SORT Special issue: Privacy in statistical databases, 2011, 25-40

Coprivacy: an introduction to the theory and applications of co-operative privacy

Josep Domingo-Ferrer*

Universitat Rovira i Virgili

Abstract

We introduce the novel concept of coprivacy or co-operative privacy to make privacy preservation attractive. A protocol is coprivate if the best option for a player to preserve her privacy is to help another player in preserving his privacy. Coprivacy makes an individual's privacy preservation a goal that rationally interests other individuals: it is a matter of helping oneself by helping someone else. We formally define coprivacy in terms of Nash equilibria. We then extend the concept to: i) general coprivacy, where a helping player's utility (*i.e.* interest) may include earning functionality and security in addition to privacy; ii) mixed coprivacy, where mixed strategies and mixed Nash equilibria are allowed with some restrictions; iii) correlated coprivacy, in which Nash equilibria are replaced by correlated equilibria. Coprivacy can be applied to any peer-to-peer (P2P) protocol. We illustrate coprivacy in P2P anonymous keyword search, in content privacy in social networks, in vehicular network communications and in controlled content distribution and digital oblivion enforcement.

MSC: 91A26, 91A40, 68P20, 94A60

Keywords: Data privacy, game theory, anonymous keyword search, content privacy in social networks, vehicular networks, content distribution, digital oblivion.

1. Introduction

The motivation of the coprivacy concept and its incipient theory presented in this paper is one of double sustainability in the information society:

1. *Privacy preservation is essential to make the information society sustainable just as environment preservation is essential to make the physical world sustainable*.

* Universitat Rovira i Virgili, UNESCO Chair in Data Privacy, Department of Computer Engineering and Mathematics, Av. Països Catalans 26, E-43007 Tarragona, Catalonia. josep.domingo@urv.cat
Received: November 2010
Accepted: March 2011

This idea, which we already introduced in Domingo-Ferrer (2009) and in the conference paper Domingo-Ferrer (2010) which this article extends, should lead to clean information and communications technologies (ICT) offering functionality with minimum invasion of the privacy of individuals. Such an invasion can be regarded as a virtual pollution as harmful in the long run to the moral welfare of individuals as physical pollution is to their physical welfare. A parallel of climate change is an information society with dwindling privacy, where everyone is scared of using any service at all. Just as people's views on environment preservation have changed (they now care about environment, they require and pay for green products, etc.) and this has forced companies to change to green, the same change is happening now regarding privacy.

2. *Privacy preservation itself should be sustainable, and be achieved as effortlessly as possible as the result of rational co-operation rather than as an expensive legal requirement*. Indeed, even if privacy was acclaimed as a fundamental right by the United Nations in article 12 of the Universal Declaration of Human Rights (1948), relying on worldwide legal enforcement of privacy is nowadays quite unrealistic and is likely to stay so in the next decades. However, unlike law, technology is global and can enforce privacy worldwide, provided that privacy is achieved as the result of rational cooperation. This is the objective of the coprivacy concept and theory presented in this paper.

Two major pollutants of privacy are privacy-unfriendly security and privacy-unaware functionality. *Privacy-unfriendly security* refers to the tendency of sacrificing privacy with the excuse of security. This is partly justified by the global threat of international terrorism. With that argument, Western states have adopted shock measures on information security. Beyond the sheer technological challenge of mass-scale communications security and analysis, a new, subtler and unaddressed challenge arises: security must be increased with minimum privacy loss for the citizens. The current trend is to sacrifice privacy for alleged security: disputably, governments track phone calls, e-mails and, as seen in the Wikileaks case, social media interactions. In the private sector, privacy-unfriendly security is also present: more and more often, biometrics is enforced on customers with the argument of fighting identity theft. *Privacy-unaware* (let alone privacy-unfriendly) *functionality* is illustrated by search engines (Google, Yahoo, etc.), social networks, Web 2.0 services (*e.g.* Google Calendar, Streetview, Latitude) and so on, which concentrate on offering enticing functionality for users while completely disregarding their privacy. At most, privacy vs third parties is mentioned, but not privacy of the user vs the service provider itself, who becomes a big brother in the purest Orwellian sense.

1.1. Contribution and plan of this paper

The environmental analogy above can be pushed further by drawing inspiration on the three "R" of environment: reducing, reusing and recycling.

Reducing Re-identifiable information must be reduced. This is the idea behind database anonymization: *e.g. k*-anonymization (Samarati 2001) by means of microdata masking methods (*e.g.*, Domingo-Ferrer, Sebé and Solanas (2008)) reduces the informational content of quasi-identifiers. Reduction is also the idea behind ring and group signatures (Chaum and Van Heyst 2006, Groth 2007), which attempt to conciliate message authentication with signer privacy by reducing signer identifiability: the larger the group, the more private is the signer. Just as in the environment there are physical limits to the amount of waste reduction, in the privacy scenario there are functionality and security limits to reduction: completely eliminating quasi-identifiers dramatically reduces the utility of a data set (functionality problem); deleting the signature in a message suppresses authentication (security problem). A useful lesson that can be extracted from reduction is *privacy graduality*: privacy preservation is not all-or-nothing, it is a continuous magnitude from no privacy to full privacy preservation.

Reusing The idea of reusing is certainly in the mind of impersonators mounting replay attacks, but it can also be used by data protectors to gain privacy. Such is the case of re-sampling techniques for database privacy: an original data set with N records is re-sampled M times with replacement (where M can be even greater than N) and the resulting data set with M records is released instead of the original one. This is the idea behind synthetic data generation via multiple imputation (Rubin 1993). Re-sampling is also the idea of the tabular protection method in (Domingo-Ferrer and Mateo-Sanz (1999). However, as it happened for reduction there are functionality limitations to data reuse: the more reuse, the less data utility.

Recycling The idea of recycling is probably more intriguing and far less explored than reducing and reusing. Adapted to the privacy context, recycling can be regarded as leveraging other people's efforts to preserve their privacy to preserve one's own privacy. The environmental analog would be to share a car with other people: we leverage the other people's wish to save fuel to save fuel ourselves. Of course, whether in the privacy or the environment scenario, there is a functionality toll to this kind of recycling: one must adjust to the needs of other people. Nonetheless, we believe that *recycling has an enormous potential in privacy preservation, as it renders privacy an attractive and shared goal, thereby making it easier to achieve and thus more sustainable*. In this spirit, we next introduce a new recycling concept, called *coprivacy*, around which this proposal is centered.

Section 2 gives some background on game theory. Section 3 gives a game-theoretic definition of coprivacy and some of its generalizations. Section 4 illustrates coprivacy in the context of peer-to-peer (P2P) anonymous keyword search. Section 5 illustrates correlated coprivacy applied to content disclosure in social networks. Section 6 shows how general coprivacy applies to vehicular networks. Section 7 sketches how coprivacy can help enforcing controlled content distribution and digital oblivion. Section 8 summarizes conclusions and open research issues. A preliminary conference version of this paper appeared in Domingo-Ferrer (2010).

2. Basics of game theory

A game is a protocol between a set of *N players*, $\{1,\ldots,N\}$. Each player i has her own *set of possible strategies*, say S_i. To play the game, each player i selects a strategy $s_i \in S_i$. We will use $s = (s_1,\ldots,s_N)$ to denote the vector of strategies selected by the players and $S = \Pi_i S_i$ to denote the set of all possible ways in which players can pick strategies.

The vector of strategies $s \in S$ selected by the players determines the outcome for each player, which can be a payoff or a cost. In general, the outcome will be different for different players. To specify the game, we need to give, for each player, a preference ordering on these outcomes by giving a complete, transitive, reflexive binary relation on the set of all strategy vectors S. The simplest way to assign preferences is by assigning, for each player, a value for each outcome representing the payoff of the outcome (a negative payoff can be used to represent a cost). A function whereby player i assigns a payoff to each outcome is called a utility function and is denoted by $u_i : S \longrightarrow \mathbb{R}$.

For a strategy vector $s \in S$, we use s_i to denote the strategy played by player i and s_{-i} to denote the $(n-1)$-dimensional vector of the strategies played by all other players. With this notation, the utility $u_i(s)$ can also be expressed as $u_i(s_i, s_{-i})$.

A strategy vector $s \in S$ is a *dominant strategy solution* if, for each player i and each alternate strategy vector $s' \in S$, it holds that

$$u_i(s_i, s'_{-i}) \geq u_i(s'_i, s'_{-i}) \tag{1}$$

In plain words, a dominant strategy s is the best strategy for each player i, independently of the strategies played by all other players.

A strategy vector $s \in S$ is said to be a *Nash equilibrium* (Nash 1951) if, for all players i and each alternate strategy $s'_i \in S_i$, it holds that

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

In plain words, no player i can change her chosen strategy from s_i to s'_i and thereby improve her payoff, assuming that all other players stick to the strategies they have chosen in s. A Nash equilibrium is self-enforcing in the sense that once the players

are playing such a solution, it is in every player's best interest to stick to her strategy. Clearly, a dominant strategy solution is a Nash equilibrium. Moreover, if the solution is strictly dominant (*i.e.* when the inequality in Expression (1) is strict), it is also the unique Nash equilibrium. See Nisan, Roughgarden, Tardos and Vazirani (2007) for further background on game theory.

3. Coprivacy and its generalizations

We introduce in this section the novel concept of coprivacy in a community of peers, whereby one peer recycles to her privacy's benefit the efforts of other peers to maintain their own privacy. Informally, there is coprivacy when the best option for a peer to preserve her privacy is to help another peer in preserving his privacy. The great advantage is that *coprivacy makes privacy preservation of each specific individual a goal that interests other individuals*: therefore, privacy preservation becomes *more attractive* and hence *easier to achieve and more sustainable*. A game-theoretic formalization of coprivacy follows.

Definition 1 (Coprivacy) *Let Π be a game with self-interested, rational peer players $P^1,\ldots,P^N$, and an optional system player P^0. Each player may have leaked a different amount of private information to the rest of players before the game starts. The game is as follows: i) P^1 selects one player P^k with $k \in \{0\}\cup\{2,\cdots,N\}$ and submits a request to P^k; ii) If $k=0$, P^0 always processes P^1's request; if $k>1$, P^k decides whether to process P^1's request (which may involve accessing the system player on P^1's behalf) or reject it. The players' strategies are $S^0=\{s_1^0\}$ (process P^1's request); $S^1=\{s_0^1,s_2^1,\cdots,s_N^1\}$, where s_j^1 means that P^1 selects P^j; for $i>1$, $S^i=\{s_1^i,s_2^i\}$, where s_1^i means processing P^1's request and s_2^i rejecting it. Game Π is said to be* coprivate *with respect to the set $U=(u_1,\cdots,u_N)$ of privacy utility functions if, for some $k>1$, a peer P^k exists such that (s_k^1,s_1^k) is a pure strategy Nash equilibrium between P^1 and P^k, that is, if the best strategy for P^1 is to request help to P^k and the best strategy for P^k is to provide the requested help.*

Note that the notions of privacy utility function and therefore of coprivacy are based on the aforementioned privacy graduality: one can have a varying degree of privacy preservation, hence it makes sense to trade it off. In the environmental analogy, coprivacy is a recycling concept which involves trading off waste reduction among players. A quantification of coprivacy follows:

Definition 2 (δ-Coprivacy) *Given $\delta\in[0,1]$, the game of Definition 1 is said to be δ-coprivate with respect to the set $U=(u_1,\cdots,u_N)$ of privacy utility functions if the probability of it being coprivate for U is at least δ.*

The following extensions of coprivacy are conceivable:

- **General coprivacy** can be defined by replacing the set U of privacy utility functions in Definition 1 with a set $\mathscr{U}$ of general utility functions for peer players P^k combining privacy preservation with security and/or functionality. In general coprivacy, the interests of peers include, in addition to privacy, functionality and/or security.
- **General δ-coprivacy** can be defined by replacing U with $\mathscr{U}$ in Definition 2.
- **Mixed coprivacy** results if one allows mixed strategies for players and replaces the requirement of pure strategy Nash equilibrium in Definition 1 by a mixed strategy Nash equilibrium. The good point of mixed coprivacy is that a theorem by Nash (Nash 1951) guarantees that any game with a finite set of players and a finite set of strategies has a mixed strategy Nash equilibrium, and is therefore *mixedly coprivate*.
- **Correlated coprivacy** results if one replaces the requirement of pure Nash equilibrium in Definition 1 by a correlated equilibrium. Indeed, the outcome of independent rational behavior by users, provided by Nash equilibria, can be inferior to a centrally designed outcome. Correlated equilibria resulting from coordination of strategies may give a higher outcome. We will illustrate this in Section 5 below. In correlated equilibria, players do not have any incentive to deviate from their corresponding equilibrium strategies. An approximation to correlated equilibria are ε-correlated equilibria, in which players have at most an incentive $\varepsilon > 0$ to deviate from their corresponding equilibrium strategies. The advantage of ε-correlated equilibria is that they can always be reached by distributed heuristics run by a set of autonomous players without centrally designed strategies.
- The above extensions can be combined to yield **mixed general coprivacy** and **correlated general coprivacy**. Since mixed coprivacy is always achievable if any mixed strategy is valid for any player, **mixed δ-coprivacy** and **mixed general δ-coprivacy** only make sense when players have boundary conditions that define a subset of feasible mixed strategies.

A *coprivate protocol* is a protocol based on a coprivate game. If a privacy preservation problem can be solved by a coprivate protocol, the advantage is that it is in a player's rational privacy interest to help other players to preserve their privacy. We next give an example to show that the coprivacy concept is latent in existing protocols. More examples of the potential of coprivacy follow in the next sections.

Example 1 (Coprivacy in anonymous communication) *The success of the well-known system Tor (`http://www.torproject.org`) for anonymous communication, made even more famous by Wikileaks, can be explained by coprivacy. As hinted in the Tor website, "each new user and relay provides additional diversity, enhancing Tor's ability to put control over your security and privacy back into your hands". Therefore, using Tor is not only good for one's own privacy, but for other people's privacy as well.*

4. Coprivacy in P2P anonymous keyword search

Private information retrieval (PIR) is normally modeled as a game between two players: a user and a database. The user retrieves some item from the database without the latter learning which item was retrieved. Most PIR protocols are ill-suited to provide PIR from a search engine or large database, not only because their computational complexity is linear in the size of the database, but also because they (unrealistically) assume active cooperation by the database in the PIR protocol.

Pragmatic approaches to guarantee some query privacy have therefore been based so far on two relaxations of PIR: standalone and peer-to-peer (P2P). In the standalone approach, a program running locally in the user's computer either keeps submitting fake queries to cover the user's real queries (TrackMeNot, Howe and Nissenbaum 2009)) or masks the real query keywords with additional fake keywords (GooPIR, Domingo-Ferrer, Solanas and Castellà-Roca 2009)). In the P2P approach, a user gets her queries submitted by other users in the P2P community; in this way, the database still learns which item is being retrieved, but it cannot obtain the real query histories of users, which become diffused among the peer users, thereby achieving anonymous keyword search. We first proposed a P2P anonymous keyword search system in Domingo-Ferrer, Bras-Amorós, Wu and Manjón (2009).

Consider a system with N peers P^1 to P^N, who are interested in querying a database DB playing the role of system player P^0. If any P^i originates a query for submission to DB, she can send the query directly to DB or ask some other peer to submit the query on P^i's behalf and return the query results.

More formally, the strategies available for a requesting P^i are:

Sii: P^i submits her query directly to DB;
Sij: P^i forwards her query to P^j, for some $j \neq i$, and requests P^j to submit the query on P^i's behalf.

When receiving P^i's query, P^j has two possible strategies:

Tji: P^j submits P^i's query to DB and returns the answer to P^i;
Tjj: P^j ignores P^i's query and does nothing.

Let $X^i(t)$ be the set of queries originated by P^i up to time t. Let $Y^i(t)$ be the set of queries submitted to DB by P^i up to time t. For each query x_r^i in $X^i(t)$, define $F^i(x_r^i,t)$ as the set of players to whom P^i has forwarded x_r^i for submission up to time t. The players in $F^i(x_r^i,t)$ can be associated relative frequencies as follows: for $j = 1$ to N with $j \neq i$, let $f^{ij}(x_r^i,t)$ be the relative frequency with which P^i has forwarded x_r^i to player P^j, up to time t.

The privacy utility function for P^i should reflect the following intuitions: (i) the more homogeneous the relative frequencies of queries in $Y^i(t)$, the more private stay

the interests of P^i vs DB; (ii) the more homogeneous the relative frequencies of peers in $F^i(x_r^i,t)$ for every $x_r^i \in X^i(t)$, the more private stay the interests of P^i vs the other peers.

Given a random variable Z taking values $z_1, z_2, \ldots, z_n$ with probabilities $p_1, p_2, \ldots, p_n$, respectively, Shannon's entropy (Shannon 1948) is a measure of uncertainty defined as

$$H(Z) = -\sum_{i=1}^{n} p_i \log_2 p_i$$

The more homogeneous the p_i, the higher is $H(Z)$: the rationale is that the outcome of Z becomes more uncertain as the p_i become more homogeneous. The maximum $H(Z)$ is reached when $p_1 = \cdots = p_n = 1/n$.

By assimilating $Y^i(t)$ and $F^i(x_r^i,t)$ to random variables and relative frequencies to probabilities, intuition (i) above can be expressed as maximizing $H(Y^i(t))$ and intuition (ii) as maximizing $H(F^i(x_r^i,t))$ for all $x_r^i \in X^i(t)$. Hence, those Shannon entropies are reasonable privacy utility functions for P^i. When P^i generates a query x_r^i at time $t+1$:

- P^i chooses *Sii* (direct submission) if $H(Y^i(t+1)) \geq H(Y^i(t))$, where $Y^i(t+1) = Y^i(t) \cup \{x_r^i\}$;
- Otherwise P^i chooses *Sik* (forwarding the query to P^k), where

$$k = \arg \max_{j \in \{1,\cdot,N\} \setminus \{i\}} H(F^i(x_r^i, t+1)). \tag{2}$$

In plain words, if direct submission decreases privacy vs DB, the query is forwarded to the player P^k vs whom the privacy loss is minimum. Note that if P^i forwards her query to a player P^j, P^i always incurs some privacy loss vs P^j, because P^j knows the query has been generated by P^i. Therefore, the best policy is to distribute the successive submissions of a certain query x_r^i as evenly as possible among the various peers. This is what the choice of k in Expression (2) attempts.

When P^k receives x_r^i, it proceeds as follows:

- P^k chooses *Tki* (submitting x_r^i) if $H(Y^k(t+1)) > H(Y^k(t))$, where $Y^k(t+1) = Y^k(t) \cup \{x_r^i\}$;
- Otherwise P^k chooses *Tkk* (ignoring x_r^i).

In plain words, P^k submits x_r^i only if doing so increases her privacy vs the DB. If P^k ignores x_r^i, then P^i will have to look for a second best player to submit x_r^i (and a third best if the second best ignores x_r^i, and so on). If, after a number of attempts to be decided by P^i, no peer is found who is willing to help, then P^i must submit x_r^i herself.

If P^i's best strategy is *Sik* and P^k best strategy is *Tki*, then (Sik, Tki) is a pure-strategy Nash equilibrium between P^i and P^k and *there is coprivacy* between P^i and P^k.

We give a detailed formalization and empirical results for the N-player P2P anonymous keyword search game in the manuscript Domingo-Ferrer and González-Nicolás (2011).

5. Correlated coprivacy in social networks

Social networks (SNs) have become an important web service with a broad range of applications: collaborative work, collaborative service rating, resource sharing, friend search, etc. Facebook, MySpace, Xing, etc., are well-known examples. In an SN, a user publishes and shares information and services.

There are two types of privacy in SNs:

- *Content privacy.* The information a user publishes clearly affects her privacy. Recently, a privacy risk score (Liu and Terzi 2009) has been proposed for the user to evaluate the privacy risk caused by the publication of a certain information. Let the information attributes published by the users in an SN be labeled from 1 to n. Then the privacy score risk of user j is

$$PR(j) = \sum_{i=1}^{n} \sum_{k=1}^{\ell} \beta_{ik} \times V(i, j, k)$$

  where $V(i, j, k)$ is the visibility of user j's value for attribute i to users which are at most k links away from j and β_{ik} is the sensitivity of attribute i vs those users.
- *Relationship privacy.* In some SNs, the user can specify how much it trusts other users, by assigning them a trust level. It is also possible to establish several types of relationships among users (like "colleague of", "friend of", etc.). The trust level and the relationship type are used to decide whether access is granted to resources and services being offered (*access rule*). The availability of information on relationships (trust level, relationship type) has increased with the advent of the Semantic Web and raises privacy concerns: knowing who is trusted by whom and to what extent discloses a lot about the user's thoughts and feelings. For a list of related abuses see Barnes (2006). In Domingo-Ferrer, Viejo, Sebé and González-Nicolás (2008), we described a new protocol offering private relationships in an SN while allowing resource access through indirect relationships without requiring a mediating trusted third party.

We focus here on content privacy in SNs. A possible privacy-functionality score for user j reflecting the utility the user derives from participating in an SN is

$$PRF(j) = \frac{\sum_{j'=1, j' \neq j}^{N} \sum_{i=1}^{n} \sum_{k=1}^{\ell} \beta_{ik} V(i, j', k) I(j, j', k)}{1 + PR(j)}$$

$$= \frac{\sum_{j'=1, j' \neq j}^{N} \sum_{i=1}^{n} \sum_{k=1}^{\ell} \beta_{ik} V(i, j', k) I(j, j', k)}{1 + \sum_{i=1}^{n} \sum_{k=1}^{\ell} \beta_{ik} V(i, j, k)}$$

where $I(j, j', k)$ is 1 if j and j' are k links away from each other, and it is 0 otherwise.

Note that:

- $PRF(j)$ decreases as the privacy score $PR(j)$ in its denominator increases, that is, as user j discloses more of her attributes.
- $PRF(j)$ increases as its numerator increases; this numerator adds up the components of privacy scores of users $j' \neq j$ due to those users disclosing attribute values to j.

The dichotomous version of the above privacy-functionality score, for the case where an attribute is simply either made public or kept secret, is:

$$PRF_2(j) = \frac{\sum_{j'=1, j' \neq j}^{N} \sum_{i=1}^{n} \beta_i V(i, j')}{1 + PR(j)}$$

$$= \frac{\sum_{j'=1, j' \neq j}^{N} \sum_{i=1}^{n} \beta_i V(i, j')}{1 + \sum_{i=1}^{n} \beta_i V(i, j)} \quad (3)$$

If we regard $PRF(j)$ as a game-theoretic utility function (Tardos and Vazirani 2007), the higher $PRF(j)$, the higher the utility for user j.

For instance, take a strategy vector $s = (s_1, \ldots, s_N)$ formed by the strategies *independently and selfishly* chosen by all users and consider the dichotomous case, that is, let the utility incurred by user j under strategy s be $u_j(s) = PRF_2(j)$. It is easy to see (and it is formally shown in Domingo-Ferrer (2010b) that rational and independent choice of strategies leads to a Nash equilibrium where no user offers any information on the SN, which results in an SN collapse. See Example 2 below.

A similar pessimistic result is known for the P2P file sharing game, in which the system goal is to leverage the upload bandwidth of the downloading peers: the dominant strategy is for all peers to attempt "free-riding", that is, to refuse to upload (Babaioff, Chuang and Feldman 2007), which causes the system to collapse.

Example 2 *The simplest version of the above game is one with two users having each one attribute, which they may decide to keep hidden (a strategy denoted by H, which implies visibility 0 for the attribute) or publish (a strategy denoted by P, which implies visibility 1). Assuming a sensitivity $\beta = 1$ for that attribute and using $u_j(s) = PRF_2(j)$, the user utilities for each possible strategy vector are as follows:*

$$u_1(H,H) = 0; u_1(H,P) = 1; u_1(P,H) = 0; u_1(P,P) = 1/2$$
$$u_2(H,H) = 0; u_2(H,P) = 0; u_2(P,H) = 1; u_1(P,P) = 1/2$$

This simple game can be expressed in matrix form:

User 2	*H*	*P*
User 1		
H	*0* *0*	*0* *1*
P	*1* *0*	*1/2* *1/2*

The above matrix corresponds to the Prisoner's Dilemma (Tardos and Vazirani 2007), perhaps the best-known and best-studied game. Consistently with our argument for the general case, it turns out that (H,H) *is a dominant strategy, because:*

$$u_1(H,P) = 1 \geq u_1(P,P) = 1/2; u_1(H,H) = 0 \geq u_1(P,H) = 0$$

$$u_2(P,H) = 1 \geq u_1(P,P) = 1/2; u_2(H,H) = 0 \geq u_2(H,P) = 0$$

The second and fourth equations above guarantee that (H,H) *is a Nash equilibrium (in fact, the only one). The Prisoner's Dilemma with* $N > 2$ *users is known as the Pollution Game (Tardos and Vazirani 2007) and corresponds to the dichotomous SN game considered above.*

The outcome of independent rational behavior by users, provided by Nash equilibria and dominant strategies, can be inferior to a centrally designed outcome. This is clearly seen in Example 2: the strategy (P, P) would give more utility than (H, H) to *both* users. However, usually no trusted third-party accepted by all users is available to enforce correlated strategies; in that situation, the problem is how User 1 (resp. User 2) can guess whether User 2 (resp. User 1) will choose P.

Using a solution based on cryptographic protocols for bitwise fair exchange of secrets would be an option, but it seems impractical in current social networks, as it would require a cryptographic infrastructure, unavailable in most SNs.

A more practical solution to this problem may be based on direct reciprocity (*i.e.* tit-for-tat) or reputation, two approaches largely used in the context of P2P file-sharing systems. We describe in Domingo-Ferrer (2010b) two correlated (actually ε-correlated) equilibrium heuristic protocols based on tit-for-tat and reputation, respectively. They are intended as "assistants" to the human user of the SN in deciding whether to disclose an attribute to another user; however, the ultimate decision belongs to the human, who may quit and renounce to reach the equilibrium.

Those heuristic protocols offer *ε-correlated general coprivacy*, referred to a utility combining privacy and functionality.

6. General coprivacy in vehicular networks

Vehicular *ad hoc* networks permitting car-to-car communication are expected to be available in cars manufactured in the near future. Several standards for VANET communication are under way both in the United States (DSRC, Dedicated Short Range Communications, IEEE 802.11p) and Europe (C2C Consortium). We argue that VANETs must provide functionality, security and privacy and are therefore an application where general coprivacy can be used:

Functionality. The main *raison d'être* of VANETs is to allow vehicles to timely disseminate announcement messages about *current* road conditions (*e.g.* icy road, traffic jam, etc.) to other vehicles, in order to improve traffic safety and efficiency.

Security. Announcement messages must be trustworthy, because false messages could seriously disrupt traffic, cause accidents and/or cause certain areas to become deserted and thus an easy prey for criminals. A posteriori security consists of punishing vehicles that have been proven to have originated false messages (*e.g.* Lin, Sun, Ho and Shen (2007)); hence, means are required to identify malicious vehicles, for example digital signatures. *A priori* security is an alternative or a complement whereby one attempts to prevent the generation of false messages (*e.g.* Raya, Aziz and Hubaux (2006)): a message is given credit only if it is has been endorsed by a number of nearby vehicles greater than a certain threshold.

Privacy. It would not be very fair if the collaboration of a driver to improve traffic safety and efficiency (functionality) by generating or endorsing announcements forced her to disclose her identity and location. Note that knowing someone's mobility pattern reveals a lot of private information: the driving style leaks information about an individual's character (nervous, calm), her whereabouts tell about her work and social habits, etc. Privacy can be added to *a posteriori* security by using pseudonyms or advanced cryptography like group signatures. Adding privacy to *a priori* security may imply vulnerability against the Sybil attack, whereby a vehicle generates a false message and takes advantage of anonymity to compute itself as many endorsements as required. We have proposed in Daza, Domingo-Ferrer, Sebé and Viejo (2009) a private *a priori* scheme based on threshold signatures which is resistant against the Sybil attack and provides irrevocable anonymity to cars generating or endorsing messages.

Security is a *must* in VANETs and cannot be traded off. Therefore *the general coprivacy that applies in vehicular networks involves a utility function combining functionality and privacy*. General coprivacy is applicable to VANETs in the following sense:

- The more privacy players allow to another player, the more announcements (functionality) they can expect from that player.

- Conversely, the more announcements players originate, the more privacy for other announcing players: indeed, the more cars originate an announcement "icy road near longitude X latitude Y", the more private the originators stay (this is the "reusing" principle mentioned above).

7. Controlled content distribution and digital oblivion

In conventional multicast transmission one sender sends the same content to a set of receivers. This precludes fingerprinting the copy obtained by each receiver (in view of redistribution control and other applications). A straightforward alternative is for the sender to separately fingerprint and send in unicast one copy of the content for each receiver. This approach is not scalable and may implode the sender.

Distributed multicast of fingerprinted content can be modeled as a coprivate protocol. Indeed, mechanism design can be used to craft a protocol such that content receivers rationally co-operate in fingerprinting and further spreading the content in a tree-like fashion. If fingerprinting at each forwarding step is anonymous Pfitzmann and Waidner 1997, Bo, Piyuan and Wenzheng 2007, Domingo-Ferrer 1999), honest receivers will stay anonymous and free from false accusation, but unlawful redistributors will be traceable.

A related problem is the lack of digital forgetting in the information society. Digital storage allows perfect and unlimited remembering. However, the right of an individual to enforce oblivion for pieces of information about her is part of her fundamental right to privacy. Enforcing expiration dates for content has been championed as a solution in Mayer-Schönberger (2009), but in a way that depends on trusted storage devices deleting the content after its expiration date. Alternative hardware approaches based on employing smart cards on the user side to process encrypted content (Domingo-Ferrer 1997) could also be envisioned, whereby the smart card would not decrypt the content after its expiration date. However, such devices do not currently exist; worse yet, placing trust in the hardware (storage devices, smart cards, etc.) to implement information protection policies has proven to be a flawed approach: *e.g.*, hardware copy prevention mechanisms for CDs and DVDs were easily bypassed.

Digital oblivion via expiration date enforcement can be reached through a coprivate protocol (Domingo-Ferrer 2011). The idea is just to fingerprint expiration dates in the content. This allows identifying and punishing whoever spreads or uses content past the expiration date. If fingerprinting is asymmetric and/or anonymous, it will not be possible to falsely accuse honest content receivers. The problem then reduces to distributed multicast of asymmetrically/anonymously fingerprinted content, which is approachable a coprivate protocols, as hinted above. With anonymous fingerprinting, honest players preserve their privacy. Therefore, the receivers must honestly contribute to the content source's privacy preservation (oblivion enforcement by anonymously fingerprinting any forwarded content with its expiration date) to preserve their own privacy. Hence, the solution is a coprivate protocol.

8. Conclusions and research directions

We have introduced in this paper the novel concept of coprivacy, as well as an incipient generalization theory on it. The main contribution of coprivacy is to make data privacy an attractive feature, especially in peer-to-peer applications:

- In many situations, players can better preserve their own privacy if they help other players in preserving theirs. We say that those situations can be handled by so-called coprivate protocols.
- In other situations, the utility of players consists of a combination of privacy plus security and/or functionality. If they can increase their own utility by helping others in increasing theirs, the situation can be handled by a generally coprivate protocol.

We have sketched the potential of coprivate protocols in very diverse areas: P2P anonymous keyword search, content disclosure in social networks, communication in vehicular networks, controlled content distribution and digital oblivion implementation.

Future research directions include developing the theory of coprivacy in the following non-exhaustive directions:

- Develop a theory of coprivacy which, given a privacy preservation problem and a parameter $\delta \in [0,1]$, can answer under which conditions a δ-coprivate game (*i.e.* protocol) that solves the problem exists.
- Elaborate a theory of general coprivacy which also takes security and functionality into account. In this generalization, the Nash or the correlated equilibrium that characterizes coprivacy is to be reached by considering utilities which combine the privacy with the security and/or the functionality obtained by the players.
- Elaborate a theory of mixed coprivacy to characterize when mixed strategies and therefore mixed coprivacy make sense for utilities about privacy, security and functionality.
- Create new cryptographic protocols to implement the privacy graduality needed in coprivacy. Specifically, *ad hoc* broadcast encryption and anonymous *ad hoc* broadcast encryption inspired in Wu, Mu, Susilo, Qin and Domingo-Ferrer (2009), (n,N)-anonymity signatures and some multiparty computation protocols for social networks are needed.

Acknowledgments and disclaimer

This work was partly funded by the Spanish Government through projects TSI2007-65406-C03-01 “E-AEGIS” and CONSOLIDER INGENIO 2010 CSD2007-0004 “ARES”, and by the Government of Catalonia through grant 2009 SGR 1135. The author is partly supported as an ICREA-Acadèmia researcher by the Government of Catalonia. He holds

the UNESCO Chair in Data Privacy, but the views expressed in this paper are his own and do not commit UNESCO.

References

Babaioff, M., Chuang. J. and Feldman, M. (2007). Incentives in peer-to-peer systems, in N. Nisan, T. Roughgarden, É. Tardos and V. V. Vazirani (eds.), *Algorithmic Game Theory*, Cambridge University Press, 593–611.

Barnes, S. B. (2006). A privacy paradox: social networking in the United States, *First Monday*, 11.

Bo, Y., Piyuan, L. and Wenzheng, Z. (2007). An efficient anonymous fingerprinting protocol, in *Computational Intelligence and Security*, Springer, LNCS 4456, 824–832.

Chaum, D. and van Heyst, E. (2006). Group signatures, in *Advances in Cryptology-Eurocrypt'91*, Springer, LNCS 547, 257–265.

Daza, V., Domingo-Ferrer, J., Sebé, F. and Viejo, A. (2009). Trustworthy privacy-preserving car-generated announcements in vehicular ad hoc networks, *IEEE Transactions on Vehicular Technology*, 58, 1876–1886.

Domingo-Ferrer, J. (1997). Multi-application smart cards and encrypted data processing, *Future Generation Computer Systems*, 13, 65–74.

Domingo-Ferrer, J. (1999). Anonymous fingerprinting based on committed oblivious transfer, in *Public Key Cryptography-PKC 99*, Springer, LNCS 1560, 43–52.

Domingo-Ferrer, J. and Mateo-Sanz, J. M. (1999). On resampling for statistical confidentiality in contingency tables, *Computers & Mathematics with Applications*, 38, 13–32.

Domingo-Ferrer, J., Sebé, F. and Solanas, A. (2008). A polynomial-time approximation to optimal multivariate microaggregation, *Computers & Mathematics with Applications*, 55, 717–732.

Domingo-Ferrer, J., Viejo, A., Sebé, F. and González-Nicolás, Ú. (2008). Privacy homomorphisms for social networks with private relationships, *Computer Networks*, 52, 3007–3016.

Domingo-Ferrer, J. (2009). The functionality-security-privacy game, in *Modeling Decisions for Artificial Intelligence-MDAI 2009*, Springer, LNCS 5861, 92–101.

Domingo-Ferrer, J., Solanas, A. and Castellà-Roca, J. (2009). $h(k)$-Private information retrieval from privacy-uncooperative queryable databases, *Online Information Review*, 33, 720–744.

Domingo-Ferrer, J., Bras-Amorós, M., Wu, Q. and Manjón, J. (2009). User-private information retrieval based on a peer-to-peer community, *Data and Knowledge Engineering*, 68, 1237–1252.

Domingo-Ferrer, J. (2010). Coprivacy: towards a theory of sustainable privacy, in *Privacy in Statistical Databases-PSD 2010*, Springer, LNCS 6344, 258–268.

Domingo-Ferrer, J. and González-Nicolás, Ú. (2011). Rational behaviour in peer-to-peer anonymous keyword search, manuscript.

Domingo-Ferrer, J. (2010). Rational privacy disclosure in social networks, in *Modeling Decisions for Artificial Intelligence-MDAI 2010*, Springer, LNCS 6408, 255–265.

Domingo-Ferrer, J. (2011). Rational enforcement of digital oblivion, in *4th International Workshop on Privacy and Anonymity in the Information Society (PAIS 2011)*, ACM Digital Library (to appear).

Groth, J. (2007). Fully anonymous group signatures without random oracles, in *Proc. of ASIACRYPT 2007*, LNCS 4833, 164–180.

Howe, D. C. and Nissenbaum, H. (2009). TrackMeNot: Resisting surveillance in web search, in *Lessons from the Identity Trail*, Oxford University Press, 409–428.

Lin, X., Sun, X., Ho, P.-H. and Shen, X. (2007). GSIS: A secure and privacy-preserving protocol for vehicular communications, *IEEE Transactions on Vehicular Communications*, 56, 3442–3456.

Liu, K. and Terzi, E. (2009). A framework for computing the privacy scores of users in online social networks, in *Proc. of ICDM 2009-The 9th IEEE International Conference on Data Mining*, 288–297.

Mayer-Schönberger, V. (2009). *The Virtue of Forgetting in the Digital Age*, Princeton University Press.

Nash, J. (1951). Non-cooperative games, *Annals of Mathematics*, 54, 289–295.

Nisan, N., Roughgarden, T., Tardos, É. and Vazirani, V. V. eds. (2007). *Algorithmic Game Theory*, Cambridge University Press.

Pfitzmann, B. and Waidner, M. (1997). Anonymous fingerprinting, in *Advances in Cryptology-EUROCRYPT 1997*, Springer, LNCS 1233, 88–102.

Raya, M., Aziz, A. and Hubaux, J.-P. (2006). Efficient secure aggregation in VANETs, in *Proc. of 3rd Intl. Workshop on Vehicular Ad Hoc Networks-VANET*, 67–75.

Rubin, D. B. (1993). Discussion on statistical disclosure limitation, *Journal of Official Statistics*, 9, 461–468.

Samarati, P. (2001). Protecting respondents' identities in microdata release, *IEEE Transactions on Knowledge and Data Engineering*, 13, 1010–1027.

Shannon, C. (1948). A mathematical theory of communication, *Bell Systems Technical Journal*, 27, 379–423 and 623–656.

Tardos, É. and Vazirani, V. V. (2007). Basic solution concepts and computational issues, in N. Nisan, T. Roughgarden, É. Tardos and V. V. Vazirani (eds.), *Algorithmic Game Theory*, Cambridge University Press, 3–28.

Wu, Q., Mu, Y., Susilo, W., Qin, B. and Domingo-Ferrer, J. (2009). Asymmetric group key agreement, in *Advances in Cryptology-EUROCRYPT 2009*, Springer, LNCS 5479, 153–170.