

Fast Algorithm for K-Truss Discovery on Public-Private Graphs

Soroush Ebadian^{1,2} and Xin Huang²

¹Sharif University of Technology

²Hong Kong Baptist University

soroushebadian@gmail.com, xinhuang@comp.hkbu.edu.hk

Abstract

In public-private graphs, users share one public graph and have their own private graphs. A private graph consists of personal private contacts that only can be visible to its owner, e.g., hidden friend lists on Facebook and secret following on Sina Weibo. However, existing public-private analytic algorithms have not yet investigated the dense subgraph discovery of k -truss, where each edge is contained in at least $k - 2$ triangles. This paper aims at finding k -truss efficiently in public-private graphs. The core of our solution is a novel algorithm to update k -truss with node insertions. We develop a classification-based hybrid strategy of node insertions and edge insertions to incrementally compute k -truss in public-private graphs. Extensive experiments validate the superiority of our proposed algorithms against state-of-the-art methods on real-world datasets.

1 Introduction

Online social networks (e.g., Facebook, Twitter, Instagram, and Sina Weibo) have become vital platforms for connecting users to share information, post daily life events, and spread influence [Kempe *et al.*, 2003; Wilder *et al.*, 2018; Zhang *et al.*, 2017; Zhang *et al.*, 2018b; Li *et al.*, 2018]. Due to privacy concerns, users tend to hide their connections, leading such private relationships not visible to other users in public but only themselves. For instance, Facebook users are likely to conceal their friend-list [Dey *et al.*, 2012]; Weibo users may prefer using the secret following feature, which hides their interested followees. Public-private graphs are developed to model this kind of social networks [Chierichetti *et al.*, 2015]. A public-private network contains a public graph which is visible and accessible to everyone; in addition, each vertex has a personal private graph only visible to its owner. Therefore, in the view of each user, the social network is a union of the public graph and its own private graph which can be significantly different for distinct users. Recently, many graph analytic tasks have been investigated on public-private networks, such as all-pairs shortest path distances, node similarities, and correlation clustering.




Figure 1: An example of public-private graph. Black solid edges are public. Blue dashed edges are private to v_5 , and red dashed edges are private to v_3 . The gray area is a 5-truss in personalized pp-graph of v_5 as g_{v_5} .

Dense subgraph discovery is a fundamental network analysis task. Many dense subgraphs have been proposed and investigated, e.g., clique, quasi-clique [Pei *et al.*, 2005], n -clan [Mokken, 1979], n -club [Mokken, 1979], and k -plex [Xiao *et al.*, 2017]. Recently, a popular notion of dense subgraphs that has been studied is k -truss. A k -truss is the largest subgraph of a graph such that each edge is contained in at least $k - 2$ triangles within this subgraph. Finding k -trusses has many useful applications such as community search [Jiang *et al.*, 2018], complex network visualization [Zhao and Tung, 2012], and task-driven team formation [Huang *et al.*, 2016]. To the best of our knowledge, finding k -truss over public-private networks has not yet been studied in the literature. In this paper, we formulate the problem of finding public-private k -truss as follows. Given a query vertex and parameter k , the problem is to find k -truss in the public-private graph owned by this query vertex.

Efficient extraction of public-private k -truss raises significant challenges. A straightforward approach is to ignore users' private edges, which can lead to inaccurate results. Another approach is to apply truss decomposition on the public-private graph of the query vertex to extract the k -truss with the given parameter k . However, this method computes k -truss from scratch, which is particularly inefficient for large-scale networks. To tackle these challenges, we develop an indexed computational paradigm to efficiently update a truss index on a public graph using the edges in a private graph, with a minimal amount of recomputation on the public graph.

To summarize, we make the following contributions:

- We formulate a new problem of finding k -truss over public-private graphs, that is finding k -truss in public-

- private graph owned by a given query vertex (Section 3).
- We analyze the structural properties of k -truss on public-private networks. Based on the observations, we develop k -truss updating algorithm using a hybrid strategy of node/edge insertions/deletions (Section 4).
- We validate the efficiency of our proposed methods through extensive experiments on real-world datasets of public-private networks (Section 5).

2 Related Work

Several essential problems of graph analysis on public-private graphs have been studied in [Chierichetti *et al.*, 2015; Archer *et al.*, 2017], such as the size of reachability tree [Cohen and Kaplan, 2007], all-pairs shortest paths [Das Sarma *et al.*, 2010], pairwise node similarities [Haveliwala, 2002], correlation clustering [Bansal *et al.*, 2004]. Moreover, the public-private model of data summarization has been investigated and solved by a fast distributed algorithm [Mirzasoleiman *et al.*, 2016]. In addition, several studies on k -truss mining have been investigated recently [Cohen, 2008; Zhang *et al.*, 2018a]. Equivalent concepts of k -truss termed as different names include triangle k -core [Zhang and Parthasarathy, 2012], k -dense community [Saito *et al.*, 2008; Gregori *et al.*, 2011], and k -mutual-friend subgraph [Zhao and Tung, 2012]. Truss decomposition is to find the non-empty k -truss for all possible k values in a graph. Algorithms of truss decomposition have also been studied in different types of graphs (e.g., directed graphs [Takaguchi and Yoshida, 2016], uncertain graphs [Zou and Zhu, 2017], and dynamic graphs [Zhang and Parthasarathy, 2012; Huang *et al.*, 2014]). In contrast to the above studies, finding k -truss over public-private networks is studied for the first time in this paper.

3 Preliminary

We consider a simple and undirected graph $G = (V, E)$ where V and E are the vertex set and edge set respectively. We define $N(v) = \{u \in V : (v, u) \in E\}$ as the set of neighbors of a vertex v , and $d(v) = |N(v)|$ as the degree of v in G . For a set of vertices $S \subseteq V$, the induced subgraph of G by S is denoted by $G[S]$, where the vertex set is S and the edge set is $E(G[S]) = \{(v, u) \in E : v, u \in S\}$.

3.1 Public-Private Graphs

We first introduce a model of public-private graph \mathcal{G} [Chierichetti *et al.*, 2015]. A public-private graph \mathcal{G} consists of one public graph and multiple private graphs. Given a public graph $G = (V, E)$, the vertex set V represents users, and the edge set E represents connections between users. For each vertex u in the public graph G , u has an associated private graph $G_u = (V_u, E_u)$, where $V_u \subseteq V$ are the users from public graph and the edge set E_u satisfies $E_u \cap E = \emptyset$. The public graph G is visible to everyone, and the private graph G_u is only visible to user u . Thus, in the view of user u , she/he can see and access the structure of graph that is the union of public graph G and its own private graph G_u , i.e., $G \cup G_u = (V, E \cup E_u)$ [Huang *et al.*, 2018]. The personalized public-private graph (a.k.a. pp-graph in short) owned by a vertex u is defined as follows.

Definition 1 (Personalized PP-Graph). *Given a public-private graph \mathcal{G} and a vertex u , the personalized pp-graph of u is denoted by g_u , where $g_u = G \cup G_u = (V, E \cup E_u)$. Here, E_u are the private edges only visible to u , and $E \cap E_u = \emptyset$.*

3.2 K-Truss

A triangle is a cycle of length 3 in graphs. Given three vertices $u, v, w \in V$, the triangle formed by u, v, w is denoted by Δ_{uvw} . The support of an edge is defined as follows.

Definition 2 (Support). *Given a subgraph $H \subseteq G$, the support of an edge $e = (u, v)$, denoted by $\text{sup}_H(e)$, is defined as the number of triangles containing edge e in H , i.e., $\text{sup}_H(e) = |\{\Delta_{uvw} : (u, v), (u, w), (v, w) \in E(H)\}|$.*

We drop the subscript and denote the support as $\text{sup}(e)$, when the context is obvious. Based on the support, we give a definition of k -truss [Wang and Cheng, 2012] as follows.

Definition 3 (K-Truss). *A k -truss H of graph G is defined as the largest subgraph of G such that every edge e has support of at least $k - 2$ in this subgraph, i.e., $\text{sup}_H(e) \geq k - 2$.*

3.3 Problem Statement

The problem of public-private k -truss discovery studied in this paper is formulated as follows.

Problem formulation: Given a public-private graph \mathcal{G} , a vertex $u \in V$, and an integer $k \geq 2$, the problem is to find the k -truss in the personalized pp-graph g_u where $g_u = G \cup G_u$.

Example 1. *Consider the public-private graph \mathcal{G} in Figure 1, a query vertex v_5 , and $k = 5$. Black edges are public. Blue edges are private to v_5 . The answer of 5-truss in personalized pp-graph g_{v_5} is the subgraph depicted in the gray region.*

4 Proposed Algorithms

This section introduces our algorithms for finding k -truss in personalized pp-graph g_u , w.r.t. a query vertex u . We first give an overview of our ideas in Section 4.1, and then present a well thorough description of technical details afterward.

4.1 Solution Overview

We consider two different ideas.

Solution 1: online search algorithm. One intuitive approach is to apply truss decomposition on pp-graph g_u to iteratively remove edges with less than $k - 2$ triangles and output the remaining graph as answers. However, such computing k -truss from scratch on g_u for each query vertex u is obviously inefficient for big graphs with a large number of vertices.

Solution 2: index-based search algorithm. Recall that personalized pp-graph $g_u = G \cup G_u$ has a public graph G and a private graph G_u only available to u . The public graph G is available to everyone, and the structure of G is identical to each query vertex u . The idea of index-based search algorithms is to construct a structural index of public graph G offline, and then online find k -truss based on the precomputed index of G and additional graph G_u . In the following, we introduce a concept of trussness, which is useful for constructing the truss-index for k -truss discovery.

Definition 4 (Trussness). *Given a subgraph $H \subseteq G$, the trussness of H is defined as $\tau(H) = \min_{e \in E(H)} \{\sup_H(e) + 2\}$. The trussness of an edge $e \in H$ is defined as the largest number k such that a connected k -truss H' contains e , i.e.,*

$$\tau_H(e) = \max_{H' \subseteq H, e \in E(H')} \tau(H').$$

We drop the subscript and denote $\tau_G(e)$ as $\tau(e)$ when the context is obvious. According to Def. 4, k -truss of G is the union of all edges e with $\tau(e) = k$. The truss-index of public graph G keeps the trussness of all edges in G . Given a truss-index of G , the remaining issue is *how to update the truss-index for pp-graph $G \cup G_u$, w.r.t. the additional G_u .*

Updating truss-index using edge insertions. A simple approach is to add edges of G_u one-by-one into G and update the truss-index accordingly, by using an existing edge-insertion algorithm [Huang *et al.*, 2014]. However, when the number of private edges E_u is large, the adaptation of edge-insertion may be inefficient. For example, consider the example graph G in Figure 1 and query vertex v_5 with 4 private edges. It invokes the edge-insertion algorithm for 4 times.

Our approach. To address the above issue, we propose a batch-update algorithm using node-insertion. The idea is to simultaneously insert a new node u with all its incident edges into graph G at the same time, and call node-insertion algorithm only once. In the above example, it needs only one node-insertion of the isolated node v_5 and all of its private edges. To handle the truss-index update with node insertions efficiently, the key is to identify the affected region in the graph precisely. We provide a theoretical analysis to define the affected scope in Section 4.2, and the detailed algorithm of node-insertion in Section 4.3. However, when u has both public and private edges, we can first remove u with its public edges, and re-insert it with both public and private edges incident to u using node-insertion. This method has significant advantages outperforming edge-insertion when private edges of u are much larger than public edges of u . On other cases, edge-insertion method may perform better. Thus, we construct a classifier to determine which algorithms of node-insertion and edge-insertion should be applied in Section 4.4.

4.2 Theoretical Analysis

In this section, we present useful rules for truss-index updating with node insertion/deletion. Consider a vertex v and the set of edges incident to v as $E(v)$. In the case of node insertion, we insert a new vertex v and its incident edges $E(v)$ into G , where $E(v) \cap E(G) = \emptyset$; in the case of node deletion, we delete vertex v and all its incident edges $E(v)$ from G , where $E(v) \subseteq E(G)$. We use $\tau(e)$ and $\hat{\tau}(e)$ to denote trussness of edge e before and after updating operation. Motivated by [Huang *et al.*, 2014], the following three updating rules hold.

Rule 1: If new node v is inserted into graph G with $\hat{\tau}(v) = \max_{e \in E(v)} \hat{\tau}(e) = l$, then $\forall e \in E(G)$ with $\tau(e) \geq l$, $\hat{\tau}(e) = \tau(e)$ holds.

Rule 2: If node v is deleted from graph G with $\tau(v) = \max_{e \in E(v)} \tau(e) = l$, then $\forall e \in E(G) \setminus E(v)$ with $\tau(e) > l$, $\hat{\tau}(e) = \tau(e)$ holds.

Rule 3: $\forall e \in E(G) \setminus E(v)$, $|\hat{\tau}(e) - \tau(e)| \leq 1$ holds.

Rules 1 and 2 hold because node v is not present at any $(l+1)$ -truss subgraph. Rule 3 holds because for each edge, at most one triangle will be formed/deformed after one node insertion/deletion, hence $\sup(e)$ will change at most by one.

In the following, we focus on node insertions. In order to apply Rule 1 for pruning, the value of $\hat{\tau}(v)$ is required. However, an exact computation of $\hat{\tau}(v)$ is costly expensive. Instead, we develop another rule based on an upper bound of $\hat{\tau}(v)$ below.

Rule 1': If node v is inserted into graph G and $\overline{\hat{\tau}(v)} \geq \hat{\tau}(v)$, then $\forall e \in E(G)$ with $\tau(e) \geq \overline{\hat{\tau}(v)}$, $\hat{\tau}(e) = \tau(e)$ holds.

To desire an upper bound of $\hat{\tau}(e)$, we need a new definition of (k, d) -neighborhood as follows.

Definition 5 ((k, d) -neighborhood). *Given a graph G , (k, d) -neighborhood of vertex v , denoted by $G_v^{k,d}$, is the maximal subgraph $H \subseteq G[N(v)]$ holding*

1. $\tau_G(e) \geq k, \forall e \in E(H)$ and
2. $d_H(u) \geq d, \forall u \in V(H)$.

Lemma 1. *Consider a new node v and its incident edges $E(v) = \{(v, w) : w \in N(v)\}$ are inserted into graph G . For each new edge $e = (v, w)$ in the new graph G_{new} , the trussness of e , $\hat{\tau}(e)$, satisfies $k_{low}(e) \leq \hat{\tau}(e) \leq k_{up}(e)$ where*

$$k_{low}(e) = \max\{k : w \in G_v^{k, k-2}\}$$

and

$$k_{up}(e) = \max\{k : w \in G_v^{k-1, k-2}\}.$$

Moreover, $|k_{up}(e) - k_{low}(e)| \leq 1$ holds.

Proof. We consider an edge $e^* = (v, w^*)$ in G_{new} . For simplification, we denote by $k_{low}(e^*) = k_l$ and $k_{up}(e^*) = k_u$.

First, we prove $\hat{\tau}(e^*) \geq k_{low}(e^*) = k_l$. To prove it, we show that there exists a k_l -truss H^* of G_{new} containing e^* . By the definition of $k_{low}(e)$, there exists a k_l -truss H of G , i.e., $\forall e \in E(H)$, $\sup_H(e) \geq k_l - 2$. Let $H^* = (V(H) \cup \{v\}, E(H) \cup \{(w, v) \mid w \in G_v^{k_l, k_l-2}\})$, which adds vertex v and v 's incident edges (w, v) with $k_{low}((w, v)) \geq k_l$ into H . For each edge $(w, v) \in E(H^*) \setminus E(H)$, w and v have at least $k_l - 2$ common neighbors in H^* by the second condition of Def. 5, indicating $\sup_{H^*}((w, v)) \geq k_l - 2$; moreover, for each edge $e \in E(H^*) \cap E(H)$, $\sup_{H^*}(e) \geq \sup_H(e) \geq k_l - 2$. As a result, H^* is a k_l -truss, and $\hat{\tau}(e^*) \geq \tau_{H^*}(e^*) \geq k_l$.

Second, we prove $\hat{\tau}(e^*) \leq k_{up}(e^*) = k_u$ by contradiction. Assume that $\hat{\tau}(e^*) \geq k_u + 1$, there exists a $(k_u + 1)$ -truss H containing e^* in G_{new} . We delete the node v and all its incident edges (v, w) from H , which leads to a new graph H^* . By Rule 3, the trussness of each edge e in H^* decreases by at most 1 after the node deletion of v , i.e., $\tau_{H^*}(e) \geq k_u$. Let the vertex set $S = V(H) \cap N(v)$. Obviously, $H^*[S] = H[S]$. For each edge e in $H^*[S]$, $\tau_G(e) \geq \tau_{H^*}(e) \geq k_u$; for each vertex w in $H[S]$, the edge (v, w) belongs to $(k_u + 1)$ -truss H , indicating w has at least $k_u - 1$ neighbors in H and also in $H[S]$, i.e., $d_{H^*[S]}(w) = d_{H[S]}(w) \geq k_u - 1$. By Def. 5, $H^*[S]$ is a $(k_u, k_u - 1)$ -neighborhood as $G_v^{k_u, k_u-1}$ in G . Thus, $\max\{k : w \in G_v^{k-1, k-2}\} \geq k_u + 1$. However, by the definition of $k_{up}(e)$, $k_u = \max\{k : w \in G_v^{k-1, k-2}\}$ and $k_u \geq k_u + 1$, which is a contradiction.

Algorithm 1 Node-Insertion Updating Algorithm

Input: $G = (V, E)$, new node v , edge set $E(v)$
Output: $\hat{\tau}(e)$ for each $e \in E \cup E(v)$

- 1: $G \leftarrow G \cup (v, E(v))$
- 2: Compute $k_{\text{low}}(e), k_{\text{up}}(e)$ for all $e \in E(v)$ by Algorithm 2
- 3: **for** e in $E(v)$ **do**
- 4: $\tau(e) \leftarrow k_{\text{low}}(e)$
- 5: **if** $k_{\text{low}}(e) < k_{\text{up}}(e)$ **then**
- 6: $L_{k_{\text{low}}(e)} \leftarrow L_{k_{\text{low}}(e)} \cup \{e\}$
- 7: **for** $e = (u, w)$ in $G_{N(v)}$ **do**
- 8: **if** $\tau(e) < \min\{k_{\text{up}}((u, v)), k_{\text{up}}((w, v))\}$ **then**
- 9: $L_{\tau(e)} \leftarrow L_{\tau(e)} \cup \{e\}$
- 10: $k_{\text{max}} \leftarrow \max\{k_{\text{up}}(e) : e \in E(v)\}$ **then**
- 11: **for** $k \leftarrow k_{\text{max}} - 1$ to 2 **do**
- 12: UpdateTrussness(k, L_k)

Third, we prove $|k_u - k_l| \leq 1$. Obviously, $k_l \leq k_u$ and $k_l \leq k_u + 1$. We next prove $k_u \leq k_l + 1$. According to the definition of $k_{\text{up}}(e)$, we have $k_u = \max\{k : w \in G_v^{k-1, k-2}\} \leq \max\{k : w \in G_v^{k-1, k-3}\}$. Moreover, $\max\{k : w \in G_v^{k-1, k-3}\} = \max\{k : w \in G_v^{k, k-2}\} + 1 = k_l + 1$ by the definition of $k_{\text{low}}(e)$, and we derive $k_u \leq k_l + 1$. As a result, $|k_u - k_l| \leq 1$. \square

Scope of Affected Edges. Let $\overline{\hat{\tau}(v)} = \max\{k_{\text{up}}(e) : e \in E(v)\}$ be an upper bound of $\hat{\tau}(v)$, and the weight of a triangle be the minimum trussness of edges within this triangle.

1. **Node Insertion.** Edge $e = (x, y) \in E(G) \cup E(v)$ with $\tau(e) < \overline{\hat{\tau}(v)}$, may have trussness increment if (v, x, y) form a triangle of weight $\tau(e)$, or e is connected to v through a series number of adjacent triangles each with weight of $\tau(e)$.
2. **Node Deletion.** Edge $e = (x, y) \in E(G) - E(v)$ with $\tau(e) \leq \max\{\tau(e) : e \in E(v)\}$ may have trussness decrement if (v, x, y) form a triangle of weight $\tau(e)$ or e is connected to v through a series number of adjacent triangles each with weight of $\tau(e)$.

4.3 Node-Insertion Updating Algorithm

We propose Algorithm 1 to update the truss-index with inserting node v and its incident edges $E(v)$ to G . Lower and upper bounds of each edge can be computed by calling Algorithm 2 (line 2). According to the scope of affected edges, first, trussness of newly added edges is set to $k_{\text{low}}(e)$, and then candidate edges for updating are found through lines 3-10. Newly added edges with $k_{\text{low}}(e) + 1 = k_{\text{up}}(e)$ might have trussness increase which are found in lines 3-6. Furthermore, any edge in $G[N(v)]$ that might get affected is found through lines 7-9. According to Rule 1', k_{max} is set to maximum of the upper bounds which results in pruning all unaffected edges which have trussness of at least k_{max} . The procedure of level-by-level updating truss-index (line 12) follows the edge-insertion algorithm [Huang *et al.*, 2014].

Algorithm 2 Node-Insertion Bound Computing Algorithm

Input: $G = (V, E)$, new node v , edge set $E(v)$,
 $type \in \{\text{low}, \text{up}\}$
Output: $\{k_{\text{type}}(e) : e \in E(v)\}$ trussness bound according to $type$

- 1: $H \leftarrow G[N(v)]; k \leftarrow 2$
- 2: **while** $H \neq \emptyset$ **do**
- 3: **while** $\exists e \in E(H)$ with $\tau_H(e) < k$ **do**
- 4: Delete edge e from H ;
- 5: **while** $\exists d_H(u) < (k - 2)$ **if** $type = \text{low}$ **else** $k - 1$ **do**
- 6: Delete vertex u and its incident edges from H ;
- 7: $k_{\text{type}}((v, u)) \leftarrow (k - 1)$ **if** $type = \text{low}$ **else** k
- 8: $k \leftarrow k + 1$

Computing $k_{\text{up}}(e)$ and $k_{\text{low}}(e)$. Algorithm 2 computes the upper bound $k_{\text{up}}(e)$ and lower bound $k_{\text{low}}(e)$ in Lemma 1. Computing lower bounds and upper bounds are almost the same, and the same codes can be used with passing a parameter $type \in \{\text{low}, \text{up}\}$ to indicate which bound to compute. Consider the case where $type$ is low and lower bound is required. It starts with an induced subgraph $G[N(v)]$ as H , which is $G_v^{2,0}$. In each iteration refines H to reach $G_v^{k, k-2}$ from $G_v^{k-1, k-3}$. Assume the $(k - 1)$ -th step has correctly stored $G_v^{k-1, k-3}$ in H . It first removes edges with trussness less than k (lines 3-4), and then removes vertices with degree less than $k - 2$ (lines 5-7), until the resulting graph H becomes $G_v^{k, k-2}$. Each node removed in this iteration is member of $G_v^{k-1, k-3}$ but not $G_v^{k, k-2}$; therefore the bound $k - 1$ finally found by the algorithm is the maximum possible value.

4.4 A Classification-based Hybrid Algorithm for Finding K-Truss in Public-Private Graphs

We present a classification-based algorithm of updating truss-index from public graph G to pp-graph $G \cup G_u$. The algorithm is outlined in Algorithm 3, which uses a hybrid strategy of updating with node/edge insertions/deletions. Specifically, we have two following strategies to update truss-index.

- **Edge-PP.** Add private edges of $E(G_u)$ one by one into G using edge-insertion algorithm [Huang *et al.*, 2014].
- **Vertex-PP.** Remove vertex u with its public edges by node-deletion algorithm, then add back u with all incident edges to obtain $G \cup G_u$ using node-insertion algorithm in Algorithm 1.




Figure 2: Win case distribution of Vertex-PP and Edge-PP for 1367 random query nodes in PP-DBLP-2013.

Algorithm 3 Hybrid-PP Algorithm

Pre-process:
Input: $G = (V, E)$, truss-index $\{\tau(e) : \forall e \in E\}$
Output: Classification model $C : D \rightarrow \{C_V, C_E\}$

- 1: Training vertex set $S \leftarrow \text{sample_nodes}(G)$
- 2: **for** v in S **do**
- 3: $T_V(v), T_E(v) \leftarrow$ Runtime of Vertex-PP and Edge-PP on v
- 4: $\mathbf{X} = [\text{feature}(v) : v \in S]$
- 5: $\mathbf{Y} = [C_V \text{ if } T_V(v) < T_E(v) \text{ else } C_E : v \in S]$
- 6: $C \leftarrow$ Classifier-Construction(\mathbf{X}, \mathbf{Y})

Query:
Input: query node u , private graph G_u , and integer k
Output: the k -truss in $G \cup G_u$

- 1: **if** $C.\text{predict}(\text{feature}(u)) = C_V$ **then**
- 2: Update index using Vertex-PP (u)
- 3: **else** Update index using Edge-PP (u)
- 4: **return** Query_KTruss(u, k) on updated index

Note that Vertex-PP can not directly add u 's private edges into G , as Lemma 1 holds only for the insertion of a completely new vertex. Both algorithms update the truss-index correctly; however, the optimal choice between Edge-PP and Vertex-PP is not straightforwardly clear, because multiple aspects affect efficiency. Determining efficiency performance requires a global knowledge of the whole graph structure. Local topological properties is not sufficient to decide which algorithm works faster. For example, Figure 2 shows the distribution of cases in which Vertex-PP or Edge-PP perform at least two times faster than the other algorithm; it consists of 1367 randomly selected nodes w.r.t their public and private degrees sampled from PP-DBLP-2013 dataset [Huang *et al.*, 2018]. It is clear by Figure 2 that the simple distinction of degree is not sufficient to make good decision upon which algorithm to use. Thus, we formulate and tackle the problem of using Edge-PP and Vertex-PP as a classification task.

Algorithm 3 predicts which updating method between Vertex-PP and Edge-PP works faster in terms of the given query, and runs that algorithm to answer the query. There are two classes C_V and C_E that each node $u \in V$ is in class C_V if Vertex-PP works faster than Edge-PP to update index from public graph G to $G \cup G_u$ and the similar for C_E . We present each node u using several features including: 1) public degree of u ; 2) private degree of u ; 3) the number of triangles containing u respectively in public graph G , private graph G_u , and pp-graph g_u ; and 4) the trussness sum of public edges;

5 Experiments

Datasets: We used four public-private graphs of PP-DBLP [Huang *et al.*, 2018] in Table 1.¹ Published articles make the public network, and ongoing collaborations form the private networks which are only known by partial authors. We also used eight real-world graphs available from SNAP [Leskovec

¹<https://github.com/samjxx/pp-data>

Name	$ V $	$ E $	$ V_{private} $	$ E_{private} $
PP-DBLP-2013	1,791,688	5,187,025	804,121	3,166,863
PP-DBLP-2014	1,791,688	5,893,083	669,138	2,491,847
PP-DBLP-2015	1,791,688	6,605,428	502,654	1,719,794
PP-DBLP-2016	1,791,688	7,378,090	257,129	719,204

Table 1: Network Statistics of Real-world Public-Private Graphs

Name	$ V $	$ E $	$ V_t $	Avg. time per node (s)		Speedup
				Node-ins.	Edge-ins.	
EmailEuAll	265K	420K	292	0.38	4.20	10.96
Wikivote	7K	103K	534	0.73	28.53	39.10
EmailEnron	36K	367K	485	0.83	24.10	29.03
Gowalla	196K	1.9M	156	2.99	102.39	34.27
WikiTalk	2.4M	5M	191	40.54	2017.26	49.76
Flickr	80K	11.8M	697	96.30	2195.57	22.80
Digg	771K	7.3M	471	142.82	5741.43	40.20
LiveJournal	4M	34.7M	377	142.63	5515.66	38.67

 Table 2: Comparison of Node-insertion and Edge-insertion methods in terms of efficiency, by inserting V_t randomly selected nodes on real-world graphs. Here $\mathbf{K} = 10^3$ and $\mathbf{M} = 10^6$.

and Krevl, 2014] shown in Table 2.

Compared Methods and Evaluated Metrics: To evaluate the efficiency of improved strategies proposed in this paper, we tested and compared four algorithms as follows.

- Truss-Decomposition: an online search approach using truss decomposition for computing k -truss index from scratch [Wang and Cheng, 2012].
- Edge-PP: an approach using the edge-insertion algorithm for updating truss index [Huang *et al.*, 2014].
- Vertex-PP: our approach using the node-deletion and node-insertion for updating truss-index in Algorithm 1.
- Hybrid-PP: our hybrid approach using both Edge-PP and Vertex-PP for updating truss-index in Algorithm 3.

After updating the index, all of the four methods mentioned above use the same query method. We compare them by reporting the running time in seconds. The less the running time is, the better the efficiency performance is. We set the parameter $k = 7$ by default. We also evaluate the methods by varying parameters k in $\{5, 7, 9, 11, 13, 15\}$.

5.1 Efficiency Evaluation on SNAP Networks

To evaluate the efficiency of the node-insertion algorithm, we conducted experiments on eight SNAP graph datasets in Table 2. Due to no available private information in these networks, we randomly generated private edges as follows. We divided nodes into 40 bins by their degree and took 50 randomly selected nodes from each bin. Bin set was defined as $\{B_1, B_2, \dots, B_{40}\}$ where $B_i = \{v : \frac{i-1}{40} < \frac{d(v)}{\Delta} \leq \frac{i}{40}\}$ with $\Delta = \max_{v \in V} \{d(v)\}$. Let $G = (V, E)$ denote the initial graph, and V_t the set of sampled vertices. All edges with at least one end in V_t was considered to be private. We ran Truss-Decomposition on the induced subgraph of vertex set $V' = V \setminus V_t$, then added each node $v \in V_t$ using Algorithm 1, and compared the running time with adding edges one by one using edge-insertion algorithm. In Table 2, our node-insertion algorithm significantly outperformed the edge-insertion method, which achieved 38.67 times speedup

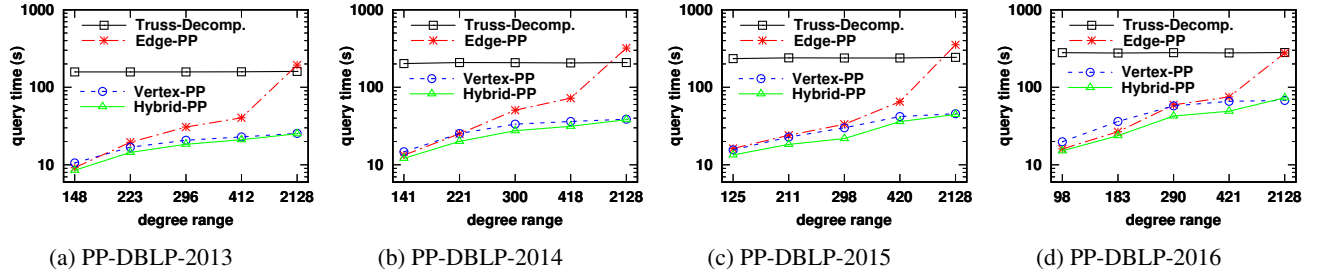
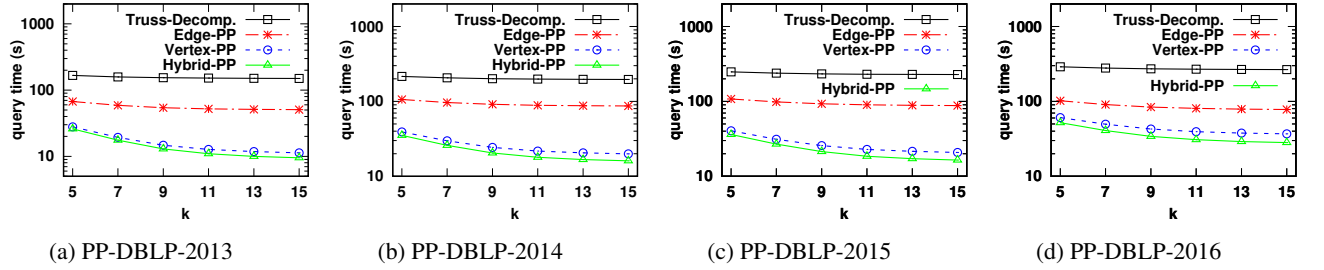


Figure 3: Average query time of different methods on PP-DBLP varied by query node degree.


 Figure 4: Average query time of different methods on PP-DBLP varied by parameter k .

on LiveJournal.

5.2 Classification Evaluation

In order to choose the proper classification method to incorporate into Hybrid-PP, we tested and compared five classifiers in terms of classification accuracy and training time, as shown in Table 3. Due to the low training time, high classification accuracy, and fast query time, Random Forest was finally used as the classifier of Hybrid-PP.

Classifier	Accuracy on PP-DBLP				Training time (s)
	2013	2014	2015	2016	
Random Forest	84.4%	87.0%	86.5%	87.8%	<1
Decision Tree	84.3%	84.6%	84.7%	87.4%	<1
SVM	84.1%	84.5%	85.5%	85.2%	1140
k-NN	76.0%	79.6%	78.2%	72.7%	<1
Degree Baseline	69.5%	69.2%	71.8%	67.2%	<1

Table 3: Accuracy of different classifiers on PP-DBLP 2013-2016.

5.3 Efficiency Evaluation on PP-DBLP Networks

We compare the efficiency of four different methods Truss-Decomposition, Edge-PP, Vertex-PP, and Hybrid-PP on real-world public-private datasets of PP-DBLP. Hybrid-PP adopted a Random Forest with 51 estimators and a maximum depth of 11 to construct a classifier. The Hybrid-PP used to answer queries on each PP-DBLP dataset was trained based on runtimes of Vertex-PP and Edge-PP on sampled nodes from the other three datasets. We first divided all nodes into 100×100 bins by their public and private degrees, and then randomly took four nodes from each bin. Bin set was $\{B_{i,j} : 1 \leq i, j \leq 100\}$ each bin defined as $B_{i,j} = \{v : \frac{i-1}{100} < \frac{d(v)}{\Delta} \leq \frac{i}{100}, \frac{j-1}{100} < \frac{d_p(v)}{\Delta_p} \leq \frac{j}{100}\}$ where

Δ is the maximum public degree, Δ_p is the maximum private degree, and $d_p(v)$ is the private degree of node v . In total 1836 ± 58 nodes were selected on each PP-DBLP datasets.

Vary Node Degree. We compared four proposed algorithms on sampled nodes. For better visualization and comparison, sampled nodes were divided into five equally sized groups by node degrees in their pp-graphs, each group taking 20% of sampled nodes, and the average query time of algorithms on each group is reported in Figure 3. On higher degree nodes, Edge-PP takes much longer than Vertex-PP, and becomes less useful for Hybrid-PP; hence the gap between Hybrid-PP and Vertex-PP decreases as Vertex-PP becomes the optimal choice for the classifier on higher degree nodes.

Vary Parameter k . The average query time varied by k is reported in Figure 4. Hybrid-PP and Vertex-PP perform much better than Edge-PP and Truss-Decomposition for all k values and datasets. Hybrid-PP is the fastest due to using Edge-PP in cases where Vertex-PP works worse than Edge-PP.

6 Conclusions

This paper studies the problem of finding k -truss on public-private graphs. We develop a novel hybrid algorithm of k -truss updating with node/edge insertions/deletions. This work opens up several interesting problems. First, developing further efficient and clever algorithms for finding k -truss in pp-graphs is important, instead of deleting and re-inserting nodes as Vertex-PP. Second, finding other kinds of dense subgraphs on public-private graph is also wide open.

Acknowledgments

This work is supported by the NSFC 61702435 and RGC 12200917. Xin Huang is the corresponding author.

References

- [Archer *et al.*, 2017] Aaron Archer, Silvio Lattanzi, Peter Likarish, and Sergei Vassilvitskii. Indexing public-private graphs. In *WWW*, pages 1461–1470, 2017.
- [Bansal *et al.*, 2004] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine learning*, 56(1-3):89–113, 2004.
- [Chierichetti *et al.*, 2015] Flavio Chierichetti, Alessandro Epasto, Ravi Kumar, Silvio Lattanzi, and Vahab Mirrokni. Efficient algorithms for public-private social networks. In *KDD*, pages 139–148, 2015.
- [Cohen and Kaplan, 2007] Edith Cohen and Haim Kaplan. Summarizing data using bottom-k sketches. In *PODS*, pages 225–234, 2007.
- [Cohen, 2008] J. Cohen. Trusses: Cohesive subgraphs for social network analysis. Technical report, National Security Agency, 2008.
- [Das Sarma *et al.*, 2010] Atish Das Sarma, Sreenivas Gollapudi, Marc Najork, and Rina Panigrahy. A sketch-based distance oracle for web-scale graphs. In *WSDM*, pages 401–410, 2010.
- [Dey *et al.*, 2012] Ratan Dey, Zubin Jelveh, and Keith Ross. Facebook users have become much more private: A large-scale study. In *IEEE International Conference on PER-COM Workshops*, pages 346–352, 2012.
- [Gregori *et al.*, 2011] Enrico Gregori, Luciano Lenzini, and Chiara Orsini. k-dense communities in the internet as-level topology. In *International Conference on Communication Systems and Networks*, pages 1–10, 2011.
- [Haveliwala, 2002] Taher H Haveliwala. Topic-sensitive pagerank. In *WWW*, pages 517–526, 2002.
- [Huang *et al.*, 2014] Xin Huang, Hong Cheng, Lu Qin, Wentao Tian, and Jeffrey Xu Yu. Querying k-truss community in large and dynamic graphs. In *SIGMOD*, pages 1311–1322, 2014.
- [Huang *et al.*, 2016] Xin Huang, Wei Lu, and Laks VS Lakshmanan. Truss decomposition of probabilistic graphs: Semantics and algorithms. In *SIGMOD*, pages 77–90, 2016.
- [Huang *et al.*, 2018] Xin Huang, Jiabin Jiang, Byron Choi, Jianliang Xu, Zhiwei Zhang, and Yunya Song. PP-DBLP: modeling and generating attributed public-private networks with dblp. In *IEEE International Conference on Data Mining Workshops*, pages 986–989, 2018.
- [Jiang *et al.*, 2018] Yuli Jiang, Xin Huang, Hong Cheng, and Jeffrey Xu Yu. Vizcs: Online searching and visualizing communities in dynamic graphs. In *ICDE*, pages 1585–1588, 2018.
- [Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146, 2003.
- [Leskovec and Krevl, 2014] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [Li *et al.*, 2018] Ye Li, Chaofeng Sha, Xin Huang, and Yanchun Zhang. Community detection in attributed graphs: an embedding approach. In *AAAI*, 2018.
- [Mirzasoleiman *et al.*, 2016] Baharan Mirzasoleiman, Morteza Zadimoghaddam, and Amin Karbasi. Fast distributed submodular cover: Public-private data summarization. In *NIPS*, pages 3594–3602, 2016.
- [Mokken, 1979] Robert J Mokken. Cliques, clubs and clans. *Quality & Quantity*, 13(2):161–173, 1979.
- [Pei *et al.*, 2005] Jian Pei, Daxin Jiang, and Aidong Zhang. Mining cross-graph quasi-cliques in gene expression and protein interaction data. In *ICDE*, pages 353–354, 2005.
- [Saito *et al.*, 2008] Kazumi Saito, Takeshi Yamada, and Kazuhiro Kazama. Extracting communities from complex networks by the k-dense method. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 91(11):3304–3311, 2008.
- [Takaguchi and Yoshida, 2016] Taro Takaguchi and Yuichi Yoshida. Cycle and flow trusses in directed networks. *Royal Society open science*, 3(11):160270, 2016.
- [Wang and Cheng, 2012] Jia Wang and James Cheng. Truss decomposition in massive networks. *PVLDB*, 5(9):812–823, 2012.
- [Wilder *et al.*, 2018] Bryan Wilder, Nicole Immorlica, Eric Rice, and Milind Tambe. Maximizing influence in an unknown social network. In *AAAI*, pages 4743–4750, 2018.
- [Xiao *et al.*, 2017] Mingyu Xiao, Weibo Lin, Yuanshun Dai, and Yifeng Zeng. A fast algorithm to compute maximum k-plexes in social network analysis. In *AAAI*, 2017.
- [Zhang and Parthasarathy, 2012] Yang Zhang and Srinivasan Parthasarathy. Extracting analyzing and visualizing triangle k-core motifs within networks. In *ICDE*, pages 1049–1060, 2012.
- [Zhang *et al.*, 2017] Fan Zhang, Ying Zhang, Lu Qin, Wenjie Zhang, and Xuemin Lin. Finding critical users for social network engagement: The collapsed k-core problem. In *AAAI*, 2017.
- [Zhang *et al.*, 2018a] Fan Zhang, Conggai Li, Ying Zhang, Lu Qin, and Wenjie Zhang. Finding critical users in social communities: The collapsed core and truss problems. *TKDE*, 2018.
- [Zhang *et al.*, 2018b] Ge Zhang, Di Jin, Jian Gao, Pengfei Jiao, Françoise Fogelman-Soulié, and Xin Huang. Finding communities with hierarchical semantics by distinguishing general and specialized topics. In *IJCAI*, pages 3648–3654, 2018.
- [Zhao and Tung, 2012] Feng Zhao and Anthony KH Tung. Large scale cohesive subgraphs discovery for social network visual analysis. In *PVLDB*, volume 6, pages 85–96, 2012.
- [Zou and Zhu, 2017] Zhaonian Zou and Rong Zhu. Truss decomposition of uncertain graphs. *Knowledge and Information Systems*, 50(1):197–230, 2017.