Efficient k-Clique Listing: An Edge-Oriented Branching Strategy

KAIXIN WANG, Nanyang Technological University, Singapore
KAIQIANG YU*, Nanyang Technological University, Singapore
CHENG LONG, Nanyang Technological University, Singapore

k-clique listing is a vital graph mining operator with diverse applications in various networks. The state-of-the-art algorithms all adopt a branch-and-bound (BB) framework with a vertex-oriented branching strategy (called VBBkC), which forms a sub-branch by expanding a partial k-clique with a *vertex*. These algorithms have the time complexity of $O(k \cdot m \cdot (\delta/2)^{k-2})$, where m is the number of edges in the graph and δ is the degeneracy of the graph. In this paper, we propose a BB framework with a new *edge-oriented branching* (called EBBkC), which forms a sub-branch by expanding a partial k-clique with two vertices that connect each other (which correspond to an *edge*). We explore various edge orderings for EBBkC such that it achieves a time complexity of $O(m \cdot \delta + k \cdot m \cdot (\tau/2)^{k-2})$, where τ is an integer related to the maximum truss number of the graph and we have $\tau < \delta$. The time complexity of EBBkC is better than that of VBBkC algorithms for $k > 3$ since both $O(m \cdot \delta)$ and $O(k \cdot m \cdot (\tau/2)^{k-2})$ are bounded by $O(k \cdot m \cdot (\delta/2)^{k-2})$. Furthermore, we develop specialized algorithms for sub-branches on dense graphs so that we can early-terminate them and apply the specialized algorithms. We conduct extensive experiments on 19 real graphs, and the results show that our newly developed EBBkC based algorithms with the early termination technique consistently and largely outperform the state-of-the-art (VBBkC based) algorithms.

CCS Concepts: • **Mathematics of computing → Graph algorithms**.

Additional Key Words and Phrases: Graph mining; branch-and-bound; k-clique listing

ACM Reference Format:
Kaixin Wang, Kaiqiang Yu, and Cheng Long. 2024. Efficient k-Clique Listing: An Edge-Oriented Branching Strategy. *Proc. ACM Manag. Data* 2, 1 (SIGMOD), Article 7 (February 2024), 26 pages. https://doi.org/10.1145/3639262

1 INTRODUCTION

Given a graph G, a k-clique is subgraph of G with k vertices such that each pair of vertices inside are connected [10]. k-clique listing, which is to list all k-cliques in a graph, is a fundamental graph mining operator that plays a crucial role in various data mining applications across different networks, including social networks, mobile networks, Web networks, and biological networks. Some significant applications include the detection of overlapping communities in social networks [30], identifying k-clique communities in mobile networks [17, 19, 31], detecting link spams in Web networks [20, 34], and discovering groups of functionally related proteins (known as modules) in gene association networks [2]. Moreover, k-clique listing serves as a key component for several

*Kaiqiang Yu is the corresponding author.

Authors' addresses: Kaixin Wang, Nanyang Technological University, Singapore, kaixin.wang@ntu.edu.sg; Kaiqiang Yu, Nanyang Technological University, Singapore, kaiqiang002@e.ntu.edu.sg; Cheng Long, Nanyang Technological University, Singapore, c.long@ntu.edu.sg.

ACM 2836-6573/2024/2-ART7
https://doi.org/10.1145/3639262

other tasks, such as finding large near cliques [39], uncovering the hierarchical structure of dense subgraphs [35], exploring k-clique densest subgraphs [15, 40], identifying stories in social media [3], and detecting latent higher-order organization in real-world networks [5]. For more detailed information on how k-clique listing is applied in these contexts, please refer to references [24] and [46].

Quite a few algorithms have been proposed for listing k-cliques [10, 12, 24, 46]. The majority of these approaches adopt a *branch-and-bound* (BB) framework, which involves recursively dividing the problem of listing all k-cliques in graph G into smaller sub-problems of listing smaller cliques in G through *branching* operations [12, 24, 46]. This process continues until each sub-problem can be trivially solved. The underlying principle behind these methods is the observation that a k-clique can be constructed by merging two smaller cliques: a clique S and an l-clique, where $|S| + l = k$. A branch B is represented as a triplet (S, g, l), where S denotes a previously found clique with $|S| < k$, g represents a subgraph induced by vertices that connect each vertex in S, and l corresponds to $k - |S|$. Essentially, branch B encompasses all k-cliques, each comprising of S and an l-clique in g. To enumerate all k-cliques within branch B, a set of sub-branches is created through branching operations. Each sub-branch expands the set S by adding one vertex from g, updates the graph g accordingly (by removing vertices that are disconnected from the newly added vertex from S), and decrements l by 1. This recursive process continues until l for a branch reduces to 2, at which point the l-cliques (corresponding to edges) can be trivially listed within g. The original k-clique listing problem on graph G can be solved by initiating the branch (S, g, l) with $S = \emptyset$, $g = G$, and $l = k$. The branching step employed in these existing methods is referred to as *vertex-oriented branching* since each sub-branch is formed by adding a vertex to the set S. We term the vertex-oriented branching BB framework for k-clique listing as VBBkC.

Existing research has focused on leveraging vertex information within the graph g to enhance performance by generating sub-branches with smaller graphs and pruning sub-branches more effectively. Specifically, when a new vertex v_i is added, the resulting sub-branch B_i can reduce the graph g to a smaller size by considering only the neighbors of v_i (e.g., removing vertices that are not adjacent to v_i). Consequently, branch B_i can be pruned if v_i has fewer than $(l - 1)$ neighbors in g. Furthermore, by carefully specifying the ordering of vertices in g during the branching process, it becomes possible to generate a group of sub-branches with compact graphs. The state-of-the-art VBBkC algorithms [24, 46] achieve a time complexity of $O(km(\delta/2)^{k-2})$, where m represents the number of edges in the graph, and δ denotes the degeneracy[1] of the graph.

In this paper, we propose to construct a branch by simultaneously including *two* vertices that connect to each other (i.e., an *edge*) from g into S. Note that these two vertices must be connected, as only then they can form a larger partial k-clique together with S. This strategy allows for the consideration of additional information (i.e., an edge or two vertices instead of a single vertex) to facilitate the formation of sub-branches with smaller graphs and pruning. When a new edge is added, the resulting branch can reduce the graph to one induced by the common neighbors of the two vertices, which is smaller compared to the graph generated by vertex-oriented branching. To achieve this, we explore an edge ordering technique based on truss decomposition [41], which we refer to as the *truss-based edge ordering*. This ordering aids in creating sub-branches with smaller graphs than those by vertex-oriented branching. Additionally, more branches can be pruned based on the information derived from the added edge. For instance, a branch can be pruned if the vertices within the newly added edge have fewer than $(l - 2)$ common neighbors in g. We term this branching strategy as *edge-oriented branching*. Consequently, the edge-oriented branching-based

[1]It is defined to be the maximum value of k such that there exists a non-empty k-core in a graph, where a k-core is a graph where each vertex has the degree at least k.

BB framework for k-clique listing is denoted as EBBkC. Our EBBkC algorithm, combined with the proposed edge ordering, exhibits a time complexity of $O(\delta m + km(\tau/2)^{k-2})$, where τ represents the truss number[2] of the graph. We formally prove that $\tau < \delta$, signifying that our EBBkC algorithm possesses a time complexity that is better than that of VBBkC algorithms [24, 46] since both $O(\delta m)$ and $O(km(\tau/2)^{k-2})$ are bounded by $O(km(\delta/2)^{k-2})$ when $k > 3$.

It is important to note that although a single branching step in our edge-oriented branching (i.e., including an edge) can be seen as two branching steps in the existing vertex-oriented branching (i.e., including two vertices of an edge via two steps), there exists a significant distinction between our EBBkC and VBBkC frameworks. In EBBkC, we have the flexibility to explore *arbitrary* edge orderings for each branching step, whereas VBBkC is inherently *constrained* by the chosen vertex ordering. For instance, once a vertex ordering is established for vertex-oriented branching, with vertex v_i appearing before v_j, the edges incident to v_i would precede those incident to v_j in the corresponding edge-oriented branching. Consequently, the existing VBBkC framework is encompassed by our EBBkC framework. In other words, for any instance of VBBkC with a given vertex ordering, there exists an edge ordering such that the corresponding EBBkC instance is equivalent to the VBBkC instance, *but not vice versa*. This elucidates why EBBkC, when based on certain edge orderings, achieves a superior time complexity compared to VBBkC.

To further enhance the efficiency of BB frameworks, we develop an *early termination* technique, which is based on two key observations. Firstly, if the graph g within a branch (S, g, l) is either a clique or a 2-plex[3], we can efficiently list l-cliques in g using a combinatorial approach. For instance, in the case of a clique, we can directly enumerate all possible sets of l vertices in g. Secondly, if the graph g is a t-plex (with $t \geq 3$), the branching process based on g can be converted to a procedure on its inverse graph, denoted as g_{inv}[4]. Since g is dense, g_{inv} would be sparse, and the converted procedure operates more rapidly. Therefore, during the recursive branching process, we can employ early termination at a branch (S, g, l) if g transforms into a t-plex, utilizing efficient algorithms to list l-cliques within g. We note that the early termination technique is applicable to all BB frameworks, including our EBBkC framework, without impacting the worst-case time complexity of the BB frameworks.

Contributions. We summarize our contributions as follows.

- We propose a new BB framework for k-clique listing problem, namely EBBkC, which is based on an edge-oriented branching strategy. We further explore different edge orderings for EBBkC such that it achieves a better time complexity than that of the state-of-the-art VBBkC for $k > 3$, i.e., the former is $O(\delta m + km(\tau/2)^{k-2})$ and the latter is $O(km(\delta/2)^{k-2})$, where τ is a number related to the maximum truss number of the graph and δ is the degeneracy of the graph and we have $\tau < \delta$. (Section 4)
- We further develop an early termination technique for boosting the efficiency of branch-and-bound frameworks including EBBkC, i.e., for branches of listing l-cliques in a dense graph (e.g., a t-plex), we develop more efficient algorithms based on combinatorial approaches (for a clique and a 2-plex) and conduct the branching process on its inverse graph (for a t-plex with $t \geq 3$), which would be faster. (Section 5)
- We conduct extensive experiments on 19 real graphs, and the results show that our EBBkC based algorithms with the early termination technique consistently and largely outperform the state-of-the-art (VBBkC based) algorithms. (Section 6)

[2]The maximum truss number of a graph defined in [41], denoted by $k_{\max}$, has the following relationship with τ: $k_{\max} = \tau + 2$.
[3]A t-plex is a graph where each vertex inside has at most t non-neighbors including itself.
[4]The inverse graph g_{inv} has the same set of vertices as g, with an edge between two vertices in g_{inv} if and only if they are disconnected from each other in g.

The rest of the paper is organized as follows. Section 2 reviews the problem and presents some preliminaries. Section 3 summarizes the existing vertex-oriented branching-based BB framework. Section 7 reviews the related work and Section 8 concludes the paper.

2 PROBLEM AND PRELIMINARIES

We consider an *unweighted* and *undirected* simple graph $G = (V, E)$, where V is the set of vertices and E is the set of edges. We denote by $n = |V|$ and $m = |E|$ the cardinalities of V and E, respectively. Given $u, v \in V$, both (u, v) and (v, u) denote the undirected edge between u and v. Given $u \in V$, we denote by $N(u, G)$ the set of neighbors of u in G, i.e., $N(u, G) = \{v \in V \mid (u, v) \in E\}$ and define $d(u, G) = |N(u, G)|$. Given $V_{sub} \subseteq V$, we use $N(V_{sub}, G)$ to denote the common neighbors of vertices in V_{sub}, i.e., $N(V_{sub}, G) = \{v \in V \mid \forall u \in V_{sub},\ (v, u) \in E\}$.

Given $V_{sub} \subseteq V$, we denote by $G[V_{sub}]$ the subgraph of G induced by V_{sub}, i.e., $G[V_{sub}]$ includes the set of vertices V_{sub} and the set of edges $\{(u, v) \in E \mid u, v \in V_{sub}\}$. Given $E_{sub} \subseteq E$, we denote by $G[E_{sub}]$ the subgraph of G induced by E_{sub}, i.e., $G[E_{sub}]$ includes the set of edges E_{sub} and the set of vertices $\{v \in V \mid (v, \cdot) \in E_{sub}\}$. Let g be a subgraph of G induced by either a vertex subset of V or an edge subset of E. We denote by $V(g)$ and $E(g)$ its set of vertices and its set of edges, respectively.

In this paper, we focus on a widely-used cohesive graph structure, namely k-clique [14], which is defined formally as below.

Definition 2.1 (k-clique [14]). *Given a positive integer k, a subgraph g is said to be a k-clique if and only if it has k vertices and has an edge between every pair of vertices, i.e., $|V(g)| = k$ and $E(g) = \{(u, v) \mid u, v \in V(g), u \neq v\}$.*

We note that 1-clique ($k = 1$) and 2-clique ($k = 2$) correspond to single vertex and single edge, respectively, which are basic elements of a graph. Therefore, we focus on those k-cliques with k at least 3 (note that 3-clique is widely known as triangle and has found many applications [23, 28]). We now formulate the problem studied in this paper as follows.

Problem 2.1 (k-clique listing [10]). *Given a graph $G = (V, E)$ and a positive integer $k \geq 3$, the k-clique listing problem aims to find all k-cliques in G.*

Hardness. The k-clique listing problem is a hard problem since the decision problem of determining whether a graph contains a k-clique is NP-hard [22] and this problem can be solved by listing all k-cliques and returning true if any k-clique is listed.

Remark. The problem of listing all 3-cliques (i.e., $k = 3$), known as *triangle listing problem*, has been widely studied [23, 28]. There are many efficient algorithms proposed for triangle listing, which run in *polynomial* time. We remark that these algorithms cannot be used to solve the general k-clique listing problem.

3 THE BRANCH-AND-BOUND FRAMEWORK OF EXISTING ALGORITHMS: VBBKC

Many algorithms have been proposed for listing k-cliques in the literature [10, 12, 24, 46]. Most of them adopt a *branch-and-bound* (BB) framework, which recursively partitions the problem instance (of listing all k-cliques in G) into several sub-problem instances (of listing smaller cliques in G) via *branching* until each of them can be solved trivially [12, 24, 46]. The rationale behind these methods is that a k-clique can be constructed by merging two smaller cliques, namely a clique S and an l-clique with $|S| + l = k$. Specifically, a branch B can be represented as a triplet (S, g, l), where

- **set** S induces a clique found so far with $|S| < k$,
- **subgraph** g is one induced by vertices that connect each vertex in S, and

Algorithm 1: The vertex-oriented branching-based BB framework: VBBkC

Input: A graph $G = (V, E)$ and an integer $k \geq 3$
Output: All k-cliques within G

```
1 VBBkC_Rec(∅, G, k);
2 Procedure VBBkC_Rec(S, g, l)
    /* Pruning                                                    */
3   if |V(g)| < l then return;
    /* Termination when l = 2                                     */
4   if l = 2 then
5     for each edge (u, v) in g do Output a k-clique S ∪ {u, v} ;
6     return;
    /* Branching when l ≥ 3                                       */
7   for each vertex v_i ∈ V(g) based on a given vertex ordering do
8     Create branch B_i = (S_i, g_i, l_i) based on Eq. (1);
9     VBBkC_Rec(S_i, g_i, l_i);
```

- **integer** l is equal to $k - |S|$.

Essentially, branch B covers all k-cliques, each consisting of S and an l-clique in g. To list all k-cliques under branch B, it creates a group of sub-branches via a branching step such that for each sub-branch, the set S is expanded with one vertex from g, the graph g is updated accordingly (by removing those vertices that are not adjacent to the vertex included in S), and l is decremented by 1. The recursive process continues until when the l for a branch reduces to 2, for which the l-cliques (which correspond to edges) can be listed trivially in g. The original k-clique listing problem on graph G can be solved by starting with the branch (S, g, l) with $S = \emptyset$, $g = G$ and $l = k$. We call the branching step involved in these existing methods *vertex-oriented branching* since each sub-branch is formed by including a *vertex* to the set S.

Consider the branching step at a branch $B = (S, g, l)$. Let $\langle v_1, v_2, \cdots, v_{|V(g)|}\rangle$ be an arbitrary ordering of vertices in g. The branching step would produce $|V(g)|$ new sub-branches from branch B. The i-th sub-branch, denoted by $B_i = (S_i, g_i, l_i)$, includes v_i to S and excludes $\{v_1, v_2, \cdots, v_{i-1}\}$ (and also those that are not adjacent to v_i). Formally, for $1 \leq i \leq |V(g)|$, we have

$$S_i = S \cup \{v_i\}, \quad g_i = \widehat{g}_i[N(v_i, \widehat{g}_i)], \quad l_i = l - 1, \tag{1}$$

where $\widehat{g}_i$ is a subgraph of g induced by the set of vertices $\{v_i, v_{i+1}, \cdots, v_{|V(g)|}\}$, i.e., $\widehat{g}_i = g[v_i, \cdots, v_{|V(g)|}]$. The branching can be explained by a recursive binary partition process, as shown in Figure 1(a). Specifically, it first divides the current branch B into two sub-branches based on v_1: one branch moves v_1 from g to S (this is the branch B_1 which will list those k-cliques in B that include v_1), and the other removes v_1 from g (so that it will list others in B that exclude v_1). For branch B_1, it also removes from g those vertices that are not adjacent to v_1 since they cannot form any k-clique with v_1. In summary, we have $g_1 = \widehat{g}_1[N(v_1, \widehat{g}_1)]$. Then, it recursively divides the latter into two new sub-branches: one branch moves v_2 from $\widehat{g}_2$ to S (this is the branch B_2 which will list those k-cliques in B that exclude v_1 and include v_2), and the other removes v_2 from $\widehat{g}_2$ (so that it will list others in B that exclude $\{v_1, v_2\}$). It continues the process, until the last branch $B_{|V(g)|}$ is formed. In summary, branch B_i will list those k-cliques in B that include v_i and exclude $\{v_1, \cdots, v_{i-1}\}$. Consequently, all k-cliques in B will be listed exactly once after branching.

We call the vertex-oriented branching-based BB framework for k-clique listing VBBkC. We present its pseudo-code in Algorithm 1. In particular, when $l = 2$, a branch (S, g, l) can be terminated

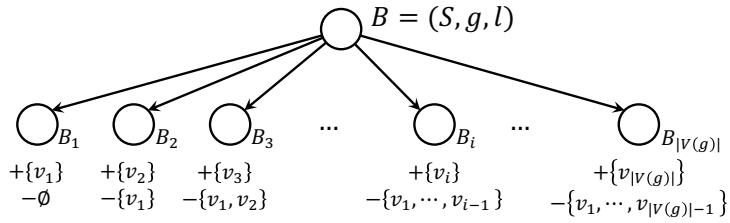

(a) Branching in VBBkC framework. The notation "+" means to include a vertex by adding it S and "−" means to exclude a vertex by removing it from the graph g.

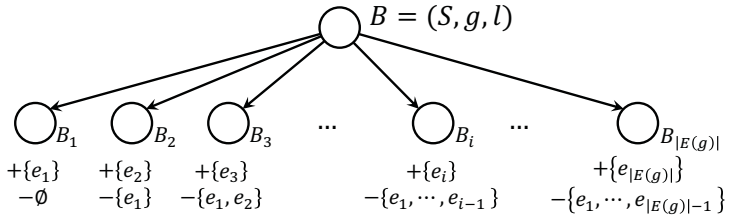

(b) Branching in EBBkC framework. The notation "+" means to include two vertices incident to an edge by adding them to S and "−" means to exclude an edge by removing it from the graph g.

Fig. 1. Illustration of VBBkC and EBBkC.

by listing each of edges in g together with S (lines 4-6). We remark that Algorithm 1 presents the algorithmic idea only while the detailed implementations (e.g., the data structures used for representing a branch) often vary in existing algorithms [12, 24, 46]. Different variants of VBBkC have different time complexities. The state-of-the-art algorithms, including DDegCol [24], DDegree [24], BitCol [46] and SDegree [46], all share the time complexity of $O(km(\delta/2)^{k-2})$, where δ is the degeneracy of the graph. Some more details of variants of VBBkC will be provided in the related work (Section 7).

4 A NEW BRANCH-AND-BOUND FRAMEWORK: EBBKC

4.1 Motivation and Overview of EBBkC

Recall that for a branch $B = (S, g, l)$, the vertex-oriented branching forms a sub-branch by moving one vertex v_i from g to S. To improve the performance, existing studies consider the information of each vertex in g towards pruning more sub-branches and/or forming sub-branches with smaller graph instances. Specifically, with the newly added vertex v_i, the produced sub-branch B_i can shrink the graph instance g as a smaller one induced by the neighbors of v_i (e.g., removing those vertices that are not adjacent to v_i). As a result, one can prune the branch B_i if v_i has less than $(l-1)$ neighbors in g. In addition, one can produce a group of sub-branches with small graph instances by specifying the ordering of vertices in g for branching.

In this paper, we propose to form a branch by *moving two vertices that connect with each other* (correspondingly, an edge) from g to S at the same time. Note that the two vertices are required to be connected since otherwise they will not form a larger partial k-clique with S. The intuition is that it would allow us to consider more information (i.e., an edge or two vertices instead of a single vertex) towards pruning more sub-branches and/or forming sub-branches with smaller graph instances. Specifically, with the newly added edge, the produced branch can shrink the graph instance as the one induced by the common neighbors of two vertices of the edge, which is smaller than that produced by the vertex-oriented branching (details can be found in Section 4.2). In addition, we

Algorithm 2: The edge-oriented branching-based BB framework: EBBkC

```
   Input: A graph G = (V, E) and an integer k ≥ 3
   Output: All k-cliques within G
1  EBBkC_Rec(∅, G, k);
2  Procedure EBBkC_Rec(S, g, l)
       /* Pruning                                                    */
3      if |V(g)| < l then return;
       /* Termination when l = 1 or l = 2                            */
4      if l = 1 then
5          for each vertex v in g do Output a k-clique S ∪ {v} ;
6          return;
7      else if l = 2 then
8          for each edge (u, v) in g do Output a k-clique S ∪ {u, v} ;
9          return;
       /* Branching when l ≥ 3                                       */
10     for each edge e_i ∈ E(g) based on a given edge ordering do
11         Create branch B_i = (S_i, g_i, l_i) based on Eq. (2);
12         EBBkC_Rec(S_i, g_i, l_i);
```

can prune more branches based on the information of the added edge, e.g., a produced branch can be pruned if the vertices in the newly added edge have less than $(l-2)$ common neighbors in g (details can be found in Section 4.3). We call the above branching strategy *edge-oriented branching*, which we introduce as follows.

Consider the branching step at a branch $B = (S, g, l)$. Let $\langle e_1, e_2, \cdots, e_{|E(g)|}\rangle$ be an *arbitrary* ordering of edges in g. Then, the branching step would produce $|E(g)|$ new sub-branches from B. The i-th branch, denoted by $B_i = (S_i, g_i, l_i)$, includes to S the two vertices incident to the edge e_i, i.e., $V(e_i)$, and excludes from g those edges in $\{e_1, e_2, \cdots, e_{i-1}\}$ (and those vertices that are disconnected from the vertices incident to e_i). Formally, for $1 \le i \le |E(g)|$, we have

$$S_i = S \cup V(e_i), \quad g_i = \overline{g}_i[N(V(e_i), \overline{g}_i)], \quad l_i = l - 2, \tag{2}$$

where $\overline{g}_i$ is a subgraph of g induced by the set of edges $\{e_i, e_{i+1}, \cdots, e_{|E(g)|}\}$, i.e., $\overline{g}_i = g[e_i, \cdots, e_{|E(g)|}]$. We note that (1) $N(V(e_i), \overline{g}_i)$ is to filter out those vertices that are not adjacent to the two vertices incident to e_i since they cannot form any k-cliques with e_i and (2) $\overline{g}_i[N(V(e_i), \overline{g}_i)]$ is the graph instance for the sub-branch induced by the vertex set $N(V(e_i), \overline{g}_i)$ in $\overline{g}_i$.

The edge-oriented branching also corresponds to a recursive binary partition process, as illustrated in Figure 1(b). Specifically, it first divides branch B into two sub-branches based on e_1: one moves e_1 from g to S (this is the branch B_1 which will list those k-cliques in B that include edge e_1), and the other removes e_1 from g (so that it will list others that exclude e_1). For branch B_1, it also removes from g those vertices that are not adjacent to the vertices in $V(e_1)$ (i.e., $g_1 = \overline{g}_1[N(V(e_1), \overline{g}_1)]$) since they cannot form any k-clique with e_1. Then, it recursively divides the latter into two new sub-branches: one moves e_2 from $\overline{g}_2$ to S (this the branch B_2 which will list those k-cliques in B that exclude e_1 and include e_2), and the other removes e_2 from $\overline{g}_2$ (so that it will list others that exclude $\{e_1, e_2\}$). It continues the process, until the last branch $B_{|E(g)|}$ is formed.

We call the *edge-oriented branching-based BB framework for k-clique listing* EBBkC. We present in Algorithm 2 the pseudo-code of EBBkC, which differs with VBBkC mainly in the branching step (lines 10-11). We note that while a branching step in our edge-oriented branching (i.e., including an

Algorithm 3: EBBkC with truss-based edge ordering: EBBkC−T

Input: A graph $G = (V, E)$ and an integer $k \geq 3$
Output: All k-cliques within G
/* Initialization and branching at $(\emptyset, G, k)$ */
1 $\pi_\tau(G) \leftarrow$ the truss-based ordering of edges in G;
2 **for** *each edge* $e_i \in E(G)$ *following* $\pi_\tau(G)$ **do**
3 Obtain S_i and g_i according to Eq. (2);
4 Initialize $VSet(e_i) \leftarrow V(g_i)$ and $ESet(e_i) \leftarrow E(g_i)$;
5 EBBkC-T_Rec$(S_i, g_i, k-2)$;
6 **Procedure** EBBkC-T_Rec(S, g, l)
7 Conduct Pruning and Termination (lines 3-9 of Algorithm 2);
/* Branching when $l \geq 3$ */
8 **for** *each edge* e *in* $E(g)$ **do**
9 $S' \leftarrow S \cup V(e)$, $g' \leftarrow (V(g) \cap VSet(e), E(g) \cap ESet(e))$;
10 EBBkC-T_Rec$(S', g', l-2)$;

edge) can be treated as two branching steps in the existing vertex-oriented branching (i.e., including two vertices of an edge via two steps), our EBBkC has a major difference from VBBkC as follows. For the former, we can explore *arbitrary* edge orderings for each branching step while for the latter, the underlying edge orderings are *constrained* by the adopted vertex ordering. For example, once a vertex ordering is decided for vertex-oriented branching with vertex v_i appearing before v_j, then the edges that are incident to v_i would appear before those that are incident to v_j for the corresponding edge-oriented branching. For this reason, the existing VBBkC framework is covered by our EBBkC framework, i.e., for any instance of VBBkC with a vertex ordering, there exists an edge ordering such that the corresponding EBBkC instance is equivalent to the VBBkC instance, *but not vice versa*. This explains why the time complexity of EBBkC based on some edge ordering is better than that of VBBkC (details can be found in Section 4.2).

In the sequel, we explore different orderings of edges in EBBkC. Specifically, with the proposed truss-based edge ordering, EBBkC would have the worst-case time complexity of $O(\delta m + km(\tau/2)^{k-2})$ with $\tau < \delta$. We note that the time complexity is better than that of the state-of-the-art VBBkC algorithms (which is $O(km(\delta/2)^{k-2})$) when $k > 3$ (Section 4.2). With the proposed color-based edge ordering, EBBkC can apply some pruning rules to improve the efficiency in practice (Section 4.3). Then, with the proposed hybrid edge ordering, EBBkC would inherit both the above theoretical result and practical performance (Section 4.4). Finally, we discuss other potential applications of EBBkC (Section 4.5).

4.2 EBBkC−T: EBBkC with the Truss-based Edge Ordering

Consider the edge-oriented branching at $B = (S, g, l)$ based on an ordering of edges $\langle e_1, e_2, \cdots, e_{|E(g)|}\rangle$. For a sub-branch B_i $(1 \leq i \leq |E(g)|)$ produced from B, we observe that the size of graph instance g_i (i.e., the number of vertices in g_i) is equal to the number of common neighbors of vertices in $V(e_i)$ in $\overline{g}_i$, which depends on the ordering of edges. Formally, we have

$$|V(g_i)| = |N(V(e_i), \overline{g}_i)|, \text{ where } \overline{g}_i = g[e_i, \cdots, e_{|E(g)|}]. \tag{3}$$

Recall that the smaller a graph instance is, the faster the corresponding branch can be solved. Therefore, to reduce the time costs, we aim to minimize the sizes of graph instances by determining the ordering of edges via the following greedy procedure.

Truss-based edge ordering. We determine the ordering of edges by an iterative process. Specifically, it iteratively removes from g the edge whose vertices have the smallest number of common neighbors (correspondingly, the smallest size of a graph instance), and then adds it to the end of the ordering. Consequently, for the i-th edge e_i in the produced ordering ($1 \leq i \leq |E(g)|$), we have

$$e_i = \min_{e \in E(g) \setminus \{e_1, \cdots, e_{i-1}\}} |N(V(e_i), g[E(g) \setminus \{e_1, \cdots, e_{i-1}\}])|. \tag{4}$$

We call the above ordering *truss-based edge ordering* of g and denote it by $\pi_\tau(g)$ since the corresponding iterative process is the same to the *truss decomposition* [9, 11, 41], which can be done in $O(\delta m)$ time, where δ is the degeneracy of the graph [9].

The EBBkC−T algorithm. When the truss-based edge ordering is adopted in EBBkC, we call the resulting algorithm EBBkC−T. The pseudo-code of EBBkC−T is presented in Algorithm 3. In particular, it only computes the truss-based edge ordering of G (i.e., $\pi_\tau(G)$) for the branching at the initial branch $(\emptyset, G, k)$ (lines 2-5). Then, for any other branching step at a following branch (S, g, l), edges in $\langle e_1, e_2, \cdots, e_{|E(g)|}\rangle$ adopt the same ordering to those used in $\pi_\tau(G)$, which could differ with the truss-based edge ordering of g. Formally, e_i comes before e_j in $\langle e_1, e_2, \cdots, e_{|E(g)|}\rangle$ (i.e., $i < j$) if and only if it does so in $\pi_\tau(G)$. To implement this efficiently, it maintains two additional auxiliary sets, i.e., $VSet(\cdot)$ and $ESet(\cdot)$ (lines 2-4). The idea is that for an edge e, all edges in $ESet(e)$ are ordered behind e in π_τ and the vertices incident to these edges are both connected with those incident to e. Therefore, the branching steps (with the introduced edge ordering) can be efficiently conducted via set intersections (line 9) for those branches following $(\emptyset, G, k)$. The correctness of this implementation can be easily verified.

Complexity. Given a branch $B = (S, g, l)$, let $\tau(g)$ be the largest size of a produced graph instance, i.e.,

$$\tau(g) = \max_{e_i \in E(g)} |V(g_i)|. \tag{5}$$

We have the following observation.

Lemma 4.1. *When applying the truss-based edge ordering $\pi_\tau(g)$ at (S, g, l), we have $\tau(g) < \delta(g)$, where $\delta(g)$ is the degeneracy of g.*

Proof. We prove by contradiction. Suppose that $\tau \geq \delta$. Since δ is defined as the largest value of k such that the k-core of g is non-empty, there must not have a $(\delta+1)$-core in g, i.e., $V(C_{\delta+1}) = \emptyset$. Here, C_k refers to a k-core. Consider the branching step at such an edge $e_i \in E(g)$ that produces the largest size of graph instance, i.e., $|V(g_i)| = \tau$. According to Eq. (4), e_i has the minimum number of common neighbors of its end points in the graph $\overline{g}_i = g[e_i, \cdots, e_{|E(g)|}]$. This means for each edge $e \in E(\overline{g}_i)$, the number of common neighbors of the end points of e is no less than τ, i.e., $|N(V(e), \overline{g}_i)| \geq \tau$, Obviously, $\overline{g}_i$ is non-empty, i.e., $V(\overline{g}_i) \neq \emptyset$. Then for each vertex $v \in V(\overline{g}_i)$, the number of its neighbors is at least $\tau + 1$, i.e., $|N(v, \overline{g}_i)| \geq \tau + 1$. Therefore, $\overline{g}_i$ is a subgraph of $(\tau + 1)$-core, i.e., $V(\overline{g}_i) \subseteq V(C_{\tau+1})$. According to the hereditary property of k-core[5] and the hypothesis $\tau \geq \delta$, we have $V(C_{\tau+1}) \subseteq V(C_{\delta+1})$, which leads to a contradiction that $V(\overline{g}_i) \subseteq V(C_{\tau+1}) \subseteq V(C_{\delta+1}) = \emptyset$. □

Based on the above result, we derive that the time complexity of EBBkC−T is better than that of the state-of-the-art algorithms, i.e., $O(km(\delta/2)^{k-2})$, which we show in the following theorem.

Theorem 4.2. *Given a graph $G = (V, E)$ and an integer $k \geq 3$, the time and space complexities of* EBBkC−T *are $O(\delta m + km(\tau/2)^{k-2})$ and $O(m + n)$, respectively, where $\tau = \tau(G)$ and it is strictly smaller than the degeneracy δ of G.*

[5]The hereditary property claims that given a graph G and two integers k and k' with $k \leq k'$, then the k'-core of G is a subgraph of the k-core of G [4].

Proof. We give a sketch of the proof and put the details in the technical report [42]. The running time of EBBkC−T consists of the time of generating the truss-based edge ordering, which is $O(\delta m)$ [9], and the time of the recursive listing procedure (lines 6-10 of Algorithm 3). Consider the latter one. Given a branch $B = (S, g, l)$, we denote by $T(g, l)$ the upper bound of time cost of listing l-cliques under such a branch. When $k \geq 3$, with different values of l, we have the following recurrences.

$$T(g,l) \leq \begin{cases} O(k \cdot |V(g)|) & l = 1 \\ O(k \cdot |E(g)|) & l = 2 \\ \sum_{e \in E(g)} \left(T(g', l-2) + T'(g') \right) & 3 \leq l \leq k-2 \end{cases} \tag{6}$$

where $T'(g')$ is the time for constructing g' given $B = (S, g, l)$ (line 9 of Algorithm 3). We show that with different values of l, $T'(g')$ satisfies the following equation.

$$\sum_{e \in E(g)} T'(g') = \begin{cases} O(\tau \cdot |E(g)|) & l = 3 \\ O(\tau^2 \cdot |E(g)|) & l > 3 \end{cases} \tag{7}$$

The reason is as follows. When given a branch $B = (S, g, l)$ with $l = 3$, for each edge, we just need to compute $V(g')$ for the sub-branch (since the termination when $l = 1$ only cares about the vertices in g'), which can be done in $O(\tau)$. When given a branch $B = (S, g, l)$ with $l > 3$, for each edge, we need to construct both $V(g')$ and $E(g')$, which can be done in $O(\tau^2)$ since there are at most $\tau(\tau - 1)/2$ edges in g. Besides, we show that given a branch $B = (S, g, l)$ and the sub-branches $B' = (S', g', l')$ produced at B, we have

$$\sum_{e \in E(g)} |E(g')| < \begin{cases} \frac{\tau^2}{4} \cdot |E(g)| & l < k \\ \frac{\tau^2}{2} \cdot |E(g)| & l = k \end{cases} \tag{8}$$

This inequality is proven in the technical report [42]. With above inequalities, we can prove the theorem by induction on l. □

VBBkC v.s. EBBkC−T (Time Complexity). The time complexity of EBBkC−T (i.e., $O(m\delta + km(\tau/2)^{k-2})$) is better than that of state-of-the-art VBBkC algorithms (i.e., $O(km(\delta/2)^{k-2})$) for $k > 3$. This is because both (1) $O(m\delta)$ and (2) $O(km(\tau/2)^{k-2})$ are bounded by $O(km(\delta/2)^{k-2})$. For (1), it is because we have $\delta < 2k(\delta/2)^{k-2}$ when $k > 3$; and for (2), it is because $\tau < \delta$. We note that for $k = 3$, the time complexity of EBBkC−T is dominated by $O(\delta m)$ since $\tau < \delta$, which is the same as that of state-of-the-art VBBkC algorithms (i.e., $O(\delta m)$) and is also the same as that of algorithms for listing triangles [23, 28] in the worst case (i.e., $O(m^{1.5})$) since $\delta < \sqrt{m}$.

Discussion on τ. By definition, τ of a graph G, i.e., $\tau(G)$, corresponds to the largest integer such that there exists a non-empty subgraph where the two vertices of each edge have at least τ common neighbors. Similar to the degeneracy δ of the graph, τ also measures the density of a graph, say, the larger the value of τ, the denser the graph. However, τ is always smaller than δ as it imposes stricter constraint on connections (i.e., the two vertices of every edge have at least τ common neighbors v.s. every vertex has at least δ neighbors). Theoretically, for a graph with n vertices, the gap between δ and τ can be as large as $n/2$. To see this, consider a complete bipartite graph with p vertices on each side, where p is a positive integer. For this graph, we have $\delta = p$ since each vertex has exactly p neighbors and $\tau = 0$ since the two vertices of each edge have no common neighbors. Practically, the ratio τ/δ is below 0.8 for many real-world graphs. For example, we have collected the statistics of τ/δ on 139 more real-world graphs [1] and found that the ratio is below 0.8 for the majority of the graphs (105 out of 139).

Remark. (1) We note that another option of designing EBBkC−T is to compute the truss-based edge ordering for each individual branch and use it for branching at the branch. However, it would

Algorithm 4: EBBkC with color-based edge ordering: EBBkC−C

Input: A graph $G = (V, E)$ and an integer $k \geq 3$
Output: All k-cliques within G

```
1 Conduct vertex coloring on G and get id(v) for each vertex in V;
2 G⃗ ← (V, E⃗) where E⃗ = {u → v | (u,v) ∈ E ∧ id(u) < id(v)};
3 EBBkC-C_Rec(∅, G⃗, k);
4 Procedure EBBkC-C_Rec(S, g⃗, l)
5     Conduct Pruning and Termination (line 3-9 of Algorithm 2);
      /* Branching when l ≥ 3                                          */
6     for each edge u → v in E(g⃗) do
7         S′ ← S ∪ {u, v} and g⃗′ ← g⃗[N⁺({u, v}, g⃗)];
8         if either of the rules of pruning applies then continue;
9         EBBkC-C_Rec(S′, g⃗′, l − 2);
```

introduce additional time cost without achieving better theoretical time complexity. Thus, we choose not to adopt this option. (2) It is worthy noting that for a truss-based edge ordering, there does not always exist a vertex ordering such that the instance of VBBkC with the vertex ordering is equivalent to the instance of EBBkC−T. We include a counter-example in the technical report [42] for illustration.

4.3 EBBkC−C: EBBkC with the Color-based Edge Ordering

While the truss-based edge ordering helps to form sub-branches with small sizes at a branch, it does not offer much power to prune the formed sub-branches - all we can leverage for pruning are some size constraints (line 3 of Algorithm 2). On the other hand, some existing studies of VBBkC have successfully adopted color-based *vertex* ordering for effective pruning [18, 45]. Specifically, consider a branch $B = (S, g, l)$. They first color the vertices in g by iteratively assigning to an uncolored vertex v the smallest color value taken from $\{1, 2, \cdots\}$ that has not been assigned to v's neighbours. Let c be the number of color values used by the coloring procedure. They then obtain a vertex orderng by sorting the vertices in a non-increasing order based on the color values $\langle v_1, v_2, \cdots, v_{|V(g)|}\rangle$ (with ties broken by node ID), i.e., for v_i and v_j with $i < j$, we have $col(v_i) \geq col(v_j)$, where $col(\cdot)$ is the color value of a vertex. As a result, they prune the sub-branch $B_i = (S_i, g_i, l_i)$, which includes v_i to S, if $col(v_i) < l$. The rationale is that since all vertices in g_i have their color values strictly smaller than $col(v_i)$ (according to Eq. (1) and the definition of the color-based vertex ordering), they do not have $l - 1$ different color values, indicating g_i does not contain any $(l-1)$-cliques, and therefore B_i can be pruned.

Recall that in Section 4.1, for an instance of VBBkC with a vertex ordering, there would exist an edge ordering such that the corresponding EBBkC instance based on the edge ordering is equivalent to the VBBkC instance. Motivated by this, we propose to adopt the edge ordering that corresponds to the color-based vertex ordering, which we call *color-based edge ordering*, for our EBBkC. One immediate benefit is that it would naturally inherit the pruning power of the color-based vertex ordering, which has been demonstrated for VBBkC [18, 45]. Furthermore, it would introduce new opportunities for pruning, compared with existing VBBkC with the color-based vertex ordering, since it can leverage the two vertices incident to an edge collectively (instead of a single vertex twice as in VBBkC) for designing new pruning rules, which we explain next.

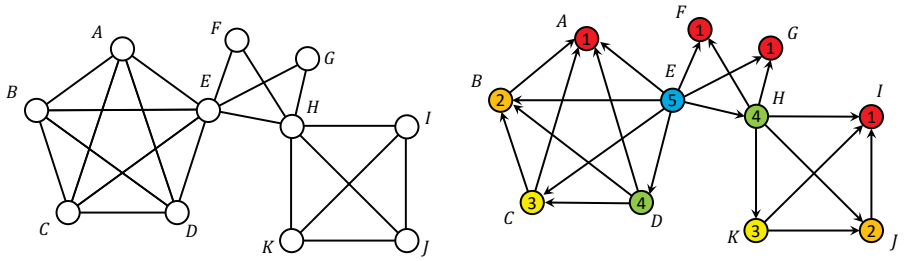

(a) An example graph G.

(b) The DAG built upon G (the numbers inside circles are the color values).

Fig. 2. Color-based edge ordering and pruning rules.

Color-based edge ordering and pruning rules. Consider the branch $B = (S, g, l)$. We first color the graph g using the graph coloring technique [18, 45] and obtain the color-based vertex ordering $\langle v_1, v_2, \cdots, v_{|V(g)|}\rangle$ such that $col(v_i) \geq col(v_j)$ for $i < j$. Then for each vertex $u \in V(g)$, let $id(u)$ be the position of a vertex u in the ordering. For each edge $e = (u, v) \in E(g)$ with $id(u) < id(v)$, we define $str(e)$ as a string, which is the concatenation of $id(u)$ and $id(v)$ (i.e., $str(e) =$'$id(u)$+$id(v)$'). Finally, we define the color-based edge ordering as the alphabetical ordering based on $str(e)$ for edges $e \in E(g)$. That is, one edge $e = (u, v)$ with $id(u) < id(v)$ comes before another edge $e' = (u', v')$ with $id(u') < id(v')$ if (1) $id(u) < id(u')$ or (2) $id(u) = id(u')$ and $id(v) < id(v')$.

Consider a branch $B = (S, g, l)$ and and a sub-branch B_i that includes edge $e_i = (u, v)$ with $col(u) > col(v)$ to S. We can apply the following two pruning rules.

- **Rule (1).** If $col(u) < l$ or $col(v) < l - 1$, we prune sub-branch B_i;
- **Rule (2).** If the vertices in produced sub-branch have less than $l - 2$ distinct color values, we prune sub-branch B_i.

We note that Rule (1) is equivalent to the pruning that VBBkC with the color-based vertex ordering conducts at two branching steps of including u and v [24]. Rule (2) is a new one, which applies only in our EBBkC framework with the color-based edge ordering. In addition, Rule (2) is more powerful than Rule (1) in the sense that if Rule (2) applies, then Rule (1) applies, but not vice versa. The reason is that the color values of u and v are sometimes much larger than the number of distinct color values in the sub-branch since both color values consider the information of their own neighbors instead of their common neighbors. For illustration, consider the example in Figure 2 and assume that we aim to list 4-cliques (i.e., $k = 4$). We focus on the edge EH. It is easy to check that Rule (1) does not apply, but Rule (2) applies since the vertices in the produced sub-branch, i.e., F and G, have only one color value. Given that it takes $O(1)$ time to check if Rule (1) applies and $O(|V(g_i)|)$ time to check if Rule (2) applies, our strategy is to check Rule (1) first, and if it does not apply, we further check Rule (2). We remark that Rule (2) can be adapted to some VBBkC based algorithms including DDegCol [24] and BitCol [46], that is, we prune sub-branch B_i, if the vertices inside have less than $l - 1$ distinct color values.

The EBBkC−C algorithm. When the color-based edge ordering is adopted in EBBkC, we call the resulting algorithm EBBkC−C. The pseudo-code of EBBkC−C is presented in Algorithm 4. With the color-based edge ordering, a directed acyclic graph (DAG), denoted by $\overrightarrow{G}$, is built for efficiently conducting the branching steps [24, 46]. Specifically, $\overrightarrow{G}$ is built upon G by orienting each edge (u, v) in E with $id(u) < id(v)$ from u to v (line 2). For illustration, consider Figure 2. Given $V_{sub} \subseteq V$, let $N^+(V_{sub}, \overrightarrow{g})$ be the common out-neighbours of the vertices in V_{sub} in $\overrightarrow{g}$. Then, given a branch $B = (S, \overrightarrow{g}, l)$[6], the edge-oriented branching at a branch with the color-based ordering can be easily

[6]Note that $\overrightarrow{g}$ is a subgraph of $\overrightarrow{G}$, whose edge orientations are the same as those of $\overrightarrow{G}$.

Algorithm 5: EBBkC with hybrid edge ordering: EBBkC−H

Input: A graph $G = (V, E)$ and an integer $k \geq 3$
Output: All k-cliques within G
/* Initialization and branching at $(\emptyset, G, k)$ */
1 $\pi_\tau(G) \leftarrow$ the truss-based ordering of edges in G;
2 **for** *each edge* $e_i \in E(G)$ *following* $\pi_\tau(G)$ **do**
3 Obtain S_i and g_i according to Eq. (2);
4 Do vertex coloring on g_i and get $id(v)$ for each vertex in $V(g_i)$;
5 $\overrightarrow{g_i} \leftarrow (V(g_i), \{u \rightarrow v \mid (u, v) \in E(g_i) \wedge id(u) < id(v)\})$;
6 EBBkC-C_Rec$(S_i, \overrightarrow{g_i}, k-2)$;

conducted by calculating the common out-neighbors of the vertices incident to an edge in $\overrightarrow{g}$ (line 7). Then, we will prune the produced branch if either of the above two rules applies (line 8).

Complexity. The time cost of EBBkC−C is $O(km(\Delta/2)^{k-2})$, which we show in the following theorem. In practice, EBBkC−C runs faster than EBBkC−T.

Theorem 4.3. *Given a graph $G = (V, E)$ and an integer $k \geq 3$, the time and space complexities of* EBBkC−C *are $O(km(\Delta/2)^{k-2})$ and $O(m+n)$, respectively, where Δ is the maximum degree of G.*

Proof. The proof is similar to that of Theorem 4.2. The difference is that the largest size of the produced graph instance in EBBkC−C can only be bounded by Δ and the time of generating the color-based edge ordering is $O(m)$, which is dominated by the cost of the recursive procedure. □

4.4 EBBkC−H: EBBkC with Hybrid Edge Ordering

Among EBBkC−T and EBBkC−C, the former has a better theoretical time complexity and the latter enables effective pruning in practice. Inspired by the hybrid vertex ordering used for DDegCol [24] and BitCol [46], we aim to achieve the merits of both algorithms by adopting both the truss-based edge ordering (used by EBBkC−T) and the color-based edge ordering (used by EBBkC−C) in the EBBkC framework. Specifically, we first apply the truss-based edge ordering for the branching step at the initial branch $(\emptyset, G, k)$. Then, for the following branches, we adopt the color-based ordering for their branching steps. We call this algorithm based on the hybrid edge ordering EBBkC−H. The pseudo-code of EBBkC−H is presented in Algorithm 5 and the implementations of branching steps are similar to those in EBBkC−T and EBBkC−C. The size of a produced problem instance for EBBkC−H is bounded by τ (due to the branching at the initial branch based on the truss-based edge ordering), and thus EBBkC−H achieves the same time complexity as EBBkC−T. In addition, EBBkC−H enables effective pruning for all branches except for the initial branch (since color-based edge coloring is adopted at these branches), and thus it runs fast in practice as EBBkC−C does.

Complexity. The time complexity of EBBkC−H is $O(\delta m + km(\tau/2)^{k-2})$, which is the same as that of EBBkC−T.

Theorem 4.4. *Given a graph $G = (V, E)$ and an integer $k \geq 3$, the time and space complexities of* EBBkC−H *are $O(\delta m + km(\tau/2)^{k-2})$ and $O(m+n)$, respectively, where $\tau = \tau(G)$ is strictly smaller than the degeneracy δ of G.*

Proof. The proof is similar to that of Theorem 4.2. Since the largest size of the produced graph instance in EBBkC−H can also be bounded by τ, it has the same worst-case time complexity as that of EBBkC−T. □

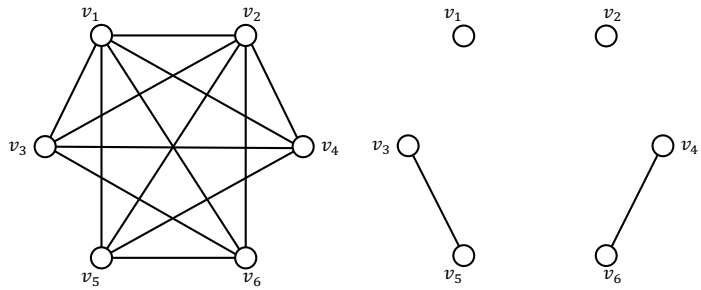

(a) An example 2-plex graph g. (b) Inverse graph g_{inv} of 2-plex g.

Fig. 3. Examples of a 2-plex and its inverse graph.

4.5 Other Potential Applications of EBBkC

Our EBBkC framework can potentially be applied to other problems than the k-clique listing problem, which we discuss as follows. First, our framework can be easily adapted to solve other clique mining tasks, including maximal clique enumeration (MCE) [13, 21, 27, 38], maximum clique search (MCS) [7, 8] and diversified top-k clique search (DCS) [43, 44]. The rationale is that our framework can explore all possible cliques in an input graph and thus can output only the desired cliques that satisfy some properties (e.g., maximality and diversity) by filtering out others. Second, our framework can be potentially extended to mining other types of *connected* dense subgraphs, e.g., connected k-plex. This is because our framework can be used to explore all possible subsets of edges by recursively including an edge, and thus the induced subgraphs will cover all possible connected dense subgraphs. Third, there are some potential benefits when adapting our framework to the above tasks. As discussed in Section 4.1, the edge-oriented branching can provide more information (i.e., an edge involving two vertices instead of one vertex) towards designing more effective pivot techniques and/or pruning rules than the existing vertex-oriented branching.

5 EARLY TERMINATION TECHNIQUE

Suppose that we are at a branch $B = (S, g, l)$, where graph g is dense (e.g., g is a clique or nearly a clique), and the goal is to list l-cliques in g (and merge them with S). Based on the EBBkC framework, we would conduct branching at branch B and form sub-branches. Nevertheless, since g is dense, there would be many sub-branches to be formed (recall that we form $|E(g)|$ sub-branches), which would be costly. Fortunately, for such a branch, we can list the l-cliques efficiently without continuing the recursive branching process of EBBkC, i.e., we can *early terminate* the branching process. Specifically, we have the following two observations.

- If g is a clique or a 2-plex (recall that a t-plex is a graph where each vertex inside has at most t non-neighbors including itself, we can list l-cliques in g efficiently in a combinatorial manner. For the former case, we can directly enumerate all possible sets of l vertices in g. For the latter case, we can do similarly, but in a bit more complex manner (details will be discussed in Section 5.1).
- If g a t-plex (with $t \geq 3$), the branching procedure based on g can be converted to that on its inverse graph g_{inv} (recall that g_{inv} has the same set of vertices as g, with an edge between two vertices in g_{inv} if and only if they are disconnected from each other in g), which is sparse, and the converted procedure would run faster (details will be discussed in Section 5.2).

We determine whether g is a t-plex for some t by checking the minimum degree of a vertex in g - if it is no less than $|V(g)| - t$, g is a t-plex; otherwise, g is not a t-plex. This can be done while constructing the corresponding branch $B = (S, g, l)$ in $O(V(g))$ time.

Algorithm 6: List k-cliques in a 2-plex: kC2Plex

Input: A branch (S, g, l) with g corresponding to a 2-plex
Output: All k-cliques within (S, g, l)

1 Partition $V(g)$ into three disjoint sets F, L and R;
2 **if** $|F| + |L| < l$ **then return**;
3 **for** $c_1 \in [\max\{0, l - |L|\}, \min\{l, |F|\}]$ *and each* c_1*-combination* F_{sub} *over* F **do**
4 **for** $c_2 \in [0, \min\{l - c_1, |L|\}]$ *and each* c_2*-combination* L_{sub} *over* L **do**
5 **for** $c_3 \leftarrow l - c_1 - c_2$ *and each* c_3*-combination* R_{sub} *over* $R \setminus \overline{N}(L_{sub}, g)$ **do**
6 **Output** a k-clique $S \cup F_{sub} \cup L_{sub} \cup R_{sub}$;

5.1 Listing k-Cliques from 2-Plex in Nearly Optimal Time

Consider a branch $B = (S, g, l)$ with g as a 2-plex. The procedure for listing k-cliques inside, called kC2Plex, utilizes the *combinatorial technique*. The rationale behind is that listing k-cliques from a large clique can be solved in the optimal time by directly enumerating all possible combinations of k vertices.

Specifically, we first partition $V(g)$ into three disjoint sets, namely F, L and R, each of which induces a clique. This can be done in two steps. First, it partitions $V(g)$ into two disjoint parts: one containing those vertices that are adjacent to all other vertices in $V(g)$ (this is F) and the other containing the remaining vertices that are not adjacent to two vertices including itself (this is $L \cup R$). Note that $L \cup R$ always involves an even number of vertices and can be regarded as a collection of pairs of vertices $\{u, v\}$ such that u is disconnected from v. Second, it further partitions $L \cup R$ into two parts by breaking each pair in $L \cup R$, that is, L and R contain the first and the second vertex in each pair, respectively. As a result, every vertex in one set connects all others within the same set and is disconnected from one vertex from the other set. Note that the partition of $L \cup R$ is not unique and can be an arbitrary one. For illustration, consider the example in Figure 3(a) and 3(b). The vertices v_1 and v_2 are adjacent to all other vertices. Thus, $F = \{v_1, v_2\}$ and $L \cup R = \{v_3, v_4, v_5, v_6\}$. One possible partition is $L = \{v_3, v_4\}$ and $R = \{v_5, v_6\}$. Below, we elaborate on how the partition $V(g) = F \cup L \cup R$ helps to speedup the k-clique listing.

Recall that the set of k-cliques in B can be listed by finding all l-cliques in g and merging each of them with S. Consider a l-clique in g. Based on the partition $V(g) = F \cup L \cup R$, it consists of three disjoint subsets of F, L and R, namely F_{sub}, L_{sub} and R_{sub}, each of which induces a small clique. Therefore, all l-cliques with the form of $F_{sub} \cup L_{sub} \cup R_{sub}$ can be found by iteratively enumerating all possible $|F_{sub}|$-combinations over F, $|L_{sub}|$-combinations over L and $|R_{sub}|$-combinations over R such that $|F_{sub}| + |L_{sub}| + |R_{sub}| = l$.

We present the pseudo-code of kC2Plex in Algorithm 6. In particular, the integers c_1, c_2 and c_3 are used to ensure the satisfaction of $|F_{sub}| + |L_{sub}| + |R_{sub}| = l$. Specifically, it first finds a c_1-clique F_{sub} from F and a c_2-clique L_{sub} from L. Recall that every vertex in L is disconnected from one vertex in R and vice versa. Hence, it removes from R those vertices that is disconnected from one vertex in L_{sub}, which we denote by $\overline{N}(L_{sub}, g)$, and this can be done efficiently in $\Theta(|L_{sub}|)$ (as verified by Theorem 5.1), and then finds a c_3-clique from $R \setminus \overline{N}(L_{sub}, g)$. Besides, when $|F| + |L| < l$, it terminates the procedure since no l-clique will be found in g (line 2).

Time complexity. We analyze the time complexity of kC2Plex as follows.

Theorem 5.1. *Given a branch $B = (S, g, l)$ with g being a 2-plex,* kC2Plex *lists all k-cliques within B in $O(|E(g)| + k \cdot c(g, l))$ time where $c(g, l)$ is the number of l-cliques in g.*

Algorithm 7: List k-cliques in a t-plex ($t \geq 3$): kCtPlex

Input: A branch (S, g, l) with g corresponding to a t-plex
Output: All k-cliques within (S, g, l)

```
1  Construct the inverse graph g_inv of g;
2  I ← set of vertices in V(g_inv) that are disconnected from all others;
3  kCtPlex_Rec(S, V(g_inv) \ I, l)
4  Procedure kCtPlex_Rec(S', C, l')
       /* Termination when l' = 0                                          */
5      if l' = 0 then
6          Output a k-clique S';
7          return;
       /* Choose all rest l' vertices from I                               */
8      if |I| ≥ l' then
9          for each l'-combination I_sub over I do
10             Output a k-clique S' ∪ I_sub;
       /* Choose at least one vertex from C                                */
11     for each v_i ∈ C do
12         Create a branch (S_i, C_i, l_i) based on Eq. (9);
13         if |C_i| + |I| ≥ l_i then kCtPlex_Rec(S_i, C_i, l_i);
```

Proof. Algorithm 6 takes $O(|E(g)|)$ for partitioning $V(g)$ by obtaining the degree of each vertex inside (line 1). For each round of lines 3-6, the algorithm can guarantee exactly one k-clique to be outputted at line 6 based on the settings of c_1, c_2 and c_3. Besides, the operation $R \setminus \overline{N}(L_{sub}, g)$ only takes $\Theta(|L_{sub}|)$ time. Specifically, we (1) maintain two arrays $L = \{u_1, u_2, \cdots\}$ and $R = \{v_1, v_2, \cdots\}$ such that u_i is not adjacent to v_i for $1 \leq i \leq |L|$, and (2) reorder R by switching $|L_{sub}|$ vertices with the same indices as those in L_{sub} to the tail of R which runs in $\Theta(|L_{sub}|)$ time, and take the first $|R| - |L_{sub}|$ vertices in R as $R \setminus \overline{N}(L_{sub}, g)$. □

Remark. We remark that kC2Plex achieves the *input-output sensitive* time complexity since the time cost depends on both the size of input $|E(g)|$ and the number of k-cliques within the branch. Besides, we note that $O(k \cdot c(g, k))$ is the optimal time for listing k-cliques within the branch and thus kC2Plex only takes extra $O(|E(g)|)$ time, i.e., kC2Plex is *nearly optimal*.

5.2 Listing k-Cliques from t-Plex with $t \geq 3$

Consider a branch $B = (S, g, l)$ with g as a t-plex with $t \geq 3$. The procedure for listing k-clique inside, called kCtPlex, differs in the way of branching (i.e., forming new branches). Specifically, it branches based on the inverse graph g_{inv} instead of g. The rationale is that since g is a t-plex and tends to be dense, its inverse graph g_{inv} would be sparse. As a result, branching on g_{inv} would run empirically faster. Below, we give the details.

Specifically, it maintains the inverse graph g_{inv} of g and represents a branch (S, g, l) by the new form of (S, C, l) where C is the set of vertices in g, i.e., $C = V(g)$. We note that the new form omits the information of edges in g, which is instead stored in g_{inv}. Consider the branching step of kCtPlex at a branch (S, C, l). Let $\langle v_1, v_2, \cdots, v_{|C|} \rangle$ be an arbitrary ordering of vertices in C. Then, the branching step would produce $|C|$ new sub-branches. The i-th branch, denoted by $B_i = (S_i, C_i, l_i)$, includes v_i to S and excludes $\{v_1, v_2, .., v_i\}$ from C. Formally, for $1 \leq i \leq |C|$, we have

$$S_i = S \cup \{v_i\},\ C_i = C \setminus \{v_1, v_2, \cdots, v_i\} \setminus N(v_i, g_{inv}),\ l_i = l - 1. \tag{9}$$

Table 1. Dataset Statistics.

Graph (Name)	$\|V\|$	$\|E\|$	Δ	δ	τ	ω
nasasrb (NA)	54,870	1,311,227	275	35	22	24
fbwosn (FB)	63,731	817,090	2K	52	35	30
wikitrust (WK)	138,587	715,883	12K	64	31	25
shipsec5 (SH)	179,104	2,200,076	75	29	22	24
socfba (SO)	3,097,165	23,667,394	5K	74	29	25
pokec (PO)	1,632,803	22,301,964	15K	47	27	29
wikicn (CN)	1,930,270	8,956,902	30K	127	31	33
baidu (BA)	2,140,198	17,014,946	98K	82	29	31
websk (WE)	121,422	334,419	590	81	80	82
citeseer (CI)	227,320	814,134	1K	86	85	87
stanford (ST)	281,904	1,992,636	39K	86	61	61
dblp (DB)	317,080	1,049,866	343	113	112	114
dielfilter (DE)	420,408	16,232,900	302	56	43	45
digg (DG)	770,799	5,907,132	18K	236	72	50
skitter (SK)	1,696,415	11,095,298	35K	111	67	67
orkut (OR)	2,997,166	106,349,209	28K	253	74	47
allwebuk (UK)	18,483,186	261,787,258	3M	943	942	944
clueweb (CW)	147,925,593	446,766,953	1M	192	83	56
wikipedia (WP)	25,921,548	543,183,611	4M	1120	426	428

Note that we need to remove from C those vertices that are not adjacent to v_i in g (since they cannot form any k-clique with v_i) and they connect to v_i in g_{inv}. Clearly, all k-cliques in (S, C, l) will be listed exactly once after branching. We note that the branching strategy we used in `kCtPlex` differs from that for `EBBkC` (and that for `VBBkC`). Specifically, the former (resp. the latter) is based on a sparse inverse graph g_{inv} (resp. a dense t-plex g) and maintains a vertex set C_i (resp. a graph instance g_i) for the produced branches. We remark that the former (correspondingly, the early stop strategy with t at least 3) runs faster than the latter in practice, as verified in our experiments.

We summarize the procedure `kCtPlex` in Algorithm 7. In particular, it also utilizes the combinatorial technique for boosting the performance (lines 8-10). Specifically, it figures out the set of vertices, denoted by I, in g_{inv}, which are disconnected from all others (line 2). Consider a k-clique in B. It may involve c vertices in I where $c \in [0, \min\{|I|, k\}]$. Hence, we remove from $V(g_{inv})$ those vertices in I for branching at line 3 while adding them back to a k-clique at lines 8-10. It is not difficult to verify the correctness. Due to the page limit, we include the time complexity analysis of `kCtPlex` in the technical report [42].

Remark. (1) With the early termination strategy, the BB algorithms retain the same time complexity provided before but run practically faster as verified in the experiments. (2) The early termination strategy is supposed to set a small threshold of t so as to apply the alternative procedures only on dense graph instances (i.e., t-plexes). We test different choices of t in the experiments; the results suggest that `EBBkC` with t being set from 2 to 5 runs comparably faster than other choices while the best one among them varies for different settings of k.

6 EXPERIMENTS

6.1 Experimental Setup

Datasets. We use 19 real datasets in our experiments, which can be obtained from an open-source network repository [33]. For each graph, we ignore the directions, weights and self-loops (if any) at

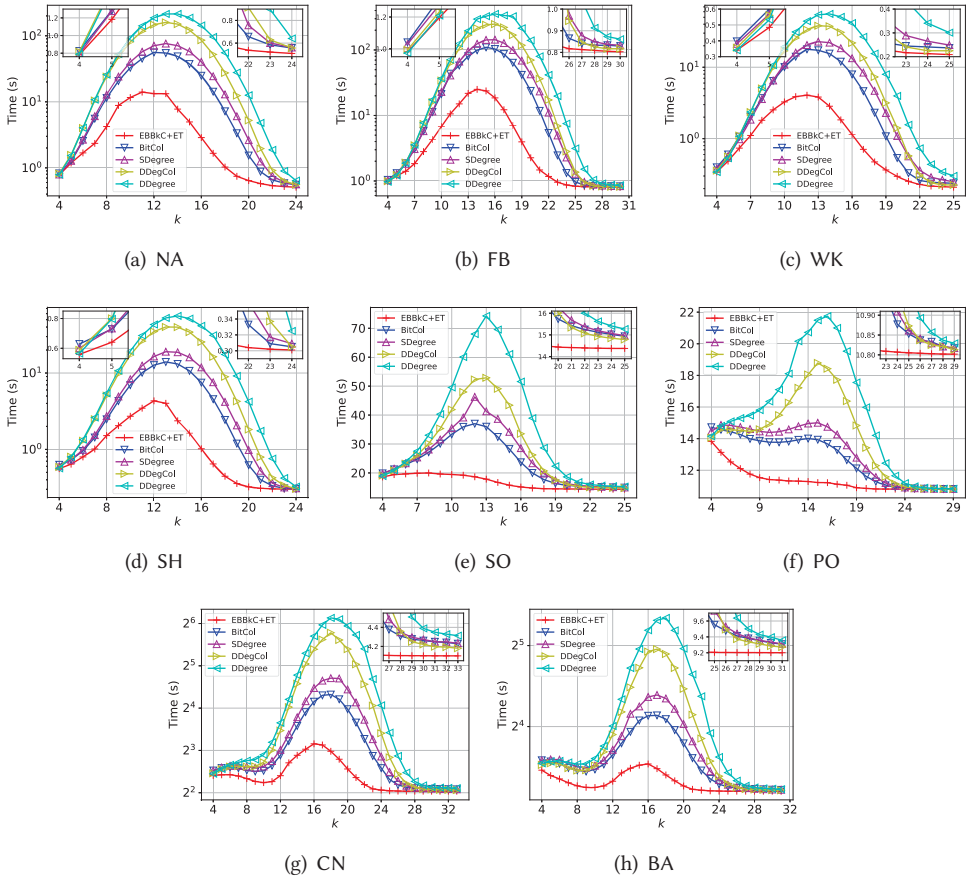

(a) NA (b) FB (c) WK

(d) SH (e) SO (f) PO

(g) CN (h) BA

Fig. 4. Comparison with baselines on the small-ω graphs, varying k from 4 to ω.

the very beginning. Following the existing study [24], we divide the real datasets into two groups based on the size of a maximum clique ω: small-ω graphs and large-ω graphs. For small-ω graphs, we list all k-cliques for all k, while for large-ω graphs, we only list k-cliques for small k values and large k values which are near ω. We collect the graph statistics and report the maximum degree Δ, the degeneracy number δ, the truss related number τ and the maximum clique size ω, which are shown in Table 1. We select four datasets, namely WK, PO, ST and OR, which are bold in the table, as default ones since they cover different size of graphs.

Baselines and Metrics. We choose `EBBkC-H` as the default edge-oriented branching-based BB framework for comparison, and denote it by `EBBkC` for brevity in the experimental results. We compare our algorithm `EBBkC+ET` with four existing algorithms, namely `DDegCol` [24], `DDegree` [24], `SDegree` [46] and `BitCol` [46] in terms of running time. Specifically, `EBBkC+ET` employs the edge-oriented branching-based BB framework with hybrid edge ordering, and applies the early-termination technique. The running times of all algorithms reported in our paper include the time costs of (1) conducting the pre-processing techniques (if any) and (2) generating orderings of vertices. For our early termination technique, we set the parameter $t = 2$ when $k \leq \tau/2$ and $t = 3$ when $\tau/2 < k \leq \omega$ for early-termination. All baselines are the state-of-the-art algorithms for k-cliques listing with vertex-oriented branching-based BB framework. Following the existing study [24], we vary k from 4 since k-clique listing problem reduces to triangle listing problem when $k = 3$ and there are efficient algorithms [23, 28] for triangle listing which run in polynomial time.

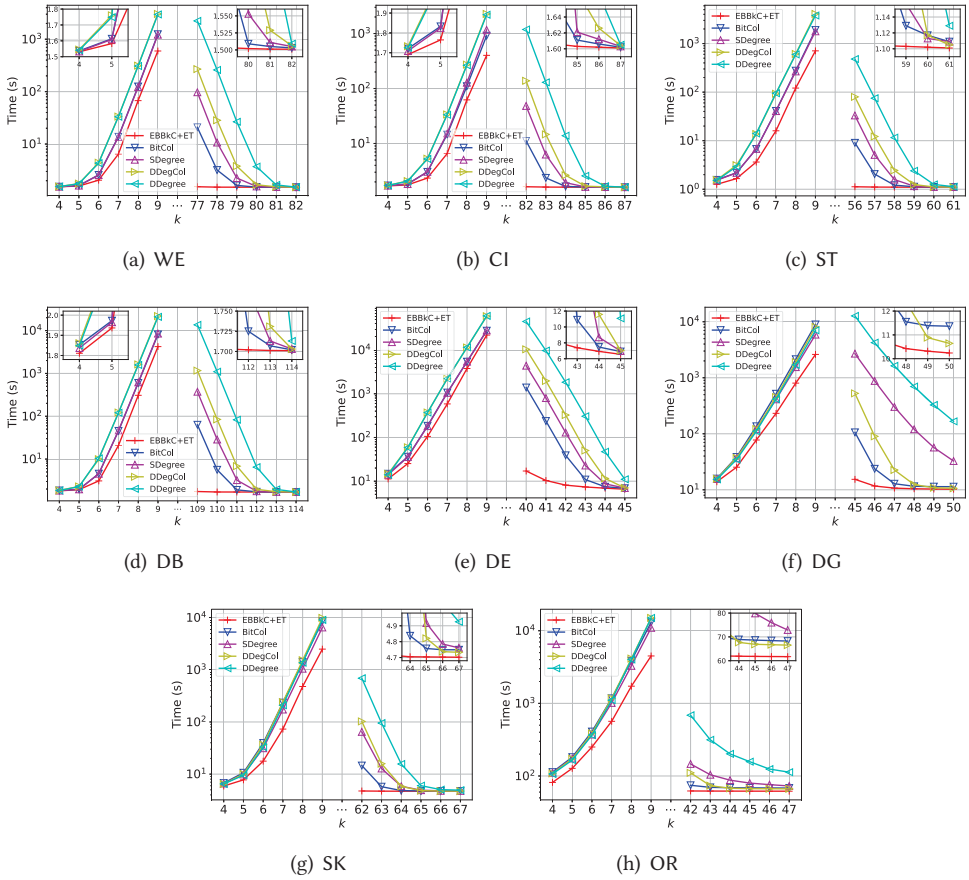

Fig. 5. Comparison with baselines on the large-ω graphs, varying k from 4 to 9 and from $\omega - 5$ to ω.

Settings. The source codes of all algorithms are written in C++ and the experiments are conducted on a Linux machine with a 2.10GHz Intel CPU and 128GB memory. We note that we have not utilized SIMD instructions for data-level parallelism in our implementation, despite the potential to further enhance the acceleration of our algorithm. We set the time limit as 24 hours (i.e., 86,400 seconds) and the running time of any algorithm that exceeds the time limit is recorded with "INF". Implementation codes and datasets can be found via this link https://github.com/wangkaixin219/EBBkC.

6.2 Experimental Results

(1) Comparison among algorithms (on small-ω graphs). Figure 4 shows the results of listing k-cliques on small-ω graphs. We observe that EBBkC+ET (indicated with red lines) runs faster than all baselines on all datasets. This is consistent with our theoretical analysis that the time complexity of EBBkC+ET is better than those of the baseline algorithms. Besides, on PO, CN and BA, we observe that the running time of EBBkC+ET first decreases when k is small. There are two possible reasons: (1) on some dataset, the number of k-cliques decreases as k increases when k is small. On BA, for example, the number of 4-cliques ($k = 4$) is nearly 28M while the number of 7-cliques ($k = 7$) is 21M; (2) as k increases, a larger number of branches can be pruned by the size constraint with the truss-based edge ordering, which provides the opportunity to run faster. To see this, we collect the

Table 2. Time for generating truss-based edge ordering and degeneracy ordering (unit: sec).

	WK	PO	ST	OR
Truss (s)	0.2	10.7	1.1	60.4
Degen. (s)	0.1	7.3	0.6	53.3

number of promising branches after the branching step with the truss-based edge ordering. On PO, there are 11M branches left when enumerating 4-cliques ($k = 4$) while there are only less than 1M branches left when enumerating 10-cliques ($k = 10$).

(2) Comparison among algorithms (on large-ω graphs). Figure 5 shows the results of listing k-cliques on large-ω graphs. We omit the results for some values of k since the time costs are beyond 24 hours due to the large number of k-cliques. We find that EBBkC+ET still runs the fastest. It is worthy noting that EBBkC+ET can greatly improve the efficiency by 1-2 orders of magnitude over the baselines when k is near the size of a maximum clique ω, e.g., EBBkC+ET runs 9.2x and 97.7x faster than BitCol on DB (when $k = 109$) and on DE (when $k = 40$), respectively. The reasons are two-fold: (1) when k is near ω, a large number of branches can be pruned by the size constraint with the truss-based edge ordering, and as a result, the remaining branches are relatively dense; (2) EBBkC+ET can quickly enumerate cliques within a dense structure, e.g., t-plex, without making branches for the search space, which dramatically reduces the running time.

(3) Ablation studies. We compare two variants of our method, namely EBBkC+ET (the full version) and EBBkC (the full version without the early termination), with two VBBkC algorithms, namely DDegCol+ (the DDegCol algorithm with Rule (2) proposed in this paper) and BitCol+ (the BitCol algorithm with Rule (2) and without SIMD-based implementations). We note that DDegCol+ and BitCol+ correspond to the SOTA VBBkC algorithms without SIMD-based implementations. We report the results in Figure 6 and have the following observations. First, DDegCol+ and BitCol+ have very similar running times, which shows that pre-processing techniques in the BitCol paper are not effective. Second, EBBkC runs clearly faster than DDegCol+ and BitCol+, which demonstrates the clear contribution of our edge-oriented BB framework (note that EBBkC and DDegCol+ differ only in their frameworks). We note that it is not fair to compare EBBkC with BitCol directly since the latter is based on SIMD-based implementations while the former is not. Third, EBBkC+ET runs clearly faster than EBBkC, which demonstrates the clear contribution of the early termination technique. For example, on dataset WK, the edge-oriented BB framework and early-termination contribute 28.5% and 71.5% to the efficiency improvements over BitCol+ when $k = 13$, respectively. We also show the time costs of generating the truss-based ordering for edges in EBBkC and those of generating the degeneracy ordering for vertices in VBBkC in Table 2. We note that while the former are slightly larger than the latter, the overall time cost of EBBkC is smaller than that of VBBkC (as shown in Figure 6).

(4) Effects of the edge ordering (comparison among EBBkC−T, EBBkC−C and EBBkC−H). For the sake of fairness, we employ all color-based pruning rules for EBBkC−C and EBBkC−H frameworks and employ the early-termination technique for all frameworks. The results are shown in Figure 7. Consider EBBkC−H and EBBkC−T. Although both frameworks have the same time complexity, EBBkC−H runs much faster since it can prune more unpromising search paths by color-based pruning rules than EBBkC−T. Consider EBBkC−H and EBBkC−C. EBBkC−H outperforms EBBkC−C since the largest sub-problem instances produced by EBBKC−H is smaller than that of EBBkC−C, which also conforms our theoretical analysis.

(5) Effect of the color-based pruning rules. Recall that in Section 4.3, we introduce two color-based pruning rules, where the first rule is adapted from the existing studies [24] and the second rule is newly proposed in this paper. Therefore, we study the effect of the second pruning rule

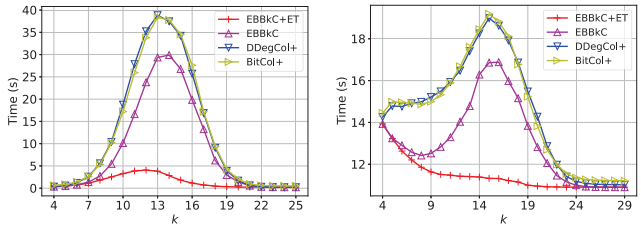

(a) WK (b) PO

Fig. 6. Ablation studies.

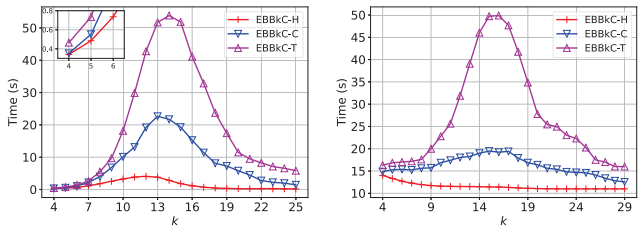

(a) WK (b) PO

Fig. 7. Effects of the edge ordering.

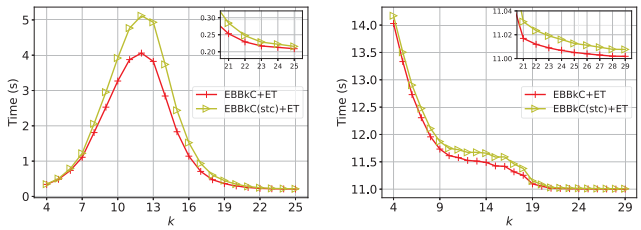

(a) WK (b) PO

Fig. 8. Effects of the pruning rules (comparison between the algorithms w/ and w/o the Rule (2)).

by making comparison between the running time of the algorithms with and without this rule, respectively. We denote the algorithm without this rule by EBBkC(stc)+ET. The results are shown in Figure 8, indicated by red lines and yellow lines. We observe that the second pruning rule brings more advantages as k increases. A possible reason is that when k is small, the graph instance g usually has more than l colors ($l \leq k$), which cannot be pruned by the second rule while as k increases, sparse graph instances can easily violate the rule and they can be safely pruned.

(6) Effects of early-termination technique (varying t). We study the effects of choosing different parameter t (in t-plex) for early-termination. We vary t in the range $\{1, 2, 3, 4, 5\}$ and report the corresponding running time of EBBkC+ET under different values of k. The results are shown in Figure 9. We have the following observations. First, for all values of k, EBBkC+ET with $t = 2$ always runs faster than that with $t = 1$, since we can list all k-cliques inside a t-plex in optimal time when $t = 1$ and $t = 2$ but we can early-terminate the recursion EBBkC_Rec in an earlier phase when $t = 2$ than that when $t = 1$ (Section 5.1). Second, when the value of k is small, EBBkC+ET with smaller t runs faster while when the value of k is large, EBBkC+ET with larger t runs faster. On WK, for example, when $k = 8$, EBBkC+ET with $t = 2$ runs the fastest; when $k = 12$ and $k = 16$,

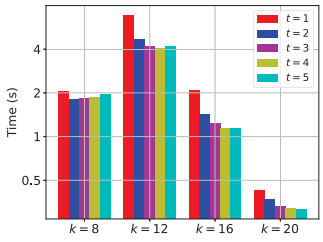

(a) Results on WK varying k

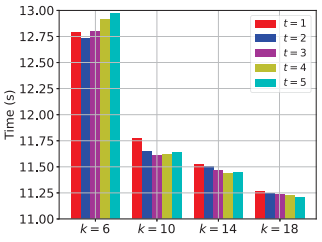

(b) Results on PO varying k

Fig. 9. Effects of early-termination technique (varying t).

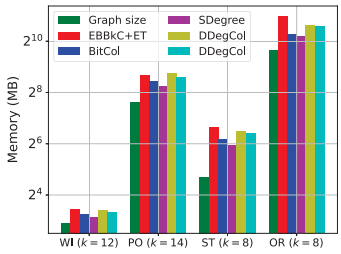

Fig. 10. Results of space costs.

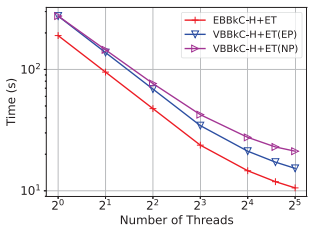

(a) Results on ST ($k = 8$)

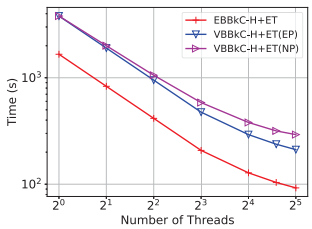

(b) Results on OR ($k = 8$)

Fig. 11. Comparison among different parallel schemes, varying the number of threads.

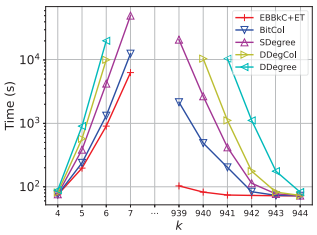

(a) UK

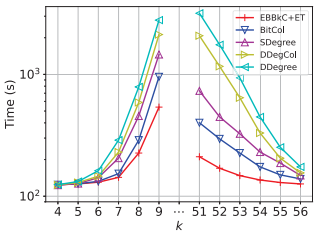

(b) CW

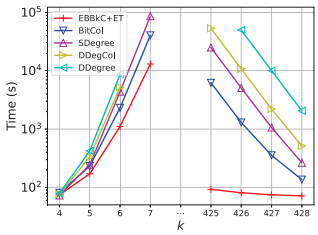

(c) WP

Fig. 12. Results on scalability test.

EBBkC+ET with $t = 4$ runs the fastest; and when $k = 20$, EBBkC+ET with $t = 5$ runs the fastest. This phenomenon conforms our theoretical analysis in Section 5 that when the value of k increases, listing k-cliques with early-termination on sparser plexes, i.e., t-plex with larger t, can also be efficient.

(7) Parallelization. We compare different algorithms in a parallel computing setting. Specifically, for EBBkC framework, since each produced sub-branch produced from $B = (\emptyset, G, k)$ can be solved independently (see line 6 of Algorithm 5), we process all of such sub-branches in parallel. For existing studies under VBBkC framework [12, 24, 46], there are two parallel schemes. One is called NodeParallel (NP for short). This strategy processes each produced sub-branch at $B = (\emptyset, G, k)$ in parallel since they are independent (see line 9 of Algorithm 1). The other is called EdgeParallel (EP for short). This strategy first aggregates the first two consecutive branching steps at $B = (\emptyset, G, k)$ as a unit, and as a result, it would produce $|E(G)|$ sub-branches, then it processes each sub-branch in parallel since they are independent [12, 24, 46]. The results of the comparison are shown in Figure 11. Consider the comparison between VBBkC+ET(NP) and VBBkC+ET(EP), indicated by blue

lines and magenta lines. The algorithm with edge-level parallelization strategy achieves a higher degree of parallelism since it would produce a number of problem instances with smaller and similar scales, which balances the computational loads across the threads. Then, we consider the comparison between EBBkC+ET and VBBkC+ET(EP), indicated by red lines and blue lines. Both algorithms can be regarded as adopting edge-level parallelization strategy but differ in the ordering of the edges. We observe that EBBkC+ET runs faster than VBBkC+ET(EP). The reason is that EBBkC uses truss-based edge ordering, which would produce even smaller problem instances than those of VBBkC+ET(EP).

(8) Space costs. We report the space costs of different algorithms in Figure 10. We have the following observations. First, the space costs of all algorithms are comparable and a few times larger than the graph size, which is aligned with the space complexity of $O(n + m)$. Second, EBBkC+ET has the space cost slightly larger than those of others since it employs extra data structures for maintaining edge ordering and conducting early-termination.

(9) Scalability test. We test the scalability of the algorithms on three large graphs under the parallel setting with 48 threads and report the results in Figure 12. All algorithms use the EP parallel scheme. EBBkC+ET outperforms other baselines consistently. In particular, on the largest graph WP, EBBkC+ET is up to about 100× faster than BitCol (when $k = 425$).

7 RELATED WORK

Listing k-cliques for arbitrary k values. Existing exact k-clique listing algorithms for arbitrary k values can be classified into two categories: backtracking based algorithms [10] and branch-and-bound based algorithms [24, 45, 46]. Specifically, the algorithm Arbo [10] is the first practical algorithm for listing all k-cliques, whose time complexity (we focus on the worst-case time complexies in this paper) is $O(km\alpha^{k-2})$, where m and α are the number of edges and the arboricity of the graph, respectively. However, it is difficult to paralleize Arbo since it involves a depth-first backtracking procedure. To solve this issue, several (vertex-oriented branching-based) branch-and-bound (BB) based algorithms are proposed, including Degree [16, 28], Degen [12], DegenCol [24], DegCol [24], DDegCol [24], DDegree [24], SDegree [46] and BitCol [46]. As introduced in Section 3, these algorithms follow the same framework but differ in how the vertex orderings are adopted and some implementation details. Specifically, Degree uses a global degree ordering, with which the size of the largest problem instance produced can be bounded by η (i.e., the h-index of the graph). Degree has a time complexity of $O(km(\eta/2)^{k-2})$. Degen uses a global degeneracy ordering, with which the size of the largest problem instance produced can be bounded by δ (i.e., the degeneracy of the graph). Degen has a time complexity of $O(km(\delta/2)^{k-2})$. Since it has been proven that $\delta \leq 2\alpha - 1$ [47], Degen is the first algorithm that outperforms Arbo theoretically and practically. However, Degen suffers from the issue that it cannot efficiently list the clique whose size is near ω (i.e., the size of a maximum clique). To solve this issue, the authors in [24] propose several algorithms, which are based on color-based vertex orderings. DegCol and DegenCol first color the graph with some graph coloring algorithms (e.g., inverse degree based [45] and inverse degeneracy based [18]) and generate the an ordering of vertices based on the color values of the vertices. While the color values of the vertices can significantly prune the unpromising search paths and make the algorithm efficient to list near-ω cliques, DegCol and DegenCol both have the time complexity of $O(km(\Delta/2)^{k-2})$, where Δ is the maximum degree of a vertex in the graph since the color-based ordering cannot guarantee a tighter size bound for the produced problem instance. To overcome this limitation, DDegCol adopts a hybrid ordering, where it first uses degeneracy ordering to branch the universal search space such that the size of each produced problem instance is bounded by δ, then it uses color-based orderings to branch each produced sub-branch. Following a similar

procedure, DDegree also combines degeneracy ordering and degree ordering to branch the universal search space and the sub-spaces, respectively. In this way, DDegCol and DDegree have the time complexity of $O(km(\delta/2)^{k-2})$. BitCol and SDegree implement DDegCol and DDegree in a more efficient way with SIMD instructions, respectively, and retain the same time complexity as that of DDegCol and DDegree. In contrast, our EBBkC algorithm is an edge-oriented branching-based BB algorithm based on edge orderings and has its time complexity (i.e., $O(m\delta + km(\tau/2)^{k-2})$) better than that of the state-of-the-art algorithms including DDegCol, DDegree, BitCol and SDegree (i.e., $O(km(\delta/2)^{k-2})$), where τ is strictly smaller than δ. Some other algorithms, e.g., MACE [26], which are originally designed for the maximal clique enumeration problem, can also be adapted to listing k-cliques problem [26, 36, 37]. These adapted algorithms are mainly based on the well-known Bron-Kerbosch (BK) algorithm [6] with some size constraints to ensure that each clique to be outputted has exactly k vertices. However, these adapted algorithm are even less efficient than Arbo theoretically and practically, e.g., MACE has a time complexity of $O(kmn\alpha^{k-2})$ [24, 36], and cannot handle large real graphs.

Listing k-cliques for special k values. There are two special cases for k-clique listing problem: when $k = 3$ and when $k = \omega$. When $k = 3$, the problem reduces to triangle listing problem [10, 23, 28]. The state-of-the-art algorithm for triangle listing problem follows a vertex ordering based framework [28], whose time complexity is $O(m\alpha)$, where α is the arboricity of the graph. When $k = \omega$, the problem reduces to maximum clique search problem [7, 25, 29, 32]. The state-of-the-art algorithm for maximum clique search problem first transforms the maximum clique problem to a set of clique finding sub-problems, then it conducts a branch-and-bound framework to iteratively check whether a clique of a certain size can be found in the sub-problem [7], whose time complexity is $O(n2^n)$. We note that these existing algorithms cannot solve the k-clique listing problem for arbitrary k's.

8 CONCLUSION

In this paper, we study the k-clique listing problem, a fundamental graph mining operator with diverse applications in various networks. We propose a new branch-and-bound framework, named EBBkC, which incorporates an edge-oriented branching strategy. This strategy expands a partial k-clique using two connected vertices (an edge), offering new opportunities for optimization. Furthermore, to handle dense graph sub-branches more efficiently, we develop specialized algorithms that enable early termination, contributing to improved performance. We conduct extensive experiments on 19 real graphs, and the results consistently demonstrate EBBkC's superior performance compared to state-of-the-art VBBkC-based algorithms. In the future, we plan to explore the potential applications of our EBBkC technique to other cohesive subgraph mining tasks including clique and connected dense subgraph mining. In addition, we will explore the possibility of adopting our algorithms to list k-cliques' counterparts in other types of graphs such as bipartite graphs.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for providing constructive feedback and valuable suggestions. This research is supported by the Ministry of Education, Singapore, under its Academic Research Fund (Tier 2 Award MOE-T2EP20221-0013, Tier 2 Award MOE-T2EP20220-0011, and Tier 1 Award (RG77/21)). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore.

REFERENCES

[1] Real Graphs. http://lcs.ios.ac.cn/~caisw/Resource/realworld%20graphs.tar.gz.

[2] Balázs Adamcsek, Gergely Palla, Illés J Farkas, Imre Derényi, and Tamás Vicsek. 2006. CFinder: locating cliques and overlapping modules in biological networks. *Bioinformatics* 22, 8 (2006), 1021–1023.

[3] Albert Angel, Nick Koudas, Nikos Sarkas, Divesh Srivastava, Michael Svendsen, and Srikanta Tirthapura. 2014. Dense subgraph maintenance under streaming edge weight updates for real-time story identification. *The VLDB journal* 23, 2 (2014), 175–199.

[4] Vladimir Batagelj and Matjaz Zaversnik. 2003. An O (m) algorithm for cores decomposition of networks. *arXiv preprint cs/0310049* (2003).

[5] Austin R Benson, David F Gleich, and Jure Leskovec. 2016. Higher-order organization of complex networks. *Science* 353, 6295 (2016), 163–166.

[6] Coen Bron and Joep Kerbosch. 1973. Algorithm 457: finding all cliques of an undirected graph. *Commun. ACM* 16, 9 (1973), 575–577.

[7] Lijun Chang. 2019. Efficient maximum clique computation over large sparse graphs. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 529–538.

[8] Lijun Chang. 2020. Efficient maximum clique computation and enumeration over large sparse graphs. *The VLDB Journal* 29, 5 (2020), 999–1022.

[9] Yulin Che, Zhuohang Lai, Shixuan Sun, Yue Wang, and Qiong Luo. 2020. Accelerating truss decomposition on heterogeneous processors. *Proceedings of the VLDB Endowment* 13, 10 (2020), 1751–1764.

[10] Norishige Chiba and Takao Nishizeki. 1985. Arboricity and subgraph listing algorithms. *SIAM Journal on computing* 14, 1 (1985), 210–223.

[11] Jonathan Cohen. 2008. Trusses: Cohesive subgraphs for social network analysis. *National security agency technical report* 16, 3.1 (2008).

[12] Maximilien Danisch, Oana Balalau, and Mauro Sozio. 2018. Listing k-cliques in sparse real-world graphs. In *Proceedings of the 2018 World Wide Web Conference*. 589–598.

[13] David Eppstein, Maarten Löffler, and Darren Strash. 2010. Listing all maximal cliques in sparse graphs in near-optimal time. In *Algorithms and Computation: 21st International Symposium, ISAAC 2010, Jeju Island, Korea, December 15-17, 2010, Proceedings, Part I 21*. Springer, 403–414.

[14] Paul Erdös and George Szekeres. 1935. A combinatorial problem in geometry. *Compositio mathematica* 2 (1935), 463–470.

[15] Yixiang Fang, Kaiqiang Yu, Reynold Cheng, Laks VS Lakshmanan, and Xuemin Lin. 2019. Efficient algorithms for densest subgraph discovery. *Proceedings of the VLDB Endowment* 12, 11 (2019), 1719–1732.

[16] Irene Finocchi, Marco Finocchi, and Emanuele G Fusco. 2015. Clique counting in mapreduce: Algorithms and experiments. *Journal of Experimental Algorithmics (JEA)* 20 (2015), 1–20.

[17] Enrico Gregori, Luciano Lenzini, and Simone Mainardi. 2012. Parallel k-clique community detection on large-scale networks. *IEEE Transactions on Parallel and Distributed Systems* 24, 8 (2012), 1651–1660.

[18] William Hasenplaugh, Tim Kaler, Tao B Schardl, and Charles E Leiserson. 2014. Ordering heuristics for parallel graph coloring. In *Proceedings of the 26th ACM symposium on Parallelism in algorithms and architectures*. 166–177.

[19] Pan Hui and Jon Crowcroft. 2008. Human mobility models and opportunistic communications system design. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 366, 1872 (2008), 2005–2016.

[20] SK Jayanthi. 2012. Clique-attacks detection in web search engine for spamdexing using k-clique percolation technique. *International Journal of Machine Learning and Computing* 2, 5 (2012), 648.

[21] Yan Jin, Bowen Xiong, Kun He, Yangming Zhou, and Yi Zhou. 2022. On fast enumeration of maximal cliques in large graphs. *Expert Systems with Applications* 187 (2022), 115915.

[22] Richard M Karp. 2010. *Reducibility among combinatorial problems*. Springer.

[23] Matthieu Latapy. 2008. Main-memory triangle computations for very large (sparse (power-law)) graphs. *Theoretical computer science* 407, 1-3 (2008), 458–473.

[24] 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).

[25] Can Lu, Jeffrey Xu Yu, Hao Wei, and Yikai Zhang. 2017. Finding the maximum clique in massive graphs. *Proceedings of the VLDB Endowment* 10, 11 (2017), 1538–1549.

[26] Kazuhisa Makino and Takeaki Uno. 2004. New algorithms for enumerating all maximal cliques. In *Scandinavian workshop on algorithm theory*. Springer, 260–272.

[27] Kevin A Naudé. 2016. Refined pivot selection for maximal clique enumeration in graphs. *Theoretical Computer Science* 613 (2016), 28–37.

[28] Mark Ortmann and Ulrik Brandes. 2014. Triangle listing algorithms: Back from the diversion. In *2014 Proceedings of the Sixteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*. SIAM, 1–8.

[29] Patric RJ Östergård. 2002. A fast algorithm for the maximum clique problem. *Discrete Applied Mathematics* 120, 1-3 (2002), 197–207.
[30] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek. 2005. Uncovering the overlapping community structure of complex networks in nature and society. *nature* 435, 7043 (2005), 814–818.
[31] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek. 2005. Uncovering the overlapping community structure of complex networks in nature and society. *nature* 435, 7043 (2005), 814–818.
[32] Bharath Pattabiraman, Md Mostofa Ali Patwary, Assefaw H Gebremedhin, Wei-keng Liao, and Alok Choudhary. 2015. Fast algorithms for the maximum clique problem on massive graphs with applications to overlapping community detection. *Internet Mathematics* 11, 4-5 (2015), 421–448.
[33] Ryan Rossi and Nesreen Ahmed. 2015. The network data repository with interactive graph analytics and visualization. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
[34] Hiroo Saito, Masashi Toyoda, Masaru Kitsuregawa, and Kazuyuki Aihara. 2007. A large-scale study of link spam detection by graph algorithms. In *Proceedings of the 3rd international workshop on Adversarial information retrieval on the web*. 45–48.
[35] Ahmet Erdem Sariyuce, C Seshadhri, Ali Pinar, and Umit V Catalyurek. 2015. Finding the hierarchy of dense subgraphs using nucleus decompositions. In *Proceedings of the 24th International Conference on World Wide Web*. 927–937.
[36] Matthew C Schmidt, Nagiza F Samatova, Kevin Thomas, and Byung-Hoon Park. 2009. A scalable, parallel algorithm for maximal clique enumeration. *Journal of parallel and distributed computing* 69, 4 (2009), 417–428.
[37] UNO Takeaki. 2012. Implementation issues of clique enumeration algorithm. *Special issue: Theoretical computer science and discrete mathematics, Progress in Informatics* 9 (2012), 25–30.
[38] Etsuji Tomita, Akira Tanaka, and Haruhisa Takahashi. 2006. The worst-case time complexity for generating all maximal cliques and computational experiments. *Theoretical computer science* 363, 1 (2006), 28–42.
[39] Charalampos Tsourakakis. 2015. The k-clique densest subgraph problem. In *Proceedings of the 24th international conference on world wide web*. 1122–1132.
[40] Charalampos Tsourakakis, Francesco Bonchi, Aristides Gionis, Francesco Gullo, and Maria Tsiarli. 2013. Denser than the densest subgraph: extracting optimal quasi-cliques with quality guarantees. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 104–112.
[41] Jia Wang and James Cheng. 2012. Truss decomposition in massive networks. *Proceedings of the VLDB Endowment* 5, 9 (2012), 812–823.
[42] Kaixin Wang, Kaiqiang Yu, and Cheng Long. 2023. Efficient k-Clique Listing: An Edge-Oriented Branching Strategy (Technical Report). https://github.com/wangkaixin219/EBBkC/blob/main/EBBkC_TR.pdf.
[43] Jun Wu, Chu-Min Li, Lu Jiang, Junping Zhou, and Minghao Yin. 2020. Local search for diversified top-k clique search problem. *Computers & Operations Research* 116 (2020), 104867.
[44] Long Yuan, Lu Qin, Xuemin Lin, Lijun Chang, and Wenjie Zhang. 2016. Diversified top-k clique search. *The VLDB Journal* 25, 2 (2016), 171–196.
[45] Long Yuan, Lu Qin, Xuemin Lin, Lijun Chang, and Wenjie Zhang. 2017. Effective and efficient dynamic graph coloring. *Proceedings of the VLDB Endowment* 11, 3 (2017), 338–351.
[46] Zhirong Yuan, You Peng, Peng Cheng, Li Han, Xuemin Lin, Lei Chen, and Wenjie Zhang. 2022. Efficient k-clique listing with set intersection speedup. *ICDE. IEEE* (2022).
[47] Xiao Zhou and Takao Nishizeki. 1994. Edge-coloring and f-coloring for various classes of graphs. In *Algorithms and Computation: 5th International Symposium, ISAAC'94 Beijing, PR China, August 25–27, 1994 Proceedings 5*. Springer, 199–207.

Received July 2023; revised October 2023; accepted November 2023