On the Tradeoff among Trust, Privacy, and Cost in Incentive-Based Networks

Alessandro Aldini¹, Alessandro Bogliolo¹, Carlos Ballester Lafuente²,
and Jean-Marc Seigneur²

¹ Department of Base Science and Fundamentals, University of Urbino, Italy

² Institute of Services Science, University of Geneva, Carouge, Switzerland

Abstract. Incentive strategies are used in collaborative user-centric networks, the functioning of which depends on the willingness of users to cooperate. Classical mechanisms stimulating cooperation are based on trust, which allows to set up a reputation infrastructure quantifying the subjective reliance on the expected behavior of users, and on virtual currency, which allows to monetize the effect of prosocial behaviors. In this paper, we emphasize that a successful combination of social and economic strategies should take into account the privacy of users. To this aim, we discuss the theoretical and practical issues of two alternative tradeoff models that, depending on the way in which privacy is disclosed, reveal the relation existing among trust, privacy, and cost.

1 Introduction

A growing trend towards autonomic user-centric architectures is giving rise to community-scale initiatives with the purpose of sharing services, among which personal hotspot and peer-to-peer are two representative examples. Members of these communities may share access to the Internet as well as network resources and user-generated contents and applications. User-centricity is reshaping the Internet value chain, and its success depends strongly on the attitude to cooperate of each actor involved [2]. Intrinsic motivations to be cooperative, such as sense of community and synergy, do not suffice to contrast typical obstacles like selfishness and, even worse, cheating. They must be integrated with extrinsic incentives, especially for communities in which users behave as *prosumers*, i.e., they combine the roles of service producers and consumers. Extrinsic motivations can be analyzed from social and economic perspectives.

From a social viewpoint, it is well-recognized that computational notions of trust support the estimation of user's trustworthiness as perceived by the community [8]. On one hand, the reputation resulting from user's behavior shall be viewed as an enabling factor for accessing the best services at the most favorable conditions. On the other hand, reputation is related to identity, thus contrasting the idea of privacy, which represents another social value that may keep the user from taking part in some kind of interaction. However, the lower the attitude to expose sensitive information is, the higher the probability of being untrusted.

Trading privacy for trust is thus a way for balancing the subjective value of what is revealed in exchange of what is obtained.

From an economic viewpoint, the capability of monetizing the effects of cooperative and prosocial behaviors is fundamental whenever trust does not represent a sufficient incentive. This is the case, e.g., of wireless communities, the highly dynamic nature of which hinders the establishment of a stable reputation infrastructure, which suffers the frequent renewal of the community members. Moreover, reputation may not provide guarantees of reciprocity, according to which the attitude to cooperation is strengthened by the perspective of future mutual interactions. While reciprocity is not perceived as an incentive in classical models based, e.g., on client-server architectures, it represents a pillar of cooperation in user-centric environments, in which community members behave as prosumers. Monetization provides a framework where virtual credits play the role of commodity money used to purchase services [7], thus replacing (or complementing) the role of trust during negotiation. While using virtual currency in place of trust can be beneficial, the maximum benefits deriving from these orthogonal incentive mechanisms are obtained when they are combined in a mixed strategy [1,17]. In other words, among the favorable conditions that can be obtained during negotiation by a trusted user (in terms, e.g., of amount of resources and, more in general, quality of service), it is quite natural to include the service cost. The relation is dual, as an effect of the marketplace is that the cost applied by a user providing a service may have an impact on her/his reputation as perceived by the buyers of the service.

In this paper, we investigate the tradeoff existing among the three dimensions that characterize the incentive strategies resulting from the discussion above, namely trust, privacy, and cost. To this aim, we discuss the theoretical benefits and the implementation issues of two models that differ for the way in which privacy is managed and traded with respect to trust and cost.

2 Incremental vs. Independent Release of Privacy

According to an established view of privacy management, sensitive information is disclosed incrementally. Whenever a user requires access to a service, a certain portion of user's identity is exposed depending on the amount of her/his personal information disclosed, while the related reputation is employed to negotiate transaction and cost to pay. The basic assumption behind the incremental model is that the amount of privacy released is irrevocable, meaning that once different pieces of sensitive information are linked and exposed by a user, it is not possible to break anymore such a link. In fact, the case in which user's identity is revealed all at once represents a limiting scenario of this model.

As an example, suppose that Alice uses a pseudonym to ask for a service without revealing to be Alice. To this aim, she discloses some information, e.g., a piece of evidence associated to the pseudonym, which is trusted enough to obtain the service at a certain cost. For privacy reasons, in order to negotiate other services with different users, Alice may use several pseudonyms, each one

characterized by its own reputation level. At a certain point, if Alice requires a service with high trust threshold, she may need to expose the link between two of her pseudonyms in order to benefit of the related reputation combination, thus revealing that the same identity is behind them. Hence, linking more evidence is a way to grant a request at the cost of increased privacy loss. However, from now on, such a link is irrevocable, in the sense that Alice cannot spend one of the two pseudonyms without spending also the other at the same time.

As opposite to this policy, we envision a model with a higher degree of freedom in the management of privacy, in which the privacy disclosure is independent of the information released in previous interactions, by breaking the irrevocability that characterizes the incremental model. Such a flexibility would allow the user to choose which (and how much) information to disclose depending solely on the service trust threshold and on the cost she/he is willing to pay for the service, without any constraint deriving from the privacy released in the past.

With respect to our example, this means that after having linked two different pseudonyms, somehow Alice would be able to break such a link in future interactions, in which she may use only one of the two separately from the other. The need for this capability could be motivated, e.g., by the fact that only one of the two pseudonyms is sufficient to negotiate certain services, or because Alice may prefer to maintain separation of identity by using two unrelated pseudonyms in two different social environments.

2.1 Theoretical Aspects

Before discussing the design of the independent model, which is the novel contribution of this paper, it is worth analyzing from the theoretical standpoint the potential benefits of this model with respect to the incremental one, in order to motivate its implementation. This is done by verifying whether the achieved flexibility of privacy influences positively the tradeoff with trust and cost.

The efficacy of mixed cooperation strategies has been demonstrated through formal methods, like game theory [12,13], and model checking [1,11]. For our purposes, we have conducted a preliminary verification based on the analysis of a real-world cooperation system for user-centric networks [3], which has been modeled and analyzed through the model checker PRISM [4]. Such a system entails a cooperation process balancing trustworthiness with service cost and is based on irrevocable *all-at-once* identity disclosure. Hence, our formal model extends this system by including the capabilities of the two different models of privacy release, in order to evaluate for each of them how trading privacy for trust influences access to services and related costs.

To summarize the comparison results, we observe that the major freedom degree of the independent model allows the user to obtain access to the same services by saving up to 30% of private information disclosure. If the objective is not only trading privacy for trust, but also cost optimization, it is worth comparing cost functions that depend on trust in different ways [3,17]. Preliminary results show that in many cases the independent model ensures lower costs when

the same average level of privacy release is considered and that such a model never induces higher costs than the incremental model.

2.2 Implementation Issues

As far as the design of the two models of privacy management is concerned, we notice that the incremental model can employ known techniques, while the independent one represents a novel approach requiring non-standard mechanisms.

The pseudonyms model by Seigneur and Jensen [14] obeys the principle of incremental release of privacy, as it is based on irrevocable linkability of pieces of evidence. In order to implement a mechanism balancing trust with privacy, they allow users to freely create pseudonyms identified by the *crypto-id*, i.e., the hash of the public key of a locally generated asymmetric cryptography key pair. Then, depending on the context, one or another pseudonym could be used to carry out actions logged as events signed with the private key of the pseudonym. If needed, one or several pseudonyms could also be linked together in order to increase the number of known actions and potentially increase the trust in the linked entity assuming that all these actions had a positive outcome.

In the following we present a technique for the independent model. As for the pseudonyms model, we use the notion of virtual identity generated by means of the *crypto-id*. As a crucial assumption, reputation and each trust association are mapped to pieces of the *crypto-id* rather than to the *crypto-id* as a whole.

Whenever issuing a service request, the sender chooses a bitmask B that is applied to her/his *crypto-id* through the bitwise AND operator in order to extract n bits of the *crypto-id* to be revealed to the request receiver. We use C_B (and the term *chunk*) to denote the result of this operation. Hence, a chunk represents a set of n bits of the *crypto-id*, of which we know value and position. The calculated chunk is transmitted to the receiver and represents a portion of the identity of the sender. Notice that the negotiated transaction is not associated to the sender directly, but is related to the chunk extracted from sender's *crypto-id*, which could be shared by several different users.

Example 1. The bitmask 01110000 identifies the same chunk for the *crypto-ids*:

$$K_1: 10010100 \quad K_2: 00010010$$

while the bitmask 00001110 does not. Now, assume that the user with *crypto-id* K_1 uses bitmask 00010010 for a certain interaction. If in a future interaction with the same receiver she/he employs the bitmask 00010000, thus revealing less information, then the receiver cannot link the two interactions to the same identity, as they could be related to different *crypto-ids*. $\square$

Since a user can spend different chunks in different transactions and can also combine chunks previously used in order to exploit a combination of their reputations, in a limiting scenario we may envision a reputation for every bit of the *crypto-id*. Hence, the reputation associated with a chunk is given by the contribution of the reputations of each bit forming the chunk. In the following, without loss of generality, we adopt such an assumption and we present a reputation system in the centralized setting.

User crypto-ids are stored in a non-public repository managed by a trusted, central authority (CA). When a user issues a service request linked to a chunk, she/he sends to the receiver an encryption (through receiver's public key) containing the chunk, a cryptographic proof allowing the CA to validate the request while preserving sender's anonymity, and transaction specific information avoiding replay attacks. Then, the receiver forwards the chunk and the proof to the CA, which performs the validity check and transmits to the receiver the reputation of the chunk. Afterwards, the objective of the CA is to update the reputation of the crypto-ids from which the chunk could be generated on the basis of the feedback reported by the receiver at the end of the transaction.

Ideally, the overall reputation associated with a crypto-id shall result from a combination of the reputations cumulated by every bit of such a crypto-id spent to expose a chunk in some interaction. However, as previously shown, a chunk is potentially shared by several different crypto-ids. Therefore, when the receiver transmits a chunk and the evaluation resulting from the transaction, the CA could not be able to infer from which crypto-id the chunk is actually originated. For this reason, we assume that the CA distributes the result of receiver's evaluation among the bits of the chunk for every crypto-id matching the chunk. Let us explain the feedback mechanism through an example.

Example 2. Let us assume that the central repository includes crypto-ids K_1 and K_2 of the previous example and the following crypto-ids:

$$K_3: 01110111 \quad K_4: 11011011$$

The four crypto-ids are associated to users U_1 , U_2 , U_3 , and U_4 , respectively. Even if in principle any reputation metric could be applied, to simplify calculations we report receiver's evaluations as unitary reputation variations and we assume that initially the reputation of each bit of every crypto-id is 0.

Firstly, U_1 negotiates a service by using the chunk identified by the bitmask 01110000, for which the receiver provides a positive feedback at the end of the transaction. Since the related chunk is shared by U_1 and U_2 , reputations are changed by the CA as follows:

$$rv(K_1): [01110000] \quad rv(K_2): [01110000]$$

where $rv(K)$ denotes the vector of the reputations of the bits forming the crypto-id K . Secondly, U_3 uses bitmask 00011100 and, again, feedback is positive. The related chunk is shared by U_3 and U_1 , for which reputations change as follows:

$$rv(K_1): [01121100] \quad rv(K_3): [00011100]$$

Thirdly, U_1 requires another service for which U_1 exhibits a chunk formed by two bits with high reputation, e.g., through the bitmask 01010000. If receiver's evaluation is positive, the CA changes the reputations as follows:

$$rv(K_1): [02131100] \quad rv(K_2): [02120000]$$

Notice that users are stimulated to use chunks of bits with high reputation in transactions in which they behave honestly. As a consequence, all the users who share these chunks and contributed to their high reputation benefit from this virtuous circle. Finally, consider a transaction with negative feedback conducted by U_4 by using bitmask 00000011. Hence:

$$rv(K_3): [000111-1-1] \quad rv(K_4): [000000-1-1]$$

□

By following the considerations concerning the uncertainty of the chunk origin, in case of a request accompanied by chunk C_B , the calculation of C_B 's reputation results from a combination (e.g., through the arithmetic mean) of the reputations of such a chunk within every crypto-id matching C_B .

With abuse of notation, we write $C_B \leq K$ to express that chunk C_B can be extracted from crypto-id K . Let $rep(C_B, K)$ denote the reputation of chunk C_B of the crypto-id K , which is calculated by combining the reputation of each bit of K contributing to C_B . Moreover, let $|C_B|$ denote the number of crypto-ids matching C_B . Then, the reputation of C_B is as follows:

$$rep(C_B) = \frac{1}{|C_B|} \cdot \sum_{C_B \leq K} rep(C_B, K)$$

Example 3. With reference to the previous example, let us consider the situation just before the third transaction, in which R_1 spends the chunk C_B identified by the bitmask $B = 01010000$. Such a chunk is shared by R_1 and R_2 . The reputations of the involved bits are 1, 2 for R_1 and 1, 1 for R_2 . Denoted by f the function used to combine the reputation of each bit of the chunk, we obtain $rep(C_B) = \frac{1}{2} \cdot (f(1, 2) + f(1, 1))$ (e.g., if f is summation then $rep(C_B) = 2.5$). $\square$

As a side effect of chunk sharing, the reputation of a chunk is the result of the behavior of all the users with crypto-ids consistent with such a chunk. In other words, the crypto-ids matching the same chunk actually benefit from the reputation (or pay the mistrust) associated to such a chunk. This aspect is crucial for the requirements of the independent model of privacy release and can be viewed as an incentive to take prosocial and honest decisions, as a high number of trustworthy chunks contribute to increase the probability of obtaining services at a reasonable cost by preserving the desired level of privacy.

Another important aspect of this model is the choice of the chunk size. If privacy is privileged and, therefore, chunks of small size are chosen, then the probability that their reputation values are influenced by a high number of community members increases, thus leading to a worse approximation of the actual reputation of the user exposing a small chunk. On the other hand, if accuracy of reputation is privileged, then the user is motivated to use chunks with a high number of bits, thus sacrificing more privacy. Therefore, a tradeoff exists between the amount of sensitive information the user is willing to invest and the accuracy of her/his reputation estimated by the request receiver in order to negotiate the transaction and the related parameters, including service cost.

The proposed implementation differs from the pseudonyms model by Seigneur and Jensen [14] since a user's pseudonym, here represented by a chunk, may also be controlled by another user. Then, actions may be carried out by several different users, without every user being able to know which other user has also control over the same chunk C_B . As shown above, it may be beneficial for the chunk if all users controlling it are trustworthy, but if at least one of them is very untrustworthy and carries out at least one illegal action linked to C_B , then the chunk becomes untrustworthy and useless for all other users. In addition, if

C_B is untrustworthy then it may also impact the trustworthiness of any chunk $C_{B'}$ refining C_B , i.e., such that $C_B \leq C_{B'}$, as one may wonder whether the user exposing $C_{B'}$ is (or is not) the same untrustworthy user who used C_B in a malicious way. This negative side effect, which can be serious in case C_B is a small chunk, can be mitigated in an implicit way by using mixed cooperation strategies based on trust and cost, or explicitly by reducing the influence upon trust of actions linked to small chunks. This can be done by weighting both the reputation calculation and the feedback evaluation by a discounting factor inversely proportional to the size of the chunk used, in order to reflect that the amount of sensitive information that is exposed in a transaction is proportional to sender's trustworthiness. An alternative, effective but severe solution consists of resorting to a CA capable of revoking blindness in case of suspicious behaviors by some chunk, in order to isolate dishonest users and repair the reputation of the chunk involved.

3 Related and Ongoing Work

The main objective of this work has been showing that trust and cost can be effectively combined while also considering privacy as a third dimension. To the best of our knowledge, such an analysis has never been conducted by joining all these aspects in the same framework. In the literature, Automated Trust Negotiation is known to have not fully resolved privacy issues [16]. Wagealla et al. [15] use trustworthiness of an information receiver to make the decision on whether private information should be disclosed or not, which is another way to envisage the relation between trust and privacy. However, it may be difficult to evaluate trustworthiness in first place without enough evidence linked with the receiving entity. The work on modelling unlinkability [10] and pseudonymity [6,9] is valuable towards founding privacy/trust trade. Moreover, the Sybil attack [5], which challenges the use of recommendations, is also worth keeping in mind when providing means to create virtual identities at will without centralized authority. Finally, the trust-privacy tradeoff can be optimized in data-centric ad-hoc networks by using incentive mechanisms [13].

The theoretical analysis conducted on a real-world case study motivates the implementation of the independent model, the applicability of which has been shown through a mechanism based on the splitting of crypto-ids into chunks and on a centralized reputation system. Since chunk sharing is the main principle behind this mechanism, we point out that, as an alternative approach, the model by Seigneur and Jensen [14] may also be turned into a scheme allowing for the sharing of a pseudonym among n users. The private key associated to a pseudonym generated by user n would be sent to the other $n - 1$ users encrypted with their public keys. Anyway, this approach raises issues related to the key exchange protocol and to the choice of the n users sharing a pseudonym.

To integrate our reputation system, we plan to design a trust model for both centralized and distributed systems. The idea is to equip every user with a structure collecting information on the set of chunks associated to completed

transactions. Such a structure is then used to manage trust towards every chunk under the assumption that the user is not aware of the set of crypto-ids from which chunks are originated.

References

1. Aldini, A.: Formal Approach to Design and Automatic Verification of Cooperation-Based Networks. *Journal On Advances in Internet Technology* 6, 42–56 (2013)
2. Aldini, A., Bogliolo, A. (eds.): *User-Centric Networking – Future Perspectives. Lecture Notes in Social Networks*. Springer (2014)
3. Bogliolo, A., Polidori, P., Aldini, A., Moreira, W., Mendes, P., Yildiz, M., Ballester, C., Seigneur, J.-M.: Virtual Currency and Reputation-Based Cooperation Incentives in User-Centric Networks. In: 8th Int. Wireless Communications and Mobile Computing Conf (IWCMC 2012), pp. 895–900. IEEE (2012)
4. Chen, T., Forejt, V., Kwiatkowska, M., Parker, D., Simaitis, A.: PRISM-games: A Model Checker for Stochastic Multi-Player Games. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013. LNCS, vol. 7795, pp. 185–191. Springer, Heidelberg (2013)
5. Douceur, J.R.: The Sybil Attack. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 251–260. Springer, Heidelberg (2002)
6. Goldberg, I.: *A Pseudonymous Communications Infrastructure for the Internet*, PhD Thesis, University of California at Berkeley (2000)
7. Greengard, S.: Social Games, Virtual Goods. *Communications of the ACM* 54(4), 19–22 (2011)
8. Jøsang, A.: Trust and Reputation Systems. In: Aldini, A., Gorrieri, R. (eds.) FOSAD 2007. LNCS, vol. 4677, pp. 209–245. Springer, Heidelberg (2007)
9. Kobsa, A., Schreck, J.: Privacy through Pseudonymity in User-Adaptive Systems. *ACM Transactions on Internet Technology* 3(2), 149–183 (2003)
10. Steinbrecher, S., Köpsell, S.: Modelling unlinkability. In: Dingledine, R. (ed.) PET 2003. LNCS, vol. 2760, pp. 32–47. Springer, Heidelberg (2003)
11. Kwiatkowska, M., Parker, D., Simaitis, A.: Strategic Analysis of Trust Models for User-Centric Networks. In: Int. Workshop on Strategic Reasoning (SR 2013). EPTCS, vol. 112, pp. 53–60 (2013)
12. Li, Z., Shen, H.: Game-Theoretic Analysis of Cooperation Incentives Strategies in Mobile Ad Hoc Networks. *IEEE Transactions on Mobile Computing* 11(8), 1287–1303 (2012)
13. Raya, M., Shokri, R., Hubaux, J.-P.: On the Tradeoff between Trust and Privacy in Wireless Ad Hoc Networks. In: 3rd ACM Conf. on Wireless Network Security (WiSec 2010), pp. 75–80 (2010)
14. Seigneur, J.-M., Jensen, C.D.: Trading Privacy for Trust. In: Jensen, C., Poslad, S., Dimitrakos, T. (eds.) iTrust 2004. LNCS, vol. 2995, pp. 93–107. Springer, Heidelberg (2004)
15. Wagealla, W., Carbone, M., English, C., Terzis, S., Nixon, P.: A Formal Model of Trust Lifecycle Management. In: Workshop on Formal Aspects of Security and Trust (FAST 2003) (2003)
16. Yu, T., Winslett, M.: A Unified Scheme for Resource Protection in Automated Trust Negotiation. In: IEEE Symp. on Security and Privacy, pp. 110–122 (2003)
17. Zhang, Y., Lin, L., Huai, J.: Balancing Trust and Incentive in Peer-to-Peer Collaborative System. *Journal of Network Security* 5, 73–81 (2007)