

The Sketching Complexity of Graph and Hypergraph Counting

John Kallaugher
UT Austin
Austin, USA
jmgk@cs.utexas.edu

Michael Kapralov
EPFL
Lausanne, Switzerland
michael.kapralov@epfl.ch

Eric Price
UT Austin
Austin, USA
ecprice@cs.utexas.edu

Abstract—Subgraph counting is a fundamental primitive in graph processing, with applications in social network analysis (e.g., estimating the clustering coefficient of a graph), database processing and other areas. The space complexity of subgraph counting has been studied extensively in the literature, but many natural settings are still not well understood. In this paper we revisit the subgraph (and hypergraph) counting problem in the sketching model, where the algorithm’s state as it processes a stream of updates to the graph is a linear function of the stream. This model has recently received a lot of attention in the literature, and has become a standard model for solving dynamic graph streaming problems.

In this paper we give a tight bound on the sketching complexity of counting the number of occurrences of a small subgraph H in a bounded degree graph G presented as a stream of edge updates. Specifically, we show that the space complexity of the problem is governed by the fractional vertex cover number of the graph H . Our subgraph counting algorithm implements a natural vertex sampling approach, with sampling probabilities governed by the vertex cover of H . Our main technical contribution lies in a new set of Fourier analytic tools that we develop to analyze multiplayer communication protocols in the simultaneous communication model, allowing us to prove a tight lower bound. We believe that our techniques are likely to find applications in other settings. Besides giving tight bounds for all graphs H , both our algorithm and lower bounds extend to the hypergraph setting, albeit with some loss in space complexity.

I. INTRODUCTION

Triangle counting is one of the most well-studied problems in streaming graph algorithms. In the standard “turnstile” version of this problem, one maintains a small-space “sketch” of a graph under a stream of insertions and deletions of edges, and at the end of the stream outputs a $(1 \pm \varepsilon)$ multiplicative approximation to the number of triangles T in the graph; unless otherwise specified, we assume ε to be a small constant.

Turnstile streaming algorithms are almost invariably constructed as *linear* sketches, where the sketch maintained is a linear function of the indicator vector of edges; this makes it easy to process insertions and

deletions. Linear sketches are also useful in other settings such as distributed computation, since sketches can be merged. There is evidence that any turnstile streaming algorithm can be efficiently implemented using linear sketches [LNW14a], [AHLW16], although these results do not quite apply to graph streams.

For worst-case graphs, counting triangles is impossible in sublinear space: $\Omega(m)$ space is required to distinguish between a graph with 0 triangles and one with $T = \Omega(m)$ triangles [BOV13]. However, the hard case is degenerate in that all the triangles share a common edge. If at most Δ_E triangles share any single edge, then this bound becomes $\Omega(m\Delta_E/T)$. In [PT12] an algorithm was given that counts triangles with

$$O\left(m\left(\frac{1}{\sqrt{T}} + \frac{\Delta_E}{T}\right)\right)$$

space, where the m/\sqrt{T} term improves upon a $m/T^{1/3}$ term in [TKMF09], [TKM11]. In [KP17] this was shown to be tight for worst-case graphs, but the hard case is again degenerate: all the triangles share a common vertex. If Δ_V bounds the maximum number of triangles to share a vertex, this bound becomes $m\sqrt{\Delta_V}/T$. The algorithm in [KP17] requires

$$\tilde{O}\left(m\left(\frac{1}{T^{2/3}} + \frac{\sqrt{\Delta_V}}{T} + \frac{\Delta_E}{T}\right)\right)$$

space. A natural question is whether this $\frac{m}{T^{2/3}}$ is necessary.

A. Linear Sketching

Suppose we have a problem of the following form: we receive a vector $v \in \mathbb{Z}^n$ as a series of updates $(v_i)_{i=1}^t$, so $v = \sum_{i=1}^t v_i$, and we want to approximate some function $f(v)$. A *linear sketch* for this problem is a linear transformation $A \in \mathbb{Z}^{n \times d}$ and a post-processing function g , so that $g(Av)$ approximates $f(v)$ (with the exact definition of “approximates” depending on the problem). The space complexity of such a sketch is the space needed to store the sketch vector, which is $\Theta(d \log n)$

bits if the maximum size of entries of v is bounded by some $M = \text{poly}(n)$ (this holds even if intermediate stages of the stream exceed M , as the sketch vector may be stored mod M).

In [LNW14a], it was shown that any turnstile streaming algorithm (an algorithm that can solve a problem of the above form when the updates v_i are allowed to be negative, but that is allowed to maintain arbitrary state) can be converted into a linear sketch with only logarithmic loss in space complexity. In [AHLW16], this result was extended to *strict* turnstile streaming algorithms, that is algorithms which require that $\sum_{i=1}^s v_i \geq 0$ for each $s \leq t$.

These results come with two important caveats. Firstly, they do not necessarily give an $O(d \log n)$ -space streaming algorithm, as neither the linear transformation A nor the post-processing function g is known to be calculable in $O(d \log n)$ space. However, our sketching lower bounds will be based on communication complexity arguments that bound the size of the sketch vector, circumventing this issue.

Secondly, [LNW14a] assumes that the turnstile algorithm in question works regardless of the value of the partial sums $\sum_{i=1}^s v_i$, while [AHLW16] only requires that these sums be non-negative. Therefore, our lower bounds do not rule out the possibility of a turnstile algorithm that requires every partial sum to form a valid graph (i.e. edges can neither be deleted before they arrive, nor can they arrive multiple times before being deleted). However, existing turnstile algorithms do not typically require this property—in particular, any sampling-based algorithm, whether adaptive or non-adaptive, can handle it by the addition of a counter to each stored edge.

B. Our Results

a) Lower bound.: We show that any linear sketching algorithm for triangle counting requires $\Omega(\frac{m}{T^{2/3}})$ space, even for constant degree graphs. Such a result is not true for the insertion-only model, where triangles can be counted in graphs with max degree d in $O(md^2/T)$ space by subsampling the edges at rate d/T and storing all subsequent edges that touch the sampled edges [JG05].

Our result generalizes to counting the number of copies of any constant-size connected subgraph H . Such problems appear, for example, in estimating the size of database joins when planning queries [AGM08]. We show that the linear sketching complexity of distinguishing between 0 and T copies of H in constant-degree graphs is at least

$$\Omega\left(\frac{m}{T^{1/\tau}}\right)$$

where τ is the fractional vertex cover number of H , the minimum value such that there exists $f : V(H) \rightarrow [0, 1]$ with $\sum_{v \in V(H)} f(v) \leq \tau$ and $f(u) + f(v) \geq 1$ for all $uv \in E(H)$.

b) Upper bound.: We also give a matching upper bound: by subsampling the vertices with probabilities dependent on their weight in the fractional vertex cover, we give an algorithm that estimates T using $O(m/T^{1/\tau})$ words of space, as long as the graph has constant degree. Additionally, the constant-degree restriction can be lifted for many graphs: if an optimal fractional vertex cover of H can place nonzero weight on every vertex (as occurs, for example, if H is a cycle) then the algorithm works for degree up to $T^{1/(2\tau)}$ graphs.

c) Hypergraph counting.: Both our upper and lower bounds extend to counting hypergraphs H , but the exponent on T no longer matches for all hypergraphs. The upper bound remains $O(m/T^{1/\tau})$, while the lower bound becomes $O(m/T^{1/\mu})$ for an exponent μ that equals the fractional vertex cover number τ on many hypergraphs but not all.

All of our results extend to $\varepsilon \ll 1$; the full statements of these results are given in Theorem 12 for the upper bound, Theorem 7 for the general lower bound, and Corollary 11 for the tight lower bound specific to non-hypergraphs.

C. Sampling and Sketching

Our upper bounds will take the form of *non-adaptive* sampling algorithms. By “sampling algorithm” we mean that the only state maintained between updates is a subset of the input edges, and by “non-adaptive” we mean that the probability of keeping an edge does not depend on which other edges have been seen so far in the stream¹.

Non-adaptive sampling algorithms may be modified into linear sketching algorithms by the use of L_0 -sampler sketches. An L_0 -sampler sketch is a linear sketch which, if v is the frequency vector of the input stream, returns a non-zero co-ordinate of v , chosen uniformly at random (more generally, an L_p -sampler samples v_i with probability proportional to $|v_i|^p$). A linear sketching algorithm for this problem was first presented in [CMR05], while [MW10] defined L_p sampling and gave algorithms for all $p \in [0, 2]$.

In [JST11], it was shown that, if the set to be sampled from has size n and the sample is required

¹Note that this does not mean that the edges are sampled *independently* of one another—for instance, if we choose a vertex at random and keep all edges incident on that vertex, the event that we keep the edge uv is independent of whether the edge vw is present in the stream, but it is not independent of the event that that we keep the edge vw .

to be within δ of uniform for some constant δ , the space required is exactly $\Theta(\log^2 n)$. In [KNP⁺17], the optimal bound in terms of n and δ was shown to be $\Theta\left(\min\left(n, \log(1/\delta) \log^2\left(\frac{n}{\log(1/\delta)}\right)\right)\right)$.

Therefore, as our sketching lower bound and sampling upper bounds match up to polylog factors, it follows that both are themselves tight up to polylog factors.

D. Our Techniques

The core of our lower bound proof is a new set of Fourier analytic techniques for analyzing multiplayer simultaneous communication protocols. Our approach is inspired by the Fourier analytic analysis of the Boolean Hidden Matching problem, but develops several new ideas that we think are likely to find applications beyond subgraph counting lower bounds. We now proceed to describe the Boolean Hidden Matching problem, the main ideas behind its analysis, and then describe our techniques.

The Boolean Hidden Matching problem of Gavinsky et al [GKK⁺07] is a two player one way communication problem where Alice holds a binary string $x \in \{0, 1\}^n$ that she compresses to a message m of s bits and sends to Bob. Bob, besides the message from Alice, gets two pieces of input: a uniformly random matching of size $n/10$ on the set $\{1, 2, \dots, n\}$, along with a vector of binary labels $w_e \in \{0, 1\}, e \in M$. In the YES case of the problem the vector w satisfies $w = Mx$, and in the NO case of the problem the vector w satisfies $w = Mx \oplus 1^{|M|}$, where we abuse notation somewhat by letting M denote the edge incidence matrix of the matching M (where each column corresponds to a vertex, and each edge to a row, with ones in the two co-ordinates corresponding to the vertices the edge is incident to).

The Boolean Hidden Matching problem and the related Boolean Hidden Hypermatching problem of Verbin and Yu [VY11] have been very influential in streaming lower bounds: streaming problems that have recently been shown to admit reductions from Boolean Hidden (Hyper)Matching include approximating maximum matching size [EHL⁺15], [AKLY15], [AKL17], approximating MAX-CUT value [KKS15], [KK15], [KKS17], subgraph counting [VY11], [KP17], and approximating Schatten p -norms [LW16], among others. Most recent streaming lower bounds (with the exception of [KKS17]) use reductions from Boolean Hidden (Hyper)Matching, without modifying the Fourier analytic techniques involved in the proof.

In this paper we develop several new Fourier analytic ideas that go beyond the Boolean Hidden Matching problem in several directions:

a) Analyzing simultaneous multiplayer communication: In the Boolean Hidden Matching problem Alice is the only player transmitting a message, but our communication problem features simultaneous communication from multiple players to a referee. We show how to use the convolution theorem from Fourier analysis to analyze the effect of combining information sent by multiple players in the one way simultaneous communication model. While the technique of combining information from two players using the convolution theorem was recently used by [KKS17] to analyze a three-player game, to the best of our knowledge our work is the first to analyze games with an arbitrary number of players in this manner.

b) Analyzing a promise version of a communication problem via Fourier analysis: While in the Boolean Hidden Matching problem Alice’s string is sampled from the uniform distribution, in our problem multiple players receive correlated binary strings (conditioned on a linear constraint over the binary field). It turns out that this specific form of conditioning lends itself naturally to a Fourier analytic approach due to the linearity of the constraints imposed on the strings, and analysing such correlated settings gives us tight bounds on the subgraph counting version of our problem.

c) Sharing M among the players: In the Boolean Hidden Matching problem, only Bob has the linear function M , while Alice must send her message based only on x . In our communication problem each (hyper)edge in H corresponds to a player, and every player receives a linear sized set of edges, together with parities of a hidden string x over these edges. Similarly to the Boolean Hidden Matching problem, these parities are either correct (YES case) or flipped simultaneously. A crucial new component, however, is the fact that instead of M being held by the recipient alone, each player holds part of it, and the parts the players hold are correlated.

Analyzing such correlations is in fact necessary even if one only wants to prove a simple lower bound on the space complexity of ‘sampling-type’ algorithms for triangle counting. We show how to analyze such correlations when H is an arbitrary hypergraph through a purely combinatorial lemma. The weights lemma (Section III) is primarily concerned with the ability of the players to co-ordinate “weight” functions. This can be used to lower bound the space complexity of sampling-based protocols—we apply it to bound the Fourier coefficients of the referee’s posterior distribution on the players’ inputs when the players send *arbitrary* messages.

E. Related Work

The past decade has seen a large amount of work on the space complexity of graph problems in the streaming model of computation (see, e.g. the recent survey by McGregor [McG17]). The semi-streaming model of computation, which allows $\tilde{O}(n)$ space to process a graph on n vertices, has been extensively studied, with space efficient algorithms known for many fundamental graph problems such as spanning trees [AGM12a], sparsifiers [AG09], [KL11], [AGM12b], [KLM⁺14], matchings [AG11], [AG13], [GKK12], [Kap13], [GO12], [HRVZ15], [Kon15], [AKLY15], spanners [AGM12b], [KW14]. Beyond the semi-streaming model, it has recently been shown that it is sometimes possible to approximate the *cost* of the solution to a graph problem in the streaming model even when the amount of space available does not suffice to store the vertex set of the graph (e.g. [KKS14], [EHL⁺15], [CCE⁺15], [Cor17], [MV18], [PS18]). The problem of designing lower bounds for graph sketches has received a lot of attention recently due to the success of graph sketching as an approach to solving dynamic graph streaming problems (e.g., [LNW14b], [AKLY15], [AKL17]). Similarly to our approach, such lower bounds normally make use of the simultaneous communication model.

a) Subgraph counting.: The streaming subgraph counting problem was introduced in [BKS02] for the case where H is a triangle. This was followed by alternative algorithms in [BFL⁺06], [JG05]. The lower bounds in [BOV13] and [KP17] were achieved by reductions to one-way communication complexity problems, the indexing problem and the Boolean Hidden Matching problem, respectively. Triangle detection has also been studied as a pure communication problem, for instance [FGO17], as well as in the adjacency-list [KMPT10], [BFL⁺06], [MVV16], multi-pass [BOV13], [CJ14], and query models [ELRS15].

Work on counting non-triangle subgraphs includes [BFLS07], which presented an algorithm for counting copies of $K_{3,3}$, [BDGL08], which studied subgraphs of size 3 and 4, [MMPS11], which studied cycles of arbitrary size, and [KMSS12], which studied arbitrary subgraphs. The problem has also been studied in the query [JSP15], [ANRD15], [PSV17] and distributed [ESBD15], [ESBD16] models.

b) Join size estimation.: The size of a database join can be viewed as a “labeled” version of hypergraph counting, where each vertex of G can only match a particular vertex (“attribute”) of H , and each hyperedge of G can only match a particular hyperedge (“relation”) of H . (Both our upper and lower bounds apply in this labeled setting.) In [AGM08] it was shown for a database

G with m hyperedges, the size of the join given by a query H can be up to $\Theta(m^\rho)$, where ρ is the fractional edge cover number of H .

This result is from a very different regime from ours because it involves very dense graphs and ours involves sparse ones. But one intriguing connection is through the $\Omega(m/T^{2/3})$ lower bound given in [KP17] for the restricted class of “triangle sampling” algorithms. Generalizing that proof for arbitrary H would use [AGM08] to get a lower bound of $\Omega(m(\frac{1}{Tn^{2\rho-|V|}})^{1/\rho})$, as opposed to our $\Omega(m/T^{1/\tau})$ bound. These are the same for some graphs, such as odd cycles, where $\rho = \tau = |V|/2$, and $\Omega(m/T^{1/\tau})$ is stronger for sparse graphs, but the two bounds are incomparable in general. It seems possible that the sample complexity for dense graphs will depend on ρ in some fashion.

II. PROOF OVERVIEW

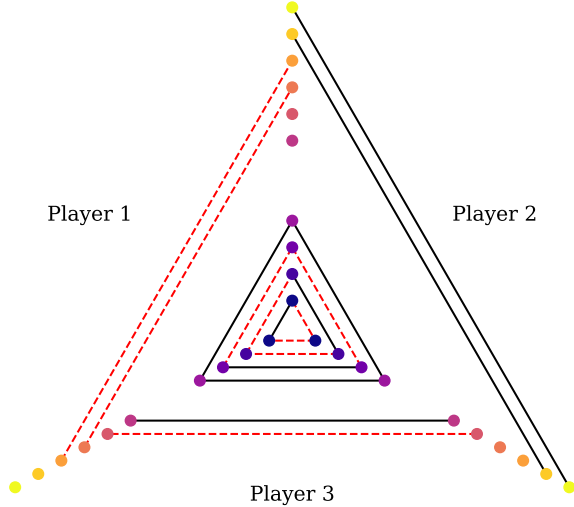
A. Lower Bound

For a fixed graph H with fractional vertex cover τ , we prove an $\Omega(n/T^{1/\tau})$ lower bound for determining whether a constant-degree graph on $\Theta(n)$ vertices has 0 or $\Theta(T)$ copies of H . For illustration, in this section we focus on the case where H is a triangle.


We consider the three-party simultaneous-message communication problem illustrated in Figure 1, where each player is associated with an edge of H . First, we construct a set of $N = \Theta(n)$ vertices V_u for each vertex $u \in H$. The player associated with edge $e = (u, v)$ receives an input consisting of n disjoint edges on $V_u \times V_v$, along with binary labels associated with each edge. We are guaranteed that the three players’ inputs collectively contain T triangles, with all the other edges disjoint. Each set of vertices are randomly permuted so that the players do not know which of their edges participate in triangles.

We also guarantee that the XOR of the labels associated with a triangle is the same for every triangle—either every triangle has an even number of 1s, or an odd number. The goal of the players is to send messages to a referee who knows the edges but not their labels, and for the referee to figure out if every triangle has an odd number of 1 labels.

We will show that a uniformly random instance of this problem requires $\Omega(n/T^{2/3})$ communication for the referee to succeed $2/3$ of the time. At the same time, it directly reduces to triangle counting: each player sketches their edges labeled 0 and sends the sketch to the referee. The referee adds the linear sketches up to get a sketch of all 0-labeled edges in the graph. This subgraph either contains zero triangles (if every triangle has an



(a) Input for triangle-counting before permutation. Each player e sees n edges with associated binary labels (pictured as solid/dashed). The edges match up into T triangles (center) and $n - T$ isolated edges (outside). The goal is to determine whether every triangle has an even number of solid edges, or an odd number.



(b) In the hard distribution, we randomly permute the vertices on top, on left, and on right. Each player sees their edges, with associated labels, in a random order; they do not see the pre-permutation vertex identities (represented by color).

Fig. 1: Lower bound instance for triangle counting

odd number of 1s) or very close to $T/4$ (otherwise), so successfully counting triangles will distinguish the two cases.

Our lower bound for the communication problem consists of two main pieces. First, we give a combinatorial lemma that bounds the players' ability to co-ordinate any assignment of "weights" to edges or subsets of edges, based on the structure of the graph. Then we use Fourier-analytic techniques to extend this to a lower bound of the communication required by any protocol for the problem.

a) Combinatorial lemma: One approach the players could take to solve the problem would be for each player to look at their n edges and pick a p fraction to send to the referee. If the referee receives a complete triangle, he can solve the problem. What is the expected number of triangles the referee receives, if the players coordinate optimally?

The naive solution where players pick independently at random would yield $p^3 T$ triangles. Vertex sampling—picking a \sqrt{p} fraction of vertices, for example those of smallest index, and only sampling edges between picked vertices—increases this to $p^{3/2} T$. In [KP17], a simple counting argument showed that "oblivious" strategies, which decide whether to sample an edge based only on the edge and not the rest of the player's input, cannot do better than this.

The combinatorial lemma we need for the Fourier-analytic proof is a stronger, generalized version of this sampling lemma. It considers players that receive some private randomness ψ_e and partially-shared randomness ϕ_u for each $u \in e$, and output an arbitrary deterministic function

$$g_e = g_e(\psi_e, (\phi_u)_{u \in e}) \in [-1, 1]$$

of their inputs. If the ϕ_u are fully independent and the ψ_e are $(|E| - 1)$ -wise independent, and

$$\max_e \mathbb{E}_{\psi, \phi} [g_e^2] \leq p$$

for some p , then we show:

$$\mathbb{E}_{\psi, \phi} \left[\prod_e g_e \right] \lesssim p^{3/2}. \quad (1)$$

To relate this to sampling, we note that the communication problem in Figure 1 can be constructed with randomness in the form above: ϕ_u contains the permutation of the vertices V_u associated with u , and ψ_e contains player e 's edge labels x_e (which are 2-wise independent) and the random order π_e in which they see their edges. Consider picking a random triangle edge $s \in [T]$ and adding to ψ_e the index of s in player e 's list. (One can show that ψ_e remains pairwise independent,

despite the shared dependence on s .) If we only allow $g_e \in \{0, 1\}$, then we can think of g_e as the event that player e samples their edge in the s th triangle. The condition on $\mathbb{E}[g_e^2]$ says that each player can pick at most a p fraction of their edges, on average over their inputs. The conclusion is that the expected fraction of triangles completely sampled is at most $p^{3/2}$.

This combinatorial lemma is different from the simple sampling lemma of [KP17] in several ways. First, it allows players to look at their entire inputs before deciding which edges to keep. Second, while the lemma of [KP17] was based on defining a fixed subset of edges to keep (so the number kept depended only on which edges were seen), in our lemma the players only need to keep a p fraction of their inputs *on average*, but the players do not have shared randomness. If they had shared randomness, there would be a trivial counterexample: with probability p every player samples every edge, giving the referee pT triangles in expectation. Without shared randomness, they can still use the correlation of their input for nontrivial algorithms: for example, for $p = 2^{1-n}$ they can send their entire input if every edge has the same label; because of the promise, if two players sample their inputs then the third is much more likely to. But this coordination is less effective than vertex sampling.

The combinatorial lemma is also more general, in ways that are important for the Fourier-analytic component of the proof. It allows for “fractional” choices of edges $g_e \in [-1, 1]$, with an ℓ_2 constraint. This allows for alternative competitive strategies—for example, placing \sqrt{p} weight on every edge also yields $p^{3/2}$ —but no strictly better ones. Additionally, the lemma will extend to cases where instead of placing weight on individual edges, the players place weight on *sets* of k triangle edges for some $k \geq 1$. These will correspond to weight k Fourier coefficients.

b) Fourier-analytic argument: Our approach for lower bounding the communication problem is inspired by [GKK⁺07]. Let $x_e \in \{0, 1\}^n$ be the player e ’s labels before permutation (i.e., from Figure 1a). We consider the referee’s posterior distribution p on the triangle parities, $x_1^{1:T} \oplus x_2^{1:T} \oplus x_3^{1:T}$. p is supported on $\{0^T, 1^T\}$, and our goal is to show that it is nearly uniform.

First, we express the referee’s posterior distribution on the labels $x = (x_e)_{e \in E}$. Let $f_e(y) = 1$ if player e ’s message to the referee is consistent with $x_e = y$, and 0 otherwise, and let $f : \{0, 1\}^{|E|n} \rightarrow \{0, 1\}$ be given by $f(x) = \prod_e f_e(x_e)$. The referee has two constraints on x : the message consistency constraint $f(x)$, and a parity constraint $q(x)$. His posterior distribution is uniform on $\text{supp}(fq)$.

The first observation we make relates the referee’s

total variation distance to the Fourier spectrum of fq . For indicator functions $g : \{0, 1\}^m \rightarrow \{0, 1\}$ it makes sense to consider the renormalized Fourier transform

$$\tilde{g}(s) := \frac{2^m}{|\text{supp}(g)|} \hat{g}(s) = \mathbb{E}_{x \in \text{supp}(g)} [(-1)^{s \cdot x}].$$

With this normalization, we observe that

$$\Delta := \|p - \mathcal{U}(\{0^T, 1^T\})\|_{TV} = \frac{1}{2} \tilde{f}q(e_1, e_1, e_1)$$

where $e_1 = (1, 0, \dots, 0) \in \{0, 1\}^n$. Using the structure of q ’s spectrum and the Fourier convolution theorem, we turn this into

$$\Delta = C \sum_{\substack{t \in \{0, 1\}^T \\ |t| \equiv 1 \pmod{2}}} \prod_e \tilde{f}_e(t0^{n-T})$$

where C is a normalising factor that is constant in expectation over x_1, x_2, x_3 . The combinatorial lemma applied to \tilde{f}_e lets us bound the sum for a fixed $|t| = k$ (in expectation over the input). The bound is, for some constant $D > 0$,

$$D \binom{T}{k} \left(\max_e \frac{1}{\binom{n}{k}} \mathbb{E} \left[\sum_{\substack{s \in \{0, 1\}^n \\ |s|=k}} \tilde{f}_e(s)^2 \right] \right)^{3/2}$$

which can be bounded in terms of the players’ c bits of communication by the KKL lemma (for small k) and Parseval’s identity (for high k). See the full version of the paper [KKP18] for statements of the bounds used. The dominant term when summing over k is $k = 1$ whence we get that the referee’s total variation distance has

$$\mathbb{E}[\Delta] \lesssim T(c/n)^{3/2}.$$

This implies the players must send at least $n/T^{2/3}$ bits to distinguish the two cases with significant probability.

c) Changes for non-triangle graphs: For counting general (hyper)graphs, the combinatorial lemma as described gives a bound of $p^{MVC_{1/2}(H)}$, where $MVC_{1/2}(H)$ is a “modified” fractional vertex cover in which weight can be placed directly on edges for half price. For odd cycles such as triangles, $MVC_{1/2}(H)$ equals the non-modified fractional vertex cover τ , giving the desired $\Omega(n/T^{1/\tau})$ bound.

For other graphs, such as the length-3 path, $MVC_{1/2}(H)$ can be less than τ leading to a suboptimal result. For these graphs we use a somewhat different proof, in which the referee is identified with a particular edge e^* in the graph. The other players’ inputs are then completely independent of one another, with no promise

on the XOR of their labels. We follow a slightly different Fourier-analytic approach that requires bounding

$$\mathbb{E} \left[\prod_{e \neq e^*} \tilde{f}_e^2 \right].$$

rather than $\mathbb{E} \left[\prod_e \tilde{f}_e \right]$. We apply the combinatorial lemma to the \tilde{f}_e^2 , on which we have an ℓ_1 constraint, giving us the bound

$$\mathbb{E} \left[\prod_{e \neq e^*} \tilde{f}_e^2 \right] \leq p^{MVC_1(H \setminus e^*)}$$

where the exponent is the non-modified fractional vertex cover of $H \setminus e^*$. This gives a lower bound of $\Omega(n/T^{1/MVC_1(H \setminus e^*)})$. For every connected (non-hyper-)graph H that is neither an odd cycle nor a single edge, this equals $\Omega(n/T^{1/\tau})$ for at least one e^* .

For graphs that *are* single edges or odd cycles, $MVC_{1/2}(H) = \tau$, and so the combination of these bounds gives $\Omega(n/T^{1/\tau})$ for every connected graph with more than one edge. For hypergraphs, the individual lower bounds still hold, but their maximum is not necessarily $\Omega(n/T^{1/\tau})$.

d) Dependence on ε : The above approach is a lower bound for distinguishing T triangles from 0 triangles. Distinguishing T triangles from $(1 - \varepsilon)T$ triangles should require more space for small ε . In the non-promise version of the proof used for graphs that are not odd cycles, we use the noise operator, an operator that takes a binary function and “noises” it by randomly flipping input bits, to get an $\varepsilon^{-2/\tau}$ dependence. In the promise version, the bound we get is only $\varepsilon^{-1/\tau}$.

B. Upper Bound

For purposes of this overview, we describe a sampling algorithm for the “labeled” version of the problem used in join size estimation, where edges and vertices in G correspond to edges and vertices in H , and we only want to count subgraphs with matching labels. Since H has constant size, we can solve the non-labeled version by trying many random labelings.

Consider a hypergraph H with minimal fractional vertex cover f , so $f(u) \in [0, 1]$ for each vertex $u \in V_H$ and $\sum_u f(u) = \tau$. Let $\chi : V_G \rightarrow V_H$ be the labels. For a parameter $p \in (0, 1)$ to be determined later, we sample each vertex $v \in V_G$ with probability $p^{f(\chi(v))}$, and we keep a hyperedge $e \in E_G$ if and only if we sample all $v \in e$.

The chance we keep any given copy of H is $\prod_{u \in H} p^{f(u)} = p^\tau$. Therefore, if we set $p = 100/T^{1/\tau}$,

the expected number of copies of H we see will be $p^\tau T \geq 100$. On the other hand, the chance we keep any single edge e is $\prod_{v \in e} p^{f(\chi(v))} \leq p$, because f covers the edge associated with e and so $\sum_{v \in e} f(\chi(v)) \geq 1$. This gives an algorithm with $O(mp) = O(m/T^{1/\tau})$ space that sees $p^\tau T \geq 100$ copies of H in expectation; from this T can be estimated.

The only tricky bit is to show that the variance of the number of sampled copies of H is small. We bound this in terms of the maximum degree of G and the maximum correlation between sampling two copies of H in G . If this correlation is 1 as can happen in general, the sampling algorithm only works for constant-degree graphs. However, if the vertex cover places at least 0.5 weight on each vertex of H , then the correlation is at most $\sqrt{p} = \Theta(1/T^{1/(2\tau)})$. This lets the algorithm work for degree $O(T^{1/(2\tau)})$ graphs. In the case of triangles, this $O(T^{1/3})$ degree bound is the correct regime for $O(m/T^{2/3})$ samples to be possible—above this threshold, the maximum number of triangles sharing a single vertex can be larger than $T^{2/3}$ and so the $\Omega(m\sqrt{\Delta_V}/T)$ lower bound of [KP17] precludes it.

III. THE WEIGHTS LEMMA

Definition 1. For a weighted hypergraph $H = (V, E)$ with weights $w : E \rightarrow [0, \infty]$, we define the λ -modified fractional vertex cover number $MVC_\lambda(H, w)$ to be:

$$MVC_\lambda(H, w) = \min_f \left(\sum_{v \in V} f(v) + \lambda \sum_{e \in E} f(e) \right)$$

over all $f : V \cup E \rightarrow [0, \infty]$ satisfying

$$\sum_{v \in e} f(v) + f(e) \geq w(e) \quad \forall e \in E.$$

When w is omitted, it is assumed that $w(e) = 1$ for all e . We note that $MVC_\lambda(H)$ equals the standard fractional vertex cover number of H whenever $\lambda \geq 1$ and H has no empty hyperedges.

Definition 2 (Totally disconnected hypergraph). For a hypergraph $H = (V_H, E_H)$ we say that H is totally disconnected if edges of H are pairwise disjoint, i.e. for every $a, b \in E$ one has $a \cap b = \emptyset$.

Lemma 3. Consider any hypergraph $H = (V, E)$ and weight function $w : E \rightarrow [0, \infty]$. Suppose that H is not totally disconnected as per Definition 2, i.e., there exist $a, b \in E$ such that $a \cap b \neq \emptyset$.

Consider any collection of random variables g_e (for $e \in E$) that can be expressed as deterministic functions of some random variables ϕ_u (for $u \in V$) that are

independent, and ψ_e (for $e \in E$) that are independent of the ϕ_u and $(|E| - 1)$ -wise independent themselves, i.e.,

$$g_e = g_e(\psi_e, (\phi_u)_{u \in e}).$$

Let $q \in \{1, 2\}$ and $0 < p < 1$, and suppose for all $e \in E$ that $|g_e| \leq 1$ always and that

$$\mathbb{E}[|g_e|^q] \leq p^{w(e)}.$$

Then

$$\mathbb{E} \left[\prod_{e \in E} g_e \right] \leq d^{|V|} p^{MVC_{1/q}(H, w)}$$

where $MVC_{1/q}(H, w)$ is the modified fractional vertex cover number per Definition 1, and d is the maximum of 1 and the greatest degree of a vertex $v \in V$.

The proof is deferred to the full version of the paper [KKP18].

IV. HYPERGRAPH COUNTING WITH A PROMISE

In both of the games that follow, we will assume that the players are deterministic. This is without loss of generality by Yao's minimax principle since the inputs are sampled from a fixed distribution.

A. Game

We will define a $|E| + 1$ -player game $\text{PromiseCounting}(H, n, T, \varepsilon)$ (with H a hypergraph, $n, T \in \mathbb{N}, \varepsilon \in \{1/T, 2/T, \dots, 1\}$), as follows: There is one referee, who receives messages from every other player. No other communication takes place. Each player besides the referee corresponds to an edge $e \in E$.

Let $N = T + (n - T)|E|$. For each edge $e \in E$, let $L_e \subset [N]$ be an n -element set containing $[T]$ and $n - T$ elements disjoint from every other L_e , so that if $a \neq b$, $L_a \cap L_b = [T]$, and $\bigcup_{e \in E} L_e = [N]$. For each $e \in E$, let $\rho_e : [n] \rightarrow L_e$ be a fixed bijection such that $\rho_e|_{[T]}$ is the identity.

An instance of $\text{PromiseCounting}(H, n, T, \varepsilon)$ is as follows:

- For each edge $e \in E$:
 - A string $x_e \in \{0, 1\}^n$.
 - A permutation π_e on L_e .
- For each vertex $v \in V$:
 - A permutation π_v on $[N]$.
- A string $\tau \in \{0^{\varepsilon T}, 1^{\varepsilon T}\}$

The players have the following promise:

$$\bigoplus_{e \in E} x_e^{1:\varepsilon T} = \tau$$

We will write X for the strings $(x_e)_{e \in E}$, Π_E for the permutations $(\pi_e)_{e \in E}$, and Π_V for the permutations $(\pi_v)_{v \in V}$. They have access to the following information:

- For each player $e \in E$:
 - $x_e \rho_e \pi_e$
 - $(\pi_v(\pi_e^{-1}(i)))_{v \in e, i \in L_e}$
- For the referee:
 - $\Pi = (\Pi_E, \Pi_V)$

Given the messages received from the players, the referee's task will be to determine whether $\tau = 0^{\varepsilon T}$ or $\tau = 1^{\varepsilon T}$.

B. Hard Instance

We will lower bound the complexity of this problem under the following hard input distribution: τ is chosen uniformly from $\{0^{\varepsilon T}, 1^{\varepsilon T}\}$, and then the strings $(x_e)_{e \in E}$ are chosen uniformly from:

$$\left\{ (x_e)_{e \in E} \in \{0, 1\}^{|E|} : \bigoplus_{e \in E} x_e^{1:\varepsilon T} = \tau \right\}$$

Every permutation π_u, π_e is chosen uniformly at random and independently of each other and the strings.

C. Lower Bound

For each player e , write $m_e(x_e \rho_e \pi_e, (\pi_v(\pi_e^{-1}(i)))_{v \in e, i \in L_e})$ for the message the player sends to the referee on seeing $x_e \rho_e \pi_e$ and $(\pi_v(\pi_e^{-1}(i)))_{v \in e, i \in L_e}$.

Theorem 4. *Let H be a connected hypergraph with more than one edge. Let $c \in [n]$. Suppose that, for all inputs (X, Π) to the game, no player sends a message of more than c bits, and suppose that $\varepsilon T \leq n/10$.*

Let $p : \{0^{\varepsilon T}, 1^{\varepsilon T}\} \rightarrow [0, 1]$ be the referee's posterior distribution on τ . Let ν be the distribution of $\mathcal{U}(\{0^{\varepsilon T}, 1^{\varepsilon T}\})$, the uniform distribution on the two-element set $\{0^{\varepsilon T}, 1^{\varepsilon T}\}$. Let $\mu = MVC_{1/2}(H)$, and let $0 < \delta < 1$.


There exists a constant γ , depending only on H , such that, if $c \leq \gamma \frac{n}{(\delta^2 \varepsilon T)^{1/\mu}}$:

$$\mathbb{E}_{X, \Pi} [||p - \nu||_{TV}] \leq \delta$$


Corollary 5. *Let H be a connected hypergraph with more than one edge. Let $c \in [n]$. Suppose that, for all inputs (X, Π) to $\text{PromiseCounting}(H, n, T, \varepsilon)$, no player sends a message of more than c bits.*

Let $\mu = MVC_{1/2}(H)$, and let $0 < \delta < 1$.

There exists a constant γ , depending only on H , such that, if $c \leq \gamma \frac{n}{(\delta^2 \varepsilon T)^{1/\mu}}$, the players succeed at the game with probability at most $1/2 + \delta$.



(a) The player's instance ignoring the permutations. The x_e are the indices of red edges, read from inside out: $x_1 = [0, 1, 1, 0, 1, 1]$, $x_2 = [1, 0, 1, 0, 0, 0]$, $x_3 = [1, 1, 0, 0, 0, 1]$



(b) The hard distribution permutes each set of vertices. The players see their edges and associated labels, but not the vertex colors (which represent the pre-permutation identities).

Player 1			Player 2			Player 3		
x	u	v	x	u	v	x	u	v
0	E	E	0	B	E	0	D	C
1	H	J	0	D	D	0	G	A
1	G	H	1	E	F	1	H	I
1	B	C	0	I	I	0	A	D
1	D	G	1	G	A	1	F	E
0	C	D	0	J	B	1	B	H

(c) Each player's input consists of their edges in (b) in a random order. u represents the vertex counterclockwise of the player, and v represents the vertex clockwise.

Fig. 2: Encoding of lower bound instance for triangle counting

The proofs are deferred to the full version of the paper [KKP18].

V. HYPERGRAPH COUNTING WITH NO PROMISE

We also study a slightly different game, in which the XOR of the strings x_e may take any value, but the referee has a “target” string that is guaranteed to either be $\bigoplus_{e \in E} x_e^{1:T}$ or $\bigoplus_{e \in E} \overline{x_e^{1:T}}$. This will give a different, and in general incomparable lower bound for subgraph counting. For further details, see the full version of the paper [KKP18].

VI. LINEAR SKETCHING LOWER BOUND

Definition 6. Let \mathcal{A} be a randomized graph streaming algorithm, and let \mathbb{S} be the set of possible states S of \mathcal{A} . We will say \mathcal{A} has composable state if, for any fixed random seed for \mathcal{A} , there is a function $c : \mathbb{S} \times \mathbb{S} \rightarrow \mathbb{S}$ such that, if S_1 is the state of \mathcal{A} after receiving the stream of edges E_1 as input, and S_2 is the state of \mathcal{A} after receiving the stream of edges E_2 as input, $c(S_1, S_2)$ is the state of \mathcal{A} after receiving the concatenation of E_1 and E_2 as input.

Theorem 7. Let $H = (V, E)$ be a (fixed) connected hypergraph with $|E| > 1$. Let $T \in \mathbb{N}$, $\varepsilon \in (1/\sqrt{T}, 1]$. Let \mathcal{A} be a graph streaming algorithm that can distinguish between graphs G presented as a stream of edges with at least T copies of H and graphs with at most $(1 - \varepsilon)T$ copies of H with probability $99/100$, provided G has no more than m edges. Let $S(m)$ be the maximum space usage of \mathcal{A} across all m -edge inputs.

Furthermore, let \mathcal{A} have composable state. Then, for all $m \geq O(T)$:

$$S(m) = \Omega \left(\max \left(\frac{m}{(\varepsilon T)^{1/\mu_2}}, \frac{m}{(\varepsilon^2 T)^{1/\mu_1}} \right) \right)$$

where $\mu_2 = MVC_{\frac{1}{2}}(H)$ and $\mu_1 = \max_{e \in E} MVC_1(H \setminus e)$, with constants that may depend on H but nothing else.

We will prove this by reductions to PromiseCounting and Counting.

Lemma 8. $\forall m \geq O(T)$, $S(m) = \Omega \left(\frac{m}{(\varepsilon T)^{1/\mu_2}} \right)$

Lemma 9. $\forall m \geq O(T)$, $\forall e^* \in E$, $S(m) = \Omega \left(\frac{m}{(\varepsilon^2 T)^{1/\mu_1}} \right)$ where $\mu_1 = MVC_1(H \setminus e^*)$.

The proofs are deferred to [KKP18].

Theorem 7 then follows directly from the previous two lemmas.

To prove this gives tight bounds (for ε constant) for all 2-uniform hypergraphs (that is, all graphs), we will need the following lemma on graph covers:

Lemma 10. *Let $G = (V, E)$ be a connected graph with $|E| > 1$. Then:*

$$\max(MVC_2(G), \max_{e \in E} MVC_1(G \setminus e)) = MVC_1(G)$$

is the standard fractional vertex cover of G .

The proof is deferred to [KKP18].

Corollary 11. *Let $H = (V, E)$ be a connected graph with $|E| > 1$. Let $\varepsilon \in (0, 1]$, $T \in \mathbb{N}$. Let \mathcal{A} be a graph streaming algorithm that can distinguish between graphs G presented as a stream of edges with T copies of H , graphs with $(1 - \varepsilon)T$ copies of H with probability 99/100, provided G has no more than m edges. Let $S(m)$ be the maximum space usage of \mathcal{A} across all m -edge inputs.*

Furthermore, let \mathcal{A} have composable state. Then:

$$\forall m \geq O(T), S(m) = \Omega\left(\frac{m}{(\varepsilon T)^{1/\tau}}\right)$$

where τ is the fractional vertex cover of H , and the constant factor may depend on H but nothing else.

VII. UPPER BOUND

Our main result is

Theorem 12. *For every hypergraph $H = (V_H, E_H)$, $\varepsilon \in (0, 1)$ there exists a sketching algorithm that, for any hypergraph $G = (V_G, E_G)$ on n vertices with degrees bounded by d , approximates the number of copies of H in G to within a $1 + \varepsilon$ multiplicative factor with probability at least 99/100 using space $s \leq C \cdot \varepsilon^{-2/\tau} \cdot m T^{-1/\tau}$, where T is the number of copies of H in G and τ is the fractional vertex cover of H and C is a constant that depends on H .*

For graphs we get a more powerful result, which allows the graph G to have higher degrees:

Theorem 13. *For every graph $H = (V_H, E_H)$ that admits a minimum vertex cover that assigns nonzero weight to every vertex, for every $\varepsilon \in (0, 1)$ there exists a sketching algorithm that, for any graph $G = (V_G, E_G)$ on n vertices with degrees bounded by $d \leq C' \varepsilon^{1/\tau} T^{1/(2\tau)}$, approximates the number of copies of H in G to within a $1 + \varepsilon$ multiplicative factor with probability 99/100 using space $C \cdot \varepsilon^{-2/\tau} \cdot m T^{-1/\tau}$, where T is the number of copies of H in G and τ is the fractional vertex cover of H , and $C, C' > 0$ are constants that depend only on H .*

This result requires the minimum vertex cover to assign nonzero weight to every vertex; this happens for cycles but not stars.

Consider the fractional vertex cover of H

$$\begin{aligned} \min \sum_{a \in V_H} x_a \\ \text{s.t. } \sum_{a \in e} x_a \geq 1 \text{ for all } e \in E_H, \end{aligned} \quad (2)$$

let $x^* \in \mathbb{R}^{V_H}$ denote an optimal solution and let τ denote its value.

Fix a mapping $\chi : V_G \rightarrow V_H$ (see Algorithm 1, line 3). For a subset $S \subseteq V_G$ we write $\chi(S) \sim H$ if the subgraph induced by S equipped with labels $\chi(S)$ contains a copy of H , i.e. for every $a \subseteq S$ one has that if $\chi(a) \in E_H$, then $a \in E_G$. Note that, if $A(H)$ is the number of automorphisms of H , the probability that a randomly chosen χ will give $\chi(S) \sim H$ is $A(H)/k^k$.

Algorithm 1 Subgraph counting by vertex sampling

- 1: **procedure** SAMPLE(H, p) \triangleright Input: hypergraph H , sampling probability p
 - 2: Compute minimum vertex cover x^* in H
 - 3: $\chi \sim UNIF([k]^{V_G})$ \triangleright Random mapping of V_G to $V_H = [k]$
 - 4: **for** $u \in V_G$ **do**
 - 5: $a \leftarrow \chi(u)$
 - 6: $X_u \leftarrow$ independent Bernoulli r.v. with mean $p^{x_a^*}$
 - 7: **end for**
 - 8: $E' \leftarrow \{e \in E_G : \chi(e) = |e| \text{ and } \prod_{u \in e} X_u = 1\}$ \triangleright Keep colorful edges only
 - 9: $Z \leftarrow k^k \cdot p^{-\tau} \cdot \sum_{S \subseteq V_G : \chi(S) \sim H} \prod_{u \in S} X_u$ \triangleright Knowing E' and χ suffices to compute Z
 - 10: **return** $Z/A(H)$
 - 11: **end procedure**
-

This result follows from a variance-based analysis. For details, see the full version [KKP18].

ACKNOWLEDGMENTS

This work was done in part while John Kallaugher and Eric Price were visiting the Simons Institute for the Theory of Computing. Michael Kapralov is supported in part by ERC Starting Grant 759471-Sublinear.

REFERENCES

- [AG09] K. Ahn and S. Guha. Graph sparsification in the semi-streaming model. *ICALP*, pages 328–338, 2009.
- [AG11] K. Ahn and S. Guha. Linear programming in the semi-streaming model with application to the maximum matching problem. *ICALP*, pages 526–538, 2011.
- [AG13] K. Ahn and S. Guha. Access to data and number of iterations: Dual primal algorithms for maximum matching under resource constraints. *CoRR*, abs/1307.4359, 2013.
- [AGM08] Albert Atserias, Martin Grohe, and Dániel Marx. Size bounds and query plans for relational joins. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 739–748. IEEE, 2008.
- [AGM12a] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. *SODA*, pages 459–467, 2012.
- [AGM12b] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketching: Sparsification, spanners, and subgraphs. *PODS*, 2012.
- [AHLW16] Yuqing Ai, Wei Hu, Yi Li, and David P Woodruff. New characterizations in turnstile streams with applications. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 50. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [AKL17] Sepehr Assadi, Sanjeev Khanna, and Yang Li. On estimating maximum matching size in graph streams. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1723–1742, 2017.
- [AKLY15] Sepehr Assadi, Sanjeev Khanna, Yang Li, and Grigory Yaroslavtsev. Tight bounds for linear sketches of approximate matchings. *CoRR*, 2015.
- [ANRD15] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, and Nick Duffield. Efficient graphlet counting for large networks. In *Proceedings of the 2015 IEEE International Conference on Data Mining (ICDM)*, ICDM '15, pages 1–10, Washington, DC, USA, 2015. IEEE Computer Society.
- [BDGL08] I. Bordini, D. Donato, A. Gionis, and S. Leonardi. Mining large networks with subgraph counting. In *2008 Eighth IEEE International Conference on Data Mining*, pages 737–742, Dec 2008.
- [BFL⁺06] Luciana S Brioli, Gereon Frahling, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Christian Sohler. Counting triangles in data streams. In *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 253–262. ACM, 2006.
- [BFLS07] Luciana S. Brioli, Gereon Frahling, Stefano Leonardi, and Christian Sohler. Estimating clustering indexes in data streams. In Lars Arge, Michael Hoffmann, and Emo Welzl, editors, *Algorithms – ESA 2007*, pages 618–632, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- [BKS02] Ziv Bar-Yossef, Ravi Kumar, and D. Sivakumar. Reductions in streaming algorithms, with an application to counting triangles in graphs. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '02*, pages 623–632, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.
- [BOV13] Vladimir Braverman, Rafail Ostrovsky, and Dan V. I. How hard is counting triangles in the streaming model? In *Automata, Languages, and Programming*, pages 244–254. Springer, 2013.
- [CCE⁺15] Rajesh Hemant Chitnis, Graham Cormode, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Andrew McGregor, Morteza Monemizadeh, and Sofya Vorotnikova. Kernelization via sampling with applications to dynamic graph streams. *CoRR*, abs/1505.01731, 2015.
- [CJ14] Graham Cormode and Hossein Jowhari. A second look at counting triangles in graph streams. *Theoretical Computer Science*, 552:44–51, 2014.
- [CMR05] Graham Cormode, S. Muthukrishnan, and Irina Rozenbaum. Summarizing and mining inverse distributions on data streams via dynamic inverse sampling. In *Proceedings of the 31st International Conference on Very Large Data Bases, VLDB '05*, pages 25–36. VLDB Endowment, 2005.
- [Cor17] The sparse awakens: Streaming algorithms for matching size estimation in sparse graphs. In *25th Annual European Symposium on Algorithms, ESA 2017, September 4-6, 2017, Vienna, Austria*, pages 29:1–29:15, 2017.
- [EHL⁺15] Hossein Esfandiari, Mohammad Taghi Hajiaghayi, Vahid Liaghat, Morteza Monemizadeh, and Krzysztof Onak. Streaming algorithms for estimating the matching size in planar graphs and beyond. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1217–1233, 2015.
- [ELRS15] Talya Eden, Amit Levi, Dana Ron, and C. Seshadhri. Approximately counting triangles in sublinear time. In *Proceedings of the 56th FOCS*, pages 614–633. IEEE, 2015.
- [ESBD15] Ethan R. Elenberg, Karthikeyan Shanmugam, Michael Borokhovich, and Alexandros G. Dimakis. Beyond triangles: A distributed framework for estimating 3-profiles of large graphs. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, pages 229–238, New York, NY, USA, 2015. ACM.
- [ESBD16] Ethan R. Elenberg, Karthikeyan Shanmugam, Michael Borokhovich, and Alexandros G. Dimakis. Distributed estimation of graph 4-profiles. In *Proceedings of the 21th International Conference on World Wide Web, WWW '16*, pages 483–493, Republic and Canton of Geneva, Switzerland, 2016. International World Wide Web Conferences Steering Committee.
- [FGO17] Orr Fischer, Shay Gershtein, and Rotem Oshman. On the multiparty communication complexity of testing triangle-freeness. *CoRR*, abs/1705.08438, 2017.
- [GKK⁺07] Dmitry Gavinsky, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald de Wolf. Exponential separations for one-way quantum communication complexity, with applications to cryptography. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, STOC '07*, pages 516–525, New York, NY, USA, 2007. ACM.
- [GKK12] A. Goel, M. Kapralov, and S. Khanna. On the communication and streaming complexity of maximum bipartite matching. *SODA*, 2012.
- [GO12] Venkatesan Guruswami and Krzysztof Onak. Superlinear lower bounds for multipass graph processing. *CCC*, 2012.
- [HRVZ15] Zengfeng Huang, Božidar Radunović, Milan Vojnović, and Qin Zhang. Communication complexity of approximate maximum matching in distributed graph data. *STACS*, 2015.
- [JG05] Hossein Jowhari and Mohammad Ghodsi. New streaming algorithms for counting triangles in graphs. In *Computing and Combinatorics*, pages 710–716. Springer, 2005.
- [JSP15] Madhav Jha, C. Seshadhri, and Ali Pinar. Path sampling: A fast and provable method for estimating 4-vertex subgraph counts. In *Proceedings of the 24th International Conference on World Wide Web, WWW '15*, pages 495–505, Republic and Canton of Geneva, Switzerland, 2015. International World Wide Web Conferences Steering Committee.
- [JST11] Hossein Jowhari, Mert Sağlam, and Gábor Tardos. Tight bounds for lp samplers, finding duplicates in streams, and related problems. In *Proceedings of the Thirtieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '11*, pages 49–58, New York, NY, USA, 2011. ACM.

- [Kap13] Michael Kapralov. Better bounds for matchings in the streaming model. *SODA*, 2013.
- [KK15] Dmitry Kogan and Robert Krauthgamer. Sketching cuts in graphs and hypergraphs. *ITCS*, 2015.
- [KKP18] J. Kallaugher, M. Kapralov, and E. Price. The Sketching Complexity of Graph and Hypergraph Counting. *ArXiv e-prints*, August 2018.
- [KKS14] Michael Kapralov, Sanjeev Khanna, and Madhu Sudan. Approximating matching size from random streams. In *25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014.
- [KKS15] Michael Kapralov, Sanjeev Khanna, and Madhu Sudan. Streaming lower bounds for approximating MAX-CUT. In *26th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015.
- [KKS17] Michael Kapralov, Sanjeev Khanna, Madhu Sudan, and Ameya Velingker. $(1 + \Omega(1))$ -Approximation to MAX-CUT requires linear space. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1703–1722, 2017.
- [KL11] Jonathan A. Kelner and Alex Levin. Spectral sparsification in the semi-streaming setting. *STACS*, pages 440–451, 2011.
- [KLM⁺14] Michael Kapralov, Yin Tat Lee, Cameron Musco, Christopher Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. *FOCS*, 2014.
- [KMPT10] Mihail N. Kolountzakis, Gary L. Miller, Richard Peng, and Charalampos E. Tsourakakis. Efficient triangle counting in large graphs via degree-based vertex partitioning. In Ravi Kumar and Dandapani Sivakumar, editors, *Algorithms and Models for the Web-Graph*, pages 15–24, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [KMSS12] Daniel M. Kane, Kurt Mehlhorn, Thomas Sauerwald, and He Sun. Counting arbitrary subgraphs in data streams. In *Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming - Volume Part II, ICALP'12*, pages 598–609, Berlin, Heidelberg, 2012. Springer-Verlag.
- [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. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 475–486. Ieee, 2017.
- [Kon15] Christian Konrad. Maximum matching in turnstile streams. *CoRR*, abs/1505.01460, 2015.
- [KP17] John Kallaugher and Eric Price. A hybrid sampling scheme for triangle counting. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1778–1797. SIAM, 2017.
- [KW14] Michael Kapralov and David Woodruff. Spanners and sparsifiers in dynamic streams. *PODC*, 2014.
- [LNW14a] Yi Li, Huy L. Nguyen, and David P. Woodruff. Turnstile streaming algorithms might as well be linear sketches. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 174–183. ACM, 2014.
- [LNW14b] Yi Li, Huy L. Nguyen, and David P. Woodruff. Turnstile streaming algorithms might as well be linear sketches. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 174–183, 2014.
- [LW16] Yi Li and David P. Woodruff. On approximating functions of the singular values in a stream. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 726–739, 2016.
- [McG17] Andrew McGregor. Graph sketching and streaming: New approaches for analyzing massive graphs. In *Computer Science - Theory and Applications - 12th International Computer Science Symposium in Russia, CSR 2017, Kazan, Russia, June 8-12, 2017, Proceedings*, pages 20–24, 2017.
- [MMPS11] Madhusudan Manjunath, Kurt Mehlhorn, Konstantinos Panagiotou, and He Sun. Approximate counting of cycles in streams. In *Proceedings of the 19th European Conference on Algorithms, ESA'11*, pages 677–688, Berlin, Heidelberg, 2011. Springer-Verlag.
- [MV18] Andrew McGregor and Sofya Vorotnikova. A simple, space-efficient, streaming algorithm for matchings in low arboricity graphs. In *1st Symposium on Simplicity in Algorithms, SOSA 2018, January 7-10, 2018, New Orleans, LA, USA*, pages 14:1–14:4, 2018.
- [MVF16] Andrew McGregor, Sofya Vorotnikova, and Hoa T. Vu. Better algorithms for counting triangles in data streams. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS '16*, pages 401–411, New York, NY, USA, 2016. ACM.
- [MW10] Morteza Monemizadeh and David P. Woodruff. 1pass relative-error lp-sampling with applications. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '10*, pages 1143–1160, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.
- [PS18] Pan Peng and Christian Sohler. Estimating graph parameters from random order streams. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2449–2466, 2018.
- [PSV17] Ali Pinar, C. Seshadhri, and Vaidyanathan Vishal. Escape: Efficiently counting all 5-vertex subgraphs. In *Proceedings of the 26th International Conference on World Wide Web, WWW '17*, pages 1431–1440, Republic and Canton of Geneva, Switzerland, 2017. International World Wide Web Conferences Steering Committee.
- [PT12] Rasmus Pagh and Charalampos E Tsourakakis. Colorful triangle counting and a mapreduce implementation. *Information Processing Letters*, 112(7):277–281, 2012.
- [TKM11] Charalampos E Tsourakakis, Mihail N Kolountzakis, and Gary L Miller. Triangle sparsifiers. *J. Graph Algorithms Appl.*, 15(6):703–726, 2011.
- [TKMF09] Charalampos E Tsourakakis, U Kang, Gary L Miller, and Christos Faloutsos. Doulion: counting triangles in massive graphs with a coin. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 837–846. ACM, 2009.
- [VY11] Elad Verbin and Wei Yu. The streaming complexity of cycle counting, sorting by reversals, and other problems. *SODA*, pages 11–25, 2011.