

Managing One's Identities in Organisational and Social Settings

Katrin Borcea-Pfzmann, Marit Hansen, Katja Liesebach, Andreas Pfzmann, Sandra Steinbrecher

Interacting in the Internet, users should be empowered to use only those subsets of their personal attributes, called partial identities, which are appropriate for the actual situation and context. Refraining from acting under few and easily linkable partial identities is a prerequisite for trustworthy privacy. Traditionally user-controlled identity management systems primarily support individuals interacting with organisations, but mainly ignore special needs which arise if individuals interact with each other. To support online communities those systems have to change.



A. Pfzmann, K. Borcea-Pfzmann, M. Hansen, S. Steinbrecher, K. Liesebach (from right).

From TU Dresden: Dipl. Medien-inf. Prof. Dr. Andreas Katja Liesebach and Pfzmann, head of Dipl.-Inform. Sandra Steinbrecher. **From ULD, Kiel:** Dipl.-Inform. Marit Hansen, head of PET division. Katrin Borcea-Pfzmann,

Research Focus: Privacy in Identity Management and Application Design¹

1 Managing Identities

Identity management (IDM) systems support persons being represented by sets of data, so-called *digital identities*, which can be managed by technical means. Depending on the situation and the context, only subsets of these attributes are needed to represent a person in the physical or the digital world, so-called *partial identities* (pID) [Pfzmann/Hansen 2007]. A person typically uses different pIDs for different situations (e.g., work, leisure activities, dealing with companies). *Pseudonyms* act as identifiers of these pIDs. Although pIDs belong to the identity of a person, this link may not be visible for arbitrary observers. Pseudonymity of pIDs encompasses the entire field between and including anonymity and identifiability [Pfzmann/Hansen 2007].

Given the growing number and variety of interactions in the online world an individual is involved in, various digital identities are being developed by these interactions. Thus, managing one's identities becomes a necessity.

User-controlled IDM systems aim at helping individuals to manage their partial identities and the corresponding user accounts. Individuals establish partial identities with online applications according to their roles and the specific context.

These IDM system run on machines trusted by the individual (usually in his possession) and under his control. Current development of user-controlled IDM systems strives for a privacy-enhancing design. An appropriate reference architecture and different prototypes have been developed in the project PRIME – Privacy and Identity Management for Europe [PRIME 2007].

Other types of IDM systems are employed by application providers to manage the users registered with them, especially for authentication, authorisation and accounting (AAA), and possibly also for profiling.

Past work on user-controlled IDM focused on individuals managing their pIDs mainly used in interactions with organisations, e.g., shops or public administrations.

2 IDM in Communities

2.1 Multilateral Interactions

If interactions may take place between arbitrary entities, we speak of *multilateral interactions*. Entities are defined as individuals or organisations or subsets of them, where organisations represent enterprises as well as public administrations. Interactions between individuals and organisations are usually modelled as 'customer ↔ business' and 'citizen ↔ administration' and follow specific workflows usually defined by the organisations. In contrast to this, multilateral interactions between individuals often do not adhere to such fixed workflows.

Communities and collaborative networks typically base on multilateral interactions between their members. From a technical perspective, most of the communities and collaborative networks are implemented in a centralised way, i.e., one or several providers operate the technical systems (e.g., mailing lists, newsgroups, or web applications such as eBay) that distribute messages to the community. Most of the currently available web-based systems already integrate IDM systems for AAA and possibly also for profiling. However, user-controlled identity management, as we focus on, usually remains unconsidered.

¹ This work was funded in part by the PRIME project which receives research funding from the

European Union's Sixth Framework Programme and the Swiss Federal Office for Education and Science.

2.2 From ARPANET to eBay

In the early days of the Internet (called ARPANET then), the Internet was a means of supporting isolated communities across distance primarily in space, but secondarily in time as well. The aim of the Internet users was mainly to exchange and thereby to share information, e.g., thoughts and knowledge. Neither security nor privacy was much of an issue: Members of the community mostly knew and trusted each other at least w.r.t. the aims of the community. With very few exceptions, each user was part of one Internet-supported community only. Already then, community members had a well-defined pID within „their“ community. Communities were self-regulating where law in general and law enforcement in particular rarely interfered.

Since the 1990s, with the Internet gaining commercial significance and getting many more people „into the net“ (cf. the eBay community), the following properties started changing:

- ◆ Since knowing each other from the offline world is no longer the norm, defining, showing, proving and checking identities becomes relevant in the online world.
- ◆ Since a growing number of online transactions involve larger amounts of money, advanced security requirements arise.
- ◆ Since the ordinary user of the Internet takes many roles and quite probably is a member of many communities, privacy becomes an issue requiring defining, developing, showing, proving and checking pIDs as well as limiting their release.
- ◆ Since technologies for providing a user's accountability, yet keeping his privacy, have been developed, linkability by multiple usages of unique pIDs should be limited. Thus, certified attributes, the validity of the attribute and its certification shall be transferable to other pIDs of the same user while maintaining privacy.

2.3 Properties

Whereas, in the case of interactions between individuals and organisations, interactions are well-defined and clearly structured, multilateral interactions in communities and collaborative networks among individuals are different.

The following paragraphs give an overview of properties of organisations, on the one hand, and individuals, on the other hand,

Property	Organisation	Individual
<i>Interests</i>	<ul style="list-style-type: none"> ◆ superordinate interests not necessarily matching the ones of its members ◆ less context-dependent ◆ do not change frequently 	<ul style="list-style-type: none"> ◆ specific (self-)interests depend on manifold purposes ◆ highly context-dependent ◆ change frequently
<i>Being addressed by law</i>	<ul style="list-style-type: none"> ◆ manifold legal obligations, e.g., from private or public law 	<ul style="list-style-type: none"> ◆ protection by Human Rights ◆ only very few legal obligations
<i>Expected „service level“ in ensuring confidentiality, integrity and availability</i>	<ul style="list-style-type: none"> ◆ guaranteeing service levels w.r.t. confidentiality, integrity, availability is necessary 	<ul style="list-style-type: none"> ◆ no professional service level expected
<i>Methods enabling interactors to estimate if others will behave as expected</i>	<ul style="list-style-type: none"> ◆ compliance with quality standards ◆ supervisory inspections ◆ centralised databases of creditworthiness of organisations ◆ internal checks and audits 	<ul style="list-style-type: none"> ◆ interactors personally know each other ◆ certified attributes indicating estimated trustworthiness ◆ reputation systems

Table 1: Properties of individuals vs. properties of organisations having impact on the design of IDM systems for multilateral interactions

that need to be considered when designing comprehensive IDM systems. Table 1 shows that the foundations of interaction for organisations and individuals significantly differ from each other:

Interests: An *organisation* providing services via the Internet primarily focuses on one or few purposes. It has superordinate interests which are mostly defined by the organisation's objective or business. These interests do not necessarily reflect the interests of its members (i.e., employees, staff). The organisation's interests are usually barely context-dependent because of its concentration on only one or at most a few services it offers. Therefore, interests of an organisation do not frequently change.

In contrast to organisations, an *individual* has a large number of (self-)interests which deeply depend on the actual contextual situation as well as on the purposes in personal life. A further property of an individual is that its interests may frequently change, depending on the context.

For example, someone is interested to share his photos with friends he spent his vacation with. But he has a strong interest that none of his colleagues gets to know these photos. Thus, he wants to keep apart the two contexts, personal life which is documented on photos and business life. However, he is willing to give recommendations to his colleagues regarding places of interests he visited during his vacation. In this case, he shares parts of his personal life with people from his business life.

Legal Foundations: Whereas *organisations* have to fulfil various legal obligations, e.g., from private or public law, only very few legal obligations apply to *individuals*. As far as individuals do not harm others and only personal matters are concerned, individuals are relatively free from legal obligations when interacting with others via Internet. They are even protected by Human Rights.

Service Level: In order to maintain services in a professional way, guaranteeing appropriate service levels is necessary for *organisations*. Concerning security requirements, this means implementation and maintenance of safeguards to provide confidentiality, integrity and availability of the data transmitted. Therefore, professional organisations need to provide a secure application environment. Furthermore, they employ system administrators as well as security and privacy officers with clear responsibilities and establish proxies for absent staff.

Obviously, *individuals* are not able to provide those mechanisms. That is, the interaction partners cannot expect from them professional service levels as individuals cannot be available all the time and usually have neither professional expertise nor the infrastructure to implement and maintain necessary technical and organisational methods to ensure confidentiality, integrity and availability of data interchanged.

Behaviour assessment: To enhance the predictability of future behaviour, specific methods are required which enable the interactors to estimate if the others will behave as expected.

In case of *organisations*, several quality standards (e.g., ISO9000 or sector-specific guidelines) exist, which organisations have to comply with. In many cases organisations are additionally inspected by supervisory authorities on a regular basis or on occasion. Furthermore, centralised databases inform on creditworthiness of organisations. Also, entries in reputation systems help to assess the trustworthiness of the interaction partner. As reputation usually is regarded as important for organisations, they often establish internal checks and audits to prevent misbehaviour resulting in bad reputation.

In contrast to organisations, *individuals* do not have to comply with specific quality standards (other than acting according to law and social norms). However, additional trust-building methods can be useful, e.g.:

- ◆ certified attributes which give hints on estimated trustworthiness or
- ◆ reputation systems that collect and aggregate interactors' experiences from former interactions.

2.4 Protections Goals

Interactors necessarily have several security requirements in common to protect the interaction against parties not involved, e.g., in terms of confidentiality, integrity and availability. In addition, interactors typically have expectations regarding the behaviour of the interaction partners that these might fulfil or not.

Fulfilment of expectation may be defined implicitly (e.g., behaviour follows social norms) or explicitly (e.g., on the basis of a contractual agreement). Of course, an IT system should aim at fulfilling the security requirements all involved interactors have agreed upon.

Following, we illustrate well-known requirements concerning the IT system for three of the most important protection goals as well as typical expectations regarding the (possible) interactors:

- Confidentiality** does not only address the
- secrecy of user data when they are transferred,
 - ◆ but also *discretion* of the interactor regarding the user data which he received and may forward to others.

Integrity does not only address

- that any modification of communicated content (including the sender's name if one is provided) can be detected by the recipient(s),
- ◆ but also that each recipient is able to estimate the sender's *bona fides* concerning the content.

Availability does not only address

- that resources are available when the user wants to use them,
- ◆ but also the *willingness* of others to interact.

The differences between the properties of individuals and organisations as well as the discussion regarding protection goals have shown that a traditional IDM system, which fits the needs of individuals addressing organisations, is not necessarily sufficient for interaction between individuals. Although numerous cryptographic primitives and building blocks help to fulfil explicit requirements of interactions on the technical level, interactors may misbehave concerning the implicit requirements. Legal enforceability helps to ensure an interactor behaving as explicitly agreed beforehand. This holds especially for professional interaction and interaction between an individual and an organisation. But many interactions between individuals may be more informal, or it may be too expensive to enforce liability.

3 Building Blocks

We have seen that in communities, individuals have specific requirements for IDM based on the necessity of managing potentially highly dynamic, multilateral interactions between individuals. Therefore, building blocks have been defined that help to find interactors who behave as expected according to the requirements (among others discretion, *bona fides* and *willingness*) or to adapt the expectations of others' behaviour according to one's knowledge about them.

In this section, we join the building blocks defined to design traditional user-controlled IDM as introduced in [Hansen et al. 2004] and the building blocks for community-supporting IDM described in [Borcea-Pfzmann et al. 2006].

3.1 Pseudonyms and Partial Identities

In Section 1 we have already introduced partial identities (pIDs) and pseudonyms. Mechanisms of making (creation of new and change of existing pID) and taking (activating a pID by using it within an interaction) pIDs are important for identity management. Obviously this is true both for the traditional approach of user-controlled IDM as well as for community-supporting IDM:

Within the context of a user-controlled IDM system, pseudonyms and pIDs are used to support control of privacy by the individual. In accordance with the context and objectives of an interaction, the individual should be able to decide, whether, to whom and for which purpose he wants to disclose his personal data. Appropriate pseudonyms have to be chosen in order to control linkability of personal data disclosed as much as possible by the IDM system.

Thereby, linkability may be understood as the most important property of pseudonyms under which an individual interacts with others. This encompasses the following characteristics:

- ◆ *anonymity*: the link between pseudonym and its holder is not known, the use of the pseudonym is limited to one transaction only, and data disclosed in different transactions are not linkable by the pseudonyms;
- ◆ different flavours of *pseudonyms in various contexts*: a pseudonym is (re-) used for interactions within a certain context, e.g., depending on the role of its holder or the relationship to the interaction partner;
- ◆ *identifiability*: the link of the pseudonym to its holder is known or can be easily established by others; the pseudonym is used as substitute for the civil identity of the holder.

3.2 Selection of Interaction Parties

As already mentioned, in case of traditional user-controlled IDM scenarios, where an individual interacts with an organisation, the establishment of the interaction environment is less of an issue. In contrast to community-oriented interactions, interactions that address an organisation typically

are clear and predetermined with respect to fixed workflows. This aspect allows the organisation to offer a technical platform which processes a large part of the communication with the individual, e.g., by offering web forms to be filled.

Finding and selecting entities for the interaction is done by an explicit choice of one interaction partners. Often organisations advertise themselves publicly.

Since collaboration and communication in a community represent direct interactions with other individuals, *community selection and formation*, i.e., finding potential communication and collaboration partners, is more complicated – in particular if it should be combined with mechanisms for user-controlled linkability. Looking at processes of the physical world, mechanisms such as reputation management, advertising etc. help. Thus, the decision to enter a community is mostly made by getting to know it through advertisements or rumours from friends (or from a friend of a friend, and so forth) which would correspond to the reputation concept and which might be driven by social relationships forming social networks.

3.3 Contexts and History

Differentiation and interpretation of contexts are used to support the individual in decision processes when selecting a pseudonym for interaction. Together with a meaningful history representation, context management aims at a usable presentation of data flows to the individual. According to [Hansen et al. 2004], history information should include extent, nature and linkability of data released in the past.

Looking at the traditional approach of interactions, it is obvious that there is a need for a rather coarse-grained context differentiation which might be geared to the particular service an individual is making use of.

Since communities themselves do not strive for just one or a few particular objectives – they rather combine many different sub-scenarios within one environment – their context differentiation usually is more fine-grained. Therefore, individuals have to increase their awareness of the need to differentiate also their privacy requirements. This can be achieved by partitioning personal data which are disclosed within different contexts. Consequently, the system has to be aware of the distinct contexts the

individuals work in and should support managing the different contexts as well as assist in selecting the appropriate pIDs [Borcea et al. 2005].

3.4 Awareness

Awareness information represents additional data sets that foster *transparency* (in the meaning of clear visibility) within the interaction environment. It is fundamental for motivating community members in participating and for efficient working. Awareness information may influence privacy-relevant decisions (e.g., consenting disclosure of personal data or configuring privacy settings). However, awareness information may be privacy-sensitive itself, so individuals may want to restrict disclosure if the information is related to themselves. In literature, there are quite a lot of approaches classifying awareness. Examples are:

- ◆ *Group awareness*: Own and other individuals' settings as well as detailed information about the configuration and the recent history of actions within the community and its members.
- ◆ *Privacy awareness*: A person's awareness about his privacy, i.e. date of transfer of personal data, which personal data was concerned, how the data was processed.
- ◆ *Informal awareness*: Implicit information which increases the feeling of being part of a real (and not an artificial) environment [Dourish/Bly 1992].
- ◆ *Context awareness*: Contextual information describing the environment the individuals are working in, e.g., place, time, utilities available, etc.

Concerning this building block, a comparison between traditional and community-supporting IDM reveals that awareness information needed for the traditional approach is limited to only little information. In these cases, the interaction environment is relatively stable and the entities interacting with each other usually do not change during a transaction. The only awareness information, which might additionally be helpful for the individuals, is technical information, e.g., the current system state.

In contrast to these traditional approaches to IDM, awareness plays a big role in communities where in particular group awareness is the most important type. It allows individuals to assess and evaluate the highly dynamic interaction environment and

to adapt to the particular situation by adjusting their behaviour – including selection of appropriate pIDs.

3.5 Access Control

A large number of online applications require *authentication and authorisation mechanisms*. These mechanisms are mostly derived from the well-known ACL (access control list) or role-based access control approaches. While in the basic ACL approach permissions to operations on objects are listed together with the indication of the according pseudonym of the authorised person, the role-based mechanism lists permissions to operations on objects together with the according roles.

As long as we consider typical interactions in a 'customer/citizen ↔ organisation' constellation and pID switching while using a specific service is not of interest for the individual, the above indicated access control approaches may be applied. However, since user-controlled IDM systems in multilateral scenarios allow for dynamically switching pIDs depending on the actual context as well as for diverse role interpretations and kinds of usage, basic ACLs and role-based access control are not suitable here. A reasonable alternative for this case may be a mechanism which is inspired by capabilities. In order to avoid linkability of different pIDs an individual employs, anonymous credentials (certified properties) can be used instead of capabilities [Franz et al. 2006]. This way, access to all kinds of objects in the environment can be controlled independently of the organisational structure at community level.

3.6 Policies – Negotiation and Enforcement

Traditionally, user-controlled IDM systems imply bilateral scenarios where the indication of strict policies on, e.g., which personal data to disclose to whom or which security mechanisms to apply for securing the transaction, allows for quite straightforward *negotiations* between two entities (or usually a client-server pair). In contrast to this traditional approach, in multilateral scenarios the individuals' personal requirements may diverge very much. Even if the individuals determine specific policies – if those policies are conditional w.r.t. the behaviours of others, the negotiation proc-

esses may become very complex and hinder the actual work.

Whereas *enforcement* of policies is important in interactions with organisations and is usually regulated by law, informal discussions among individuals hardly are legally bound, neither inside nor outside the community. Instead, conceptual designs and implementations of negotiation and enforcement mechanisms for complex scenarios are open research questions.

3.7 Trust Management

Trust is an essential aspect of interactions. Based on trust the behaviour of others in interactions can be estimated. Trust usually is built up on information about and from the interactors distributed, e.g., by the following means:

Successful organisations typically build up reputation and benefit from word-of-mouth advertising that induces certain trust among the customers. Further, the organisations may advertise audits they received and quality standards that they comply with. In return, in order to assess a customer, an organisation may rely on centralised databases indicating the customer's creditworthiness.

The best known means for trust management is the use of *reputation systems*. Reputation systems manage information about past behaviour of interaction partners. Based on this information, individuals can get a clue how others might interact in the future. Reputation systems do not make expensive accountability measures obsolete (like, e.g., digital signatures under agreements made), but aim to reduce the cases where expensive legal enforceability using these measures might become necessary.

Unfortunately, reputation systems raise further privacy aspects since reputation means to give away some privacy against the benefit of linking reputation to a pseudonymous peer. Privacy-enhancing measures that help to prevent this are outlined in [Steinbrecher 2006]: Each peer uses several pseudonyms in parallel and only for a limited time. To maintain the same level of

reputation, the transfer of reputation between pseudonyms is needed.

3.8 Workflows

By adapting the organisation of personal lives to new technical possibilities, individuals more and more relocate administrative tasks (e.g., time management or arrangements) to the technical level. Whereas *workflows* between organisations and individuals are quite specified, the formalisation of workflows in personal lives is much more complex and highly dynamic. The assisting system must consider specific privacy and IDM concerns of its users. Therefore, generic building blocks of workflows should be offered. When considering privacy aspects, workflows should not allow for recognition of an individual in case he re-uses the same building block of workflows using different pIDs. Therefore, it might be reasonable to allow export of building blocks to other individuals to increase the anonymity set. Of course, in this case the building blocks have to be accordingly sanitised to avoid linkability options. Users need a framework for composition and modification to be able to individually construct their own workflows from the introduced building blocks. Furthermore, the building blocks should be designed to support these operations.

4 Conclusion

How to build user-controlled identity management systems for 'customer ↔ business' and 'citizen ↔ administration' applications is widely developed. However, these systems do not cover needs from users interacting within communities. Thus, supporting privacy-enhancing IDM within communities deserves attention, research and development.

Acknowledgements

We thank Stefan Berthold and Stefanie Pötzsch for their valuable comments and fruitful discussions to improve this text.

References

- [Borcea et al. 2005] K. Borcea, H. Donker, E. Franz, K. Liesebach, A. Pfitzmann, H. Wahrig: Intra-Application Partitioning of Personal Data, *Proceedings of Workshop on Privacy-Enhanced Personalization (PEP 2005)*, Edinburgh, UK, 2005, <http://www.isr.uci.edu/pep05/papers/borcea-pep.pdf>.
- [Borcea-Pfitzmann et al. 2006] K. Borcea-Pfitzmann, M. Hansen, K. Liesebach, A. Pfitzmann, S. Steinbrecher: What user-controlled identity management should learn from communities, *Information Security Technical Report*, Vol. 11, No. 3, Elsevier, 2006, pp. 119-128.
- [Dourish/Bly 1992] P. Dourish, S. Bly: Portholes: Supporting Awareness in a distributed work group, *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*, 1992, pp. 541-547, <http://www.ics.uci.edu/~jpd/publications/1992/chi92-portholes.pdf>.
- [Franz et al. 2006] E. Franz, A. Böttcher, H. Wahrig, K. Borcea-Pfitzmann: Access Control in a Privacy-Aware eLearning Environment, *Proceedings of AREs 2006, Workshop on Security in eLearning (SEL)*, Vienna, April 2006.
- [Hansen et al. 2004] M. Hansen, P. Berlich, J. Camenisch, S. Clauß, A. Pfitzmann and M. Waidner: Privacy-enhancing identity management, *Information Security Technical Report*, Volume 9, Issue 1, January-March 2004, pp. 35-44.
- [Pfitzmann/Hansen 2007] A. Pfitzmann, M. Hansen: Anonymity, Unlinkability, Unobservability, Pseudonymity, and Identity Management – A Consolidated Proposal for Terminology, Working Paper v0.29. July 2007, http://dud.inf.tu-dresden.de/Anon_Terminology.shtml.
- [PRIME 2007] PRIME Project – R. Leenes, J. Schallaböck, M. Hansen (Eds.): PRIME White Paper V2, June 2007, https://www.prime-project.eu/prime_products/whitepaper/.
- [Steinbrecher 2006] S. Steinbrecher: Design options for privacy-respecting reputation systems within centralized Internet communities, *Proceedings of 21st IFIP International Information Security Conference „Security and Privacy in Dynamic Environments”*, May 22-24, 2006, Karlstad, Sweden, IFIP 201, Springer, 2006.