Definitions Matter: Reconciling Differential and Adversarial Privacy

Dan Suciu

Joint work with Vibhor Rastogi

University of Washington

Example: AOL Search Data Fiasco

[August 4th 2006]

User Id	Search Query
4417749	"landscapers in Lilburn, GA"
4417749	"people with last name Arnold"
4417749	"numb fingers"
44 7749	"dog that urinates on everything"

Anonymized Search Logs for Research

End user (Thelma): opt-in/opt-out decision

Data owner (CIO): publish/no-publish decision

Needed: formal privacy guarantees

Other Examples of Private Data

facebook

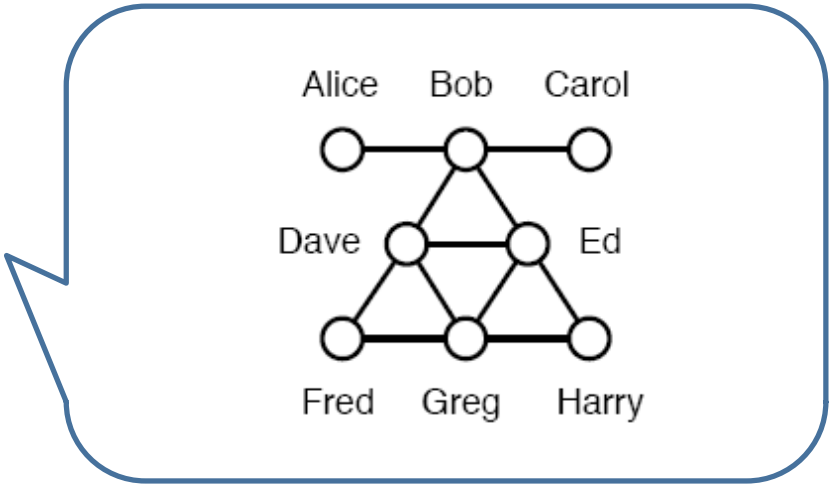

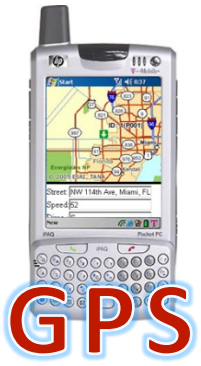

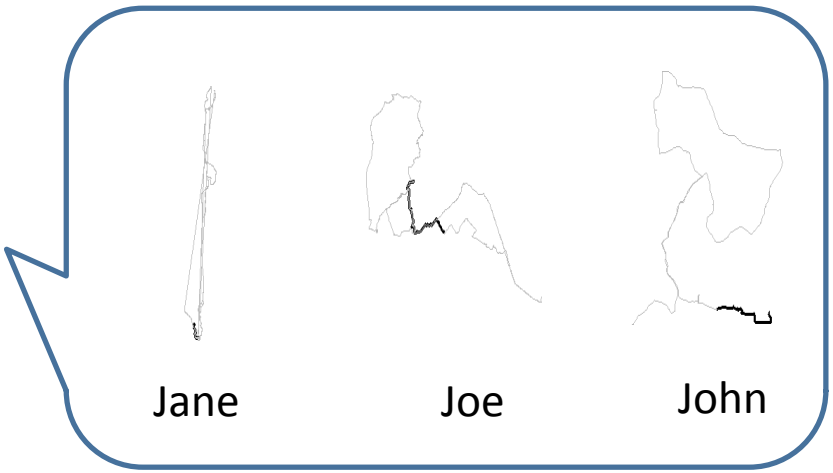

Other Examples of Private Data

Understand social trends

How many triangles in the graph?

Deny: ***Who are Bob's friends?***

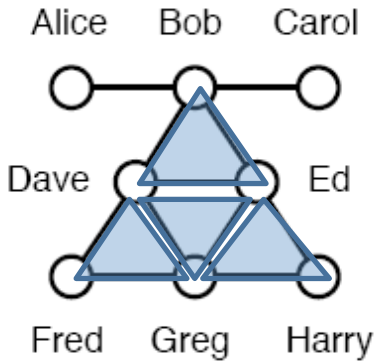

No known techniques with privacy guarantees !

Analyze traffic & congestion

How many drivers take route I-90?

Deny: ***Where does Joe go after work?***

Outline

- Background: Privacy definitions
- Central result: Differential vs. Adversarial privacy
- Applications: New techniques with formal privacy
- Conclusions

Background: Differential privacy [Dwork 06]

State-of-the-art privacy definition

Name	Age	Zip	Disease
John	15	98104	Flu
…..	….	…..	….
Smith	27	98105	Cancer
Alice	11	98105	Flu

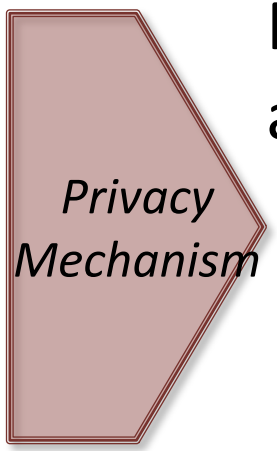

Possible perturbed answers:

163 prob = p_1 p_1'

175 prob = p_2 p_2'

. . . . .

Query: ***count # of flu patients with age < 10***

Current correct answer = **156**

New correct answer = **157**

Differential Privacy [Dwork 06] Output probabilities change very little with addition/removal of a single tuple

Background: Differential privacy [Dwork 06]

max change in query answer on adding/removing a single tuple

Standard Algorithm [Dwork et. al. 06]

- Add random noise according to ***query sensitivity***

Example:

SELECT-COUNT-WHERE queries:

'# of flu patients with age < 10'

Answer can change by at most 1
Query sensitivity is 1
Small noise required

Name	Age	Zip	Disease
John	15	98104	Flu
…..	….	…..	….
Smith	27	98105	Cancer
Alice	11	98105	Flu

Summary of Differential Privacy

- Designed for the end-user
  - Opt-in/opt-out decision
  - Property of the security mechanism, not of attackers
- Simple privacy mechanism based on query sensitivity

Background: Adversarial Privacy

Models the attacker explicitly

Prior belief

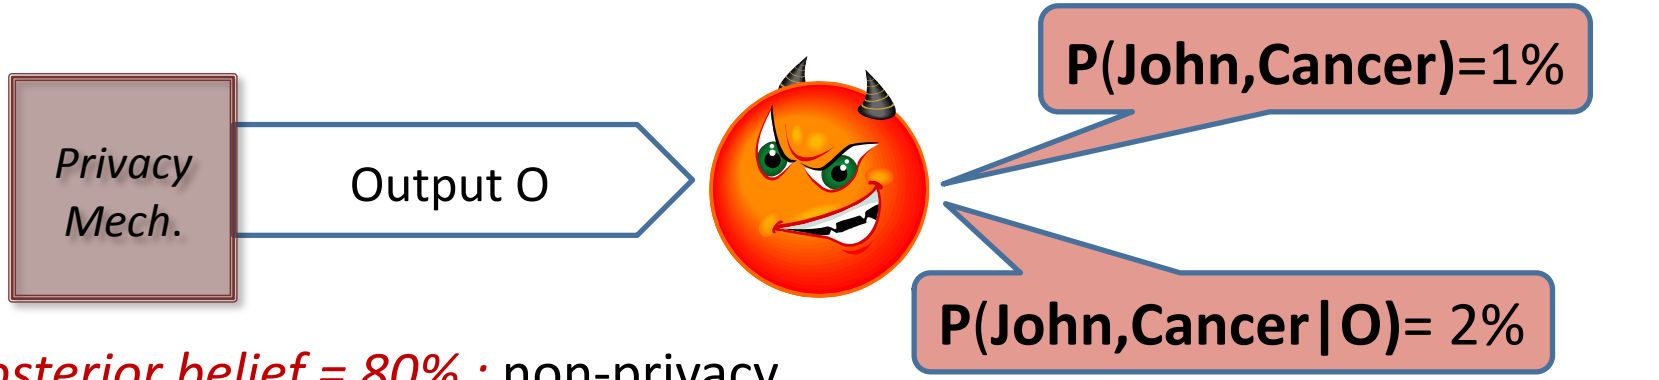

Posterior belief = 80% : non-privacy

Posterior belief = 1.2% : privacy

Posterior belief

Class of attackers

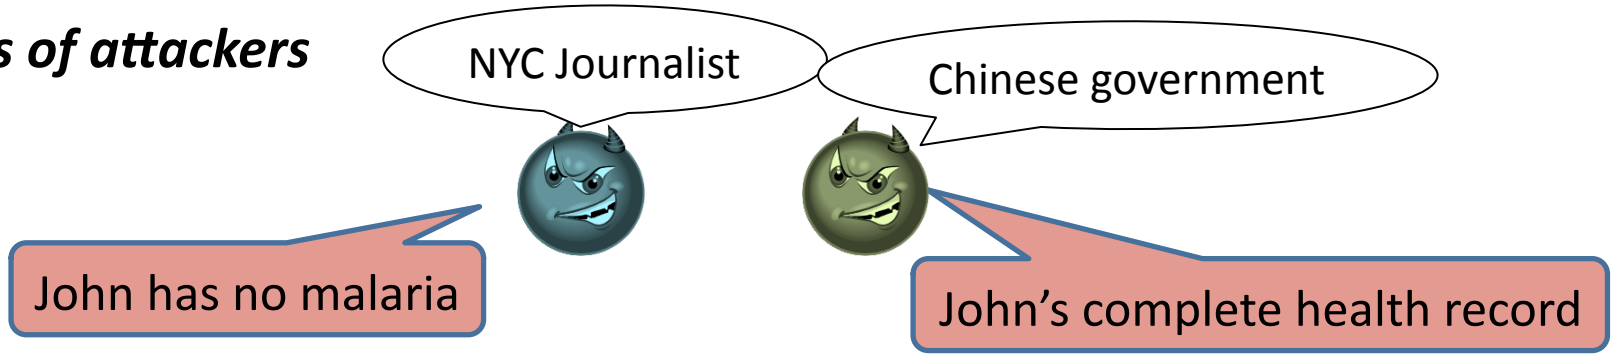

Adversarial privacy [Evfimievski 02] For all attackers in a ***class***, his knowledge changes little after learning the query's answer

Summary of Adversarial Privacy

- Designed for the data owner
  - Publish / don't publish decision
  - Stated in terms of an attacker, not the privacy mechanism
  - Applies equally well to group secrets

Outline

- Background: Privacy definitions
- Central result: Differential vs. Adversarial privacy
- Applications: New techniques with formal privacy
- Conclusions

Differential vs. Adversarial Privacy

Against which class of attackers does differential privacy protect?

Differential vs. Adversarial Privacy

Against which class of attackers does differential privacy protect?

Equivalence Result [Rastogi. et. al. 09]

Characterized the class of ***DP attackers*** for which*:

adversarial privacy against ***DP attckers*** = ***differential privacy***

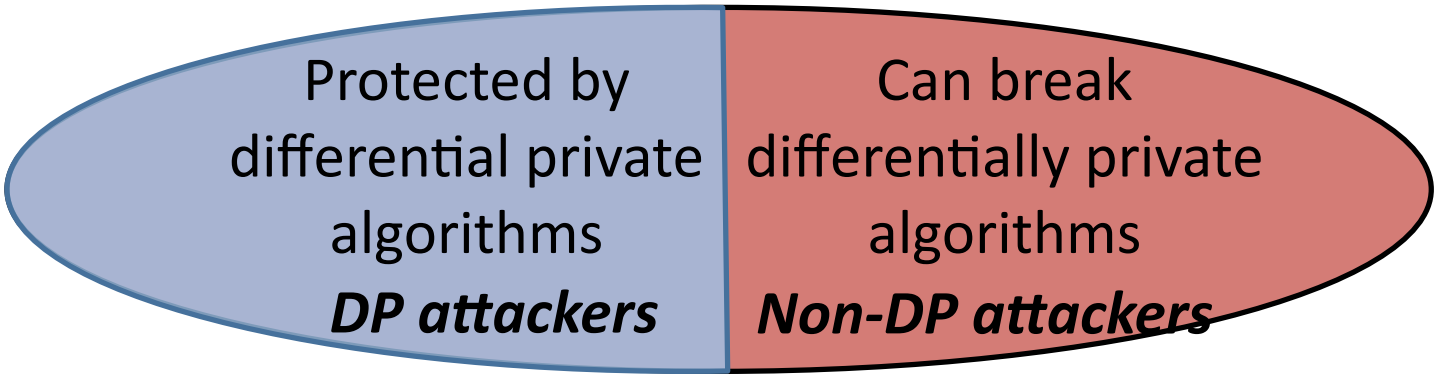

DP & *Non-DP* Attackers

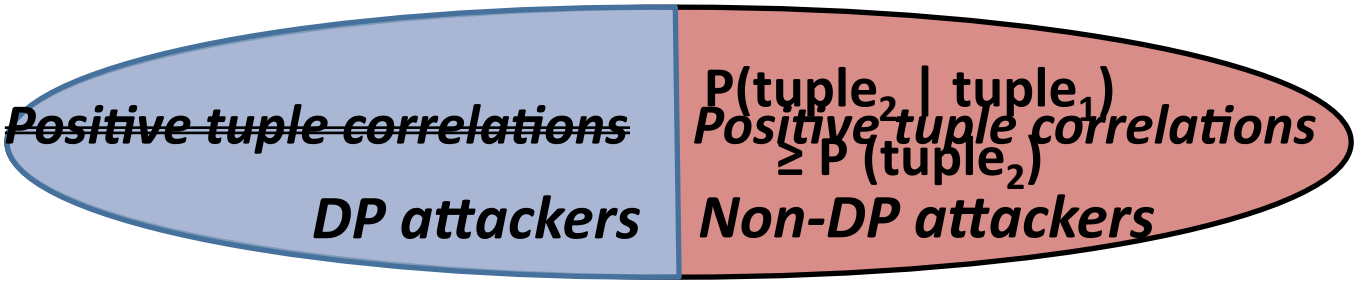

Formal definition of DP attackers

An attacker with prior belief distribution **P** is ***DP*** iff:

1. **P** is *log-submodular*
2. Every marginalization of **P** is also log-submodular
3. **P** is *planar*: exists n, P(DB)=0 if |DB| ≠ n

Non-DP Attackers: Positive Tuple Correlations

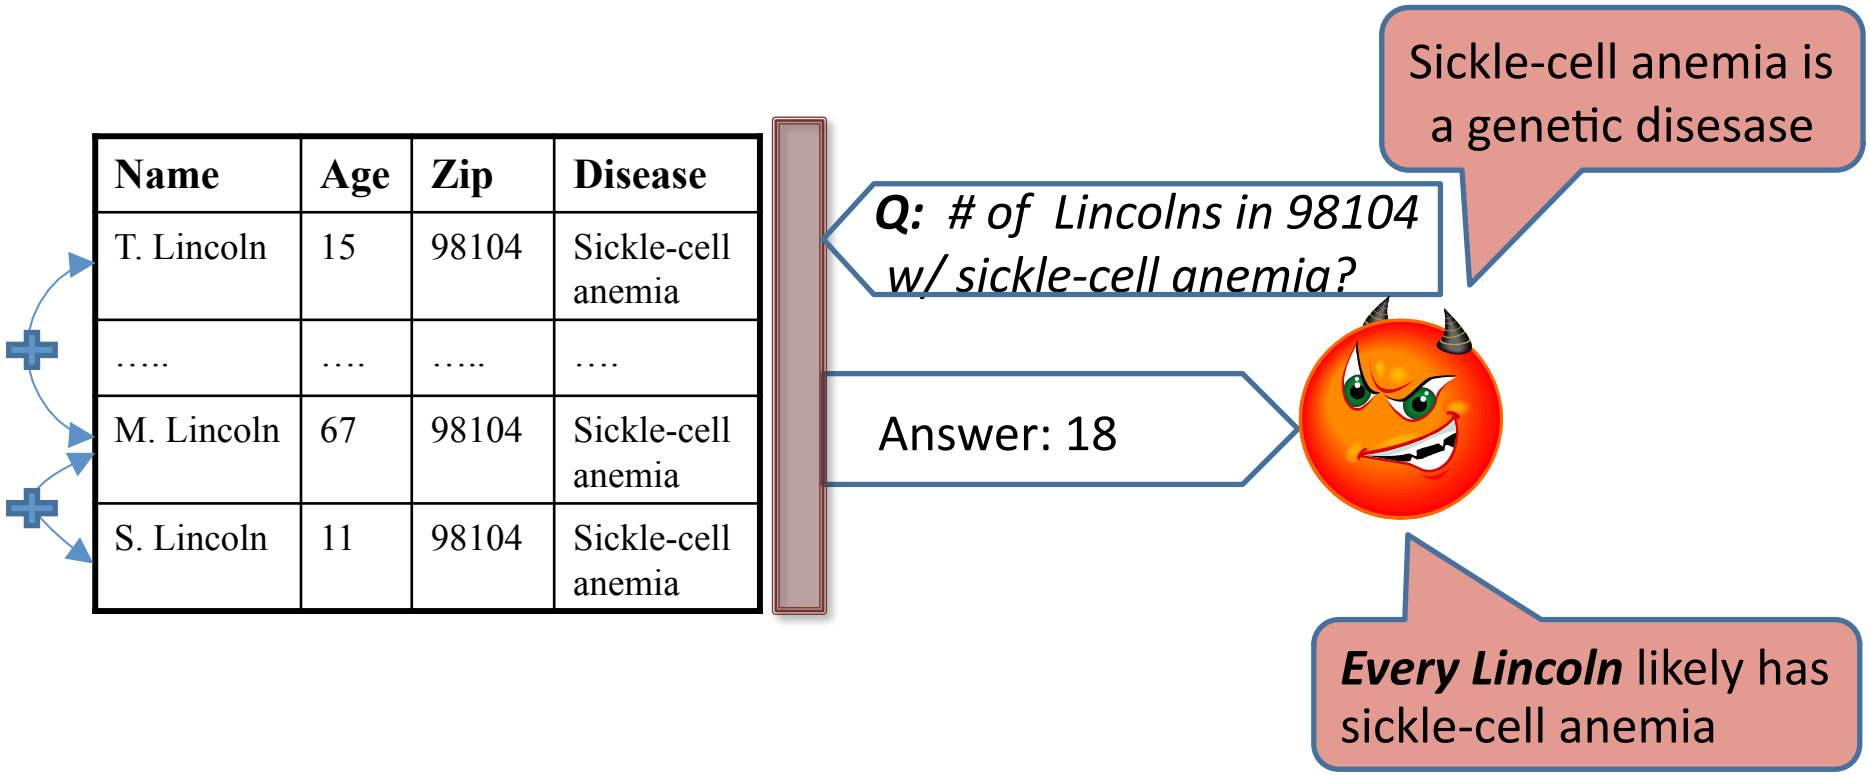

- In differential privacy, noise added according to query sensitivity
- That noise insufficient to hide the presence or absence of 20 tuples

Outline

- Background: Privacy definitions
- Central result: Differential vs. Adversarial privacy
- Applications: New techniques with formal privacy
- Conclusions

Applications: Analysis of Social Networks

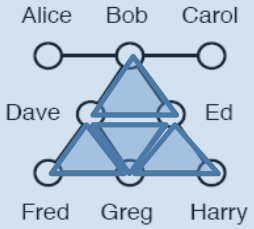

Algorithm for analyzing

social networks [PODS 09]

Algorithm for analyzing

time-series data [SIGMOD 10]

Central Result [PODS 09]: Connection between

Differential Privacy [Dwork 06] & ***Adversarial Privacy*** [Evfimievski 02]

Applications: Analysis of Social Networks

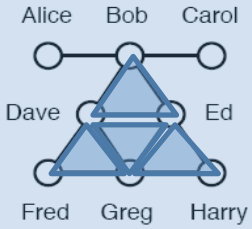

Algorithm for analyzing

social networks [PODS 09]

Error in response
for Differential Privacy

Common Graph Queries	Error
# of triangles	Θ(n)
# of paths of length 2	Θ(n)
# of cycles of length i	$\Theta(n^{i-2})$
# of nodes at dist. i	Θ(n)
# of cliques of size i	$\Theta(n^{i-2})$

Motif Analysis [Newman 03]

To understand the structure of the social network graph

Differential privacy adds too much noise

DP is too strong for social networks.

Relaxing Differential Privacy

- Use its connection with adversarial privacy

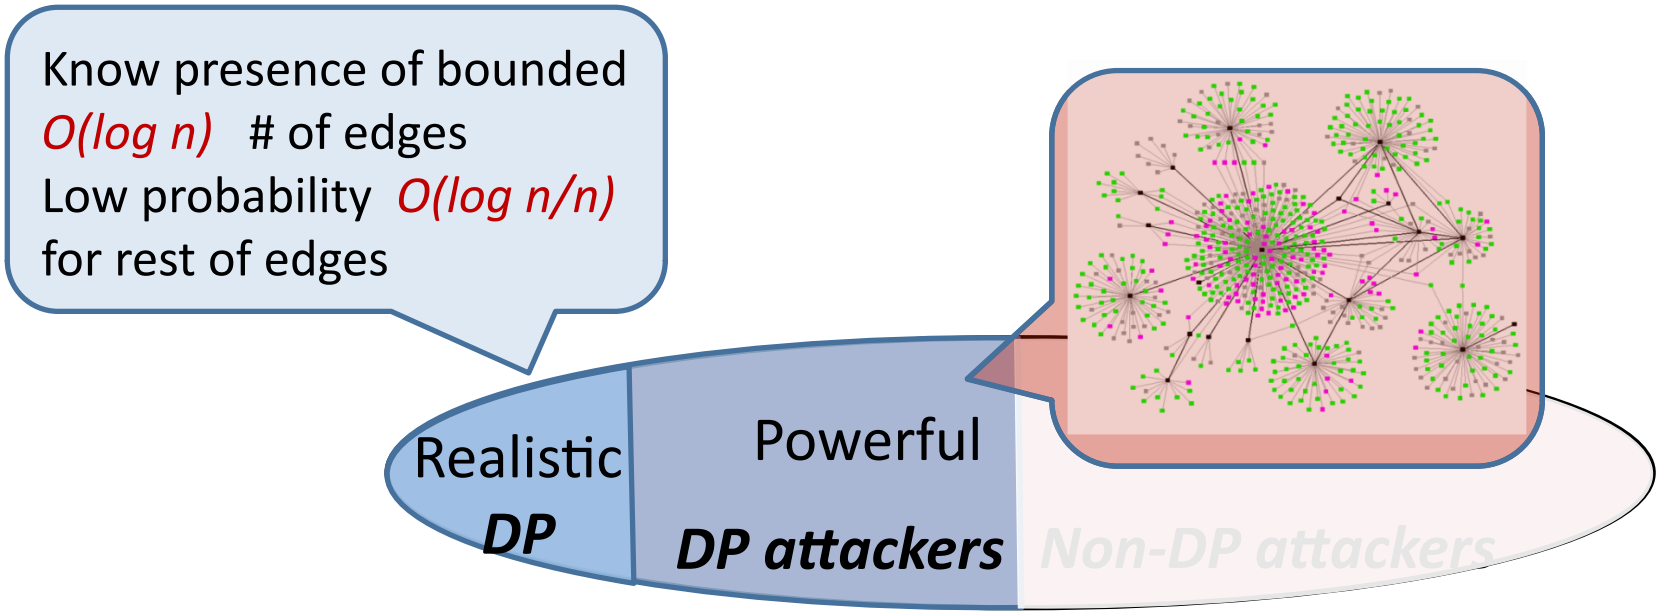

adversarial privacy against DP attackers = realistic DP = adversarial privacy against realistic DP attackers

Our Algorithm

Modification of the standard differentially private algorithm

Standard Algorithm [Dwork et. al. 06]

- Add random noise according to ***query sensitivity***

- Random noise has an expected error of Θ (query sensitivity)

Example '*# of triangles*': query sensitivity = n-2
expected error = Θ(n)

Our Algorithm

Modification of the standard differentially private algorithm

Expected change in query answer on adding/removing a single tuple

Our Algorithm **[Rastogi et. al. 09]**

- Return completely random answer with negligible probability
- Otherwise add random noise according to ***adversarial sensitivity***

Example '*# of triangles*': ***adversarial sensitivity*** = $O(\log^2 n)$

expected error = $O(\log^2 n)$

We bound ***adversarial sensitivity*** for a large class of queries

Proof based on new concentration results

Examples for Improved utility

Graph Query	**Standard** [Dwork06]	**Ours** [R. et. al. 09]
# of triangles	$\Theta(n)$	$\Theta(\log^2 n)$
# of 2-length paths	$\Theta(n)$	$\Theta(\log n)$
# of i-node cycles	$\Theta(n^{i-2})$	$\Theta(\log^{i-1} n)$
# of nodes at dist. i from a node v	$\Theta(n)$	$\Theta(\log^{i-1} n)$
# of i-node cliques	$\Theta(n^{i-2})$	$\Theta(\log^{i^2} n)$
# of conn. subgraphs of i nodes & j edges	$\Theta(n^{i-2})$	$\Theta(\log^{i-1} n)$

- **Standard** satisfies differential privacy = adversarial privacy against ***all DP*** attackers
- **Ours** satisfies adversarial privacy against only ***realistic DP*** attackers

The Class of Queries (Details….)

- Answer **count(q)** where **q :- R1, R2, …**

Density(q) = |Atoms(q)| / |Vars(q)|

Definition q is **dense**, if **Density(h(q)) >= 1**, forall homomorphism **h**

The **derivative** of **q** is:

∂(q(Ri)) = replace Ri with constants, then remove Ri

Definition q is **stable** if forall k >=1, $\partial^k(q)$ is dense

Theorem Adversarial sensitivity is bounded for all stable queries and "realistic" DP adversaries

Summary

- Differential privacy not possible for join queries

- Adversarial privacy possible against “relaxed DP” adversary for all “stable” queries

Applications: Analysis of Time-Series Data [Rastogi & Nath: SIGMOD 2010]

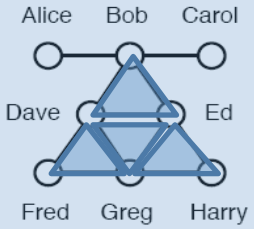

Algorithm for analyzing

social networks [PODS 09]

Algorithm for analyzing

time-series data [SIGMOD 10]

Equivalence Result [PODS 09]: Connects

Differential Privacy [Dwork 06] & ***Adversarial Privacy*** [Evfimievski 02]

Applications: Analysis of Time-Series Data [Rastogi & Nath: SIGMOD 2010]

Example: GPS traces

Name	Age	Location	Time
Alice	25	Dreese Labs	5 PM
Alice	25	Scott Labs	5:02 PM
Alice	25	Smith Labs	5:03 PM
Bob	32	Dreese Labs	5:35 PM

Lots of positive correlations!

Algorithm for analyzing ***time-series data*** [SIGMOD 10]

Differential privacy does not protect against positive correlations

Differential privacy too weak for time-series data!

Outline

- Background: Privacy definitions
- Central result: Differential vs. Adversarial privacy
- Applications: New techniques with formal privacy
- Conclusions

Conclusion

- Formal privacy guarantees required
- Differential privacy:
  - Opt-in / opt-out decision for end user
  - Property of privacy mechanism
- Adversarial privacy:
  - Publish / no-publish decision for data owner
  - Property of the class of attackers
- Fundamental equivalence theorem for ***DP attackers***
- Applications to:
  - Analysis of Social Networks
  - Analysis of time-series data

Thank You!