

Finding a Concise, Precise, and Exhaustive Set of Near Bi-Cliques in Dynamic Graphs

Hyeonjeong Shin
KAIST

Seoul, South Korea
hyeonjeong1@kaist.ac.kr

Taehyung Kwon
KAIST

Seoul, South Korea
taehyung.kwon@kaist.ac.kr

Neil Shah
Snap Inc.

Seattle, Washington, USA
nshah@snap.com

Kijung Shin
KAIST

Seoul, South Korea
kijungs@kaist.ac.kr

ABSTRACT

A variety of tasks on dynamic graphs, including anomaly detection, community detection, compression, and graph understanding, have been formulated as problems of identifying constituent (near) bi-cliques (i.e., complete bipartite graphs). Even when we restrict our attention to maximal ones, there can be exponentially many near bi-cliques, and thus finding all of them is practically impossible for large graphs. Then, two questions naturally arise: (Q1) What is a “good” set of near bi-cliques? That is, given a set of near bi-cliques in the input dynamic graph, how should we evaluate its quality? (Q2) Given a large dynamic graph, how can we rapidly identify a high-quality set of near bi-cliques in it? Regarding Q1, we measure how concisely, precisely, and exhaustively a given set of near bi-cliques describes the input dynamic graph. We combine these three perspectives systematically on the Minimum Description Length principle. Regarding Q2, we propose CUTNPEEL, a fast search algorithm for a high-quality set of near bi-cliques. By adaptively re-partitioning the input graph, CUTNPEEL reduces the search space and at the same time improves the search quality. Our experiments using six real-world dynamic graphs demonstrate that CUTNPEEL is (a) *High-quality*: providing near bi-cliques of up to 51.2% **better quality** than its state-of-the-art competitors, (b) *Fast*: up to 68.8× **faster** than the next-best competitor, and (c) *Scalable*: scaling to graphs with 134 **million edges**. We also show successful applications of CUTNPEEL to graph compression and pattern discovery.

CCS CONCEPTS

• Information systems → Data mining; • Theory of computation → Dynamic graph algorithms.

KEYWORDS

Bi-clique; Dynamic Graph; Graph Compression; Pattern Discovery

ACM Reference Format:

Hyeonjeong Shin, Taehyung Kwon, Neil Shah, and Kijung Shin. 2022. Finding a Concise, Precise, and Exhaustive Set of Near Bi-Cliques in Dynamic Graphs. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM '22)*, February 21–25, 2022, Tempe, AZ, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3488560.3498390>

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

WSDM '22, February 21–25, 2022, Tempe, AZ, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9132-0/22/02...\$15.00

<https://doi.org/10.1145/3488560.3498390>

1 INTRODUCTION

Daily activities on the Web generate an enormous amount of data. Especially, many activities, such as e-mail communications, web surfing, online purchases, result in data in the form of graphs, such as e-mail networks, IP-IP communication networks, and user-business bipartite graphs. Most real-world graphs, including the aforementioned ones, evolve over time, and thus each of them is represented as a *dynamic graph*, i.e., a sequence of graphs over time.

A *bi-clique* is a complete bipartite graph. That is, a bi-clique consists of two disjoint sets of nodes where every node in a set is adjacent to every node in the other set. We use the term *near bi-clique* to denote a bipartite graph “close” to a bi-clique with few or no missing edges. We intentionally leave the concept near bi-clique flexible without rigid conditions (e.g., conditions in [7, 29, 41]).

Bi-cliques are a fundamental concept in graph theory, and a variety of tasks on static and dynamic graphs are formulated as problems of finding (near) bi-cliques in them. Examples include:

- **Anomaly Detection** [18, 38–40]: (Near) bi-cliques in real-world dynamic graphs signal anomalies, such as network intrusion [38], spam reviews [40], edit wars on Wikipedia [38], and retweet boosting on Weibo [18].
- **Community Detection** [3, 23, 26]: Relevant web pages often do not reference each other [23], and thus finding (near) cliques can be ineffective for detecting them. Instead, related web pages can be detected by finding (near) bi-cliques formed between them and users visiting them [26]. Near bi-cliques were also searched for detecting temporal communities in phone-call networks [3].
- **Lossless Graph Compression** [3, 22, 35]: A static or dynamic graph \mathcal{G} can be described concisely but losslessly by (a) (near) bi-cliques and (b) the difference between \mathcal{G} and the graph described by the (near) bi-cliques [3, 22, 35].
- **Phylogenetic Tree Construction** [34]: A phylogenetic tree shows the evolutionary interrelationships among multiple species that are thought to have a common ancestor. Identifying bi-cliques formed between genes and species containing the genes is helpful for accurate tree construction.

Additionally, (near) bi-cliques have been used for better understanding protein-protein interactions [7] and stock markets [41].

Due to the importance of (near) bi-cliques, a great number of search algorithms have been developed. Most of them aim to enumerate all maximal (near) bi-cliques that maximize certain criteria within a static or dynamic graph [2, 7, 10, 12, 20, 26, 27, 29, 33, 41]. However, as the number of such subgraphs can be exponential in the number of nodes, finding all is practically impossible for large graphs and often unnecessary with many insignificant ones (e.g., too small or highly overlapped ones). Thus, several algorithms




Figure 1: Strengths of CUTNPEEL. See Section 5 for details. (a) CUTNPEEL finds high-quality near bi-cliques (i.e., those with low cost in Eq. (5)) rapidly. (b) CUTNPEEL scales to a graph with 134 million edges, near-linearly with the number of edges. (c) CUTNPEEL detects a long-lasting near bi-clique between four relevant venues (Int. Test Conf., Asian Test Symp., IEEE VLSI Test Symp., and VLSI Design Conf.) and seven authors contributing to them. (d) CUTNPEEL achieves the best lossless compression.

aim to find one or a predefined number of (near) bi-cliques that maximize certain criteria in a static or dynamic graph [18, 38–40].

This paper focuses on finding a “good” set of near bi-cliques since finding all (maximal) near bi-cliques is infeasible for large graphs and finding one or few of them is insufficient for many applications. Then, two questions naturally arise:

- (1) What is a “good” set of near bi-cliques within a graph? How should we measure the quality of a set of near bi-cliques?
- (2) How can we rapidly detect a high-quality set of near bi-cliques, especially in large dynamic graphs?

Regarding the first question, we measure how *concisely, precisely, and exhaustively* a given set of near bi-cliques describes the input dynamic graph \mathcal{G} . That is, a high-quality set of near bi-cliques in \mathcal{G} consists of a *tractable number* of subgraphs *close to bi-cliques* that together *cover a large portion* of \mathcal{G} . We take these three perspectives into account simultaneously in a systematic way based on the Minimum Description Length principle [15, 17].

Regarding the second question, we propose CUTNPEEL (**C**utting a**N**d **P**eeling), a fast search algorithm for a high-quality set of near bi-cliques in a dynamic graph. As its name suggests, CUTNPEEL consists of the “cutting” step and the “peeling” step. Specifically, it first partitions the input dynamic graph to reduce the search space, and then it finds near bi-cliques within each partition one by one by a top-down search. Thanks to our adaptive re-partitioning scheme, surprisingly, the cutting step not only reduces the search space but also improves the quality of near bi-cliques by guiding the top-down search in the peeling step.

We demonstrate the effectiveness of CUTNPEEL, using six real-world dynamic graphs, though comparisons with four state-of-the-art algorithms [3, 35, 38, 40] for finding near bi-cliques in dynamic graphs. In summary, CUTNPEEL has the following advantages:

- **High Quality:** CUTNPEEL provides near bi-cliques of up to 51.2% better quality than the second best method (Figure 1a).
- **Speed:** CUTNPEEL is up to 68.8× faster than the competitors that is the second best in terms of quality (Figure 1a).
- **Scalability:** Empirically, CUTNPEEL scales near-linearly with the size of the input graph (Figure 1b).
- **Applicability:** Using CUTNPEEL, we achieve the best lossless compression and discover meaningful patterns (Figure 1c-d).

Reproducibility: The source code and datasets used in this paper can be found at <https://github.com/hyeonjeong1/cutnpeel>.

In Section 2, we introduce some basic concepts. In Section 3, we formally define the problem of finding near bi-cliques in dynamic graphs. In Section 4, we present our proposed algorithm CUTNPEEL.

Table 1: Table of frequently-used symbols

Symbol	Definition
$\mathcal{G} = (\mathcal{G}^{(t)})_{t \in \mathcal{T}}$	input dynamic graph
$\mathcal{S}, \mathcal{D}, \mathcal{T}$	sets of source nodes, destination nodes, and timestamps in \mathcal{G}
$\mathcal{I} = \mathcal{S} \cup \mathcal{D} \cup \mathcal{T}$	set of objects in \mathcal{G}
\mathcal{E}	set of edges in \mathcal{G}
$\tilde{\mathcal{I}} = \tilde{\mathcal{S}} \cup \tilde{\mathcal{D}} \cup \tilde{\mathcal{T}}$	subset of the object set \mathcal{I} where $\tilde{\mathcal{S}} \subseteq \mathcal{S}$, $\tilde{\mathcal{D}} \subseteq \mathcal{D}$, and $\tilde{\mathcal{T}} \subseteq \mathcal{T}$
$\mathcal{G}_{\tilde{\mathcal{I}}}$	near bi-clique composed of $\tilde{\mathcal{I}}$
$\mathcal{E}_{\tilde{\mathcal{I}}}$	set of edges in $\mathcal{G}_{\tilde{\mathcal{I}}}$
$\mathcal{B}_{\tilde{\mathcal{I}}}$	exact bi-clique composed of $\tilde{\mathcal{I}}$
$ \mathcal{B}_{\tilde{\mathcal{I}}} $	number of edges in $\mathcal{B}_{\tilde{\mathcal{I}}}$
$\hat{\mathcal{B}}$	set of near bi-cliques
$\phi(\hat{\mathcal{B}}, \mathcal{G})$	objective function to be minimized
$\mathcal{B}, \mathcal{R}, \mathcal{M}$	sets of exact bi-cliques, remaining edges, and missing edges
$\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$	approximate saving in the objective due to $\mathcal{G}_{\tilde{\mathcal{I}}}$
$\theta(t)$	threshold in each t -th iteration
α, T	decrement rate of thresholds and the number of iterations
$\cup_i \{\mathcal{G}_i\}$	set of partitions in \mathcal{G}'

In Section 5, we provide our experimental results. In Section 6, we discuss related studies. In Section 7, we conclude the paper.

2 BASIC CONCEPTS

In this section, we introduce some basic concepts used throughout the paper. We list some frequently-used notations in Table 1.

Dynamic graphs: A *dynamic graph* $\mathcal{G} = (\mathcal{G}^{(t)})_{t \in \mathcal{T}}$ is a sequence of graphs over timestamps \mathcal{T} . We denote \mathcal{G} at each time $t \in \mathcal{T}$ by $\mathcal{G}^{(t)} = (\mathcal{S}^{(t)} \cup \mathcal{D}^{(t)}, \mathcal{E}^{(t)})$, where $\mathcal{S}^{(t)}$ is the set of source nodes, $\mathcal{D}^{(t)}$ is the set of destination nodes, and $\mathcal{E}^{(t)}$ is the set of edges. We let $\mathcal{S} := \bigcup_{t \in \mathcal{T}} \mathcal{S}^{(t)}$ be the set of all source nodes and let $\mathcal{D} := \bigcup_{t \in \mathcal{T}} \mathcal{D}^{(t)}$ be the set of all destination nodes. For simplicity, we assume \mathcal{S}, \mathcal{D} and \mathcal{T} are disjoint. That is, unipartite graphs are treated as bipartite graphs (i.e., a source and the same node as a destination are treated as different nodes). We use $\mathcal{I} := \mathcal{S} \cup \mathcal{D} \cup \mathcal{T}$ to denote the set of *objects* in \mathcal{G} . We use $e = (s, d, t) \in \mathcal{S} \times \mathcal{D} \times \mathcal{T}$ to indicate the edge $(s, d) \in \mathcal{E}^{(t)}$ at time t , and we use $\mathcal{E} := \{(s, d, t) \in \mathcal{S} \times \mathcal{D} \times \mathcal{T} : (s, d) \in \mathcal{E}^{(t)}\}$ to indicate the set of all edges in \mathcal{G} . For any subset $\tilde{\mathcal{I}} \subseteq \mathcal{I}$ of objects, we use $\mathcal{G}_{\tilde{\mathcal{I}}}$ to denote the subgraph of \mathcal{G} induced by $\tilde{\mathcal{I}}$ and use $\mathcal{E}_{\tilde{\mathcal{I}}} := \{(s, d, t) \in \mathcal{E} : s \in \tilde{\mathcal{I}}, d \in \tilde{\mathcal{I}}, t \in \tilde{\mathcal{I}}\}$ to denote the edges in it.

Bi-cliques in dynamic graphs: A *bi-clique* is a complete bipartite graph. That is, a graph $\tilde{\mathcal{G}} = (\tilde{\mathcal{S}} \cup \tilde{\mathcal{D}}, \tilde{\mathcal{E}})$ with two disjoint node sets $\tilde{\mathcal{S}}$ and $\tilde{\mathcal{D}}$ is a bi-clique if every node in $\tilde{\mathcal{S}}$ is adjacent to every node $\tilde{\mathcal{D}}$, i.e., if $\tilde{\mathcal{E}} = (\tilde{\mathcal{S}} \times \tilde{\mathcal{D}})$. A *temporal bi-clique* is a bi-clique that appears once or repeatedly over time. For any set $\tilde{\mathcal{I}}$ of objects, we use $\mathcal{B}_{\tilde{\mathcal{I}}}$



Figure 2: The input dynamic graph $\mathcal{G} = (\mathcal{G}^{(t)})_{t \in \mathcal{T}}$ is described by exact bi-cliques \mathcal{B} , missing edges \mathcal{M} , and remaining edges \mathcal{R} .

to denote the temporal bi-clique formed by the objects, and we use $|\mathcal{B}_{\tilde{I}}|$ to denote the number of edges in it. That is, if $\tilde{I} = \tilde{\mathcal{S}} \cup \tilde{\mathcal{D}} \cup \tilde{\mathcal{T}}$ where $\tilde{\mathcal{S}} \subseteq \mathcal{S}$, $\tilde{\mathcal{D}} \subseteq \mathcal{D}$, and $\tilde{\mathcal{T}} \subseteq \mathcal{T}$, then $\mathcal{B}_{\tilde{I}} = (\tilde{\mathcal{G}}^{(t)})_{t \in \tilde{\mathcal{T}}}$ where $\tilde{\mathcal{G}}^{(t)} = (\tilde{\mathcal{S}} \cup \tilde{\mathcal{D}}, \tilde{\mathcal{S}} \times \tilde{\mathcal{D}})$ for all $t \in \tilde{\mathcal{T}}$, and $|\mathcal{B}_{\tilde{I}}| = |\tilde{\mathcal{S}}| \cdot |\tilde{\mathcal{D}}| \cdot |\tilde{\mathcal{T}}|$.

Near bi-cliques in dynamic graphs: We use the term *near bi-clique* to denote a bipartite graph “close” to a bi-clique with few or no of missing edges. Specifically, a subgraph $\tilde{\mathcal{G}} = (\tilde{\mathcal{S}} \cup \tilde{\mathcal{D}}, \tilde{\mathcal{E}})$ is a near bi-clique if it satisfies $\tilde{\mathcal{E}} \approx (\tilde{\mathcal{S}} \times \tilde{\mathcal{D}})$. Similarly, a dynamic graph $\tilde{\mathcal{G}} = (\tilde{\mathcal{G}}^{(t)})_{t \in \tilde{\mathcal{T}}}$ with objects $\tilde{I} = \tilde{\mathcal{S}} \cup \tilde{\mathcal{D}} \cup \tilde{\mathcal{T}}$ and edges $\tilde{\mathcal{E}} \subseteq \tilde{\mathcal{S}} \times \tilde{\mathcal{D}} \times \tilde{\mathcal{T}}$ is a *near temporal bi-clique* if it satisfies $\tilde{\mathcal{E}} \approx (\tilde{\mathcal{S}} \times \tilde{\mathcal{D}} \times \tilde{\mathcal{T}})$. Since this paper focuses on dynamic graphs, from now on, we will use the term **(near) bi-cliques** to refer to **(near) temporal bi-cliques** for ease of explanation.

3 PROBLEM FORMULATION

In this section, we present how we measure the quality of a set of near bi-cliques. Then, we define the problem of finding a high-quality set of near bi-cliques in a dynamic graph.

3.1 Quality of a Set of Near Bi-Cliques

In this subsection, we describe three sub-goals that any “good” set of near bi-cliques should pursue. We also define cost functions that we aim to minimize to achieve the sub-goals. The cost functions are commonly in bits so that we can directly compare and eventually balance them. To this end, given a dynamic graph \mathcal{G} and a set $\hat{\mathcal{B}}$ of near bi-cliques in it, we consider the following three sets:

- **Exact bi-cliques \mathcal{B} :** the set of minimal exact bi-cliques that contain the near bi-cliques in $\hat{\mathcal{B}}$. That is, $\mathcal{B} := \bigcup_{\mathcal{B}_{\tilde{I}} \in \hat{\mathcal{B}}} \{\mathcal{B}_{\tilde{I}}\}$.
- **Missing edges \mathcal{M} :** the set of *missing edges*, which are included in exact bi-cliques in \mathcal{B} but not in \mathcal{G} .
- **Remaining edges \mathcal{R} :** the set of *remaining edges* in \mathcal{G} that do not belong to any near bi-clique in $\hat{\mathcal{B}}$.

As illustrated in Figure 2, the bi-cliques in \mathcal{B} may contain missing edges, and the near bi-cliques in $\hat{\mathcal{B}}$ are obtained by simply removing the missing edges from each bi-clique in \mathcal{B} .

Sub-goal 1. (Preciseness): The members of a “good” set of near bi-cliques should actually be close to exact bi-cliques. That is, the number of missing edges \mathcal{M} contained in the bi-cliques in \mathcal{B} should be minimized. To pursue preciseness, we aim to minimize Eq. (1), which corresponds to the number of bits to encode \mathcal{M} .

$$\phi_P(\mathcal{M}) := |\mathcal{M}| \cdot \mathcal{L}_e, \quad (1)$$

where $\mathcal{L}_e := \log_2 |\mathcal{S}| + \log_2 |\mathcal{D}| + \log_2 |\mathcal{T}|$ corresponds to the number of bits to encode one missing edge. We assume that the coordinate list (COO) format is used, and thus we list s, d and t to encode an edge $e = (s, d, t)$. Additionally, we assume that $\log_2 |\mathcal{S}|$ bits are required to encode a member of a set \mathcal{S} .

Sub-goal 2. (Exhaustiveness): The members of a “good” set of near bi-cliques should cover a large portion of \mathcal{G} . Equivalently, the number of remaining edges \mathcal{R} uncovered by any near bi-cliques in $\hat{\mathcal{B}}$ should be minimized. To achieve exhaustiveness, we aim to minimize Eq. (2), i.e., the number of bits to encode \mathcal{R} .

$$\phi_E(\mathcal{R}) := |\mathcal{R}| \cdot \mathcal{L}_e. \quad (2)$$

Sub-goal 3. (Conciseness): Note that Eq. (1) and Eq. (2) can be trivially minimized by enumerating all edges, which are 1×1 bi-cliques, in \mathcal{G} . However, larger bi-cliques are more likely to be useful in applications mentioned in Section 1, so larger ones should be preferred over trivial ones in a “good” set of near bi-cliques. To design a cost function in bits, we note that larger bi-cliques can be used to encode the edges in them more concisely with fewer bits. Specifically, all edges in a bi-clique $\mathcal{B}_{\tilde{I}}$ where $\tilde{I} = \tilde{\mathcal{S}} \cup \tilde{\mathcal{D}} \cup \tilde{\mathcal{T}}$ can be encoded together by listing the objects that form it, requiring only $O(\tilde{I}) = O(|\tilde{\mathcal{S}}| + |\tilde{\mathcal{D}}| + |\tilde{\mathcal{T}}|)$ bits, instead of $O(|\tilde{\mathcal{S}}| \cdot |\tilde{\mathcal{D}}| \cdot |\tilde{\mathcal{T}}|)$ bits for encoding each edge separately. This saving in bits increases as bi-cliques become larger, as formalized in Lemma 1.

Lemma 1. If a bi-clique $\mathcal{B}_{\tilde{I}}$ is strictly bigger¹ than a bi-clique $\mathcal{B}_{\tilde{I}'}$, then the number of bits per edge is smaller in $\mathcal{B}_{\tilde{I}}$ than in $\mathcal{B}_{\tilde{I}'}$, i.e.,

$$\mathcal{L}_I(\mathcal{B}_{\tilde{I}}) / |\mathcal{B}_{\tilde{I}}| < \mathcal{L}_I(\mathcal{B}_{\tilde{I}'}) / |\mathcal{B}_{\tilde{I}'}|. \quad (3)$$

PROOF. See Appendix A [37]. \square

Based on this observation, we aim to minimize Eq. (4), which is the number of bits to encode the edges contained in the bi-cliques in \mathcal{B} , to pursue conciseness.

$$\phi_C(\mathcal{B}) := \sum_{\mathcal{B}_{\tilde{I}} \in \mathcal{B}} \mathcal{L}_I(\mathcal{B}_{\tilde{I}}), \quad (4)$$

where $\mathcal{L}_I(\mathcal{B}_{\tilde{I}}) := (|\tilde{I}| + 1) \cdot \log_2 |\tilde{I}|$ is the number of bits for encoding all edges in a bi-clique $\mathcal{B}_{\tilde{I}}$ together.² Note that in Eq. (4), edges that belong to multiple bi-cliques are encoded multiple times, increasing the cost. Thus, Eq. (4) penalizes highly-overlapped bi-cliques, aligning with our pursuit of conciseness.

Total cost and the MDL principle: Recall that the three cost functions (i.e., Eq. (1), Eq. (2), and Eq. (4)) are measured commonly in bits and thus directly comparable. Since \mathcal{B} , \mathcal{M} , and \mathcal{R} together exactly describe the input dynamic graph \mathcal{G} , as shown in Figure 2, minimizing their sum aligns with the Minimum Description Length (MDL) principle [17]. Thus, the *total cost* for balancing our three sub-goals (i.e., preciseness, exhaustiveness, and conciseness) is defined as:

$$\phi(\hat{\mathcal{B}}, \mathcal{G}) := \phi_P(\mathcal{M}) + \phi_E(\mathcal{R}) + \phi_C(\mathcal{B}), \quad (5)$$

where ϕ can be a function of $\hat{\mathcal{B}}$ and \mathcal{G} since \mathcal{B} , \mathcal{M} , and \mathcal{R} are directly obtained from $\hat{\mathcal{B}}$ and \mathcal{G} .

¹The number of object of each type in $\mathcal{B}_{\tilde{I}}$ is greater than or equal to that in $\mathcal{B}_{\tilde{I}'}$, and the number of object of a type in $\mathcal{B}_{\tilde{I}}$ is strictly greater than that in $\mathcal{B}_{\tilde{I}'}$.

²Additional $\log_2 |\tilde{I}|$ bits are used to encode the number of objects in $\mathcal{B}_{\tilde{I}}$.

3.2 Problem Definition

Using the quality measure in Eq. (5), we formalize our problem as:

Problem 1 (Finding a Concise, Precise, and Exhaustive Set of Near Bi-Cliques in a Dynamic Graph).

- (1) **Given:** a dynamic graph $\mathcal{G} = (\mathcal{G}^{(t)})_{t \in \mathcal{T}}$,
- (2) **Find:** near bi-cliques $\hat{\mathcal{B}}$
- (3) **To minimize:** $\phi(\hat{\mathcal{B}}, \mathcal{G})$.

4 PROPOSED ALGORITHM

In this section, we propose CUTNPEEL, a fast algorithm for finding a high-quality set of near bi-cliques in dynamic graphs. We first present its preliminary version PEEL. Then, we present CUTNPEEL.

4.1 PEEL: Preliminary Algorithm

In this subsection, we present PEEL, a preliminary version of our proposed algorithm CUTNPEEL. We provide an overview of PEEL and then describe how it finds a single near bi-clique in detail. After that, we analyze its complexity and lastly discuss its limitations.

4.1.1 Overview (Algorithm 1). PEEL repeatedly finds near bi-cliques one by one using a top-down search called PEELONE (line 3). Specifically, PEELONE returns objects $\tilde{\mathcal{I}}$, and the near bi-clique $\mathcal{G}_{\tilde{\mathcal{I}}}$ formed by them in the original input graph \mathcal{G} is added to the output set $\hat{\mathcal{B}}$ (line 6). Whenever PEEL finds a near bi-clique, the edges in it are removed from the input graph to prevent PEEL from finding the same near bi-clique (line 7), and we use \mathcal{G}' to denote the dynamic graph with the remaining edges. PEEL stops searching when $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$ in Eq. (6) is zero or negative for found objects $\tilde{\mathcal{I}}$ (lines 4-5) and returns $\hat{\mathcal{B}}$ as the final output (line 8).

$$\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}') := \phi(\emptyset, \mathcal{G}') - \phi(\{\mathcal{G}'_{\tilde{\mathcal{I}}}\}, \mathcal{G}'), \quad (6)$$

where $\phi(\emptyset, \mathcal{G}')$ is the total cost $\phi(\hat{\mathcal{B}}, \mathcal{G})$ when $\hat{\mathcal{B}} = \emptyset$ and $\mathcal{G} = \mathcal{G}'$; and $\phi(\mathcal{G}'_{\tilde{\mathcal{I}}}, \mathcal{G}')$ is that when $\hat{\mathcal{B}} = \{\mathcal{G}'_{\tilde{\mathcal{I}}}\}$ and $\mathcal{G} = \mathcal{G}'$. We use $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$ in Eq. (6) to approximate $\phi(\hat{\mathcal{B}}, \mathcal{G}) - \phi(\hat{\mathcal{B}} \cup \{\mathcal{G}'_{\tilde{\mathcal{I}}}\}, \mathcal{G})$, i.e., the amount of saving in the total cost due to a new near bi-clique $\mathcal{G}'_{\tilde{\mathcal{I}}}$. $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$ can be computed in $O(1)$ time, as shown in Lemma 2, while providing a preciseness guarantee in Theorem 1.

Lemma 2. Given the numbers of objects of each type and edges in \mathcal{G}' and $\mathcal{G}'_{\tilde{\mathcal{I}}}$, $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$ can be computed in $O(1)$ time.

PROOF. If $\tilde{\mathcal{I}} = \tilde{\mathcal{S}} \cup \tilde{\mathcal{D}} \cup \tilde{\mathcal{T}}$ where $\tilde{\mathcal{S}} \subseteq \mathcal{S}$, $\tilde{\mathcal{D}} \subseteq \mathcal{D}$, and $\tilde{\mathcal{T}} \subseteq \mathcal{T}$, then Eqs. (1)-(6) imply Eq. (7), where every term is given.

$$\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}') = (2 \cdot |\mathcal{E}'_{\tilde{\mathcal{I}}}| - |\tilde{\mathcal{S}}| \cdot |\tilde{\mathcal{D}}| \cdot |\tilde{\mathcal{T}}|) \cdot \mathcal{L}_e - (|\tilde{\mathcal{I}}| + 1) \cdot \log_2 |\mathcal{I}|. \quad (7)$$

Theorem 1 (Preciseness of PEEL). The density of every near bi-clique $\mathcal{G}_{\tilde{\mathcal{I}}} \in \hat{\mathcal{B}}$ obtained by Algorithm 1 is greater than 0.5, i.e.,

$$|\mathcal{E}_{\tilde{\mathcal{I}}}| / |\mathcal{B}_{\tilde{\mathcal{I}}}| > 0.5. \quad (8)$$

PROOF. By lines 4 and 5 of Algorithm 1, Eq. (9) holds.

$$\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}') > 0, \quad \forall \mathcal{G}_{\tilde{\mathcal{I}}} \in \hat{\mathcal{B}}. \quad (9)$$

Eq. (7), Eq. (9), and $|\mathcal{E}_{\tilde{\mathcal{I}}}| \geq |\mathcal{E}'_{\tilde{\mathcal{I}}}|$ imply Eq. (8). \square

Algorithm 1: Overview of PEEL

Input: dynamic graph \mathcal{G} with objects \mathcal{I} and edges \mathcal{E}
Output: near bi-cliques $\hat{\mathcal{B}}$

```

1 initialize  $\mathcal{G}'$  with objects  $\mathcal{I}$  and edges  $\mathcal{E}' = \mathcal{E}$ 
2 while true do
3    $\tilde{\mathcal{I}} \leftarrow \text{PEELONE}(\mathcal{G}')$  ▷ Algorithm 2
4   if  $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}') \leq 0$  then
5     break
6    $\hat{\mathcal{B}} \leftarrow \hat{\mathcal{B}} \cup \{\mathcal{G}_{\tilde{\mathcal{I}}}\}$ 
7    $\mathcal{E}' \leftarrow \mathcal{E}' - \mathcal{E}'_{\tilde{\mathcal{I}}}$ 
8 return  $\hat{\mathcal{B}}$ 
```

Algorithm 2: PEELONE: top-down search of a bi-clique

Input: dynamic graph \mathcal{G}' with objects \mathcal{I} and edges \mathcal{E}'
Output: objects $\tilde{\mathcal{I}}_{\max}$ that form a near bi-clique

```

1  $\tilde{\mathcal{I}} \leftarrow \mathcal{I}$ ;  $s_{\max} \leftarrow -\infty$ ;  $t_{\max} \leftarrow 0$ 
2 for  $t = 1, \dots, |\mathcal{I}|$  do
3   if  $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}') > s_{\max}$  then
4      $s_{\max} \leftarrow \text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$ ;  $t_{\max} \leftarrow t$ 
5    $i_t \leftarrow i \in \tilde{\mathcal{I}}$  with minimum  $\rho(\tilde{\mathcal{I}}, i)$ 
6    $\tilde{\mathcal{I}} \leftarrow \tilde{\mathcal{I}} \setminus \{i_t\}$ 
7  $\tilde{\mathcal{I}}_{\max} \leftarrow \{i_{t_{\max}}, i_{t_{\max}+1}, \dots, i_{|\mathcal{I}|}\}$ 
8 return  $\tilde{\mathcal{I}}_{\max}$ 
```

4.1.2 Finding one near bi-clique (Algorithm 2). We describe PEELONE, which is a subroutine of PEEL for finding a single near bi-clique. Starting from the set \mathcal{I} of all objects in \mathcal{G}' (line 1), PEELONE repeatedly removes an object (line 6) until an empty set is left. When choosing the object to be removed, PEELONE chooses one with the sparsest connectivity for a near bi-clique to remain. Specifically, it chooses an object $i \in \mathcal{I}$ that minimizes $\rho(\tilde{\mathcal{I}}, i)$ in Eq. (10) (line 5).

$$\rho(\tilde{\mathcal{I}}, i) := \frac{|\mathcal{E}_{\tilde{\mathcal{I}}}| - |\mathcal{E}_{\tilde{\mathcal{I}} \setminus \{i\}}|}{|\mathcal{B}_{\tilde{\mathcal{I}}}| - |\mathcal{B}_{\tilde{\mathcal{I}} \setminus \{i\}}|}, \quad (10)$$

where $|\mathcal{B}_{\tilde{\mathcal{I}}}| - |\mathcal{B}_{\tilde{\mathcal{I}} \setminus \{i\}}|$ is the number of potential edges adjacent to i , and $|\mathcal{E}_{\tilde{\mathcal{I}}}| - |\mathcal{E}_{\tilde{\mathcal{I}} \setminus \{i\}}|$ is the number of existing edges among them. As the set $\tilde{\mathcal{I}}$ of remaining objects changes, CUTNPEEL tracks $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$ in Eq. (6) (lines 3-4). Lastly, as its final output, PEELONE returns $\tilde{\mathcal{I}}$ when $\text{Saving}(\tilde{\mathcal{I}}, \mathcal{G}')$ is maximized (lines 7-8).

4.1.3 Relation with [38]. PEEL is based on the top-down greedy search framework [38], which is originally designed for detecting dense tensors (i.e., multi-dimensional arrays). PEEL uses the novel selection functions in Eq. (6) and Eq. (10), instead of those suggested in [38],³ and this change significantly improves the quality of detected near bi-cliques, as shown empirically in Section 5.2 and Appendix D [37] (compare the relative cost of M-Zoom [38] in Figure 3 and that of PEEL in Figure 6a in [37]). Specifically, the selection functions suggested in [38] do not guarantee Eq. (8), and empirically, using them gives near bi-cliques that lack preciseness (see Figure 4 in Section 5.2).

³For example, $|\mathcal{E}'_{\tilde{\mathcal{I}}}|/|\tilde{\mathcal{I}}|$ and $|\mathcal{E}'_{\tilde{\mathcal{I}}}|/(|\mathcal{B}_{\tilde{\mathcal{I}}}|)^{1/3}$.

4.1.4 Complexity. We present the time and space complexities of PEEL in Theorems 2 and 3. For simplicity, we assume $|I| = O(|\mathcal{E}|)$.

Theorem 2 (Time Complexity of PEEL). The time complexity of Algorithm 1 is $O(|\mathcal{E}|^2 \cdot \log |I|)$.

PROOF. Eq. (7) and Eq. (9) imply $|\mathcal{E}'_{\tilde{I}}| \geq 1$ for all $\mathcal{G}_{\tilde{I}} \in \hat{\mathcal{B}}$. This and $\sum_{\mathcal{G}_{\tilde{I}} \in \hat{\mathcal{B}}} |\mathcal{E}'_{\tilde{I}}| \leq |\mathcal{E}|$ imply $|\hat{\mathcal{B}}| \leq |\mathcal{E}|$. That is, the number of detected near bi-cliques is at most $|\mathcal{E}|$. From Lemma 2 and the aforementioned connection with [38], it can be shown that finding one near bi-clique takes $O(|\mathcal{E}| \cdot \log |I|)$ time (see Theorem 1 of [38]). Thus, the total time complexity is $O(|\mathcal{E}|^2 \cdot \log |I|)$. \square

Theorem 3 (Space Complexity of PEEL). The space complexity of Algorithm 1 is $O(|\mathcal{E}| \cdot \log |I|)$.

PROOF. Eq. (7) and Eq. (9) imply Eq. (11).

$$|\mathcal{E}'_{\tilde{I}}| \cdot \mathcal{L}_e > (|\mathcal{B}_{\tilde{I}}| - |\mathcal{E}'_{\tilde{I}}|) \cdot \mathcal{L}_e + (|\tilde{I}| + 1) \cdot \log_2 |I|, \forall \mathcal{G}_{\tilde{I}} \in \hat{\mathcal{B}}, \quad (11)$$

where the right hand side upper bounds the number of bits required for storing $\mathcal{G}_{\tilde{I}}$. Eq.(11) and $\sum_{\mathcal{G}_{\tilde{I}} \in \hat{\mathcal{B}}} |\mathcal{E}'_{\tilde{I}}| \leq |\mathcal{E}|$ imply that storing the output $\hat{\mathcal{B}}$ requires $O(|\mathcal{E}| \cdot \mathcal{L}_e) = O(|\mathcal{E}| \cdot \log |I|)$ bits. Additionally, storing \mathcal{G} and \mathcal{G}' requires $O(|\mathcal{E}| \cdot \mathcal{L}_e) = O(|\mathcal{E}| \cdot \log |I|)$ bits, and storing $\{i_t\}_{t=1}^{|I|}$ in Algorithm 2 requires $O(|I| \cdot \log |I|)$ bits. Thus, the total space complexity is $O(|\mathcal{E}| \cdot \log |I|)$. \square

4.1.5 Limitations. Finding a large number of near bi-cliques using PEEL suffers from high computational cost since PEEL needs to access all remaining edges (i.e., \mathcal{E}') to detect each near bi-clique. Moreover, PEEL often fails to detect bi-cliques formed by objects whose connectivity outside the bi-cliques are sparse. This is because such objects are likely to be removed in the early stage of PEELONE, which prioritizes objects with dense connectivity.

4.2 CUTNPEEL: Proposed Algorithm

In this subsection, we present CUTNPEEL, our proposed algorithm for Problem 1. We first discuss the main ideas behind CUTNPEEL. Then, we present an outline of CUTNPEEL. After that, we describe its components in detail. Lastly, we compare it with PEEL.

4.2.1 Main ideas. In order to address the two aforementioned limitations of PEEL, CUTNPEEL first partitions the input dynamic graph and then finds near bi-cliques within each partition. Partitioning reduces the search space, and as a result, it reduces the computational cost for finding each near bi-clique. Partitioning is also helpful to detect bi-cliques formed by objects with sparse global connectivity (i.e., sparse connectivity outside the bi-cliques). This is because CUTNPEEL prioritizes objects based on their connectivity within partitions, instead of their global connectivity. By reducing the search space, partitioning may also have a harmful impact on search accuracy. In order to minimize such a harmful impact, CUTNPEEL uses adaptive re-partitioning while employing randomness.

4.2.2 Overview (Algorithm 3). Given an input dynamic graph \mathcal{G} , the maximum number of iteration T , and the threshold decrement rate α , CUTNPEEL returns near bi-cliques $\hat{\mathcal{B}}$ in \mathcal{G} . As in PEEL, CUTNPEEL starts a search from \mathcal{G} (line 1), and as it detects near bi-cliques, removes the edges belonging to the near bi-cliques from \mathcal{G}' (lines 15

Algorithm 3: Overview of CUTNPEEL

Input:

- (1) dynamic graph \mathcal{G} with objects $I = \mathcal{S} \cup \mathcal{D} \cup \mathcal{T}$ and edges \mathcal{E}
- (2) maximum number of iterations $T (\geq 1)$
- (3) threshold decrement rate $\alpha (\in (0, 1))$

Output: near bi-cliques $\hat{\mathcal{B}}$

```

1 initialize  $\mathcal{G}'$  with objects  $I$  and edges  $\mathcal{E}' = \mathcal{E}$ 
2  $\theta(1) \leftarrow \infty$ 
3 for  $t = 1, \dots, T$  do
4    $\theta(t+1) \leftarrow 0$ 
5   for each object set  $\mathcal{V} \in \{\mathcal{S}, \mathcal{D}, \mathcal{T}\}$  do
6      $\cup_i \{\mathcal{G}'_{(i)}\} \leftarrow \text{CUT}(\mathcal{G}', \mathcal{V})$  ▷ Algorithm 4
7     for each partition
8        $\mathcal{G}'_{(k)} = (I - \mathcal{V} \cup \mathcal{V}_{(k)}, \mathcal{E}'_{(k)}) \in \cup_i \{\mathcal{G}'_{(i)}\}$  do
9         while true do ▷ Algorithm 2
10           $\tilde{I} \leftarrow \text{PEELONE}(\mathcal{G}'_{(k)})$ 
11          if  $\text{Saving}(\tilde{I}, \mathcal{G}'_{(k)}) < \theta(t)$  then
12             $\theta(t+1) \leftarrow \max(\text{Saving}(\tilde{I}, \mathcal{G}'_{(k)}) \cdot \alpha, \theta(t+1))$ 
13            break the while loop
14          else
15             $\hat{\mathcal{B}} \leftarrow \hat{\mathcal{B}} \cup \{\mathcal{G}_{\tilde{I}}\}$ 
16             $\mathcal{E}'_{(k)} \leftarrow \mathcal{E}'_{(k)} - \mathcal{E}'_{\tilde{I}}$ 
17             $\mathcal{E}' \leftarrow \mathcal{E}' - \mathcal{E}'_{\tilde{I}}$ 
18   if  $\theta(t+1) = 0$  then
19     break
20 return  $\hat{\mathcal{B}}$ 

```

and 16), where \mathcal{G}' denotes the graph with remaining edges. At each iteration, CUTNPEEL partitions \mathcal{G}' into subgraphs based on an object set $\mathcal{V} \in \{\mathcal{S}, \mathcal{D}, \mathcal{T}\}$ using CUT (line 6), which is described in detail later. CUT employs randomness so that different partitions are obtained at different iterations. Within each partition $\mathcal{G}'_{(k)}$, CUTNPEEL repeatedly detects near bi-cliques using PEELONE (line 9) as in PEEL. Since the current partitions may not be ideal, PEELONE stops finding near bi-cliques within a partition if $\text{Saving}(\tilde{I}, \mathcal{G}'_{(k)})$ in Eq. (6) is less than a threshold $\theta(t)$ for a found near bi-clique $\mathcal{G}_{\tilde{I}}$ (lines 10-12). CUTNPEEL terminates (line 19) if $\theta(t)$ reaches 0 (line 17) or the number of iterations reaches T .

4.2.3 Adaptive thresholding. The threshold $\theta(t)$ decreases adaptively as iterations proceed. Specifically, the threshold $\theta(t+1)$ at the $(t+1)$ -th iteration is set to $\alpha \in (0, 1)$ of the maximum $\text{Saving}(\tilde{I}, \mathcal{G}'_{(k)})$ value, among those less than $\theta(t)$, at the t -th iteration. The threshold balances exploration and exploitation. Specifically, in early iterations, CUTNPEEL puts more emphasis on exploration (i.e., exploring near bi-cliques in partitions in later iterations) by accepting near bi-cliques under strict conditions (i.e., with large $\theta(t)$). In later iterations, it puts more emphasis on exploitation (i.e., finding near bi-cliques in current partitions) by accepting near bi-cliques under lenient conditions (i.e., with small $\theta(t)$).

4.2.4 Cut for partitioning (Algorithm 4). Below, we propose CUT, a dynamic-graph partitioning algorithm used by CUTNPEEL. When designing CUT, we pursue the following goals:

- **G1.** Speed: CUT should be fast since it is performed repeatedly.

Algorithm 4: CUT: partitioning \mathcal{G} into subgraphs

Input:

- (1) dynamic graph \mathcal{G}' with objects $\mathcal{I} = \mathcal{S} \cup \mathcal{D} \cup \mathcal{T}$ and edges \mathcal{E}'
- (2) given a set of objects $\mathcal{V} \in \{\mathcal{S}, \mathcal{D}, \mathcal{T}\}$

Output: disjoint subgraphs of graph, $\cup_i \{\mathcal{G}'_{(i)}\}$

```

1 Generate a uniform random 1-to-1 function  $h : \mathcal{I} \rightarrow \{1, \dots, |\mathcal{I}|\}$ 
2 for each object  $v \in \mathcal{V}$  do
3    $f(v) \leftarrow \infty$ 
4   for each edge  $e \in \mathcal{E}'_I - \mathcal{E}'_{I \setminus \{v\}}$  do
5      $f(v) \leftarrow \min(g_v(e), f(v))$  ▷ see Eq. (12) for  $g_v(e)$ 
6 partition  $\mathcal{V}$  into  $\cup_i \{\mathcal{V}_{(i)}\}$  based on  $f(\cdot)$  values
7 for each  $\mathcal{V}_{(k)} \in \cup_i \{\mathcal{V}_{(i)}\}$  do
8    $\mathcal{G}'_{(k)} \leftarrow (\mathcal{I} - \mathcal{V} \cup \mathcal{V}_{(k)}, \mathcal{E}'_{\mathcal{I} - \mathcal{V} \cup \mathcal{V}_{(k)}})$ 
9 return  $\cup_i \{\mathcal{G}'_{(i)}\}$ 

```

- **G2.** Locality: CUT should locate each near bi-clique inside a partition, to make it detectable, rather than splitting it into partitions.
- **G3.** Randomness: CUT should employ randomness so that more possibilities can be explored by repeatedly performing it.

CUT partitions a given dynamic graph \mathcal{G}' based on a given object set $\mathcal{V} \in \{\mathcal{S}, \mathcal{D}, \mathcal{T}\}$. It first generates a uniform random 1-to-1 function $h : \mathcal{I} \rightarrow \{1, \dots, |\mathcal{I}|\}$ for the set \mathcal{I} of all objects (line 1), bringing randomness for **G3**. Using the function h , we compute a function $f : \mathcal{V} \rightarrow \mathbb{Z}_{>0}$ (lines 2-5), based on which \mathcal{V} is partitioned (line 6). Specifically, for each object $v \in \mathcal{V}$, $f(v)$ is defined as:

$$f(v) := \min\{g_v(e)\}_{e \in \mathcal{E}'_I - \mathcal{E}'_{I \setminus \{v\}}},$$

where $\mathcal{E}'_I - \mathcal{E}'_{I \setminus \{v\}}$ is the set of edges adjacent to v and

$$g_v(e = (s, d, t)) := \begin{cases} h(d) \cdot |\mathcal{I}| + h(t) & \text{if } v = s \\ h(t) \cdot |\mathcal{I}| + h(s) & \text{if } v = d \\ h(s) \cdot |\mathcal{I}| + h(d) & \text{if } v = t. \end{cases} \quad (12)$$

Our design of $f(\cdot)$ is inspired by min-hashing [6], and two objects in \mathcal{V} are more likely to have the same $f(\cdot)$ value as their connectivity to other objects are more similar. That is, CUT pursues **G2** by making objects with similar connectivity be likely to belong to the same partition. Lastly, edges in \mathcal{G}' are partitioned according to the partitions of their incident objects in \mathcal{V} (lines 7 and 8). Regarding **G1**, CUT takes linear time as formalized in Theorem 4.

Theorem 4 (Time Complexity of CUT). The time complexity of Algorithm 4 is $O(|\mathcal{I}| + |\mathcal{E}'|)$.

PROOF. Generating $h(i)$ for all $i \in \mathcal{I}$ takes $O(|\mathcal{I}|)$ time. Computing $f(v)$ for all $v \in \mathcal{V}$ takes $O(|\mathcal{V}| + |\mathcal{E}'|) = O(|\mathcal{I}| + |\mathcal{E}'|)$ time. Partitioning \mathcal{V} takes $O(|\mathcal{V}|) = O(|\mathcal{I}|)$ time, and partitioning \mathcal{E} takes $O(|\mathcal{E}'|)$ time. Hence, it takes $O(|\mathcal{I}| + |\mathcal{E}'|)$ time in total. \square

Comparison with PEEL: Below, we assume $|\mathcal{I}| = O(|\mathcal{E}|)$ for ease of explanation. For each near bi-clique, CUTNPEEL needs to access only the remaining edges within a partition (e.g., $\mathcal{E}'_{(k)}$), while PEEL needs to access all remaining edges (i.e., \mathcal{E}'). However, a partition can be as large as the entire dynamic graph. Thus, the worst-case time complexity of CUTNPEEL is $O(|\mathcal{E}|^2 \cdot \log |\mathcal{I}|)$, as in PEEL, if we assume $T = O(|\mathcal{E}| \cdot \log |\mathcal{I}|)$ so that performing CUT T times takes $O(T(|\mathcal{I}| + |\mathcal{E}'|)) = O(|\mathcal{E}|^2 \cdot \log |\mathcal{I}|)$ time. Compared to PEEL, CUTNPEEL additionally stores $h(i)$ for all $i \in \mathcal{I}$ and $f(v)$ for all $v \in$

Table 2: Summary of the six real-world dynamic graphs.

Name	Description	# of Objects	# of Edges
Enron [36]	sender / receiver / time [week]	140 / 144 / 128	11, 568
Darpa [25]	Src IP / Dst IP / time [day]	9, 484 / 23, 398 / 57	140, 069
DDoS [44]	Src IP / Dst IP / time [second]	9, 312 / 9, 326 / 3, 954	22, 844, 324
DBLP [11]	author / venue / time [year]	418, 236 / 3, 566 / 49	1, 325, 416
Yelp [47]	user / business / time [month]	552, 339 / 77, 079 / 134	2, 214, 201
Weeplaces [46]	user / place / time [month]	15, 793 / 971, 308 / 92	3, 970, 922

\mathcal{V} , which take only $O(|\mathcal{I}|) = O(|\mathcal{E}|)$ space. Regarding preciseness, Eq. (8) still holds for every near bi-clique $\mathcal{G}_{\hat{\mathcal{I}}} \in \hat{\mathcal{B}}$ from CUTNPEEL.

Empirically, however, CUTNPEEL scales near linearly with the size of the input dynamic graph, as shown experimentally in Section 5.3. Moreover, as shown in Appendix D [37], CUTNPEEL detects better near bi-cliques faster than PEEL.

5 EXPERIMENTS

In this section, we review our experiments to answer Q1-Q3.

- **Q1. Search Quality & Speed:** How rapid is CUTNPEEL? Are identified near bi-cliques precise, exhaustive, and concise?
- **Q2. Scalability:** How does the running time of CUTNPEEL scale with respect to the size of the input graph?
- **Q3. Application 1 - Compression:** How concisely can we represent dynamic graphs using the outputs of CUTNPEEL?

5.1 Experiment Specifications

Machines: We used a machine with a 3.8GHz AMD Ryzen 3900X CPU and 128GB RAM for the scalability test in Section 5.3 and a machine with 3.7GHz i9-10900K CPU and 64GB RAM for the others.

Datasets: We used six real-world dynamic graph datasets that are briefly described in Table 2.

Competitors and parameters: We implemented CUTNPEEL in Java and set $T = 80$ and $\alpha = 0.8$ unless otherwise stated. We compared CUTNPEEL with four competitors: TIMECRUNCH [35], COM2 [3], M-ZOOM [38], and D-CUBE [40]. For the competitors, we used the official implementations provided by the authors, which are in Java (COM2, M-ZOOM, and D-CUBE) or MATLAB (TIMECRUNCH). For TIMECRUNCH, we used the step-wise heuristic, which performed best in [35], and excluded chain subgraphs, which cannot be considered as bi-cliques, from its outputs. For D-CUBE, we set $\theta = 1$, which performed best in [40]. For M-ZOOM, and D-CUBE, we used the density function that minimized the cost function $\phi(\hat{\mathcal{B}}, \mathcal{G})$ on each dataset. For all competitors, we report the results at the iteration when the cost function $\phi(\hat{\mathcal{B}}, \mathcal{G})$ was minimized.

5.2 Q1. Search Quality & Speed

In this subsection, we compare the considered algorithms in terms of speed and the quality of output near bi-cliques.

Evaluation metric: For the quality, we use the total cost $\phi(\hat{\mathcal{B}}, \mathcal{G})$ relative to the encoding cost of the input dynamic graph, i.e.,

$$\text{Relative Cost}(\mathcal{B}, \mathcal{G}) := \frac{\phi(\hat{\mathcal{B}}, \mathcal{G})}{|\mathcal{E}| \cdot \mathcal{L}_e}, \quad (13)$$

where the denominator is the number of bits to individually encode every edge in the input graph \mathcal{G} . Since the denominator is a constant for a given \mathcal{G} , Eq. (13) is proportional to the total cost $\phi(\hat{\mathcal{B}}, \mathcal{G})$. Thus,




Figure 3: CUTNPEEL is fast while providing high-quality near bi-cliques. o.o.t.: out of time (> 6 hours). R.C: relative cost in Eq. (13) (the lower, the better). CUTNPEEL yields near bi-cliques with up to 51.2% better quality, up to 68.8 \times faster, than the competitors with the second best quality. We report the means of five trials, and error bars indicate ± 1 standard deviation.

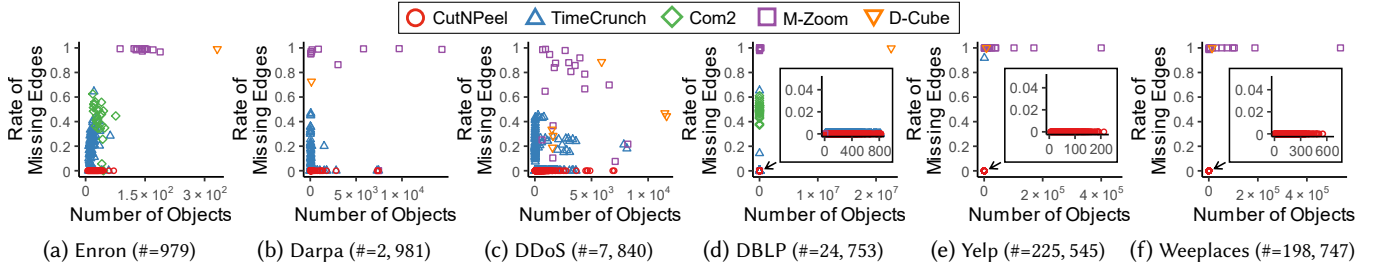


Figure 4: CUTNPEEL detects subgraphs close to bi-cliques. Each plot shows the size and preciseness of the near bi-cliques detected by the considered algorithms in each dataset. In parentheses, the number of near bi-cliques detected by CUTNPEEL in each dataset is reported. The near bi-cliques detected by CUTNPEEL tend to be precise with smaller ratios of missing edges. The number of near bi-cliques detected by CUTNPEEL is tractable and much smaller than the number of edges.

as explained in Section 3.1, the lower Eq. (13) is, the more concise, precise, and exhaustive detected near bi-cliques are.

Summary of comparison: As seen in Figure 3, CUTNPEEL gave near bi-cliques with the highest quality in all considered datasets. Specifically, the near bi-cliques detected by CUTNPEEL had up to 51.2% **better quality** in terms of the relative cost. Moreover, CUTNPEEL was up to 68.8 \times **faster** than TIMECRUNCH, which detected near bi-cliques with the second highest quality in most cases.

Detailed analysis: In Figure 5, we divide the relative cost in Eq. (13) into three portions related to preciseness (i.e., $\phi_p(M)$ in Eq. (1)), exhaustiveness (i.e., $\phi_e(R)$ in Eq. (2)), and conciseness (i.e., $\phi_c(B)$ in Eq. (4)), respectively. The total cost was lowest in CUTNPEEL, and the near bi-cliques detected by it were especially superior in exhaustiveness and preciseness. While $\phi_c(B)$ in CUTNPEEL was relatively high with many near bi-cliques, individual near bi-cliques detected by CUTNPEEL were precise with few missing edges, as shown in Figure 4, where we compare the size and preciseness of the near bi-cliques detected by the considered algorithms. Noticeably, the near bi-cliques detected by CUTNPEEL are precise with extremely small ratios of missing edges, and their count is tractable and especially much smaller than the number of edges in the input graph. M-ZOOM and D-CUBE tended to detect fewer larger subgraphs with higher ratios of missing edges than the other algorithms.

5.3 Q2. Scalability

In this subsection, we test the scalability of CUTNPEEL. To this end, we created Erdős-Rényi random graphs with various sizes and measured the running time of CUTNPEEL on them. The number of objects of each type was 0.1% of the number of edges, and the largest graph had $2^{27} = 134, 217, 728$ edges. As seen in Figure 1b in

Section 1, the running time of CUTNPEEL scaled near linearly with the number of edges and objects in the input dynamic graph.

5.4 Q3. Application: Compression

In this subsection, we show a successful application of CUTNPEEL to lossless graph compression. We consider COM2 and TIMECRUNCH as competitors, which are, to the best of our knowledge, the only lossless compression methods for the entire history (rather than the current snapshot [21]) of a dynamic graph.

The relative cost in Eq. (13) is the ratio of the number of bits for encoding the input graph losslessly using detected near bi-cliques and the number of bits for encoding each edge separately. Thus, it can naturally be interpreted as the compression rate. Alternatively, the input graph can be encoded using near bi-cliques as suggested in [3], and the compression rate can be computed accordingly. Thus, we measured it and Eq. (13) for all considered methods. Additionally, for TIMECRUNCH, we measured the compression rate in [35], which is based on an encoding method not applicable to the others.

As seen in Table 3, CUTNPEEL achieved the best compression in all considered datasets. Specifically, CUTNPEEL achieved up to 64.9% better compression rates than the second best method.

Extra experiments: In the supplementary document [37], we review extra experiments for Q4-Q6:

- **Q4. Ablation Study:** How much do adaptive thresholds and partitioning contribute to the performance of CUTNPEEL?
- **Q5. Parameter Analysis:** How do the decrement rate α and the iteration number T affect the performance of CUTNPEEL?
- **Q6. Application 2 - Pattern Discovery:** What interesting patterns does CUTNPEEL detect in real-world data?




Figure 5: CUTNPEEL provides near bi-cliques with the best quality, and they are especially superior in exhaustiveness and preciseness. See Appendix B [37] for full results.

6 RELATED WORK

In this section, we review studies of finding dense (near) (bi-)cliques in static and dynamic graphs.

6.1 Finding (near) (bi-)cliques in static graphs

Below, we focus on finding dense subgraphs in static graphs.

Enumerating exact (bi-)cliques: Finding (bi-)cliques (i.e., complete subgraphs), especially, enumerating all maximal (bi-)cliques, which are not a subset of any other clique, has been extensively studied [2, 9, 12, 14, 20, 26, 27, 33, 43]. The problem of enumerating all maximal (bi-)cliques was shown to be NP-hard [13, 31].

Enumerating near (bi-)cliques: Considerable attention has been paid to near (bi-)cliques with constraints on minimum degree [5], edge density [1, 45], and diameter [4, 32]. Notably, k -cores [5] for all k can be found in linear time [5], while enumerating (maximal) quasi-cliques [45], (maximal) k -clubs [32], or (maximal) k -plexes [28, 48] is NP-hard. Similarly, near bi-cliques include a subgraph where (a) every node in one node set is adjacent to at least $(1 - \epsilon)$ of the nodes in the other node set [29], (b) every node in both node sets is adjacent to at least $(1 - \epsilon)$ of the nodes in the counterpart node set [7], and (c) every node in both node sets is not adjacent to at most ϵ nodes in the counterpart node set [41]. Regardless of the definitions, enumerating (maximal) near bi-cliques is NP-hard since their number can be exponential in the number of nodes.

Finding a partial set of near (bi-)cliques: There also have been extensive studies on finding one or a predefined number of near (bi-)cliques. A representative example is to find a subgraph that maximizes average degree [8, 16, 19] or the ratio between the number of contained (bi-)cliques and the number of nodes [30, 42]. VoG [22] searches for a partial set of near (bi-)cliques and chain subgraphs by which the input graph is summarized.

6.2 Finding near bi-cliques in dynamic graphs

Below, we focus on finding dense subgraphs and their occurrences over time in dynamic graphs.

Finding a predefined number of near bi-cliques: Shin et al. [38, 40] and Jiang et al. [18] studied the problem of finding a given number of dense subtensors [18, 38, 40] in a tensor (i.e., a multi-dimensional array). For example, starting from the entire tensor, M-ZOOM [38] and D-CUBE [40] greedily remove entries so that one of the proposed density measures is minimized. Since a dynamic graph is naturally expressed as a 3-way tensor, the above algorithms can be used for identifying dense subgraphs in a dynamic graph.

Table 3: CUTNPEEL consistently achieved the best lossless compression. o.o.t.: out of time (> 6 hours). See Appendix C [37] for full results with standard deviations.

Metric*	Method	Compression Rates in % (the lower, the better)					
		Enron	Darpa	DDoS	DBLP	Yelp	Weeplaces
Eq. (13)	CUTNPEEL	52.5	32.3	13.1	53.7	77.1	55.9
	TIMECRUNCH	87.7	37.2	31.5	74.2	100.6	o.o.t.
	COM2	95.5	100	o.o.t.	99.9	o.o.t.	o.o.t.
[35]**	TIMECRUNCH	86.3	36.7	16.3	79.7	100.0	o.o.t.
[3]	CUTNPEEL	48.3	27.1	8.0	52.4	78.0	52.4
	TIMECRUNCH	79.3	34.6	22.8	63.9	100.2	o.o.t.
	COM2	58.3	202.8	o.o.t.	57.5	o.o.t.	o.o.t.

* Different metrics are based on different encoding methods.

** The encoding method used in [35] is not applicable to CUTNPEEL and COM2.

Summarizing dynamic graphs using near bi-cliques: COM2 [3] and TIMECRUNCH [35] aim to concisely describe the input dynamic graph using dense subgraphs, including near bi-cliques. COM2 uses rank-1 CP decomposition to obtain the score (i.e., corresponding value of the factor matrices) of each object and sorts the objects of each type based on the score. Objects are added one by one to a near bi-clique greedily until the description length does not decrease, and the above step is repeated for detecting multiple near bi-cliques. TIMECRUNCH aims to identify (near) (bi-)cliques and chain subgraphs. To this end, TIMECRUNCH decomposes the input graph at each timestamp into such subgraphs using SLASH-BURN [24], and “stitch” some of the found ones over timestamps based on the description length. Among these candidates, some are selected using multiple heuristics.

7 CONCLUSION

In this work, we consider the problem of finding a concise, precise, and exhaustive set of near bi-cliques in a dynamic graph. We formulate the problem as an optimization problem whose objective combines the three aspects (i.e., conciseness, preciseness, and exhaustiveness) in a systematic way based on the MDL principle. Our algorithmic contribution is to design CUTNPEEL for the problem. Compared to a widely-used top-down greedy search, CUTNPEEL reduces the search space and at the same time improves search accuracy, through a novel adaptive re-partitioning scheme. We summarize the strengths of CUTNPEEL as follows:

- **High Quality:** Compared to its best competitors, CUTNPEEL finds near bi-cliques of up to *51.2% better quality* (Figures 3-4).
- **Speed:** CUTNPEEL is up to *68.8× faster* than the competitor with the second best quality (Figure 3).
- **Scalability:** CUTNPEEL scales to graphs with up to *134 million edges*, near-linearly with the size of the input graph (Figure 1b).
- **Applicability:** Using CUTNPEEL, we achieve up to *64.9% better compression* than the best competitor (Table 3), and we spot interesting patterns (e.g., relevant conferences in Figure 1c).

Acknowledgements: This work was supported by Samsung Electronics Co., Ltd., National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. NRF-2020R1C1C1008296), and Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)).

REFERENCES

- [1] James Abello, Mauricio GC Resende, and Sandra Sudarsky. 2002. Massive quasi-clique detection. In *LATIN*. Springer.
- [2] 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.
- [3] Miguel Araujo, Spiros Papadimitriou, Stephan Günnemann, Christos Faloutsos, Prithwish Basu, Ananthram Swami, Evangelos E Papalexakis, and Danai Koutra. 2014. Com2: fast automatic discovery of temporal (‘comet’) communities. In *PAKDD*.
- [4] Balabhaskar Balasundaram, Sergiy Butenko, and Illya V Hicks. 2011. Clique relaxations in social network analysis: The maximum k-plex problem. *Operations Research* 59, 1 (2011), 133–142.
- [5] Vladimir Batagelj and Matjaz Zaversnik. 2003. An $O(m)$ algorithm for cores decomposition of networks. *arXiv preprint cs/0310049* (2003).
- [6] Andrei Z Broder, Moses Charikar, Alan M Frieze, and Michael Mitzenmacher. 2000. Min-wise independent permutations. *JCSS* 60, 3 (2000), 630–659.
- [7] Dongbo Bu, Yi Zhao, Lun Cai, Hong Xue, Xiaopeng Zhu, Hongchao Lu, Jingfen Zhang, Shiwei Sun, Lunjiang Ling, Nan Zhang, et al. 2003. Topological structure analysis of the protein–protein interaction network in budding yeast. *Nucleic Acids Research* 31, 9 (2003), 2443–2450.
- [8] Moses Charikar. 2000. Greedy approximation algorithms for finding dense components in a graph. In *APPROX*.
- [9] Norishige Chiba and Takao Nishizeki. 1985. Arboricity and subgraph listing algorithms. *SIAM Journal on computing* 14, 1 (1985), 210–223.
- [10] Apurba Das and Srikanta Tirthapura. 2018. Incremental maintenance of maximal bicliques in a dynamic bipartite graph. *TMSCS* 4, 3 (2018), 231–242.
- [11] The dblp computer science bibliography. 2021. *DBLP Data*. <https://dblp.uni-trier.de/db/>
- [12] Vânia MF Dias, Celina MH De Figueiredo, and Jayme L Swarcfiter. 2005. Generating bicliques of a graph in lexicographic order. *Theoretical Computer Science* 337, 1-3 (2005), 240–248.
- [13] David Eppstein. 1994. Arboricity and bipartite subgraph listing algorithms. *Inform. Process. Lett.* 51, 4 (1994), 207–211.
- [14] David Eppstein, Maarten Löffler, and Darren Strash. 2010. Listing all maximal cliques in sparse graphs in near-optimal time. In *ISAAC*.
- [15] Esther Galbrun. 2020. The minimum description length principle for pattern mining: A survey. *arXiv preprint arXiv:2007.14009* (2020).
- [16] Andrew V Goldberg. 1984. *Finding a maximum density subgraph*. University of California Berkeley.
- [17] Peter D Grünwald. 2007. *The minimum description length principle*. MIT press.
- [18] Meng Jiang, Alex Beutel, Peng Cui, Bryan Hooi, Shiqiang Yang, and Christos Faloutsos. 2015. A general suspiciousness metric for dense blocks in multimodal data. In *ICDM*.
- [19] Samir Khuller and Barna Saha. 2009. On finding dense subgraphs. In *ICALP*.
- [20] Kyle Kloster, Blair D Sullivan, and Andrew van der Poel. 2019. Mining maximal induced bicliques using odd cycle transversals. In *SDM*.
- [21] Jihoon Ko, Yunbum Kook, and Kijung Shin. 2020. Incremental Lossless Graph Summarization. In *KDD*.
- [22] Danai Koutra, U Kang, Jilles Vreeken, and Christos Faloutsos. 2014. VoG: Summarizing and understanding large graphs. In *SDM*.
- [23] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew Tomkins. 1999. Trawling the web for emerging cyber-communities. *Computer Networks* 31, 11-16 (1999), 1481–1493.
- [24] Yongsub Lim, U Kang, and Christos Faloutsos. 2014. Slashburn: Graph compression and mining beyond caveman communities. *TKDE* 26, 12 (2014), 3077–3089.
- [25] Richard P Lippmann, David J Fried, Isaac Graf, Joshua W Haines, Kristopher R Kendall, David McClung, Dan Weber, Seth E Webster, Dan Wyszogrod, Robert K Cunningham, et al. 2000. Evaluating intrusion detection systems: The 1998 DARPA off-line intrusion detection evaluation. In *DISCEX*.
- [26] Guimei Liu, Kelvin Sim, and Jinyan Li. 2006. Efficient mining of large maximal bicliques. In *DaWaK*.
- [27] Kazuhisa Makino and Takeaki Uno. 2004. New algorithms for enumerating all maximal cliques. In *SWAT*.
- [28] Benjamin McClosky and Illya V Hicks. 2012. Combinatorial algorithms for the maximum k-plex problem. *Journal of Combinatorial Optimization* 23, 1 (2012), 29–49.
- [29] Nina Mishra, Dana Ron, and Ram Swaminathan. 2004. A new conceptual clustering framework. *Machine Learning* 56, 1-3 (2004), 115–151.
- [30] Michael Mitzenmacher, Jakub Pachocki, Richard Peng, Charalampos Tsourakakis, and Shen Chen Xu. 2015. Scalable large near-clique detection in large-scale networks via sampling. In *KDD*.
- [31] John W Moon and Leo Moser. 1965. On cliques in graphs. *Israel journal of Mathematics* 3, 1 (1965), 23–28.
- [32] Esmaeel Moradi and Balabhaskar Balasundaram. 2018. Finding a maximum k-club using the k-clique formulation and canonical hypercube cuts. *Optimization Letters* 12, 8 (2018), 1947–1957.
- [33] Willard V Quine. 1955. A way to simplify truth functions. *The American mathematical monthly* 62, 9 (1955), 627–631.
- [34] Michael J Sanderson, Amy C Driskell, Richard H Ree, Oliver Eulenstein, and Sasha Langley. 2003. Obtaining maximal concatenated phylogenetic data sets from large sequence databases. *Molecular biology and evolution* 20, 7 (2003), 1036–1042.
- [35] Neil Shah, Danai Koutra, Tianmin Zou, Brian Gallagher, and Christos Faloutsos. 2015. Timecrunch: Interpretable dynamic graph summarization. In *KDD*.
- [36] Jitesh Shetty and Jafar Adibi. 2004. The Enron email dataset database schema and brief statistical report. *Information sciences institute technical report, University of Southern California* 4, 1 (2004), 120–128.
- [37] Hyeonjeong Shin, Taehyung Kwon, Neil Shah, and Kijung Shin. 2021. *Finding a Concise, Precise, and Exhaustive Set of Near Bi-Cliques in Dynamic Graphs (Supplementary Document)*. <https://github.com/hyeonjeong1/cutnpeel>
- [38] Kijung Shin, Bryan Hooi, and Christos Faloutsos. 2018. Fast, accurate, and flexible algorithms for dense subtensor mining. *TKDD* 12, 3 (2018), 1–30.
- [39] Kijung Shin, Bryan Hooi, Jisu Kim, and Christos Faloutsos. 2017. Densealrt: Incremental dense-subtensor detection in tensor streams. In *KDD*.
- [40] Kijung Shin, Bryan Hooi, Jisu Kim, and Christos Faloutsos. 2021. Detecting Group Anomalies in Tera-Scale Multi-Aspect Data via Dense-Subtensor Mining. *Frontiers in Big Data* 3 (2021), 58.
- [41] Kelvin Sim, Jinyan Li, Vivekanand Gopalkrishnan, and Guimei Liu. 2009. Mining maximal quasi-bicliques: Novel algorithm and applications in the stock market and protein networks. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 2, 4 (2009), 255–273.
- [42] Charalampos Tsourakakis. 2015. The k-clique densest subgraph problem. In *WWW*.
- [43] Shuji Tsukiyama, Mikio Ide, Hiromu Ariyoshi, and Isao Shirakawa. 1977. A new algorithm for generating all the maximal independent sets. *SIAM J. Comput.* 6, 3 (1977), 505–517.
- [44] The CAIDA UCSD. 2021. *DDoS Attack 2007*. https://www.caida.org/catalog/datasets/ddos-20070804_dataset
- [45] Takeaki Uno. 2010. An efficient algorithm for solving pseudo clique enumeration problem. *Algorithmica* 56, 1 (2010), 3–16.
- [46] Weeplaces. 2021. *Weeplaces Data*. <https://www.yongliu.org/datasets.html>
- [47] Yelp. 2021. *Yelp Data*. <https://www.kaggle.com/yelp-dataset/yelp-dataset>
- [48] Yi Zhou, Jingwei Xu, Zhenyu Guo, Mingyu Xiao, and Yan Jin. 2020. Enumerating maximal k-plexes with worst-case time guarantee. In *AAAI*.