Privacy-Preserving Assessment of Social Network Data Trustworthiness

Chenyun Dai[1]   Fang-Yu Rao[1]   Traian Marius Truta[2]   Elisa Bertino[1]

[1]Department of Computer Science
Purdue University
West Lafayette, IN
{daic, raof, bertino}@purdue.edu

[2]Department of Computer Science
Northern Kentucky University
Highland Heights, KY
trutat1@nku.edu

Abstract—**Extracting useful knowledge from social network datasets is a challenging problem. To add to the difficulty of this problem, privacy concerns that exist for many social network datasets have restricted the ability to analyze these networks and consequently to maximize the knowledge that can be extracted from them. This paper addresses this issue by introducing the problem of data trustworthiness in social networks when repositories of anonymized social networks exist that can be used to assess such trustworthiness. Three trust score computation models (absolute, relative, and weighted) that can be instantiated for specific anonymization models are defined and algorithms to calculate these trust scores are developed. Using both real and synthetic social networks, the usefulness of the trust score computation is validated through a series of experiments.**

Keywords-social network; privacy; trustworthiness

I. INTRODUCTION

Social networks have been studied by various research communities for more than fifty years [11]. However, the advent of the online social networks and the wide adoption of such networks by our society have significantly increased the importance of obtaining useful information from those networks. Extracting useful knowledge from social network datasets proves to be a difficult problem and social network mining is currently identified as one of the most challenging problems in data mining research [31]. To add to the difficulty of this problem, privacy concerns exist for many social network datasets. Such concerns have resulted in limited accessibility to social network data and thus in reducing the quantity and quality of the knowledge that could be extracted from these datasets. Such knowledge may have important applications, such as disease spreading in epidemiology, emergency management, protection from cyber-attacks, etc.

While large online social networks such as Facebook and LinkedIn are well known and gather millions of users, small social networks are today becoming increasingly common. Currently, such small niche social networks such as GoFISHn and GoHUNTn are considered as the new trend in online social network usage [20]. Many corporations already use existing social networks to connect to their customers. Seeing the increasing usage of small social networks, such companies will likely start to create in-house online social networks where they will own the data shared by customers. Nowadays, for many services (insurance, airline miles, travel sites, etc.), users have individual accounts on company websites. However, there is no network structure connecting accounts of different users, and therefore the relationships that may exist among such users are not efficiently used by the company. The benefits that can be obtained from adding relationships among customers are significant and include increased possibilities to expand the customer base, increased usage of provided services by the current customers, better marketing opportunities, and so on. Future breakthroughs in social network mining will also expand these opportunities. A challenge is that users must have an incentive to connect among themselves in a company-owned social network. This is not a trivial problem and will likely be a difficult challenge to address. However, by using incentives, the users will start to connect to their friends or acquaintances in order to get better deals. For instance, an insurance company may use incentives such as 10% savings on their car insurance costs if a customer registers on his/her social networks site and recommends a minimum number of friends. Next, the amount of savings can be increased based on how many of his/her friends will buy insurance from the same company. Such incentives could also be used to motivate a user to complete his/her profile, and this would allow the insurance company to have a wealth of information about its users that could potentially be used to increase its business.

It can be easily seen that such local social networks have many benefits for the organizations that own them. However, the users' main motivation for joining and providing the required information is to get the desired service at a discount price or any other incentive associated with the use of this company-owned social networking site. Therefore, it is expected that users will be less likely to provide only accurate information in their profiles (due to privacy concerns or because of other advantages that could be obtained by partially faking profile information). An example of a possible advantage that could be obtained is as follows. A user can report his marital status as single although he is married. The reason of such reporting is that his wife may be under 25 years old and adding her in the profile may result in the insurance agency include her in the insurance policy and therefore in increasing the auto insurance rate. Other examples include misreporting of

address, age, and so on. However, due to incentives and since relationships are approved by both participants, links in such social network will likely be accurate. This possibility of faking part of profile data will diminish the utility of the data. The organizations that own such data will benefit from it if they can assess the trustworthiness of such data and if identify possible fake information. Unfortunately, due to privacy regulations, large social network datasets that could potentially be used to verify local information are not available in clear due to privacy concerns. However, we can expect that anonymized social network datasets be available and they can be used to determine the trustworthiness of local data.

To summarize, in our framework, one company will create and maintain a local social network (usually of their customers). This company also has access to one or more anonymized social network datasets that contain the nodes (customers) present in the local social network. Based on these anonymized networks, this company will compute a trust score of their customers and based on these score values will decide to make additional verifications regarding the validity of the reported data or take other actions.

The main contributions of this paper are as follows:

- To introduce the problem of data trustworthiness in social networks when repositories of anonymized social networks exist. To our knowledge there is no prior work that addresses data trustworthiness in social networks.

- To present three trust score computation models (absolute, relative, and weighted) that can be instantiated for specific anonymization models.

- To introduce algorithms to calculate the trust score for nodes (customers) from the local graph for an existing social network anonymization model.

- To illustrate the validity of our trust score computation through experiments. We used both real and synthetic social networks in our experiments.

The remaining of the paper is structured as follows. Sections 2 and 3 define the problem and introduce the notion of trust score for customer profile information, respectively. Section 4 summarizes the social network anonymization model introduced by Campan and Truta [3] and presents an algorithm used to compute the trust score when such an anonymized social network is available. Experimental results of real and synthetic social network datasets are reported in Section 5. Section 6 reviews related work, whereas Section 7 concludes the paper and outlines future research directions.

II. PROBLEM DEFINITION

We assume that a company has created its own social network. Since this network is usually obtained from its own customers that willingly share their data with the company, we call such a company *data owner*. We use the term *local social network* to refer to the company-owned network. We model this local social network as a graph $G = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the set of nodes and $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ is the set of edges. Each node represents an individual entity such as a customer and each edge represents an existing relation between two nodes. Each node has an associated profile represented by a set of attributes. This set of attribute contains *identifier*, *quasi-identifier*, and occasionally *sensitive* attributes [3] that are supposed to be known by the data owner. We assume that all relationships in this local social network are binary. Moreover, we represent all relationships via unlabeled undirected edges. We use X or Y to represent individual nodes, and X_i, $i = 1 \dots n$, to represent all the nodes in $\mathcal{N}$, where $n = |\mathcal{N}|$. We use the notation $X.A$ to refer to the attribute A's value for the node X.

We assume that the owner of the local social network has access to one or more anonymized social networks. An anonymized social network is provided by an external organization (such as Facebook or LinkedIn) that protects the identity and the sensitive information in the social network data by using an anonymization process. We assume that there are s such anonymized social networks available. We represent these networks as $\mathcal{AG}_j = (\mathcal{AN}_j, \mathcal{AE}_j)$ $(j = 1 \dots s)$. Each such anonymized social network is created by the external organization, owner of the social network, from an original graph. We label the corresponding original graphs as $G_j = (\mathcal{N}_j, \mathcal{E}_j)$. It is worth noting that these graphs are large compared to the local social network.

We initially assume that each anonymized social network contains all nodes from $\mathcal{N}$. Moreover, we assume that all edges from the local social network are present in the underlying social networks from which the anonymized social networks were created. In other words, $\mathcal{N} \subseteq \mathcal{N}_j$ and $\mathcal{E} \subseteq \mathcal{E}_j$, for all $j = 1 \dots s$. Since individuals from the local social network G are present in each anonymized social networks $\mathcal{AG}_j$, the set of attributes from $\mathcal{N}$ and the set of attributes from $\mathcal{AN}_j$ are not disjoint. However, due to anonymization, all identifier attributes are removed from $\mathcal{AN}_j$. Also, the sensitive attributes are not shared between $\mathcal{N}$ and $\mathcal{AN}_j$ (due to privacy concerns and/or regulations the sensitive data that might be owned by the owner of the local social network is not available to the external organization that create the anonymized networks). Based on these assumptions, only a subset of quasi-identifier attributes is common.

We also assume that the owner of the local social network trusts the validity of some of the attributes shared with the anonymized networks, while consider other attributes less reliable. The data owner wants to determine a trust score of nodes information with respect to those non-trusted attributes and to perform additional verification if this score is low. As we discuss later, the data owner will compare the trust scores for all nodes and will select a percentage of the lowest scores for this additional verification. Without limiting the applicability of our approach, we assume that we have only one target attribute, labeled B, which may contain misreported / non-trusted information. When more attributes are non-trusted, we can compute the trust score for one attribute at a time. In order

to assess the trustworthiness of values for this attribute B, the attribute must exist in each anonymized social networks (otherwise the anonymized social network is not useful and will not be considered). We denote the other trusted attributes existing in both the anonymized social network and the local social network with $A_1, A_2,..., A_q$, where q is the number of shared attributes. Note that this set depends on the selected anonymized social network; in other words, the set of common attributes is specific for each anonymized social network and changes for each selected anonymized social network (the number q also changes). Since in our analysis we use only one anonymized social network at the time (see Figure 1), for simplicity we use the same notation for each set of common attributes between local and anonymized social network.
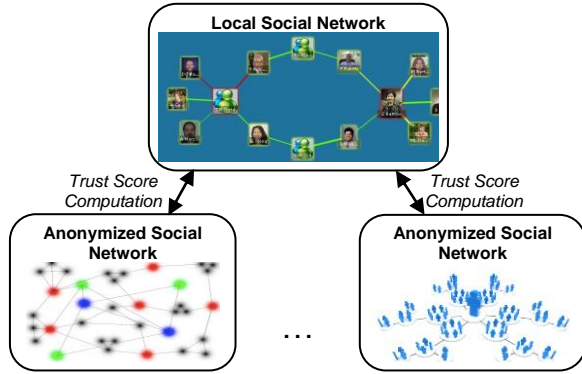

Figure 1. Trust Score Computation Framework

III. TRUST SCORES

We use the notation $TS(X.B)$ (*trust score*) to denote the trustworthiness for the attribute B's value for the node X. For this measure we use all available s anonymized social networks. To obtain this measure, we use the *intermediary trust scores* that we compute for each anonymized social network. We use the notation $TS_j(X.B)$ ($j = 1 .. s$) when $\mathcal{AG}_j$ is used in this intermediary measure.

We compute such an intermediary trust score, $TS_j(X.B)$, by matching a node X from the local social network to nodes from an anonymized social network. We consider in this matching, the node attribute's information (that is, the values of attributes $B, A_1,..., A_q$) and the graph structure. The approach used to compute such score is not unique. For instance, we can consider the trust score as the percentage of nodes from the anonymized social network that could potentially be X. We refer to this approach as *absolute trust score*. A second approach first computes how many nodes from the anonymized network can be X when only the trusted attributes $A_1,..., A_q$ and the graph structure are used. We then find the subset of those nodes that match the value of the B attribute as well (note that a non-generalized value will match its ancestors on the value generalization hierarchy). The number of those nodes divided by the number of nodes that matches X based only on trusted values and graph structure is our second measure of trust. We refer to this measure as *relative trust score*. Our last approach to compute an intermediary trust score includes a weight that depends on how the values of attribute B are

published in the anonymized social network. In most anonymized networks, generalization [26, 27] is used to anonymize the quasi-identifier attributes, and in this case we would like to differentiate between cases when a specific value (such as the exact name of a city) or a generalized value (such as the name of the country) is used. We thus extend the relative trust score computation approach by assigning a higher weight to matches of X with anonymized nodes that contain more specific information for B. More precisely, the weight associated with a specific value is 1, and the weight decreases when the amount of generalization increases. For example, considering the attribute *city*, the weight associated with a single value like *Chicago* is 1 and the weight associated with a generalized value like *Illinois* is 1/10 assuming that there are 10 cities in Illinois in the value generalization hierarchy used for this attribute. We refer to this approach as *weighted trust score*.

Before we formally represent these three measures in Definitions 1 − 3, we introduce the following notations:

- n_j: the number of nodes in the anonymized graph $\mathcal{AG}_j$.

- $matched\_B(X, \mathcal{AG}_j)$: the number of nodes in $\mathcal{AG}_j$ that can potentially be X. All attributes ($B, A_1,..., A_q$) and the graph structure are used in this determination.

- $matched\_No\_B(X, \mathcal{AG}_j)$: the number of nodes in $\mathcal{AG}_j$ that can potentially be X. Only trusted attributes ($A_1, .., A_q$) and the graph structure are used in this determination.

- $matched\_weighted\_B(X, \mathcal{AG}_j)$: each node in $\mathcal{AG}_j$ that can potentially be X will contribute with a weight in the interval $(0, 1]$. This weight is based on the amount of generalization for the attribute B's value in the anonymized graph (see previous paragraphs for a discussion of weights). All such weights are added for the final result. All attributes ($B, A_1,..., A_q$) and the graph structure are used in this determination. It can be easily noticed that for any X, $matched\_weighted\_B(X, \mathcal{AG}_j) \leq matched\_B(X, \mathcal{AG}_j)$.

Definition 1 (*Intermediary Absolute Trust Score*): The intermediary absolute trust score for a node X with respect to an anonymized network $\mathcal{AG}_j$, denoted by $ATS_j(X.B)$, is defined as: $ATS_j(X.B) = \frac{matched\_B(X, AG_j)}{n_j}$

Definition 2 (*Intermediary Relative Trust Score*): The intermediary relative trust score for a node X with respect to an anonymized network $\mathcal{AG}_j$, denoted by $RTS_j(X.B)$, is defined as: $RTS_j(X.B) = \frac{matched\_B(X, AG_j)}{matched\_No\_B(X, AG_j)}$

Definition 3 (*Intermediary Weighted Trust Score*): The intermediary weighted trust score for a node X with respect to an anonymized network $\mathcal{AG}_j$, denoted by $WTS_j(X.B)$, is defined as: $WTS_j(X.B) = \frac{matched\_weighted\_B(X, AG_j)}{matched\_No\_B(X, AG_j)}$

The techniques for compuing those three intermediary trust score measures depend on the specific techniques used for the social network anonymization, and we defer their presentation to Sections 4.

The range for any such intermediary trust score measure is between 0 and 1. The value of 0 is obtained when X does not match any node from the anonymized social network. The value of 1 is obtained in different situations depending of the used intermediary trust score measure. For the absolute trust score the value of 1 means that all nodes in the anonymized network can potentially be X. For the relative trust score a value of 1 is obtained when the use of the B attribute value will not limit the number of matches the node X has in the anonymized graph. For the weighted trust score the value of 1 is obtained when the relative trust score measure is 1 and the attribute B has specific values (non-generalized) in each matched node in the anonymized social network.

Definition 4 (*Trust Score*): The trust score for a node X with respect to a non-trusted quasi-identifier attribute B value is defined as the average of intermediary trust score values computed for all anonymized social networks if all such values are strictly greater than zero and 0 if one such intermediary value is 0. Considering s anonymized social networks are available, the trust score is as:

$$TS(X.B) = \begin{cases} 0, & \text{if } \exists\, j \text{ in } 1..s \text{ such that } TS_j(X.B) == 0 \\ \underset{j=1..s}{\text{avg}}\, TS_j(X.B), & \text{otherwise} \end{cases}$$

In Definition 4, we can use any intermediary trust score measure (absolute, relative, and weighted) and consequently we will obtain three different measures for the total trust score. We use $ATS(X.B)$, $RTS(X.B)$, and $WTS(X.B)$ to denote those three total trust score measures. The following properties hold.

Property 1: $ATS(X.B)$, $RTS(X.B)$, and $WTS(X.B)$ take values in the interval [0, 1] for all nodes X.

Property 2: If any trust score measure is 0 for a given node X, then the other two trust score values are also 0.

Property 3: For any node X, $ATS(X.B) \leq RTS(X.B)$.

Property 4: For any node X, $WTS(X.B) \leq RTS(X.B)$.

Due to space limitations, we omit the proofs for Property 1 to 4. Please refer to our technical report [35] for more details.

The data owner computes the trust score for local nodes, either for all nodes or for a subset that it has already identified through other means as being less trusted with respect to their self-reported B attribute value. Of particular importance are nodes with a trust score of 0. Based on our problem assumptions such nodes have fake values for the attribute B. The data owner can extract the set of possible values for the attribute B that such a node X may have from each anonymized social network. This is executed by considering all B attribute values from all the nodes from each anonymized network that match a node X when only

trusted attributes $A_1,..., A_q$ and the graph structure are used. Therefore a set of possible values for attribute B is obtained for each anonymized network. The last step in this procedure is to intersect all these sets of possible values. Again, due to the assumption that the underlying graphs contain only valid information, the set of possible B attribute values obtained by such intersection will never be empty. For those nodes that have a positive trust score, the data owner is in general not certain of the correctness of the B attribute value. The strategy for those nodes is to sort them based on trust score values and to consider a predefined number (or percentage) of nodes that have the lowest trust score values for a human analysis. These nodes have definitely a higher risk of misreported information.

The weighted trust score measure has an interesting property. If the intermediary weighted trust score for a node X and an anonymized network $\mathcal{AG}_j$ is 1, then *matched_ No_B*(X, $\mathcal{AG}_j$) is equal to *matched_ weighted_B*(X, $\mathcal{AG}_j$). This means that the set of nodes from $\mathcal{AG}_j$ that matches X is the same whether the attribute B value of the node X is used or not. Moreover for each node from this set the corresponding attribute B value must be identical to the attribute B value of node X because if this value would be generalized the weight associated with the computation of *matched_ weighted_B*(X, $\mathcal{AG}_j$) would be strictly less than 1 and the intermediary weighted trust score could not be 1. This means that the data owner knows with absolute certainty that the attribute B value for node X is correct, and thus node X is fully trustworthy. It is easy to notice that the other two trust measures do not have this interesting property. This observation leads us to the updated version of weighted trust score definition (see Definition 5).

Definition 5 (*Weighted Trust Score*): The weighted trust score for a node X with respect to a non-trusted quasi-identifier attribute B value is defined as follows:

$$WTS(X.B) = \begin{cases} 0, & \text{if } \exists\, j \text{ in } 1..s \text{ such that } WTS_j(X.B) == 0 \\ 1, & \text{if } \exists\, j \text{ in } 1..s \text{ such that } WTS_j(X.B) == 1 \\ \underset{j=1..s}{\text{avg}}\, WTS_j(X.B), & \text{otherwise} \end{cases}$$

IV.   TRUST SCORES FOR K-ANONYMOUS CLUSTERED SOCIAL NETWORK

A. Anonymization Model

As our approach depends on the specific anonymization approach used, we first succinctly present the social network anonymization model introduced in [3]. Consider an initial social network to be anonymized. Using a grouping strategy, one can partition the nodes from this network into pairwise disjoint clusters. The goal is that any two nodes from any cluster are indistinguishable based on their relationships and quasi-identifier attributes values. To achieve these objectives, Campan and Truta developed intra-cluster and inter-cluster edge generalization techniques that were used for generalizing the social network structure. They also used generalization [26, 27] for quasi-identifier attributes values and each cluster will have its profile replaced by the generalization information of that cluster. This generalization information is defined next.

Definition 6 (*Generalization Information of a Cluster*): Let $cl = \{X_1, X_2,\dots, X_u\}$ be a cluster of tuples corresponding to nodes selected from IN, QN be the set of numerical quasi-identifier attributes, and QC be the set of categorical quasi-identifier attributes. The generalization information of cl w.r.t. quasi-identifier attribute set $QI = QN \cup QC$ is the "tuple" $gen(cl)$, having the schema QI, where:

- For each categorical attribute $C \in QI$, $gen(cl)[C]$ is equal to the lowest common ancestor in $\mathcal{H}_{Cj}$ of $\{ X_1.C, \dots, X_u.C \}$. We denote by $\mathcal{H}_C$ the hierarchies (domain and value) associated with the categorical quasi-identifier attribute C.

- For each numerical attribute $N \in QI$, $gen(cl)[N]$ is equal to the interval $[min\{X_1.N,\dots, X_u.N \}, max\{ X_1.N,\dots,X_u.N \}]$.

For a cluster cl, its generalization information $gen(cl)$ is the tuple having as value for each quasi-identifier attribute, numerical or categorical, the most specific common generalized value for all that attribute values from cl tuples. In an anonymized graph, each tuple from cluster cl will have its quasi-identifier attributes values replaced by $gen(cl)$.

The notion of *k*-anonymous anonymized social network is fully specified in the following two definitions.

Definition 7 (*Anonymized Social Network*): Given an initial social network, modeled as a graph $IG = (IN, IE)$, and a partition $S = \{cl_1, cl_2, \dots , cl_v\}$ of the nodes set IN, $\bigcup_{j=1}^{v} cl_j = IN$; $cl_i \cap cl_j = \varnothing$; $i, j \in \{1, 2, \dots, v\}$, $i \neq j$; the corresponding anonymized social network AG is defined as $AG = (AN, AE)$, where:

- $AN = \{Cl_1, Cl_2, \dots , Cl_v\}$, Cl_j is a node corresponding to the cluster $cl_j \in S$ and is described by the "tuple" $gen(cl_j)$ (the generalization information of cl_j, w.r.t. quasi-identifiers) and the intra-cluster generalization pair $(|cl_j|, |IE_{clj}|)$;

- $AE \subseteq AN \times AN$; $(Cl_i, Cl_j) \in AE$ iif $Cl_i, Cl_j \in AN$ and $\exists\ X \in cl_i, Y \in cl_j$, such that $(X, Y) \in IE$. Each edge $(Cl_i, Cl_j) \in AE$ is labeled with the inter-cluster generalization value $|IE_{cli,clj}|$.

By construction, all nodes from a cluster cl collapsed into the generalized node Cl are indistinguishable from each other. In order to satisfy the *k*-anonymity property for an anonymized social network each cluster from partition S must have the size at least k.

Definition 8 (*K-Anonymous Social Network*): An anonymized social network $AG = (AN, AE)$, where $AN = \{Cl_1, Cl_2, \dots , Cl_v\}$, and $Cl_j = [(|cl_j|, |IE_{clj}|)]$, $j = 1, \dots, v$ is *k*-anonymous iff $|cl_j| \geq k$ for all $j = 1, \dots, v$.

Campan and Truta developed a social network anonymization algorithm called *Sangreea*, which creates one cluster at the time. In each new formed cluster, nodes are included that are as similar as possible, both in terms of their quasi-identifier attribute values, and in terms of their neighborhood structure. We present an example of application of this anonymization model.

Example 1: Consider the social network IG_1 in Figure 2. IG_1 contains nine nodes, described by the quasi-identifier attribute: *sex*, *age*, and *city*. The *age* quasi-identifier is numerical, *sex* and *city* are categorical. The attribute *sex* can take two values (*M* and *F*) and these values can be generalized only to person (*P*). For this dataset, the attribute *city* takes only four values (*Chicago*, *Detroit*, *Miami*, and *Seattle*). The cities *Chicago* and *Detroit* have the value *Midwest* as their direct ancestor in the value generalization hierarchy. *Miami*, *Seattle*, and *Midwest* have *US* as their ancestor.


|	Sex	Age	City
X_1	M	20	Detroit
X_2	F	25	Detroit
X_3	F	30	Detroit
X_4	M	24	Chicago
X_5	M	32	Chicago
X_6	F	35	Miami
X_7	F	33	Seattle
X_8	F	40	Detroit
X_9	F	45	Detroit

Figure 2. The Social Network IG_1

Figure 3. The *k*-Anonymous Social Network AG_1


Figure 4. The *k*-Anonymous Social Network AG_2

In Figures 3 and 4, two 3-anonymous social networks, AG_1 and AG_2, are given. For generating the anonymized network in Figure 3, the structure was considered more important in the determination of clusters, while for the anonymized network in Figure 4, the quasi-identifiers attribute values were given priority. In each cluster we represent the generalization information of that cluster, followed by a pair of numbers that represent the number of nodes and the number of inter-cluster edges. Outside each cluster we represent the set of original nodes that were

clustered together. The edges' weight represents the number of edges between nodes from the two connected clusters.

B. Intermediary Trust Scores Computations

In this section we present a method to compute the intermediary trust scores when the *k*-anonymous clustered model is used for anonymizing the social networks. We explain our approach through the following example which continues Example 1 from the previous subsection.

Example 1 – Cont.: Suppose that the data owner has a local social network G_1 shown in Figure 5. The data owner has access to both *k*-anonymous social networks AG_1 and AG_2 in Figures 3 and 4. The data owner fully trusts the values of the *age* and *sex* attributes, and wants to assess the trustworthiness of the *city* attribute. Note that all assumptions made in Section 2 are satisfied. Suppose now that three nodes in G_1 contain fake information for the attribute *city*. These fake values are underlined and italicized in Figure 5. The data owner does not have access to IG_1, and nodes labels are used just for illustration purposes.

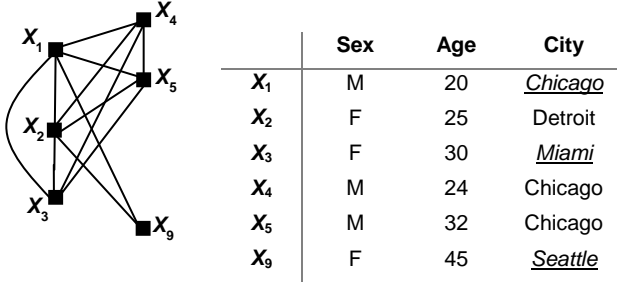|	Sex	Age	City
X_1	M	20	*Chicago*
X_2	F	25	Detroit
X_3	F	30	*Miami*
X_4	M	24	Chicago
X_5	M	32	Chicago
X_9	F	45	*Seattle*

Figure 5.   The Local Social Network G_1

In order to compute the trust scores, the first step is to match each node from the local graph to clusters from each anonymized graph. We perform two types of matches. First, we use the graph structure and the *sex* and *age* attributes values. Second, we also add the non-trusted attribute *city* to this match process. Algorithms for these procedures are detailed later in this section. After these matches, we can directly compute the intermediary absolute and relative trust scores according to Definitions 1 and 2. For the weighted trust score, we define a set of weights as follows: for a single value in a cluster the weight is 1; for *Midwest*, the weight is 1/2 (there are two possible values, Detroit or Chicago); for US the weight is 1/4 (again, due to four possible *city* values). Now we have all the tools to compute the intermediary trust scores. Table 1 and Table 2 show the intermediary trust score computations with respect to AG_1 and AG_2 respectively. Table 3 shows the results of the combined trust scores computed according to Definitions 4 (for absolute and relative trust scores) and 5 (for weighted trust score). For two nodes (X_1 and X_9) the trust score is 0; based on this trust score the data owner is able to determine that the *city* attribute values reported by these two nodes are not correct. However, for X_3, the trust scores values are positive, and the data owner does not know for certain if the values are correct or not. The data owner will use the strict positive values as follows. It will order them increasingly

and will consider the ones with smaller values for additional verification. The data owner may select a specified percentage of values for this additional verification, or may consider a trust score value as a threshold. It is important to notice that the trust score values are relative to each other and the availability of more anonymized networks as well as the quality of those anonymized networks have a major impact on the values of those trust scores.

TABLE I.   THE INTERMEDIARY TRUST SCORES FOR AG_1

	Clusters matched using City	Clusters matched no City	matched_B	matched_No_B	matched_weighted_B	ATS_1	RTS_1	WTS_1
X_1	∅	cl_1	0	3	0	0	0	0
X_2	cl_1	cl_1	3	3	3	1/3	1	1
X_3	cl_2	cl_1, cl_2	3	6	3/4	1/3	1/2	1/8
X_4	cl_2	cl_1, cl_2	3	6	3/4	1/3	1/2	1/8
X_5	cl_2	cl_2	3	3	3/4	1/3	1	1/4
X_9	cl_3	cl_3	3	3	3/4	1/3	1	1/4

TABLE II.   THE INTERMEDIARY TRUST SCORES FOR AG_2

	Clusters matched using City	Clusters matched no City	matched_B	matched_No_B	matched_weighted_B	ATS_2	RTS_2	WTS_2
X_1	cl_1	cl_1	3	3	3/2	1/3	1	1/2
X_2	cl_2	cl_2	3	3	3/4	1/3	1	1/4
X_3	cl_2	cl_2, cl_3	3	6	3/4	1/3	1/2	1/8
X_4	cl_1	cl_1	3	3	3/2	1/3	1	1/2
X_5	cl_1	cl_1	3	3	3/2	1/3	1	1/2
X_9	∅	cl_3	0	3	0	0	0	0

TABLE III.   THE TRUST SCORES FOR NODES FROM G_1

	ATS	RTS	WTS
X_1	*0*	*0*	*0*
X_2	1/3	1	5/8
X_3	*1/3*	*1/2*	*1/8*
X_4	1/3	3/4	5/16
X_5	1/3	1	3/4
X_9	*0*	*0*	*0*

We now introduce an algorithm for *matched_No_B*(X, AG) that we used in our trust score computations.

```
Algorithm Compute_Matched_No_B(X, AG) is
  Input
    G = (N, E) – the local social network;
    X – a node from G;
    AG = (AN, AE) – a clustered anonymized social
        network of a social network G'=(N', E');
    N ⊆ N';  E ⊆ E';
    All information from G' is accurate;
    All information from G is accurate with the
      possible exception of B attribute values;
  Output
    Matched_No_B(X, AG);

A Part  CL1 = ∅;
        For each cl ∈ AN
          If (X ∈ gen(cl)) then
            /*  using A1, .., Aq attributes) */
            CL1 = {cl} ∪ CL1;
B Part  CL2 = ∅;
        For each cl ∈ CL1
          If (check_degree(X, cl) == true) then
            CL2 = {cl} ∪ CL2;
```

```
CL₃ = ∅;
For each cl ∈ CL₂
    /* Neighbors_cl – is a set of clusters that
       contain all clusters where neighbors of
       a node from cl are; this includes cl. */
  Neighbors_cl = ∅;
    If (|IE_cl| > 0) then
     Neighbors_cl = {cl} ∪ Neighbors_cl;
     For each cl' ∈ AN
        If (|IE_cl,cl'| > 0) then
        Neighbors_cl = {cl'} ∪ Neighbors_cl;
     For each neighbor Y of X
        /* Y can be a neighbor of a node from cl
           using degree and attribute's values
        (A₁, .., A_q attributes only) */
     For each cl' ∈ Neighbors_cl
        If (check_degree(Y, cl') == true) and
           (Y ∈ gen(cl')) then
          CL₃ = {cl} ∪ CL₃;
          Break;
 Return the number of nodes that are contained in
 clusters from CL₃;
End Algorithm.
```

(left margin label: Part C)

In parts B and C of the above algorithm in order to select only matching clusters based on possible degree of their nodes (function *check_degree*) we use the following property.

Property 4: Given an anonymized social network AG (cfr. Definition 7), the degree of a node X that belongs to a cluster cl_j with $| cl_j | \geq 2$ is bounded by [$min\_deg(X)$, $max\_deg(X)$] where:

$$min\_deg(X) = max\left(0, \left|IE_{cl_j}\right| - \frac{(|cl_j| - 1) \cdot (|cl_j| - 2)}{2}\right)$$
$$+ \sum_{\substack{i,j=1,v \\ i \neq j}} max\left(0, \left|IE_{cl_i cl_j}\right| - (|cl_j| - 1) \cdot |cl_i|\right)$$

$$max\_deg(X) = min\left((|cl_j| - 1), \left|IE_{cl_j}\right|\right) + \sum_{\substack{i,j=1,v \\ i \neq j}} min\left(|cl_i|, \left|IE_{cl_i cl_j}\right|\right)$$

Proof: Due to space limitations, we omit the proof here.

In the *Compute_Matched_No_B* algorithm we consider a node X from the local graph that matches a cluster from the anonymized graph. Since the local graph is a subgraph of the original graph from which the anonymized graph was created, the degree of X in the local graph is less than or equal to the degree of X in the original graph. Therefore the $min\_deg(X)$ from Property 4 may not be a lower bound of X and we use value 0 instead in the function *check_degree*. On the other hand, $max\_deg(X)$ from Property 4 is used in this function.

The algorithms for computing *matched_B*(X, AG) and *matched_weighted_B*(X, AG) are similar. For *matched_ B*(X, AG), for each cluster cl from CL_3 found by *Compute_matched_No_B*(X, AG) we test if node X belongs to *gen*(cl) using the B attribute. The function returns the number of nodes that are contained only in the clusters that satisfy this test. For *matched_weighted_B*(X, AG), in addition to this new test we also weigh each node based on the amount of generalization for the attribute B in the containing cluster.

The weight associated with a node will be always in the interval (0, 1]. The value of 1 is obtained when the value of the B attribute from the node is identical to the B attribute value from the matching cluster. In other words, this value was not generalized in that cluster. The data owner can define such weights on the value generalization hierarchy for the attribute B (and if such a hierarchy does not exist the data owner can create one for this purpose [4]). There are just two properties that these weights must follow. First, all the leaf values must have a weight of 1. Second, each parent node must have a lower weight than that of all its children.

The running time for part A and part B in *Compute_Matched_No_B(X, AG)* is proportional to the number of clusters, $\lceil |N|/k \rceil$, in the anonymized graph AG. Thus the complexity of part A and part B is O($|N|$). For part C, in the worst case, the number of neighbors that are checked for a given node X is $|N|$-1 and the number of clusters is $\lceil |N|/k \rceil$. The complexity of part C is O($|N|^2$) and the overall complexity of this algorithm is O($|N|^2$).

V. EXPERIMENTAL RESULTS

In this section we evaluate the effectiveness and efficiency of our trust model and algorithms. We used two datasets in the experiments: the synthetic AB datasets (for Albert-Barabasi); and the Enron datasets [10]. For each dataset we selected 5,000 and 10,000 nodes as our initial datasets. To generate the synthetic dataset, we used the Albert-Barabasi model [2] to generate a random graph that follows the power-law distribution. For the Enron datasets the numbers of edges are 67,283 (for 5,000 nodes) and 110,615 (for 10,000 nodes). For the AB datasets the numbers of edges are 97,503 and 197,248 (for 5,000 and 10,000 nodes, respectively).

To fit the need of our experiments, we extended the synthetic and real datasets by introducing the following attribute: *age*, *sex*, *marital status*, and *city*. We use a simple program to generate the attribute values for the synthetic and real social networks described above. To do this, we randomly select a single node and assign to this node values for all four attributes. Next, we run the *breath first search* (BFS) algorithm starting from this source node. Each time when a new node v is first discovered by BFS, we generate its *age* and *city* according to its parent u. For the *age* attribute, we define 6 possible intervals: [18, 27], [28, 37], [38, 47], [48, 57], [58, 67], and [68, 77]. We consider the probability that v is within the same *age* interval as its parent u to be 0.7. The probability that v is in any other interval used by our algorithm is 0.3/5. After the *age* interval of node v is decided, its actual *age* will be a randomly chosen integer within that interval. For *gender* attribute, a node is assigned the *Male* value with probability 0.5 and the *Female* value with probability 0.5. To generate the values of *marital status* attribute for a node v, if v is of *age* greater than or equal to 70, with probability 0.5, its *marital status* will be *Widowed*, and with probability 0.5, its *marital status* will be randomly generated as one of one of the following values: *Divorced*, *Never-married*, *Separated*, *Married-civ-spouse*, *Married-spouse-absent*, *Married-AF-spouse*. If v is of *age* less than 70, then v's *marital status*

will be generated as one of the values above with equal probability. For the *city* attribute, we consider 20 values as shown in Table 4. Similarly, the probability that v is in the same city as its parent u is 0.7. The probability that v is in any other 19 cities is 0.3/19.

For each experiment, from an original dataset (*Enron* or *AB*; each with size 5,000 or 10,000) we generate one *local dataset* and one or more *remote datasets*. The local dataset is generated by taking 20% of the synthetic dataset/real dataset and modifying some of the records in them to create fake/incorrect values. In our experiments, the *city* attribute is the only attribute that contains fake values, in other words, it is the non-trusted attribute. The distance between two leaf nodes is defined as the height of their lowest common ancestor. For example, the distance between *Tampa* and *Miami* is 1, since their lowest common ancestor is *FL* and its height is 1. For the same reason, the distance between *Boston* and *Tampa* is 2 and distance between *Miami* and *Detroit* is 3. To generate the incorrect values, we take a portion (referred as p% in Table 5) of nodes in the local dataset and modify them according to a parameter (m) that controls distances. Among the candidates that share same distance between the original values, we randomly pick one.

TABLE IV.   HIERARCHY STRUCTURE OF THE CITY ATTRIBUTE

City	East	MA	Boston
		FL	Tampa, Miami
	Midwest	MI	Detroit
		IL	Chicago
		IN	Lafayette, Indianapolis, Bloomington
		OH	Cincinnati, Columbus, Cleveland
		KY	Lexington, Louisville, Frankfort, Newport
	Mountain	CO	Aspen
	West	CA	Sacramento, Riverside
		WA	Seattle, Redmond

For remote datasets, we first take the same portion (20%) of the original dataset as local dataset without injecting any fake values. And then randomly pick remaining nodes in the original dataset until the size of the remote datasets reaches 80% of the original ones. In such a way, the remote datasets always contain the corresponding local one.

Table 5 lists the parameters used in the experiments. p controls the percentage of fake values injected into the local dataset. m is the magnitude parameter of the fake values. For instance, when m=3, all distances between fake values and their original ones are 3. When a *city* attribute value does not have a sibling node with m=1 or 2 (e.g., *Boston* and *Aspen*), we will not modify its value and skip that node. k is the k-anonymity parameter for generating anonymization graphs. s is the number of nodes in the original dataset. In all our experiments when a parameter value is not specified, the one denoted in boldface in Table 5 is used.

TABLE V.   EXPERIMENT PARAMETERS

p	10%, **20%**, 30%, 40%, 50%
m	1, 2, **3**
k	3, **5**, 10, 15, 20
s	**5000**, 10000

In our experiments we use two quantities to measure the effectiveness of our trust computation algorithms: (1) the ratio of average score of unmodified nodes (*real*) to the average score of modified nodes (*fake*), denoted as ***Score Ratio*** in the figure; and (2) the ***Recall*** that is the percentage of modified nodes that have been classified as "fake". To classify the local nodes, we simply rank the nodes with respect to their *ATS*, *RTS*, and *WTS* and choose either p% or 1.5p% nodes with the lowest trust scores as possible fake nodes. The data owner will then perform an additional investigation regarding these nodes to determine which of the determined nodes have fake values for the target attribute. For example, if p = 10% and size of the local dataset is 1000, we will classify 100 or 150 nodes with the lowest trust scores as fake nodes. It is worth noting that when no trust score is used and we do a random sampling to choose those 100 or 150 nodes, only p% (10 or 15 if p = 10%) of nodes have fake values.

Due to space limitation, all figures we reported below are based on *Enron* dataset with size 5,000. We observe almost identical trends in the experimental results on *Enron* dataset with size 10,000 and on synthetic *AB* datasets with sizes 5,000 and 10,000.

In the experiments reported in Figure 6, one k-anonymous clustered social network is used with the local graph to compute the trust scores. We observe the effects of changing p on the score ratio and the recall. We notice that the *WTS* score ratio is the highest among the three score ratios regardless of the p value. As for the recall, we observe that all three scores are effective and, among them, *WTS* performs the best. Both *ATS* and *RTS* have more than 70% recall and *WTS* has more than 80% recall when p=10% and it increases when p grows. When p=10%, our approach can increase the recall by more than 7 times when *ATS* or *RTS* are used and more than 8 times when *WTS* is used.

In the experiments reported in Figure 7, one k-anonymous clustered social network is used and p is set to 20%. We can see the effects of changing m on the score ratio and the recall. When m increases, all the score ratios increases, which is as expected because the larger the magnitude of error we inject into a node, the more easily this fake node will be detected. For the recall, it is obvious that recall increases with m and that *WTS* performs the best among those three scores. We can also observe that even for m = 1, the recall values are over 40% (*ATS*), 50% (*RTS*) and 60% (*WTS*) (the baseline recall is 20%).

In the experiments reported in Figure 8, we test how different k values impact the trust score computation results. One k-anonymous clustered social network with different k values is used. Both score ratio and recall decrease when k increases. This is due to the fact that larger k values increase information loss. We can also observe that *WTS* is more affected by k values. This is because *WTS* takes attribute weights into consideration. When the data is more generalized, it is more difficult for *WTS* to distinguish fake values from real ones, since the weight associated with each generalized values is very small.

We also test the running time of our trustworthiness computation algorithms for different privacy models, values of k, and dataset sizes with 5,000 (1,000 local nodes and 4,000 remote nodes) and 10,000 nodes (2,000 local nodes and 8,000 remote nodes). We observe that the running time first increases with k and this is because when k grows it introduces more neighbors. However, when k reaches a threshold the running time becomes stable and it starts to decrease slightly since the number of neighbors is no longer increasing. It is worth mentioning that both algorithms are much more efficient than the anonymization algorithms used to create anonymized social networks. For example, when $k = 20$, the clustered-based anonymization algorithm 5,197 seconds (~ 1.44 hours) to anonymize a 8,000 nodes social network while it only takes the *Compute_Matched_No_B* algorithm less than 100 seconds to compute the trust scores.

Figure 6. One k-anonymous clustered social network, change p, $m = 3$, $s = 5,000$, Enron dataset


Figure 7. One k-anonymous clustered social network, change m, $p = 20\%$, $s = 5,000$, Enron dataset

Figure 8. One k-anonymous clustered social network, change k, $p = 20\%$, $m = 3$, $s = 5,000$, Enron dataset

VI. RELATED WORK

Computing trust for social networks based on external anonymized networks to our knowledge has not yet been investigated. The closest areas to our work are *social network privacy* and *trust (reputation) in social networks*. Other areas that are of interest include *graph isomorphism* and *approximate graph matching*. Social network privacy is a growing area by itself and there are a large number of papers in this field. Many recent studies including the two anonymization models used in this paper focus on protecting vertex identities in published social networks [3, 5, 14, 19, 28, 30, 32, 33, etc.]. Most of these works focus on a *k*-anonymity like approach, while others use a randomization approach [32]. A different privacy model is proposed by Gehrke at al. [12] defined as an extension of differential privacy [9]. The focus of such model is on releasing information about the structure of the social network instead of an anonymized network. De-anonymizing users in social network is also analyzed in

several papers [1, 23, 14]. Backstrom et al. introduce passive and active attacks for de-anonymization of social networks [1]. Narayanan and Smatikov propose an algorithm that using the social network topologies aims at de-anonymizing the social network [23]. Hay et al. [14] use structural queries such as subgraph and hub fingerprint queries to de-anonymize social networks.

Trust in social networks is in general considered at two levels: *global* and *local*. In global level, the trustworthiness of each node is computed based upon complete graph information, while in the local level, trust is computed with respect to the perspective of each specific user. Local trust models include a reputation system based on maximum flows [18], a system based on weighted paths [22], and a system based on spreading activation [34]. Global trust is in general computed based on social ties and high trust is typically associated with influential/authoritative nodes in the network. Methods such as Katz's index [16], PageRank [24], and HITS [17] are used to find such relevant nodes in a

social network. Some recently developed websites [15, 8] also use the same idea of reputation system to detect and remove fake profiles from social networks. To our knowledge, there is no global trust computation approach for a local social network that is owned by an organization as in the problem that we address in this paper.

Graph isomorphism and exact subgraph matching are well-known NP-problems [6, 29]. While improvements to the original backtracking algorithm proposed by Ullman [29] exist, they are not efficient for large networks [7]. More efficient approximate graph matching methods exist [13, 25]. Unfortunately, they try to minimize node mismatches [13] or edge mismatches [25] and they are not applicable to our trust score computation.

VII. Conclusions and Future Work

In this paper, we have shown the importance of taking into account privacy when computing trustworthiness of social network data in applications such as deciding insurance rate. We have proposed three trust score computation models (absolute, relative, and weighted) and algorithms for computing them based on an anonymization model. An extensive experimental evaluation on both real and synthetic data has demonstrated the effectiveness and efficiency of our approach.

In future work, we plan to investigate how to extend our techniques to multiple non-trusted attributes and unknown non-trusted attribute. We also plan to develop new approaches similar to attributes prediction [21] to see whether attribute values of a local node match our prediction. However, our goal is to create prediction models based on remote anonymized datasets and use those models to determine the trustworthiness of local datasets.

VIII. References

[1] Backstrom, L., Dwork, C., and Kleinberg J. M. 2007. Wherefore art thou r3579x?: Anonymized Social Networks, Hidden Patterns, and Structural Steganography. In *Proceedings of the WWW*.

[2] Barabasi, A.-L. and Albert, R. 1999. Emergence of scaling in random networks. *SCIENCE*, 286:509.

[3] Campan, A. and Truta, T. M. 2008. A Clustering Approach for Data and Structural Anonymity in Social Networks. In *PinKDD'08*.

[4] Campan, A., Cooper, N., and Truta, T. M. 2011. On-the-Fly Generalization Hierarchies for Numerical Attributes Revisited. In *Secure Data Management Workshop*, 18-32.

[5] Cheng, J., Fu, A., and Liu, J. 2010. K-isomorphism: Privacy Preserving Network Publication against Structural Attacks. In *SIGMOD'10*.

[6] Cook, S. A. 1971. The Complexity of Theorem-proving Procedures. In *Proceedings of the ACM Symposium on Theory of Computing*, 151-158.

[7] Cordella, L. P., Foggia, P, Sansone, C., and Vento, M. 2004. A (Sub)Graph Isomorphism Algorithm for Matching Large Graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 26, No. 10, 1367-1372.

[8] Duedil, www.duedil.com. 2011.

[9] Dwork, C., McSherry, F., Nissim, K., Smith, A. 2006. Calibrating Noise to Sensitivity in Private Data Analysis.In *Proceedings of the Theory of Cryptography*, 265-284.

[10] ENRON Dataset, Available at http://snap.stanford.edu/data.

[11] Freeman, L. 2006. The Development of Social Network Analysis. *Vancouver: Empirical Press*.

[12] Gehrke, J., Lui, E., and Pass, R. (2011). Towards Privacy for Social Networks: A Zero-Knowledge Based Definition of Privacy. In *Proceedings of the Theory of Cryptography Conference*, 432-449.

[13] Gori, M., Maggini, M., and Sarti, L. 2005. Exact and Approximate Graph Matching Using Random Walks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 27, No. 7, 1100-1111.

[14] Hay, M., Miklau, G., Jensen, D., Towsley, D., and Weis, P. Resisting Structural Re-identification in Anonymized Social Networks. In *Proceedings of the VLDB*, 102-114.

[15] Honestly, www.honestly.com. 2010.

[16] Katz, L. 1953. A New Status Index Derived from Sociometric Analysis. *Psychometrika*, Vol. 18, 39-43.

[17] Kleinberg, J. M. 1999. Authoritative Sources in a Hyper-linked Environment. *Journal of the ACM*, Vol. 6(5), 604-632.

[18] Levien, R. and Aiken, A. 1998. Attack-Resistant Trust Metrics for Public Key Certification. In *Proceedings of the USENIX Security Symposium*, 229-242.

[19] Liu, K, and Terzi, E. 2008. Towards Identity Anonymization on Graphs. In *SIGMOD'08*, 93-106.

[20] MacMillan, D. 2010. The Big Trend in Social Network Sites. *BussinesWeek*, September 9, 2010.

[21] Mislove, A., Viswanath, B., Gummadi, K. P., and Druschel, P. 2010. You Are Who You Know: Inferring User Profiles in Social Networks. In *Proceedings of the ACM Conference on Web Search and Data Mining*.

[22] Mui, L., Mohtashemi, M., and Halberstadt, A. 2002. A Computational Model of Trust and Reputation. In *Proceedings of the Hawaii International Conference on System Science (HICSS)*.

[23] Narayanan, A. and Shmatikov, V. 2008. Robust De-anonymization of Large Sparse Datasets. In *Proceedings of the IEEE Symposium on Security and Privacy*, 111-125.

[24] Page, L., Brin, S., Motwani, R., and Winograd, T.1998. The PageRank Citation Ranking: Bringing Order to the Web. *Stanford Digital Library Technologies Project*, Technical Report.

[25] Pedarsani, P. and Grossglauser, M. 2011. On the Privacy of Anonymized Networks. *In Proceedings of the KDD*, 1235-1243.

[26] Samarati, P. 2001. Protecting Respondents Identities in Microdata Release. *IEEE Transactions on Knowledge and Data Engineering*, Vol. 13, No. 6, 1010-1027.

[27] Sweeney, L. 2002. K-Anonymity: A Model for Protecting Privacy. *International Journal on Uncertainty, Fuzziness, and Knowledge-based Systems*, Vol. 10, No. 5, 557-570.

[28] Tai, C. H., Yang D. N. Yu, P. S., and Chen, M. S. 2011. Structural Diversity for Privacy in Publishing Social Networks. In *Proceedings of the SDM*.

[29] Ullman, J. R. 1976. An Algorithm for Subgraph Isomorphism. *Journal of the ACM*, Vol. 23, No. 1, 31-42.

[30] Wu. W., Xiao, Y., Wang, W., He, Z., and Wang, Z. 2011. K-symmetry Model for Identity Anonymization in Social Networks. In *Proceedings of the EDBT*.

[31] Yang, Q. and Wu, X. 2006. 10 Challenging Problems in Data Mining Research. *International Journal of Information Technology and Decision Making*, Vol. 5, No. 4, 597-604.

[32] Ying, X. and Wu, X. 2008. Randomizing Social Networks: A Spectrum Preserving Approach. In *Proceedings of the SDM*.

[33] Zhou, B. and Pei, J. 2008. Preserving Privacy in Social Networks against Neighborhood Attacks. In *Proceedings of the ICDE*.

[34] Ziegler, C. N. and Lausen, G. 2005. Propagation Models for Trust and Distrust in Social Networks. *Information Systems Frontiers*, Vol. 4, No. 5, 337-358.

[35] Dai, C., Rao, F., Truta, T. M. and Bertino, E. 2012. Privacy-Preserving Assessment of Social Network Data Trustworthiness. *Technical Report*.