Counteracting active attacks in social network graphs

Rolando Trujillo-Rasua

University of Luxembourg, SnT

joint work with Sjouke Mauw and Bochuan Xuan

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs
- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs
- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

This talk is about privacy

Socialization and privacy seems to be in conflict

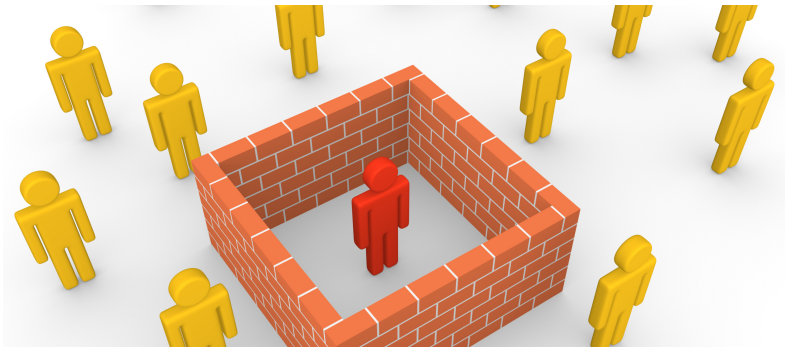

Social zones

https://www.youtube.com/watch?v=n_PoqT3qg5o

This talk is about privacy

Socialization and privacy seems to be in conflict

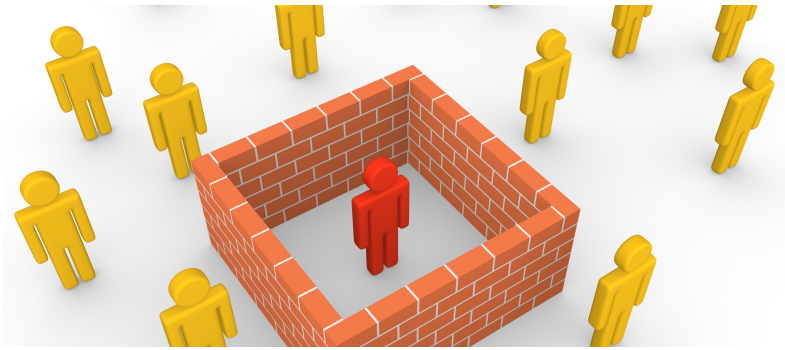

Social zones

`https://www.youtube.com/watch?v=n_PoqT3qg5o`

Which privacy zone Google belongs to?

- Personalize services, ergo saving time
- Define what is important for you and others like you, ergo fast learning from similar peers

But

- How is data being disseminated? Do we have any control over our data? based on the analysis of private information?
- The American Target chain case. How much is too much?

Which privacy zone Google belongs to?

- Personalize services, ergo saving time
- Define what is important for you and others like you, ergo fast learning from similar peers

But

- How is data being disseminated? Do we have any control over our data? based on the analysis of private information?
- The American Target chain case. How much is too much?

Personal data

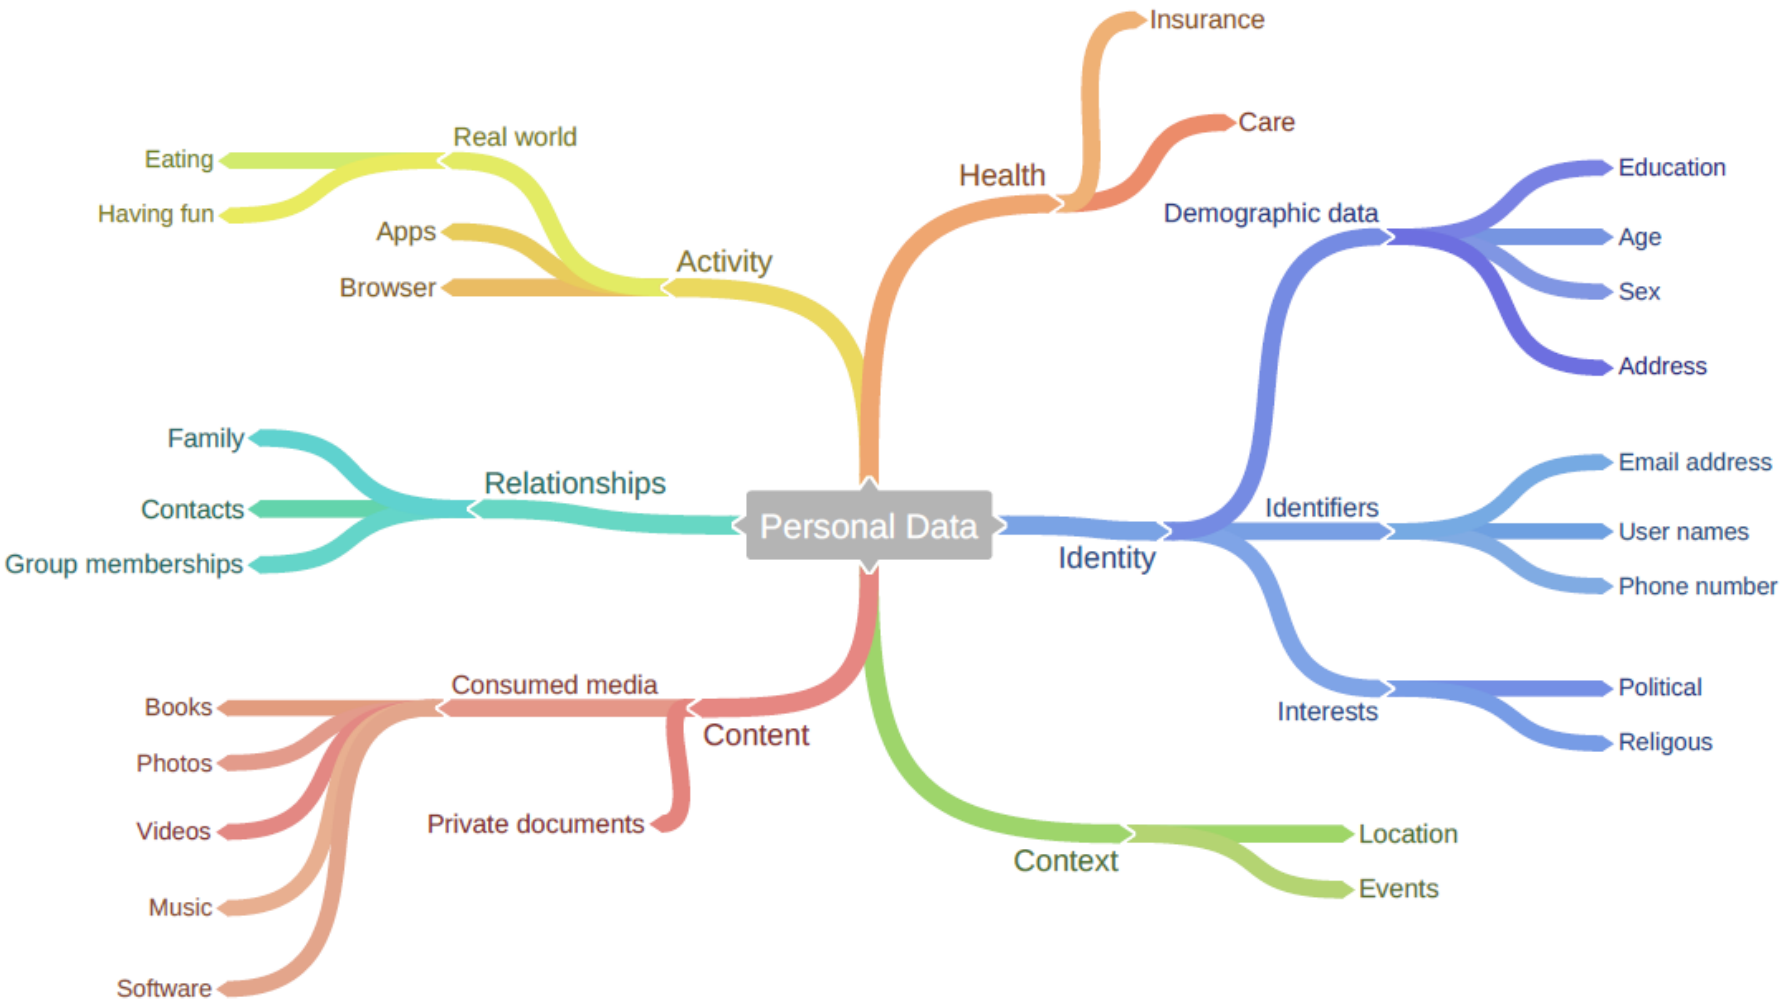

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs

- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

Anonymity: relational database in canonical form

Birth	Postcode	Illness
1975	4350	fever
1955	4350	HIV
1955	5432	flu
1955	5432	HIV
1975	4350	flu
1975	4350	fever

In 2002 Sweeney estimated that 87% of the population in United States can be uniquely identified by combining seemingly innocuous attributes such as gender, date of birth and zip code.

Anonymity: microdata

Birth	Postcode	Illness
1975	4350	fever
1955	4350	HIV
1955	5432	flu
1955	5432	HIV
1975	4350	flu
1975	4350	fever

Anonymity: microdata

Birth	Postcode	Illness
*	4350	fever
*	4350	HIV
1955	5432	flu
1955	5432	HIV
1975	4350	flu
1975	4350	fever

Anonymity: relational database

Table: **Left**: Patient data. **Right**: Potential route of transmission.

Id	Birth	Postcode	Illness
1	1975	4350	fever
2	1955	4350	HIV
3	1955	5432	flu
4	1955	5432	HIV
5	1975	4350	flu
6	1975	4350	fever

Patient 1	Patient 2
1	2
1	3
1	5
3	4
3	5
4	6

Anonymity: relational database

Table: **Left**: Patient data. **Right**: Potential route of transmission.

Id	Birth	Postcode	Illness
1	*	4350	fever
2	*	4350	HIV
3	1955	5432	flu
4	1955	5432	HIV
5	1975	4350	flu
6	1975	4350	fever

Patient 1	Patient 2
1	2
1	3
1	5
3	4
3	5
4	6

Different types of data require different anonymization techniques.

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs
- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

Where do social graphs come from?

Many datasets can be represented as graphs:

- Friendship in online social network
- Financial transactions
- Email communication
- Romantic relationships

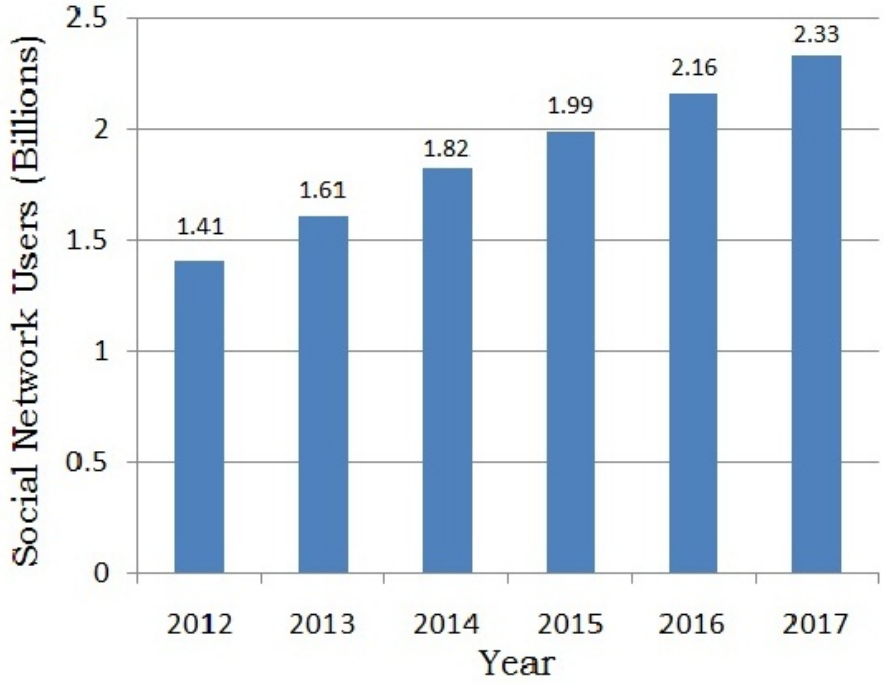

http://www.emarketer.com/Article/India-Leads-Worldwide-Social-Networking-Growth/1010396

What can we infer purely from link structure?

- Popularity
- Centrality
- Introvert vs. Extrovert
- Leadership potential
- Communities

And more, we can also use knowledge from other sources of information.

Can we anonymize graphs?

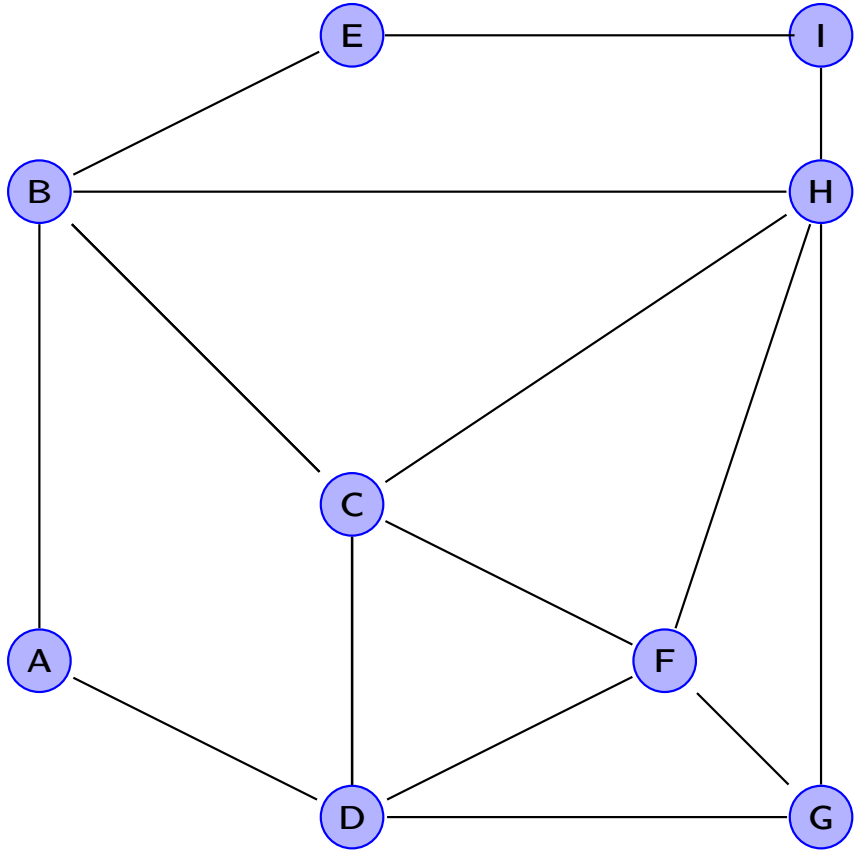

A social graph.

Can we anonymize graphs?

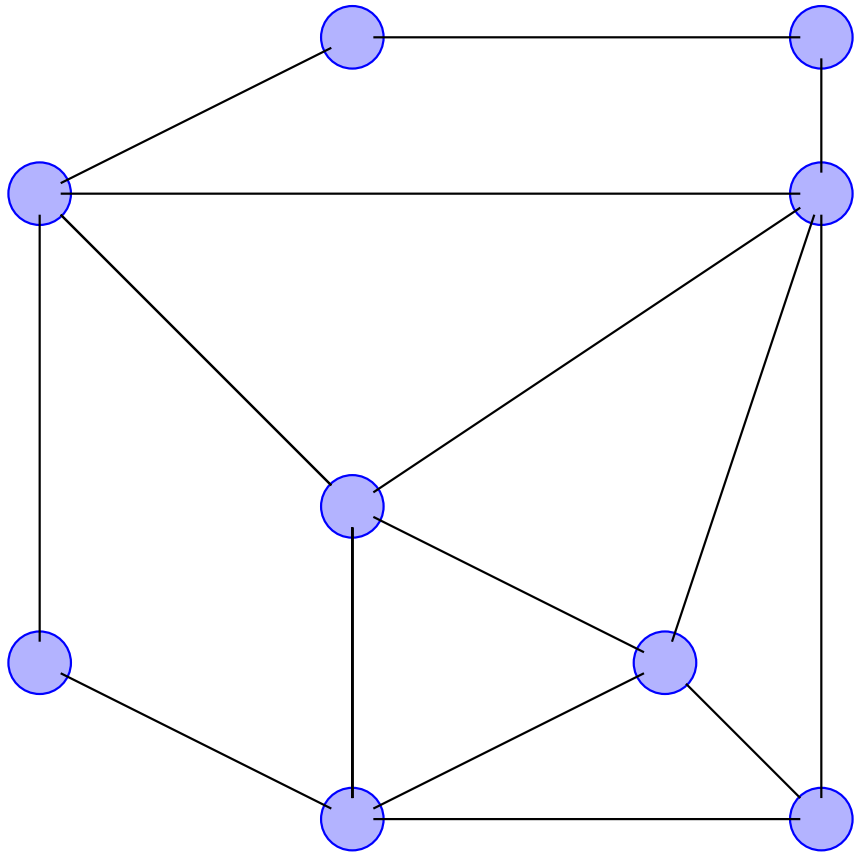

The graph is anonymized and published.

Can we anonymize graphs?

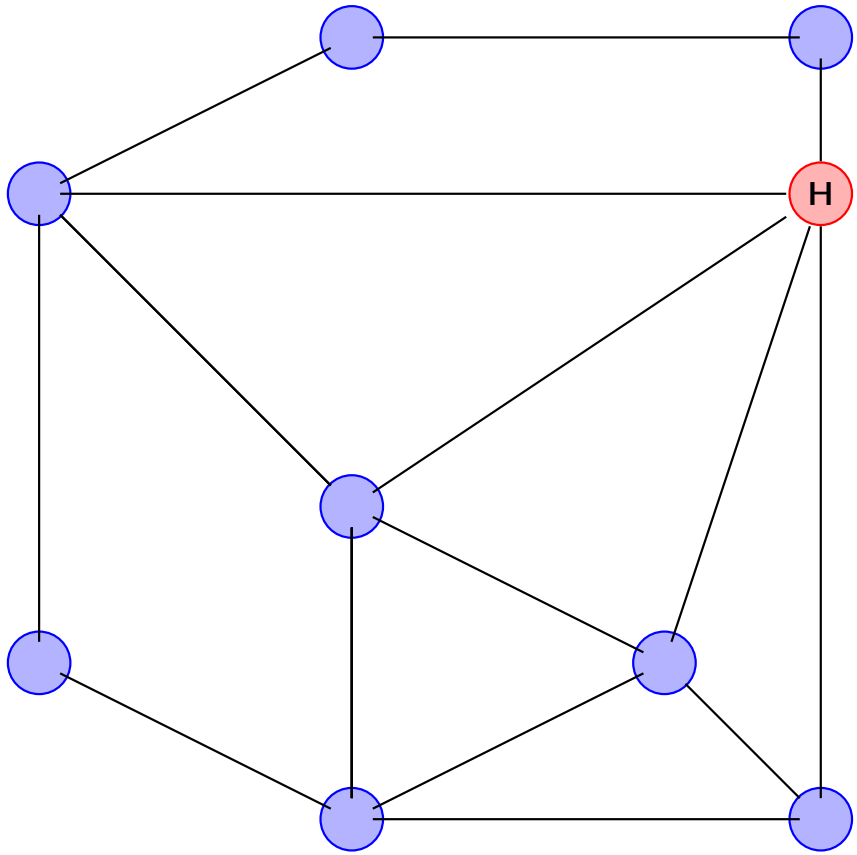

Knowing the number of links of the target node (5) the adversary can re-identify H.

Can we anonymize graphs?

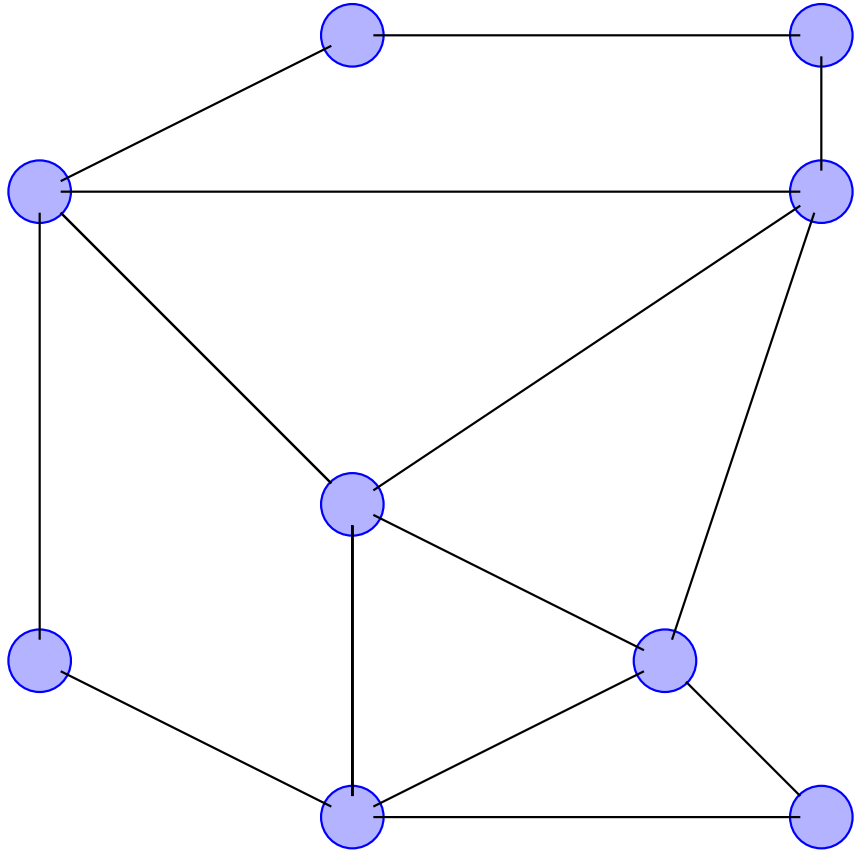

A graph satisfying 4-degree anonymity.

Adversary knowledge and privacy notions

Adversary knowledge	**Anonymity concept**
Vertex degree	k-degree anonymity (2008)
Vertex's neighbourhood	k-neighbourhood anonymity (2008)
Full graph	k-automorphism (2009)

How hard is to keep a secret.

https://www.youtube.com/watch?v=d6gMrLb5ljU

Adversary knowledge and privacy notions

Adversary knowledge	**Anonymity concept**
Vertex degree	k-degree anonymity (2008)
Vertex's neighbourhood	k-neighbourhood anonymity (2008)
Full graph	k-automorphism (2009)

How hard is to keep a secret.

`https://www.youtube.com/watch?v=d6gMrLb5ljU`

Active attacks (Backstrom et al. 2009)

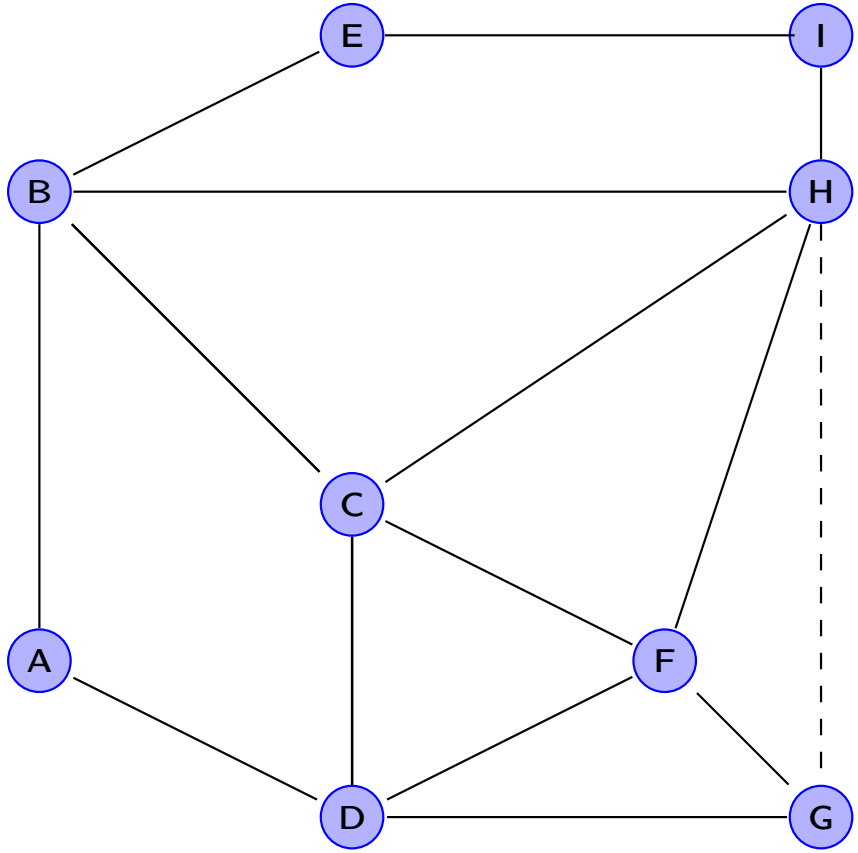

A social graph.

Active attacks (Backstrom et al. 2009)

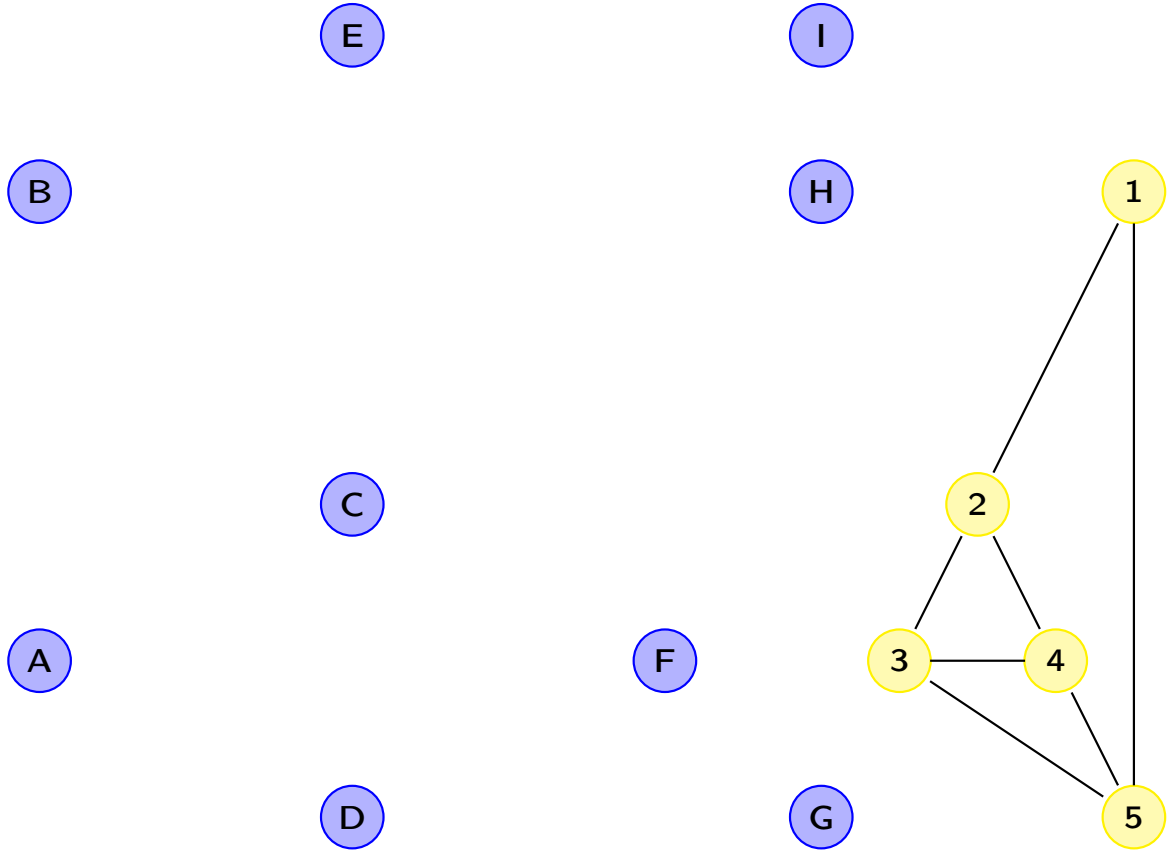

An adversary adds 5 nodes with random connections.

Active attacks (Backstrom et al. 2009)

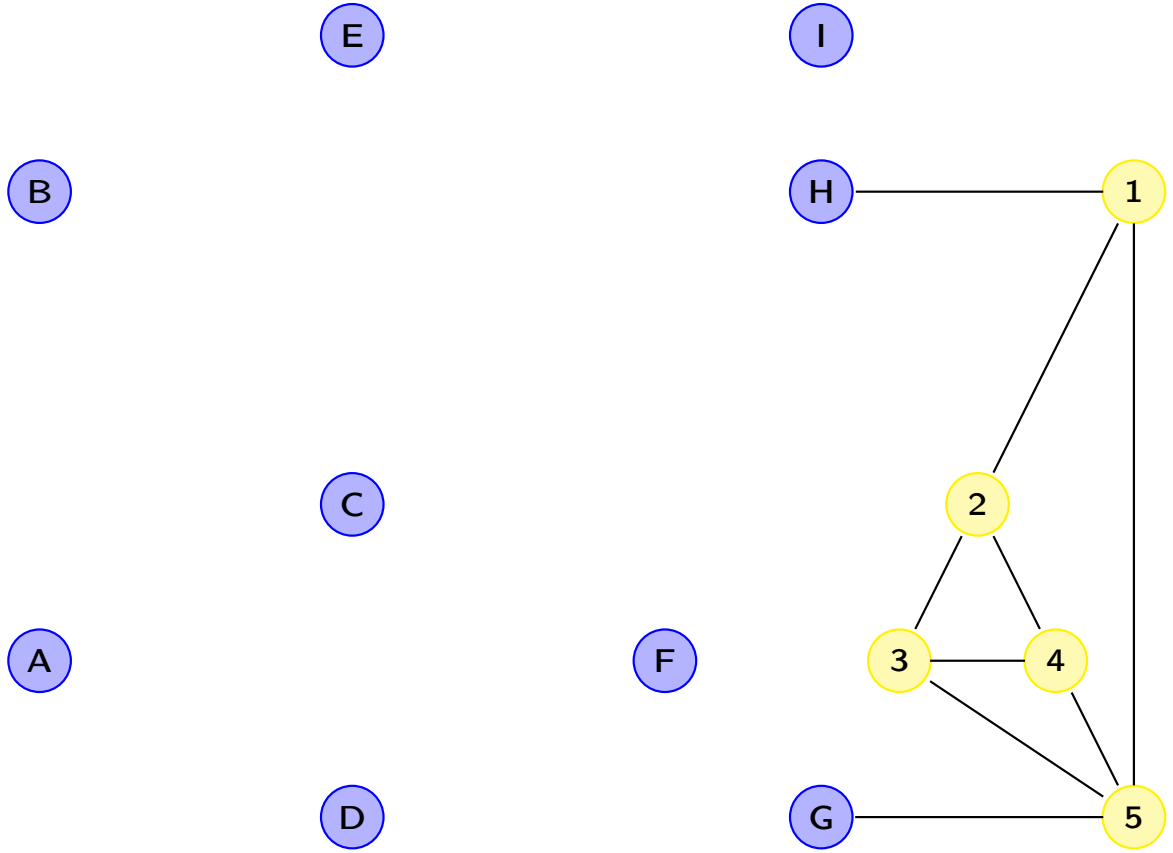

Links with the target nodes are established.

Active attacks (Backstrom et al. 2009)

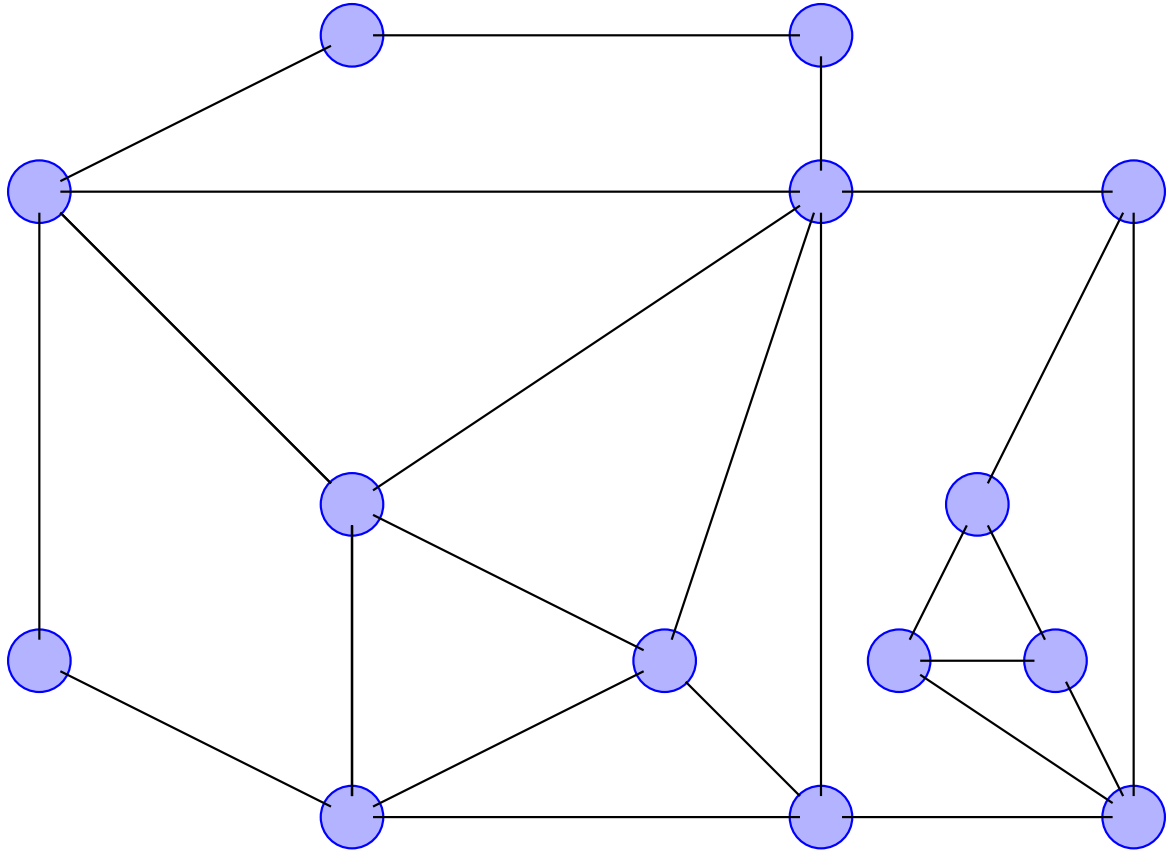

The graph is anonymized and published.

Active attacks (Backstrom et al. 2009)

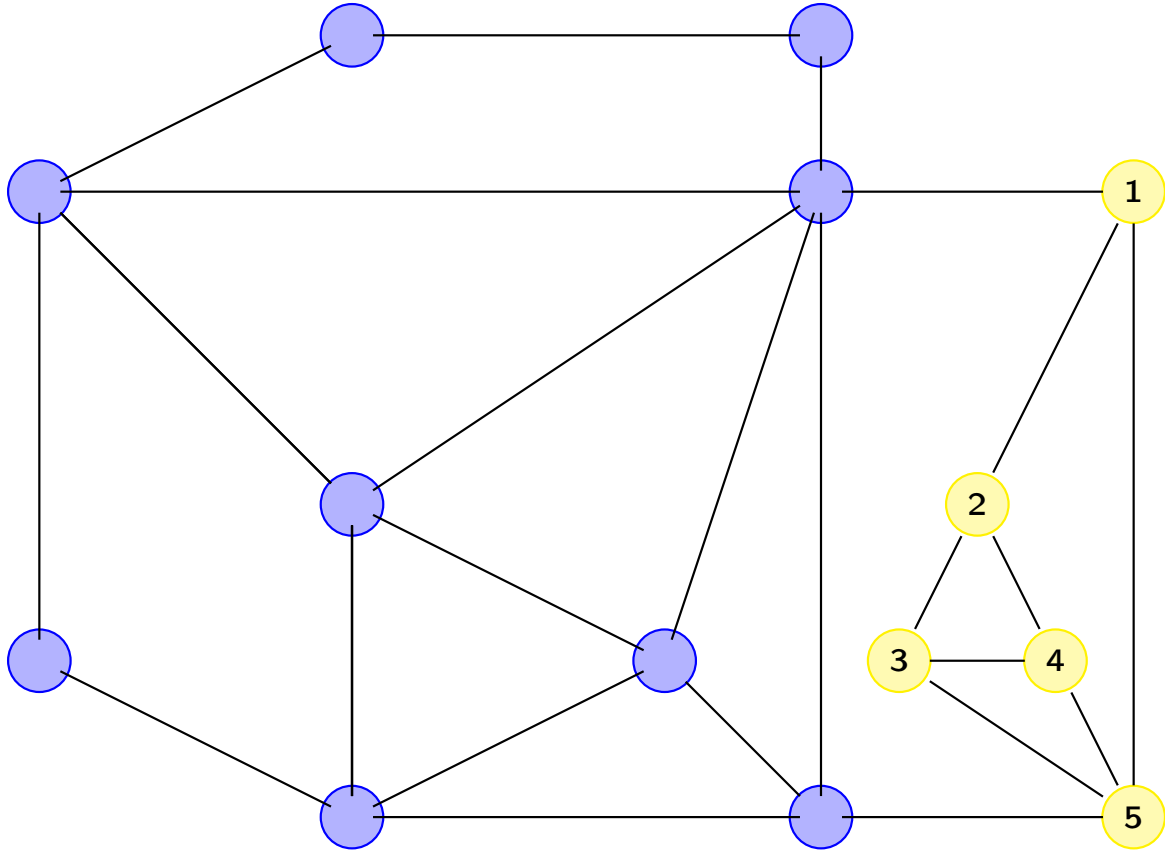

The adversary's subgraph is found.

Active attacks (Backstrom et al. 2009)

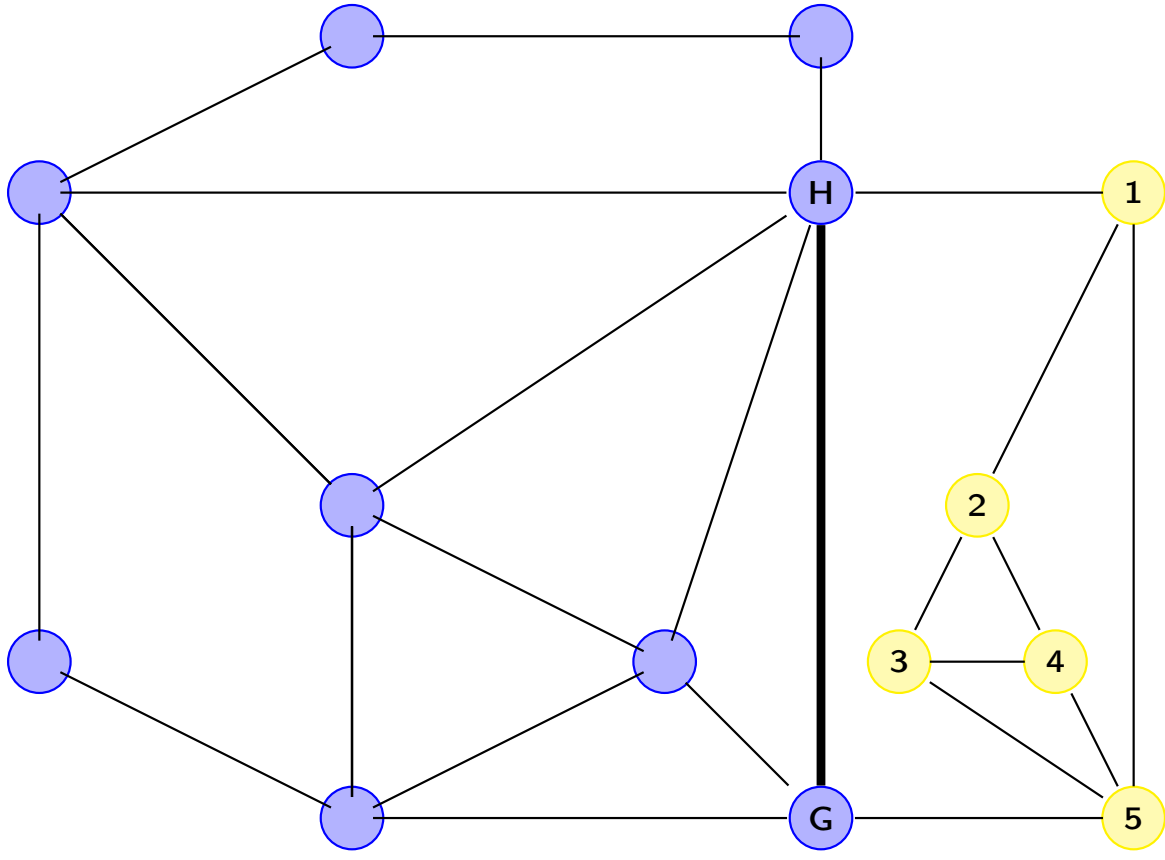

Link between H and G is confirmed.

(Wei et al. 2014) extended this attack to reach arbitrary nodes

How can we quantify privacy against active attacks?

Active attacks (Backstrom et al. 2009)

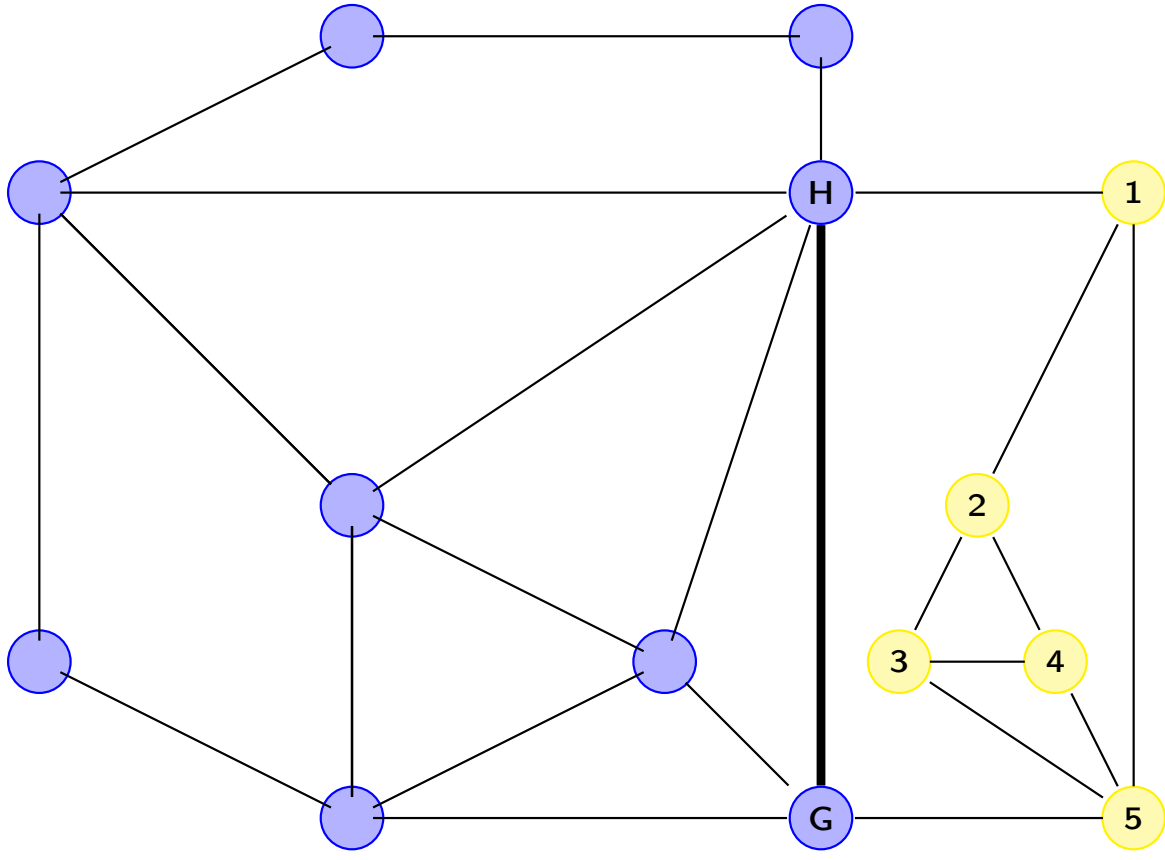

Link between H and G is confirmed.

(Wei et al. 2014) extended this attack to reach arbitrary nodes

How can we quantify privacy against active attacks?

Active attacks (Backstrom et al. 2009)

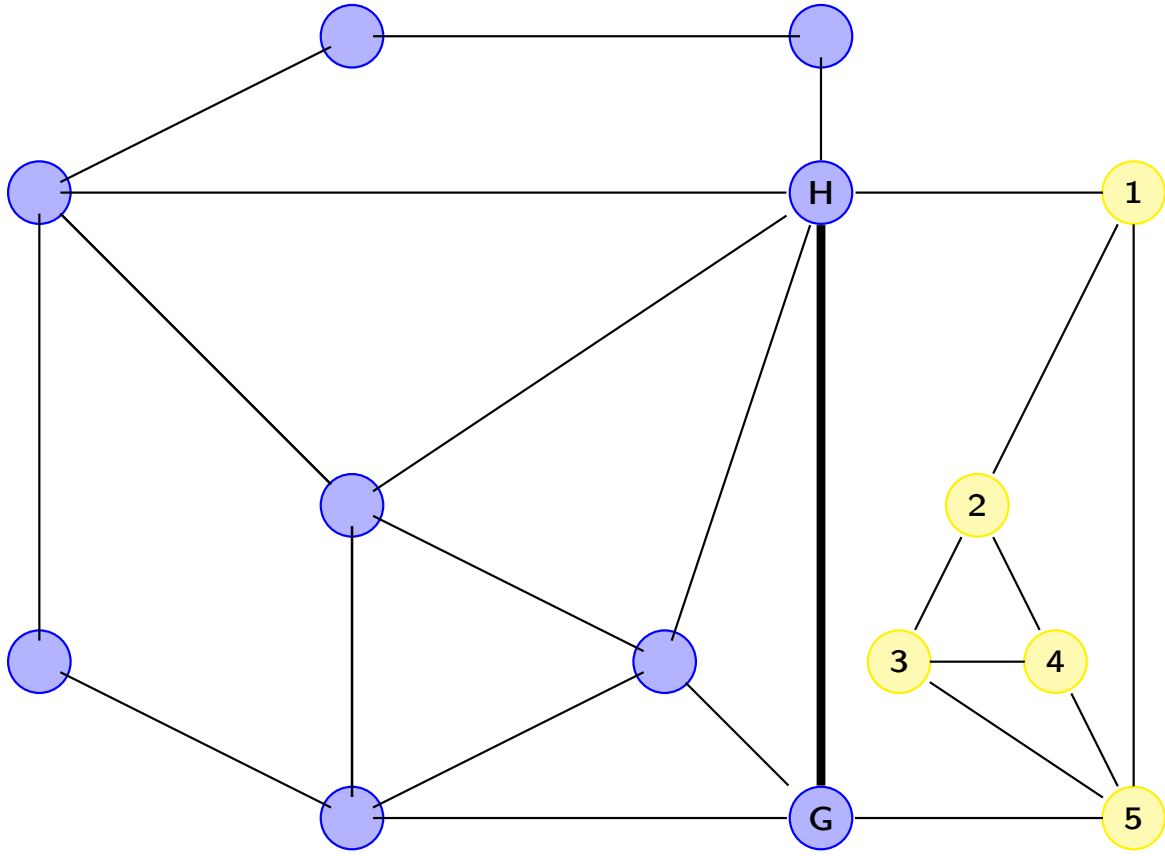

Link between H and G is confirmed.

(Wei et al. 2014) extended this attack to reach arbitrary nodes

How can we quantify privacy against active attacks?

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs

- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

Can k-anonymity be used to prevent active attacks?

A couple of concepts from graph theory.

Resolving set: ordered subset S of vertices such that all vertices have distinct vectors of distances to the vertices in S.

Can k-anonymity be used to prevent active attacks?

A couple of concepts from graph theory.

Resolving set: ordered subset S of vertices such that all vertices have distinct vectors of distances to the vertices in S.

Can k-anonymity be used to prevent active attacks?

A couple of concepts from graph theory.

Resolving set: ordered subset S of vertices such that all vertices have distinct vectors of distances to the vertices in S.

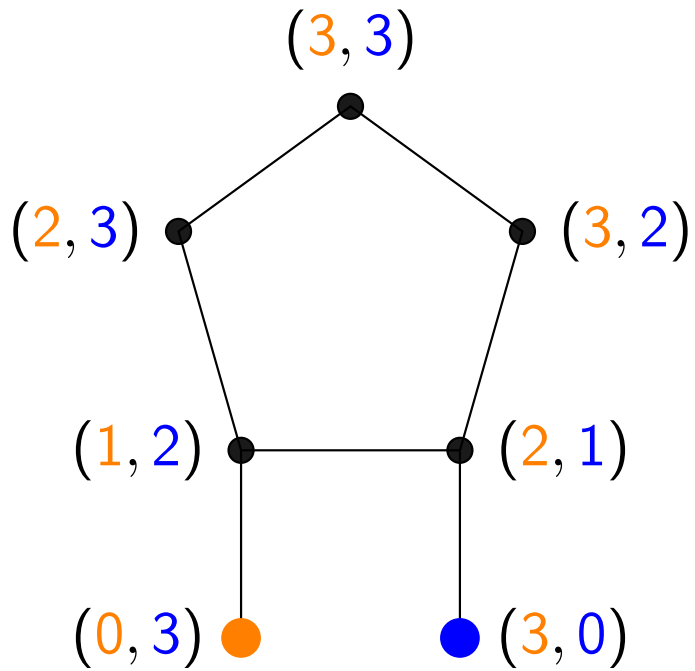

Can k-anonymity be used to prevent active attacks?

- Resolving sets have been used as a model in several applications: navigation of robots in networks, representation of chemical compounds, pattern recognition, mastermind games, amongst others.
- *Metric dimension*: minimum cardinality of a resolving set.
- A resolving set has the ability to uniquely identify every vertex in a graph, as an adversary expects to do

Can k-anonymity be used to prevent active attacks?

- Resolving sets have been used as a model in several applications: navigation of robots in networks, representation of chemical compounds, pattern recognition, mastermind games, amongst others.
- *Metric dimension*: minimum cardinality of a resolving set.
- A resolving set has the ability to uniquely identify every vertex in a graph, as an adversary expects to do

Can k-anonymity be used to prevent active attacks?

- Resolving sets have been used as a model in several applications: navigation of robots in networks, representation of chemical compounds, pattern recognition, mastermind games, amongst others.
- *Metric dimension*: minimum cardinality of a resolving set.
- A resolving set has the ability to uniquely identify every vertex in a graph, as an adversary expects to do

Can k-anonymity be used to prevent active attacks?

- Resolving sets have been used as a model in several applications: navigation of robots in networks, representation of chemical compounds, pattern recognition, mastermind games, amongst others.
- *Metric dimension*: minimum cardinality of a resolving set.
- A resolving set has the ability to uniquely identify every vertex in a graph, as an adversary expects to do

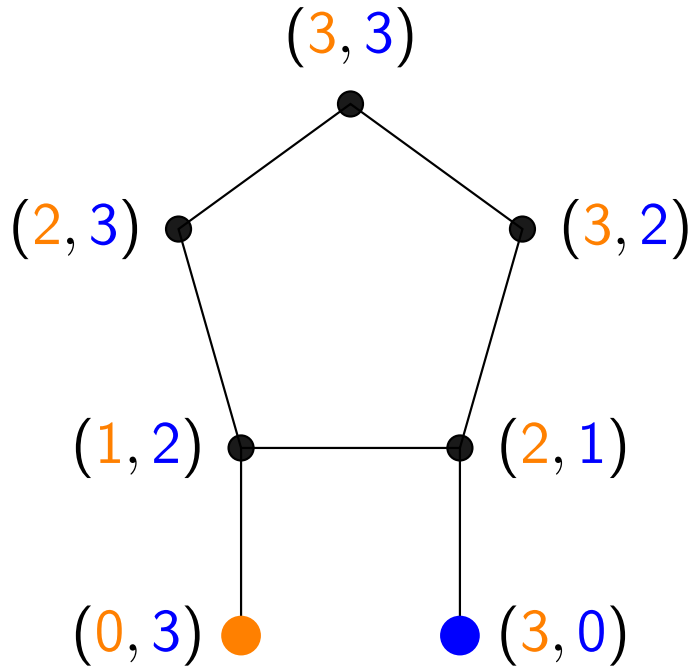

k-antiresolving sets ($k \geq 1$)

k-antiresolving set: Let $G = (V, E)$ be a simple connected graph and let $S = \{u_1, \cdots, u_t\}$ be a subset of vertices of G. The set S is called a *k-antiresolving set* if k is the greatest positive integer such that for every vertex $v \in V - S$ there exist at least $k - 1$ different vertices $v_1, \cdots, v_{k-1} \in V - S$ with $r(v|S) = r(v_1|S) = \cdots = r(v_{k-1}|S)$.

A 2-antiresolving set

k-antiresolving sets ($k \geq 1$)

k-antiresolving set: Let $G = (V, E)$ be a simple connected graph and let $S = \{u_1, \cdots, u_t\}$ be a subset of vertices of G. The set S is called a *k-antiresolving set* if k is the greatest positive integer such that for every vertex $v \in V - S$ there exist at least $k - 1$ different vertices $v_1, \cdots, v_{k-1} \in V - S$ with $r(v|S) = r(v_1|S) = \cdots = r(v_{k-1}|S)$.

A 2-antiresolving set

k-antiresolving sets ($k \geq 1$)

k-antiresolving set: Let $G = (V, E)$ be a simple connected graph and let $S = \{u_1, \cdots, u_t\}$ be a subset of vertices of G. The set S is called a *k-antiresolving set* if k is the greatest positive integer such that for every vertex $v \in V - S$ there exist at least $k - 1$ different vertices $v_1, \cdots, v_{k-1} \in V - S$ with $r(v|S) = r(v_1|S) = \cdots = r(v_{k-1}|S)$.

A 2-antiresolving set

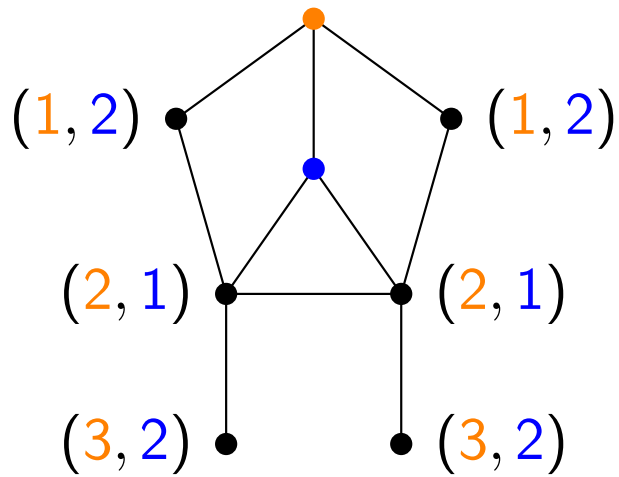

Relating k-antiresolving sets to active attacks

- The attacker controls a set of nodes S in the graph
- The attacker is assumed to know the metric representation (distances) of the target vertices to S
- So, if S is a k-antiresolving set the adversary cannot uniquely re-identify any node in the network with probability higher than $1/k$

Relating k-antiresolving sets to active attacks

- The attacker controls a set of nodes S in the graph
- The attacker is assumed to know the metric representation (distances) of the target vertices to S
- So, if S is a k-antiresolving set the adversary cannot uniquely re-identify any node in the network with probability higher than $1/k$

Relating k-antiresolving sets to active attacks

- The attacker controls a set of nodes S in the graph
- The attacker is assumed to know the metric representation (distances) of the target vertices to S
- So, if S is a k-antiresolving set the adversary cannot uniquely re-identify any node in the network with probability higher than $1/k$

Relating k-antiresolving sets to active attacks

- The attacker controls a set of nodes S in the graph
- The attacker is assumed to know the metric representation (distances) of the target vertices to S
- So, if S is a k-antiresolving set the adversary cannot uniquely re-identify any node in the network with probability higher than $1/k$

Relating k-antiresolving sets to active attacks

The challenge is that potentially any subset can be regarded as the set of attacker nodes. Moreover,

Proposition (Trujillo-Rasua and Yero, 2015)

Given subset of vertices X, we denote $\sim_X$ to the relation satisfying that $u \sim_X v$ if and only if u and v have the same metric representation w.r.t. X. Let S be a subset of vertices and $S' \subset S$, then for every pair of vertices u and v it holds that $u \sim_S v \implies u \sim_{S'} v$.

Definition ((k, ℓ)-anonymity)

G meets (k, ℓ)-*anonymity* with respect to active attacks, if k is the smallest positive integer such that the k-metric antidimension of G is lower than or equal to ℓ, where the *k-metric antidimension* is the minimum cardinality of a k-antiresolving set.

Relating k-antiresolving sets to active attacks

The challenge is that potentially any subset can be regarded as the set of attacker nodes. Moreover,

Proposition (Trujillo-Rasua and Yero, 2015)

Given subset of vertices X, we denote $\sim_X$ to the relation satisfying that $u \sim_X v$ if and only if u and v have the same metric representation w.r.t. X. Let S be a subset of vertices and $S' \subset S$, then for every pair of vertices u and v it holds that
$u \sim_S v \implies u \sim_{S'} v$.

Definition ((k, ℓ)-anonymity)

G meets (k, ℓ)-*anonymity* with respect to active attacks, if k is the smallest positive integer such that the k-metric antidimension of G is lower than or equal to ℓ, where the *k-metric antidimension* is the minimum cardinality of a k-antiresolving set.

Relating k-antiresolving sets to active attacks

The challenge is that potentially any subset can be regarded as the set of attacker nodes. Moreover,

Proposition (Trujillo-Rasua and Yero, 2015)

Given subset of vertices X, we denote $\sim_X$ to the relation satisfying that $u \sim_X v$ if and only if u and v have the same metric representation w.r.t. X. Let S be a subset of vertices and $S' \subset S$, then for every pair of vertices u and v it holds that
$u \sim_S v \implies u \sim_{S'} v$.

Definition ((k, ℓ)-anonymity)

G meets (k, ℓ)-*anonymity* with respect to active attacks, if k is the smallest positive integer such that the k-metric antidimension of G is lower than or equal to ℓ, where the *k-metric antidimension* is the minimum cardinality of a k-antiresolving set.

A simple-yet-unrealistic example

Complete graph K_{12}. Consider three attacker nodes ($\ell = 3$) at most.

Problem: Find smallest k, such that $adim_k(K_{12}) \leq 3$

A simple-yet-unrealistic example

Complete graph K_{12}. Consider three attacker nodes ($\ell = 3$) at most.

Problem: Find smallest k, such that $adim_k(K_{12}) \leq 3$

A simple-yet-unrealistic example

Complete graph K_{12}. Consider three attacker nodes ($\ell = 3$) at most.

Problem: Find smallest k, such that $adim_k(K_{12}) \leq 3$

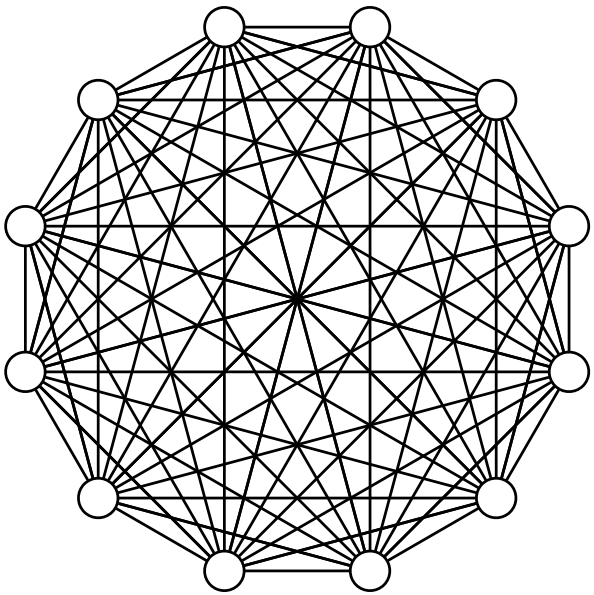

A simple-yet-unrealistic example

Complete graph K_{12}. Consider three attacker nodes ($\ell = 3$) at most.

Problem: Find smallest k, such that $adim_k(K_{12}) \leq 3$

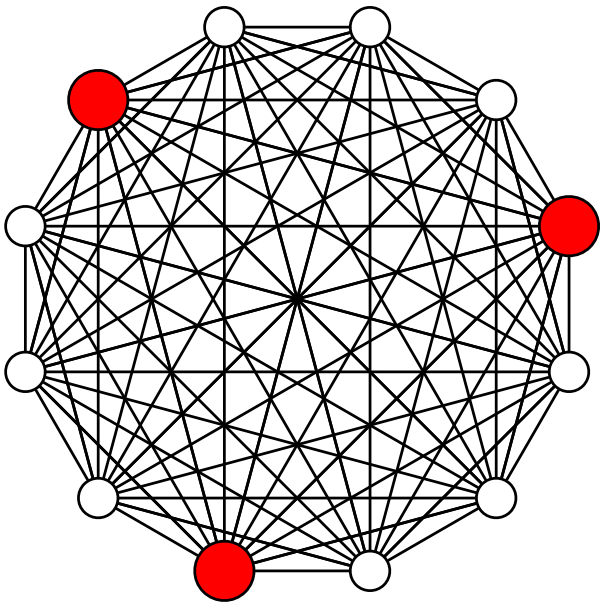

A simple-yet-unrealistic example

Complete graph K_{12}. Consider three attacker nodes ($\ell = 3$) at most.

Problem: Find smallest k, such that $adim_k(K_{12}) \leq 3$

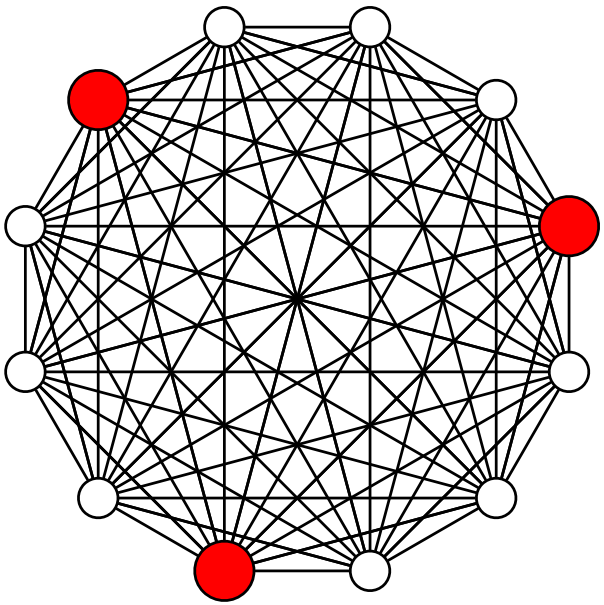

Any subset of vertices of cardinality 3 is a 9-antiresolving set. In general, every subset of vertices of cardinality ℓ is an $(n-\ell)$-antiresolving set in a complete graph. So, the complete graph K_n meets $(n-\ell, \ell)$-anonymity, in particular, $adim_9(K_{12}) = 3$ and K_{12} satisfies $(9, 3)$-anonymity.

Computing (k, ℓ)-anonymity

Ideally, we would like to compute $adim_i(G)$ for every $i \in \{1, \ldots, N\}$ where $N = |V(G)|$ until we find i such that $adim_i(G) \leq \ell$. However, (DasGupta et al. 2016) have proven that the k-metric antidimension problem is NP-hard.

So, transforming a graph into a (k, ℓ)-anonymous graph seems challenging.

Computing (k, ℓ)-anonymity

Ideally, we would like to compute $adim_i(G)$ for every $i \in \{1, \ldots, N\}$ where $N = |V(G)|$ until we find i such that $adim_i(G) \leq \ell$. However, (DasGupta et al. 2016) have proven that the k-metric antidimension problem is NP-hard.

So, transforming a graph into a (k, ℓ)-anonymous graph seems challenging.

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs
- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

Counteracting active attacks: a transformation approach

Observation

- *In terms of offered privacy, (k, ℓ)-anonymity forms a lattice where $(1,1)$-anonymity is the minimum*

Counteracting active attacks: a transformation approach

Observation

- *In terms of offered privacy, (k, ℓ)-anonymity forms a lattice where $(1, 1)$-anonymity is the minimum*

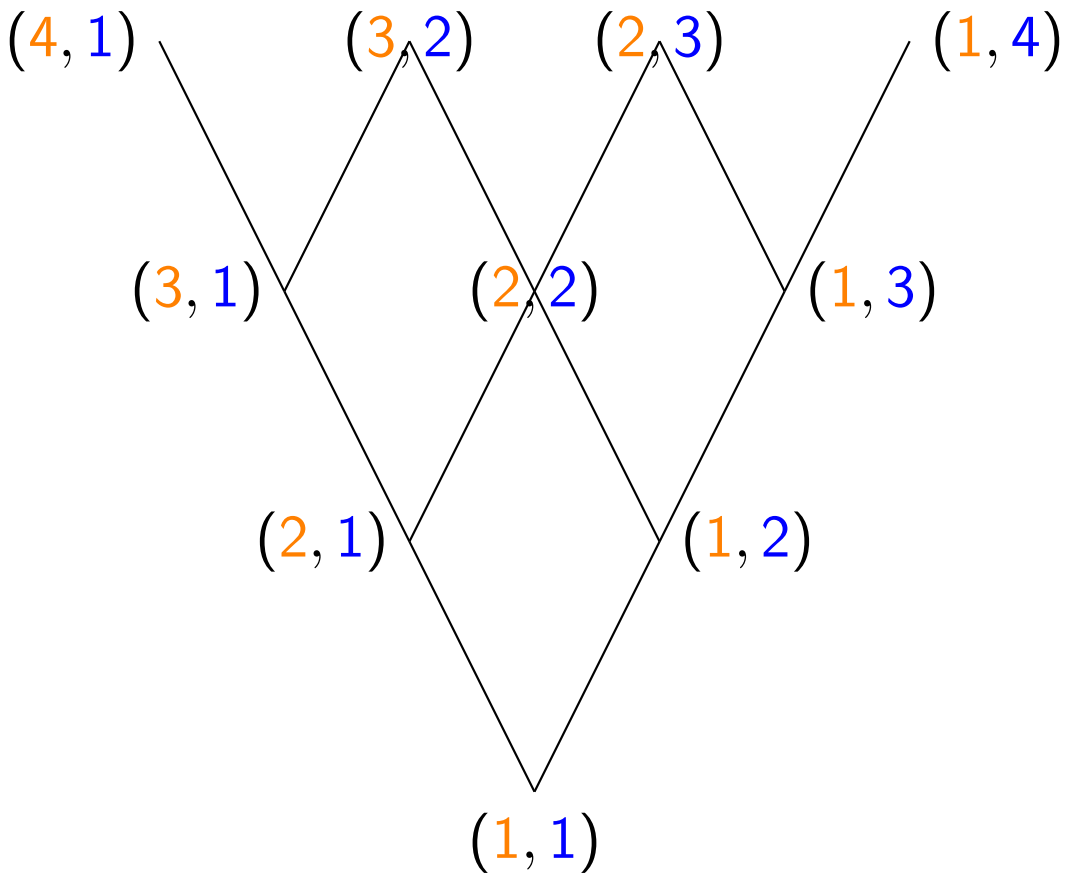

Counteracting active attacks: a transformation approach

Observation

- *And, real life social graphs are typically* $(1, 1)$*-anonymous*
- *This leads to the question whether it is possible to define privacy-preserving transformation techniques to transform a graph* G *into another graph* G' *such that* G' *is not* $(1, 1)$*-anonymous.*

Counteracting active attacks: a transformation approach

Proposition (Mauw et al. 2016)

If G contains a 1-antiresolving set, say $\{v\}$, then there exists a vertex u such that $d(v,u) \neq d(v,w)$ for every $w \in V - \{v,u\}$. We call such a vertex u a 1-resolvable vertex, in particular, we say that u is 1-resolvable by $\{v\}$. It follows that G contains a 1-resolvable vertex if and only if G is $(1,1)$-anonymous.

How to find them?

Lemma (Mauw et al. 2016)

Let $\{v\}$ be a 1-antiresolving set in G, and let $v_1 \cdots v_m$ be an eccentricity path of v, i.e., $v_1 = v$. For every vertex u that is 1-resolvable by $\{v\}$ there exists $i \in \{1, \ldots, m\}$ such that $u = v_i$.

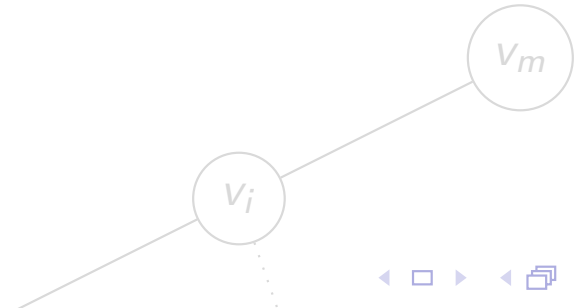

Counteracting active attacks: a transformation approach

How to find them?

Lemma (Mauw et al. 2016)

Let $\{v\}$ be a 1-antiresolving set in G, and let $v_1 \cdots v_m$ be an eccentricity path of v, i.e., $v_1 = v$. For every vertex u that is 1-resolvable by $\{v\}$ there exists $i \in \{1, \ldots, m\}$ such that $u = v_i$.

Counteracting active attacks: a transformation approach

How to find them?

Lemma (Mauw et al. 2016)

Let $\{v\}$ be a 1-antiresolving set in G, and let $v_1 \cdots v_m$ be an eccentricity path of v, i.e., $v_1 = v$. For every vertex u that is 1-resolvable by $\{v\}$ there exists $i \in \{1, \ldots, m\}$ such that $u = v_i$.

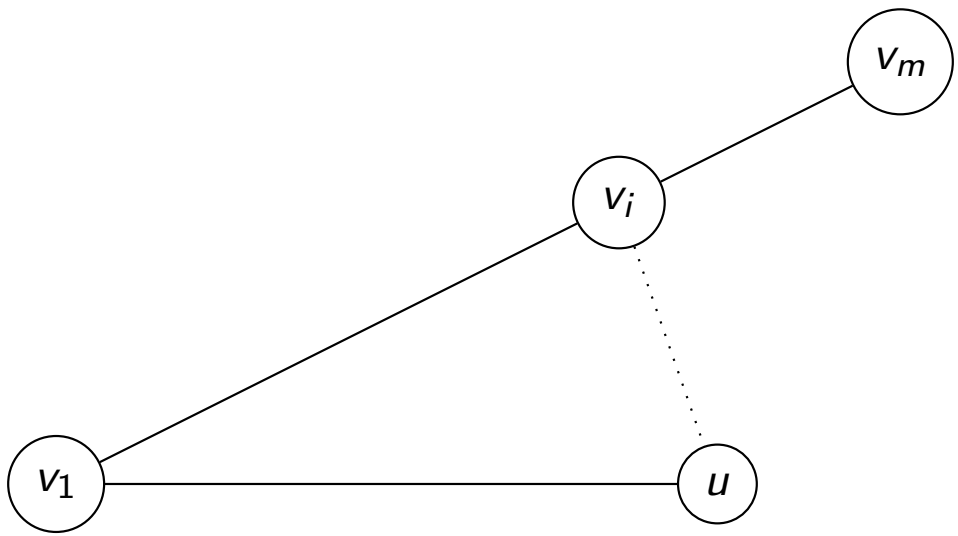

Counteracting active attacks: a transformation approach

How to get rid off 1-resolvable vertices?

Proposition (Mauw et al. 2016)

A cycle graph C_n of odd order doesn't contain 1-resolvable vertices. Indeed, it satisfies $(2,1)$-anonymity.

Counteracting active attacks: a transformation approach

How to get rid off 1-resolvable vertices?

Proposition (Mauw et al. 2016)

A cycle graph C_n of odd order doesn't contain 1-resolvable vertices. Indeed, it satisfies $(2,1)$-anonymity.

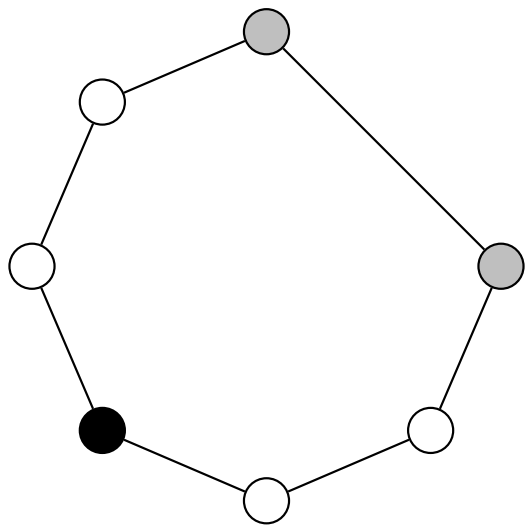

Counteracting active attacks: a transformation approach

Theorem (Mauw et al. 2016)

Let $G = (V, E)$ be a simple connected graph, $\{v\}$ a 1-antiresolving set, and G' the graph resulting from a v-transformation in G. Let S be the set of vertices in G contained in an eccentricity path of v in G. Every $w \in S$ is not 1-resolvable by $\{v\}$ in G'.

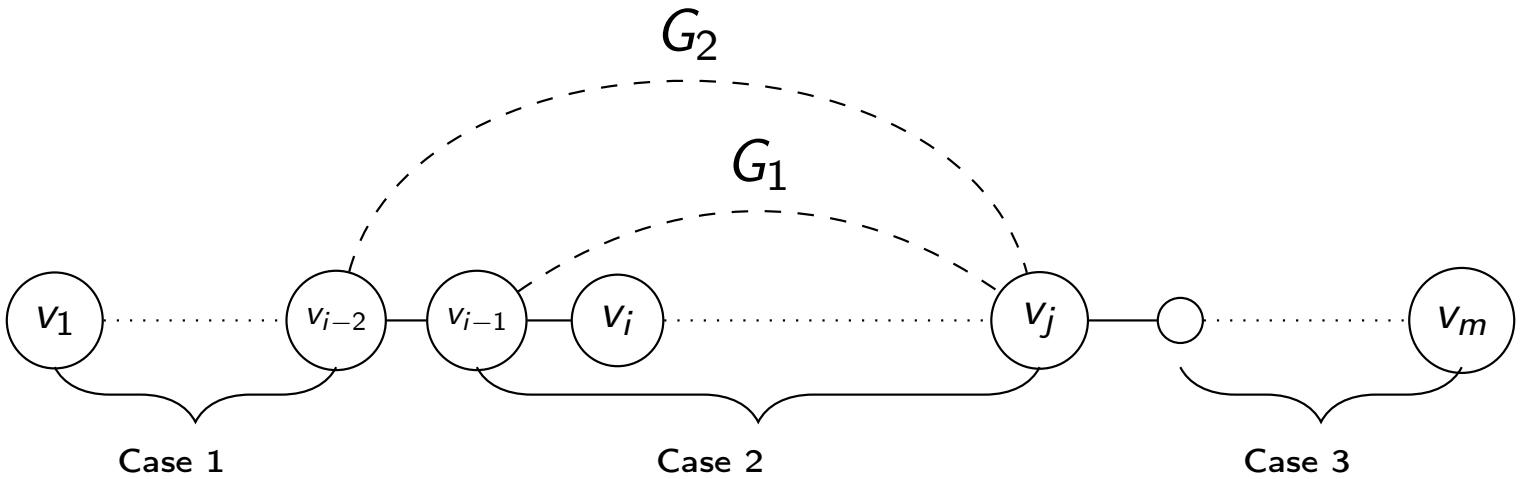

Counteracting active attacks: a transformation approach

Remark: a v-transformation may create new 1-resolvable vertices.

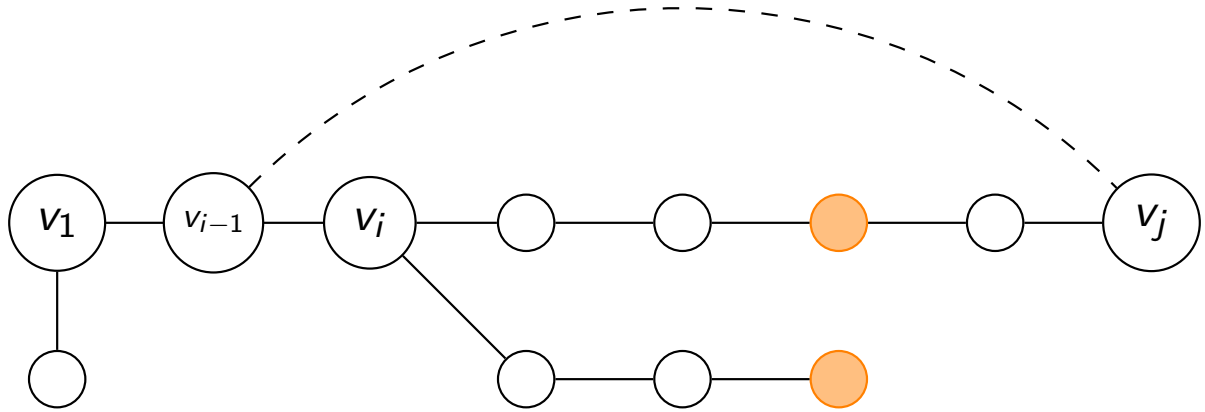

Nevertheless: successive application of v-transformations converge.

It can be found a tight upper bound on the number of v-transformations required to anonymize a $(1, 1)$-anonymous social graph.

Counteracting active attacks: a transformation approach

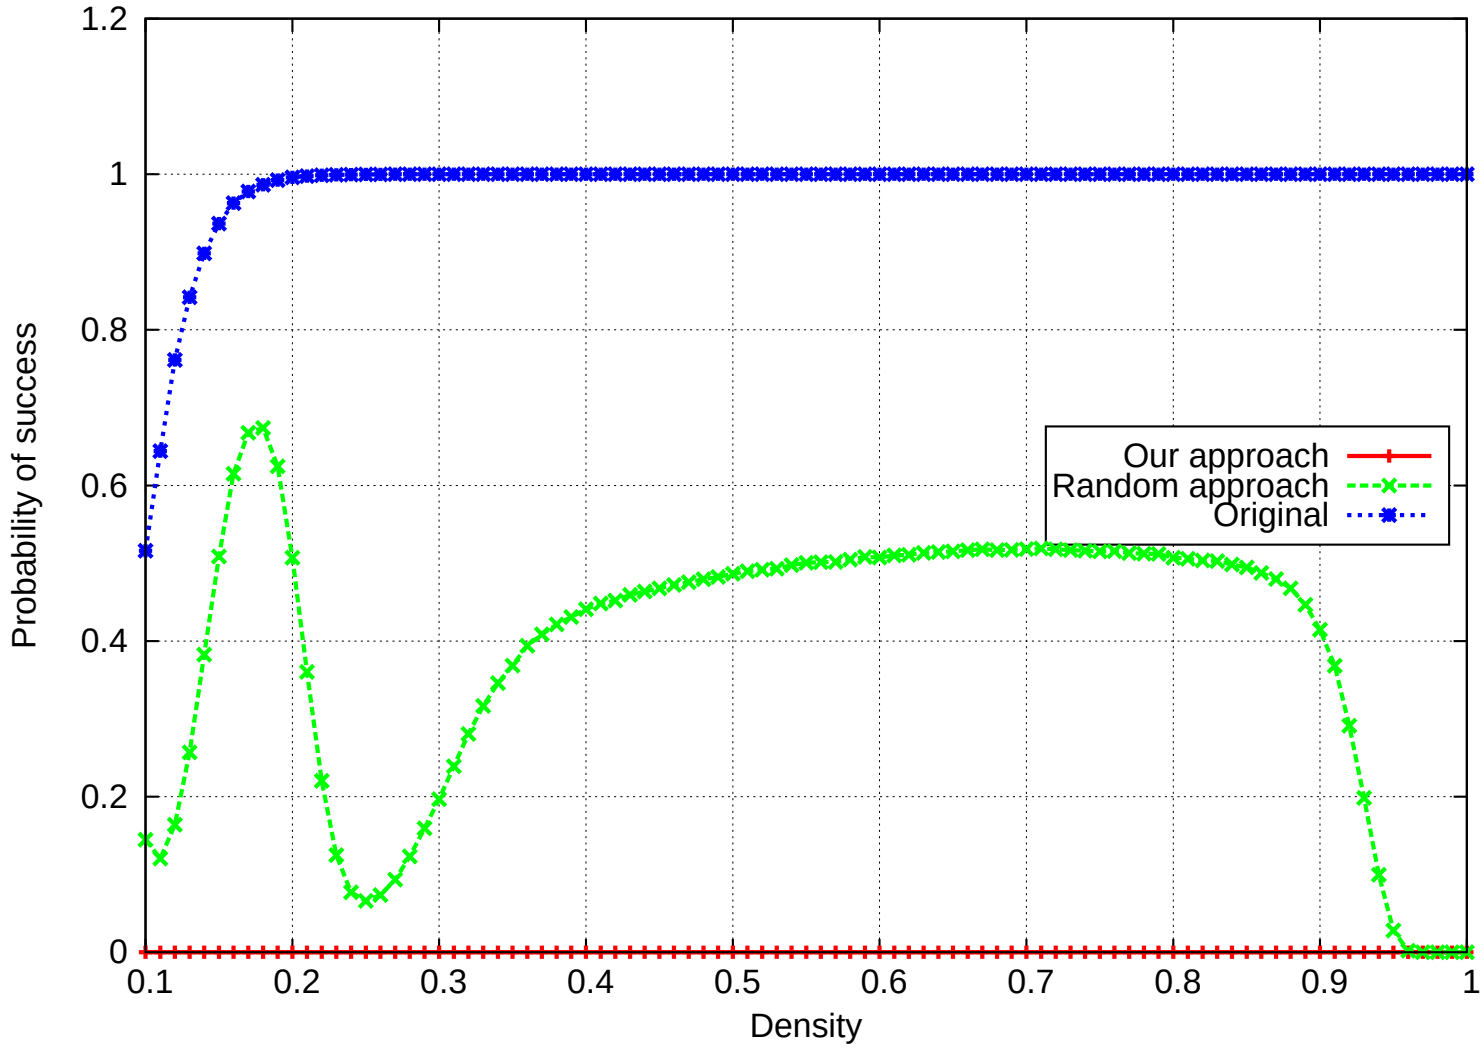

Counteracting active attacks: a transformation approach

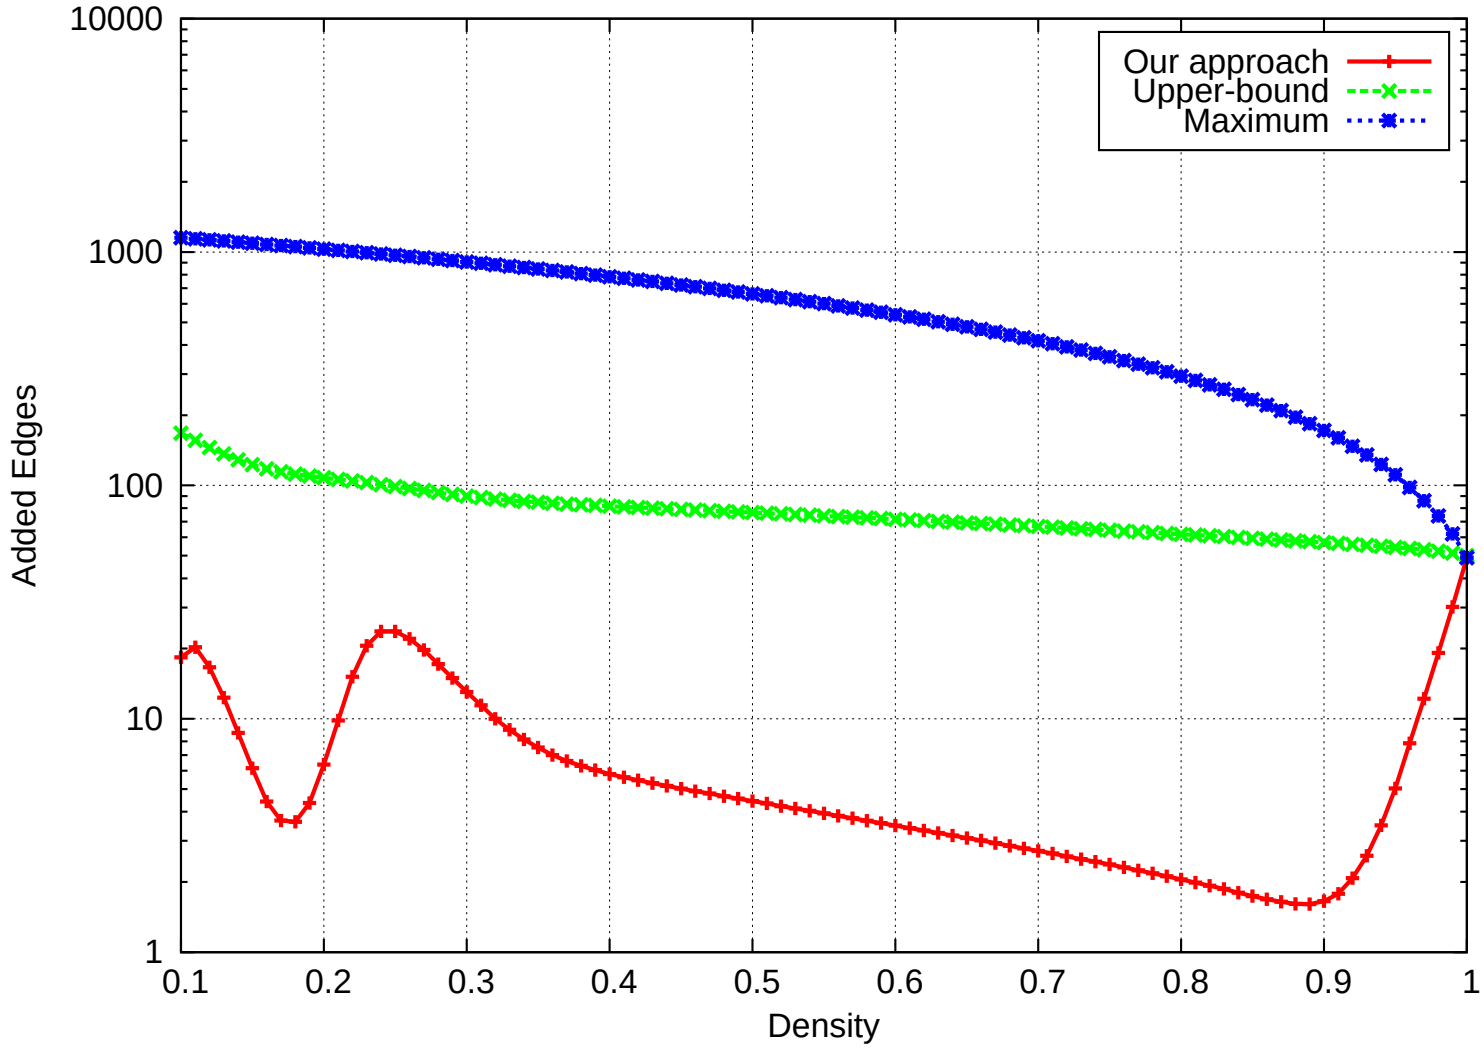

Counteracting active attacks: a transformation approach

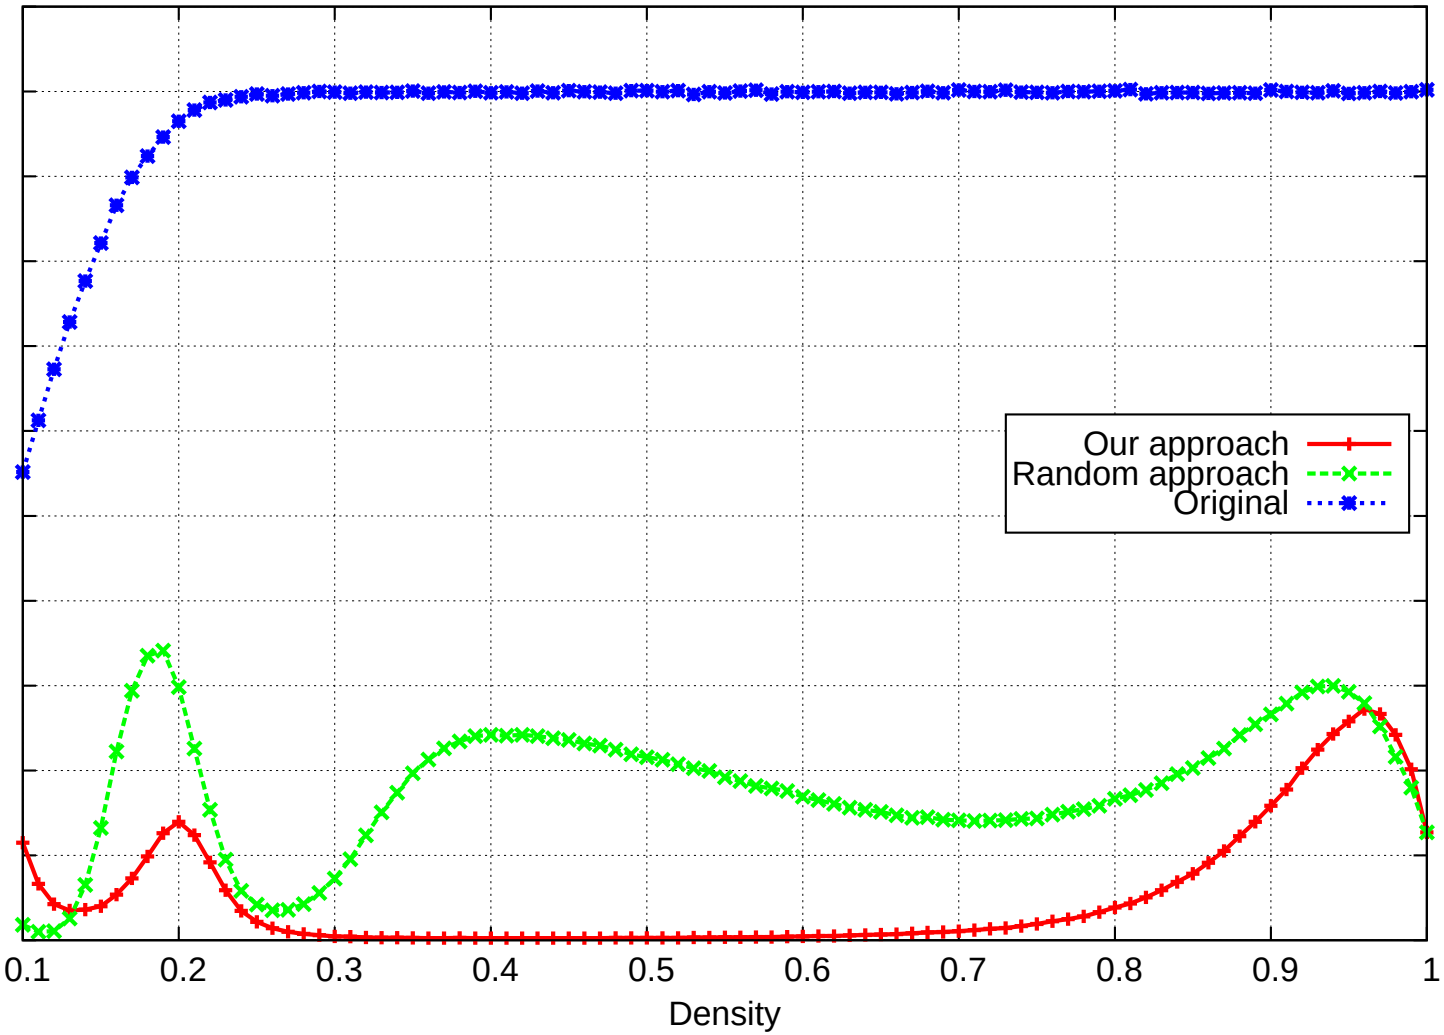

Counteracting active attacks: a transformation approach

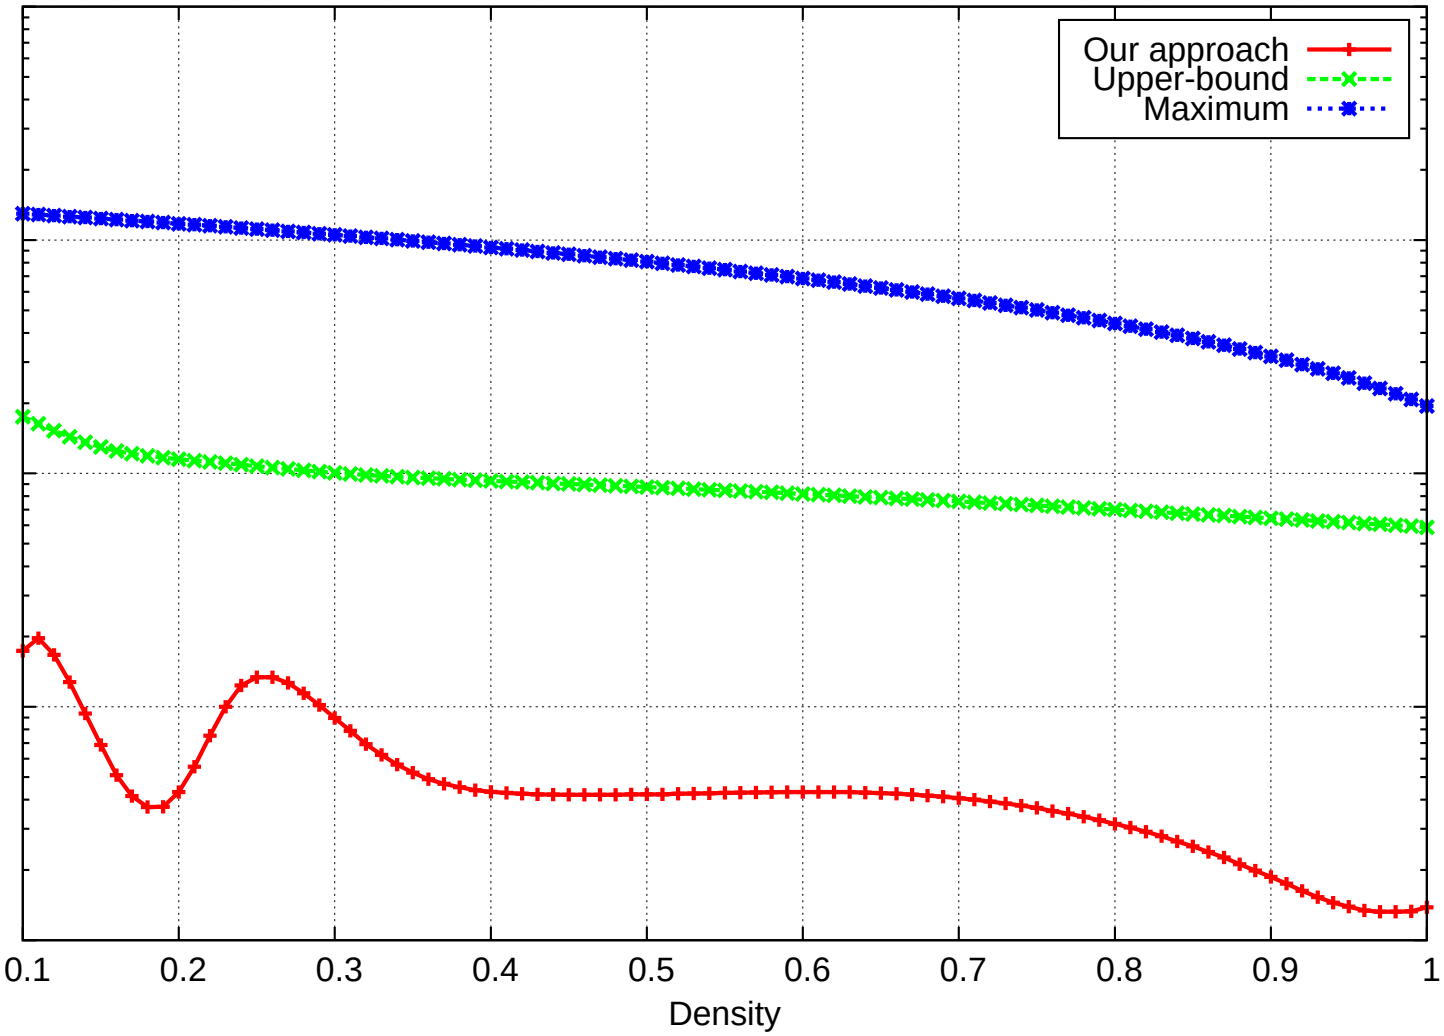

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs
- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

Can we do better?

- Indeed, at very little cost.

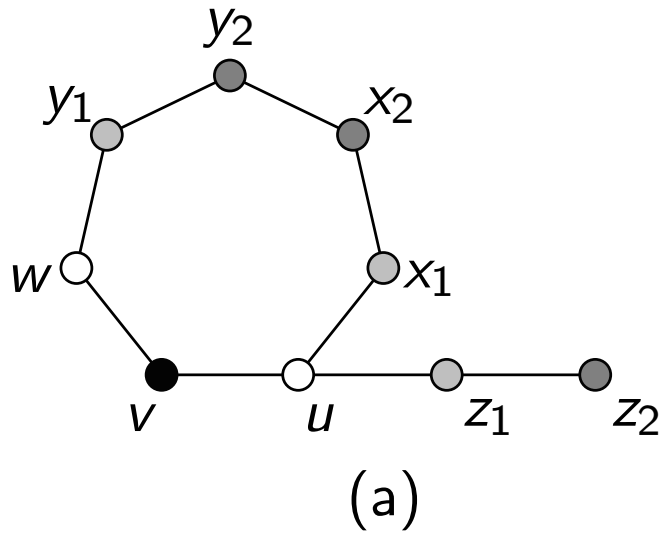

(a)

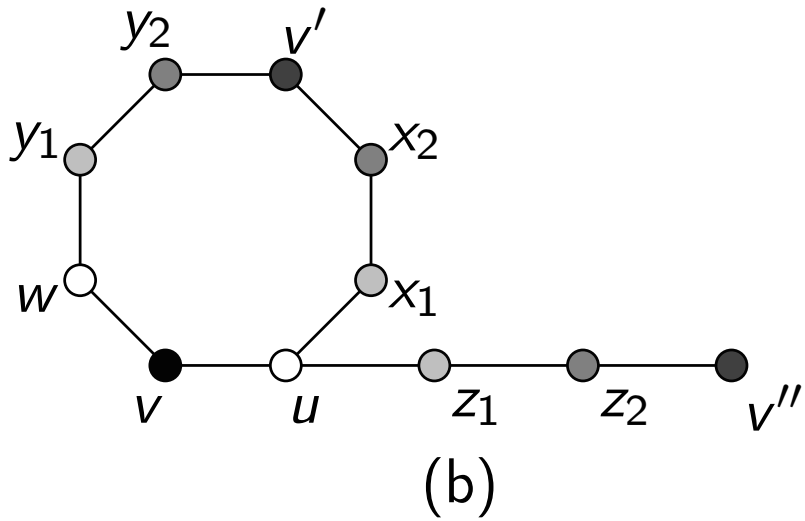

(b)

- What about link prediction
- Isn't (k, ℓ)-anonymity too strong?
  - Can an adversary know beyond the presence or not of a link?
  - Do we really need to protect against all possible 1-antiresolvings?
- Can we revert anonymization?

Can we do better?

- Indeed, at very little cost.

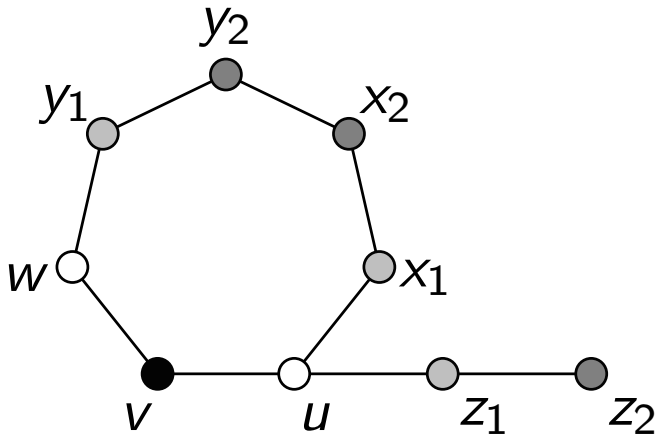

(a)

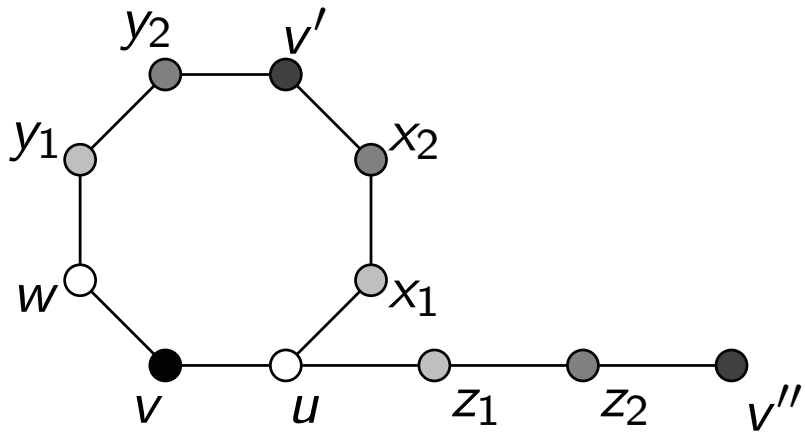

(b)

- What about link prediction
- Isn't (k, ℓ)-anonymity too strong?
  - Can an adversary know beyond the presence or not of a link?
  - Do we really need to protect against all possible 1-antiresolvings?
- Can we revert anonymization?

Outline

The privacy problem: a brief introduction

Privacy-friendly dissemination of personal data

Publication of social graphs
- Passive attacks
- Active attacks

(k, ℓ)-anonymity: a privacy measure

Transforming $(1, 1)$-anonymous graphs: a privacy goal

Future work

Conclusions

Conclusions

- Social graph anonymization is part of a major issue in modern society, that is, how to protect user's privacy in online social networks
- While most anonymization approaches address passive attacks only, we propose the first privacy-preserving transformation method against active attacks
- It remains an open question the relation between (k, ℓ)-anonymity and other privacy notions, such as k-neighbourhood anonymity
- Further empirical evaluation ought to be performed

The end

Thanks.