Differential Privacy vs Quantitative Information Flow

Kostas Chatzikokolakis

CNRS, INRIA, Ecole Polytechnique

joint work with
Mário Alvim, Miguel Andrés, Pierpaolo Degano, Catuscia Palamidessi

SRM seminar
Apr 3, 2012

Outline

1. Quantitative information flow
2. Differential privacy
3. QIF and Differential Privacy in the same context
4. Results
   - A general bound
   - Application to the leakage
   - Application to the utility

Outline

Information Flow

System

Secret Information Public Observables

Problem: Leakage of secret information via public observables

Information Flow

Programs

High variable values Low variable values

Problem: Leakage of secret information via public observables

Information Flow

Side channel attacks

Encryption keys Encryption time

Problem: Leakage of secret information via public observables

Information Flow

Anonymity protocols

Senders Public protocol events

Problem: Leakage of secret information via public observables

Information Flow

Ideally: No leak

Non-interference [Goguen & Meseguer'82]

In practice: there is almost always some leak

Intrinsic to the problem

Side channels

Information Flow

Intrinsic leak

out := OK
for $i = 1, \ldots, N$ **do**
  if $x_i \neq K_i$ **then**
    out := FAIL

  end if
end for

Side channel

out := OK
for $i = 1, \ldots, N$ **do**
  if $x_i \neq K_i$ **then**
    { out := FAIL
    exit() }
  end if
end for

Quantitative Information Flow

Goal: quantify the notion of information leakage

Most recent proposals use information theoretic approaches

Convergence of different fields: information flow, side channel analysis, anonymity protocols, . . .

Systems as Information-Theoretic channels

Channels are noisy: outputs are produced by multiple inputs and each input can generate multiple outputs

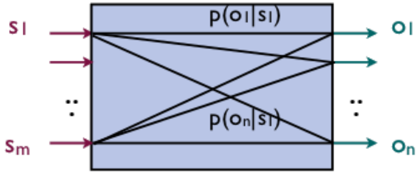

$p(o_j|s_i)$: probability to observe o_j given the input s_i

Systems as Information-Theoretic channels

Channels are characterized by their matrix of conditional probabilities

	o_1	$\dots$	o_n
s_1	$p(o_1 \mid s_1)$	$\dots$	$p(o_n \mid s_1)$
$\vdots$	$\vdots$		
s_m	$p(o_1 \mid s_m)$		$p(o_n \mid s_m)$

A prior distribution on the secrets models the attacker's side information

Useful concepts from information theory

Entropy $H(S)$

the attacker's initial uncertainty about the secret (difficulty to guess)

Conditional entropy $H(S|O)$

the attacker's uncertainty after observing the output

$Leakage = H(S) - H(S|O)$

Several notions of entropy (how we measure the attacker's success)

- Shannon entropy
- Min-entropy
- Guessing entropy
- ...

Min-entropy

[Rényi 61], [Smith 09]

One-try attacks

questions of the form: "is $S = s$?"

Measure of success:

$H_\infty(S) = -\log \max_s p(s)$

Leakage:

$I_\infty(S;O) = H_\infty(S) - H_\infty(S|O)$

$C_\infty = \max I_\infty$ over all input distributions

$= \log \sum_o \max_s p(o|s)$

Min-entropy

C_∞ is small when the difference between the rows is small

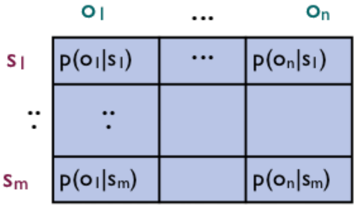

$C_\infty = 0$ iff $p(o|s) = p(o|s')$ for all o, s, s'

Outline

Statistical queries

- Database: a collection of individuals each having a value from a set $\mathcal{V}$
- Goal: publish the result of a statistical query. eg: average salary
- Problem: the query reveals information about a user's value
  - Databases can be dynamic, rows might be added/deleted
  - Sometimes even the participation in the database should be hidden

Statistical queries

Name/Id	age	weight	sex	epilepsy	...
Mario Rossi	65	82	M	yes	...
Daniele Bianchi	35	120	M	yes	...
Lucia Verdi	40	45	F	no	...
...	...	...	...	...	...

- We want to reveal global information:
  - How many people have epilepsy ?
  - What is the average age and weight of men who have epilepsy ?
- While protecting individual information:
  - Does Daniele Bianchi have epilepsy ?
  - What is the name of the last record inserted in the database ?

Statistical queries

Name/Id	age	weight	sex	epilepsy	...
Mario Rossi	65	82	M	yes	...
Daniele Bianchi	35	120	M	yes	...
Lucia Verdi	40	45	F	no	...
...	...	...	...	...	...

- How many men have epilepsy ? 2
- What is the average age / weight of men who have epilepsy ? 50 / 101

insertion of a new record

Name/Id	age	weight	sex	epilepsy	...
Mario Rossi	65	82	M	yes	...
Daniele Bianchi	35	120	M	yes	...
Lucia Verdi	40	45	F	no	...
Sergio Neri	20	140	M	yes	...
...	...	...	...	...	...

- How many men have epilepsy ? 3
- What is the average age / weight of men who have epilepsy ? 40 / 114

We can deduce the exact age / weight of the new record

Differential privacy

- Ideally: any information obtained from the database should be obtainable without it
- This is impossible [Dwork 06]
- Differential Privacy:
  - adding a user (or modifying his value) should have negligible affect on the query's result

Differential privacy

u: number of users

$\mathcal{V}$: set of values, possibly containing an "absence" value $\emptyset$

$\mathcal{V}^u$: set of all databases (u-tuples of values in $\mathcal{V}$)

$\langle 1, 4, 5, 2 \rangle \qquad \langle 1, 4, 5, 9 \rangle \qquad \langle 2, 9, 6, 3 \rangle$

adjacency relation: $D \sim D'$ iff they differ in exactly one value

Differential privacy

Differential Privacy

$$Pr[\mathcal{K}(D) = o] \leq e^{\epsilon}\ Pr[\mathcal{K}(D') = o] \qquad \forall D \sim D', o$$

Equivalently

$$Pr[\mathcal{K}(D) = o] \leq e^{\epsilon\ d(D,D')}\ Pr[\mathcal{K}(D') = o] \qquad \forall D, D', o$$

Equivalently

Let $D^i = \{D' \in V^u | D'_j = D_j\ \forall j \neq i\}$

$$Pr[D \mid o, D^i] \leq e^{\epsilon} Pr[D \mid D^i] \qquad \forall D, i, o$$

Differential privacy

Differential Privacy

$$Pr[\mathcal{K}(D) = o] \leq e^{\epsilon}\ Pr[\mathcal{K}(D') = o] \qquad \forall D \sim D', o$$

Equivalently

$$Pr[\mathcal{K}(D) = o] \leq e^{\epsilon\ d(D,D')}\ Pr[\mathcal{K}(D') = o] \qquad \forall D, D', o$$

Equivalently

Let $D^i = \{D' \in V^u | D'_j = D_j\ \forall j \neq i\}$

$$Pr[D\,|\,o, D^i] \leq e^{\epsilon} Pr[D\,|\,D^i] \qquad \forall D, i, o$$

Differential privacy

Differential Privacy

$$Pr[\mathcal{K}(D) = o] \leq e^{\epsilon}\ Pr[\mathcal{K}(D') = o] \qquad \forall D \sim D', o$$

Equivalently

$$Pr[\mathcal{K}(D) = o] \leq e^{\epsilon\ d(D,D')}\ Pr[\mathcal{K}(D') = o] \qquad \forall D, D', o$$

Equivalently

Let $D^i = \{D' \in V^u | D'_j = D_j\ \forall j \neq i\}$

$$Pr[D \mid o, D^i] \leq e^{\epsilon} Pr[D \mid D^i] \qquad \forall D, i, o$$

Achieving Differential privacy

Typical approach: oblivious mechanisms

compute the real answer $f(D)$ to the query, then add noise

the noise depends only on the real answer

Achieving Differential privacy

Example: Laplacian mechanism

Global sensitivity: $\Delta_f = \max_{D \sim D'} |f(D) - f(D')|$

Draw $\mathcal{K}(D)$ from a laplacian distribution with mean $f(D)$ and variance Δ_f / ϵ

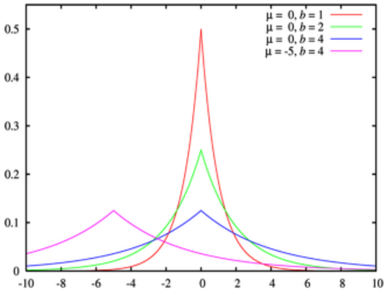

Utility

- The reported answer is only useful if it provides information about the real answer
- gain function $g(i, j)$
  - how much we gain when we believe i and the real answer is j
- we define the utility as the expected gain
- it depends on both the gain function and the prior distribution
- Goal: find optimal mechanisms for different types of queries

Outline

Statistical queries as noisy channels

Input: the database X

Output: the reported answer Z

Probabilistic behaviour due to the added noise

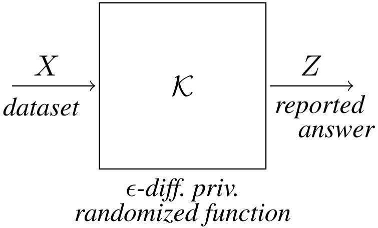

Statistical queries as noisy channels

Something new: a graph structure on the inputs

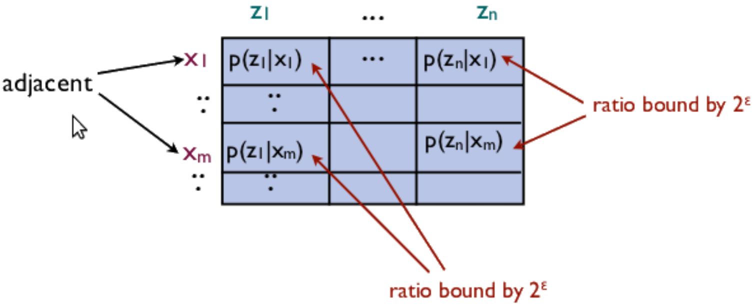

Diff. privacy requires rows to be similar, but only adjacent ones

Oblivious queries

The noise only depends on the real answer

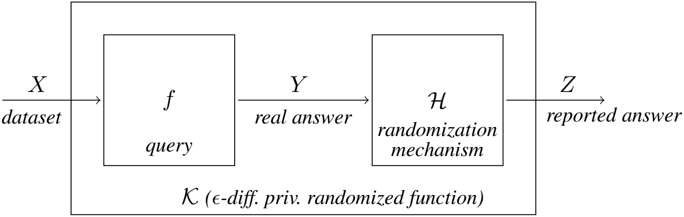

Oblivious queries

The noise only depends on the real answer

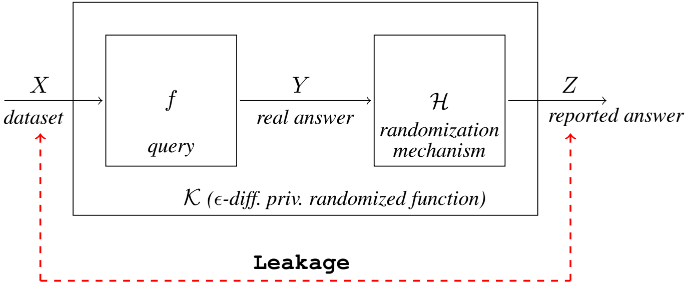

Oblivious queries

The noise only depends on the real answer

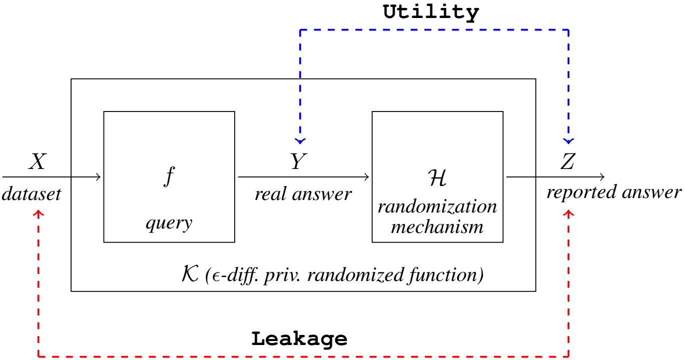

Leakage and utility

Leakage: $I_\infty(X; Z)$

Utility: $\mathcal{U} = 2^{-H_\infty(Y|Z)}$

for the binary gain function $g(i,j) = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$

Questions:

Does ϵ-d.p. impose a bound on the leakage?

Does ϵ-d.p. impose a bound on the utility?

How to construct an ϵ-d.p. mechanism with maximal utility?

Outline

A general bound for symmetric graphs

- consider a channel $X \to Z$
  - and a graph structure $(X, \sim)$ on its inputs
  - s.t. ϵ-d.p. is satisfied
- different graphs impose different bounds on the leakage

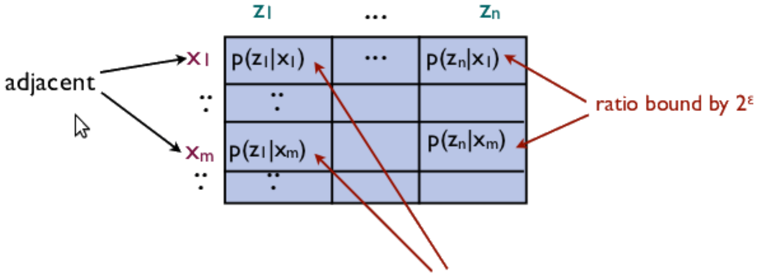

A general bound for symmetric graphs

We consider two families of graphs:

- vertex transitive:
  for all vertices v, w there exists an automorphism mapping v to w
- distance regular:
  for all vertices v and w at distance i the number of vertices adjacent to w and at distance j from v is the same

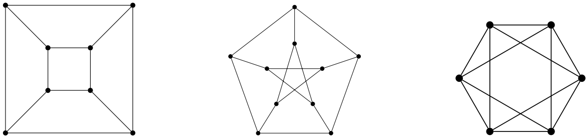

A general bound for symmetric graphs

Theorem

Assuming that $(X, \sim)$ is distance regular or vertex transitive+, and that it satisfies ϵ-d.p., we have

$$H_\infty(X|Y) \leq -\log \frac{1}{\sum_d \frac{n_d}{e^{\epsilon\, d}}}$$

where n_d is the number of nodes at distance d from a fixed node r.

Application to the leakage

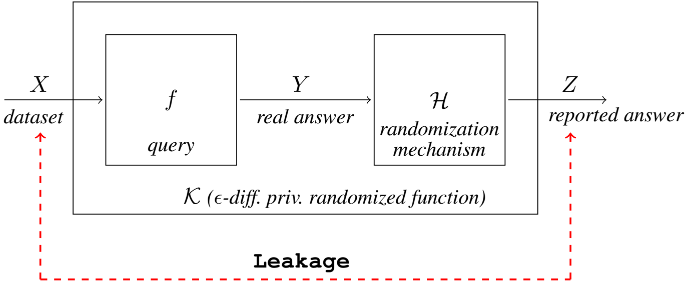

Channel from X to Z, inputs are databases

Application to the leakage

consider the set of databases $\mathcal{V}^u$ with the corresp. adjacency relation

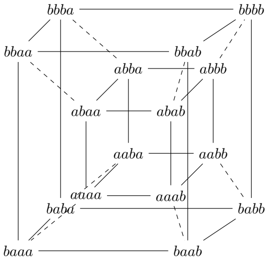

$(\mathcal{V}^u, \sim)$ is both distance-regular and vertex-transitive

moreover $n_d = \binom{u}{d}(v-1)^d$

Application to leakage

Theorem

Let $v = |\mathcal{V}|$. If $\mathcal{K}$ satisfies ϵ-d.p. then:

$$I_\infty(X;Z) \leq -u \log_2 \frac{v\, e^\epsilon}{v - 1 + e^\epsilon}$$

The bound is strict.

A stronger bound can be proven for the leakage of a single individual

Application to the utility

- Channel from Y to Z, inputs are real answers
- induced graph: the adjacency relation on X induces one on Y
  - $y \sim y' \quad \text{iff} \quad x \sim x', f(x) = y, f(x') = y'$
- the graph $(Y, \sim)$ depends on the actual query f

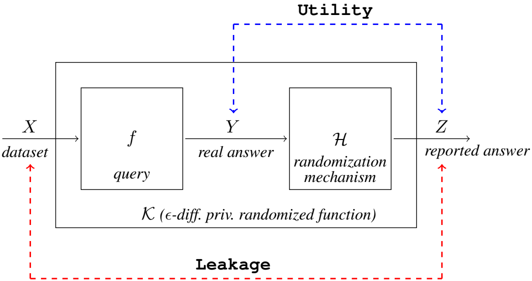

Two results from the litarature

- The geometric mechanism is universally optimal for counting queries (i.e. the induced graph is a path graph)

$$p(j|i) = c_j \alpha^{-|i-j|} \quad \text{where} \quad c_j = \begin{cases} \frac{\alpha}{\alpha+1} & j = 1 \text{ or } j = n \\ \frac{\alpha-1}{\alpha+1} & 1 < j < n \end{cases}$$

- For all other graphs no universally optimal mechanism exists

Application to the utility

Theorem

Assuming that $(Y, \sim)$ is distance regular or vertex transitive+, and that it satisfies ϵ-d.p., we have

$$\mathcal{U} \leq \frac{1}{\sum_d \frac{n_d}{e^{\epsilon\, d}}}$$

Constructing an optimal mechanism

we construct a matrix $\mathcal{H}$ as follows:

$$\mathcal{H}_{i,j} = \frac{c}{e^{\epsilon\, d(i,j)}} \qquad\qquad c = \frac{1}{\sum_d \frac{n_d}{e^{\epsilon\, d}}}$$

this is a valid matrix that

- satisfies ϵ-d.p
- has utility $\mathcal{U} = \frac{1}{\sum_d \frac{n_d}{e^{\epsilon\, d}}}$

so under the symmetry assumptions on $(Y, \sim)$ it has optimal utility

Example

Consider a database with electoral information where each row corresponds to a voter and contains the following three fields:

- Id : a unique (anonymized) identifier assigned to each voter;
- City: the name of the city where the user voted; one of $\{A, B, C, D, E, F\}$
- Candidate: the name of the candidate the user voted for.

Query: "What is the city with the greatest number of votes for a given candidate?".

Every two answers are adjacent, i.e. the graph structure of the answers is a complete graph.

Example

The optimal matrix is

In/Out	A	B	C	D	E	F
A	2/7	1/7	1/7	1/7	1/7	1/7
B	1/7	2/7	1/7	1/7	1/7	1/7
C	1/7	1/7	2/7	1/7	1/7	1/7
D	1/7	1/7	1/7	2/7	1/7	1/7
E	1/7	1/7	1/7	1/7	2/7	1/7
F	1/7	1/7	1/7	1/7	1/7	2/7

Related work

Barthe & Köpf have been independently working on the same problem

They provide the first bounds on information leakage imposed by differential privacy [CSF 2011]

Differences of our approach

- different technique, based on graph symmetries
- improved bound
- we also consider bounds on the utility

Ongoing work

- A generalization of min-entropy leakage by considering the attacker's gain function (CSF'12)
- This can lead to a closer correspondance with differential privacy
- Extend the optimality results to more general families of graphs, including path graphs
- Optimality results for classes of gain functions and prior distributions

Questions?