

Fully Dynamic Four-Vertex Subgraph Counting

Kathrin Hanauer  

Faculty of Computer Science, University of Vienna, Austria

Monika Henzinger  

Faculty of Computer Science, University of Vienna, Austria

Qi Cheng Hua

Faculty of Computer Science, University of Vienna, Austria

Abstract

This paper presents a comprehensive study of algorithms for maintaining the number of all connected four-vertex subgraphs in a dynamic graph. Specifically, our algorithms maintain the number of paths of length three in deterministic amortized $\mathcal{O}(m^{\frac{1}{2}})$ update time, and any other connected four-vertex subgraph which is not a clique in deterministic amortized update time $\mathcal{O}(m^{\frac{2}{3}})$. Queries can be answered in constant time. We also study the query times for subgraphs containing an arbitrary edge that is supplied only with the query as well as the case where only subgraphs containing a vertex s that is fixed beforehand are considered. For length-3 paths, paws, 4-cycles, and diamonds our bounds match or are not far from (conditional) lower bounds: Based on the OMv conjecture we show that any dynamic algorithm that detects the existence of paws, diamonds, or 4-cycles or that counts length-3 paths takes update time $\Omega(m^{1/2-\delta})$.

Additionally, for 4-cliques and all connected induced subgraphs, we show a lower bound of $\Omega(m^{1-\delta})$ for any small constant $\delta > 0$ for the amortized update time, assuming the static combinatorial 4-clique conjecture holds. This shows that the $\mathcal{O}(m)$ algorithm by Eppstein et al. [9] for these subgraphs cannot be improved by a polynomial factor.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis; Theory of computation → Dynamic graph algorithms

Keywords and phrases Dynamic Graph Algorithms, Subgraph Counting, Motif Search

Digital Object Identifier 10.4230/LIPIcs.SAND.2022.18

Related Version *Full Version*: <https://arxiv.org/abs/2106.15524> [11]

Funding This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 101019564, “The Design of Modern Fully Dynamic Data Structures (MoDynStruct)”), as well as from the Austrian Science Fund (FWF) and netIDEE SCIENCE project P 33775-N.



Acknowledgements The authors want to thank Leonhard Paul Sidl for careful proofreading.

1 Introduction

Detecting or counting subgraphs is an important question in social network analysis, where dense subgraphs usually represent communities, as well as in telecommunication network surveillance, and computational biology. This can also be seen in a recent study by Sahu et al. [21]: finding and counting fixed subgraphs was the fourth most popular graph computation in practice, only superseded by finding connected components, computing shortest paths, and answering queries about the degree of neighbors. Furthermore the same study showed that the dynamic setting is important in practice as 65% of the graphs were dynamic. Thus, the goal of this paper is to advance the study of subgraph counting problems in dynamic graphs.



© Kathrin Hanauer, Monika Henzinger, and Qi Cheng Hua;
licensed under Creative Commons License CC-BY 4.0

1st Symposium on Algorithmic Foundations of Dynamic Networks (SAND 2022).

Editors: James Aspnes and Othon Michail; Article No. 18; pp. 18:1–18:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Algorithmic problems in dynamic graphs are usually modeled by the following data structure question. Given a potentially non-empty initial graph and a fixed subgraph pattern \mathcal{P} (such as a k -clique) maintain a data structure that allows the following updates to the current graph G :

- $\text{Insert}(u, v)$: Insert the edge $\{u, v\}$ into G .
- $\text{Delete}(u, v)$: Delete the edge $\{u, v\}$ from G .
- $\text{Query}()$: Return the number of subgraphs of pattern \mathcal{P} in G .

Given a subgraph pattern \mathcal{P} that is not a clique, there are two variants of this problem: One variant, called *induced subgraph counting*, counts a subgraph if it is exactly equivalent to \mathcal{P} and does *not* contain any additional edges. In the *non-induced* version, a subgraph is counted if it contains \mathcal{P} and potentially additional edges (but *not* additional vertices). Eppstein and Spiro [10] studied subgraph counting for all possible connected three-vertex patterns in both the induced and the non-induced variant and gave a dynamic algorithm with amortized update time $\mathcal{O}(h)$, where h is the h -index of G , i.e., the maximum number such that the graph contains h vertices of degree at least h . Note that h is $\mathcal{O}(\sqrt{m})$, where m always is the current number of edges in the graph.

There are six connected graphs on four vertices, which we refer to as *length-3 path* P_3 , *claw* $\text{K}_{1,3}$, *paw* P_3 with an additional edge, *4-cycle* C_4 , *diamond* $\text{K}_4 - \text{e}$, and *4-clique* K_4 . Eppstein et al. [9] extended the method of [10] to maintain counts of any (induced and non-induced) connected four-vertex subgraph in amortized time $\mathcal{O}(h^2) = \mathcal{O}(m)$ and $\mathcal{O}(mh^2)$ space. This paper contains a comprehensive study of the complexity of dynamically counting all possible connected four-vertex subgraphs. We present new improved dynamic algorithms and give the first conditional lower bounds.

Upper bounds. We show how to maintain the number of any connected four-vertex non-induced subgraph that is not a clique (such as a paw, a 4-cycle, or a diamond) in update time $\mathcal{O}(m^{2/3})$ and at most $\mathcal{O}(nm)$ space. For graphs with an h -index larger than $\mathcal{O}(m^{1/3})$, our algorithms are hence faster than the $\mathcal{O}(h^2)$ algorithm by Eppstein et al. [9]. Besides, our data structure can also be used to count all s -triangles, i.e., triangles that contain a fixed vertex s , in $\mathcal{O}(m^{1/2})$ update time, constant query time, and $\mathcal{O}(n)$ space, and likewise for s -length-3 paths. The update time is in $\mathcal{O}(m^{2/3})$ for s -4-cycles, s -paws, and s -diamonds, with $\mathcal{O}(n^2)$ space, and $\mathcal{O}(m)$ for s -4-cliques with constant space. For $\varepsilon \in [0, 1]$, our data structure supports queries on the number of triangles containing an arbitrary vertex or edge in $\mathcal{O}(\min(m^{2\varepsilon}, n^2))$ or $\mathcal{O}(\min(m^{1-\varepsilon}, n))$ worst-case time, respectively, with an update time of $\mathcal{O}(m^{\max(\varepsilon, 1-\varepsilon)})$ or $\mathcal{O}(m^\varepsilon)$, respectively, and $\mathcal{O}(n^2)$ space. We also show how to maintain length-3 paths in time $\mathcal{O}(m^{1/2})$, but this result was already stated in [9]. See Table 1 for an overview. All our algorithms are deterministic and the running time bounds are amortized unless stated otherwise.

Lower bounds. We also give the first conditional lower bounds for counting various four-vertex subgraphs based on two popular hypotheses: the Online Boolean Matrix-Vector Multiplication (OMv) conjecture [12] and the Combinatorial k -Clique hypothesis. In the OMv conjecture we are given a Boolean $n \times n$ matrix M that can be preprocessed. Then, an online sequence of vectors v_1, \dots, v_n is presented and the goal is to compute each Boolean product Mv_i (using conjunctions and disjunctions) before seeing the next vector v_{i+1} .

► **Conjecture 1 (OMv).** *For any constant $\delta > 0$, there is no $\mathcal{O}(n^{3-\delta})$ -time algorithm that solves OMv with error probability at most $1/3$ in the word-RAM model with $\mathcal{O}(\log n)$ bit words.*

■ **Table 1** Upper and conditional lower bounds on the time per update and query for counting different subgraphs, where $\delta > 0$ is an arbitrarily small constant and $h \in \mathcal{O}(\sqrt{m})$. Update times are amortized, query times are worst-case. Results in blue are new or improved.

* Read: For polynomial preprocessing time and $\mathcal{O}(\cdot)$ query time, the update time is $\Omega(\cdot)$.

† The previous space complexity for 3-cycles and length-3 paths was $\mathcal{O}(mh)$ with amortized update time $\mathcal{O}(h)$ [10] and $\mathcal{O}(m)$ with amortized update time $\mathcal{O}(m^{\frac{1}{2}})$ for 3-cycles [14]; for (other) 4-vertex subgraphs, it was $\mathcal{O}(mh^2)$ with amortized update time $\mathcal{O}(h^2)$ [9].

^a Thm 4, ^b Thm 5, ^c Cor. 19, ^d Thm 20, ^e Thm 21, ^f Cor. 25, ^g Thm 26, ^{\alpha} [10], ^{\beta} [9], ^{\zeta} [7].

Subgraph	Lower Bounds*		Update Time		Query Time		Space†
	Update	Query	ours	previous	all	$e \in E$	
<i>Non-induced subgraphs and s-subgraphs</i>							
connected, $n = 4$			$\mathcal{O}(m)^a$		$\mathcal{O}(1)$	$\mathcal{O}(1)^a$	$\mathcal{O}(1)/\mathcal{O}(m)^a$
claw \wedge	$\Omega(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)^{\alpha\beta}$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$
length-3 path \sphericalcap	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{1}{2}})^b$	$\mathcal{O}(h)^{\alpha\beta}$	$\mathcal{O}(1)$	$\mathcal{O}(m^{\frac{1}{2}})^b$	$\mathcal{O}(\min(n^2, m^{1.5}))^b$
paw \sphericalcap	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{2}{3}})^b$	$\mathcal{O}(h^2)^\beta$	$\mathcal{O}(1)$	$\mathcal{O}(m^{\frac{2}{3}})^b$	$\mathcal{O}(n^2)^b$
3-cycle \triangle	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{1}{2}})^c$	$\mathcal{O}(h)^\alpha$	$\mathcal{O}(1)$	$\mathcal{O}(m^{\frac{1}{2}})^c$	$\mathcal{O}(\min(n^2, m^{1.5}))^c$
4-cycle \diamond	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{2}{3}})^b$	$\mathcal{O}(h^2)^\beta$	$\mathcal{O}(1)$	$\mathcal{O}(m^{\frac{2}{3}})^b$	$\mathcal{O}(n^2)^b$
k -cycle, $k \geq 5$	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$					
diamond \diamond	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{2}{3}})^b$	$\mathcal{O}(h^2)^\beta$	$\mathcal{O}(1)$	$\mathcal{O}(m^{\frac{2}{3}})^b$	$\mathcal{O}(\min(nm, m^{\frac{5}{3}}))^b$
4-clique \boxtimes	$\Omega(m^{1-\delta})^e$	$\mathcal{O}(m^{2-\delta})^e$	$\mathcal{O}(m)^\zeta$	$\mathcal{O}(h^2)^\beta$	$\mathcal{O}(1)$	$\mathcal{O}(m)^\zeta$	$\mathcal{O}(1)$
<i>s-subgraphs</i>							
s-claw \wedge	$\Omega(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)^d$		$\mathcal{O}(1)^d$		$\mathcal{O}(1)^d$
s-length-3-path \sphericalcap	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{1}{2}})^d$		$\mathcal{O}(1)^d$		$\mathcal{O}(n)^d$
s-paw \sphericalcap	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{2}{3}})^d$		$\mathcal{O}(1)^d$		$\mathcal{O}(n^2)^d$
s-3-cycle \triangle	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{1}{2}})^d$		$\mathcal{O}(1)^d$		$\mathcal{O}(n)^d$
s-4-cycle \diamond			$\mathcal{O}(m^{\frac{2}{3}})^d$		$\mathcal{O}(1)^d$		$\mathcal{O}(n^2)^d$
s- k -cycle, $k \geq 5$, odd	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$					
s-diamond \diamond	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m^{\frac{2}{3}})^d$		$\mathcal{O}(1)^d$		$\mathcal{O}(n^2)^d$
s-4-clique \boxtimes	$\Omega(m^{\frac{1}{2}-\delta})^g$	$\mathcal{O}(m^{1-\delta})^g$	$\mathcal{O}(m)^d$		$\mathcal{O}(1)^d$		$\mathcal{O}(1)^d$
<i>Induced subgraphs</i>							
connected, $n = 4$	$\Omega(m^{1-\delta})^e$	$\mathcal{O}(m^{2-\delta})^e$	$\mathcal{O}(m)^f$	$\mathcal{O}(h^2)^\beta$	$\mathcal{O}(1)$		$\mathcal{O}(1)^f$

Based on the OMv conjecture we show that *detecting* (with probability at least $2/3$ in the word-RAM model with $\mathcal{O}(\log n)$ bit words) the existence of (non-induced) paws, diamonds, 4-cliques, or k -cycles for any $k \geq 3$ in a graph with edge insertions and deletions takes amortized update time $\Omega(m^{1/2-\delta})$ or query time $\Omega(m^{1-\delta})$ if only polynomial preprocessing time is allowed. This lower bound applies also to the worst-case update time of any insertions-only or deletions-only algorithm. Note that this lower bound does not only apply to *counting* the number of such subgraphs but already to *detecting* whether such a subgraph exists. Let s be a fixed vertex in the graph. The same lower bounds apply to algorithms that detect whether a diamond, 4-clique, or k -cycle with odd k containing s exists. Finally, we also show a lower bound for *counting* the number of length-3 paths and length-3 s -paths. We remark that the conditional lower bounds for (s -)3-cycles were already known before [12].

We also use the Combinatorial k -Clique hypothesis which is defined as follows and has become popular in recent years (e.g. [20, 1, 5, 4]).

► **Conjecture 2** (Combinatorial k -Clique). *For any constant $\delta > 0$, for an n -vertex graph there is no $\mathcal{O}(n^{k-\delta})$ time combinatorial algorithm for k -clique detection with error probability at most $1/3$ in the word-RAM model with $\mathcal{O}(\log n)$ bit words.*

Let $\delta > 0$ be a small constant. Based on the 4-clique conjecture we show that (with probability at least $2/3$ in the word-RAM model with $O(\log n)$ bit words) there does not exist a combinatorial algorithm that counts *any connected induced four-vertex subgraph* in a dynamic graph with amortized update time $O(n^{4-2\delta}/m)$, which is $O(m^{1-\delta})$, and query preprocessing time $O(n^{4-2\delta})$. This bound applies also to any insertions-only algorithm. The bound can be extended to any k -clique with $k > 4$ showing that the amortized update time is $\Omega(n^{k-2\delta}/m)$ with $\Omega(n^{k-2\delta})$ preprocessing and query time.

Technical contribution. For the upper bounds we extend and improve upon Eppstein et al. [9] both with respect to running time and space. The high-level idea is as follows: We partition the vertices into (few) *high-degree* and (many) *low-degree* vertices and then maintain for each vertex, vertex pair, or vertex triple certain information in a data structure such as the number of certain paths up to length 3 that contain low-degree vertices. When an edge $\{u, v\}$ is updated, four-vertex subgraphs that contain u, v , and two other low-degree vertices can be quickly counted using the information in the data structure. On the other side, subgraphs that contain a high-degree vertex in addition to u and v can often be counted “from scratch” after each update as there are few high-degree vertices. The more challenging case is the situation where relationships involving two or more high-degree vertices in the subgraph need to be checked or maintained. How to deal with this depends on the subgraph to count. For diamonds, e.g., this requires to keep certain information about triples of vertices.

For the conditional lower bounds based on the combinatorial k -clique conjecture we first directly deduce the lower bound for incremental 4-clique counting. Then we use the fact that (a) we have a lower bound for 4-cliques, (b) we developed algorithms with $\mathcal{O}(m^{2/3})$ update time for all non-induced subgraphs, and (c) there exist “counting formulas” that allow to compute the number of any induced subgraph based on the number of 4-cliques and the number of non-induced subgraphs. Thus, if the number of an induced subgraph pattern could be computed in $\mathcal{O}(m^{1-\delta})$ time per update for some small $\delta > 0$ then we could use the corresponding counting formula and our algorithms for non-induced subgraphs to dynamically maintain the number of 4-cliques, contradicting our dynamic lower bound for 4-cliques.

For the conditional lower bounds based on the OMv conjecture we construct for each subgraph pattern \mathcal{P} based on an 1-uMv instance (which is a variant of OMv) a suitable graph based on \mathcal{P} with $\mathcal{O}(n)$ vertices and $\mathcal{O}(n^2)$ edges such that detecting the existence of the (non-induced) version of \mathcal{P} in the graph equals finding the answer for the 1-uMv instance. Then the lower bound follows as in [12]. The challenge is to construct such a graph. We show how to do this for detecting non-induced (s)-paws, (s)-diamonds, (s)-4-cliques, and (s)- k -cycles for $k \geq 3$ and for counting non-induced length-3 (s)-paths.

Our paper gives in Section 2 the preliminaries, in Section 3 our new algorithms, and in Section 4 our lower bounds. Some proofs had to be omitted due to space and are given in the full version [11].

2 Preliminaries

Basic Definitions

We consider an undirected dynamic graph $G = (V, E)$ and use n to denote the number of vertices and m for the current number of edges. Two vertices $u \neq v$ are *adjacent* if there is an edge $e = \{u, v\} \in E$. In this case, u and v are *incident* to e . The *neighborhood* $N(v)$ of

a vertex v is defined as $\{u \mid \{u, v\} \in E\}$ and v 's *degree* is $\deg(v) = |N(v)|$. As a shorthand notation to exclude just one vertex, we use $N_{\bar{w}}(v) = N(v) \setminus \{w\}$ and $\deg_{\bar{w}}(v) = |N_{\bar{w}}(v)|$, i.e., $\deg_{\bar{w}}(v) = \deg(v) - 1$ if $w \in N(v)$ and $\deg_{\bar{w}}(v) = \deg(v)$ otherwise. A k -*path* (also length- k path) is a sequence of distinct edges $\langle \{v_0, v_1\}, \{v_1, v_2\}, \dots, \{v_{k-1}, v_k\} \rangle$ of length k , where $v_i \neq v_j$ for all $0 \leq i, j \leq k$. A k -*cycle* is a k -path where as an only exception the first vertex equals the last, i.e., $v_0 = v_k$. A 3-cycle is also called *triangle*. A *claw* is a graph consisting of a vertex x , called the *central* vertex, and three edges incident to it. A *paw* is a graph consisting of a triangle together with an additional edge attached to one of the vertices of the triangle. This vertex is called the *central* vertex of the paw and the additional edge the *arm*. A *diamond* is a 4-cycle with a *chord*, i.e., an additional edge connecting one of the two pairs of non-adjacent vertices, which creates two triangles sharing the chordal edge. A k -*clique* is the complete graph K_k on k vertices.

A graph $G' = (V', E')$ is a *subgraph* of G if $V' \subseteq V$ and $E' \subseteq E$. A subgraph $G' = (V', E')$ is said to be *induced* if $E' = \{\{u, v\} \in E \mid u, v \in V'\}$. The term *non-induced subgraph* or just *subgraph* without an adjective refers to all subgraphs, induced and not. For a static graph \mathcal{P} , also called *pattern*, we denote by $\mathbf{c}(G, \mathcal{P})$ the number of (non-induced) subgraphs of G that are isomorphic to \mathcal{P} , and by $\mathbf{c}_I(G, \mathcal{P})$ the number of induced subgraphs that are isomorphic to \mathcal{P} , in each case divided by the number of automorphisms of \mathcal{P} . For a vertex $s \in V$, we further denote by $\mathbf{c}(G, \mathcal{P}, s)$ the number of non-induced subgraphs of G that are isomorphic to \mathcal{P} and contain s . We denote by \mathbb{P}_k the set of connected subgraphs on k vertices. There are six connected graphs in \mathbb{P}_4 , which we refer to as *length-3 path* \curvearrowright , *claw* \wedge , *paw* \curvearrowleft , *4-cycle* \diamond , *diamond* \diamond , and *4-clique* \boxtimes .

Given a pattern \mathcal{P} , we study the problem of *maintaining* the number of occurrences of \mathcal{P} as a subgraph or induced subgraph of G , called the *non-induced subgraph count* $\mathbf{c}(G, \mathcal{P})$ or the *induced subgraph count* $\mathbf{c}_I(G, \mathcal{P})$, respectively, of \mathcal{P} in G , and the analogous problem of *maintaining* the count of non-induced subgraphs containing a specific predefined vertex $v \in V$ or edge $e \in E$, $\mathbf{c}(G, \mathcal{P}, v)$ or $\mathbf{c}(G, \mathcal{P}, e)$, respectively. Unless stated otherwise, *maintaining* a count implies that we can retrieve it by a query in constant time. We also study the closely related problem of *querying* the number of subgraphs containing a specific vertex or edge that is given only with the query.

Further Related Work

Detecting (or counting) subgraphs, also known as the *subgraph isomorphism* problem, generalizes the *clique* or *Hamiltonian cycle* problem and is hence \mathcal{NP} -hard. Nevertheless, it can be solved efficiently if the subgraphs to detect or count are restricted.

Static algorithms. Algorithms counting numbers of subgraphs and solving related problems have been studied extensively for static graphs. Alon, Yuster and Zwick [2] developed an algorithm to count the number of triangles and other circles (up to seven vertices) in a graph in time $\mathcal{O}(n^\omega)$, where n denotes the number of vertices and $\omega < 2.373$ [19] is the fast matrix multiplication exponent, i.e., the smallest value such that two $n \times n$ matrices can be multiplied in $\mathcal{O}(n^\omega)$ time. Kloks, Kratsch and Müller [16] showed how to compute the number of 4-cliques in time $\mathcal{O}(m^{(\omega+1)/2})$ (m denotes the number of edges) and the number of any other subgraph of size 4 in time $\mathcal{O}(n^\omega + m^{(\omega+1)/2})$. There is also a large body of work on parallel algorithms for counting subgraphs (see e.g. [18]) and to develop approximate algorithms in the streaming setting (see e.g. [6, 3]).

Dynamic algorithms. A more recent development is counting subgraph numbers for dynamic graphs. Kara et al. [14] provided an algorithm for counting triangles in amortized time $\mathcal{O}(\sqrt{m})$ per update and $\mathcal{O}(m)$ space, which can also enumerate them with constant time delay. Dhulipala et al. [7] extended it to a batch-dynamic parallel algorithm with $\mathcal{O}(\Delta\sqrt{\Delta+m})$ amortized work and $\mathcal{O}(\text{polylog}(\Delta+m))$ depth *w.h.p.* for a batch of Δ updates. Based on a static algorithm to enumerate cliques, they show how to obtain a dynamic algorithm for maintaining the number of k -cliques for a fixed $k > 3$ with expected $\mathcal{O}(\Delta(m+\Delta)\alpha^{k-4})$ work and $\mathcal{O}(\log^{-2} n)$ depth *w.h.p.* per update and $\mathcal{O}(m+\Delta)$ space, where $\alpha \in \mathcal{O}(\sqrt{m})$ is the arboricity of the graph. They also give a parallel fast matrix multiplication algorithm with $\mathcal{O}(\min(\Delta m^{(2k-1)\omega_p/(3\omega_p+3)}, (\Delta+m)^{2(k+1)\omega_p/(3\omega_p+3)}))$ amortized work and $\mathcal{O}(\log(\Delta+m))$ depth, with parallel matrix multiplication constant ω_p . Eppstein et al. [9] also count all three-vertex subgraphs in directed graphs in amortized time $\mathcal{O}(h)$. For specific graph classes, namely *bounded expansion* graphs, Dvorak and Tuma [8] gave a different algorithm for maintaining counts for arbitrary graph patterns \mathcal{P} of k vertices the number of induced subgraphs of pattern \mathcal{P} in amortized time $\mathcal{O}(\log^{(k^2-k)/2-1} n)$ and in amortized time $\mathcal{O}(n^\epsilon)$ for any constant $\epsilon > 0$ in *no-where dense* graphs.

In recent subsequent work [13], it was shown that counting 4-cycles is hard also in random graphs, i.e., with $\Omega(m^{1/2-\delta})$ update time or $\Omega(m^{1-\delta})$ query time.

3 New Counting Algorithms for Subgraphs on Four Vertices

Counting the number of claws [10] and 4-cliques is fairly straightforward [7]. The latter extends to 4-vertex subgraphs in general, where all work is either done during the updates or queries.

► **Observation 3** ([10]). *Let $G = (V, E)$ be a dynamic graph and \mathcal{P} be the claw \wedge . Then, $\mathbf{c}(G, \mathcal{P}) = \sum_{v \in V, \deg(v) \geq 3} \binom{\deg(v)}{3}$. The count can be initialized in $\mathcal{O}(n)$ time and maintained in constant time and space. We can query $\mathbf{c}(G, \mathcal{P}, \{u, v\})$ for an arbitrary edge $\{u, v\} \in E$ in constant time.*


► **Theorem 4.** *Let $G = (V, E)$ be a dynamic graph and $\mathcal{P} \in \mathbb{P}_4$. We can*

- (i) *maintain $\mathbf{c}(G, \mathcal{P})$ in worst-case $\mathcal{O}(m)$ update time and constant space.*
- (ii) *query $\mathbf{c}(G, \mathcal{P}, \{u, v\})$ for an arbitrary edge $\{u, v\} \in E$ in worst-case $\mathcal{O}(m)$ time, with constant update time and space.*
- (iii) *query $\mathbf{c}(G, \mathcal{P}, \{u, v\})$ for an arbitrary edge $\{u, v\} \in E$ in worst-case constant time, with $\mathcal{O}(m)$ worst-case update time and $\mathcal{O}(m)$ space.*

For connected 4-vertex subgraphs other than the claw and the clique, we can invest space to achieve speedups in running time. In the following, we present our data structure and show how to update it efficiently. We then use different parts of the data structure to count different subgraphs. Specifically, we prove the following result:

► **Theorem 5.** *Let G be a dynamic graph and $\mathcal{P} \in \mathbb{P}_4$. We can maintain $\mathbf{c}(G, \mathcal{P})$ in*

- (i) *amortized $\mathcal{O}(\sqrt{m})$ update time with $\mathcal{O}(\min(m^{1.5}, n^2))$ space and query $\mathbf{c}(G, \mathcal{P}, e)$ for an arbitrary edge $e \in E$ in worst-case $\mathcal{O}(\sqrt{m})$ time if \mathcal{P} is the length-3 path \curvearrowright ,*
- (ii) *amortized $\mathcal{O}(m^{2/3})$ update time with $\mathcal{O}(n^2)$ space and query $\mathbf{c}(G, \mathcal{P}, e)$ for an arbitrary edge $e \in E$ in worst-case $\mathcal{O}(m^{2/3})$ time if \mathcal{P} is the paw ∇ or the 4-cycle \diamond ,*
- (iii) *amortized $\mathcal{O}(m^{2/3})$ update time with $\mathcal{O}(\min(nm, m^{5/3}))$ space and query $\mathbf{c}(G, \mathcal{P}, e)$ for an arbitrary edge $e \in E$ in worst-case $\mathcal{O}(m^{2/3})$ time if \mathcal{P} is the diamond \diamond .*



■ **Figure 1** Subgraph structures of \mathcal{D}_ε . Small and filled vertices have low degree, large and empty vertices high degree, medium-sized and shaded vertices can have either high or low degree. Anchors are marked by arrows.

Our algorithm makes use of a standard technique in dynamic graph algorithms that partitions vertices into *high-degree* and *low-degree* vertices. We adapt it to our needs as follows: Let m_0 be the number of edges of G at construction or when recomputing from scratch and let $M = 2m_0$. A recomputation from scratch and re-initialization of the partition is triggered whenever the current number of edges $m < \lfloor \frac{M}{4} \rfloor$ or $m \geq M$. Initially and at each recomputation from scratch, a vertex $v \in V$ is classified as *high-degree* and added to partition \mathcal{H} if $\deg(v) \geq \theta$ and otherwise as *low-degree* and added to partition $\mathcal{L} := V \setminus \mathcal{H}$, for some threshold θ . As G evolves, a high-degree vertex v is reclassified as low and moved to \mathcal{L} only if $\deg(v) < \frac{1}{2}\theta$. Vice-versa, a low-degree vertex v is reclassified as high and moved to \mathcal{H} only if $\deg(v) \geq \frac{3}{2}\theta$. We call such a partition $(\mathcal{H}, \mathcal{L})$ a *dynamic vertex partition with threshold θ* . If $\theta = M^\varepsilon$ for some $\varepsilon \in [0, 1]$, we call the partition an ε -partition.

► **Theorem 6** (ε -Partition [15]). *Let $\varepsilon \in [0, 1]$ and consider an ε -partition $(\mathcal{H}, \mathcal{L})$ for a dynamic graph $G = (V, E)$. Then, $|\mathcal{H}| \in \mathcal{O}(m^{1-\varepsilon})$. The partition can be constructed in $\mathcal{O}(n)$ time and maintained in amortized constant time per update with amortized $\mathcal{O}(m^{-\varepsilon})$ changes to the partition per update and $\Omega(m)$ updates between two recomputations from scratch. The required space is $\mathcal{O}(n)$.*

Data Structure \mathcal{D}_ε

We assume that the algorithm can access the degree of a vertex v in constant time and, for each pair of vertices u, v determine in constant time whether $\{u, v\} \in E$. In addition, we maintain the following data structure \mathcal{D}_ε or a subset of it, if we are not only interested in counting some specific subgraphs on four vertices. All subgraph structures that are part of \mathcal{D}_ε are non-induced. See Figure 1 for visualizations.

- an ε -partition $(\mathcal{H}, \mathcal{L})$
- For each vertex $v \in V$: $\text{vLV}[v]$: the number of 2-paths $\langle \{v, x\}, \{x, y\} \rangle$ with $x \in \mathcal{L}$
- For each vertex $v \in \mathcal{H}$: $\text{t}[v]$: the number of 3-cycles $\langle \{v, x\}, \{x, y\}, \{y, v\} \rangle$
- For each distinct, unordered pair of vertices $u, v \in V$:
 - $\text{uLv}[u, v]$: the number of 2-paths $\langle \{u, x\}, \{x, v\} \rangle$ with $x \in \mathcal{L}$
 - $\text{uLLv}[u, v]$: the number of length-3 paths $\langle \{u, x\}, \{x, y\}, \{y, v\} \rangle$ with $x, y \in \mathcal{L}$
 - $\text{cLV}[u, v]$: the number of claws with a central vertex $x \in \mathcal{L}$, $u \neq x \neq v$
 - $\text{pLL}[u, v]$: the number of paws with a central vertex $x \in \mathcal{L}$, $u \neq x \neq v$, u or v at the other end of the arm, and fourth vertex $y \in \mathcal{L}$
- For each distinct, unordered pair of vertices $u, v \in \mathcal{H}$: $\text{uHv}[u, v]$: the number of length-2 paths $\langle \{v, x\}, \{x, v\} \rangle$ with $x \in \mathcal{H}$
- For each distinct, unordered triple of vertices $u, v, w \in V$: $\text{cL}[u, v, w]$: the number of claws with a fourth vertex $x \in \mathcal{L}$ at the center

18:8 Fully Dynamic Four-Vertex Subgraph Counting

For each auxiliary subgraph whose count is maintained by the data structure, we call vertices of the set that acts as key *anchors*, e.g., u and v are anchors for the 2-paths $\langle\{u, x\}, \{x, v\}\rangle$ with $x \in \mathcal{L}$, which are counted by $\mathbf{uLv}[u, v]$. We use hash tables with $\mathcal{O}(1)$ amortized access time and only store non-zero counts. Note that \mathbf{uLv} , \mathbf{uLLv} , \mathbf{pLL} , and \mathbf{cL} correspond to s_2 , s_3 , s_5 , and s_7 , respectively, in the algorithm by Eppstein et al. [9], whereas \mathbf{vLV} , \mathbf{t} , and \mathbf{cLV} are modifications of s_0 , s_1 , and s_4 , and \mathbf{uHv} has no equivalent at all. However, Eppstein et al. [9] use a different partitioning scheme, where there are at most $\mathcal{O}(h)$ vertices of degree $\Omega(h)$, whereas in our case, there are at most $\mathcal{O}(m^{1-\varepsilon})$ vertices of degree $\Omega(m^\varepsilon)$, which requires a different running time analysis also for the common auxiliary counts.

We generally assume that in case of an edge insertion, the auxiliary counts are updated immediately *before* the counts of interest, and in reverse order for an edge deletion. The update of the ε -partition can either happen first or last (but not in between). We also assume that we start with an empty graph and all counts are initialized to zero.

Maintaining the Data Structure \mathcal{D}_ε

Given a dynamic graph $G = (V, E)$ and $\varepsilon \in [0, 1]$, we show how the components of the data structure \mathcal{D}_ε can be updated after an edge insertion or deletion and if a vertex changes partition. We start with a helper lemma:

► **Lemma 7.** *Let \mathbf{aux} be an auxiliary subgraph count in \mathcal{D}_ε with worst-case update time $\mathcal{E}_{\mathbf{aux}}$ after an edge insertion or deletion, worst-case update time $\mathcal{V}_{\mathbf{aux}}$ after a vertex changes partition, and $\mathcal{S}_{\mathbf{aux}}$ space. Then, \mathcal{D}_ε with \mathbf{aux} can be maintained in amortized update time $\mathcal{O}(\mathcal{E}_{\mathbf{aux}} + \mathcal{V}_{\mathbf{aux}} \cdot m^{-\varepsilon})$ with $\mathcal{O}(n + \mathcal{S}_{\mathbf{aux}})$ space.*

Proof. By Theorem 6, the ε -partition can be maintained in $\mathcal{O}(n)$ space and such that there are $\Omega(m)$ updates between two recomputations of the partition from scratch. After each such complete repartitioning, we set \mathbf{aux} to zero and re-insert all edges one-by-one. The total recomputation time hence is $\mathcal{O}(m \cdot \mathcal{E}_{\mathbf{aux}})$ and amortization over $\Omega(m)$ edge updates results in an amortized edge update time of $\mathcal{O}(\mathcal{E}_{\mathbf{aux}})$. By Theorem 6, there are amortized $\mathcal{O}(m^{-\varepsilon})$ vertices changing partition per edge update, hence the claim follows. ◀

As the insertion and deletion operations are entirely symmetric and only differ in whether a certain amount is added or subtracted from the stored counts, we only give the details for edge insertions in the following. Similarly, we only consider the case that a vertex v changes from \mathcal{L} to \mathcal{H} ; the other case is symmetric. Note that if v is about to change partitions, $\deg(v) \in \Theta(m^\varepsilon)$.

► **Lemma 8.** *\mathcal{D}_ε with \mathbf{vLV} can be maintained in amortized $\mathcal{O}(m^\varepsilon)$ update time and $\mathcal{O}(n)$ space.*

Proof. Let $\{u, v\}$ be the newly inserted edge. If $u \in \mathcal{L}$ ($v \in \mathcal{L}$), increase $\mathbf{vLV}[w]$ by one for each $w \in N_{\bar{v}}(u)$ ($w \in N_{\bar{u}}(v)$) and increase $\mathbf{vLV}[v]$ by $\deg(u) - 1$ ($\mathbf{vLV}[u]$ by $\deg(v) - 1$). This takes $\mathcal{O}(m^\varepsilon)$ time.

If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} , this affects all length-2 paths where v is the central, low-degree vertex. For each neighbor $w \in N(v)$, decrease $\mathbf{vLV}[w]$ by $\deg(v) - 1$. The running time is $\mathcal{O}(\deg(v)) = \mathcal{O}(m^\varepsilon)$.

As each vertex may be adjacent to at least one low-degree vertex, the space requirement is $\mathcal{O}(n)$. By Lemma 7, \mathcal{D}_ε with \mathbf{vLV} can hence be maintained in $\mathcal{O}(m^\varepsilon)$ amortized update time and $\mathcal{O}(n)$ space. ◀

► **Lemma 9.** \mathcal{D}_ε with \mathbf{uLv} can be maintained in amortized $\mathcal{O}(m^\varepsilon)$ time per update and $\mathcal{O}(\min(m^{1+\varepsilon}, n^2))$ space.

Proof. Let $\{u, v\}$ be the newly inserted edge. If $u \in \mathcal{L}$ ($v \in \mathcal{L}$), increment $\mathbf{uLv}[v, w]$ ($\mathbf{uLv}[u, w]$) by one for each $w \in N_{\bar{v}}(u)$ ($w \in N_{\bar{u}}(v)$). This takes $\mathcal{O}(m^\varepsilon)$ time.

If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} , this affects all length-2 paths where v is the central, low-degree vertex. For each pair of distinct neighbors $x, y \in N(v)$, decrease $\mathbf{uLv}[x, y]$ by 1. The running time is $\mathcal{O}(\deg(v)^2) = \mathcal{O}(m^{2\varepsilon})$.

Each edge may be incident to at least one low-degree vertex v and form $\mathcal{O}(m^\varepsilon)$ length-2 paths with the other edges incident to v . The space requirement hence is $\mathcal{O}(\min(m^{1+\varepsilon}, n^2))$. By Lemma 7, \mathcal{D}_ε with \mathbf{uLv} can hence be maintained in $\mathcal{O}(m^\varepsilon + m^{2\varepsilon}m^{-\varepsilon}) = \mathcal{O}(m^\varepsilon)$ amortized update time and $\mathcal{O}(\min(m^{1+\varepsilon}, n^2))$ space. ◀

► **Lemma 10.** \mathcal{D}_ε with \mathbf{t} can be maintained in amortized $\mathcal{O}(m^{\max(1-\varepsilon, \varepsilon)})$ time per update and $\mathcal{O}(\min(m^{1+\varepsilon}, n^2))$ space.

Proof. Let $\{u, v\}$ be the newly inserted edge. For each $h \in \mathcal{H}$, increment $\mathbf{t}[h]$ by one if h is adjacent to both u and v . If $u \in \mathcal{H}$ ($v \in \mathcal{H}$): Increment $\mathbf{t}[u]$ ($\mathbf{t}[v]$) by one for each $h \in \mathcal{H}$ that is adjacent to both u and v , and increment $\mathbf{t}[u]$ ($\mathbf{t}[v]$) by $\mathbf{uLv}[u, v]$. This takes $\mathcal{O}(|\mathcal{H}|) = \mathcal{O}(m^{1-\varepsilon})$ time.

If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} , then for each pair of distinct neighbors $x, y \in N(v)$ such that $\{x, y\} \in E$, we increase $\mathbf{t}[v]$ by one. Otherwise, if v changes from \mathcal{H} to \mathcal{L} , set $\mathbf{t}[v] := 0$. The running time is $\mathcal{O}(\deg(v)^2) = \mathcal{O}(m^{2\varepsilon})$.

By Lemma 9, \mathcal{D}_ε with \mathbf{uLv} can be maintained in amortized $\mathcal{O}(m^\varepsilon)$ update time and $\mathcal{O}(\min(m^{1+\varepsilon}, n^2))$ space. As $|\mathcal{H}| \in \mathcal{O}(m^{1-\varepsilon})$, the space requirement for \mathbf{t} is $\mathcal{O}(\min(m^{1+\varepsilon}, n^2))$. By Lemma 7, \mathcal{D}_ε with \mathbf{t} can be maintained in $\mathcal{O}(m^{1-\varepsilon} + m^{2\varepsilon}m^{-\varepsilon}) + \mathcal{O}(m^\varepsilon) = \mathcal{O}(m^{\max(1-\varepsilon, \varepsilon)})$ amortized update time and $\mathcal{O}(\min(m^{1+\varepsilon}, n^2))$ space. ◀

► **Lemma 11.** \mathcal{D}_ε with \mathbf{uLLv} can be maintained in amortized $\mathcal{O}(m^{2\varepsilon})$ time per update and $\mathcal{O}(\min(m^{1+2\varepsilon}, n^2))$ space.

Proof. Let $\{u, v\}$ be the newly inserted edge. If $u \in \mathcal{L}$ ($v \in \mathcal{L}$), we count the length-3 paths starting/ending with $\{u, v\}$ as follows: For each low-degree neighbor $w \in N_{\bar{v}}(u) \cap \mathcal{L}$ ($w \in N_{\bar{u}}(v) \cap \mathcal{L}$), increment $\mathbf{uLLv}[v, x]$ ($\mathbf{uLLv}[u, x]$) by one for each $x \in N(w) \setminus \{u, v\}$. This takes $\mathcal{O}(m^{2\varepsilon})$ time. If both $u \in \mathcal{L}$ and $v \in \mathcal{L}$, we additionally count the length-3 paths having $\{u, v\}$ as centerpiece in $\mathcal{O}(\deg(u)^2) = \mathcal{O}(m^{2\varepsilon})$ time: For each pair of distinct vertices x, y with $x \in N_{\bar{v}}(u)$, $y \in N_{\bar{u}}(v)$, increment $\mathbf{uLLv}[x, y]$ by one.

If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} , we iterate over all pairs of distinct vertices y, w , where $y \in N_{\bar{v}}(x)$ for some low-degree neighbor $x \in N(v) \cap \mathcal{L}$ and $w \in N_{\bar{x}}(v)$, and decrease $\mathbf{uLLv}[w, y]$ by one. As v has $\mathcal{O}(m^{2\varepsilon})$ pairs of neighbors and each low-degree neighbor has in turn $\mathcal{O}(m^\varepsilon)$ neighbors, the running time is in $\mathcal{O}(m^{3\varepsilon})$.

Each edge may be incident to two low-degree vertices and hence form $\mathcal{O}(m^{2\varepsilon})$ length-3 paths with the other edges incident to the end vertices. The space requirement hence is $\mathcal{O}(\min(m^{1+2\varepsilon}, n^2))$. By Lemma 7, \mathcal{D}_ε with \mathbf{uLLv} can be maintained in $\mathcal{O}(m^{2\varepsilon} + m^{3\varepsilon}m^{-\varepsilon}) = \mathcal{O}(m^{2\varepsilon})$ amortized update time and $\mathcal{O}(\min(m^{1+2\varepsilon}, n^2))$ space. ◀

► **Lemma 12.** \mathcal{D}_ε with \mathbf{cLv} can be maintained in amortized $\mathcal{O}(m^{2\varepsilon})$ time per update and $\mathcal{O}(n \min(n, m^{2\varepsilon}))$ space.

Proof. Let $\{u, v\}$ be the newly inserted edge. If $u \in \mathcal{L}$ ($v \in \mathcal{L}$), we update the number of claws where u (v) is the central vertex as follows: For each pair of distinct neighbors $x, y \in N_{\bar{v}}(u)$ ($x, y \in N_{\bar{u}}(v)$), increment $\mathbf{cLv}[x, y]$ by one. This accommodates for the claws

18:10 Fully Dynamic Four-Vertex Subgraph Counting

where $v(u)$ is not an anchor vertex and takes $\mathcal{O}(m^{2\varepsilon})$ time. For the other case, increment $\text{cLV}[v, w]$ ($\text{cLV}[u, w]$) by $\deg(u) - 2$ ($\deg(v) - 2$) for each $w \in N_{\bar{v}}(u)$ ($w \in N_{\bar{u}}(v)$) in $\mathcal{O}(m^\varepsilon)$ time.

If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} , we decrease $\text{cLV}[x, y]$ by $\deg(v) - 2$ for each pair of distinct neighbors $x, y \in N(v)$ in total $\mathcal{O}(\deg(v)^2) = \mathcal{O}(m^{2\varepsilon})$ time.

As each vertex may be adjacent to at least one low-degree vertex and we store the count for all pairs, the space requirement is in $\mathcal{O}(n^2)$. On the other hand, each low-degree vertex has at most $\mathcal{O}(m^{2\varepsilon})$ neighbors that can serve as anchors, which yields a space requirement of $\mathcal{O}(nm^{2\varepsilon})$. By Lemma 7, \mathcal{D}_ε with cLV can be maintained in $\mathcal{O}(m^{2\varepsilon} + m^{2\varepsilon}m^{-\varepsilon}) = \mathcal{O}(m^{2\varepsilon})$ amortized update time and $\mathcal{O}(n \min(n, m^{2\varepsilon}))$ space. \blacktriangleleft

► **Lemma 13.** \mathcal{D}_ε with pLL can be maintained in amortized $\mathcal{O}(m^{2\varepsilon})$ time per update and $\mathcal{O}(\min(n^2, nm^{2\varepsilon}))$ space.

Proof. Let $\{u, v\}$ be the newly inserted edge.

If $u \in \mathcal{L}$ ($v \in \mathcal{L}$): First, we update all paws where $u(v)$ is the central vertex and $v(u)$ is the anchor vertex at the arm: For each ordered pair of distinct neighbors $x, y \in N_{\bar{v}}(u)$ ($x, y \in N_{\bar{u}}(v)$) such that $x \in \mathcal{L}$ and $\{x, y\} \in E$, increment $\text{pLL}[v, y]$ ($\text{pLL}[u, y]$) by one. This can be done in $\mathcal{O}(m^{2\varepsilon})$ time. Second, we update all paws where $u(v)$ is the central vertex and $v(u)$ is the anchor vertex in the triangle: For each ordered pair of distinct neighbors $x, y \in N_{\bar{v}}(u)$ ($x, y \in N_{\bar{u}}(v)$) such that $x \in \mathcal{L}$ and $\{x, v\} \in E$ ($\{x, u\} \in E$), increment $\text{pLL}[v, y]$ ($\text{pLL}[u, y]$) by one. This again can be done in $\mathcal{O}(m^{2\varepsilon})$ time. Third, we update all paws where $u(v)$ is the non-anchor, non-central vertex in the triangle and $v(u)$ is the anchor vertex in the triangle: For each neighbor $x \in N_{\bar{v}}(u)$ ($x \in N_{\bar{u}}(v)$) with $x \in \mathcal{L}$ and $\{v, x\} \in E$ ($\{u, x\} \in E$), increment $\text{pLL}[v, y]$ ($\text{pLL}[u, y]$) by one for each $y \in N(x) \setminus \{u, v\}$. The running time is in $\mathcal{O}(m^{2\varepsilon})$, as $u, x \in \mathcal{L}$.

If both $u \in \mathcal{L}$ and $v \in \mathcal{L}$, we update all paws where $\{u, v\}$ connects the central vertex to the non-anchor vertex in the triangle: For each ordered pair of distinct neighbors $x, y \in N_{\bar{v}}(u)$ with $\{x, v\} \in E$ and each ordered pair of distinct neighbors $x, y \in N_{\bar{u}}(v)$ with $\{x, u\} \in E$, increment $\text{pLL}[x, y]$ by one. The running time is in $\mathcal{O}(\deg(u)^2 + \deg(v)^2) = \mathcal{O}(m^{2\varepsilon})$.


If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} : For all paws where v was the central vertex, we iterate over all unordered pairs of neighbors $y, z \in N(v)$ and every neighbor $x \in N(v) \cap \mathcal{L}$ such that $y \neq x \neq z$. If $\{x, y\} \in E$, we decrease $\text{pLL}[y, z]$ by one, and if $\{x, z\} \in E$, we also decrease $\text{pLL}[y, z]$ by one. For all paws where v was the low-degree, non-central vertex in the triangle, we iterate over all pairs of distinct neighbors $x, y \in N(v)$ such that $\{x, y\} \in E$ and $x \in \mathcal{L}$, and, for each $z \in N(x) \setminus \{v, y\}$, decrease $\text{pLL}[y, z]$ by one. In this case, $\{x, z\}$ forms the arm. As $\deg(x) \in \mathcal{O}(m^\varepsilon)$ in the second case, the total running time is $\mathcal{O}(\deg(v)^3 + \deg(v)^2 \cdot m^\varepsilon) = \mathcal{O}(m^{3\varepsilon})$.

The argument for the space requirement is the same as for to cLV and $\mathcal{O}(n \min(n, m^{2\varepsilon}))$ by Lemma 12. By Lemma 7, \mathcal{D}_ε with pLL can be maintained in $\mathcal{O}(m^{2\varepsilon} + m^{3\varepsilon}m^{-\varepsilon}) = \mathcal{O}(m^{2\varepsilon})$ amortized update time and $\mathcal{O}(n \min(n, m^{2\varepsilon}))$ space. \blacktriangleleft

► **Lemma 14.** \mathcal{D}_ε with uHv can be maintained in amortized $\mathcal{O}(m^{\max(1-\varepsilon, \varepsilon)})$ time per update and $\mathcal{O}(n + \min(n^2, m^{2-2\varepsilon}))$ space.

Proof. Let $\{u, v\}$ be the newly inserted edge. If $u, v \in \mathcal{H}$, we iterate over all $h \in \mathcal{H} \setminus \{u, v\}$. If h is adjacent to $u(v)$, increment $\text{uHv}[h, v]$ ($\text{uHv}[u, h]$), respectively, by one. The running time is $\mathcal{O}(|\mathcal{H}|) = \mathcal{O}(m^{1-\varepsilon})$.

If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} (analogously vice-versa): For each pair of distinct high-degree neighbors $x, y \in N(v) \cap \mathcal{H}$, increment $\text{uHv}[x, y]$ by one in total $\mathcal{O}(\deg(v)^2) = \mathcal{O}(m^{2\varepsilon})$ time. *Only if v changes from \mathcal{L} to \mathcal{H} :* For every high-degree neighbor $w \in N(v) \cap \mathcal{H}$, we



■ **Figure 2** Counting the number of diamonds that contain an edge $\{u, v\}$. Green, dashed edges belong to paths that are considered via auxiliary counts, whereas dotted edges are edges whose presence is looked up by the algorithm. As before, small and filled vertices have low degree, large and empty vertices high degree, medium-sized and shaded vertices can have either high or low degree.

iterate over all $h \in \mathcal{H}$ and increase $\text{uHv}[v, h]$ by one if $\{w, h\} \in E$ in total $\mathcal{O}(\deg(v) \cdot |\mathcal{H}|) = \mathcal{O}(m^\varepsilon \cdot m^{1-\varepsilon}) = \mathcal{O}(m)$ time. Only if v changes from \mathcal{H} to \mathcal{L} , we set $\text{uHv}[v, h] := 0$ for each $h \in \mathcal{H}$ in total $\mathcal{O}(|\mathcal{H}|) = \mathcal{O}(m^{1-\varepsilon})$ time. The overall time is hence $\mathcal{O}(m^{\max(2\varepsilon, 1)})$.

There are $\mathcal{O}(\min(n, m^{1-\varepsilon}))$ high-degree vertices, which results in $\mathcal{O}(\min(n^2, m^{2-2\varepsilon}))$ pairs of anchor vertices. \mathcal{D}_ε with uHv can hence be maintained in $\mathcal{O}(m^{1-\varepsilon} + m^{\max(2\varepsilon, 1)}m^{-\varepsilon}) = \mathcal{O}(m^{\max(1-\varepsilon, \varepsilon)})$ amortized update time and $\mathcal{O}(n + \min(n^2, m^{2-2\varepsilon}))$ space by Lemma 7. ◀

► **Lemma 15.** \mathcal{D}_ε with cL can be maintained in amortized $\mathcal{O}(m^{2\varepsilon})$ time per update and $\mathcal{O}(\min(n^3, nm^{3\varepsilon}, m^{1+2\varepsilon}))$ space.

Proof. Let $\{u, v\}$ be the newly inserted edge. If $u \in \mathcal{L}$ ($v \in \mathcal{L}$), increment $cL[v, x, y]$ ($cL[u, x, y]$) by one for each pair of distinct neighbors $x, y \in N_{\bar{v}}(u)$ ($x, y \in N_{\bar{u}}(v)$). This takes $\mathcal{O}(m^{2\varepsilon})$ time.

If a vertex $v \in \mathcal{L}$ changes to \mathcal{H} : For each triple of distinct neighbors $x, y, z \in N(v)$, we decrease $cL[x, y, z]$ by one in total $\mathcal{O}(\deg(v)^3) = \mathcal{O}(m^{3\varepsilon})$ time.

As each edge may be incident to a low-degree vertex v , the number of triples with non-zero count for cL is in $\mathcal{O}(\min(n^3, nm^{3\varepsilon}, m^{1+2\varepsilon}))$. By Lemma 7, \mathcal{D}_ε with cL can be maintained in $\mathcal{O}(m^{2\varepsilon} + m^{3\varepsilon-\varepsilon}) = \mathcal{O}(m^{2\varepsilon})$ amortized update time and $\mathcal{O}(\min(n^3, m^{1+2\varepsilon}))$ space. ◀

Non-Induced Subgraph Counts

We are now ready to prove Theorem 5 and show for each connected subgraph on four vertices how to count it using the data structure \mathcal{D}_ε .

► **Lemma 16.** Let $G = (V, E)$ be a dynamic graph, $\varepsilon \in [0, 1]$, and \mathcal{P} be the diamond \diamond . We can query $c(G, \mathcal{P}, \{u, v\})$ for an arbitrary edge $\{u, v\} \in E$ in $\mathcal{O}(\min(m^{\max(1-\varepsilon, 2\varepsilon)}, n^2))$ worst-case time if we maintain the data structure \mathcal{D}_ε with auxiliary counts uLv , pLL , uHv , and cL .

Proof. Edge $\{u, v\}$ can either be the chord of the diamond or be part of the 4-cycle. See Figure 2 for an illustration.

For the first case, where $\{u, v\}$ is the chord: (a) If $u, v \in \mathcal{H}$, we can obtain the number of length-2 paths p between u and v as $p = \text{uLv}[u, v] + \text{uHv}[u, v]$. As each pair of length-2 paths forms a diamond with $\{u, v\}$, the total number of diamonds is $\binom{p}{2}$. (b) Otherwise, $\{u, v\} \cap \mathcal{L} \neq \emptyset$. W.l.o.g., $u \in \mathcal{L}$. We then iterate over all distinct, unordered pairs of neighbors $x, y \in N_{\bar{v}}(u)$ in $\mathcal{O}(\deg(u)^2) = \mathcal{O}(\min(m^{2\varepsilon}, n^2))$ time. For each such pair with $\{x, v\}, \{y, v\} \in E$, we count one diamond.

For the second case, where $\{u, v\}$ is part of the cycle, we distinguish between the degrees of the other two vertices. (a) The number of diamonds where the other two vertices have low degree is given by $\text{pLL}[u, v]$. Note that either u or v is incident to the chord. (b) The number of diamonds where the other vertex incident to the chord has low degree and

18:12 Fully Dynamic Four-Vertex Subgraph Counting

the fourth vertex has high degree can be obtained by iterating over all $h \in \mathcal{H} \setminus \{u, v\}$ in $\mathcal{O}(|\mathcal{H}|) = \mathcal{O}(\min(m^{1-\varepsilon}, n))$ time. If either $\{h, u\} \in E$ or $\{h, v\} \in E$, we have $\text{cL}[u, v, h]$ more diamonds. If both $\{h, u\}, \{h, v\} \in E$, we add $2\text{cL}[u, v, h]$ to the number of diamonds. (c, d) The number of diamonds where the other vertex incident to the chord has high degree can be obtained as follows: (c) If $u \in \mathcal{H}$ ($v \in \mathcal{H}$), the number of diamonds where the chord is incident to u (v) can be obtained by iterating over all $h \in \mathcal{H} \setminus \{u, v\}$ in $\mathcal{O}(|\mathcal{H}|) = \mathcal{O}(\min(m^{1-\varepsilon}, n))$ time. For each such vertex h , we check whether $\{u, h\}, \{v, h\} \in E$ and add $\text{uLv}[u, h] + \text{uHv}[u, h] - 1$ ($\text{uLv}[v, h] + \text{uHv}[v, h] - 1$) to the count. The correction by 1 is necessary because the auxiliary counts also contain the path $\langle \{u, v\}, \{v, h\} \rangle$ ($\langle \{h, u\}, \{u, v\} \rangle$). (d) If $u \in \mathcal{L}$ ($v \in \mathcal{L}$), we iterate over all high-degree neighbors $h \in N_{\bar{v}}(u) \cap \mathcal{H}$ ($h \in N_{\bar{u}}(v) \cap \mathcal{H}$) and in each case over all $x \in N(u) \setminus \{v, h\}$ ($x \in N(v) \setminus \{u, h\}$) in total $\mathcal{O}(\min(m^{2\varepsilon}, n^2))$ time and count one diamond each if $\{h, x\}, \{h, v\} \in E$ ($\{h, x\}, \{h, u\} \in E$). ◀

► **Lemma 17.** *Let $G = (V, E)$ be a dynamic graph and \mathcal{P} be the diamond \diamond . We can maintain $\mathbf{c}(G, \mathcal{P})$ in amortized $\mathcal{O}(m^{2/3})$ update time and $\mathcal{O}(\min(nm, m^{5/3}))$ space. We can query $\mathbf{c}(G, \mathcal{P}, e)$ for an arbitrary edge $e \in E$ in worst-case $\mathcal{O}(m^{2/3})$ time.*

Proof. After an edge $\{u, v\}$ was inserted or before an edge $\{u, v\}$ is removed, the number of diamonds containing it can be obtained in $\mathcal{O}(\min(m^{\max(1-\varepsilon, 2\varepsilon)}, n^2))$ time worst-case time by Lemma 16 if \mathcal{D}_ε with auxiliary counts uLv , pLL , uHv , and cL is maintained. By Lemma 9, Lemma 13, Lemma 14, and Lemma 15, this can be done in amortized $\mathcal{O}(m^\varepsilon + m^{\max(1-\varepsilon, 2\varepsilon)} + m^{2\varepsilon}) = \mathcal{O}(m^{\max(1-\varepsilon, 2\varepsilon)})$ time and $\mathcal{O}(\min(n^3, \max(n^2, nm^{3\varepsilon}, m^{2-2\varepsilon}, m^{1+2\varepsilon})))$ space. Together with the cost for the query, this yields a total amortized update time of $\mathcal{O}(m^{\max(1-\varepsilon, 2\varepsilon)}) = \mathcal{O}(m^{2/3})$ for $\varepsilon = \frac{1}{3}$ and $\mathcal{O}(\min(nm, m^{5/3}))$ space. By Lemma 16, the worst-case time to query $\mathbf{c}(G, \mathcal{P}, e)$ for an arbitrary edge $e \in E$ then is $\mathcal{O}(m^{2/3})$. ◀

Queries with Vertices and Edges and Non-Induced s -Subgraph Counts

With similar techniques, we can count non-induced triangles containing a specified vertex or edge as well as maintain s -subgraph counts for patterns with up to four vertices.

► **Theorem 18.** *Let $G = (V, E)$ be a dynamic graph, \mathcal{P} be the 3-cycle \triangle , and $\varepsilon \in [0, 1]$. We can query $\mathbf{c}(G, \mathcal{P}, a)$ for an arbitrary vertex or edge a in*

- (i) *worst-case $\mathcal{O}(\min(m^{2\varepsilon}, n^2))$ time with $\mathcal{O}(m^{\max(\varepsilon, 1-\varepsilon)})$ amortized update time and $\mathcal{O}(\min(n^2, m^{1+\varepsilon}))$ space if $a \in V$,*
- (ii) *worst-case time $\mathcal{O}(\min(m^{1-\varepsilon}, n))$ with $\mathcal{O}(m^\varepsilon)$ amortized update time and $\mathcal{O}(\min(n^2, m^{1+\varepsilon}))$ space if $a \in E$.*

► **Corollary 19.** *Let $G = (V, E)$ be a dynamic graph and \mathcal{P} the 3-cycle \triangle . We can maintain $\mathbf{c}(G, \mathcal{P})$ with an amortized update time of $\mathcal{O}(\sqrt{m})$ and $\mathcal{O}(\min(n^2, m^{1.5}))$ space and query $\mathbf{c}(G, \mathcal{P}, e)$ for $e \in E$ arbitrary in worst-case $\mathcal{O}(\sqrt{m})$ time. We can query $\mathbf{c}(G, \mathcal{P}, v)$ for arbitrary $v \in V$ in worst-case $\mathcal{O}(m^{2/3})$ time with an amortized update time of $\mathcal{O}(m^{2/3})$ and $\mathcal{O}(\min(n^2, m^{4/3}))$ space.*

► **Theorem 20.** *Let $G = (V, E)$ be a dynamic graph, $s \in V$, and \mathcal{P} be a connected subgraph. We can maintain the non-induced s -subgraph count $\mathbf{c}(G, \mathcal{P}, s)$ in*

- (i) *worst-case constant update time and constant space if \mathcal{P} is the claw λ ,*
- (ii) *amortized update time $\mathcal{O}(\sqrt{m})$ and $\mathcal{O}(n)$ space if \mathcal{P} is the 3-cycle \triangle or the length-3 path ρ ,*
- (iii) *amortized update time $\mathcal{O}(m^{2/3})$ and $\mathcal{O}(n^2)$ space if \mathcal{P} is the paw \blacktriangleleft , the 4-cycle \diamond , or the diamond \diamond ,*
- (iv) *worst-case update time $\mathcal{O}(m)$ and constant space if \mathcal{P} is the 4-clique \boxtimes .*

4 Lower Bounds

We give new lower bounds for detecting and counting induced and non-induced subgraphs.

Induced Subgraph Counts

Our results for counting induced subgraphs on four vertices are conditioned on the combinatorial k -clique conjecture:

► **Theorem 21.** *Let G be a dynamic graph, $\mathcal{P} \in \mathbb{P}_4$, and let $\gamma > 0$ be a small constant. There is no incremental or fully dynamic combinatorial algorithm with preprocessing time $\mathcal{O}(m^{2-\gamma})$ for maintaining $\mathbf{c}_I(G, \mathcal{P})$ in amortized update time $\mathcal{O}(m^{1-\gamma})$ and query time $\mathcal{O}(m^{2-\gamma})$, unless the k -clique conjecture fails.*

Before we turn to the proof, we recall the following relations between subgraph counts.

► **Lemma 22** ([9]). *For each pair $P, P' \in \mathbb{P}_4$, the non-induced subgraph count $\mathbf{c}(P, P') = 1$ if $P = P'$ and otherwise nonzero only in the following cases:*

$$\begin{array}{llll} \mathbf{c}(\neg\triangleleft, \neg\triangleleft) = 2, & \mathbf{c}(\diamond, \neg\triangleleft) = 4, & \mathbf{c}(\diamond, \triangleleft) = 6, & \mathbf{c}(\boxtimes, \neg\triangleleft) = 12, \\ \mathbf{c}(\neg\triangleleft, \wedge) = 1, & \mathbf{c}(\diamond, \wedge) = 2, & \mathbf{c}(\boxtimes, \wedge) = 4, & \mathbf{c}(\diamond, \neg\triangleleft) = 4, \\ \mathbf{c}(\boxtimes, \neg\triangleleft) = 12, & \mathbf{c}(\diamond, \diamond) = 1, & \mathbf{c}(\boxtimes, \diamond) = 3, & \mathbf{c}(\boxtimes, \diamond) = 6. \end{array}$$

► **Proposition 23** ([17]). *Let G, \mathcal{P} be graphs and let k be the number of vertices of \mathcal{P} . Then, $\mathbf{c}(G, \mathcal{P}) = \sum_{P \in \mathbb{P}_k} \mathbf{c}_I(G, P) \cdot \mathbf{c}(P, \mathcal{P})$.*

► **Lemma 24.** *Let $G = (V, E)$ be a graph. The following relationships between induced and non-induced subgraph counts hold:*

$$\begin{array}{l} \mathbf{c}_I(G, \boxtimes) = \mathbf{c}(G, \boxtimes) \\ \mathbf{c}_I(G, \diamond) = \mathbf{c}(G, \diamond) - 6\mathbf{c}_I(G, \boxtimes) \\ \mathbf{c}_I(G, \triangleleft) = \mathbf{c}(G, \triangleleft) - \mathbf{c}(G, \diamond) + 3\mathbf{c}_I(G, \boxtimes) \\ \mathbf{c}_I(G, \neg\triangleleft) = \mathbf{c}(G, \neg\triangleleft) - 4\mathbf{c}(G, \diamond) + 12\mathbf{c}_I(G, \boxtimes) \\ \mathbf{c}_I(G, \wedge) = \mathbf{c}(G, \wedge) - \mathbf{c}(G, \neg\triangleleft) + 2\mathbf{c}(G, \diamond) - 4\mathbf{c}_I(G, \boxtimes) \\ \mathbf{c}_I(G, \neg\triangleleft) = \mathbf{c}(G, \neg\triangleleft) - 2\mathbf{c}(G, \diamond) - 4\mathbf{c}(G, \triangleleft) + 6\mathbf{c}(G, \diamond) - 12\mathbf{c}_I(G, \boxtimes) \end{array}$$


Proof. Let $\mathcal{P} \in \mathbb{P}_4$. The statement follows from Proposition 23 and Lemma 22 by substituting equations and solving them for $\mathbf{c}_I(G, \mathcal{P})$ in the order as listed. ◀

► **Corollary 25.** *Let $G = (V, E)$ be a dynamic graph and $\mathcal{P} \in \mathbb{P}_4$. We can maintain $\mathbf{c}_I(G, \mathcal{P})$ with $\mathcal{O}(m)$ worst-case update time and constant space.*

Proof of Theorem 21. First consider the case that $\mathcal{P} = \boxtimes$. Suppose there is an incremental or fully dynamic algorithm \mathcal{A} that maintains $\mathbf{c}_I(G, \mathcal{P})$ in time $\mathcal{O}(m^{1-\gamma})$ with query time $\mathcal{O}(m^{2-\gamma})$ for some $\gamma > 0$. Construct an algorithm \mathcal{A}' for static 4-clique detection as follows: Run \mathcal{A} on an initially empty graph, insert all edges one-by-one in total $\mathcal{O}(m^{2-\gamma})$ time, and query the result in $\mathcal{O}(m^{2-\gamma})$ time. As $\mathcal{O}(m^{2-\gamma}) = \mathcal{O}(n^{4-2\gamma})$ this contradicts Conjecture 2.

For the remaining five induced four-vertex subgraphs, let \mathcal{P} be such a subgraph. We construct a deterministic algorithm \mathcal{A}' for static 4-clique detection as follows: \mathcal{A}' executes the above operations for our non-induced subgraph counting algorithm from Sect. 3, which can maintain the number of all connected subgraphs on four vertices with $\mathcal{O}(m^{2/3})$ amortized update time and $\mathcal{O}(1)$ query time by Theorem 5. It thus takes $\mathcal{O}(m^{5/3})$ time in total to compute $\mathbf{c}(G, \mathcal{P})$ for all $\mathcal{P}' \in \mathbb{P}_4 \setminus \{\boxtimes\}$. Assume by contradiction that there exists an algorithm

18:14 Fully Dynamic Four-Vertex Subgraph Counting



■ **Figure 3** Construction of $G_{M,g,h}$ and G for $u^\top M v$ for (s) -5-cycle and (s) -diamond detection.

\mathcal{A}^* that maintains $\mathbf{c}_I(G, \mathcal{P})$ in update time $\mathcal{O}(m^{1-\gamma})$ and query time $\mathcal{O}(m^{2-\gamma})$. Then \mathcal{A}' also executes the same operations with \mathcal{A}^* to compute $\mathbf{c}_I(G, \mathcal{P})$. Using the formula for \mathcal{P} in Lemma 24, \mathcal{A}' can solve the static 4-clique detection problem in time $\mathcal{O}(m^{\max(5/3, 2-\gamma)}) \subseteq \mathcal{O}(n^{4-\delta})$ time for $\delta = \min(2\gamma, 2/3) > 0$, a contradiction to Conjecture 2. ◀

Non-Induced Subgraph Counts

In this section, we give new lower bounds for detecting (and thus counting) cycles of arbitrary length, paws, diamonds, and 4-cliques¹, as well as counting length-3 paths. Our results are based on the OMv conjecture. In [12] it is proven that instead of reducing from OMv directly it suffices to reduce from the following 1-uMv version: For any positive integer parameters n_1, n_2 , given an $n_1 \times n_2$ matrix M , there is no algorithm with preprocessing time polynomial in n_1 and n_2 that computes for an n_1 -dimensional vector u and an n_2 -dimensional vector v the product $u^\top M v$ in time $\mathcal{O}(n_1 n_2^{1-\delta} + n_1^{1-\delta} n_2)$ for any small constant $\delta > 0$ with error probability of at most $\frac{1}{3}$ in the word-RAM model with $\mathcal{O}(\log n)$ bit words.

► **Theorem 26.** *Let G be a partially dynamic graph and let \mathcal{P} be the paw \triangleleft , the diamond \diamond , the 4-clique \boxtimes , or a k -cycle with $k \geq 3$. On condition of Conjecture 1, there is no partially dynamic algorithm to maintain whether $\mathbf{c}(G, \mathcal{P}) > 0$ with polynomial preprocessing time and worst-case update time $\mathcal{O}(m^{1/2-\delta})$ and query time $\mathcal{O}(m^{1-\delta})$ with an error probability of at most $1/3$ for any $\delta > 0$. This also holds for fully dynamic algorithms with amortized update time and for paws, diamonds, 4-cliques, or odd k -cycles containing a specific vertex s , as well as for maintaining the number of length-3 paths $\not\sim$ and the number of length-3 paths containing a specific vertex s .*

Our constructions build on the following graph (see Figure 3 for an example).

► **Definition 27** ($G_{M,g,h}$). *Given a matrix $M \in \{0, 1\}^{n_1 \times n_2}$ and two integers $g, h \geq 0$, we denote by $G_{M,g,h} = (\bigcup_{0 \leq p \leq g} L^{(p)} \cup \bigcup_{0 \leq q \leq h} R^{(q)}, E_L \cup E_R \cup E_M)$ the $(g+h+2)$ -partite graph with*

$$L^{(p)} = \{l_1^{(p)}, \dots, l_{n_1}^{(p)}\}, \quad 0 \leq p \leq g; \quad E_L = \{(l_i^{(p)}, l_i^{(p+1)}) \mid 1 \leq i \leq n_1 \wedge 0 \leq p < g\}$$

$$R^{(q)} = \{r_1^{(q)}, \dots, r_{n_2}^{(q)}\}, \quad 0 \leq q \leq h; \quad E_R = \{(r_j^{(q)}, r_j^{(q+1)}) \mid 1 \leq j \leq n_2 \wedge 0 \leq q < h\}$$

and $E_M = \{(l_i^{(g)}, r_j^{(h)}) \mid M_{ij} = 1\}$. $G_{M,g,h}$ has $(g+1) \cdot n_1 + (h+1) \cdot n_2$ vertices and at most $n_1 n_2 + g \cdot n_1 + h \cdot n_2$ edges. All vertices x in $L^{(p)}$ for $1 \leq p < g$ and in $R^{(q)}$ for $1 \leq q < h$ have $\deg(x) = 2$, whereas all vertices y in $L^{(0)}$ and in $R^{(0)}$ have $\deg(y) = 1$.

For convenience, we set $L := L^{(0)}$, $R := R^{(0)}$, $l_i := l_i^{(0)}$, and $r_j := r_j^{(0)}$.

► **Observation 28.** *Every cycle in $G_{M,g,h}$ is even and has length at least 4.*

¹ Theorem 21 applies also to 4-cliques, but it is based on a different assumption.

Let s be a fixed vertex in the graph. The s - k -cycle detection problem requires the algorithm to detect whether a k -cycle containing s exists. We use the same notation for the other subgraphs as well.

► **Lemma 29.** *Given a partially dynamic algorithm \mathcal{A} for one of the problems listed below, one can solve 1-uMv with parameters n_1 and n_2 by running the preprocessing step of \mathcal{A} on a graph with $\mathcal{O}(m + \sqrt{m} \cdot k)$ edges and $\Theta(\sqrt{m} \cdot k)$ vertices, and then making $\mathcal{O}(\sqrt{m})$ insertions (or $\mathcal{O}(\sqrt{m})$ deletions) and 1 query, where m is such that $n_1 = n_2 = \sqrt{m}$. The problems are*

- | | |
|--|--|
| (a) (s) - k -Cycle Detection for odd k | (d) (s) - k -Clique Detection for $k = 4$ |
| (b) (s) -Paw Detection | (e) (s) -length- k Path Counting for $k = 3$ |
| (c) (s) -Diamond Detection | (f) k -Cycle Detection |

Proof of Case (c). We only prove the decremental case. Consider a 1-uMv problem with $n_1 = n_2 = \sqrt{m}$. Given M , we construct the tripartite graph G from $G_{M,1,0}$ by adding to it a vertex s and connecting it by an edge to every vertex in $G_{M,1,0}$. Thus, the total number of edges is at most $n_1 n_2 + 3n_1 + n_2 = \mathcal{O}(m)$. Once u and v arrive, we delete $\{s, l_i\}$ and $\{s, l_i^{(1)}\}$ iff $u_i = 0$ and delete $\{r_j, s\}$ iff $v_j = 0$. See Figure 3 for an example.

Consider the case that G contains a diamond with chord e . As every triangle must be incident to s by Observation 28, $e = \{s, x\}$ for some $x \in L \cup L^{(1)} \cup R$. Furthermore, all vertices in L have degree at most two and x has degree at least three, so $x \notin L$. If $x = r_j \in R$, then by construction, x must have two neighbors $l_i^{(1)}, l_{i'}^{(1)} \in L^{(1)}$ such that there are edges $\{s, l_i^{(1)}\}$ and $\{s, l_{i'}^{(1)}\}$ in G . Again by construction, there are thus also edges $\{s, l_i\}$ and $\{s, l_{i'}\}$ and a diamond $\{s, l_i, l_{i'}^{(1)}, x = r_j\}$ with chord $\{s, l_i^{(1)}\}$. As each vertex in $L^{(1)}$ is adjacent to exactly one vertex in L , every diamond must contain a vertex $r \in R$ and the edge $\{r, s\}$. Hence, we have $u^\top M v = 1$ iff there is a diamond in G iff there is a diamond incident to s . In total, we need to do $2n_1 + n_2 = \mathcal{O}(\sqrt{m})$ updates and 1 query. ◀

5 Conclusion

Our focus in this work was especially on non-induced and induced four-vertex subgraphs. We gave improved both upper and lower bounds for detecting or counting four-vertex subgraphs in the dynamic setting, thereby closing the gap (w.r.t. improvements by a polynomial factor) for counting non-induced length-3 paths and narrowing it considerably for non-induced paws, 4-cycles, and diamonds. For counting induced subgraphs, we showed that the update time of the algorithm by Eppstein et al. [9] cannot be improved by a polynomial factor, but that a better space complexity can be achieved in the worst case.

Many of our lower bounds also apply to subgraphs with more than four vertices, but to the best of our knowledge, only algorithms for cliques have been considered here so far. Hence, besides closing the gap for four-vertex subgraphs, the complexity of detecting and counting subgraphs with five or more vertices would be an interesting field for future work.

We also investigated the complexity of querying the number of subgraphs containing a specific edge, as such queries are relevant, e.g., to measure the similarity of graphs via histograms. This can similarly be done for vertices. As a by-product of our results for four-vertex subgraphs, we showed for 3-cycles that such vertex queries can be answered in $\mathcal{O}(m^{\frac{2}{3}})$. The complexity of vertex queries for larger subgraphs remains an open question.

Further interesting lines to follow regard the complexity of approximate counting, counting all approximately densest subgraphs, as well as the complexity of *enumerating* subgraphs.

Our work was strongly motivated also by the practical relevance of counting subgraphs in the dynamic setting. For this reason, we consider an experimental evaluation of dynamic subgraph counting algorithms a relevant and very interesting task for future work.

References

- 1 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is valiant’s parser. *SIAM Journal on Computing*, 47(6):2527–2555, 2018.
- 2 N. Alon, R. Yuster, and U. Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, March 1997. doi:10.1007/BF02523189.
- 3 Ziv Bar-Yossef, Ravi Kumar, and D. Sivakumar. Reductions in streaming algorithms, with an application to counting triangles in graphs. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA ’02*, pages 623–632, USA, 2002. Society for Industrial and Applied Mathematics.
- 4 Thiago Bergamaschi, Monika Henzinger, Maximilian Probst Gutenberg, Virginia Vassilevska Williams, and Nicole Wein. New techniques and fine-grained hardness for dynamic near-additive spanners. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1836–1855. SIAM, 2021. doi:10.1137/1.9781611976465.110.
- 5 Karl Bringmann, Nick Fischer, and Marvin Künnemann. A fine-grained analogue of schaefer’s theorem in p: Dichotomy of $\exists k\forall$ -quantified first-order graph properties. In *34th Computational Complexity Conference*, pages 1–27. Schloss Dagstuhl, 2019.
- 6 Luciana S. Buriol, Gereon Frahling, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Christian Sohler. Counting triangles in data streams. In *Proceedings of the Twenty-Fifth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS ’06*, pages 253–262, New York, NY, USA, 2006. Association for Computing Machinery. doi:10.1145/1142351.1142388.
- 7 Laxman Dhulipala, Quanquan C. Liu, Julian Shun, and Shangdi Yu. Parallel batch-dynamic k-clique counting. *CoRR*, abs/2003.13585, 2020. arXiv:2003.13585.
- 8 Zdeněk Dvořák and Vojtěch Tůma. A dynamic data structure for counting subgraphs in sparse graphs. In Frank Dehne, Roberto Solis-Oba, and Jörg-Rüdiger Sack, editors, *Algorithms and Data Structures*, pages 304–315, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- 9 David Eppstein, Michael T. Goodrich, Darren Strash, and Lowell Trott. Extended dynamic subgraph statistics using h-index parameterized data structures. *Theor. Comput. Sci.*, 447:44–52, 2012. doi:10.1016/j.tcs.2011.11.034.
- 10 David Eppstein and Emma S. Spiro. The h-index of a graph and its application to dynamic subgraph statistics. *J. Graph Algorithms Appl.*, 16(2):543–567, 2012. doi:10.7155/jgaa.00273.
- 11 Kathrin Hanauer, Monika Henzinger, and Qi Cheng Hua. Fully dynamic four-vertex subgraph counting. *CoRR*, abs/2106.15524, 2021. arXiv:2106.15524.
- 12 Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, pages 21–30, 2015.
- 13 Monika Henzinger, Andrea Lincoln, and Barna Saha. The complexity of average-case dynamic subgraph counting. In *Proceedings of the Thirty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Alexandria, Virginia, USA, January 9-12, 2022*. SIAM, 2022. to appear.
- 14 Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Counting Triangles under Updates in Worst-Case Optimal Time. In Pablo Barcelo and Marco Calautti, editors, *22nd International Conference on Database Theory (ICDT 2019)*, volume 127 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:18, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ICDT.2019.4.
- 15 Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Maintaining triangle queries under updates. *ACM Trans. Database Syst.*, 45(3):11:1–11:46, 2020. doi:10.1145/3396375.

- 16 T. Kloks, D. Kratsch, and H. Müller. Finding and counting small induced subgraphs efficiently. In Manfred Nagl, editor, *Graph-Theoretic Concepts in Computer Science*, pages 14–23, Berlin, Heidelberg, 1995. Springer Berlin Heidelberg.
- 17 William L Kocay. Some new methods in reconstruction theory. In *Combinatorial Mathematics IX*, pages 89–114. Springer, 1982.
- 18 Tamara G. Kolda, Ali Pinar, Todd Plantenga, C. Seshadhri, and Christine Task. Counting triangles in massive graphs with mapreduce. *SIAM Journal on Scientific Computing*, 36(5):S48–S77, 2014. doi:10.1137/13090729X.
- 19 F. Le Gall. Powers of tensors and fast matrix multiplication. In K. Nabeshima, K. Nagasaka, F. Winkler, and Á. Szántó, editors, *International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23-25, 2014*, pages 296–303. ACM, 2014. doi:10.1145/2608628.2608664.
- 20 Andrea Lincoln, Virginia Vassilevska Williams, and Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1236–1252. SIAM, 2018.
- 21 Siddhartha Sahu, Amine Mhedhbi, Semih Salihoglu, Jimmy Lin, and M. Tamer Özsu. The ubiquity of large graphs and surprising challenges of graph processing: extended survey. *VLDB J.*, 29(2-3):595–618, 2020. doi:10.1007/s00778-019-00548-x.