TEco: an integration model to augment the Web with a trust area for inter-pares interactions.

Gennaro Costagliola, Vittorio Fuccella, Fernando A. Pascuccio
Dipartimento di Informatica, University of Salerno
Via Giovanni Paolo II, 84084 Fisciano (SA), Italy
{gencos, vfuccella, fpascuccio}@unisa.it

Abstract

We propose an integrated and modular model called TEco. It is a Web-based trust area in which, through the integration of various systems, users interact with a greater degree of trust. In particular, users: own a Trusted Digital Identity to authenticate keeping anonymity (when required); establish Inter-Pares Interactions based on contracted agreements and knowing each other's reputation; can be the owners of the information they produce and protect their privacy. We discuss the feasibility of the model, the compatibility with the current Web and the things to do for putting it into practice.

Keywords: *Trust area, TEco, Privacy, Anonimity, Single Sign-On, Trust and Reputation, User-centric.*

1 Introduction

Despite its enormous success and indisputable usefulness, the Web is not exempt from problems that should be addressed to make it even more secure and reliable. In fact, some issues, such as the uncertainty of the identities, the almost complete lack of privacy and of guarantees on the reliability of the counterparts, i.e. the lack of trust among people, may limit its potential development [1]. Other issues are the lack of control and ownership of the information regarding a person or a company; the lack of specific information about service providers (e.g., reliability, quality, punctuality, etc. I.e, their reputation); the exploitation of anonymity to perform malicious actions [2].

In recent years many studies have focused on the development of new protocols and methodologies to allow unambiguous identification of the user, the ability to keep anonymity, to protect privacy and to authenticate once to access many services (Single Sign-On) [3, 4]. However, the aforementioned issues were almost always addressed individually.

DOI reference number: 10.18293/DMS2015-051

Generally, organizations (i.e. companies, academics, etc.) authenticate users and grant them roles through Identity Providers. The Identity Management systems, instead, adopting a user-centric paradigm, rely on the user rather than on the service provider for the control of digital identities [5, 4, 6]. The control of their own identities allows users to decide which information to share with others and under which conditions.

Based on the above reasons, our objective is the design of a comprehensive framework aimed at providing a trust area in the Web that combines the online and offline world smoothly and seamlessly, including the best solutions in a single model.

Our integrated and modular model is called *TEco* (acronym of Trust Ecosystem). Here, *ecosystem* means (see [7, 8]) a loosely coupled, domain clustered environment where each species conserves the environment, is proactive and responsive for its own benefits. In our case, species are the entities (e.g., users and online services) which preserve the environment and comply with fixed rules, are proactive and responsive as each of them, using a reward-punishment mechanism (feedback), contribute to the success of the system and, consequently, to their own benefit.

Digital ecosystems, as emphasized in [9], "can play the role of a unification *'umbrella'* over significant, challenging and visionary computing approaches that emerge in parallel". In this sense TEco will also act as a "field of comparison" and facilitate scientific communication in the sector.

The remainder of this paper is organized as follows: the next section summarizes some works related to ours; in Section 3 we introduce TEco; in Section 4 we discuss some practical issues related to the implementation of the model; lastly, in Section 5 we draw some conclusions and outline future work.

2 Related work

In the literature, we can find non-integrated solutions for:

- Identity Management systems (IdMs) and Single Sign-On (SSO);
- Trust and Reputation Management systems (TRMs);
- Anonimity and privacy protection;

Nevertheless, to the best of our knowledge, there are no studies in the literature that have faced all the aforementioned problems in a comprehensive and systematic view.

In their survey on IdMs [6], Torres et al. point out that the use of IdMs, which also enable SSO, may help to solve the new challenges related to security and privacy protection. Conversely, authors in [10] highlight some of their weaknesses, especially the impossibility for a user to decide which personal information to share with every service provider or to obtain information on their reliabilty. To address these issues, the authors propose techniques to integrate IdMs with Reputation Management Systems, which provide information on the past behavior of the service providers [2].

Other studies, driven by the emerging of new technologies, are focused on the next generation Internet, termed **Future Internet** [11]. Nevertheless, they do not provide a common view on what the FutureInternet is and mostly consider its network infrastructure, termed **Future Network** [6].

Despite the importance of many of the problems faced, such as infrastructural ones, we believe that even other aspects deserve attention, such as the relationship between digital identities and reputation and other little investigated sectors: the respect of user rights and the possibility for users to keep control of their data.

It is worth noting that Microsoft introduced a *Trust Ecosystem*, more narrowly defined as an environment that engenders trust and accountability between people and businesses [12]. In that system, users have several *Windows CardSpace* to access a service provider without having to authenticate [13]. Despite the similar name, our model includes more features and differs substantially from the one introduced by Microsoft, as we will show in the following.

3 Trust Ecosystem

The TEco system can be accessed by users (individuals and legal persons) and online services. All of them are considered as "entities" which interact with each other "at par" with no distinction between client and server, user and provider, services and humans. Following a user-centric paradigm, TEco was built by integrating different innovative systems to provide the following features:

- **Trusted Digital Identity**: every digital identity corresponds to an individual who is identified with "certainty" still keeping anonymity and privacy;
- **Content Management**: users are the owners of the information they produce and can manage such information autonomously;
- **Reputation Management**: it is possible to obtain reliable and updated reputation information about all users;
- **Interaction Agreement**: interactions are always based on a contract agreed between the parties, that have equal bargaining power;

The coexistence of these features makes TEco a trust area. In fact, users can mutually trust, as they are all identifiable, their reputation is known and while interacting, they can bargain conditions with law effectiveness. Furthermore, depending on their needs and the demands of others, users can decide which information to disseminate, protecting their privacy or maintaining complete anonymity.

3.1 Trusted Digital Identity

In the current Web, each entity has a **Digital Identity**, which can be defined as the digital representation of the information known about a specific individual or organization [3]. In TEco, in addition, each digital identity corresponds to an entity in the offline world whose identity is verified with certainty. To this end, an entity is required to register at TEco providing its own unique identifier. For individuals, this can be the identifier used by the governments for tracking their citizens as the National Identification Number. For corporate bodies (companies, organizations, associations, etc.) it can be their VAT number. To complete the registration to TEco it is therefore necessary that an entity proves to be the owner of the provided identifier. For instance, individuals could complete the registration at the Municipal Registry Office and legal persons at the Registry of Companies. Once the registration is completed, the entity will possess a **Trusted digital Identity (TId)** in TEco. The TId will correspond to an account associated with all the information available of the requester and its identifier. The requester is the only owner of the access credentials for that account. As the TEco is only accessible to certified digital identities, online services must have a TId too. In this case, the owner of the domain name must certify the association between provider's TId and URL of the service (used as its unique identifier). The whole registration process is handled by the **Identity Management Systems (IDMs)** which assign and manage identities and belong to a **Federated Identity Management (FIdM)**. In general, a Federation can be defined as the set of agreements, policies,

standards and technologies to achieve its objective [3, 5]. The purpose of FIdM, is to allow entities belonging to different IdMs to be identified from all others, regardless the used authentication system (e.g. Kantara Initiative [1], Liberty Alliance [2], Shibboleth [14], Kerberos [15], etc.).

The IDMs are the only ones to know the association between offline world entities and their TId. For this reason, an entity must possess a **Web Alter Ego (WAE)** to interact in TEco, i.e. an alternative identity to present itself to others. Based on his/her needs, an individual can create different WAEs (e.g., as a researcher, as a chess player, etc.), choosing for each WAE which information to show among those associated to his/her own TId. Each WAE is completely independent from the others and is seen by counterparts as a separate entity. In fact, a counterpart cannot relate all the WAEs belonging to the same identity. This safeguards an entity's privacy, since it can use one of its WAEs without worrying that its true identity is revealed or that one of its WAEs is associated to others (in the following, we will see how this can be guaranteed through the use of temporary identifiers).

A registered entity to access TEco must logon at the IdM which manages its TId through the planned identification procedure (e.g. based on username/password, biometric data, etc.). Then, the entity receives from the IdM the list of all its own temporary identifiers, referred to as *TempWAEs*, specifically generated. Each of them uniquely identifies a specific WAE and allows the entity to interact within TEco without logging on to any specific service. This enables an SSO authentication. While the entity is "connected" to TEco, the TempWAEs are regenerated and sent back by the IdM to the entity periodically according to predefined security criteria or upon an entity's explicit request. It should be noted that the regeneration of the identifiers does not require a new logon. The TempWAEs' validity expires as the entity "disconnects", by logging out after an indefinite time. Besides identity management, TEco also provides a reputation system based on several **Reputation Management Systems (RMSs)**, each responsible to collect, aggregate and disseminate data on the reputation of the entities [2]. The RMSs belong to a **Federated Reputation Management System (FRMS)**, which manages their interaction. The integration of FRMS and FIdM provides the users with a high level of mutual trust. In fact, they are encouraged to take appropriate behavior because they know they are identified with certainty and their past behavior is known to all. The greater mutual trust increases the *social capital*, intended as the richness of the interactions between members, which itself affects the reputation system encouraging an active and honest participation and thus increasing its effectiveness [16]. The FIdM assigns each entity a reference RMS which is also involved in managing the reputation of all its WAEs. At the end of an interaction, an entity is required to leave an anonymous feedback on the counterparts to its reference RMS. The latter, in turn, according to the times and rules set by the federation, sends the feedback to the reference RMS of the recipient entity. An entity can request the reputation of the other entities to its own reference RMS, which obtains it through the federation. It is worth recalling that, being independent, each WAE has its own reputation independently from others. Since a good reputation requires time, this reduces the proliferation of WAEs (see "newbies" in [17]).

[1] www.kantarainitiative.org

[2] www.projectliberty.org

3.2 Inter Pares Interaction

In the current Web, the users share information and request services by establishing interactions. In TEco, any interaction is always based on a contract agreed between the parties. We refer to the interaction as **Inter Pares Interaction** (in the following referred to only as "*TEco Interaction*") and to the contract as **Negotiated Interaction Agreement** (in the following referred to only as "negotiated agreement"). The negotiated agreement is composed of two parts: the first, preliminary and fixed, contains the principles and general conditions that oversee any interaction in TEco (e.g., to respect owners' constraints on the data, not to maliciously alter reputation, etc.). The second part is subject to negotiation and contains a list of **Agreement's Terms** (in the following referred to only as "term"), i.e. constraints and preferences established in a formal language, that the parties agree to comply with. If some constraints in the *negotiated agreement* are not respected by one of the parts, as terms of a contract with the force of law, can be asserted in judicial offices. Since, as stated in [17], an entity interacts with the others in a given context and assuming a specific role, an **Interaction Context/Role (ICR)** in the *negotiated agreement* will also be mandatorily negotiated. For instance, the consultation of a website is a typical "interaction" between end-user and website owner, where the ICR for the user is "content visualization/reader".

The *negotiated agreement* is established through a phase of **Negotiation of the Agreement** (in the following referred to only as "negotiation"), in which each party sends the other its contract proposal, called **Interaction Agreement** (in the following referred to only as "agreement"), composed of the list of *terms* that a party intends to include in the second part. During negotiation, each *term* can be modified or accepted to reach the *negotiated agreement* in its final form. If all parties agree, the *negotiated agreement* can be changed at any time. Clearly, an entity that does not conclude the phase of *negotiation* can not take part in the interaction.

The *terms, agreements e negotiated agreements* are

defined through a formal language. This allows the entity to participate to the *TEco interaction* through a **Web Agent**, which suggests or takes decisions on the basis of its acquired experience (self-learning), on the type of entity (e.g., individual) and on the context/role (e.g., e-learning/instructor). For instance, in the case of an individual, human intervention may be required during bargaining. In the context/role e-commerce/seller, the *negotiation* phase of the seller is automatically handled by the Web Agent and the human intervention is not required, unless expressly prescribed by the seller. It should also be pointed out that an *agreement* can be defined by including only standard *terms* that are stored in an archive at the FIDM which also manages an archive of default *agreements*. A new *agreement* is created by choosing the *terms* from a list of standard ones through an appropriate GUI. In order to simplify the *negotiation* phase, while logging on to TEco, the entities receive (similarly to WAEs) lists of predefined *terms* and *agreements* from the FIdM. This way, they can set an *agreement* for each WAE choosing it from the default ones. For instance, the entity could select a WAE called "*Web surfing*" associated to an *agreement* called "*High Privacy*", requiring counterparties not to request private information such as the *home address*.

Figure 1 shows how two entities establish a *TEco interaction* (the schema can be extended to more than two entities). We use the following notation: **TempWAE** for the temporary identifier of a entity's Web Alter Ego; **PermWAE** for the permanent one. It is worth recalling that permanent identifiers are never disclosed to entities. As shown in the figure, an interaction is composed of the following steps:

- Step 1. *Ann* requests an interaction to a service provider (SP) providing the WAE (*TempWAEa*) with which she intends to identify herself and her *agreement*;
- Step 2. The SP requests to its reference RMS (*RMSp* in the figure) the reputation associated to *TempWAEa* and related to the context/role (*ICRa*) provided by *Ann* in her *agreement*;
- Step 3. *RMSp* requests to the FIdM the permanent identifier (*PermWAEa*) associated to *TempWAEa*;
- Step 4. Once obtained the *PermWAEa*, *RMSp* checks if it has the reputation associated to *PermWAEa* in the context/role *ICRa*. If not, *RMSp* requests it to the FRMS.
- Step 5. Then, *RMSp* returns to SP the reputation of *TempWAEa* in *ICRa*. It is worth noting that SP receives the reputation of Ann's WAE knowing only her temporary identifier.
- Step 6. SP decides, based on the received reputation, whether to accept TEco interaction request. If so, SP sends *Ann* the WAE with which it intends to interact (*TempWAEp*) and its own *agreement*. Otherwise, it sends a message of rejection and abandons the interaction.
- Steps 7-11. The same actions performed in Steps 2 - 6 on *SP*'s side are now executed on *Ann*'s side. Step 11 opens the negotiation phase which ends with the negotiated agreement.

The *TEco interaction* was schematically shown in sequential steps in order to facilitate the exposure but, actually, some steps may be performed in parallel (e.g., the negotiation phase). As previously mentioned, the parties may express a feedback on counterparts at the end of the interaction.

Nevertheless, to prevent malicious attacks and improve the reputation system, TEco adopts some important countermeasures already described in [17]. After the negotiated agreement is established and before starting a *TEco interaction*, an entity's Web Agent sends to its *reference RMS* a list of pairs ICR-WAE, each referred to an entity it is going to interact with. The RMS, in turn, sends back to the entity the **Interaction Token** with which the interaction will be uniquely identified for a predetermined time interval. This token allows the RMS to accept only feedbacks to and from entities that indeed took part to the interaction and to make sure that the interaction indeed took place. Therefore, every feedback must include the *interaction token* and the TempWAEs of both the judging and the judged entities. This assures that a feedback is expressed once for each entity involved in an interaction.

Furthermore, the RMS could release encrypted reputation data with date and time of encryption. This ensures data integrity and authenticity. This also speeds up the reputation retrieval, since entities may store the encrypted reputation data and share them with counterparts without querying the FRMS. Counterparts may decide whether to query the FRMS on the basis of both the certification date and the reputation of the entity (too old data may be untrustworthy). We remark that the presence of a contract having the force of law strongly discourages illicit proactices, as they can be prosecuted.

3.3 Content Management Framework

As mentioned before, one of the objectives of TEco is to ensure that the entities are direct owners of the information they produce. To this aim, an important role is played by the **Content Management Framework (CMF)**, which manages all data (text, multimedia, WAE's attributes, etc.) related to the entities. Whenever a new content is created

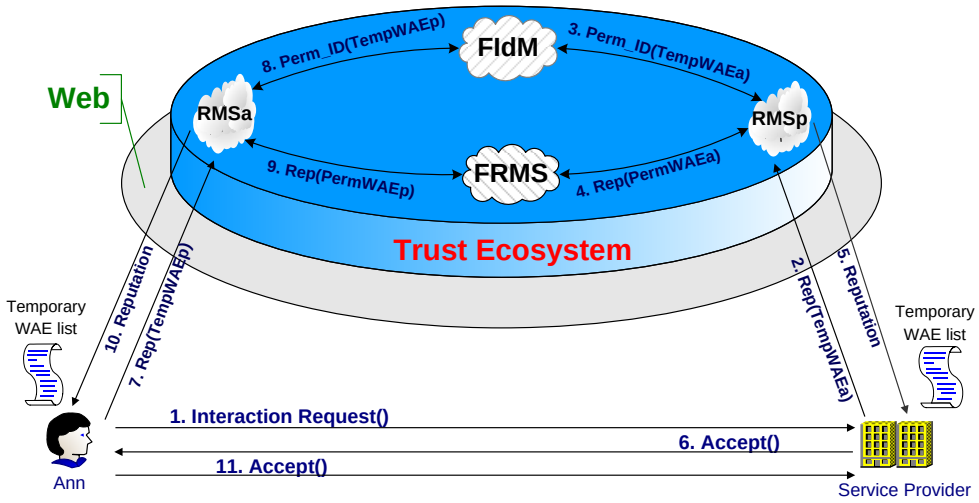

Figure 1: Interaction between entities.

at a service provider, the CMF automatically creates a link between the content and the producing entity. These data are physically stored at the SP or in personal cloud storage systems, local hard disks, etc. In any case, they are property of the entity that produced them, which can decide which *access rights* to grant to other entities (reading, modification, deletion, duplication, disclosure, etc.). Furthermore, the time validity of each *privilege* can also be established. Therefore, contrary to what normally happens in the current Web where the users are deprived of these rights, in TEco the users have management and responsibility of their own data. The CMF ensures that the rights set by the content owner, with any changes, are disclosed to the other entities and that they respect such rights. To this end, the CMS tracks the use of contents by entities and, in case of infringement of privileges, it requests the FRMS to lower the reputation of the infringing entity and, in extreme cases, that it is excluded from TEco. Another duty of the CMF is to "certify" with legal value the publication of a content on the Web. This feature is strongly felt by many users in the current Web. Let us consider, for example, the case of a university that has issued a call for a research grant. The CMF must certify: 1) that the URL of the content is accessible at any time (**W**here); 2) the date and time (timestamp) of the publication (**W**hen); 3) the integrity of the content (**W**hat); 4) the authenticity of the publication, i.e., that it comes directly and truly by the entity (**W**ho). If the content is modified after publication, the CMF certifies the four ***W*** for all the previous versions, which are still stored (*versioning*) and made public. As for RMSs, the CMFs are also part of a federation, called **Federated Content Management Frameworks (FCMF)**, which manages their interaction.

4 Discussion

The TEco system is an incremental model which enhances the current Web without replacing it. This is one of its strengths as it requires no upheavals in infrastructures. Furthermore, it does not compel users to adapt to new rules or new software. TEco can be developed in parallel with the Web, leaving the users free to choose between a deregulated area and a trust area, exactly as in the offline world. To make it applicable it is necessay that the systems described so far, federated Identity Management system (IdM), federated Reputation Management System (RMS) and federated Content Management Framework (CMF) are implemented and integrated bearing in mind the characteristics described in this section. To obtain the permanent identifier associated to a TempWAE (see Fig. 1 - Steps 3 and 8) from the federated IdM, the federated RMS must use a specific communication protocol that may be similar to the protocols used in Internet to resolve domain names. This protocol must ensure that the association between permanent and temporary identifiers is known only to the two federations. The implementation of federated RMS also requires that an ontology of the Contexts/Roles and, for each of them, the Main Features are identified, as explained in [17], to which the reader should refer for further details.

It is necessary to define a formal language for the specification of Agreement Terms and Interaction Agreements and to define Standard Agreement Terms and some predefined Interaction Agreements. It is also crucial for the success of TEco the implementation of an efficient Web Agent that facilitates entities in all activities related to TEco. In particular, it could include a plugin which works during Web navigation (e.g., as done in [18]). This plugin would allow

an entity to request a *TEco Interaction* by simply entering the address of the website in the browser and specifying the alter ego s/he intends to use. It will then be the Web Agent to handle the request by interacting with the service providers (see Fig. 1 - Steps 1 and 6). A *TEco Interaction* will be established if and only if the service provider, which also owns a TId, accepts the request. In both cases, the navigation would continue normally, except that a *TEco Interaction* will enable all the benefits of TEco (negotiated agreement, SSO, reputation, etc.) and a browser icon will indicate that the transaction is performed in the trust area (as in https). A protocol for negotiation of the agreement is also necessary to allow the Web Agents to perform it autonomously.

As already mentioned, the federated CMF has to manage all the contents and information related to a TId. This may be done by associating a tuple *[name, url_of_value, is_certified]* to each content, where: *name* represents the attribute name, which can be standard (e.g. *date_of_birth*) or user-defined (e.g. *preferred_wine*); *url_of_value* indicates where the attribute value is located (e.g. at a Municipal Registry Office, a link to a *Google+* post, etc.); *is_certified* indicates whether the attribute is declared by the entity or the url is referred to a *certified* value. Whenever a new attribute of an entity is declared, a new record will be added in the CMF. For instance, following the achievement of the PhD in computer science, a new record like *[PhD, www.unisa.it/ PascuccioFA/CSPhD, true]* will be added for the corresponding TId. To simplify the handling of content for an entity, its Web Agent could support the user during the creation of new contents. For instance, when a user publishes in a blog, his/her Web Agent suggests a default repository (the user can chose another one) in which to store the data and then sends a link to the content to the blog. If the content is already present in the CMF, the user can simply choose the link without rewriting the text. This would be totally transparent to the user, who would only compose the content through an appropriate GUI, while all other activities would be carried out independently by the Web Agent.

5 Conclusions and Future Work

In this work we discussed some critical issues related to the current Web and proposed an overall solution called TEco, which defines a trust area in the Web, where users can move and safely interact with a greater degree of mutual trust. We showed how in TEco entities can: be identified through a Trusted Digital Identity; keep anonymity and protect their privacy through the use of a Web alter ego; perform Single Sign-On authentication; establish inter pares interactions tying counterparts to comply with specific and agreed conditions; know the reputation of counterparts and have complete control of their data. We also discussed how it can be implemented through the integration of some existing and new systems and how this enhances the current Web without upheavals.

The work is still preliminary. In the future we will continue to work on TEco taking into account the contributions received by the scientific community. In addition, we will develop the communication protocols among all subsystems and the formal languages to define the Agreement Terms and the Interaction Agreements. Lastly, we will develop a prototypical Web Agent with a basic expertise to enable the testing of TEco.

References

[1] S. Srinivasan and R. Barker. Global analysis of security and trust perceptions in web design for e-commerce. *Int. J. of Inf. Security and Privacy*, 6(1):1–13, 2012.

[2] F. Hendrikx, K. Bubendorfer, and R. Chard. Reputation systems: A survey and taxonomy. *Journal of Parallel and Distributed Computing*, 75:184–197, 2015.

[3] E. Bertino, F. Paci, and N. Shang. Digital identity protection - concepts and issues. *ARES '09*, pages lxix–lxxviii.

[4] G. D. Tormo, G. L. Millán, and G. M. Pérez. Definition of an advanced identity management infrastructure. *Int. Jour. of Infor. Security*, 12(3):173–200, 2012.

[5] Y. Cao and L. Yang. A survey of identity management technology. *ICITIS '10*, pages 287–293.

[6] J. Torres, M. Nogueira, and G. Pujolle. A survey on identity management for the future network. *IEEE Comm. Surveys and Tutorials*, 15(2):787–802, 2013.

[7] H. Boley and E. Chang. Digital ecosystems: Principles and semantics. In *DEST '07*, pages 398–403.

[8] E. Chang and M. West. Digital Ecosystems A Next Generation of the Collaborative Environment. *iiWAS '06*.

[9] E. Pournaras and S.J. Miah. From metaphor towards paradigm - a computing roadmap of digital ecosystems. *DEST '12*, pages 1–6.

[10] G. D. Tormo, F. G. Mármol, and G. M. Pérez. Towards the integration of reputation management in openid. *Computer Standards and Interfaces*, 2013.

[11] S. Paul, J. Pan, and R. Jain. Architectures for the future networks and the next generation internet: A survey. *Computer Communications*, 34(1):2–42, 2011.

[12] B. Gates. Bill Gates: Microsoft's Security Vision and Strategy. *RSA 2006*.

[13] H. Jo, H. Jin Lee, K. Chun, and H. Park. Interoperability and anonymity for id management systems. *ICACT 2009*, 02:1257–1260.

[14] R. L. Morgan, S. Cantor, S. Carmody, W. Hoehn, and K. Klingenstein. Federated Security: The Shibboleth Approach. *EDUCAUSE Quarterly*, 27(4):12–17, 2004.

[15] B. C. Neuman and T. Ts'o. Kerberos: An authentication service for computer networks. *Communications Magazine, IEEE*, 32(9):33–38, 1994.

[16] W. Sherchan, S. Nepal, and C. Paris. A Survey of Trust in Social Networks. *ACM Comp.Surv.*, 45(4):47:1–47:33, 2013.

[17] G. Costagliola, V. Fuccella, and F. A. Pascuccio. Towards a Trust, Reputation and Recommendation Meta Model. *JVLC*, 25(6):850–857, 2014.

[18] G. Costagliola, R. Esposito, V. Fuccella, and F. Gioviale. An architecture for user-centric identity, profiling and reputation services. *DMS 2009*, pages 170–173.