Reputation - from Social Perception to Internet Security

Ehud Gudes

Ben-Gurion University, Beer-Sheva, 84105, Israel
ehud@cs.bgu.ac.il

Abstract. Reputation is a concept that we use in many aspects of our social life and as part of our decision making process. We use reputation in our interaction with people or companies we do not know and we use it when we buy merchandise or reserve a room in a hotel. However, reputation plays also an important role in the internet society and enables us to establish trust which is essential for interaction in the virtual world. Reputation has several important aspects such as Aggregation, Identity and Transitivity which make it applicable in completely different domains. In this presentation we show the use of these aspects in several different domains and demonstrate it with our own previous and current research on reputation.

*A good name is more desirable than great riches;
to be esteemed is better than silver or gold.*
Proverbs 22:1

1 Introduction

Reputation is a key concept in our social life. Many of our day to day decisions such as which book to buy or which physician to consult with are based on Trust. This trust is based either on our own direct experience or when such direct experience is lacking, on other people (whose opinion we value) direct experience. However when no such direct or indirect experience is available we tend to rely on an aggregated opinion of a large set of people or a community which is manifested as Reputation. Reputation plays also a major role in virtual communities and social networks. Attempts to tarnish reputation in social networks have caused much damage to people in recent years (several cases of suicide have been reported as a result of tarnished reputation). So maintaining a good online reputation becomes a critical issue for both people and businesses. The existence of easily accessible virtual communities makes it both possible and legitimate to communicate with total strangers. Such interaction however must be based on trust which is usually based on personal experience. When such experience is not readily available, one often relies on reputation. Thus, computing reputation to capture a community's viewpoint is an important challenge.

Reputation has become a key component of several commercial systems such as E-bay [3]. Also, quite a few models for trust and reputation were developed.

Different models use different conceptual frameworks including simple average of ratings, bayesian systems, belief models [11] which enable the representation of uncertainty in rating, flow models in which the concept of transitive trust is central such as Eigen-trust [13] and Page-rank [16] and group-based models such as the Knot model [7]. In this presentation we discuss three important aspects of reputation and show how they are used in different domains. While the first two domains we discuss involve reputation of real-life users, the third domain deals with abstract entities, internet domains, yet computing and using reputation in this domain is similar to its use in the social domain.

The first aspect we deal with is the use of reputation as part of an **Identity**. In the social domains, reputation is an important part of a person identity, and the identity of a person determines its permitted actions. An expert programmer may gain more access rights to an open source code managed by some company, as her reputation increases. Such rights may be review or modify code at different levels. Our first domain then is the Authorization domain and the use of reputation for fine-grained access control. In Section 2 we present some models which use reputation as part of a user identity and consider it in making access control decisions.

The second aspect we examine is **Aggregation**. Most reputation computational models use some form of aggregation of ratings to compute the reputation [12]. However, such aggregation is usually done within a single community. In real-life, users may be active in several communities and to protect their privacy, users may use different identities in different communities. 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. Another shortcoming is the inability of one community to learn about the dishonest behavior of some member as identified by other communities. Thus the need arises to aggregate reputation from multiple communities. We developed the Cross-Community Reputation (CCR) model for the sharing of reputation knowledge across virtual communities [5, 6, 9]. The CCR model is aimed at leveraging reputation data from multiple communities to obtain more accurate reputation. It enables new virtual communities to rapidly mature by importing reputation data from related communities. The use of Aggregation in the CCR model is discussed in Section 3.

The third aspect we discuss is **Transitivity**, an important property of trust which has implications on the computation of reputation. It enables us to compute reputation not only from our own experience or our friends experience but also from our "friends of friends" experience, etc. Several flow models for computing reputation while practicing the transitivity property, have been published, including Eigen-trust [13] and Page-rank [16]. Our unique contribution here is in transferring these ideas to the computation of Internet domains reputation. Today's internet world is full of threats and malware. Hackers often use various domains to spread and control their malware. The detection of these misbehaving domains is difficult since there is no time to collect and analyze traffic data in real-time, thus their identification ahead of time is very important. We use

the term *domain reputation* to express a measure of our belief that a domain is benign or malicious. Computing domain reputation by using the Transitivity property and a Flow algorithm was investigated by us [15] and will be discussed in Section 3.

2 Identity Reputation and Access Control

Conventional access control models like role based access control are suitable for regulating access to resources by known users. However, these models have often found to be inadequate for open and decentralized multi-centric systems where the user population is dynamic and the identity of all users are not known in advance. For such systems, there must be, in addition to user authentication, some trust measure associated with the user. Such trust measure can be represented by the user reputation as one attribute of its identity. Chakraborty and Ray [2] presented TrustBAC, a trust based access control model. It extends the conventional role based access control model with the notion of trust levels. Users are assigned to trust levels instead of roles based on a number of factors like user credentials, user behavior history, user recommendation etc. Trust levels are assigned to roles which are assigned to permissions as in role based access control. In Trustbac, when the reputation of a user decreases because of past actions, its assignment to the original role may not be valid anymore and a new role with less permissions is assigned by the system. An example of such scenario in the digital library domain is given in [2]. The switching of roles may not be desirable in all cases. In a medical domain for example, a physician with less reputation may not lose its role as "doctor" but may lose instead some of her permissions. This dynamic assignment of permissions for the same role, based on the user reputation may be much more flexible and can prevent the proliferation of too many roles. In [14] we define this dynamic model formally and show a detailed example of its operation in the software development domain. The main observation of this is that when one considers reputation as part of the user identity, one can support much more flexible role-based models without the need to increase significantly the number of roles in the system.

3 Aggregation and Cross Community Reputation

In this section we briefly describe the way reputation is aggregated from several communities using the CCR model [9, 5]. The CCR model defines the major stages required to aggregate the reputation of a community member with the reputation of that member in other communities. The first stage determines the confidence one community has in another as a precondition for receiving reputation information from the latter. The second stage involves the conversion of reputation values from the domain values of one community to those of the other. In the third stage, a matching procedure is carried out between the sets of attributes used by the participating communities to describe reputation.

As an example, suppose there are two sport communities in which a commentator is active, one for Basketball, the other for Football. Assume that Bob a commentator likes to import (and aggregate) his reputation from the football community into the basketball community. The first stage considers the general confidence that basketball community members have for reputation computed in the football community. The second stage considers the statistical distribution of reputation values in the two communities and apply the required transformation (e.g., a very good rating in one community may only be considered "good" in the other). The third stage maps the specific attributes that are used to compute the reputation in the two communities (e.g., the attribute "prediction accuracy" in the football community may be partially mapped to the attribute "general reliability" in the basketball community). A detailed mathematical model which explains the process of the mapping and aggregation of CCR, is described in [5]. The CCR model was implemented as the TRIC software. TRIC is concerned primarily with aggregating different reputation mechanisms across communities and with protecting user rights to privacy and control over data during this aggregation. The CCR computation process [5] begins when a *requesting community* that wishes to receive CCR data regarding one of its users, sends a request to relevant *responding communities*. Communities that have reputation data of the user and are willing to share the information reply with the relevant reputation data. The received data is aggregated and assembled into an object containing the CCR data of the user in the context of the requesting community. This process is illustrated in Figure 1.

Fig. 1: Request for CCR scenario: (1): A requesting community sends TRIC a request for the CCR of a community member; (2): TRIC compiles a request and (3) submits it to all potential responding communities; (4): Responding communities submit a reputation object of the member at subject; (5): TRIC processes all reputation objects and compiles a CCR object; (6): TRIC sends the CCR object to the requesting 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. First 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. 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 issue of Privacy within the CCR model is discussed in [8].

4 Transitivity and Computing Domains Reputation

As was discussed earlier, computing domain reputation and identifying suspicious domains is a very important problem in Internet security today. Our approach to the problem [15] uses a graph of domains and IPs which is constructed from mapping information available in DNS log records. The Domain Name Service (DNS) maps domain names to IP addresses and provides an essential service to applications on the internet. Many botnets use a DNS service to locate their next Command and Control (C&C) site. Therefore, DNS logs have been used by several researchers to detect suspicious domains and filter their traffic if necessary. We take the famous expression *Tell me who your friends are and I will tell you who you are*, motivating many social trust models, into the internet domains world. Thus a domain that is related to malicious domains is more likely to be malicious as well. This Transitivity property motivates the use of a Flow algorithm. Although DNS data was used by several researchers before to compute domain reputation (see [1]), in [15] we present a new approach by applying a flow algorithm on the DNS graph to obtain the reputation of domains and identify potentially malicious ones. Computing reputation for domains raises several new difficulties:

- Rating information if exists, is sparse and usually binary, a domain is labeled either "white" or "black"
- Static sources like blacklists and whitelists are often not up-to-date
- There is no explicit concept of trust between domains which makes it difficult to apply a flow or a transitive trust algorithm.
- Reputation of domains is dynamic and changes very fast

These difficulties make the selection of an adequate computational model for computing domain reputation a challenging task. Our approach is based on a

flow algorithm, commonly used for computing trust in social networks and virtual communities. We are mainly inspired by two models: the Eigentrust model [4] which computes trust and reputation by transitive iteration through chains of trusting users and the model by Guha et al. [10] which combines the flow of trust and distrust. The motivation for using a flow algorithm is the assumption that IPs and domains which are neighbors of malware generating IPs and domains, are more likely to become malware generating as well. We construct a graph which reflects the topology of domains and IPs and their mappings and relationships and use a flow model to propagate the knowledge received in the form of black list, to label domains in the graph as malicious or suspected domains. Although we do not claim that every domain (or IP) connected to a malicious domain in our graph is malicious, our research hypothesis is that such domains(IPs) have a higher probability to become malicious. Our preliminary experimental results support this hypothesis.

The main input to the flow algorithm is the Domains/IPs graph. This graph is built from the following sources: 1) A-records: a database of successful mappings between IPs and domains, collected from a large ISP over several months. These mapping basically construct the edges between Domains and IPs. 2) Whois: a query and response protocol that is widely used for querying databases that store the registered users or assigners of an Internet resource. This database groups IPs which have similar characteristics and is therefore the base for IP to IP edges. In addition there are Domain to Domain edges which are related to similarity between domain names. 3) Feed-framework: a list of malicious domains which is collected over the same period of time as the collected A-records. This list is used as the initial "malicious" domains set. 4) Alexa: Alexa database ranks websites based on a combined measure of page views and unique site users. The initial "benign" domains is derived from this list. 5) VirusTotal: a website that provides free checking of domains for viruses and other malware. We use it to test our results as will be described below. The most difficult part in constructing the Domain/IP graph is assigning the weight on the edges, since the weight is proportional to the amount of flow on the edge. We tested several methods to assign weights which consider topologies of the graph and other factors, see [15]. Once the DNS graph is built and the sets of "benign" and "malicious" domains are extracted, the algorithm can be performed. The entire process is depicted in Figure 2.

The flow algorithm models the idea that each IP or domain distribute their reputation to IPs or domains connected to them. This is done iteratively and the reputation in each iteration is added to the total reputation of a domain or IP, with some attenuation factor. The attenuation factor is a means to reduce the amount of reputation one vertex can gain from a vertex that is not directly connected to it by transitivity. The flow algorithm is executed separately to propagate good reputation and bad reputation and then the two reputation values are combined in several manners resulting with several variations of the algorithm (see details in [15].)

Fig. 2: The process for computing the score: (1) Create the graph and assign weights represented as matrix; (2) Create the initial vector used for propagation; (3) Combine the matrix and the vector to execute the flow algorithm; (4) Get the final scores.

The important contribution of these algorithms is their ability to correctly predict future malicious domains. Although not all malicious domains are identified, a significant amount is discovered. In one of the experiments we used DNS logs over a 3 months period from which a large Domain-IP graph was constructed (close to one million nodes!), and the flow algorithm was applied to it. The results were that out of the top 1000 highly suspected domains, 30% were found to be known malicious (using VirusTotal), while in a random set of 1000 domains only 0.9% were known as malicious.

5 Conclusions

Reputation is a key concept in making decisions in our social life. In this paper we have discussed three key aspects of reputation: Identity, Aggregation and Transitivity which are important when migrating the concept of reputation from one domain to another. This was shown by briefly reviewing several research papers of ours. The main conclusion is that reputation plays a major role in a wide range of domains beside the social arena domain.

References

- [1] Antonakakis, M., Perdisc, R., Dagon, D., Lee, W., Feamster, N.: Building a dynamic reputation model for dns. In: USENIX Security Symposium. pp. 273–290 (2010)
- [2] Chakraborty, S., Ray, I.: Trustbac: integrating trust relationships into the rbac model for access control in open systems. In: Proceedings of the 11th ACM symposium on Access control models and technologies (SACMAT '06). pp. 49–58. ACM, New York, NY, USA (2006)

- [3] Dellarocas, C.: Analyzing the economic efficiency of ebay-like online reputation reporting mechanisms. In: ACM Conference on Electronic Commerce. pp. 171–179 (2001)
- [4] D.Kamvar, S., Schlosser, M.T., GarciaMolina, H.: The eigentrust algorithm for reputation management in p2p networks. In: WWW. pp. 640–651 (2003)
- [5] Gal-Oz, N., Grinshpoun, T., Gudes, E.: Sharing reputation across virtual communities. *Journal of Theoretical and Applied Electronic Commerce Research* 5(2), 1–25 (2010)
- [6] Gal-Oz, N., Grinshpoun, T., Gudes, E., Meisels, A.: Cross-community reputation: Policies and alternatives. In: Proceedings of the International Conference on Web Based Communities (IADIS - WBC2008) (2008)
- [7] Gal-Oz, N., Gudes, E., Hendler, D.: 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). pp. 167–182. Trondheim, Norway (Jun 2008)
- [8] Gal-Oz, N., Grinshpoun, T., Gudes, E.: Privacy issues with sharing reputation across virtual communities. In: Proceedings of the 2011 International Workshop on Privacy and Anonymity in Information Society, PAIS 2011, Uppsala, Sweden, March 2011. p. 3 (2011)
- [9] Grinshpoun, T., Gal-Oz, N., Meisels, A., Gudes, E.: 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). pp. 34–41. IEEE (2009)
- [10] Guha, R., Kumar, R., Raghavan, P., Tomkins, A.: Propagation of trust and distrust. In: WWW. pp. 403–412 (2004)
- [11] Jøsang, A., Ismail, R.: The beta reputation system. In: Proceedings of the 15th Bled Electronic Commerce Conference. vol. 160, pp. 17–19 (2002)
- [12] Jøsang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. *Decision Support Systems* 43(2), 618–644 (2007)
- [13] Kamvar, S., Schlosser, M., Garcia-Molina, H.: The eigentrust algorithm for reputation management in p2p networks. In: Proceedings of the 12th international conference on World Wide Web(WWW2003). pp. 640–651. ACM (2003)
- [14] Lavi, T., Gudes, E.: A dynamic reputation based rbac model. Report, The Open University Raanana Israel (2015)
- [15] Mishsky, I., Gal-Oz, N., Gudes, E.: A flow based domain reputation model. Report, Ben-Gurion University, Beer-Sheva, Israel (2015)
- [16] Parreira, J.X., Donato, D., Michel, S., Weikum, G.: Efficient and decentralized pagerank approximation in a peer-to-peer web search network. In: Proceedings of the 32nd international conference on Very large data bases. pp. 415–426 (2006)