
Secure Friend Discovery via Privacy-Preserving and Decentralized Community Detection

Pili Hu
Sherman S.M. Chow
Wing Cheong Lau

HUPILI@IE.CUHK.EDU.HK
SMCHOW@IE.CUHK.EDU.HK
WCLAU@IE.CUHK.EDU.HK

Department of Information Engineering, The Chinese University of Hong Kong

Abstract

The problem of secure friend discovery on a social network has long been proposed and studied. The requirement is that a pair of nodes can make befriending decisions with minimum information exposed to the other party. In this paper, we propose to use community detection to tackle the problem of secure friend discovery. We formulate the first privacy-preserving and decentralized community detection problem as a multi-objective optimization. We design the first protocol to solve this problem, which transforms community detection to a series of Private Set Intersection (PSI) instances using Truncated Random Walk (TRW). Preliminary theoretical results show that our protocol can uncover communities with overwhelming probability and preserve privacy. We also discuss future works, potential extensions and variations.

1 Introduction

One important function provided by social network is friend discovery. The problem of finding people of the same attribute/ interest/ community has long been studied in the context of social network. For example, profile-based friend discovery can recommend people who have similar attributes/ interests; topology-based friend discovery can recommend people from the same community.

One special requirement of algorithms operating on social network is that it must be privacy-preserving. For example, social network nodes may be willing to share their attributes/ interests with people having similar profile; Or they may be willing to share their raw connections with

people in the same community. However, it is unfavourable to leak those private data to arbitrary strangers. Towards this end, the friend discovery routine should only expose minimal necessary information to involved parties.

In the current model of large-scale OSNs, service providers like Facebook play a role of Trusted-Third-Party (TTP). The friend discovery is accomplished as follows: 1) Every node (user) give his/her profile and friend list to TTP; 2) TTP runs any sophisticated social network mining algorithm (e.g. link prediction, community detection) and returns the friend recommendations to only related users. The mining algorithm can be a complex one involving node-level attributes, network topology, or both. Since TTP has all the data, the result can be very accurate. This model is commercially viable and successfully deployed in large-scale. However, recent arise of privacy concern motivates both researchers and developers to pursue other solutions. Decentralized Social Network (DSN) like Diaspora¹ has recently been proposed and implemented. Since it is very difficult to design, implement and deploy a DSN (Datta et al., 2010), much research attention was focused on system issues. We envision that the DSN movement will gradually grow with user's increasing awareness of privacy. In fact, Diaspora, the largest DSN up-to-date, has already accumulated 1 million users. With the decentralized infrastructure established, next question is: can we support accurate friend discovery under the constraint that each node only observes partial information of the whole social network? Note that the whole motivation of DSN is that single service provider can not be fully trusted, so the TTP approach can not be re-used. Towards this end, the computation procedure must be decentralized.

One common approach in literature to achieve decentralized and privacy-preserving friend discovery is to transform it into a set matching problem. For the first type, it is natural to represent one's attributes/ interests/ social activities in form of a set (Zhang et al., 2012). For the second

type, one straightforward way is to represent one’s friend (neighbour) list in form of a set (Nagy et al., 2013). In this way, both profile matching and common friend detection become a set intersection problem. There exists one useful crypto primitive called Private Set Intersection (PSI). Briefly and roughly speaking, given two sets \mathcal{W}_1 and \mathcal{W}_2 held by two node v_1 and v_2 , PSI protocol can compute $|\mathcal{W}_1 \cap \mathcal{W}_2|$ without letting either v_1 or v_2 know other party’s raw input. Researchers have proposed PSI schemes based on commutative encryption (Agrawal et al., 2003), oblivious polynomial evaluation (Freedman et al., 2004) oblivious pseudorandom function (Freedman et al., 2005), indexing message encoding (Manulis et al., 2010), hardware (Hazay & Lindell, 2008) or generic construction (Huang et al., 2012) using garbled circuit (Yao, 1982). The aforementioned privacy-preserving profile matching/ common friend detection protocols are variants of PSI protocols in terms of output, adversary model, security requirement and efficiency.

One major drawback of all the above works is that they can not fully utilize the topology of a social network. Firstly, profile is just node-level information and not always available on every social network. On the contrary, topology (connections/ friendship relations) is the fundamental data available on social networks. Secondly, common friend is just one topology-based approach and it only works for nodes within 2-hops. In fact, our previous investigation showed that common friend heuristic has a moderate precision and low recall for discovering community-based friendship (Hu & Lau, 2013). This result is unsurprising because a community can easily span multiple hops. Towards this end, we focus on extending traditional secure friend discovery beyond 2-hops via community detection. Note that topology-only community detection (Clauset et al., 2004) (Blondel et al., 2008) (Raghavan et al., 2007) (Leung et al., 2009) (Agarwal & Kempe, 2008) (Coscia et al., 2012) (Soundarajan & Hopcroft, 2013) is a classical problem under centralized and non privacy-preserving setting, i.e. a single-party possesses the complete social graph and does arbitrary computation. Although one can translate those algorithms into a privacy-preserving and decentralized protocol using generic garbled circuit construction (Yao, 1982), the computation and communication cost renders it impractical in the real world. To design an efficient scheme, we need to consider community detection accuracy and privacy preservation as a whole. A tradeoff among accuracy, privacy and efficiency can also be made when necessary.

To summarize, this paper made the following contributions:

- We proposed and formulated the first *privacy-preserving* and *decentralized* community detection problem, which largely improves the recall of topology-based friend dis-

covery on Decentralized Social Networks.

- We designed the first protocol to solve this problem. The protocol transforms the community detection problem to a series of Private Set Intersection (PSI) instances via Truncated Random Walk (TRW). Preliminary results show that the protocol can uncover communities with overwhelming probability and preserve privacy.
- We propose open problems and discuss future works, extensions and variations in the end.

2 Related Work

First type of related work is Private Set Intersection (PSI) as they are already widely used for secure friend discovery. Second type of related work is topology-based graph mining. Although our problem is termed “community detection”, the most closely related works are actually topology-based Sybil defense. This is because previous community detection problems are mainly considered under the centralized scenario. On the contrary, Sybil defense scheme sees wide application in P2P system, so one of the root concern is decentralized execution. Note, there exist some distributed community detection works but they can not be directly used because nodes exchange too much information. For example (Hui et al., 2007) allow nodes to exchange adjacency lists and intermediate community detection results, which directly breaks the privacy constraint that we will formulate in following sections. Due to space limit, a detailed survey of related work is omitted. Interested readers can see community detection surveys (Fortunato, 2010)(Xie et al., 2013) and Sybil detection surveys (Yu, 2011)(Alvisi et al., 2013).

3 Problem Formulation

The notion of community is that intra-community is dense and inter-community linkage is sparse. In this section, we first review classical community detection formulations under centralized scenario and our previous formulation under decentralized scenario. Then we formulate the privacy-preserving version. To make the problem amenable to theoretical analysis, we consider a Community-Based Random Graph (CBRG) model in the last part.

3.1 Previous Community Detection Formulations

Classical community detection is formulated as a clustering problem. That is, given the full graph $G = (V, E)$, partition the vertex set into K subsets S_1, S_2, \dots, S_K (a partitioning), such that $\cap_{i=1}^K S_i = \emptyset$ and $\cup_{i=1}^K S_i = V$. A quality metric $Q(\{S_1, \dots, S_K\})$ is defined over the partitions and a community detection algorithm will try to find a partitioning that maximize or minimize Q depending on its nature. This is for non-overlapping community detection

and one can simply remove the constraint $\cap_{i=1}^K S_i = \emptyset$ to get the overlapping version. Note that Q is only an artificial surrogate to the axiomatic notion of community. The maximum Q does not necessarily corresponds to the best community. However, the community detection problem becomes tractable via well-studied optimization frameworks by assuming a form of Q e.g. Modularity, Conductance. Most classical works are along this line mainly due to the lack of ground-truth data at early years.

Now consider the decentralized scenario. One node (observer) is limited to its local view of the whole graph. It is unreasonable to ask for a global partitioning in terms of sets of nodes. The tractable question to ask is: whether one node is in the same community as the observer or not? This gives a binary classification formulation of community detection (Hu & Lau, 2013). The result of community detection with respect to a single observer can be represented as a length- $|V|$ vector. Stacking all those vectors together, we can get a community encoding matrix (Zhong et al., 2014):

$$M_{i,j} = \begin{cases} 1 & \exists S_k, s.t. v_i \in S_k, v_j \in S_k \\ 0 & \text{else} \end{cases}$$

This matrix representation is subsumed by partitioning representation in general case. If restricted to non-overlapping case, the two representations are equivalent. Since M encodes all pair-wise outcome, it is immediately useful for friend discovery application. In what follows, we will define accuracy and privacy in terms of how well M can be learned by nodes or adversary.

3.2 Privacy-Preserving Community Detection

In this initial study, we focus on non collusive passive adversary. That is, DSN nodes all execute our protocol faithfully but they are curious to infer further information from observed protocol sequence. We use a single non-collusive sniff-only adversary to capture this notion. The system components are as follows:

- Graph: $G = (V, E)$. The connection matrix is denoted as C , where $C_{i,j} = 1$ if $(v_i, v_j) \in E$; otherwise, $C_{i,j} = 0$. The ground-truth community encoding matrix is denoted as M^g , which is unknown to all parties at the beginning. For simplicity of discussion, we assume the nodes identifiers, i.e. V , is public information.
- Nodes: $v_1, \dots, v_{|V|} \in V$. A node's initial knowledge is its own direct connections, i.e. $N(v_i) = \{v_j | (v_j, v_i) \in E\}$. Nodes are fully honest. Their objective is to maximize the accuracy of detecting M . Eventually, a node v_i can get full row (column) in M denoted by $M_{i,:}$ ($M_{:,i}$). Depending on the protocol choice, relevant cells in M can be made available immediately or on-demand.
- Adversary: A . It can passively sniff on one node $v_a \in V$. A will observe all protocol sequence related with a , in-

cluding initial knowledge $N(v_a)$ and the community detection result $M_{a,:}$. A 's objective is to maximize successful rate in guessing M^g and C , using any Probabilistic Polynomial Algorithms (PPA). Note, the full separation of Nodes and Adversary is for ease of discussion. In real DSN, this passive attacker can be a curious user who wants to infer more information of the network.

As protocol designer, our objectives are:

- Accurately detect community after execution of the protocol, i.e. making M and M^g as close as possible.
- Limit the successful rate of adversary's guessing of M^g and C , under the condition that A gets the protocol sequence on node v_a and makes best guess via PPA.

One can see that our problem is multi-objective in nature. The accuracy part is a maximization problem and the privacy part is a min-max problem. Formal definition is given in Eq. 1.

In this formulation, "Protocol" is an abstract notation of the protocol specification, not protocol execution sequence. I_a is the information observed by adversary, which is dependent on Protocol. $\text{Succ}(B^1, B^2, R)$ is the measure of successful rate with symbols defined as follows:

- $B^1, B^2 \in \{0, 1\}^{|V| \times |V|}$ are two $\{0, 1\}$ matrix in the same size as M and C .
- $R \subseteq V \times V$ is the challenge relations.
- To measure how close are the two matrix over the challenge set, we use the successful rate:

$$\text{Succ}(B^1, B^2, R) = \Pr \left\{ B_{i,j}^1 = B_{i,j}^2 | (v_i, v_j) \stackrel{\$}{\leftarrow} R \right\}$$

That is, how likely a randomly selected pair of nodes from R will have the same value in B^1 and B^2 .

For the accuracy part, we define the challenge relation as $V \times V$ because we want the result to be accurate for all nodes. For the privacy part, we define the challenge relation as $R_a^C = R_a^M = (V - U(a)) \times (V - U(a))$, where $U(a)$ denotes the set of nodes in the same community as a . The reason to exclude nodes from the same community is obvious. Since adversary will get $M_{a,:}$ after protocol execution, it already knows the community membership of $U(a)$. Given the knowledge of community, one can make more intelligent guess of the connections. This is made clear in later discussions.

3.3 Community-Based Random Graph (CBRG) Generation Model

Before proceed, we remark that the problem defined in Eq. 1 is hard even without the privacy-preserving objective. In other words, the community detection problem (accuracy) has not been fully solved even under the TTP scenario. To improve the accuracy, researchers have already used

$$\max_{\text{Find Protocol, } M = \text{Protocol}(G)} \left(\text{Succ}(M, M^g, V \times V), - \left(\max_{\substack{\text{Algo} \in \text{PPA}, \\ a \stackrel{\$}{\leftarrow} V, \\ C^A, M^A \leftarrow \text{Algo}(\text{Protocol}, I_a)}} \left(\begin{array}{l} \text{Succ}(C^A, C, R_a^C), \\ \text{Succ}(M^A, M^g, R_a^M) \end{array} \right) \right) \right) \quad (1)$$

$$M^g = \begin{bmatrix} 1 & & & \\ 1 & 1 & & \\ 0 & 0 & 1 & \\ 0 & 0 & 1 & 1 \end{bmatrix}, \mathbb{E}[C] = \begin{bmatrix} p & & & \\ p & p & & \\ q & q & p & \\ q & q & p & p \end{bmatrix}$$

Figure 1. Illustration of community-based random graph generation. $K = 2, c = 2$

heavy mathematical programming tools, try to incorporate more side information, develop problem-specific heuristics, or perform heavy-duty parameter tuning. To make our problem amenable to theoretical analysis, we consider a Community-Based Random Graph (CBRG) model in this paper. Let M^g be the ground-truth community encoding matrix. We generate the random connection matrix as follows: 1) $\Pr\{C_{i,j} = 1\} = p$ if $M_{i,j}^g = 1$ (v_i and v_j are in the same community); $\Pr\{C_{i,j} = 1\} = q$ otherwise. There are K communities and each of size c , so the total number of vertices is $|V| = Kc$. We denote such a random graph as $\text{CBRG}(K, c, p, q)$. One example ground-truth community encoding matrix and the expected connection matrix are illustrated in Fig. 1.

4 Proposed Scheme

In this section, we present our protocol and main results.

4.1 Protocol Design

Our protocol involves the two stages:

- Pre-processing is done via Truncated Random Walk. Every node send out W random walkers, $w_1^{v_i}, \dots, w_W^{v_i}$, with time-to-live (TTL) values $l_1^{v_i}, \dots, l_W^{v_i}$ initially set to L . Upon receiving a Random Walker (RW) w , the node records the ID of w , deducts its TTL l , and sends it to a random neighbour if $l > 0$. At the end of this stage, each node v_i accumulated a set of random walker IDs \mathcal{W}_i . With proper parameters W and L , the truncated random walker issued by v_i will more likely reach other nodes in the same community as v_i . So by inspecting the intersection size of \mathcal{W}_i and \mathcal{W}_j , we can answer whether v_i and v_j are in the same community. This essentially transforms the community detection problem to a set intersection problem.
- To uncover the relevant cells in pairwise community encoding matrix M , we only need to perform Privacy Set Intersection (PSI) on two sets. PSI schemes differ in their flavours: 1) reveal intersection set (PSI-Set); 2) reveal intersection size (PSI-Cardinality); 3) reveal whether inter-

section size is greater than a threshold (PSI-Threshold). We use the 3rd type PSI in our construction, which can be implemented by adapting (Zhang et al., 2012). In what follows, we just assume existence of such a crypto primitive: it computes $\mathbb{I}[|\mathcal{W}_i \cap \mathcal{W}_j| \geq T]$ without leaking extra information.

One can see that the scheme is decentralized by design. We only need to argue its community detection accuracy and the privacy-preserving property.

4.2 Summary of Theoretical Guarantees

The intuition of our proof is as follows:

- Truncated Random Walk will be mostly limited to one community, if the axiomatic notion of ‘‘community’’ holds. More precisely, as long as p is enough larger than $\beta_1 = (K - 1)q$, there will be enough difference in intersection size for nodes coming from the same and different communities. In this case, we can set proper threshold to ensure low error rate.
- Observe two facts about privacy objective: 1) most protocol sequence the adversary observed comes from its own community; 2) we exclude A ’s community from challenge relations. In order to make better-than-priori guesses, A at least need to observe some other nodes from protocol sequence. The number of nodes from $V - U(a)$ can be observed is limited. Even if we assume adversary can make good use of the information (captured by coefficient $\gamma_M, \gamma_C \in [0, 1]$), this small advantage is averaged out over a large challenge relation set.

The detailed proof is omitted and the main results are summarized in the following theorem.

Theorem 1 *Our protocol guarantees:*

- *False Positive Rate:*

$$\Pr\{|\mathcal{W}_i \cap \mathcal{W}_j| \geq T_1 | M_{i,j}^g = 0\} \leq \frac{\phi W L (L + 1)^2}{2(K - 1)T_1}$$

- *False Negative Rate: ($\mu = cWP$)*

$$\Pr\{|\mathcal{W}_i \cap \mathcal{W}_j| \leq T_2 | M_{i,j}^g = 1\} \leq e^{-\mu(1 - T_2/\mu)^2/2}$$

- *Adversary’s advantage:*

$$\text{Adv}(M^A, M^g, R_a^M) \leq \gamma_M \frac{4W(L + 1)}{(K - 1)c}$$

$$\text{Adv}(C^A, C^g, R_a^C) \leq \gamma_C \frac{4W(L + 1)}{(K - 1)c}$$

In the theorem, $\text{Adv}(B^1, B^2, R) = \text{Succ}(B^1, B^2, R) - \text{Prior}(B^1, B^2, R)$. $\text{Prior}(B^1, B^2, R)$ denotes the probability to make successful guess based on mere prior information of B^2 . For example, suppose B^2 contains 1 as majority, i.e. $\Pr \left\{ B_{i,j}^2 = 1 \mid i, j \stackrel{\$}{\leftarrow} R \right\} = P > 0.5$. The best guess is to let $B_{i,j}^1 = 1, \forall i, j \in R$. One can show that the success probability is P and this strategy is optimal if no other information is available. Due to the specifics of our problem, adversary can make more intelligent guesses than random $\{0, 1\}$ bit. Towards this end, the advantage is defined with respect to successful rate of this priori-based strategy.

4.3 One Instantiation

Due to the specifics of the problem, both accuracy and privacy guarantees are parameterized. To give an intuitive view of what can be achieved, consider one instantiation of CBRG: $K = 100$ (# of communities), $c = 500$ (# of nodes in one community), $p = 0.5$ (intra-community edge generation probability), $\beta_1 = q(K - 1) = 0.05$, $q = 0.0005$ (inter-community edge generation probability).

We can set protocol parameters as follows: $W = 100$ (# of RWs issued by one node), $L = 3$ (length of RW) and $T = 61$ (threshold of intersection size). This gives us following accuracy and privacy guarantees:

- False Negative Rate: $\leq 1.9 \times 10^{-22}$
- False Positive Rate: ≤ 0.066
- Advantage for guessing M : $\leq 0.032 \times \gamma_M$
- Advantage for guessing C : $\leq 0.032 \times \gamma_C$

One can see that our proposed protocol can accurately detect community and preserve privacy given proper parameters. Note first that above W and L are casually selected by heuristics, which have not been jointly optimized. Note second that the FPR and FNR can be exponentially reduced by repeated experiments, which only maps to a linear increase in W . The example in this section is only to demonstrate the effectiveness of our protocol and a full exploration of design space is left for future work.

5 Conclusion, Discussion and Future Work

We formulated the privacy-preserving community detection problem in this paper as a multi-objective optimization. We proposed a protocol based on Truncated Random Walk (TRW) and Private Set Intersection (PSI). We have proven that our protocol detects community with overwhelming probability and preserves privacy. Exploration of the design space and thorough experimentation on synthesized/ real graphs are left for future work. In following parts of this early report, we discuss several simpler candidate protocols and how they fail to meet our objective.

This help to demonstrate the rationale of our formulation and protocol design.

5.1 Simpler But Weaker Protocols

Suppose we change the protocol such that v_i and v_j first exchange \mathcal{W}_i and \mathcal{W}_j and then run any intersection algorithm separately. After uncovering all related cells in M , adversary knows $\mathcal{W}_i, \forall i = 1, \dots, |V|$. A can directly calculate $|\mathcal{W}_i \cap \mathcal{W}_j|, \forall i, j$. This allows adversary to guess M perfectly. From the community membership, A can further infer links because intra-community edge generation probability and inter-community generation probability are different. This already allows better guess than using global prior of C . Furthermore, inferring links from measurements is a classical well-studied topic called Network Tomography. A can actually re-organize \mathcal{W}_i 's into a list of size- L sets, each representing the nodes traversed by a RW. Researchers have shown that links can be inferred from this co-occurrence data with good accuracy, e.g. NICO (Rabbat et al., 2008).

Another natural thought to protect non-common set elements is via hashing. Suppose there exists a cryptographic hash $h(\cdot)$. We define $\mathcal{H}_i = \{h(w) \mid w \in \mathcal{W}_i\}$. Now, two nodes just compare \mathcal{H}_i and \mathcal{H}_j in the community uncover stage. This can protect true identities of the RWs if their ID space is large enough. However, it does not prevent adversary from intelligent guess of M and C . Methods noted in previous paragraph can also be used in this case.

In our protocol, we used the PSI-Threshold version. That is, given \mathcal{W}_i and \mathcal{W}_j , the two parties know nothing except for the indicator $I[|\mathcal{W}_i \cap \mathcal{W}_j| \geq T]$. Two weaker and widely studied variations are: PSI-Cardinality and PSI-Set. Consider PSI-Set. The adversary now only knows elements in the intersection. Based on his own \mathcal{W}_a and PSI-Set protocol sequence, he can get $\mathcal{W}_i \cap \mathcal{W}_j \cap \mathcal{W}_a, \forall i, j$. A can calculate the probability that a RW w traverses both v_i and v_j conditioned on w traverses v_a . Based on this information, A can adjust threshold T_1 and T_2 to accurately detect communities. The derivation is similar to our protocol in this paper but more technically involved, which is also left as future work. The bottom line is that PSI-Set leaks enough information for more intelligent guesses. As for PSI-Cardinality, we are not sure at present what an adversary can do with $|\mathcal{W}_i \cap \mathcal{W}_a|, \forall i$. Since the two variants leak more information and might be potentially exploited, we use PSI-Threshold in our protocol.

5.2 Open Problems

Following are some open problems of privacy-preserving community detection:

- If we allow a small fraction of nodes to collude, how to de-

fine a reasonable security game? What privacy-preserving result can we achieve?

- Current scheme requires all nodes to re-run the protocol, if there is any change in the topology, e.g. new node joins or new friendship (connection) is formed. Is it possible to find a privacy-preserving community detection scheme that can be incrementally updated?
- The privacy preservation of our proposed protocol is dependent on graph size. One root cause is that we only leveraged crypto primitives in the Private Set Intersection (PSI) part. The simulation of Truncated Random Walk (TRW) is done in a normal way. Since random walk is a basic construct in many graph algorithms, it is of interest know how (whether or not) nodes can simulate Random Walk in a decentralized and privacy preserving fashion.

REFERENCES

- Agarwal, G. and Kempe, D. Modularity-maximizing graph communities via mathematical programming. *The European Physical Journal B-Condensed Matter and Complex Systems*, 66(3):409–418, 2008.
- Agrawal, R., Evfimievski, A., and Srikant, R. Information sharing across private databases. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*, pp. 86–97. ACM, 2003.
- Alvisi, Lorenzo, Clement, Allen, Epasto, Alessandro, Lattanzi, Silvio, and Panconesi, Alessandro. SoK: The evolution of sybil defense via social networks. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pp. 382–396. IEEE, 2013.
- Blondel, V.D., Guillaume, J.L., Lambiotte, R., and Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008.
- Clauset, A., Newman, M.E.J., and Moore, C. Finding community structure in very large networks. *Physical review E*, 70(6):066111, 2004.
- Coscia, M., Rossetti, G., Giannotti, F., and Pedreschi, D. DEMON: a local-first discovery method for overlapping communities. In *ACM SIGKDD*, 2012.
- Datta, A., Buchegger, S., Vu, L.H., Strufe, T., and Rzađca, K. Decentralized online social networks. *Handbook of Social Network Technologies and Applications*, pp. 349–378, 2010.
- Fortunato, S. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.
- Freedman, M., Nissim, K., and Pinkas, B. Efficient private matching and set intersection. In *Advances in Cryptology-EUROCRYPT 2004*, pp. 1–19. Springer, 2004.
- Freedman, Michael J, Ishai, Yuval, Pinkas, Benny, and Reingold, Omer. Keyword search and oblivious pseudo-random functions. In *Theory of Cryptography*, pp. 303–324. Springer, 2005.
- Hazay, Carmit and Lindell, Yehuda. Constructions of truly practical secure protocols using standard smartcards. In *Proceedings of the 15th ACM conference on Computer and communications security*, pp. 491–500. ACM, 2008.
- Hu, Pili and Lau, Wing Cheong. Community classification on decentralized social networks based on 2-hop neighbourhood information. In *IEEE ICNP*, 2013.
- Huang, Yan, Evans, David, and Katz, Jonathan. Private set intersection: Are garbled circuits better than custom protocols. In *Network and Distributed System Security Symposium (NDSS). The Internet Society*, 2012.
- Hui, P., Yoneki, E., Chan, S.Y., and Crowcroft, J. Distributed community detection in delay tolerant networks. In *Proceedings of 2nd ACM/IEEE international workshop on Mobility in the evolving internet architecture*, pp. 7. ACM, 2007.
- Leung, I.X.Y., Hui, P., Lio, P., and Crowcroft, J. Towards real-time community detection in large networks. *Physical Review E*, 79(6):066107, 2009.
- Manulis, Mark, Pinkas, Benny, and Poettering, Bertram. Privacy-preserving group discovery with linear complexity. In *Applied Cryptography and Network Security*, pp. 420–437. Springer, 2010.
- Nagy, Marcin, De Cristofaro, Emiliano, Dmitrienko, Alexandra, Asokan, N, and Sadeghi, Ahmad-Reza. Do i know you?: efficient and privacy-preserving common friend-finder protocols and applications. In *Proceedings of the 29th Annual Computer Security Applications Conference*, pp. 159–168. ACM, 2013.
- Rabbat, M.G., Figueiredo, M.A.T., and Nowak, R.D. Network inference from co-occurrences. *IEEE Transactions on Information Theory*, 54(9):4053–4068, 2008.
- Raghavan, U.N., Albert, R., and Kumara, S. Near linear time algorithm to detect community structures in large-scale networks. *Physical Review E*, 76(3):036106, 2007.
- Soundarajan, Sucheta and Hopcroft, John E. Use of local group information to identify communities in networks. *TKDD*, 2013.
- Xie, Jierui, Kelley, Stephen, and Szymanski, Boleslaw K. Overlapping community detection in networks: the state of the art and comparative study. *ACM Computing Surveys*, 45(4):1–37, 2013.
- Yao, A.C. Protocols for secure computations. In *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, pp. 160–164, 1982.
- Yu, H. Sybil defenses via social networks: a tutorial and survey. *ACM SIGACT News*, 42(3):80–101, 2011.
- Zhang, R., Zhang, Y., Sun, J.S., and Yan, G. Fine-grained private matching for proximity-based mobile social networking. In *Infocom*, 2012.
- Zhong, Xiang, Hu, Pili, and Lau, Wing Cheong. Scalable and robust community detection via proximity-based cut and merge. In *IEEE ICC*, 2014.