Non-interactive privacy

Sébastien Gambs
Université de Rennes 1 - Inria / IRISA
CIDRE research team

sgambs@irisa.fr

25 March 2015

Interactive privacy

- Most of the privacy enhancing technologies involve a distributed computation between two (or more) participants based on cryptographic techniques.
- Communication model: two-way communication with several rounds of communication.
- View of cryptography: a distributed computation is said to be private if no other information is learnt from the protocol that the output of the computation itself.

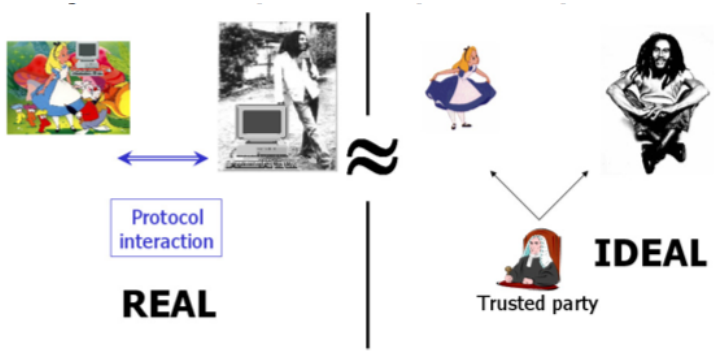

- Shortcoming: cryptography does not say anything about the fact that the output itself can leak a lot information about the input of the participants.

Drawbacks of interactive privacy

1. The users need to be online to participate.
2. A lot of security issues to take into account (malicious participant, eavesdropper, security of cryptographic primitives, side channels, ...) $\Rightarrow$ complexity of defining a strong and realistic adversary model
3. Difficulty of reasoning about inferences that can be done in case of multiple computations on the same private data.
4. Additional assumptions about the communication setting : synchronicity, availability of broadcast channel or private channel between each pair of participants, ...

Going distributed has a first step to protect privacy

- Distributed collaborative social platform : each user is characterized by a profile, which is a list of items he has tagged/liked.

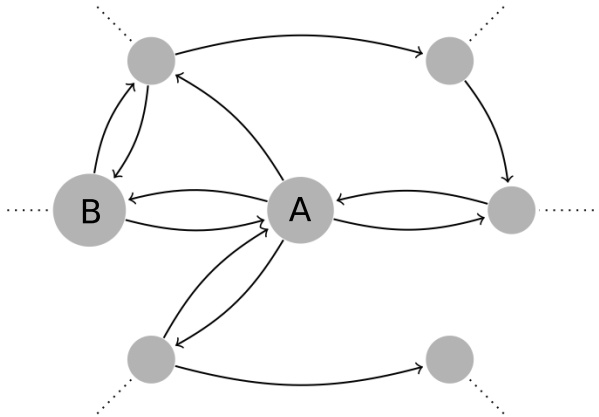

- Example of application : decentralized personalized search (Bertier, Frey, Guerraoui, Kermarrec and Leroy 10)
- Main challenge : using the personal information while preserving privacy.

Privately computing the similarity between two users (OPODIS'11)

- Joint work with Anne-Marie Kermarrec and Mohammad Alaggan (Inria).
- Main objective of this work : computing similarity between users in a privacy-preserving manner.
- Profiles can be represented as binary vectors $\in \{0,1\}^n$, in which n is the size of the domain of possible items.
- Cosine similarity : $\text{Cos\_Sim}(A,B) = \frac{|A \cap B|}{\sqrt{|A| \times |B|}}$.
- Cryptographic approach : use *bipartite secure computation* and *private intersection protocols* to compute the similarity without exchanging the profiles in clear.

First solution : threshold similarity

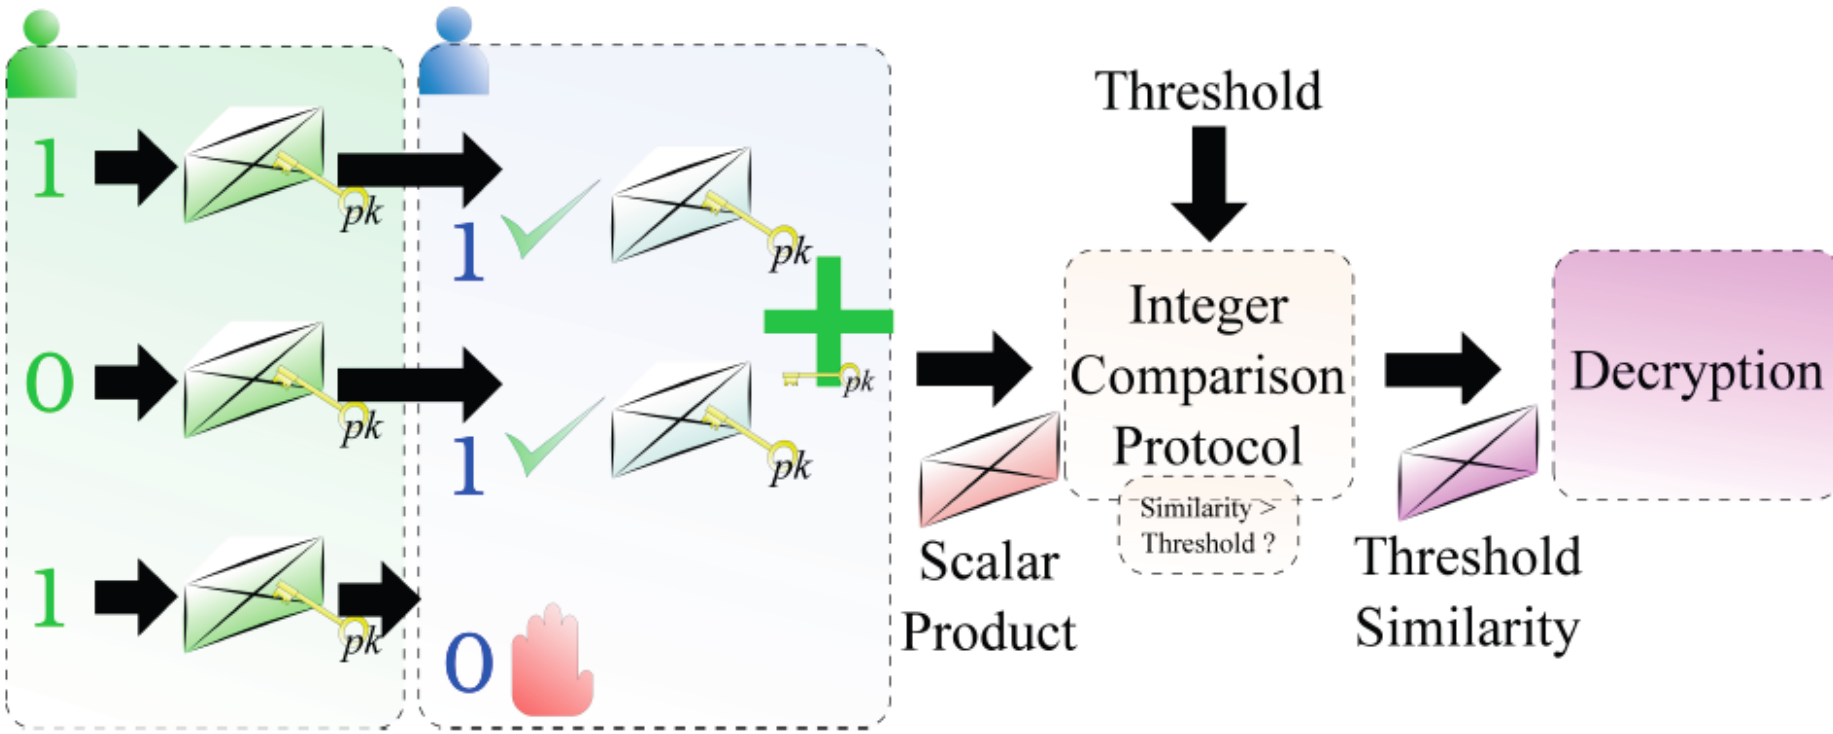

Can be implemented by a combination of an homomorphic scheme such as Paillier's cryptosystem (99) and an integer comparison protocol (Nishide and Ohta 07).

Possible inferences from the output of the similarity computation

- Full disclosure : if the similarity is 100% $\Rightarrow$ both profiles are identical.
- Partial disclosure : Cos_Sim$(\{5, 8\}, B) > 0 \Rightarrow$ B contains either 5 or 8.
- De-anonymization : "84% of subscribers present in the Netflix dataset can be uniquely identified if the adversary knows 6 out of 8 movies outside the top 500" (Narayanan and Shmatikov 08).
- Equivalently Cos_Sim $> \frac{6}{\sqrt{8 \times b}}$ (> 0.15 if $b < 200$).

Differential privacy: principle (Dwork 06)

- Recent privacy notion developed within the community of private data analysis.

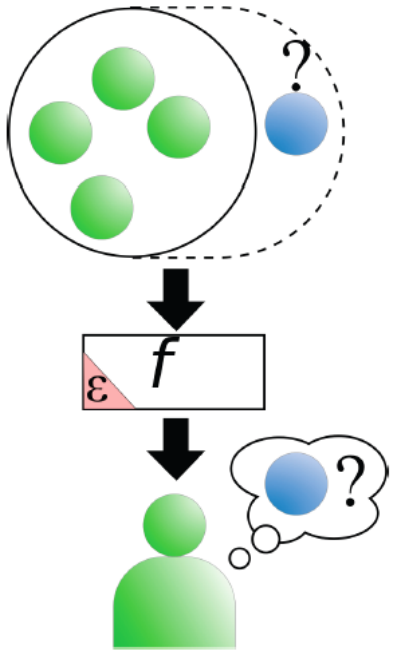

- Basically ensures that whether or not an item is in the profile of an individual does not influence too much the output.
- Give strong privacy guarantees that hold independently of the auxiliary knowledge of the adversary.

Differential privacy: definition

- Differential privacy (Dwork 06): A randomized function K gives $\epsilon-$*differential privacy* if for all possible inputs X_1 and X_2 differing in a most one element, and all $S \subseteq Range(K)$,

$$\Pr[K(X_1) \in S)] \leq \exp(\epsilon) \times \Pr[K(X_2) \in S)] \qquad (1)$$

  The probability is taken over all the coin tosses of K.
- ϵ is a public privacy parameter.
- Typical value: 0.01, 0.1 or even 0.25.
- Properties:
  - Composition: the application of k ϵ-differentially private mechanisms leads to a $k\epsilon$-differentially private mechanism.
  - Postprocessing does not hurt privacy.

Sensitivity

- The sensitivity measures how much the output of a function can change with respect to a small change in the input.
- Global sensitivity (Dwork 06): For $f : D^n \rightarrow R$, the (global)*sensitivity* of f is

$$GS(f) = max_{X_1,x_2} \|f(X_1) - f(X_2)\|_1 \tag{2}$$

for all X_1,X_2 differing in at most one element.

- Example: two profiles S_1 and S_2 are *neighbours* if they are the same up to a particular item.
- The sensitivity of the Hamming distance (computed between two binary vectors) is one.

Laplacian mechanism

- Achieves ϵ-differential privacy by adding noise directly proportional to GS(f)

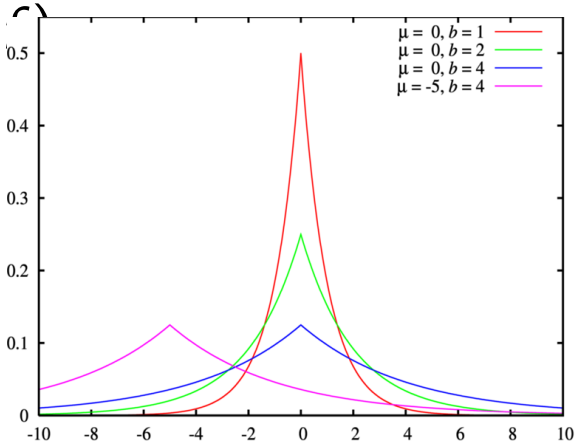

- Theorem (Dwork 06): For $f : D^n \rightarrow R$, a randomized function K achieves ϵ-differential privacy if it releases on input x

$$K(x) = f(x) + \text{Lap}(\frac{\text{GS}(f)}{\epsilon}) \tag{3}$$

for GS(f) the sensitivity of the function f and Lap is a randomly generated noise according to the Laplacian distribution parametrized by $\frac{\text{GS}(f)}{\epsilon}$.

Second solution: differentially-private threshold similarity

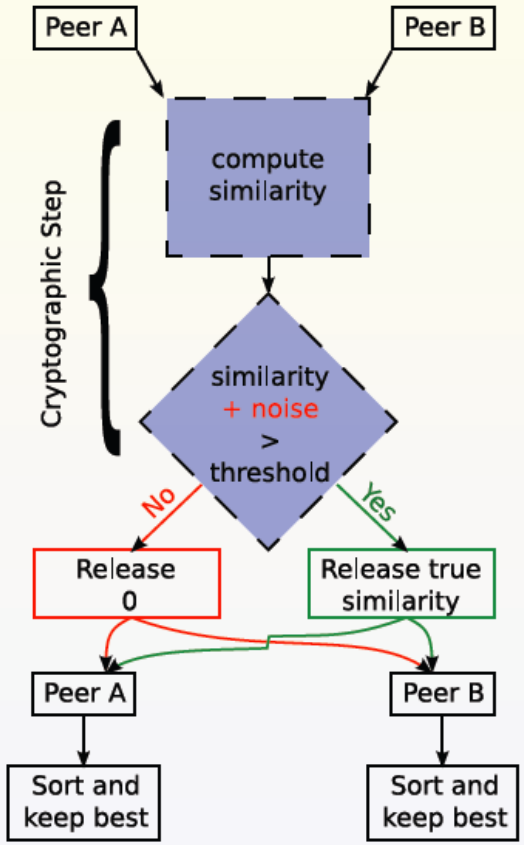

Alaggan, M., Gambs, S., Kermarrec, A.M.: Private similarity computation in distributed systems: From cryptography to differential privacy. OPODIS11.

Non-interactive privacy

- Communication model: one-way communication.
- Main advantages of non-interactivity :
  - Once the data is released, the user can go offline.
  - Can be used an unbounded number of times compared to interactive mechanisms.
  - Simple adversary model that mainly involve the inferences that can be performed out of the data released.
  - Decrease the trust assumptions.
- Principal limit : the utility can be less than for an interactive mechanism.

Bloom filter (Bloom 70)

- Compact representation of profiles.

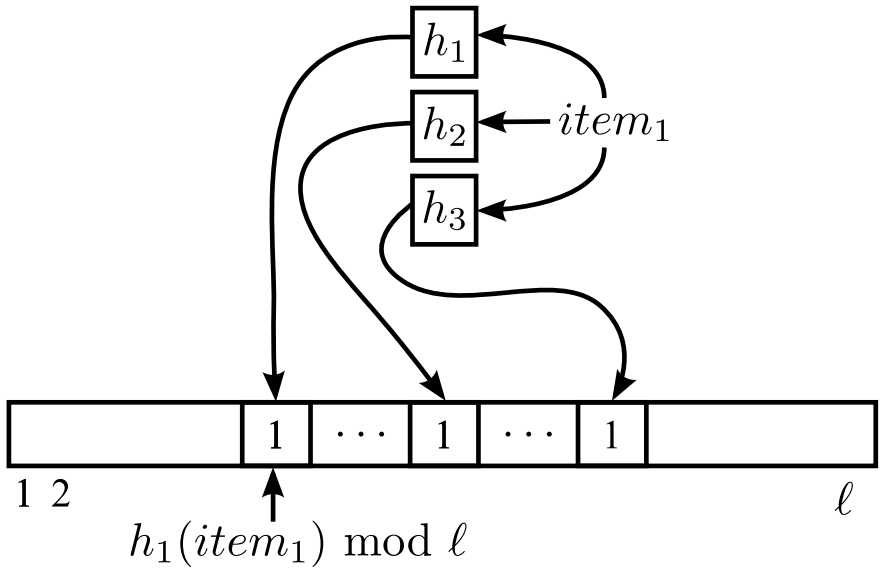

- Two operations : *add* and *contain*.
- Possibility of false positives (but no false negatives).
- Main privacy issue : could be exhaustively queried to reconstruct approximately the profile.

BLIP mechanism (SSS 12)

BLIP (BLoom-then-flIP) mechanism (in a nutshell):

- Step 1 : representation of the profile as a Bloom filter.
- Step 2 : flip each bit of the Bloom filter with probability p chosen in order to ensure differential privacy while maintaining a high level of utility.

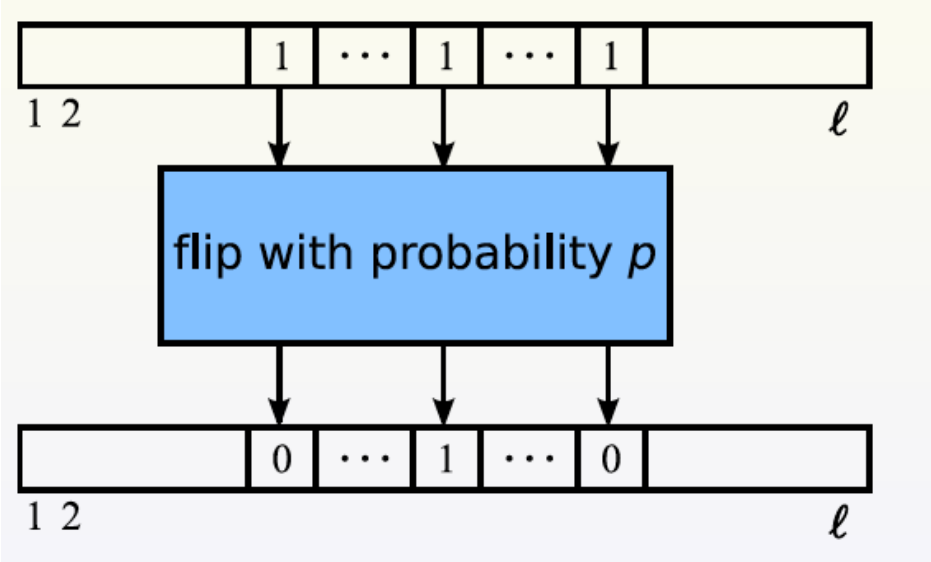

- Main privacy issue : could be exhaustively queried to reconstruct approximately the profile.

Main results

Theorem (Privacy)

Flipping each bit with probability $p = 1/(1 + e^{\epsilon/k})$*, in which* k *is the number of hash functions, satisfies differential privacy for items. Moreover, this probability is optimal (no lower flipping probability can satisfy differential privacy for this particular value of* ϵ*).*

Theorem (Utility)

Given a flipped Bloom filter and non-flipped Bloom filter, the additive error of their scalar product is $\Theta(\sqrt{l})$ *with constant probability, in which* l *is the size of the Bloom filter.*

Experimental results (utility on Digg dataset)

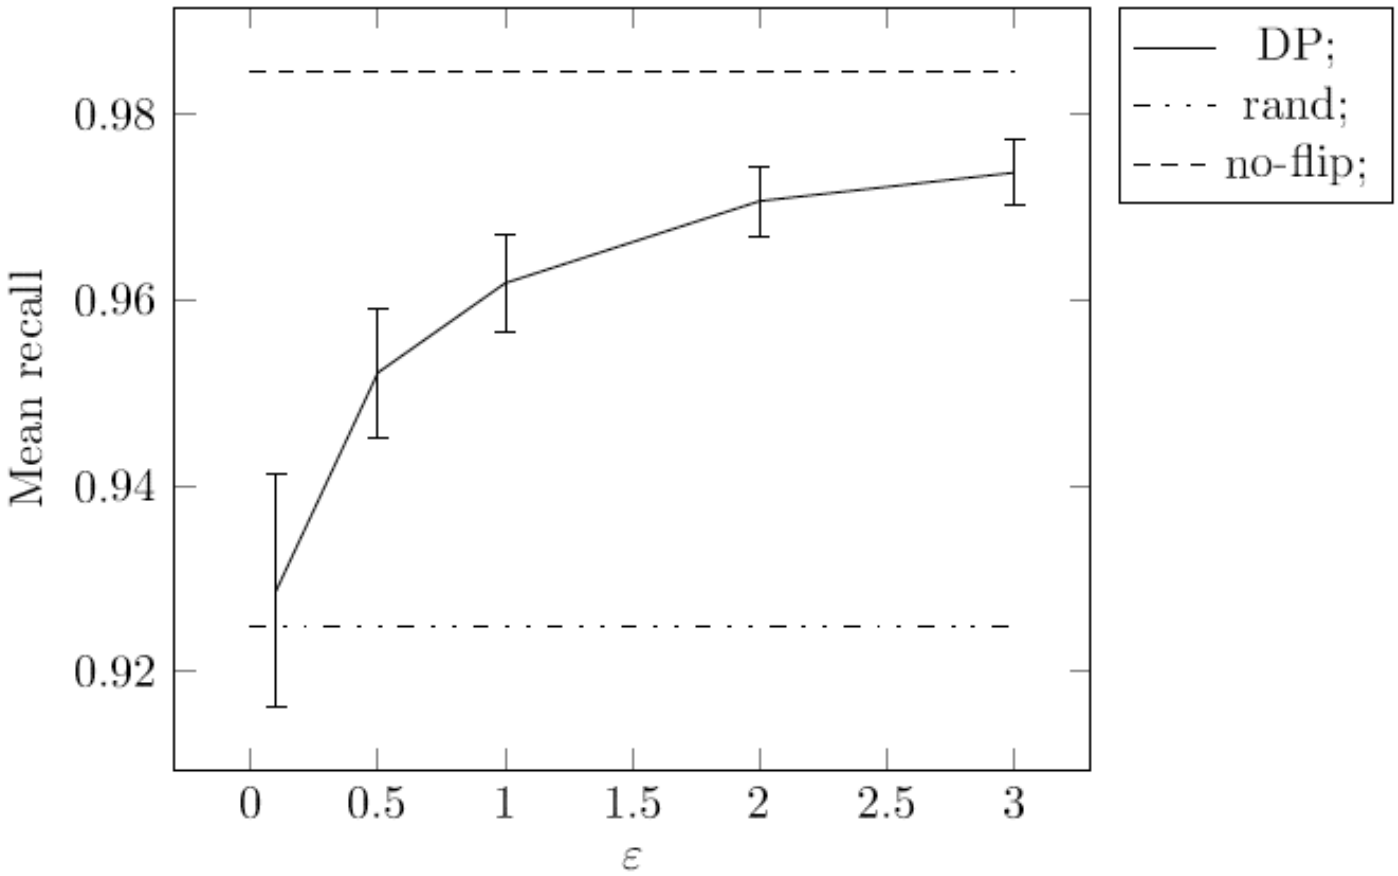

Recall : percent of search items found in the collective set of his neighbors.

Possible applications

1. Private similarity computation in P2P systems.
   - Main advantage over previous techniques : offline computation of the similarity, unlimited number of uses of the perturbed profile (no exhaustion of the privacy budget).
   - The concerned peer can publish his Blipped profile and then disappear.
2. Possibility to perform clustering using a centralized entity while providing privacy guarantees to the individuals.
   - The central entity only sees the perturbed version of the profile and cannot reconstruct the original one.
3. Non-interactive "matchmaking" protocols.
4. ...

Profile reconstruction attack

- Objective: determine for which value of ϵ, the adversary is not able to reconstruct the private profile.
- Gives an upper bound to ϵ.

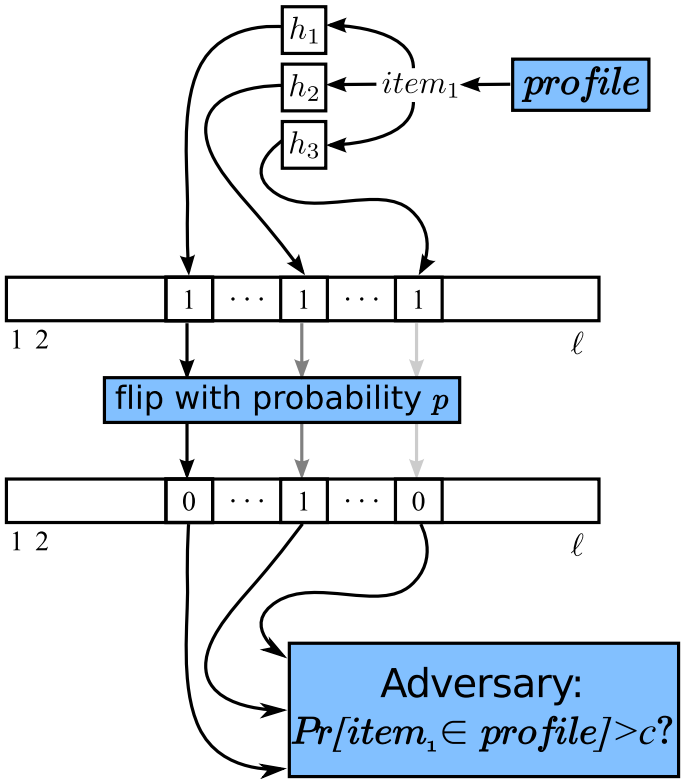

Profile reconstruction attack (experimental results)

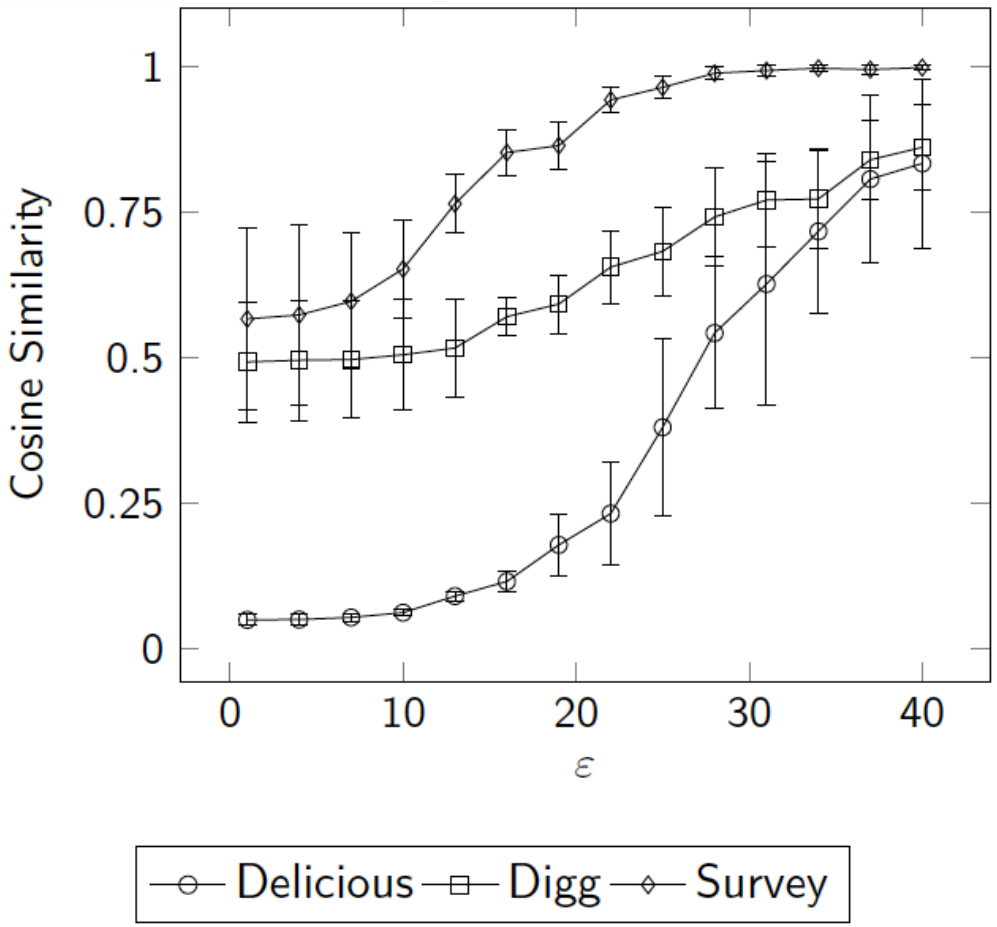

Privacy : similarity between reconstructed and original profile.

Challenging differential privacy: the case of non-interactive mechanisms (ESORICS'14)

- Join work with Raghavendran Balu and Teddy Furon (Linkmedia).
- Objective of this work : development of generic inference attacks to assess the privacy level offered by non-interactive differentially-private mechanisms.
- Other mechanism studied : differentially-private version of Johnson-Lindenstrauss transformation (Kenthapadi, Korolova, Mironov and Mishra 12).

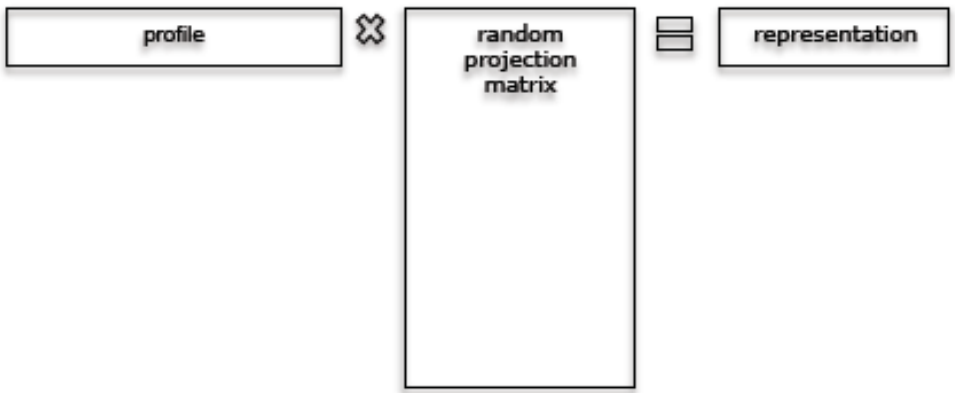

Single and joint decoders

- Single decoder : the adversary infers the presence of a single item.
- This procedure is repeated for each possible item.
- Joint decoder : the adversary performs a joint test to determine if a subset of c items is contained in the true profile.
- By testing all the subset of c items, the adversary can find the true profile (not tractable in practice).
- Theoretical analysis :The efficiency of the decoders can be analyzed in terms of the mutual information disclosed by the structure published about the presence of a particular item.

Profile reconstruction attack

Implementation of the joint decoder through a Monte Carlo Markov chain approach.

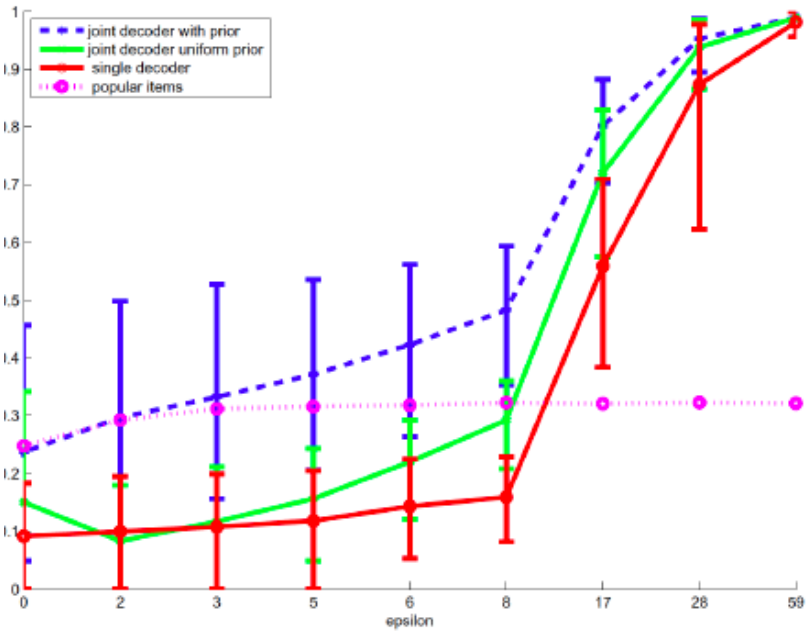

MovieLens

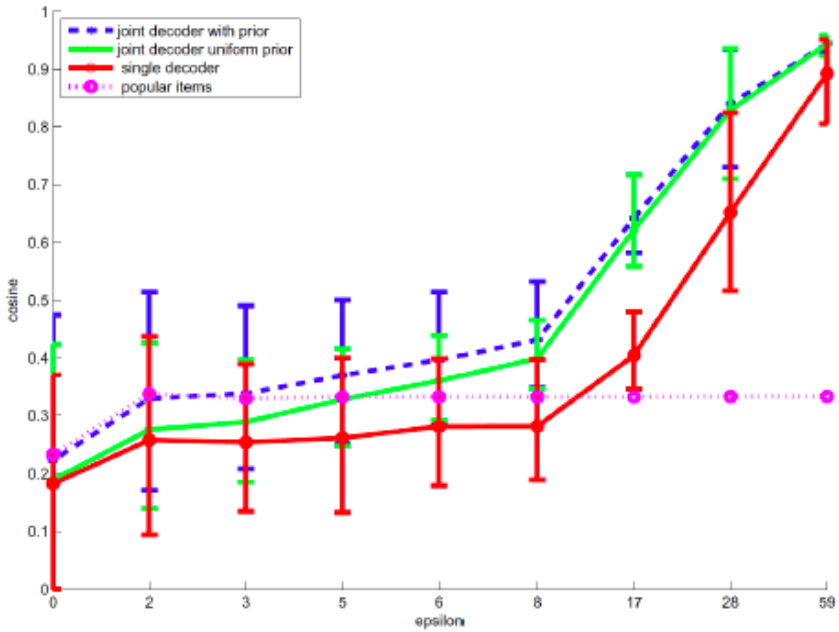

Digg

Presence of a single item

Measure of success of the attack : mean average precision of top-k ranked items.

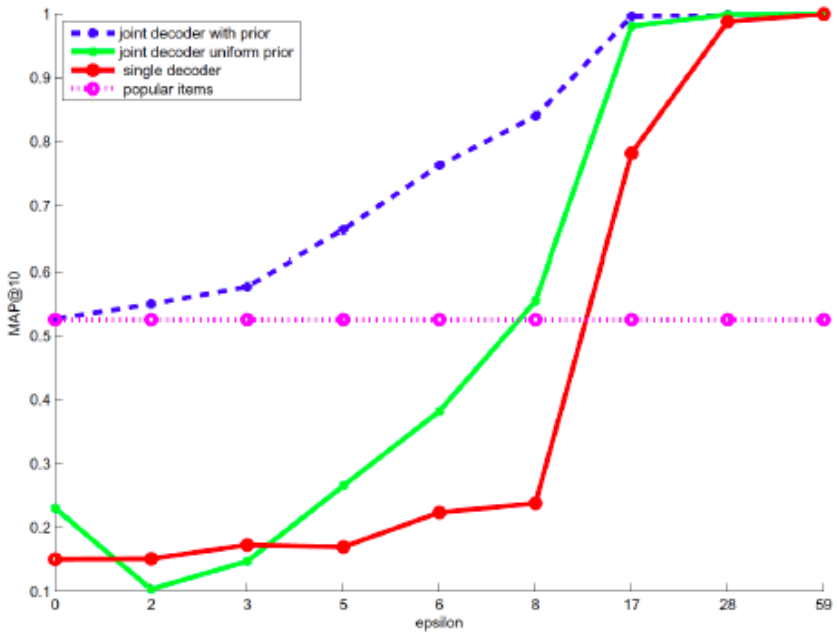

MovieLens

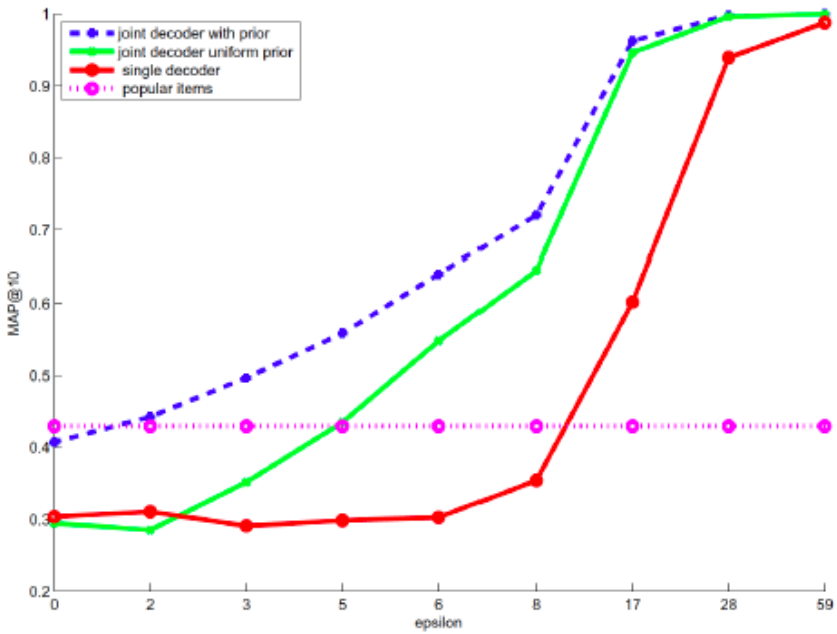

Digg

Conclusion

Theoretical analysis and experimental results demonstrate that joint decoding is more powerful than single decoding.
The attacks developed help to :

- Understand the privacy guarantees of differentially-private mechanisms.
- Experimentally tune the parameter ϵ.
- Compare different non-interactive mechanisms.

Sanitization of Call Detail Records (CDRs)

- Ongoing work with Stan Matwin and Mohammed Tuhin (Dalhousie university) and Mohammad Alaggan (Helwan university).
- Main objective : performing data mining on CDRs data.

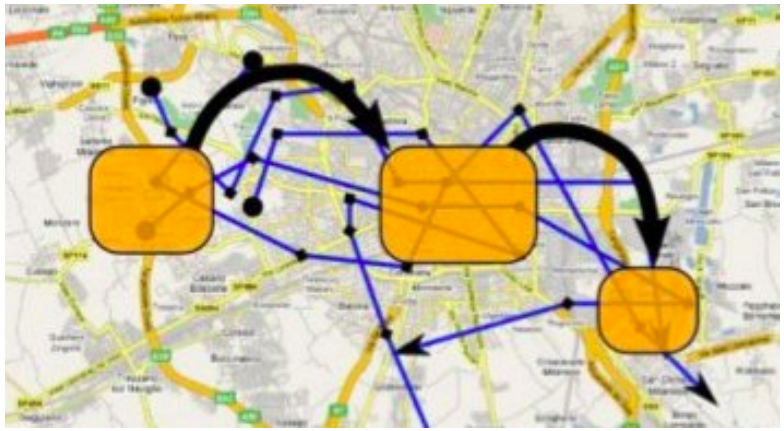

- Basic operations : counting the number of individuals and intersection between sets (possible with BLIP).
- Other possible non-interactive mechanisms for location data : aggregation of trajectories, mobility model for a group of the population, sketches measuring global properties of the dataset, synthetic dataset, ...

Thanks for your attention

Questions?

Profile distinguishing game

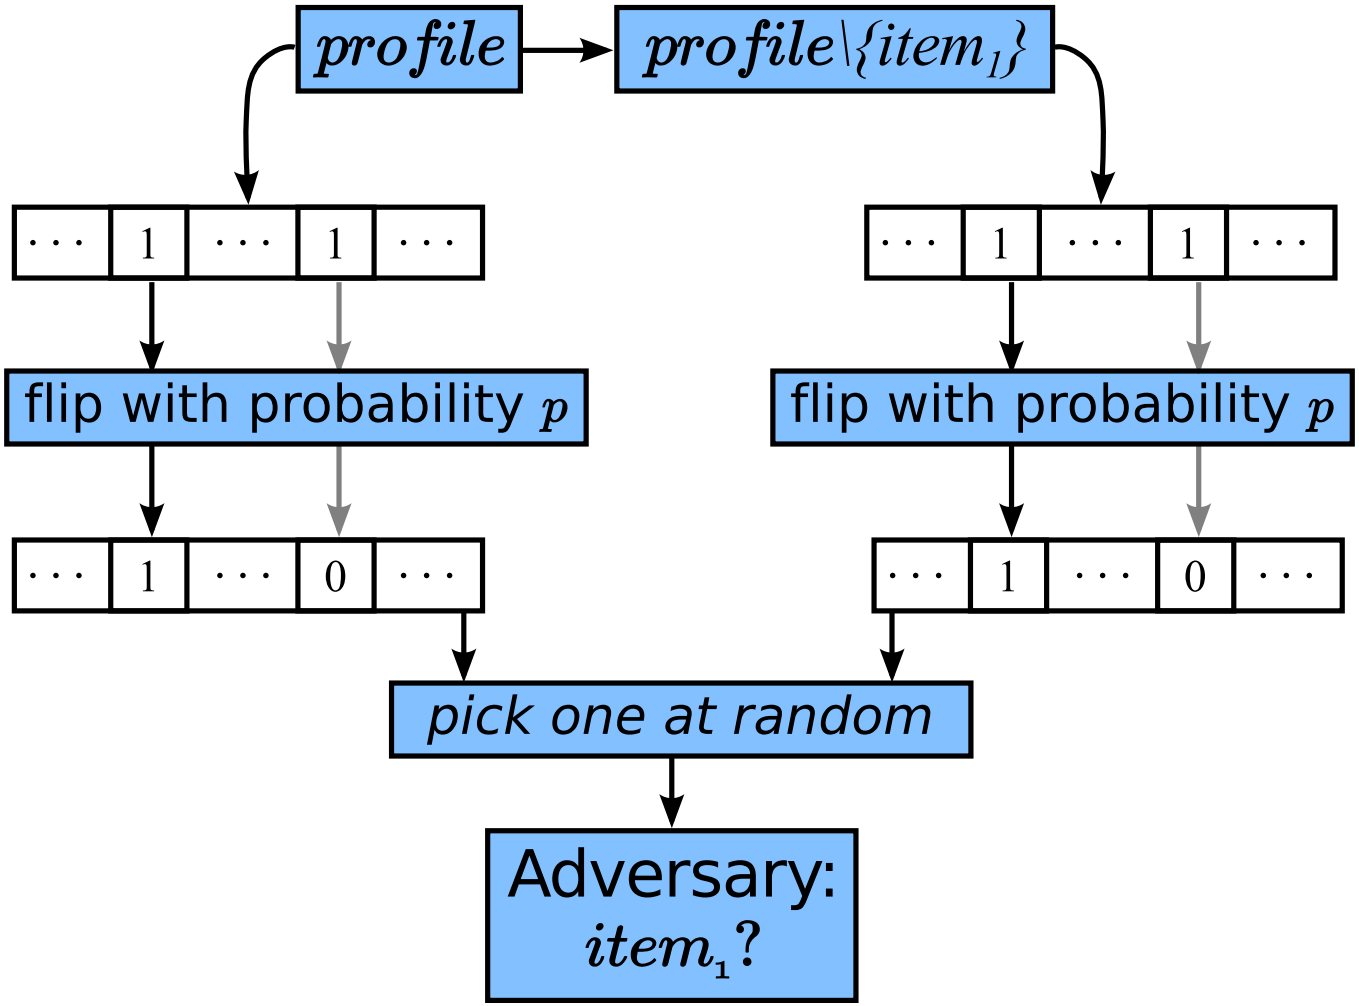

Profile distinguishing game