

Renato Iannella

Principal Scientist

National ICT Australia (NICTA)

<renato@nicta.com.au>

A FRAMEWORK FOR THE POLICY-ORIENTED WEB IN SOCIAL NETWORKS

Abstract: With so much user-generated content being shared by these social networks, there is the real danger that the implicit sharing policy rules that communities have developed over time will be lost in translation in the new digital communities. 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 expression, transparency, conflict, and accountability will be the driving force that enables Social Networks to appreciate new interoperable policy regimes

Keywords: Policy-Oriented Web, Social Networks.

1. Introduction

Of all the ICT technologies of the past decades, the Web has certainly had the biggest and most profound impact on the world. It has successfully made the Internet a common appliance and broken down technical barriers to enable end users with the capability to access and share services and build community and social interactions. The web has been successful because it consists of a set of open, simple and scalable technologies, including identifiers (URI), the application layer protocol (HTTP) and the content markup language (HTML). Over time these standards have been refined; HTML, for example, has been reinvented as the XML compliant XHTML standard. In addition, completely new standards, such as RDF, have been introduced with varying degrees of success and uptake.

Alongside these web standards, a range of tools has emerged to simplify the tasks of creating and publishing content on the web. The advent of cheap web storage, weblog applications, wikis, camera phones and so forth has made it trivial and inexpensive for people and organisations to share their knowledge, opinions, photos and other content with the world. The Web is also becoming

the de-facto way of life for many; “as online connections between people become increasingly intertwined with real-world interests, social networking methods are moving toward simulating real-life social interactions” [12] and we need to address technology support for the implications of these social interactions [2].

Such implications are that sometimes personal and corporate information was never supposed to be shared with the entire world. Content owners currently lack a standard means to control how, where and by whom their content is used. Although web service providers (e.g., Facebook and Flickr) may implement their own privacy and rights management systems, they are generally not interoperable nor scalable. When content is taken from one of these systems and placed in another, any policy information that was attached to that content is lost.

All organisations need to collaborate and all organisations have different and dynamic rules and guidelines that govern these relationships. This varies from a consumer buying music for their iPod to a medical officer needing to analyse a number of reports and provide a recommendation to the Minister for a new health directive. In all cases, information and services are provided from autonomous organisations, some open, some closed, with different rules for access, different obligations, different expectations of service levels, and different consequences for non-conformance. In essence, the interactions are driven by dynamic policies.

Today’s current infrastructure, in particular, the web and service-oriented architectures, are not positioned to address the business and technical issues dealing with the “policy-oriented architecture”. Specifically, they deal at the lower layers of infrastructure, not the semantically richer information layers. For example, they can deal with a quality-of-service policy that needs to guarantee a high bit-rate transfer, but balk at a Privacy policy that needs the names removed from medical records before the summary information can be distributed to medical researchers who in turn can only publish the results in an open access journal, which must not allow commercial advertising next to the paper. This is not only a case of policy integration and harmonisation, but of policy integrity.

2. The Policy-Oriented Web (POW) Framework

The key motivation of the Policy-Oriented Web (POW) regime is a focus on collaborative web interactions - typified by today's Social Networks. POW is taking the Web away from today's simple browser/server transfer of information to a more sophisticated model based on predefined and dynamic information reflecting interactions between people and resources (ie web objects) across time and space. A web object is the smallest possible resource (or information bit) that needs to be managed by a POW infrastructure [3]. It will have needs that require policy management to enable it to successfully be part of a collaborative web interaction.

The web object will be supported by the following POW policy-types:

- Obligations that manage the rules or expectations of web objects, such as privacy, and rights management.
- Profiles that manage the characteristics of web objects, such as personal characteristics, identity, and role-based management.
- Preferences that manages the changing context of web objects, such as priority and social rules management.
- Presence that manages the dynamic environment of web objects, such as location and time management

In effect, the POW Framework is escalating all Web Objects up to independent "first class" Web citizens. No longer will a web object simply be a bit of data on a server ready to be transferred to a client. Web objects will exhibit behaviour and state and manage their own interactions with other web objects via the collaborative web infrastructure as show in Figure 1.

The collaborative POW infrastructure also moves away from the typical "layered architecture" towards a framework where objects - and their obligations, profile, preferences, and presence - interact at the same level and with the same priority. Web objects are user-definable and will help address real problems faced today on the Web.

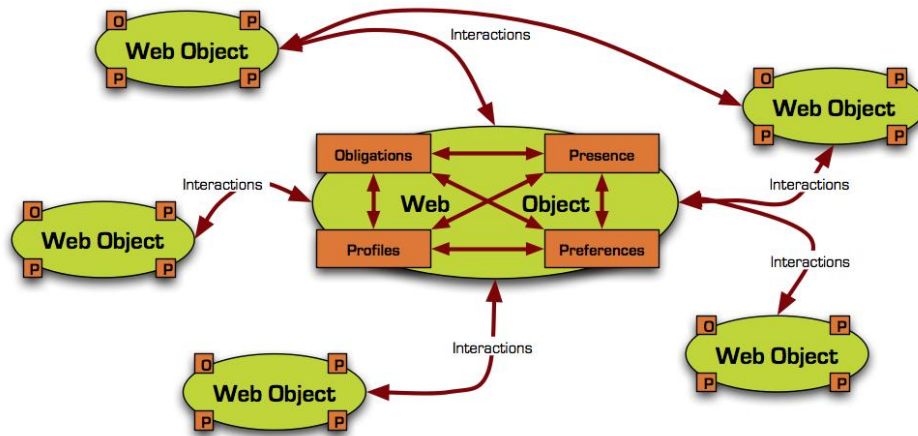


Figure 1. Collaborative Policy-Oriented Web Framework

3. Future Scenario - Nikky in 2020

Nikky is the Chief Scientist at NICTA and has just been invited to keynote the 29th International World Wide Web Conference in Singapore on the history of the Policy-Oriented Web. The incoming email from the conference committee automatically triggers Nikky's POW-enabled Organiser software (POWER) to make the appropriate arrangements for the trip. Nikky prefers Tiger airways and the Sheraton hotel chain and has loyalty programs with them.

POWER registers Nikky with the WWW2020 Conference Web Site as it trusts this site with Nikky's identity profile. Next, POWER searches the global flights services and finds two options to get to Singapore, one with Tiger, the other with Tesna airways. Even though she prefers Tiger, the Tesna flights are more convenient for her and are booked with her preferred aisle seating. POWER also tries to enroll Nikky in the Tesna FF program so she can earn points for this trip. The Tesna FF program states that members will be subject to promotional emails from time to time. Nikky's privacy preferences stipulates that she does not wish to enroll in such programs and POWER does not enroll her.

Luckily, the conference is being held at the Singapore Sheraton and POWER books a non-smoking room for the trip. Unfortunately, the costs of the room is greater than the per-night recommendation in the NICTA Travel Policy. This policy is flexible and allows the supervisor to override this under some circumstances. Her boss gets a request to override the policy which he agrees to. The room is booked and her Loyalty card details also registered.

Nikky begins work on her presentation for the Conference using her POW-enabled presentation software (POWTALK). She looks through all the archives from the old POW project and extracts some key images and text. She copies an image from an old presentation (circa 2010) that was licensed for non-profit use only for 2 years. POWTALK communicates with the license owner and negotiates a one-year extension for the use of the image. Her boss approves the purchase.

Nikky is on the plane traveling to Singapore. The Tesna POW-enabled inflight entertainment network (POWTEN) knows she likes satirical comedies (from POWER) and lines up a marathon of Woody Allen movies for her to watch. Unfortunately, her plane is diverted to Kuala Lumpur due to bad weather and eventually arrives in Singapore hours later.

Nikky arrives and is in the taxi on her way to the Hotel. POWER notices that she is in Singapore now and that her Hotel booking was only guaranteed to 6PM that day. POWER immediately contacts the Hotel and informs them that she will be arriving in 20 minutes. But they have already released the room. POWER finds another hotel close by the Sheraton and books a room there and tells Nikky. POWER then informs the Taxi driver's GPS Navigator of the change of destination.

Nikky's talk is a super success. The conference organisers publish her slides on the conference web site under the NICTA intellectual property terms. POWTALK notices this publication and checks that the new extended license they obtained did allow for the re-publishing of the image by a third-party. It did not, so when others access her slides, the licensed image is not shown.

Nikky travels home. Her favourite movies are back on POWTEN and her plane arrives safely back home on time.

The Key Points being addressed behind this scenario include:

- Profile Management - trusted sharing of your personal details
- Preference Management - modifying situations with additional information
- Obligations Management - ensuring privacy details are not shared and contract terms are fulfilled
- Presence Management - realtime changes to situations

4. Research Challenges Ahead

The four policy types (Obligations, Profiles, Preferences, and Presence) all have common challenges to meet future deployment realities. These ensuing POW research activities can be clustered into four areas:

- Policy Expression
- Policy Transparency
- Policy Conflict
- Policy Accountability

Policy Expression will focus on representing the interactions and dependencies between the critical policy languages for collaboration and sharing, such as Privacy, Rights, and Identity management. A policy language is a mechanism to declare a set of rules or statements that capture and express the requirements of individuals and organisations from a corporate, legal, best practices, and/or social perspective. Currently, these existing individual policy languages are missing an overall framework and architecture allowing the combination of different policy languages to interoperate and provide an accountable, enforceable, flexible and trusted experience for the web community.

One of the first steps is to define the entities of the POW Information Model as this holds the key to interoperability across the various policy language layers [10]. Figure 2 shows an example model based on the analysis of a number of different policy languages and includes:

- 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”.
- Resource - these are the resources (web objects) 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”.
- 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”.

The POW Information Model also include a filter layer which narrows the Policy with:

- Constraint - conditions that will limit an Action of the policy.
- Context - changes to the Policy depending on certain external conditions.
- Obligation - requirements that must be met by a Party in order to satisfy the policy.

The end goal is to develop a framework in which policy languages can coexist and share semantics and to support a wide range of policy functions. This poses fundamental challenges with each language having its own specific goals and context, but having a common impact on web objects when taken together. In particular, how they are harmonised to support open Policy Transparency of distributed policy information, conformance, and behaviour. An unintended consequence of transparency is that it sometimes “intensifies privacy concerns rather than engender trust” [5] and these needs serious rethinking, redesign, and new approaches.

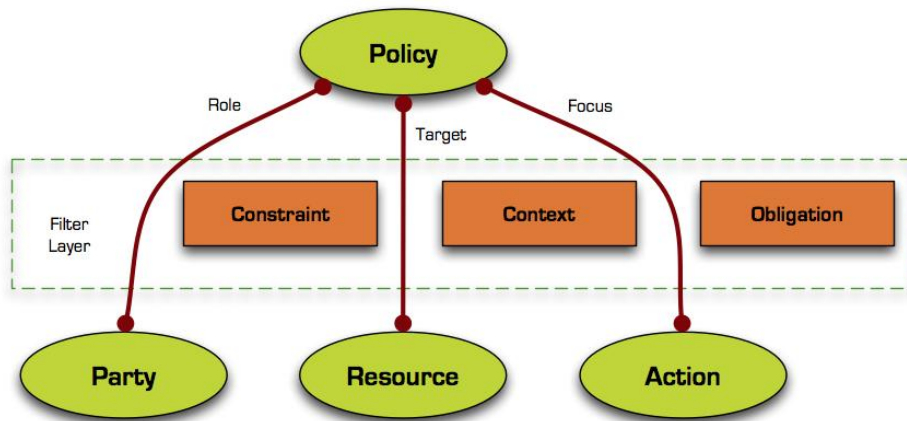


Figure 2. Policy-Oriented Web Information Model

Policy Conflict is a major challenges of this activity to resolve policy contradictions. For example, it may happen that a user’s privacy policy states that a service may not track the user’s browsing history, while the service’s rights policy says that the user may view the content provided the user’s browsing history can be tracked. The first problem is identifying that a conflict exists. The second problem is what to do when the conflict is identified. The appropriate resolution may involve input from the user, and therefore the research must take into account issues surrounding human interaction with these policies.

Policy Accountability is another major challenge to detect policy breaches so that some modicum of accountability can be instilled in the web. This is a departure from the current approach of attempting to provide policy enforcement. Most attempts to provide enforcement on the web (for example, traditional digital rights management for multimedia content) have ended in failure, and are not well accepted by the web community. The accountability mechanism may support recording policy breaches that the user is completely aware of and the consequences of their actions and is a more “relaxed cooperative policy enforcement regime to not discourage users” [6].

The overall research challenges to provide this open POW framework include:

- How to develop an overarching framework for policy language management including an abstract model of the concepts relevant to the POW.
- How to enable accountability and transparency for policy use and breaches.
- How to deal with inconsistencies across policies by enabling policy conflicts to be detected and finding the appropriate means for resolving that conflict.
- How to visualise large inter-related complex policy sets and allowing policy interactions to be scrutinised.
- How to enable semantic interoperability (e.g., via ontologies) of concepts across the realms of identity, privacy, rights management and related policy types.
- How to process policies using different reasoning techniques to infer implied policies rules.

The impact of this research will also be far ranging as the POW framework will provide the policy community with a flexible mechanism to represent their policy language with the added benefit of policy consistency checking and accountability metrics. The uses of such a framework would be wide-scale across many domains and sectors and will also require international consensus on standardising the interoperability aspects of POW [15].

5. Current Scenario - Social Networks

We looked at two popular Social Networks (FaceBook and Flickr) and reviewed the types of policy conditions (or constraints) that can be applied to user-generated content. Figure 3 shows examples of these policy conditions from these two Social Networks for the user's privacy and permissions settings. Similar screens can also be seen and consistent with Professional Networks, such as LinkedIn and Plaxo. We then looked at the policy implications from Social Networks using the four POW Challenges.

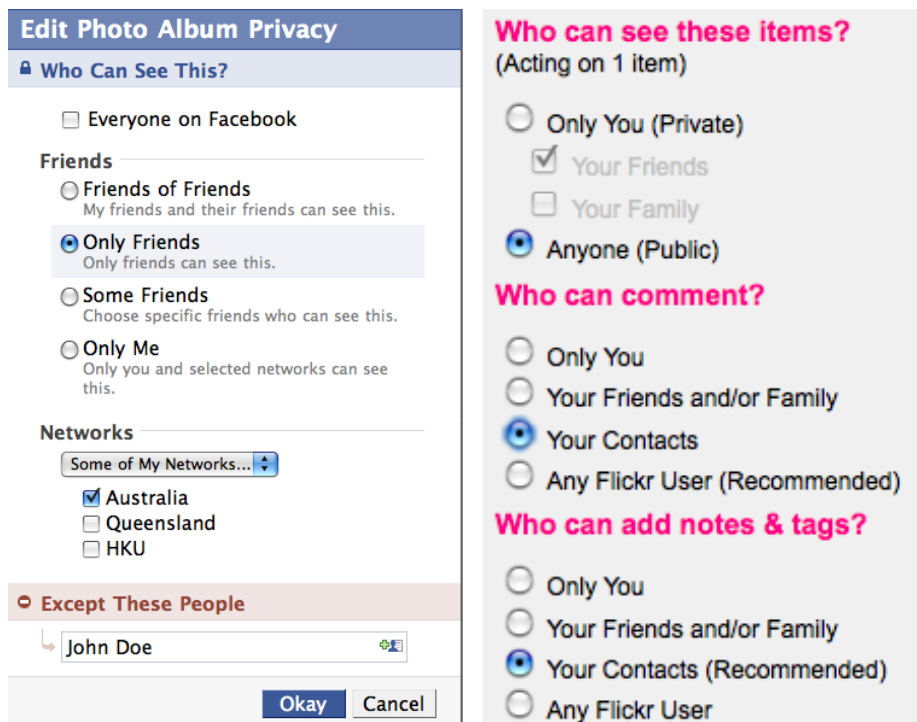


Figure 3. FaceBook (left) and Flickr (right) Privacy and Permissions Settings

5.1.Policy Expression

What is clear from Figure 3 is that the policy decision points are focussed on constraining who the end user party is or what group they are a member of. This is a different approach from past requirements (from the transaction-based DRM world) that was more interested in temporal, spatial and count-based constraints.

For Social Networks, the content owner can specify these types of limitations for who can access their content:

- Only the content owner (ie no one else)
- Specific (named) friends and colleagues (both allowed and not allowed)
- Family members

- All direct friends or colleagues
- Your second level friends or colleagues (ie friends of friends)
- All Groups (that the content owner is a member of) or some Groups
- Everyone (ie public)

Policy languages, such as ODRL Version 2.0 [18] will need to be expressive enough to handle these types of conditions. We have begun to look at the extensions needed by ODRL to meet this need [17]. One immediate outcome has identified the need for a generic “anyone” party and membership semantics that need to be part of the ODRL Core Metadata [19]. This “anyone” party construct can then act as a constraint for smaller subsets of people and groups so that Duties, for example, can be specified.

The expression implication for Policy languages is a need for a “person” entity as well as a role indicator (such as friend, family, contact etc) and a group or membership mechanism.

5.2. Policy Transparency

In Social Networks it is important that there is awareness by all users of policies related to their actions (or inactions) and any consequences. In the examples shown in Figure 3, the owner of the content is selecting the policy options, but the end recipients (their friends and colleagues etc) are not typically made fully aware of these constraints since they will simply not see the content. For example, if the Flickr photo owner only allows “Commenting” to “Your Friends and Family”, then their full list of “Contacts” will simply not know they have been disallowed. In the same case, what is the difference between adding a Comment or a Note to a photo that requires a separation of the two in Flickr, and why conflate “notes and tags” together? What if you wanted your family to add tags and your friends to add comments?

More interestingly, if a Facebook user allows “Friends of Friends” to see some content, then the friend-of-a-friend will not know why they have been granted access - just that they have. This could be very confusing if you then take away that right. The end recipient should be made aware that they have access to this content because they are a friend of a friend.

There are many challenges in this area that deal with how to best inform the end recipient why they have access rights to some content, and why they do not to other content. Policy management systems need to not only inform the user, but know of all the relevant policies related to that user. This implies some federation mechanisms across policy management systems. What is also important is making the end recipient aware of any consequences from not abiding by these policies.

5.3. Policy Conflict

Policy conflict becomes relevant when interoperability between Social Networks is considered. Today, there is very little of this occurring at the policy level. That is, you cannot take all your Flickr photos - together with their policy descriptions - and import them into Facebook - and preserve the policy conditions.

The interoperability issue can only be addressed by first having a common or consistent policy semantics that are supported by various Social Network providers. There needs to be common policy language that can be the *lingua-franca* of Social Network portability. Even with this in place, the concepts used in each Social Network will need to be harmonised to allow for the mapping between them. Conflicts will occur so there needs to be a way to resolve these, and most may not be machine-resolvable.

For example, in Figure 3 we see that Facebook has the concept of “Friends of Friends” and Flickr does not. So how can we translate a policy that allows content to be seen by Friends of Friends to a service that does not have this concept? Even harder, and more important, is to recognise that there is even a potential conflict point between the two service providers.

Policy conflict will be a long and tedious task to solve. Semantic Web technologies and other more formal logical approaches may be the answer, but much experimentation will be required.

5.4. Policy Accountability

Given that policy interoperability will be a difficult task solely for machines to manage, there is clearly a better model that needs to be involve less “enforcement” and more “cooperation” with users. Social Networks policy

management should mirror the real world in which laws are made, people break laws, but they are aware of the consequences of their actions.

It would therefore be more feasible in the long term to allow Policies that are difficult to resolve (eg because of unknown semantics or conflicts) to be addressed by the end user. This means that the service can inform the user of the concern/conflict and allow them to make a decision as to what to do next. The process would also make it clear that their actions (either way) would be recorded for future accountability purposes. So the model is to trust the user first and if there is continual abuse of that trust, there is a mechanism to followup and resolve the issue.

Take the above example of “friends of friends” not being supported in Flickr (from a FaceBook policy). In this case the service would say that it does not support/understand this concept but will tell the end user that they can access the content - if they are a friend of Billie’s friends. If they continue, then the service can record that fact for any future followup. The alternative, would be simply to decline access and that would then have greater detrimental impact on the overall social experience.

6. Summary

“A new generation of Policy-Aware Web technology can hold the key for providing open, distributed and scaleable information access on the World Wide Web” [1] argues Tim Berners-Lee, the inventor of the Web. Today’s decentralised web, with inflexible access specification tools has hampered the full development of the Web as a social and corporate information sharing space. Without prior information sharing policies and only coarse access control (eg an entire web site, versus one specific web object), the reluctance to share information remains constant. Complexity can also stop progress of a new technology, and the POW will need to offer simplicity and efficiency to gain the support of the masses, like Web 2.0 Social Networks [11].

The POW is needed now as we are facing real serious issues in Social Networks [16] with privacy violations [14] and rights issues [13] and these use cases can help drive the POW roadmap in the right direction towards success.

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 [7] 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 [8] nor the way a user's privacy decisions are based on the relationship with the provider as well as other individuals [9].

The Web has evolved from the ground up following a "ladder of authority" [4] - a sequence of specifications that incrementally build greater levels of infrastructure for the Web. The greatest challenge the Web faces today is to capitalise on the efforts that have gone on previously in the development of specific policy languages - and to build the next layers of infrastructure to support a policy-oriented web. This will transform the web from an information delivery system into an information interaction system that will meet the emerging needs of the web community. In particular, it will address communities that have based their professional and societal life on the functions and structure of the web - and are now looking for greater infrastructure support for common services. Just like URIs, HTML and HTTP, the web now needs new reliable structures for policy management: the Policy-Oriented Web.

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] Tim Berners-Lee, Daniel J. Weitzner, Jim Hendler, & Dan Connolly. Creating a Policy-Aware Web: Discretionary, Rule-based Access for the World Wide Web. In Elena Ferrari and Bhavani Thuraisingham, editors, Web and Information Security. IRM Press, 2006
- [2] James Hendler, Nigel Shadbolt, Wendy Hall, Tim Berners-Lee, & Daniel Weitzner. Web Science: An Interdisciplinary Approach to Understanding the Web. Communications of the ACM, July 2008

- [3] Tim Berners-Lee, Wendy Hall, James A. Hendler, Kieron O'Hara, Nigel Shadbolt and Daniel J. Weitzner. A Framework for Web Science. Foundations and Trends in Web Science. Vol. 1, No 1 (2006) 1–130
- [4] Nigel Shadbolt, Wendy Hall, & Tim Berners-Lee. The Semantic Web Revisited. IEEE Intelligent Systems. May/June 2006
- [5] Irene Pollach. What's Wrong With Online Privacy Policies? Communications of the ACM, September 2007
- [6] Bonatti, P.A. & Duma, C. & Fuchs, N. & Nejdil, W. & Olmedilla, D. & Peer, J. & Shahmehri, N. Semantic Web Policies - A Discussion of Requirements and Research Issues. Proc. of the European Semantic Web Conference (ESWC 2006). LNCS 4011, Springer, pp. 712-724
- [7] Felt, A. & Evans, D. Privacy Protection for Social Networking Platforms. Web 2.0 Security and Privacy at the 2008 IEEE Symposium on Security and Privacy. Oakland, California, USA, 18-21 May 2008
- [8] Chew, M. & Balfanz, D. & Laurie, B. (Under)Mining Privacy in Social Networks. Web 2.0 Security and Privacy at the 2008 IEEE Symposium on Security and Privacy. Oakland, California, USA, 18-21 May 2008
- [9] Grandison, T. & Maximilien, E.M. Towards Privacy Propagation in the Social Web. Web 2.0 Security and Privacy at the 2008 IEEE Symposium on Security and Privacy. Oakland, California, USA, 18-21 May 2008
- [10] Iannella, R. & Henricksen, K. & Robinson, R. A Policy Oriented Architecture for the Web: New Infrastructure and New Opportunities. W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement, Ispra, Italy, 17-18 October 2006.
- [11] Mikroyannidis, A. Toward a Social Semantic Web. IEEE Computer, November 2007, pp 113-115
- [12] Breslin, J. & Decker, S. The Future of Social Networks on the Internet. IEEE Internet Computing, Nov/Dec 2007, pp 86-90
- [13] Cohen, N. Use My Photo? Not Without Permission. New York Times, 1 October 2007
<<http://www.nytimes.com/2007/10/01/technology/01link.html?ex=134897>

7600&en=182a46901b23f450&ei=5124&partner=permalink&exprod=permalink> (Accessed 5 Dec 2008)

- [14] ABC Media Watch, Filleting Facebook. Australian Broadcasting Corporation (ABC), 29 October 2007
<<http://www.abc.net.au/mediawatch/transcripts/s2074079.htm>> (Accessed 5 Dec 2008)
- [15] Policy Languages Interest Group (PLING). W3C Interest Group
<<http://www.w3.org/Policy/pling>> (Accessed 5 Dec 2008)
- [16] Iannella, R. Industry Challenges for Social and Professional Networks. W3C Workshop on the Future of Social Networking, 15-16 January 2009, Barcelona <<http://www.w3.org/2008/09/msnws/papers/nicta-position-paper.pdf>>
- [17] Governatori, G. & Iannella, R. Modelling and Reasoning Languages for Social Networks Policies. Thirteenth IEEE International EDOC Conference. 31 August - 4 September 2009, Auckland, New Zealand (to appear)
- [18] Guth, S. & Iannella, R. (eds). ODRL Version 2.0 Core Model. Draft Specification, 6 March 2009. ODRL Initiative <<http://odrl.net/2.0/DS-ODRL-Model.html>>
- [19] Guth, S. & Iannella, R. (eds). ODRL Version 2.0 Core Metadata. Working Draft, 12 February 2009. ODRL Initiative <<http://odrl.net/2.0/WD-ODRL-Core-Metadata.html>>