

Massively Parallel Algorithms for Small Subgraph Counting

Amartya Shankha Biswas* Talya Eden†

Quanquan C. Liu‡ Slobodan Mitrović§ Ronitt Rubinfeld*

July 19, 2022

Abstract

Over the last two decades, frameworks for distributed-memory parallel computation, such as MapReduce, Hadoop, Spark and Dryad, have gained significant popularity with the growing prevalence of large network datasets. The Massively Parallel Computation (MPC) model is the de-facto standard for studying graph algorithms in these frameworks theoretically. Subgraph counting is one such fundamental problem in analyzing massive graphs, with the main algorithmic challenges centering on designing methods which are both scalable and accurate.

Given a graph $G = (V, E)$ with n vertices, m edges and T triangles, our first result is an algorithm that outputs a $(1 + \varepsilon)$ -approximation to T , with asymptotically *optimal round and total space complexity* provided any $S \geq \max(\sqrt{m}, n^2/m)$ space per machine and assuming $T = \Omega(\sqrt{m/n})$. Our result gives a quadratic improvement on the bound on T over previous works. We also provide a simple extension of our result to counting *any* subgraph of k size for constant $k \geq 1$. Our second result is an $O_\delta(\log \log n)$ -round algorithm for exactly counting the number of triangles, whose total space usage is parametrized by the *arboricity* α of the input graph. We extend this result to exactly counting k -cliques for any constant k . Finally, we prove that a recent result of Bera, Pashanasangi and Seshadhri (ITCS 2020) for exactly counting all subgraphs of size at most 5 can be implemented in the MPC model in $\tilde{O}_\delta(\sqrt{\log n})$ rounds, $O(n^\delta)$ space per machine and $O(m\alpha^3)$ total space.

In addition to our theoretical results, we simulate our triangle counting algorithms in real-world graphs obtained from the Stanford Network Analysis Project (SNAP) database. Our results show that both our approximate and exact counting algorithms exhibit improvements in terms of round complexity and approximation ratio, respectively, compared to two previous widely used algorithms for these problems.

1 Introduction

Estimating the number of small subgraphs, cliques in particular, is a fundamental problem in computer science, and has been extensively studied both theoretically and from an applied perspective. Given its importance, the task of counting subgraphs has been explored in various computational settings, e.g., sequential [AYZ97, Vas09, BHKK09], distributed and parallel [SV11, PT12, KPP+14, PSKP14, LQLC15], streaming [BYKS02, KMSS12, BC17, MVV16], and sublinear-time [ELRS17, ABG+18, AKK19, ERS20]. There are usually two perspectives from which subgraph counting is studied: first, optimizing the running time (especially relevant in the

*CSAIL, MIT, Cambridge, MA, USA, {asbiswas,slobo}@mit.edu, ronitt@csail.mit.edu

†Boston University/CSAIL, MIT, Cambridge, MA, USA, teden@mit.edu

‡Northwestern University, Evanston, IL, USA, {quanquan}@northwestern.edu

§UC Davis, Davis, CA, smitrovic@ucdavis.edu

sequential and sublinear-time settings) and, second, optimizing the space or query requirement (relevant in the streaming, parallel, and distributed settings). In each of these perspectives, there are two, somewhat orthogonal, directions that one can take. The first is *exact* counting. However, in most scenarios, algorithms that perform exact counting are prohibitive, e.g., they require too much space or too many parallel rounds to be implementable in practice.

Hence, the second direction of obtaining an *estimate/approximation* on the number of small subgraphs is both an interesting theoretical problem and of practical importance. If $H_{\#}$ is the number of subgraphs isomorphic to H , the main question in approximate counting is whether we can design algorithms that, under given resource constraints, provide approximations that concentrate well. This concentration is usually parametrized by $H_{\#}$ (and potentially some other parameters). In particular, most known results do not provide a strong approximation guarantee when $H_{\#}$ is very small, e.g., $|H_{\#}| = O(1)$. So, the main attempts in this line of work is to provide an estimation that concentrates well while imposing as small a lower bound on $H_{\#}$ as possible.

Due to ever increasing sizes of data stores, there has been an increasing interest in designing scalable algorithms. The *Massively Parallel Computation* (MPC) model is a theoretical abstraction of popular frameworks for large-scale computation such as MapReduce [DG08], Hadoop [Whi12], Spark [ZCF+10] and Dryad [IBY+07]. MPC gained significant interest recently, most prominently in building algorithmic toolkits for graph processing [GSZ11, LMSV11, BKS13, ANOY14, BKS14, HP15, AG15, RVW16, IMS17, CLM+20, Ass17, ABB+19, GGK+18, HLL18, BFU18, ASW18, BEG+18, BDH+19, BBD+19, BHH19, ASZ19, ASW19, GLM19, GKMS19, GU19, LMOS19, ILMP19, CFG+19, GKU19, GNT20]. Efficiency of an algorithm in MPC is characterized by three parameters: round complexity, the space per machine in the system, and the number of machines/total memory used. Our work aims to design efficient algorithms with respect to all three parameters and is guided by the following question:

How does one design efficient massively parallel algorithms for small subgraph counting?

1.1 The MPC Model

In this paper, we are working in the Massively Parallel Computation (MPC) model introduced by [KSV10, GSZ11, BKS13]. The model operates as follows. There exist \mathcal{M} machines that communicate with each other in synchronous rounds. The graph input is initially distributed across the machines in some organized way such that machines know how to access the relevant information via communication with other machines. During each round, the machines first perform computation locally without communicating with other machines. The computation done locally can be unbounded (although the machines have limited space so any reasonable program will not do an absurdly large amount of computation). At the end of the round, the machines exchange messages to inform the computation for the next round. The total size of all messages that can be received by a machine is upper bounded by the size of its local memory, and each machine outputs messages of sufficiently small size that can fit into its memory. If N is the total size of the data and each machine has S words of space, we are interested in the settings when S is sublinear in N . We use *total space* to refer to $\mathcal{M} \cdot S$, which is the total space that is available across all the machines.

1.2 Our Contributions

1.2.1 Triangle Counting

We provide a number of results for triangle counting in both the approximate and exact settings. Let $G = (V, E)$ be a graph with n vertices, m edges and T triangles. First we study the question

Problem	Work	MPC Rounds	Space Per Machine	Total Space	ALB
Exact Triangle Counting	[SV11]	2	$O(\sqrt{m})$	$O(m^{3/2})$	-
	[SV11]	1	$o(m)$	$\omega(m)$	-
	[CC11]	$O(n)$	$O(n)$	$O(m)$	-
	folklore	$O(\log n)$	$\Omega(\alpha^2)$	$O(m\alpha)$	-
	Ours	$\mathbf{O}_\delta(\log \log n)$	$\mathbf{O}(\mathbf{n}^\delta)$	$\mathbf{O}(\mathbf{m}\alpha)$	-
Approximate Triangle Counting	[PT12]	$O(1)$	$\Omega(m)$	$O(m)$	$\Omega(d_{avg})$
	[SPK13]	$O(1)$	$O(n^\delta)$	$O(m)$	$\Omega(\sum_{v \in V} \deg(v)^2)$
	Ours	$\mathbf{O}(1)$	$\tilde{\mathbf{O}}(\mathbf{n})$	$\tilde{\mathbf{O}}(\mathbf{m})$	$\mathbf{\Omega}(\sqrt{\mathbf{d}_{avg}})$

Table 1: Summary of our main MPC triangle counting results compared to previous work. Our results are **bolded**. “ALB” refers to the approximation lower bound on the number of triangles required to obtain a $(1 + \varepsilon)$ -approximation, with high probability. α is the arboricity of the input graph and is generally small (logarithmic) in real-world networks. Parameter $\delta > 0$ is *any* constant.

of approximately counting the number of triangles under the restriction that the round and total space complexities are essentially *optimal*, i.e., $O(1)$ and $\tilde{O}(m)$, where \tilde{O} hides $O(\text{poly } \log n)$ factors, respectively. Here and throughout, we use O_δ and O_ε to hide factors of δ and ε , respectively, where we consider constant factors of $\delta, \varepsilon > 0$ in this paper.

Our algorithm is surprisingly simple with a more complicated analysis, but improves on the previous best-known result by giving a $(1 + \varepsilon)$ -approximation, with high probability, while achieving a *quadratic* improvement on the number of triangles required to ensure this approximation. The specific bounds are given in [Table 1](#).

Theorem 1.1. *Let $G = (V, E)$ be a graph with n vertices, m edges, and let T be the number of triangles in G . Assuming*

$$(i) \ T = \tilde{\Omega}\left(\sqrt{\frac{m}{S}}\right), \quad (ii) \ S = \tilde{\Omega}\left(\max\left\{\frac{\sqrt{m}}{\varepsilon}, \frac{n^2}{m}\right\}\right),$$

there exists an MPC algorithm, using \mathcal{M} machines, each with local space S , and total space $\mathcal{M}S = \tilde{O}_\varepsilon(m)$, that outputs a $(1 \pm \varepsilon)$ -approximation of T , with high probability, in $O(1)$ rounds.

For $S = \Theta(n \log n)$ (specifically, $S > 100n \log n$) in [Theorem 1.1](#), we derive the following corollary.

Corollary 1.2. *Let G be a graph and T be the number of triangles it contains. If $T \geq \sqrt{d_{avg}}$, then there exists an MPC algorithm that in $O(1)$ rounds with high probability outputs a $(1 + \varepsilon)$ -approximation of T . This algorithm uses a total space of $\tilde{O}(m)$ and space $\tilde{O}(n)$ per machine. d_{avg} is the average degree of the vertices in the graph.*

There is a long line of work on computing approximate triangle counting in parallel computation [[Coh09](#), [TKMF09](#), [SV11](#), [YK11](#), [PT12](#), [KMPT12](#), [PC13](#), [SPK13](#), [AKM13](#), [PSKP14](#), [KPP+14](#), [JS17](#), [DLSY21](#)] and references therein. Despite this progress, and to the best of our knowledge, on one hand, each MPC algorithm for exact triangle counting either requires strictly super-polynomial in m total space, or the number of rounds is super-constant (as seen in [Table 1](#)). On the other hand, the best-known, classic algorithm for approximate triangle counting by Pagh and Tsourakakis [[PT12](#)] requires $T \geq d_{avg}$ even when the space per machine is $\Theta(n)$. We design an algorithm that has essentially optimal total space and round complexity, while at least quadratically improving the requirement on T .

Furthermore, since the amount of messages sent and received by each machine is bounded by $O(n)$, by [BDH18], our algorithm directly implies an $O(1)$ -rounds algorithm in the CONGESTED-CLIQUE model¹ under the same restriction $T = \Omega(\sqrt{m/n})$. The best known (to our knowledge) triangle approximation algorithm for general graphs in this model, is an $O(n^{1/3}/T^{2/3})$ -rounds algorithm by [DLP12]. The best-known previous bound only results in constant round complexity when $T = \Omega(\sqrt{n})$.

Corollary 1.3. *Given a graph $G = (V, E)$ with T triangles, if $T = \Omega(\sqrt{m/n})$, then there exists a $O(1)$ -rounds algorithm in the CONGESTED-CLIQUE model that gives a $(1 + \varepsilon)$ -approximation of T with high probability.*

The second question we consider is the question of exact counting, for which we present an algorithm whose total space depends on the arboricity of the input graph. The arboricity of a graph (roughly) equals the average degree of its densest subgraph. The class of graphs with bounded arboricity includes many important graph families such as planar graphs, bounded degree graphs and randomly generated preferential attachment graphs. In addition, many real-world graphs exhibit bounded arboricity [GG06, ELS13, SERF18], making this property important also in practical settings. For many problems, a bound on the arboricity of the graph allows for much more efficient algorithms and/or better approximation ratios [AG09, ELS13].

Specifically for the task of subgraph counting, in a seminal paper, Chiba and Nishizeki [CN85] prove that triangle enumeration can be performed in $O(m\alpha)$ time, and assuming 3SUM-hardness this result is optimal up to dependencies in $O(\text{poly log } n)$ [Pat10, KPP16]. Many applied algorithms also rely on the property of having bounded arboricity in order to achieve better space and time bounds, e.g., [SW05, CC11, Lat08]. Our main theorem with respect to this question is the following.

Theorem 1.4. *Let $G = (V, E)$ be a graph with n vertices, m edges and arboricity α . COUNT-TRIANGLES(G) takes $O_\delta(\log \log n)$ rounds, $O(n^\delta)$ space per machine for any $\delta > 0$, and $O(m\alpha)$ total space.*

It is interesting to note that our total space complexity matches the time complexity (both upper and conditional lower bounds) of combinatorial² triangle counting algorithms in the sequential model [CN85, Pat10, KPP16]. The best-known previous algorithm in this setting is the folklore algorithm of placing each vertex and its out-neighbors in the same machine and counting the incident triangles. Such an approach requires $O(\log n)$ rounds and $\Omega(\alpha^2)$ space per machine (summarized in Table 1). We prove the above theorem in Section 4.

1.2.2 Clique Counting

All of our above triangle counting results can be extended to k -clique counting. In Section 5, we prove that our exact triangle counting result can be extended to exactly counting k -cliques for any constant k :

Theorem 1.5. *Let $G = (V, E)$ be a graph with n vertices, m edges and arboricity α . COUNT-CLIQUE(G) takes $O_\delta(\log \log n)$ rounds, $O(n^\delta)$ space per machine for any $\delta > 0$, and $O(m\alpha^{k-2})$ total space.*

¹A distributed model where nodes communicate with each other over a complete network using $O(\log n)$ bit messages [LPSP05].

²Combinatorial algorithms, usually, refer to algorithms that do not rely on fast matrix multiplication.

We can improve on the total space usage if we are given machines where the memory for each individual machine satisfies $\alpha < n^{\delta'/2}$ where $\delta' < \delta$. In this case, we obtain an algorithm that counts the number of k -cliques in G using $O(n\alpha^2)$ total space and $O_\delta(\log \log n)$ communication rounds.

Furthermore, our approximate triangle counting results can be extended to counting *any* subgraph of size K where K is constant. Specifically, we obtain the following result:

Theorem 1.6. *Let $G = (V, E)$ be a graph with n vertices, m edges, and let B be the number of occurrences of a subgraph H with K vertices in G . If $B \geq d_{avg}^{K/2-1}$, then there exists an MPC algorithm that gives a $(1+\varepsilon)$ -approximation of B in $O(1)$ rounds, total space $\tilde{O}(m)$, and $\tilde{O}(n)$ space per machine, with high probability. Here, d_{avg} is the average degree of the vertices in the graph.*

1.3 Other Small Subgraphs

Finally, in [Section 7](#), we consider the problem of exactly counting subgraphs of size at most 5, and show that the recent result of Bera, Pashanasangi and Seshadhri [[BPS20](#)] for this question in the sequential model, can be implemented in the MPC model. Ours is the first result for counting any arbitrary subgraph of size at most 5 in $\text{poly}(\log n)$ rounds in the MPC model. Here too, our total space complexity matches the time complexity of the sequential model algorithm. It is an interesting open question whether our results can be extended to more general subgraphs following the results of [[Bre19](#), [BPS21](#)]. [Section 9](#) summarizes the difficulties of implementing these algorithms in the MPC model and we present this question as interesting future work.

Theorem 1.7. *Let $G = (V, E)$ be a graph with n vertices, m edges, and arboricity α . The algorithm of BPS for counting the number of occurrences of a subgraph H over $k \leq 5$ vertices in G can be implemented in the MPC model in $O_\delta(\sqrt{\log n} \log \log m)$ rounds, with high probability. The space requirement per machine is $O(n^{2\delta})$ and the total space is $O(m\alpha^3)$.*

1.4 Related Work

There has been a long line of work on small subgraph counting in massive networks in the MapReduce model whose results translate to the MPC model. We first describe the works for *exact* triangle and k -clique counting. [[SV11](#)] first designed an algorithm for triangle counting, but their approach requires a super-linear total space of $O(m^{3/2})$. Another work, [[AFU13](#)], shows how to count small subgraphs by using b^3 machines, each requiring $O(m/b^2)$ space per machine. Hence, it uses a total space of $O(mb)$. Therefore, this approach either requires super-linear total space or almost $O(m)$ space per machine. [[SV11](#)] were the first to achieve constant number of rounds in MPC, where they design two algorithms. The first of those algorithms, that runs in 2 rounds, requires $O(\sqrt{m})$ space per machine and total space $O(m^{3/2})$. Their second algorithm requires only one round for exact triangle counting, total space $O(\rho m)$ and space per machine $O(m/\rho^2)$. Therefore, for this algorithm to work with polynomially less than space m per machine, it has to allow for a total space that is polynomially larger than m . [[CC11](#)] focus on algorithms that require a total space of $O(m)$. In the worst case, their algorithm performs $O(|E|/S)$ MPC rounds to output the exact count where S is the maximum space per machine. [[FFF15](#)] extended and provided new algorithms for clique counting but they also require $\Omega(m^{3/2})$ total space.

[[TKMF09](#), [AKM13](#)] designed randomized algorithms for *approximate* triangle counting also in the MapReduce model (whose results, again, can be translated rather straightforwardly to the MPC model). Their approach first sparsifies the input graph by sampling a subset of edges, and executes some of the known algorithms for triangle counting on the sampled subgraph. Denoting

their sampling probability by p , their approach outputs a $(1 + \varepsilon)$ -approximate triangle count with probability at most $1 - 1/(\varepsilon^2 p^3 T)$.³ To contrast this result with our approach, consider a graph G where $m = \Theta(n^2)$. Let G' be the edge-sparsified graph as explained above. To be able to execute the first algorithm of [SV11] on G' such that the total space requirement is $O(m)$, one can verify that it is needed to set $p = \Theta(n^{-2/3})$. This in turn implies that the result in [TKMF09, AKM13] outputs the correct approximation with constant probability only if $T = \Omega(n^2)$. An improved lower-bound can be obtained by using the second algorithm of [SV11]. By balancing out ρ and p and for $S = O(n)$, one can show that the sparsification results in a constant probability of success for $T = \Omega(n)$. On the other hand, for $S = O(n)$, our approach obtains the same guarantee even when $T = \Theta(\sqrt{d_{avg}(G)}) = \Theta(\sqrt{n})$.

The best-known algorithm of [PT12] is a randomized algorithm for approximate triangle counting based on graph partitioning. The graph is partitioned into $1/p$ pieces, where p is at least the ratio of the maximum number of triangles sharing an edge and T . When all the triangles share one edge, then $p \geq 1$, and hence such an approach would require the space per machine to be $\Omega(m)$. Furthermore, this approach requires the number of triangles to be lower bounded by $T = \Omega(d_{avg})$. Another more recent work of [SPK13] uses wedge sampling and provides a $(1 + \varepsilon)$ -approximation of the triangle count in $O(1)$ rounds when T is a constant fraction of the sum of squares of degrees. The comparison of our bounds with these previous results are summarized in Table 1.

Other related work. Subgraph counting (primarily triangles) was also extensively studied in the streaming model, see [BYKS02, KMSS12, BOV13, JSP13, MVV16, BC17, AKK19] and references therein. This culminated in a result that requires space $\tilde{O}(m^{3/2}/(T\varepsilon^2))$ to estimate the number of triangles within a $(1 + \varepsilon)$ -factor. In the semi-streaming setting it is assumed that one has $\tilde{O}(n)$ space at their disposal. This result fits in this regime for $T \geq m^{3/2}/n = d_{avg} \cdot m^{1/2}$. As a reminder, our MPC result requires $T \geq \sqrt{d_{avg}}$ when $S = \tilde{O}(n)$.

In a celebrated result, [AYZ97] designed an algorithm for triangle counting in the sequential settings that runs in $O(m^{2\omega/(\omega+1)})$ time, where ω is the best-known exponent of matrix multiplication. Since then, several important works have extended this result to k -clique counting [EG04, Vas09]. In the work-depth (shared-memory parallel processors) model, several results are known for this problem. There has been significant work on practical parallel algorithms for the case of triangle counting (e.g. [AKM13, SV11, PC13, PSKP14, ST15] among others). There is even an annual competition for parallel triangle counting algorithms [Gra]. For counting $k = 4$ and $k = 5$ cliques, efficient practical solutions have also been developed [ANR⁺17, DAH17, ESD16, HD14, PSV17]. [DBS18] recently implemented the Chiba-Nishizeki algorithm [CN85] for k -cliques in the parallel setting; although, their work does not achieve polylogarithmic depth. Even more recently, [SDS20] enumerated k -cliques in the work-depth model in $O(m\alpha^{k-2})$ expected work and $O(\log^{k-2} n)$ depth with high probability, using $O(m)$ space. Among other distinctions from our setting, the work-depth model assumes a shared, common memory.

In the CONGESTED-CLIQUE model, [CHKK⁺19] present an $\tilde{O}(n^{1-2/\omega}) = \tilde{O}(n^{0.158})$ rounds algorithm for matrix multiplication, implying the same complexity for exact triangle counting. [DLP12] present an algorithm for approximate triangle counting in general graphs whose expected running time is $O(n^{1/3}/T^{2/3})$. They also present an $O(\alpha^2/n)$ -rounds algorithm for bounded arboricity graphs.

³The actual probability is even smaller and also depends on pairs of triangles that share an edge.

2 Preliminaries

Counting Duplicates. We make use of interval trees for certain parts of our paper to count the number of repeating elements in a sorted list, given bounded space per machine. We use the interval tree implementation given by [GSZ11] to obtain our count duplicates algorithm in the MPC model. We prove the following theorem in the MPC model regarding our count duplicates tree implementation. The proofs of the following claims are given in [Appendix A](#).

Theorem 2.1. *Given a sorted list of N elements implemented on processors where the space per processor is S and the total space among all processors is $O(N)$, for each unique element in the list, we can compute the number of times it repeats in $O(\log_S N)$ communication rounds.*

We also use the following two new MPC primitives in proving our bounds. These primitives may be of use in other algorithms beyond the scope of our paper.

Lemma 2.2. *Given two sets of tuples Q and C (both of which may contain duplicates), for each tuple $q \in Q$, we return whether $q \in C$ in $O(|Q \cup C|)$ total space and $O_\delta(1)$ rounds given machines with space $O(n^\delta)$ for any $\delta > 0$.*

Lemma 2.3. *Given a machine M that has space $O(n^{2\delta})$ for any $\delta > 0$ and contains data of $O(n^\delta)$ words, we can generate x copies of M , each holding the same data as M , using $O(M \cdot x)$ machines with $O(n^\delta)$ space each in $O(\log_{n^\delta} x)$ rounds.*

3 Overview of Our Techniques

3.1 Exact Triangle Counting

Let $G = (V, E)$ be a graph with n vertices, m edges and arboricity at most α . We tackle the task of exactly counting the number of triangles in G in $O_\delta(\log \log n)$ rounds using the following ideas. In each round i , we partition the vertices into low-degree vertices A_i and high-degree vertices, according to a degree threshold γ_i , which grows doubly exponentially in the number of rounds. We then count the number of triangles incident to the set of low degree vertices A_i . Each low-degree vertex $v \in A_i$ sends a list of its neighbors to all its neighbors. Then, any neighbor u of v that detects a common neighbor w to u and v , adds the triangle (u, v, w) to the list of discovered triangles.

Once all triangles incident to the vertices in A_i are processed, we remove this set from the graph and continue with the now smaller graph. This removal of the already processed vertices allows us to handle larger and larger degrees from step to step while using a total space of $O(m\alpha)$. This behavior also leads to the $O_\delta(\log \log n)$ round complexity, as after this many rounds all vertices are processed. The key insight in our proof that we maintain $O(m\alpha)$ total space *even when we increase the degree threshold doubly exponentially*. Such insight allows us to obtain our improved number of rounds while maintaining the same total space as the previous folklore algorithm. Finally, we achieve improved space per machine to $O(n^\delta)$ for any constant $\delta > 0$ via a number of new MPC primitives. Our algorithm and its analysis are provided in [Section 4](#). We provide extensions of our triangle counting algorithm to k -cliques in [Section 5](#).

3.2 Approximate Triangle Counting

Our work reduces approximate triangle counting to exact triangle counting in multiple (randomly chosen) induced subgraphs of the original graph. In our work, and in contrast to prior approaches (e.g., [PT12]), the induced subgraphs on different machines might *overlap* in both vertices and edges.

This enables us to obtain better concentration bounds compared to prior work, but also brings many challenges. Surprisingly, our algorithm is very simple (with a more complicated analysis), but is able to achieve a better lower bound on the number of triangles required to achieve a $(1 + \varepsilon)$ -approximation with high probability.

The high level idea is that each machine M_i samples a subset of vertices V_i by including each vertex in V_i with probability \hat{p} . Then, each machine computes the induced subgraph $G[V_i]$ and the number of triangles in that subgraph. The total number of triangles seen across all the machines is used as an estimator. We repeat in parallel this sampling process $O(\log n)$ times and return the median of the estimates. The main challenge this approach raises is: *How do we efficiently collect overlapping induced subgraphs?* (Indeed, approximate triangle counting, even when the number of triangles is $O(1)$, can be reduced to counting the number of edges in *sparse* induced subgraphs with the total size of subgraphs being $\tilde{O}(m)$.) We now describe how to handle this task in our case.

Computing induced overlapping subgraphs. It is unclear how to compute the induced subgraph on each machine in $O(1)$ rounds without exceeding the total allowed space of $\tilde{O}(m)$. This task becomes easier if the subgraphs are disjoint. For example, such an issue is avoided when the graph is *partitioned* across machines as in the algorithm of Pagh and Tsourakakis [PT12] since there is one copy of each vertex among all the machines. This is not the case for our algorithm.

The trivial strategy of sampling vertices into the machines and querying for all possible edges between any pair of two vertices takes total space at least $\sum_{i=1}^{\mathcal{M}} X_i^2$ where X_i is the number of vertices sampled to each machine i . In general, this approach requires much larger than $\tilde{O}(m)$ space. We tackle this challenge by using a *globally known* hash function $h : V \times [\mathcal{M}] \rightarrow \{0, 1\}$, to indicate whether vertex v is sampled in the i^{th} machine. By requiring that the hash function is known to all machines, we can efficiently compute which edges to send to each machine, i.e., which edges belong to the subgraph $G[V_i]$. However, in order for all machines to be able to compute the hash function, the hash function has to use limited space. Hence, we cannot hope for a fully independent function, rather we can only use an $(S/\log n)$ -wise independent hash function. Still, we manage to show that we are able to handle the dependencies introduced by the hash function, even if we allow as little as $O(\log n)$ -independence.

We present our approximate algorithm in Section 6 and give an extension of this algorithm to *any* subgraph of size K in Section 6.6.

3.3 Counting k -cliques and 5-subgraphs.

We use similar techniques for both problems of exactly counting the number of k -cliques and of exactly counting subgraphs up to size 5. See Section 5.1 for details on the former task, and Section 7 for details on the latter. Our final result is the first MPC algorithm for counting any *arbitrary* subgraph H of size at most 5 in $\text{poly}(\log n)$ MPC rounds.

Let H denote the subgraph of interest. We say that a subgraph that can be mapped to a subset of H of size i is a i -*subcopy* of H . Our main contribution in this section is a new MPC procedure that in each round, tries to extend i -subcopies of H to $(i + 1)$ -subcopies of H by increasing the total space by a factor of at most α . This is possible by ordering the vertices in H such that each vertex has at most $O(\alpha)$ outgoing neighbors so that in each iteration only α possible extensions should be considered per each previously discovered subcopy.

Challenges. The major challenge we face here is dealing with *finding* and *storing* copies of small (constant-sized) subgraphs in individual machines. This is a challenge due to the fact that an entire neighborhood of a vertex v may not fit on one machine (recall that we have no restrictions on how large the constant δ in $O(n^\delta)$ machine size can be). Thus, we cannot compute all such small subgraphs on one machine. However, if not done carefully, computing small subgraphs across many

machines could potentially result in many rounds of computation (since we potentially have to try all combinations of vertices in a neighborhood). We solve this issue by formulating a new MPC procedure (Lemma 2.3) in which we carefully duplicate neighborhoods of vertices across machines. The detailed analysis of our algorithm is given in Section 7.

4 Exact Triangle Counting in $O(m\alpha)$ Total Space

In this section we describe our algorithm for (exactly) counting the number of triangles in graphs $G = (V, E)$ of arboricity α and prove Theorem 1.4, restated here, in Section 4.2. We first provide an overview of our algorithm and its challenges.

Theorem 4.1. *Let $G = (V, E)$ be a graph over n vertices, m edges and arboricity α . COUNT-TRIANGLES(G) takes $O_\delta(\log \log n)$ rounds, $O(n^\delta)$ space per machine for some constant $0 < \delta < 1$, and $O(m\alpha)$ total space.*

Importantly, unlike previous methods, we *do not* need to assume knowledge of the arboricity of the graph α as input into our algorithm. The arboricity only shows up in our space bound as a property of the graph but we do not need to have knowledge of its value as we run the algorithm. The folklore algorithm shown in Table 1 requires an assumption of an upper bound on α since in order to achieve $O(\log n)$ rounds, we must count triangles incident to and remove all vertices with degree less than or equal to 2α in each round. The procedure gets stuck if we remove vertices with degree c where $c < \alpha$ in each round because there exists an induced subgraph with degree at least α in a graph with arboricity α . One can estimate the arboricity of the graph using $O(\log n)$ additional rounds or an $O(\log n)$ additional factor in space. Our algorithm does not require this additional step.

In this section, we assume that individual machines have space $\Theta(n^\delta)$ where δ is some constant $0 < \delta < 1$. Given this setting, there are several challenges associated with this problem.

Challenge 4.2. *The entire subgraph neighborhood of a vertex may not fit on a single machine. This means that all triangles incident to a particular vertex cannot be counted on one machine. Even if we are considering vertices with degree at most α , it is possible that $\alpha > n^\delta$. Thus, we need to have a way to count triangles efficiently when the neighborhood of a vertex is spread across multiple machines.*

The second challenge is to avoid over-counting.

Challenge 4.3. *When counting triangles across different machines, over-counting the triangles might occur, e.g., if two different machines count the same triangle. We need some way to deal with duplicate counting of the triangles to obtain the exact count of the triangles.*

We deal with the above challenges in our procedures below. We assume in our algorithm that each vertex can access its neighbors in $O(1)$ rounds of communication; such can be ensured via standard MPC techniques. Let $d_Q(v)$ be the degree of v in the subgraph induced by vertex set Q , i.e. in $G[Q]$. Our main algorithm consists of the following COUNT-TRIANGLES(G) procedure.

Algorithm 1. Count-Triangles($G = (V, E)$)

- 1: Let Q_i be the set of vertices not yet processed by iteration i . Initially set $Q_0 \leftarrow V$.
- 2: Let T be the current count of triangles. Set $T \leftarrow 0$.
- 3: **for** $i = 0$ to $i = \lceil \log_{3/2}(\log_2(n)) \rceil$ **do**

```

4:  $\gamma_i \leftarrow 2^{(3/2)^i}$ .
5: Let  $A_i$  be the list of vertices  $v \in Q_i$  where  $d_{Q_i}(v) \leq \gamma_i$ . Set  $Q_{i+1} \leftarrow Q_i \setminus A_i$ .
6: parfor  $v \in A_i$  do
7:   Retrieve the list of neighbors of  $v$  and denote it by  $L_v$ .
8:   Send each of  $v$ 's neighbors a copy of  $L_v$ .
9: end parfor
10: parfor  $w \in Q_i$  do
11:   Let  $\mathcal{L}_w = \bigcup_{v \in (N(w) \cap A_i)} L_v$  be the union of neighbor lists received by  $w$ .
12:   Set  $T \leftarrow T + \text{FIND-TRIANGLES}(w, \mathcal{L}_w)$ . ▷ Algorithm 2
13: end parfor
14: Return  $T$ .

```

Round compression is a technique formulated by [Ona18, CLM⁺20] that randomly partitions the vertices in a graph across machines where each machine then stores the induced subgraph induced by the partition. Then, a problem (e.g. maximum matching) is solved locally in each induced subgraph in each machine. The solutions in each machine allows one to remove certain vertices, reducing the degree of the remaining graph. In each round compression step, the maximum degree of the graph drops by a polynomial factor. This degree reduction then allows for more aggressive sampling in the next round compression step. This leads to $O(\log \log \Delta)$ round compression steps until the maximum degree is $\text{poly}(\log n)$; in this case, the remaining graph can be placed on a single machine.

Our algorithm, although similar, is simpler than the round compression technique. We do not require sampling since vertices are assigned to machines by degree, deterministically. The crux of our argument is showing that allowing for total space in terms of the arboricity α leads to a simpler and deterministic argument. Furthermore, for this specific problem, we also do not need to place the induced subgraph on one machine. In the next section, we show an implementation that allows us to operate in the sublinear space per machine regime. We hope our algorithm and analysis will lead to other deterministic algorithms for bounded arboricity graphs in sublinear space per machine and $O(\log \log n)$ rounds.

4.1 MPC Implementation Details

In order to implement $\text{COUNT-TRIANGLES}(G)$ in the MPC model, we define our $\text{FIND-TRIANGLES}(w, \mathcal{L})$ procedure and provide additional details on sending and storing neighbor lists across different machines. We define *high-degree* vertices to be the set of vertices whose degree is $> \gamma$ and *low-degree* vertices to be ones whose degree is $\leq \gamma$ (for some γ defined in our algorithm). We now define the function $\text{FIND-TRIANGLES}(w, \mathcal{L})$ used in the above procedure:

Algorithm 2. Find-Triangles(w, \mathcal{L}_w)

- 1: Sort all elements in $(\mathcal{L}_w \cup (N(w) \cap Q_i))$ lexicographically, using the procedure given in Lemma 4.3 of [GSZ11]. Let this sorted list of all elements be S .
- 2: Let T denote the corrected^a number of duplicates in S using Theorem 2.1.
- 3: Return T .

^aSome care must be taken here to avoid over-counting, since a distinct triangle can show up as several counted duplicates. See Algorithm 3 for details.

Allocating machines for sorting. Since each $v \in Q_i$ could have multiple neighbors whose

degrees are $\leq \gamma$, the total size of all neighbor lists v receives could exceed their allowed space $\Theta(n^\delta)$. Thus, we allocate $O\left(\frac{\gamma d_{Q_i}(v)}{n^\delta}\right)$ machines for each vertex $v \in Q_i$ to store all neighbor lists that v receives.

The complete analysis for [Theorem 4.1](#) is given in [Section 4.2](#).

We provide two additional extensions of our triangle counting algorithm to counting k -cliques:

Theorem 4.4. *Given a graph $G = (V, E)$ with arboricity α , we can count all k -cliques in $O(m\alpha^{k-2})$ total space, $O_\delta(\log \log n)$ rounds, on machines with $O(n^{2\delta})$ space for any $0 < \delta < 1$.*

We can prove a stronger result when we have some bound on the arboricity of our input graph. Namely, if $\alpha = O(n^{\delta'/2})$ for any $\delta' < \delta$, then we obtain the following result:

Theorem 4.5. *Given a graph $G = (V, E)$ with arboricity α where $\alpha = O(n^{\delta'/2})$ for any $\delta' < \delta$, we can count all k -cliques in $O(n\alpha^2)$ total space and $O_\delta(\log \log n)$ rounds, on machines with $O(n^\delta)$ space for any $0 < \delta < 1$.*

The proofs of these theorems are provided in [Section 5](#).

4.2 Detailed Analysis

In this section we give the full details and analysis of algorithm [Algorithm 1](#) given in [Section 4](#), for exactly counting the number of triangles in the graph.

We first provide a detailed version of [Algorithm 2](#) that also takes into account over counting due to the fact that each triangle might be counted by several endpoints, and then continue to prove the main theorem of this section, [Theorem 1.4](#).

4.2.1 Details about finding duplicate elements using [Theorem 2.1](#)

`FIND-TRIANGLES(w, \mathcal{L}_w)` finds triangles by counting the number of duplicates that occur between elements in lists. [Theorem 2.1](#) provides a MPC implementation for finding the count of all occurrences of every element in a sorted list. Provided a sorted list of neighbors of $v \in Q_i$ and neighbor lists in \mathcal{L}_v , this function counts the number of intersections between a neighbor list sent to v and the neighbors of v . Every intersection indicates the existence of a triangle. As given, `FIND-TRIANGLES(w, \mathcal{L}_w)` (see [v Algorithm 2](#)) returns a 6-approximation of the number of triangles in any graph. We provide a detailed and somewhat more complicated algorithm `FIND-TRIANGLES-EXACT(w, \mathcal{L}_w)` that accounts for over-counting of triangles and returns the exact number of triangles.

Since [Theorem 2.1](#) returns the total count of each element, we subtract the value returned by 1 to obtain the number of intersections. Finally, each triangle containing one low-degree vertex will be counted twice, each containing two low-degree vertices will be counted 4 times, and each containing three low-degree vertices will be counted 6 times. Thus, we need to divide the counts by 2, 4, and 6, respectively, to obtain the exact count of unique triangles.

Algorithm 3. Find-Triangles-Exact(w, \mathcal{L}_w)

- 1: Set the number of triangles $T_i \leftarrow 0$.
- 2: Sort all elements in $(\mathcal{L}_w \cup (N(w) \cap Q_i))$ lexicographically using the procedure given in Lemma 4.3 of [\[GSZ11\]](#). Let this sorted list of all elements be S .
- 3: Count the duplicates in S using [Theorem 2.1](#).
- 4: **parfor** all $v \in N(w)$ **do**

```

5:   Let  $R$  be the number of duplicates of  $v$  returned by Theorem 2.1.
6:   if  $d_{Q_i}(v) > \gamma_i$  and  $d_{Q_i}(w) > \gamma_i$  then
7:     Increment  $T_i \leftarrow T_i + \frac{R-1}{2}$ .
8:   else if ( $d_{Q_i}(v) > \gamma_i$  and  $d_{Q_i}(w) \leq \gamma_i$ ) or ( $d_{Q_i}(v) \leq \gamma_i$  and  $d_{Q_i}(w) > \gamma_i$ ) then
9:     Increment  $T_i \leftarrow T_i + \frac{R-1}{4}$ .
10:  else
11:    Increment  $T_i \leftarrow T_i + \frac{R-1}{6}$ .
12: end parfor
13: Return  $T_i$ .

```

Substituting FIND-TRIANGLES-EXACT in COUNT-TRIANGLES finds the exact count of triangles in graphs with arboricity α using $O(m\alpha)$ total space.

4.2.2 Proof of [Theorem 1.4](#)

First, all proofs below assume we start at a cutoff of $\gamma = 4\alpha$. Because we increase the cutoff bound doubly exponentially, we can reach such a bound in $O(\log \log \alpha)$ rounds. Thus, in the following proofs, we ignore all rounds before we get to a round where $\gamma \geq 4\alpha$. Before proving the theorem, we provide several useful lemmas stating that the number of vertices and edges remaining at the beginning of each iteration is bounded.

Lemma 4.6. *At the beginning of iteration i of COUNT-TRIANGLES, given $\gamma_i = 2^{(3/2)^i} \cdot (2\alpha)$ as stated in [Algorithm 1](#), the number of remaining vertices $N_i = |Q_i|$ is at most $\frac{n}{2^{2 \cdot ((3/2)^i - 1)}}$.*

Proof. Let N_i be the number of vertices in Q_i at the beginning of iteration i . Since the subgraph induced by Q_i must have arboricity bounded by α , we can bound the total degree of Q_i ,

$$\sum_{v \in Q_i} d_{Q_i}(v) < 2\alpha|Q_i| = 2N_i\alpha.$$

At the end of the iteration, we only keep the vertices in $Q_{i+1} = \{v \in Q_i \mid d_{Q_i}(v) > \gamma_i\}$. If we assume that $|Q_{i+1}| > \frac{N_i}{\gamma_i/(2\alpha)}$, then we obtain a contradiction since this implies that

$$\sum_{v \in Q_{i+1}} d_{Q_i}(v) > |Q_{i+1}| \cdot \gamma_i > 2N_i\alpha > \sum_{v \in Q_i} d_{Q_i}(v).$$

Then, the number of remaining vertices follows directly from the above by induction on i with base case $N_1 = n$,

$$N_i \leq \frac{N_{i-1}}{\gamma_i/(2\alpha)} = \frac{N_{i-1}}{2^{(3/2)^{i-1}}} \leq \frac{n}{\prod_{j=0}^{i-1} 2^{(3/2)^j}} = \frac{n}{2^{2 \cdot ((3/2)^i - 1)}}.$$

□

We can show a similar statement for the number of edges that remain at the start of the i^{th} iteration.

Lemma 4.7. *At the beginning of iteration i of COUNT-TRIANGLES, given γ_i , the number of remaining edges m_i is at most $m_i \leq \frac{m}{2^{2 \cdot ((3/2)^i - 1)}}$.*

Proof. The number of vertices remaining at the beginning of iteration i is given by $|Q_i|$. Thus, because the arboricity of our graph is α , we can upper bound m_i by

$$m_i \leq |Q_i|\alpha.$$

Then, we can also lower bound the number of edges at the beginning of iteration $i - 1$ since the vertices that remain at the beginning of round i are ones which have greater than γ_{i-1} degree,

$$m_{i-1} \geq \frac{1}{2} \sum_{v \in Q_{i-1}} d_{Q_{i-1}}(v) \geq \frac{1}{2}|Q_i|\gamma_{i-1}.$$

Thus, we conclude that $m_i \leq \frac{2\alpha m_{i-1}}{\gamma_{i-1}}$. By induction on i with base case $m_0 = m$, we obtain,

$$m_i \leq 2\alpha \left(\frac{m_{i-1}}{\gamma_{i-1}} \right) \leq \frac{m}{\prod_{j=0}^{i-2} 2(3/2)^j} = \frac{m}{2^{2 \cdot ((3/2)^{i-1} - 1)}}.$$

□

The above lemmas allows us to bound the total space used by the algorithm.

Lemma 4.8. *COUNT-TRIANGLES(G) uses $O(m\alpha)$ total space when run on a graph G with arboricity α .*

Proof. The total space the algorithm requires is the sum of the space necessary for storing the neighbor lists sent by all vertices with degree $\leq \gamma_i$ and the space necessary for all vertices to store their own neighbor lists. The total space necessary for each vertex to store its own neighbor list is $O(m)$.

Now we compute the total space used by the algorithm during iteration i . The number of vertices in Q_i at the beginning of this iteration is at most $N_i \leq \frac{n}{2^{2 \cdot ((3/2)^{i-1} - 1)}}$ by Lemma 4.6. Each vertex v with $d_{Q_i}(v) \leq \gamma_i$, makes $d_{Q_i}(v)$ copies of its neighbor list ($N(v) \cap Q_i$) and sends each neighbor in $N(v) \cap Q_i$ a copy of the list. Thus, the total space required by the messages sent by v is $d_{Q_i}(v)^2 \leq \gamma_i^2$. v sends at most one message of size $d_{Q_i}(v) \leq \gamma_i$ along each edge (v, w) for $w \in N(v) \cap Q_i$. Then, by Lemma 4.7 the total space required by all the low-degree vertices in round i is at most (as at most two messages are sent along each edge):

$$2m_i \cdot \gamma_i < \frac{m}{2^{2 \cdot ((3/2)^{i-1} - 1)}} \cdot \left[2^{(3/2)^i} (2\alpha) \right] = 16m\alpha.$$

□

We are now ready to prove Theorem 1.4.

Proof of Theorem 1.4. By Lemma 4.6, the number of vertices remaining in Q_i at the beginning of iteration i is $\frac{n}{2^{2 \cdot ((3/2)^{i-1} - 1)}}$. This means that the procedure runs for $O(\log \log n)$ iterations before there will be no vertices. For each of the $O(\log \log n)$ iterations, COUNT-TRIANGLES(G) uses $O_\delta(1)$ rounds of communication for the low-degree vertices to send their neighbor lists to their neighbors. The algorithm then calls FIND-TRIANGLES-EXACT(w, \mathcal{L}_w) on each vertex $w \in Q_i$ (in parallel) to find the number of triangles incident to w and vertices in $A_i \subseteq Q_i$. FIND-TRIANGLES-EXACT(w, \mathcal{L}_w) requires $O(\log_{n^\delta}(m\alpha)) = O(1/\delta)$ rounds by Lemma 4.3 of [GSZ11] and Theorem 2.1. Therefore, the total number of rounds required by COUNT-TRIANGLES(G) is $O\left(\frac{\log \log n}{\delta}\right) = O_\delta(\log \log n)$. □

5 Extensions to Exact k -Clique Counting in Graphs with Arboricity α

In this section, we briefly provide two algorithms for exact counting of k -cliques (where k is constant) in graphs with arboricity α . The first is an extension of our exact triangle counting result given in Section 4. The second is a query-based algorithm where the neighborhood of a low-degree vertex is constructed on a single machine via edge queries. In this case, the triangles incident to any given low-degree vertex can be counted on the same machine.

5.1 Exact k -Clique Counting

Exact k -Clique Counting in $O(m\alpha^{k-2})$ Total Space and $O_\delta(\log \log n)$ Rounds We extend our algorithm given in Section 4 to exactly count k -cliques (where k is constant) in $O(m\alpha^{k-2})$ total space and $O_\delta(\log \log n)$ rounds. Given a graph $G = (V, E)$ with arboricity α , the idea behind the algorithm is the following: let $G_i = (V_i \cup V, E_k \cup E)$ be a graph where each vertex $v \in V_i$ corresponds to an i -clique in G . Let $K(u)$ denote the $K_i \in G$ represented by $u \in V_k$. An edge (u, v) exists in E_i iff $u \in V_i, v \in V$ and $K(u) \cup \{v\}$ is an $(i+1)$ -clique in G . We construct the G_k graphs iteratively, starting with $G_1 = G$. Then, given G_{i-1} , we recursively construct G_i by using our exact triangle counting algorithm. Once we have G_{k-2} , we obtain our final count of the number of k -cliques by running our exact triangle counting algorithm one last time. The total space used is dominated by running the triangle counting algorithm on G_{k-2} , which uses $O(m\alpha^{k-2})$ total space. Since we run the triangle counting algorithm $O(k)$ times and k is a constant, the total number of rounds of communication necessary is $O_\delta(\log \log n)$ rounds. This detailed algorithm is given below.

Below, we describe our $O(n\alpha^{k-1})$ total space, $O(\log \log n)$ rounds exact k -clique counting algorithm that can be run on machines with space $O(n^\delta)$. Calling $\text{COUNT-}k\text{-CLIQUES}(G, k, k')$ for any given graph $G = (V, E)$ returns the number of k -cliques in G .

Algorithm 4. k -Clique-Counting($G = (V, E), k, k'$)

```

1: if  $k \leq 1$  then
2:   Return  $(|N|, G)$ 
3: else
4:    $(x, G_{k-1}) \leftarrow \text{COUNT-}k\text{-CLIQUES}(G, k-1, k')$ 
5:    $T \leftarrow \text{ENUMERATE-TRIANGLES}(G_{k-1})$ . Let  $T$  be the set of all enumerated triangles.
6:   Initialize sets  $V_k \leftarrow \emptyset$  and  $E_k \leftarrow \emptyset$ .
7:   parfor  $t \in T$  do
8:     Let  $K(t)$  represent the set of vertices in  $V$  composing the clique represented by  $t \in T$ .
9:     parfor  $v \in K(t)$  do
10:      Let  $v'(S)$  be a vertex  $v$  representing a set of vertices  $S$ . In other words,  $K(v) = K(v'(S)) = S$ .
11:       $V_k \leftarrow V_k \cup v'(K(t) \setminus v)$ .
12:       $E_k \leftarrow E_k \cup (v, v'(K(t) \setminus v))$ .
13:     end parfor
14:   end parfor
15:   if  $k = k' - 2$  then
16:     Return  $|T|$ .
17:   else

```

18: Return $(|V_k|, G_k(V \cup V_k, E \cup E_k))$.

Algorithm 5. Triangle-Enumeration $(G = (V, E))$

1: Let the set of enumerated triangles to be $T \leftarrow \emptyset$.
2: Let Q_i be the set of vertices that have not yet been processed by iteration i . Initially set $Q_0 \leftarrow V$.
3: **parfor** $i = 0$ to $i = \lceil \log_{3/2}(\log_2(n)) \rceil$ **do**
4: $\gamma_i \leftarrow 2^{(3/2)^i} \cdot 2\alpha$.
5: Let A_i be the list of vertices $v \in Q_i$ where $d_{Q_i}(v) \leq \gamma_i$. Set $Q_{i+1} \leftarrow Q_i \setminus A_i$.
6: Use [Lemma 5.1](#) to enumerate the set of triangles incident to A_i . Let this set be T_i .
7: $T \leftarrow T \cup T_i$.
8: **end parfor**
9: Return T .

5.2 MPC Implementation

To implement COUNT- k -CLIQUES in the MPC model, we must be able to create the graph G_2, \dots, G_{k-1} efficiently in our given space and rounds. The crux of this algorithm is the procedure for enumerating all triangles given a set A of vertices in G where $d(v) \leq \gamma$ for all $v \in A$. To do the triangle enumeration, we prove [Lemma 5.1](#) which can enumerate all such triangles incident to A in $O(m\gamma)$ total space, $O_\delta(1)$ rounds given machines with space $O(n^{2\delta})$.

Lemma 5.1. *Given a graph G , a constant integer $k \geq 2$, and a subset $A \subseteq G$ of vertices such that for every $v \in A$, $d(v) \leq \gamma$, we can generate all triangles in G that are incident to vertices in A in $O_\delta(1)$ rounds, $O(n^{2\delta})$ space per machine, and $O(m\gamma)$ total space.*

Proof. Let R be the set of machines holding the edges incident to A . Here too, similarly to the proof of [Lemma 7.4](#), it will be easier to think of each machine M as a set of n^δ parts, so that each edge, incident to a vertex in A , resides on a single part. We duplicate each such part, holding some neighbor of A , α times, using [Lemma 2.3](#). (We will actually duplicate machines, but, again, think of the duplicated machines as a collection of duplicated parts.) By [Lemma 2.3](#), this takes $O(\log_{n^\delta} \alpha) = O_\delta(1)$ rounds. Fix some vertex $v \in A$ and assume that $u \in N(v)$ resides on part $P_i(v)$. After the duplication step, there are α copies of each part. We denote these copies $P_{i,1}(v), \dots, P_{i,\alpha}(v)$. All parts $P_{i,j}(v)$ where $j \in [\alpha]$ and $v \in A$ then asks for v 's i -th neighbor in $O(1)$ rounds of communication. Now, each part $P_{i,j}(v)$ creates $O(1)$ edge queries to check whether its vertices form a triangle. All of the queries generated by all parts can be answered in parallel using [Lemma 2.2](#) in $O_\delta(1)$ rounds. Then each part that discovered a triangle incident to v adds it to a list K . Now we sort the list K and remove any duplicated triangles, so that the list only holds a single copy of every clique incident to some vertex in A . The total round complexity is $O_\delta(1)$ due to the duplications, sorting, and answering the queries. The space per machine is $O(n^{2\delta})$ and the total memory is $O(m\alpha)$ as each machine was duplicated α times. \square

Using [Lemma 5.1](#), we can now prove the space usage and round complexity of ENUMERATE-TRIANGLES.

Lemma 5.2. *Given a graph $G = (V, E)$ with arboricity α , ENUMERATE-TRIANGLES(G) uses $O(n\alpha^2)$ total space, $O_\delta(\log \log n)$ rounds on machines with $O(n^{2\delta})$ space.*

Proof. By Lemma 4.6, the number of vertices remaining in Q_i at the beginning of the i -th iteration of ENUMERATE-TRIANGLES is at most $\frac{n}{2^{2 \cdot ((3/2)^i - 1)}}$. By Lemma 5.1, the total space usage of enumerating all triangles incident to A_i is $O(m\gamma_i) = O\left(m \cdot \left(2^{(3/2)^i} \cdot 2\alpha\right)\right)$. The summation of the space used for all i is then:

$$\sum_{i=0}^{\lceil \log_{3/2}(\log_2(n)) \rceil} \left(\frac{n}{2^{2 \cdot ((3/2)^i - 1)}} \right) \cdot \left(2^{(3/2)^i} \cdot 2\alpha \right) = O(n\alpha^2).$$

The number of rounds required by this algorithm is $O(\log \log n) \cdot O_\delta(1) = O_\delta(\log \log n)$. \square

Given the total space usage and number of rounds required by ENUMERATE-TRIANGLES, we can now prove the total space usage and number of rounds required by COUNT- k -CLIQUES. But first, we show that for any graph $G = (V, E)$ with arboricity α , all graphs G_1, \dots, G_{k-1} created by COUNT- k -CLIQUES has arboricity $O(\alpha)$ for constant k .

Lemma 5.3. *Given a graph $G = (V, E)$ with arboricity α as input to ENUMERATE-TRIANGLES, all graphs G_1, \dots, G_{k-1} generated by the procedure have arboricity $O(\alpha)$ for constant k .*

Proof. We prove this lemma via induction. In the base case, $G_1 = G$ and so G_1 has arboricity α . Now we assume that G_i for $i \in [k-1]$ has arboricity $O(\alpha)$ (for constant i) and show that G_{i+1} has $O(\alpha)$ arboricity. Suppose that G_i has arboricity $c\alpha$ for some constant c . We prove via contradiction that the arboricity of G_{i+1} is upper bounded by $3(i+1)c\alpha$. Suppose for the sake of contradiction that the arboricity of G_{i+1} is greater than $3(i+1)c\alpha$. Then, there must exist a subgraph, $G_{i+1}[V']$ for some vertex set, V' , of G_{i+1} that contains greater than $3(i+1)c\alpha|V'|$ edges (by definition of arboricity). We now convert this subgraph $G_{i+1}[V']$ to a subgraph in G_i . Every vertex in V' maps to at most i pairs of vertices in G_i connected by an edge. Every edge in $G_{i+1}[V']$ maps to at least 1 edge. Thus, the subgraph in G_i that $G_{i+1}[V']$ maps to contains at most $2i|V'|$ vertices and at least $3(i+1)c\alpha|V'|$ edges. This implies, by the definition of arboricity, that the arboricity of G_i is $\geq \frac{3(i+1)c\alpha|V'|}{2i|V'|} > c\alpha$, a contradiction. Hence, the arboricity of G_{i+1} is at most $3(i+1)c\alpha$. And we have proven that the arboricity of G_{i+1} is $O(\alpha)$ for constant k . By induction, all graphs G_1, \dots, G_{k-1} have arboricity $O(\alpha)$. \square

Now we prove our final theorem of the space and round complexity of COUNT- k -CLIQUES.

Proof of Theorem 4.4. The number of i -cliques in a graph with arboricity α is at most $O(m\alpha^{i-2})$. Thus, by Lemma 5.2 and Lemma 5.3, COUNT- k -CLIQUES during the i -th call uses $O(m\alpha^i)$ total space, $O_\delta(\log \log n)$ rounds. Thus, COUNT- k -CLIQUES uses $O(m\alpha^{k-2})$ space, $O_\delta(\log \log n)$ rounds given machines with $O(n^{2\delta})$ space to count k -cliques given that the procedure terminates on the $(k-2)$ -th iteration. \square

5.3 Exact k -Clique Counting in $O(n\alpha^2)$ Total Space and $O_\delta(\log \log n)$ Rounds

We can improve on the total space usage if we are given machines where the memory for each individual machine satisfies $\alpha < n^{\delta'/2}$ where $\delta' < \delta$. In this case, we obtain an algorithm that counts the number of k -cliques in G using $O(n\alpha^2)$ total space and $O_\delta(\log \log n)$ communication rounds.

The entire neighborhood of any vertex with degree $\leq n^{\delta'/2}$ can fit on one machine. Suppose that $\alpha < n^{\delta'/2}$ where $\delta' < \delta$, then, there will always exist vertices that have degree $\leq n^{\delta'/2}$. Our algorithm proceeds as follows:

Algorithm 6. Count-Cliques($G = (V, E)$)

```

1: Let  $Q_i$  be the set of vertices that have not yet been processed by iteration  $i$ . Initially set  $Q_0 \leftarrow V$ .
2: Let  $C$  be the current count of cliques. Set  $C \leftarrow 0$ .
3: for  $i = 0$  to  $i = \lceil \log_{3/2}(\log_2(n)) \rceil$  do
4:    $\gamma_i \leftarrow 2^{(3/2)^i} \cdot 2\alpha$ .
5:   Let  $A_i$  be the list of vertices where  $d_{Q_i}(v) \leq \min(cn^{\delta/2}, \gamma_i)$  for some constant  $c$ .
6:   Set  $Q_{i+1} \leftarrow Q_i \setminus A_i$ .
7:   parfor  $v \in A_i$  do
8:     Retrieve all neighbors of  $v$ . Let this list of  $v$ 's neighbors be  $L_v$ .
9:     Query for all pairs  $u, v \in L_v$  to determine whether edge  $(u, v)$  exist. Retrieve all edges that exist.
10:    Count the number of triangles  $T_v$  incident to  $v$ , accounting for duplicates.
11:     $T \leftarrow T + T_v$ .
12:   end parfor

```

5.4 MPC Implementation Details

Accounting for Duplicates We account for duplicates by counting for each iteration i how many triangles on each machine contains 1, 2 or 3 vertices which have degree $\leq \min(cn^{\delta/2}, \gamma_i)$ (again we call these vertices low-degree). We multiply the count of triangles which have $t \geq 2$ low-degree vertices by $\frac{1}{t}$ to correct for over-counting due to multiple low-degree vertices performing the count on the same triangle. Each machine can retrieve the degrees of vertices in it in $O_\delta(1)$ rounds and such information can be stored on the machine given sufficiently small constant c in COUNT-CLIQUE.

Proof of Theorem 4.5. Since we are considering vertices with degree at most $\min(cn^{\delta/2}, \gamma_i)$, by Lemma 4.8, the total space used by our algorithm during any iteration i is

$$N_i \cdot \left(\min \left(cn^{\delta/2}, \gamma_i \right) \right)^2 < 16n\alpha^2.$$

By Lemma 2.2, we query for whether each of the $\min \left(cn^{\delta/2}, \gamma_i \right)^2$ potential edges on each machine is an edge in G in parallel using $O(n\alpha^2)$ total space and $O_\delta(1)$ rounds.

If $\gamma_i < cn^{\delta/2}$ for all iterations i , then by Theorem 4.1, the number of communication rounds required by COUNT-CLIQUE is $O_\delta(\log \log n)$. If, on the other hand, $cn^{\delta/2} < \gamma_i$, then the number of vertices remaining in Q_i decreases by a factor of $cn^{\delta/2}$ every round. Thus, the number of rounds required in this case is $O\left(\frac{2+\delta'}{\delta}\right)$. Since we assume δ' and δ are constants, the number of communication rounds needed by this algorithm is $O_\delta(\log \log n)$. \square

6 Approximate Triangle Counting in General Graphs

In this section we provide our algorithm for estimating the number of triangles in general graphs (see Algorithms 7 and 9) and hence prove Theorem 1.1.

Theorem 1.1. *Let $G = (V, E)$ be a graph with n vertices, m edges, and let T be the number of triangles in G . Assuming*

$$(i) T = \tilde{\Omega}\left(\sqrt{\frac{m}{S}}\right),$$

$$(ii) S = \tilde{\Omega}\left(\max\left\{\frac{\sqrt{m}}{\varepsilon}, \frac{n^2}{m}\right\}\right),$$

there exists an MPC algorithm, using \mathcal{M} machines, each with local space S , and total space $\mathcal{M}S = \tilde{O}_\varepsilon(m)$, that outputs a $(1 \pm \varepsilon)$ -approximation of T , with high probability, in $O(1)$ rounds.

The rationale behind the lower bound constraints in [Theorem 1.1](#) will become clear when we discuss the challenges and analysis (formally presented in the following sections).

6.1 Overview of the Algorithm and Challenges

Our approach is to use the collection of machines to repeat the following experiment multiple times in parallel. Each machine M_i samples a subset of vertices V_i , and then counts the number of triangles \hat{T}_i seen in each induced graph $G[V_i]$. We then use the sum \hat{T} of all \hat{T}_i 's as an unbiased estimator (after appropriate scaling) for the number of triangles T in the original graph.

Algorithm 7. Approximate-Triangle-Counting($G=(V,E)$)

```

1:  $R \leftarrow 0$ 
2: parfor  $i \leftarrow 1 \dots \mathcal{M}$  do
3:   Let  $V_i$  be a random subset of  $V$  ▷ See Section 6.2 for details about the sampling
4:   if size of  $G[V_i]$  exceeds machine space  $S$  then
5:     Ignore this sample and set  $\hat{T}_i \leftarrow 0$ 
6:   else
7:     Let  $\hat{T}_i$  be the number of triangles in  $G[V_i]$ 
8:      $R \leftarrow R + 1$ 
9:   end parfor
10: Let  $\hat{T} = \sum_{i=1}^{\mathcal{M}} \hat{T}_i$ 
11: return  $\frac{1}{\hat{p}^3 R} \hat{T}$ 

```

Moving forwards, for the most part, we will focus on a specific machine M_i containing V_i (a single experiment). We list the main challenges in the analysis of this algorithm, along with the sections that describe them.

1. [Section 6.2](#): The induced subgraph $G[V_i]$ fits into the memory S of M_i (thus allowing us to count the number of triangles in $G[V_i]$ in one round).
2. [Section 6.3](#): We can efficiently (in one round) collect all the edges in the induced subgraph $G[V_i]$. This involves presenting an MPC protocol such that the number of messages *sent and received* by any machine is at most the *space per machine* S .
3. [Section 6.4](#) With high constant probability, the number of messages sent and received by each machine M_i is at most S .
4. [Section 6.5](#): With high *constant* probability (of at least 0.9), the sum of triangles across all machines, \hat{T} , is close to its expected value. Then, repeating the algorithm polylogarithmic number of times with only a polylogarithmic increase in total space, and by using the median trick, allows us to get a high probability bound. The specifics are discussed in [Section 6.5.1](#).

In each of the following sections, we first present a high level overview of the challenges that we need to solve and then follow these high-level descriptions with detailed proofs.

6.2 Challenge (1): Ensuring That $G[V_i]$ Fits on a Single Machine

Ensuring that edges fit on a machine: Our algorithm constructs V_i by including each $v \in V$ with probability \hat{p} , which implies that the expected number of edges in $G[V_i]$ is $\hat{p}^2 m$. Since we have to ensure that each induced subgraph $G[V_i]$ fits on a single machine, we obtain the constraint $\hat{p}^2 m = O(S)$. Concretely, we achieve this by defining:

$$\hat{p} \stackrel{\text{def}}{=} \frac{1}{10} \cdot \sqrt{\frac{S}{mk}}, \quad (1)$$

where the parameter $k = O(\log n)$ will be exactly determined later (See [Section 6.3](#)).

Ensuring that vertices fit on a machine: In certain regimes of values of n and m , the *expected number of vertices* ending up in an induced subgraph – $\hat{p}n$, may exceed the space limit S . Avoiding this scenario introduces an additional constraint $\hat{p}n = O(S) \iff S = \Omega(kn^2/m)$.

Getting a high probability guarantee: As discussed above, the value of $\hat{p} = \tilde{\Theta}_\varepsilon(\sqrt{S/m})$ is chosen specifically so that the *expected number of edges* in the induced subgraphs $G[V_i]$ is $\hat{p}^2 m \leq \Theta(S)$, thus using all the available space (asymptotically). In order to guarantee that this bound holds *with high probability* (see [Section 6.4.1](#)), we require additional constraints on the space per machine $S = \tilde{\Omega}_\varepsilon(\sqrt{m})$. We remark that this lower bound $S = \tilde{\Omega}_\varepsilon(\sqrt{m})$ is essentially saying that $\mathcal{M} = \tilde{O}_\varepsilon(\sqrt{m})$, i.e. the *space per machine* is much larger than the *number of machines*. This is a realistic assumption as in practice we can have machines with 10^{11} words of local random access memory, however, it is unlikely that we also have as many machines in our cluster.

Lower Bound on space per machine: Combining the above two constraints, we get:

$$S > \max \left\{ 15 \frac{\sqrt{mk}}{\varepsilon}, \frac{100kn^2}{m} \right\} \implies S = \tilde{\Omega}_\varepsilon \left(\max \left\{ \sqrt{m}, \frac{n^2}{m} \right\} \right) \quad (2)$$

Note that [Eq. \(2\)](#) always allows *linear space per machine*, as long as $m = \Omega(n)$. The following sections, [Sections 6.4.1](#) and [6.4.2](#) present a detailed analysis, showing that the number of vertices and edges in each subgraph is at most S *with high probability*. In this high-level overview of the challenges, we defer a detailed analysis of these bounds to the later sections ([Sections 6.4.1](#) and [6.4.2](#)) since the formal proof of these bounds also require a discussion of [Section 6.3](#).

6.3 Challenge (2): Using k -wise Independence to Compute the Induced Subgraph $G[V_i]$ in MPC

For each sub-sampled set of vertices V_i , we need to compute $G[V_i]$, i.e. we need to send all the edges in the induced subgraph $G[V_i]$ to the machine M_i . Let Q_u denote the set of all machines containing u . Each edge (u, w) then needs to be sent to all machines that contain both u and w , $Q_u \cap Q_w$. Naively, one could try to send the sets Q_u and Q_w to the edge $e = (u, w)$, for all $e \in E$. However, this strategy could result in Q_v being replicated $d(v)$ times. Since the expected size of Q_v is $|Q_v| = \hat{p}\mathcal{M}$ the total expected memory usage of this strategy would be $\sum_{v \in V} |Q_v| \cdot d(v) = \tilde{\Theta}_\varepsilon(m \cdot \hat{p}\mathcal{M}) = \tilde{\omega}_\varepsilon(m)$, since $\hat{p} = \tilde{\Theta}(1/\sqrt{\mathcal{M}})$. This defies our goal of optimal total memory.

Instead, we address this challenge by using globally known hash functions to sample the vertices on each machine. That is, we let $h : V \times [\mathcal{M}] \rightarrow \{0, 1\}$ (formally presented in [Definition 6.1](#)) be a hash function known globally to all the machines. Then we can compute the induced subgraphs $G[V_i]$ as follows.

Algorithm 8. Compute-Induced-Subgraphs

```

1:  $Q_v \leftarrow \{i \in [\mathcal{M}] \mid h(v, i) = 1\}$ .
2:  $Q_w \leftarrow \{i \in [\mathcal{M}] \mid h(w, i) = 1\}$ .
3: parfor  $i \in Q_v \cap Q_w$  do
4:   Send  $e$  to machine  $M_i$ , containing  $V_i$ .
5: end parfor

```

Definition 6.1. The hash function $h(v, i)$ indicates whether vertex v is sampled in V_i or not. Specifically, $h : V \times [\mathcal{M}] \rightarrow \{0, 1\}$ such that $\mathbb{P}[h(v, i) = 1] = \hat{p}$ for all $v \in V$ and $i \in [\mathcal{M}]$. Recall that \mathcal{M} is the number of machines, and $\hat{p} = \frac{1}{10} \cdot \sqrt{\frac{S}{mk}}$ is the sampling probability set in Eq. (1).

Using limited independence. Ideally, we would want a perfect hash function, which would allow us to sample the V_i 's i.i.d. from the uniform distribution on V . However, since the hash function needs to be known globally, it must fit into each of the machines. This implies that we *cannot* use a fully independent perfect hash function. Rather, we *can* use one that has a high level of independence. Specifically, given that the space per machine is S , we can have a globally known hash function h that is k -wise independent⁴ for any $k < \Theta(S/\log n)$. In fact, we can get away with as little as $(6 \log n)$ -wise independence (i.e., $k = 6 \log n$). Recalling Eq. (1), this also fixes the sampling probability to be $\hat{p} = \sqrt{S/600m \log n}$.

6.3.1 Showing Concentration for the Triangle Count

In the subsequent proofs, we will use the following assumptions from within Theorem 1.1 (note that we added specific constants).

$$T \geq 10\sqrt{\frac{mk}{S}} \quad S \geq \max \left\{ 15\frac{\sqrt{mk}}{\varepsilon}, \frac{100kn^2}{m} \right\} \quad \mathcal{M} = \frac{2000mk}{\varepsilon^2 S} \quad (3)$$

Note that we set the number of machines to a specific value, instead of lower bounding it. This is acceptable, because we can just ignore some of the machines.

Algorithm 7 outputs an estimate on the number of triangles in G (Line 11). It is not hard to show that in expectation this output equals T *even with limited independence* as discussed above. The main challenge is to show that this output also concentrates well around its expectation. Specifically, we show the following claim.

Lemma 6.2. Ignore Line 4 of Algorithm 7. Let \hat{T} be as defined on Line 10 and $\mathcal{M} = \frac{20}{\varepsilon^2 \hat{p}^2}$ be as defined in Eq. (3), and assume that $T \geq 1/\hat{p}$. Then, the following hold:

- (A) $\mathbb{E}[\hat{T}] = \hat{p}^3 \cdot R \cdot T$, and
- (B) $\mathbb{P}\left[|\hat{T} - \mathbb{E}[\hat{T}]| > \varepsilon \mathbb{E}[\hat{T}]\right] < \frac{1}{10}$.

We will prove Property (B) of the claim by applying Chebyshev's inequality, for which we need to compute $\text{Var}[\hat{T}]$. Let $\Delta(G)$ be the set of all triangles in G . For a triangle $t \in \Delta(G)$, let $\hat{T}_{i,t} = 1$

⁴A k -wise independent hash function is one where the hashes of *any* k distinct keys are guaranteed to be independent random variables (see [WC81]).

if $t \in V[G_i]$, and $\hat{T}_{i,t} = 0$ otherwise. Hence, $\hat{T}_i = \sum_{t \in \Delta(G)} \hat{T}_{i,t}$. We begin by deriving $\mathbb{E}[\hat{T}]$ and then proceed to showing that $\text{Var}[\hat{T}] = \sum_{i=1}^R \text{Var}[\hat{T}_i]$. After that we upper-bound $\text{Var}[\hat{T}_i]$ and conclude the proof by applying Chebyshev's inequality.

Deriving $\mathbb{E}[\hat{T}]$. Let t be a triangle in G . Let \hat{T}_t be a random variable denoting the total number of times t appears in $G[V_i]$, for all $i = 1 \dots R$. Given that $\mathbb{P}[u \in V_i] = \hat{p}$, we have that $\mathbb{P}[t \in G[V_i]] = \hat{p}^3$. Therefore, $\mathbb{E}[\hat{T}_t] = R \cdot \hat{p}^3$.

Since $\hat{T} = \sum_{t \in \Delta(G)} \hat{T}_t$, we have

$$\mathbb{E}[\hat{T}] = \sum_{t \in \Delta(G)} \mathbb{E}[\hat{T}_t] = \hat{p}^3 \cdot R \cdot T. \quad (4)$$

This proves Property (A) of this claim.

Decoupling $\text{Var}[\hat{T}]$. To compute variance, one considers the second moment of a given random variable. So, to compute $\text{Var}[\hat{T}]$, we will consider products $\hat{T}_{i,t_1} \cdot \hat{T}_{j,t_2}$. Each of those products depend on at most 6 vertices. Now, given that we used a 6-wise independent function (see [Section 6.3](#)) to sample vertices in each V_i , one could expect that $\text{Var}[\hat{T}_i]$ and $\text{Var}[\hat{T}_j]$ for $i \neq j$ behave like they are independent, i.e., one could expect that it holds $\text{Var}[\hat{T}] = \sum_{i=1}^R \text{Var}[\hat{T}_i]$. As we show next, it is indeed the case. We have

$$\begin{aligned} \text{Var}[\hat{T}] &= \mathbb{E}[\hat{T}^2] - \mathbb{E}[\hat{T}]^2 \\ &= \mathbb{E}\left[\left(\sum_{i=1}^R \sum_{t \in \Delta(G)} \hat{T}_{i,t}\right)^2\right] - \left(\sum_{i=1}^R \sum_{t \in \Delta(G)} \mathbb{E}[\hat{T}_{i,t}]\right)^2 \end{aligned} \quad (5)$$

Consider now \hat{T}_{i,t_1} and \hat{T}_{j,t_2} for $i \neq j$ and some $t_1, t_2 \in \Delta(G)$ not necessarily distinct. In the first summand of (5), we will have $\mathbb{E}[2\hat{T}_{i,t_1} \cdot \hat{T}_{j,t_2}]$. The vertices constituting t_1 and t_2 are 6 *distinct* copies of some (not necessarily all distinct) vertices of V . Since they are chosen by applying a 6-wise independent function, we have $\mathbb{E}[2\hat{T}_{i,t_1} \cdot \hat{T}_{j,t_2}] = 2\mathbb{E}[\hat{T}_{i,t_1}] \cdot \mathbb{E}[\hat{T}_{j,t_2}]$.

On the other hand, the second summand of (5) also contains $2\mathbb{E}[\hat{T}_{i,t_1}] \cdot \mathbb{E}[\hat{T}_{j,t_2}]$, which follows by direct expansion of the sum. Therefore, all the terms $\mathbb{E}[2\hat{T}_{i,t_1} \cdot \hat{T}_{j,t_2}]$ in $\text{Var}[\hat{T}]$ for $i \neq j$ cancel each other. So, we can also write $\text{Var}[\hat{T}]$ as

$$\begin{aligned} \text{Var}[\hat{T}] &= \sum_{i=1}^R \mathbb{E}\left[\left(\sum_{t \in \Delta(G)} \hat{T}_{i,t}\right)^2\right] - \sum_{i=1}^R \left(\sum_{t \in \Delta(G)} \mathbb{E}[\hat{T}_{i,t}]\right)^2 \\ &= \sum_{i=1}^R \text{Var}[\hat{T}_i]. \end{aligned} \quad (6)$$

Therefore, to upper-bound $\text{Var}[\hat{T}]$ it suffices to upper-bound $\text{Var}[\hat{T}_i]$.

Upper-bounding $\text{Var} [\hat{T}_i]$. We have

$$\begin{aligned}
\text{Var} [\hat{T}_i] &= \mathbb{E} \left[\left(\sum_{t \in \Delta(G)} \hat{T}_{i,t} \right)^2 \right] - \left(\sum_{t \in \Delta(G)} \mathbb{E} [\hat{T}_{i,t}] \right)^2 \\
&\leq \mathbb{E} \left[\left(\sum_{t \in \Delta(G)} \hat{T}_{i,t} \right)^2 \right] \\
&= \mathbb{E} \left[\sum_{t \in \Delta(G)} \hat{T}_{i,t}^2 \right] + \mathbb{E} \left[\sum_{t_1, t_2 \in \Delta(G); t_1 \neq t_2} \hat{T}_{i,t_1} \cdot \hat{T}_{i,t_2} \right]. \tag{7}
\end{aligned}$$

Since each $\hat{T}_{i,t}$ is a 0/1 random variables, $\hat{T}_{i,t}^2 = \hat{T}_{i,t}$. Let $t_1 \neq t_2$ be two triangles in $\Delta(G)$. Let k be the number of distinct vertices they are consisted of, which implies $4 \leq k \leq 6$. Then, observe that $\mathbb{E} [\hat{T}_{i,t_1} \cdot \hat{T}_{i,t_2}] = \hat{p}^k \leq \hat{p}^4$. We now have all ingredients to upper-bound $\text{Var} [\hat{T}_i]$. From (7) and our discussion it follows

$$\text{Var} [\hat{T}_i] \leq T\hat{p}^3 + T^2\hat{p}^4 \leq 2T^2\hat{p}^4, \tag{8}$$

where we used our assumption that $T \geq 1/\hat{p}$.

Finalizing the proof. From (6) and (8) we have

$$\text{Var} [\hat{T}] \leq 2RT^2\hat{p}^4.$$

So, from Chebyshev's inequality and (4) we derive

$$\begin{aligned}
\mathbb{P} \left[|\hat{T} - \mathbb{E} [\hat{T}]| > \varepsilon \mathbb{E} [\hat{T}] \right] &< \frac{\text{Var} [\hat{T}]}{\varepsilon^2 \mathbb{E} [\hat{T}]^2} \\
&\leq \frac{2RT^2\hat{p}^4}{\varepsilon^2 \hat{p}^6 R^2 T^2} \\
&= \frac{2}{\varepsilon^2 \hat{p}^2 R}.
\end{aligned}$$

Hence, for $R \geq \frac{20}{\varepsilon^2 \hat{p}^2}$ we get the desired bound.

6.4 Challenge (3): Showing that, with high constant probability, the size of the sent/received messages is bounded.

We need to show that the number of edges sent and received by any machine M_i is at most S with high constant probability. To this end, we partition the vertex set V into V_{light} and V_{heavy} by picking a threshold degree τ for the vertices. Following this, we define *light edges* as ones that have both end-points in V_{light} , and conversely, any edge with at least one end-point in V_{heavy} is designated as *heavy*. In order for the protocol to succeed, the following must hold:

- (A) The number of *light edges* concentrates (see Section 6.4.1).
- (B) The number of *heavy edges* concentrates (see Section 6.4.2).

(C) The number of sent messages is at most S (see [Section 6.4.3](#)).

The first two items ensure that each machine M_i *receives* at most S messages, and the last item ensures that each machine *sends* at most S messages. Given the above, we proceed to address the last challenge.

6.4.1 Bounding the Number of Light Edges *Received* by a Machine

We will now bound the probability that any of the induced subgraphs *does not fit* on a machine. To that end, we set a degree threshold $\tau = \frac{k}{\hat{p}}$, and define the set of *light* vertices V_{light} to be the ones with degree less than τ . All other vertices are *heavy*, and we let them comprise the set V_{heavy} .

Fix a machine M_i . We prove that, with probability at least $9/10$, the number of edges in $G[V_i]$ is upper bounded by S .

We start with analyzing the contribution of the light vertices to the induced subgraphs. We first consider the simpler case of bounding the number of edges in $G[V_i]$ that have both end-points in V_{light} . We refer to such edges as *light edges* and denote them by E_{light} . For every edge $e \in E_{light}$, we define a random variable $Z_e^{(i)}$ as follows.

$$Z_e^{(i)} = \begin{cases} 1 & \text{if } e \in G[V_i], \\ 0 & \text{otherwise.} \end{cases}$$

We let $Z^{(i)}$ be the sum over all random variables Z_e^i , $Z^i = \sum_{e \in E_{light}} Z_e^i$, and we let m_ℓ denote the total number of edges with *light* endpoints in the original graph G , i.e., $m_\ell = |E_{light}|$.

We prove the following lemma.

Lemma 6.3. *With probability at least $9/10$, for every $i \in [\mathcal{M}]$, $G[V_i]$ contains at most $\frac{1}{4}S$ light edges.*

Proof. Fix a machine M_i , and let $Z = Z^i$ be as defined in the previous paragraph.

$$\mathbb{E}[Z] = \mathbb{E} \left[\sum_{e \in E_{light}} Z_e \right] = m_\ell \hat{p}^2 \leq m \cdot \frac{S}{100mk} = \frac{S}{100k} \leq \frac{S}{100}.$$

As Z_e are $\{0, 1\}$ random variables, we also have $\mathbb{E}[Z] = \mathbb{E} \left[\sum_{e \in E_{light}} Z_e^2 \right]$. Now we upper-bound the

variance.

$$\begin{aligned}
\text{Var}[Z] &= \mathbb{E} \left[\left(\sum_{e \in E_{light}} Z_e \right)^2 \right] - \mathbb{E} \left[\sum_{e \in E_{light}} Z_e \right]^2 \\
&\leq \sum_{e \in E_{light}} \mathbb{E}[Z_e^2] + \sum_{\substack{e_1, e_2 \in E_{light} \\ e_1 \neq e_2}} 2 \cdot \mathbb{E}[Z_{e_1} Z_{e_2}] \\
&\quad - \sum_{\substack{e_1, e_2 \in E_{light} \\ e_1 \neq e_2}} 2 \cdot \mathbb{E}[Z_{e_1}] \mathbb{E}[Z_{e_2}] \\
&= m_\ell \cdot \hat{p}^2 + \sum_{\substack{e_1, e_2 \in E_{light} \\ e_1 \neq e_2}} 2 \cdot \mathbb{E}[Z_{e_1} Z_{e_2}] \\
&\quad - \sum_{\substack{e_1, e_2 \in E_{light} \\ e_1 \neq e_2}} 2 \cdot \mathbb{E}[Z_{e_1}] \mathbb{E}[Z_{e_2}] \\
&\leq m_\ell \cdot \hat{p}^2 + \sum_{e_1 \text{ and } e_2 \text{ intersect}} 2 \cdot \mathbb{E}[Z_{e_1} Z_{e_2}] \\
&\leq m_\ell \cdot \hat{p}^2 + \left(\sum_{v \in V_{light}} d(v)^2 \right) \cdot \hat{p}^3 \\
&\leq m_\ell \cdot \hat{p}^2 + \left(\sum_{v \in V_{light}} d(v) \right) \cdot \frac{k}{\hat{p}} \cdot \hat{p}^3 \\
&\leq 3m_\ell \cdot \hat{p}^2 \cdot k \leq 3m \cdot \frac{S}{100mk} \cdot k < \frac{S}{30}
\end{aligned}$$

□

We can now use Chebyshev's inequality to conclude that

$$\begin{aligned}
\mathbb{P} \left[|Z^{(i)} - \mathbb{E}[Z^{(i)}]| > S/\sqrt{3} \right] &\leq \frac{\text{Var}[Z^{(i)}]}{S^2/3} \\
\implies \mathbb{P} \left[Z^{(i)} > 3S/4 \right] &\leq \frac{3}{30S} = \frac{1}{10S}
\end{aligned}$$

Finally, we can use union bound over all \mathcal{M} machines to upper bound the probability that, *any* of the $Z^{(i)}$ values exceeds $3S/4$ (using the the constraints described in [Eq. \(3\)](#) to simplify).

$$\frac{\mathcal{M}}{10S} = \frac{2000mk}{\varepsilon^2 S} \cdot \frac{1}{10S} \leq \frac{200mk}{\varepsilon^2} \cdot \frac{1}{(15\sqrt{mk}/\varepsilon)^2} = \frac{200mk}{\varepsilon^2 S^2},$$

Therefore, with probability at least $9/10$, none of the induced subgraphs $G[V_i]$ will contain more than $3S/4$ light edges.

6.4.2 Bounding the Number of Heavy Edges *Received* by a Machine

Next, we turn our attention to the edges that have at least one endpoint in V_{heavy} (we call such edges *heavy*). We will show that for each $v \in V_{heavy} \cap V_i$, the number of edges contributed by v

concentrates around its expectation.⁵ In this section, we will use $2m_h$ to denote the total degree of all the heavy vertices i.e. $2m_h = \sum_{v \in V_{heavy}} d(v)$.

Let $Z_w^{(v)}$ be the $\{0, 1\}$ indicator random variable for $w \in V_i$ conditioned on the event that $v \in V_i \cap V_{heavy}$. We use this conditioning on v being present, because, in its absence, the number of edges contributed by v , can be zero with probability $(1 - \hat{p})$, i.e. this naive estimator would not concentrate around its expectation.

Let $Z^{(v)}$ be the sum of all $Z_w^{(v)}$ for $w \in N(v)$. For a particular v , the $Z_w^{(v)}$ variables are k -wise independent, which allows us to use the following lemma to bound $Z^{(v)}$. In what follows, we will omit the super-script (v) for the sake of convenience.

Lemma 6.4. *If Z_1, Z_2, \dots, Z_n are k -wise independent $\{0, 1\}$ random variables with $\mathbb{E}[Z_i] = p$ and $k \leq np$, then for $Z = \sum_i Z_i$ we have*

$$\mathbb{P}[Z > 3np] \leq 2^{-k}.$$

Proof. To prove the claim, we will re-write $\mathbb{P}[\sum Z_i > 3np]$, as the probability that the number of size k subsets of $\{Z_1, Z_2, \dots, Z_n\}$ that are all equal to 1 is larger than $\binom{3np}{k}$.

$$\begin{aligned} & \mathbb{P}[Z > 3np] \\ &= \mathbb{P}\left[|\{T : T \subseteq [n], |T| = k, \text{ and } Z_i = 1 \forall i \in T\}| > \binom{3np}{k}\right] \\ &\leq \frac{\mathbb{E}[|\{T : T \subseteq [n], |T| = k, \text{ and } Z_i = 1 \forall i \in T\}|]}{\binom{3np}{k}} \\ &= \frac{\binom{n}{k} \cdot p^k}{\binom{3np}{k}} \leq \left(\frac{n}{3np - k} \cdot p\right)^k \leq \left(\frac{np}{2np}\right)^k = 2^{-k} \end{aligned}$$

where to obtain $3np - k \geq 2np$ we used our assumption that $k \leq np$. □

Since v is heavy, there are at least τ variables in the sum $Z^{(v)} = \sum_{w \in N(v)} Z_w^{(v)}$. Additionally, we know that $\mathbb{E}[Z_w^{(v)}] = \hat{p}$ and $k \leq \tau \hat{p}$. Thus, we obtain the following corollary from [Lemma 6.4](#):

Corollary 6.5. *For any vertex $v \in V_{heavy} \cap V_i$, we get $\mathbb{P}[Z^{(v)} > 3d(v) \cdot \hat{p}] < 2^{-k}$, or explicitly*

$$\mathbb{P}[N(v) \cap V_i > 3d(v)\hat{p} \mid v \in V_i \text{ and } d(v) > \tau] < 2^{-k} = \frac{1}{n^6}$$

Corollary 6.6. *With high probability $1 - \frac{1}{n^5}$, we ensure that for all $v \in V_{heavy}$, $Z^{(v)} \leq 3 \cdot \mathbb{E}[Z^{(v)}]$*

The important point is that the sum of $Z^{(v)}$ (over all $v \in V_i$) is an upper bound on m_h – the number of heavy edges in $G[V_i]$. In order to bound this sum, we define random variables W_v for each $v \in V_{heavy}$ as follows:

$$W_v = \begin{cases} \frac{d(v)}{n} & \text{if } v \in V_i \\ 0 & \text{otherwise} \end{cases}$$

We also define W to be the sum of all W_v , thus implying $\mu = \mathbb{E}[W] = \sum_{v \in V_{heavy}} \hat{p} \cdot \frac{d(v)}{n} \leq \frac{2\hat{p}m_h}{n}$.

⁵Intuitively, this is because v has high degree, and therefore the number of its sampled neighbors ($|N(v) \cap V_i|$) will concentrate.

Theorem 6.7. (Theorem 5 from [SSS95]) If W is the sum of k -wise independent random variables, each of which takes values in the interval $[0, 1]$, and $\delta \geq 1$, then:

$$k < \lfloor \delta \mu e^{-1/3} \rfloor \implies \mathbb{P}[|W - \mu| > \delta \mu] \leq e^{-\lfloor k/2 \rfloor}$$

Corollary 6.8. $\mathbb{P}\left[W > \frac{4\hat{p}mk}{n}\right] \leq e^{-\lfloor k/2 \rfloor}$

Proof. We can use the fact the random variables W_v are k -wise independent to apply [Theorem 6.7](#). First, we ensure that $k < \lfloor \delta \mu e^{-1/3} \rfloor$, that we achieve by setting $\delta = \frac{mk}{m_h}$.

Recall that m_h is the number of heavy edges (ones with at least one heavy end-point), and m is the total number of edges in the original graph G .

$$\delta = \frac{mk}{m_h} \implies \delta \mu e^{-1/3} = \frac{mk \cdot 2\hat{p}m_h}{m_h \cdot n} \cdot e^{-1/3} > \frac{\hat{p}mk}{n} \implies \delta \mu e^{-1/3} > k$$

In the last step, we used the fact that $S > 100kn^2/m$ from [Eq. \(3\)](#), to imply that $\hat{p}m/n > 1$. Therefore, we can now apply [Theorem 6.7](#) to conclude:

$$\begin{aligned} & \mathbb{P}[|W - \mu| > \delta \mu] \leq e^{-\lfloor k/2 \rfloor} \\ \implies & \mathbb{P}\left[W > \mu + \frac{2\hat{p}mk}{n}\right] \leq e^{-\lfloor k/2 \rfloor} \\ \implies & \mathbb{P}\left[W > \frac{4\hat{p}mk}{n}\right] \leq e^{-\lfloor k/2 \rfloor} \end{aligned}$$

In the second step, we used the fact that $\mu = \mathbb{E}[W] = \sum_{v \in V_{heavy}} \hat{p} \cdot \frac{d(v)}{n} \leq \frac{2m\hat{p}}{n}$. □

Now we are finally ready to upper bound the number of heavy edges in $G[V_i]$. With high probability (using [Corollary 6.5](#)), the following holds:

$$\begin{aligned} \#(\text{heavy edges in } G[V_i]) & \leq \sum_{v \in V_{heavy}} \mathbb{P}[v \in V_i] \cdot (3d(v)\hat{p}) \\ & \leq \sum_{v \in V_{heavy}} W_v \cdot n \cdot (3\hat{p}) = 3n\hat{p} \cdot W \\ & \leq 12\hat{p}^2 mk = \frac{12S}{100} < \frac{S}{8} \end{aligned}$$

Theorem 6.9 (Heavy edges). *With high probability, the number of edges in $G[V_i]$ that have some endpoint with degree larger than τ is at most $S/8$.*

Combining this result with [Theorem 6.9](#), we conclude the following:

Theorem 6.10. *With probability at least $9/10$, the maximum number of edges in any of the $G[V_i]$ s (where $i \in [R]$) does not exceed S , and hence [Algorithm 7](#) does not terminate on [Line 4](#).*

6.4.3 Upper-Bounding the Number of Messages *Sent* by any Machine

Recalling [Algorithm 8](#), we note that the number of messages *received* by the machine containing V_i , is equal to the number of edges in $G[V_i]$. Therefore, the last section essentially proved that the number of messages (edges) *received* by a particular machine is upper-bounded by S . Conversely,

in this section, we will justify that the number of messages *sent* by any machine is $O(S)$. Since the number of edges stored in a machine is $\leq S$, it suffices to show that for each edge e , [Algorithm 8](#) sends only $O(1)$ messages (each message is a copy of the edge e).

Let $Z_i^{(e)}$ be the $\{0, 1\}$ indicator random variable for $e \in G[V_i]$, and let $Z^{(e)}$ be the sum of $Z_i^{(e)}$ for all $i \in [\mathcal{M}]$. Here, $Z^{(e)}$ represents the number of messages that are created by edge e . Additionally we make $r = S\mathcal{M}/m = O_\varepsilon(\log n)$ copies of each edge e , and ensure that all replicates reside on the same machine. We distribute the $Z^{(e)}$ messages evenly amongst the replicates, so that each replica is only responsible for $Z^{(e)}/r$ messages.

Since all replicates are on the same machine, this last step is purely conceptual, but it will simplify our argument, by allowing us to charge the outgoing messages to each replicate (as opposed to each edge). Our goal will be show that each replicate is responsible for only $O(1)$ messages, which is the same as showing that w.h.p. $Z^{(e)}/r = O(1)$.

Clearly $\mu = \mathbb{E}[Z^{(e)}] = \hat{p}^2 \cdot \mathcal{M} = \frac{S\mathcal{M}}{100mk}$. This allows us to apply [Lemma 6.4](#) with $\delta = \frac{100e^{1/3}mk^2}{S\mathcal{M}}$

$$\mathbb{P}\left[Z^{(e)} > \delta\mu\right] \leq e^{-\lfloor k/2 \rfloor} = \frac{1}{n^3} \implies \mathbb{P}\left[\frac{Z^{(e)}}{r} > \frac{e^{1/3}k}{r}\right] \leq \frac{1}{n^3}$$

Using the assumption (from [Eq. \(3\)](#)) that $\mathcal{M} > 2000mk/S \implies r > 2000k$, we see that with high probability, the number of messages sent by any replicate is bounded above by $e^{1/3}/2000 \leq 1$. So, the number of messages sent from any machine is bounded by S with high probability.

6.5 Challenge (4): \hat{T} is close to its expected value

In this section, we provide merely a brief discussion of this challenge for intuition, and we fully analyze the approximation guarantees of our algorithm in [Section 6.3.1](#). That analysis also makes clear the source of our advertised lower-bound on T for which an estimated count concentrates well.

Lower Bound on Number of Triangles. In order to output *any* approximation (note that we are ignoring all factors of ε and $O(\text{poly } \log n)$ here) to the triangle count, we must see $\Omega(1)$ triangles amongst all of the induced subgraphs on all the machines. The expected number of triangles in a specific induced $G[V_i]$ is \hat{p}^3T , and therefore, the expected number of triangles overall is $\hat{p}^3T\mathcal{M}$ which must be $\Omega(1)$ for some setting of T . Since we set \hat{p} such that $\hat{p}^2m = \Theta(S)$, this gives that $\hat{p}^2 = O(S/m)$ which implies $\hat{p}^2 \cdot \mathcal{M} = \hat{p}^2 \cdot (m/S) = \Theta(1)$. This then immediately implies that to show that \hat{p}^3T is $\Omega(1)$, we need only show that $\hat{p} \cdot T$ is $\Omega(1)$. Specifically, we show in [Lemma 6.2](#) that when $T > 1/\hat{p}$, we can obtain a $(1 \pm \varepsilon)$ -approximation. To get some intuition for this lower bound on T , note that, in the linear memory regime, when $S = \Theta(n)$, this translates to $T > \sqrt{d_{avg}}$, where d_{avg} is the average degree of G .

$$T > \frac{1}{\hat{p}} = \tilde{\Theta}\left(\sqrt{\frac{m}{S}}\right) \xrightarrow{\text{for } S=\tilde{\Theta}(n)} T > \tilde{\Theta}\left(\sqrt{d_{avg}}\right).$$

6.5.1 Getting the High Probability Bound

By building on [Lemma 6.2](#) and [Algorithm 7](#), we design [Algorithm 9](#) that outputs an approximate triangle counting with high probability, as opposed with only constant success probability. It is important to note that in the below algorithm, all $O(\log n)$ *independent iterations* ([Line 3](#)) are done *in parallel*, simultaneously, not sequentially.

Algorithm 9. Approximate Triangle Counting

```

1: function APPROX-TRIANGLES-MAIN( $G = (V, E)$ )
2:   Let  $I \leftarrow 100 \cdot \log n$ .
3:   parfor  $i \leftarrow 1 \dots I$  do  $\triangleright$  Perform all  $I$  iterations in parallel simultaneously in  $O(1)$  rounds.
4:     Let  $Y_i$  be the output of Algorithm 7 invoked on  $G$ . We assume that each invocation of Algorithm 7 uses fresh randomness compared to previous runs.
5:   end parfor
6:   Let  $\mathcal{Y}$  be the list of all  $Y_i$ , for  $i = 1 \dots I$ .
7:   Sort  $\mathcal{Y}$  in non-decreasing order.
8:   return the median of  $\mathcal{Y}$ 

```

We have the following guarantee for [Algorithm 9](#).

Theorem 6.11. *Let Y be the output of [Algorithm 9](#). Then, with high probability it holds*

$$|Y - T| \leq \varepsilon T.$$

In the proof of this theorem we use the following concentration bound.

Theorem 6.12 (Chernoff bound). *Let X_1, \dots, X_k be independent random variables taking values in $[0, 1]$. Let $X \stackrel{\text{def}}{=} \sum_{i=1}^k X_i$ and $\mu \stackrel{\text{def}}{=} \mathbb{E}[X]$. Then, for any $\delta \in [0, 1]$ it holds $\mathbb{P}[X \leq (1 - \delta)\mu] \leq \exp(-\delta^2\mu/2)$.*

Proof of [Theorem 6.11](#). The proof of this theorem is essentially the so-called ‘‘Median trick’’. We provide full proof here for completeness.

Let Y_i be as defined on [Line 4](#) of [Algorithm 9](#). By [Theorem 6.10](#), with probability at most $1/10$ [Algorithm 7](#) terminates due to creating too big subgraphs. If we ignore [Line 4](#) of [Algorithm 7](#), then by Property (A) of [Lemma 6.2](#) we have $\mathbb{E}[Y_i] = T$. Y_i significantly deviates from its expectation if [Algorithm 7](#) terminates on [Line 4](#) or if the estimate Y_i is simply off. Define a 0/1 variable Z_i which equals 1 iff $|Y_i - T| \leq \varepsilon T$. By union bound on Property (B) of [Lemma 6.2](#) and [Theorem 6.10](#), we have $\mathbb{P}[Z_i = 1] \geq 1 - 1/10 - 1/10 = 4/5$. Also, following [Line 4](#) of [Algorithm 9](#) we have that all Z_i are independent.

Let $Z = \sum_{i=1}^I Z_i$. We have that $\mathbb{E}[Z] \geq \frac{4}{5}I$, implying that in expectation at least $4/5$ fraction on Z -variables are 1. We now bound the probability that at least $2/5$ of these variables equal 0, i.e., at most $3/5$ of them equal 1. Since Z -variables are independent, for this we can use [Theorem 6.12](#), obtaining

$$\mathbb{P}\left[Z \leq \frac{3}{5}I\right] \leq \mathbb{P}\left[Z \leq \left(1 - \frac{1}{5}\right)\mathbb{E}[Z]\right] \leq \exp(-\mathbb{E}[Z]/50).$$

Given that $I = 100 \cdot \log n$ (see [Line 2](#) of [Algorithm 9](#)), we derive that $\mathbb{P}\left[Z \leq \frac{3}{5}I\right] < n^{-1}$. This now implies that with probability at least n^{-1} the output of [Algorithm 9](#) is some Y_j such that $Z_j = 1$. This completes the analysis. \square

6.6 Showing Concentration for the K -Subgraph Count

Using similar analysis to the previous section, in this section, we show the expectation and concentration bound of our subgraph counting algorithm for any subgraphs consisting of K nodes where K is constant. Let this subgraph be H .

Lemma 6.13. *Let \hat{B} be the count of subgraph H (with K vertices) in $G[V_i]$ and $\mathcal{M} = \frac{20}{\varepsilon^2 \hat{p}^{K-1}}$ be as defined in [Eq. \(3\)](#), and assume that $B \geq 1/\hat{p}$. Then, the following hold:*

(A) $\mathbb{E} [\hat{B}] = \hat{p}^K \cdot R \cdot B$, and

(B) $\mathbb{P} [|\hat{B} - \mathbb{E} [\hat{B}]| > \varepsilon \mathbb{E} [\hat{B}]] < \frac{1}{10}$.

Proof. We first prove [Item A](#). The probability that a particular K -vertex occurrence h of H appears in machine M_i is $\mathbb{P} [h \in G[V_i]] = \hat{p}^K$. There are B number of occurrences of H in G . Thus, the expected number of occurrences of H in machine M_i is $\mathbb{E} [\hat{B}_{M_i}] = \hat{p}^K \cdot B$. Since there are R machines (which did not exceed the memory limit), the expected number of occurrences of H in all R machines is $\mathbb{E} [\hat{B}] = \sum_{i \leq R} \mathbb{E} [\hat{B}_{M_i}] = \hat{p}^K R B$.

We now prove [Item B](#). Let $H(G)$ be the set of occurrences of H in G . Let $\hat{B}_{i,h}$ be a random variable where $\hat{B}_{i,h} = 1$ if h , a particular occurrence of H in G , is in machine i ; otherwise, $\hat{B}_{i,h} = 0$. Then,

$$\text{Var} [\hat{B}] = \mathbb{E} [\hat{B}^2] - \mathbb{E} [\hat{B}]^2 \quad (9)$$

$$= \mathbb{E} \left[\left(\sum_{i=1}^R \sum_{h \in H(G)} \hat{B}_{i,h} \right)^2 \right] - \left(\sum_{i=1}^R \sum_{h \in H(G)} \mathbb{E} [\hat{B}_{i,h}] \right)^2. \quad (10)$$

First, consider random variables \hat{B}_{i,h_1} and \hat{B}_{j,h_2} ; for each $i \neq j \in [R]$ and each $h_1, h_2 \in H(G)$, there exists a term in the first summand of [Section 6.6](#) containing $\mathbb{E} [2\hat{B}_{i,h_1}\hat{B}_{j,h_2}]$. The vertices constituting h_1 and h_2 are $2K$ distinct copies of some not necessarily distinct copies of vertices in V . Suppose we use a $2k$ -wise independent hash function, then we have $\mathbb{E} [2\hat{B}_{i,h_1} \cdot \hat{B}_{j,h_2}] = 2\mathbb{E} [\hat{B}_{i,h_1}] \cdot \mathbb{E} [\hat{B}_{j,h_2}]$. We see that this term also shows up in the second summand of [Section 6.6](#). Hence, the terms cancel for each $i \neq j$ and we can simplify [Section 6.6](#) to the following.

$$\text{Var} [\hat{B}] = \sum_{i=1}^R \mathbb{E} \left[\left(\sum_{h \in H(G)} \hat{B}_{i,h} \right)^2 \right] - \sum_{i=1}^R \left(\sum_{h \in H(G)} \mathbb{E} [\hat{B}_{i,h}] \right)^2 = \sum_{i=1}^R \text{Var} [\hat{B}_i]. \quad (11)$$

Now, what remains is to upper bound $\text{Var} [\hat{B}_i]$. Using the same approach as in the previous section with the observation that any two distance occurrences h_1 and h_2 must contain $K+1 \leq k \leq 2K$ distinct vertices. This means that $\mathbb{E} [\hat{B}_{i,h_1} \cdot \hat{B}_{i,h_2}] = \hat{p}^k \leq \hat{p}^{K+1}$. Then, we can bound

$$\text{Var} [\hat{B}_i] \leq 2B^2 \hat{p}^{K+1} \quad (12)$$

(assuming $B \geq 1/\hat{p}$).

Then, from [Section 6.6](#) and [Section 6.6](#), we get the bound on the variance to be $\text{Var} [\hat{B}] \leq 2RB^2 \hat{p}^{K+1}$.

By Chebyshev's inequality and [Item A](#), we compute

$$\mathbb{P} [|\hat{B} - \mathbb{E} [\hat{B}]| > \varepsilon \mathbb{E} [\hat{B}]] < \frac{\text{Var} [\hat{B}]}{\varepsilon^2 \mathbb{E} [\hat{B}]^2} \leq \frac{2B^2 \hat{p}^{K+1}}{\varepsilon^2 \hat{p}^{2K} R^2 B^2} = \frac{2}{\varepsilon^2 \hat{p}^{K-1} R^2}. \quad (13)$$

When setting $R \geq \frac{20}{\varepsilon^2 \hat{p}^{K-1}}$, we obtain [Item B](#). \square

Using [Lemma 6.13](#), we can obtain the following theorem about counting occurrences of *any* subgraph H with K vertices.

Theorem 6.14. *Let $G = (V, E)$ be a graph over n vertices, m edges, and let B be the number of occurrences of subgraph H with K vertices in G . Assuming*

$$(i) \ B = \tilde{\Omega} \left(\left(\frac{m}{S} \right)^{K/2-1} \right), \quad (ii) \ S = \tilde{\Omega} \left(\max \left\{ \frac{\sqrt{m}}{\varepsilon}, \frac{n^2}{m} \right\} \right),$$

there exists an MPC algorithm, using \mathcal{M} machines, each with local space S , and total space $\mathcal{M}S = \tilde{O}_\varepsilon(m)$, that outputs a $(1 \pm \varepsilon)$ -approximation of B , with high probability, in $O(1)$ rounds.

Corollary 6.15. *Let $G = (V, E)$ be an input graph and B be the number of occurrences of subgraph H with K vertices in G . If $B \geq d_{avg}^{K/2-1}$, then there exists an MPC algorithm that in $O(1)$ rounds with high probability outputs a $(1 + \varepsilon)$ -approximation of B . This algorithm uses a total space of $\tilde{O}(m)$ and space $\tilde{\Theta}(n)$ per machine. d_{avg} is the average degree of the vertices in the graph.*

7 Counting subgraphs of size at most 5 in bounded arboricity graphs

In this section, we present a procedure that for every subgraph H such that $|H| \leq 5$, counts the exact number of occurrences of H in G in $O(\sqrt{\log n})$ rounds and $O(m\alpha^3)$ total memory, where as before, α is an upper bound on the arboricity of G ⁶. The procedure is based on a recent paper by Bera, Pashanasangi and Seshadhri [[BPS20](#)] (henceforth BPS) which presented an $O(m\alpha^3)$ time and space algorithm for the same task in the sequential model. We will start by a short description of the BPS result, and then continue to explain how to implement it in the MPC model.

7.1 The BPS algorithm

BPS generalize the ideas of Chiba and Nisheziki [[CN85](#)] for counting constant-size-cliques and 4-cycles in the classical sequential model to counting all subgraphs of up to 5 nodes in $O(n + m\alpha^3)$ time. Let H be the subgraph in question. The main idea of BPS is as follows. The algorithm starts by computing a degeneracy ordering of G , which is an acyclic orientation of G , denoted \vec{G} , where each vertex has at most $O(\alpha)$ outgoing neighbors. The idea is then to consider all acyclic orientations of H (up to isomorphisms), and for each such acyclic orientation \vec{H} , count the number of occurrences of \vec{H} in \vec{G} , as described next. The algorithm computes what is referred to as a largest directed rooted tree subgraph (DRTS) of \vec{H} , denoted \vec{T} . That is, the DRTS \vec{T} is a largest (in number of vertices) tree that is contained in \vec{H} such that all of the edges are directed away from the root of \vec{T} . Given a DRTS \vec{T} , proceed by looking for all copies of \vec{T} in \vec{G} . Once a copy of \vec{T} is found, it needs to be verified whether it can be extended to a copy of \vec{H} in \vec{G} . This verification is based on the observation that for any directed subgraph \vec{H} on at most 5 vertices, and for every largest directed rooted tree \vec{T} of \vec{H} , the complement of \vec{T} in \vec{H} is a collection of rooted paths and stars⁷. Therefore, all potential completions of a copy of \vec{T} to \vec{H} in \vec{G} can be computed and hashed in time $O(m \cdot \text{poly}(\alpha))$. See figure below for an illustration of a possible \vec{H} and its DRTS \vec{T} (adapted from [[BPS20](#)]). Hence, whenever a copy of \vec{T} is discovered in \vec{G} , it can be verified in

⁶Strictly speaking, we will have $\alpha \leq 5\alpha(G)$ but as this does not affect the asymptotic bounds, it is easier to just relate to it as the exact arboricity.

⁷This does not hold for subgraphs H that are stars, but stars can be dealt with differently.

$O(1/\delta)$ rounds whether this copy can be extended to \vec{H} . Since all copies of \vec{T} can be enumerated in $O(m\alpha^3)$ time, the overall algorithm takes $O(m\alpha^3)$ time.




Figure 1: From left to right: A subgraph H ; a possible directed copy of H ; the DRTS in green, and its complement with respect to H in red. Based on a figure from BPS [BPS20].

7.2 Implementation in the MPC model

Notation 7.1 (Outgoing neighbors and out-degree). Let $\vec{G} = (V, \vec{E})$ be a directed graph. For a vertex $v \in V$, We denote by $N^+(v)$ its set of outgoing neighbors, and by $d^+(v) = |N^+(v)|$ its outgoing degree or out-degree.

Definition 7.2 (Degeneracy and degeneracy ordering.). A degeneracy ordering of a graph G , is an ordering obtained by repeatedly removing a minimum degree vertex and all the edges incident to this vertex. A vertex u precedes a vertex v in this ordering, $u \prec v$, if u was removed before v . The degeneracy of a graph G is then the maximum outgoing degree over all vertices in a degeneracy ordering of G .

Theorem 7.3 (Thm 2 in [GLM19]). Given a graph G with arboricity α , it outputs, with high probability, an orientation of G , \vec{G} , where each vertex in \vec{G} has out-degree at most $O(\alpha)$. The algorithm performs $O(\sqrt{\log n} \cdot \log \log n)$ rounds, uses $\tilde{O}(n^\delta)$ space per machine, for an arbitrary constant $\delta \in (0, 1)$, and the total memory is $O(\max\{m, n^{1+\delta}\})$.

The following is a key lemma.

Lemma 7.4. Let \vec{G} be a directed graph over m edges such that each vertex has out-degree at most α . Let \vec{T} be a directed rooted tree of size $t \geq 2$. We can list all copies of \vec{T} in G in $O(1/\delta)$ rounds, $O(n^{2\delta})$ space per machine, and $O(m \cdot \alpha^{t-2})$ total memory.

Proof. Let a_1, \dots, a_t denote the vertices of \vec{T} , where a_1 is the root, and a_i is the i^{th} vertex with respect to the BFS ordering of \vec{T} . Let \vec{T}_i denote $\vec{T}[\{a_0, \dots, a_i\}]$.

We prove the claim by induction on t . For $t = 2$, all edges in G are copies of \vec{T} , so the claim holds trivially.

Assume that the claim holds for i , and we now prove it for $i + 1$. By the assumption, in $O(1/\delta)$ rounds and $O(m\alpha^{i-2})$ total memory, all copies of \vec{T}_i can be listed. We will show that we can use these copies to find all copies of \vec{T}_{i+1} in $O(1/\delta)$ rounds and $O(m\alpha^{i-1})$ memory. Recall that we have machines with $O(n^{2\delta})$ memory. We will divide the copies among the machines, so that each machine only holds $O(n^\delta)$ copies. Let M be some machine containing copies $\tau_1, \dots, \tau_{n^\delta}$ of \vec{T}_i . It will be easier to think of M as a collection of n^δ constant memory parts, each holding a single copy of \vec{T}_i . Consider a specific copy τ of \vec{T}_i and let P_τ denote the part storing that copy. Let a_p denote the vertex in \vec{T} that is the parent of a_{i+1} , and let u denote the vertex in τ that is mapped to a_p . We would like to create all tuples (τ, w) , where $w \in N^+(u) \setminus \tau$ and w can be mapped to a_{i+1} . In order

to achieve this we duplicate P_τ for α times, to get copies $P_{\tau,1}, \dots, P_{\tau,\alpha}$. Each part $P_{(\tau,i)}$ then asks u for its i^{th} neighbor w , and then checks if τ can be extended to \vec{T}_{t_1+1} using w . If (τ, w) is a copy of \vec{T}_{i+1} , then the part creates the tuple (τ, w) . All the the duplications above can be done in parallel to all copies of \vec{T} residing on a single machine, so that in total each machine is duplicated α time. Since each machine has $O(n^\delta)$ information, and $O(n^{2\delta})$ space, by [Lemma 2.3](#), this process takes $O(\log_{n^\delta} \alpha) = O(1/\delta)$ rounds. Furthermore, as each machine is duplicated α times, the amount of total memory increases by a factor of α .

Hence, at the end of the process, all copies of \vec{T} are generated, the round complexity is $O(1/\delta)$, and the total memory is $O(m\alpha^{t-2})$. \square

For a directed graph \vec{G} , we consider the following lists of key-value pairs, as described in [Lemma 15](#) in [\[BPS20\]](#).

- \mathcal{HM}_1 : $((u, v), 1)$ for all $(u, v) \in E(\vec{G})$.
- \mathcal{HM}_2 : $(S, \ell) \forall S \subseteq V(\vec{G})$ such that $1 \leq |S| \leq 4$ and ℓ is the number of vertices u such that $S \subseteq N^+(u)$.
- \mathcal{HM}_3 : $((S_1, S_2, \ell)) \forall S_1, S_2 \subseteq V(\vec{G})$, where $1 \leq |S_1 \cup S_2| \leq 3$, and ℓ is the number of edges $e = (u, v) \in E(\vec{G})$ such that $S_1 \subseteq N^+(u)$ and $S_2 \subseteq N^+(v)$.

Lemma 7.5. *Let \vec{G} be a directed graph with m edges, such that for every $v \in V(\vec{G})$, $d^+(v) \leq \alpha$. The lists $\mathcal{HM}_1, \mathcal{HM}_2$ and \mathcal{HM}_3 can be computed in $O(1/\delta)$ rounds and $O(m\alpha^3)$ total memory.*

Proof. In order to create \mathcal{HM}_1 , each vertex u simply adds for each $v \in N^+(u)$ the pair $((u, v), 1)$ to the list. Clearly this can be done in $O(1)$ rounds, and $O(m)$ total memory.

We now consider \mathcal{HM}_2 . Let $s = |S|$ denote the size of the requested set. Fix s , and let \vec{T} be a DRT which consists of a root and s outgoing neighbors. By [Lemma 7.4](#), we can generate all copies of \vec{T} in $O(1/\delta)$ rounds, and $O(m \cdot \alpha^{s-2}) = O(m \cdot \alpha^2)$ total memory. From each copy (v, u_1, \dots, u_s) of \vec{T} , we create a tuple $(\{u_1, \dots, u_s\}, 1)$ and add it to a temporary list \mathcal{HM}'_2 . Finally, we use [Theorem 2.1](#) to sort this list and aggregate the counts of each set $S = \{u_1, \dots, u_s\}$, so that for every S we create the tuple (S, ℓ) and add it to \mathcal{HM}_2 , where ℓ is the number of occurrences of the tuple $(S, 1)$ in \mathcal{HM}'_2 . By [Theorem 2.1](#), this process takes $O(\log_{n^\delta} m \cdot \alpha^2) = O(1/\delta)$ rounds.

\mathcal{HM}_3 is constructed similarly. Fix some s_1 and s_2 such that $1 \leq s_1 + s_2 \leq 3$, and consider the corresponding DRT \vec{T} . That is, \vec{T} is a DRT with a vertex u with s_1 outgoing neighbors, where one of the neighbors has s_2 additional outgoing neighbors. This is a DRT over $|S_1| + |S_2| + 1 \leq 4$ vertices, so by [Lemma 7.4](#), we can generate all copies in $O(1/\delta)$ rounds, and $O(m \cdot \alpha^2)$ total memory. From the list of all copies we can generate \mathcal{HM}_3 , similarly to as described for \mathcal{HM}_2 , in $O(1/\delta)$ rounds. \square

Theorem 7.6. *Let $G = (V, E)$ be a graph with $n = |V|$ and $m = |E|$. There is an algorithm for counting the number of occurrences of any given subgraph H over $k \leq 5$ vertices in G with high probability, with round complexity $O(\sqrt{\log n} + 1/\delta)$, $O(n^{2\delta})$ memory per machine, and $O(m\alpha^3)$ total memory.*

Proof. If H is a k -star, then the number of occurrence of H in G is simply $\sum_{v \in V} \binom{d(v)}{k}$ where $\binom{d(v)}{k} = 0$ for $k > d(v)$, which can be computed in $O(1)$ rounds. Hence, we assume that H is not a star.

The first step in the algorithm of BPS is to direct the graph G according to the degeneracy ordering (see [Definition 7.2](#)). We achieve this using the algorithm of [\[GLM19\]](#) described in [Theorem 7.3](#). Note that the algorithm of [\[GLM19\]](#) returns an approximate degeneracy ordering, but as the degeneracy of a graph is at most twice the arboricity, it holds that each vertex has out-degree $O(\alpha)$.

Given the ordering of \vec{G} , the algorithm continues by considering all orientations \vec{H} of H (up to isomorphisms). For each \vec{H} it computes the maximal rooted directed tree, DRT, of \vec{H} , denoted \vec{T} . As H is of constant size, this can be computed in $O(1)$ rounds on a single machine.

The next step is to find all copies of \vec{T} in \vec{G} . By [Lemma 7.4](#), this can be implemented in $O(1/\delta)$ rounds, $O(n^{2\delta})$ space per machine, and $O(m\alpha^2)$ total memory.

Now, for each copy of \vec{T} in \vec{G} it needs to be verified if the copy can be completed to a copy of \vec{H} in \vec{G} . By Lemma 16 in [\[BPS20\]](#), this can be computed in if given query access to $\mathcal{HM}_1, \mathcal{HM}_2$ and \mathcal{HM}_3 , as defined in [Section 7.2](#). That is, it can be determined if a copy τ of \vec{T} using $O(|H|^2) = O(1)$ queries to the lists $\mathcal{HM}_1, \mathcal{HM}_2$ and \mathcal{HM}_3 . By [Lemma 7.5](#), these lists can be generated in $O(1/\delta)$ rounds, and $O(m\alpha^2)$ total memory. For $i \in [1..3]$, let Q_i denote the set of all queries to list \mathcal{HM}_i . By [\[GSZ11\]](#), all queries Q_i to \mathcal{HM}_i can be answered in time $O(1/\delta)$.

Finally, by Lemma 16 in [\[BPS20\]](#), each v can use the answers to its queries to compute the number of copies of \vec{H} it can be extended to. Therefore, by summing over all vertices and over all possible orientations of H , and taking into account isomorphisms, we can compute the number of occurrences of H in \vec{G} . The total round complexity is dominated by computing the approximate arboricity orientation of G and the sorting operations. Therefore the round complexity is $O(\sqrt{\log n} \log \log n + 1/\delta)$. The space per machine is $O(n^{2\delta})$, and the total memory over all machines is $O(m\alpha^3)$. \square

8 Experiments

We performed experiments using our algorithms given in [Sections 4](#) and [6](#). Our code [\[exp20\]](#) simulates the algorithms described in these sections as well as the MPC procedures we use as subroutines. In the implementation of the algorithm of [Section 6](#), our algorithm achieves a better approximation *on all tested graphs* than the best-known previous algorithm. In the implementation of the exact algorithm of [Section 4](#), our algorithm achieves fewer number of MPC rounds than the baseline algorithm. Given that the focus of this paper is on our theoretical contributions, we include these experiments only as a proof-of-concept below. All real-world graphs on which we performed our experiments can be found in the Stanford Large Network Dataset Collection (SNAP) [\[LK14\]](#). We leave as interesting future directions implementing and testing our algorithms in massively parallel software frameworks, such as Apache Hadoop and others, on much larger graphs.

We performed experiments using our algorithms given in [Sections 4](#) and [6](#). Our code [\[exp20\]](#) simulates the algorithms described in these sections as well as the MPC procedures we use as subroutines. In the implementation of the algorithm of [Section 6](#) we output the approximation factor we achieve using our algorithm versus the amount of space per machine. In the implementation of the exact algorithm of [Section 4](#) we output the number of MPC rounds necessary to execute the algorithm versus the arboricity bound we pass into it. Given that the focus of this paper is on our theoretical contributions, we include these experiments only as a proof-of-concept. We leave as interesting future directions implementing and testing our algorithms in massively parallel software frameworks, such as Apache Hadoop and others, on much larger graphs.

All real-world graphs on which we performed our experiments can be found in the Stanford

Large Network Dataset Collection (SNAP) [LK14].

We test our algorithms against datasets described in Table 2.

File	Number of Vertices (n)	Number of Edges (m)	Number of Triangles (T)
email-Eu-core	1005	25571	105461
ego-Facebook	4039	88234	1612010
feather-lastfm-social	7624	27806	40433
ca-GrQc	5242	14496	48260
musae-twitch (DE)	9498	153138	603088
ca-HepTh	9877	25998	28339
oregon1_010519	11051	22724	17677
ca-HepPh	12008	118521	3358499
email-Enron	36692	183831	727044

Table 2: All datasets can be found in the Stanford Large Network Dataset (SNAP) Collection [LK14]. This table shows the number of vertices, edges, and exact number of triangles in each of these graphs.

Results for Section 4. The first set of experiments were performed using our new exact triangle count algorithm provided in Section 4. The results are shown in Fig. 2. Our experiments were performed on five datasets: **oregon1_010519**, **email-Enron**, **ca-HepTh**, **ca-GrQc**, **email-Eu-core**. We compare against the baseline algorithm (labeled “-base” in the figures) of removing (and counting) the vertices with degree at most the degeneracy of each graph during each round. We measure the number of rounds our algorithm takes against the amount of space per machine (indicated by the different colors of the bars) and our *initial* setting of our degree bound. Recall that since the degree bound of our algorithm for removal of vertices grows doubly exponentially, we can set our initial degree bound to be smaller than the degeneracy of the graph. Note that for the baseline algorithm, we cannot do this since the degree bound remains the same (and hence, a smaller degree bound than the degeneracy will cause the algorithm to terminate at some point with vertices still left in the graph with degree greater than the bound). Our initial degree bound settings are shown on the x-axis while the y-axis shows the number of rounds. The machine space bounds are in terms of the *number of nodes* of the graph that can fit in each machine. These numbers are shown in the legend indicating the colors of the bars. Due to its usefulness in the baseline algorithm, we consider the degeneracy of the graphs instead of the arboricity. But as we noted previously, the degeneracy of the graph is at most within a constant factor of 2 of the arboricity of the graph.

We see in Fig. 2 that our algorithms result in less (or the same number of) rounds than the baseline algorithm for all cases given the same space and degree bound except the 2225 case for the **email-Enron** dataset. This confirms our theoretical analysis of the asymptotic number of rounds of our algorithm, taking $O_\delta(\log \log n)$ compared to the $O_\delta(\log n)$ rounds we expect the baseline algorithm to take. The anomaly with the single 2225 case might be due to the larger constant factor (derived from MPC sort and find duplicate) of having to sort more items per round in our case compared to the baseline case. We also see that due to the doubly exponential nature of our algorithm, we are able to also count triangles using much lower initial degree bounds than the degeneracy of the graph. Even for such degree bounds, the number of rounds necessary is still often (much) less than the number of rounds necessary for the baseline algorithm. This provides the advantage of being able to use our algorithm for real-world graphs without first needing to determine their degeneracy value.

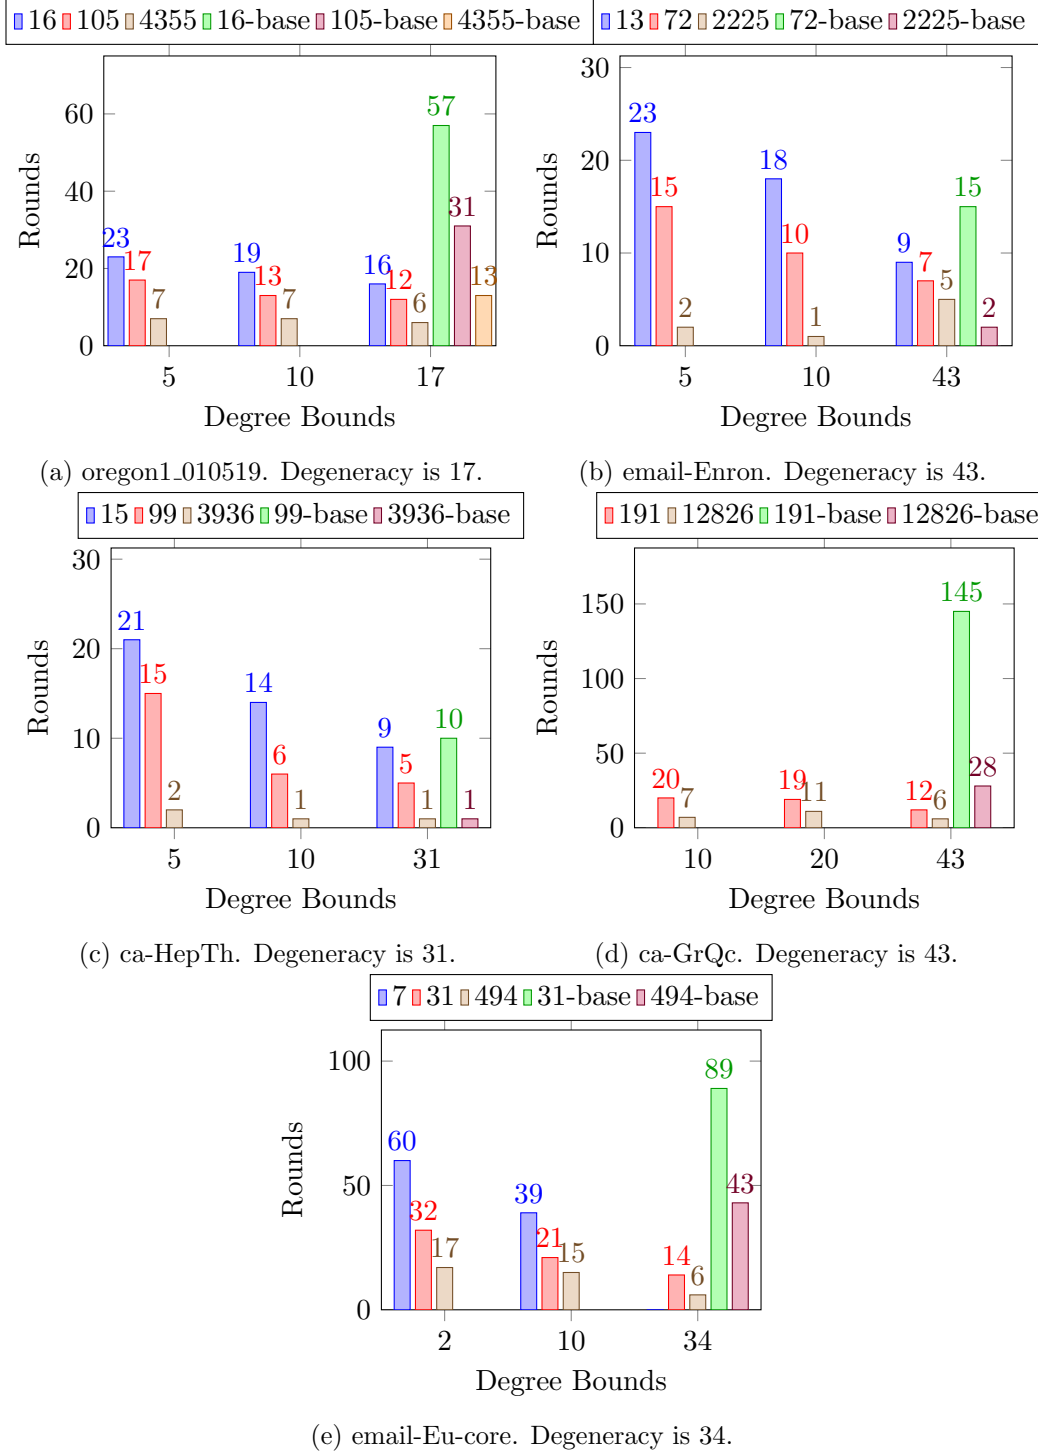


Figure 2: This set of graphs shows the results of our experiments using our exact counting algorithm described in Section 4. We test on five datasets labeled under each plot. In each of these graphs, we compare against the number of rounds required by the MPC algorithm that removes, in each round, only vertices with degree at most the degeneracy of the graph α . Each color represents a different space per machine, which is represented in terms of the number of nodes that can fit in each machine. The colors (green, red, yellow) labeled with “-base” represent our baseline algorithm results.

File	δ	Partition Approximation	Our Approximation
ego-Facebook	0.5	0.62	1.31
feather-lastfm-social	0.5	5.41	1.08
ca-GrQc	0.5	4.53	1.64
ca-HepPh	0.5	0.66	1.22
ego-Facebook	0.75	0.75	0.97
ca-GrQc	0.75	5.82	0.82
ca-HepPh	0.75	5.90	0.86
musae-twitch (DE)	0.75	0.74	0.95
oregon1_010519	0.75	0.60	0.71

Table 3: The approximation factors obtained when running our algorithm given in Section 6 against our implementation of the partition algorithm given in Algorithm 1 and Algorithm 2 of [PT12]. We perform the algorithms on machines of size $2m^\delta \cdot \log n$. The approximation factor is calculated by C/T where C is the triangle count returned by either algorithm and T is the actual count of the triangles in the graph.

Results for Section 6. The results of experiments for our approximation algorithm described in Section 6 are given in Table 3. We further compare our approximations against our implementation of the partition algorithm given in Algorithms 1 and 2 of [PT12]. In the implementation of our algorithm, due to the (sequential) time constraints of simulating our algorithms, we do not use a k -wise independent hash function, as such functions require too much time to compute. Hence, for our experiments, we use standard pseudorandom functions given in programming packages (specifically `numpy` in Python3). For each of the experiments the space per machine is $2m^\delta \cdot \log n$, where δ is specified in Table 3. The same space per machine is used for both algorithms. The total space used for both algorithms is $2m \cdot \log m$. For the implementation of our algorithm, we set the probability of sampling to be $\frac{1}{5} \cdot \sqrt{\frac{S}{M \cdot k}}$ where we set $k = \max(2, \lfloor 6.5 \cdot \log n \rfloor)$. We chose to test these algorithms on these specific δ values because $\delta = 0.5, 0.7$ represent $\tilde{O}(n)$ and $\tilde{o}(m)$, respectively. Because the theoretical guarantees of our algorithm relies on some specific constraints on T and S , we wanted to see how our algorithm performs on real-world networks. We use the median-of-means trick for the concentration for both algorithms.

As Table 3 shows, compared to the partition algorithm, our algorithm obtains a better approximation ratio for all datasets and for all machine spaces. This follows from our theoretical analysis as we ensure $(1 + \varepsilon)$ -approximations on the number of triangles in each graph using $\tilde{O}(n)$ space and with a quadratically smaller constraint on the number of actual triangles in the graph than all other previous work. Thus, we show that practically, on real-world graphs, our algorithm obtains better approximations on the number of triangles given smaller space per machine.

9 Open Questions

There are many interesting open questions that result from our study; among these open questions include improving the bounds presented in our algorithm: the round complexity and total space usage in our exact algorithms and the space per machine in our approximation algorithms. In addition to these questions, we also discuss two additional open questions with a larger research scope.

Small subgraph counting counting for a broader class of small subgraphs Two recent works of [Bre19, BPS21] extend the result of [BPS20] to a broader set of small subgraphs in the sequential model. However, their results depend crucially on a *DAG tree decomposition* which is non-trivial to implement in the MPC model. Furthermore, even given this DAG tree decomposition, their approach requires iterating through the tree from the leaf level by level up the tree. Such a procedure when implemented in the MPC model requires number of rounds that is $O(\text{depth})$ where *depth* is the depth of the tree. The depth may not be $\text{poly}(\log n)$. In order to obtain efficient MPC implementation of these new algorithms, we must find novel solutions to the above two challenges.

Counting in the AMPC model A new (stronger) model of MPC, called the *adaptive* MPC model, was recently introduced by [BDE⁺21]. The AMPC model allows access to a *shared* distributed hash table at the end of every round; additionally, the algorithms are allowed *adaptive* access to this hash table. Such a model has shown to be very practical and have led to improvements in the number of rounds over previous MPC algorithms. Such a model seems to be quite relevant to our work since one of the main challenges in our approximation algorithms is to find the set of edges to give to each machine. (Such a challenge may no longer exist given a shared-memory distributed hash table.) We leave as an interesting open question to obtain better, more round efficient approximate triangle counting algorithms in the AMPC model.

Triangle Counting in $O(1)$ Rounds in Sparse Graphs For sparse graphs where $m = \tilde{O}(n)$, our approximation algorithm requires $\tilde{\Omega}(n)$ space per machine which means that (almost) the entire graph can fit on one machine. This naturally leads to an interesting open question for whether we can obtain an approximate or exact triangle counting algorithm in $O(1)$ rounds in sparse graphs while using *sublinear* space per machine (n^δ space for any constant $\delta > 0$).

References

- [ABB⁺19] Sepehr Assadi, MohammadHossein Bateni, Aaron Bernstein, Vahab Mirrokni, and Cliff Stein. Coresets meet EDCS: algorithms for matching and vertex cover on massive graphs. In *Proceedings 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2019.
- [ABG⁺18] Maryam Aliakbarpour, Amartya Shankha Biswas, Themis Gouleakis, John Peebles, Ronitt Rubinfeld, and Anak Yodpinyanee. Sublinear-time algorithms for counting star subgraphs via edge sampling. *Algorithmica*, 80(2):668–697, 2018.
- [AFU13] Foto N Afrati, Dimitris Fotakis, and Jeffrey D Ullman. Enumerating subgraph instances using map-reduce. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, pages 62–73. IEEE, 2013.
- [AG09] Noga Alon and Shai Gutner. Linear time algorithms for finding a dominating set of fixed size in degenerated graphs. *Algorithmica*, 54(4):544–556, 2009.
- [AG15] Kook Jin Ahn and Sudipto Guha. Access to data and number of iterations: Dual primal algorithms for maximum matching under resource constraints. pages 202–211, 2015.
- [AKK19] Sepehr Assadi, Michael Kapralov, and Sanjeev Khanna. A simple sublinear-time algorithm for counting arbitrary subgraphs via edge sampling. In *ITCS*, volume 124 of *LIPICs*, pages 6:1–6:20. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2019.

- [AKM13] Shaikh Arifuzzaman, Maleq Khan, and Madhav Marathe. Patric: A parallel algorithm for counting triangles in massive networks. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 529–538, 2013.
- [ANOY14] Alexandr Andoni, Aleksandar Nikolov, Krzysztof Onak, and Grigory Yaroslavtsev. Parallel algorithms for geometric graph problems. pages 574–583, 2014.
- [ANR⁺17] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, Nick G. Duffield, and Theodore L. Willke. Graphlet decomposition: framework, algorithms, and applications. *Knowl. Inf. Syst.*, 50(3):689–722, 2017.
- [Ass17] Sepehr Assadi. Simple round compression for parallel vertex cover. *arXiv preprint arXiv:1709.04599*, 2017.
- [ASS⁺18] A. Andoni, Z. Song, C. Stein, Z. Wang, and P. Zhong. Parallel graph connectivity in log diameter rounds. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 674–685, Oct 2018.
- [ASW18] Sepehr Assadi, Xiaorui Sun, and Omri Weinstein. Massively parallel algorithms for finding well-connected components in sparse graphs. *arXiv preprint arXiv:1805.02974*, 2018.
- [ASW19] Sepehr Assadi, Xiaorui Sun, and Omri Weinstein. Massively parallel algorithms for finding well-connected components in sparse graphs. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 461–470. ACM, 2019.
- [ASZ19] Alexandr Andoni, Clifford Stein, and Peilin Zhong. Log diameter rounds algorithms for 2-vertex and 2-edge connectivity. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPICs*, pages 14:1–14:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [AYZ97] Noga Alon, Raphael Yuster, and Uri Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997.
- [BBD⁺19] Soheil Behnezhad, Sebastian Brandt, Mahsa Derakhshan, Manuela Fischer, MohammadTaghi Hajiaghayi, Richard M. Karp, and Jara Uitto. Massively parallel computation of matching and MIS in sparse graphs. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 481–490. ACM, 2019.
- [BC17] Suman K Bera and Amit Chakrabarti. Towards tighter space bounds for counting triangles and other substructures in graph streams. In *34th Symposium on Theoretical Aspects of Computer Science*, 2017.
- [BDE⁺19] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Lacki, and Vahab S. Mirrokni. Near-optimal massively parallel graph connectivity. In *60th IEEE Annual*

Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019, pages 1615–1636, 2019.

- [BDE⁺21] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Lacki, Vahab Mirrokni, and Warren Schudy. Massively parallel computation via remote memory access. *ACM Transactions on Parallel Computing*, 8(3):1–25, 2021.
- [BDH18] Soheil Behnezhad, Mahsa Derakhshan, and MohammadTaghi Hajiaghayi. Semireduce meets congested clique. *arXiv preprint arXiv:1802.10297*, 2018.
- [BDH⁺19] Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Marina Knittel, and Hamed Saleh. Streaming and massively parallel algorithms for edge coloring. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPICs*, pages 15:1–15:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [BEG⁺18] Mahdi Boroujeni, Soheil Ehsani, Mohammad Ghodsi, MohammadTaghi Hajiaghayi, and Saeed Seddighin. Approximating edit distance in truly subquadratic time: quantum and MapReduce. pages 1170–1189, 2018.
- [BFU18] Sebastian Brandt, Manuela Fischer, and Jara Uitto. Breaking the linear-memory barrier in MPC: Fast MIS on trees with n^ϵ memory per machine. *arXiv preprint arXiv:1802.06748*, 2018.
- [BHH19] Soheil Behnezhad, MohammadTaghi Hajiaghayi, and David G. Harris. Exponentially faster massively parallel maximal matching. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1637–1649. IEEE Computer Society, 2019.
- [BHKK09] Andreas Bjöklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. Counting paths and packings in halves. *Algorithms - ESA 2009*, page 578–586, 2009.
- [BKS13] Paul Beame, Paraschos Koutris, and Dan Suciu. Communication steps for parallel query processing. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, pages 273–284, 2013.
- [BKS14] Paul Beame, Paraschos Koutris, and Dan Suciu. Skew in parallel query processing. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 212–223, 2014.
- [BOV13] Vladimir Braverman, Rafail Ostrovsky, and Dan Vilenchik. How hard is counting triangles in the streaming model? In *International Colloquium on Automata, Languages, and Programming*, pages 244–254. Springer, 2013.
- [BPS20] Suman K. Bera, Noujan Pashanasangi, and C. Seshadhri. Linear Time Subgraph Counting, Graph Degeneracy, and the Chasm at Size Six. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:20, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

- [BPS21] Suman K Bera, Noujan Pashanasangi, and C Seshadhri. Near-linear time homomorphism counting in bounded degeneracy graphs: the barrier of long induced cycles. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2315–2332. SIAM, 2021.
- [Bre19] Marco Bressan. Faster subgraph counting in sparse graphs. In *14th International Symposium on Parameterized and Exact Computation (IPEC 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [BYKS02] 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*, pages 623–632. Society for Industrial and Applied Mathematics, 2002.
- [CC11] Shumo Chu and James Cheng. Triangle listing in massive networks and its applications. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 672–680, 2011.
- [CFG⁺19] Yi-Jun Chang, Manuela Fischer, Mohsen Ghaffari, Jara Uitto, and Yufan Zheng. The complexity of $(\Delta+1)$ coloring in congested clique, massively parallel computation, and centralized local computation. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 471–480. ACM, 2019.
- [CHKK⁺19] Keren Censor-Hillel, Petteri Kaski, Janne H Korhonen, Christoph Lenzen, Ami Paz, and Jukka Suomela. Algebraic methods in the congested clique. *Distributed Computing*, 32(6):461–478, 2019.
- [CLM⁺20] Artur Czumaj, Jakub Lacki, Aleksander Madry, Slobodan Mitrovic, Krzysztof Onak, and Piotr Sankowski. Round compression for parallel matching algorithms. *SIAM J. Comput.*, 49(5), 2020.
- [CN85] Norishige Chiba and Takao Nishizeki. Arboricity and subgraph listing algorithms. *SIAM Journal on computing*, 14(1):210–223, 1985.
- [Coh09] Jonathan Cohen. Graph twiddling in a mapreduce world. *Computing in Science & Engineering*, 11(4):29–41, 2009.
- [DAH17] V. S. Dave, N. K. Ahmed, and M. Hasan. PE-CLoG: Counting edge-centric local graphlets. In *IEEE International Conference on Big Data*, pages 586–595, 2017.
- [DBS18] Maximilien Danisch, Oana Balalau, and Mauro Sozio. Listing k-cliques in sparse real-world graphs*. In *Proceedings of the 2018 World Wide Web Conference, WWW ’18*, page 589–598, Republic and Canton of Geneva, CHE, 2018. International World Wide Web Conferences Steering Committee.
- [DG08] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- [DLP12] Danny Dolev, Christoph Lenzen, and Shir Peled. “Tri, Tri again”: Finding triangles and small subgraphs in a distributed setting. *Distributed Computing*, page 195–209, 2012.

- [DLSY21] Laxman Dhulipala, Quanquan C. Liu, Julian Shun, and Shangdi Yu. Parallel batch-dynamic k -clique counting. In Michael Schapira, editor, *2nd Symposium on Algorithmic Principles of Computer Systems, APOCS 2020, Virtual Conference, January 13, 2021*, pages 129–143. SIAM, 2021.
- [EG04] Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed parameter clique and dominating set. *Theor. Comput. Sci.*, 326(1–3):57–67, October 2004.
- [ELRS17] Talya Eden, Amit Levi, Dana Ron, and C Seshadhri. Approximately counting triangles in sublinear time. *SIAM Journal on Computing*, 46(5):1603–1646, 2017.
- [ELS13] David Eppstein, Maarten Löffler, and Darren Strash. Listing all maximal cliques in large sparse real-world graphs. *ACM Journal of Experimental Algorithms*, 18(3):364–375, 2013.
- [ERS20] Talya Eden, Dana Ron, and C. Seshadhri. Faster sublinear approximation of the number of k -cliques in low-arboricity graphs. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1467–1478, 2020.
- [ESBD16] Ethan R. Elenberg, Karthikeyan Shanmugam, Michael Borokhovich, and Alexandros G. Dimakis. Distributed estimation of graph 4-profiles. In *International Conference on World Wide Web (WWW)*, pages 483–493, 2016.
- [exp20] <https://github.com/qqliu/mpc-triangle-count-exact-approx-simulations>, 2020.
- [FFF15] Irene Finocchi, Marco Finocchi, and Emanuele G Fusco. Clique counting in mapreduce: Algorithms and experiments. *Journal of Experimental Algorithmics (JEA)*, 20:1–20, 2015.
- [GG06] Gaurav Goel and Jens Gustedt. Bounded arboricity to determine the local structure of sparse graphs. In *Graph-Theoretic Concepts in Computer Science, 32nd International Workshop, WG 2006, Bergen, Norway, June 22-24, 2006, Revised Papers*, pages 159–167, 2006.
- [GGK⁺18] Mohsen Ghaffari, Themis Gouleakis, Christian Konrad, Slobodan Mitrović, and Ronitt Rubinfeld. Improved massively parallel computation algorithms for mis, matching, and vertex cover. arXiv:1802.08237, 2018.
- [GKMS19] Buddhima Gamlath, Sagar Kale, Slobodan Mitrovic, and Ola Svensson. Weighted matchings via unweighted augmentations. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 491–500. ACM, 2019.
- [GKU19] Mohsen Ghaffari, Fabian Kuhn, and Jara Uitto. Conditional hardness results for massively parallel computation from distributed lower bounds. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1650–1663. IEEE Computer Society, 2019.

- [GLM19] Mohsen Ghaffari, Silvio Lattanzi, and Slobodan Mitrović. Improved parallel algorithms for density-based network clustering. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2201–2210, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- [GNT20] Mohsen Ghaffari, Krzysztof Nowicki, and Mikkel Thorup. Faster algorithms for edge connectivity via random 2-out contractions. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1260–1279. SIAM, 2020.
- [Gra] GraphChallenge. <http://graphchallenge.mit.edu/>.
- [GSZ11] Michael T. Goodrich, Nodari Sitchinava, and Qin Zhang. Sorting, searching, and simulation in the mapreduce framework. In *Proceedings of the 22Nd International Conference on Algorithms and Computation, ISAAC’11*, pages 374–383, Berlin, Heidelberg, 2011. Springer-Verlag.
- [GU19] Mohsen Ghaffari and Jara Uitto. Sparsifying distributed algorithms with ramifications in massively parallel computation and centralized local computation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1636–1653. SIAM, 2019.
- [HD14] Tomaz Hocevar and Janez Demsar. A combinatorial approach to graphlet counting. *Bioinformatics*, pages 559–65, 2014.
- [HLL18] Nicholas J. A. Harvey, Christopher Liaw, and Paul Liu. Greedy and local ratio algorithms in the MapReduce model. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 43–52, New York, NY, USA, 2018. ACM.
- [HP15] James W Hegeman and Sriram V Pemmaraju. Lessons from the congested clique applied to MapReduce. *Theoretical Computer Science*, 608:268–281, 2015.
- [IBY⁺07] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. Dryad: distributed data-parallel programs from sequential building blocks. In *ACM SIGOPS operating systems review*, volume 41, pages 59–72. ACM, 2007.
- [ILMP19] Giuseppe F. Italiano, Silvio Lattanzi, Vahab S. Mirrokni, and Nikos Parotsidis. Dynamic algorithms for the massively parallel computation model. In Christian Scheideler and Petra Berenbrink, editors, *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019, Phoenix, AZ, USA, June 22-24, 2019*, pages 49–58. ACM, 2019.
- [IMS17] Sungjin Im, Benjamin Moseley, and Xiaorui Sun. Efficient massively parallel methods for dynamic programming. pages 798–811, 2017.
- [JS17] Shweta Jain and C Seshadhri. A fast and provable method for estimating clique counts using turán’s theorem. In *Proceedings of the 26th International Conference on World Wide Web*, pages 441–449, 2017.

- [JSP13] Madhav Jha, Comandur Seshadhri, and Ali Pinar. A space efficient streaming algorithm for triangle counting using the birthday paradox. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 589–597, 2013.
- [KMPT12] Mihail N Kolountzakis, Gary L Miller, Richard Peng, and Charalampos E Tsourakakis. Efficient triangle counting in large graphs via degree-based vertex partitioning. *Internet Mathematics*, 8(1-2):161–185, 2012.
- [KMSS12] Daniel M Kane, Kurt Mehlhorn, Thomas Sauerwald, and He Sun. Counting arbitrary subgraphs in data streams. In *International Colloquium on Automata, Languages, and Programming*, pages 598–609. Springer, 2012.
- [KPP⁺14] Tamara G Kolda, Ali Pinar, Todd Plantenga, C Seshadhri, and Christine Task. Counting triangles in massive graphs with mapreduce. *SIAM Journal on Scientific Computing*, 36(5):S48–S77, 2014.
- [KPP16] Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Higher lower bounds from the 3sum conjecture. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1272–1287. SIAM, 2016.
- [KSV10] Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for MapReduce. pages 938–948, 2010.
- [Lat08] Matthieu Latapy. Main-memory triangle computations for very large (sparse (power-law)) graphs. *Theoretical computer science*, 407(1-3):458–473, 2008.
- [LK14] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [LMOS19] Jakub Lacki, Slobodan Mitrovic, Krzysztof Onak, and Piotr Sankowski. Walking randomly, massively, and efficiently. *CoRR*, abs/1907.05391, 2019.
- [LMSV11] Silvio Lattanzi, Benjamin Moseley, Siddharth Suri, and Sergei Vassilvitskii. Filtering: a method for solving graph problems in MapReduce. pages 85–94, 2011.
- [LPSP05] Zvi Lotker, Boaz Patt-Shamir, Elan Pavlov, and David Peleg. Minimum-weight spanning tree construction in $o(\log \log n)$ communication rounds. *SIAM Journal on Computing*, 35(1):120–131, 2005.
- [LQLC15] Longbin Lai, Lu Qin, Xuemin Lin, and Lijun Chang. Scalable subgraph enumeration in mapreduce. *Proceedings of the VLDB Endowment*, 8(10):974–985, 2015.
- [MNV16] 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*, pages 401–411, 2016.
- [Ona18] Krzysztof Onak. Round compression for parallel graph algorithms in strongly sublinear space. *CoRR*, abs/1807.08745, 2018.
- [Pat10] Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 603–610, 2010.

- [PC13] Ha-Myung Park and Chin-Wan Chung. An efficient mapreduce algorithm for counting triangles in a very large graph. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 539–548, 2013.
- [PSKP14] Ha-Myung Park, Francesco Silvestri, U Kang, and Rasmus Pagh. Mapreduce triangle enumeration with guarantees. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 1739–1748, 2014.
- [PSV17] Ali Pinar, C. Seshadhri, and Vaidyanathan Vishal. ESCAPE: Efficiently counting all 5-vertex subgraphs. In *International Conference on World Wide Web (WWW)*, pages 1431–1440, 2017.
- [PT12] Rasmus Pagh and Charalampos E Tsourakakis. Colorful triangle counting and a mapreduce implementation. *Information Processing Letters*, 112(7):277–281, 2012.
- [RVW16] Tim Roughgarden, Sergei Vassilvitskii, and Joshua R. Wang. Shuffles and circuits: (on lower bounds for modern parallel computation). pages 1–12, 2016.
- [SDS20] Jessica Shi, Laxman Dhulipala, and Julian Shun. Parallel clique counting and peeling algorithms. *CoRR*, abs/2002.10047, 2020.
- [SERF18] Kijung Shin, Tina Eliassi-Rad, and Christos Faloutsos. Patterns and anomalies in k-cores of real-world graphs with applications. *Knowledge and Information Systems*, 54(3):677–710, 2018.
- [SPK13] Comandur Seshadhri, Ali Pinar, and Tamara G Kolda. Triadic measures on graphs: The power of wedge sampling. In *Proceedings of the 2013 SIAM International Conference on Data Mining*, pages 10–18. SIAM, 2013.
- [SSS95] Jeanette P Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff–hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995.
- [ST15] J. Shun and K. Tangwongsan. Multicore triangle computations without tuning. In *2015 IEEE 31st International Conference on Data Engineering*, pages 149–160, April 2015.
- [SV11] Siddharth Suri and Sergei Vassilvitskii. Counting triangles and the curse of the last reducer. In *Proceedings of the 20th international conference on World wide web*, pages 607–614, 2011.
- [SW05] Thomas Schank and Dorothea Wagner. Finding, counting and listing all triangles in large graphs, an experimental study. In *Experimental and Efficient Algorithms*, pages 606–609. 2005.
- [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, 2009.
- [Vas09] Virginia Vassilevska. Efficient algorithms for clique problems. *Information Processing Letters*, 109(4):254–257, 2009.

- [WC81] Mark N Wegman and J Lawrence Carter. New hash functions and their use in authentication and set equality. *Journal of computer and system sciences*, 22(3):265–279, 1981.
- [Whi12] Tom White. *Hadoop: The definitive guide*. ” O’Reilly Media, Inc.”, 2012.
- [YK11] Jin-Hyun Yoon and Sung-Ryul Kim. Improved sampling for triangle counting with mapreduce. In *International Conference on Hybrid Information Technology*, pages 685–689. Springer, 2011.
- [ZCF⁺10] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: Cluster computing with working sets. *HotCloud*, 10(10-10):95, 2010.

A Preliminaries

A.1 Proof of [Theorem 2.1](#)

Using the construction of the interval tree defined in [\[GSZ11\]](#) that has branching factor $d = \mathcal{M}/2$ we perform the following to count the number of times each element repeats in our sorted list of N elements. To initialize the tree, each leaf of the tree contains exactly one of the elements in the sorted list of elements where leaf v_i contains element x_i of the list. Let the height of the tree be L , the leaves of the tree be at level $L - 1$ and the root be at level 0. Then, the rest of the algorithm proceeds in two phases:

1. **Bottom-up phase:** For each level $\ell = L - 1$ up to 0:
 - (a) For each node v on level ℓ :
 - i. If v is a leaf, it sends its value x_i to its parent $p(v)$.
 - ii. If v is a vertex in level $L - 2$, let $(x_i, x_{i+1}, \dots, x_{i+j})$ where $j < d$ be values obtained from its leaf children from left to right. Let $c(x_i)$ be the count of element x_i among the values obtained from the children of v . The counts are computed locally on the machine storing v . Then, v sends $x_i, c(x_i), x_{i+j}, c(x_{i+j})$ to its parent $p(v)$.
 - iii. If v is a non-leaf node on level $\ell < L - 2$, let $x_a, c(a), x_b, c(b), \dots$ be the values of elements obtained from its children and their counts. v updates the counts of all elements received. For example, if $x_a = x_b$, v updates $c(a)$ and $c(b)$ to be $c(a) + c(b)$. Let x_{left} be the first element received from v ’s leftmost leaf and x_{right} be the second element received from v ’s rightmost leaf. Then, send these elements and their updated counts, $x_{left}, c(x_{left}), x_{right}$, and $c(x_{right})$, to its parent $p(v)$.
2. **Top-down phase:** For each level $\ell = 0$ down to $\ell = L - 1$:
 - (a) For each node v at level ℓ :
 - i. If v is the root, then it computed and stored in its memory new repeating counts for the values it received from its children: $x_a, c(x_a), x_b, c(x_b), \dots$. It sends the new counts and values to its respective child that sent it the value originally (e.g. $x_{left}, c(x_{left})$ to v_{left}). Intuitively, this updates the child’s count of values with values that are not in its subtree.
 - ii. If v is not the root and is a non-leaf node, it receives the values from its parents for its leftmost and rightmost child counts. Given the set of values it stored from its

children it updates the counts with counts of values received from its parents. This allows for the counts to reflect values not in its subtree. Then, it sends the updated counts to its children.

- iii. If v is a leaf, it receives values $x_i, c(x_i)$ from its parent. $c(x_i)$ is then the number of times x_i occurs in the sorted list.

The above procedure uses $O(d)$ space per processor and $O(L)$ rounds of communication. Since $L = O(\log_d(N))$ and $d = \mathcal{M}/2$, the number of rounds of communication that is necessary is $O(\log_{\mathcal{M}} N)$.

A.2 Proof of Lemma 2.2

We first create the following tuples in parallel to represent tuples in Q and C , respectively. For each tuple $q \in Q$, we create the tuple $(q, 1)$. For each tuple $c \in C$, we create the tuple $(c, 0)$. Let F denote the set of tuples $(c, 0)$ and $(q, 1)$. First, we sort the tuples in F lexicographically (where 0 comes before 1) [GSZ11]. Then, we use the predecessor primitive given in (e.g. [GSZ11, ASS⁺18], Appendix A of [BDE⁺19]) to determine the queries $q \in Q$ that are in C . Given the sorted F , we use the predecessor algorithm of [BDE⁺19] to determine for each $(q, 1)$ tuple, the first tuple that appears before it that has value 0. Suppose this tuple is $(c, 0)$. Then, if $q = c$, then the queried tuple q is in C . For all tuples $q \in Q$, we can then return in parallel whether $q \in C$ also. Both the sorting and the predecessor queries take $O(|Q \cup C|)$ total space and $O_\delta(1)$ rounds.

A.3 Proof of Lemma 2.3

Let M be some machine with n^δ information and $O(n^{2\delta})$ space. We create the x duplicates by repeatedly duplicating each machine M_j^i to n^δ machines $M_{n^\delta \cdot j}^{i+1}, \dots, M_{n^\delta \cdot j + n^\delta - 1}^{i+1}$, starting with $M_0^0 = M$. Therefore, after $\ell = \log_{n^\delta} x$ rounds this process terminates, and the required duplicates is the set of machines M_1^ℓ to M_x^ℓ .