



# Scalable Approximate Biclique Counting over Large Bipartite Graphs

Jingbang Chen\*  
chenjb@cuhk.edu.cn  
CUHK-Shenzhen & SLAI

Weinuo Li\*  
liweinuo@zju.edu.cn  
Zhejiang University

Yingli Zhou\*  
yinglizhou@link.cuhk.edu.cn  
CUHK-Shenzhen

Hangrui Zhou  
zhouhr23@mails.tsinghua.edu.cn  
Tsinghua University

Qiuyang Mang  
qiuyangmang@link.cuhk.edu.cn  
CUHK-Shenzhen

Can Wang  
wcan@zju.edu.cn  
Zhejiang University

Yixiang Fang  
fangyixiang@cuhk.edu.cn  
CUHK-Shenzhen

Chenhao Ma<sup>†</sup>  
machenhao@cuhk.edu.cn  
CUHK-Shenzhen

## ABSTRACT

Counting  $(p, q)$ -bicliques in bipartite graphs is crucial for a variety of applications, from recommendation systems to cohesive subgraph analysis. Yet, it remains computationally challenging due to the combinatorial explosion to exactly count the  $(p, q)$ -bicliques. In many scenarios, e.g., graph kernel methods, however, exact counts are not strictly required. To design a scalable and high-quality approximate solution, we novelly resort to  $(p, q)$ -broom, a special spanning tree of the  $(p, q)$ -biclique, which can be counted via graph coloring and efficient dynamic programming. Based on the intermediate results of the dynamic programming, we propose an efficient sampling algorithm to derive the approximate  $(p, q)$ -biclique count from the  $(p, q)$ -broom counts. Theoretically, our method offers unbiased estimates with provable error guarantees. Empirically, our solution outperforms existing approximation techniques in both accuracy (up to  $8\times$  error reduction) and runtime (up to  $50\times$  speedup) on nine real-world bipartite networks, providing a scalable solution for large-scale  $(p, q)$ -biclique counting.

## PVLDB Reference Format:

Jingbang Chen, Weinuo Li, Yingli Zhou, Hangrui Zhou, Qiuyang Mang, Can Wang, Yixiang Fang, and Chenhao Ma. Scalable Approximate Biclique Counting over Large Bipartite Graphs. PVLDB, 19(4): 576 - 588, 2025. doi:10.14778/3785297.3785301

## PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/lwn16/Biclique>.

## 1 INTRODUCTION

The bipartite graph stands as a cornerstone in graph mining, comprising two distinct sets of vertices where edges exclusively link

vertices from different sets. This concept serves as a powerful tool for modeling relationships across various real-world domains, including recommendation networks [7], collaboration networks [2], and gene coexpression networks [15]. For example, in recommendation networks, users and items typically constitute two distinct vertex types. The interactions between vertices form a bipartite graph, with edges representing users' historical purchase behaviors.

In the realm of bipartite graph analysis, counting  $(p, q)$ -bicliques has attracted much attention. A  $(p, q)$ -biclique is a complete subgraph between two distinct vertex sets of size  $p$  and  $q$ . The  $(2, 2)$ -bicliques, also known as butterflies, have played an important role in the analysis of bipartite networks [20, 27, 28]. Furthermore, there are more scenarios in which  $p, q$  is not fixed to 2. In general, counting  $(p, q)$ -bicliques serves as a fundamental operator in many applications, including cohesive subgraph analysis [3] and information aggregation in graph neural networks [32], and densest subgraph mining [21]. In higher-order bipartite graph analysis, the clustering coefficient is the ratio between the counts of  $(p, q)$ -bicliques and  $(p, q)$ -wedges, where counting  $(p, q)$ -wedges can be reduced to counting  $(p, q)$ -bicliques [34].

Despite its importance, counting  $(p, q)$ -bicliques is very challenging due to its exponential increase with respect to  $p$  and  $q$  [32]. For example, in one of the graph datasets Twitter, with fewer than  $2 \times 10^6$  edges, there are more than  $10^{13}$   $(5, 4)$ -bicliques and more than  $10^{18}$   $(6, 3)$ -bicliques within. Therefore, enumeration-based counting methods including BCList++ [32], EPivoter [34] that produce the exact solution are not scalable. On the other hand, in many applications of biclique counting, an approximate count is often sufficient. In **graph kernel methods** [26], motifs serve as the basis for defining similarity measures between graphs in tasks like classification and anomaly detection. Since graph kernels depend on relative similarities rather than exact counts, approximate counts can preserve kernel performance while significantly reducing the computational overhead. This efficiency facilitates similarity computations in various fields:

- **Bioinformatics:** Zhang et al. [36] proposed algorithms to enumerate maximal bicliques in bipartite graphs derived from functional genomics (genes vs. gene-sets), enabling the identification of co-regulatory modules. A high number of bicliques in certain regions often indicates dense regulatory relationships, which can

\*The first three authors contributed equally to this research.

<sup>†</sup>Chenhao Ma is the corresponding author.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 19, No. 4 ISSN 2150-8097. doi:10.14778/3785297.3785301

be helpful for detecting such modules. However, full enumeration quickly becomes infeasible as graphs grow larger or the number of bicliques explodes. This limitation motivates our work on scalable approximate biclique counting, which aims to estimate biclique numbers (or those of particular sizes) accurately without exhaustive enumeration.

- **Cybersecurity:** In e-commerce fraud detection, bicliques also serve as key indicators of suspicious activity. For example, Wu et al. [30] show that in Taobao’s large transaction network, click-farming attacks can be modeled as large bicliques. Similarly, Ban [1] introduces the Maximal Half-Isolated Biclique problem to detect synchronized fraudulent behavior in customer–product graphs. These studies highlight that high biclique counts—or the presence of unusually large bicliques—in specific regions of a bipartite graph can be strong signals of fraud.
- **Data Mining:** In recommendation systems [27], biclique counting helps uncover dense substructures among users or items, such as groups of users with shared interests or items commonly co-purchased. In large-scale environments such as e-commerce or streaming platforms, allowing a small error margin for counting can significantly reduce the runtime cost while still preserving the effectiveness of the downstream tasks.

From this perspective, we can see that developing algorithms that produce approximate answers to  $(p, q)$ -biclique counting is more practical and more applicable to large-scale data. Ye et al. propose a sampling-based method called EP/Zz++ for approximate counting  $(p, q)$ -bicliques [34]. However, its precision is not satisfactory as  $p, q$  increases. For example, when querying  $(5, 9)$ -bicliques in the data set Twitter, EP/Zz++ could produce an answer of the error ratio of more than 200%.<sup>1</sup> This shows the need to develop a better algorithm for approximate  $(p, q)$ -biclique counting, aiming to improve scalability and accuracy.

In this paper, we propose a new sampling-based method that produces a high-accuracy approximate counting of  $(p, q)$ -bicliques. We first adapt the coloring trick in this problem, inspired by several counting methods of  $k$ -cliques in general graphs [18, 33]. Then, we design a special type of subgraph that has a strong correlation with  $(p, q)$ -biclique, named  $(p, q)$ -brooms. The  $(p, q)$ -broom is a special type of spanning tree in a  $(p, q)$ -biclique, which has a fixed structure that can be taken advantage of when counting the amount. Using dynamic programming, we can count this motif efficiently and precisely. It also naturally adapts to the coloring. Finally, we develop a sampling method that takes advantage of the coloring and the intermediate result of dynamic programming. Compared to the h-zigzag pattern proposed by EP/Zz++, our  $(p, q)$ -brooms have a structure that is closer to  $(p, q)$ -cliques, and they usually have a smaller amount in the graph. As a result, it provides a better error guarantee when sampling  $(p, q)$ -bicliques from them. Our algorithm is still heuristic since it includes a graph-coloring process, and the sampling time analysis contains a ratio that could be exponential. However, these processes are shown to be very efficient in practice. We empirically evaluate our algorithm against state-of-the-art exact and approximate counting algorithms on nine

real-world datasets. All experimental results demonstrate the effectiveness of our algorithm. Besides the mathematical analysis, the high-level idea here is the sparsification. Since  $(p, q)$ -bicliques in real-world datasets contain realistic meanings, finding a sparsifier that is likely to preserve the information within bicliques is a promising approach to reducing the computational cost in practice. On the other hand, the mathematical analysis of unbiasedness and sampling time ensures that we can use this algorithm in practice. We can set the sampling time to a large number, allowing us to withstand the overall cost. Or we can iteratively increase the sampling time and terminate while the result shows a convergence trend.

Our contributions are summarized as follows:

- We design  $(p, q)$ -broom, a special type of spanning tree in a  $(p, q)$ -biclique, which can be efficiently counted via dynamic programming.
- Based on the dynamic programming result of counting  $(p, q)$ -brooms, we develop an efficient sampling-based algorithm that computes a high-accuracy counting of  $(p, q)$ -bicliques in bipartite graphs.
- Mathematical analysis shows that our algorithm is unbiased and obtains a provable error guarantee.
- Extensive experiments show that our algorithm consistently outperforms state-of-the-art approximate algorithms, with up to  $8\times$  reduction in approximation errors and up to  $50\times$  speed-up in running time.
- Additional experiments further support the effectiveness and novelty of our algorithm, which includes empirical results such as an ablation study on the newly designed coloring method and an experiment on the hyper-parameters.

**Outline.** The rest of the paper is organized as follows. We review the related work in Section 2, and introduce notations and definitions in Section 3. Section 4 presents our sampling-based algorithm. Experimental results are given in Section 5 and Section 6. We present major experimental results in Section 5 and provide several detailed analyses. We conclude in Section 7.

## 2 RELATED WORKS

*Biclique Counting.* The biclique counting problem focuses on enumerating  $(p, q)$ -bicliques within bipartite graphs, with particular emphasis on the  $(2, 2)$ -biclique, commonly known as a butterfly—a fundamental structural pattern with significant real-world applications. Given the importance of butterfly motifs, researchers have developed numerous algorithms for their efficient enumeration and counting [24, 27, 28]. The state-of-the-art method utilizes the vertex priority and cache optimization [28]. Recent advances have emerged in multiple directions: parallel computing techniques that optimize both memory and time efficiency [31], I/O-efficient methods [29] that minimize disk operations, and approximation strategies that provide accurate counts while reducing computational overhead. The field has further expanded to address more complex graph types, including uncertain graphs with probabilistic edges [38] and temporal graphs incorporating time-varying relationships [6]. The  $(p, q)$ -biclique counting is a general version of butterfly counting, which recently gained much attention [32, 34]. The state-of-the-art methods are introduced and analyzed in the latter (See Section 3.1).

<sup>1</sup>Let  $C$  and  $\hat{C}$  be the exact and approximate counts of the  $(p, q)$ -biclique, respectively. The estimation error ratio is defined as  $\frac{|C-\hat{C}|}{C}$ .

*k*-clique Counting. The evolution of *k*-clique counting algorithms began with Chiba and Nishizeki’s backtracking enumeration method, which was later enhanced through more efficient ordering-based optimizations, including degeneracy ordering (Danisch et al. [9]) and color ordering (Li et al. [18]). While these algorithms perform efficiently for small *k*, their performance deteriorates with increasing *k*. Pivoter [13], introduced by Jain and Seshadhri, marked a significant advancement by employing pivoting techniques from maximal clique enumeration, enabling combinatorial counting instead of explicit enumeration. However, Pivoter’s performance can degrade on large, dense graphs. To address scalability challenges, researchers have developed sampling-based approaches, including TuranShadow [14] and coloring-based sampling techniques [33].

*Other Motif Counting.* Beyond *k*-clique and  $(p, q)$ -biclique, numerous works are focusing on counting general motifs, such as four-cycles and other subgraph patterns [4, 12, 19, 25, 37, 39]. These methods can be broadly divided into two categories: traditional counting approaches [4, 19, 25] and learning-based methods [12, 37]. The former is typically based on the sampling and enumeration techniques used to obtain the approximate and exact count of the motif, respectively. In contrast, learning-based approaches offer an alternative approximate solution by leveraging machine learning techniques. Beyond this primary categorization, subgraph counting can be further classified by scope: global counting determines subgraph frequencies across the entire graph, while local counting focuses on specific nodes or edges.

*Dynamic Programming in Subgraph Problems* In our proposed algorithm, dynamic programming is used to compute the number of special types of spanning trees. In many recent works for subgraph matching algorithms [10, 16], they also use dynamic programming to search a data graph for embeddings of a path tree obtained from a query graph. Although we all fix an ordering for these trees and conduct dynamic programming according, our method focuses and works differently: We use both the intermediate and final results to ensure the following sampling process can be effective. Furthermore, our transition formula and recurrence relations are significantly different. The key idea of our method is to take advantage of the bipartite structure. We split the tree into several “prefix” subtrees and identify them by the number of edges and the last edge that adds in. In this way, we manage to compute the number of all trees as well as the number of all “prefix” tree.

### 3 PRELIMINARIES

Throughout the paper, we study the bipartite graphs. A bipartite graph is an undirected graph  $G = ((U, V), E)$ , where  $U$  and  $V$  are disjoint sets of vertices and each edge  $(u, v) \in E$  satisfies  $u \in U, v \in V$ . To denote the corresponding vertex and edge sets of a certain  $G$ , we may use  $U(G)$ ,  $V(G)$ , and  $E(G)$ , respectively. For each vertex  $u \in U(G)$ , we denote its neighbor set as  $N(u, G) = \{v \in V \mid (u, v) \in E\}$ . Similarly, for a vertex  $v \in V(G)$ , its neighbor set is  $N(v, G) = \{u \in U \mid (u, v) \in E\}$ . For any vertex  $x \in U \cup V$ , its degree  $d(x)$  is the size of its neighbor set. Now we define the  $(p, q)$ -biclique.

*Definition 3.1 (Biclique).* Given a bipartite graph  $G = ((U, V), E)$ , a  $(p, q)$ -biclique is a subgraph  $G' = ((U', V'), E') \in G$  where  $|U'| = p$ ,  $|V'| = q$  and  $E' = \{(u, v) \mid u \in U', v \in V'\}$ .

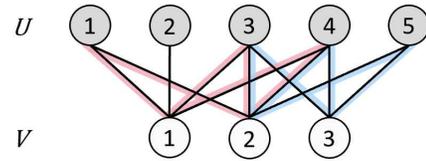


Figure 1: An illustrative example of two  $(3, 2)$ -bicliques.

Lastly, we define the problem that we mainly study:

**PROBLEM 1 (BICLIQUE COUNTING).** Given a bipartite graph  $G$  and two integers  $p$  and  $q$ , compute the number of  $(p, q)$ -bicliques in  $G$ .

For example, in Figure 1, there are two  $(3, 2)$ -bicliques: **(1)**  $U' = \{1, 3, 4\}$ ,  $V' = \{1, 2\}$ ; **(2)**  $U' = \{3, 4, 5\}$ ,  $V' = \{2, 3\}$ . Their edges are highlighted with pink lines and blue lines, respectively. Therefore, the number of  $(3, 2)$ -bicliques is 2. It is also known that Problem 1 is NP-Hard [32], which means there is no algorithm that runs in polynomial time. Throughout this paper, we assume  $p, q \geq 2$ .

### 4 ALGORITHM

In this section, we propose a new algorithm called *Colored Broom-based Sampling (CBS)* that solves Problem 1 approximately.

*1. Coloring.* Inspired by methods for solving *k*-clique counting problems on general graphs [18, 33], we start by assigning a color to every vertex on the graph (Section 4.1). Our coloring method is simple, and produces assignments with a small number of colors empirically. Although graph coloring has already been used in *k*-clique problems on general graphs, it is a new attempt to utilize it on bipartite graphs for biclique problems. On the other hand, if we directly apply the traditional coloring, it becomes 2-coloring and cannot ensure the distinct color assignment for every existing biclique. In general, coloring can be considered as an optimization aiming to reduce the number of accountable brooms. Intuitively, each  $(p, q)$ -biclique must obtain at least  $p$  or  $q$  colors due to the definition. In turn, increase the probability such that the sampled broom is a biclique. It also allows us to construct a dynamic programming implemented by following the color indices.

We construct a new coloring condition and prove its effectiveness in result. Since the graph-coloring problem is NP-Complete, there should be no theoretically optimal guarantee. In our empirical study, we have conducted an ablation study (Section 6.1) showing the effectiveness of the coloring process as well as slightly different coloring strategies.

*2. Broom Pattern Counting.* Instead of counting bicliques directly, our next step is to calculate the number of a special motif called “broom” that we design (Section 4.2). It is a special type of spanning trees of  $(p, q)$ -bicliques. Our first insight is that a spanning tree captures a subgraph better than a simple path. Then, since in the real-world dataset,  $(p, q)$ -bicliques contain realistic meanings and we also hope the sparsified motif can preserve this inner information. Therefore, we try to allocate all  $p + q - 1$  edges to each vertex as evenly as possible. In turn, we have the following mathematical definition of  $(p, q)$ -brooms that indicates this purpose:

*Definition 4.1 (Broom).* Given a bipartite graph  $G((U, V), E)$ , two vertex sets  $U' \in U$  and  $V' \in V$  of size  $p$  and  $q$ , respectively. A

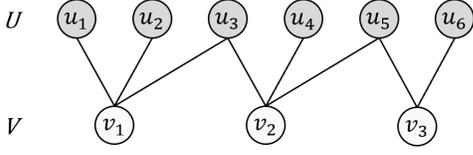


Figure 2: An illustrative example of a (6, 3)-broom.

vertex ordering is also given:  $\{u_1, u_2, \dots, u_p\}$  and  $\{v_1, v_2, \dots, v_q\}$ . The corresponding  $(p, q)$ -broom is the subgraph  $G'((U', V'), E') \in G$  that satisfies

$$E' = \left\{ (u_i, v_{\lfloor \frac{(i-1)(q-1)}{p-1} + 1 \rfloor}) \mid 1 \leq i \leq p \right\} \\ \cup \left\{ (u_{\lceil \frac{(i-1)(p-1)}{q-1} \rceil}, v_i) \mid 2 \leq i \leq q \right\}.$$

In Figure 2, we provide an illustration of a (6, 3)-broom. We can observe that such a subgraph is, in fact, a tree, which contains  $p + q - 1$  edges and is connected. As shown in Figure 2, we can clearly see that each vertex  $v_i$  in  $V$  connects to the 2–3 vertices  $u_j$  in  $U$  that  $i$  and  $j$  are close. Note that the defined ordering is only used to characterize the pattern, providing a better visualization, and it does not have to align with the original vertex index. Visually speaking, a  $(p, q)$ -broom is the "skeleton" of a  $(p, q)$ -biclique. The general idea here is to design a way to sparsify the dense biclique while trying best to preserve its identity. This broom pattern is very different from EP/Zz++ [34]'s  $h$ -zigzag patterns, as they only use a simple path. We then empirically show that such pattern's count can lead to a much better approximation of biclique counting. This distribution design limits the total number of existing brooms, which makes the error analysis tight. Despite its complicated structures, we show in Section 4.2 that its amount can be computed precisely and efficiently using a dynamic programming process with respect to color assignment.

3. *Counting by Sampling Brooms.* Lastly, by sampling, we extract the quantity relationship between the  $(p, q)$ -brooms and  $(p, q)$ -bicliques in the graph and eventually compute our approximate answer (Section 4.3). Shown in Section 5, our algorithm produces high-accuracy approximate counting with faster runtime in real-world datasets. We also prove unbiasedness and provide an error guarantee for our algorithm.

## 4.1 Coloring

The color assignment should guarantee that for any  $(p, q)$ -clique  $H$ , there are no two vertices in  $U(H)$  that share the same color. Similarly, the same is true for any two vertices in  $V(H)$ . The way to achieve this is to relax the condition as "limiting same-color-endpoint-wedges  $\leq q$ ", which is easy to keep track of.

Our coloring algorithm is shown in Algorithm 1. We run the coloring process `CalColor` for the vertices in  $U$  and  $V$  separately (Line 1, 2). Note that here, we consider the neighbor symmetrically. After preprocessing the input graph  $G$  with Algorithm 1, each vertex  $x \in U \cup V$  is colored, and we denote its color as  $c(x)$ .

$S$  denotes the set of vertices that has not been colored. Whenever  $S$  is not empty (Line 7), we will try coloring some vertices with a new color, denoted as  $\kappa$ . For brevity, we initialize  $\kappa$  to be 1 (Line 6), and whenever we need a new color, we increase  $\kappa$  by 1 (Line 16).  $Cnt$

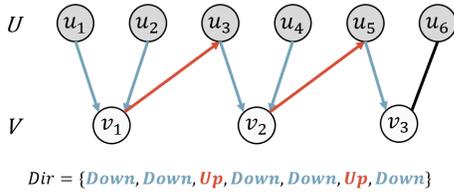
is a temporary array that is used to guarantee the coloring is legal. Specifically, when we try coloring with color  $\kappa$ ,  $Cnt[u]$  stores the number of  $v$  that satisfies  $v \in N(u, G)$  and there exists  $w \in N(v, G)$  such that  $w \neq u$  and  $c(w) = \kappa$ . In each round of coloring, we start by initializing  $Cnt$  to be all zeros (Line 8). We use  $S_{next}$  to store vertices that have to be colored in the next round, initialized as  $\emptyset$  (Line 9). We enumerate vertices in  $S$  in a randomized order (Line 10). For each enumerated vertex  $u$ , we first set its color to be  $\kappa$  (Line 11) and update  $Cnt$  correspondingly (Line 12). If there exists a vertex  $x$  such that  $Cnt[x] \geq q$  and  $c(x) = \kappa$  (Line 13), then there could exist a  $(p, q)$ -biclique with two vertices with the same color. Therefore, such coloring is illegal, and we cannot color  $u$  with  $\kappa$ . In such case, we first revert the changes in  $Cnt$  (Line 14) and add  $u$  into  $S_{next}$  (Line 15). In the end, we have tried every vertex in  $S$  at least once and have deferred some vertices' coloring to upcoming rounds. We assign  $S_{next}$  to  $S$  and continue (Line 17).

To see that the runtime of this coloring process is dominated by the total time of the total number of colors, i.e.  $\kappa$ , we need to show that querying, updating, and reverting  $cnt$  can be done in  $O(|N(u, G)|)$  each operation. Here we provide one possible implementation about how to achieve  $|N(u, G)|$  for a single update as follows: The idea is to use the lazy propagation trick. We only maintain  $cnt$  from vertices that are (will be) colored  $\kappa$ , and each round we maintain a status  $status[v] = (t, w)$  for  $V$  during each round coloring  $U$ . Specifically,  $t = 0$  indicates  $v$  has no neighbor of color  $\kappa$ ;  $t = 1$  means exactly one neighbor  $w$  of color  $\kappa$ , and  $t = 2$  means at least two neighbors are colored  $\kappa$ . Whenever we color  $u$  to  $\kappa$ , we need to check all its neighbor  $v \in N(u, G)$ . Let  $status[v] = (t, w)$ . If  $t \geq 1$ , we update  $Cnt[u]$ . If  $t = 1$ , we update  $Cnt[w]$ . After it, we update  $status[v]$  accordingly. In all, for each  $u$ , we update its neighbor  $v$  with  $O(1)$  operations. Therefore, the total cost for a single update is  $O(|N(u, G)|)$ . In all, the time complexity is  $O(\kappa|E|)$ . Note that  $\kappa$  is naturally bounded by  $|U(H)|$ . In fact, shown in Table 1, on real-world datasets,  $\kappa$  is not large and our coloring algorithm runs very efficiently in turn. The only question left is how such an algorithm guarantees a correct color assignment. We prove it in the following lemma.

LEMMA 4.2. *After executing `Coloring`( $G, p, q$ ), for any  $(p, q)$ -clique  $H$  in  $G$ , there are no two vertices  $x, y$  either both in  $U(H)$  or in  $V(H)$  that share the same color, i.e.  $c(x) = c(y)$ .*

PROOF. Let  $U(H) = \{u_1, u_2, \dots, u_p\}$ ,  $V(H) = \{v_1, v_2, \dots, v_q\}$ . Because of symmetry, we only prove the case of  $U$  and  $p$ . We assume that there are two vertices  $u_1, u_2 \in U(H)$  with the same color, i.e.  $c(u_1) = c(u_2) = k$ . WLOG, when we color vertices with the  $k$ -th color, we first set  $c(u_1) \leftarrow k$ . Then we will have  $Cnt[u_2] \geq q$  since  $\{v_1, v_2, \dots, v_q\}$  are  $u_2$ 's neighbors and they all have neighbor  $u_1$  such that  $c(u_1) = k$ . So  $u_2$  cannot be colored in this round. This contradicts the assumption.  $\square$

Randomization is a classical way to produce a solution in graph coloring. We have conducted an additional experiment in Section 6.1 using three different ordering strategies: ascending order, descending order, and random order. The results indicate that there are no significant differences in efficiency or error between the three strategies. There is also another ablation study on the effectiveness



**Figure 3: An example of edges' directions in a  $(6, 3)$ -broom, where edges with different directions are represented by Up and Down in  $Dir$  and colored in red and blue, respectively.**

---

**Algorithm 1: CBS: Coloring**

---

**Input:**  $G$ : a bipartite graph;  $p, q$ : two parameters.  
**Output:**  $G'$ : a colored bipartite graph.

```

1 CalcColor( $G, U(G), p, q$ );
2 CalcColor( $G, V(G), q, p$ );
3  $G' \leftarrow G$ ;
4 return  $G'$ ;
5 Procedure CalcColor( $G, S, p, q$ )
6   Initialize  $\kappa \leftarrow 1$ ;
7   while  $|S| > 0$  do
8     Initialize an array  $Cnt$  with zeros;
9      $S_{next} \leftarrow \emptyset$ ;
10    /* enumerate vertices in a random order */
11    for  $u \in S$  do
12       $c(u) \leftarrow \kappa$ ;
13      Update  $Cnt$ ;
14      if  $\exists x, Cnt[x] \geq q$  and  $c(x) = \kappa$  then
15        Revert  $Cnt$ 's change;
16         $S_{next} \leftarrow S_{next} \cup \{u\}$ ;
17     $\kappa \leftarrow \kappa + 1$ ;
18     $S \leftarrow S_{next}$ ;

```

---

of adapting the coloring process in counting bicliques. As shown in Section 6.1, with a slight increase in runtime, the coloring process enhances the accuracy in turn.

## 4.2 Pattern Counting

After coloring, we enhance the motif we are counting with a condition of distinct colors. However, it is still hard to count bicliques directly. Instead, we count the number of  $(p, q)$ -brooms. Since we assign a concrete integer for each color in Algorithm 1 for any  $(p, q)$ -biclique, to avoid overcounting, we assume its vertex ordering of  $U'$  and  $V'$  is following the color from small to large. With this exact ordering, by Definition 4.1, there is exactly one  $(p, q)$ -broom subgraph in each  $(p, q)$ -biclique. Note that the color of these  $(p, q)$ -brooms is also unique, which is convenient for counting. Algorithm 2 is able to compute the number of  $(p, q)$ -brooms of this type. We then show how to use this value to approximate the biclique count.

The whole process is a dynamic programming approach, where its state is represented by the total number of edges accumulated so far and the last edge being tracked. To well define this dynamic programming state, we need to specify a few more properties in this type of  $(p, q)$ -broom:

**A.** Any  $(p, q)$ -broom  $H$  is connected and contains exactly  $p + q - 1$  edges. That is to say, it is actually a tree.

**B.** If we sort all edges  $(u, v) \in E(H)$  by the increasing order of the tuple  $(c(u), c(v))$ , there is only three types of edges:

- (1) The edge with the maximum tuple.
- (2) Edges sharing an endpoint in  $U(H)$  with the next edge.
- (3) Edges sharing an endpoint in  $V(H)$  with the next edge.

Here, for sorting, the comparison between two tuples  $(c(u), c(v))$  and  $(c(u'), c(v'))$  is first done by verifying if  $c(u) < c(u')$ . If  $c(u) = c(u')$ , we then check if  $c(v) < c(v')$ . Therefore, the maximum tuple refers to the last element after sorting in this manner. If For (2) and (3), we signal them with a "direction" as either Up or Down. For example, as shown in the  $(6, 3)$ -broom in Figure 3,  $(u_1, v_1)$  is with the smallest partial order. Its direction is Down since it shares the same endpoint  $v_1$  with the next edge  $(u_2, v_1)$ . Similarly,  $(u_3, v_2)$ ,  $(u_4, v_2)$ ,  $(u_5, v_3)$  are also with direction Down, highlighted with blue arrows.  $(u_3, v_1)$ ,  $(u_5, v_2)$  are with direction Up, highlighted with reversed red arrows.  $(u_6, v_3)$  has no direction since it is the edge with the maximum tuple.

After fixing the structure and order of edges by **B**, we can simply use the number of edges and the last edge to represent the current structure of the growing broom. The rest is enumeration and direct transition. The details are as follows.

To begin with, we sort  $U(G), V(G)$  by the increasing order of colors  $c(u), c(v)$  (Line 1), and sort  $E(G)$  by the increasing order of the color tuple  $(c(u), c(v))$  (Line 2).

Let  $Dp[len][pre]$  denote the number of possible ways to build a growing broom with the current number of edges  $len$  and the last edge  $pre$ , which follows the fixed structure and order of edges by **B**. The objective is to correctly computed all states for  $1 \leq len \leq p + q - 1$  (there are  $p + q - 1$  edges in total besides the first one) and  $pre \in E(G)$ . We use a 2D table to store the results, initialized to 0 (Line 3). The base case contains the case where each  $(u, v)$  acts as the first edge in the broom: we set  $Dp[1][u, v]$  to be 1 (Line 4, 4). We iterate this dimension from small to large to accumulate the growing broom from 1.. $t$ -th edges to 1.. $t + 1$ -th edges (Line 5).

We use  $Dir[t]$  to denote the  $t$ -th edge's direction in the  $(p, q)$ -broom. When  $Dir[t]$  is Up (Line 6), the  $t + 1$ -th edge and the  $t$ -th edge share a common point in  $U$ . Therefore, for each edge  $(u, v)$  (Line 7),  $Dp[t + 1][u, v]$  should accumulate the number of ways coming from any  $Dp[t][u, w]$  such that  $w$  is in  $N(u, G)$  and  $c(w) < c(v)$  (Line 8). Similarly, when  $Dir[t]$  is Down (Line 9), the  $t + 1$ -th edge and the  $t$ -th edge share a common point in  $V$ . In this case,  $Dp[t + 1][u, v]$  should accumulate the number of ways coming from any  $Dp[t][w, v]$  such that  $w$  is in  $N(v, G)$  and  $c(w) < c(u)$  (Line 10, 11). After all computations, we return  $B$  as the sum of all  $Dp[p + q - 1][\cdot]$  and the  $Dp$  table as the results (Line 12).

By using the prefix sum technique, for each iteration, we can finish all transition computations in  $O(|E|)$ . There are  $O(p + q)$  rounds. Therefore, the total time complexity is  $O((p + q)|E|)$ .

## 4.3 Approximate Counting via Sampling

After `CountingIndex`( $G, p, q$ ), we have accumulated  $B$   $(p, q)$ -brooms. To approximate the quantity relation between our computed  $(p, q)$ -brooms and  $(p, q)$ -bicliques, we can do the following sampling:

- (1) Set two counters  $cnt_{broom}$  and  $cnt_{biclique}$ .
- (2) Uniformly sample a  $(p, q)$ -broom  $P$  from all  $B$  ones and increase  $cnt_{broom}$  by 1.

---

**Algorithm 2:** CBS: CountingIndex

---

**Input:**  $G$ : a colored bipartite graph;  $p, q$ : two parameters.  
**Output:**  $Dp$ : a 2D array;  $B$ : the total number of  $(p, q)$ -brooms in  $G$ .

- 1 Sort  $U(G), V(G)$  by the increasing order of  $c(u), c(v)$ ;
- 2 Sort  $E(G)$  by the increasing order of  $(c(u), c(v))$ ;
- 3 Initialize  $Dp$  with zeros;
- 4 **for**  $(u, v) \in E(G)$  **do**  $Dp[1][(u, v)] \leftarrow 1$  ;
- 5 **for**  $t \leftarrow 1$  **to**  $p + q - 2$  **do**
- 6     **if**  $Dir[t] = Up$  **then**
- 7         **for**  $(u, v) \in E(G)$  **do**
- 8              $Dp[t+1][(u, v)] \leftarrow \sum_{w \in N(u, G), c(w) < c(v)} Dp[t][(u, w)]$ ;
- 9         **else**
- 10             **for**  $(u, v) \in E(G)$  **do**
- 11                  $Dp[t+1][(u, v)] \leftarrow \sum_{w \in N(v, G), c(w) < c(u)} Dp[t][(w, v)]$ ;
- 12  $B \leftarrow \sum_{(u, v) \in E(G)} Dp[p+q-1][(u, v)]$ ;
- 13 **return**  $Dp, B$ ;

---

(3) If the induced subgraph by  $U(P)$  and  $V(P)$  is a  $(p, q)$ -biclique, increase  $cnt_{biclique}$  by 1.

After applying this sampling process sufficiently many times, we can roughly approximate the number of  $(p, q)$ -bicliques by  $cnt_{biclique}/cnt_{broom} \times B$ . However, such a sampling process is too inefficient. We now propose Algorithm 3 to accelerate this process, utilizing the computed dynamic programming results from CountingIndex( $G, p, q$ ). Recall that  $Dp[t][(u, v)]$  represents the number of ways to build a growing  $(p, q)$ -broom til the  $t$ -th edge, which is  $(u, v)$ . To begin with, we start by sampling the last edge of the  $(p, q)$ -broom, denoted as  $(u_{last}, v_{last})$  (Line 1). Here, the sampling should follow the weight distribution of  $\{Dp[p+q-1][\cdot]\}$ , indicating how many  $(p, q)$ -brooms end with each edge. The following process is similar to reverting the dynamic programming process. Instead of building the  $(p, q)$ -broom from the 1-st edge to the  $(p+q-1)$ -th edge, we do it reversely. We use  $(u_{cur}, v_{cur}), U', V'$  to denote the current growing  $(p, q)$ -broom. They are initialized as  $(u_{last}, v_{last}), \{u_{last}\}, \{v_{last}\}$  respectively (Line 2, 3). Specifically, since  $(u_{last}, v_{last})$  is the  $(p+q-1)$ -th edge, we are now adding the  $(p+q-2)$ -th edge til the 1-th edge gradually (Line 5).

Assume we are processing the  $i$ -th edge now, and we know the  $(i+1)$ -th edge is  $(u_{cur}, v_{cur})$ . We initialize an edge set  $S$  to store the candidate edge for the  $i$ -th edge. Based on the direction signaling, we can know whether it should share the common endpoint in  $U$  or  $V$ . If  $Dir[i]$  is Up (Line 7), then we should find all  $w \in N(u_{cur}, G)$  such that  $c(w) < c(v_{cur})$  and  $U' \subseteq N(w, G)$ . We assign  $S$  to be  $\{(u_{cur}, w)\}$  (Line 8). The first condition is for the color order, and the second condition indicates that adding  $w$  can guarantee that the induced subgraph is biclique. Similarly, if  $Dir[i]$  is Down (Line 9), we find all  $w \in N(v_{cur}, G)$  such that  $c(w) < c(u_{cur})$  and  $V' \subseteq N(w, G)$ . Then we assign  $S$  to be  $\{(w, v_{cur})\}$  (Line 10).

The idea to accelerate the sampling is, by building the  $(p, q)$ -broom, we strictly guarantee that it will correspond to a  $(p, q)$ -biclique while keeping track of the probability of sampling out this  $(p, q)$ -broom, which can be computed through the dynamic programming table. In the beginning, we initialize  $ans$  to be 1 (Line 4). Then, for the  $i$ -th edge, the probability contribution of sampling

---

**Algorithm 3:** CBS: Sampling

---

**Input:**  $G$ : a colored bipartite graph;  $p, q$ : two parameters;  
 $Dp$ : a 2D array;  $B$ : the total number of  $(p, q)$ -brooms in  $G$ .  
**Output:**  $ans$ : estimate number of  $(p, q)$ -clique in  $G$ .

- 1 Sample the last edge  $(u_{last}, v_{last})$  in  $E(G)$  following the weight distribution of brooms;
- 2  $(u_{cur}, v_{cur}) \leftarrow (u_{last}, v_{last})$ ;
- 3  $U' \leftarrow \{u_{last}\}, V' \leftarrow \{v_{last}\}$ ;
- 4  $ans \leftarrow 1$ ;
- 5 **for**  $i \leftarrow p + q - 2$  **to** 1 **do**
- 6     Initialize set  $S \leftarrow \emptyset$ ;
- 7     **if**  $Dir[i] = Up$  **then**
- 8          $S \leftarrow \{(u_{cur}, w) \mid w \in N(u_{cur}, G), c(w) < c(v_{cur}), U' \subseteq N(w, G)\}$ ;
- 9     **else**
- 10          $S \leftarrow \{(w, v_{cur}) \mid w \in N(v_{cur}, G), c(w) < c(u_{cur}), V' \subseteq N(w, G)\}$ ;
- 11      $ans \leftarrow ans \times \frac{\sum_{(u, v) \in S} Dp[i][(u, v)]}{Dp[i+1][(u_{cur}, v_{cur})]}$ ;
- 12     **if**  $ans = 0$  **then break** ;
- 13     Sample the next edge  $(u_{next}, v_{next})$  in  $S$  following the weight distribution of growing brooms;
- 14      $(u_{cur}, v_{cur}) \leftarrow (u_{next}, v_{next})$ ;
- 15      $U' \leftarrow U' \cup \{u_{next}\}, V' \leftarrow V' \cup \{v_{next}\}$ ;
- 16 **return**  $ans \times B$ ;

---

it out from  $S$  should be the number of the growing broom ends with  $(u, v) \in S$  divided by the total number of possible brooms at this step, which is  $\frac{\sum_{(u, v) \in S} Dp[i][(u, v)]}{Dp[i+1][(u_{cur}, v_{cur})]}$ . We multiply  $ans$  by this value (Line 11). If  $ans$  becomes 0, then there is no possible  $(p, q)$ -biclique from the current sampled  $(p, q)$ -broom. We return with 0 in this case (Line 12). Now we are ready to sample the  $i$ -th edge from  $S$  (Line 13). We do so following the normalized distribution of possible brooms from this step ( $\{Dp[i][\cdot]\}$ ). To continue the process, we assign  $(u_{cur}, v_{cur}), U', V'$  to  $(u_{next}, v_{next}), U' \cup \{u_{next}\}, V' \cup \{v_{next}\}$ , respectively (Line 14, 15).

In the end, we successfully sample a  $(p, q)$ -broom that corresponds to one  $(p, q)$ -biclique and compute the probability of sampling it out. We return this probability by multiplying  $B$  as the approximate count of  $(p, q)$ -bicliques (Line 16). It is not hard to see that the time complexity is dominated by the cost of updating  $S$ , which requires enumerating all neighbors. The loop only lasts  $O(p+q)$ . Therefore, the time complexity is  $O((p+q) \times \Delta)$ , where  $\Delta$  denotes the maximal degree in  $G$ .

#### 4.4 Overall Algorithm

We provide Algorithm 4 as a black box to call all three subroutines properly and output the estimated answer of the biclique count. We start by coloring (Line 1) and pattern counting (Line 2). Then by the input sampling size parameter  $T$ , we repeatedly call Sampling( $G, p, q, Dp, B$ ) and aggregate the return to  $\hat{C}$ . The output approximate  $(p, q)$ -biclique count is the average of all  $T$  attempts (Line 7). Note that as a common trick, we will execute core-reduction for the original graph [32]: For any query  $(p, q)$ , we split the graph and reduce the query pair to  $(p-1, q)$ . Now, we give a theoretical analysis of CBS in the following.

---

**Algorithm 4: CBS: Main**


---

**Input:**  $G$ : a bipartite graph;  $p, q$ : two parameters;  $T$ : sampling times.

**Output:**  $\hat{C}$ : estimate number of  $(p, q)$ -clique.

- 1  $G \leftarrow \text{Coloring}(G, p, q)$ ;
  - 2  $Dp, B \leftarrow \text{CountingIndex}(G, p, q)$ ;
  - 3  $\hat{C} \leftarrow 0$ ;
  - 4 **for**  $i \leftarrow 1$  **to**  $T$  **do**
  - 5      $\hat{C} \leftarrow \hat{C} + \text{Sampling}(G, p, q, Dp, B)$ ;
  - 6  $\hat{C} \leftarrow \hat{C}/T$ ;
  - 7 **return**  $\hat{C}$ ;
- 

**4.4.1 Unbiasedness.** We now show that the CBS method is unbiased. Let  $(u_{cur}, v_{cur})$ ,  $U'$ , and  $V'$  represent the current growing  $(p, q)$ -broom. For clarity, we introduce the following definitions:

- $\mathcal{F}_{U', V'}$ : The total value to be multiplied into  $ans$ , defined as the product of the last  $i = (p + q) - (|U'| + |V'|)$  fractions.
- $\mathcal{B}_{U', V'}$ : The number of  $(p, q)$ -brooms  $H$  such that  $U', V'$  are the maximal tuples of  $U(H)$  and  $V(H)$ .
- $C_{U', V'}$ : The number of  $(p, q)$ -bicliques  $H'$  such that  $U', V'$  are the maximal tuples of  $U(H')$  and  $V(H')$ .

After that, we prove the following lemma by mathematical induction.

**LEMMA 4.3.** *For any growing  $(p, q)$ -broom  $U', V'$ ,  $\mathbb{E}[\mathcal{F}_{U', V'}] = C_{U', V'}/\mathcal{B}_{U', V'}$ .*

**PROOF.** (a) When  $|U'| + |V'| = p + q$ , we have  $\mathcal{B}_{U', V'} = 1$ . In this case,  $\mathcal{F}_{U', V'} = C_{U', V'} = 1$  if  $U', V'$  forms a  $(p, q)$ -biclique; otherwise  $\mathcal{F}_{U', V'} = C_{U', V'} = 0$ . Hence,  $\mathbb{E}[\mathcal{F}_{U', V'}] = C_{U', V'}/\mathcal{B}_{U', V'}$  trivially holds in this case.

(b) Suppose that the lemma holds for all growing  $(p, q)$ -brooms with  $|U'| + |V'| = k + 1$ . For any broom with  $|U'| + |V'| = k$ , let  $P = \frac{\sum_{(u,v) \in S} Dp[i][u,v]}{Dp[i+1][u_{cur}, v_{cur}]}$ . We then have

$$\mathbb{E}[\mathcal{F}_{U', V'}] = P \sum_{(u,v) \in S} \frac{Dp[i][u,v]}{\sum_{(u,v) \in S} Dp[i][u,v]} \mathbb{E}[\mathcal{F}_{U' \cup \{u\}, V' \cup \{v\}}].$$

Since  $u_{cur} = u$  or  $v_{cur} = v$  always holds, we have that  $|U' \cup \{u\}| + |V' \cup \{v\}| = |U'| + |V'| + 1 = k + 1$ . According to the definition of  $Dp$  and  $\mathcal{B}$ , We can derive that  $Dp[i][u,v] = \mathcal{B}_{U' \cup \{u\}, V' \cup \{v\}}$ , and thus

$$\begin{aligned} \mathbb{E}[\mathcal{F}_{U', V'}] &= \frac{\sum_{(u,v) \in S} Dp[i][u,v] \frac{C_{U' \cup \{u\}, V' \cup \{v\}}}{\mathcal{B}_{U' \cup \{u\}, V' \cup \{v\}}}}{Dp[i+1][u_{cur}, v_{cur}]} \\ &= \frac{\sum_{(u,v) \in S} C_{U' \cup \{u\}, V' \cup \{v\}}}{\mathcal{B}_{U', V'}} = \frac{C_{U', V'}}{\mathcal{B}_{U', V'}}, \end{aligned}$$

where the last equation is because  $C_{U' \cup \{u\}, V' \cup \{v\}} = 0$  for all pairs of  $(u, v) \notin S$ .  $\square$

Then we can derive the following theorem:

**THEOREM 4.4.** *Let  $ans_i$  be the value of  $ans$  in the  $i$ -th sampling. Let  $\hat{C} = \frac{1}{T} \sum_{i=1}^T ans_i \times B$ . Then  $\hat{C}$  is an unbiased estimator of the number of  $(p, q)$ -bicliques in  $G$ .*

**PROOF.** We first have the following induction by Lemma 4.3:

$$\begin{aligned} &\mathbb{E}[ans_i \times B] \\ &= B \times \sum_{(u,v) \in E(G)} \frac{Dp[p+q-1][u,v]}{\sum Dp[p+q-1][u,v]} \mathbb{E}[\mathcal{F}_{\{u\}, \{v\}}] \\ &= \sum_{(u,v) \in E(G)} Dp[p+q-1][u,v] \times \mathbb{E}[\mathcal{F}_{\{u\}, \{v\}}] \\ &= \sum_{(u,v) \in E(G)} C_{\{u\}, \{v\}} = C_{0,0}. \end{aligned}$$

Then, we can derive that:  $\mathbb{E}[\hat{C}] = \frac{1}{T} \sum_{i=1}^T \mathbb{E}[ans_i \times B] = C_{0,0}$ .

Note that  $C_{0,0}$  is the number of  $(p, q)$ -bicliques in  $G$  according to the definition. Therefore,  $\hat{C}$  is an unbiased estimator of the number of  $(p, q)$ -bicliques in  $G$ .  $\square$

Based on Theorem 4.4, we have obtained an unbiased estimator  $\hat{C}$  of the number of  $(p, q)$ -bicliques in Algorithm 4.

**4.4.2 Error Analysis.** we now analyze the estimation error of our sampling algorithms. Our analysis relies on the classic Hoeffding's inequalities, which are shown below.

**LEMMA 4.5 (HOEFFDING'S INEQUALITY, [8, 11]).** *For the random variables  $X_i \in [0, M], 1 \leq i \leq n$ , we let  $X = \sum_{i=1}^n X_i$ . Then for  $\epsilon > 0$ , we have*

$$\begin{aligned} \Pr(X \geq (1 + \epsilon)\mathbb{E}[X]) &\leq \exp\left(-\frac{2\epsilon^2\mathbb{E}[X]^2}{nZ^2}\right), \\ \Pr(X \leq (1 - \epsilon)\mathbb{E}[X]) &\leq \exp\left(-\frac{2\epsilon^2\mathbb{E}[X]^2}{nZ^2}\right). \end{aligned}$$

Based on Lemma 4.5, we can derive the estimation error of our algorithms as shown in the following theorem.

**THEOREM 4.6.** *Let  $C, B$  be the number of  $(p, q)$ -bicliques,  $(p, q)$ -brooms, respectively. Let  $\hat{C} = \frac{1}{T} \sum_{i=1}^T ans_i \times B$  denote the estimated number of  $(p, q)$ -bicliques, where  $ans_i$  is the return result in the  $i$ -th sampling. Then,  $\hat{C}$  is a  $(1 + \epsilon)$  approximation of  $C$  with probability  $(1 - \alpha)$  if  $T \geq \frac{B^2}{2\epsilon^2 C^2} \ln\left(\frac{2}{\alpha}\right)$ .*

**PROOF.** We can derive  $ans_i \times B \in [0, B]$  since the values multiplied to  $ans$  are all probabilities in  $[0, 1]$ . And we have  $\mathbb{E}[\hat{C}T] = CT$  based on Theorem 4.4. Given a positive value  $\epsilon$ , applying Lemma 4.5 by plugging in  $n = T, X_i = ans_i \times B, X = \sum_{i=1}^n X_i = \hat{C}T$ , we then have

$$\begin{aligned} \Pr(X \geq (1 + \epsilon)\mathbb{E}[X]) &= \Pr(\hat{C}T \geq (1 + \epsilon)CT) \leq \exp\left(-\frac{2\epsilon^2(CT)^2}{TB^2}\right), \\ \Pr(X \leq (1 - \epsilon)\mathbb{E}[X]) &= \Pr(\hat{C}T \leq (1 - \epsilon)CT) \leq \exp\left(-\frac{2\epsilon^2(CT)^2}{TB^2}\right). \end{aligned}$$

Further, we have  $\Pr\left(\frac{|\hat{C}-C|}{C} \geq \epsilon\right) \leq 2 \exp\left(-\frac{2\epsilon^2 C^2 T}{B^2}\right)$ .

Let  $2 \exp\left(-\frac{2\epsilon^2 C^2 T}{B^2}\right) \leq \alpha$ , we can derive that  $T \geq \frac{B^2}{2\epsilon^2 C^2} \ln\left(\frac{2}{\alpha}\right)$ .  $\square$

**4.4.3 Time and Space Complexity.** The time complexity for each of the three parts in CBS are already discussed separately. We now give the following theorem as a summary.

Table 1: Datasets used in experiments.

Graphs (Abbr.)	Category	U	V	E	max $\kappa$
github (GH)	Authorship	56,519	120,867	440,237	508
StackOF (SO)	Rating	545,195	96,678	1,301,942	290
Twitter (Wut)	Interaction	175,214	530,418	1,890,661	2933
IMDB (IMDB)	Affiliation	685,568	186,414	2,715,604	158
Actor2 (Actor2)	Affiliation	303,617	896,302	3,782,463	189
Amazon (AR)	Rating	2,146,057	1,230,915	5,743,258	155
DBLP (DBLP)	Authorship	1,953,085	5,624,219	12,282,059	126
Epinions (ER)	Rating	120,492	755,760	13,668,320	13200
Wikipedia-edits-de (DE)	Authorship	1,025,084	5,910,432	129,885,939	118356

THEOREM 4.7. *The time complexity of CBS is  $O(\kappa|E| + (p + q)|E| + T \cdot (p + q) \cdot \Delta)$ , where  $\Delta$  denotes the maximal degree in  $G$ . Specifically, to reach a  $(1 + \epsilon)$ -approximation with probability  $(1 - \alpha)$ , we require  $T \geq \frac{B^2}{2\epsilon^2 C^2} \ln(\frac{2}{\alpha})$  where  $C, B$  denotes the number of  $(p, q)$ -bicliques and  $(p, q)$ -brooms.*

PROOF. The first two terms follow from the analysis in Section 4.1 and Section 4.2. The third term follows from the fact that we need to conduct the sampling process (Algorithm 3) for  $T$  times, and each time costs  $(p + q) \cdot \Delta$ . The approximation guarantee comes directly from Theorem 4.6.  $\square$

We now provide the memory usage analysis as follows. We have also conducted an empirical experiment on the memory consumption in Section 6.4.

THEOREM 4.8. *The space complexity of CBS is  $O(|U| + |V| + (p + q) \cdot |E|)$ .*

PROOF. In Algorithm 1, since we need to maintain  $Cnt$  and  $S_{next}$  which is at most  $|U| + |V|$ , the memory usage is  $O(|U| + |V|)$ . In Algorithm 2, the memory usage is dominated by the size of the DP table, thus  $O((p + q) \cdot |E|)$ . Algorithm 3 is for sampling a  $(p, q)$ -biclique, and  $|S|$  is at most  $\max \deg(u)$ , thus the memory usage is  $O(\max \deg(u))$ . In all, the total space complexity is  $O(|U| + |V| + (p + q) \cdot |E|)$ .  $\square$

The EP/Zz++ uses the heuristic framework to combine EPivoter (for exact counting on small portion of graphs) and ZigZag++. In the analysis of ZigZag++, its time and space complexity is  $O((p + q)|E| + (p + q)^2|V| + T \cdot \Delta \cdot (p + q)^2)$  and  $O((p + q)|E|)$ . Our memory usage is similar to each other. If we consider  $T$  and  $\kappa$  as a constant, our time complexity has some advantage as we don't have the term of  $(p + q)^2$ . However, it is obvious that both our algorithms' performance are highly determined by the size of  $T$ .

From Theorem 4.6, the sample size  $T$  is mainly determined by  $(\frac{B}{C})^2$ . That is to say, we need a larger sample size  $T$  to ensure high accuracy with a larger  $(\frac{B}{C})^2$ , where  $B$  is the number of  $(p, q)$ -brooms,  $C$  is the number of  $(p, q)$ -bicliques. As shown in Section 6.3,  $(\frac{B}{C})^2$  is usually small comparing to the dominating term of EP/Zz++  $(\frac{Z}{\rho^2})$  [34]. Therefore, our algorithm generally does not need a large  $T$  to achieve good accuracy in real-world datasets. It can also be concluded that CBS guarantees an equivalent error ratio with significantly fewer sampling times and higher precision with the same sampling times compared to EP/Zz++.

## 5 EVALUATION

### 5.1 Experimental Setting

*Datasets.* We use nine real datasets from different domains, which are available at SNAP [23], Laboratory of Web Algorithmics [22], and Konect [17]. Table 1 shows the statistics of these graphs.

*Methods.* We compare our method with the baselines from highly related works, summarized as follows:

- BCList++ [32]: the biclique listing-based algorithm, which is based on the Bron-Kerbosch algorithm [5]. The key idea of BCList++ is to iteratively enumerate all  $(p, q)$ -bicliques containing each vertex through a node expansion process. Specifically, it employs an ordering-based search paradigm, where for each vertex, only its higher-order neighbors are considered during enumeration. Additionally, a graph reduction technique is applied to reduce the search space before biclique counting.
- EPivoter [34]: the state-of-the-art algorithm for exact biclique counting, which relies on the edge-based pivot technique. In EPivoter, an edge-based search framework is introduced. Unlike BCList++, it iteratively selects edges from the candidate set (i.e., the edge set used for biclique expansion) to expand the current biclique. Specifically, during the enumeration process, in each branch, it first select one edge as the pivot edge, and based on this edge, vertices can be grouped into four disjoint groups. By storing the entire enumeration tree along with the four vertex sets, biclique counting can then be efficiently performed using combination counting.
- EP/Zz++ [34]: the state-of-the-art algorithm for approximate counting. It first partitions the graph into two regions: a dense region (a subgraph containing high-degree vertices) and a sparse region (a subgraph containing low-degree vertices). For the sparse region, EP/Zz++ utilizes EPivoter for exact counting, while for the dense region, it proposes a zigzag path-based sampling algorithm for approximate counting. Specifically, it leverages the fact that a  $(p, q)$ -biclique must contain a fixed number of  $(\min\{p, q\})$ -zigzag paths. Based on this property, the algorithm first counts the number of  $h$ -zigzag paths and then samples  $T$  such paths, where  $h = \min\{p, q\}$ . Finally, it estimates the number of  $h$ -bicliques based on their proportion with the sampled paths.
- CBS: our proposed approximation algorithm (Section 4).

Notice that we compare EP/Zz++ instead of EP/Zz [34], since the former one is a better version of the latter in terms of efficiency and accuracy. For completeness, we also conduct some initial experiment on EP/Zz. Specifically, on the two largest datasets, ER and DE, EP/Zz failed to produce results within  $10^5$  seconds for any values of  $p$  and  $q$  ( $3 \leq p, q \leq 9$ ). Thus, it is not applicable to scalable datasets in practice. Thus, we focus our evaluation primarily on EP/Zz++ in our main experiments.

There is also another potential baseline introduced in [35]. The method takes over 1000 seconds to run on the DBLP dataset, while our algorithm requires only around 1 second. We have performed a primary initial evaluation based on the conclusions reported in the article. They claim to outperform BCList++ and EPivoter by two orders of magnitude in runtime. However, both baselines fail to return results on the ER and DE datasets within  $10^5$  seconds. Even based on their own estimation, their method remains significantly

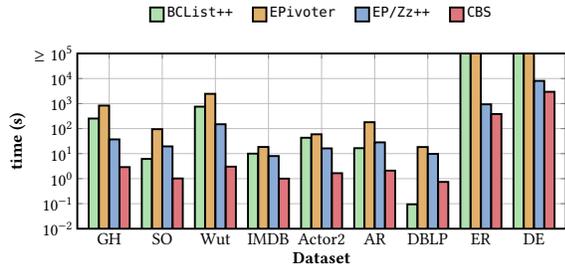


Figure 4: Average runtime of different biclique counting algorithms for all  $3 \leq p, q \leq 9$ .

slower than ours in practice. As a result, we do not include it as a baseline in our empirical study. We implement all the algorithms in C++ and run experiments on a machine having an Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz and 512GB of memory, with Ubuntu installed.

**Parameter Settings.** In our experiments, we following the existing work [34] setting  $T=10^5$  as default values. We define the estimation error as  $\frac{|\hat{C}-C|}{C}$ , where  $C$  denotes the exact count of bicliques, and  $\hat{C}$  represents its approximated value. To ensure the reliability of our results, we run each approximation algorithm 10 times, with the reported error representing the average across these executions. For some datasets, the exact count of bicliques cannot be computed (e.g., ER and DE), we evaluate the estimation error differently by using the average approximated biclique count from 10 runs as the exact count (i.e.,  $C$ ).

## 5.2 Overall Comparison Results

In this section, we compare CBS with three competitors (introduced in Section 5.1), w.r.t. overall efficiency, accuracy, and the effect of  $p$  and  $q$ , and the number of samples to demonstrate the superior of our algorithm.

**1. Efficiency of All Algorithms.** Figure 4 depicts the average running time of all the biclique counting algorithms on nine datasets for counting all  $3 \leq p, q \leq 9$  bicliques. We make the following observations and analysis:

- (1) Our method CBS is up to two orders of magnitude faster than all competitors, this is mainly because our algorithm has a better theoretical guarantee.
- (2) On almost all datasets, CBS is at least  $10\times$  faster than all methods, except DBLP dataset. On this dataset, BCList++ achieves the best performance, as graph reduction significantly reduces  $|U|$  and  $|V|$ , allowing it to perform efficiently without extra steps. Meanwhile, CBS still outperforms EPivoter and EP/Zz++.
- (3) On the two largest datasets, ER and DE, the exact algorithms fail to terminate in  $10^5$  seconds, while both approximate algorithms successfully complete the task. Our method, CBS, demonstrates at least 3 times faster performance compared to EP/Zz++.

Note that besides Figure 4, we choose to simply compare the sampling time with EP/Zz++ in Figure 6 and Figure 8, which is for exposing the significant improvement in the sampling efficiency.

**2. Accuracy of All Algorithms.** Figure 5 illustrates the average error rates of CBS and EP/Zz++ across all datasets for biclique sizes ranging from 3 to 9 (i.e.,  $3 \leq p, q \leq 9$ ). We can see that our algorithm

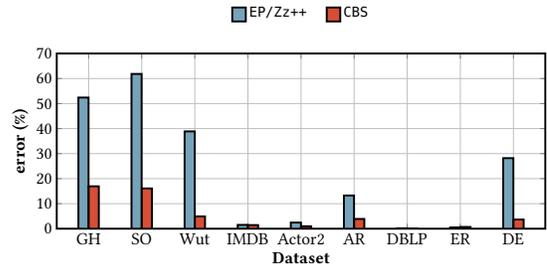


Figure 5: Average error of EP/Zz++ and CBS for all  $3 \leq p, q \leq 9$ .

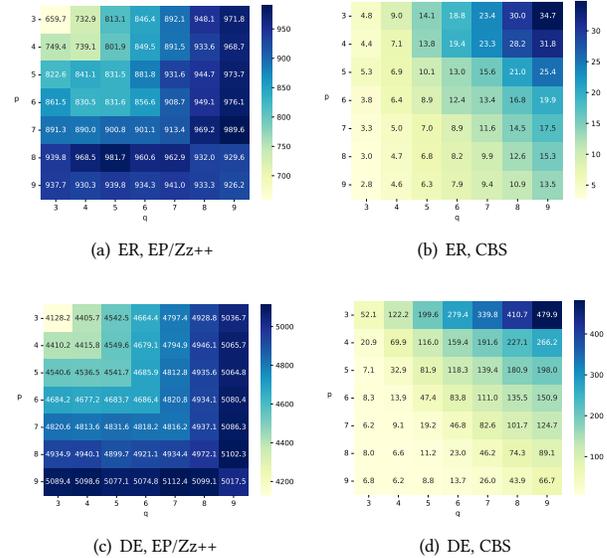
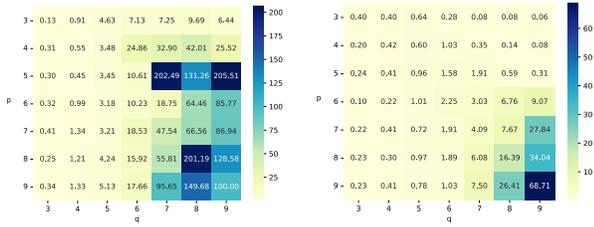


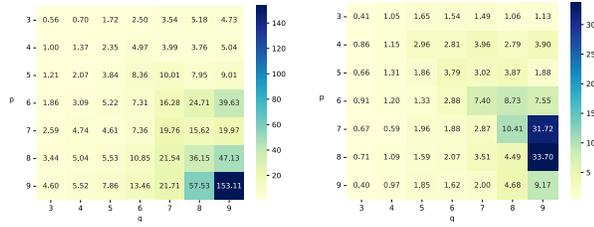
Figure 6: The heat-map of sampling time of EP/Zz++ and CBS with varying  $p$  and  $q$  (s).

demonstrates up to a  $8\times$  reduction in error compared to EP/Zz++, thanks to our carefully designed coloring scheme and  $(p, q)$ -broom-based sampling technique. Note that for the first three smaller datasets, EP/Zz++ exhibits significant errors, with at least 38.9% inaccuracy, rendering its results unreliable. In contrast, our method, CBS, maintains a maximum error rate of 16.9% under the same number of sampling rounds. Our algorithm also consistently obtains an outstanding accuracy in larger datasets. A minor exception occurs in ER that our algorithm obtains a slightly worse error ratio in only one case  $(p, q) = (4, 9)$ . To compensate for this, since our sampling time is more than 70 times faster compared to EP/Zz++ (Figure 6 (a)(b)), by simply increasing the sampling times, we can decrease the error ratio while still outperforming the baseline. Initially, we sampled  $10^5$  times. By sampling  $5 \times 10^5$ ,  $10^6$  and  $7 \times 10^6$ , our error ratio decreases to 1.06%, 1.01%, 0.34%, respectively. These results all outperform EP/Zz++ in both error ratio and sampling time.

Combining these observations with the performance results shown in Figure 4, we can conclude that our algorithm demonstrates an even greater advantage when considering the trade-off between accuracy and efficiency. This indicates that to achieve the same error rate, our algorithm would likely exhibit an even more substantial performance advantage over EP/Zz++.



(a) Wut, EP/Zz++ (b) Wut, CBS



(c) AR, EP/Zz++ (d) AR, CBS

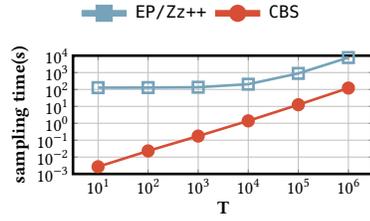
**Figure 7: The heat-map of estimation errors of EP/Zz++ and CBS with varying  $p$  and  $q$  (%).**

3. *Effect of  $p$  and  $q$ .* Figures 6 and 7 illustrate the sampling time and error rate, respectively, of CBS and EP/Zz++ across various  $p$  and  $q$  values. In each figure, rows represent  $p$  values and columns represent  $q$  values. Each cell displays the sampling time (in Figure 6) or the estimation error (in Figure 7) for counting the corresponding  $(p, q)$ -biclques. Clearly, our method, CBS, outperforms EP/Zz++ by up to two orders of magnitude in both sampling time and accuracy. For instance, with  $p = 5$  and  $q = 3$  on the ER dataset, EP/Zz++ requires 822.6 seconds for sampling, whereas our algorithm completes this stage in just 5.3 seconds.

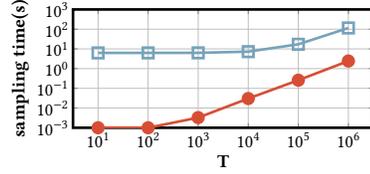
In terms of accuracy, on the Wut dataset with  $p = 5$  and  $q = 7$ , EP/Zz++ has an estimation error of 202.49, while our algorithm achieves a significantly lower error of 1.91. This observation aligns with our analysis in Section 4. As a minor observation in Figure 6, the sampling time of CBS is more affected by  $q$  than by  $p$  while EP/Zz++ does not exhibit this characteristic. This observation indeed demonstrates the difference between EP/Zz++ and our algorithm to some levels. Specifically, EP/Zz++ needs to sample a zig-zag path with  $h = \min(p, q)$  to approximate the number of  $(p, q)$ -biclques, therefore its sampling time increases in a staircase manner w.r.t  $\min(p, q)$ . On the other hand, our algorithm samples brooms, which is different from EP/Zz++, this could be why we exhibit different characteristics. In our case, one of the possible reasons is that we start "growing" a broom from the left side during the sampling process.

4. *Effect of  $T$ .* We evaluate the effect of sample numbers on sampling time and error rate. The results on three datasets are shown in Figures 8 and 9. Based on these results, we observe the following:

- (1) For sampling time, our algorithm consistently outperforms EP/Zz++, particularly when the sample numbers are relatively low (e.g., from  $10^1$  to  $10^3$ ).
- (2) Our algorithm consistently produces more accurate results than EP/Zz++, regardless of whether the sample size is high or low.

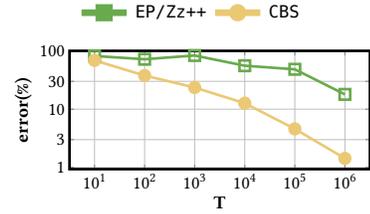


(a) ER

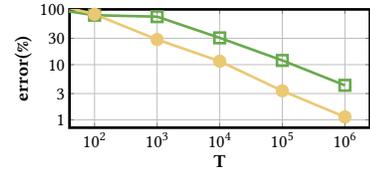


(b) AR

**Figure 8: Average sampling time of EP/Zz++ and CBS with varying  $T$ .**



(a) Wut



(b) AR

**Figure 9: Average error of EP/Zz++ and CBS with varying  $T$ .**

- (3) Even with very few samples, our algorithm achieves acceptable solutions. For instance, on the AR dataset, CBS requires only  $10^3$  samples to achieve a solution with a 30% error rate, whereas EP/Zz++ requires over  $10^4$  samples.

## 6 DETAILED ANALYSIS

In this section, to provide a better understanding of our algorithm, besides the primary evaluation in Section 5, we now extensively evaluate and analyze CBS from different angles:

- An ablation study on the effect of our coloring technique, showing that it can significantly improve accuracy without highly impacting sampling time.
- An ablation study on the coloring technique conducted using different ordering strategies.
  - A study on the time cost of each step in CBS, including the initialization and sampling time, compared to the baseline EP/Zz++.
  - A study on the statistical of the hyper-parameters that determines the error ratio, which provides evidence that our algorithm CBS consistently requires fewer samples than EP/Zz++.

$T$	AR				Wut			
	BS	CBS	ACBS	DCBS	BS	CBS	ACBS	DCBS
$10^2$	72.75	81.63	98.50	78.99	60.60	37.70	87.07	33.91
$10^3$	43.33	28.47	32.44	32.43	22.03	23.50	24.48	21.16
$10^4$	13.08	11.54	11.55	10.62	13.51	12.64	13.41	12.89
$10^5$	4.37	3.35	3.92	3.40	5.11	4.59	4.44	3.98
$10^6$	1.35	1.13	1.19	1.06	1.71	1.43	1.67	1.42

Figure 10: Average error of different coloring techniques in AR and Wut with varying  $T$ .

- A study of the memory consumption of our algorithm.
- An additional scalability experiment conducted by randomly sampling different percentages of vertices in the same graph.
- An additional efficiency experiment with small  $p, q$  aligning with the empirical study in [34].

### 6.1 Ablation Study on the Coloring Technique

We conduct an ablation study on the effect of our coloring technique. To begin with, in order to evaluate the effect of our coloring technique, we design a new variant of CBS by removing the vertex coloring step, denoted by BS. Then, note that in our proposed coloring methods, we decide to traverse the neighbors in a random order manner. A question is raised that if we follow different order of neighbors, would the performance be significantly affect? To answer this question, we design three more variants of CBS by using two different ordering strategies in coloring: the ascending order and descending order of degrees, denoting as ACBS and DCBS.

To compare the inaccuracy, we first run these variants on the two datasets AR and Wut with varying parameters of  $T$ , and report the results in Figure 10. Compared to BS (without the coloring), CBS generally achieves lower error with the same number of samples, aligning with our previous analysis. Although the coloring technique initially slightly increases the sampling time, the gap diminishes as the number of sampling iterations increases. This demonstrates that in high-accuracy scenarios, our coloring technique can significantly enhance accuracy without significantly impacting the sampling time. Compared to ACBS and DCBS, there is no significant difference in error in the three strategies. This confirms that our algorithm is robust to the choice of coloring order.

In Figure 11, we run all four algorithms on the dataset ER with varying values of  $T$  to compare their efficiency. We can see that without the coloring (BS), the total runtime does not decrease much. On the other hand, the average runtime of CBS, ACBS, and DCBS is similar. This again confirms the robustness of our coloring method. Combining with the accuracy result, we can see that the coloring method is effective, as with a slight increase in runtime, the accuracy is enhanced. Overall, our ablation study demonstrates that the coloring process is vital in counting bicliques.

### 6.2 Time Cost of Different Stages

The two approximation algorithms, CBS and EP/Zz++, both require an initialization stage to precompute some auxiliary data for sampling. Specifically, in CBS, we need to assign a color number for each vertex (i.e., coloring) and count the number of  $(p, q)$ -brooms in bipartite graph, while in EP/Zz++, it needs to count the number of  $h$ -zigzag paths bipartite graph, where  $h = \min\{p, q\}$ . In Figure 12, we report the initializing and sampling times for these algorithms

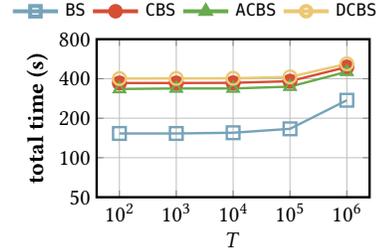


Figure 11: Average runtime of different coloring techniques in ER with varying  $T$ .

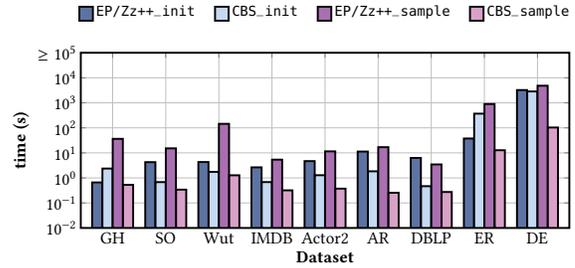


Figure 12: Average initializing and sampling time of EP/Zz++ and CBS for all  $3 \leq p, q \leq 9$  ( $T = 10^5$ ).

across all datasets. We observe that on more than half of the datasets, our algorithm requires less initializing and sampling time, which indicates that the number of  $(p, q)$ -brooms in the bipartite graph is typically less than  $h$ -zigzag paths (using in EP/Zz++). In terms of sampling time, our algorithm is up to two orders of magnitude faster than EP/Zz++. While on two datasets, GH and ER, EP/Zz++ is 10 times faster in initializing, CBS achieves 100 times lower sampling time, making the total time (i.e., initializing plus sampling) of our algorithm still highly efficient.

### 6.3 Statistical of the Hyper-parameters

Recall that in CBS and EP/Zz++, when  $\epsilon$  is fixed, their required sample sizes are proportional to  $\frac{B^2}{C^2}$  and  $\frac{Z^2}{\rho^2}$ , respectively. Therefore, the size of these two values in the graph data determines the efficiency of the two algorithms. We conduct an experiment to compare these two values in real-world data. Specifically, as shown in Table 2, we report the values of  $\frac{B^2}{C^2}$  and  $\frac{Z^2}{\rho^2}$  on the Amazon dataset for varying  $p$  and  $q$ . Similar trends are observed across other datasets. Based on Table 2, the value of  $\frac{B^2}{C^2}$  is consistently smaller than  $\frac{Z^2}{\rho^2}$  for all pairs of  $p$  and  $q$ , up to  $100\times$ . Therefore, our algorithm CBS consistently requires fewer samples than EP/Zz++. This also explains why CBS achieves the same or even lower estimation error with a smaller sample size. For instance, CBS requires up to  $88\times$  fewer samples than EP/Zz++ on  $p = 5$  and  $q = 9$ , showing its efficiency in reducing sampling overhead while maintaining accuracy.

### 6.4 Memory Consumption

The theoretical guarantee of the memory usage is given in Theorem 4.8. We now present the actual memory consumption on each dataset in Table 3 (in Megabytes). Specifically, we report the average memory usage on each dataset for  $3 \leq p, q \leq 9$ . As we can observe, the memory usage is similar to all other baselines for most datasets. However, our memory consumption increases for the two

Table 2: The value of  $\frac{B^2}{C^2}$  and  $\frac{Z^2}{\rho^2}$  with varying  $p, q$  (Amazon).

$(p, q)$	$\frac{B^2}{C^2}$	$\frac{Z^2}{\rho^2}$
(3,4)	1.40E+03	1.79E+03
(3,5)	3.97E+03	1.69E+04
(3,9)	4.55E+04	2.85E+05
(4,5)	5.41E+04	5.98E+04
(4,8)	9.27E+04	2.48E+05
(5,3)	3.96E+02	6.17E+03
(5,6)	4.04E+05	1.20E+06
(5,9)	3.77E+04	3.34E+06
(6,4)	5.73E+03	5.80E+04
(6,7)	6.42E+06	5.32E+07
(7,4)	8.07E+03	1.54E+05
(7,7)	7.15E+06	2.44E+07
(8,4)	1.36E+04	3.16E+05
(8,8)	4.74E+07	1.04E+09
(9,4)	2.36E+04	5.32E+05
(9,9)	2.92E+08	8.01E+10

Table 3: Memory consumption of different counting algorithms (in MBs).

dataset	BCList++	EPivoter	EP/Zz++	CBS
GH	21.3	880.8	21.3	462.1
SO	40.9	901.0	73.7	120.8
Wut	52.2	911.0	69.8	312.8
IMDB	66.5	926.0	106.1	164.2
Actor2	88.3	946.9	103.1	268.5
AR	152.7	1004.5	295.8	461.0
DBLP	319.5	1166.3	444.5	718.2
ER	X	X	402.7	46600.6
DE	X	X	1744.8	195153.0

large datasets. This is in line with our analysis, as it is dominated by  $(p+q)|E|$ . Our algorithm requires to store the intermediate result for retrieving the broom structure. In result, the memory is mainly for storing the dynamic programming table.

## 6.5 Scalability Experiment

We consider counting bicliques approximately, mainly because all exact counting methods are highly inefficient in large graph data. Therefore, it is essential to verify the scalability of our algorithm thoroughly. Besides the main experiment in Section 5, we now conduct an additional graph scalability experiment that evaluates performance under dynamic modifications, which is done by randomly removing a percentage of edges. Specifically, we construct several graphs from DE as follows:

- (1) Randomly sample 60%, 80%, 100% of the vertices on both sides and keep the induced graph.
- (2) Execute our algorithm on  $p = 6, q = 8, T = 10^5$ .
- (3) Repeat (1) and (2) 10 times, compute the average runtime for initialization and sampling.

The three constructed graphs are DE60, DE80, and DE100. Shown in Table 4, we can observe that as the size of the graph increases, the sampling time also increases linearly while the error ratio is within a certain range. This shows that CBS’s scalability, which is expected to be particularly effective on large datasets.

Table 4: Scalability experiment.

	init time(s)	sample time(s)	error(%)
DE60%	298.6	40.8	3.76
DE80%	1216.7	95.3	2.49
DE100%	2478.1	135.5	3.19

Table 5: Average runtime (s) of EPivoter and CBS for all  $3 \leq p, q \leq 5$ .

Dataset	EPivoter	CBS
GH	805.4	3.4
SO	95.0	1.4
Wut	2429.4	4.3
IMDB	18.4	1.6
Actor2	53.4	2.8
AR	180.6	3.7
DBLP	18.5	1.9
ER	INF	364.4
DE	INF	3782.1

## 6.6 Efficiency Comparison on Small $p, q$

One of our major baselines, EPivoter is designed for exact counting on small  $p, q$  (e.g.,  $p, q \leq 5$ ). However, our major comparison experiment is conducted for all combinations of  $p, q \leq 9$ . To bridge the gap, we provide an additional experiment using the same parameter settings as EPivoter, i.e.,  $p, q \leq 5$ . Shown in Table 5, for ER and DE, EPivoter can not terminate in  $10^5$  seconds. For other datasets, we observe that our method maintains its performance advantage even when  $p$  and  $q$  are small, highlighting its robustness to parameter settings. These additional results confirm that our method performs consistently well regardless of the parameter choices.

## 7 CONCLUSION

In this paper, we tackle the problem of  $(p, q)$ -biclique counting in large-scale bipartite graphs, crucial for applications such as recommendation systems and cohesive subgraph analysis. To address scalability and accuracy issues in existing methods, we propose a novel sampling-based algorithm that leverages  $(p, q)$ -brooms, special spanning trees within  $(p, q)$ -bicliques. By utilizing the graph coloring method and dynamic programming, our method efficiently approximates  $(p, q)$ -biclique counts with unbiased estimates and provable error guarantees. Experimental results on nine real-world datasets show that our approach outperforms state-of-the-art methods, achieving up to  $8\times$  error reduction and  $50\times$  speed-up. Interesting future work includes extending our method to dynamic bipartite graphs with evolving structures and exploring its application to counting motifs/cliques in heterogeneous information networks.

## ACKNOWLEDGMENTS

Chen hao Ma was partially supported by NSFC under Grant 62302421, Basic and Applied Basic Research Fund in Guangdong Province under Grant 2025A1515010439, and the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong, Shenzhen. Weinuo Li and Can Wang were partially supported by the National Natural Science Foundation of China (No. 62476244), the Starry Night Science Fund of Zhejiang University Shanghai Institute for Advanced Study, China (Grant No: SN-ZJU-SIAS-001). Parts of this work was done while Jingbang was still at University of Waterloo.

## REFERENCES

- [1] Yikun Ban. 2018. On finding dense subgraphs in bipartite graphs: Linear algorithms. *arXiv preprint arXiv:1810.06809* (2018).
- [2] Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. 2013. Copycatch: stopping group attacks by spotting lockstep behavior in social networks. In *Proceedings of the 22nd international conference on World Wide Web*. 119–130.
- [3] Stephen P Borgatti and Martin G Everett. 1997. Network analysis of 2-mode data. *Social networks* 19, 3 (1997), 243–269.
- [4] Marco Bressan, Flavio Chierichetti, Ravi Kumar, Stefano Leucci, and Alessandro Panconesi. 2018. Motif counting beyond five nodes. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 12, 4 (2018), 1–25.
- [5] Coenraad Bron and Joep Kerbosch. 1973. Finding all cliques of an undirected graph (algorithm 457). *Commun. ACM* 16, 9 (1973), 575–576.
- [6] Xinwei Cai, Xiangyu Ke, Kai Wang, Lu Chen, Tianming Zhang, Qing Liu, and Yunjun Gao. 2024. Efficient Temporal Butterfly Counting and Enumeration on Temporal Bipartite Graphs. *Proc. VLDB Endow.* 17, 4 (mar 2024), 657–670. <https://doi.org/10.14778/3636218.3636223>
- [7] Hongxu Chen, Hongzhi Yin, Tong Chen, Weiqing Wang, Xue Li, and Xia Hu. 2020. Social boosted recommendation with folded bipartite network embedding. *IEEE Transactions on Knowledge and Data Engineering* 34, 2 (2020), 914–926.
- [8] Herman Chernoff. 1952. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics* (1952), 493–507.
- [9] Talya Eden, Dana Ron, and C Seshadhri. 2020. Faster sublinear approximation of the number of k-cliques in low-arboricity graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1467–1478.
- [10] Myoungji Han, Hyunjoon Kim, Geonmo Gu, Kunsoo Park, and Wook-Shin Han. 2019. Efficient subgraph matching: Harmonizing dynamic programming, adaptive matching order, and failing set together. In *Proceedings of the 2019 international conference on management of data*. 1429–1446.
- [11] Wassily Hoeffding. 1994. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding* (1994), 409–426.
- [12] Wenzhe Hou, Xiang Zhao, and Bo Tang. 2024. LearnSC: An Efficient and Unified Learning-Based Framework for Subgraph Counting Problem. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 2625–2638.
- [13] Shweta Jain and C Seshadhri. 2020. The power of pivoting for exact clique counting. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 268–276.
- [14] Shweta Jain and C Seshadhri. 2020. Provably and efficiently approximating near-cliques using the Turán shadow: PEANUTS. In *Proceedings of The Web Conference 2020*. 1966–1976.
- [15] Mehdi Kaytoue, Sergei O Kuznetsov, Amedeo Napoli, and Sébastien Duplessis. 2011. Mining gene expression data with pattern structures in formal concept analysis. *Information Sciences* 181, 10 (2011), 1989–2001.
- [16] Hyunjoon Kim, Yunyoung Choi, Kunsoo Park, Xuemin Lin, Seok-Hee Hong, and Wook-Shin Han. 2021. Versatile equivalences: Speeding up subgraph query processing and subgraph matching. In *Proceedings of the 2021 international conference on management of data*. 925–937.
- [17] Konect. 2006. Konect. <http://konect.cc/networks/>.
- [18] Ronghua Li, Sen Gao, Lu Qin, Guoren Wang, Weihua Yang, and Jeffrey Xu Yu. 2020. Ordering Heuristics for k-clique Listing. *Proc. VLDB Endow.* (2020).
- [19] Chenhao Ma, Reynold Cheng, Laks VS Lakshmanan, Tobias Grubenmann, Yixiang Fang, and Xiaodong Li. 2019. Linc: a motif counting algorithm for uncertain graphs. *Proceedings of the VLDB Endowment* 13, 2 (2019), 155–168.
- [20] Qiuyang Mang, Jingbang Chen, Hangrui Zhou, Yu Gao, Yingli Zhou, Richard Peng, Yixiang Fang, and Chenhao Ma. 2024. Efficient Historical Butterfly Counting in Large Temporal Bipartite Networks via Graph Structure-aware Index. *arXiv preprint arXiv:2406.00344* (2024).
- [21] Michael Mitzenmacher, Jakub Pachocki, Richard Peng, Charalampos Tsourakakis, and Shen Chen Xu. 2015. Scalable large near-clique detection in large-scale networks via sampling. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 815–824.
- [22] Laboratory of Web Algorithmics. 2013. Laboratory of Web Algorithmics Datasets. <http://law.di.unimi.it/datasets.php>.
- [23] Stanford Network Analysis Project. 2009. SNAP. <http://snap.stanford.edu/data/>.
- [24] Seyed-Vahid Sanehi-Mehri, Ahmet Erdem Sariyuce, and Srikanta Tirathapura. 2018. Butterfly counting in bipartite networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2150–2159.
- [25] Comandur Seshadhri and Srikanta Tirathapura. 2019. Scalable subgraph counting: The methods behind the madness. In *Companion Proceedings of The 2019 World Wide Web Conference*. 1317–1318.
- [26] Aida Sheshbolouki and M Tamer Özsu. 2022. sGrapp: Butterfly approximation in streaming graphs. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 4 (2022), 1–43.
- [27] Jia Wang, Ada Wai-Chee Fu, and James Cheng. 2014. Rectangle counting in large bipartite graphs. In *2014 IEEE International Congress on Big Data*. IEEE, 17–24.
- [28] Kai Wang, Xuemin Lin, Lu Qin, Wenjie Zhang, and Ying Zhang. 2019. Vertex Priority Based Butterfly Counting for Large-scale Bipartite Networks. *PVLDB* (2019).
- [29] Zhibin Wang, Longbin Lai, Yixue Liu, Bing Shui, Chen Tian, and Sheng Zhong. 2023. I/O-Efficient Butterfly Counting at Scale. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–27.
- [30] Chunqi Wu, Jingdong Li, Zhao Li, Ji Zhang, and Pan Tang. 2024. Accelerating Maximal Bicliques Enumeration with GPU on large scale network. *Future Generation Computer Systems* 161 (2024), 601–613.
- [31] Yifei Xia, Feng Zhang, Qingyu Xu, Mingde Zhang, Zhiming Yao, Lv Lu, Xiaoyong Du, Dong Deng, Bingsheng He, and Siqi Ma. 2024. GPU-based butterfly counting. *The VLDB Journal* (2024), 1–25.
- [32] Jianye Yang, Yun Peng, and Wenjie Zhang. 2021. (p, q)-biclique counting and enumeration for large sparse bipartite graphs. *Proceedings of the VLDB Endowment* 15, 2 (2021), 141–153.
- [33] Xiaowei Ye, Rong-Hua Li, Qiangqiang Dai, Hongzhi Chen, and Guoren Wang. 2022. Lightning fast and space efficient k-clique counting. In *Proceedings of the ACM Web Conference 2022*. 1191–1202.
- [34] Xiaowei Ye, Rong-Hua Li, Qiangqiang Dai, Hongchao Qin, and Guoren Wang. 2023. Efficient biclique counting in large bipartite graphs. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–26.
- [35] Xiaowei Ye, Rong-Hua Li, Longlong Lin, Shaojie Qiao, and Guoren Wang. 2024. Fast Biclique Counting on Bipartite Graphs: A Node Pivot-based Approach. *arXiv preprint arXiv:2412.16485* (2024).
- [36] Yun Zhang, Charles A Phillips, Gary L Rogers, Erich J Baker, Elissa J Chesler, and Michael A Langston. 2014. On finding bicliques in bipartite graphs: a novel algorithm and its application to the integration of diverse biological data types. *BMC bioinformatics* 15, 1 (2014), 110.
- [37] Kangfei Zhao, Jeffrey Xu Yu, Qiyan Li, Hao Zhang, and Yu Rong. 2023. Learned sketch for subgraph counting: a holistic approach. *The VLDB Journal* 32, 5 (2023), 937–962.
- [38] Alexander Zhou, Yue Wang, and Lei Chen. 2021. Butterfly counting on uncertain bipartite graphs. *Proceedings of the VLDB Endowment* 15, 2 (2021), 211–223.
- [39] Yingli Zhou, Yixiang Fang, Chenhao Ma, Tianci Hou, and Xin Huang. 2024. Efficient Maximal Motif-Clique Enumeration over Large Heterogeneous Information Networks. *Proceedings of the VLDB Endowment* 17, 11 (2024), 2946–2959.