



Efficient Maximal Frequent Group Enumeration in Temporal Bipartite Graphs

Yanping Wu
University of Technology Sydney
yanping.wu@student.uts.edu.au

Renjie Sun
East China Normal University
renjie.sun@stu.ecnu.edu.cn

Xiaoyang Wang
The University of New South Wales
xiaoyang.wang1@unsw.edu.au

Dong Wen
The University of New South Wales
dong.wen@unsw.edu.au

Ying Zhang
University of Technology Sydney
ying.zhang@uts.edu.au

Lu Qin
University of Technology Sydney
lu.qin@uts.edu.au

Xuemin Lin
Shanghai Jiaotong University
xuemin.lin@sjtu.edu.cn

ABSTRACT

Cohesive subgraph mining is a fundamental problem in bipartite graph analysis. In reality, relationships between two types of entities often occur at some specific timestamps, which can be modeled as a temporal bipartite graph. However, the temporal information is widely neglected by previous studies. Moreover, directly extending the existing models may fail to find some critical groups in temporal bipartite graphs, which appear in a unilateral (i.e., one-layer) form. To fill the gap, in this paper, we propose a novel model, called maximal λ -frequency group (MFG). Given a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$, a vertex set $V_S \subseteq V$ is an MFG if *i*) there are no less than λ timestamps, at each of which V_S can form a (τ_U, τ_V) -biclique with some vertices in U at the corresponding snapshot, and *ii*) it is maximal. To solve the problem, a filter-and-verification (FilterV) method is proposed based on the Bron-Kerbosch framework, incorporating novel filtering techniques to reduce the search space and array-based strategy to accelerate the frequency and maximality verification. Nevertheless, the cost of frequency verification in each valid candidate set computation and maximality check could limit the scalability of FilterV to larger graphs. Therefore, we further develop a novel verification-free (VFree) approach by leveraging the advanced dynamic counting structure proposed. Theoretically, we prove that VFree can reduce the cost of each valid candidate set computation in FilterV by a factor of $\mathcal{O}(|V|)$. Furthermore, VFree can avoid the explicit maximality verification because of the developed search paradigm. Finally, comprehensive experiments on 15 real-world graphs are conducted to demonstrate the efficiency and effectiveness of the proposed techniques and model.

PVLDB Reference Format:

Yanping Wu, Renjie Sun, Xiaoyang Wang, Dong Wen, Ying Zhang, Lu Qin, and Xuemin Lin. Efficient Maximal Frequent Group Enumeration in Temporal Bipartite Graphs. PVLDB, 17(11): 3243 - 3255, 2024. doi:10.14778/3681954.3681997

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. Proceedings of the VLDB Endowment, Vol. 17, No. 11 ISSN 2150-8097. doi:10.14778/3681954.3681997

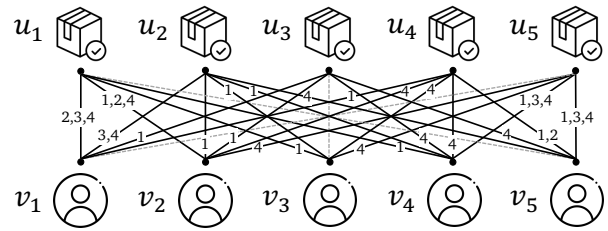


Figure 1: Customer-product temporal bipartite graph (dotted lines denote the edges at $t = 5$ for presentation simplicity)

PVLDB Artifact Availability:

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

1 INTRODUCTION

Bipartite graphs are widely used to model the complex relationships between two types of entities, e.g., author-paper networks [9, 40], customer-product networks [8, 20, 29, 45, 51], and patient-disease networks [6, 21, 50]. As a fundamental problem in graph analysis, cohesive subgraph mining is broadly investigated. To analyze the properties of bipartite graphs, many cohesive subgraph models have been proposed, such as (α, β) -core [24], bitruss [40] and biclique [29, 36, 47]. Among these models, biclique, which requires every pair of vertices from different vertex sets to be mutually connected, has gained widespread popularity due to its unique features and diverse applications. However, existing models on bipartite graphs primarily focus on static graphs, disregarding the temporal aspect of relationships in real-world applications. For instance, Figure 1 shows a customer-product network. The timestamps on an edge between two vertices indicate when a customer purchased the corresponding product. In a patient-disease network, patients and diseases can be represented by two disjoint vertex sets, and the timestamps on an edge represent when the patient suffered from the disease. The above cases can be modeled as a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$, where each edge $e \in \mathcal{E}$ can be represented as a tuple (u, v, t) , indicating that the interaction between vertex $u \in U$ and vertex $v \in V$ occurs at timestamp t (e.g., [6, 12]). The bipartite graph with all the edges at timestamp t is called a snapshot of \mathcal{G} , denoted by G_t .

Analyzing the properties of temporal bipartite graphs is essential to reveal more sophisticated semantics. In the literature, many subgraph models are defined for unipartite scenarios. For instance, Li et al. [23] propose a k -core based model to capture the persistence of a community within the time interval. Qin et al. [32] design a clique based model to find the subgraphs periodically occurring in the temporal graphs. Qin et al. [33] propose the (l, δ) -maximal bursting core, where each vertex has an average degree no less than δ during a period of length no less than l . Although temporal subgraph search and enumeration problems have been extensively studied on temporal unipartite graphs, the temporal bipartite graph case is still under-explored. Due to the involvement of two distinct types of entities, existing models on temporal unipartite graphs are not suitable for capturing important patterns in temporal bipartite graphs. Only a few recent works consider temporal bipartite graphs, such as (α, β) -core based persistent community search [22] and temporal butterfly counting and enumeration [9]. Unfortunately, the aforementioned temporal models are mainly established on interval-based or periodic-based constraints, which fail to capture the real scenarios occurring irregularly at non-consecutive timestamps. Additionally, there is no existing work considering the *unilateral* (i.e., one-layer) frequency of the occurred group in temporal bipartite graphs, which is an important factor in identifying practical groups. For example, in a temporal customer-product bipartite graph, a group of users who frequently act together (e.g., purchase the same items at different timestamps) may show strong connections. Existing works about temporal subgraph mining are inadequate for the above scenario, highlighting the need to define the exclusive model tailored for temporal bipartite graphs.

To fill the gap, in this paper, we propose a novel model, called *Maximal λ -Frequency Group* (MFG), to characterize the unilateral patterns in temporal bipartite graphs. In our model, we leverage *biclique* to measure the cohesiveness of bipartite subgraphs, due to its diverse applications such as social recommendation [25] and anomaly detection [29]. Specifically, given two size constraints τ_U , τ_V and a frequency constraint λ , a vertex set $V_S \subseteq V$ in a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$ is an MFG if *i*) there are no less than λ timestamps, at each of which V_S can form a (τ_U, τ_V) -biclique with some vertices in U in the corresponding snapshot graph, and *ii*) it is maximal. A (τ_U, τ_V) -biclique $S = (U_S, V_S, E_S)$ is a biclique with $|U_S| \geq \tau_U$ and $|V_S| \geq \tau_V$. Note that, an MFG V_S is unilateral, i.e., only consists of vertices from V . In addition, the vertices in U that form (τ_U, τ_V) -bicliques with V_S can be different for different timestamps.

EXAMPLE 1.1. *Reconsider the temporal customer-product bipartite graph in Figure 1 with $\tau_U = 2$, $\tau_V = 2$ and $\lambda = 2$. Suppose a company plans to promote a new product to a group of customers with similar interests. Directly extending the static graph model (i.e., biclique) to temporal bipartite graphs may fail to retrieve useful patterns. For instance, if we treat it as a static graph, i.e., ignore the timestamp information, the whole graph itself is a biclique. All five customers will be grouped together and treated equally since they all purchased all the products. If we consider it as a temporal bipartite graph, we still cannot find practical results by applying the frequent (τ_U, τ_V) -biclique model, which is the (τ_U, τ_V) -biclique occurring in at least λ snapshots. For the MFG model, $\{v_2, v_3, v_4\}$ is the returned result, since these customers*

frequently act together, i.e., it exists in biclique $\{u_1, u_2, v_2, v_3, v_4\}$ at timestamp $t = 1$ and in biclique $\{u_4, u_5, v_2, v_3, v_4\}$ at timestamp $t = 4$. As discussed before, customers within a group, who frequently act together, are more likely to share similar preferences or behavioral patterns. Identifying these customer-specified communities is crucial to improve the performance of downstream tasks such as product recommendation and enhance customer engagement.

Besides biclique, many other cohesive subgraph models are proposed in the literature for bipartite graph analysis, such as degree-based model (α, β) -core and butterfly-based model bitruss. All these models have various applications in different domains [24, 40]. In this paper, we focus more on the strong connections among entities (e.g., co-purchase) in the target layer. Thus, we choose biclique, which is the most cohesive model. Recall the example in Figure 1, if we directly replace the biclique constraint with $(2,2)$ -core in our model, all users $\{v_1, v_2, v_3, v_4, v_5\}$ will be returned as a group. In this paper, we employ the frequency constraint (i.e., λ), since frequently appearing patterns are usually worthy of attention and may represent the important concept in the environment [5, 43, 44, 46, 48], such as frequent co-purchasing behavior discussed above. Besides, compared with the existing temporal models that mainly focus on interval-based or periodic-based timestamp constraints, our model can better capture the real-world patterns occurring frequently. Moreover, we employ the maximality constraint, since any subset of an MFG with size no less than τ_V is also a λ -frequency group. Without the maximality constraint, it may generate many redundant results. In this paper, we aim to enumerate all MFGs from a temporal bipartite graph.

In addition to the application of customer analysis mentioned above, MFG can find many other applications in different domains. For instance, a temporal bipartite graph is a suitable data structure to model the data of patients' diagnostic records, where the vertices correspond to patients (i.e., U) and health conditions (i.e., V), and the links indicate the presence of a diagnosis at the corresponding timestamp [6, 21, 50]. By mining MFGs from the patient-condition temporal bipartite graph, we can find the combinations of health conditions that frequently and simultaneously appear in multiple patients. The results can provide data support for the study of multimorbidity, facilitating diagnosis and prevention [34, 39]. In Section 5, we present two case studies on real-world datasets to illustrate the effectiveness of our model.

Challenges and our approaches. To the best of our knowledge, we are the first to propose and investigate the maximal λ -frequency group (MFG) enumeration problem in temporal bipartite graphs. We prove the hardness of counting MFGs. In the literature, maximal biclique enumeration is the most relevant problem to ours (e.g., [3, 11, 47, 49]). However, the introduction of temporal and unilateral aspects significantly complicates the problem. Naively, we can enumerate all bicliques over each snapshot and post-process the intermediate results. Due to the hardness of the biclique enumeration problem in static bipartite graphs, treating each timestamp separately is time-consuming. Moreover, since the correlation among vertices varies over time, considering temporal and cohesive aspects simultaneously in algorithm design is nontrivial. In previous studies, the Bron-Kerbosch (BK) framework is widely used for biclique enumeration (e.g., [3, 11]). It iteratively adds vertices

from the candidate set to expand the current result in a DFS manner for biclique enumeration. By extending the BK framework, in our problem, we need to further check the frequency and maximality constraints for each candidate group. Since MFG focuses on the unilateral vertex set and the number of bicliques in each timestamp is large, it means numerous candidate groups will be generated while only very few of them will belong to the final results. Thus, the huge time cost to apply the naive frequency verification on each candidate group and eliminate the non-maximal results presents a unique challenge for our problem.

To address the challenges, in this paper, we first propose a filter-and-verification (FilterV) approach by leveraging the BK framework. Generally, FilterV maintains a recursive search tree and traverses in a depth-first manner. In each iteration, FilterV iterates over all the candidate vertices (i.e., vertex-oriented search paradigm), verifies the frequency constraint and obtains the *valid candidate set*. Note that, for each vertex in the valid candidate set, the group composed of it and the current processing set still meets the frequency requirement. If no more vertices can be added to the current processing set to form a new frequent group, FilterV terminates the current search branch and checks the maximality of the vertex set. Since there could be many vertices that cannot be involved in any MFGs, a novel structure $(\tau_V, \tau_U, \lambda)$ -core is first designed to reduce the search space. Then, a filter strategy is proposed to first efficiently prune the candidate set before the examination, which can reduce the unnecessary call of frequency verification. To further accelerate the frequency verification of a given vertex set, we present an elaborate array-based verification strategy. The frequency check method is also employed to accelerate the maximality verification by avoiding the numerous set comparisons.

Even though FilterV can remarkably accelerate the MFG enumeration procedure, we need to compute the valid candidate set for each current processing result during the search. The overall cost significantly increases with the size of candidate set and the workload of maximality verification, which may hinder its scalability to larger graphs. As shown in Table 1, whose details can be found in Section 4, the components of the valid candidate set computation and maximality verification take up a majority of the overall execution time. Therefore, if we can reduce or even avoid the cost of frequency and maximality verification to some extent, the overall performance can be significantly improved. Motivated by this, we further develop a novel verification-free (VFree) approach. Instead of iterating over vertices during the candidate set computation and maximality verification, we develop a timestamp-oriented search paradigm. That is, VFree iterates through the timestamps to obtain the valid candidate set using the advanced dynamic counting structures proposed, where the unpromising timestamp can be skipped and common neighbor information can be carried forward in the subsequent search process. Theoretically, we prove that VFree can significantly reduce each valid candidate set computation cost in FilterV by a factor of $\mathcal{O}(|V|)$. Additionally, by integrating the developed search paradigm and dynamic counting techniques, VFree can avoid explicit maximality verification.

Contributions. The main contributions of the paper are summarized as follows.

- To capture the properties of temporal bipartite graphs, we conduct the first research to propose and investigate the maximal λ -frequency group enumeration problem. (Section 2)
- To solve the problem, we introduce a filter-and-verification framework. Novel $(\tau_V, \tau_U, \lambda)$ -core structure and candidate filtering rule are developed to shrink the search space. Advanced array-based method is proposed to accelerate the computation of valid candidate set and maximality verification. (Section 3)
- To overcome the frequency verification cost and scale for larger networks, we further develop a verification-free framework by leveraging the dynamic counting structure proposed. The framework can also avoid the explicit verification of maximality based on the propounded search paradigm. (Section 4)
- Extensive experiments are conducted on 15 real-world graphs to demonstrate the performance of proposed techniques and model. Compared with the baseline, the optimized method can achieve up to three orders of magnitude speedup. (Section 5)

Note that, due to the limited space, all the proofs, detailed experiment setup and partial experiment results can be found in the full version [42].

2 PRELIMINARY AND PROBLEM DEFINITION

2.1 Preliminary

Let $\mathcal{G} = (U, V, \mathcal{E})$ denote an undirected temporal bipartite graph, where U and V are two disjoint vertex sets, i.e., $U \cap V = \emptyset$, and \mathcal{E} is the set of temporal edges. (u, v, t) denotes a temporal edge between $u \in U$ and $v \in V$, where t is the interaction timestamp between u and v . Without loss of generality, we use $\mathcal{T} = \{t_1, t_2, \dots, t_{|\mathcal{T}|}\}$ to represent the set of timestamps, i.e., $\mathcal{T} = \{t | (u, v, t) \in \mathcal{E}\}$ ¹. Given a temporal bipartite graph \mathcal{G} , its corresponding *static bipartite graph*, i.e., by ignoring all the timestamps on edges, is denoted by $G = (U, V, E)$, where $E = \{(u, v) | (u, v, t) \in \mathcal{E}\}$. We can extract a series of snapshots $\{G_1, G_2, \dots, G_{|\mathcal{T}|}\}$ from \mathcal{G} based on the timestamps. Specifically, given a timestamp $t \in \mathcal{T}$, its corresponding snapshot is a bipartite graph $G_t = (U_t, V_t, E_t)$, where $U_t = \{u | (u, v, t) \in \mathcal{E}\}$, $V_t = \{v | (u, v, t) \in \mathcal{E}\}$, and $E_t = \{(u, v) | (u, v, t) \in \mathcal{E}\}$.

DEFINITION 2.1 (STRUCTURAL NEIGHBOR (S-NEIGHBOR)). *Given a vertex $u \in \mathcal{G}$, the s-neighbor set of u is the set of vertices connected to u in G , denoted by $N(u, G)$, i.e., $N(u, G) = \{v | (u, v) \in E\}$. $d(u, G)$ denotes its structural degree (s-degree), i.e., $d(u, G) = |N(u, G)|$.*

DEFINITION 2.2 (MOMENTARY NEIGHBOR (M-NEIGHBOR)). *Given a vertex $u \in \mathcal{G}$ and a timestamp $t \in \mathcal{T}$, the m-neighbor set of u at t is the set of vertices connected to u in G_t , denoted by $\Gamma(u, t)$, i.e., $\Gamma(u, t) = \{v | (u, v) \in E_t\}$. We use $\delta(u, t)$ to denote its momentary degree (m-degree) at t , i.e., $\delta(u, t) = |\Gamma(u, t)|$.*

EXAMPLE 2.1. *Figure 2 shows a temporal bipartite graph with six snapshots G_1 to G_6 . The s-neighbors of u_1 are $\{v_1, v_2, v_3, v_4, v_5\}$ and the s-degree $d(u_1, G)$ is 5. The m-neighbor of u_1 at $t = 1$ is $\{v_3\}$ and the corresponding m-degree $\delta(u_1, 1)$ is 1.*

Given a static bipartite graph $G = (U, V, E)$, a biclique $S = (U_S, V_S, E_S)$ is a complete subgraph of G , where $U_S \subseteq U$, $V_S \subseteq V$, and for each pair of vertices $u \in U_S$ and $v \in V_S$, we have $(u, v) \in E_S$.

¹We use the same setting as the previous studies for the timestamp, which is the integer, since the UNIX timestamps are integers in practice [33, 48].

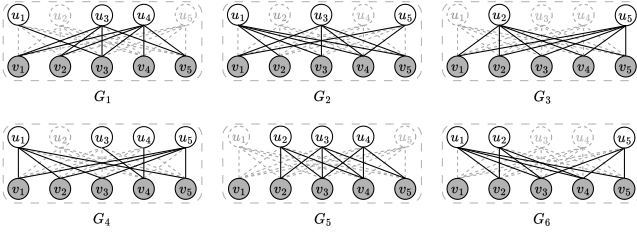


Figure 2: A temporal bipartite graph \mathcal{G} with six timestamps (G_1 - G_6 are the corresponding snapshots and solid lines denote the edges in each snapshot)

DEFINITION 2.3 ((τ_U, τ_V)-BICLIQUE). Given a static bipartite graph G and two positive integers τ_U and τ_V , a (τ_U, τ_V) -biclique $S = (U_S, V_S, E_S)$ is a biclique of G with $|U_S| \geq \tau_U$ and $|V_S| \geq \tau_V$.

2.2 Problem Definition

In this paper, we aim to retrieve the frequent vertex set in unilateral layer, e.g., the customer set in Figure 1. **For presentation simplicity, we assume the frequent vertex set is from V .** Before introducing the λ -frequency group, we first define the support timestamp for the unilateral vertex set below.

DEFINITION 2.4 (SUPPORT TIMESTAMP). Given a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$, a subset $V_S \subseteq V$, a snapshot $G_t = (U_t, V_t, E_t)$ of the timestamp $t \in \mathcal{T}$, and two positive integers τ_U, τ_V , we say t is a support timestamp of V_S , if i) $V_S \subseteq V_t$, and ii) V_S can form a (τ_U, τ_V) -biclique with a subset of vertices in U_t , i.e., V_S is included in a (τ_U, τ_V) -biclique of G_t .

DEFINITION 2.5 (λ -FREQUENCY GROUP). Given a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$, and three positive integers τ_U, τ_V and λ , a λ -frequency group is a subset of vertices $V_S \subseteq V$ where there are at least λ support timestamps in \mathcal{T} for V_S .

DEFINITION 2.6 (MAXIMAL λ -FREQUENCY GROUP (MFG)). Given a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$ and three positive integers τ_U, τ_V and λ , a λ -frequency group V_S is maximal if there is no other λ -frequency group V'_S that is a superset of V_S .

Problem definition. Given a temporal bipartite graph \mathcal{G} and three positive integers τ_U, τ_V and λ , in this paper, we aim to find all the maximal λ -frequency groups (MFGs) in \mathcal{G} .

EXAMPLE 2.2. Considering the temporal bipartite graph shown in Figure 2, suppose $\tau_U = \tau_V = 2$ and $\lambda = 3$. There are three MFGs in the graph, i.e., $V_{S1} = \{v_1, v_2, v_3, v_5\}$ with 3 support timestamps $\{t_1, t_3, t_4\}$, $V_{S2} = \{v_2, v_3, v_4\}$ with 3 support timestamps $\{t_3, t_5, t_6\}$ and $V_{S3} = \{v_3, v_4, v_5\}$ with 4 support timestamps $\{t_2, t_3, t_5, t_6\}$.

Problem properties. Based on the definition of MFG, Lemmas 2.1 and 2.2 can be immediately obtained. The proofs are omitted. Then, we show the hardness of our problem in Theorem 2.1.

LEMMA 2.1 (STRUCTURAL PROPERTY). Given a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$ and three positive integers τ_U, τ_V and λ , any MFG $V_S \subseteq V$ must be contained in a maximal (τ_U, τ_V) -biclique of the static bipartite graph G of \mathcal{G} .

LEMMA 2.2 (ANTIMONOTONE PROPERTY). Given a vertex set $V_S \subseteq V$, i) if V_S satisfies the frequency constraint, any subset of V_S also

satisfies the constraint; ii) if V_S does not meet the frequency constraint, any superset of V_S is not frequent.

THEOREM 2.1. The problem of counting all MFGs is #P-complete.

Due to the space limitation, the detailed proof for Theorem 2.1 and other omitted proofs can be found in the full version [42].

3 FILTER-AND-VERIFICATION APPROACH

In the literature, the closest problem to ours is the maximal biclique enumeration problem (e.g., [3, 11, 49]), where most studies are based on the Bron-Kerbosch (BK) framework. It maintains a recursion search tree and traverses in a depth-first manner.

Baseline method. Motivated by Lemmas 2.1 and 2.2, a reasonable approach for our problem is to employ the BK framework by jointly considering the frequency and maximality constraints. Specifically, we operate on three dynamically changing vertex sets (U_S, V_S, C_V) . V_S is the current result. U_S is the common s-neighbors of all the vertices in V_S . C_V is the candidate set. In each iteration, we select a vertex from the candidate set C_V to expand V_S , and update the corresponding U_S . If V_S satisfies the frequency constraint, we continue to expand it. If no other vertex can be added into V_S to form a new frequent group, we terminate the current search branch and check the maximality of V_S by comparing it with the existing found results. After enumerating through each search branch, all the MFGs are returned. This algorithm is referred to as BK-ALG.

Limitations. Although BK-ALG can correctly return all the MFGs for a given temporal bipartite graph, we find that directly extending the BK framework is inefficient due to the following two drawbacks. The first drawback is the huge search space. The search space of BK-ALG is the whole graph \mathcal{G} , and it needs to iterate through all the vertices in C_V in each branch, which may involve many unpromising vertices that cannot exist in any MFG. The second drawback is the cumbersome frequency constraint check. Similar to biclique enumeration in static graphs, in BK-ALG, C_V maintains the candidate vertices and we need to ensure the frequency constraint during the search, which is computationally expensive.

To address these limitations, in this section, we propose a novel filter-and-verification (FilterV) algorithm. In the following, we first introduce the search framework of FilterV (Section 3.1). For *drawback 1*, we develop novel graph and candidate set filtering techniques to dramatically shrink search space (Section 3.2). For *drawback 2*, an advanced array-based algorithm is presented to facilitate the frequency (Section 3.3) and maximality (Section 3.4) verification.

3.1 Framework Overview

Hereafter we present an overview of our filter-and-verification (FilterV) framework. We call each search branch that fails to find an MFG an invalid branch. Recall the search branch (U_S, V_S, C_V) . To reduce even avoid the search cost on invalid search branches, instead of directly adding each candidate vertex from C_V into V_S , we first perform a verification procedure on C_V to generate the **valid candidate set** $C_V^* \subseteq C_V$ for V_S . That is, in subsequent branches, the new set obtained by adding any candidate vertex from C_V^* to the current processing set V_S still can meet the frequency requirement.

FilterV framework. Motivated by the above idea, the pseudocode of FilterV framework is presented in Algorithm 1. It first applies the

Algorithm 1: Filter-and-Verification (FilterV)

Input : $\mathcal{G} = (U, V, \mathcal{E})$: a temporal bipartite graph,
 τ_U, τ_V : size constraints, λ : frequency constraint

Output : \mathcal{R} : all the MFGs

```
1  $\mathcal{G} \leftarrow \text{GFCore}(\mathcal{G}, \tau_U, \tau_V, \lambda)$ ; /* graph filter: Algorithm 2 */
2  $\mathcal{R} \leftarrow \emptyset$ ;
3 EnumMFG( $U, \emptyset, V$ );
4 Procedure EnumMFG( $U_S, V_S, C_V$ )
5  $C_V \leftarrow$  filter candidate set  $C_V$ ; /* candidate filter: Lemma 3.2 */
6  $C_V^* \leftarrow \emptyset$ ; /* compute valid candidate set */
7 for each  $v \in C_V$  do
8   if CheckFRE( $U_S \cap N(v, G), V_S \cup \{v\}, \lambda$ ) then /* Algorithm 3 */
9      $C_V^* \leftarrow C_V^* \cup \{v\}$ ;
10 if  $|U_S| < \tau_U \vee |V_S| + |C_V^*| < \tau_V$  then
11   return;
12 if  $C_V^* = \emptyset$  then
13   Check the maximality for  $V_S$ ; /* Section 3.4 */
14   if  $V_S$  is maximal then
15      $\mathcal{R} \leftarrow \mathcal{R} \cup \{V_S\}$ ;
16 for each  $v \in C_V^*$  do
17    $C_V \leftarrow C_V \setminus \{v\}$ ;
18   EnumMFG( $U_S \cap N(v, G), V_S \cup \{v\}, C_V^*$ );
```

graph filtering technique (Algorithm 2 in Section 3.2) to shrink the given temporal bipartite graph. In the following, FilterV invokes the procedure EnumMFG to enumerate all the MFGs. Similar as the BK method, EnumMFG maintains three sets U_S , V_S and C_V , which are initialized as U , \emptyset and V . In EnumMFG, we first try to filter the candidate set C_V (line 5) and then compute the valid candidate set C_V^* for V_S (lines 7-9). The branch is terminated if it violates the (τ_U, τ_V) -biclique size constraints. We check the maximality of V_S when the valid candidate set is empty (lines 12-15). If C_V^* is not empty, we process each vertex $v \in C_V^*$ to expand V_S , and continue search on the updated V_S and U_S (lines 16-18).

Discussion. To compute the valid candidate set C_V^* for V_S in lines 7-9, we need to check the frequency of each vertex set $V_S \cup \{v\}$ for $v \in C_V$. Given the vertex set V_S and the checking vertex $v \in C_V$, a naive method to check the frequency of $V_S \cup \{v\}$ is to compute the common m -neighbors of $V_S \cup \{v\}$ at each timestamp. Specifically, for each timestamp $t \in \mathcal{T}$, we check whether there exists no less than τ_U common m -neighbors of all the vertices in $V_S \cup \{v\}$. If it satisfies the constraint, t can contribute to the frequency for $V_S \cup \{v\}$. If the number of such timestamps for $V_S \cup \{v\}$ is no less than λ , v is the valid candidate vertex for V_S and it can be added into C_V^* . After checking all vertices in C_V , we can return C_V^* . Due to the large scale of candidate vertices and the inefficiency of the naive frequency checking method, the above process is very time-consuming. To speed up the computation of the valid candidate set, we designed novel filtering strategies and efficient verification techniques.

3.2 Filtering Rules

Graph filter. Given a temporal bipartite graph \mathcal{G} , many unpromising vertices cannot exist in any MFGs. Therefore, we propose a novel graph structure to filter the search space. Before presenting the details, we first introduce the concept of (α, β) -core [24].

DEFINITION 3.1 ((α, β)-CORE). *Given a static bipartite graph G and two positive integers α and β , the subgraph $S = (U_S, V_S, E_S)$ is the (α, β) -core of G , denoted by $\text{Core}(G)$, if it satisfies: i) the degree*

of each vertex in U_S is at least α and the degree of each vertex in V_S is at least β in S , and ii) any supergraph of S cannot satisfy i).

To compute the (α, β) -core, we can iteratively remove the vertices that violate the degree constraint with time complexity $\mathcal{O}(|E|)$. Based on the definition and properties of MFG, we can obtain that every vertex of MFG must be contained in the temporal bipartite graph $\{\text{Core}(G_1), \text{Core}(G_2), \dots, \text{Core}(G_{|\mathcal{T}|})\}$. To further model the frequency property of the vertex, we present a new frequent cohesive subgraph model, called $(\tau_V, \tau_U, \lambda)$ -core, which will be applied to prune unpromising vertices before enumerating all the MFGs.

DEFINITION 3.2 ((τ_V, τ_U, λ)-CORE). *Given a temporal bipartite graph \mathcal{G} and three positive integers τ_U , τ_V and λ , the induced subgraph (U_S, V_S, E_S) of \mathcal{G} is the $(\tau_V, \tau_U, \lambda)$ -core if it meets the following conditions: i) each vertex $u \in U_S$ can be included in the (τ_V, τ_U) -core of at least one snapshot, ii) each vertex $v \in V_S$ can be included in the (τ_V, τ_U) -core of at least λ snapshots, and iii) there is no supergraph of (U_S, V_S, E_S) that satisfies i) and ii).*

Note that, in (τ_U, τ_V) -biclique, τ_U and τ_V restrict the number of vertices in U and V . But, in Definition 3.2, the parameters restrict the number of neighbors, so the order is changed. Based on the analysis, we can directly obtain the connection between the $(\tau_V, \tau_U, \lambda)$ -core and the proposed MFG model, whose details can be found in the following lemma.

LEMMA 3.1. *Given a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$ and three positive integers τ_U , τ_V and λ , if a subset $V_S \subseteq V$ is an MFG, then V_S must be contained in the $(\tau_V, \tau_U, \lambda)$ -core of \mathcal{G} .*

To derive the $(\tau_V, \tau_U, \lambda)$ -core, we can iteratively delete the vertex that violates the degree constraint or the frequency constraint. Since the deletion of one vertex may cause its neighbors to violate the constraints in cascade, we can iteratively prune the graph until all the remaining vertices in \mathcal{G} meet the constraints. Details of computing $(\tau_V, \tau_U, \lambda)$ -core are shown in Algorithm 2. At first, we use $s[w]$ to count the number of timestamps when the vertex w has enough m -neighbors (lines 1-5). Then we process all the vertices at each timestamp in lines 6-11. Specifically, for each vertex w , if $s[w]$ violates the frequency constraint or the m -degree constraint at t , we remove this vertex at t and invoke the procedure CorePrune to update the graph. Details of CorePrune are shown in lines 16-29. For the processing vertex w at timestamp t in line 16, we set $\delta(w, t)$ as 0 and traverse all m -neighbors of w at t , i.e., $\Gamma(w, t)$. For each vertex $x \in \Gamma(w, t)$, we first check its m -degree. If it is larger than 0, we reduce its m -degree by 1 (lines 19-20). Then, if x violates the m -neighbor constraint at t , we invoke CorePrune for (x, t) in lines 21-22. For the vertex w with $s[w] > 0$, we need to update $s[w]$ for it in lines 23-29. We first reduce $s[w]$ by 1. If $s[w]$ violates the constraint, we set $s[w]$ as 0 and invoke CorePrune for w at all timestamps (lines 25-29). After updating the graph, we remove all unsatisfied vertices at each snapshot (lines 12-14). Finally, we return all the updated snapshots as the reduced graph in line 15. Similar as (α, β) -core, the time complexity is bounded by $\mathcal{O}(|E|)$.

Candidate set filter. Recall the search branch with processed vertex sets (U_S, V_S, C_V) , where we iteratively add one vertex from C_V into V_S and check its frequency. If we can efficiently skip a batch of vertices in C_V without compromising any results, we can reduce

Algorithm 2: GFCore($\mathcal{G}, \tau_U, \tau_V, \lambda$)

Input : $\mathcal{G} = (U, V, \mathcal{E})$: a temporal bipartite graph,
 τ_U, τ_V : size constraints, λ : frequency constraint

Output : $(\tau_V, \tau_U, \lambda)$ -core of \mathcal{G}

```
1 for each  $w \in \mathcal{G}$  do  $s[w] = 0$ ;
2 for each  $t \in \mathcal{T}$  do
3   for each  $w \in \mathcal{G}$  do
4     if  $w \in U \wedge \delta(w, t) > 0$  then  $s[w]++$ ;
5     if  $w \in V \wedge \delta(w, t) > 0$  then  $s[w]++$ ;
6 for each  $t \in \mathcal{T}$  do
7   for each  $w \in \mathcal{G}$  do
8     if  $w \in U \wedge 0 < \delta(w, t) < \tau_V$  then
9       CorePrune( $w, t$ );
10    if  $w \in V \wedge (0 < \delta(w, t) < \tau_U \vee 0 < s[w] < \lambda)$  then
11      CorePrune( $w, t$ );
12 for each  $t \in \mathcal{T}$  do
13   for each  $w \in \mathcal{G}$  do
14     if  $\delta(w, t) = 0$  then  $G_t \leftarrow G_t \setminus \{w\}$ ;
15 return  $\{G_1, G_2, \dots, G_{|\mathcal{T}|}\}$ ;
16 Procedure CorePrune( $w, t$ )
17  $\delta(w, t) = 0$ ;
18 for each  $x \in \Gamma(w, t)$  do
19   if  $\delta(x, t) > 0$  then
20      $\delta(x, t) = \delta(x, t) - 1$ ;
21     if  $(x \in U \wedge \delta(x, t) < \tau_V) \vee (x \in V \wedge \delta(x, t) < \tau_U)$  then
22       CorePrune( $x, t$ );
23 if  $s[w] > 0$  then
24    $s[w] = s[w] - 1$ ;
25   if  $(w \in U \wedge s[w] < 1) \vee (w \in V \wedge s[w] < \lambda)$  then
26      $s[w] = 0$ ;
27     for each  $t \in \mathcal{T}$  do
28       if  $\delta(w, t) > 0$  then
29         CorePrune( $w, t$ );
```

many unnecessary calls of frequency check. To achieve this, we propose the following rule to quickly filter the candidate vertex set.

LEMMA 3.2 (CANDIDATE SET FILTERING RULE). *Given a temporal bipartite graph $\mathcal{G} = (U, V, \mathcal{E})$ and $v \in V$, we use $T(v)$ to denote the set of timestamps when v has more than τ_U m -neighbors, i.e., $T(v) = \{t \in \mathcal{T} \mid \delta(v, t) \geq \tau_U\}$. Then, for the current processing vertex sets V_S , we can skip a candidate vertex $v' \in C_V$, if $|\cap_{v \in V_S \cup \{v'\}} T(v)| < \lambda$.*

3.3 Frequency Verification

Recall the computation of the valid candidate set C_V^* for V_S , whose main cost is the frequency verification for all the vertex sets $V_S \cup \{v\}$, where $v \in C_V$. In addition, as discussed before, the naive frequency verification method is very time-consuming. Therefore, reducing the cost of frequency verification is crucial for optimizing the performance of algorithm. Motivated by this, in this section, we design a novel array-based structure to speed up the processing. The detailed method is presented in Algorithm 3.

CheckFRE algorithm. Algorithm 3 has three input parameters, i.e., U_S, V_S and λ , which corresponds to line 8 in Algorithm 1. It returns *true* if V_S satisfies the frequency constraint. Otherwise, it returns *false*. The algorithm employs two array structures, i.e., Reborn Array (RA) and Update Array (UA), both of which have a length of $|\mathcal{T}|$. Initialization (lines 1-2). We initialize λ' as 0 to

Algorithm 3: CheckFRE(U_S, V_S, λ)

Input : U_S : the common s -neighbors of all the vertices in V_S ,
 V_S : the checking vertex set, λ : frequency constraint

Output : frequency verification result *true/false*

```
1  $\lambda' = 0$ ;
2 for each  $t \in \mathcal{T}$  do  $UA[t] = 0$ ; /* Update Array */
3 for each  $u \in U_S$  do
4   for each  $t \in \mathcal{T}$  do  $RA[t] = 0$ ; /* Reborn Array */
5   for each  $v \in N(u, G) \wedge v \in V_S$  do
6     for each  $t \in \mathcal{T}_{(u,v)}$  do
7        $RA[t]++$ ; /* count  $u$ 's  $m$ -neighbors in  $V_S$  at  $t$  */
8   for each  $t \in \mathcal{T}$  do
9     if  $RA[t] = |V_S|$  then
10       $UA[t]++$ ; /* count common  $m$ -neighbors of  $V_S$  at  $t$  */
11 for each  $t \in \mathcal{T}$  do
12   if  $UA[t] \geq \tau_U$  then
13      $\lambda'++$ ; /* count support timestamp for  $V_S$  */
14   if  $\lambda' = \lambda$  then return true;
15 return false;
```

count the frequency for the checking vertex set and all the elements in UA as 0. Then, we process each vertex $u \in U_S$ iteratively (lines 3-10). Count u 's m -neighbors in V_S at t (lines 5-7). For each processed vertex u , we use RA to count its m -neighbors in V_S at each timestamp, whose elements are initialized as 0 (line 4). $\mathcal{T}_{(u,v)}$ is the set of timestamps associated with edge (u, v) . After processing all m -neighbors of u , the RA for u is constructed. Count common m -neighbors of V_S at t (lines 8-10). Then, for each element $RA[t]$ in RA, we check whether $RA[t]$ equals $|V_S|$. If $RA[t] = |V_S|$, it means u connects all vertices in V_S at t , and we increase the corresponding element $UA[t]$ in Update Array by 1. $UA[t]$ represents the number of the common m -neighbors of all vertices in V_S at t . After processing all vertices in U_S , we obtain the final UA. Count support timestamp for V_S (lines 11-14). Then, we check the value of each element in UA. If there is an element that is no less than τ_U , we add the number of frequency λ' by 1. We return *true* if λ' equals λ , which means that V_S is frequent. Otherwise, we return *false* (line 15).

EXAMPLE 3.1. *Reconsider the graph in Figure 2 with $\tau_U = \tau_V = 2$ and $\lambda = 3$. Figure 3 displays the checking process for the vertex set $V_S = \{v_1, v_2, v_3, v_5\}$, and $U_S = \{u_1, u_2, u_3, u_4, u_5\}$. Based on Algorithm 3, we traverse each vertex in U_S and record its m -degree in V_S at each timestamp. Specifically, when traversing u_1 , we record its m -degree in V_S at each timestamp as "1, 3, 0, 4, 0, 3" in the Reborn Array. Among these m -degrees, only one is equal to $|V_S| = 4$ (the fourth element of the Reborn Array), so we add 1 in the fourth element of the Update Array, i.e., $UA[4] = 1$, indicating that there is one vertex connecting all vertices in V_S at $t = 4$. After traversing u_1 , we clear Reborn Array and maintain Update Array. Similarly, we traverse all the vertices in U_S and check all the elements of final Update Array. We find that there are three elements of the final Update Array that are no less than $\tau_U = 2$, i.e., $UA[1] = 2, UA[3] = 2, UA[4] = 2$, which correspond to the timestamp $t = 1, t = 3$ and $t = 4$, respectively. This means that there are three timestamps at which the vertices in V_S have no less than two common m -neighbors, i.e., the frequency of $\{v_1, v_2, v_3, v_5\}$ is 3. Therefore, this vertex set satisfies the frequency constraint.*

Discussion. To compute the valid candidate set, we can iteratively invoke Algorithm 3 to examine the frequency of the newly obtained

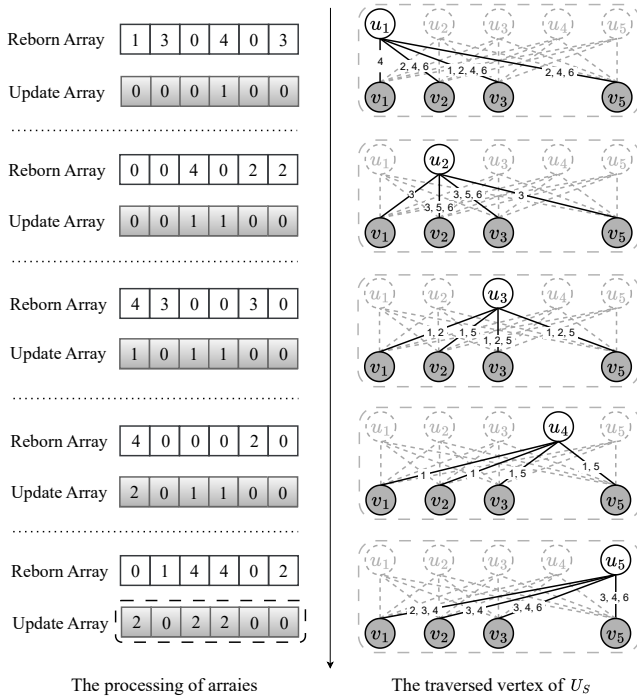


Figure 3: Frequency verification example for vertex set $\{v_1, v_2, v_3, v_5\}$ with $\tau_U = \tau_V = 2$ and $\lambda = 3$

vertex set, i.e., line 8 in Algorithm 1. After checking all the vertices in C_V , we obtain the valid candidate set C_V^* . The corresponding time complexity is shown in Theorem 3.1.

THEOREM 3.1. *Based on Algorithm 3, the time complexity of computing the valid candidate set C_V^* is $O(|V| \cdot d_{\max}(u) \cdot d_{\max}(v) \cdot |T|)$, where $d_{\max}(u)$ and $d_{\max}(v)$ are the largest s -degree of the vertices in U and V , respectively.*

3.4 Maximality Verification

Due to the properties of MFG, the traditional maximality checking technique for the maximal biclique enumeration [3, 11, 49] cannot be used in our problem. To check the maximality of obtained vertex set V_S , a naive method is to compare it with all the obtained results. If V_S is the subset of an existing result, i.e., not maximal, we can skip it. Otherwise, for these vertex sets that are the subset of V_S , we remove them from the currently found result set and add V_S into the result set. However, this method requires extensive computation, since it involves numerous set comparisons. In this section, we introduce a new maximality checking technique. Specifically, we use the vertex set X_V to store the vertices that are previously processed and can be included in at least one λ -frequency group in the current branch, i.e., adding v into X_V after line 17 in Algorithm 1. Based on X_V , we present the details of checking method below.

LEMMA 3.3 (MAXIMALITY VERIFICATION). *Given the current processing tuple (U_S, V_S, C_V, X_V) , V_S is a λ -frequency group. V_S is an MFG iff i) $C_V^* = \emptyset$ and ii) any $v \in X_V$ cannot form a λ -frequency group with V_S , i.e., $\nexists v \in X_V$ s.t. $V_S \cup \{v\}$ is a λ -frequency group.*

By applying Lemma 3.3, we can eliminate the enormous comparisons to determine the containment relationship between two

Table 1: Comparison of FilterV and VFree in computing valid candidate set and checking maximality on D14

$(\tau_U, \tau_V, \lambda)$	(8,4,8)	(9,5,8)	(10,6,6)	(10,6,10)
FilterV-CM (%)	88.26%	88.52%	85.05%	86.68%
FilterV-CM (s)	899.30s	702.27s	617.14s	248.64s
VFree-CM (s)	63.80s	28.78s	26.65s	9.04s

vertex sets. The remaining issue is how to efficiently check whether there exists a vertex $v \in X_V$ such that $V_S \cup \{v\}$ is a λ -frequency group. To achieve this, we first apply the filtering rule (Lemma 3.2) to efficiently shrink X_V . Then we apply the verification method (Algorithm 3) to check the frequency of the vertex set obtained by adding each remaining vertex of X_V separately into V_S .

4 VERIFICATION-FREE APPROACH

FilterV is significantly faster than BK-ALG, even if we equip BK-ALG with graph filtering technique (i.e., BK-ALG+ in the experiment). However, FilterV still suffers from some limitations, due to its search philosophy. When computing the valid candidate set C_V^* and checking maximality, FilterV needs to iterate over each vertex in C_V and X_V separately for frequency verification, which could be time-consuming. In addition, due to the vertex-oriented search paradigm, FilterV cannot effectively utilize the shared information among different computations. For instance, the invalid timestamp information cannot be inherited by the subsequent computations. In Table 1, we report the execution time of computing valid candidate set and checking maximality within FilterV (i.e., FilterV-CM (s)), and the percentage FilterV-CM (%) of overall execution time on a network D14 with more than 60 million edges (the dataset details can be found in the full version [42]). As we can observe, the two components, i.e., compute the valid candidate set and maximality verification, take up a majority of the overall performance. Therefore, if we can reduce the frequency and maximality verification cost or even avoid such cost to some extent, the overall performance can be significantly improved.

This motivates us to develop a strategy without verification, i.e., derive the valid candidate set directly. Specifically, in this section, we develop a timestamp-oriented search paradigm. It iterates through the timestamps to obtain the valid candidate set C_V^* using the dynamic counting structures proposed (Section 4.1), where the unpromising timestamp can be skipped and common neighbor information can be carried forward. Then, we present the verification-free algorithm (VFree) in Section 4.2. Theoretically, VFree can significantly reduce each valid candidate set computation cost in FilterV by a factor of $O(|V|)$ (details are shown in Theorem 4.2). Besides, based on the search paradigm proposed in VFree, we can avoid explicit maximality verification. In Table 1, VFree-CM (s) denotes the execution time of computing the valid candidate set and verifying maximality in VFree, which is much faster than that in FilterV.

4.1 Valid Candidate Set Computation

General idea. In the verification-free framework, we process timestamps sequentially to compute the valid candidate set C_V^* . The procedure on one timestamp t consists of four steps, where V_S is the current processing vertex set.

- *Step 1: ascertain from U .* Obtain the common m -neighbors of V_S in snapshot G_t and store them into $cand_U$.
- *Step 2: termination check.* If $|cand_U| < \tau_U$, stop processing t and move to the next timestamp, since V_S cannot form a (τ_U, τ_V) -biclique in G_t . Otherwise, t is a **survived timestamp**.
- *Step 3: reverse-ascertain from V .* Find all the vertices in $V_t \setminus V_S$ that connect at least τ_U vertices in $cand_U$ and store them into $cand_V$.
- *Step 4: survived timestamp update.* Increase the survived timestamp count of $V_S \cup \{v\}$ by 1 for $v \in cand_V$.

EXAMPLE 4.1. Reconsider the graph in Figure 2 with $\tau_U = \tau_V = 2$. Suppose $V_S = \{v_1, v_2\}$ and $t = 1$. Step 1) We store the common m -neighbors of V_S in G_1 to $cand_U$, i.e., $cand_U = \{u_3, u_4\}$. Step 2) Since $|cand_U| \geq \tau_U = 2$, $t = 1$ is a survived timestamp for V_S . Step 3) We proceed to examine the common m -neighbors of $cand_U$. Besides V_S , both v_3 and v_5 connect $\tau_U = 2$ vertices in $cand_U$. Thus, $cand_V = \{v_3, v_5\}$, and we increase the survived timestamp count of $V_S \cup \{v_3\}$ and $V_S \cup \{v_5\}$ by 1, respectively in Step 4.

Based on the above procedure, we iterate through all the timestamps. It is easy to verify that, after processing all the timestamps, C_V^* is the subset of $cand_V$, where for each $v \in C_V^*$, its survived timestamp count is no less than λ . In the following, we present the details about how to i) enable the inheritance of invalid timestamp information and ii) accelerate the neighbor computation.

Timestamp inheritance. To enable the inheritance of timestamps, we further maintain a timestamp set C_T for V_S , which stores all the survived timestamps when computing C_V^* of V_S . Then, according to Lemma 4.1, we only need to check the timestamps in C_T when processing the subsequent branches of V_S .

LEMMA 4.1 (TIMESTAMP SKIPPING RULE). Given the processing vertex set V_S and computed C_T , if $t \notin C_T$, we can skip the processing at the timestamp t in the subsequent search branches of V_S , i.e., the branch by adding any vertex into V_S .

Neighbor computation acceleration. To compute and maintain the neighbor information, we design three dynamic counting data structures as follows.

- A two-dimensional array, denoted by $cnt_U[t][u]$, to count the number of m -neighbors of $u \in U$ in V_S at $t \in C_T$.
- A two-dimensional array, denoted by $cnt_V[t][v]$, to count the number of m -neighbors of $v \in V \setminus V_S$ in $cand_U$ at $t \in C_T$.
- An one-dimensional array, denoted by $cnt_T[v]$, used to record the number of survived timestamps for $V_S \cup \{v\}$, where $v \in C_V$.

For a given timestamp $t \in C_T$ and the current processing vertex set V_S , if $cnt_U[t][u] = |V_S|$, it means u is the common m -neighbors of all the vertices in V_S in G_t . Therefore, u can be stored in $cand_U$ (correspond to Step 1). If $|cand_U| < \tau_U$, we can skip the timestamp (Step 2). Similarly, if $cnt_V[t][v] \geq \tau_U$, it means v connects to at least τ_U vertices in $cand_U$ at t and can be stored in $cand_V$ (Step 3). Then, we increase $cnt_T[v]$ by 1, which means t is a survived timestamp for $V_S \cup \{v\}$ (Step 4).

EXAMPLE 4.2. Reconsider the graph in Figure 2 with $\tau_U = \tau_V = 2$ and $\lambda = 3$. Suppose the current processing set $V_S = \{v_1\}$. To compute the valid candidate set, Figure 4(a) visualizes the results of three data structures after processing timestamp $t = 1$. When $V_S = \{v_1\}$, we have $C_V = \{v_2, v_3, v_4, v_5\}$ and $C_T = \{1, 2, 3, 4, 5, 6\}$,

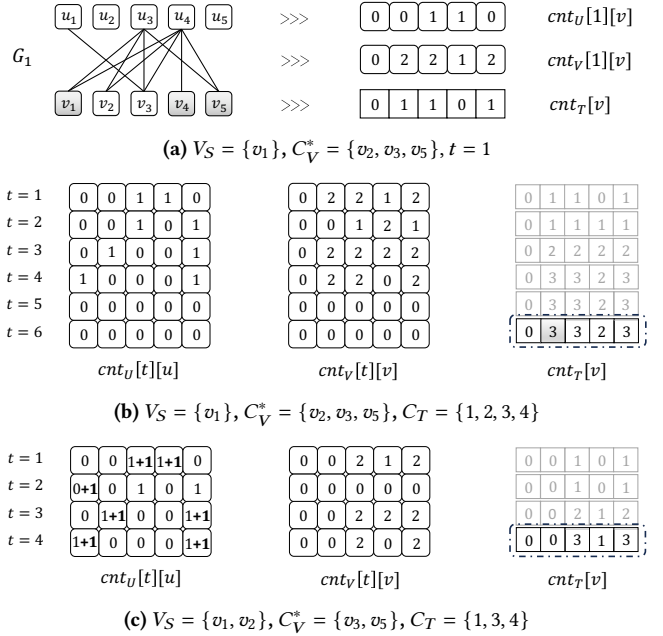


Figure 4: Running example of computing valid candidate set using VFree with $\tau_U = \tau_V = 2$ and $\lambda = 3$

since it is the first vertex explored. The three data structures are initialized with 0. The m -neighbors of v_1 at $t = 1$ are u_3 and u_4 . Thus, we increase $cnt_U[1][3]$ and $cnt_U[1][4]$ by 1, respectively. As $cnt_U[1][3] = cnt_U[1][4] = |V_S| = 1$, we have $cand_U = \{u_3, u_4\}$. Since $|cand_U| \geq \tau_U = 2$, $t = 1$ is a survived timestamp. Next, we process the m -neighbors of u_3 at $t = 1$, i.e., $\{v_1, v_2, v_3, v_5\}$. Thus, we assign a value of 1 to $cnt_V[1][2]$, $cnt_V[1][3]$, $cnt_V[1][5]$ for vertex v_2, v_3 and v_5 , respectively. Similarly, we process the m -neighbors of u_4 and increase the corresponding count. We have $cnt_V[1][2] = 2$, $cnt_V[1][3] = 2$, $cnt_V[1][4] = 1$ and $cnt_V[1][5] = 2$. Finally, since $cnt_V[1][2] = cnt_V[1][3] = cnt_V[1][5] = 2$, we have $cand_V = \{v_2, v_3, v_5\}$ and increment $cnt_T[2]$, $cnt_T[3]$ and $cnt_T[5]$ by 1.

After processing each timestamp in C_T , we get the final $cnt_T[v]$. For a vertex $v' \in V \setminus V_S$, if $cnt_T[v'] \geq \lambda$, it means v' is in the valid candidate set of V_S .

EXAMPLE 4.3. Following Example 4.2, Figure 4(b) shows the final results of the three data structures after processing all the timestamps in C_T . As shown, $cnt_T[v_2] = cnt_T[v_3] = cnt_T[v_5] = 3 \geq \lambda$. Therefore, we have the valid candidate set $C_V^* = \{v_2, v_3, v_5\}$, and its survived timestamp set is $\{1, 2, 3, 4\}$. The first column of $cnt_V[t][v]$ is 0. This is because v_1 is in the current processing vertex set V_S , and we do not record the value.

In the search branch of V_S , V_S is continuously expanded by adding new vertices into V_S . It means that the values in $cnt_U[t][u]$ are non-decreasing. Thus, we can incrementally maintain the data structure $cnt_U[t][u]$, i.e., $cnt_U[t][u]$ can be inherited in the subsequent search to avoid duplicated computation. $cnt_V[t][v]$ and $cnt_T[v]$ are temporally maintained when computing the valid candidate set C_V^* for a V_S , and will be reset after each C_V^* computation.

EXAMPLE 4.4. Following Example 4.3, suppose we add v_2 to expand the vertex set $\{v_1\}$, and compute the valid candidate set C_V^* for $V_S =$

$\{v_1, v_2\}$. Figure 4(c) shows the final results of the three data structures. The bold values in $cnt_U[t][u]$ denote the incremental computations. Note that, since the survived timestamps of $\{v_1\}$ is $C_T = \{1, 2, 3, 4\}$, we can skip the processing of $t = 5$ and 6. Similarly, to compute C_V^* , we need to iterate through C_T . The difference is that, when updating $cnt_U[t][u]$, we only need to increment the values based on the newly added vertex v_2 instead of V_S . For instance, at $t = 1$, the m -neighbors of v_2 is $\{u_3, u_4\}$. Therefore, we increment $cnt_U[1][3]$ and $cnt_U[1][4]$ by 1, respectively. After processing all the timestamps in C_T , We have $C_V^* = \{v_3, v_5\}$ by checking $cnt_T[v]$.

4.2 VFree Algorithm

By incorporating the above techniques, we present our verification-free (VFree) approach, whose details are shown in Algorithm 4. *Initialization (lines 1-7)*. \mathcal{R} is used to store all the MFGs. We also use the graph filtering technique (Algorithm 2 in Section 3.2) to reduce the search space. Then, we reassign the id of each vertex in V in ascending order of the structural degree. Ties are randomly broken if vertices have the same structural degree. Note that, in VFree, we process vertices of V in the order of vertex id to ensure the maximality of the result, i.e., avoid explicit maximality verification. That is, for any given vertex id setting, as long as we process the vertices in ascending order of the id, the algorithm's properties and correctness can be ensured. We initialize cnt_U , cnt_V and cnt_T in lines 4-7. In line 8, we invoke the procedure VerifyFreeMFG to enumerate all MFGs, where we initialize V_S with \emptyset , C_V with V and C_T with \mathcal{T} . Generally, V_S is the current processing vertex set, C_V is the valid candidate set and C_T is the survived timestamp set for V_S . Details of the procedure VerifyFreeMFG are shown in lines 9-43.

In lines 11-36, we compute the valid candidate set for $V_S' = V_S \cup \{v\}$, where the *notRepeat* flag in lines 31 and 35 is used for later maximality check. We initialize C_V^* , C_T' and $cand_V$ with \emptyset (line 11), and perform the following operations in turn under each timestamp $t \in C_T$ (lines 12-30). *Step 1: ascertain from U (lines 13-17)*. We first initialize $cand_U$ with \emptyset to denote the common m -neighbor set of all the vertices in V_S' at t and $visit_V$ with \emptyset to help maintain the information in $cnt_V[t][v]$. Recall that $cnt_V[t][v]$ and $cnt_T[v]$ are temporally maintained, and $cnt_U[t][u]$ can be inherited. $visit_V$ can help restore the corresponding information in $cnt_V[t][v]$ instead of repeated initialization. For each m -neighbor $u \in \Gamma(v, t)$, we increase $cnt_U[t][u]$ by 1 denoting that u is connected with v at t (line 16). If $cnt_U[t][u] = |V_S'|$, which means u connects all the vertices in V_S' , we push u into $cand_U$ (line 17). *Step 2: termination check (lines 18-19)*. After processing all the m -neighbors of v , if $|cand_U|$ is less than τ_U , which means that t is not a survived timestamp for V_S' , we can skip t . Otherwise, we push t into C_T' that records the survived timestamps for V_S' . Then we need to explore all the vertices in $cand_U$ in lines 20-30. *Step 3: reverse-ascertain from V (lines 20-29)*. For each vertex $u' \in cand_U$, we need to traverse its m -neighbors $\Gamma(u', t)$ iteratively. Specifically, if its m -neighbor $v' \in \Gamma(u', t)$ exists in V_S' , we skip the current vertex (line 22). If $v' \notin visit_V$, which means that it is the first time to visit v in the current search, we set $cnt_V[t][v']$ with 1 and push v' into $visit_V$ (lines 23-24). Otherwise, we add $cnt_V[t][v']$ by 1 (lines 25-26). If $cnt_V[t][v'] = \tau_U$, which means v' has no less than τ_U common m -neighbors with all vertices in V_S' at timestamp t (i.e., v' connects at least τ_U vertices

Algorithm 4: Verification-Free (VFree)

```

Input   :  $\mathcal{G} = (U, V, \mathcal{E})$ : a temporal bipartite graph,
            $\tau_U, \tau_V$ : size constraints,  $\lambda$ : frequency constraint
Output :  $\mathcal{R}$ : all the MFGs
1  $\mathcal{R} \leftarrow \emptyset$ ;
2  $\mathcal{G} \leftarrow \text{GFCore}(\mathcal{G}, \tau_U, \tau_V, \lambda)$ ;           /* graph filter: Algorithm 2 */
3 Reassign vertex id of  $V$  in ascending order of the structural degree;
4 for each  $t \in \mathcal{T}$  do
5   for each  $u \in U$  do  $cnt_U[t][u] = 0$ ;
6   for each  $v \in V$  do  $cnt_V[t][v] = 0$ ;
7 for each  $v \in V$  do  $cnt_T[v] = 0$ ;
8 VerifyFreeMFG( $\emptyset, V, \mathcal{T}$ );
9 Procedure VerifyFreeMFG( $V_S, C_V, C_T$ )
10 for each  $v \in C_V$  do
11    $V_S' \leftarrow V_S \cup \{v\}, C_V^* \leftarrow \emptyset, C_T' \leftarrow \emptyset, cand_V \leftarrow \emptyset$ ;
12   for each  $t \in C_T$  do
13      $cand_U \leftarrow \emptyset$ ;
14      $visit_V \leftarrow \emptyset$ ;           /* store vertices for the first verify */
15     for each  $u \in \Gamma(v, t)$  do
16        $cnt_U[t][u] = cnt_U[t][u] + 1$ ;
17       if  $cnt_U[t][u] = |V_S'|$  then  $cand_U.push(u)$ ;
18   if  $|cand_U| < \tau_U$  then continue;
19    $C_T'.push(t)$ ;           /* survived timestamp set for  $V_S'$  */
20   for each  $u' \in cand_U$  do
21     for each  $v' \in \Gamma(u', t)$  do
22       if  $v' \in V_S'$  then continue;
23       if  $v' \notin visit_V$  then
24          $cnt_V[t][v'] = 1, visit_V \leftarrow visit_V \cup \{v'\}$ ;
25       else
26          $cnt_V[t][v'] = cnt_V[t][v'] + 1$ ;
27       if  $cnt_V[t][v'] = \tau_U$  then
28         if  $cnt_T[v'] = 0$  then
29            $cand_V.push(v')$ ;
30          $cnt_T[v'] = cnt_T[v'] + 1$ ;
31    $notRepeat \leftarrow true$ ;
32   for each  $v' \in cand_V$  do
33     if  $cnt_T[v'] < \lambda$  then  $cnt_T[v'] = 0, continue$ ;
34      $cnt_T[v'] = 0$ ;
35     if  $v' < v$  then  $notRepeat \leftarrow false$ ;
36     else  $C_V^*.push(v')$ ;           /* valid candidate set for  $V_S'$  */
37   if  $|V_S'| + |C_V^*| \geq \tau_V \wedge |C_T'| \geq \lambda$  then
38     sort  $C_V^*$  based on vertex id; /* ensure the processing order */
39     VerifyFreeMFG( $V_S', C_V^*, C_T'$ );
40     if  $|C_V^*| = 0 \wedge notRepeat$  then  $\mathcal{R} \leftarrow \mathcal{R} \cup \{V_S'\}$ ; /* MFG */
41   for each  $t \in C_T$  do
42     for each  $u \in \Gamma(v, t)$  do
43        $cnt_U[t][u] = cnt_U[t][u] - 1$ ;           /* update  $cnt_U$  */

```

in $cand_U$), we check whether $cnt_T[v']$ equals 0 or not (lines 27-28). As discussed before, $cnt_T[v']$ denotes the number of survived timestamps for the set $V_S' \cup \{v'\}$. If $cnt_T[v'] = 0$, we push v' into $cand_V$ (lines 28-29). *Step 4: survived timestamp update (line 30)*. We add the number of survived timestamps for v' (i.e., $cnt_T[v']$) by 1. *Valid candidate set computation (lines 31-36)*. After processing all the timestamps in C_T , we can obtain the final $cnt_T[v']$ for each vertex $v' \in cand_V$. If $cnt_T[v'] < \lambda$, we set $cnt_T[v']$ as 0 and skip v' . If $v' < v$, we set *notRepeat* as *false*. If $v' \geq v$, we push v' into C_V^* as the valid candidate vertex for V_S' .

We recursively invoke VerifyFreeMFG if the size constraint is satisfied (lines 37-39). C_V^* is sorted to ensure the processing order of vertices in V . We restore $cnt_U[t][u]$ by reducing 1 for each m -neighbor u of v to utilize it in the next iteration (lines 41-43). Unlike

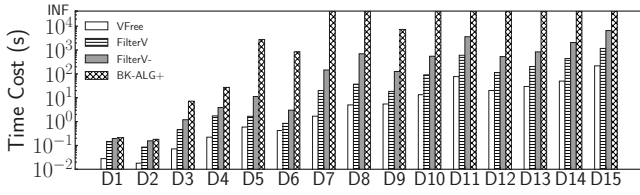


Figure 5: Response time on all the datasets

FilterV, which explicitly conducts maximality verification through expensive computation (i.e., Section 3.4), in VFree, if $C_V^* = 0$ and *notRepeat* is true, it means we find an MFG that can be added to the result set \mathcal{R} (line 40). The correctness is shown in Theorem 4.1.

THEOREM 4.1 (ALGORITHM CORRECTNESS). *Given a temporal bipartite graph \mathcal{G} , three positive integers τ_U , τ_V and λ , Algorithm 4 can return all the MFGs in \mathcal{G} .*

THEOREM 4.2 (TIME COMPLEXITY ANALYSIS). *The time complexity of computing the valid candidate set C_V^* for a vertex set V_S based on Algorithm 4 is $O(d_{\max}(u) \cdot d_{\max}(v) \cdot |\mathcal{T}|)$.*

According to Theorems 3.1 and 4.2, VFree can reduce each valid candidate set computation cost in FilterV by a factor of $O(|V|)$.

5 EXPERIMENT

In this section, we conduct the experiment on 15 real-world datasets to verify the performance of proposed techniques. More details about experiment setup can be found in the full version [42].

5.1 Efficiency Evaluation

Exp-1: Experiments over all the datasets. Figure 5 reports the response time of BK-ALG+, FilterV-, FilterV and VFree over all the datasets with the default settings. FilterV- is faster than BK-ALG+ due to the optimization in the BK framework, i.e., compute the valid candidate set first before the exploration. FilterV further improves FilterV- because of the filtering and verification techniques developed. VFree outperforms the other algorithms with a significant margin. This is because *i)* VFree is time-oriented when computing the valid candidate, and the dynamic counting structures can extraordinarily speedup the computation; *ii)* the search paradigm also significantly reduces the cost of the maximality verification. For instance, on the dataset D14, BK-ALG+ fails to complete the computation within 12 hours. FilterV- and FilterV return the result in 2081 seconds and 445 seconds, respectively. VFree can return the result in 50 seconds. On the largest dataset D15 with more than 70 million edges, the response time of VFree is 218 seconds. In datasets D5, D6 and D9, where BK-ALG+ can finish in a reasonable time, VFree can achieve up to three orders of magnitude speedup.

Exp-2: Response time by varying parameters. In Figures 6(a) to 6(f), we report the response time of BK-ALG+, FilterV-, FilterV and VFree on the two largest datasets D14 and D15 by varying parameters τ_U , τ_V and λ , respectively. Note that, in these experiments, we use default settings for the other unchanged parameters. As shown, VFree is faster than the other three algorithms under all the parameter settings. Besides, the response time of all the algorithms decreases when the parameters increase. This is because more vertices can be skipped and larger search space can be pruned by the graph filtering technique due to the tighter constraints. BK-ALG+

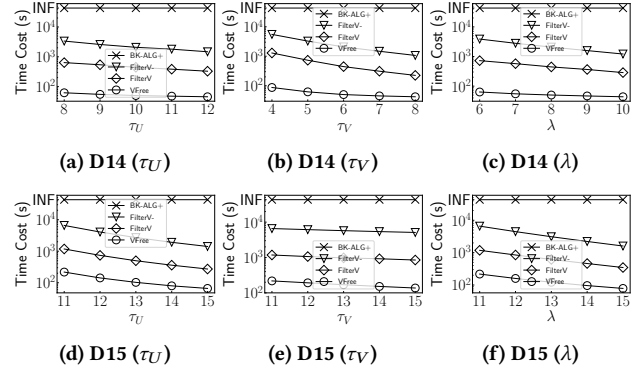


Figure 6: Response time by varying τ_U , τ_V and λ

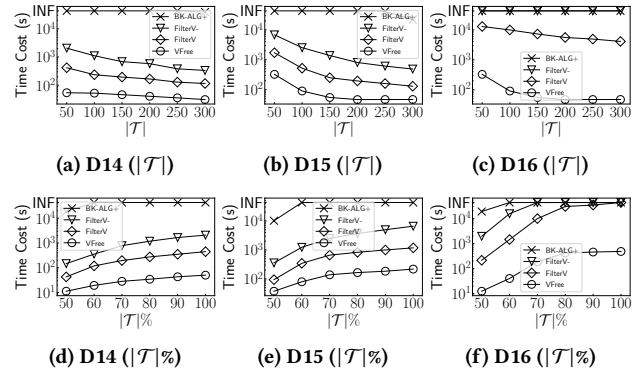


Figure 7: Response time by varying $|\mathcal{T}|$ and $|\mathcal{T}\%$

fails to return the results in a reasonable time on most datasets even with large parameters.

Exp-3: Response time by varying $|\mathcal{T}|$ and $|\mathcal{T}\%$. In this experiment, we report the response time of BK-ALG+, FilterV-, FilterV and VFree by varying the settings of time span on D14, D15 and D16. D16² is a larger graph (with over 256 million edges) employed to further demonstrate the performance of proposed techniques, and default parameters for D16 are $\tau_U = 10$, $\tau_V = 15$ and $\lambda = 10$. We conduct two sets of experiments. *i)* For each dataset, we generate 6 temporal bipartite graphs by varying $|\mathcal{T}|$, i.e., setting different time scales. The results are shown in Figures 7(a)-7(c). As observed, VFree outperforms the other algorithms by a significant margin and scales well. For example, when $|\mathcal{T}| = 300$ in D16, BK-ALG+ and FilterV- cannot return results within 12 hours. The response time of FilterV and VFree are 4105.32 seconds and 46.86 seconds, respectively. The response time of all the algorithms decreases when $|\mathcal{T}|$ increases. This is because, as $|\mathcal{T}|$ increases, the number of edges in each snapshot decreases, leading to more space pruned based on the cohesive constraint. *ii)* For each dataset, we keep $|\mathcal{T}|$ unchanged as the default value, i.e., 66 for D14, 67 for D15 and 134 for D16, and generate 6 temporal bipartite graphs by covering the edges in the first 50%-100% timestamps, i.e., $|\mathcal{T}\%$. The results are shown in Figures 7(d)-7(f). As observed, VFree outperforms the other algorithms. The response time of all algorithms grows with the increase of $|\mathcal{T}\%$ due to the larger search space. The above

²<http://konect.cc/networks/yahoo-song/>

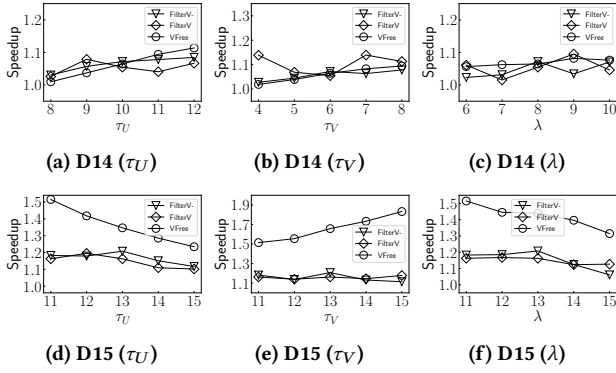


Figure 8: Evaluation of the vertex ID reorder technique

experiments demonstrate that our proposed algorithms are scalable towards different temporal settings.

Exp-4: Evaluation of the ID reorder technique. To evaluate the impact of the ID reorder technique in VFree, we conduct the experiments on D14 and D15 by varying τ_U , τ_V and λ , respectively. For FilterV-, FilterV and VFree, we report the *speedup ratio* of ID reorder technique, which is calculated by $\frac{\text{response time of algorithms without ID reorder}}{\text{response time of algorithms with ID reorder}}$. The results are shown in Figure 8. Note that, we omit BK-ALG+ here, since it cannot finish on D14 and D15 within 12 hours even with the ID reorder technique. As observed, the ID reorder technique can improve the efficiency of algorithms. For instance, the speedup for VFree is up to 1.8x.

We also conduct the experiments to evaluate the scalability, graph filtering technique, candidate filtering rule and the verification technique, whose details can be found in the full version [42].

5.2 Effectiveness Evaluation

Exp-5: Multimorbidity detection. To demonstrate the effectiveness of proposed model, we conduct a case study on D1 (MIMIC-III) for potential multimorbidity detection. MIMIC-III is a real clinical dataset representing relationships between patients and health conditions, where the timestamp denotes the time of diagnosis [1, 2, 15]. By applying our model MFG, we can model the situation where multiple health conditions appear for different patients at multiple times simultaneously. The partially identified MFGs are shown in Table 2. For instance, ‘SEPSIS’ and ‘PNEUMONIA’ are highly correlated, since pneumonia is a common cause of sepsis. Moreover, we compare MFG with two variants based on existing models, i.e., *i*) maximal frequent (τ_U, τ_V) -biclique (MFB) and *ii*) maximal static group (MSG). MFB is the maximal (τ_U, τ_V) -biclique with frequency at least λ , i.e., appearing in at least λ snapshots. MSG is the maximal group included in (τ_U, τ_V) -biclique of the corresponding static graph. As shown in Table 2, these two models cannot obtain practical results. Specifically, we cannot find any results by applying MFB model due to the tight constraint. For MSG, the identified groups may be too large to provide practical information due to neglect of temporal aspect.

Exp-6: Fraud detection. In this case study, we demonstrate the application for a fraud detection task in face of the camouflage attack [18] on D9 (Amazon Ratings) compared with MFB and MSG. Considering a practical scenario with a random camouflage attack

Table 2: Case study on D1 ($\tau_U = \tau_V = 2$ and $\lambda = 6$)

Model	Partial results
MFG	{SEPSIS, PNEUMONIA}, {GASTROINTESTINAL BLEED, LOWER GI BLEED}, {ASTHMA, COPD EXACERBATION, PNEUMONIA}, {UPPER GI BLEED, LOWER GI BLEED}, ...
MSG	{CHRONIC OBST PULM DISEASE, CHRONIC OBSTRUCTIVE PULMONARY, RESPIRATORY FAILURE, PNEUMONIA, COPD EXACERBATION, ASTHMA}, {HYPERTENSIVE EMERGENCY, HYPERTENSIVE URGENCY, ABDOMINAL PAIN, DIABETIC KETOACIDOSIS}, ...
MFB	N/A

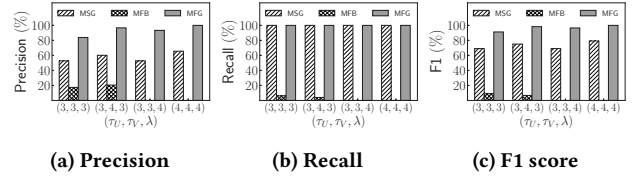


Figure 9: Case study on D9

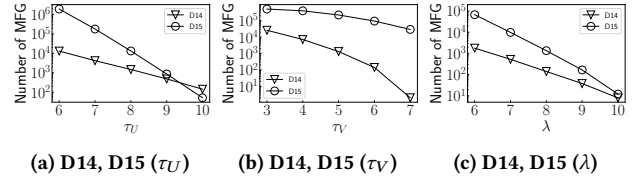


Figure 10: Number of MFGs by varying parameters

on D9, we randomly select five timestamps and choose 2K products at each selected timestamp. Specifically, at each of the selected timestamps, we introduce a fraudulent block containing 1K fake users, 1.5K fake comments and 0.5K camouflage comments to the data. The same fake users are used across selected timestamps. Besides, at each selected timestamp, fake comments (resp. camouflage comments) are generated randomly between these injected fake users and the chosen products (resp. real products). We categorize all users identified within the detected subgraphs as fake items and all others as real one, and report the precision, recall and F1 score in Figure 9. As observed, MFG can achieve better results compared to the others, which demonstrates the effectiveness of our proposed model. For MFB, it cannot find results under some settings, and we omit the corresponding value in the figure. For MSG, it usually has high recall but low precision since it would find large groups involving many users.

Exp-7: Number of MFGs. In this experiment, we report the number of returned MFGs on D14 and D15 by varying different parameters. The results are shown in Figure 10. The settings for other unchanged parameters follow the default values. The number of MFGs decreases with the increase of three parameters due to the tighter constraint required. As shown, even with large parameters, we can still find MFG patterns, which demonstrates the applicability of our model in different settings for real scenarios.

6 RELATED WORK

Graphs are widely used to capture the relationships between entities [10, 37, 41]. In the literature, many subgraph models have

been proposed for temporal unipartite graphs (e.g., [22, 30, 32, 33]). Ma et al. [30] study the connected temporal subgraph whose aggregated graph has the maximum cohesive density on temporal weighted graphs. They develop algorithms based on the assumption that the weights of all edges are increasing or decreasing in the same direction. Li et al. [23] define a novel k -core based model to capture the persistence of a community. They first propose a novel temporal graph reduction method and then develop a novel branch and bound algorithm with several powerful pruning rules to find the result. [32] leverages the clique model and investigates the σ -periodic k -clique enumeration problem. The authors first prune the input temporal graph based on two novel relaxed periodic clique models. Then, they propose a graph transformation technique and efficiently enumerate the results on the transformed graph. Qin et al. [31] propose a novel concept named (μ, τ, ϵ) -stable core, to characterize the stable core nodes of the clusters. They call a node u is (μ, τ, ϵ) -stable core if it has no less than μ neighbors that are simultaneously similar to itself in at least τ snapshots of the temporal graph. Using the similar ideas of [32], they first propose the weak and strong cores to significantly prune the unpromising nodes, and then identify all the stable clusters from the remaining graph. There are also some studies that focus on bursting subgraph mining, which are established on the time-interval based constraint. In [33], Qin et al. study the problem of mining bursting cores, where the (l, δ) -maximal bursting core model is developed. They propose a novel dynamic programming algorithm to speedup the calculation of the segment density. Zhang et al. [48] propose a frequency bursting pattern in temporal graphs. It tries to model the interactive behaviors that accumulate their frequencies the fastest. Chu et al. [13] aim to find the top- k density bursting subgraphs, each of which is the subgraph that accumulates their density at the fastest speed in the temporal graph. In [28], a new metric named T-cohesiveness is proposed by jointly considering the time and topology dimensions. For temporal bipartite graphs, Li et al. [22] study the community search problem, and a persistent community model is proposed based on (α, β) -core, which has different semantics from ours. If we consider a temporal graph as a set of snapshots, the problems of frequent subgraph mining and multi-layer cohesive subgraph mining are correlated. Specifically, the goal of frequent subgraph mining is identifying frequently appeared subgraphs from a collection of graphs (e.g., [19, 26]). As for multi-layer graph mining, many models are proposed to pinpoint structures, such as the d -coherent core and firm truss [7, 17, 52]. Besides the above studies, there are some works focusing on dynamic graphs (e.g., [14, 27, 38]). For instance, Tang et al. [38] propose a novel (θ, k) -core reliable community in the weighted dynamic networks. That is, a k -core with each edge weight no less than the weight threshold θ spans over a period of time. To find the most reliable local community with the highest reliability score, they first filter the unpromising edges from the graph. Then they develop an index structure and devise an index-based dynamic programming search algorithm. Although many studies have been conducted over temporal or dynamic graphs, it is non-trivial to extend the existing solutions to our problem.

As the most cohesive structure in bipartite graphs, the biclique model has attracted significant attention. The problem of maximal biclique enumeration has been widely studied in static bipartite graphs [3, 4, 11, 25, 49]. Most of the existing biclique enumeration

algorithms conduct the search by expanding the vertices from one side. Then, we can intersect their neighbors to form the corresponding biclique [3, 11, 49]. In [16], the authors reduce the maximal biclique enumeration problem to the maximal clique enumeration problem. In [49], Zhang et al. remove the unpromising candidates from the search branches by employing the branch-and-bound framework. In [3], Abidi et al. develop a novel pivot-based technique to block non-maximal search branches. In [11], the authors accelerate the process of maximal biclique enumeration by proposing the concepts of unilateral coreness for individual vertices, unilateral order for each vertex set and unilateral convergence for a large sparse bipartite graph. Generally, in the unilateral core, for all vertices on one side of it, the number of their two-hop neighbors within must be no less than k . The concepts of unilateral coreness, order and convergence are proposed based on the concept of unilateral core. However, these concepts are orthogonal to our model and have different semantics. In [47], the authors propose a dense bipartite subgraph model, which considers similarity between vertices from the same side on static bipartite graphs. To the best of our knowledge, although there are a few works considering the one-layer properties in bipartite graphs, no existing works study the unilateral frequency model. There are also some studies that consider the biclique problem in various bipartite graphs, such as signed bipartite graph [35, 36]. As we can see, few studies consider the case of temporal bipartite graphs. Moreover, due to the unilateral property of our model, the existing studies cannot be extended to solve our problem efficiently.

7 CONCLUSION

Temporal bipartite graph is an important data structure to model many real-world applications. To analyze the properties of temporal bipartite graphs, in this paper, we propose a novel model, named maximal λ -frequency group, by considering both unilateral cohesiveness and temporal aspect. We first introduce a filter-and-verification method by extending the BK framework. Novel filtering techniques and array-based checking method are developed. To further improve the performance, a verification-free approach is proposed based on advanced dynamic counting strategy, which can significantly reduce the cost of valid candidate set computation and avoid explicit maximality verification. Experiments over 15 real-world datasets confirm the efficiency and effectiveness of proposed techniques and model.

ACKNOWLEDGMENTS

This work is partially supported by ARC DP230101445, ARC DP240101322, ARC DE240100668, ARC DP210101393, ARC LP210301046, ARC FT200100787, NSFC U2241211, NSFC U20B2046, and 23H020101910. Renjie Sun and Xiaoyang Wang are the corresponding authors.

REFERENCES

- [1] Goldberger A, Amaral L, Glass L, Hausdorff J, Ivanov PC, Mark R, Mietus JE, Moody GB, Peng CK, and Stanley HE. 2000. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation* 101 (23) (2000), e215–e220.
- [2] Johnson A, Pollard T, and Mark R. 2016. MIMIC-III Clinical Database (version 1.4). *PhysioNet* (2016).
- [3] Aman Abidi, Rui Zhou, Lu Chen, and Chengfei Liu. 2020. Pivot-based Maximal Biclique Enumeration. In *IJCAI*. 3558–3564.
- [4] Gabriela Alexe, Sorin Alexe, Yves Crama, Stephan Foldes, Peter L Hammer, and Bruno Simeone. 2004. Consensus algorithms for the generation of all maximal bicliques. *Discrete Applied Mathematics* 145, 1 (2004), 11–21.
- [5] Cigdem Aslay, Muhammad Anis Uddin Nasir, Gianmarco De Francisci Morales, and Aristides Gionis. 2018. Mining frequent patterns in evolving graphs. In *CIKM*. 923–932.
- [6] Furqan Aziz, Victor Roth Cardoso, Laura Bravo-Merodio, Dominic Russ, Samantha C Pendleton, John A Williams, Animesh Acharjee, and Georgios V Gkoutos. 2021. Multimorbidity prediction using link prediction. *Scientific Reports* 11, 1 (2021), 16392.
- [7] Ali Behrouz, Farnoosh Hashemi, and Laks VS Lakshmanan. 2022. FirmTruss Community Search in Multilayer Networks. *Proceedings of the VLDB Endowment* 16, 3 (2022), 505–518.
- [8] Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. 2013. Copycatch: stopping group attacks by spotting lockstep behavior in social networks. In *WWW*. 119–130.
- [9] Xinwei Cai, Xiangyu Ke, Kai Wang, Lu Chen, Tianming Zhang, Qing Liu, and Yunjun Gao. 2024. Efficient Temporal Butterfly Counting and Enumeration on Temporal Bipartite Graphs. *Proceedings of the VLDB Endowment* 17, 4 (2024), 657–670.
- [10] Chen Chen, Yanping Wu, Renjie Sun, and Xiaoyang Wang. 2023. Maximum Signed θ -Clique Identification in Large Signed Graphs. *IEEE Trans. Knowl. Data Eng.* 35, 2 (2023), 1791–1802.
- [11] Lu Chen, Chengfei Liu, Rui Zhou, Jiajie Xu, and Jianxin Li. 2022. Efficient maximal biclique enumeration for large sparse bipartite graphs. *Proceedings of the VLDB Endowment* 15, 8 (2022), 1559–1571.
- [12] Xiaoshuang Chen, Kai Wang, Xuemin Lin, Wenjie Zhang, Lu Qin, and Ying Zhang. 2021. Efficiently answering reachability and path queries on temporal bipartite graphs. *Proceedings of the VLDB Endowment* 14, 10 (2021), 1845–1858.
- [13] Lingyang Chu, Yanyan Zhang, Yu Yang, Lanjun Wang, and Jian Pei. 2019. Online density bursting subgraph detection from temporal graphs. *Proceedings of the VLDB Endowment* 12, 13 (2019), 2353–2365.
- [14] Daniel J DiTursi, Gaurav Ghosh, and Petko Bogdanov. 2017. Local community detection in dynamic networks. In *ICDM*. 847–852.
- [15] Johnson Alistair EW, Pollard Tom J, Shen Lu, Lehman Li wei H, Feng Mengling, Ghassemi Mohammad, Moody Benjamin, Szolovits Peter, Anthony Celi Leo, and Mark Roger G. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data* 3, 1 (2016), 1–9.
- [16] Alain Gély, Lhouari Nourine, and Bachir Sadi. 2009. Enumeration aspects of maximal cliques and bicliques. *Discrete applied mathematics* 157, 7 (2009), 1447–1459.
- [17] Farnoosh Hashemi, Ali Behrouz, and Laks VS Lakshmanan. 2022. FirmCore Decomposition of Multilayer Networks. In *WWW*. 1589–1600.
- [18] Bryan Hooi, Hyun Ah Song, Alex Beutel, Neil Shah, Kijung Shin, and Christos Faloutsos. 2016. Fraudar: Bounding graph fraud in the face of camouflage. In *SIGKDD*. 895–904.
- [19] Chuntao Jiang, Frans Coenen, and Michele Zito. 2013. A survey of frequent subgraph mining algorithms. *The Knowledge Engineering Review* 28, 1 (2013), 75–105.
- [20] Junghoon Kim, Kaiyu Feng, Gao Cong, Diwen Zhu, Wenyuan Yu, and Chunyan Miao. 2022. ABC: attributed bipartite co-clustering. *Proceedings of the VLDB Endowment* 15, 10 (2022), 2134–2147.
- [21] Sanjukta Krishnagopal, Rainer von Coelln, Lisa M Shulman, and Michelle Girvan. 2020. Identifying and predicting Parkinson’s disease subtypes through trajectory clustering via bipartite networks. *PLoS one* 15, 6 (2020), e0233296.
- [22] Mo Li, Zhiran Xie, and Linlin Ding. 2023. Persistent Community Search Over Temporal Bipartite Graphs. In *ADMA*. 324–339.
- [23] Rong-Hua Li, Jiao Su, Lu Qin, Jeffrey Xu Yu, and Qiangqiang Dai. 2018. Persistent community search in temporal networks. In *ICDE*. 797–808.
- [24] Boge Liu, Long Yuan, Xuemin Lin, Lu Qin, Wenjie Zhang, and Jingren Zhou. 2019. Efficient (α, β) -core Computation: an Index-based Approach. In *WWW*. 1130–1141.
- [25] Guimei Liu, Kelvin Sim, and Jinyan Li. 2006. Efficient mining of large maximal bicliques. In *International Conference on Data Warehousing and Knowledge Discovery*. 437–448.
- [26] Muiy Liu and Pan Li. 2022. SATMargin: Practical Maximal Frequent Subgraph Mining via Margin Space Sampling. In *WWW*. 1495–1505.
- [27] Xuanming Liu, Tingjian Ge, and Yinghui Wu. 2019. Finding densest lasting subgraphs in dynamic graphs: A stochastic approach. In *ICDE*. 782–793.
- [28] Yunkai Lou, Chaokun Wang, Tiankai Gu, Hao Feng, Jun Chen, and Jeffrey Xu Yu. 2021. Time-topology analysis. *Proceedings of the VLDB Endowment* 14, 13 (2021), 3322–3334.
- [29] Bingqing Lyu, Lu Qin, Xuemin Lin, Ying Zhang, Zhengping Qian, and Jingren Zhou. 2020. Maximum biclique search at billion scale. *Proceedings of the VLDB Endowment* 13, 9 (2020), 1359–1372.
- [30] Shuai Ma, Renjun Hu, Luoshu Wang, Xuelian Lin, and Jinpeng Huai. 2017. Fast computation of dense temporal subgraphs. In *ICDE*. 361–372.
- [31] Hongchao Qin, Rong-Hua Li, Guoren Wang, Xin Huang, Ye Yuan, and Jeffrey Xu Yu. 2020. Mining stable communities in temporal networks by density-based clustering. *IEEE Transactions on Big Data* 8, 3 (2020), 671–684.
- [32] Hongchao Qin, Rong-Hua Li, Guoren Wang, Lu Qin, Yurong Cheng, and Ye Yuan. 2019. Mining periodic cliques in temporal networks. In *ICDE*. 1130–1141.
- [33] Hongchao Qin, Rong-Hua Li, Ye Yuan, Guoren Wang, Lu Qin, and Zhiwei Zhang. 2022. Mining Bursting Core in Large Temporal Graphs. *Proceedings of the VLDB Endowment* 15, 13 (2022), 3911–3923.
- [34] Ingmar Schäfer, Hanna Kaduszkiewicz, Hans-Otto Wagner, Gerhard Schön, Martin Scherer, and Hendrik Van Den Bussche. 2014. Reducing complexity: a visualisation of multimorbidity by combining disease clusters and triads. *BMC Public Health* 14 (2014), 1–14.
- [35] Renjie Sun, Chen Chen, Xiaoyang Wang, Wenjie Zhang, Ying Zhang, and Xuemin Lin. 2023. Efficient maximum signed biclique identification. In *ICDE*. 1313–1325.
- [36] Renjie Sun, Yanping Wu, Chen Chen, Xiaoyang Wang, Wenjie Zhang, and Xuemin Lin. 2022. Maximal balanced signed biclique enumeration in signed bipartite graphs. In *ICDE*. 1887–1899.
- [37] Renjie Sun, Yanping Wu, Xiaoyang Wang, Chen Chen, Wenjie Zhang, and Xuemin Lin. 2024. Efficient Balanced Signed Biclique Search in Signed Bipartite Graphs. *IEEE Trans. Knowl. Data Eng.* 36, 3 (2024), 1069–1083.
- [38] Yifu Tang, Jianxin Li, Nur Al Hasan Haldar, Ziyu Guan, Jia-Jie Xu, and Chengfei Liu. 2022. Reliable Community Search in Dynamic Networks. *Proceedings of the VLDB Endowment* 15, 11 (2022), 2826–2838.
- [39] Davide L Vetrano, Albert Roso-Llorach, Sergio Fernández, Marina Guisado-Clavero, Concepción Violán, Graziano Onder, Laura Fratiglioni, Amaia Calderón-Larrañaga, and Alessandra Marengoni. 2020. Twelve-year clinical trajectories of multimorbidity in a population of older adults. *Nature Communications* 11 (2020), 3223.
- [40] Kai Wang, Xuemin Lin, Lu Qin, Wenjie Zhang, and Ying Zhang. 2022. Towards efficient solutions of bitruss decomposition for large-scale bipartite graphs. *The VLDB Journal* 31, 2 (2022), 203–226.
- [41] Yiqi Wang, Long Yuan, Zi Chen, Wenjie Zhang, Xuemin Lin, and Qing Liu. 2023. Towards efficient shortest path counting on billion-scale graphs. In *ICDE*. 2579–2592.
- [42] Yanping Wu, Renjie Sun, Xiaoyang Wang, Dong Wen, Ying Zhang, Lu Qin, and Xuemin Lin. 2024. Efficient Maximal Frequent Group Enumeration in Temporal Bipartite Graphs. *arXiv preprint arXiv:2407.03954* (2024).
- [43] Guizhen Yang. 2004. The complexity of mining maximal frequent itemsets and maximal frequent patterns. In *SIGKDD*. 344–353.
- [44] Guizhen Yang. 2006. Computational aspects of mining maximal frequent patterns. *Theoretical Computer Science* 362, 1–3 (2006), 63–85.
- [45] Yixing Yang, Yixiang Fang, Maria E Orlowska, Wenjie Zhang, and Xuemin Lin. 2021. Efficient bi-triangle counting for large bipartite networks. *Proceedings of the VLDB Endowment* 14, 6 (2021), 984–996.
- [46] Yi Yang, Da Yan, Huanhuan Wu, James Cheng, Shuigeng Zhou, and John CS Lui. 2016. Diversified temporal subgraph pattern mining. In *SIGKDD*. 1965–1974.
- [47] Kai Yao, Lijun Chang, and Jeffrey Xu Yu. 2022. Identifying similar-bicliques in bipartite graphs. *Proceedings of the VLDB Endowment* 15, 11 (2022), 3085–3097.
- [48] Qianzhen Zhang, Deke Guo, Xiang Zhao, Long Yuan, and Lailong Luo. 2023. Discovering Frequency Bursting Patterns in Temporal Graphs. In *ICDE*. 599–611.
- [49] Yun Zhang, Charles A Phillips, Gary L Rogers, Erich J Baker, Elissa J Chesler, and Michael A Langston. 2014. On finding bicliques in bipartite graphs: a novel algorithm and its application to the integration of diverse biological data types. *BMC bioinformatics* 15, 1 (2014), 1–18.
- [50] Qi Zhao, Yingjuan Yang, Guofei Ren, Erxia Ge, and Chunlong Fan. 2019. Integrating bipartite network projection and KATZ measure to identify novel CircRNA-disease associations. *IEEE transactions on nanobioscience* 18, 4 (2019), 578–584.
- [51] Alexander Zhou, Yue Wang, and Lei Chen. 2021. Butterfly counting on uncertain bipartite graphs. *Proceedings of the VLDB Endowment* 15, 2 (2021), 211–223.
- [52] Rong Zhu, Zhaonian Zou, and Jianzhong Li. 2018. Diversified coherent core search on multi-layer graphs. In *ICDE*. 701–712.