



# Trust and Reputation Management

**Ling Liu**  
*Georgia Institute of Technology*

**Weisong Shi**  
*Wayne State University*

Internet applications have evolved from centralized and private computing platforms to distributed and collaborative networked computing systems. Collaboration is a fundamental Internet computing capability, and collaborative computing represents a significant step toward the continued promotion of openness and social intelligence as Internet systems and applications evolve. However, trust and reputation management is essential for establishing a healthy and efficient collaboration among a network of participants and players that might not have sufficient prior knowledge about each other.

Over the past decade, research in reputation-based trust systems has made remarkable progress in both theoretical foundation and practical deployment in real-world applications. A respectable number of e-commerce companies, such as Amazon, eBay, and NetFlix, have deployed reputation-based trust in ranking their products and suppliers. Such rankings are capitalized as effective incentives and low-cost mechanisms, letting e-commerce

companies enhance product marketing and sale management. Reputation-based trust systems typically build trust on system participants' social ratings, which take into account both the participant's (or computing node's) current social trust score and history of behavior in early transactions or interactions with others. Thus, applications often employ reputation-based trust to help identify and avoid malicious players, minimize the threat of dishonest or manipulative behavior, and protect a networked computing system from possible misuses and abuses.

Trust and reputation management research is highly interdisciplinary, involving researchers from networking and communication, data management and information systems, e-commerce and service computing, artificial intelligence, and game theory, as well as the social sciences and evolutionary biology. Researchers have studied some trust and reputation management problems, such as aggregation algorithms or trust-based recommender systems, in different contexts. For example, some trust models and trust-

aggregation algorithms are established based on transactions performed between participants or players in networked systems, whereas other models establish trust and reputation based on direct trust and trust-inference relationships. However, most research results to date cover only particular aspects of the whole problem space. Few have addressed how to combine and use trust and reputations established from different communities and in different contexts to further enhance the reliability and robustness of distributed, Internet-scale, data-intensive applications. It's thus not only timely but also important to "connect the dots" of previous efforts from different communities. In particular, we must study trust and reputation systems in their entirety and clearly distinguish the fundamental issues of trust and reputation management from domain-specific and application-dependent problems. Concretely, we identify three areas that present interesting and important technical challenges in effectively utilizing trust and reputation management for future Internet computing applications.

### Trust and Reputation Representation

Most existing systems represent a subject's reputation or trust as a single numerical value, often normalized between 0 and 1. This approach is simple to implement but not powerful enough in certain scenarios. For example, a seller on eBay could have earned a very good reputation on its electronic products, but might have only average performance on other products. Using a single value to represent this seller's reputation might be too coarse-grained to be useful for consumers interested in different types of products. In this case, a vector-based reputation/trust representation might be more appropriate.

Furthermore, how do we distinguish two subjects that have the same trust value when one is a newcomer and the other a misbehavior? Some researchers have recently proposed using *uncertainty* to describe the extent to which a subject can't predict another's behavior. By formally modeling such uncertainty, we might be able to separate these two types of subjects in a system. A key challenge in modeling such uncertainty is having an in-depth understanding of the various factors that introduce uncertainty and their inherent interactions and dependencies. The experiences of numerous Internet systems and applications, includ-

ing Amazon and eBay, demonstrate that trust models establishing trust based only on positive experiences have been more successful in practical deployment than those that combine positive experiences with negative ones. Positive experiences are more easily and widely accepted by many in the social and collaborative computing and communication environments. A key challenge is to study methods and techniques for obtaining and using negative experiences in a more positive manner. Meeting this challenge not only helps to strengthen trust models' attack resilience but also encourages the broader participation of social ratings and feedback, which is critical for building, managing, and utilizing trust and reputation. In short, finding effective methods to model reputation-based trust and identifying ways to leverage these rankings in real-world applications is an active and important research direction.

### Recommendation Aggregation

In real life, many people trust recommendations (also known as ratings) from reliable information sources such as their family members, friends, and experts with good reputations. Following this analogy to human social networks, researchers have shown great interest in employing recommendations in the computing world. However, whether these recommendations have the same positive effect they do in real-world social circles is an open question, due to some fundamental differences between the computing world and human society. First, relationships in the computing world are much more unpredictable and dynamic – a computing node's (or network participant's) reputation changes much more frequently in terms of both quality and lifetime.

Autonomous nodes in a networked computing system typically play two main roles: service provider and rater. The dynamics of quality can affect both roles unpredictably. On one hand, a provider's quality might not be absolutely consistent, and such inconsistency could make honest raters give inconsistent ratings themselves, thus causing the system to punish those raters unfairly. This is particularly true when the system uses the level of oscillations in ratings to detect dishonest or malicious raters. On the other hand, frequent changes in a rater's quality can make it difficult for system participants to get accurate information and

discover good service providers. Furthermore, system participants might vary in both service provider and rater quality more frequently than people do in real-world contexts.

In addition to the quality dimension, the lifetime of a node can also influence how consistent its reputation-based trust scores are. Autonomy in open computing environments lets nodes move between being online to offline at any time, which makes a node's lifetime more dynamic than a person's participation in a real-world social circle (that is, people don't move in and out of their social circles as frequently in the real world as they do online).

Another fundamental difference between the computing world and human society is the diameter of interaction. Humans have a much smaller diameter of interaction in the real world than they do in a computer environment because human sociability is quite stable over time. In the computing world, nodes or participants might need to interact with much larger set of players to acquire certain services. Also, such interactions are efficiently supported through overlay network topology maintenance protocols and message routing protocols. Most people in human society keep a much smaller circle of contacts and select those small groups based on a specific interaction context – that is, people within the same geographical region or organizational boundary interact socially more often. These differences bring uncertainty to whether a recommendation rating can play the same positive role in terms of trust and reputation and have the same effect on collaboration in an open computing environment as in human society. Understanding the impact of recommendations on trust inference is critical for effectively leveraging recommendations in trust and reputation management.

### Attack-Resilient Reputation Systems

In a trust and reputation system where participants don't have sufficient prior knowledge about each other, a node's decisions during interactions with other players are greatly affected by those players' established trustworthiness in the network. Computing nodes can use a trust and reputation system to evaluate other nodes in terms of interaction efficiency and accountability and reduce risks involved in interacting with unknown or unfamiliar nodes. However, any system can be vul-

nerable to certain attacks or compromised by malicious or dishonest attempts to manipulate reputation. Adversaries might exploit vulnerabilities in the trust system itself and launch attacks, eroding community trust and thus causing detrimental effect on nodes engaging in interactions. Examples of such attacks include replaying, unfair ratings, fake feedback, collusion, discrimination, self-promoting, whitewashing, and denial of service. Most existing trust and reputation systems employ various defense mechanisms against certain attacks. Although researchers have proposed and implemented several such defense techniques, existing systems to date typically address only a very limited set of attacks and rarely provide comprehensive solutions. We believe that the ability to design a comprehensive reputation system that's resilient to an entire collection of attacks is not only an open problem but also a big challenge.

### In this Issue

Following two consecutive international workshops on Trust and Reputation Management in massively Distributed Computing Systems (TRAM 2007 and TRAM 2008), held in conjunction with the 2007 and 2008 IEEE International Conference on Distributed Computing Systems (ICDCS), respectively, we invited researchers and practitioners to submit articles to this special issue that describe research efforts and experiences concerning the model, design, and potential applications of trust and reputation. We selected four articles that best address the challenges we see in this area.

"Trusted Cloud Computing with Secure Resources and Data Coloring," by Kai Hwang and Deyi Li, suggests using layered trust-overlay networks over cloud-based data centers to implement reputation systems. The authors propose a reputation-based trust management scheme augmented by data and software watermarking techniques to address trust and dependability issues in cloud computing.

In "The Impact of Churn on Uncertainty Decay in P2P Reputation Systems," Marc Sánchez-Artigas examines churn's impact on the perception of newcomers and identifies key factors that help uncertainty to decay slowly. Using a simple transactional model, the author illustrates the issue of newcomer slowness to diminish uncertainty in the presence of churn

on a decentralized P2P reputation system and makes several suggestions for handling newcomers in trust and reputation management.

The article "Advanced Feedback Management for Internet Auction Reputation Systems," by Tomasz Kaszuba and Albert Hupa, presents an approach that considers types of complaints and the connections between them to manage and learn from user feedback and endorse reputation systems in the context of Internet auctions. They design a taxonomy of compliant types for sellers based on a real-world dataset and then propose rating complaint types to build a reputation system.

Finally, "A Formal-Semantics-Based Calculus of Trust," by Jingwei Huang and David Nicol, addresses an essential issue in trust models – the formal semantics of trust – by conceptualizing trust from social studies, formalizing key trust concepts, modeling uncertainty, and applying quantified trust in trust decision making and risk analysis. They demonstrate their work in the context of PGP (Pretty Good Privacy).

**T**rust and reputation management as a multidisciplinary field can benefit from careful integration and exploitation of advances in artificial intelligence, distributed computing, information systems, knowledge discovery, knowledge modeling and engineering, social sciences, and economics. With the fast growth of collaborative Internet applications, such as BitTorrent for file sharing, YouTube for video sharing, and Facebook and Twitter for social networking, we envision that trust and reputation systems will play an increasingly important role in establishing effective cooperation among distributed Internet application participants. The Internet computing community at large can benefit from trust and reputation management researchers' continued efforts to address several important open issues, including how to represent trust and reputation appropriately, establish attack-resilient trust among nodes that don't have prior knowledge about one another, utilize historical data in reputation management, deal with sparsity and malicious manipulation of reputation-based trust, and exploit reputation-based trust inference for building more reliable, large-scale Internet applications. We believe that this special issue will enrich readers' understanding of the bene-

fits and challenges in trust and reputation management for future Internet applications. We trust that you will enjoy reading it. □

#### Acknowledgments

We express our gratitude to the authors of all submitted articles and the reviewers for their contribution to this special issue. We thank Michael Rabinovich, associate editor in chief, for his support as well as the content editor of IEEE Internet Computing, who made this issue possible.

---

**Ling Liu** is a professor in the College of Computing at the Georgia Institute of Technology. Her research interests include Internet computing, data management, distributed systems, and information security and privacy. Liu has a PhD in computer science from Tilburg University. Contact her at [lingliu@cc.gatech.edu](mailto:lingliu@cc.gatech.edu).

---

**Weisong Shi** is an associate professor of computer science at Wayne State University. His research interests include computer systems, mobile computing, and wireless sensor networks. Shi has a PhD in computer engineering from the Chinese Academy of Sciences. Contact him at [weisong@wayne.edu](mailto:weisong@wayne.edu).



*Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.*