gonen et al. generalize these techniques to estimate the count of k -stars in sublinear time [29]. Eden et al. [19] further extended and simplified all these results, and show connections between this problem and the graph degeneracy. They also build on the basic ELRS framework. Eden and Rosenbaum [21] provide an algorithm for sampling edges almost uniformly, and our clique sampler uses some of their ideas to sample high-degree vertices. Dasgupta et al. [15] and Chierichetti et al. [10] consider sublinear algorithms (for average degree and related problems) in a weaker model where uniform random vertices are not allowed. In practical settings, we can only “crawl” a graph, which translates to performing random walks. Their results typically require some assumption about the mixing time of the input graph G . Again, we believe this is an interesting direction for future work, to consider weaker query models but stronger assumptions on graph structure.

Other work on sublinear algorithms for estimating graph parameters (in the standard query model) include results on the minimum weight spanning tree [7, 13, 14], maximum matching [41, 58] and minimum vertex cover [30, 37, 41–43, 58].

²We note that [9] also present a general algorithm for listing k -cliques for any k , but this algorithm and its analysis are more involved.

2 PRELIMINARIES

We consider simple undirected graphs over a set V of n vertices. It is convenient to think of the graph edges as ordered pairs, so that every edge is considered from both endpoints. We say that the ordered edge (u, v) originates from the vertex u . We denote the set of all ordered edges by E and let $m \triangleq |E|$. We use the following notations.

- $d(u)$: the degree of a vertex u (the number of edges originating from u). Note that $\sum_{u \in V} d(u) = m$.
- $E(S), m(S)$: $E(S) \triangleq \{(u, v) \mid u \in S\}$ and $m(S) \triangleq |E(S)| = \sum_{u \in S} d(u)$.
- $d_S(u)$: for any vertex u and set of vertices S , $d_S(u)$ is the number of neighbors of u in S .
- $C_k, c_k(u)$: C_k is the number of k -cliques in the given graph. For $u \in V$, $c_k(u)$ is the number of k -cliques that u participates in. Note that $C_k = \frac{1}{k} \cdot \sum_{u \in V} c_k(u)$.

We use $<$ to denote a total order over the graph vertices such that for every two vertices u and v , if $d(u) < d(v)$, then $u < v$, and if $d(u) = d(v)$, then the order between u and v is determined in an arbitrary but fixed manner (e.g., by vertex id).

Let $[r] \triangleq \{1, \dots, r\}$ and let $(1 \pm \alpha)^t \cdot x$ denote the interval $[(1 - \alpha)^t \cdot x, (1 + \alpha)^t \cdot x]$.

We make use of the following version of Chernoff's inequality [8]. Let χ_i for $i = 1, \dots, m$ be random variables taking values in $[0, B]$, such that for every i , $\text{Exp}[\chi_i] = p$. Then

$$\Pr \left[\frac{1}{m} \sum_{i=1}^m \chi_i > (1 + \gamma)p \right] < \exp \left(-\frac{\gamma^2 pm}{3B} \right) \quad \text{and} \\ \Pr \left[\frac{1}{m} \sum_{i=1}^m \chi_i < (1 - \gamma)p \right] < \exp \left(-\frac{\gamma^2 pm}{2B} \right).$$

The proof of the following claim is similar to the proof of [9] for their exact clique enumeration algorithm (and we include it here for the sake of completeness).

CLAIM 3. *For every graph G with m (ordered) edges and C_k k -cliques, $C_k \leq m \cdot \binom{\sqrt{m}}{k-2}$.*

PROOF. Let D be the DAG obtained by orienting edges in G according to $<$. Let $d^+(v)$ be the out-degree of vertex v in D . Observe that $\max_v \{d^+(v)\} \leq \sqrt{m}$. (All $d^+(v)$ out-neighbors of v have degree at least $d(v) \geq d^+(v)$. Thus, $d^+(v) \leq \sqrt{m}$.) The number of k -cliques where v is the lowest vertex according to $<$ is at most $\binom{d^+(v)}{k-1}$. Thus, $C_k \leq \sum_v \binom{d^+(v)}{k-1} \leq \binom{\sqrt{m}}{k-2} \sum_v d^+(v) = m \binom{\sqrt{m}}{k-2}$. \square

3 THE ALGORITHM

The main algorithm for approximating the number of k -cliques is presented in Subsection 3.2 and named Approximate-Cliques. It takes the following parameters.

- \bar{m} : This is assumed to be a fairly precise estimate of the number of (ordered) edges m , and can be obtained using [28] (in expected time $O\left(\frac{n}{\sqrt{m}}\right) \cdot \text{poly}(\log n, 1/\epsilon)$).
- \bar{C}_k : This is assumed to be a constant-factor estimate of C_k , which is obtained by geometric search (as shown in Subsection 3.7).
- ϵ : The main approximation parameter. We set $\bar{\epsilon} = \epsilon/5$.
- δ : The failure parameter. We set $\bar{\delta} = \delta/4$.

3.1 Sociable Vertices and the Assignment of Cliques to Vertices

The notion of *sociable* vertices, defined next, is critical in reducing the variance of the output of our algorithm.

DEFINITION 4 (SOCIABLE AND SHY VERTICES). *We say that a vertex u is sociable if $c_k(u) > k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k}$ or if $d(u) > 4\bar{m} / (\bar{\epsilon}\bar{C}_k)^{1/k}$. If $c_k(u) \leq \frac{1}{4}k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k}$ and $d(u) \leq 4\bar{m} / (\bar{\epsilon}\bar{C}_k)^{1/k}$, then we say that u is shy.*

Note that a vertex u may be neither sociable nor shy. This is the case if $d(u) \leq 4\bar{m} / (\bar{\epsilon}\bar{C}_k)^{1/k}$ and $\frac{1}{4}k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k} < c_k(u) \leq k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k}$.

The following claim, whose proof follows directly from Definition 4, shows that we can ignore cliques that do not contain shy vertices.

CLAIM 5. *If $\bar{m} \in [(1 - \bar{\epsilon})m, m]$ and $\bar{C}_k > C_k/4$, then at most $\bar{\epsilon}C_k$ k -cliques consist solely of vertices that are not shy.*

PROOF. For every vertex u that is not shy, either

$$c_k(u) > \frac{1}{4}k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k} \quad \text{or} \quad d(u) > 4\bar{m} / (\bar{\epsilon}\bar{C}_k)^{1/k}.$$

Recall that $\sum_{u \in V} c_k(u) = k \cdot C_k$. Therefore, if $\bar{C}_k \geq C_k/4$, then there are at most

$$\frac{k \cdot C_k}{\frac{1}{4}k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k}} \leq (\bar{\epsilon}C_k)^{1/k} / 2$$

vertices of the former type, and if $\bar{m} \geq (1 - \epsilon)m$, then there are at most

$$\frac{m}{4\bar{m} / (\bar{\epsilon}\bar{C}_k)^{1/k}} \leq (\bar{\epsilon}C_k)^{1/k} / 2$$

of the latter type. This implies that there are at most $(\bar{\epsilon}C_k)^{1/k}$ vertices that are not shy, and it follows that the number of k -cliques for which all of their vertices are not shy is at most $\bar{\epsilon}C_k$. \square

DEFINITION 6 (AN APPROPRIATE PARTITION). *We say that a partition $P = (V_0, V_1)$ of V is appropriate if every shy vertex (as defined in Definition 4) is in V_0 and every sociable vertex is in V_1 (and any other vertex can be either in V_0 or V_1).*

We next specify the assignment of cliques to vertices.

DEFINITION 7 (ASSIGNING CLIQUES). *Fix a partition $P = (V_0, V_1)$.*

- *Assignment of cliques:* We assign each k -clique $K = \{u_1, \dots, u_k\}$ to the vertex u_i that is the first (according to $<$) vertex of K in V_0 . If all of K 's vertices are in V_1 , then K is not assigned to any vertex.
- $\alpha_P(u), \alpha_P(S)$: We denote the number of k -cliques assigned to u (for this P) by $\alpha_P(u)$. For a set S of vertices, $\alpha_P(S) = \sum_{u \in S} \alpha_P(u)$.

The following is a corollary of Claim 5, Definition 6 and Definition 7.

COROLLARY 8. *For every partition $P = (V_0, V_1)$ of V it holds that $\alpha_P(V) \leq C_k$. Furthermore, if $P = (V_0, V_1)$ is appropriate, $\bar{m} \geq (1 - \epsilon)m$ and $\bar{C}_k \geq C_k/4$, then $\alpha_P(V) \in [(1 - \bar{\epsilon})C_k, C_k]$.*

Another distinction between types of vertices that will play a central role in our analysis is the following.

DEFINITION 9 (HIGH-DEGREE AND LOW-DEGREE VERTICES). We say that a vertex u is a *high-degree vertex* if $d(u) > 2\sqrt{\bar{m}}$ and otherwise we say it is a *low-degree vertex*.

3.2 The Main Algorithm and the Procedures It Uses

In this subsection we present our main algorithm and the corresponding main theorem. Our algorithm invokes several procedures, which are provided in the following subsections. Here we shortly describe all procedures and state the main claim regarding each of them. Building on these claims we give a proof sketch of the main theorem (the complete proof appears in Subsection 3.6).

Approximate-Cliques. This is the main algorithm, and it is provided in Figure 1. The algorithm begins by constructing two random multisets, S and T . The multiset S is obtained by simply selecting vertices uniformly (independently) at random. The multiset T is constructed by a procedure *Sample-Degrees-Typical*. We show that with high probability, S and T have certain desired properties (where the correctness of subsequent steps of the algorithm relies on these properties).

In Step 5, the algorithm calls two procedures: *Sample-a-Clique* and *Is-sociable*. The heart of the algorithm is the procedure *Sample-a-Clique* that either returns a k -clique $\{u, v, w_1, \dots, w_{k-2}\}$ where $u \in S$ or returns *fail*. The procedure *Is-sociable* distinguishes between sociable and shy vertices (as defined in Definition 4). It is used in order to decide for each k -clique that is output in the previous step, whether it is assigned to u (as defined in Definition 7).

Note that if the sample size s (defined in Step 2) is larger than n , then the algorithm can simply set $S = V$. Similarly, if q (the number of iterations in Step 5) is larger than \bar{m} , then the algorithm can query upfront all edges incident to S and their neighbors so that it never performs more than $\min\{m, \bar{m}\}$ queries. (If it views more than \bar{m} edges, then it can abort.) Finally, we may assume that $\epsilon > 1/\bar{m}^{k/2}$, since otherwise we are required to output the exact number of k -cliques in the graph (recall that by Claim 3, $C_k < m^{k/2}$), and thus can simply invoke the exact enumeration algorithm of [9].

The main theorem of our paper is the following (where the second item in the theorem is used by the geometric search algorithm for C_k).

THEOREM 10. Consider an invocation of Algorithm *Approximate-Cliques*($n, k, \bar{m}, \bar{C}_k, \epsilon, \delta$).

- (1) If $\bar{m} \in [(1 - \bar{\epsilon})m, m]$ and $\bar{C}_k \in [C_k/4, C_k]$, then with probability at least $1 - \delta$, *Approximate-Cliques* returns a value \hat{C}_k such that $\hat{C}_k \in (1 \pm \epsilon) \cdot C_k$.
- (2) If $\bar{m} \in [(1 - \bar{\epsilon})m, m]$ and $\bar{C}_k > C_k$, then with probability at least $\epsilon/4$, *Approximate-Cliques* returns a value \hat{C}_k such that $\hat{C}_k \leq (1 + \epsilon) \cdot C_k$.
- (3) If $\bar{C}_k \leq \bar{m}^{k/2}$, then the expected query complexity and running time of *Approximate-Cliques* are $O\left(\frac{n}{\bar{C}_k^{1/k}} + \frac{\max\{m, \bar{m}\} \cdot \bar{m}^{(k-2)/2}}{\bar{C}_k} \cdot \frac{C_k}{\bar{C}_k}\right) \cdot \text{poly}(\log n, 1/\epsilon, \log(1/\delta), k)$.³ The number of queries is always

³In the second additive term there is actually a dependence on $2^k/(k-2)!$, which we ignored for the sake of simplicity.

Approximate-Cliques ($n, k, \bar{m}, \bar{C}_k, \epsilon, \delta$)

- (1) Let $\bar{\epsilon} = \epsilon/5$ and $\bar{\delta} = \delta/4$.
- (2) Let S be a multiset of $s = \frac{700 \cdot k \cdot n \cdot \ln(1/\bar{\delta})}{\bar{\epsilon}^{2+1/k} \cdot \bar{C}_k^{1/k}}$ vertices chosen uniformly at random.
- (3) Query the degree of each vertex in S and set up a data structure $D(S)$ that supports sampling a uniform edge in $E(S)$ in constant time.
- (4) Invoke *Sample-Degrees-Typical*($n, k, \bar{m}, \bar{C}_k, \bar{\epsilon}, \bar{\delta}$). If the procedure returned *fail*, then **return fail**. Otherwise, let $(T, m(T), D(T))$ be its output.
- (5) For $i = 1$ to $q = \frac{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}{(1-\bar{\epsilon})^3 \cdot (k-2)! \cdot \bar{C}_k \cdot (s/n)} \cdot \frac{10 \ln(1/\bar{\delta})}{\bar{\epsilon}^2}$ do:
 - (a) Invoke *Sample-a-Clique*($S, T, m(S), m(T), D(S), D(T), k, \bar{m}$).
 - (b) If the procedure returned *fail* then set $\chi_i = 0$. Otherwise, let $K_i = (u_i, v_i, w_{i,1}, \dots, w_{i,k-2})$ be the k -tuple returned and do the following.
 - (i) Query the degree and invoke *Is-sociable*($x, T, m(T), D(T), \bar{m}, \bar{C}_k, n, k, \bar{\epsilon}, \bar{\delta}$) on each vertex $x \in K_i$.
 - (ii) If u_i is the first vertex (according to $<$) in K_i for which *Is-sociable* returned *shy*, then set $\chi_i = 1$. Otherwise, set $\chi_i = 0$.
- (6) **Return** $\hat{C}_k = \frac{m(S)(2\sqrt{\bar{m}})^{k-2}}{(k-2)! \cdot (s/n)} \cdot \frac{1}{q} \sum_{i=1}^q \chi_i$.

Figure 1: The main algorithm for computing a $(1 \pm \epsilon)$ -estimate of the number of k -cliques in a graph (given a constant factor estimate of this number).

upper bounded by $O\left(\frac{n}{\bar{C}_k^{1/k}}\right) \cdot \text{poly}(\log n, 1/\epsilon, \log(1/\delta), k) + \min\{m, \bar{m}\}$.

Sample-Degrees-Typical. This procedure is described in Figure 4. Its goal is to output a degrees-typical multiset, which is defined below. The procedure itself is quite simple; it repeats the process of sampling a uniform random multiset for a sufficient number of times in order to achieve this condition with high probability.

DEFINITION 11. We say that a multiset T of size t is *degrees-typical* if $m(T) \leq \frac{t}{n} \cdot 4\bar{m}$ and for every high-degree vertex $w \in V$,

$$d_T(w) \in \left(1 \pm \frac{\bar{\epsilon}}{k}\right) \cdot \frac{t}{n} \cdot d(w).$$

In Subsection 3.5 we prove the following lemma regarding the correctness and running time of the procedure *Sample-Degrees-Typical*.

LEMMA 12. Consider an invocation of *Sample-Degrees-Typical*($n, k, \bar{m}, \bar{C}_k, \bar{\epsilon}, \bar{\delta}$). The procedure either returns *fail*, or returns a multiset T together with $m(T)$ and a data structure $D(T)$ that supports selecting a uniform edge in $E(T)$ in time $O(1)$.

Let $\gamma = \min\{1/(4\bar{m}^{k/2}), \bar{\delta}\}$. If $\bar{m} \in [(1 - \bar{\epsilon})m, m]$, then with probability at least $1 - \gamma$, the procedure returns a triple $(T, m(T), D(T))$ such that the multiset T is degrees-typical.

The running time of *Sample-Degrees-Typical* is $O\left(\frac{n}{\sqrt{\bar{m}}} \cdot \frac{k^2 \cdot \log(n/\gamma) \cdot \log(1/\gamma)}{\bar{\epsilon}^2}\right)$.

Sample-a-Clique. As mentioned earlier, this is the most important and novel aspect of our algorithm. Given any multiset of vertices S , the procedure Sample-a-Clique produces cliques incident to S with roughly uniform probability. The procedure is given in Figure 2.

DEFINITION 13. Let $C(S)$ denote the set of k -tuples $(u, v, w_1, \dots, w_{k-2})$ that have the following properties: (1) the subgraph induced by $\{u, v, w_1, \dots, w_{k-2}\}$ is a k -clique; (2) $u \in S$; (3) $v < w_j$ for every $j \in [k-2]$.

By the above definition, each clique containing a vertex $u \in S$ is associated with the edge (u, v) in the clique such that $v < w$ for every other vertex $w \neq u$ in the clique, and the clique has exactly $(k-2)!$ corresponding tuples in $C(S)$. In Subsection 3.3 we prove the following lemma regarding the correctness and running time of the procedure Sample-a-Clique.

LEMMA 14. Let T be a degrees-typical multiset and let S be any multiset. For any fixed k -tuple $K \in C(S)$, the probability that an invocation of Sample-a-Clique($S, T, m(S), m(T), D(S), D(T), k, \bar{m}$) returns K is in $(1 \pm \bar{\epsilon}) \cdot \frac{1}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}$. The running time of Sample-a-Clique is $O(k^2)$.

Is-sociable. This procedure decides if a vertex u is sociable or not. This is done by multiple independent invocations of Sample-a-Clique for the set $S = \{u\}$. The procedure is provided in Figure 3 in Subsection 3.4, where we also prove the following lemma regarding the correctness and running time of the procedure Is-sociable.

LEMMA 15. Let T be a degrees-typical multiset. If *Is-sociable*($u, T, m(T), D(T), \bar{m}, \bar{C}_k, n, k, \bar{\epsilon}, \bar{\delta}$) is invoked with a sociable vertex u , then with probability at least $1 - \bar{\delta}/n$, the procedure returns sociable, and if u is shy, then with probability at least $1 - \bar{\delta}/n$, the procedure returns shy. The running time of *Is-sociable* is $O\left(\frac{\bar{m}^{k/2}}{\bar{C}_k} \cdot \frac{k \cdot 2^k \cdot \log(n/\bar{\delta})}{(k-2)! \cdot \bar{\epsilon}^{2-1/k}}\right)$.

Proof sketch of the first item in Theorem 10. The full proof of Theorem 10 appears in Subsection 3.6. Here we provide a proof sketch for the case that $\bar{C}_k \in [C_k/4, C_k]$, relying on Lemmas 12–15.

By the first premise of this item, $\bar{m} \in [(1 - \bar{\epsilon})m, m]$. Hence, by Lemma 12, with high probability the multiset T is degrees-typical. From this point on we condition on this event.

Consider (as a thought experiment) invoking *Is-sociable* on all vertices with the degrees-typical multiset T . Based on these invocations, we can define a partition $P_{is} = (V_0, V_1)$, where V_0 contains all vertices for which *Is-sociable* returns shy and V_1 contains all vertices for which *Is-sociable* returns sociable. By Lemma 15 and a union bound over all vertices, we get that with probability at least $1 - \bar{\delta}$, P_{is} is an appropriate partition (as defined in Definition 6). Conditioned on P_{is} being appropriate (and using our assumptions on \bar{m} and \bar{C}_k), by Corollary 8 we have that $\alpha_{P_{is}}(V) \in [(1 - \bar{\epsilon})C_k, C_k]$.

Now consider the selection of the multiset S . We show that the size s of this sample ensures that with high probability $\alpha_{P_{is}}(S)$ is close to its expected value, $\frac{s}{n} \cdot \alpha_{P_{is}}(V)$, so that $\alpha_{P_{is}}(S) \in (1 \pm \bar{\epsilon})^2 \cdot \frac{s}{n} \cdot C_k$. We condition on this event as well. Since T is degree-typical, by Lemma 14, whenever we invoke Sample-a-Clique (in

Step 5a) it returns each k -tuple in $C(S)$ with probability approximately $\frac{1}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}$. Observe that if Sample-a-Clique returns a k -tuple $K_i = (u_i, v_i, w_{i,1}, \dots, w_{i,k-2})$, then in Steps 5(b)i and 5(b)ii the algorithm determines whether the corresponding k -clique is assigned to u_i according to P_{is} and sets χ_i to 1. Therefore, $\Pr[\chi_i = 1]$ is approximately $\frac{(k-2)! \cdot \alpha_{P_{is}}(S)}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}$ (recall that for each k -clique that a vertex $u \in S$ participates in, there are $(k-2)!$ corresponding k -tuples in $C(S)$). By the setting of the number of invocations, q , of Sample-a-Clique, with high probability, the sum of the χ_i 's is close to its expected value, and the output of the algorithm is as claimed.

3.3 Sampling a Clique

In this subsection we provide the procedure Sample-a-Clique and prove Lemma 14. The procedure first samples a uniform edge (u, v) in $E(S)$. It then tries to construct a k -clique that this edge participates in by selecting $k-2$ additional vertices. More precisely, it tries to construct a k -tuple $(u, v, w_1, \dots, w_{k-2})$ in $C(S)$. Recall that such a k -tuple satisfies: $u \in S$, $\{u, v, w_1, \dots, w_{k-2}\}$ induces a k -clique, and $v < w_j$ for each $j \in [k-2]$. To this end the procedure repeats the following $k-2$ times. If v is a low-degree vertex, then it selects a uniform neighbor w of v , and maintains it with probability $d(v)/2\sqrt{\bar{m}}$. If v is a high-degree vertex, then the procedure tries to sample a random high-degree vertex. It does so by first sampling a uniform edge $(x, y) \in E(T)$ and if y is a high-degree neighbor of v , performing rejection sampling according to the degree of y . We prove that conditioned on T being degrees-typical, we obtain each k -tuple in $C(S)$ with almost equal probability.

Sample-a-Clique ($S, T, m(S), m(T), D(S), D(T), k, \bar{m}$)

- (1) Sample a uniform edge $e = (u, v)$ in $E(S)$ (using the data structure $D(S)$).
- (2) For $j = 1$ to $k-2$ do:
 - (a) If $d(v) \leq 2\sqrt{\bar{m}}$ (v is a low-degree vertex), then:
 - (i) Uniformly select a neighbor w_j of v .
 - (ii) Keep w_j with probability $\frac{d(v)}{2\sqrt{\bar{m}}}$ and with probability $1 - \frac{d(v)}{2\sqrt{\bar{m}}}$ **return fail**.
 - (b) Else (v is a high-degree vertex):
 - (i) Sample a random edge (x, y) in $E(T)$ (using the data structure $D(T)$).
 - (ii) If $d(y) \leq 2\sqrt{\bar{m}}$, **return fail**.
 - (iii) With probability $\frac{m(T)}{d(y) \cdot \frac{t}{n} \cdot 2\sqrt{\bar{m}}}$ set $w_j = y$, otherwise, **return fail**.
- (3) For every pair of vertices in $\{u, v, w_1, \dots, w_{k-2}\}$ query if there is an edge between the two vertices.
- (4) If the subgraph induced by $\{u, v, w_1, \dots, w_{k-2}\}$ is a k -clique and $v < w_j$ for every $j \in [k-2]$ then **return** $K = (u, v, w_1, \dots, w_{k-2})$. Otherwise **return fail**.

Figure 2: A procedure for sampling a k -clique incident to S with almost uniform probability.

PROOF OF LEMMA 14: Let $(a, b, z_1, \dots, z_{k-2})$ be a k -tuple in $C(S)$. Recall that by the definition of $C(S)$ we have that $a \in S$, the subgraph induced by $\{a, b, z_1, \dots, z_{k-2}\}$ is a k -clique and $b < z_j$ for every $j \in [k-2]$. If Sample-a-Clique does not return fail, then its output is a tuple $(u, v, w_1, \dots, w_{k-2})$ in $C(S)$. The probability that the procedure returns the tuple (a, b, z_1, \dots, z_k) is

$$\Pr\left[(u, v) = (a, b) \text{ and } \forall j \in [k-2], w_j = z_j\right] \\ = \Pr[(u, v) = (a, b)] \cdot \Pr\left[\forall j \in [k-2], w_j = z_j \mid (u, v) = (a, b)\right].$$

Clearly, $\Pr[(u, v) = (a, b)] = \frac{1}{m(S)}$, so it remains to compute the probability that $w_j = z_j$ for each $j \in [k-2]$, conditioned on $(u, v) = (a, b)$.

If $b = v$ is a low-degree vertex, then the vertices w_1, \dots, w_{k-2} are uniformly selected random neighbors of v . For each $j \in [k-2]$, the probability that $w_j = z_j$ and that the procedure did not return fail due to rejection sampling is $\frac{1}{d(v)} \cdot \frac{d(v)}{2\sqrt{m}} = \frac{1}{2\sqrt{m}}$. Therefore, $\Pr[\forall j \in [k-2], w_j = z_j \mid (u, v) = (a, b)] = 1/(2\sqrt{m})^{k-2}$.

Otherwise (b is a high-degree vertex), since $b < z_j$ for each $j \in [k-2]$, the vertices z_1, \dots, z_{k-2} are also high-degree vertices. In this case (conditioned on $(u, v) = (a, b)$), the procedure tries to sample $k-2$ high-degree vertices by sampling edges originating from the vertices of T . We next prove that, in Step 2(b)iii, any specific high-degree vertex is sampled with probability in $(1 \pm \frac{\bar{\epsilon}}{k}) \cdot \frac{1}{2\sqrt{m}}$.

Since T is degrees-typical (as defined in Definition 11), $m(T) \leq \frac{t}{n} \cdot 4\bar{m}$. This implies that for every high-degree vertex z ,

$$\frac{m(T)}{d(z) \cdot \frac{t}{n} \cdot 2\sqrt{m}} \leq \frac{\frac{t}{n} \cdot 4\bar{m}}{2\sqrt{m} \cdot \frac{t}{n} \cdot 2\sqrt{m}} = 1.$$

Thus, Step 4 is valid. Since T is degrees-typical, for every high-degree vertex z , $d_T(z) \in (1 \pm \frac{\bar{\epsilon}}{k}) \cdot \frac{t}{n} \cdot d(z)$.

For any vertex y , the probability of obtaining an edge in Step 2(b)i with y as an endpoint is $d_T(y)/m(T)$. Therefore, for each $j \in [k-2]$, the probability that $w_j = z_j$ in Step 4 is $\frac{d_T(z_j)}{m(T)} \cdot \frac{m(T)}{d(z_j) \cdot \frac{t}{n} \cdot 2\sqrt{m}} \in (1 \pm \frac{\bar{\epsilon}}{k}) \cdot \frac{1}{2\sqrt{m}}$. It follows that the probability that the procedure returns any specific k -tuple in $C(S)$ is $(1 \pm \bar{\epsilon}) \cdot \frac{1}{m(S)(2\sqrt{m})^{k-2}}$.

It remains to bound the running time of the procedure. Given the data structure $D(S)$, it takes time $O(1)$ to sample an edge in $E(S)$, and similarly it takes time $O(1)$ to sample an edge in $D(T)$. The procedure samples a single edge in $E(S)$ and possibly $k-2$ edges in $E(T)$. Adding the time to perform queries on all pairs of vertices in $\{u, v, w_1, \dots, w_{k-2}\}$ (in addition to a degree query on each of these vertices), the total running time is $O(k^2)$. \square

3.4 Is-Sociable

The procedure Is-sociable determines (with high success probability) whether a given vertex u is sociable or shy. For vertices that are neither sociable nor shy, it can answer arbitrarily. Recall that by Definition 4, the distinction between a sociable u and a shy u involves bounds on both $d(u)$ and $c_k(u)$. These will be critical in bounding the running time of Is-sociable. The procedure basically invokes Sample-a-Clique repeatedly to check if $c_k(u)$ is larger than the specified threshold.

Is-Sociable $(u, T, m(T), D(T), \bar{m}, \bar{C}_k, n, k, \bar{\epsilon}, \bar{\delta})$

- (1) Query the degree of u and if $d(u) > 4\bar{m}/(\bar{\epsilon}\bar{C}_k)^{1/k}$ then **return** sociable.
- (2) For $i = 1$ to $r = \frac{d(u) \cdot (2\sqrt{m})^{k-2}}{(k-2)! \cdot (k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k})} \cdot \frac{15 \ln(n/\bar{\delta})}{\bar{\epsilon}^2}$ do:
 - (a) Invoke Sample-a-Clique($\{u\}, T, d(u), m(T), D(\{u\}), D(T), k, \bar{m}$).
 - (b) If a k -tuple (corresponding to a k -clique) was returned, then set $\chi_i = 1$. Otherwise (the procedure returned fail), set $\chi_i = 0$.
- (3) Let $\hat{c}_k(u) = \frac{d(u) \cdot 2\sqrt{m}}{r} \cdot \sum_{i=1}^r \chi_i$.
- (4) If $\hat{c}_k(u) \geq \frac{1}{2} \cdot (k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k})$, then **return** sociable, otherwise, **return** shy.

Figure 3: A procedure for determining with high probability whether a given vertex is sociable or not.

PROOF OF LEMMA 15: Consider any fixed vertex u . Let $\chi = \frac{1}{r} \sum_{i=1}^r \chi_i$, where χ_1, \dots, χ_r are as defined in Step 2b of Is-sociable. Note that $S = \{u\}$, $m(S) = m(\{u\}) = d(u)$ and $|C(S)| = C(\{u\}) = (k-2)! \cdot c_k(u)$. By Lemma 14 and the assumption that T is degrees-typical, $\text{Exp}[\chi] \in (1 \pm \bar{\epsilon}) \cdot \frac{(k-2)! \cdot c_k(u)}{d(u) \cdot (2\sqrt{m})^{k-2}}$.

First consider the case that u is sociable. By Definition 4, either $d(u) > 4\bar{m}/(\bar{\epsilon}\bar{C}_k)^{1/k}$ or $c_k(u) \geq k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k}$ (or both). Clearly, if $d(u) > 4\bar{m}/(\bar{\epsilon}\bar{C}_k)^{1/k}$, then the procedure returns sociable. Hence, assume that $c_k(u) \geq k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k}$. By Chernoff's inequality and the setting of

$$r = \frac{d(u) \cdot (2\sqrt{m})^{k-2}}{(k-2)! \cdot (k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k})} \cdot \frac{15 \ln(n/\bar{\delta})}{\bar{\epsilon}^2}$$

in Step 2 of the procedure,

$$\Pr\left[\frac{1}{r} \sum_{i=1}^r \chi_i < (1 - \bar{\epsilon}) \cdot \text{Exp}[\chi]\right] < \exp\left(-\frac{\bar{\epsilon}^2 \cdot \text{Exp}[\chi] \cdot r}{3}\right) \leq \frac{\bar{\delta}}{n}.$$

It follows that if T is typical and u is sociable, then with probability at least $1 - \bar{\delta}/n$,

$$\hat{c}_k(u) \geq (1 - \bar{\epsilon})^2 \cdot c_k(u) \geq \frac{1}{2} \cdot k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k},$$

causing the procedure to return sociable.

Now consider the case that u is shy. By Definition 4,

$$c_k(u) \leq \frac{1}{4} \cdot k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k},$$

implying that

$$\text{Exp}[\chi] < (1 + \bar{\epsilon}) \cdot \frac{\frac{(k-2)!}{4} \cdot (k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k})}{d(u) \cdot (2\sqrt{m})^{k-2}}.$$

Hence, by Chernoff's inequality, and by the setting of r ,

$$\Pr\left[\frac{1}{r} \sum_{i=1}^r \chi_i > \frac{3}{2} \cdot \left(\frac{(1 + \bar{\epsilon})}{4} \cdot k \cdot (50\bar{C}_k)^{1-1/k}/\bar{\epsilon}^{1/k}\right)\right] < \frac{\bar{\delta}}{n}.$$

Therefore, if T is degrees-typical and u is shy, then with probability at least $1 - \bar{\delta}/n$, the algorithm returns shy.

It remains to bound the running time of the procedure. By Lemma 14, the procedure Sample-a-Clique runs in time $O(k^2)$.

Crucially, Sample-a-Clique is invoked only if $d(u) \leq 4\bar{m}/(\bar{\epsilon}\bar{C}_k)^{1/k}$. Since

$$r = \frac{d(u) \cdot (2\sqrt{\bar{m}})^{k-2}}{(k-2)! \cdot (k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k})} \cdot \frac{15 \ln(n/\bar{\delta})}{\bar{\epsilon}^2},$$

the running time of the procedure is

$$O(r \cdot k^2) = O\left(\frac{\bar{m}^{k/2}}{\bar{C}_k} \cdot \frac{k \cdot 2^k \cdot \log(n/\bar{\delta})}{(k-2)! \cdot \bar{\epsilon}^{2-1/k}}\right),$$

as claimed. \square

3.5 Sampling a Degrees-Typical Set

In this subsection we provide the (simple) procedure Sample-Degrees-Typical and prove Lemma 12 regarding its correctness and running time. As in the case of the choice of the sample S by Approximate-Cliques, here too if the sample size t is larger than n , then the algorithm can simply set $T = V$.

Sample-Degrees-Typical ($n, k, \bar{m}, \bar{C}_k, \bar{\epsilon}, \bar{\delta}$)

- (1) $\gamma = \min\{1/(4\bar{m}^{k/2}), \bar{\delta}\}$.
- (2) For $i = 1$ to $\log(2/\gamma)$ do:
 - (a) Let T_i be a multiset of $t = \frac{10n \cdot \ln(2n/\gamma^2)}{(\bar{\epsilon}/k)^2 \cdot 2\sqrt{\bar{m}}}$ vertices chosen uniformly, at random.
 - (b) Query the degrees of all the vertices in T_i and compute $m(T_i)$.
- (3) Let T be the first set T_i such that $m(T_i) \leq \frac{t}{n} \cdot 4\bar{m}$. If no such set exists, then **return fail**. Else, set up a data structure $D(T)$ that supports sampling a uniform edge in $E(T)$ in constant time.
- (4) **Return** ($T, m(T), D(T)$).

Figure 4: The procedure for sampling the multiset T .

PROOF OF LEMMA 12: For each iteration i , consider the selection of the multiset T_i . For any fixed high-degree vertex $u \in V$ and for $j = 1, \dots, t$, let $\chi_j(u) = 1$ if the j^{th} vertex in T_i is a neighbor of u and let $\chi_j(u) = 0$ otherwise. By the definition of $\chi_1(u), \dots, \chi_t(u)$ and the premise of the lemma regarding \bar{m} , $\text{Exp}\left[\frac{1}{t} \sum_{j=1}^t \chi_j(u)\right] = \frac{d(u)}{n} \geq \frac{\sqrt{\bar{m}}}{n}$. By Chernoff's inequality and the setting of $t = \frac{10n \cdot \ln(2n/\gamma^2)}{(\bar{\epsilon}/k)^2 \cdot 2\sqrt{\bar{m}}}$ in Step 3,

$$\Pr\left[\left|\frac{1}{t} \sum_{j=1}^t \chi_j(u) - \frac{d(u)}{n}\right| > \frac{\bar{\epsilon}}{k} \cdot \frac{d(u)}{n}\right] < 2 \exp\left(-\frac{(\bar{\epsilon}/k)^2 \cdot t \cdot \frac{d(u)}{n}}{3}\right) < \frac{\gamma^2}{2n}.$$

By taking a union bound over all high-degree vertices, it holds that with probability at least $1 - \gamma^2/2$, for every high-degree vertex $u \in V$, $d_{T_i}(u) = \sum_{j=1}^t \chi_j(u) \in (1 \pm (\bar{\epsilon}/k)) \cdot \frac{t}{n} \cdot d(u)$. By taking a union bound over the $\log(2/\gamma)$ sampled multisets T_i , it holds that with probability at least $1 - \frac{1}{2}\gamma^2 \log(2/\gamma) > 1 - \gamma/2$, for each of the multisets T_i and for every high-degree vertex $u \in V$, $d_{T_i}(u) \in (1 \pm \frac{\bar{\epsilon}}{k}) \cdot \frac{t}{n} \cdot d(u)$.

We now turn to bounding the probability that none of the multiset T_i satisfies $m(T_i) \leq \frac{t}{n} \cdot 4\bar{m}$. By the definition of $m(\cdot)$, for every T_i we have that $\text{Exp}[m(T_i)] = \frac{t}{n} \cdot m$. By the assumption that $\bar{m} \in [(1-\bar{\epsilon})m, m]$ and by Markov's inequality, $\Pr[m(T_i) > \frac{t}{n} \cdot 4\bar{m}] \leq \Pr[m(T_i) > \frac{t}{n} \cdot 2m] < \frac{1}{2}$. Hence, the probability that $m(T_i) > \frac{t}{n} \cdot 4\bar{m}$ for every $i = 1, \dots, \log(2/\gamma)$ is at most $\gamma/2$.

By combining the two failure probabilities, we get that with probability at least $1 - \gamma$, the algorithm returns a multiset T that is degrees-typical (as defined in Definition 11).

Finally, by performing a preprocessing step that takes $\Theta(t)$ time it is possible to build a data structure $D(T)$ that allow for sampling each vertex $u \in S$ with probability proportional to $d(u)/m(T)$ (see e.g., [38, 56, 57]). This in turn implies that, using $D(T)$, it is possible to sample a uniform edge in $E(T)$ in constant time.

The running time of Sample-Degrees-Typical is

$$O(t \cdot \log(1/\gamma)) = O\left(\frac{n}{\sqrt{\bar{m}}} \cdot \frac{k^2 \cdot \log(n/\gamma) \cdot \log(1/\gamma)}{\bar{\epsilon}^2}\right),$$

and the proof is complete. \square

3.6 Proof of Theorem 10

In this subsection we prove Theorem 10. We first define the notion of a *cliques-typical* multiset.

DEFINITION 16. We say that a multiset S of size s is *cliques-typical* with respect to a partition P if $\alpha_P(S) \in (1 \pm \bar{\epsilon})^2 \cdot \frac{s}{n} \cdot C_k$.

We establish a simple claim regarding the multiset S selected by our algorithm (appropriate partitions are as defined in Definition 6).

CLAIM 17. Consider the multiset S sampled in Step 2 of Algorithm Approximate-Cliques. For any fixed partition P of V , $\text{Exp}[\alpha_P(S)] \leq \frac{s}{n} \cdot C_k$. Furthermore, if P is appropriate, $\bar{m} \geq (1 - \epsilon)m$ and $\bar{C}_k \in [C_k/4, C_k]$, then with probability at least $1 - \bar{\delta}$ (over the choice of S), the multiset S is *cliques-typical*.

PROOF. Recall that by Definition 7, $\alpha_P(S)$ is the number of k -cliques assigned to the vertices of S , and that by Corollary 8, for every partition P , $\alpha_P(V) \leq C_k$. Hence, $\text{Exp}[\alpha_P(S)] = \frac{s}{n} \cdot \alpha_P(V) \leq \frac{s}{n} \cdot C_k$.

We turn to consider the case that the partition P is appropriate, $\bar{m} \geq (1 - \epsilon)m$ and $\bar{C}_k \in [C_k/4, C_k]$. By Corollary 8, $\alpha_P(V) \in [(1 - \bar{\epsilon})C_k, C_k]$. By Definition 7, k -cliques are only assigned to vertices that are not sociable, implying that for every vertex $u \in V$, $\alpha_P(u) \leq c_k(u) \leq k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k}$. Hence, by the multiplicative Chernoff bound and the setting of $s = \frac{700 \cdot k \cdot n \cdot \ln(1/\bar{\delta})}{\bar{\epsilon}^{2+1/k} \cdot \bar{C}_k^{1/k}}$,

$$\begin{aligned} \Pr\left[\left|\frac{1}{s} \sum_{u \in S} \alpha_P(u) - \frac{\alpha_P(V)}{n}\right| > \bar{\epsilon} \cdot \frac{\alpha_P(V)}{n}\right] &< 2 \exp\left(-\frac{\bar{\epsilon}^2 \cdot \frac{\alpha_P(V)}{n} \cdot s}{3 \cdot k \cdot (50\bar{C}_k)^{1-1/k} / \bar{\epsilon}^{1/k}}\right) \\ &\leq 2 \exp\left(-\frac{\bar{\epsilon}^{2+1/k} \cdot \frac{(1-\bar{\epsilon})C_k}{n} \cdot \frac{700 \cdot k \cdot n \cdot \ln(1/\bar{\delta})}{\bar{\epsilon}^{2+1/k} \cdot \bar{C}_k^{1/k}}}{150 \cdot \bar{C}_k^{1-1/k}}\right) < \bar{\delta}. \end{aligned}$$

Therefore, if P is appropriate, $\bar{m} \geq (1 - \epsilon)m$ and $\bar{C}_k \in [C_k/4, C_k]$, with probability at least $1 - \bar{\delta}$,

$$\alpha_P(S) \in (1 \pm \bar{\epsilon}) \cdot \frac{s}{n} \cdot \alpha_P(V) \in (1 \pm \bar{\epsilon})^2 \cdot \frac{s}{n} \cdot C_k,$$

which implies that S is cliques-typical by Definition 16. \square

3.6.1 The case $\bar{C}_k \in [C_k/4, C_k]$. We start with the first item in the theorem. Recall that by the first premise of this item, $\bar{m} \in [(1 - \bar{\epsilon})m, m]$. By Lemma 12, with probability at least $1 - \min\{1/(4\bar{m}^{k/2}), \bar{\delta}\} \geq 1 - \bar{\delta}$ the procedure Sample-Degrees-Typical returns a multiset T that is degrees-typical. We henceforth condition on this event.

Since T is degrees-typical, by Lemma 15, an invocation of the procedure Is-sociable with a shy vertex u returns shy with probability at least $1 - \bar{\delta}/n$, and similarly an invocation with a sociable vertex u returns sociable with probability at least $1 - \bar{\delta}/n$. Consider (as a thought experiment) running the procedure Is-sociable on all the vertices in the graph, and letting V_0 be the set of vertices for which the procedure returned shy, $V_1 = V \setminus V_0$ and $P_{is} = (V_0, V_1)$. Conditioned on the event that T is degrees-typical, by Lemma 15 and by taking a union bound over all the vertices in V , with probability at least $1 - \bar{\delta}$ the partition P_{is} is appropriate (as defined in Definition 6). Suppose we fix the random coins used by Is-sociable on all vertices to a setting that indeed induces an appropriate partition P_{is} , and assume that all calls made by the algorithm to Is-sociable are answered consistently with P_{is} . (The probability that P_{is} is not appropriate is taken into account in the total failure probability of the algorithm.) Conditioned on P_{is} being appropriate (and using our assumptions on \bar{m} and \bar{C}_k), by Corollary 8 we have that $\alpha_{P_{is}}(V) \in [(1 - \bar{\epsilon})C_k, C_k]$.

Now consider the multiset S sampled in Step 2 of the algorithm. By Claim 17, with probability at least $1 - \bar{\delta}$, S is cliques-typical with respect to P_{is} . That is, $\alpha_{P_{is}}(S) \in (1 \pm \bar{\epsilon})^2 \cdot \frac{s}{n} \cdot C_k$. We condition on this event as well. Since T is degree-typical, by Lemma 14, whenever we invoke Sample-a-Clique (in Step 5a) it returns each k -tuple in $C(S)$ with probability in $(1 \pm \bar{\epsilon}) \cdot \frac{1}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}$.

By the description of Steps 5(b)i and 5(b)ii, χ_i is set to 1 only if the procedure Sample-a-Clique returns a k -tuple $K = (u, v, w_1, \dots, w_{k-2})$ in $C(S)$ that is assigned to u according to P_{is} . For each k -clique assigned to a vertex $u \in S$, the number of corresponding k -tuples in $C(S)$ is $(k-2)!$. Therefore, for a cliques-typical multiset S ,

$$\begin{aligned} \text{Exp}[\chi_i] &\in (1 \pm \bar{\epsilon}) \cdot \frac{\alpha_{P_{is}}(S) \cdot (k-2)!}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}} \\ &\in (1 \pm \bar{\epsilon})^3 \cdot \frac{\frac{s}{n} \cdot C_k \cdot (k-2)!}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}. \end{aligned} \quad (1)$$

By the multiplicative Chernoff bound and by the setting of $q = \frac{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}{(1 - \bar{\epsilon})^3 \cdot (k-2)! \cdot \bar{C}_k \cdot (s/n)} \cdot \frac{10 \ln(1/\bar{\delta})}{\bar{\epsilon}^2}$ in Step 5 of the algorithm,

$$\Pr[|\chi - \text{Exp}[\chi]| > \bar{\epsilon} \cdot \text{Exp}[\chi]] < 2 \exp\left(-\frac{\bar{\epsilon}^2 \cdot \text{Exp}[\chi] \cdot q}{3}\right) < \bar{\delta}.$$

Therefore, if $\bar{m} \in [(1 - \bar{\epsilon})m, m]$, $\bar{C}_k \in [C_k/4, C_k]$, T is degrees-typical, P_{is} is appropriate and S is cliques-typical, then with probability at least $1 - \bar{\delta}$,

$$\chi \in (1 \pm \bar{\epsilon})^4 \cdot \frac{\frac{s}{n} \cdot C_k \cdot (k-2)!}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}} \implies \hat{C}_k \in (1 \pm \bar{\epsilon})^4 \cdot C_k \in (1 \pm \epsilon)C_k,$$

where we have used the fact that $\bar{\epsilon} = \epsilon/5$. By taking a union bound over all bad events, $C_k \in (1 \pm \epsilon)C_k$ with probability at least $1 - 4\bar{\delta} > 1 - \delta$ (since $\bar{\delta} = \delta/4$).

3.6.2 The case $\bar{C}_k > C_k$. We now prove the second item of the theorem. As in the first item, since $\bar{m} \in [(1 - \bar{\epsilon})m, m]$, by Lemma 12 with probability at least $1 - \min\{1/(4\bar{m}^{k/2}), \bar{\delta}\}$ the multiset T is degrees-typical. Conditioned on T being degrees-typical, an invocation of the procedure Sample-a-Clique returns each k -tuple in $C(S)$ with probability in $(1 \pm \epsilon) \cdot \frac{1}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}$.

In this case, since $\bar{C}_k > C_k$, the setting of r in the procedure Is-sociable is not sufficiently large to ensure that the procedure is accurate, and that the partition P_{is} is appropriate (with high probability). Similarly, $\alpha_{P_{is}}(S)$ might not be close to its expected value. Therefore, we only use an upper bound on the expected value of $\alpha_{P_{is}}(S)$, as explained next.

Consider the following random variables. For a multiset S and for each $i \in [q]$ let $\chi_i(S)$ denote the random variable χ_i conditioned on S (where q and χ_i are as defined in Step 5 of the algorithm). Let $Y_i(S) = \frac{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}{(k-2)!(s/n)} \cdot \chi_i(S)$, and let $\hat{C}_k(S)$ denote the value of the random variable \hat{C}_k (which is the output of the algorithm), conditioned on S . We use the notation S-c to denote the random coins of the procedure Sample-a-Clique. By the above discussion and by the setting of \hat{C}_k in Step 6, if T is degrees-typical, then

$$\begin{aligned} \text{Exp}[\hat{C}_k] &= \text{Exp}_S \left[\text{Exp}_{S-c} [\hat{C}_k(S)] \right] \\ &= \text{Exp}_S \left[\text{Exp}_{S-c} \left[\frac{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}{(k-2)! \cdot \frac{s}{n}} \cdot \frac{1}{q} \sum_{i=1}^q \chi_i(S) \right] \right] \\ &= \text{Exp}_S \left[\text{Exp}_{S-c} \left[\frac{1}{q} \sum_{i=1}^q Y_i(S) \right] \right] \\ &= \text{Exp}_S [\text{Exp}_{S-c} [Y_1(S)]] . \end{aligned} \quad (2)$$

The last equality holds simply because all the $Y_i(S)$ variables are equally distributed (for each fixed S). We note that q also depends on S (since it is a function of $m(S)$), but this does not effect our analysis, and hence this dependence is not explicit. As in the analysis of the case that $\bar{C}_k \in [C_k/4, C_k]$, since for each k -clique assigned to a vertex $u \in S$, the number of corresponding k -tuples in $C(S)$ is $(k-2)!$, we have that if T is degrees-typical, then

$$\Pr[\chi_i(S) = 1] \leq (1 + \bar{\epsilon}) \cdot \frac{\alpha_{P_{is}}(S) \cdot (k-2)!}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}.$$

By the definition of the $Y_i(S)$ variables and Equation (2), it follows that

$$\text{Exp}[\hat{C}_k] \leq \text{Exp}_S \left[(1 + \bar{\epsilon}) \cdot \frac{\alpha_{P_{is}}(S)}{s/n} \right] = (1 + \bar{\epsilon}) \cdot \frac{n}{s} \cdot \text{Exp}_S [\alpha_{P_{is}}(S)].$$

By Claim 17, for any partition P , $\text{Exp}[\alpha_P(S)] \leq \frac{s}{n} \cdot C_k$. Therefore, if T is degrees-typical, then $\text{Exp}[\hat{C}_k] \leq (1 + \bar{\epsilon})C_k$. Finally, by

Markov's inequality (and recalling that $\bar{\epsilon} = \epsilon/5$),

$$\Pr \left[\widehat{C}_k > (1 + \epsilon/2)(1 + \bar{\epsilon})C_k \right] < 1 - \epsilon/2.$$

As noted in Subsection 3.2, we can assume that $\bar{\epsilon} \geq 1/\bar{m}^{k/2}$. Since T is not degrees-typical with probability at most $\min\{1/\bar{m}^{k/2}, \bar{\delta}\}$, by taking a union bound and by the setting of $\bar{\epsilon} = \epsilon/5$, it holds that with probability at least $\epsilon/2 - \min\{1/(4\bar{m}^{k/2}), \bar{\delta}\} > \epsilon/4$, the algorithm returns a value \widehat{C}_k such that $\widehat{C}_k \leq (1 + \epsilon/2)(1 + \bar{\epsilon})C_k \leq (1 + \epsilon)C_k$.

3.6.3 The expected query complexity and running time. By Lemma 12 and the assumption that $\bar{C}_k \leq \bar{m}^{k/2}$, the invocation of the procedure Sample-Degrees-Typical takes $O\left(\frac{n}{\bar{C}_k^{1/k}} \cdot \frac{k^2 \cdot \log(n/\gamma) \cdot \log(1/\gamma)}{\bar{\epsilon}^2}\right)$ time, for $\gamma = \Theta(\min\{1/\bar{m}^{k/2}, 1/\bar{\delta}\})$. The sampling of S and the computation of $m(S)$ and $D(S)$ take time $O(s) = O\left(\frac{k \cdot n}{\bar{C}_k^{1/k}} \cdot \frac{\ln(1/\bar{\delta})}{\bar{\epsilon}^{2+1/k}}\right)$. By Lemma 14, each invocation of Sample-a-Clique in Step 5a takes time $O(k^2)$, and Step 5(b)ii takes constant time. Therefore, excluding the invocations of the Is-sociable procedure, the running time of the for loop in Step 5 is $O(k^2 \cdot q)$. Since $\text{Exp}[m(S)] = \frac{s}{n} \cdot m$ and by the setting of q , it follows that the expected running time of the for loop is $O\left(\frac{m \cdot \bar{m}^{(k-2)/2}}{\bar{C}_k} \cdot \frac{k^2 \cdot \ln(1/\bar{\delta})}{(k-2)! \cdot \bar{\epsilon}^2}\right)$. It remains to bound the running time resulting from the invocations of the procedure Is-sociable.

We first bound the expected number of invocations of Is-sociable when the multiset T is degrees-typical. Let I denote the number of invocations of Is-sociable. Similarly to the analysis of the case $\bar{C}_k > C_k$, let $z_i(S)$ be a 0/1 random variable that is defined as follows: $z_i(S) = 1$ if (and only if) the procedure Sample-a-Clique returned a clique in the i^{th} step of the for loop in Step 5a of the algorithm, conditioned on the set S . As before, let $S\text{-c}$ denote the random coins of Sample-a-Clique. Here it is actually relevant that q depends on S , and therefore we use the notation $q(S)$. By the definition of I ,

$$\begin{aligned} \text{Exp}[I] &= \text{Exp}_S \left[\text{Exp}_{S\text{-c}} \left[\sum_{i=1}^{q(S)} z_i(S) \right] \right] \\ &= \text{Exp}_S [q(S) \cdot \text{Exp}_{S\text{-c}} [z_1(S)]] \\ &= \text{Exp}_S [\text{Exp}_{S\text{-c}} [q(S) \cdot z_1(S)]] . \end{aligned}$$

Recall that if T is degrees-typical, then by Lemma 14, an invocation of Sample-a-Clique returns each k -tuple in $C(S)$ with probability in $(1 \pm \bar{\epsilon}) \cdot \frac{1}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}$. Hence,

$$\Pr_{S\text{-c}}[z_1(S) = 1] \leq (1 + \bar{\epsilon}) \frac{|C(S)|}{m(S) \cdot (2\sqrt{\bar{m}})^{k-2}}.$$

By the setting of $q(S)$,

$$\text{Exp}_{S\text{-c}}[q(S) \cdot z_1(S)] \leq \frac{(1 + \bar{\epsilon})|C(S)|}{(1 - \bar{\epsilon})^3 \cdot (k-2)! \cdot \bar{C}_k \cdot \frac{s}{n}} \cdot \frac{10 \ln(1/\bar{\delta})}{\bar{\epsilon}^2}.$$

Since $\text{Exp}[|C(S)|] = \frac{s}{n} \cdot (k-2)! \cdot k \cdot C_k$, it follows that $\text{Exp}[I] \leq \frac{C_k}{\bar{C}_k} \cdot \frac{50k \ln(1/\bar{\delta})}{\bar{\epsilon}^2}$. Therefore, by Lemma 15, if T is degrees-typical, then

the expected running time resulting from all of the invocations of the Is-sociable procedure is $O\left(\frac{\bar{m}^{k/2}}{\bar{C}_k} \cdot \frac{C_k}{\bar{C}_k} \cdot \frac{k^2 \cdot 2^k \cdot \log(n/\bar{\delta}) \log(1/\bar{\delta})}{(k-2)! \cdot \bar{\epsilon}^{2-1/k}}\right)$.

If T is not degrees-typical then we can no longer upper bound the success probability of Sample-a-Clique, implying that the number of invocations of Is-sociable can be $\Theta(q)$. However, since T is not degrees typical with probability at most $\frac{1}{\bar{m}^{k/2}}$, this does not affect the expected running time resulting from the invocations of Is-sociable. The remaining steps take constant time, and therefore the total running time of the algorithm is

$$\begin{aligned} O\left(\frac{n}{\bar{C}_k^{1/k}} \cdot \frac{k^2 \cdot \log(n/\gamma) \cdot \log(1/\gamma)}{\bar{\epsilon}^{2+1/k}}\right) \\ + \frac{\min\{m, \bar{m}\} \cdot \bar{m}^{(k-2)/2}}{\bar{C}_k} \cdot \frac{C_k}{\bar{C}_k} \cdot \frac{k^2 \cdot 2^k \cdot \log(n/\bar{\delta}) \log(1/\bar{\delta})}{(k-2)! \cdot \bar{\epsilon}^{2-1/k}} \end{aligned}$$

which is

$$O\left(\frac{n}{\bar{C}_k^{1/k}} + \frac{\max\{m, \bar{m}\} \cdot \bar{m}^{(k-2)/2}}{\bar{C}_k} \cdot \frac{C_k}{\bar{C}_k}\right) \cdot \text{poly}(\log n, 1/\epsilon, \log(1/\bar{\delta}), k),$$

since $\gamma = \Theta(\min\{1/\bar{m}^{k/2}, \bar{\delta}\})$.

Finally, as discussed in the beginning of Subsection 3.2, if q is greater than \bar{m} then the algorithm may query beforehand for all the edges incident to S and their neighbors, and if it views more than \bar{m} edges then it can abort. Hence, the number of queries is always bounded by $O\left(\frac{n}{\bar{C}_k^{1/k}}\right) \cdot \text{poly}(\log n, 1/\epsilon, k) + \min\{m, \bar{m}\}$.

3.7 The Search Algorithm

In this subsection we describe an algorithm that returns an estimate of the number of k -cliques in a graph G without prior knowledge on the number of edges or k -cliques in the graph. We prove this by establishing a more general claim:

THEOREM 18. Let $\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V})$ be an algorithm that is given parameters $\bar{v}, \epsilon, \delta$ and possibly an additional set of parameters denoted \vec{V} , for which the following holds.

- (1) If $\bar{v} \in [v/4, v]$, then with probability at least $1 - \delta$, \mathcal{A} returns a value \widehat{v} such that $\widehat{v} \in (1 \pm \epsilon)v$.
- (2) If $\bar{v} > v$, then \mathcal{A} returns a value \widehat{v} , such that with probability at least $\epsilon/4$, $\widehat{v} \leq (1 + \epsilon)v$.
- (3) The expected running time of \mathcal{A} , denoted $E_{rt}(\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V}))$, is monotonically non-increasing with \bar{v} and furthermore, if $\bar{v} < v$, then $E_{rt}(\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V})) \leq E_{rt}(\mathcal{A}(v, \epsilon, \delta, \vec{V})) \cdot (v/\bar{v})^\ell$ for some constant $\ell > 0$.
- (4) The maximal running time of \mathcal{A} is D .

Then there exists an algorithm \mathcal{A}' that, when given an upper bound B on v , a parameter ϵ and a set of parameters \vec{V} , returns a value X such that the following holds.

- (1) $\mathcal{A}'(B, \epsilon, \vec{V})$ returns a value X such that $X \in (1 \pm \epsilon)v$ with probability at least $4/5$.

- (2) The expected running time of $\mathcal{A}'(B, \epsilon, \vec{V})$ is $E_{rt}(\mathcal{A}(v, \epsilon, \delta, \vec{V})) \cdot \text{poly}(\log B, 1/\epsilon, \ell)$ for $\delta = \Theta\left(\frac{\epsilon}{2^{\ell}(\ell + \log \log(B))}\right)$.
- (3) The maximal running time of \mathcal{A}' is $D \cdot \text{poly}(\log B, 1/\epsilon, \ell)$.

The algorithm \mathcal{A}' referred to by the theorem is provided in Figure 5. We note that the algorithm and the proof of Theorem 18 are a direct generalization of the search algorithm Estimate and the proof of Theorem 12 in [17]. However, this generalization may be useful as a “black box” in future work. We assume that $\epsilon \leq 1/4$, since otherwise we can simply set ϵ to $1/4$.

- $\mathcal{A}'(\mathcal{A}, B, \epsilon, \vec{V})$
- (1) Let $\bar{v} = B$.
 - (2) Let $\delta' = 1/(5 \cdot 2^{\ell})$.
 - (3) While $\bar{v} \geq 1$ do:
 - (a) For $\bar{v} = B, B/2, \dots, \bar{v}/2$ do:
 - (i) Let $r = (4/\epsilon) \cdot \ln(2 \log^2(B)/\delta')$ and let $\delta = \delta'/(2r)$.
 - (ii) Let $X_{\bar{v}}$ be the minimum value returned over r invocations of $\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V})$.
 - (iii) If $X_{\bar{v}} \geq (1 + \epsilon)\bar{v}$ then **return** \bar{v} .
 - (b) Let $\bar{v} = \bar{v}/2$.
 - (4) **Return** fail.

Figure 5: The search algorithm.

The following definition will be useful in the proof of the theorem.

DEFINITION 19. We say that a value X is a good estimate of v if $X \in (1 \pm \epsilon)v$.

PROOF OF THEOREM 18: Our search algorithm has two nested loops running with decreasing values of “guesses” for the value of v . The outer loop runs with \bar{v} , which is our current guess for the value of v , and the purpose of the inner for loop is to enhance the success probability of the algorithm when \bar{v} “passes” the good guess of v and runs with values smaller than $v/4$. We provide the full details subsequently, and start by considering only the outer while loop. Namely, imagine for now that instead of the for loop in Step 3a, we have the command “Let $\bar{v} = \bar{v}$ ” and the rest of the algorithm is as described in Figure 5.

First consider iterations of the while loop for which $\bar{v} > v$. By Item 2 in the properties of Algorithm \mathcal{A} , for values \bar{v} such that $\bar{v} > v$, the probability that $\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V})$ returns a value such that $\hat{v} > (1 + \epsilon)v$ is at most $1 - \epsilon/4$. Hence, the probability that the minimum value returned over r invocations, that is, $X_{\bar{v}}$, satisfies $X_{\bar{v}} > (1 + \epsilon)v$, is at most $(1 - \epsilon/4)^r = (1 - \epsilon/4)^{(4/\epsilon) \cdot \ln(2 \log^2(B)/\delta')} < \frac{\delta'}{2 \log^2(B)}$. It follows that for each value $\bar{v} > v$, with probability at least $1 - \frac{\delta'}{2 \log^2(B)}$,

$$X_{\bar{v}} < (1 + \epsilon)v < (1 + \epsilon)\bar{v}.$$

This implies that for each value $\bar{v} > v$, with probability at least $1 - \frac{\delta'}{2 \log^2(B)}$, the algorithm \mathcal{A}' will continue to run with $\bar{v} = \bar{v}/2$.

Now consider values of \bar{v} such that $\bar{v} \in [v/4, v]$. By Item 1 in the properties of Algorithm \mathcal{A} , if $\bar{v} \in [v/4, v]$, then with probability at least $1 - \delta'$, $\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V})$ returns a value \hat{v} , such that $\hat{v} \in (1 \pm \epsilon)v$. By the setting $\delta = \delta'/(2r)$, and by taking a union bound, with probability at least $1 - \delta'/2$, all the r invocations return a good estimate of v , implying that $X_{\bar{v}} \in (1 \pm \epsilon)v$ (i.e., $X_{\bar{v}}$ is a good estimate of v). Note that in such a case,

$$X_{\bar{v}} \geq (1 - \epsilon)v \geq (1 - \epsilon) \cdot 2\bar{v} \geq (1 + \epsilon)\bar{v}.$$

(Here we have used the assumption that $\epsilon \leq 1/4$.) Hence, once \mathcal{A}' reaches a value $\bar{v} \in [v/4, v/2]$, with probability at least $1 - \delta'/2$, it returns a value $X_{\bar{v}}$ that is a good estimate of v .

Finally, if \mathcal{A}' reaches values $\bar{v} < v/4$, we no longer have a guarantee on the probability that \mathcal{A} returns a value \hat{v} or on the quality of the estimate \hat{v} . Hence, we also have an inner for loop so that whenever we halve the guess \bar{v} we first run with all the values B, \dots, \bar{v} . This ensures that even if the algorithm did not return a value when running with a good guess $\bar{v} \in [v/4, v/2]$, we can still bound the probability that it will continue to run with decreasing values of \bar{v} .

We now consider the original algorithm with both the outer while loop and the inner for loop as described in Figure 5. There are at most $\log(B)$ invocations of the while loop with values $\bar{v} > v$, implying that there are at most $\log^2(B)$ invocations of the for loop in Step 3(a) with a value $\bar{v} > v$. Therefore, by the above analysis, the probability that \mathcal{A}' returns a value that is not a good estimate of v in these invocations is at most $\delta'/2$. Hence, with probability at least $1 - \delta'/2$ we will reach an invocation in which $\bar{v} \in [v/4, v/2]$ and $\bar{v} \in [v/4, v/2]$, for which with probability at least $1 - \delta'/2$ the algorithm \mathcal{A} returns a value $X_{\bar{v}}$ that is a good estimate of v . By taking a union bound and by the setting of δ' , the algorithm \mathcal{A}' returns a value that is a good estimate of v with probability at least $1 - \delta' > 4/5$.

We turn to analyze the running time of \mathcal{A}' . By Item 3, for values of \bar{v} such that $\bar{v} \geq v$, the expected running time of $\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V})$ is $E_{rt}(\mathcal{A}(v, \epsilon, \delta, \vec{V}))$. By the above analysis, once \mathcal{A}' reaches a value $\bar{v} < v/4$, since we run with all values $\bar{v} = B, \dots, \bar{v}$, the probability that the algorithm will halve the guess to $\bar{v} = \bar{v}/2$ is at most δ' . Hence, the probability of invoking \mathcal{A} with a value $\bar{v} = v/2^z$ is at most $(\delta')^z$. By Item 3 in the properties of Algorithm \mathcal{A} , when running with values $\bar{v} < v$, it holds that $E_{rt}(\mathcal{A}(\bar{v}, \epsilon, \delta, \vec{V})) \leq E_{rt}(\mathcal{A}(v, \epsilon, \delta, \vec{V})) \cdot (v/\bar{v})^{\ell}$. Therefore, the expected running time of \mathcal{A}' is upper bounded by

$$\begin{aligned} & \log^2(B) \cdot r \cdot E_{rt}(\mathcal{A}(v, \epsilon, \delta, \vec{V})) \\ & + \sum_{z=1}^{\log B} (\delta')^z \cdot 2^{\ell \cdot z} \cdot E_{rt}(\mathcal{A}(v, \epsilon, \delta, \vec{V})) \\ & = E_{rt}(\mathcal{A}(v, \epsilon, \delta, \vec{V})) \cdot \text{poly}(\log B, 1/\epsilon, \ell), \end{aligned}$$

where the first term is due to the invocations in which $\bar{v} \geq v$ and the second is due to invocations in which $\bar{v} < v$, and the equality is due to the setting of $\delta' < 1/2^\ell$. Therefore, Item 2 of the theorem holds. The proof of Item 3 is immediate. \square

The following is a corollary of Theorem 10 and Theorem 18 and is a restatement of Theorem 1.

COROLLARY 20. *There exists an algorithm that, given $n, k \geq 3$, and query access to a graph G , returns a value X such that $X \in (1 \pm \epsilon)C_k$ with probability at least $2/3$. The expected running time of the algorithm is $O\left(\frac{n}{C_k^{1/k}} + \frac{m^{k/2}}{C_k \cdot (k-2)!}\right) \cdot \text{poly}(\log n, 1/\epsilon, k)$, and its expected query complexity is $O\left(\frac{n}{C_k^{1/k}} + \min\left\{\frac{m^{k/2}}{C_k \cdot (k-2)!}, m\right\}\right) \cdot \text{poly}(\log n, 1/\epsilon, k)$.*

PROOF. We start by obtaining an estimate \bar{m} of m such that with probability at least $1 - \min\{1/n, 1/m^{k/2}\}$ it holds that $\bar{m} \in [(1 - \epsilon/4)m, m]$. This can be done by invoking the algorithm of Goldreich and Ron [28] for estimating the number of edges $\Theta(\log(n^{2k}))$ times and taking the median value returned. Next we invoke \mathcal{A}' (Approximate-cliques, $B = \min\{n^k, \bar{m}^{k/2}\}, \epsilon, \vec{V}$) with $\vec{V} = (n, k, \bar{m})$ and return the value X returned by \mathcal{A}' .

By Theorem 10, if $\bar{m} \in [(1 - \epsilon/5)m, m]$, then Approximate-cliques satisfies the conditions required from Algorithm \mathcal{A} in Theorem 18, with $\ell = 2$ and $D = n + m^{k/2}$ (since this is the maximal running time when $\bar{C}_k = 1$). Hence, by Theorem 18 and by the union bound, \mathcal{A}' returns a value X such that with probability at least $4/5 - 1/m^k > 2/3$, it holds that $X \in (1 \pm \epsilon)C_k$.

By [28], the first step of approximating the number of edges \bar{m} with success probability at least $1 - \min\{1/n, 1/m^k\}$ takes time $O\left(\frac{n}{\sqrt{m}}\right) \cdot \text{poly}(\log n, 1/\epsilon, k)$ which by Claim 3 is at most $O\left(\frac{n}{C_k^{1/k}}\right) \cdot \text{poly}(\log n, 1/\epsilon, k)$. When $\bar{m} \in [1 - (\epsilon/5)m, m]$, by Item 3 in Theorem 10, and by Item 2 in the properties of Algorithm \mathcal{A}' , the expected running time of the algorithm is

$$O\left(\frac{n}{C_k^{1/k}} + \frac{m^{k/2}}{C_k}\right) \cdot \text{poly}(\log n, 1/\epsilon, k).$$

Since $\bar{m} \notin [(1 - \epsilon/5)m, m]$ with probability at most $\min\{1/n, 1/m^{k/2}\}$ and the maximal running time of Approximate-cliques is at most $n + m^{k/2}$ up to $\text{poly}(\log n, 1/\epsilon, k)$ factors, this event does not effect the expected query complexity.

By Item 3 in Theorem 10, Algorithm Approximate-cliques never performs more than $O\left(\frac{n}{C_k^{1/k}}\right) \cdot \text{poly}(\log n, 1/\epsilon, \log(1/\delta), k) + \min\{m, \bar{m}\}$ queries. It follows that the expected query complexity is

$$O\left(\frac{n}{C_k^{1/k}} + \min\left\{\frac{m^{k/2}}{C_k}, m\right\}\right) \cdot \text{poly}(\log n, 1/\epsilon, k),$$

as claimed. \square

REFERENCES

- [1] A. Abboud, A. Backurs, and V. V. Williams. 2015. If the Current Clique Algorithms are Optimal, So is Valiant's Parser. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, 98–117.
- [2] K. J. Ahn, S. Guha, and A. McGregor. 2012. Graph sketches: sparsification, spanners, and subgraphs. In *Proceedings of the Symposium on Principles of Database Systems (PODS)*, 5–14.
- [3] H. Avron. 2010. Counting triangles in large graphs using randomized matrix trace estimation. In *Workshop on Large-scale Data Mining: Theory and Applications (LDMTA)*, Vol. 10, 10–9.
- [4] L. Becchetti, P. Boldi, C. Castillo, and A. Gionis. 2008. Efficient semi-streaming algorithms for local triangle counting in massive graphs. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, 16–24.
- [5] J. W. Berry, B. Hendrickson, R. A. LaViolette, and C. A. Phillips. 2011. Tolerating the Community Detection Resolution Limit with Edge Weighting. *Physical Review E* 83, 5 (May 2011), 056119.
- [6] R. S. Burt. 2004. Structural Holes and Good Ideas. *Amer. J. Sociology* 110, 2 (2004), 349–399. <http://www.jstor.org/stable/10.1086/421787>
- [7] B. Chazelle, R. Rubinfeld, and L. Trevisan. 2005. Approximating the Minimum Spanning Tree Weight in Sublinear Time. *SIAM J. Comput.* 34, 6 (2005), 1370–1379.
- [8] H. Chernoff. 1952. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics* (1952), 493–507.
- [9] N. Chiba and T. Nishizeki. 1985. Arboricity and subgraph listing algorithms. *SIAM J. Comput.* 14, 1 (1985), 210–223.
- [10] F. Chierichetti, A. Dasgupta, R. Kumar, S. Lattanzi, and T. Sarlos. 2016. On Sampling Nodes in a Network. In *Conference on the World Wide Web (WWW)*, 471–481.
- [11] S. Chu and J. Cheng. 2011. Triangle listing in massive networks and its applications. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, 672–680.
- [12] J. S. Coleman. 1988. Social Capital in the Creation of Human Capital. *Amer. J. Sociology* 94 (1988), S95–S120. <http://www.jstor.org/stable/2780243>
- [13] A. Czumaj, F. Ergün, L. Fortnow, A. Magen, I. Newman, R. Rubinfeld, and C. Sohler. 2005. Approximating the Weight of the Euclidean Minimum Spanning Tree in Sublinear Time. *SIAM J. Comput.* 35, 1 (2005), 91–109.
- [14] A. Czumaj and C. Sohler. 2009. Estimating the Weight of Metric Minimum Spanning Trees in Sublinear Time. *SIAM J. Comput.* 39, 3 (2009), 904–922.
- [15] A. Dasgupta, R. Kumar, and T. Sarlos. 2014. On Estimating the Average Degree. In *Conference on the World Wide Web (WWW)*, ACM, 795–806.
- [16] J. P. Eckmann and E. Moses. 2002. Curvature of Co-links Uncovers Hidden Thematic Layers in the World Wide Web. *Proceedings of the National Academy of Sciences* 99, 9 (2002), 5825–5829. <https://doi.org/10.1073/pnas.032093399>
- [17] T. Eden, A. Levi, D. Ron, and C. Seshadhri. 2015. Approximately counting triangles in sublinear time. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, 614–633.
- [18] T. Eden, D. Ron, and C. Seshadhri. 2017. On Approximating the Number of k -cliques in Sublinear Time. *CoRR abs/1707.04858v1* (2017). [arXiv:1707.04858v1](http://arxiv.org/abs/1707.04858v1)
- [19] T. Eden, D. Ron, and C. Seshadhri. 2017. Sublinear Time Estimation of Degree Distribution Moments: The Degeneracy Connection. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, 7:1–7:13. <https://doi.org/10.4230/LIPIcs.ICALP.2017.7>
- [20] T. Eden and W. Rosenbaum. 2017. Lower Bounds for Approximating Graph Parameters via Communication Complexity. *CoRR abs/1709.04262* (2017). [arXiv:1709.04262](http://arxiv.org/abs/1709.04262)
- [21] T. Eden and W. Rosenbaum. 2018. On Sampling Edges Almost Uniformly. In *1st Symposium on Simplicity in Algorithms, SOSA 2018, January 7-10, 2018, New Orleans, LA, USA*, 7:1–7:9. <https://doi.org/10.4230/OASIcs.SOSA.2018.7>
- [22] F. Eisenbrand and F. Grandoni. 2004. On the complexity of fixed parameter clique and dominating set. *Theoretical Computer Science* 326, 1-3 (2004), 57–67.
- [23] D. Eppstein, M. Löffler, and D. Strash. 2013. Listing All Maximal Cliques in Large Sparse Real-World Graphs. *ACM Journal of Experimental Algorithmics* 18 (2013), 3–1. <https://doi.org/10.1145/2543629>
- [24] U. Feige. 2006. On sums of independent random variables with unbounded variance and estimating the average degree in a graph. *SIAM J. Comput.* 35, 4 (2006), 964–984.
- [25] I. Finocchi, M. Finocchi, and E. G. Fusco. 2015. Clique Counting in MapReduce: Algorithms and Experiments. *ACM Journal of Experimental Algorithmics* 20 (2015), 1–7. <https://doi.org/10.1145/2794080>
- [26] B. Foucault Welles, A. Van Deventer, and N. Contractor. 2010. Is a friend a friend?: Investigating the structure of friendship networks in virtual worlds. In *CHI Extended Abstracts on Human Factors in Computing Systems*, 4027–4032.
- [27] O. Goldreich. 2017. *Introduction to Property Testing*. Cambridge University Press.
- [28] O. Goldreich and D. Ron. 2008. Approximating average parameters of graphs. *Random Structures and Algorithms* 32, 4 (2008), 473–493.
- [29] M. Gonen, D. Ron, and Y. Shavitt. 2011. Counting stars and other small subgraphs in sublinear-time. *SIAM Journal on Discrete Mathematics* 25, 3 (2011), 1365–1411.

- [30] A. Hassidim, J. A. Kelner, H. N. Nguyen, and K. Onak. 2009. Local graph partitions for approximation and testing. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 22–31.
- [31] P. W. Holland and S. Leinhardt. 1970. A method for detecting structure in sociometric data. *Amer. J. Sociology* 76 (1970), 492–513.
- [32] M. O. Jackson, T. Rodriguez-Barraquer, and X. Tan. 2012. Social Capital and Social Quilts: Network Patterns of Favor Exchange. *American Economic Review* 102, 5 (2012), 1857–1897.
- [33] S. Jain and C. Seshadhri. 2017. A Fast and Provable Method for Estimating Clique Counts Using Turán's Theorem. In *Conference on the World Wide Web (WWW)*. 441–449.
- [34] H. Jowhari and M. Ghodsi. 2005. New streaming algorithms for counting triangles in graphs. In *Proceedings of the International Conference Computing and Combinatorics (COCOON)*. Springer, 710–716.
- [35] D. M. Kane, K. Mehlhorn, T. Sauerwald, and H. Sun. 2012. Counting Arbitrary Subgraphs in Data Streams. In *International Colloquium on Automata, Languages, and Programming (ICALP)*. 598–609.
- [36] M. N. Kolountzakis, G. L. Miller, R. Peng, and C. E. Tsourakakis. 2012. Efficient triangle counting in large graphs via degree-based vertex partitioning. *Internet Mathematics* 8, 1-2 (2012), 161–185.
- [37] S. Marko and D. Ron. 2009. Approximating the distance to properties in bounded-degree and general sparse graphs. *ACM Transactions on Algorithms* 5, 2 (2009), 22.
- [38] G. Marsaglia, W. W. Tsang, J. Wang, et al. 2004. Fast generation of discrete random variables. *Journal of Statistical Software* 11, 3 (2004), 1–11.
- [39] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. 2002. Network motifs: simple building blocks of complex networks. *Science* 298, 5594 (2002), 824–827.
- [40] J. Nešetřil and S. Poljak. 1985. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae* 26, 2 (1985), 415–419.
- [41] H. N. Nguyen and K. Onak. 2008. Constant-time approximation algorithms via local improvements. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*. 327–336.
- [42] K. Onak, D. Ron, M. Rosen, and R. Rubinfeld. 2012. A near-optimal sublinear-time algorithm for approximating the minimum vertex cover size. In *Proceedings of the Symposium on Discrete Algorithms (SODA)*. 1123–1131.
- [43] M. Parnas and D. Ron. 2007. Approximating the minimum vertex cover in sublinear time and a connection to distributed algorithms. *Theoretical Computer Science* 381, 1-3 (2007), 183–196.
- [44] Alejandro Portes. 2000. Social Capital: Its Origins and Applications in Modern Sociology. In *Knowledge and Social Capital*, Eric L. Lesser (Ed.). Butterworth-Heinemann, Boston, 43 – 67. <https://doi.org/10.1016/B978-0-7506-7222-1.50006-4>
- [45] T. Schank and D. Wagner. 2005. Approximating Clustering Coefficient and Transitivity. *Journal of Graph Algorithms and Applications* 9 (2005), 265–275. Issue 2.
- [46] T. Schank and D. Wagner. 2005. Finding, Counting and Listing All Triangles in Large Graphs, an Experimental Study. In *Experimental and Efficient Algorithms*. 606–609.
- [47] C. Seshadhri, T. G. Kolda, and A. Pinar. 2012. Community structure and scale-free collections of Erdős-Rényi graphs. *Physical Review E* 85, 5 (May 2012), 056109. <https://doi.org/10.1103/PhysRevE.85.056109>
- [48] C. Seshadhri, A. Pinar, and T. G. Kolda. 2013. Fast Triangle Counting through Wedge Sampling. arXiv:1202.5230. In *Proceedings of the International Conference on Data Mining (ICDM)*, Vol. 4. 5. <http://arxiv.org/abs/1202.5230>
- [49] S. Suri and S. Vassilvitskii. 2011. Counting triangles and the curse of the last reducer. In *Proceedings of the International Conference on World Wide Web (WWW)*. 607–614. <https://doi.org/10.1145/1963405.1963491>
- [50] K. Tangwongsan, A. Pavan, and S. Tirthapura. 2013. Parallel Triangle Counting in Massive Streaming Graphs. In *Proceedings of the International Conference on Information and Knowledge Management (CIKM)*. ACM, 781–786.
- [51] C. E. Tsourakakis. 2008. Fast counting of triangles in large real networks without counting: Algorithms and laws. In *International Conference on Data Mining (ICDM)*. 608–617.
- [52] C. E. Tsourakakis. 2015. The K-clique Densest Subgraph Problem. In *Proceedings of the International Conference on World Wide Web (WWW)*. 1122–1132. <https://doi.org/10.1145/2736277.2741098>
- [53] C. E. Tsourakakis, U. Kang, G. L. Miller, and C. Faloutsos. 2009. Doulion: counting triangles in massive graphs with a coin. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining (SIGKDD)*. 837–846.
- [54] C. E. Tsourakakis, M. N. Kolountzakis, and G. L. Miller. 2011. Triangle Sparsifiers. *Journal of Graph Algorithms and Applications* 15, 6 (2011), 703–726.
- [55] V. Vassilevska. 2009. Efficient algorithms for clique problems. *Inform. Process. Lett.* 109, 4 (2009), 254–257.
- [56] A. J. Walker. 1974. New fast method for generating discrete random numbers with arbitrary frequency distributions. *Electronics Letters* 10, 8 (1974), 127–128.
- [57] A. J. Walker. 1977. An efficient method for generating discrete random variables with general distributions. *ACM Trans. Math. Software* 3, 3 (1977), 253–256.
- [58] Y. Yoshida, M. Yamamoto, and H. Ito. 2009. An improved constant-time approximation algorithm for maximum. In *Proceedings of the Symposium on Theory of Computing (STOC)*. 225–234.