Privacy Preserving Social Network Publication Against Mutual Friend Attacks

Chongjing Sun†, Philip S. Yu‡, Xiangnan Kong‡ and Yan Fu†
†*Web Science Center, School of Computer Science and Engineering*
University of Electronic Science and Technology of China, Chengdu, China, 611731
‡*Department of Computer Science, University of Illinois at Chicago, Chicago, IL 60612*
Email:chingsun00@gmail.com, psyu@cs.uic.edu, xkong4@cs.uic.edu, fuyan@uestc.edu.cn

***Abstract*—Publishing social network data for research purposes has raised serious concerns for individual privacy. There exist many privacy-preserving works that can deal with different attack models. In this paper, we introduce a novel privacy attack model and refer it as a mutual friend attack. In this model, the adversary can re-identify a pair of friends by using their number of mutual friends. To address this issue, we propose a new anonymity concept, called *k*-NMF anonymity, i.e., *k*-anonymity on the number of mutual friends, which ensures that there exist at least *k*-1 other friend pairs in the graph that share the same number of mutual friends. We devise algorithms to achieve the *k*-NMF anonymity while preserving the original vertex set in the sense that we allow the occasional addition but no deletion of vertices. Further we give an algorithm to ensure the *k*-degree anonymity in addition to the *k*-NMF anonymity. The experimental results on real-word datasets demonstrate that our approach can preserve the privacy and utility of social networks effectively against mutual friend attacks.**

***Keywords*-privacy-preserving; social network; mutual friend**

I. Introduction

With the advance on mobile and Internet technology, more and more information is recorded by social network applications, such as Facebook and Twitter. The relationship information in social networks attracts researchers from different academic fields. As a consequence, more and more social network datasets were published for research purposes [1]. The published social network datasets may incur the privacy invasion of some individuals or groups. With the increasing concerns on the privacy, many works have been proposed for the privacy-preserving social network publication [2], [3].

Tai and Yu proposed the friendship attack model [4], which addressed the issue that an attacker can find out not only the degree of a person, but also the degree of his friend. It solves the attacks based on the degrees of two connected vertices. But it is not sufficient to just protect against the friendship attack as there are more information available on the social network. For example, the graph in Fig. 2(a) is a k^2-degree anonymized graph with $k = 2$. If an attacker can obtain the number of mutual friends between two connected vertices, he still can identify (D, F) from other friend pairs, as only (D, F) has 2 mutual friends. This will be explained in more details later. In most social networking sites, such as Facebook, Twitter, and LinkedIn, the adversary can easily get the number of mutual friends of two individuals linked by a relationship. As shown in Figure 1, one can directly see mutual friend list shared with one of his friends on Facebook. Usually, the adversary can get the friend lists of two individuals from Facebook, such as the friend list in Figure 1, and then get the number of mutual friends by intersecting their friend lists.

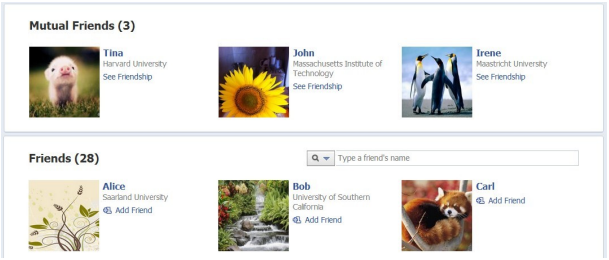

Figure 1: Friend lists on Facebook

In this paper, we introduce a new relationship attack model based on the number of mutual friends of two connected individuals, and refer it as a *mutual friend attack*. Figure 2 shows an example of the mutual friend attack. The original social network *G* with vertex identities is shown in Figure 2(b), and can be naively anonymized as the network *G'* shown in Figure 2(c) by removing all individuals' names. The *number* on each edge in *G'* represents the number of mutual friends of the two end vertices. Alice and Bob are friends, and their mutual friends are Carl, Dell, Ed and Frank. So the number of mutual friends of Alice and Bob is 4. After obtaining this information, the adversary can uniquely re-identify the edge (D, E) is $(Alice, Bob)$. Also, $(Alice, Carl)$ can be uniquely re-identified in *G'*. By combining $(Alice, Bob)$ and $(Alice, Carl)$, the adversary can uniquely re-identify individuals Alice, Bob and Carl. This simple example illustrates that it is possible for the

Yan Fu is the corresponding author.

This research work was supported in part by the National Natural Science Foundation of China under Grant No.61003231 and No.60973120, the research funds for central universities under grant No. ZYGX2012J085, and US NSF through grants CNS-1115234, DBI-0960443, and OISE-1129076.

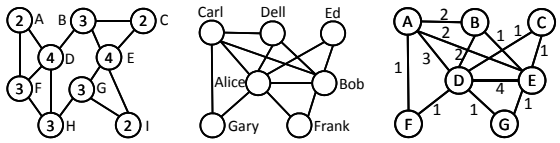

(a) (b) (c)

Figure 2: Mutual friend attack in a social network

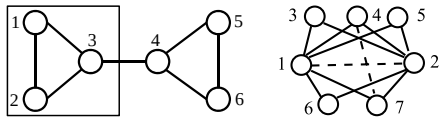

(a) (b)

Figure 3: Examples of the k-NMF anonymization

adversary to re-identify an edge between two individuals and maybe indeed identify the individuals when he can get the number of mutual friends of individuals. Note that we do not consider the mutual friend number of two nodes if they are not connected. For convenience, we say the *number of mutual friends of two nodes* connected by an edge e as the *number of mutual friends of* e.

In order to protect the privacy of relationship from the mutual friend attack, we introduce a new privacy-preserving model, *k*-anonymity on the number of mutual friends (*k*-NMF Anonymity). For each edge *e*, there will be at least *k*-1 other edges with the same number of mutual friends as *e*. It can be guaranteed that the probability of an edge being identified is not greater than 1/*k*. We propose algorithms to achieve the *k*-NMF anonymity for the original graph while preserving the original vertex set in the sense that we allow the occasional addition but no deletion of vertices. By preserving the original vertex set, various analysis on the anonymized graph, such as identifying vertices providing specific roles like centrality vertex, influential vertex, gateway vertex, outlier vertex, etc., will be more meaningful. The experimental results on real datasets show that our approaches can preserve much of the utility of social networks against mutual friend attacks.

Related Work. Backstorm et al. [5] pointed out that simply removing identities of vertices cannot guarantee privacy. Many works have been done to prevent the vertex re-identification with the vertex degree. Liu et al. [6] studied the *k*-degree anonymization which ensures that for any node *v* there exist at least *k*-1 other vertices in the published graph with the same degree as *v*. Tai et al. [4] introduced a friendship attack, in which the adversary uses the degrees of two end vertices of an edge to re-identify victims. Associated with community identity for each vertex, in [7] they proposed the k-structural diversity anonymization, which guarantees the existence of at least *k* communities containing vertices with the same degree for each vertex. As these works only focus on the vertex degree, they cannot achieve the *k*-NMF anonymity, which focuses on the number of common neighbors of two vertices.

Many works have also been done to prevent the vertex re-identification based on the subgraph structural information. Zhou and Pei [8] proposed a solution to battle the adversary's 1-neighborhood attacks. Cheng et al. [9] proposed the *k*-isomorphism model, which disconnects the original graph into k-isomorphic subgraph. To protect against multiple structural attacks, Zou et al. [10] proposed the *k*-automorphism model, which converts the original network into a k-automorphic network. But it does not prevent the mutual friend attack. The network in Figure 3(a) satisfies the 2-automorphism, but the edge $(3, 4)$ is not protected under the mutual friend attack. This is because the edge $(3, 4)$ does not have mutual friends while all the others have one. Wu et al. [11] proposed the k-symmetry model, which gets a k-automorphic network by orbit copying. All these algorithms need to introduce many new vertices and adjust many edges to achieve their targets. Therefore, the utility of the original graph will be decreased too much. In any case, these works are aimed at different types of attack model from ours as illustrated in Figure 3(a).

Hay et al. [12] proposed a generalizing method for anonymizing a graph, which partitions the vertices and summarizes the graph at the partition level. Other works focus on the problem of link disclosure, which decides whether there exists a link between two individuals. It is different from the relationship re-identification introduced in this section.

Challenges. As the k-NMF anonymity model is more complicated than the k-degree anonymity model, more challenges need to be handled. First, adding or removing a different edge may affect a different number of edges on their mutual friends. In the k-degree anonymity model, the adversary attacks using the degree of the vertex. Adding an edge only increase the degrees of the two end vertices of this edge. In the k-NMF anonymity model, the adversary attacks using the number of mutual friends. Adding an edge can increase the numbers of mutual friends of many edges. In Figure 2(b), adding an edge between Dell and Frank will affect the NMFs of (Dell, Alice), (Frank, Alice), (Dell, Bob), (Frank, Bob), and (Dell, Frank). Second, we need to provide a criterion on choosing where to add or delete the edge while considering the utility of the graph. Since we aim to preserve the vertex set, we cannot add a vertex to connect an edge. In fact, we map the k-NMF anonymization problem into an edge anonymization problem in contrast to the vertex anonymization problem in the k-degree anonymization. Edges are anonymized one by one. Adding or deleting an edge should not destroy the anonymization of the already anonymized edges. To anonymize an edge, we can get many candidate edge operations and need to choose the best one. Besides, we need to consider the impact of the newly added edges on the number of mutual friends.

Contributions. Our contributions can be summarized as

follows. (1) We introduce the k-NMF problem and formulate it as an edge weight anonymization problem where the edge weight is the NMF of the two end vertices. (2) We explore the geometry property of the graph to devise effective anonymization algorithms while preserving the vertex set to achieve better utility. (3) For the edge addition, we use the breadth-first manner to preserve utility. We also introduce the maximum mutual friend criterion to break the tie on selecting candidate vertex to connect. (4) For the edge deletion, we explore the triangle linking property to delete edges between vertices already belonging to a triangle connection in the network to avoid repeated re-anonymization of edges. (5) We devise an algorithm which can anonymize the k-NMF anonymized graph to simultaneously satisfy the k-degree anonymity, while preserving the vertex set. (6) The empirical results on real datasets show that our algorithms perform well in anonymizing the real social networks.

The rest of the paper is organized as follows. We define the problem and design algorithms to solve it in section 2 and 3. We conduct the experiments on real data sets and conclude in Section 4 and 5.

II. Problem definition

In this paper, we model a social network as an undirected simple graph $G(V, E)$, where V is a set of vertices representing the individuals, and $E \subseteq V \times V$ is the set of edges representing the relationship of individuals.

Definition 1. ***The NMF of an edge.*** For an edge e between two vertices v_1 and v_2 in a graph $G(V, E)$, i.e., $v_1, v_2 \in V$, $e \in E$ and $e = (v_1, v_2)$, the number of mutual friends of the edge e is the number of mutual friends of v_1 and v_2.

Let $\boldsymbol{f}$ be the *number sequence of mutual friends* for G, in which entries are sorted in descending order, i.e., $\boldsymbol{f}_1 \geq \boldsymbol{f}_2 \geq ... \geq \boldsymbol{f}_m$. Let $\boldsymbol{l}$ be the list of edges corresponding to $\boldsymbol{f}$, i.e., $\boldsymbol{f}_i$ is the NMF of the edge $\boldsymbol{l}_i$. For example, in Figure 4(c), $\boldsymbol{f} = \{2, 2, 2, 2, 1, 1, 1, 1\}$, and $\boldsymbol{l} = \{(v_1, v_3), (v_2, v_3), (v_3, v_4), (v_3, v_5), (v_3, v_4), (v_3, v_5), (v_1, v_2), (v_1, v_4), (v_2, v_5), (v_4, v_5)\}$. Similar to the power law distribution of the vertex degree [13], the NMF also has the same property [14].

Property 1. Scale free distribution of NMFs [14]**.** *The NMFs of edges in the large social network often have a scale-free distribution, which means that the distribution follows a power law or at least asymptotically.*

Definition 2. ***Mutual friend attack***. Given a social network $G(V, E)$ and the anonymized network $G'(V', E')$ for publishing. For an edge $e \in E$, the adversary can get the number $\boldsymbol{f}_e$ of mutual friends of e. Mutual Friend Attack will identify all *candidate edges* $e' \in E'$ with the number $\boldsymbol{f}_{e'}$ of mutual friends as $\boldsymbol{f}_e$.

Suppose that the candidate edge set of an edge e is $E'_e = \{e' | e' \in E', \boldsymbol{f}_{e'} = \boldsymbol{f}_e\}$. An adversary re-identifies the edge e with high confidence if the number of candidate edges is too small. Hence, we set a threshold k to make sure that for each edge $e \in E$, the number of candidate edges is no less than k, i.e., $|E'_{\boldsymbol{f}_e}| \geq k$. We define the k-anonymous sequence before defining the k-NMF anonymous graph.

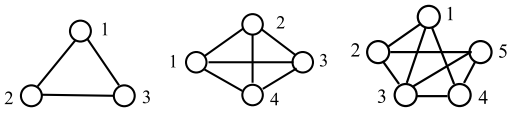

(a) 3-NMF (b) 6-NMF (c) 4-NMF

Figure 4: Examples of k-NMF anonymous graph

Definition 3. ***k-anonymous sequence***[6]. A sequence vector $\boldsymbol{f}$ is k-anonymous, if for any entry with value as v, there exist at least $k-1$ other entries with value as v.

Definition 4. ***k*-NMF**. A graph $G'(V', E')$ is k-NMF anonymous if the number sequence $\boldsymbol{f}'$ of mutual friends of edges in G' is a k-anonymous sequence.

Definition 4 states that for each edge $e \in E$, the number of candidate edges in G' is no less than k. Consider the graphs in Figure 4 as an example. There are three edges in Figure 4(a), and the NMFs of all these edges are equal to 1. Hence, this graph is a 3-NMF anonymous graph. As the six edges in the graph of Figure 4(b) have 2 mutual friends, this graph is a 6-NMF anonymous graph. The graph in Figure 4(c) has four edges $(v_1, v_3), (v_2, v_3), (v_3, v_4), (v_3, v_5)$ with the NMF as 2, and the NMFs of other four edges are equal to 1. Hence, this graph is a 4-NMF anonymous graph. Some properties on the number of mutual friends in the graph are described as follows.

Proposition 1. *Given a graph $G(V, E)$, the number of mutual friends of an edge $e \in E$ is equal to the number of triangles containing e in G .*

Take the graph in Figure 4(c) as an example. The mutual friends of vertices v_2 and v_3 are v_1 and v_5, so the number of mutual friends of the edge $e = (v_2, v_3)$ is 2. It is equal to the number of triangles containing e. These triangles are (v_1, v_2, v_3) and (v_2, v_3, v_5).

Proposition 2. *Let $G(V, E)$ be a graph and f be the number sequence of mutual friends of edges in G, where $|E| = m$. Then $\sum_{i=1}^{m} \boldsymbol{f}_i = 3n_{\triangle}$, where $n_{\triangle}$ is the number of triangles in G and $\boldsymbol{f}_i$ is the number of mutual friends of the i-th edge.*

Different from the degree sequence in previous work [6], which can maintain the number of entries in the sequence, the number sequence of mutual friends will have more entries added into it when new edges are added into the graph. Besides, according to Propositions 1 and 2, the number of mutual friends is related to the number of triangles in the graph. Therefore, adding one edge will affect the NMF of many edges, and adding a different edge may affect the NMF of a different number of edges. This can be illustrated by an example shown in Figure 3(b). After we add the edge $(1, 2)$, the NMFs of all ten edges increase by one. If we add the edge $(4, 7)$, only the NMFs of edges $(1, 4), (1, 7), (2, 4)$, and

$(2, 7)$ increase by one. Therefore, one cannot anonymize a graph by simply minimizing the number of changed edges.

Anonymized Triangle Preservation Principle (ATPP). In our algorithms, we anonymize the edges in the graph one by one. An *anonymized triangle* is a triangle with some edges already anonymized in the process of the graph anonymizing. The *Anonymized Triangle Preservation Principle* aims to preserve the anonymized triangles containing already anonymized edges. It means that we neither create some additional anonymized triangles via edge addition nor destroy any via edge deletion.

Creating (destroying) a triangle containing an already anonymized edge by edge addition (deletion) will increase (decrease) the NMF of this edge, indeed destroy the anonymization of this edge. This leads to repeatedly anonymization of this edge. By preserving the anonymized triangles, we can avoid this problem during the anonymization process.

Definition 5. *k*-NMF anonymization problem. Given a graph $G(V, E)$ and an integer k, the problem is to anonymize the graph G to a k-NMF anonymous graph G' with edge addition and deletion, such that the vertex set of the original graph G is preserved.

III. k-NMF ANONYMIZATION APPROACH

In the above section, we found that changing one edge may affect the NMFs of other edges. To handle this challenge, we utilize the scala free distribution property shown in Property 1, and introduce the principle of preserving the anonymized triangles. By exploring the geometry property of the graph, we devise two effective anonymization algorithms to preserve the utility while satisfying the k-NMF anonymity.

A. Algorithm ADD

In this subsection, we aim to anonymize the original graph only by edge addition. We organize edges into groups, and anonymize the edges in the same group to have the same NMF. The k-anonymity requires there exist at least k edges in a group. Property 1 states that the NMFs of edges in large social networks follow a scala free distribution. Hence, only a small number of edges have a high NMF. We first anonymize these edges, and many edges with low NMF do not need to be processed.

Suppose the original graph is $G(V, E)$ and the gradually anonymized graph is $G'(V', E')$. Initially, we sort the NMF sequence $\boldsymbol{f}$ in descending order and construct the corresponding edge list $\boldsymbol{l}$ as described in Section II. We mark all edges as "unanonymized", and then anonymize the edges one by one. Iteratively, we start a new group GP with the group NMF, g_f, equal to the NMF of the first unanonymized edge in $\boldsymbol{l}$. Then we select the edges with NMF equal to g_f and mark them as "anonymized". We iteratively select the first unanonymized edge in $\boldsymbol{l}$ and anonymize it by adding edges to increase its NMF to g_f. After anonymizing this edge, we mark it as "anonymized" and put it into GP. Adding new edges affects the NMF of some other edges, and these new edges will be added into $\boldsymbol{f}$ and $\boldsymbol{l}$. Hence we resort the sequences $\boldsymbol{f}$ and $\boldsymbol{l}$ after each edge is anonymized. Algorithm 1 shows the detailed description of the ADD algorithm. Next, we consider when we start another new group.

1) Group edges: An intuitive method, named **Intuit-Group**, starts another group when the number of edges in the group GP is equal to k. Alternatively, to consider the anonymization cost, we propose a greedy method to decide when we start another group after $|GP| \geq k$, named **GreedyGroup**. Suppose that $\boldsymbol{f}^{(u)} \subseteq \boldsymbol{f}$ is the NMF sequence corresponding to the unanonymized edge list $\boldsymbol{l}^{(u)} \subseteq \boldsymbol{l}$. Notice that $\boldsymbol{f}^{(u)}$ and $\boldsymbol{l}^{(u)}$ are dynamically updated with $\boldsymbol{f}$ and $\boldsymbol{l}$ after anonymizing each edge. Similar to the consideration in [6], after putting k edges into GP, GreedyGroup iteratively checks whether it should merge the edge $\boldsymbol{l}_1^{(u)}$ into GP or start another group. The decision is made according to the following two costs based on the number of added mutual friends in Eq.(1) and Eq.(2).

$$C_{merge} = (g_f - \boldsymbol{f}_1^{(u)}) + I(\boldsymbol{f}_2^{(u)}, \boldsymbol{f}_{k+1}^{(u)}) \tag{1}$$

$$C_{new} = I(\boldsymbol{f}_1^{(u)}, \boldsymbol{f}_k^{(u)}) \tag{2}$$

where $I(\boldsymbol{f}_i^{(u)}, \boldsymbol{f}_j^{(u)}) = \sum_{l=i}^{j}(\boldsymbol{f}_i^{(u)} - \boldsymbol{f}_l^{(u)})$.

For Eq.(1), we put $\boldsymbol{l}_1^{(u)}$ into GP. $\boldsymbol{l}_1^{(u)}$ has $\boldsymbol{f}_1^{(u)}$ mutual friends, so we need to add $g_f - \boldsymbol{f}_1^{(u)}$ mutual friends for anonymizing $\boldsymbol{l}_1^{(u)}$. To satisfy k-anonymity, we need to put at least k edges into a new group GP'. Hence we put edges $\boldsymbol{l}_2^{(u)}, ..., \boldsymbol{l}_{k+1}^{(u)}$ into GP'. As we only adding edges, the group NMF of GP' is the maximum NMF among $\boldsymbol{f}_2^{(u)}, ..., \boldsymbol{f}_{k+1}^{(u)}$, i.e., $\boldsymbol{f}_2^{(u)}$. To anonymize $\boldsymbol{l}_i^{(u)}$, $\boldsymbol{f}_2^{(u)} - \boldsymbol{f}_i^{(u)}$ mutual friends need to be added. For Eq.(2), we put $\boldsymbol{l}_1^{(u)}, ..., \boldsymbol{l}_k^{(u)}$ into a new group GP', and the group NMF of GP is $\boldsymbol{f}_1^{(u)}$. Hence C_{merge} is the cost for anonymizing $k+1$ edges while C_{new} is for k edges. So if C_{merge} is less than C_{new}, we anonymize $\boldsymbol{l}_1^{(u)}$ and merge it into GP, and check the next unanonymized edge. Otherwise we start another new group with $\boldsymbol{l}_1^{(u)}$.

For each edge e, GreedyGroup looks ahead at $O(k)$ other edges to decide whether merging e with this group or starting a new group. Therefore, the time complexity of GreedyGroup is $O(k|E|)$.

2) Cleanup-operation: In each iteration of the ADD algorithm, it checks the number of unanonymized edges, n_u. If $n_u < 2k$, the ramaining edges are put into a group; and if $n_u < k$, $k - n_u$ edges needed to be added following the ATPP, so these k edges can form a group. New vertices will be added into the graph if the ATPP cannot be satisfied.

Next, we anonymize the edges E_u in this group. Usu-

Algorithm 1 The ADD Algorithm (GreedyGroup)

Input: Original graph $G(V, E)$, k
Output: k-NMF anonymized graph $G'(V', E')$
Initialization: $G' = G$, and mark all edges as "unanonymized". Compute and sort the sequences $\boldsymbol{f}$ and $\boldsymbol{l}$. $\boldsymbol{f}^{(u)} = \boldsymbol{f}, \boldsymbol{l}^{(u)} = \boldsymbol{l}, G_f = \emptyset$

```
while l^(u) ≠ ∅ do
    if |l^(u)| < 2k then do cleanup-operation and break.
    GP = {e|e ∈ l^(u) and f_e^(u) = f_1^(u)}; g_f = f_1^(u); G_f = G_f ∪ g_f.
    Mark any e ∈ GP as "anonymized"; update f^(u) and l^(u).
    while |GP| < k or (|GP| ≥ k and C_merge ≤ C_new) do
        Anoymize l_1^(u) by BFSEA. GP = GP ∪ l_1^(u), update l^(u) and f^(u).
    end while
end while
return G'(V', E').
```

ally, we set the group NMF as the largest NMF among unanonymized edges, denoted as g_f. Then we sum the difference as $sd = \sum_{e \in E_u}(g_f - \boldsymbol{f}_e)$, where $\boldsymbol{f}_e$ is the number of mutual friends of the edge e. If $sd >= k/2$, then we add sd nodes and $2 \cdot sd$ edges into the graph. That means that for each unanonymized edge e, we add $g_f - \boldsymbol{f}_e$ vertices and link them with the two end vertices of e. As all the newly added $(2 \cdot sd >= k)$ edges have only one mutual friend, they can form a new group. Then we mark the new edges as "anonymized" and achieve the task. If $sd < k/2$, then we enlarge the group NMF $g_f = g_f + 1$, and repeat the above process. By the *clean-up operation*, we can successfully anonymize the original network at the last step of the anonymization process.

B. BFS-based Edge Anonymization(BFSEA)

In this section, we consider how to anonymize an edge by edge addition while preserving the utility. There are three challenges to increase the NMF of an edge via adding edges. First, the added edge should not affect the NMF of already anonymized edges. Secondly, the added edge should minimize the effect on the utility of the graph. Thirdly, the NMF of the newly added edges should not disrupt the current anonymization process which is progressing in descending order of the NMF value.

Before anonymizing an edge (u, v), the ADD algorithm has created some anonymized groups and got a set G_f containing the group NMFs of these groups. Let g_f be the NMF of the current group GP, and we put g_f into G_f. Anonymizing the edge (u, v) means that we need to increase the NMF of (u, v) to the current group NMF g_f, i.e. we need to create some new triangles containing this edge. Then we try to find some candidate vertices and add new edges to create new triangles. Considering the utility of the graph, we find the candidate vertices based on the Breadth First Search (BFS).

From the nodes u and v, *BFS-based Edge Anonymization* traverses the graph in a breadth-first manner. For the i-hop neighbors of u and v, represented by $neig_i(u)$ and $neig_i(v)$, *edge anonymization* finds the candidate vertices from $neig_i(u) \cup neig_i(v)$ and iteratively link the best one with u or v to create a new triangle. We formulize the NMF of the edge (u, v) as $nmf(u, v)$.

1) Candidates generation: We search the candidate vertices for edge (u, v) in a BFS manner. In the i-hop neighbors of u and v, many vertices cannot be the candidate vertices as violating the ATPP. The vertices w need to satisfy the following conditions to be the candidates in the set CV_i.

a) $w \in neig_i(u) \cup neig_i(v)$.
b) $(w, u, v) \neq \triangle$.
c) $\forall x \in \{u, v\}$ *and* $z \in V'$, if $(w, x) \notin E', (w, z) \in E'$ *and* $(x, z) \in E'$, then (x, z) and (w, z) must be unanonymized.

Condition b) states that (u, v, w) is not a complete triangle, which needs to add edges to create a new triangle. This mainly focus on the case when $i = 1$, where w may links with both u and v. Condition c) follows the ATPP, which guarantees that there will be no effect on the already anonymized edges.

2) Candidates selection: After getting all the candidate vertices satisfying the conditions, we can add new edges between u,v and $w \in CV_i$ to increase the NMF of (u, v). We iteratively select a vertex from CV_i to increase the NMF of (u, v) until $nmf(u, v)$ reaches g_f or CV_i is empty. If $nmf(u, v) = g_f$, this edge is anonymized successfully.

In each iteration, we need to select the best one which can preserve the most utility of the graph. Based on the link prediction theory [15], we select the candidate vertex w_{max} which guarantees that $nmf'(w_{max}, u) + nmf'(w_{max}, v)$ is maximum, where $nmf'(w, x)$ is defined in Eq.3.

$$nmf'(w, x) = \begin{cases} 0 & (x, w) \in E' \\ nmf(w, x) & otherwise \end{cases} \quad (3)$$

Where $x \in \{u, v\}$. This is referred to as the *maximum mutual friend criterion* for adding edges. The more mutual friends between the two vertices, the less impact the edge addition will have on the utility of the graph.

The selection criteria described in the Eq.3 only can be used for the candidates in the 1-hop and 2-hop neighbors. For all the candidates w in the i-hop neighbors with $i \geq 3$, the NMF of (x, w) is 0. In this situation, we randomly select a candidate vertex w_{max} from CV_i.

As we anonymize edges in descending order of NMF, we must consider the different situations on the NMF of the new edge (x, w_{max}). In the situation $nmf(x, w_{max}) >= g_f$, if $nmf(x, w_{max})$ is not equal to any $g_f^{'} \in G_f$, (x, w_{max}) cannot be added into the graph. This is because we cannot anonymize this edge in descending order anonymization. Otherwise, we add (x, w_{max}) and mark it as "anonymized". We put this edge into the group with NMF equal to $g_f^{'}$. If $nmf(x, w_{max}) < g_f$, add (x, w_{max}) and mark it as "unanonymized".

3) Candidates dynamic removal: After a new triangle was created with the vertex $w_{max} \in CV_i$, we need to consider the effect of this triangle on the other candidate

vertices in CV_i. To ensure the linking u or v with vertices in CV_i follows the ATPP, some vertices will be dynamically removed from CV_i.

If $w \in CV_i$ connected with the selected vertex w_{max} and the edge (w, w_{max}) is anonymized, then we remove w from CV_i. This is because adding either (w, u) or (w, v) creates a new triangle containing (w, w_{max}), and destroys the anonymization of (w, w_{max}).

For any vertex $w \in CV_i$ with (w, w_{max}) is unanonymized, if (w_{max}, x) is anonymized and $(w, x) \notin E'$, then we remove w from CV_i. This is because if we select this w as a new maximum vertex, we need to add (w, x) to create a triangle containing (u, v), meanwhile created a triangle containing (w_{max}, x). This destroyed the anonymization of the edge (w_{max}, x).

4) Edge anonymization: From the nodes u and v, *BFS-based Edge Anonymization* traverses the graph in a breadth-first manner. The *BFSEA* iteratively generates a candidate set CV_i from the i-hop neighbors of u and v, where i increases from 1 to ∞. After getting the candidate set CV_i, *BFSEA* iteratively selects the best one from CV_i by *candidates selection* and creates a triangle to increase the NMF of (u, v), then updates the CV_i by the *candidates dynamic removal*. These operations will break when the NMF of (u, v) reaches the current group NMF g_f or no more candidate vertex can be found from the whole graph.

If $nmf(u, v)$ reaches the current g_f, i.e. (u, v) is anonymized successfully, we mark it as "anonymized". If the *BFSEA* cannot successfully anonymize this edge, adding new vertices can achieve the task. Linking one new vertex with the end vertices of this edge can increase the NMF of this edge by 1. The newly added edges have only one mutual friend, and will be anonymized at the last step of the anonymization algorithm. The above scenario is a pathological case that rarely occurs as in our experiments, no new vertices were added in all cases.

By the breadth-first manner, the *BFSEA* first link u or v with w from the 1-hop neighbors. Thus after (x, w) is added, the shortest path length (SPL) between $x \in \{u, v\}$ and w will only decrease to 1 from 2 with little effect to the utility. Then we gradually increase the value of i, and link u and v with w from the i-hop neighbors, which decreases the SPL between x and w from i to 1. Hence, we prefer the candidates from i-hop neighbors with smaller i value, i.e. breadth-first manner, which can have less effect to the utility of the graph.

To get the $neig_i(u)$ and $neig_i(v)$ for every i, we execute the *Breadth-First Search* with the time complexity as $O(|V| + |E|)$. When $i = 1$, we need to compute the $neig_i(u) \cap neig_i(v)$ to ensure $(w, u, v) \neq \triangle$ stated in the *candidates generation*, and the time complexity is $O(|V|)$. When $i \leq 2$, to get the best candidate from CV_i, we compute the $nmf(w, x)$, $x \in \{u, v\}$, with the time complexity as $O(|V|)$. Hence, for each candidate set, the total running time of the NMF computation is $O(|V|^2)$. When $i \geq 3$, we randomly select a candidate from CV_i to create a triangle, and the time complexity is $O(1)$. Hence, the time complexity of *candidates selection* is $O(|V|^2)$. Therefore, the time complexity of *BFSEA* is $O(|V|^2)$.

As there are $O(|E|)$ edges need to be anonymized, the time complexity of the ADD algorithm is $O(|E||V|^2)$.

C. Algorithm ADD&DEL

Usually, anonymization combining edge deletion with addition will remove or add fewer edges than only applying edge addition. Indeed, it can improve the utility of the anonymized graph. Before introducing the ADD&DEL algorithm, we discuss the method on how to anonymize an edge by edge deletion.

Edge-deletion. For an unanonymized edge (u, v), the algorithm finds any candidate edge (x, w), where x is u or v, which satisfies the following conditions.

a) Both (u, w) and (v, w) exist and are unanonymized.
b) For any vertex z linked with x and w, edges (x, z) and (w, z) are still unanonymized.
c) If both (u, w) and (v, w) satisfy condition b), we choose the one with fewer mutual friends.

Condition c) is the reverse of the maximum mutual friend criterion for adding edge. The fewer the mutual friends, the weaker the relationship. Hence dropping the edge has less impact to the utility. After (x, w) is deleted, the shortest path length between x and w will only increase to 2 from 1 with little effect to the utility. Condition a) and b) follows the anonymized triangle preservation principle to guarantee that there will be no effect on the already anonymized edges.

For an unanonymized edge (u, v), *edge-deletion* initially finds all candidate edges satisfying the edge-deletion conditions, and then puts them into the set CE. During each iteration, the edge $e_{min} \in CE$ with the least mutual friends will be removed from the graph and the set CE. The algorithm stops when the NMF of (u, v) reaches the group NMF g_f or CE becomes an empty set. If CE is empty and the NMF of (u, v) is not equal to g_f, the anonymization of (u, v) is unsuccessful; Otherwise, we successfully anonymize this edge and mark it as "anonymized".

The *edge-deletion* is the reverse of the methods in ADD algorithms. The running time mainly costs on the computing of mutual friends, so the complexity of *edge-deletion* is $O(|V|^2)$.

The ADD&DEL Algorithm. This algorithm is shown in Algorithm 2, which anonymizes the graph by edge addition and deletion. Similar to the ADD algorithm, ADD&DEL checks the number of unanonymized edges with NMF equal to the NMF of the first unanonymized edge in sorted sequence $\boldsymbol{l}^{(u)}$. If there are more than k edges, we put them into this group

Algorithm 2 The ADD&DEL Algorithm

Input: Original graph $G(V, E)$, k
Output: k-NMF anonymized graph $G'(V', E')$
Initialization: $G' = G$, and mark all edges as "unanonymized". Compute and sort the sequences $\boldsymbol{f}$ and $\boldsymbol{l}$. $\boldsymbol{f}^{(u)} = \boldsymbol{f}, \boldsymbol{l}^{(u)} = \boldsymbol{l}, G_f = \emptyset$

```
while l^(u) ≠ ∅ do
  if |l^(u)| < 2k then do cleanup-operation and break.
  EE = {e|e ∈ l^(u) and f_e^(u) = f_1^(u)};
  if |EE| ≥ k, then new group GP = EE, and mark any e ∈ GP as
    anonymized, G_f = G_f ∪ f_1^(u), update l^(u) and f^(u) and continue.
  GP = ∅, g_f = round(mean(f_1^(u), ..., f_k^(u))). Record all initial info.
  while f_1^(u) ≥ g_f do
    Anonymize l_1^(u) by edge-deletion.
    if anonymize failed, then roll back to initial info, and g_f = g_f + 1;
    else mark l_1^(u) as anonymized and GP = GP ∪ l_1^(u); update f^(u)
    and l^(u).
  end while
  G_f = G_f ∪ g_f.
  while |GP| < k do
    Anonymize l_1^(u) by BFSEA. GP = GP ∪ l_1^(u), update l^(u) and f^(u).
  end while
end while
return G'(V', E').
```

and start another group. Otherwise, we need to anonymize edges to form this group. To gradually anonymize edges and create this group, we initially set the group NMF, g_f, as the mean value of NMFs of the first k unanonymized edges. We record *all initial information* before anonymizing this group. For the unanonymized edge with NMF greater than g_f, we use edge-deletion to anonymize it. If we cannot successfully anonymize this edge, we set $g_f = g_f + 1$ and roll back to *all initial information*. For the unanonymized edge with NMF less than g_f, we apply the ADD algorithm to anonymize it. We gradually anonymize unanonymized edges in sorted sequence $\boldsymbol{l}^{(u)}$ until this group has k edges, and start another group.

In the ADD&DEL algorithm, an edge will be anonymized by either Edge-deletion or methods of the ADD algorithm. Therefore, the time complexity of anonymizing an edge is $O(|V|^2)$, and the time complexity of the ADD&DEL algorithm is $O(|E||V|^2)$.

D. k_1-degree Anonymization Based on k_2-NMF Anonymization

In this subsection, we propose the KDA algorithm on anonymizing the k_2-NMF anonymized graph G' to satisfy k_1-degree anonymity. To maintain the k_2-NMF anonymity of G', the KDA algorithm does not change the NMF of edges in G' when performing anonymization. Proposition 1 stated that the NMF of an edge is related on the number of triangles in which this edge participate, so we anonymize the graph G' without adding new triangles, i.e., the anonymized triangle preservation principle. We can connect two vertices with shortest path length (SPL) no less than three to guarantee that no new triangles will be introduced. Then the NMF of newly added edge is zero. As the degree distribution of the social network follows the power law [13], we only need to anonymize these vertices with large degrees.

The KDA algorithm is similar to the ADD algorithm. The unanonymized vertices are sorted in descending order of their degrees. We gradually group and anonymize them only by edge addition. The vertices in the same group have same degree. To start a new group, KDA set the group degree g_d as the greatest degree of unanonymized vertices. If there are less than k vertices in this group, we anonymize the unanonymized vertices in descending order of their degrees. If this group has more than k vertices, we compute the C_{merge} and C_{new} for the next unanonymized vertex, and decide whether put it into this group or start a new group.

Suppose that the i-hop neighbors of vertex u is $neig_i(u)$. To anonymize the unanonymized vertex u, KDA iteratively and randomly select an unanonymized vertex w_{max} from $neig_3(u)$ and connect u and w_{max}. If the vertex u cannot be anonymized, KDA update the $neig_3(u)$ based on the newly added edges and repeat the above process. If u still cannot be anonymized, we select the candidate vertex from $neig_4(u), neig_5(u)$ and so on until u is anonymized.

When anonymizing a vertex, the KDA algorithm searches the graph in a breadth-first manner to get the candidate vertices. In the worst case, the KDA searches the whole graph and the time complexity is $O(|E| + |V|)$. As there are $O(V)$ vertices needed to be anonymized, the time complexity of the KDA algorithm is $O(|E||V| + |V|^2)$ in the worst case.

IV. Experimental Results

In this section, we conduct experiments on real data sets to evaluate the performance of the proposed graph anonymization algorithms.

A. Datasets

We conduct our experiments on three real datasets: ACM, Cora, and Brightkite. All datasets are preprocessed into simple undirected graphs without self-loop and multiple edges. We also remove the isolated vertices from the graph.

ACM: This dataset was extracted from ACM digital library. We extracted papers published in 12 conference proceedings on computer science before the year 2011. We derive a graph describing the citations between papers. If one paper cites another paper, an undirected edge will connect both corresponding vertices. The graph includes 7,315 vertices and 16,203 edges.

Cora: This dataset is composed of a number of scientific papers on computer science [16]. We extract the collaborations between authors to derive the graph. If two authors had co-authored some papers they would be connected. After we removed the authors without any collaboration, the graph contains 14,076 vertices and 72,871 edges.

Brightkite: This dataset shows the friendships between users in the social network Brightkite over the period of

Table I: The numbers of vertices violating k-degree anonymization and edges violating k-NMF anonymization

	ACM		Cora		Brightkite	
k	k-deg	k-NMF	k-deg	k-NMF	k-deg	k-NMF
5	54	28	141	106	266	93
10	75	28	267	179	533	129
15	103	43	408	277	705	285
20	137	62	446	349	795	393
25	162	62	584	488	891	598
30	221	62	752	575	1088	762
50	262	99	1142	733	1425	1297
100	526	226	1472	1350	2326	2578

April 2008 to October 2010. The graph consists of 58,228 nodes and 214,078 edges, and is available at the SNAP [1].

B. Mutual Friend Attack in Real Data

In the k-degree anonymization model, the adversary re-identifies a vertex using the degree of this vertex. In the k-NMF anonymization model, the adversary re-identifies an edge using the NMF of this edge. We compare both attacks on the real datasets listed in Subsection IV-A, and show the results in Table I. We removed all labels in three datasets. From Table I, we can see that the number of edges violating k-NMF anonymity can be sizable when we set k from 5 to 100. It is a very easy way for an adversary to take the mutual friend attack. k-NMF anonymization problem can be seen as a parallel of the k-degree anonymization problem.

C. Evaluating k-NMF Anonymization Algorithms

We evaluate the performance of the Greedy and Intuitive ADD algorithms and the ADD&DEL algorithm by measuring the average clustering coefficient, average path length, betweenness centrality and the ratios of edges change. Figures 5-8 show the results, where ADD-Int and ADD-Gre stand for the ADD algorithm with IntuitGroup and GreedyGroup respectively. ADD&DEL stands for the ADD&DEL algorithm.

Average Clustering Coefficient (CC): We first compare the average clustering coefficients of the anonymized graphs with the original graph, and the results are shown in Figure 5. The CC values on datasets ACM and Brightkite increase when k increases, but decreased on dataset Cora when k increases. Hence no clear trend on CC change can be concluded, however the average clustering coefficients derived by our three methods deviate slightly from the original values on three datasets. The ADD&DEL performs better than the two ADD algorithms in Figure 5, and the ADD algorithm with GreedyGroup looks slightly better than the algorithm with IntuitGroup.

Average Path Length (APL): Figure 6 shows the average path lengths for the anonymized graphs and the original graphs on three datasets. The APL of the graph anonymized by the ADD&DEL algorithm is very close to the APL of the original graph. By adding and deleting edges, the ADD&DEL algorithm can preserve more utility than the ADD algorithm. Besides, the differences of APL between the graphs anonymized by our methods and the original graphs are very small, and the largest difference value is 0.8 when k is set as 100 on the dataset Cora.

Betweenness Centrality (BC): All the plots of the average betweenness centralities are very similar to the plots of the *APL*. Hence we show the distribution of betweenness centralities of all vertices in Figure 7. Due to space constraints, we only show the results on Cora. The sub-figures in Figures 7(a), 7(b) and 7(c) enlarge the details on the frequency varied from 0 to 100. Clearly, in Figures 7(a) and 7(b), ADD&DEL performs better than the ADD algorithm with GreedyGroup, and shows little sensitivity to the value of k while ADD with GreedyGroup degrades as k increases. Also Figure 7(c) shows that ADD&DEL performs better than the ADD algorithms.

Percentages of edges changed: As there is no vertex addition occurred in all cases considered under ADD and ADD&DEL which do not perform node deletion operations, we consider the edge changes. Figure 8 shows the edge changes on the original graphs. The changes on ADD&DEL includes the ratios of edges added and removed. The ADD&DEL algorithm changed fewest edges, and the ADD algorithm with GreedyGroup added fewer edges than the algorithm with IntuitGroup.

From the above evaluation, we can see that our algorithms can preserve the utility of the original graph effectively. Among them, ADD&DEL performs better than the ADD algorithm, and GreedyGroup performs better than IntuitGroup.

D. Evaluating the KDA Algorithm

In this subsection, we evaluate the performance of the KDA algorithm in Section III-D , and compare it with the classic k-degree anonymization algorithm in [6].

Since there are no new triangles formed after the KDA algorithm adds new edges, the clustering coefficient decreases a little bit as k increases as shown in Figures 9(a), 10(a) and 11(a). Our algorithm performs better than the classic k-degree anonymization on this measure. Since new edges are added into the graph, the APL value decreases a little bit as k increases as shown in Figures 9(b), 10(b), and 11(b). As we consider the k-NMF anonymity, the classic k-degree anonymization performs a little better than our algorithm on the APL measure. But when the APL of the graph is large, our algorithm can perform better than the classic k-degree anonymization as shown in Figure 9(b). The results show that our algorithm performs well on preserving the utility while protecting the privacy by carefully exploring the graph property. The classic k-degree anonymization makes less effort on this except minimizing the number of edges added. Figures 9(c), 10(c) and 11(c) show the distributions of betweenness centrality of graphs anonymized by the KDA algorithm when we set k_{deg} as 10, 20 and 30. The distributions of the anonymized graphs are very similar

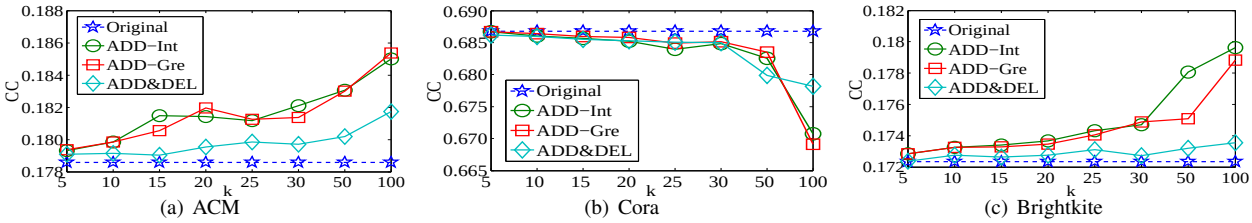

(a) ACM (b) Cora (c) Brightkite

Figure 5: Clustering coefficients

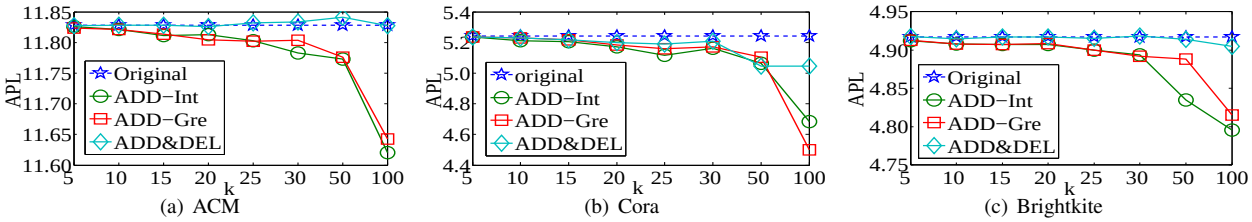

(a) ACM (b) Cora (c) Brightkite

Figure 6: Average path lengths

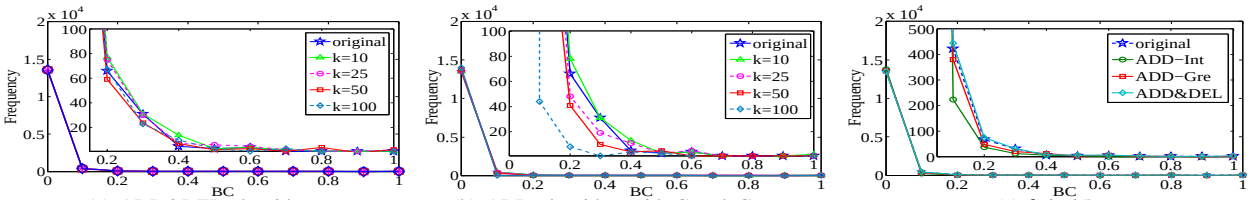

(a) ADD&DEL algorithm (b) ADD algorithm with GreedyGroup (c) k is 25

Figure 7: Betweenness centrality distributions on Cora

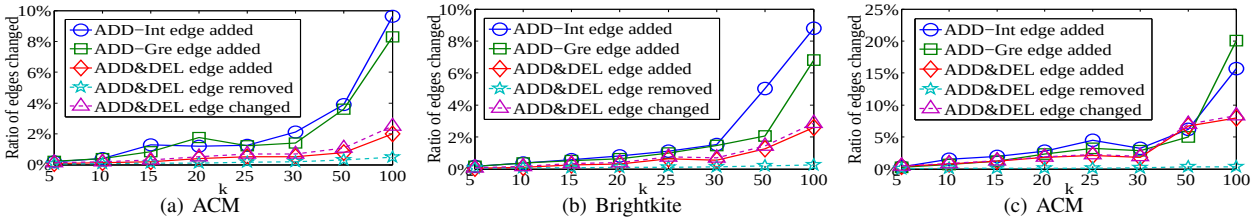

(a) ACM (b) Brightkite (c) ACM

Figure 8: Edge changes

to the distributions of the original graphs especially for the ACM and Brightkite datasets. It shows that the KDA algorithm can preserve much of the utility of the graph anonymized by the k-NMF algorithms.

V. Conclusions

In this paper, we have identified a new problem of k-anonymity on the number of mutual friends, which protects against the mutual friend attack in the social network publication. To solve this problem, we designed two heuristic algorithms which consider the utility of the graph. We also devised an algorithm to ensure the k-degree anonymity based on the k-NMF anonymity. The experimental results demonstrate that our approaches can ensure the k-NMF anonymity while preserve much of the utility in the original social networks.

References

[1] Stanford Network Analysis Project. Available online: http://snap.stanford.edu/index.html. 2007.

[2] B. Zhou, J. Pei, and W. Luk, "A brief survey on anonymization techniques for privacy preserving publishing of social network data," *ACM SIGKDD Explorations Newsletter*, vol. 10, no. 2, pp. 12–22, 2008.

[3] X. Wu, X. Ying, K. Liu, and L. Chen, "A survey of privacy preservation of graphs and social networks," *Managing and Mining Graph Data*, vol. 40, pp. 421–453, 2010.

[4] C.-H. Tai, P. S. Yu, D.-N. Yang, and M.-S. Chen, "Privacy preserving social network publication against friendship attacks," in *Proc. of KDD, San Diego, CA*, 2011, pp. 1262–1270.

[5] L. Backstrom, C. Dwork, and J. Kleinberg, "Wherefore art thou r3579x? anonymized social networks, hidden patterns and structural steganography," in *Proc. of WWW, Banff, Alberta*, 2007, pp. 181–190.

[6] K. Liu and E. Terzi, "Towards identity anonymization on graphs," in *Proc. of SIGMOD, Vancouver, BC*, 2008, pp. 93–106.

[7] C.-H. Tai, P. S. Yu, D.-N. Yang, and M.-S. Chen, "Structural diveristy for pricacy in publishing soical networks," in *Proc. of SDM, Mesa, AZ*, 2011, pp. 35–46.

[8] B. Zhou and J. Pei, "Preserving privacy in social networks against neighborhood attacks," in *Proc. of ICDE, Cancun, Mexico*, 2008, pp. 506–515.

[9] J. Cheng, A. W. Fu, and J. Liu, "K-isomorphism: privacy preserving network publication against structural attacks," in *Proc. of SIGMOD, Indianapolis, IN*, 2010, pp. 459–470.

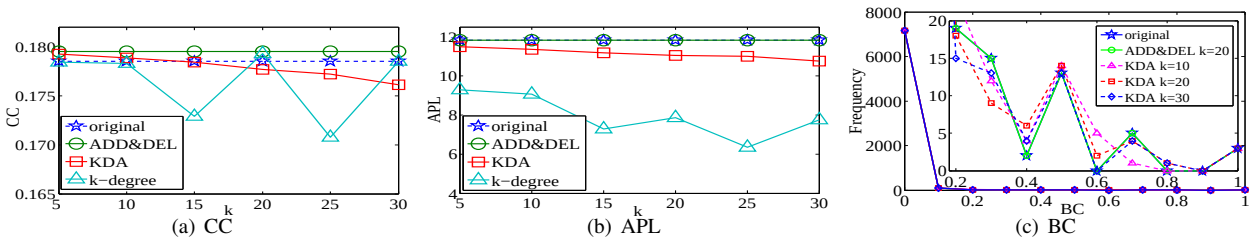

(a) CC (b) APL (c) BC

Figure 9: k-degree anonymization on 20-NMF anonymized graph of ACM

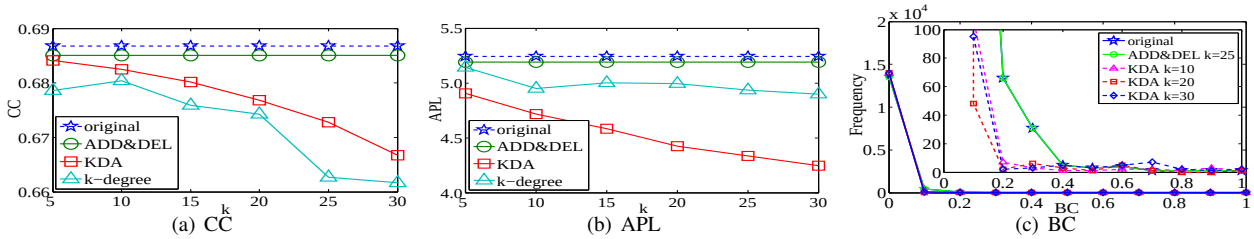

(a) CC (b) APL (c) BC

Figure 10: k-degree anonymization on 25-NMF anonymized graph of Cora

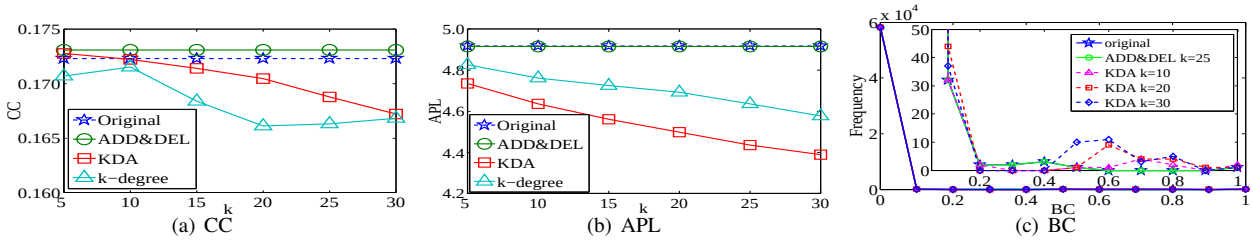

(a) CC (b) APL (c) BC

Figure 11: k-degree anonymization on 25-NMF anonymized graph of Brightkite

[10] L. Zou, L. Chen, and M. T. Özsu, "K-automorphism: a general framework for privacy preserving network publication," *Journal Proceedings of the VLDB Endowment*, vol. 2, no. 1, pp. 946–957, 2009.

[11] W. Wu, Y. Xiao, W. Wang, Z. He, and Z. Wang, "K-symmetry model for identity anonymization in social networks," in *Proc. of EDBT, Lausanne, Switzerland*, 2010, pp. 111–122.

[12] M. Hay, G. Miklau, D. Jensen, D. Towsley, and P. Weis, "Resisting structural re-identification in anonymized social networks," *The VLDB Journal*, vol. 19, no. 6, pp. 797–823, 2008.

[13] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power law relationships of the internet topology," in *Proc. of SIGCOMM, Cambridge, MA*, 1999, pp. 251–262.

[14] V. Zlatic, D. Garlaschelli, and G. Caldarelli, "Complex networks with arbitrary edge multiplicities," *Physics*, vol. 97, pp. 8–11, 2011.

[15] D. Liben-Nowell and J. Kleinberg, "The link-prediction problem for social networks," *Journal of the American society for information science and technology*, vol. 58, no. 7, pp. 1019–1031, 2007.

[16] A. K. McCallum, K. Nigam, J. Rennie, and K. Seymore, "Automating the construction of internet portals with machine learning," *Information Retrieval Journal*, vol. 3, pp. 127–163, 2000.