MEASURING LOCAL TOPOLOGICAL ANONYMITY IN SOCIAL NETWORKS

Gábor György Gulyás and Sándor Imre

Dept. of Telecommunications (BME)

{gulyasg, imre}@hit.bme.hu

PinSoDa: Privacy in Social Data Workshop

in conjunction with the 11th IEEE International Conference on Data Mining (ICDM 2012)

Brussels

December 10, 2012

Anonymous exports and private information?

Auxiliary information, G_{src}
(a public crawl, e.g., Flickr)

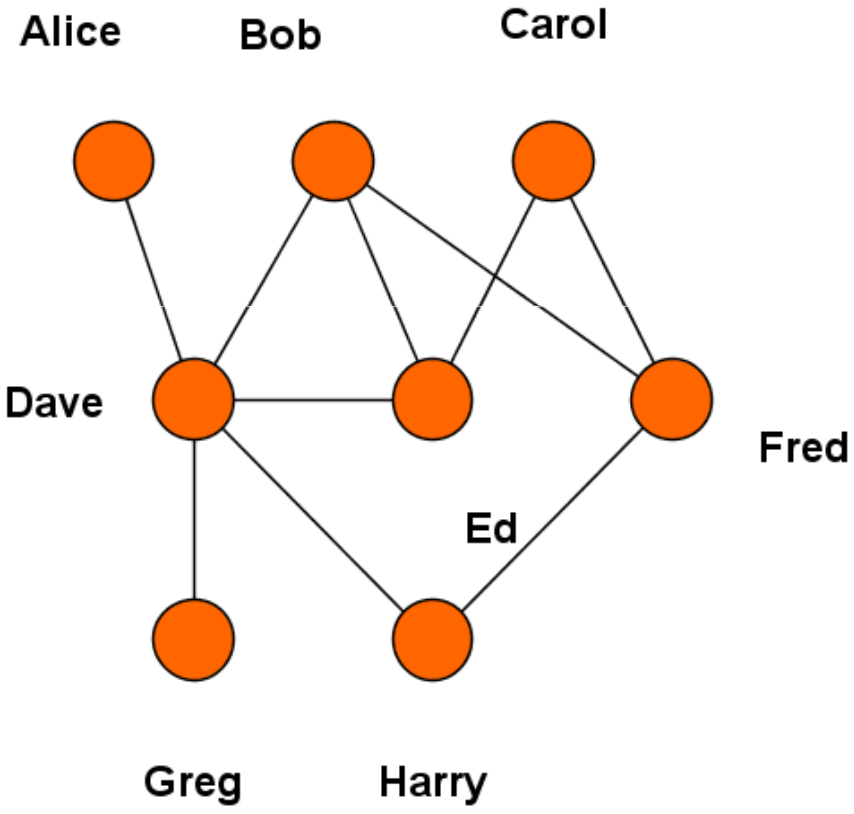

Anonimized graph, G_{tar}
(anonimized export, e.g., Twitter)

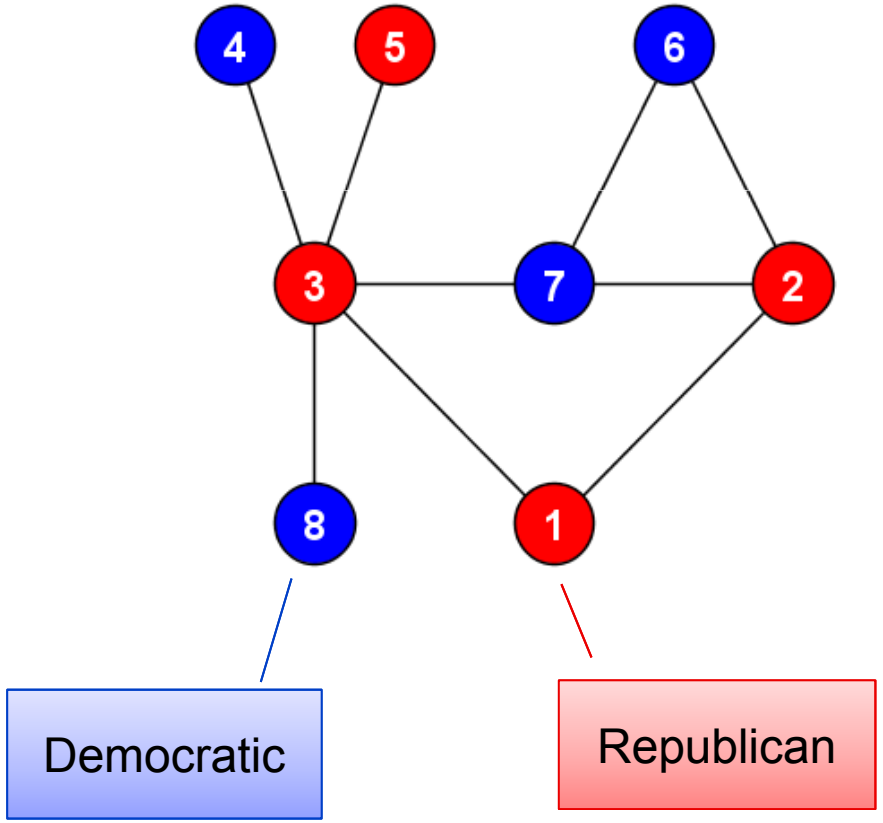

Primary attack types

Active attacks[1]

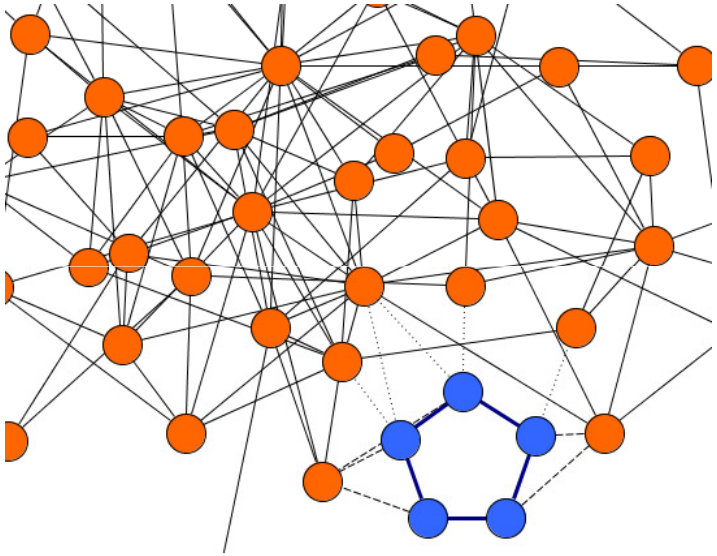

Passive attacks[2]

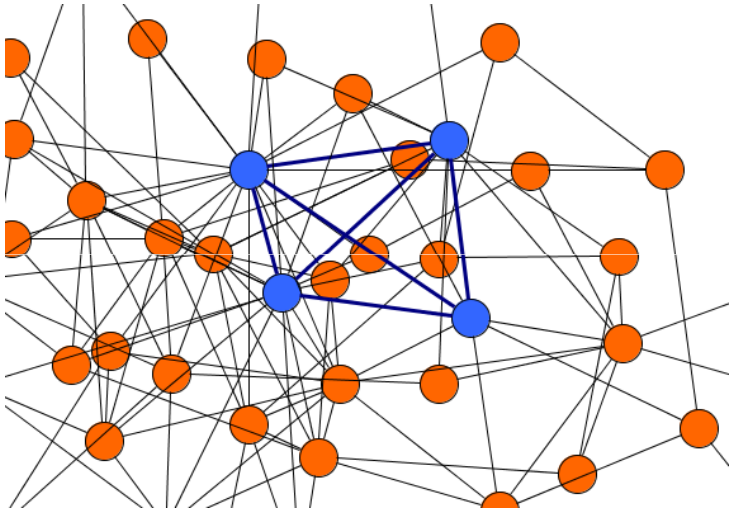

[1] Backstrom et al.: Wherefore Art Thou R3579X? Anonymized Social Networks, Hidden Patterns, and Structural Steganography (2007)
[2] Narayanan & Shmatikov: De-anonymizing social networks (2009)

Primary attack types

Active attacks[1]

Passive attacks[2]

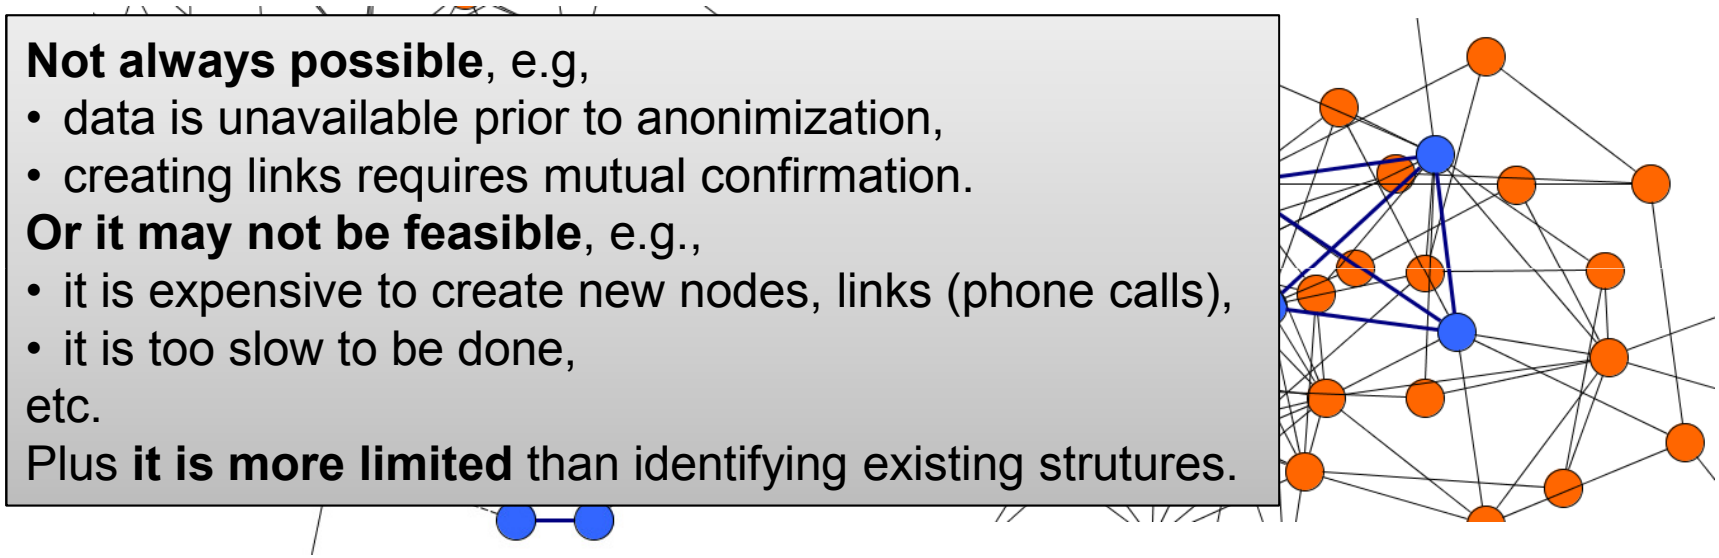

Not always possible, e.g,
- data is unavailable prior to anonimization,
- creating links requires mutual confirmation.

Or it may not be feasible, e.g.,
- it is expensive to create new nodes, links (phone calls),
- it is too slow to be done,

etc.

Plus **it is more limited** than identifying existing strutures.

[1] Backstrom et al.: Wherefore Art Thou R3579X? Anonymized Social Networks, Hidden Patterns, and Structural Steganography (2007)

[2] Narayanan & Shmatikov: De-anonymizing social networks (2009)

Re-identification example

G_{src}:

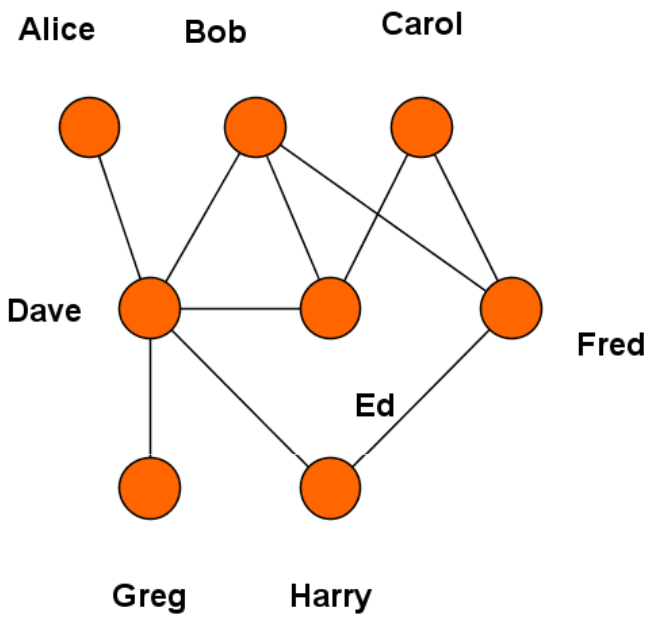

G_{tar}:

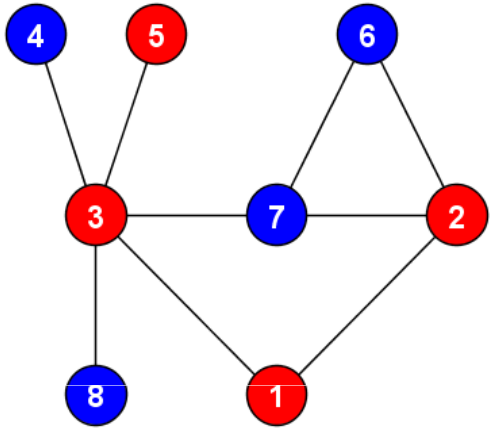

- The attacker uses node degrees.
  Anonymity sets:
  {D} {A, G} {C, H} {B, E, F}
- Dave is *globally* unique:
  Dave $\leftrightarrow$ 3

- But what about Harry?
  - He is in {H, C}
  - Relatively to Dave, i.e., $(D,*) \in E$, anon. sets:
    {A, G}, {B, E}, {H}
  - Harry is *locally* unique:
    Harry $\leftrightarrow$ 1

Global re-identification phase

(a.k.a. seed identification phase)

G_{src}:

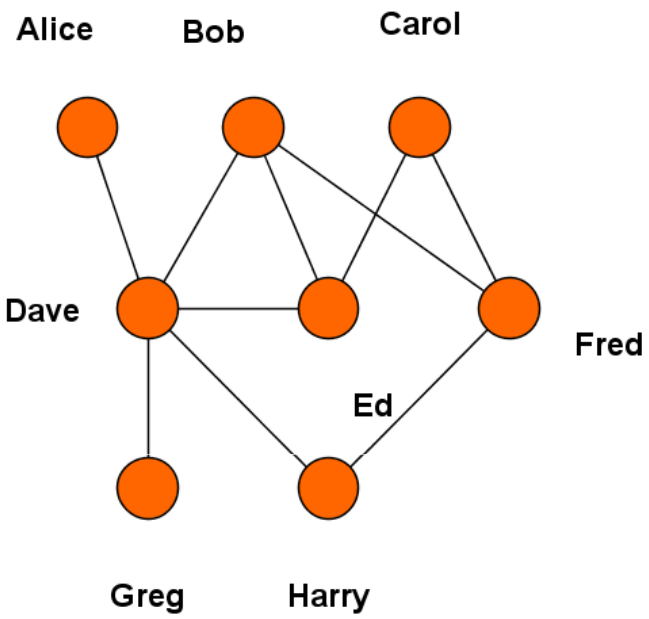

G_{tar}:

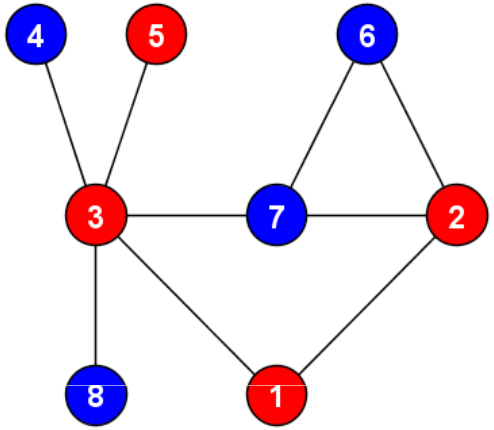

Measuring anonymity:

	A	B	C	D	E	F	G	H
$d(v_i)$	1	3	2	5	3	3	1	2
$A(v_i)$	1/2	2/3	1/2	0	2/3	2/3	1/2	1/2

Anonymity sets: {D} {A, G} {C, H} {B, E, F}

Problems:

- Not feasible for large networks
- Limited de-anonymization rate
- Most of the nodes are not globally outstanding

Local re-identification phase

(a.k.a. propagation identification phase)

G_{src}:

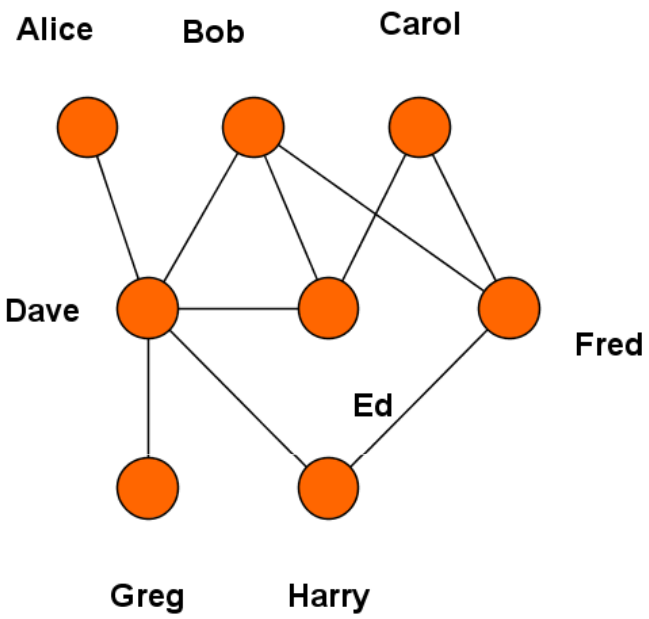

G_{tar}:

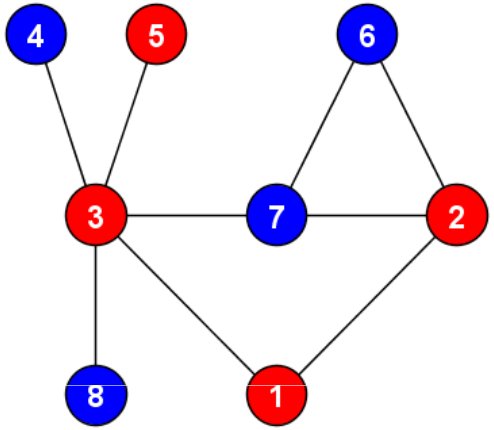

- Problem: anonymity sets depend on seed locations
- How to measure a priori anonymity?

- Local Topological Anonymity (LTA)
  - User: privacy status estimation
  - Data providers (and attackers): estimation of the success of an attack

How do propagation phases work?

- State-of-the-art algorithm:
  Narayanan, A., Shmatikov, V.: De-anonymizing social networks. In: 30th IEEE Symposium on Security and Privacy, pp. 173-187, IEEE Computer Society, Washington (2009)

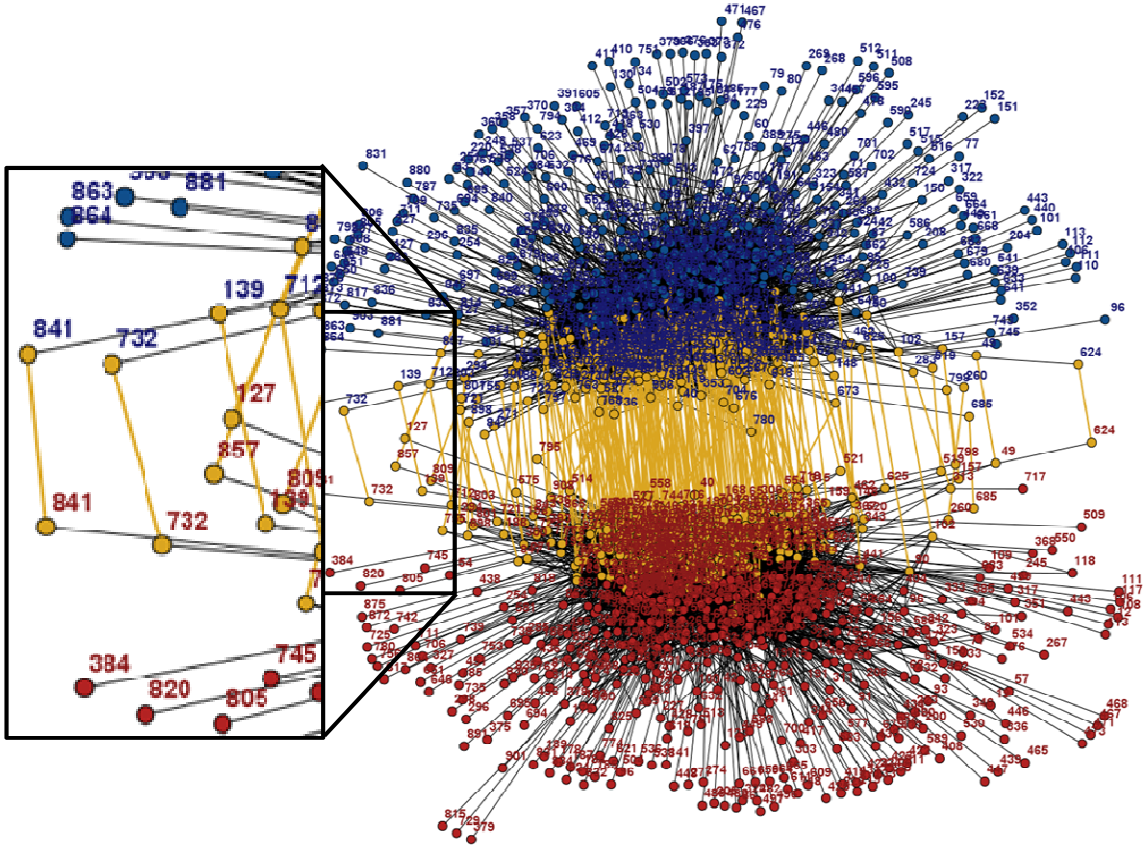

How do propagation phases work? (2)

- Need seeds as an initial mapping ($G_{src} \rightarrow G_{tar}$)
- Round based: tries to extend mapping in each round
  - Unmapped source nodes are structurally compared to unmapped targed nodes
  - Comparison involves their mapped neighbors and their degree values

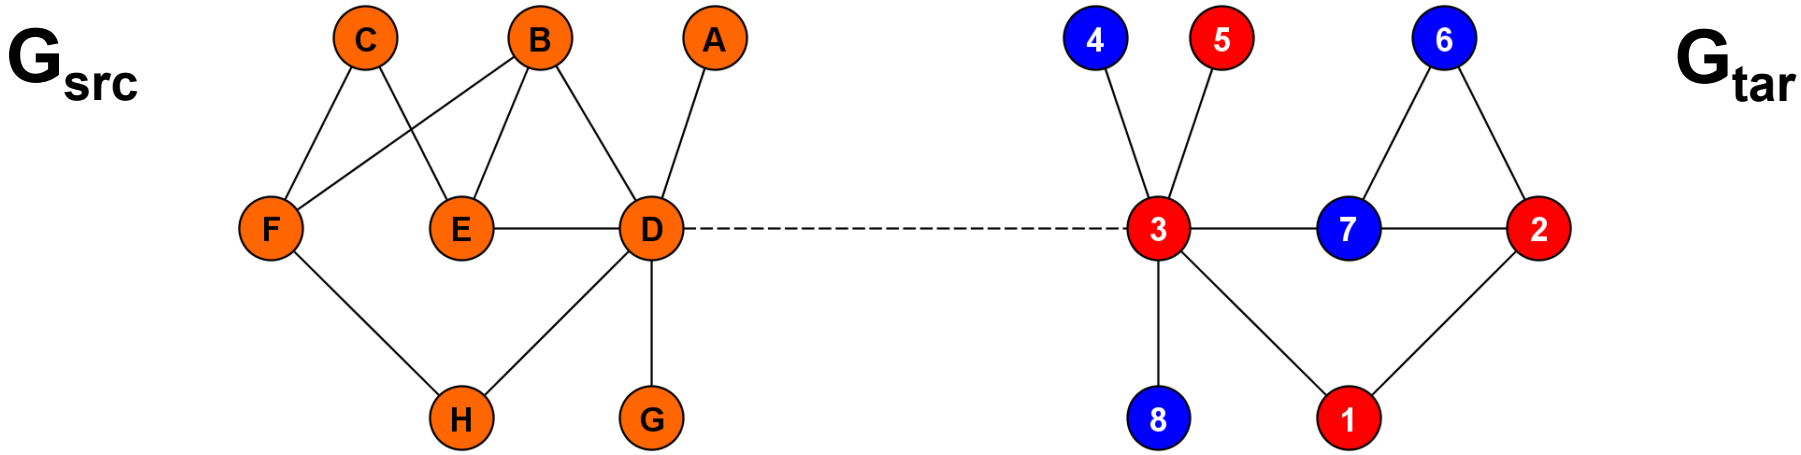

How do propagation phases work? (3)

- Need seeds as an initial mapping ($G_{src} \rightarrow G_{tar}$)
- Round based: tries to extend mapping in each round
  - Unmapped source nodes are structurally compared to unmapped targed nodes
  - Comparison involves their mapped neighbors and their degree values

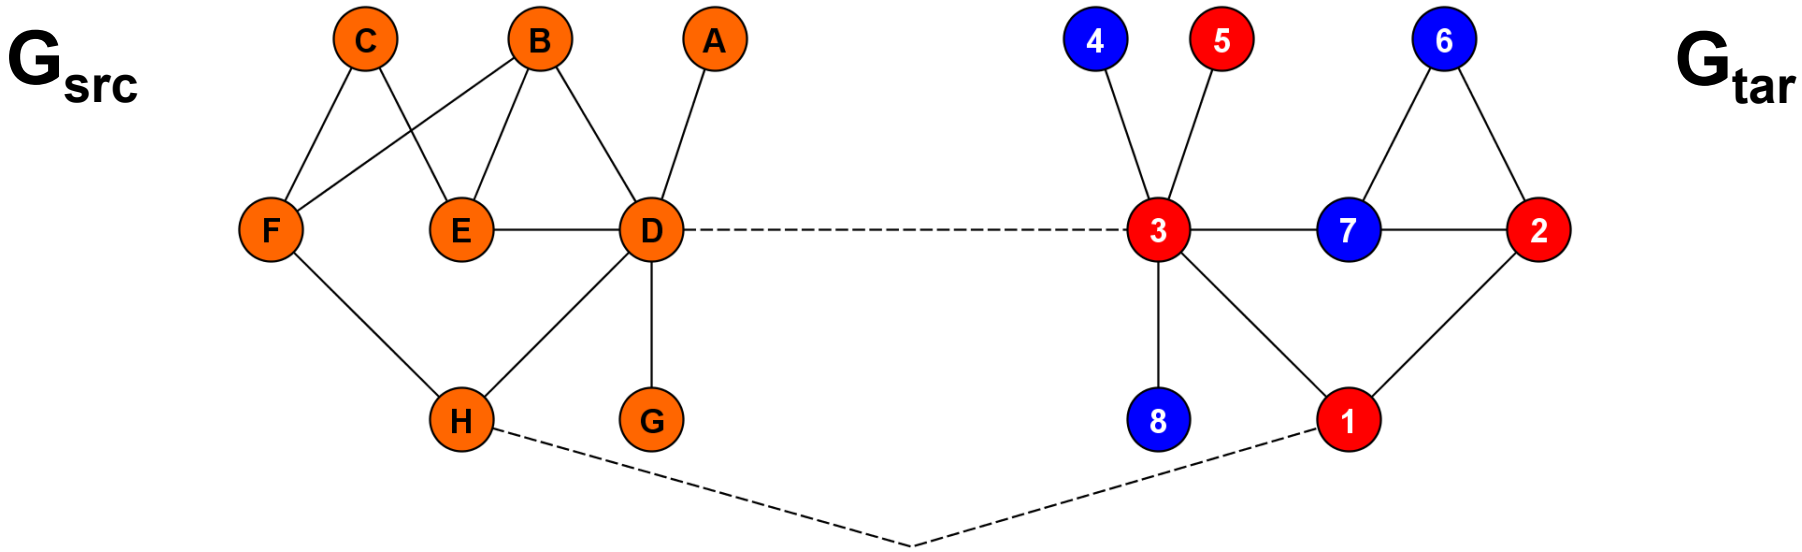

How do propagation phases work? (4)

- Need seeds as an initial mapping ($G_{src} \rightarrow G_{tar}$)
- Round based: tries to extend mapping in each round
  - Unmapped source nodes are structurally compared to unmapped targed nodes
  - Comparison involves their mapped neighbors and their degree values

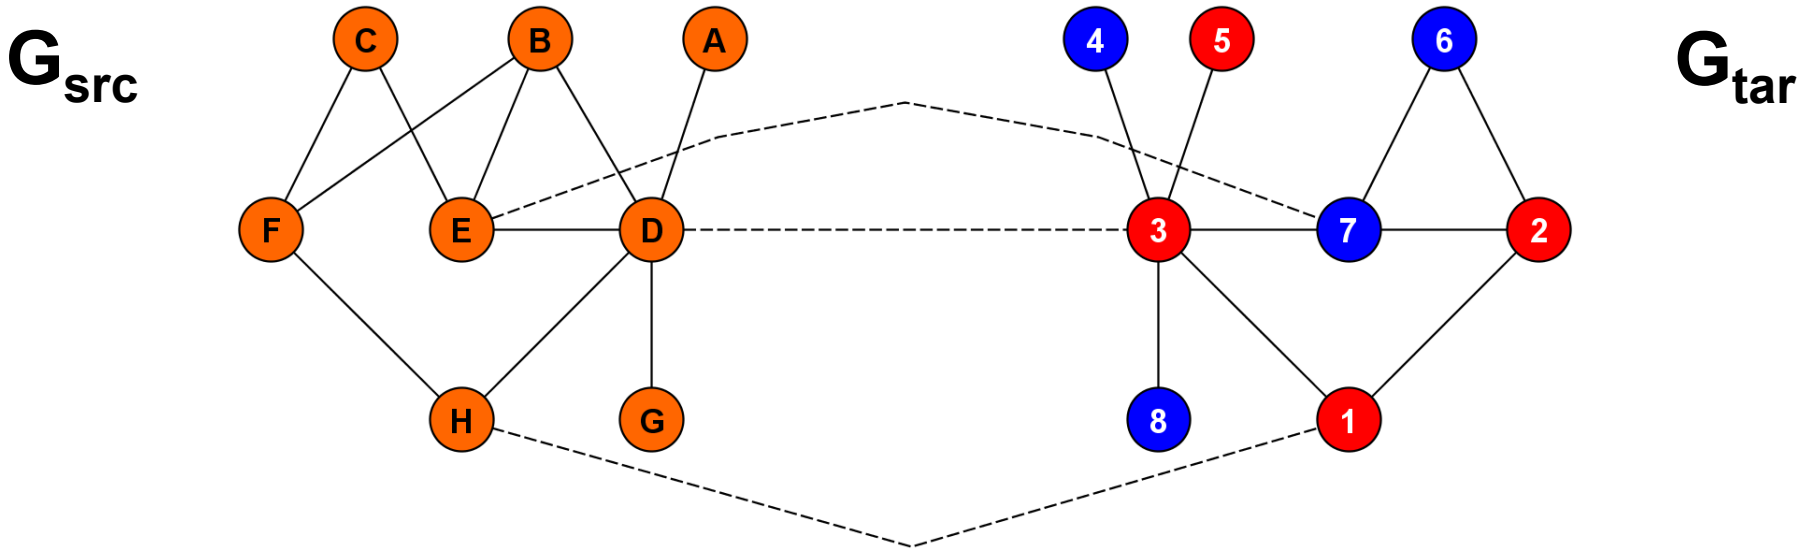

How do propagation phases work? (5)

- Future algorithms are likely to share these principles
- Node comparison yields success if
  - a source node has an instance in the target graph,
  - the source node and its target instance are similar enough,
  - and the target instance is outstanding to its „competitors”.

  ⇒ this property can be captured by an a priori anonymity measure!

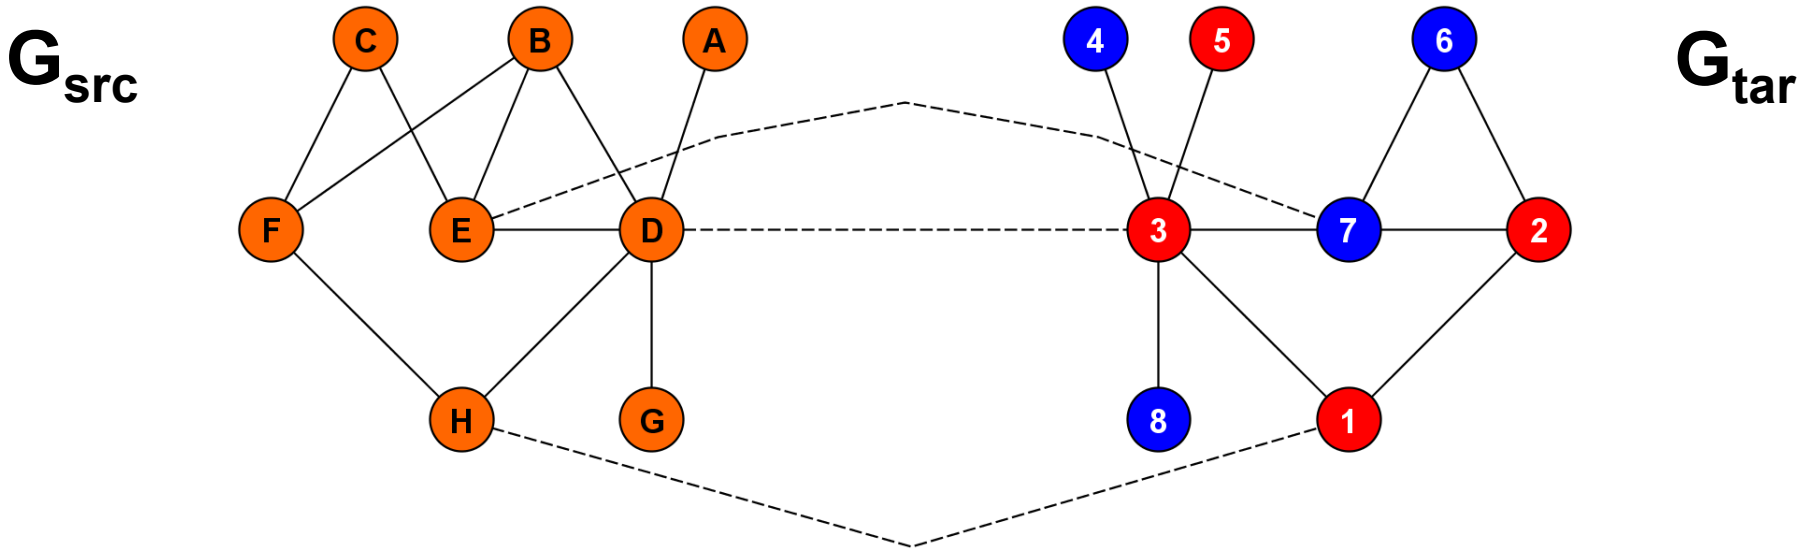

Local Topological Anonymity

- Principle: how v_i is structurally hidden in its 2-neighborhood
  - i.e., how similar v_i is to its neighbors of neighbors
- Proposed metrics:

$$LTA_A(v_i) = \sum_{\forall v_k \in V_i^2} \frac{sim(v_i, v_k)}{|V_i^2|}$$

$$LTA_B(v_i) = \sum_{\forall v_k \in V_i^2} \frac{sim(v_i, v_k)}{\max(|V_i|, 2)}$$

$$LTA_C(v_i) = \sum_{\forall v_k \in V_i^2} \frac{sim(v_i, v_k)}{|V_i^2| \cdot \max(\sigma_{deg}(\Delta V_i^2), 1)}$$

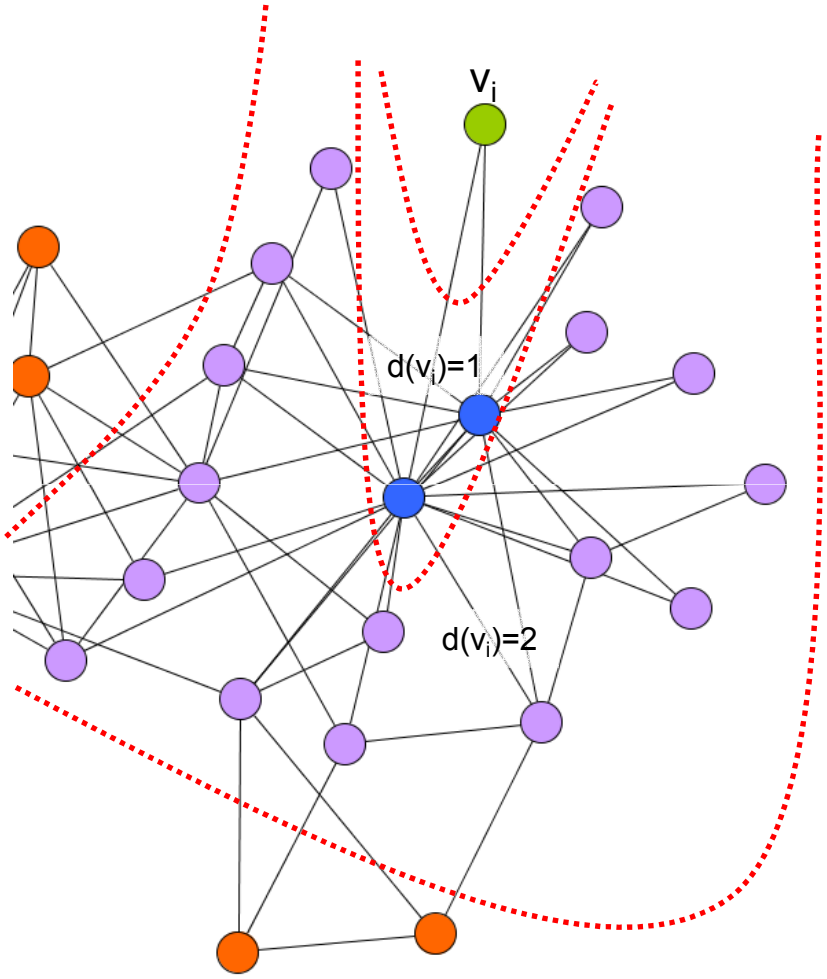

Local Topological Anonymity (2)

- Similarity metric?
  - State-of-the-art attack is based on cosine similarity
  - CosSim produced best results in the comparison of similarity metrics
    - (our comparison & Spertus et al, 2005)
- Simulations: CosSim()

$$\mathrm{CosSim}(v_i, v_k) = \frac{|V_i \cap V_k|}{\sqrt{|V_i| \cdot |V_k|}}$$

  - For the state-of-the-art attack
  - Other attack → different metric

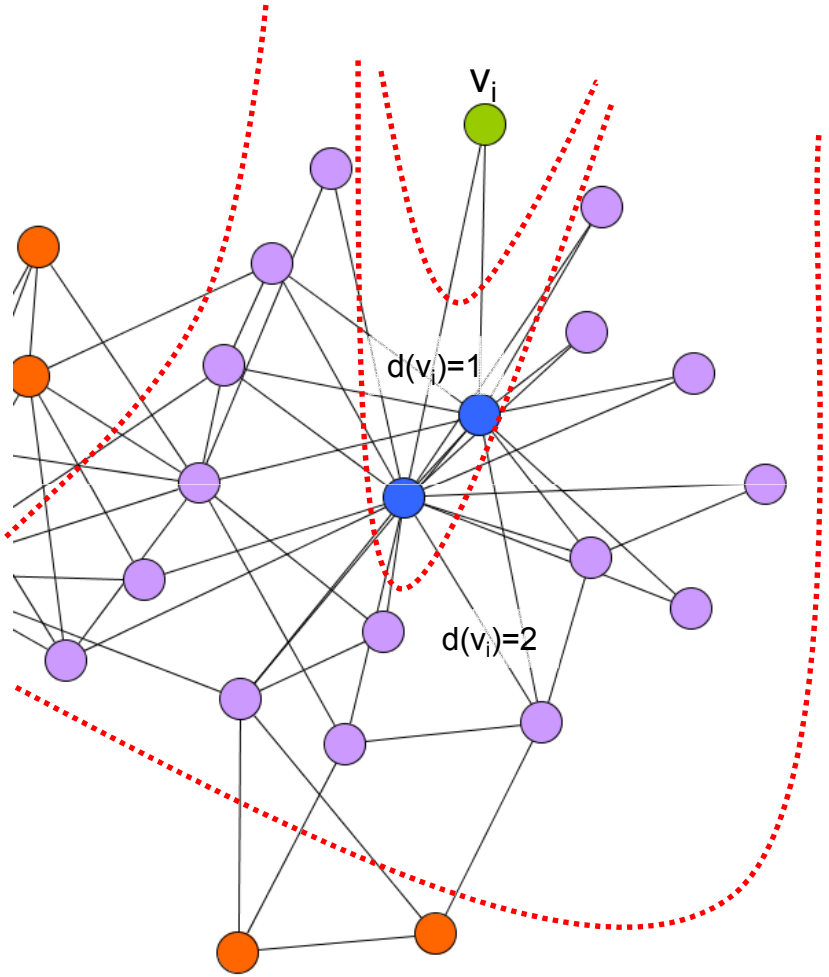

Spertus et al.: Evaluating similarity measures: a large-scale study in the orkut social network. (2005)

Visual comparison for small nets

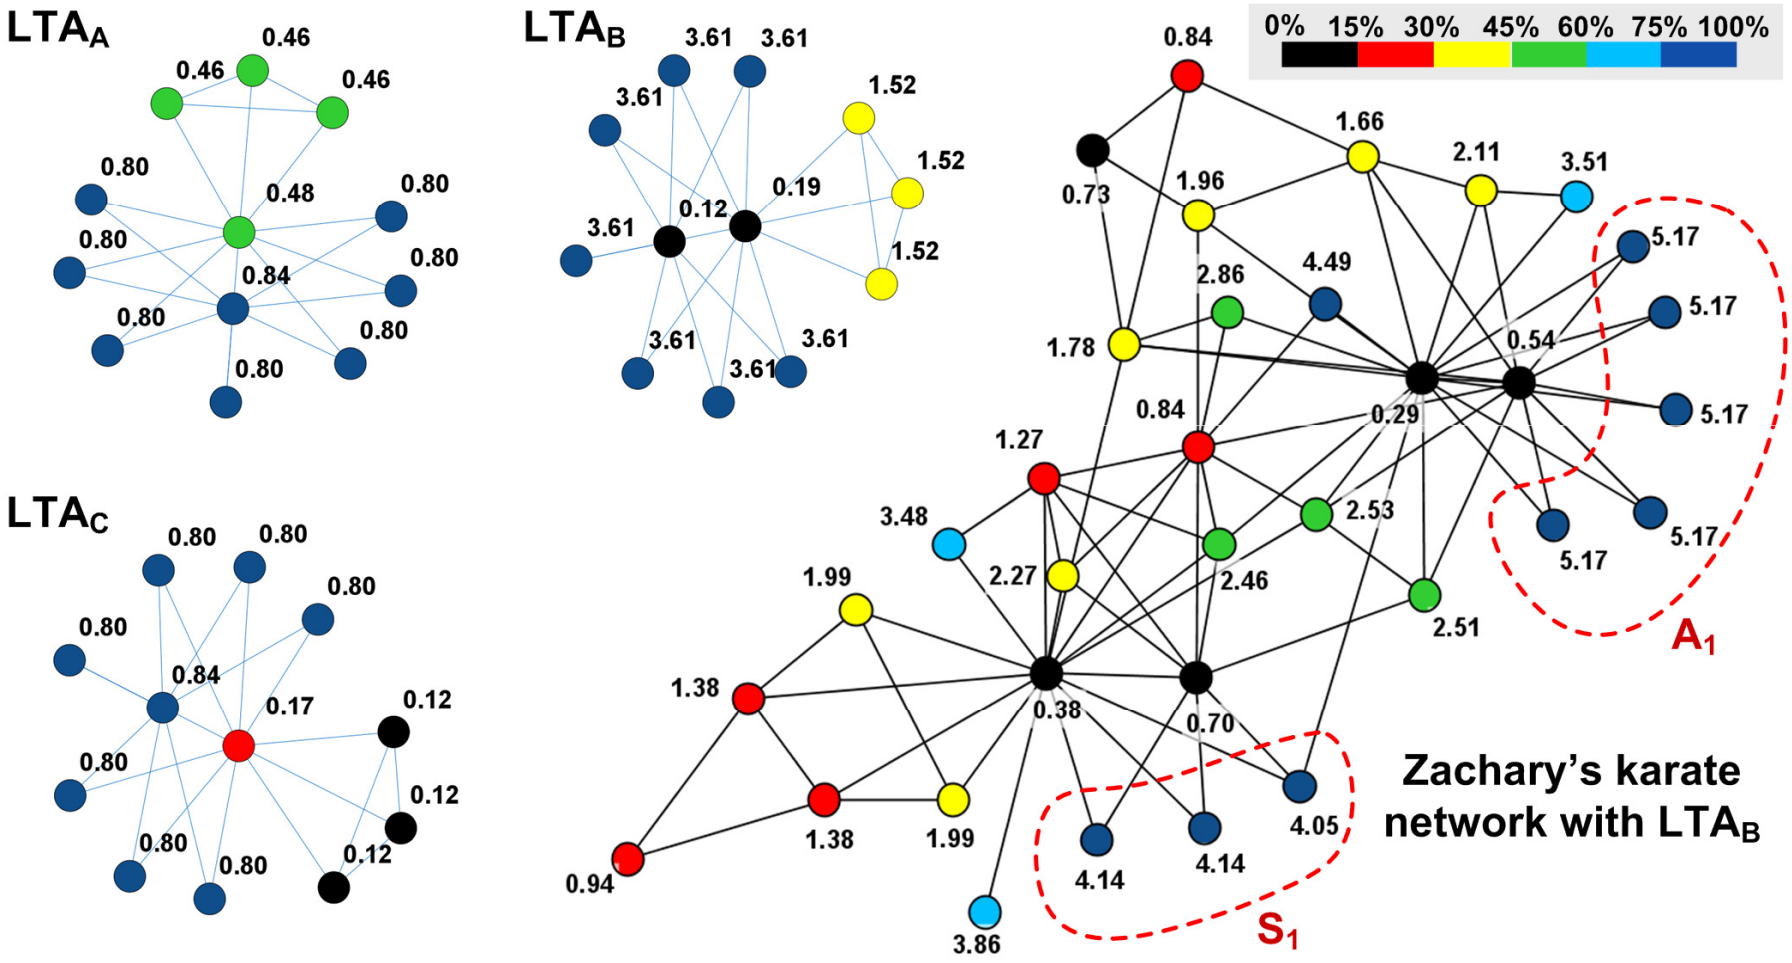

Evaluation methodology

- Simulational evaluation: attack results vs. LTA prediction
  - State-of-the-art attack (**)
  - 10 rounds to avoid seed dependencies (e.g., location)
  - Strong attacker
- Dataset sources: Slashdot, Wikivote, Epinions (*) LiveJournal (our crawl)
- Realistic test data (**)
  - Given overlap factors: $\alpha_V = Jacc(V_{src}, V_{tar})$ $\alpha_E = Jacc(E_{src}, E_{tar})$

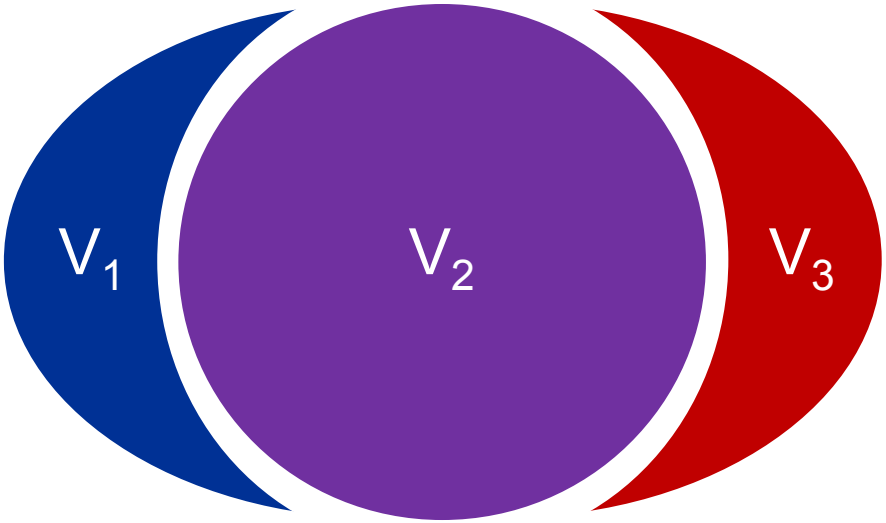

* Source: http://snap.stanford.edu/data/index.html
** Narayanan & Shmatikov: De-anonymizing social networks (2009)

Dataset generation

Network	Nodes	Edges	Density	Diam.	Avg. path length
EP-1000	1,000	29,509	0.0591	4	2.1746
WV-1000	910	9,407	0.0227	5	2.7708
SD-1000	1,104	10,348	0.0170	5	2.4295
LJ-1000	1,033	10,521	0.0197	4	2.5608
LJ-10K	10,056	231,416	0.0046	6	2.8291
WV-Full	7,115	100,762	0.0040	7	3.2475

Dataset generation (2)

Network	Nodes	Edges	Density	Diam.	Avg. path length
EP-1000	1,000	29,509	0.0591	4	2.1746
WV-1000	910	9,407	0.0227	5	2.7708
SD-1000	1,104	10,348	0.		
LJ-1000	1,033	10,521	0.		
LJ-10K	10,056	231,416	0.		
WV-Full	7,115	100,762	0.		

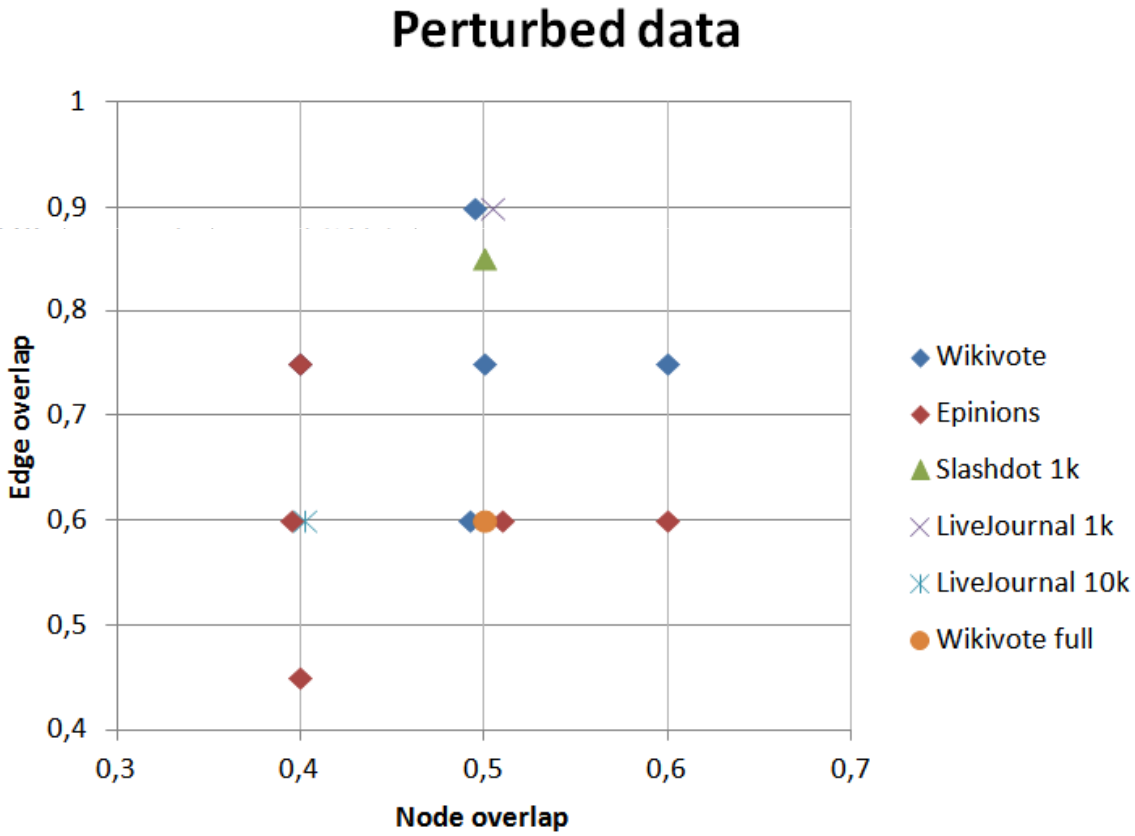

Simulational results of LJ-10KvA

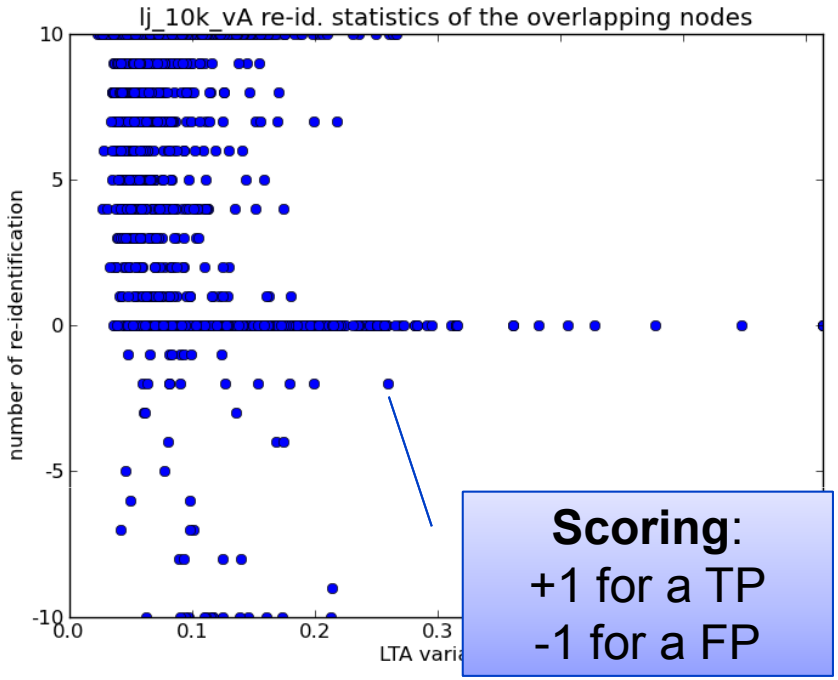

Scoring:
+1 for a TP
-1 for a FP

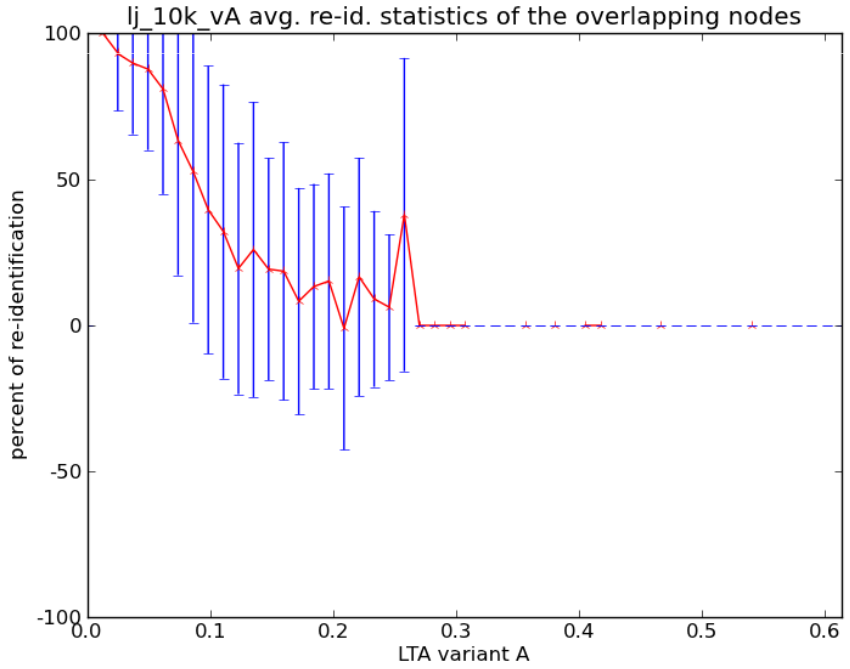

Simulational results of LJ-10KvA (2)

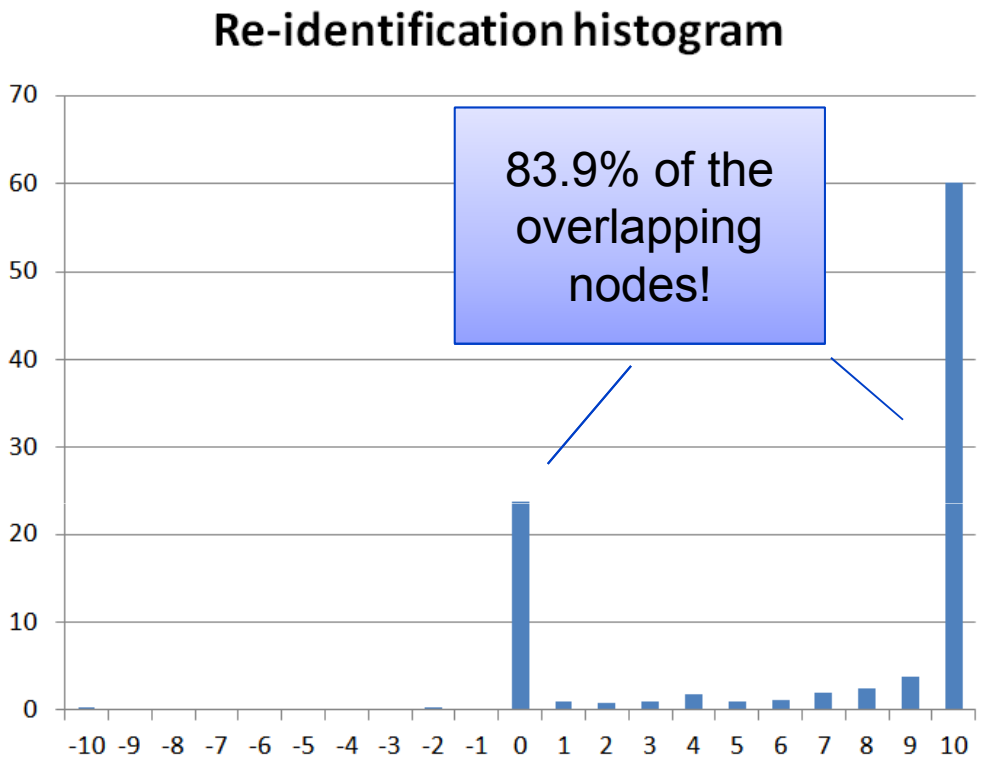

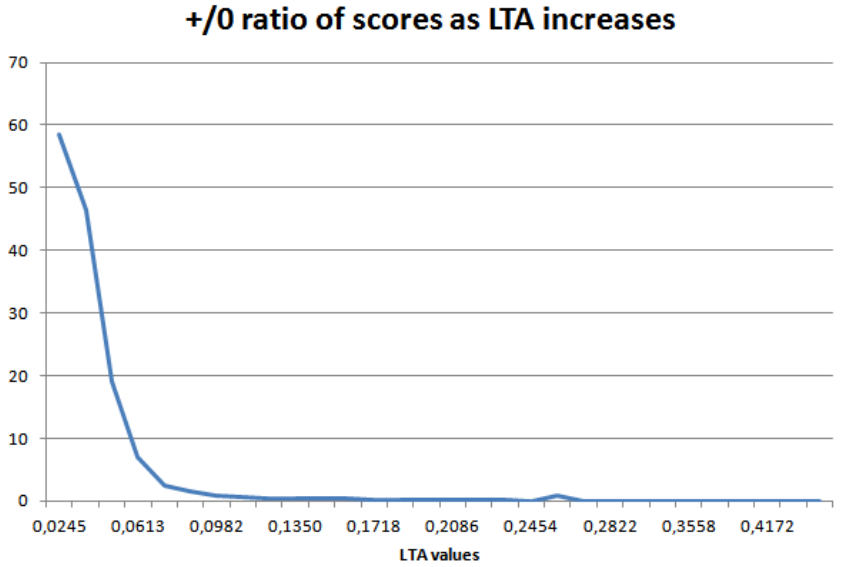

Pearson correlation of LTA and results

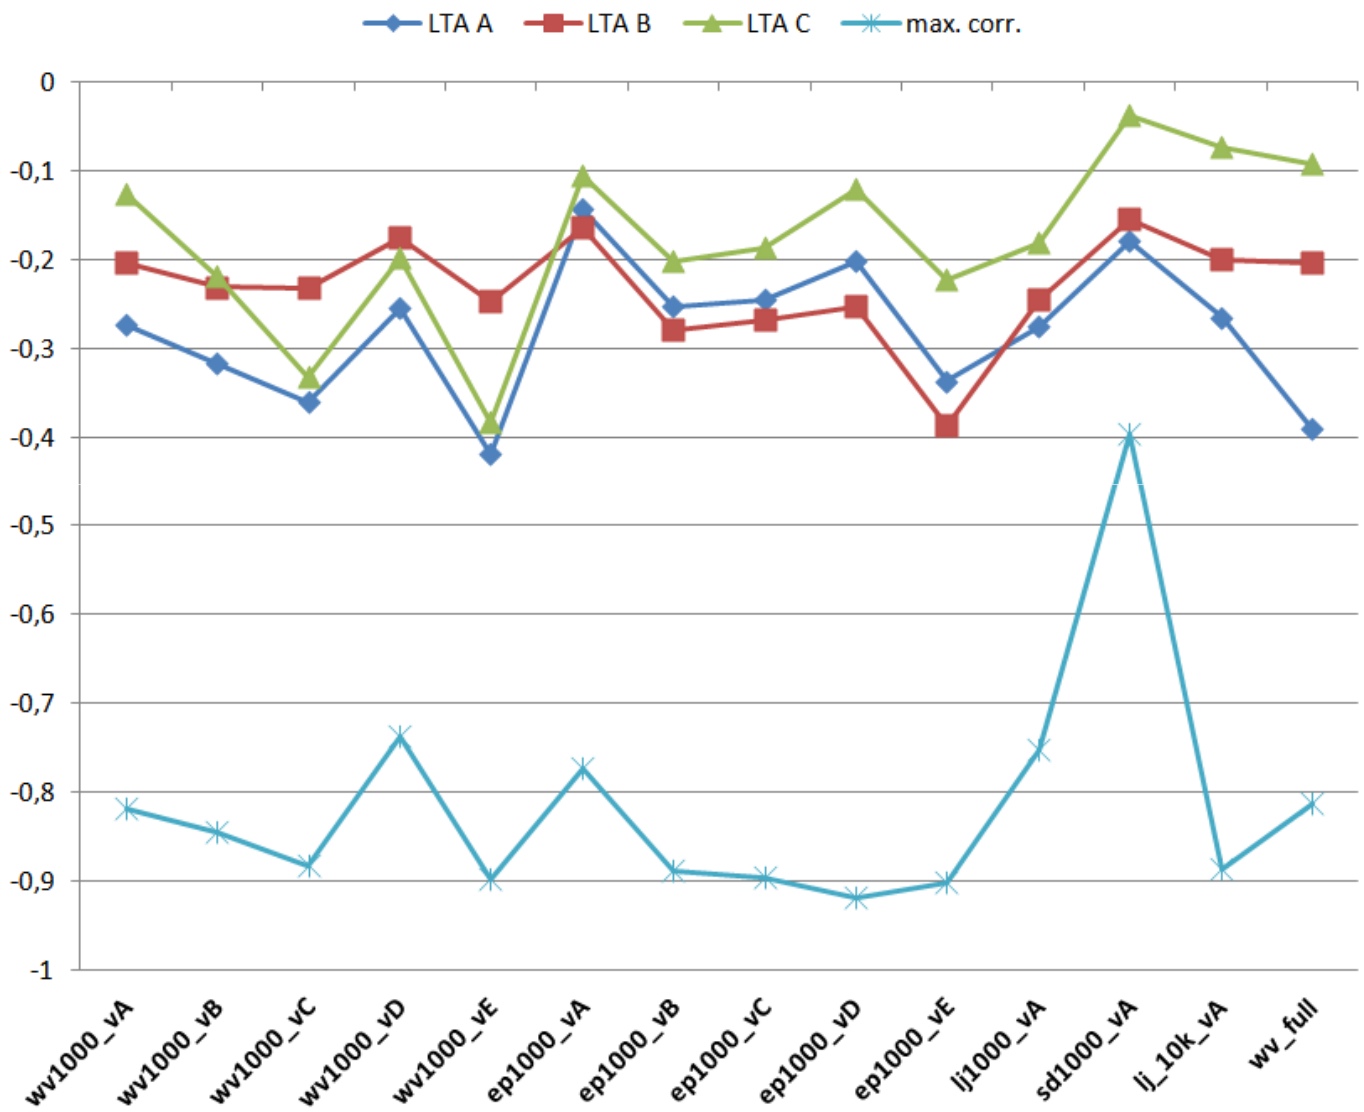

$avg(LTA_A)$=-0.27945
$avg(LTA_B)$=-0.23164
$avg(LTA_C)$=-0.17742

Corrected LTA evaluation

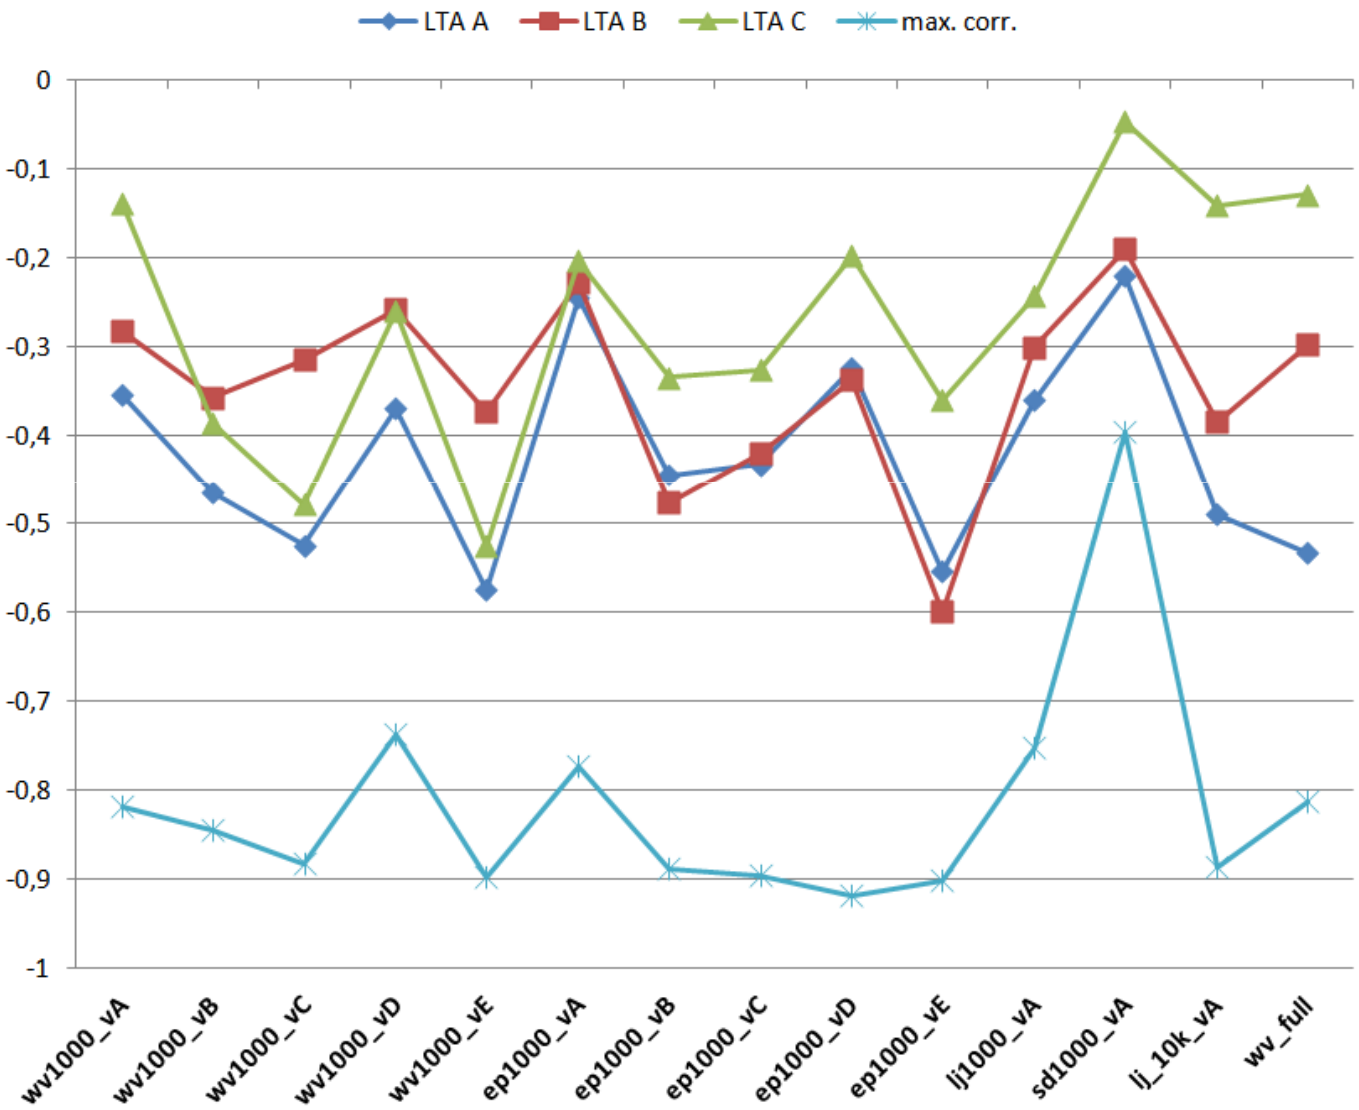

$avg(LTA_A)$=-0.42133
$avg(LTA_B)$=-0.34466
$avg(LTA_C)$=-0.26988

Future work

- Attacker perspective: network level LTA predictions?
- Directed networks?
- Improving measures?
- Combined global + local metrics?
- Further LTA analysis
  - Structural dependency
- LTA testing for other algorithms
  - E.g., seed-and-grow

Questions?

THANK YOU FOR YOUR ATTENTION!

Department of Telecommunications

Gábor György Gulyás
assistant research fellow
Dept. of Telecommunications (BME)
gulyasg@hit.bme.hu