Guest Editorial
Special Issue on Privacy and Trust Management in Cloud and Distributed Systems

WITH the continuous drive towards availability of data and services anytime anywhere, privacy risks have significantly increased. We are now witnessing vast amounts of privacy-sensitive data being collected, sometimes unintentionally, in a myriad of different types of networks. Unauthorized disclosure, modification, usage, or denial of access to these data may result in high human and financial costs. In the distributed computing environments, trust plays a crucial role in mitigating such privacy risks by guaranteeing meaningful interactions, data sharing, and communications. When the actors (machine or human) in the networks do not know how to trust a piece of data or other actors, they can be either naïve or paranoid. The naïve actors will be exposed to privacy risks and security threats, whereas the paranoid actors will ignore opportunities and services otherwise available. Trust management is, thus, an enabling technology for security and privacy enhancement. While privacy preservation and trust management are already challenging problems, it is imperative to explore how privacy-oriented and trust-oriented approaches can work together and bring new solutions to protecting information sharing and critical cyber infrastructure. Furthermore, there are questions about whether existing trust models and privacy preserving schemes are robust against various types of malicious attacks.

Privacy and trust can be meaningfully achieved only in the context of an application domain. Privacy and trust are playing critical roles in ensuring reliability, usability, and security of many emerging cloud and distributed system applications. Recent advances in cloud computing promise essentially an unlimited amount of resources to individual users, prompting a gradual migration of all types of personal digital content to the cloud. Social media networks, from collaborative projects like Wikipedia, blogs and microblogs like Twitter, to content communities like YouTube and social networking sites like Facebook, enable the creation and exchange of user-generated content and dramatically expand the information flow from enterprises to individuals. Small- and large-scale sensor networks, from global positioning systems, video surveillance networks, smart grid to the umbrella class of the Internet of Things, capture a wide range of analog information including locations, movements, appearance, power usage, grid stability, and other environmental measurements. Diverse sensitive information is available in and transmitted across different complex heterogeneous networks. Privacy preservation and trust management remain open questions in all of these applications. The codevelopment of privacy assurance, trust management, usability, and scalability of networked systems is essential to the healthy growth of such distributed system applications.

I. Papers in This Issue

The goals of this Special Issue are to amass state-of-the-art work across various efforts in privacy preservation and trust management in distributed systems, and to provide readers a single point of reference for future research. As a result of the broad nature of this effort, we have received 93 submissions and all of them were reviewed in accordance with the transactions policy. Thirteen regular papers have been accepted in this Special Issue covering three major areas including privacy enhanced technology, trust and reputation, as well as applications in cloud computing environments.

A. Privacy Enhanced Technology

Even though research on privacy enhanced technology (PET) began more than 20 years ago, the increase in diversity among content types from textual data to location information and the migration of platforms from centrally managed systems to distributed environments provide ample challenges to PET researchers. Six papers are selected under this category, addressing topics including theoretical foundation of privacy, content-specific PETs, and information sharing in distributed environments.

All PETs must be based on a strong theoretical foundation on privacy guarantees, which are often at odds with the desired utility of the information system. In the first paper titled "Utility-Privacy Tradeoff in Databases: An Information-Theoretic Approach," Sankar, Rajagopalan, and Poor introduce a rigorous information-theoretic framework, in the context of both categorical and numerical data sources, for analytically quantifying the tradeoff between the conflicting objectives of maximizing data utility and data privacy.

The next three papers address algorithmic designs in protecting privacy of various types of content. In the second paper

Digital Object Identifier 10.1109/TIFS.2013.2259431

titled "Automatic General-Purpose Sanitization of Textual Documents," Sánchez, Batet, and Viejo propose a simple but effective information-content-based scheme that automatically redacts sensitive data from documents without adversely impacting the semantics conveyed in these documents. Location privacy is the subject for the third paper titled "A Novel Privacy Preserving Location-Based Service Protocol with Secret Circular Shift for k-NN Search" by Lien *et al.* They propose using secret circuit shift and encrypted-domain operations to effectively hide the location information of the user. Gao *et al.* in their paper "TrPF: A Trajectory Privacy-Preserving Framework for Participatory Sensing" go beyond pure location privacy. They present a graph-theoretic mix-zones model to protect trajectories of participators in a sensing network. Effectiveness of their proposed model is evaluated with simulations in terms of privacy-level and information-loss metrics.

The last two papers in this category have to do with how sensitive information can be shared and brokered in distributed environments. In the fifth paper titled "Enforcing Secure and Privacy-Preserving Information Brokering in Distributed Information Sharing," Li *et al.* propose a privacy-preserving information brokering system in a distributed information sharing scenario. Using a novel automaton segmentation scheme, in-network access control, and query segment encryption, the authors demonstrate that their system can resist several formally modeled attackers and empirically verify the efficiency of their system. The sixth paper titled "Distributed Architecture With Double-Phase Microaggregation for the Private Sharing of Biomedical Data in Mobile Health," by Solanas, Martínez-Ballesté, and Mateo-Sanz, demonstrates that double-phase multivariate microaggregation can provide highly desirable privacy and efficiency characteristics in the context of healthcare services provisioned through mobile devices.

B. *Trust and Reputation*

In distributed computing environments, trust plays a crucial role in mitigating the risk by guaranteeing meaningful interactions while reputation provides the basis of evaluating trust in a scalable and public manner. In this Special Issue, there are four papers addressing trust modeling and reputation management in various applications.

Automatic tagging of images is increasingly being used in social networks but a wrong or a spam tag can easily damage the integrity and reliability of the system. In the seventh paper titled "Comparative Study of Trust Modeling for Automatic Landmark Tagging," Ivanov *et al.* adopt a user-trust model based on social feedback from users of a photo-sharing system. The authors show that by propagating tags based on the trust model built on users' tagging behavior, a larger number of tags can be propagated with high fidelity. "LDTS: A Lightweight and Dependable Trust System for Clustered Wireless Sensor Networks," by Li, Zhou, and Du, presents a lightweight and dependable trust system (LDTS) for clustered wireless sensor networks. With a combination of trust decision-making, evaluation, and aggregation schemes, the authors demonstrate that their proposed system demands less memory and communication overheads than others even in the presence of malicious, selfish, and faulty clutter heads.

The next two papers are on reputation systems. In the ninth paper titled "Securing Online Reputation Systems through Trust Modeling and Temporal Analysis," Liu *et al.* propose to detect unfair ratings and malicious users in online rating systems. The proposed scheme protects online rating systems from a new angle: the combination of time domain anomaly detection and Dempster-Shafer theory-based trust computation. It achieves a significantly better performance over the state-of-the-art in terms of identifying items under attack, detecting malicious users who insert dishonest ratings, and recovering reputation scores. In the last paper under this category titled "A Decentralized Privacy Preserving Reputation Protocol for the Malicious Adversarial Model," Hasan *et al.* address an interesting and important problem in reputation systems that has not received the attention it deserves, namely the privacy protection of users providing honest feedback in order to protect them from retaliation. The solution to this problem is based on a decentralized, cryptographic system with low message complexity that is robust against a number of attacks.

C. *Applications in Cloud Computing*

Cloud Computing is undeniably the distributed-computing platform of choice in the post-PC era. Despite its popularity, privacy and trust constantly rank among the top concerns when it comes to migrating sensitive data to the cloud. Many submissions to this Special Issue aim at mitigating these concerns, and three papers have been selected for publication.

In the eleventh paper titled "Using Mussel-Inspired Self-Organization and Account Proxies to Obfuscate Workload Ownership and Placement in Clouds," Rice, Phoha, and Robinson propose an approach for mitigating two attacks in a cloud environment, namely coresidence profiling and public-to-private IP mapping for performing unauthorized surveillance and/or data extraction attacks. The solution is based on a bio-inspired mechanism, namely mussel-inspired clustering of users with similar preferences and workload characteristics that increases the effort required for an adversary to carry out a directed attack against a target set.

In "Towards Trustworthy Resource Scheduling in Clouds," Abbadi and Ruan propose a novel cloud resource scheduler, which takes into account the trustworthiness of the infrastructure. The achievement in this paper is the successful integration between trusted computing and the OpenStack cloud computing platform. The protocol leverages trusted computing modules to establish that machines have been booted into a known configuration, and can be verified by remote attestation. This property is leveraged directly by the scheduler and the authors show how this can help cloud providers and users establish trust in the operation of the infrastructure.

In the last paper titled "CAM: Cloud-Assisted Privacy Preserving Mobile Health Monitoring," Lin *et al.* address privacy concerns in mobile healthcare technologies. The authors leverage cryptographic privacy preserving techniques to allow certain computations to be performed in a cloud computing environment, without revealing any privacy-sensitive information through the infrastructure. In particular, they look at branching programs, which take as input various health data and provide recommendations for how patients can improve their overall health posture.

Acknowledgment

The guest editors would like to thank all the authors for submitting their work. This Special Issue would not have been possible without the hard work and valuable time by the referees, to whom the guest editors are greatly indebted. They would also like to thank C.-C. Jay Kuo, Editor-in-Chief of IEEE Transactions on Information Forensics and Security for his guidance in developing this Special Issue, and Deborah Tomaro, Coordinator of IEEE Signal Processing Society Publications, for her excellent support at various stages of the review process.

Sen-Ching Samson Cheung, *Guest Editor*
University of Kentucky,
Lexington, KY 40506 USA

Yan Lindsay Sun, *Guest Editor*
University of Rhode Island,
Kingston, RI 02881 USA

Karl Aberer, *Guest Editor*
École Polytechnique Fédérale de Lausanne
Lausanne, CH-1015 Switzerland

Jayant Haritsa, *Guest Editor*
Indian Institute of Science
Bangalore, 560012 India

Bill Horne, *Guest Editor*
Hewlett-Packard Laboratories
Princeton, NJ 08540 USA

Kai Hwang, *Guest Editor*
University of Southern California,
Los Angeles, CA 90089 USA