Encyclopedia of Information Science and Technology

Second Edition

Mehdi Khosrow-Pour
Information Resources Management Association, USA

Volume VII
R-S

Information Science REFERENCE

INFORMATION SCIENCE REFERENCE
Hershey · New York

Director of Editorial Content: Kristin Klinger
Director of Production: Jennifer Neidig
Managing Editor: Jamie Snavely
Assistant Managing Editor: Carole Coulson
Cover Design: Lisa Tosheff
Printed at: Yurchak Printing Inc.

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue, Suite 200
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com/reference

and in the United Kingdom by
Information Science Reference (an imprint of IGI Global)
3 Henrietta Street
Covent Garden
London WC2E 8LU
Tel: 44 20 7240 0856
Fax: 44 20 7379 0609
Web site: http://www.eurospanbookstore.com

Library of Congress Cataloging-in-Publication Data

Encyclopedia of information science and technology / Mehdi Khosrow-Pour, editor. -- 2nd ed.
p. cm.
Includes bibliographical references and index.
Summary: "This set of books represents a detailed compendium of authoritative, research-based entries that define the contemporary state of knowledge on technology"--Provided by publisher.
ISBN 978-1-60566-026-4 (hardcover) -- ISBN 978-1-60566-027-1 (ebook)
1. Information science--Encyclopedias. 2. Information technology--Encyclopedias. I. Khosrowpour, Mehdi, 1951-
Z1006.E566 2008
004'.03--dc22
2008029068

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

Security and Privacy in Social Networks

Barbara Carminati
Università degli Studi dell'Insubria, Italy

Elena Ferrari
Università degli Studi dell'Insubria, Italy

Andrea Perego
Università degli Studi dell'Insubria, Italy

INTRODUCTION

Web-based social networks (WBSNs) are online communities that allow users to publish resources (e.g., personal data, annotations, blogs) and to establish relationships, possibly of a different type ("friend," "colleague," etc.) for purposes that may concern business, entertainment, religion, dating, and so forth. In the last few years, the usage and diffusion of WBSNs has been increasing, with about 300 Web sites collecting the information of more than 400 million registered users. As a result, the "net model" is today used more and more to communicate, share information, make decisions, and 'do business' by companies and organizations (Staab et al., 2005).

Regardless of the purpose of a WBSN, one of the main reasons for participating in social networking is to share and exchange information with other users. Recently, thanks to the adoption of Semantic Web technologies such as FOAF and other RDF-based vocabularies (Brickley & Miller, 2005; Davis & Vitiello, 2005; Golbeck, 2004), accessing and disseminating information over multiple WBSNs has been made simpler (Ding, Zhou, Finin, & Joshi, 2005). If this has been quite a relevant improvement towards an easier sharing of information, it makes more urgent that content owners have control over information access. In fact, making available possibly sensitive and private data and resources implies that they can be used by third parties for purposes different from the intended ones. As a matter of fact, users' personal data and resources are regularly exploited not only by companies for marketing purposes, but also by governments and institutions for tracking persons' behaviors and opinions, and in the worst case, by online predators (Barnes, 2006).

It is then a challenging issue to devise security mechanisms for social networks, able to protect private information and regulate access to shared resources. In this article, besides providing an overview of the characteristics of the WBSN environment and its protection requirements, we illustrate the current approaches and future trends to social network security, with particular attention paid to the emerging technologies related to the so-called Web 2.0.

BACKGROUND

Usually, a social network is defined as a *small-world network* (Watts, 2003), consisting of a set of individuals (persons, groups, organizations) connected by personal, work, or trust relationships. Social networking is then a quite broad and generic notion, which in the Web context might be applied to any kind of virtual community. For instance, users registered to a Web service, such as Web mail, online journals, or newspapers requiring a subscription, can be considered as a social network. In the following, we adopt the definition provided by Golbeck (2005), according to which an online community's Web site can be considered a Web-based social network only if it satisfies the following conditions:

- Relationships are explicitly specified by its members, and not inferred from existing interactions (e.g., a mailing list can be used to infer implicit relationships).
- Relationships are stored and managed by using technologies, such as database management systems, allowing relationship analysis and regulating access and retrieval of relationship data.
- Members are able to access relationship information, at least partially.

Born in the late 1990s, in the last few years WBSNs gained increasing interest and diffusion. Although the first and most successful ones, such as MySpace, Friendster, and Facebook, were formerly designed for entertainment and socialization purposes, they are currently establishing themselves as a business model, through which institutions and organizations can set up a collaborative environment for specific purposes, and where it is possible to share resources at an intra- and inter-organizational level. Due to the great amount of collected data, WBSNs are currently the subject of great interest for statistical analysis (Wasserman & Faust, 1994; Freeman, 2004), since they may provide useful information not only to social researchers, but also for marketing purposes.

WBSNs may provide different kinds of services, ranging from information and contact sharing, to collaborative

rating, collaborative work environments, and so on. However, independently from the specific purposes of a WBSN, members' relationships are the core information on which all the provided services are based. In fact, they can be used not only to create connections among people sharing similar interests, but also to customize WBSN services themselves. This is particularly true in WBSNs supporting collaborative rating: in such a context, ratings may be given different weights, depending on the relationships existing between WBSN members. For instance, it may be the case that a given WBSN member m_1 considers more relevant (or trustworthy) the opinions of member m_2 than, say, those of member m_3. For this purpose, some WBSNs allow their members not only to specify personal relationships (e.g., "friend of," "colleague of") but also to establish *trust* relationships, which express how much they trust the other members either with respect to a specific topic (*topical trust*) or in general (*absolute trust*). For a thorough discussion on trust relationships and how they can be used, we refer the reader to the work by Golbeck and Hendler (2006).

As far as security is concerned, current WBSNs enforce simple protection mechanisms, which only allow their members to label given information as public or private, or to make it available to WBSN members with whom there exists a direct relationship of a given type (friend, colleague, etc.). However, these solutions on one hand may dramatically reduce the possibility of sharing information, which is the basic function of a WBSN, and on the other hand, they do not necessarily grant the required protection to personal information. In fact, giving to WBSN members just the choice of stating whether a given resource is public or private may result in hiding a huge amount of information. Moreover, it may frequently happen that WBSN members make publicly available resources that are accessed by people different from the ones they intended—the most typical case is a student publishing photos or blogs in recreational WBSNs, without considering that they can be accessed by his or her teachers.

Additionally, personal information and relationships among WBSN members must be protected when WBSN data are analyzed by data mining tools, that is, tools capable of analyzing massive datasets of personal information with the purpose of extracting models of social and commercial interest.

SECURITY AND PRIVACY REQUIREMENTS IN SOCIAL NETWORKS

In this section we consider the security and privacy issues related to WBSNs from two different points of view. First, we discuss the privacy-preserving techniques adopted to allow statistical analysis on social network data without compromising WBSN members' privacy, and then we illustrate the current approaches aimed at enforcing privacy protection when performing access control.

Privacy-Preserving Social Network Analysis

Data collected by WBSNs are an important source for social and marketing analysis, which may provide useful information on the evolution of a social community, collaborative problem solving, information distribution, and so on. Additionally, they can also be used to optimize social network services and customize them with respect to users' preferences and interests. However, when analyzing WBSN data for statistical purposes, it is necessary to avoid as much as possible disclosing private information about WBSN members.

So far, this issue has been addressed by anonymizing the network graph according to two main strategies, namely, *node anonymization* and *edge perturbation*. The former strategy aims at hiding members' identities by labeling the corresponding network nodes with random identifiers (naïve anonymization). In case nodes are associated with attributes which can be used to identify the corresponding user, the possibility of using techniques based on k-anonymity (Sweeney, 2002) has been discussed−see, e.g., Zheleva and Getoor (2007). By contrast, edge perturbation performs a set of random edge deletions and insertions, which prevent an attacker from inferring the identity of network nodes based on the existing relationships but, at the same time, preserve the utility of the graph for network analysis.

It has been noticed that the proposed solutions to node anonymization do not grant total privacy protection. In particular, Backstrom, Dwork, and Kleinberg (2007) carried out an extensive analysis of the possible attacks, and argued that the most effective strategies for privacy protection are those based on *interactive* techniques. According to this approach, the anonymized network graph is not disclosed; rather it is analyzed by the social network management system itself upon submission of a query, and then the result is perturbed by adding noise to the real answer.

By contrast, edge perturbation, when combined with node anonymization, grants a greater degree of protection. Examples of how such techniques are applied are provided by Frikken and Golle (2006), Hay, Miklau, Jensen, Weis, and Srivastava (2007), and Zheleva and Getoor (2007). In particular, Hay et al. (2007) report experimental results which show that random edge deletions and insertions grant graph anonymity when the perturbation affects a percentage of graph edges ranging from 5% to 10%. By contrast, a perturbation rate greater than 10% dramatically increases information loss, thus making useless the results obtained by

analyzing the perturbed graph. Although Zheleva and Getoor (2007) do not provide experimental results, they enhance the edge perturbation strategy by considering the different possible methods according to which it can be performed, and by evaluating the obtained perturbed graph with respect to information loss and link re-identification attacks.

Note, however, that graph anonymization is based on the assumption that the only information that can be obtained by an attacker is the one publicly released by the social network service. By contrast, this strategy is useless when applied to social networks, as most WBSNs are, to which any Web user can register, and where each member has a total or partial view of the network graph. In such a case, attackers can infer the network structure and members' identity with more or less accuracy by using techniques like *node bribing*, that is, by obtaining access to the partial view of the WBSN graph of one or more of its members, as illustrated by Korolova, Motwani, Nabar, and Xu (2008). The authors argue that it is possible to reduce the effectiveness of such attacks by limiting the neighborhood visibility of a member (his or her *lookahead* ℓ) to his or her neighbors ($\ell = 0$), and to the neighbors of his or her neighbors ($\ell = 1$). By contrast, in case $\ell > 1$, the possibility of obtaining correct information on the WBSN graph exponentially increases.

In conclusion, available privacy-preserving techniques, both those based on graph anonymization and those limiting WBSN members' lookahead, have the goal of preserving users' privacy when network data are analyzed through data mining tools. An additional issue is to enable a WBSN user to state which information should be public or private, and which members are authorized to access it. In this respect, current WBSNs enforce very naïve default protection mechanisms which cannot be personalized by WBSN members. We elaborate more on this issue in the next section.

Privacy-Aware Access Control

WBSN resources have protection requirements that cannot be enforced by simple mechanisms, as those currently adopted by WBSNs. An access control model for WBSNs should therefore take into account the specific characteristics of the application domain, in order to devise the most suitable access control strategies. In the following, we first discuss the main requirements for an access control mechanism tailored to WBSNs. Then, we survey the solutions proposed so far.

According to the traditional approach, access control requirements are expressed by *authorizations,* which in their basic representation are tuples of the form $\langle s, p, o \rangle$, where s is the subject authorized to access object o under privilege p (Bertino & Sandhu, 2005). However, such an approach is not suitable for dynamic and distributed environments, as WBSNs are, since a member may be required to update the authorizations applying to his or her resources whenever he or she knows new members, or if relationships he or she participates in are revoked. In such a scenario, it is preferable to *intensionally* denote authorized members by specifying the *requirements* they must satisfy to access a given resource. According to this strategy, whenever any modification to the state of the WBSN structure occurs, the set of authorized members will dynamically change, without the need to modify the existing authorizations.

So far, a variety of access control models have been proposed, which denote authorized users in terms of their characteristics, and not only by their identities. The role-based model (Ferraiolo, Kuhn, & Chandramouli, 2003) is the most popular one; others are those based on credentials (e.g., Winslett, Ching, Jones, & Slepchin, 1997; Agarwal, Sprick, & Wortmann, 2004) or certificates (e.g., Thompson et al., 1999; Palomar, Estevez-Tapiador, Hernandez-Castro, & Ribagorda, 2006). An analogous approach can be applied to WBSNs. In fact, WBSN members usually publish resources having in mind a specific audience consisting of, for example, their friends or colleagues. Therefore, in a WBSN context, *relationships* can be used to intensionally denote authorized members.

The enforcement of relationship-based access control requires addressing two main issues. First, it must be possible to verify the authenticity and reliability of information about relationships, in order to avoid security attacks based on forging faked relationships. Second, relationship information may have privacy protection requirements, and thus mechanisms should be enforced to regulate their disclosure.

A further requirement is related to the support of content-based access control (Adam, Atluri, Bertino, & Ferrari, 2002). Actually, the practice of 'tagging' resources is currently diffused among WBSN members, and content analysis is another possible solution to enforce content-based access control. However, since resource rating is performed by each single member and content analysis gives only probabilistic results about the actual content of a resource, strategies should be devised in order to obtain accurate and unambiguous descriptions, usable for access control purposes.

Finally, access permissions should take into account the possible operations to be performed on WBSN resources. Besides the traditional 'read' privilege, in collaborative environments WBSN members may be authorized to modify/delete a resource or add content to it. In such a case, it may be useful to support different types of 'write' privileges, such as 'modify' (authorized members can modify existing content or add new content), 'delete', and 'append' (authorized members can only add content, but not modify existing content). Additionally, when supporting 'write' privileges, it is important that any modification performed on a resource can be associated with the member who performed it. This means that supporting different privilege types requires enforcing accountability in the WBSN framework.

A last issue to be addressed concerns the access control architecture to be adopted. According to the traditional ap-

proach, access control is enforced on the side of the content provider. However, this solution may not be suitable to WBSNs, which may have millions of registered members and, as a consequence, the WBSN management service would be a bottleneck to the whole system.

As far as we are aware, the only two proposals of an access control mechanism based on WBSN relationships are the ones by Carminati, Ferrari, and Perego (2006) and Hart, Johnson, and Stent (2007).

In the proposal by Carminati et al. (2006), access control requirements are expressed by *access conditions,* which denote authorized members not only in terms of relationship types (e.g., friend, colleague), but also with respect to the relationship *depth* and *trust level*. The depth of a relationship corresponds to the distance between two members, considering only the edges labeled with a given relationship type. Thanks to this, it is possible to specify authorizations stating that a given resource can be accessed only by the friends of Alice, or by the friends of Alice's friends. By contrast, the trust level denotes how much confidence a member has on the fact that another given member does not reveal protected information to unauthorized members.

As far as access control enforcement is concerned, Carminati et al. (2006) adopt the rule-based approach proposed by Weitzner, Hendler, Berners-Lee, and Connolly (2006). More precisely, access authorizations are expressed by Horn-like clauses (rules), and it is the requestor who is in charge of demonstrating to the content provider of being authorized to access a given resource, by providing a proof of the corresponding access rules. WBSN resources and the corresponding access rules are managed by the resource owner, whereas relationship certificates are stored in a central directory, stored and managed by the WBSN management system. Whenever an access control request is submitted, the resource owner sends back to the requestor the set of associated access rules. The requestor then contacts the WBSN management system, in order to retrieve the relationship certificates concerning the relationships denoted by the received access rules. Then, he or she computes a proof, if any, demonstrating that he or she satisfies the rules. The resource owner sends the resource to the requestor only in case the provided proof is valid.

Also the access control model proposed by Hart et al. (2007) in their position paper uses existing WBSN relationships to denote authorized members, but only the direct relationships they participate in are considered, and the notion of trust level is not used in access authorizations. In addition, differently from Carminati et al. (2006), resources are not denoted by their identity, but based on their content. Information about resources' content is derived based on users' tags and content analysis techniques. Hart et al. (2007) do not provide any information about access control enforcement.

Both the approaches by Carminati et al. (2006) and Hart et al. (2007) assume that relationships are public. Later, Carminati et al. (2007) have extended their earlier research (Carminati et al., 2006) by proposing a privacy-aware access control mechanism, where the existing relationships are protected by a set of rules, called *distribution rules.* Such rules are used to regulate the distribution of relationship certificates to authorized members. Carminati et al. (2007) also address the issue of protecting relationship information that may be inferred by access rules, when enforcing access control. In fact, if an access rule states that, in order to be able to access a given resource, the requestor must be a friend of Alice, it is possible to infer that Alice participates in at least one relationship of type *friendOf,* otherwise no member will be able to access that resource. In order to deal with this issue, WBSN members are equipped with a set of group keys (Rafaeli & Hutchinson, 2003), called *relationship keys,* used to encrypt access conditions. More precisely, each WBSN member *m* holds a key for each type of relationship he or she participates in. These keys are shared by all the WBSN members in his or her social network group, that is, all the WBSN members connected to *m* by paths labeled with those relationship types. Whenever *m* receives an access request to a resource he or she owns, he or she does not send the corresponding access rules in plaintext. Rather, each access condition in the access rule is encrypted with the corresponding relationship key. For instance, if an access condition puts a constraint on relationships of type *friendOf, m* will encrypt it with the key corresponding to that relationship type. As a consequence, the requestor will be able to read that access condition only when he or she belongs to the same group of type *friendOf m* participates in.

It is important to note that all the approaches we have described so far support 'read' privileges only. Of course, they can be extended with other types of access modes, but enforcing accountability would require a relevant extension to the access control mechanisms described above.

FUTURE TRENDS

WBSN security and privacy is quite a new and challenging research area, and as such, the proposals discussed in this article are just a starting point. It is then difficult, given the state of the art, to provide an exhaustive summary of all the possible future trends and research directions. However, some general considerations can be done on the main open issues with respect to the topics discussed in the previous sections.

First of all, it is clear that privacy-preserving social network analysis and privacy-aware access control address WBSN privacy from two different points of view: the former, from an *external* perspective—that is, the one of an

analyst carrying on social network analysis; the latter, from an *internal* perspective—that is, the point of view of WBSN members themselves. In privacy-preserving data mining, the goal is to provide, on average, an acceptable degree of privacy (e.g., by using anonymization techniques) to all the WBSN members to whom the data refer. By contrast, in privacy-aware access control, each WBSN member can explicitly state his or her privacy and/or access control requirements—for instance, some members may have stricter privacy requirements than others. This means that potential conflicts between social network analysis tools and WBSN privacy requirements may arise. Therefore, in the future it is desirable that these two research directions find some common points, in order to proceed towards the definition of a comprehensive framework, able to address all the privacy and security requirements of WBSNs. It must also be taken into account that social network analysis is carried out based on the assumption that the social network management system is able to release periodically, or upon demand, the network graph (or a perturbed version of it). This implies that the existing relationships must be stored in a central repository, accessible by the social network management system itself. However, this is not always the case. For instance, privacy protection mechanisms enforced in a WBSN might adopt approaches according to which relationship information is stored by WBSN members themselves, to avoid that the social network management system infers private information from the existing relationships. Therefore, privacy-preserving data mining tools must also take into account the different architectures according to which access control is enforced.

As far as privacy-aware access control is concerned, we argued in the previous sections that, when adopting a relationship-based approach to specify access control requirements, it is necessary at the same time that access to relationship information is regulated by proper protection mechanisms. The strategy proposed by Carminati et al. (2007) addresses this issue, but other solutions are also possible. For instance, instead of assuming that relationship information is directly distributed by the WBSN members involved in them, as in Carminati et al. (2007), an alternative is to support negotiations and privacy policies, similar to those provided by P3P (Cranor et al., 2006) and trust negotiation mechanisms. According to this approach, relationship information is held by WBSN members and released upon request after having verified whether the requestor satisfies given privacy protection policies, and/or whether he or she can be considered trustworthy about the use he or she will make of such information and the protection he or she can assure to it.

Content-based access authorizations are one of the other open issues. By using content-based access control, it is possible to simplify the task of policy specification as well as to express access control requirements related to the semantics of the protected objects. However, applying it to distributed environments such as WBSNs, where any member can use any vocabulary and any language (either standard or user defined) for describing resources, might make such strategy ineffective for access control purposes. Using content analysis tools has similar drawbacks, since, independently of the efficiency and effectiveness of the adopted tools, it may happen that a given resource is incorrectly described, thus granting unauthorized access to it or denying access to authorized members. Finally, the trade-off between accuracy and complexity in describing resources must be taken into account. Inaccurate and ambiguous descriptions are useless for access control purposes, but evaluating too complex descriptions may have computational costs that make unfeasible, in practice, the enforcement of content-based authorizations.

We think that a solution to this issue must satisfy two main requirements. First, resource descriptions should be encoded by using standard schemes, and the vocabularies used for describing resources must enforce semantic interoperability. Second, mechanisms should be devised that are able to confirm the actual validity of a description.

As far as the former issue is concerned, a possible solution might be provided by the outcome of the work currently carried on by the W3C working group named, "Protocol for Web Description Resources" (POWDER, 2007), which aims at defining a standard metadata format for describing the content/characteristics of a group of resources. In addition, POWDER aims at granting the accountability of such descriptions, referred to as *description resources* (DRs, for short), thus making any Web user able to verify their trustworthiness. Finally, DRs provide a simple mechanism for enforcing semantic interoperability. In fact, any Web user describing a resource can state that such description, independently of how it is specified, is equivalent to one or more given DRs released by other users.

However, POWDER DRs by themselves do not ensure the trustworthiness of resource descriptions. A possible solution is to use a content analyzer to validate the description provided by a given user. However, the results of a content analyzer are reliable when applied to resources all belonging to a given content domain, which is not the case of WBSNs. An alternative is to use social networking itself in order to validate resource descriptions, by exploiting *collaborative rating*. According to this strategy, WBSN members, on one side, can express their opinions on the trustworthiness of a description, and on the other side, can specify their personal descriptions of the same resource. The result is that, for the same resource, more descriptions are available, whereas a description is associated with ratings stating whether it is trustworthy. Given the huge population of WBSNs, it is possible to collect a data set having a size suitable to perform statistical analysis, which can provide a more accurate measure of how much the claims made by a given

description can be trusted. Such an approach is currently under development in the framework of the QUATRO Plus EU project (*http://www.quatro-project.org*).

Support for different types of access privileges is another of the issues not addressed by Carminati et al. (2006, 2007) and Hart et al. (2007). As we mentioned, a key issue is the support for accountability, in order to be able to identify who performed which access operation on which resource. This is extremely important for 'write' operations, especially in collaborative environments where the members of the working group should be able to identify, for instance, who inserted/modified/deleted given portions of a shared document. Finally, it is worth noting that future WBSNs may rely on architectures different from the current one, where the WBSN management service is in charge of running almost all the supported services. In fact, from the privacy protection and access control approaches proposed by Carminati et al. (2006, 2007), it comes out that a decentralized architecture grants a more accurate protection to WBSN data. In such a scenario, WBSN members themselves store and manage their personal data, relationships, and resources, and are in charge of carrying on most of the tasks concerning relationship establishment/revocation and the enforcement of privacy and access control policies. By contrast, the WBSN management system provides just basic services, such as user registration, and may be used as a common space from which it is possible to access all the information WBSN members wish to share publicly. Such decentralized architectures pose challenging research issues with respect to security and privacy protection as well as efficiency.

CONCLUSION

With the increasing diffusion and usage of online social networks, protecting personal data and resources of their members is becoming a fundamental issue. Contributions to this research area are currently very limited, and can be grouped into two main classes: on one side, anonymization techniques able to protect the privacy of social network members when performing social network statistical analysis, and on the other side, privacy-aware access control mechanisms, making social network members able to regulate access to their data, relationships, and resources by, at the same time, protecting the privacy of their relationships. The proposed solutions are far from addressing all the privacy and security requirements of social networks, and not all the potential approaches have been investigated. Although it is difficult to predict with enough precision the possible evolution of this new research area, it is very likely that the enforcement of security and privacy mechanisms for social networks, more sophisticated than the ones currently available, may have two main relevant results. First, it might lead to the development of new security paradigms able to address the distributed nature of social networks. Moreover, it might determine a dramatic modification of current online social networks into a decentralized architecture, where the management of social network information, and of the social network itself, will be carried out collectively by its members inside a collaborative framework.

REFERENCES

Adam, N.R., Atluri, V., Bertino, E., & Ferrari, E. (2002). A content-based authorization model for digital libraries. *IEEE Transactions on Knowledge and Data Engineering, 14*(2), 296-315.

Agarwal, S., Sprick, B., & Wortmann, S. (2004). Credential-based access control for Semantic Web services. *Proceedings of the AAAI Spring Symposium on Semantic Web services*. Retrieved from http://www.aifb.uni-karlsruhe.de/WBS/sag/papers/Agarwal_Sprick_Wortmann-CredentialBasedAccessControlForSemanticWebServices-AAAI_SS_SWS-04.pdf

Backstrom, L., Dwork, C., & Kleinberg, J. (2007). Wherefore art thou R3579X? Anonymized social networks, hidden patterns, and structural steganography. *Proceedings of the 2007 World Wide Web Conference.*

Barnes, S.B. (2006, September). A privacy paradox: Social networking in the United States. *First Monday, 11*(9). Retrieved from *http://www.firstmonday.org/issues/issue11 9/barnes*

Bertino, E., & Sandhu, R. (2005). Database security—concepts, approaches, and challenges. *IEEE Transactions on Dependable and Secure Computing, 2*(1), 2-19.

Brickley, D., & Miller, L. (2005, July). *FOAF vocabulary specification* (RDF vocabulary specification). Retrieved from http://xmlns.com/foaf/0.1

Carminati, B., Ferrari, E., & Perego, A. (2006). Rule-based access control for social networks. *Proceedings of the OTM 2006 Workshops* (pp. 1734-1744). Berlin: Springer-Verlag.

Carminati, B., Ferrari, E., & Perego, A. (2007). Private relationships in social networks. *Proceedings of the ICDE 2007 Workshops* (pp. 163-171). IEEE CS Press.

Cranor, L., Dobbs, B., Egelman, S., Hogben, G., Humphrey, J., Langheinrich, M. et al. (2006, November). *The platform for privacy preferences 1.1 (P3P1.1) specification* (W3C working group note). *Proceedings of the World Wide Web Consortium.* Retrieved from http://www.w3.org/TR/P3P11

Davis, I., & Vitiello, E., Jr. (2005, August). *RELATIONSHIP: A vocabulary for describing relationships between people*

(RDF vocabulary specification). Retrieved from http://purl.org/vocab/relationship

Ding, L., Zhou, L., Finin, T., & Joshi, A. (2005). How the Semantic Web is being used: An analysis of FOAF documents. *Proceedings of the 38th Annual Hawaii International Conference on System Sciences* (HICSS'05) (p. 113.3). IEEE CS Press.

Ferraiolo, D.F., Kuhn, D.R., & Chandramouli, R. (Eds.). (2003). *Role-based access control.* Norwood MA: Artech House.

Freeman, L.C. (2004). *The development of social network analysis: A study in the sociology of science.* BookSurge.

Frikken, K.B., & Golle, P. (2006). Private social network analysis: How to assemble pieces of a graph privately. *Proceedings of the 5th ACM Workshop on Privacy in Electronic Society* (WPES 2006) (pp. 89-98).

Golbeck, J.A. (2004). *The Trust ontology* (OWL vocabulary). Retrieved from http://trust.mindswap.org/ont/trust.owl

Golbeck, J.A. (2005). *Computing and applying trust in Web-based social networks.* Unpublished Doctoral Dissertation, University of Maryland, USA. Retrieved from http://trust.mindswap.org/papers/GolbeckDissertation.pdf

Golbeck, J.A., & Hendler, J. (2006). Inferring binary trust relationships in Web-based social networks. *ACM Transactions on Internet Technology, 6*(4), 497-529.

Hart, M., Johnson, R., & Stent, A. (2007). More content—less control: Access control in the Web 2.0. *Proceedings of the Web 2.0 Security & Privacy 2007 Workshop.* Retrieved from http://seclab.cs.rice.edu/w2sp/2007/papers/paper-193-z6706.pdf

Hay, M., Miklau, G., Jensen, D., Weis, P., & Srivastava, S. (2007, March). *Anonymizing social networks.* Technical Report No. 07-19, University of Massachusetts Amherst, USA. Retrieved from http://www.cs.umass.edu/~mhay/papers/hay-et-al-tr0719.pdf

Korolova, A., Motwani, R., Nabar, S.U., & Xu, Y. (2008). Link privacy in social networks. *Proceedings of the 24th International Conference on Data Engineering* (ICDE'08).

Palomar, E., Estevez-Tapiador, J.M., Hernandez-Castro, J.C., & Ribagorda, A. (2006). Certificate-based access control in pure P2P networks. *Proceedings of the 6th IEEE International Conference on Peer-to-Peer Computing* (P2P'06) (pp. 177-184). IEEE CS Press.

POWDER. (2007). *Protocol for Web Description Resources working group.* Retrieved from http://www.w3.org/2007/powder

Rafaeli, S., & Hutchinson, D. (2003). A survey of key management for secure group communication. *ACM Computing Surveys, 35*(3), 309-329.

Staab, S., Domingos, P., Mika, P., Golbeck, J., Ding, L., Finin, T. W. et al. (2005). Social networks applied. *IEEE Intelligent Systems, 20*(1), 80-93.

Sweeney, L. (2002). *k*-anonymity: A model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems, 10*(5), 557-570.

Thompson, M., Johnston, W., Mudumbai, S., Hoo, G., Jackson, K., & Essiari, A. (1999). Certificate-based access control for widely distributed resources. *Proceedings of the 8th USENIX Security Symposium.* Retrieved from http://dsd.lbl.gov/~mrt/papers/AkentiUsenixSec.pdf

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications* (vol. 8). Cambridge: Cambridge University Press.

Watts, D.J. (2003). *Small worlds: The dynamics of networks between order and randomness.* Princeton, NJ: Princeton University Press.

Weitzner, D.J., Hendler, J., Berners-Lee, T., & Connolly, D. (2006). Creating a policy-aware Web: Discretionary, rule-based access for the World Wide Web. In E. Ferrari & B. Thuraisingham (Eds.), *Web and information security* (pp. 1-31). Hershey, PA: Idea Group.

Winslett, M., Ching, N., Jones, V.E., & Slepchin, I. (1997). Using digital credentials on the World Wide Web. *Journal of Computer Security, 5*(3), 255-266.

Zheleva, E., & Getoor, L. (2007). Preserving the privacy of sensitive relationships in graph data. *Proceedings of the 1st ACM SIGKDD International Workshop on Privacy, Security, and Trust in KDD* (PinKDD 2007). Retrieved from http://www-kdd.isti.cnr.it/pinkdd07/Zheleva PinKDD07.pdf

KEY TERMS

Edge Perturbation: Graph anonymization technique aimed at hiding the actual social network relationships by performing a set of random edge deletions/insertions in the network graph.

Graph Anonymization: Technique aimed at hiding private information about social network members when performing social network analysis. Node anonymization and edge perturbation are the two main graph anonymization techniques currently used.

Node Anonymization: Graph anonymization technique aimed at hiding social network members' identities by labeling the corresponding nodes with random identifiers (naïve anonymization), or, in case nodes are associated with attributes which can be used to identify the corresponding user, by using techniques based on k-anonymity (Sweeney, 2002).

Privacy-Aware Access Control: In the context of social networks, denotes an access control paradigm where access control requirements of social network members are enforced without disclosing private information about the relationships they participate in.

Relationship Trust Level: In a social network, denotes the value associated with a trust relationship, providing a measure of how much a given member considers another member trustworthy. Depending on the purpose for which it is used, this notion may have different meanings. For instance, in a collaborative rating environment, it denotes how much a given member trusts the opinions of another member with respect to a specific topic (*topical trust*) or in general (*absolute trust*) (Golbeck & Hendler, 2006). By contrast, in an access control context, it has some similarities to the notion of *security level* used in mandatory access control models (Carminati et al., 2006, 2007).

Relationship-Based Access Control: An access control paradigm specifically tailored to social networks, according to which social network members authorized to access a given resource are denoted in terms of the relationships they must participate in to get the access.

Social Network: A *small-world network* (Watts, 2003) consisting of a set of individuals (persons, groups, organization) connected by personal, work, or trust relationships. Usually modeled as a graph, where nodes correspond to social network members, whereas edges denote the relationships existing between them.

Social Network Analysis: A discipline aimed at collecting statistical data from the analysis of social network topology (Wasserman & Faust, 1994; Freeman, 2004).

Social Network Relationship: A relationship concerning two members of a social network. In WBSNs, besides personal/work relationships (e.g., friend/colleague), also trust relationships may be supported which denote how much a one member trusts another. In the graph representation of a social network, relationships are usually denoted by edges, labeled with a relationship type and/or a relationship trust level.

Web-Based Social Network: A Web-based system that allows its registered members to establish relationships with other members and to share different types of information (e.g., personal data, contacts, multimedia resources). A precise, but not normative definition of Web-based social network has been provided by Golbeck (2005).