Towards E-Society Policy Interoperability

Renato Iannella

NICTA

Level 5, Axon Building #47

Staff House Rd, St Lucia, QLD, 4072, Australia

renato@nicta.com.au

Abstract. The move towards the Policy-Oriented Web is destined to provide support for policy expression and management in the core web layers. One of the most promising areas that can drive this new technology adoption is e-Society communities. With so much user-generated content being shared by these social networks, there is the real danger that the implicit sharing rules that communities have developed over time will be lost in translation in the new digital communities. This will lead to a corresponding loss in confidence in e-Society sites. The Policy-Oriented Web attempts to turn the implicit into the explicit with a common framework for policy language interoperability and awareness. This paper reports on the policy driving factors from the Social Networks experiences using real-world use cases and scenarios. In particular, the key functions of policy-awareness - for privacy, rights, and identity - will be the driving force that enables the e-Society to appreciate new interoperable policy regimes.

Keywords: E-Society, Policy-Oriented Web, Policy Languages, Social Networks, Interoperability, Privacy, Rights.

1 Introduction

The e-Society has been a long term dream that the ICT community, amongst others, have moved towards with new technologies over the past decade. The engagement of citizens in e-Societies has enabled greater participation and opportunities for communities to offer “information commons” [1] for digital interactions. Today, we clearly have this dream realised with Social Networks. Social Networks - via the innovative use of Web 2.0 features - have also taken the ICT community by surprise with such rapid uptake and widespread content sharing.

Social Networks attempt to mimic and support normal society interactions and experiences. In many cases, these seem to be working well, such as keeping friends and family in contact and sharing status information. However, the wide-spread sharing of personal and corporate information within Social Networks (eg photos, documents) have an impact on policy support, such as privacy and rights management decisions. These issues have now become more relevant as Social Networks have empowered the end user to share even more private content with increasing global reach. Additionally, the providers that offer these services have an immense database of personal information at their disposal.

Generally, Social Networks “provide complex and indeterminate mechanisms to specific privacy and other policies for protecting access to personal information, and allow information to be shared that typically would not follow social and professional norms” [2]. There have been numerous attempts to solve this problem in the past but none have been really successful, nor applicable to the Social Networks community. A new approach is required to manage seamless policy interaction for the e-Society masses. The “Policy-Oriented Web” is an emerging idea to bring greater policy management technologies to the core web infrastructure. This will enable policies to inter-operate across Social Network service providers.

In this paper we present e-Society use cases from Social Networks to highlight the driver for the adoption of new interoperable policy technologies. We then present an information model for the Policy-Oriented Web and show some example representations. Finally, we look at related works and conclude with how e-Society - via Social Networks - can lead to greater interoperability opportunities for policies across the wider Web.

2 E-Society Use Case: Social Networks

Social Networks, like FaceBook, Flickr, LinkedIn, Xing, YouTube, and MySpace, have been phenomenally successful. They have achieved this by providing simple yet user empowering features that digitally support the online social experience. In particular, the relative ease of sharing content with close colleagues and friends has driven Social Networks participation. However, this experience can have serious repercussions if the implicit arrangements under which content is shared are not known explicitly, or worse, are not respected.

Two recent examples have highlighted these negative experiences. The first was the use of photographs from Flickr in a commercial advertising program [4]. In this case, the image of a person was used by Virgin Mobile in billboard advertising. They had taken the image from Flickr as the photo owner (the person’s friend) had selected a Creative Commons license that allowed commercial usage. This highlighted two issues; understanding the implications of commercial usage, and publishing images of your friends on public websites. The photo owner had assumed that commercial usage may have enabled him to participate in the financial rewards (it didn’t). His friend who appeared in the photos also had no idea her image was being used, until it was too late (she was not impressed). The lack of understanding the requirement for “model release” permission in the license policy also contributed to this situation.

The second example involved photos from FaceBook being used by the mainstream media to report on the death of a defence force trooper [5]. The media had used his personal photos from his FaceBook profile - including photos of his family - to print in the national newspapers. At no time did they seek permission to reproduce these images. In some of the media responses to this issue, the assumption was stated that since the photos were on the Internet anyway, they were deemed “public domain” and you could basically do whatever you like with the images with little recourse.

Both of these cases involve sharing of photos on Social Networks and highlight challenges to owners and end users on the right level of respect for use of such content. To investigate this issue further, we looked at the processes for sharing photos on Facebook.

Like many Web 2.0 Social Networks, Facebook requires the account owner to certify - implicitly - that they have the right to distribute uploaded photos and that it does not violate the Facebook terms and conditions (see Figure 1). The latter is an eight page document of dense 'legal-ese' wording that not only is unlikely to be read by account owners, but rarely would be understood by the layperson.

Upload Photos – Billie

[Edit Photos](#) [Add More](#) [Organize](#) [Edit Info](#) [Delete](#) [Back to Album](#)

I certify that I have the right to distribute these photos and that they do not violate the [Terms of Use](#).

You can upload up to 60 photos per album.

Fig. 1. FaceBook Photo Upload Certification

Facebook allows you to add photos in named Albums to your account. You can then decide on who can see these photo albums with some simple image privacy controls. Figure 2 shows the options available including; Everyone, My Networks and Friends, Friends of Friends, and Only Friends.

Edit Photo Album Privacy

Who Can See This?

☐ Everyone on Facebook

Friends

☐ Friends of Friends
My friends and their friends can see this.

☒ Only Friends
Only friends can see this.

☐ Some Friends
Choose specific friends who can see this.

☐ Only Me
Only you and selected networks can see this.

Networks

Some of My Networks... [+](#)

☒ Australia

☐ Queensland

☐ HKU

Except These People

John Doe [+](#)

[Save](#) [Cancel](#)

Fig. 2. FaceBook Photo Options

Fig. 3. FaceBook Photo Privacy Options

However, when you choose the Customize options, additional detailed privacy controls are available (see Figure 3). Now you can be very specific, such as indicating which individual friends can see the photo (“Some Friends”) and who cannot (“Except These People”). You can also specify specific Networks of friends as well.

At this stage, when an end user - be they public, in your network, or a friend - sees your photo, they have the usual file manipulation controls in their web browser to “Save Image As” to the local disk (see Figure 4). Obviously the photo is now out of the reach of Facebook’s privacy control mechanisms and can now be forwarded to anyone via email and other means, or printed in national newspapers, or plastered on billboards. So the privacy controls that we had carefully crafted in Facebook are now no longer available outside the domain of this Social Network.

Clearly the reason for this overriding of the Facebook privacy policy is the fact the a standard Web Browser has no knowledge of the policy and any embedded image in a web page can “normally” be saved to local disk. If we could design an enhancement to Facebook - if not all Social Networks - then we would consider a simple mechanism that informs the end user that the photo has some restrictions attached. We don’t envisage an “enforcement” mechanism, as this would not be consistent with the ethos of Social Networks, but an “accountability” mechanism would be sufficient and appropriate. This would allow, for example, images to be cached by the browser (for efficiency) but not explicitly saved outside the browser environment.

Figure 5 shows a hypothetical dialog that could appear instead of the “Save Image As” dialog (as shown in Figure 4). The key point is that this dialog - albeit simple - informs the end user of the privacy rules pertaining to the image and allows them to honour this (i.e. to cancel the request) or to continue with the file download, but being warned that this may be recorded for accountability purposes.

The image in Figure 4 is a picture of my cat Billie, and she is not too concerned about her image being published on FaceBook. The issue becomes really compounded, as we have already seen, when sharing pictures of your family, friends and

Fig. 4. Web Browser “Same Image As...” Menu

Fig. 5. New Save Dialog

colleagues. Facebook includes a feature whereby you can annotate photos and indicate the names of the people in the photo. These can be existing Facebook members or non-members. The image would then show their names (with a mouse-over their face) and, for members, would link to their profile.

Figure 6 shows an example of selecting my colleagues faces in a photo and assigning them to their Facebook identity. (Note that the images and names have been deliberately blurred in Figure 6 to protect their privacy). As with the photo of Billie, this photo can also be downloaded and shared bypassing the Facebook privacy policy. It also poses greater threat as my two colleagues in the photo also do not wish their image to be used for any other purpose than a corporate image of the project team.

Fig. 6. FaceBook - Photo Friends Tagging

However, we do now have the new possibility of checking the individual's policy needs since we have identified all the people in the photo. We could automatically notify each of them and ask if they would allow their friend (ie me) to publish the photo of them in his photo album and under the privacy policy I have designated. For

example, I could allow the photo for complete public access, or limited to a network, or my friends. They could then respond based on this. This “policy negotiation” could also be automated to allow quicker responses, based on an individuals own privacy policy.

To summarise, Social Networks, like Facebook and others, have a tremendous opportunity now to look towards simple, yet powerful, policy support to match the community expectations when sharing content. The emerging Policy-Oriented Web can exploit these use cases as the driver to develop new web infrastructure. Future Web 2.0 services can be built upon this new web infrastructure to provide fair and accountable content sharing services.

3 The Policy-Oriented Web

The Policy-Oriented Web, also sometimes referred to as the Policy-Aware Web, is an emerging field that aims to address the need to manage multiple and conflicting policies in the future distributed service-oriented world. This will increase connectivity across disparate web systems and services as they can achieve a new level of automated interoperability, guided by declarative policies that can adapt to different contexts and environments. In an earlier position paper [6] we outlined the major key strategic challenges posed by the Policy-Oriented Web. This included the need for a unified model that can adequately represent policies. Such a unified model - based on various policy requirements - will capture the core concepts and structures common to all policies. The model should also provide the framework for addressing even deeper policy-specific challenges such as the evaluation, enforcement, and reasoning of policies, and how to deal with inconsistencies across policies.

We have developed a preliminary semantic information model based on the analysis of three types of existing policy languages; privacy, rights, and identity. Specifically, we analysed the P3P [8], ODRL [7] and XACML [14] languages and reviewed their information features, structures, and relationships to determine the commonalities across these policy languages. These three were chosen as they represent the most used languages for privacy (P3P), rights (ODRL) and access control (XACML). Each lacked the complete structure to be a general policy language on their own. For example, P3P lacks mechanisms to link to multiple parties, ODRL lacks negation, and XACML lacks inheritance.

The resultant Policy information model (shown in Figure 7) contains three primary classes that express the policy semantics:

- Action - these are the activities involved in the policy. The related Focus class indicates which aspects of the Action drive the policy, such as “Allow” or “Deny” or “Exclusive”.
- Resource - these are the resources/content involved in the policy. The related Target class indicates which aspects of the Resource are relevant to the policy, such as “One” or “Any” or “All”.
- Party - these are the people and organisations involved in the policy. The related Role class indicates which role the Party plays in relation to the policy, such as “Licensee” or “Consumer”.

These three classes were found to be the core components from the policy languages analysed. Supporting these three classes are the following classes:

- Act - identifies specific acts that can be performed.
- Object - identifies specific entities.
- Function - identifies comparative operators.

A Policy can also include two other classes that modify the behaviour of Actions and Parties:

- Constraint - conditions that will limit an Action of the policy. This can cover a range of options from the fundamental (such as numeric, date ranges, geospatial) to the more complex (such as a particular purpose or domain use).
- Obligation - requirements that must be met by a Party in order to satisfy the policy. This can also cover a wide range from the fundamental (such as payment) to the more complex (such as being tracked for usage).

Both the Constraint and Obligation classes are supported by the Operation class. The Operation class links instances of Act, Object, and Function classes to uniquely express the required operation. The Operation will enable reasoning services over policies as it will contain the fundamental data for policy expressions.

We have represented the Policy-Oriented Web Information Model (Figure 7) in RDF and RDF Schema (modeled using the Protégé tool). It was quite challenging in converting a typical information model (such as Figure 7) and mapping it into an RDF model. In some cases it was not clear how to best represent the information artifact, such as Focus and Deny, into the triple-based model of RDF.

Fig. 7. Policy Information Model

The example RDF Schema snippet below shows the “hasObject” property with “Object” range and domain of “Operation”, “Party”, and “Resource” classes.

```
<rdf:Property rdf:about="urn:policy:10:hasObject">
  <rdfs:range rdf:resource="urn:policy:10:Object"/>
  <rdfs:domain rdf:resource="urn:policy:10:Operation"/>
  <rdfs:domain rdf:resource="urn:policy:10:Party"/>
  <rdfs:domain rdf:resource="urn:policy:10:Resource"/>
</rdf:Property>
```

We have chosen just RDF and RDF Schema to keep the first iteration of the Policy-Oriented Web as “simple” as possible without over complicating the expression structures. This will be important to meet the technical needs of the Web 2.0 and Social Networks communities. We plan to also use RDFa as another encoding direct into HTML pages.

Others have proposed semantic expression of specific policy languages in the richer OWL language for rights [11, 12] and privacy [13] policies. We envisage the future where use of such advanced semantic web languages will also be supported in the Web 2.0 platform of technologies. In the first iteration of the Policy-Oriented Web, we believe that basing it on the RDF language is the best compromise. The Policy-Oriented Web information model could also be expressed in OWL for more advanced reasoning and ontological features needed by high-end communities. However, this may lead to an unnecessarily complex language and lessen the appeal to the wider communities. Nonetheless, there should be support to extend the language for those requiring these features and this will, at least, guarantee some level of interoperability.

4 Policy–Oriented Web For E-Society

If we now revisit some of the scenarios presented in Section 2, we can start to see how to apply the Policy-Oriented Web to the specific use cases for Social Networks. In particular, the two use cases of view control over photos, and publishing photos of your friends and colleagues.

4.1 View Rights

In this use case, we need to support the ability to define the scope of users who can view the photos in a photo album. Using Facebook, as an example, we need to support:

- all the public,
- all or some of your friends,
- all or some of your networks, and
- disallow one or more friends.

These would be expressed as Constraints in the Policy-Oriented Web model as part of a policy instance. The example below shows the RDF/XML snippet that would express that only people in your “Australia” network can view the photos.

```
<p:Constraint rdf:ID="view-aust-network">
  <p:hasAction rdf:resource="policy:render"/>
  <p:hasOperation rdf:resource="group-australia-network"/>
</p:Constraint>

<p:Operation rdf:ID="group-australia-network">
  <p:hasAct rdf:resource="policy:group"/>
  <p:hasFunction rdf:resource="policy:equal"/>
  <p:hasObject rdf:resource="network-australia"/>
</p:Operation>

<p:Object rdf:ID="network-australia"
  p:hasIdentity="urn:facebook:network:australia"/>
```

These expressions capture the unique identifier for the Facebook Australia Network, and allows viewing (policy:render) for this network (policy:group) using the unique identifiers from the policy model semantics.

4.2 Friend’s Privacy

In this use case, we need to allow your friends and colleagues that appear in your photos the ability to state whether they approve their image being published, including the scope of users who can view the photos. Using Facebook, as an example, we need to support the scenario of a friend tagging them on a photo and allow them to opt-out if they do not agree.

This use case is more complex in that it requires negotiation between the owner of the photo and the friends in the photo. We will not discuss the intricacies of policy negotiation [10] in this paper, but highlight this as a requirement for the future Policy-Oriented Web.

Typically, we would see that users would have a default privacy policy as part of their account profile. This policy would express their preferences on how their image can be used in social network photo albums. Their privacy policy would then be compared to the “view rights policy” that one of their friends is proposing. If there is conflict then this would stop the publication (the default action) and the user may be asked to “consider” the policy and confirm/deny it manually.

For example, the below RDF/XML snippet shows a privacy-policy in which the user has denied viewing (render) for any resources containing their image for any Facebook Network.

```

<p:Policy rdf:ID="myPrivacy">
  <p:hasResource rdf:resource="images-of-me"/>
  <p:hasAction rdf:resource="view-deny"/>
  <p:hasOperation rdf:resource="group-network"/>
</p:Policy>

<p:Operation rdf:ID="group-network">
  <p:hasAct rdf:resource="policy:group"/>
  <p:hasFunction rdf:resource="policy:equal"/>
  <p:hasObject rdf:resource="network"/>
</p:Operation>

<p:Object rdf:ID="network"
  p:hasIdentity="urn:facebook:network:all"/>

<p:Action rdf:ID="viewDeny" p:hasFocus="deny" >
  <p:hasAct rdf:resource="policy:render"/>
</p:Action>

```

The same model can be used to deny “public” and “friend” access (with appropriate identifiers from FaceBook). Conversely, if the user was happy to allow access for any network, friend, or public, then the “hasFocus” can simply be changed to “allow”.

4.3 Toward Interoperability

Returning to the two real-world use cases described in Section 2, both of these should have been avoided with an appropriate policy expression and accountability across Social Networks and platforms. Today, however, even if the correct rights/privacy/access criteria was selected under the controlled Social Network environment, the lack of policy support at the operating system level (including the web browser) hinders policy conformance. This is one of the greatest challenges for the Policy-Oriented Web; to become pervasive across all platforms and

Who can see these items?
(Acting on 1 item)

- ☐ Only You (Private)
- ☒ Your Friends
- ☐ Your Family
- ☒ Anyone (Public)

Who can comment?

- ☐ Only You
- ☐ Your Friends and/or Family
- ☒ Your Contacts
- ☐ Any Flickr User (Recommended)

Who can add notes & tags?

- ☐ Only You
- ☐ Your Friends and/or Family
- ☒ Your Contacts (Recommended)
- ☐ Any Flickr User

Fig. 8. Flickr Photo Permissions

services to enable any application to depend on open and interoperable policy-support services.

Looking back at the Flickr case, a rights policy could express that “Your Friends” in your photo have not given permission for their image to be reproduced (outside this specific Social Network). Figure 8 shows the permissions from Flickr for photos, which includes these constraints. In the FaceBook case, a similar policy could express that the family photos are not reproducible outside of Facebook.

Notice that the fundamental differences between Flickr (Figure 8) and FaceBook (Figure 3) include some permissions (such as excluding named people and “networks” versus “groups”) but are also similar in other respects. This means that at one level interoperability across these two Social Networks is possible if they share (and reuse) some of the core policy constructs. However, if one used “policy:group” and the other defined their own “flickr:family” then there will be some issues to overcome. More significant will be the lack of support for some features (eg exclusion of people) that only one Social Network supports.

The Flickr options (see Figure 8) also includes more permissions than just view (render). Specifically, they also allow for “commenting” and “tagging”. However, when you look deeper at the Facebook implementation of “view” it does also allow commenting on photos. This implies that if you translate “view” from Flickr to Facebook then you must not allow “commenting” and “tagging”, unless they are also specified. This leads to issues of conflict detection across these policies.

For example, this Act:

```
<p:Act p:hasIdentity="urn:facebook:view"/>
```

is similar to:

```
<p:Act p:hasIdentity="urn:flickr:view"/>
```

```
<p:Act p:hasIdentity="urn:flickr:comment"/>
```

```
<p:Act p:hasIdentity="urn:flickr:tag"/>
```

However, the reverse is not true.

These two sets of Acts could be in direct conflict if not used correctly. As they currently stand, the process in determining this conflict may involve prior knowledge, most likely via humans mapping the two core parts of the different policy language ontologies, and building the conflict detection into the software application logic.

The longer term aim is to support services for ontology mappings to help automate this process. Ontology mapping is extremely difficult to generalise but significant research efforts are showing early promises [9]. We see this as a key feature of the Policy-Oriented Web and a future research challenge.

The challenges that lay ahead are for the Social Network communities to develop the common vocabularies (ontologies) for the policy expressions. This will enable a policy in Facebook to be supported in Flickr, for example. A greater challenge is the support in different platforms, like web browsers, to be aware that policies are attached to content. This is the long term goal of the Policy-Oriented Web.

5 Related Work

Requirements for any new area of work are always important. A number of research goals in the area of semantic policies include [15]:

- lightweight knowledge representation to reduce the effort for policy-oriented frameworks for specific communities,
- incorporation of controlled natural language syntax for expressing policy rules, and a
- relaxed cooperative policy enforcement regime to not discourage users.

Others [16] indicate that the primary requirement is viewing policies from the privacy and business perspective so as to enable compatibility across the enterprise. Previous international workshops on the Semantic Policies [17] [18] presented many papers on emerging requirements for the policy-oriented web, including trust and negotiation mechanisms. However, very few deal with e-Society and Social Networks as the driver and consider the policy requirements from that context.

There are some efforts now appearing on an initial functional architecture for the policy-oriented web. These include the three basic capabilities of [19]:

- policy transaction logs to enable the assessment of past policy decisions, either in real-time or for post-processing,
- policy language framework that enables a shared policy vocabulary to evolve over time from overlapping communities on the web, and
- policy reasoning tools to enable policies to be evaluated and decisions made to assist the user.

There is also relevant work on privacy and identity management in the PRIME Project [20] and POEM Project [35] that has developed detailed enterprise architectures that could be generalised to support policy management tasks within a Social Networks context.

Some frameworks [21] are grounded on XML technology and define architectures consisting of policy management tools, policy databases, policy decision points, and policy enforcement points. Others follow this idea and extend the policy architecture based on a role-based access control model [22] or view-based access control [23] and a trading services model [24]. Frameworks also classify policies into high-level and low-level [25] to reflect and support different enforcement capabilities.

There is a significant body of work that reviews and compares different Privacy languages (such as as EPAL, P3P, XACML) and supporting frameworks [13] [26] [27] [28] [29]. Their general conclusion is that a common approach in the future will simplify policy analysis and reduce inconsistencies and promote policy reuse across communities and enhance such policy protection on the web. We have found that our Policy information model (see Figure 8) moves towards this goal, and provides more relevant policy-semantics (over existing languages) to express such policies. For example, Parties and roles, and dual-focus Actions provide clearer semantic and functions more relevant to Social Network requirements.

The application of semantic web technologies to structured policy languages (eg XACML) has shown how its expressive power can easily accommodate such

transformation and extensions but highlight several aspects for future research [30] and specific needs for a policy language for defining security requirements [31].

Investigation of privacy support in Social Networks has found that third-party access to user information (eg via open APIs such as OpenSocial) as potentially compromising [32] to users as the conformance to the user's policy is solely at the discretion of the third-party. Others found that the user model used for Privacy is not consistent with what is implemented by the Social Network providers [33] nor the way a user's privacy decisions are based on the relationship with the provider as well as other individuals [34].

We have also been working on extending the ODRL rights expression language [36] to accommodate more general policy features. We expect that this will be the basic model that could be widely deployed given its success in the mobile community.

6 Conclusion

We have seen that Social Networks have become an "overnight" phenomenon - backed up by Web 2.0 technologies - and provide rich user experiences. We have also shown that some of these experiences are not socially (or legally) acceptable. This is a golden opportunity for the Policy-Oriented Web to play a more significant role in the e-Society. The core area would be to better express the semantics of policies covering the access to user-generated content, and users personal preferences.

We have also defined the basis for a flexible information model that can underpin the Policy-Oriented Web and promote it as a new platform that will enable pervasive policy management across Web 2.0. We have shown some examples of applying the Policy-Oriented Web language to some use cases from real issues dealing with Social Networks. The current model is not complete and we expect that there will be a number of enhancements that can be applied to this preliminary semantic model with additional use cases, but the key idea is that we can begin to articulate the core concepts, classes, and relationships for a policy language framework. Future research areas will include policy conflict detection and accountability.

These are just the first steps in bringing policy-supportive technologies to the e-Society communities. These communities thrive on "simple" technologies that address their needs. The Policy-Oriented Web - as a semantic policy platform - will need to be integrated into the Web 2.0 style of technologies. This means more work is needed on the user interfaces for policy interactions and the integration with existing Web 2.0 platforms and deployment technologies. The end result should see the Policy-Oriented Web supporting more of the e-Society needs and evolving into a more user-focussed technology platform.

Acknowledgments

NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program and the Queensland Government.

References

- [1] Qui, X.: Citizen Engagement: Driving Force of E-Society Development. In: Wang, W. (ed.) IFIP International Federation for Information Processing. Integration and Innovation Orient to E-Society, Integration and Innovation Orient to E-Society, vol. 252, 2, pp. 540–548. Springer, Boston
- [2] Iannella, R.: Industry Challenges for Social and Professional Networks. In: W3C Workshop on the Future of Social Networking, Barcelona (January 15-16, 2009), <http://www.w3.org/2008/09/msnws/papers/nicta-position-paper.pdf>
- [3] Weitzner, D.G.: Profiling, and Privacy. IEEE Internet Computing, 95–97 (November/December 2007)
- [4] Cohen, N.: Use My Photo? Not Without Permission. New York Times (October 1, 2007), http://www.nytimes.com/2007/10/01/technology/01link.html?ex=1348977600&en=182a46901b23f450&ei=5124&partner=permalink&exp_r=permalink
- [5] ABC Media Watch, Filleting Facebook. Australian Broadcasting Corporation (ABC) (October 29, 2007), <http://www.abc.net.au/mediawatch/transcripts/s2074079.htm>
- [6] Iannella, R., Henriksen, K., Robinson, R.: A Policy Oriented Architecture for the Web: New Infrastructure and New Opportunities. In: W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement, Ispra, Italy, October 17-18 (2006)
- [7] Iannella, R. (ed.): Open Digital Rights Language, Version 1.1 Specification. ODRL Initiative (September 19, 2002), <http://odrl.net/1.1/ODRL-11.pdf>, <http://www.w3.org/TR/odrl/>
- [8] Wenning, R., Schunter, M. (eds.): The Platform for Privacy Preferences 1.1 (P3P1.1) Specification. W3C Working Group Note (November 13, 2006), <http://www.w3.org/TR/P3P11/>
- [9] Euzenat, J., Shvaiko, P.: Ontology Matching. Springer, Berlin (2007)
- [10] Arnab, A., Hutchison, A.: DRM Use License Negotiation using ODRL v2.0. In: Proceedings 5th International Workshop for Technology, Economy, and Legal Aspects of Virtual Goods and the 3rd International ODRL Workshop, Koblenz, Germany, October 11-13 (2007)
- [11] Hu, Y.J.: Semantic-Driven Enforcement of Rights Delegation Policies via the Combination of Rules and Ontologies. In: Workshop on Privacy Enforcement and Accountability with Semantics, International Semantic Web Conference 2007, Busan Korea (2007)
- [12] García, R., Gil, R.: An OWL Copyright Ontology for Semantic Digital Rights Management. In: IFIP WG 2.12 & WG 12.4 International Workshop on Web Semantics, Montpellier, France (November 2006)
- [13] Kolari, P., Ding, L., Shashidhara, G., Joshi, A., Finin, T., Kagal, L.: Enhancing Web privacy protection through declarative policies. In: Sixth IEEE International Workshop on Policies for Distributed Systems and Networks, June 6-8, pp. 57–66 (2005)
- [14] OASIS eXtensible Access Control Markup Language (XACML), Version 2.0. OASIS Standard (February 1, 2005), <http://docs.oasis-open.org/xacml/2.0/XACML-2.0-OS-NORMATIVE.zip>

- [15] Bonatti, P.A., Duma, C., Fuchs, N., Nejd, W., Olmedilla, D., Peer, J., Shahmehri, N.: Semantic Web Policies - A Discussion of Requirements and Research Issues. In: Sure, Y., Domingue, J. (eds.) ESWC 2006. LNCS, vol. 4011, pp. 712–724. Springer, Heidelberg (2006)
- [16] Kolari, P., Finin, T., Yesha, Y., Lyons, K., Hawkins, J., Perelgut, S.: Policy Management of Enterprise Systems: A Requirements Study. In: Proceedings of the Seventh IEEE International Workshop on Policies for Distributed Systems and Networks, POLICY 2006 (2006)
- [17] Semantic Web and Policy Workshop, Galway, Ireland (November 7, 2005)
- [18] 2nd International Semantic Web Policy Workshop (SWPW 2006), Athens, GA, USA, November 5–9 (2006)
- [19] Weitzner, D.J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J., Sus, G.J.: Information Accountability. MIT Computer Science and Artificial Intelligence Laboratory Technical Report MIT-CSAIL-TR-2007-034 (June 13, 2007)
- [20] Casassa-Mont, M., Crosta, S., Krieglstein, T., Sommer, D.: PRIME Architecture V2 (March 29, 2007),
https://www.prime-project.eu/prime_products/reports/arch/pub_del_D14.2.c_ec_WP14.2_v1_Final.pdf
- [21] Clemente, F.J.G., Perez, G.M., Skarmeta, A.F.G.: An XML-Seamless Policy Based Management Framework. In: Gorodetsky, V., Kutenko, I., Skormin, V.A. (eds.) MMM-ACNS 2005. LNCS, vol. 3685, pp. 418–423. Springer, Heidelberg (2005)
- [22] Bhatti, R., Ghafoor, A., Bertino, E., Joshi, J.B.D.: X-GTRBAC: an XML-based policy specification framework and architecture for enterprise-wide access control. ACM Transactions on Information and System Security (TISSEC) 8(2), 187–227 (2005)
- [23] Koch, M., Parisi-Presicce, F.: UML specification of access control policies and their formal verification. Software and Systems Modeling 5(4), 429–447 (2006)
- [24] Lamparter, S., Ankolekar, A., Studer, R., Oberle, D., Weinhardt, C.: A policy framework for trading configurable goods and services in open electronic markets. In: Proceedings of the 8th International Conference on Electronic Commerce, Fredericton, New Brunswick, Canada, August 14–16 (2006)
- [25] Pretschner, A., Hilty, M., Basin, D.: Distributed usage control. Commun. ACM 49, 9 (2006)
- [26] Tonti, G., Bradshaw, J., Jeffers, R., Montanari, R., Suri, N., Uszok, A.: Semantic Web Languages for Policy Representation and Reasoning: A Comparison of Kaos, Rei, and Ponder. In: Fensel, D., Sycara, K., Mylopoulos, J. (eds.) ISWC 2003. LNCS, vol. 2870, pp. 419–437. Springer, Heidelberg (2003)
- [27] Anderson, A.: A comparison of two Privacy Policy Languages: EPAL and XACML. Sun Microsystems Labs technical report (2005),
http://research.sun.com/techrep/2005/sml_i_tr-2005-147/
- [28] Ardagna, C., Damiani, E., De Capitani di Vimercati, S., Fugazza, C., Samarati, P.: Offline Expansion of XACML Policies Based on P3P Metadata. In: Lowe, D.G., Giedke, M. (eds.) ICWE 2005. LNCS, vol. 3579, pp. 363–374. Springer, Heidelberg (2005)
- [29] Jensen, C., Tullio, J., Potts, C., Mynatt, E.D.: STRAP: A Structured Analysis Framework for Privacy. Georgia Institute of Technology Technical Report GIT-GVU-05-02, <http://hdl.handle.net/1853/4450>
- [30] Damiani, E., De Capitani di Vimercati, S., Fugazza, C., Samarati, P.: Extending Policy Languages to the Semantic Web. In: Koch, N., Fraternali, P., Wirsing, M. (eds.) ICWE 2004. LNCS, vol. 3140, pp. 330–343. Springer, Heidelberg (2004)

- [31] Kagal, L., Finin, T., Joshi, A.: A Policy Based Approach to Security for the Semantic Web. In: Fensel, D., Sycara, K., Mylopoulos, J. (eds.) ISWC 2003. LNCS, vol. 2870, pp. 402–418. Springer, Heidelberg (2003)
- [32] Felt, A., Evans, D.: Privacy Protection for Social Networking Platforms. In: Web 2.0 Security and Privacy at the 2008 IEEE Symposium on Security and Privacy, Oakland, California, USA, May 18-21 (2008)
- [33] Chew, M., Balfanz, D., Laurie, B.: (Under) Mining Privacy in Social Networks. In: Web 2.0 Security and Privacy at the 2008 IEEE Symposium on Security and Privacy, Oakland, California, USA, May 18-21 (2008)
- [34] Grandison, T., Maximilien, E.M.: Towards Privacy Propagation in the Social Web. In: Web 2.0 Security and Privacy at the 2008 IEEE Symposium on Security and Privacy, Oakland, California, USA, May 18-21 (2008)
- [35] Kaiser, M.: Toward the Realization of Policy-Oriented Enterprise Management. *IEEE Computer*, 57–63 (November 2007)
- [36] Guth, S., Iannella, R. (eds.): ODRL Version 2.0 Core Model. Draft Specification (March 6, 2009), <http://odrl.net/2.0/DS-ODRL-Model.html>