Usage Control Requirements in Mobile and Ubiquitous Computing Applications

Manuel Hilty[1], Alexander Pretschner[1], Christian Schaefer[2], Thomas Walter[2]
[1]Information Security Group, ETH Zurich, 8092 Zurich, Switzerland
{manuel.hilty, alexander.pretschner}@inf.ethz.ch
[2]DoCoMo Euro-Labs, 80687 Munich, Germany
{schaefer, walter}@docomolab-euro.com

Abstract— **Usage control is concerned with control over data after its release to third parties, and includes requirements such as "this data must be deleted after 30 days" and "this data may be used for statistical purposes only". With the ultimate goal of respective policy languages that go beyond current technology for DRM of intellectual property, we provide (1) a catalog of requirements that are specific to the domains of mobile and ubiquitous computing applications, and (2) a classification of mobile and ubiquitous computing application scenarios that involve usage control, including a description of different kinds of involved data.**

Keywords- Usage Control, Policy Languages, Mobile Computing, Ubiquitous Computing, Requirements.

I. INTRODUCTION

The emergence of mobile and ubiquitous computing drastically increases the amount of personal and other sensitive data that is stored and processed by computers. Controlling access to and usage of sensitive data is important for the acceptance of such upcoming technologies.

Access control [19] is concerned with the problem of granting access to data on the grounds of information available at the very moment of deciding whether or not to release the data. The respective conditions include the requester's membership in a certain role or certain clearance attributes. Usage control [5], [22], [23] generalizes this notion to what must or must not happen with the data in the future. In this vein, usage control is concerned with requirements such as "this data must be deleted within 30 days" or "this data must not be further distributed". In general, requirements of this kind can neither be enforced, nor can their fulfillment be observed. However, with an ever-increasing amount of (personal) data, the specification, negotiation and enforcement of usage control requirements becomes a serious challenge. In general, usage control seems particularly relevant in the contexts of privacy, of intellectual property rights (digital rights management, DRM), and of military data.

For economic reasons, there are lots of activities in the area of DRM for intellectual property, including mechanisms [1], [27]; policy languages [11], [31], [30]; architectures [15], [16], [21], [9]; and interoperability frameworks [8]. As it turns out, mobile and ubiquitous computing applications are also concerned with data that is not directly content-related and yet privacy sensitive, while the—deliberately incompletely—cited literature exclusively focuses on DRM for content. In particular, dedicated policy languages for the specification of usage control requirements in the domains of mobile and ubiquitous computing applications *that embrace both DRM for intellectual property and privacy issues* do not exist. A first step towards such languages obviously is the understanding of the respective requirements. This paper provides a catalog of these requirements, the result of a joint project between the authors' institutions. Requirements were elicited by conducting structured interviews and analyzing project reports.

Contribution. We are not aware of any catalogs of usage control requirements in mobile and ubiquitous computing applications, and neither do we know of any general description of scenarios that involve usage control requirements particular to the described domains. We see our non-comprehensive requirements catalog and the scenario classification as a possible basis for any work in usage control policy languages—for mobile and ubiquitous computing applications—that extend beyond DRM, and also for respective system architectures.

Overview. The remainder of this paper is structured as follows. Section II sets the scene by describing a simple system model in which usage control is to be exercised. Section III sketches the setup of our study, describes the encountered requirements, reports on different kinds of data relevant for usage control, and provides a general schema for scenarios that involve usage control. Related work is discussed in Section IV. Section V presents our conclusions.

II. ABSTRACT SYSTEM MODEL

Throughout this paper, the concepts of usage control will be applied to the following simple system model. Data providers handle requests issued by data consumers. Data subjects are those subjects that the data in question is associated with. When access is granted by the provider, a copy of the requested data object is sent to the consumer, together with certain conditions—called obligations—on the future usage of this data object. These conditions contain restrictions on storage, distribution, aggregation and processing of the data, possibly defined by explicitly allowed purposes. Over time, the roles of data providers and data consumers may change. For instance, a

data subject equipped with a mobile device may send GPS data to a network operator. In this case, the data subject is the data provider and the network operator is the data consumer. A location-based service provider (a data consumer) may then ask for this particular data; this makes the network operator a data provider.

Data providers can specify usage control policies that govern access to and usage of their data. For instance, the owner of the mobile device in the above example can specify that the location data the mobile device gives to the network operator may only be forwarded to certain service providers, or should only be stored for a certain amount of time.

Obligations immediately raise the question of enforcement. Sometimes they can be enforced by relying on DRM technology or trusted platform mechanisms at the data consumer's site. Sometimes they cannot, but in some special cases it might be possible to at least observe the violation of obligations, e.g. by requesting audits or consulting trusted logging mechanisms. In these cases, compensating actions such as decreasing some trustworthiness attributes or taking legal actions can be taken [4]. This is similar to regular law enforcement: pedestrians cannot in general be prevented from jaywalking, but they can be penalized when caught. Finally, some obligations cannot be rendered observable, in which case the data provider has to trust the consumer.

III. Requirements Analysis

A. Methodology

The elicitation of usage control requirements was done by taking into account various information sources. In addition to surveying the relevant literature, we conducted structured interviews with employees of DoCoMo Euro-Labs, and analyzed mobile and ubiquitous applications being developed in the projects MOSQUITO [18], Remember&Find [8] and MobiLife [17] (cf. the following sub-sections) and some other applications (Section III.B.4).

B. Mobile and Ubiquitous Computing Applications

We studied articles and project documentations and conducted interviews with employees of the involved project partners. The following projects were the main sources of input.

1) MOSQUITO

The MOSQUITO project (Mobile Workers' Secure Business Applications in Ubiquitous Environments) [18], a part of the European Union's 6th framework programme, is concerned with developing a security framework for business applications running on top of mobile and ubiquitous computing systems. One focus is on context-based access control, where information gathered from context sensors is used for making access control decisions.

Three business scenarios are defined for MOSQUITO: "Fair Deal" features mobile and ad hoc IT applications between employees from different companies; "Healthcare Emergency" features the adaptation of electronic health applications and of their supporting monitoring and communicating devices, in particular in emergency situations and ad hoc scenarios. "Paperless Car" contains the integrated support of vehicle administrative tasks with public administrations, insurance companies or car manufacturers through mobile and ad hoc communication.

The following usage control aspects were found within these scenarios.

- In "Fair Deal", employees of different companies may exchange business information. Before doing so, non-disclosure agreements (NDAs) must be signed. While MOSQUITO is not concerned with the technical enforcement of such NDAs, one could imagine supporting the deletion or non-distribution of exchanged blueprints by technical means.
- In "Healthcare Emergency", health information on a person's mobile healthcare system can be accessed by doctors (for evaluating the person's health situation) and by pharmacies (for verifying subscriptions). It could be desirable to include policies that tell a pharmacy to delete the received information after a certain time.
- In "Paperless Car", black boxes built into cars might store information that can be used for accident analysis later, and even communicate with other vehicles and infra-structure. Many usage control problems are foreseeable in this context, for example what insurances may or may not do with information gained about the driver's driving habits.

2) Remember&Find

The Remember&Find project [8] develops mechanisms for finding personal items that are equipped with sensor technology such as RFID or Bluetooth. Such personal items can be registered with one's mobile phone, which remembers when it sensed the item the last time. If the owner does not find the item anymore, the mobile phone tells the last remembered location where the item can then be searched with the respective sensor technology. Searches can also be delegated to other phones in the suspected surrounding of the item. The following usage-control problems arise in this application.

- When delegating a search to others, the user tells the network that he is looking for the respective item. However, he might require that the network does not convey his identity to the search helpers, or maybe even that the search helpers delete the item ID again afterwards.
- It is foreseeable that a search request involves a description of the item, not an ID. Here, the user might not want the network to use the gained information to create user profiles, or give that information to other parties.

3) MobiLife

The MobiLife (Mobile Life) project [17], an Integrated Project the European Union's 6th framework programme, is to bring advances in mobile applications and services within the reach of users in their everyday life. This is achieved by innovating and deploying new applications and services based on the evolving capabilities of the 3G systems and beyond. The project addresses, with a strong user-centric view, problems related to different end-user devices, available communication

networks, interaction modes, applications and services. MobiLife, in particular, deals also with context information where context is the situation or state of an entity (i.e. person, place, physical or computational object), and context information is understood as characterizes context – or a part of it – of said entity.

Context information can come from various sources: location sensors (e.g., GPS receivers or network cells), calendars (e.g., which may indicate whether the user is in a meeting), sensors that recognize other nearby devices, clocks, and many other sensors such as noise or temperature sensors. Much of this context information is produced by the user's own mobile phone. When giving such information to a server-based service (e.g., a tracking system of a company site or a map provider), the user might want to specify what the data may be used for, how long it may be retained, or to whom it may be given at which level of detail. One important characteristic of context data, specifically location data, is that it can be used and displayed at different levels of detail. For example, a user's location can be represented in exact coordinates, by the city he is in, or only the country.

Context information is sometimes processed by third parties (e.g., location-based service providers). In this case, the network operator has to make sure that the third party does not use the data for purposes other than specified, and that it does not retain the information longer than envisioned. This is a classical usage control scenario [23].

4) Issues in Other Applications

Through the conducted interviews and the examined documents, we also collected information about mobile and ubiquitous applications other than the ones mentioned above. The following aspects of usage control were found in these areas.

- In the newest proposal of the session instantiation protocol (SIP) for video conferences [24], there are options that require the other party to neither publish nor record the video feed. The other party must accept these requirements before receiving the feed.
- For data that is neither payload data nor data gathered by the user's own sensors, the user cannot directly specify his usage control requirements (unless the sensors are equipped with special policy negotiation capabilities). In these cases, the requirements may come from subscription contracts, laws, or self-regulation of the data provider (e.g. a mobile network operator). The latter is especially important as the public can be very sensitive with regard to the use of personal information, and therefore the data providers can prevent the materialization of threats to the company's reputation by treating personal data carefully.
- In the future, trusted platform technology might increasingly be included in mobile phones. This could be leveraged for obligation enforcement.
- In roaming scenarios, the usage control aspects can become quite complicated because more actors are involved. Delegation of rights then becomes an important issue.
- Usability is a special concern in connection with usage control. By introducing usage control policies, more requirements have to be specified. And if it is the user's task to specify these requirements, guaranteeing usability might become difficult.
- It is also not clear who pays for usage control mechanisms. This is an important question to keep in mind, but is outside of the scope of this study.

C. Analysis

As a consequence, in the context of mobile and ubiquitous computing applications, there appear to be at least six different scenarios that involve different instantiations of usage control as schematically depicted in Figure 1. These do not include usage control scenarios involving customer contact or billing data because these scenarios are not particular to the domain of mobile and ubiquitous computing applications.

In this figure, the arrows represent data flow, and the other elements represent sensors, active tags, mobile devices, operators, and service providers. In the following it is assumed that some data subject S wants to control usage of its own personal data. "My sensor" is a sensor (GPS, noise, vital parameters, gyroscope) under control of S. Active tags include signaling devices such as RFIDs. "My mobile" is a mobile device under control of S. "My operator" is the network operator that S has a contract with, and "some operator" is any other operator (involved in roaming scenarios). "Some service provider" is any service provider including location-based services not under governance of the provider, profiling services under governance of the provider, or the provider's marketing department.

1. This scenario describes a situation where sensor data is collected by the mobile device of S. The sensor may or may not be part of the device itself but is assumed to be under control of S, e.g. a device taking vital parameters. The collected data is forwarded to the operator who forwards it to some service provider (note that the data may also be sent directly to some service provider). There is a need to define usage control requirements in the mobile device or, by means of default policies, at the operator's site.
2. This scenario describes a situation where an active element such as an RFID (which is worn by S) sends data (which hence relates to S) to a reading device that is not under control of S (and also includes the symmetrical case where sensor data from a "foreign" subject is collected by a reading device under control of S). In this case, the reading device and the active tag are physically separated. Usage control requirements might be defined at the active tag's side (leading to smart active tags equipped with policy negotiation engines), at the reading device, and at the operator's side.
3. This scenario describes the situation where data is "created" at the operator's site. Among other things, this is the case for connection data and cell location information. Usage control requirements can be formulated at the mobile device's site which wants to restrict the use of any

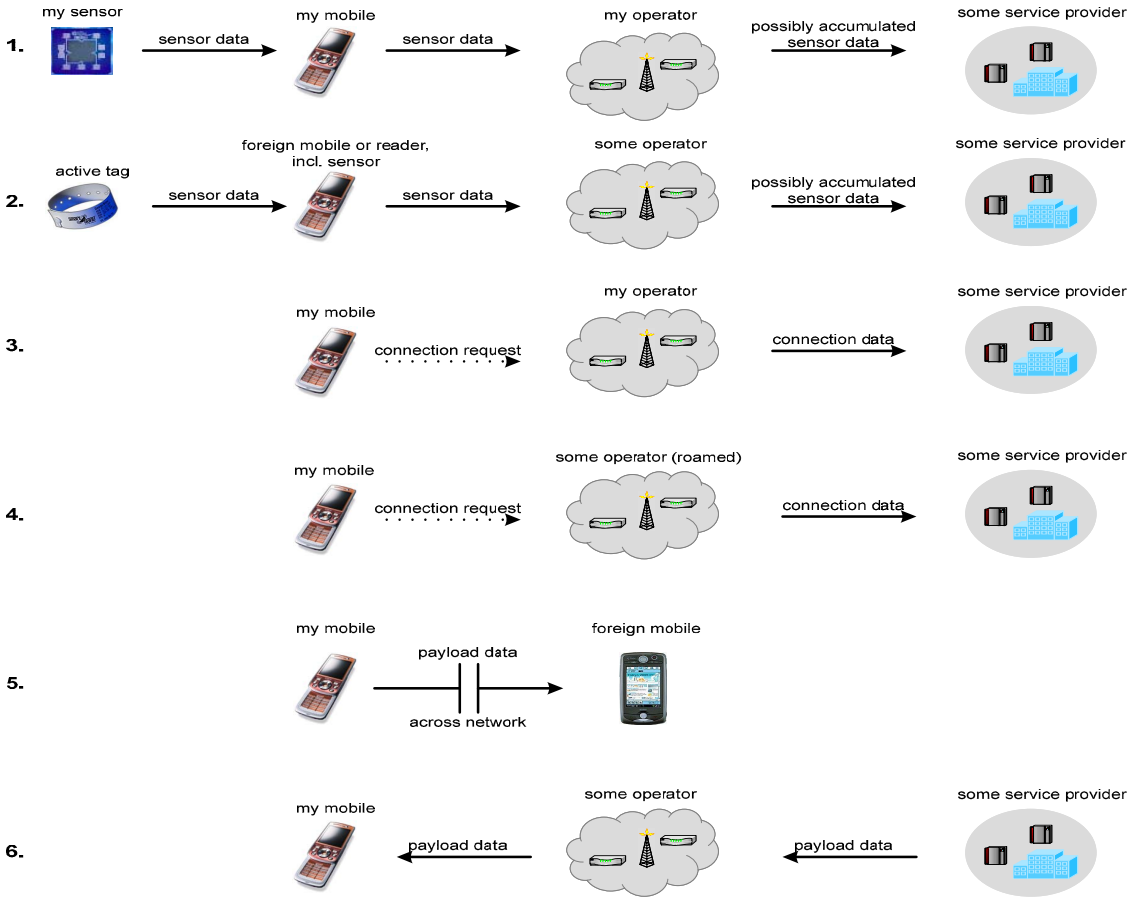

Figure 1: Data flows with need for usage control

data that is created at the operator's site, and also at the site of the operator when it releases data to a service provider. (Note that scenarios 2 and 3 share some obvious commonalities.)

4. This scenario is similar to the third, but roaming is involved. Connection data is created and handled by a network operator that has no direct contracts with S.
5. This scenario describes the situation where payload data such as video or audio streams are passed from end user to end user (sensor data can also be considered to be payload data). The above example of video data not to be video projected is an example of the kind of usage control requirements that might have to be formulated at the site of the mobile device.
6. This scenario describes a classical DRM scenario where payload data such as audio files are transferred from a service provider to the mobile device. Usage control policies can be formulated both by the service provider and the operator.

We observe that the type of data and actors involved in a scenario can make a big difference. In the following, we show the requirements derived from the scenarios.

1) Policy language issues

Some requirements encountered in this study are typical usage control requirements that also occur in many other areas. For example, all requirements within the MOSQUITO project are of this type. However, we noticed a few interesting aspects of usage control in a mobile or ubiquitous setting that are not of the same relevance in non-mobile scenarios. Particularly, they impose some constraints on the expressiveness of policy languages, usability, and available hardware as further exemplified hereafter:

- Often, the origin of data matters. For example, if location data comes from a GPS receiver in the user's mobile phone, the user can theoretically decide about each request for accessing this data, and impose requirements about its usage. When the location data comes from the network operator who detects in which network cell the phone currently is, or from a ubiquitous gateway sensing the phone, then the data is collected by the data provider itself. In this case, the user does not have direct control over how much of this data is gathered, and may not even be aware of it. Therefore, there must be policies specifying the rules in connection with this data.
- If the data comes directly from the user, there must be a possibility for the user to specify preferences for how the network operator must treat the data. A language for

specifying such preferences must be easy to use to satisfy the usability requirements mentioned above. To achieve this, the relevant requirements that must be expressed must be studied carefully in order to tailor the expressiveness of the language.

- When data that is subject to usage control requirements is sent from the network operator to a single mobile phone (e.g., search requests in Remember&Find, or multimedia data from a subscription service), then there exist possibilities of controlling the use of the data through trusted modules included in the phones (e.g., [1]). An interesting problem here is how server-side usage control policies can be mapped to policies or rights objects for these trusted platforms.

It seems inevitable to cater for the consistent interplay of different policies at different sites. For instance, it must be specified how the rights that are granted to the operator by means of establishing respective policies using the mobile device carry over to the service provider.

2) Policy enforcement at multiple sites

While the simple system model of Section II suggests that usage control policies are defined at the data provider's site and possibly enforced at the consumer's site, this scenario breakdown shows that usage control in the context of mobile and ubiquitous computing applications is, in reality, much more complex, simply because the roles of the involved entities may change.

Policies must be defined and enforced at multiple sites, including devices with special usability constraints: Sensors; Mobiles; Network operators; and Service providers.

The issue of where the data originates is even more complicated in roaming scenarios. The roaming scenario evidently poses particular problems: rights granted to the operator that S has contracted must be consistently propagated to respective other operators. To understand this problem in its full complexity, the legal situation has to be carefully examined as well.

3) Support for roles

Roles of the involved entities may change. For example a mobile phone may be a data provider to his mobile network operator but at the same time a data consumer of sensor data.

4) Support for different types of data

Data bound to usage control requirements can be more than context information as mentioned in many of the above examples. The types of data encountered are:

- payload data such as audio or video streams that are transferred from one end user to another or from a service provider to an end user (e.g. MP3 files);
- sensor data such as GPS coordinates that is collected by the owner's mobile device and sent to the operator who, in turn, might transfer it to context-based service providers;
- connection data that is collected at the network operator's site and that is used for billing and may be used for profiling purposes;
- sensor data collected by parties other than the signal emitter's owner such as RFID transponder signals; and
- customer contact and billing information.

5) Support for the following restrictions

Except for the first kind, the above data occurs in "raw" and "aggregated" forms, e.g. profiles. Usage control requirements in this domain boil down to

- restrictions on the use of data that a user provides to the network (e.g., context data);
- restrictions on the use of data that a service provider provides to the user (e.g., multimedia data);
- restrictions on the use of data that a user provides to another user (e.g., in video conferences); or
- restrictions on the use of data that a network operator provides to service providers.

IV. Related Work

Usage control is an extension of access control [19]. In earlier work [13], we defined some key concepts in the area of usage control policies like provisions and obligations. A subsequent article [23] provides an overview of specification and enforcement aspects in usage control.

Much work has been done in defining policy languages for specific sub-domains of usage control such as privacy policies or digital rights management. Examples for privacy policy languages include P3P [29], which can be used to specify the privacy practices of web sites, and EPAL [3], which is more general and precise than P3P. XrML [31] and MPEG-REL [30] are XML-based policy languages for DRM applications that also cover ongoing usage. ODRL is an XML-based rights definition language that can also be used for mobile computing applications. There also exists a privacy profile of XACML [20]. All these languages focus on either DRM or privacy but not both. Rei [14] seems general enough to cover both domains but, on the downside, does not provide explicit abstractions for these domains, and for the different types of data.

UCON [22] is an access control extension that adds the notion of ongoing usage, which is the usage of digital data during an access. UCON assumes data providers and data consumers reside on the same machine which facilitates control. Enforcement mechanisms in the area of digital rights management include trusted computing [25] and software-based mechanisms as surveyed in [28]. Ubiquitous computing applications are not taken into account.

V. Conclusions

Since ubiquitous computing scenarios are not a reality yet, usage control requirements are of a necessarily speculative nature. However, among many others, a recent special issue of the Communications of the ACM (September 2005) on RFID technology suggests that adequately dealing with privacy concerns may well turn out to be an enabling factor of ubiquitous computing technology. In terms of mobile applications, personal data including connection, billing and

contact data are of utmost sensitivity even today, as related recent incidents with unintentional disclosure of such data impressively convey [26]. In a similar vein, the need for DRM technology [12] as expressed by digital data providers (movies, music) is an instance of the more general problem of usage control.

In terms of general usage control, there is a wide variety of open research problems. Policy languages need to be defined that allow the expression of access control requirements, obligations and enforcement strategies that cater for both DRM and privacy. We are well aware that our catalog of requirements is but a first step. This includes the difficult problem of making precise the notion of allowed purposes and of describing which enforcement mechanisms can be used for which kind of requirements. Further, more general, server-side and client-side architectures for enforcing usage control requirements on different devices remain to be conceived and implemented.

These research problems do not exclusively relate to mobile and ubiquitous computing applications. However, all of them appear relevant for a telecommunication operator once any kind of data has been stored at the operator's databases. As mentioned in the end of Section III.C, usage control policies for sensors and reading devices as well as the consistency of policies in sensors, mobile devices, the operator's data bases and the service providers seems to be a problem that is especially relevant in the domain of mobile/ubiquitous computing. Once these problems have been solved, the problem of usage control in roaming scenarios needs to be addressed. Finally, research into all kinds of enforcement mechanisms in mobile devices seems to be justified in the context of usage control.

All these problems are very hard problems. There exist client-side DRM mechanisms that can be embedded into mobile phones, as for example the architectures described in [2], [16]. However, such mechanisms are usually tailored to enforce one specific type of requirement. In the example presented in, [2], this is the payment for listening to a song. Based on our current research on policy languages, we would like to find out what kind of obligations can possibly be enforced with such a conceptual DRM architecture and how the respective obligations can be specified in the corresponding rights objects.

REFERENCES

[1] 4C Entity, LLC. Content-Protection for Recordable Media Specification—SD Memory Card Book, SD-Video Part. Revision 0.95, 2005

[2] H. Aono, R. Hoshino, and S. Hongo. Tamper-Resistant Charging Technology for a Seamless Environment. NTT DoCoMo Technical Journal 6(2):31-37, 2004.

[3] P. Ashley, S. Hada, G. Karjoth, C. Powers, and M. Schunter. Enterprise Privacy Authorization Language (EPAL). Research Report 3485, IBM Research, 2003.

[4] M. Backes, B. Pfitzmann, and M. Schunter. A toolkit for managing enterprise privacy policies. In Proc. ESORICS, 2003.

[5] C. Bettini, S. Jajodia, X. S. Wang, and D. Wijesekera. Provisions and obligations in policy rule management. J. Network and System Mgmt., 11(3):351–372, 2003.

[6] European Union. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of per-sonal data and on the free movement of such data, Official Journal L 281, 31/11/1995 P. 0031 – 0050.

[7] European Union. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of pri-vacy in the electronic communications sector, Official Journal L 201, 31/7/2002 P. 0037 – 0047.

[8] Coral Consortium. Interoperable Media&Home Networks, 2006. Available from www.coral-interop.org

[9] G. Fernando, T. Jacobs, V. Swaminathan. Project DreaM—An Architectural Overview, 2005. Available from openmediacommons.org

[10] C. Frank, C. Roduner, C. Noda, M. Sgroi, W. Kellerer. Interfacing the Real World with Ubiquitous Gateways. European Workshop on Wireless Sensor Networks – Poster Session, Zurich, 2006.

[11] R. Iannella (ed.). Open Digital Rights Language version 1.0. Available from oderl.net

[12] J.Irwin. Digital Rights Management: The Open Mobile Alliance DRM Specifications. In Information Security – Technical Report, Elsevier, Vol. 9 No. 4, 2004.

[13] M. Hilty, D. Basin, A. Pretschner. On Obligations. In Proc. ESORICS 2005.

[14] L. Kagal, T. Finin, A. Joshi. A Policy Language for a Pervasive Computing Environment. Proc. 4th IEEE Intl. Workshop on Policies for Distributed Systems and Networks, 2003.

[15] Marlin Developer Community. Marlin architecture overview, 2006. Available from www..marlin-community.com

[16] Melodeo Inc. PachyDRM Digital Rights Management—Technical Overview v3.1, 2005. Avalaible from www.pachydrm.org

[17] MobiLife consortium. Mobile Life. http://www.ist-mobilife.org/.

[18] MOSQUITO consortium. Mobile Workers' Secure Business Applications in Ubiquitous Environments. http://www.mosquito-online.org/.

[19] B. W. Lampson. Protection. Fifth Princeton Symposium on Information Science and Systems, Princeton, NJ, March 1971.

[20] OASIS. eXtensible Access Control Markup Language (XACML), 2005. v2.0.

[21] Open Mobile Alliance. DRM Architecture, 2006. Available from www.openmobilealliance.org

[22] J. Park and J. Sandhu. The UCON ABC Usage Control Model. ACM Transactions on Information and Systems Security, 7:128-174, 2004.

[23] A. Pretschner, M. Hilty, and D. Basin. Distributed Usage Control. Communications of the ACM, September 2006, to appear.

[24] R. Shacham, H. Schulzrinne, W. Kellerer and S. Thakolsri. Use of the SIP Preconditions Framework for Media Privacy. IETF Internet-Draft, draft-schacham-sip-media-privacy-01, November 11, 2005.

[25] S. W. Smith. Trusted Computing. Springer, 2005.

[26] B. Sullivan. Bad-news data letters put consumers to stray. http://www.msnbc.msn.com/id/9581522/, Oct. 2005; last accessed in Jan. 2006.

[27] SVP. SVP Open Content Protection System—Technical Overview, 2005. Available from www.svpalliance.org

[28] P. van Oorschot. Revisiting software protection. In IST'03, Springer Verlag, 2003.

[29] W3C. The Platform for Privacy Preferences 1.1 (P3P 1.1) Specification, Working Draft, 2005.

[30] X. Wang, T. DeMartini, B. Wragg, M. Paramasivam and C. Barlas. The MPEG-21 Rights Expression Language and Rights Data Dictionary. IEEE Transactions on Multimedia 7(3):408—417, June 2005

[31] X. Wang, G. Lao, T. DeMartini, H. Reddy, M. Nguyen, and E. Valenzuela. XrML – eXtensible rights Markup Language. In XMLSEC '02. ACM Press, 2002.