The Power of Pivoting for Exact Clique Counting

Shweta Jain

University of California, Santa Cruz
Santa Cruz, CA, USA
sjain12@ucsc.edu

C. Seshadhri

University of California, Santa Cruz
Santa Cruz, CA, USA
sesh@ucsc.edu

ABSTRACT

Clique counting is a fundamental task in network analysis, and even the simplest setting of 3-cliques (triangles) has been the center of much recent research. Getting the count of k -cliques for larger k is algorithmically challenging, due to the exponential blowup in the search space of large cliques. But a number of recent applications (especially for community detection or clustering) use larger clique counts. Moreover, one often desires *local* counts, the number of k -cliques per vertex/edge.

Our main result is PIVOTER, an algorithm that exactly counts the number of k -cliques, for all values of k . It is surprisingly effective in practice, and is able to get clique counts of graphs that were beyond the reach of previous work. For example, PIVOTER gets all clique counts in a social network with a 100M edges within two hours on a commodity machine. Previous parallel algorithms do not terminate in days. PIVOTER can also feasibly get local per-vertex and per-edge k -clique counts (for all k) for many public data sets with tens of millions of edges. To the best of our knowledge, this is the first algorithm that achieves such results.

The main insight is the construction of a Succinct Clique Tree (SCT) that stores a compressed unique representation of all cliques in an input graph. It is built using a technique called *pivoting*, a classic approach by Bron-Kerbosch to reduce the recursion tree of backtracking algorithms for maximal cliques. Remarkably, the SCT can be built without actually enumerating all cliques, and provides a succinct data structure from which exact clique statistics (k -clique counts, local counts) can be read off efficiently.

KEYWORDS

Social network analysis; clique counting; local clique counting

ACM Reference Format:

Shweta Jain and C. Seshadhri. 2020. The Power of Pivoting for Exact Clique Counting. In *The Thirteenth ACM International Conference on Web Search and Data Mining (WSDM '20)*, February 3–7, 2020, Houston, TX, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3336191.3371839>

1 INTRODUCTION

Subgraph counting (also known as motif counting, graphlet counting) is a fundamental algorithmic problem in network analysis, widely applied in domains such as social network analysis,

bioinformatics, cybersecurity, and physics (refer to tutorial [26] and references within). One of the most important cases is that of *clique counting*. A k -clique is a complete subgraph on k vertices, and has great significance in network analysis (Chap. 11 of [14] and Chap. 2 of [15]). Indeed, just the special case of $k = 3$ (triangle counting) has a rich history in modern network science. General clique counting has received much attention in recent times [1, 10, 13, 16, 17, 19, 21]. There is a line of recent work on exploiting clique counts for community detection and dense subgraph discovery [4, 18, 24, 30, 31, 36].

Despite much effort on this problem, it has been challenging to get scalable algorithms for clique counting. There is a large literature for counting 3-cliques (triangles) and some of these methods have been extended to counting cliques upto size 5 [1, 17, 19, 21]. However, practical algorithms for counting cliques beyond size 5 have proven to be much harder, and the reason for this is combinatorial explosion. Essentially, as k increases, the number of k -cliques blows up. For large graphs, some recent practical algorithms have succeeded in counting up to (around) 10-cliques [10, 13, 16]. They either use randomized approximation or parallelism to speed up their counting. Besides the obvious problem that they do not scale for larger k , it is difficult to obtain more refined clique counts (such as counts for every vertex or every edge).

1.1 Problem Statement

We are given an undirected, simple graph $G(V, E)$. For $k \geq 3$, a k -clique is a set of k vertices that induce a complete subgraph (it contains all edges among the k vertices). We will denote the number of k -cliques as C_k . For a vertex $v \in V$, we use $c_k(v)$ to denote the number of k -cliques that v participates in. Analogously, we define $c_k(e)$ for edge $e \in E$.

We focus on the following problems, in increasing order of difficulty. We stress that k is *not* part of the input, and we want results for all values of k .

- Global clique counts: Output, $\forall k \geq 3, C_k$.
- Per-vertex clique counts: Output, $\forall k, \forall v \in V$, the value $c_k(v)$.
- Per-edge clique counts: Output, $\forall k, \forall e \in E$, the value $c_k(e)$.

The per-vertex and per-edge counts are sometimes called *local counts*. In clustering applications, the local counts are used as vertex or edge weights, and are therefore even more useful than global counts [4, 18, 24, 30, 31, 36].

Challenges: Even the simplest problem of getting global clique counts subsumes a number of recent results on clique counting [10, 13, 16]. The main challenge is combinatorial explosion: for example, the web-Stanford web graph with 2M edges has *3000 trillion* 15-cliques. These numbers are even more astronomical for

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

WSDM '20, February 3–7, 2020, Houston, TX, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6822-3/20/02...\$15.00

<https://doi.org/10.1145/3336191.3371839>

Figure 1: Fig. 1i shows the comparison of time taken (in seconds) by PIVOTER for 4 of our largest graphs to count *all* k -cliques with the time taken by kClist40 (the parallel version of the state of the art algorithm kClist that uses 40 threads) to count the number of k -cliques, where k is the maximum clique size in each graph. For Stanford, BerkStan, as-skitter, orkut, the maximum clique sizes were 61, 201, 67 and 51 resp. PIVOTER terminated for most graphs in minutes, (except for orkut, for which it took about 2 hours) whereas kClist40 had not terminated even after 3 days, giving a speedup of 100x to 10000x. Fig. 1i also shows the time taken by PIVOTER to obtain the per-vertex and per-edge k -clique counts. They were within a factor of the time taken to obtain global k -clique counts. Fig. 1ii and Fig. 1iii shows the frequency distribution of k -cliques i.e. for every number r on the x-axis, the y-axis shows the number of vertices that participate in r k -cliques, for $k \in [5, 10, 15, 20, 25]$ for as-skitter and web-Stanford graphs.

larger graphs. Any method that tries to enumerate is doomed to failure.

Amazingly, recent work by Danisch-Balalau-Sozio uses parallel algorithms to count beyond trillions of cliques. But even their algorithm fails to get all global clique counts for a number of datasets. Randomized methods have been used with some success, but even they cannot estimate all clique counts [13, 16].

Local counting, for all k , is even harder, especially given the sheer size of the output. Parallel methods would eventually need to store local counts for every subproblem, which would increase the overall memory footprint. For local counts, sampling would require far too many random variables, each of which need to be sampled many times for convergence. (We give more explanation in §1.3.)

This raises the main question:

Is there a scalable, exact algorithm for getting all global and local cliques counts, on real-world graphs with millions of edges?

To the best of our knowledge, there is no previous algorithm that can solve these problems on even moderate-sized graphs with a few million edges.

1.2 Main contributions

Our main contribution is a new practical algorithm PIVOTER for the global and local clique counting problems.

Exact counting without enumeration: Current methods for exact clique counting perform an *enumeration*, in that the algorithm explicitly “visits” every clique. Thus, this method cannot scale to counting larger cliques, since the number of cliques is simply too large. Our main insight is that the method of *pivoting*, used to reduce recursion trees for maximal clique enumeration [7, 12], can be applied to counting cliques of all sizes.

Succinct Clique Trees through Pivoting: We prove that pivoting can be used to construct a special data structure called the *Succinct Clique Tree* (SCT). The SCT stores a unique representation of all cliques, but is much smaller than the total number of cliques. It can also be built quite efficiently. Additionally, given the tree, one can easily “read off” the number of k -cliques and various local counts in the graph. Remarkably, we can get all counts without storing the entire tree and the storage required at any point is linear in the number of edges.

Excellent practical performance: We implement PIVOTER on a commodity machine. For global clique counting, PIVOTER is able to process graphs of up to tens of millions of edges in *minutes*. Previous results either work only for small values of k (typically up to 10) or take much longer. Consider Fig. 1i, where the time of PIVOTER is compared with that of kClist (the state of the art parallel algorithm for clique counting) [10]. In the instances shown kClist did not terminate even after running for 3 days. By contrast, for the largest com-orkut social network with more than 100M edges, PIVOTER gets all values of C_k within two hours. (Typically, in this time, kClist gets k -clique counts only up to $k \leq 13$.)

Feasible computation of local counts: PIVOTER is quite efficient for per-vertex counts, and runs in at most twice the time for global counts. The times for local clique counting are given in Fig. 1i. Even for the extremely challenging problem of per-edge counts, in most instances PIVOTER gets these numbers in a few hours. (For the com-orkut social network though, it takes a few days.)

This allows us to get data shown in Fig. 1ii and Fig. 1iii, that plots the frequency distribution of k -cliques. (In other words, for every number r , we plot the number of vertices that participate in r k -cliques.) As mentioned earlier, this information is used for

dense subgraph discovery [24, 30]. To the best of our knowledge, this is the first algorithm that is able to get such information for real-world graphs.

1.3 Related Work

Subgraph counting has an immensely rich history in network science, ranging from applications across social network analysis, bioinformatics, recommendation systems, graph clustering (we refer the reader to the tutorial [26] and references within). We only describe work directly relevant to clique counting.

The simplest case of clique counting is *triangle counting*, which has received much attention from the data mining and algorithms communities. Recent work has shown the relevance of counts of large subgraphs (4, 5 vertex patterns) [5, 23, 27, 32, 35]. Local clique counts have played a significant role in a flurry of work on faster and better algorithms for dense subgraph discovery and community detection [4, 24, 30, 31]. The latter results define the “motif conductance”, where cuts are measured by the number of subgraphs (not just edges) cut. This has been related to higher order clustering coefficients [35, 36]. These quantities are computed using local clique counts, underscoring the importance of these numbers.

The problem of counting cliques (and variants such as counting maximal cliques) has received much attention both from the applied and theoretical computer science communities [3, 8, 9, 33]. Classic techniques like color-coding [6, 37] and path sampling [17, 25, 34] have been employed for counting cliques up to size 5.

For larger cliques, Finocchi-Finocchi-Fusco gave a MapReduce algorithm that uses orientation and sampling techniques [13]. Jain and Seshadhri use methods from extremal combinatorics to give a fast sampling algorithm [16], that is arguably the fastest approximate clique counter to date. In a remarkable result, Danisch-Balalau-Sozio gave a parallel implementation (kClist) of a classic algorithm of Chiba-Nishizeki, which is able to enumerate upto trillions of cliques [10]. For exact counting, we consider kClist as the state of the art. Despite the collection of clever techniques, none of these methods really scale beyond counting (say) 10-cliques for large graphs.

Why local counting is hard: Note that either parallelism or sampling is used to tame the combinatorial explosion. Even though (at least for small k), one can enumerate all cliques in parallel, local counting requires updating a potentially global data structure, the list of all $c_k(v)$ or $c_k(e)$ values. To get the benefits of parallelism, one would either have to duplicate a large data structure or combine results for various threads to get all local counts. While this may be feasible, it adds an extra memory overhead.

Sampling methods typically require some overhead for convergence. For local counts, there are simply too many samples required to get accurate values for (say) all $c_k(v)$ values. For these reasons, we strongly believe that new ideas were required to get efficient local counting.

Maximal clique enumeration: Extremely relevant to our approach is a line of work of maximal clique enumeration. A *maximal clique* is one that is not contained in a larger clique. Unlike the combinatorial explosion of k -cliques, maximal cliques tend to be much fewer. The first algorithm for this problem is the classic Bron-Kerbosch backtracking procedure from the 70s [2, 7].

They also introduced an idea called *pivoting*, that prunes the recursion tree for efficiency. Tomita-Tanaka-Takahashi gave the first theoretical analysis of pivoting rules, and showed asymptotic improvements [29]. Eppstein-Löffler-Strash combined these ideas with orientation methods to give a practical and provably fast algorithm for maximal clique enumeration [11, 12]. An important empirical observation of this line of work is that the underlying recursion tree created with pivoting is typically small for real-world graphs. This is the starting point for our work.

2 MAIN IDEAS

Inspired by the success of maximal clique enumeration through pivoting, we design the Succinct Clique Tree (SCT) of a graph for clique counting.

To explain the SCT, it is useful to begin with the simple backtracking algorithm for listing all cliques. For any vertex v , let $N(v)$ denote the neighborhood of v . Any clique containing v is formed by adding v to a clique contained in $N(v)$. Thus, we can find all cliques by this simple recursive procedure: for all v , recursively enumerate all cliques in $N(v)$. For each such clique, add v to get a new clique. It is convenient to think of the recursion tree of this algorithm. Every node of the tree (corresponding to a recursive call) corresponds to a subset $S \subseteq V$, and the subtree of calls enumerates all cliques contained in S . A call to S makes a recursive call corresponding to every $s \in S$, which is over the set $N(s) \cap S$ (the neighbors of s in S). We can label every edge of the tree (call them *links* to distinguish from edges of G) with a vertex, whose neighborhood leads to the next recursive call. It is not hard to see that the link labels, along any path from a root (that might not end at a leaf), give a clique. Moreover, every clique has such a representation.

Indeed, every permutation of clique forms such a path. A simple and classic method to eliminate multiple productions of a clique is *acyclic orientations*. Simply orient the graph as a DAG, and only make recursive calls on out-neighborhoods. Typically, an orientation is chosen by degeneracy/core decomposition or degree orderings, so that out-neighborhood sizes are minimized. This is a central technique in all recent applied algorithms on clique counting [10, 13, 16]. Yet it is not feasible to construct the recursion tree to completion, and it is typically truncated at some depth (≤ 10) for large graphs.

Is it possible to somehow “compress” the tree, and get a unique (easily accessible) representation of all cliques?

The power of pivoting: We discover a suprising answer, in pivoting. This was discovered by Bron-Kerbosch in the context of *maximal* cliques [7]. We describe, at an intuitive level, how it can be applied for global and local clique counting. For the recursive call at S , first pick a *pivot* vertex $p \in S$. Observe that the cliques in S can be partitioned into three classes as follows. For clique C contained in S : (i) $p \in C$, (ii) $C \subset N(p)$, (iii) C contains a non-neighbor of p . There is 1-1 correspondence between cliques of type (i) and (ii), so we could hope to only enumerate type (ii) cliques.

Thus, from a recursive call for S , we make recursive calls to find cliques in $N(p) \cap S$, and $N(u) \cap S$ for every *non-neighbor* u of p in S . We avoid making recursive calls corresponding to vertices in $N(p)$. This gives the main savings over the simple backtracking procedure.

The natural choice of p is the highest degree vertex in the graph induced on S . The recursion tree obtained is essentially the SCT. We stress that this is quite different from the Bron-Kerbosch recursion tree. The BK algorithm also maintains a set of excluded vertices since it only cares for maximal cliques. This excluded set is used to prune away branches that cannot be maximal; moreover, the pivots in BK are potentially chosen from outside S to increase pruning. The SCT is constructed in this specific manner to ensure unique clique representations, which the BK tree does not provide.

The SCT is significantly smaller than recursion trees that use degeneracy orientations (which one cannot feasibly construct). In practice, it can be constructed efficiently for graphs with tens of millions of edges. As before the nodes of the SCT are labeled with subsets (corresponding to the recursive calls), and links are labeled with vertices (corresponding to the vertex whose neighborhood is being processed). Abusing notation, in the following discussion, we refer to a path by the set of link labels in the path.

How can we count all cliques using the SCT? Every root to leaf path in the tree corresponds to a clique, but not all cliques correspond to paths. This is distinct from the standard recursion tree discussed earlier, where every clique corresponds to a path from the root. Indeed, this is why the standard recursion trees (even with degeneracy orientations) are large.

We prove the following remarkable “unique encoding” property. Within any root to leaf path T , there is a subset of links P corresponding to the pivot calls. Every clique C in the graph can be uniquely expressed as $(T \setminus P) \cup Q$ for some $Q \subseteq P$ (for a specific path T). The uniqueness is critical for global and local counting, since we can simply write down formulas to extract all counts. Thus, the SCT gives a unique encoding for every clique in the graph.

Intuitively, the source of compression can be seen in two different ways. The simplest way is to see that pivoting prunes the tree, because recursive calls are only made for a subset of vertices. But also, not every clique is represented by (the link labels of) a path from the root. Thus, there are far fewer paths in the SCT. The final algorithm is quite simple and the main work was coming up with the above insight. Despite this simplicity, it outperforms even parallel methods for exact clique counting by orders of magnitude.

Our main theorem follows. Basically, clique counts can be obtained in time proportional to the size of the SCT. All the technical terms will be formally defined in §3.

THEOREM 2.1. *Let G be an input graph with n vertices, m edges, and degeneracy α . Let $SCT(G)$ be the Succinct Clique Tree of graph G .*

The procedure $PIVOTER(G)$ correctly outputs all global and local counts. For global and per-vertex counts, the running time is $O(\alpha^2 |SCT(G)| + m + n)$. For per-edge counts, the running time is $O(\alpha^3 |SCT(G)| + m + n)$. The storage cost is $O(m + n)$.

Empirically, we observe that the SCT is quite small. In the worst-case, $|SCT(G)| = O(n^{3^{\alpha/3}})$, which follows from arguments by Eppstein-Löffler-Strash [12] and Tomita-Tanaka-Takahashi [29] (an exponential dependence is necessary because of the NP-hardness of maximum clique). Details on this are in the full paper.

3 PRELIMINARIES

We start with the mathematical formalism required to describe the main algorithm and associated proofs. The input is a simple, undirected graph $G = (V, E)$, where $|V| = n$ and $|E| = m$. It is convenient to assume that G is connected. We use vertices to denote the elements of V (the term *nodes* will be used for a different construct). We use the following notation for neighborhoods.

- $N(v)$: This is the neighborhood of v .
- $N(S, v)$: For any subset of vertices S , we use $N(S, v)$ to denote $N(v) \cap S$. Alternately, this is the neighborhood of v in S .

We will use *degeneracy orderings* (or core decompositions) to reduce the recursion tree. This is a standard technique for clique counting [9, 10, 13, 16]. This ordering is obtained by iteratively removing the minimum degree vertex, and can be computed in linear time [20]. Typically, one uses this ordering to convert G into a DAG. The largest *out-degree* is the graph degeneracy, denoted α . We state this fact as a lemma, which is considered a classic fact in graph theory and network science.

LEMMA 3.1. [20] *Given a graph $G = (V, E)$, there is a linear time algorithm that constructs an acyclic orientation of G such that all outdegrees are at most α .*

The most important construct we design is the *Succinct Clique Tree* (SCT) T . The SCT stores special node and link attributes that are key to getting global and local clique counts, for all values of k . The construction and properties of the SCT are given in the next section. Here, we list out technical notation associated with the SCT T .

Formally, T is a tree where nodes are labeled with subsets of V , with the following properties.

- The root is labeled V .
- Parent labels are strict supersets of child labels.
- Leaves are labeled with the empty set $\emptyset$.

An important aspect of T are *link labels*. A link label is a pair with a vertex of V and a “call type”. The label is of the form (v, p) or (v, h) , where p is shorthand for “pivot” and h for “hold”. For a link label $(v, \cdot)$ of the link (S, S') (where $S \supset S'$ is the parent), v will be an element of S .

Consider a root to leaf path T of T . We have the following associated set of vertices. It is convenient to think of T as a set of tree links.

- $H(T)$: This is the set of vertices associated with “hold” call types, among the links of T . Formally, $H(T)$ is $\{v | (v, h) \text{ is label of link in } T\}$.
- $P(T)$: This is the set of vertices with “pivot” calls. Formally $P(T)$ is $\{v | (v, p) \text{ is label of link in } T\}$.

We now describe our algorithm. We stress that the presentation here is different from the implementation. The following presentation is easier for mathematical formalization and proving correctness. The implementation is a recursive version of the same algorithm, which is more space efficient. This is explained in the proof of Theorem 2.1.

4 BUILDING THE SCT

We give the algorithm to construct the SCT. We keep track of various attributes to appropriately label the edges. The algorithm

will construct the SCT T in a breadth-first manner. Every time a node is processed, the algorithm creates its children and labels all the new nodes and links created.

Algorithm 1: SCTBUILDER(G)

 Output: SCT of G

- 1 Find degeneracy orientation of G , and let $N^+(v)$ denote the outneighborhood of a vertex v .
 - 2 Initialize tree T with root labeled V .
 - 3 For every $v \in V$, create a child of root with node label $N^+(v)$. Set the edge label to (v, h) .
 - 4 Insert all these child nodes into a queue Q .
 - 5 While Q is non-empty:
 - 6 Dequeue to get node γ . Let node label be S .
 - 7 If $S = \emptyset$, continue.
 - 8 Find $p \in S$ with largest $N(S, p)$ value.
 - 9 Create child node of γ with vertex label $N(S, p)$. Add this node to T and set the link label (of the new link) to (p, p) . Also, add this node to Q .
 - 10 Let $S \setminus (p \cup N(p)) = \{v_1, v_2, \dots, v_\ell\}$ (listed in arbitrary order).
 - 11 For each $i \leq \ell$: create child node of γ labeled $N(S, v_i) \setminus \{v_1, v_2, \dots, v_{i-1}\}$. Add this node to T and set link label to (v_i, h) . Also add this node to Q .
 - 12 Return T .
-

As mentioned earlier, the child of the node labeled S has one child corresponding to the pivot vertex p , and children for all non-neighbors of p . Importantly, we label each “call” with p or h . This is central to getting unique representations of all the cliques.

Now for our main theorem about SCT.

THEOREM 4.1. *Every clique C (in G) can be uniquely represented as $H(T) \cup Q$, where $Q \subseteq P(T)$ and T is a root to leaf path in T . (Meaning, for any other root to leaf path $T' \neq T$, $\forall Q \subseteq P(T')$, $C \neq H(T') \cup Q$.)*

We emphasize the significance of this theorem. Every root to leaf path T represents a clique, given by the vertex set $H(T) \cup P(T)$. Every clique C is a subset of potentially many such sets; and there is no obvious bound on this number. So one can think of C “occurring” multiple times in the tree T . But Theorem 4.1 asserts that if we take the labels into account ($H(T)$ vs $P(T)$), then there is a *unique* representation or “single occurrence” of C .

PROOF. (of Theorem 4.1) Consider a node γ of T labeled S . We prove, by induction on $|S|$, that every clique $C \subseteq S$ can be expressed as $H(T) \cup Q$, where T is a path from γ to a leaf, and $Q \subseteq P(T)$. The theorem follows by setting γ to the root.

The base case is vacuously true, since for empty S , all relevant sets are empty. Now for the induction. We will have three cases. Let p be the pivot chosen in Step 8. (If S is the root, then there is no pivot. We will directly go to Case (iii) below.)

Case (i): $p \in C$. By construction, there is a link labeled (p, p) to a child of γ . Denote the child β . The child β has label $N(S, p)$. Observe that $C \setminus p$ is a clique in $N(S, p)$ (since by assumption, C is a clique in S .) By induction, there is a unique representation $C \setminus p = H(T) \cup Q$, for path T from the child node to a leaf and

$Q \subseteq P(T)$. Moreover there cannot be a representation of C by a path rooted at β , since $N(S, p) \not\subseteq p$. Consider the path T' that contains T and starts from γ . Note that $H(T') = H(T)$ and $P(T') = P(T) \cup p$. We can express $C = H(T') \cup (Q \cup p)$, noting that $Q \cup p \subseteq P(T')$. This proves the existence of a representation. Moreover, there is only one representation using a path through β .

We need to argue that no other path can represent C . The pivoting is critical for this step. Consider any path rooted at γ , but not passing through β . It must pass through some other child, with corresponding links labeled (v_i, h) , where v_i is a *non-neighbor* of p . Since $C \ni p$, a non-neighbor v_i cannot be in C . Moreover, for any path $\hat{T}$ passing through these other children, $\hat{T}$ must contain some non-neighbor. Thus, $\hat{T}$ cannot represent C .

Case (ii): $C \subseteq N(S, p)$. The argument is essentially identical to the one above. Note that $C \setminus p = C$, and by induction $C \setminus p$ has a unique representation using a path through β . For uniqueness, observe that C does not contain a non-neighbor of p . The previous argument goes through as is.

Case (iii): C contains a non-neighbor of p . Recall that $S \setminus (N(p) \cup p)$ (the set of non-neighbors in S) is denoted $\{v_1, v_2, \dots, v_\ell\}$. Let i be the smallest index i such that $v_i \in C$. For any $1 \leq j \leq \ell$, let $N_j := N(S, v_j) \setminus \{v_1, v_2, \dots, v_{j-1}\}$. Observe that for all j , there is a child labeled N_j . Moreover, all the link labels have h , so for path T passing through N_j , $H(T) \ni v_j$. Thus, if T can represent C , it cannot pass through N_j for $j < i$. Moreover, if $j > i$, then $N_j \not\subseteq v_i$ and no path passing through this node can represent C .

Hence, if there is a path that can represent C , it must pass through N_i . Note that $C \setminus v_i$ is a clique contained in N_i . By induction, there is a unique path T rooted at N_i such that $C \setminus v_i = H(T) \cup Q$, for $Q \subseteq P(T)$. Let T' be the path that extends T to γ . Note that $H(T') = H(T) \cup v_i$, so $C = H(T') \cup Q$. The uniqueness of T implies the uniqueness of T' . $\square$

5 GETTING GLOBAL AND LOCAL COUNTS

The tree T is succinct and yet one can extract fine-grained information from it about all cliques.

Algorithm 2: PIVOTER(G)

 Output: Clique counts of G

- 1 Let $T = \text{SCTBuilder}(G)$.
 - 2 Initialize all clique counts to zero.
 - 3 For every root to leaf path T in T :
 - 4 For every $0 \leq i \leq |P(T)|$, increment $C_{|H(T)|+i}$ by $\binom{|P(T)|}{i}$.
 - 5 For every $v \in H(T)$ and every $0 \leq i \leq |P(T)|$, increment $C_{|H(T)|+i}(v)$ by $\binom{|P(T)|}{i}$.
 - 6 For every $v \in P(T)$ and every $0 \leq i \leq |P(T)| - 1$, increment $C_{|H(T)|+i+1}(v)$ by $\binom{|P(T)|-1}{i}$.
 - 7 Output the sets of values $\{C_k\}$ and $\{c_k(v)\}$.
-

One can also get per-edge counts, but there are three cases. We omit that for space reasons, and give details in the full version. (Per-edge counts can be obtained in $O(\alpha^3 |SCT(G)|)$. Again, this proof is in the full version.)

The storage complexity of the algorithm, as given, is potentially $O(\alpha^2 |SCT(G)|)$, since this is required to store the tree. In the proof of Theorem 2.1, we explain how to reduce the storage.

PROOF. (of Theorem 2.1) **Correctness:** By Theorem 4.1, a root to leaf path T of T represents exactly $2^{P(T)}$ different cliques, with $\binom{P(T)}{i}$ of size $|H(T)| + i$. Moreover, over all T , this accounts for all cliques in the graph. This proves the correctness of global counts.

Pick a vertex $v \in H(T)$. For every subset of $P(T)$, we get a different clique containing v (that is uniquely represented by Theorem 4.1). This proves the correctness of Step 5. For a vertex $v \in P(T)$, we look at all subsets containing v . Equivalently, we get a different represented clique containing v for every subset of $P(T) \setminus v$. This proves the correctness of Step 6.

Running time: Consider the procedure $SCTBuilder(G)$. Note that the size of T is at least n , so we can replace any running time dependence on n by $|T|$. The degeneracy orientation can be found in $O(m + n)$ [20]. For the actual building of the tree, the main cost is in determining the pivot and constructing the children of a node. Suppose a non-root node labeled S is processed. The above mentioned steps can be done by constructing the subgraph induced on S . This can be done in $O(|S|^2)$ time. Since this is not a root node, $|S| \leq \alpha$ (this is the main utility of the degeneracy ordering). Thus, the running time of $SCTBuilder(G) = O(\alpha^2 |T|) = O(\alpha^2 |SCT(G)|)$.

Now we look at PIVOTER. Note that the subsequent counting steps do *not* need the node labels in T ; for all path T , one only needs $P(T)$ and $H(T)$. The paths can be looped over by a DFS from the root. For each path, there are precisely $|P(T)| + 1$ updates to global clique counts, and at most $|H(T) \cup P(T)| \times (|P(T)| + 1)$ updates to local clique counts. The length of T is at most α , and thus both these quantities are $O(\alpha^2)$. Thus, the total running time is $O(\alpha^2 |SCT(G)|)$.

Storage cost: Currently, PIVOTER is represented through two parts: the construction of $SCT(G)$ and then processing it to get clique counts. Conceptually, this is cleaner to think about and it makes the proof transparent. On the other hand, it requires storing $SCT(G)$, which is potentially larger than the input graph. A more space efficient implementation is obtained by combining these steps.

We do not give full pseudocode, since it is somewhat of a distraction. (The details can be found in the code.) Essentially, instead of constructing $SCT(G)$ completely in breadth-first manner, we construct it depth-first through recursion. This will loop over all the paths of T , but only store a single path at any stage. The updates to the clique counts are done as soon as any root to leaf path is constructed. The total storage of a path is the storage for all the labels on a path. As mentioned earlier in the proof of Theorem 2.1, all non-root nodes are labeled with sets of size at most α . The length of the path is at most α , so the total storage is $O(\alpha^2)$. A classic bound on the degeneracy is $\alpha \leq \sqrt{2m}$ (Lemma 1 of [9]), so the storage, including the input, is $O(m + n)$. $\square$

Parallel version of PIVOTER: While this is not central to our results, we can easily implement a parallel version of PIVOTER for *global* clique counts. We stress that our aim was not to delve into complicated parallel algorithms, and merely to see if there was a way to parallelize the counting involving minimal code changes.

The idea is simple, and is an easier variant of the parallelism in kClist [10]. Observe that the children of the root of $SCT(G)$ correspond to finding cliques in the sets $N^+(v)$, for all v . Clique counting in each of these sets can be treated as an independent problem, and can be handled by an independent thread/subprocess. Each subprocess maintains its own array of global clique counts. The final result aggregates all the clique counts. The change ends up being a few lines of code to the original implementation.

Note that this becomes tricky for local counts. Each subprocess cannot afford (storage-wise) to store an entire copy of the local count data structure. The aggregation step would be more challenging. Nonetheless, it should be feasible for each subprocess to create local counts for $N^+(v)$, and appropriately aggregate all counts. We leave this for future work.

6 EXPERIMENTAL RESULTS

Preliminaries: All code for PIVOTER is available here: <https://bitbucket.org/sjain12/pivoter/>. We implemented our algorithms in C and ran our experiments on a commodity machine equipped with a 1.4GHz AMD Opteron(TM) processor 6272 with 8 cores and 2048KB L2 cache (per core), 6144KB L3 cache, and 128GB memory. We performed our experiments on a collection of social networks, web networks, and infrastructure networks from SNAP [49]. The graphs are simple and undirected (for graphs that are directed, we ignore the direction). A number of these graphs have more than 10 million edges, and the largest has more than 100 million edges. Basic properties of these graphs are presented in Tab. 1.

The data sets are split into two parts, in Tab. 1. The upper part are instances feasibly solved with past work (notably kClist40 [10]), while the lower part has instances that cannot be solved with previous algorithm (even after days). We give more details in §6.1.

Competing algorithms: We compare with (what we consider) are the state of the art clique counting algorithms: Turán-Shadow (TS) [16] and kClist40 [10].

kClist40: This algorithm by Danisch-Balalau-Sozio [10] uses degeneracy orientations and parallelization to enumerate all cliques. The kClist40 algorithm, to the best of our knowledge, is the only existing algorithm that can feasibly compute all global counts for some graphs. Hence, our main focus is runtime comparisons with kClist40.

We note that the implementation of kClist40 visits every clique, but only updates the (appropriate) C_k . While it could technically compute local counts, that would require more expensive data structure updates. Furthermore, there would be overhead in combining the counts for independent threads, and it is not immediately obvious how to distribute the underlying data structure storing local counts. As a result, we are unaware of any algorithm that computes local counts (at the scale of dataset in Tab. 1).

We perform a simple optimization of kClist40, to make counting faster. Currently, when kClist40 encounters a clique, it enumerates every smaller clique contained inside it. For the purpose of counting though, one can trivially count all subcliques of a clique using formulas. We perform this optimization (to have a fair comparison with kClist40), and note significant improvements in running time.

In all our runs, for consistency, we run kClist with 40 threads. Note that we compare the *sequential* PIVOTER with the *parallel* kClist40.

TS: This is an approximate clique counting algorithm for k upto 10 [16]. It mines dense subgraphs (shadows) and samples cliques within the dense subgraphs to give an estimate. For fast randomized estimates, it is arguably the fastest algorithm. It runs significantly faster than a sequential implementation of kClist, but is typically comparable with a parallel implementation of kClist. It requires the entire shadow to be available for sampling which can require considerable space.

6.1 Running time and comparison with other algorithms

Running time for global counting: We show the running time results in Tab. 1. For most of the graphs, PIVOTER was able to count all k -cliques in seconds or minutes. For the largest com-orkut graph, PIVOTER ran in 1.5 hours. This is a huge improvement on the state of the art. For the “infeasible” instances in Tab. 1, we do not get results even in two days using previous algorithms. (This is consistent with results in Table 2 of [10], where some of the graphs are also listed as “very large graphs” for which clique counting is hard.)

A notable hard instance is com-1j where PIVOTER is unable to get all clique counts in a day. Again, previous work also notes this challenge, and only gives counts of 7-cliques. We can get some partial results for com-1j, as explained later.

Feasible local counting: Notably, PIVOTER can get per-vertex counts in less than twice the time of global clique counting. Thus, we get results for more graphs in a few minutes, and can process the com-orkut graph within 3 hours. We consider this a significant achievement, given the combinatorial explosion of clique counting.

PIVOTER is also able to get per-edge clique counts, though it can take an order of magnitude more time than global clique counting. Note that for obtaining the per-vertex and per-edge k -clique counts, the result data structure can become extremely large. Indeed, most of the time is spent in updating the data structure, rather than in constructing the SCT. Nonetheless, for all but the as-skitter and com-orkut graph, it runs in minutes.

Comparison with state of the art: We only focus on the “infeasible” instances of Tab. 1. For all the other instances, both PIVOTER and kClist40 get results within two minutes. For space considerations, we do not report all the running times for such instances. It is worth noting that the sequential PIVOTER is comparable to the parallel kClist40 (when they both terminate).

In Tab. 2, we report times on TS and kClist40 on the hard datasets. We are unable to get all values of C_k using either of these two methods. We run these algorithms for up to 100 times the running time of PIVOTER or two days, whichever is shorter. We try to count the largest feasible clique count.

Let us focus on kClist40, where we cannot go beyond counting 13-cliques (we note that this is consistent with results reported in [10]). Notably, in the BerkStan graph, kClist40 needs more than 2 days to count 13-cliques, while PIVOTER gets all clique counts in a minute. As mentioned earlier, clique counting on the large com-orkut graph is done in a few hours by PIVOTER, while even counting 13-cliques takes kClist40 more than two days.

TS also does not scale well for larger cliques and PIVOTER is faster than TS. For example, for the Stanford graph, TS required 230 seconds to estimate the number of 13-cliques whereas PIVOTER obtained all k -clique counts in 5 seconds. Similar trends are observed with other graphs.

Parallel global clique counting: As mentioned in §5, we do a simple parallelization of the global clique counting of PIVOTER using 30 threads. It gives moderate benefits for most instances, and about a factor two speedup for large instances. For the challenging com-1j instances, the effect is much more dramatic. We are able to count 7-cliques in an hour using the parallel PIVOTER, while the sequential version takes more than a day.

Performance on com-1j. This is a particularly challenging graph. The sequential version of PIVOTER for counting all k -cliques did not terminate within a day, so we used the parallel version of our algorithm to show a comparison for global counts upto $k = 10$. We can truncate the SCT to get cliques of some fixed size. Tab. 3 shows the results. Even for this graph, the parallel version of PIVOTER is faster than kClist40 for $k = 7$ and beyond. kClist40 did not terminate after six days, for $k = 9$ and beyond. We note the astronomical number of 10-cliques ($> 10^{19}$), which makes enumeration infeasible, but PIVOTER was able to get the exact count.

Size of $SCT(G)$: In Fig. 2i, we plot the number of nodes of $SCT(G)$ as a function of the number of edges in G . We observe that for most graphs, the size is quite close to m , explaining why PIVOTER is efficient.

6.2 Demonstrations of PIVOTER

Global and local cliques have numerous applications. It is outside the scope of this work for detailed demonstrations, but we show a few examples in this section.

As mentioned earlier, local clique counts are an important aspect of graph processing. In Fig. 1ii and Fig. 1iii, we plot the per-vertex clique distributions, also called the *graphlet* degree distribution in bioinformatics [22] for the as-skitter and web-Stanford graphs. We choose values of $k = 5, 10, 15, 20, 25$. Then, we plot the function $f_k(b)$ that is the number of vertices that participate in b k -cliques. We notice interesting trends. While the as-skitter graph has a nicely decaying f_k function, there is much more noise in web-Stanford. It would be interesting to design models that can capture such behavior in the local clique counts.

In Fig. 2ii, we plot the C_k values for a number of graphs. We notice, for example, that the soc-pokec network has a “flatter” distribution of C_k for some of the initial values, while the com-orkut graph looks much closer to a binomial distribution. The latter suggests that the bulk of cliques are coming from the maximum clique in the com-orkut graph, but not so in the soc-pokec graph.

In Fig. 2iii, we plot the k -clique counts (vs k) for two different versions across time for the DBLP citation network [28]. Interestingly, despite the later version only having less than twice as many edges, the clique distribution (plotted in semilog) has a much bigger difference. It appears that the graph is becoming significantly dense in certain part. This sort of analysis may help in understanding dynamic graphs.

Graph	Vertices	Edges	Degen	Max clique	PIVOTER (C_k)	PIVOTER ($c_k(v)$)	PIVOTER ($c_k(e)$)	PIVOTER (C_k parallel)
Feasible by previous algorithms								
dblp-v5	1.56E+06	2.08E+06	15	10	7	7	8	19
dblp-v7	3.67E+06	4.18E+06	19	12	15	16	19	34
amazon0601	4.03E+05	2.44E+06	10	11	4	5	6	4
web-Google	8.76E+05	4.32E+06	44	44	8	9	15	9
youtube	1.13E+06	2.99E+06	51	17	7	8	11	9
cit-Patents	3.77E+06	1.65E+07	64	11	40	41	53	46
soc-pokec	1.63E+06	2.23E+07	47	29	68	75	93	44
Not feasible for previous algorithms								
Stanford	2.82E+05	1.99E+06	71	61	5	5	38	3
BerkStan	6.85E+05	6.65E+06	201	201	25	26	237	9
as-skitter	1.70E+06	1.11E+07	111	67	120	200	9245	75
com-orkut	3.07E+06	1.17E+08	253	51	5174	8802	99389	3441
com-lj	4.00E+06	3.47E+07	360	-	-	-	-	108000*

Table 1: Table shows the sizes, degeneracy, maximum clique size, and the time taken (in seconds) by PIVOTER to obtain global k -clique counts, per-vertex and per edge k -cliques counts for all k . *For the com-lj graph, we were not able to get all k -clique counts in 1 day so we tested for the maximum k we could count in about a day. PIVOTER was able to count the number of 9-cliques in 30 hours whereas kClist40 had not terminated even after 6 days.

Figure 2: Fig. 2i shows the number of nodes in the SCT vs the number of edges (m) for different graphs. The running time of PIVOTER is directly proportional to the SCT size which seems to be roughly linear in the number of edges. Fig. 2ii shows the trends in clique counts for a number of graphs. For some of the graphs, the complete distribution of their clique counts has been obtained for the first time. Fig. 2iii shows the trends in the clique counts of 2 different versions over time of the dblp graph.

Graph	k=13, TS	k=13, kClist40	all k , PIVOTER
Stanford	230	12600	5
BerkStan	1198	> 172800	25
as-skitter	798	12480	120
com-orkut	> 28800	> 172800	5174

Table 2: Time taken in seconds by the state-of-the-art randomized (TS, short for TuránShadow) and parallel (kClist40) algorithms. Note that PIVOTER obtains all k -clique counts for these graphs in a fraction of the time taken by other methods to count just 13-cliques.

k	k-cliques	kClist40	PIVOTER
7	4.49E+15	2.2 hours	1.2 hours
8	1.69E+16	42.5 hours	6.4 hours
9	5.87E+17	> 6 days	30 hours
10	1.89E+19	> 6 days	5.9 days

Table 3: Table shows the time taken to count k -cliques for com-lj graph. For $k=9$, PIVOTER terminated in about 30 hours where kClist40 had not terminated in 6 days.

7 FUTURE WORK

We provide an exact clique counting algorithm that counts all k -cliques in a fraction of the time of other state-of-the-art parallel algorithms. One of the key ideas is the use of pivoting to create the SCT, and succinct representation of all the cliques of the graph. The success of [10] in using parallelization for clique counting suggests combining their ideas with our pivoting techniques. We may be able to come up with an efficient parallel building of the SCT that is much faster than our current implementation. Indeed, the results on the com-1j graph suggest that even PIVOTER has its limits for real data.

An orthogonal approach would be to exploit the sampling techniques in the Turán-Shadow algorithm [16]. For many subgraph counting problems, randomization has been the key to truly practical algorithms. We believe that PIVOTER could be made faster with these ideas.

Moreover, it also gives per-edge and per-vertex k -clique counts. This is the first time that k -clique counts are known for many of the graphs we experimented with and this will open doors for further use of cliques in generation and analysis of graphs.

ACKNOWLEDGMENTS

Shweta Jain and C. Seshadhri acknowledge the support of NSF Awards CCF-1740850, CCF-1813165, and ARO Award W911NF1910294.

REFERENCES

- [1] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, and Nick Duffield. 2015. Efficient Graphlet Counting for Large Networks. In *Proceedings of International Conference on Data Mining (ICDM)*.
- [2] E. A. Akkoyunlu. 1973. The enumeration of maximal cliques of large graphs. *SIAM J. Comput.* 2 (1973), 1–6.
- [3] Noga Alon, Ráphy Yuster, and Uri Zwick. 1994. Color-coding: A New Method for Finding Simple Paths, Cycles and Other Small Subgraphs Within Large Graphs. In *Symposium on the Theory of Computing (STOC)*. 326–335. <https://doi.org/10.1145/195058.195179>
- [4] A. Benson, D. F. Gleich, and J. Leskovec. 2016. Higher-order organization of complex networks. *Science* 353, 6295 (2016), 163–166.
- [5] Jonathan W. Berry, Bruce Hendrickson, Randall A. LaViolette, and Cynthia A. Phillips. 2011. Tolerating the community detection resolution limit with edge weighting. *Phys. Rev. E* 83 (May 2011), 056119. Issue 5. <https://doi.org/10.1103/PhysRevE.83.056119>
- [6] Nadja Betzler, René van Bevern, Michael R. Fellows, Christian Komusiewicz, and Rolf Niedermeier. 2011. Parameterized Algorithmics for Finding Connected Motifs in Biological Networks. *IEEE/ACM Trans. Comput. Biology Bioinform.* 8, 5 (2011), 1296–1308.
- [7] Coen Bron and Joep Kerbosch. 1973. Algorithm 457: Finding All Cliques of an Undirected Graph. *Commun. ACM* 16, 9 (Sept. 1973), 575–577. <https://doi.org/10.1145/362342.362367>
- [8] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. 2004. Linear FPT reductions and computational lower bounds. In *Symposium on the Theory of Computing (STOC)*, László Babai (Ed.). ACM, 212–221. <https://doi.org/10.1145/1007352.1007391>
- [9] Norishige Chiba and Takao Nishizeki. 1985. Arboricity and subgraph listing algorithms. *SIAM J. Comput.* 14 (1985), 210–223. Issue 1. <https://doi.org/10.1137/0214017>
- [10] Maximilien Danisch, Oana Denisa Balalau, and Mauro Sozio. 2018. Listing k -cliques in Sparse Real-World Graphs. In *World Wide Web (WWW)*. 589–598. <https://doi.org/10.1145/3178876.3186125>
- [11] David Eppstein, Maarten Löffler, and Darren Strash. 2010. Listing all maximal cliques in sparse graphs in near-optimal time. In *International Symposium on Algorithms and Computation*. Springer, 403–414.
- [12] David Eppstein, Maarten Löffler, and Darren Strash. 2013. Listing All Maximal Cliques in Large Sparse Real-World Graphs. *ACM Journal of Experimental Algorithmics* 18 (2013). <https://doi.org/10.1145/2543629>
- [13] Irene Finocchi, Marco Finocchi, and Emanuele G. Fusco. 2015. Clique Counting in MapReduce: Algorithms and Experiments. *ACM Journal of Experimental Algorithmics* 20 (2015). <https://doi.org/10.1145/2794080>
- [14] Robert A. Hanneman and Mark Riddle. 2005. *Introduction to social network methods*. University of California, Riverside. <http://faculty.ucr.edu/~hanneman/nettext/>.
- [15] Matthew O. Jackson. 2010. *Social and Economic Networks*. Princeton University Press.
- [16] Shweta Jain and C Seshadhri. 2017. A Fast and Provable Method for Estimating Clique Counts Using Turán's Theorem. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 441–449.
- [17] M. Jha, C. Seshadhri, and A. Pinar. 2015. Path Sampling: A Fast and Provable Method for Estimating 4-Vertex Subgraph Counts. In *World Wide Web (WWW)*. 495–505.
- [18] Zhenqi Lu, Johan Wahlström, and Arye Nehorai. 2018. Community detection in complex networks via clique conductance. *Scientific reports* 8, 1 (2018), 5982.
- [19] Dror Marcus and Yuval Shavitt. 2010. Efficient Counting of Network Motifs. In *ICDCS Workshops*. IEEE Computer Society, 92–98.
- [20] David W Matula and Leland L Beck. 1983. Smallest-last ordering and clustering and graph coloring algorithms. *Journal of the ACM (JACM)* 30, 3 (1983), 417–427.
- [21] Ali Pinar, C Seshadhri, and Vaidyanathan Vishal. 2017. Escape: Efficiently counting all 5-vertex subgraphs. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1431–1440.
- [22] Natasa Przulj. 2007. Biological network comparison using graphlet degree distribution. *Bioinformatics* 23, 2 (2007), 177–183.
- [23] Rahmtin Rotabi, Krishna Kamath, Jon M. Kleinberg, and Aneesh Sharma. 2017. Detecting Strong Ties Using Network Motifs. In *World Wide Web (WWW)*. 983–992. <https://doi.org/10.1145/3041021.3055139>
- [24] Ahmet Erdem Sariyüce, C. Seshadhri, Ali Pinar, and Ümit V. Çatalyürek. 2015. Finding the Hierarchy of Dense Subgraphs using Nucleus Decompositions. In *World Wide Web (WWW)*. ACM, 927–937.
- [25] C. Seshadhri, Ali Pinar, and Tamara G. Kolda. 2014. Wedge sampling for computing clustering coefficients and triangle counts on large graphs. *Statistical Analysis and Data Mining* 7, 4 (2014), 294–307. <https://doi.org/10.1002/sam.11224>
- [26] C. Seshadhri and Srikanta Tirathapura. 2019. Scalable Subgraph Counting: The Methods Behind The Madness: WWW 2019 Tutorial. In *Proceedings of the Web Conference (WWW)*.
- [27] Ann Sizemore, Chad Giusti, and Danielle S. Bassett. 2016. Classification of weighted networks through mesoscale homological features. *Journal of Complex Networks* 10.1093 (2016).
- [28] Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. 2008. ArnetMiner: Extraction and Mining of Academic Social Networks. In *KDD'08*. 990–998.
- [29] Etsuji Tomita, Akira Tanaka, and Haruhisa Takahashi. 2006. The Worst-Case Time Complexity for Generating All Maximal Cliques. *Theoretical Computer Science* 363, 1 (2006), 28–42. https://doi.org/10.1007/978-3-540-27798-9_19
- [30] Charalampos E. Tsourakakis. 2015. The K -clique Densest Subgraph Problem. In *Proceedings of the Conference on World Wide Web WWW*. 1122–1132. <https://doi.org/10.1145/2736277.2741098>
- [31] Charalampos E. Tsourakakis, Jakub Pachocki, and Michael Mitzenmacher. 2017. Scalable Motif-aware Graph Clustering. In *World Wide Web (WWW)*. 1451–1460. <https://doi.org/10.1145/3038912.3052653>
- [32] Johan Ugander, Lars Backstrom, and Jon M. Kleinberg. 2013. Subgraph frequencies: mapping the empirical and extremal geography of large graph collections. In *WWW*. 1307–1318.
- [33] Virginia Vassilevska. 2009. Efficient algorithms for clique problems. *Inform. Process. Lett.* 109, 4 (2009), 254 – 257. <https://doi.org/10.1016/j.ipl.2008.10.014>
- [34] Pinghui Wang, Junzhou Zhao, Xiangliang Zhang, Zhenguo Li, Jiefeng Cheng, John C. S. Lui, Don Towsley, Jing Tao, and Xiaohong Guan. 2018. MOSS-5: A Fast Method of Approximating Counts of 5-Node Graphlets in Large Graphs. 30, 1 (2018), 73–86. <https://doi.org/10.1109/TKDE.2017.2756836>
- [35] Hao Yin, Austin R. Benson, and Jure Leskovec. 2018. Higher-order clustering in networks. *Phys. Rev. E* 97 (2018), 052306. <https://doi.org/10.1145/3289600.3290991>
- [36] Hao Yin, Austin R. Benson, and Jure Leskovec. 2019. The Local Closure Coefficient: A New Perspective On Network Clustering. 303–311. <https://doi.org/10.1145/3289600.3290991>
- [37] Z. Zhao, G. Wang, A. Butt, M. Khan, V. S. Anil Kumar, and M. Marathe. 2012. SAHAD: Subgraph Analysis in Massive Networks Using Hadoop. In *Proceedings of International Parallel and Distributed Processing Symposium (IPDPS)*. 390–401.