The final version of this paper appears in the proceedings of
the 2nd International Conference on Trust Management
published by Springer as
Lecture Notes in Computer Science.
(http://www.springerlink.com/openurl.asp?genre=article&issn=0302-9743&volume=2995&spage=93)

Trading Privacy for Trust

Jean-Marc Seigneur
Trinity College Dublin
Jean-Marc.Seigneur@cs.tcd.ie

Christian Damsgaard Jensen
Technical University of Denmark
Christian.Jensen@imm.dtu.dk

Abstract. Both privacy and trust relate to knowledge about an entity. However, there is an inherent conflict between trust and privacy: the more knowledge a first entity knows about a second entity, the more accurate should be the trustworthiness assessment; the more knowledge is known about this second entity, the less privacy is left to this entity. This conflict needs to be addressed because both trust and privacy are essential elements for a smart working world. The solution should allow the benefit of adjunct trust when entities interact without too much privacy loss. We propose to achieve the right trade-off between trust and privacy by ensuring minimal trade of privacy for the required trust. We demonstrate how transactions made under different pseudonyms can be linked and careful disclosure of such links fulfils this right trade-off.

1 Introduction

Privacy can be seen as a fundamental human right "to be left alone" [2] or a basic need (according to Maslow's hierarchy of needs [12]) for a private sphere protected against others. Regardless of the definition, different mechanisms have been proposed to protect the privacy of people in the online world. The most common mechanisms are either legislative or technological, depending on whether privacy is seen a right which should be protected by law or a need which should be supported by the devices that are used to access the online world. In this paper we focus on the technological aspects of privacy protection, especially techniques to control the dissemination of personal information.

Information becomes personal when it can be linked back to an individual or when it, in some way, allows two individuals to be linked together. This means that control of the dissemination of personal information can be exercised through preventing, or at least limiting, linkability of information to individuals. This is illustrated in Figure 1, where a user Alice performs some transactions with another user Bob (neither Alice nor Bob needs to be actual users, but could be clients, servers or part of the computing infrastructure).

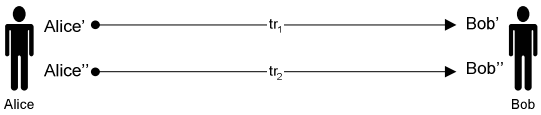

Figure 1: Linkability of transactions

In Figure 1, Alice performs two transactions tr_1 and tr_2 with Bob. In order to protect the privacy of Alice[1], it is important that Bob, or anyone who eavesdrops on their communication, is unable to link either transaction tr_1 or tr_2 directly to Alice's real-world identity. However, it is equally important to prevent Bob from linking the two transactions to each other, since this would allow him to compile a comprehensive profile of the other party, which could eventually identify Alice. Moreover, the violation of Alice's privacy would be increased dramatically if any future transaction tr_x can be linked to Alice, since this would allow Bob to link the full profile to Alice and not just tr_x. However, trust is based on knowledge about the other party [7], which directly contradicts the prevention of linkability of information to users, so perfect privacy protection, i.e., preventing actions to be linked to users, prevents the formation, evolution and exploitation of trust in the online world.

In the human world, trust exists between two interacting entities and is very useful when there is uncertainty in result of the interaction. The requested entity uses the level of trust in the requesting entity as a mean to cope with uncertainty, to engage in an action in spite of the risk of a harmful outcome. Trust can be seen as a complex predictor of the entity's future behaviour based on past evidence. In the literature, divergent trust definitions are proposed but it is argued that they can fit together [13].

Interactions with uncertain result between entities also happen in the online world. So, it would be useful to rely on trust in the online world as well. The goal of a computational trust/risk-based security framework (TSF) is to provide trust in the online world. Researchers are working both theoretically and practically towards the latter goal. Others have shown how trust can be formalized as a computational concept [7, 11]. The aim of the SECURE project [1, 14] is an advanced TSF formally grounded and usable. The basic components of a TSF (depicted in Figure 2) should expose a decision-making component that is called when a requested entity has to decide what action should be taken due to a request made by another entity, the requesting entity.

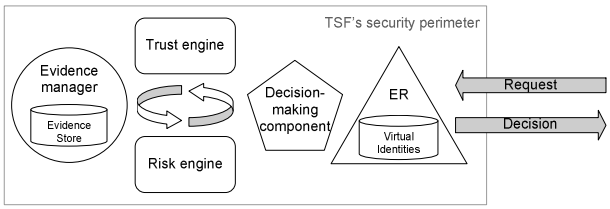

Figure 2: High-level view of a TSF

In order to take this decision, two sub-components are used:

- a trust engine that can dynamically assess the trustworthiness of the requesting entity based on pieces of evidence (e.g., observation or recommendation [19])
- a risk engine that can dynamically evaluate the risk involved in the interaction and choose the action that would maintain the appropriate cost/benefit

[1] The rights/needs to privacy of Alice and Bob are symmetrical, so it may be equally important to prevent Alice from knowing that the two transactions were performed with the same entity.

In the background, another component is in charge of gathering evidence (e.g., recommendations, comparisons between expected outcomes of the chosen actions and real outcomes…) This evidence is used to update risk and trust information. Thus, trust and risk follow a managed life-cycle. In the remainder of the paper, we use TSF in its broad sense: any TSF can be used (even though the TSF being developed in the SECURE project is an example of an advanced TSF).

Recalling the process of trust formation makes apparent the fact that privacy is at stake in trust-based systems. In order to be able to trust another entity, the first step is to establish the level of trust in that entity[2], which is the result of an analysis of the existing knowledge and evidence. Thus, trust relies on profiling, where more information is better, because it allows the likely behaviour of the other entity to be more accurately predicted. Any link with the real-world identity of the user changes this information into sensitive personally identifiable information (PII). From a privacy point of view, a first technological line of defence may be to use virtual identities – pseudonyms (mapping to principals in SECURE). The ordinary definition of a pseudonym is "a fictitious name used when the person performs a particular social role"[3]. Ian Goldberg underlined that any transaction engaged by a person reveals meta-content, especially information about the identity of the person. He defined "the nymity of a transaction to be the amount of information about the identity of the participants that is revealed" and gave a continuum, called the "Nymity Slider", with different levels of nymity: verynimity (e.g., government id), persistent pseudonymity (e.g., pen names), linkable anonymity (e.g., prepaid phone cards), unlinkable anonymity (e.g., anonymous remailers). He also pointed out that it makes sense to associate reputation with persistent pseudonyms. In a TSF, the minimum requirement is a local reference for the formation of trust, which is in turn managed by other components in the TSF. According to the privacy protection principle of "collection limitation" [10], data collection should be strictly restricted to mandatory required data for the purpose of the collection.

Our requirement is to establish the trustworthiness of entities and not their real-world identity. This is why pseudonymity, the level of indirection between trust and the real-world entity, is necessary. Transaction pseudonyms [8] (i.e., a pseudonym used for only one transaction) and anonymity cannot be effectively used because they do not allow linkability between transactions as required when building trust. In the following, we consider a model where linkability of different transactions with a specific pseudonym is achieved by using the APER [15] Entity Recognition (ER) scheme for transactions between the two principals. There are two roles distinguished in APER, the recogniser and the claimant (though any party can take on any role). The approach is for the claimant to send claims, i.e., digitally signed messages, and for the recogniser to be able to recognise the claimant on the basis of correctly signed claims. A principal, i.e., a pseudonym, is an APER claimant who is recognised using a digital signature and who sends APER claims. When an entity makes a request, which requires a trusting decision from another entity, the requesting entity sends an

[2] In this paper, we use the following terms as synonyms: *level of trust* and *trustworthiness*. In a TSF, they are represented as a trust value. This is different than *trust*, which is the concept.

[3] Definition from WordNet Dictionary: http://www.hyperdictionary.com/search.aspx?define=pseudonym

APER claim that tells the requested entity which pseudonym is claimed. So, transactions are linked through asymmetric key digital signature validation (which provides a level of confidence in recognition called APERLevel1) using the same key. The requested entity can refer to a specific pseudonym (e.g., in order to get recommendations about a specific pseudonym) by specifying the Public Key (Pub) claimed by the requesting pseudonym.

The next section describes a scenario where it makes sense to trade privacy for trust. A model for privacy/trust trade is given in Section 3. This model is applied at the level of virtual identities in Section 4. Section 5 surveys related work and we draw conclusions.

2 Scenario

As an example, the following figure depicts the scenario where Alice plans to spend her holidays in SunnyVillage. Normally Alice works and lives in RainyTown. She will take the plane and relax for two weeks in this village where she has never been but that some of her friends recommended.

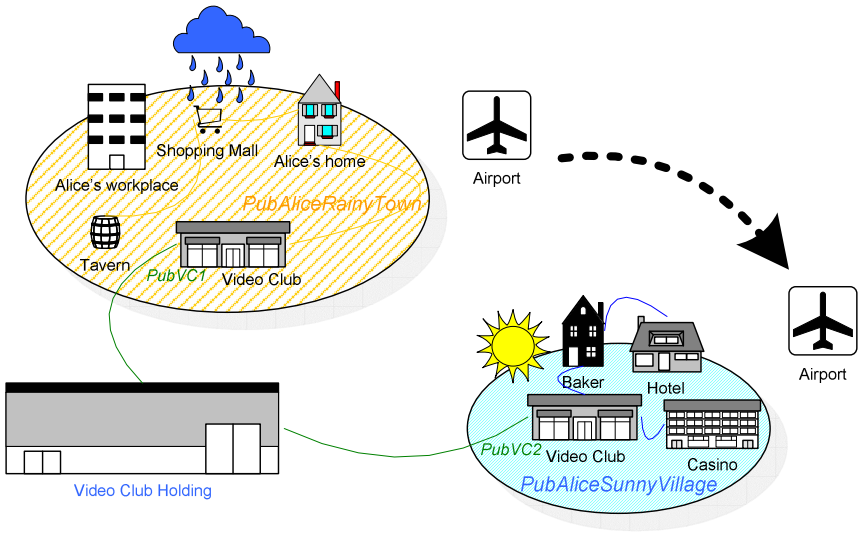

Figure 3: Alice's smart world

She will have to pay to enjoy some of her leisure activities, which could be enhanced if collaboration with other local entities is allowed. We assume that Alice uses an e-purse. So, an e-purse is associated with Public Key (Pub) / Private Key (Pri) pairs: a Pub becoming a pseudonym for Alice. An e-purse has also an embedded TSF, which takes care of trust decision-making and management. Similarly, a vendor's cashier-machine can be recognised with a Pub and run a TSF. For example, exchange of Alice's trustworthiness in being a good payer in the neighbourhood would let her pay without being asked real-world credentials (e.g., a passport); credit may also become viable. Vendors would also benefit from trust calculation adjunct. The video shop of

SunnyVillage, having to deal with passing customers, would be reassured to take a lower risk if payment with electronic coins is combined with the level of trust in the customer. Nevertheless, Alice also wishes to be left alone and have different social profiles in different places. Alice has indeed two pseudonyms automatically selected according to location: one in RainyTown (PubAliceRainyTown) and one in SunnyVillage (PubAliceSunnyVillage). This offers better protection for her privacy than having one pseudonym. Even though the video club holding spans both domains, SunnyVillage's video club cannot obviously link PubAliceRainyTown and PubAliceSunnyVillage by comparing keys known by RainyTown's video club. The latter would not be true with a unique Pub for Alice's e-purse.

However, trust, as with privacy, is dynamic and evolving interaction after interaction. Privacy is a constant interaction where information flows between parties [5, 17]. Privacy expectations vary [5, 17] and depend on context [8]. We have demonstrated a prototype where privacy disclosure policies can be based on context [17], especially location. Depending on what people can get based on their trustworthiness, they may be willing to disclose more of their private data in order to increase trust. There is a need for contextual privacy/trust trade. Let us assume that the trustworthiness of people for being good payers is managed by the TSF of the vendor's cashier-machine. Recalling the scenario in Figure 3, if Alice arrives in SunnyVillage's video club for the first time, her e-purse will exhibit PubAliceSunnyVillage when she wants to pay for the large video display that she wants to rent. Since no direct observation, i.e., a previous experience with PubAliceSunnyVillage, is available, PubVC2 (the SunnyVillage video club cashier's Pub) will ask for recommendations from its neighbors (e.g., PubBaker). However, Alice's trust obtained through recommendations is not enough to commit the renting transaction. Alice really wants the display, so she is now disposed to give up some of her privacy in order to exhibit enough trust. In fact, SunnyVillage's video club is held by a holding of video clubs, which has a video club in RainyTown. The following example of contextual privacy/trust trade is started. The list of Pubs owned by the holding is sent to Alice's e-purse, which finds that PubVC1 of RainyTown's video club is a known entity. Alice has noticed that she could link PubAliceRainyTown and PubAliceSunnyVillage in order to reach the necessary level of trust. Although Alice now knows that what she has done in RainyTown is potentially exposed to both areas, i.e., RainyTown and SunnyVillage, she agrees to present herself as the owner of both keys (i.e., pseudonyms).

3 Privacy/Trust Trade Model

We start by an informal summary of the model. When true knowledge[4] about an entity increases:

- The evaluation of its trustworthiness is more accurate and if this entity is indeed truly trustworthy, its trustworthiness increases[5].

[4] By true knowledge, we mean knowledge which cannot be refuted (i.e., it cannot be a lie, noise information or revised).

- Its privacy decreases and it is almost a one-way function[6] because privacy recovery is hard to achieve [16].

Knowledge is composed of evidence. A piece of evidence ev may be any statement about some entity(ies), especially: a transaction tr, an observation[7] obs (i.e., evaluated outcome of a transaction [6]), a recommendation rec (i.e., locally discounted[8] observation of a recommending external entity)… The nymity of evidence is the amount of information about the identity of the entity that is revealed. The trustworthiness assessment impact, called tai of evidence, is the amount of information that can be used for assessing the trustworthiness of the entity, which is represented as a trust value.

There are different levels of nymity. So we assume that there is a partial order between nymity levels, called Privacy Asset Order (PAO). The Nymity Slider is one example of such ordering. We present another example of PAO below:

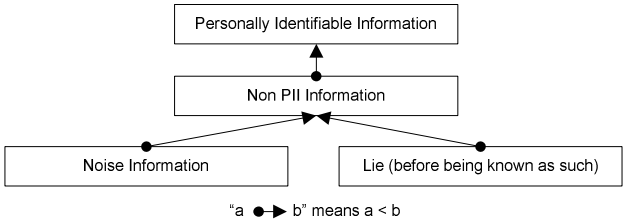

Figure 4: Privacy Asset Order example

Similarly, evidence may be more or less useful for trustworthiness assessment. So we assume that there is a partial order between tai levels, called Trustworthiness Assessment Impact Order (TAIO). An example of TAIO is:

---

[5] We do not mean that the trustworthiness increases in all possible trust dimensions (but at least it increases in the dimension where the knowledge is useful/relevant, e.g., propensity to be a good payer).

[6] On Goldberg's Nymity Slider, it is "easy to change the transaction to have a higher position on the slider" and "extremely difficult to move a transaction down the slider (towards unlinkable anonymity)" [4].

[7] It is sometime difficult to find out when the observation should be made because it is not clear whether the action is finished or not. It may be solved by having a kind of dynamic observation, i.e., a piece of evidence which varies through time as well.

[8] By discounted, we mean that the trustworthiness of the recommender is taken into account. The final value, which is used locally, may be different than the recommended one. For example, a recommender with trust value of 0.6 on a [0,1] scale giving a recommendation of 0.8 provides the discounted trust value: 0.6*0.8.

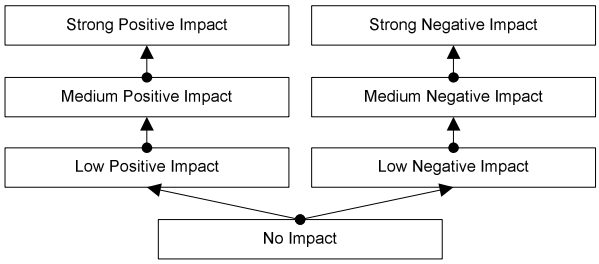

Figure 5: Trustworthiness Assessment Impact Order example

A piece of evidence of PII nymity is more likely to have a strong positive impact tai, especially when it is assumed that the real-world identity can be sued. However, one non-PII evidence may have low positive impact and another one strong positive impact.

We provide a mechanism that can link n pieces of evidence ev_i for i=1,…,n and represented by:

$$link(ev_1, ev_2, ..., ev_n)$$

The result of link is a new piece of evidence with a new tai level as well as a new nymity level. Sometimes, linking of evidence is implicit (i.e., the requesting entity cannot keep secret that two pieces of evidence are linked) and it is redundant to make it explicit (i.e., the requesting entity discloses to other entities that two pieces of evidence are indeed linked). For example, if two events ev_2 and ev_3 are implicitly linked, then explicitly linking ev_1 and ev_2 is equivalent to explicitly linking ev_1, ev_2 and ev_3: $link(ev_1,ev_2) = link(ev_1,ev_2,ev_3)$.

It is needed to recognise entities and it is useful to know what piece of evidence is linked to a specific entity for the recognition of entities. An APER virtual identity vi (i.e., pseudonym) is recognised by a public key Pub, which can be seen as evidence. However, presenting a public-key is meaningless until you link it to the current (or a previous) transaction by signing something with it, i.e., providing linkability. In our case, after the first transaction, the requested entity links the transaction with the pseudonym Pub: $link(tr_1,Pub)$. Then, after the second transaction, the requested entity does: $link(tr_1, Pub, tr_2)$ and so on. Thus, the pseudonym links a set of pieces of evidence together. If each transaction is non-PII/low positive impact and Pub considered as non-PII/no impact, the resulting evidence is: two low positive impacts from a tai point of view and three non-PII from a nymity point of view.

If not enough evidence is available under the chosen pseudonym, evidence not linked to this pseudonym may improve trustworthiness and allow the requesting entity to be granted the request. The entity may be willing to disclose further evidence to the requested entity in spite of potential increased privacy loss. So, a protocol for disclosing to the requested entity that some evidence can be linked is needed. We present such a protocol, called the privacy/trust trade process (depicted in Figure 6). In this process, the requested entity makes the decision that not enough evidence is available for granting and this fact should be disclosed to the requesting entity. So, after step 2, the requesting entity knows the tai of evidence that should be obtained.

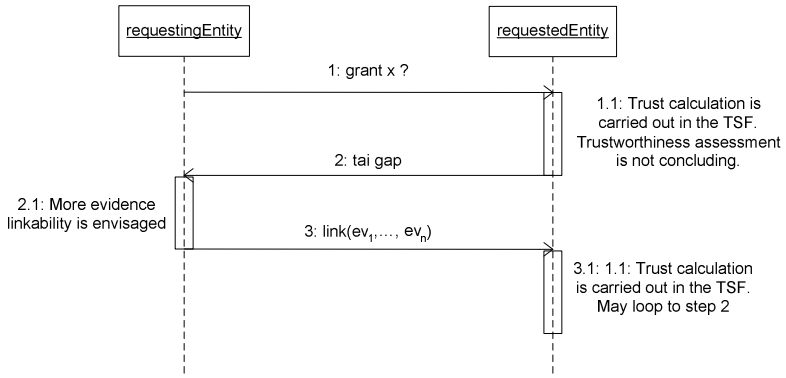

Figure 6: Privacy/trust trade sequence diagram

In step 2.1, different potential evidence can be envisaged to be linked by the requesting entity. The choice of evidence should be based on the following principle:

The "Minimal Linkability" principle: No more evidence than needed should be linked.

The latter principle is a variant of the "Need-To-Know" principle. One of the reasons is that more trust implies more knowledge given out, thus less chance for privacy. Some thresholds should be set concerning the acceptable evidence that should be disclosed in step 3. Without such thresholds, an attacker may ask to retrieve all evidence (i.e., knowledge), which is what we want to prevent by using pseudonyms. If the user must confirm that some evidence can be linked, more care has to be taken into account. It is known that users can easily agree to sell privacy in stressed circumstances without thinking of the consequences [18], which are often irrevocable since privacy recovery is hard [16]. Alice, in order to get quick access to the large video display, may regret to present her full profile to the video club due to this small benefit compared to life-long spam messages. One way to prevent such abuse may be the existence of a broker where reasonable trades are listed (this also reduces interoperability issues). In practice, it may require an exchange of messages with trusted third parties to decide whether the trade is fair (within the current market price) or not. We propose to introduce another partial order to cope with such abusive trade attack. The utility of a transaction is represented on a utility partial order (UO). An example UO may be:

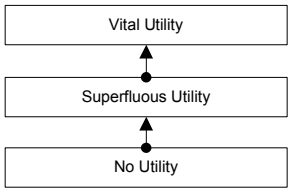

Figure 7: Utility Order example

During a trade process, tai, nymity and utility must be balanced. Alice under the pseudonym Pub requests Bob to grant the transaction tr_x of utility u from Alice's point of view. In step 1.1, if Pub had done two previous transactions tr_1 and tr_2 with Bob, Bob's TSF checks if the trustworthiness given by this previous evidence is

enough to grant tr_x. In this case, the trustworthiness assessment is not concluding, so the TSF computes the z tai of evidence missing, called tai gap. Alice's TSF is noticed that z tai of evidence is missing. In step 2.1, Alice's TSF does the following 2-step algorithm, called link selection engagement (liseng) algorithm:

1. *Search link of evidence expected to fill the tai gap but minimizing nymity*: As an example, we assume that the TSF cannot guarantee that all recommenders of Pub can exhaustively be found and queried in a timely manner. All transactions directly done between Alice and Bob should have been taken into account by Bob's TSF. However, Alice has done 2 transactions with Charles, tr_{1r} and tr_{2r}. We assume that these two transactions may not have been recommended by Charles to Bob in the first round. We end up with one set[9]: link(tr_{1r}, tr_{2r}, Pub). Alice has done transactions with other people than Charles and Bob but tr_{1r} and tr_{2r} fills the tai gap and adding more transactions would increase nymity.
2. *Check that nymity of the selected link of evidence is reasonable compared to the utility*: if yes engage in further trade steps; else abort the trade. We assume that each utility level is associated with a maximum nymity threshold. This check corresponds to a cost/benefit analysis. So, the risk engine of the TSF should be responsible for carrying out this analysis. The tai gap message may be treated as a request from the requested entity to the requesting entity. If the trustworthiness of the requesting entity in keeping private information for personal use only is available, it is possible to have finer PAO. A level may be: PII-information kept for personal use. For example, this level happens when users subscribe to privacy policies specifying that their private information will not be disclosed to third parties. The consequence of detecting breached privacy policies is lower trustworthiness. In this case, the check also uses the trust engine as in the standard decision-making process of a TSF.

A difficult aspect of the liseng algorithm is to take into account the sequencing of interactions. Pieces of evidence revealed before the current interaction can impact the selection as well as future pieces of evidence due to the combination of pieces of evidence. For example, for two candidates ev_1 and ev_2 with same tai but different nymity ($nymity_1 < nymity_2$), in the scope of this specific interaction, ev_1 should be chosen. However, if a future interaction links ev_3 with $nymity_{link(ev1,ev3)} > nymity_{link(ev2,ev3)}$, the choice becomes more difficult.

By allowing any entity to make recommendations we directly support a change of identity, where evidence can be transferred and linked to the new identity through a recommendation, without explicitly linking the two identities. This limits the extent of the profile that can be built for a given virtual identity, thereby reducing the violation of privacy resulting from a single transaction being linked to the real-world identity of a user. So, in step 3, a list of pseudonyms owned by the requesting entity could be sent back as potential new recommenders. If the requested entity has not already used these pseudonyms as recommenders, it would do so. However, the tai of evidence

[9] There are two choices to retrieve the recommendations rec_1 and rec_2 associated with tr_{1r} and tr_{2r}: either Alice's TSF contacts Charles to get the signed recommendations and passes them back to Alice, or Bob's TSF contacts Charles to get the signed recommendations.

provided by these entities would be discounted by the recommendation process. This is why it may be more beneficial to make the link between some pseudonyms explicit as explained in the next section.

4 Linking Evidence on Multiple Virtual Identities

In the above privacy/trust trade model, we said that a virtual identity vi is a set of linked pieces of evidence, indeed vi is the result of linking evidence with its own nymity and tai. In our example implementation, evidence is linked through digital signature validation. In this case, it is possible to link virtual identities as it is possible to link any other piece of evidence. For example, we may have link(Pub', vi) = link(Pub', tr_1, …, tr_i, …, tr_n, Pub) with tr_i being all n transactions linked to Pub. It is worth noticing that we also implicitly link all m transactions tr'_j linked to Pub': link(Pub', vi) = link(tr'_1, …, tr'_j, …, tr'_m, Pub', tr_1, …, tr_i, …, tr_n, Pub) = link(vi', vi). In our payment scenario [17], customers are given the possibility to generate pseudonyms on demand in order to protect their privacy. However, due to the resulting division of evidence between virtual entities, it takes more time for these virtual entities to reach the same trustworthiness than for a unique virtual identity. So, customers can link virtual identities during trust calculation in the privacy/trust trade process (depicted in Figure 6).

This new prospect for linking evidence allows us to envisage new linked evidence in step 2.1 of Figure 6. So, in step 3, a list of pseudonyms owned by the requesting entity could be sent back as potential new evidence of the form: link(Pub_1,…, Pub_i,…,Pub_n) with Pri_i known by the requesting entity for all i. In step 1 of the liseng algorithm (using the example we presented in Section 3 when describing this algorithm), another choice may be to use two transactions, tr_3 and tr_4, that Alice under the pseudonym Pub' did with Bob: the resulting link can be specified with more or less explicit linked evidence depending on what can be implicitly linked. For example, if the TSF does not guarantee that all transactions done under a specific pseudonym can be available in a timely manner (especially for recommendations), the explicit link should be longer: link(tr_3, tr_4, Pub', Pub). If any transaction is guaranteed to be known by all entities[10], it would be sufficient with a link of this type: link(Pub', Pub). Anyway, the first choice that we had, link(tr_{1r}, tr_{2r}, Pub) has low nymity because the implicit link appears somewhere in clear and can be established if other legitimate means are used. From a tai point of view, both give the same tai if each transaction gives the same tai and linking two keys is not acknowledged further. However, link(tr_3, tr_4, Pub', Pub) has potentially high nymity (it is intuitively higher than link(tr_{1r}, tr_{2r}, Pub) because Pub' could be used in another context and/or in the future whilst still being linked). Then, the link between two virtual identities is permanent and cannot be easily undone (e.g., as explained at the end of this section, when we link two keys, we use the fact that an entity cryptographically shows the ownership of both private keys of the two pseudonyms). It is important to note that transactions are often temporary, while linking transaction and/or virtual identities is permanent. This must be taken into account when estimating the utility of a given transaction.

[10] It is a strong assumption to guarantee global propagation of information. This assumption is not realistic in most scenarios (e.g., when random disconnection is possible).

We emphasize that care should be taken when linked evidence on multiple virtual identities is assessed. The most important requirement is to avoid counting the same evidence twice when it is presented as part of two different pseudonyms or overcounting overlapping evidence. In some cases, passing recommendations in the form of a simple trust value, instead of all supporting information[11], does not fulfil the later requirement. Assessing evidence may require analysis and comparison of each piece of evidence to other pieces of evidence. For example, let assume that: we have the relation depicted in Figure 8 and we know the trust values of two virtual identities vi_1 and vi_2, tvi_1 and tvi_2 respectively.

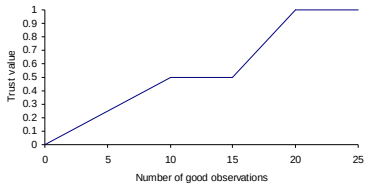

Figure 8: Example relation between observations and trust values

If $tiv_1 = 0.5$, whatever value tiv_2 is, we cannot compute the combined trust value without knowing the number of good observations, which is at a level of evidence deeper[12] than the level of trust values. In fact, assessing linked evidence requires great care and implementations may vary depending on the complexity of trust-lifecycle [19] and trust dynamics [6]. When recommendations are used, previous self-recommendations (i.e., recommendations from virtual identities belonging to the same entity) are also not easy to take into account. If this is part of a low cost mechanism for introducing new pseudonyms, it may be correct to simply discard the recommendations in the calculation. Another choice might be to consider such recommendations as evidence of untrustworthiness. Let vi_1 and vi_2 be two pseudonyms of the same entity. At the first interaction with the requested entity, vi_2 is used as a recommender for vi_1 due to the recommendation rec_{21}. So, the entity has now $link(vi_1,tr_1,rec_{21})$ for trustworthiness assessment of vi_1. At the second interaction, vi_1 discloses $link(vi_1,vi_2)$. Logically, the tai of rec_{21} needs to be revised, e.g., by discarding rec_{21} in the tai of the resulting evidence.

We shortly propose our view for a group of entities. A group may consist of a number of entities, the exact number of entities being unknown as well as the virtual identities of the entities part of this group. In this case, it is valid to assume that trust should be formed and built as if the group of entities would be indeed one conceptual virtual identity. For example, if a group signature scheme is used to sign and send messages on behalf of the entire team. In addition to the fact that powerful entity recognition could discern entities from such conceptual virtual identity, we see another case

[11] We agree that only passing the trust value may improve performance and may be better from a privacy point of view than all evidence information but it may also decrease interoperability as highlighted here, may show how another entity computes trust from evidence which may help to mount attacks and may reveal feelings towards other entities which may not be welcome.

[12] With this case of relation, it is also insufficient to only transfer the trust value in recommendations.

where a different approach would be welcome, especially when collaboration is from many-to-one entities. If two or more already known virtual identities make a specific request under an explicit group (i.e., the different members are known), the group should not be considered as a completely new virtual identity for several reasons (e.g., past history may show untrustworthiness or it may simply be unfair and inefficient to rebuild trust from scratch). Thus, a mechanism is needed to assess evidence from many virtual identities.

Combining levels of trust in entities is also very important when the ER process is used. The outcome of ER [15] can be a set of n principals p (i.e., virtual entity or pseudonym) associated with a level of confidence in recognition lcr:

$$\sum_{i=1}^{n}(p_i, lcr_i), e.g.\{(Pub_1, APERLevel1), (Pub_2, APERLevel1)\}$$

The above example is when an APER claim is signed by two keys[13] and both signatures are valid. It may be because both keys are indeed pseudonyms for the same entity or two entities decided to form a group and sign the claim as one entity. However, we envision that ER can be more proactive and uses evidence not directly provided by the requesting entities to compute a probability distribution of recognised entities. A range of methods can be used to compute this distribution (e.g., using fuzzy logic or Bayes). A person among n persons enters a building which is equipped with a biometric ER scheme. The outcome of recognition demonstrates hesitation between two persons: p_2 and p_3 are recognized at 45% and 55% respectively. So, all other principals are given 0%. We have:

$$OutcomeOfRecognition = \sum_{i=1}^{n} lcr_i\, p_i = 0*p_1 + 0{,}45*p_2 + 0{,}55*p_3 + \ldots + 0*p_i + \ldots + 0*p_n$$

If the level of trust in an entity is given by a value between [0,1], let say that p_2 is 0.5 and p_3 is 0.6. We then apply our simplest end-to-end trust model [15]:

End-to-end trust = aFunctionOf (Confidence In Recognition, Trust In Entity)
End-to-end trust = Level Of Confidence * Trust In Entity
End-to-end trust = 0.45 * 0.5 + 0.55 * 0.6

Once again, we assess evidence on different entities and care should be taken during the assessment.

Finally, we propose the following implementation[14] to carry out the privacy/trust trade process when pseudonyms are linked. Let p_1 be the requesting entity and p_2 the requested entity, they exchange APER claims with special keywords in Ctxt:

1: $p_1 \rightarrow p_2$: [GRANTX]p_1
2: $p_2 \rightarrow p_1$: [TAIGAP,HINT]p_2
3: $p_1 \rightarrow p_2$: [LINK]$p_1,\ldots,p_i,\ldots$

In step 2, HINT is optional and may contain hints for optimizing the liseng on the requesting entity's side. In fact, it may say which recommenders have been used for the first round of the trustworthiness assessment. It would then be known that it is useless to send back a link for the same recommenders. In our scenario, the HINT

[13] We restrain from using other technical trust cues (e.g., key length and algorithm)

[14] We use the notation: X is the special keyword used in the Ctxt of a claim, p is a principal; $p_1 \rightarrow p_2$ means that an APER Claim is sent from p_1 to p_2; [X]$p_1,\ldots,p_i,\ldots,p_n$ means that X is signed by several private keys, e.g., p_i's Pri

consists of a list of other pseudonyms (video clubs) owned by the video club holding company. Then, the liseng should try to link evidence to these pseudonyms. In step 3, the LINK lists other Pubs that are linked to p_1 and the claim must be signed by the Pri of each listed Pub. For example, in Alice's scenario, we have:

1: $p_1 \rightarrow p_2$: [GRANTX("rent large video display")]PubAliceSunnyVillage
2: $p_2 \rightarrow p_1$: [TAIGAP("strong positive impact"),HINT("PubVC1")]PubVC2
3: $p_1 \rightarrow p_2$: [LINK("PubAliceSunnyVillage,PubAliceRainyTown")]PubAliceSunny Village,PubAliceRainyTown

Concerning the liseng, the provided hint allows the requesting entity's TSF to search straightaway evidence that can be linked to PubVC1 and find the link with PubAliceRainyTown.

5 Related Work

Although automated trust negotiation (ATN) [20] is argued to establish trust between strangers, the approach considerably differs from the TSF's approach described in Section 1 (e.g., as used in SECURE). The method consists of iteratively disclosing digital credentials between two entities. Through this sequence of bilateral credential disclosures, trust is incrementally founded. The notion of trust formation and assessment based on past experience does not explicitly appear in ATN. However, the notion of negotiation underlined the importance of the "Minimal Linkability" principle and that care should be taken when more trust is asked before choosing to disclose linked evidence. In ATN, revocation is based on certificate revocation whereas in TSF-like approach the trustworthiness may be decreased without the use of certificates. In fact, certificates could be seen as another type of evidence and included in the list of evidence of our privacy/trust model. Revocation implies that a piece of evidence based on a credential also varies over time. It is beyond the scope of the paper to fully study credentials but the following points are worth mentioning. First of all, credentials can be redundant. The issue appears when virtual identities are combined. Patient ID could be linked with another credential (e.g., Driver License) as well as Employee ID. However, when Patient ID is linked to Employee ID, the logic would be that Driver License should be counted once. Winslett encourages more work on the issue of multiple virtual identities and this paper is a contribution on this topic. Also, ATN is known to have not fully resolved privacy issues [21]. In our approach, it is possible to use pseudonyms and to stop using a specific pseudonym. This has the effect to break too much evidence accumulation.

Another type of evidence that can be used in our privacy/trust trade model is reputation. By reputation, we mean that a piece of evidence on the trustworthiness of another entity is given by a supposed large number of entities but unknown. Again, it is not clear how reputation should be combined if the goal is to avoid overcounting overlapping evidence.

Wagella et al. [19] 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, as highlighted in this paper, it may be difficult to evaluate trustworthiness in first place without enough evidence linked with the receiving entity.

The work on modelling unlinkability [9] and pseudonymity [4, 8] is valuable towards founding privacy/trust trade. Previous work on pseudonym credential system should be useful to formally prove (in future work) that an entity really owns different private keys. The Sybil attack [3], which challenges the use of recommendations, is also worth keeping in mind when providing means to create virtual identities at will without centralized authority.

6 Conclusion

There is an inherent conflict between trust and privacy because both depend on knowledge about an entity but in the opposite ways. Although trust allows us to accept risk and engage in actions with a potential harmful outcome, a computational TSF must take into account that humans need (or have the right to) privacy. Trust is based on knowledge about the other entity: the more evidence about past behaviour is known, the better the prediction of future behaviour will be. This is why we propose to use pseudonymity as a level of indirection, which allows the formation of trust without exposing the real-world identity.

However, depending on what benefits can be reaped through trustworthiness, people may be willing to trade part of their privacy for increased trustworthiness: hence, contextual privacy/trust trade is needed. We propose a model for privacy/trust trade based on linkability of pieces of evidence. If insufficient evidence is available under the chosen pseudonym, more evidence may be linked to this pseudonym in order to improve trustworthiness and grant the request. We present a protocol for explicitly disclosing to the requested entity that some evidence can be linked. Some thresholds should be set concerning the acceptable evidence that should be disclosed. This is why we introduce the liseng algorithm to ensure that the Minimal Linkability principle is taken into account. During a trade process, tai, nymity and utility must be balanced.

We then explain that it may be more beneficial to make the link between some pseudonyms explicit (e.g., to avoid discounted evidence or reduce the time to reach trustworthiness due to division of evidence between virtual identities). We show how we implemented this on top of the APER scheme.

We emphasize that care should be taken when linked evidence on multiple virtual identities is assessed, especially when pseudonyms are linked during the privacy trade/process but also when groups and the outcome of entity recognition result in a set of possible principals (as defined in ER).

As levels of privacy asset, trust assessment impact and utility are key metrics to carry out minimal linkability, we are trying to enhance the trade in our prototype with real metrics on privacy loss and trust gain extracted from localized payment transactions.

This work is sponsored by the European Union, which funds the IST-2001-32486 SECURE project and the IST-2001-34910 iTrust Working Group.

7 References

[1] V. Cahill, et al., "Using Trust for Secure Collaboration in Uncertain Environments", in *Pervasive Computing*, vol. 2(3), IEEE, 2003.

[2] T. M. Cooley, "A Treatise on the Law of Torts", Callaghan, Chicago, 1888.

[3] J. R. Douceur, "The Sybil Attack", in *Proceedings of the 1st International Workshop on Peer-to-Peer Systems*, 2002.

[4] I. Goldberg, "A Pseudonymous Communications Infrastructure for the Internet", PhD Thesis, University of California at Berkeley, 2000.

[5] X. Jiang, J. I. Hong, and J. A. Landay, "Approximate Information Flows: Socially Based Modeling of Privacy in Ubiquitous Computing", in *Proceedings of Ubicomp 2002*, pp. 176-193, Springer-Verlag, 2002.

[6] C. M. Jonker and J. Treur, "Formal Analysis of Models for the Dynamics of Trust based on Experiences", in *Proceedings of the Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 1999.

[7] A. Jøsang, "The right type of trust for distributed systems", in *Proceedings of the 1996 New Security Paradigms Workshop*, ACM, 1996.

[8] A. Kobsa and J. Schreck, "Privacy through Pseudonymity in User-Adaptive Systems", in *ACM Transactions on Internet Technology*, vol. 3 (2), 2003.

[9] S. Köpsell and S. Steinbrecher, "Modeling Unlinkability", in *Proceedings of the Third Workshop on Privacy Enhancing Technologies*, 2003.

[10] M. Langheinrich, "Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems", in *Proceedings of Ubicomp 200*, Springer, 2001.

[11] S. Marsh, "Formalising Trust as a Computational Concept", PhD Thesis, Department of Mathematics, University of Stirling, 1994.

[12] A. H. Maslow, "Motivation and Personality", Harper, 1954.

[13] D. McKnight and N. L. Chervany, "The Meanings of Trust", MISRC 96-04, University of Minnesota, 1996.

[14] SECURE project, Website, http://secure.dsg.cs.tcd.ie.

[15] J.-M. Seigneur, S. Farrell, C. D. Jensen, E. Gray, and Y. Chen, "End-to-end Trust Starts with Recognition", in *Proceedings of the First International Conference on Security in Pervasive Computing*, LNCS, Springer, 2003.

[16] J.-M. Seigneur and C. D. Jensen, "Privacy Recovery with Disposable Email Addresses", in *Special Issue on "Understanding Privacy", December 2003*, vol. 1(6), IEEE Security&Privacy, 2003.

[17] J.-M. Seigneur and C. D. Jensen, "Trust Enhanced Ubiquitous Payment without Too Much Privacy Loss", in *Proceedings of SAC 2004*, ACM, 2004.

[18] S. Spiekermann, J. Grossklags, and B. Berendt, "E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior", in *Proceedings of the 3rd Conference on Electronic Commerce*, ACM, 2001.

[19] W. Wagealla, M. Carbone, C. English, S. Terzis, and P. Nixon, "A Formal Model of Trust Lifecycle Management", in *Proceedings of FAST2003*, 2003.

[20] M. Winslett, "An Introduction to Trust Negotiation", Proceedings of the First International Conference on Trust Management, LNCS, Springer, 2003.

[21] T. Yu and M. Winslett, "A Unified Scheme for Resource Protection in Automated Trust Negotiation", in *Proceedings of the IEEE Symposium on Security and Privacy*, 2003.