Tutorial: Trust and Reputation in and Across Virtual communities

Nurit Gal-Oz
Telekom Innovation Laboratories
Department of Computer Science
Ben-Gurion University
Beer-Sheva, Israel
Department of Computer Science
Sapir Academic College, Israel
galoz@cs.bgu.ac.il

Ehud Gudes
Department of Computer Science
Ben-Gurion University
Beer-Sheva, Israel
ehud@cs.bgu.ac.il

ABSTRACT

Trust and Reputation systems have become key enablers of positive interaction experiences on the Web. These systems accumulate information regarding activities of people or peers in general, to infer their reputation in some context or within a virtual community. Reputation information improves the quality of interactions between peers and reduces the effect of fraudulent members. In this tutorial we motivate the use of trust and reputation systems and survey some of the important models introduced in the past decade. Among these models, we present our work on the knot model, which deals with communities of strangers. Special attention is given to the way existing models tackle attempts to attack reputation systems. In a dynamic world, a person or a service may be a member of multiple communities and valuable information can be gained by sharing reputation of members among communities. In the second part of the tutorial, we present the CCR model for sharing reputation across virtual communities and address major privacy concerns related to it. In the third part of our talk, we discuss the use of reputation systems in other contexts, such as domain reputation for fighting malware, and outline our research directions on this subject.

Categories and Subject Descriptors

K.4.1 [**Public Policy Issues**]: Privacy; K.4.4 [**Electronic Commerce**]: Security; H.2.0 [**General**]: Security, integrity, and protection; H.5.3 [**Group and Organization Interfaces**]: Web-based interaction

General Terms

Trust and Reputation

EDBT/ICDT '13, Mar 18-22 2013, Genoa, Italy
ACM 978-1-4503-1597-5/13/03.

Keywords

Reputation models, Virtual Communities, Cross-Community Reputation

1. OVERVIEW

The evolution of virtual communities in the past decade has transformed them into the new real-life platform for sharing information and opinions, and for consulting with experts. Like all other areas of our life, we aim to make them a better place to be. The existence of easily accessible virtual communities makes it both possible and legitimate to communicate with total strangers. We can now anonymously interact with other virtual community members whom we do not really know in ways that break the boundaries and limitations of the real world. However in the real world, people consider interacting or getting a service from others (person, company, web site, etc.), based on the trust they have for them. Such trust may be built from personal experience of interaction with the target entity or from the experiences of other members, which are trusted by the service requesting person. When such experiences are not readily available, one often relies on reputation, which is an aggregated perception of the community on the trustiness of the target entity. Thus, basing trust on reputation is quite common and computing reputation to capture a community's viewpoint is an important challenge.

Trust and reputation systems have major roles in improving our overall experience in virtual communities. These systems have become essential precautionary components for community users to help regulate their communication with total strangers. Reputation has also become a key component of several commercial systems such as E-bay [2]. In the past decade, quite a few models for trust and reputation were developed. Different models use different conceptual frameworks including simple summation or average of ratings; bayesian systems, belief models such as the Beta reputation model [11] which enables the representation of uncertainty in rating; flow models in which the concept of transitive trust is central such as Eigen-trust [13] and Pagerank [15]; group-based models in which trust is computed based on a subset of the community members such as the knot model [6].

The first part of the tutorial is dedicated to surveying the above major trust based reputation models. A major concern of the various models is their ability to detect fraud and

identify users trying to artificially increase or decrease the reputation of other users. We discuss the ability of the models to overcome attacks such as selfish peers and malicious attacks of ballot spoofing or sybil attacks.

Part of the survey includes the presentation of the knots-based model developed by the tutorial presenters. We developed the *Knot model* for obtaining trust-based reputation in communities of strangers. The knot model identifies groups of trustees we call knots. By definition, knot members are the most capable of providing reputation information to other members within their knot. Each knot may potentially represent different view points and therefore, may assign different levels of reputation to the same person or service (local reputation within knots). Using this approach, we can deal with heterogeneous communities where trust scores assigned to one may be distributed in a multi modal manner that makes any single reputation value meaningless. Thus, the advantage of knots is amplified when users do not share the popular opinion. The principle underlying the knot-aware model is that "less is more": the use of relatively small, but carefully selected subsets of the overall community's reputation data yields better results than those obtained from the full data set.

Our viewpoint in this tutorial which is reflected in the knot model, is to have the individual at the center. We are motivated by the desire to provide a safe and reliable environment that on the one hand allows users to preserve their privacy and have control over their private information. On the other hand our goal is also to enable users to rely on other people, or at least to understand the extent to which they are reliable, and to receive reputation information as close to their own point of view as possible. A person's reputation can therefore be viewed as part of one's identity. If reliable environments are enabled, individuals may use the reputation they gained in one place to promote themselves in other places, be they virtual or real.

In real-life users are active in several communities, each concerned with possibly different aspects of their lives. To protect their privacy, users may use different identities in different communities. Taking this global perspective, a major shortcomings is that user efforts to gain a good reputation in one community are not utilized in other communities they are active in. The quality of a community as a reputation provider is limited when the reputations of new users or inactive users is required. Therefore the need for transferring and sharing reputation between communities or their combination arises.

We developed the Cross-Community Reputation (CCR) model for the sharing of reputation knowledge across virtual communities [5, 8]. Despite its importance, cross community reputation has scarcely been addressed by other research. The CCR model is aimed at leveraging reputation data from multiple communities to obtain more accurate reputation. It also deals with the differences in the reputation attributes and reputation computation model between the various communities. It enables new virtual communities to rapidly mature by importing reputation data from related communities. At the same time, users do not have to build their reputations from scratch when joining a new community.

One of the important goals associated with sharing reputation between communities is dealing with privacy. Within the CCR model, we identified three major privacy concerns that are not present or that are less significant in single community domains. Unlinkability is a primary concern raised by the CCR model. Although we aim to compute a user's CCR from several communities, we provide the means to do so without compromising the user's anonymity in each community and while upholding the requirement of unlinkability between the communities. Controlling the dissemination of reputation information is another privacy requirement. We present a policy-based approach that enables both the users and the communities to have control over the dissemination of reputation data. The third privacy issue we address is the tradeoff between privacy and trust. There is an inherent conflict between trust and privacy. The more information disclosed about users, the more we can trust their opinions and intentions. However, this scenario leads to a simultaneous decrease in their privacy. Understanding this tradeoff is a valuable resource that helps the user make rational decisions. We suggest the transparency measure for evaluating CCR objects. To attain a high transparency rank, members are encouraged to disclose their reputation-related information whenever it is clear that disclosing their information is preferable and more valuable to them than the potential impairment of their privacy.

The second part of the tutorial presents the CCR model, discusses in detail the relevant privacy issues and their possible solutions, and briefly describes TRIC, a reference infrastructure for computing and exchanging reputation. TRIC is concerned primarily with integrating different reputation mechanisms across communities and protecting user rights to privacy and control over data In TRIC we focus on devising some multifunctional architectural guidelines that can be implemented in more than one way.

The third part of the tutorial discusses challenges facing computing reputation in other contexts e.g., computing the reputation of domains to identify malicious ones and fight possible cyber attacks.

Detecting malicious domains in real time is difficult, however it is common to collect information based on past behavior of domains and assign bad reputation score to misbehaving domains in the form of black lists. In the same manner high reputation score can be assigned to legitimate domains (e.g.,pre-trusted domains) in the form of white lists. However, since attackers are also aware of the common white and black lists they tend to attack domains which are not in any of the lists. Therefore one needs to evaluate the reputation of domains not appearing in such lists. A domain's reputation can be inferred from understanding its relations to other domains and IPs for which reputation is known. Relations must reflect some sort of trust between any two related entities. A reputation score assigned to a domain can be viewed as the probability that the domain name is a legitimate one.

Although the basic idea of inferring reputation from trust relations is relevant in this setting, the technical challenges and the computational models required are quite different then those related to virtual communities. We outline the differences and some of the major characteristics of malicious domains and their connections to other domains. We briefly review one model for assigning reputation to domains which is based on statistical features of related IPs and domains [1]. We close the tutorial by presenting some open issues for further research.

2. GOAL

Database security is a growing concern especially when it comes to unauthorized access to data and data based services. The growing amount of data collected and the various channels for sharing the data across the net, make it vulnerable to a wide range of attacks. When the data is shared by community members another concern is added and this time from the users perspective, related to the trust they have in the data provider. While authentication and access control are considered "hard security" mechanism for databases, trust and reputation systems have become prime "soft security" mechanism for online communities.

Extended database technology today includes accessing various resources such as web services and web applications, and protecting the clients of these services, and the services themselves from various malwares and other attacks is essential. Trust and reputation are often used to address user protection aspects that are not covered by traditional security mechanisms. The tutorial goal is to provide an overview of the major models in the field and open up some new and challenging issues like cross-community reputation and Domain reputation.

3. TUTORIAL OUTLINE

In this tutorial we cover of the following topics:

1. Background and Motivation [16, 19]
   (a) Virtual Communities
   (b) Trust, Reputation
   (c) Properties of Trust
2. Trust and Reputation computation Models [12, 14]
   (a) Aggregation models
   (b) Baysian and Belief Models [11]
   (c) Flow models [13, 15]
   (d) Group based models and the Knot Model [6, 18]
   (e) Analyzing the models in terms of accuracy, utility and ability to detect fraud
3. Reputation across communities [3, 5, 8, 7, 9, 10]
   (a) Motivation and challenges
   (b) The CCR (cross community reputation) Model
   (c) Privacy concerns of sharing reputation information and their possible solutions
   (d) The TRIC (trust and reputation infra-structure) system [4]
4. Domain Reputation and trust in other contexts and open issues [1, 17]

4. TARGET AUDIENCE

The audience expected are the typical attendees of the EDBT conference. No special background is expected except for general knowledge of artificial intelligence basic terms and models. Even for people familiar with some of the better known reputation models, the information on the knots and CCR models is probably new.

5. REFERENCES

[1] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster. Building a dynamic reputation system for dns. In *USENIX Security Symposium*, pages 273–290, 2010.

[2] C. Dellarocas. Analyzing the economic efficiency of ebay-like online reputation reporting mechanisms. In *ACM Conference on Electronic Commerce*, pages 171–179, 2001.

[3] N. Gal-Oz, T. Grinshpoun, and E. Gudes. Sharing reputation across virtual communities. *Journal of Theoretical and Applied Electronic Commerce Research*, 5(2):1–25, 2010.

[4] N. Gal-Oz, T. Grinshpoun, and E. Gudes. TRIC: An infrastructure for trust and reputation across virtual communities. In *Proceedings of the ICIW The Fifth International Conference on Internet and Web Applications and Services (ICIW'2010)*, pages 34–41, 2010.

[5] N. Gal-Oz, T. Grinshpoun, E. Gudes, and A. Meisels. Cross-community reputation: Policies and alternatives. In *Proceedings of the International Conference on Web Based Communities (IADIS - WBC2008)*, 2008.

[6] N. Gal-Oz, E. Gudes, and D. Hendler. A robust and knot-aware trust-based reputation model. In *Proceedings of the 2nd Joint iTrust and PST Conferences on Privacy, Trust Management and Security (IFIPTM'08)*, pages 167–182, Trondheim, Norway, June 2008.

[7] N. Gilboa, N. Gal-Oz, and E. Gudes. Schemes for orivately computing trust and reputation. In *Proceedings of the 4th Joint iTrust and PST Conferences on Privacy, Trust Management and Security (IFIPTM'10)*, pages 1–16, Morioka, Japan, 2010.

[8] T. Grinshpoun, N. Gal-Oz, A. Meisels, and E. Gudes. CCR: A model for sharing reputation knowledge across virtual communities. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology(WI'09)*, pages 34–41. IEEE, 2009.

[9] T. Grinshpoun, N. Gal-Oz, A. Meisels, and E. Gudes. Ccr: A model for sharing reputation knowledge across virtual communities. In *Proceedings of the 2009 IEEE/WIC/ACM International Conference on Web Intelligence (WI'09)*, pages 34–41, 2009.

[10] E. Gudes, N. Gal-Oz, and A. Grubshtein. Methods for computing trust and reputation while preserving privacy. In *Proceedings of the 23rd Annual IFIP WG 11.3 Working Conference on Data and Applications Security XXIII*, pages 291–298, 2009.

[11] A. Jøsang and R. Ismail. The beta reputation system. In *Proceedings of the 15th Bled Electronic Commerce Conference*, volume 160, pages 17–19, 2002.

[12] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, 2007.

[13] S. Kamvar, M. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th international conference on World Wide Web(WWW2003)*, pages

640–651. ACM, 2003.

[14] M. Kinateder, E. Baschny, and K. Rothermel. Towards a generic trust model - comparison of various trust update algorithms. In *iTrust*, Lecture Notes in Computer Science, pages 177–192. Springer Berlin / Heidelberg, 2005.

[15] J. X. Parreira, D. Donato, S. Michel, and G. Weikum. Efficient and decentralized pagerank approximation in a peer-to-peer web search network. In *Proceedings of the 32nd internatiProceedings of the 32nd international conference on Very large data basesonal conference on Very large data bases*, pages 415–426, 2006.

[16] H. Rheingold. *The virtual community*. MIT Press Cambridge, MA, 2000.

[17] S. Sinha, M. Baile, and F. Jahanian. Shades of grey: On the effectiveness of reputation-based blacklists. In *MALWARE*, 2008.

[18] H. Tian, S. Zou, W. Wang, and S. Cheng. A group based reputation system for p2p networks. *Autonomic and trusted computing*, pages 342–351, 2006.

[19] E. Wenger. *Communities of practice: Learning, meanings, and identity*. Cambridge University Press, 2007.