

Some Trust Issues in Social Networks and Sensor Networks

Krishnaprasad Thirunarayan, Pramod Anantharam,
Cory A. Henson, and Amit P. Sheth
*Kno.e.sis Center, Department of Computer Science and Engineering
Wright State University, Dayton, OH 45435
t.k.prasad@wright.edu, {pramod,cory,amit}@knoesis.org*

ABSTRACT

Trust and reputation are becoming increasingly important in diverse areas such as search, e-commerce, social media, semantic sensor networks, etc. We review past work and explore future research issues relevant to trust in social/sensor networks and interactions. We advocate a balanced, iterative approach to trust that marries both theory and practice. On the theoretical side, we investigate models of trust to analyze and specify the nature of trust and trust computation. On the practical side, we propose to uncover aspects that provide a basis for trust formation and techniques to extract trust information from concrete social/sensor networks and interactions. We expect the development of formal models of trust and techniques to glean trust information from social media and sensor web to be fundamental enablers for applying semantic web technologies to trust management.

KEYWORDS: Trust management, Trust metrics, Social networks, Sensor networks, Semantic web technologies, Reputation

1. INTRODUCTION

Trust and reputation are becoming increasingly important in diverse areas such as search, e-commerce, social media, semantic sensor networks, etc. While trust is local, personal and subjective, reputation is global and objective. In search, trust information can be useful in ranking result sets. In e-commerce, trust is fundamental to clients for conducting transactions with remote sellers. In social networks, it is important to determine whom to trust and on

what topic, before making a decision. In sensor networks, the trend is towards using large numbers of cheap low-quality sensors rather than a few expensive high-fidelity sensors, and relying on the middleware for aggregating, mediating, and determining trusted sensors and trustworthy sensor data. Thus, both humans and machines (collectively called *agents*) use some form of trust to make informed and reliable decisions, or resolve conflicts, before acting. As agents providing critical content and services continue to become distributed and remote from the agents that consume them, and as miscreants attempt to corrupt, subvert or attack existing infrastructure, the issue of trust aggregation, propagation, inference, and update (collectively called *management*) will continue to remain significant. Unfortunately, there is neither a universal notion of trust that is applicable to all domains, nor an explicit description of how one arrives at trust information in many situations, much less its automation.

Even if we confine ourselves to interactions among agents in narrow domains, there are still several fundamentally different notions of trust. In order to provide a concrete example to better appreciate the research issues to be addressed later, consider the following adaptation of examples from Josang et al [15] as described in [26].

1.1. Example : Trust Network

Alice may trust Bob for recommending a good car mechanic because of Bob's experiences with car problems. Bob may trust Dick to be a good car mechanic because of Bob's past experiences with Dick. On the basis of this, Alice may infer trust in Dick to be a good car mechanic upon recommendation from Bob.

Let us say that *trust scope* captures the domain/context/task/function for which the trust relationship is applicable. In general, an agent a_1 may trust another agent a_2 for agent a_2 's ability to provide good recommen-

dations in a trust scope because agent a2 is knowledgeable in that trust scope. In this case, we say agent a1 has *referral trust* in agent a2 in the *trust scope*. Similarly, an agent a1 may trust another agent a2 for agent a2's ability to perform certain task/function in a trust scope. In this case, we say agent a1 has *functional trust* in agent a2 in the given *trust scope*.

Alice may also trust Charlie for recommending a good car mechanic, and Charlie may have a negative recommendation of Dick.

An agent a1 may distrust another agent a2 because of agent a2's inability to perform certain task/function in a trust scope. In this case, we say agent a1 has *nonfunctional trust* in agent a2 in the given *trust scope*.

Furthermore, Alice may trust Bob over Charlie for recommending a good car mechanic, possibly because of Bob's extensive experience with car problems. So, even if Bob and Charlie provide conflicting recommendations on Dick, Alice may prefer Bob's recommendations over Charlie's.

Thus, for a given *trust scope*, each agent a0 may have differential referral trust among its neighbors a1, a2, a3, ..., which can be formalized using a local partial ordering relationship among neighbors of agent a0. Recall that the partial order enables us to model incomparable trust, that is, it is not necessary to be able to, or to force, a linear total order on neighboring agents with respect to trust. This ordering can be used to resolve conflicts, minimizing ambiguity. Note that presence of ambiguity requires further investigation for resolution in order to permit decision or action.

Subsequently, Alice's direct experience with Dick's incompetence as a car mechanic may lead Alice to distrust Dick, irrespective of the recommendations of Bob and Charlie.

In general, an explicit functional (nonfunctional) trust edge between agents a1 and a2 always *overrides* conflicting trust inference paths involving a sequence of referral trust edges terminated by a nonfunctional (functional) trust edge. In other words, a claim supported by trust edge embodying strict knowledge always overrides a counter claim sanctioned by longer trust paths embodying defeasible knowledge. Thus, our approach can use local partial ordering and overriding to enable trust personalization.

Alice can have referral trust in Eric due to his knowledge about good car mechanics without having a functional trust in Eric about his being a good car mechanic. Similarly, Alice can have functional trust in Eric about being a good car

mechanic without having a referral trust in Eric, possibly due to conflict of interest and competitive spirit that may lead Eric to be less than candid.

Thus, we permit agent a1 to have referral (functional) trust in agent a2 for a trust scope without having a functional (referral) trust in agent a2 for the same trust scope. That is, referral trust and functional trust are not coupled in general.

Even though Alice trusts Bob over Charlie for recommending a good car mechanic, Alice may trust Charlie over Bob for recommending a good baby sitter.

Effectively, the local partial ordering relationship on referral trust among neighboring agents a1, a2, a3, ..., of agent a0 can depend on *trust scope*, and can be different for different trust scopes.

Furthermore, our proposal enables setting *majority thresholds* for deriving functional/nonfunctional trust conclusions (for example, a majority functional trust threshold of 4 to model requiring 4 out of 5 stars, or 4 positive referrals for every negative referral on amazon.com), and remaining ambivalent if the thresholds are not crossed.

Complementing this is social media and interactions that may further provide a basis for trust formation using content analysis of informal exchanges. Popular social networking websites such as Facebook, Orkut, and Myspace, etc. link friends and provide a closely-knit forum to discuss events and opinions, get suggestions, and generally stay connected. Popular websites such as Amazon and Ebay enable sharing of product reviews and experiences with vendors. Popular microblogging websites such as Twitter enable sharing of short messages and observations (140 characters) with interested parties. Machine-based sensors can provide raw numeric data while citizen sensors can provide semantically rich symbolic information.

In Section 2, we provide brief review of existing works on trust in social networks. We also summarize our *local* framework for representing and reasoning with referral trust and functional trust that uses local partial orders to resolve conflicts and aggregate functional trust. Note that we do not catalog situation-specific issues that influence the acquisition of partial orderings, however. In Sections 3, 4, and 5, we present research issues relevant to social networks and sensor networks. Specifically, we discuss developing formal trust models, and determining and abstracting trust information from concrete data. In Section 6, we summarize our conclusions.

2. RELATED WORKS

Traditionally, trust between a pair of agents is represented as a real number between 0 and 1. This enables abstracting and unifying trust garnered from various different sources, and simplifies trust computation, such as via aggregation and propagation. However, this has its shortcomings in that it does not provide explicit semantic justification for computed trust values. Furthermore, according to Guha et al [11]: *While continuous-valued trusts are mathematically clean [23], from the standpoint of usability, most real-world systems will in fact use discrete values at which one user can rate another.* In [25, 26], we explored “realistic” models of trust based on partially ordered discrete values to specify relative trust. On the positive side, this approach provides a natural and robust representation of *relative* trust information that an agent (aggregator) has. On the negative side, it is non-trivial to automatically glean and (potentially nonmonotonically) evolve trust information on the basis of agent networks and interactions.

A number of work focus on Epinions.com dataset, where users can add other users to their “Web of Trust”, i.e., reviewers whose reviews and ratings they have consistently found to be valuable and to their “Block list”, i.e., authors whose reviews they find consistently offensive, inaccurate, or in general not valuable. Richardson et al [23] starts with Epinions user trust graph, *synthetically* generate real-valued user trust values and statement belief information using user quality parameter and user reviews data, to study the relationship between user quality and trust propagation. Massa and Hayes [20] makes a case for distinguishing (referential) hyperlinks into two categories: positive endorsement links and negative criticizing links. PageRank algorithm [4] is run on Epinions user trust graph with various combination of trust and distrust edges, to analyze the effect of added expressiveness on user rankings. Guha et al [11] encodes trust and distrust information as 1 and -1, and define four different atomic operations for propagating trust: direct propagation, co-citation, transpose trust and trust-coupling. These operations are captured via matrix operations.

Artz and Gil [1] surveys existing models of trust, different definitions of trust, trust metrics, and their specific determination using policies or reputation. Massa and Avesani [20, 22] analyze the variation in average trust values for different equivalence classes of users, determined on the basis of path length. In contrast with majority of the approaches in the literature, our work *prefers* claims supported via highly trusted edges over conflicting claims via less trusted edges, in arriving at a conclusion, rather than averaging them.

The works of Golbeck et al [9, 10, 17] infer trust using weighted average of trust information from the neighbors of the source agent, while Bintzios et al [2] infers trust from neighbors of the target agent. In general, it is hard to make quantitative and qualitative comparisons of these trust inference algorithms because they present frameworks and algorithms that contain user tunable parameters and aggregation functions.

Josang et al [15] presents a novel approach to trust that discriminates between two different forms of trust and reasoning with them. They use subjective logic to reason with opinions represented as 4-tuples (belief, disbelief, uncertainty, base rate). Even though this provides an expressive representational framework, it is unclear how quantitative opinions are obtained initially and what semantics to associate with “fine-grain” numeric values, compared to acquiring and reasoning with relative, binary trust information.

Huang and Fox [14] introduce two different types of trust: *trust in belief* and *trust in performance* that roughly parallels referral trust and functional trust respectively. Within a specific context, the former is transitive, while the latter is not. The notion of *context* is analogous to trust scope. They formalize the ontology of trust using situation calculus, distinguishing between *direct trust* and *indirect (social networks based) trust* and explicitly encoding trust formation rules. Our works [25, 26] capture semantics of trust using a set-based, model-theoretic approach. Essentially, direct trust is supported by an explicit trust edge, while indirect trust is supported by “unpreempted” paths in the trust network. (We do not separate out system trust (due to a certifying agency or standards body) from other forms of trust.)

In relation to our works, Wang et al [28] discuss three orthogonal aspects pertaining to trust networks: (1) two ways of determining referral trust — one based on *similarity* and another based on *truthfulness* between agents; (2) characterization of small world networks; and (3) approaches to encourage/remunerate agents to share recommendations and ratings.

Carroll et al [5, 6] discuss extensions to RDF to enable incorporation of trust and provenance information into Semantic Web. They also present different kinds of trust mechanisms (reputation-based, context-based and content-based) and several example trust policies that shed light on the nature of trust and that can be used to codify trust rules.

Finally, we summarize our recent works on trust. Thirunarayan et al [26] develops a computational model of referral trust and (non)functional trust among agents that

abstracts weights on edges through local partial ordering on edges w.r.t. trust scopes, and propagates it via local distributed computation. It is focussed on qualitative information that is natural and more readily available than on quantitative information. The binary approach also permits derivation of conclusions that permit taking action or decision. Our approach is robust with respect to redundant edges obtained by replacing a node with a pair of synonymously named connected nodes. The discretization of trust values, trust scope dependent partial ordering, and trust aggregation via least-upper bound operation, enables us to readily see the semantic consequences of the trust network and the computational properties such as locality, convergence, etc. Our approach differs from popular works (such as [11, 20, 23, 15]) as follows:

- We distinguish both referral trust and (non)functional trust among agents implicitly as *discrete* values.
- Our approach is sensitive to *local, relative* ordering of trust values rather than their magnitudes.
- We distinguish between *direct* trust and *inferred* trust, letting direct information override conflicting inferred information.
- We regard *equal* or *incomparable* evidence in support and against functional trust in an agent as ambiguous trust, and represent the ambiguity explicitly.

The trust networks in Thirunarayan et al [26] differs from trust-distrust-belief networks in Thirunarayan and Verma [25], in that it distinguishes between different kinds of trust and the local partial orders are parameterized w.r.t. trust scopes.

3. TRUST IN SOCIAL NETWORKS

It is not sufficient to present an abstract definition of trust. Formal computational models of trust are necessary in order to understand and apply trust information in practice.

3.1. Structure of Trust Networks

Currently, a trust network consists of nodes and edges, where nodes represent agents and edges represent trust relationships between agents. However, at the next level of detail, there are at least two extreme interpretations of these edges prevalent in the literature: In one, an edge from agent a_1 to agent a_2 is interpreted as agent a_1 trusts agent a_2 in all contexts (e.g., [10, 25]), while in another, each edge is additionally required to carry an edge label e_1 that specifies the context for which the trust relationship holds (e.g.,

[15, 26]). That is, the agent a_1 trusts agent a_2 only as far as the context (or trust scope) e_1 is concerned. The former approach provides a very compact representation and a powerful mechanism to deal with incomplete information. In many cases, such trust networks have been employed in a limited setting, making the domain implicit (e.g., movie recommendations). (In general, such trust edges can also be exploited to remedy data sparsity problems in recommender systems based on collaborative filtering [22, 19].) On the other hand, it is not sound to assume that when one agent trusts another agent, it is without further qualifications. The latter approach involving trust scopes provides more natural and expressive representation of the context by making it explicit for each trust edge.

In future, *we propose to explore trust networks that integrate these two approaches by introducing the notion of exceptions*. We want to be able to express knowledge such as 'Normally agent a_1 trusts agent a_2 ' on all topics and then provide specific exceptional contexts in which agent a_1 may prefer to rely on other agents different from agent a_2 . This approach may enable one to resolve conflicts automatically whereas the earlier approach may have treated it as ambiguous, thus promoting both expressive adequacy and notational efficacy. In terms of the topology of the trust network, we will consider tree-structured networks, followed by directed-acyclic graphs, followed by graphs with cycles, to better understand the relationship between the expressive power of the network and its consequences on the computational complexity.

The vocabulary used for trust scopes can be structured as a hierarchy. For instance, if agent a_1 trusts agent a_2 to be a good car mechanic, it is reasonable to assume that agent a_1 will also trust agent a_2 for replacing engine oil, changing automobile tires, etc. That is, one can organize edge labels that signify trust scopes using an inheritance hierarchy.

To apply this in practice, either the trust network needs to be constructed explicitly by the users of the system as exemplified by Epinions dataset, Movie Recommendation Systems, etc., or automatically gleaned from available social network datasets such as Facebook, Twitter, etc. Viljanen [27] discusses factors that can influence trust formation, such as awareness of identity, competence, business value, history, etc. Guha et al [11] discuss how to infer trust based on atomic operations such as transitivity, co-citation, transpose trust, trust coupling, etc. There are fundamental differences between traditional trust and online trust – traditional cues for trust and reputation are missing in the online world but communicating and sharing trust information is very convenient [16]. This leads us to several important research questions of practical signifi-

cance: *What aspects of the online interactions can underlie trust formation? Can we verbalize the intuitions behind them and eventually formalize its extraction from available data? How do we track online entities? Are these aspects of interactions and identities robust or can they be subverted easily?* For example, in Facebook, the *friend*-relationship between pairs of people can be thought of as “strong” bidirectional trust edge, while in Twitter, the *follower*-relationship can be thought of as “weak” unidirectional trust edge. Furthermore, textual analysis of informal communication among friends in Facebook may provide basis for creating edge labels – evolving trust edges of the first kind to trust edges of the second kind. In Twitter, tweets (short messages) posted by an user is delivered to all his/her followers, who can then re-tweet it. Re-tweets, which include information such as original author, count, etc., can be analyzed to glean trust in the author and trust-worthiness in the tweet. Note also that people have distinct identities on different networks and may have multiple identities on the same network. *For web sites that host consumer reviews (e.g., Amazon.com), can we devise content specific text analysis methods to complement manual ranking of reviews, and eventually determining trust in the reviewers?*

Eventually, trust models and data analytics discussed above can provide the foundation for standardization and application of semantic web technologies. Specifically, the trust network can be rendered in RDF, and rules and SPARQL queries can be used to formalize the necessary reasoning.

3.2. Nature of Trust

The form of trust values associated with an agent, and what it stands for, has evolved to account for desired abstraction and computational convenience.

- An agent can be given a global trust value (more appropriately called reputation), which is a real number.
- The trust value associated with an agent can depend on trust scope (e.g., authority on a topic), and various agents can be ranked on this basis (e.g., Pagerank).
- The trust values associated with a target agent can depend on the source agent, that is, ranking can be personalized (e.g., [23]).
- The trust values associated with a set of neighboring agents can depend on the source agent and the trust scope. In the most general case, this may be materialized as a partial order on a discrete set of trust values customized for each agent and for each trust scope (e.g., [26]).

- The trust values can be gleaned from a history of interactions, optionally using trust paths, by tracking the number of positive experiences and negative experiences, or through the notion of opinions to allow for uncertainty (e.g., [15]).

Furthermore, these values may be based on a local “myopic” view for scalability, or on paths for generality/expressiveness. In the local approach, how do we capture trust discounting over a path in a sound way? In the path-based approach, can we summarize path length information (e.g., an integer value) and the nature of agents on the path (e.g., experts, friends, rumours, etc) through additional attributes? That is, instead of reflecting basis for trust in the final trust values, how can we keep justification summaries explicit and separate for consumption by trust aggregation functions?

Thus, the fundamental research problem is: *How do we balance computational complexity with semantic clarity, that is, how do we craft a quantitative approach that properly reflects qualitative analysis.*

To apply this in practice, we need users to manually provide absolute or relative trust information explicitly, or to automatically determine relative trust information from available data. For example, in Epinions, the users rate other users. In Facebook, analysis of conversations between friends may provide a basis for determining relative trust relationship. *The basic difficulty lies in coming up with a machine-consumable analog that can be manually populated or extracted automatically in a consistent and natural way.*

Eventually, trust information can be incorporated in the RDF encoding of trust networks and semantic web technologies can be used to formalize the necessary reasoning such as is done for trustworthiness of data based on provenance information in tSPARQL [12].

3.3. Trust Ontology

There is no universal trust relationship. In fact, in addition to trust being context dependent, the trust itself can be of different kinds, for a given context. For instance, one may trust someone for providing good recommendations (referral trust or trust in belief), while another may trust someone for carrying out actions (functional trust or trust in performance) [14, 15]. See Section .

Orthogonal to this is the purpose for which we are defining the trust relationship. For instance, one may additionally

qualify trust as identity trust, provision trust, access trust, delegation trust, etc., to indicate the purpose [16].

We expect the kind of trust and the trust purpose to determine if and when trust information can be propagated. For example, referral trust paths can be used to derive indirect trust with trust purpose influencing the dependence of trust on path length. At the other extreme, paths involving functional trust edges or distrust edges do not make sense, and so they are not transitive. Furthermore, personal (direct) experience carries more weight than secondhand referrals in the case of conflicts. In fact, Josang et al [16] recommends relying only on direct neighbors of the source node for referrals, while Bintzios et al [2] and Kuter and Golbeck [18] recommend relying only on direct neighbors of the target node (because of their firsthand experience). Formal analysis of various trust will also lead to understanding subtleties and avoiding potential capricious behavior on trust edge updates.

Can we develop upper-level (domain-independent) ontologies for different kinds of trust (cf., referral vs functional) and for different trust purposes (cf., purchasing, reviews, etc.) that can be extended to various domains via trust scopes?

To apply this in practice, we need to understand how to determine these aspects from concrete data, and formalize them in RDF and guide querying/reasoning.

4. TRUST IN SENSOR NETWORKS

With the trend towards deploying large sensor networks built from cheap, low-quality sensors, middleware is being relied upon for determining trust in sensors, for self-configuring sensor network, and eventually, supplying trustworthy sensor data. Protecting the communication channels against eavesdroppers and adversaries is necessary but not sufficient. Sensor nodes not only use direct observations but also exchange trust information with other nodes, or central reputation repository. According to Fernandez-Gago et al [7], a node that is uncooperative or disappears off and on, or provides false or delayed response even infrequently, should not be trusted.

Approaches to determining trust in a sensor can be classified into three categories [1]:

Reputation-based Trust: Trust in a sensor is determined by considering its past behavior over time.

Redundancy in a sensor network is essential to determine when a sensor in a network is behaving abnormally. In homogeneous sensor networks, spatio-

temporal locality can be used to retrieve relevant sensor data to be fed to outlier detection algorithms for flagging abnormal sensors. (E.g., temperature sensors in close proximity should yield readings within some tolerance range.) Sensor reputation can be built by aggregating results over time. In contrast, in heterogeneous sensor networks, we need complex domain models capturing correlations among different observations for flagging abnormal sensors.

Trust in a sensor has been modeled as beta probability distribution function with parameters (a,b) (gleaned from total number of correct observations (a-1) and erroneous observations (b-1)) satisfactorily, for the following reasons:

Computational Ease: This approach requires us to retain just two values (a,b) to summarize the history, and as new data arrives, requires us to incrementally update only one value after checking whether the new data is an outlier or not.

Conceptual Simplicity: This approach does not require prior initialization, and the variation of probability density function in response to new data seems intuitively appealing and sufficiently expressive.

In fact, trust in a sensor is equated to the statistical expectation associated with Beta(a,b). In order to determine personalized trust associated with a target node, the source node can integrate its direct experience with the indirect experiences provided by other nodes [8]. To stave off bad-mouthing attacks (cf. E-commerce analogy: Sellers collude with buyers to give bad ratings to others.) and ballot stuffing attacks (cf. E-commerce analogy: Sellers collude with buyers to give it unfairly good ratings), the source node can weight the feedback on the target node given by an intermediate node by the trust in the intermediate node. Aging of trust values allows robustness with respect to sleeper attacks in which apparently trusted agent defects.

In practice, *automatic detection of abnormal sensors and abnormal sensor data, and distinguishing faulty behavior from malicious behavior are all non-trivial.* For example, consider abnormalities such as stuck-at-zero fault, or transient, non-deterministic behavior under low power conditions, where the observations are still within bounds though incorrect. Similarly, as in real life with human moles/defectors, if apparently trusted sensor suddenly shows abnormal behavior, it

is not easy to flag that purely on the basis of past behavior.

Policy-based Trust: Trust in a sensor is determined by how well its specification/characteristics satisfies explicitly stated constraints, or if it has a dependable third-party certification. In theory, this approach can detect abnormal behavior that cannot be caught by a reputation-based system.

Evidence-based Trust: Trust in a sensor is evolved by seeking/verifying corroborating evidence in its proper behavior [13]. Observations (and hypotheses) are more trusted if they can be verified through empirical evidence. Sensors are more trusted if their observations are trusted. (Active perception in sensors context can further enable actionable intelligence by narrowing set of explanations to one, and use a small set of always-on sensors to bootstrap, and selectively turn-on additional sensors on demand in a resource (e.g., power) constrained environment.)

5. INTEGRATING SOCIAL NETWORKS AND SENSOR NETWORKS

Agents carrying sensor-based devices can benefit from the integration of social networks and sensor networks [3]. For instance, alerts can be triggered for a face-to-face meeting among friends in a social network who are in close proximity, by utilizing their dynamically publicized GPS locations via mobile phones or taxi (cf. Citysense, Brighkite, etc). Signs of abnormal activity from health monitoring devices can be used to trigger alerts to family and friends on a social network in addition to health care providers. Road/bridge traffic monitoring services and weather monitoring services can be used to gauge arrival delays dynamically in the context of family get-togethers over national holidays. *There are still significant technical, social, and privacy challenges to overcome, before sensor-enabled social networks applications go mainstream.*

6. CONCLUSIONS

Previously, we developed a framework for describing semantics of trust in social networks containing referral trust, (non)functional trust, trust scopes, etc., by exploiting and adapting many evidence-based insights [25, 26]. In future, we may explore extending current local trust framework by using a richer language of trust annotations to accommodate different kinds of trusts, trust scope labels with additional structure, trust path length (propagation horizon), and more expressive trust measures (subsuming real-valued trust), etc.

There are interesting parallels between social networks and sensor networks. Agents (resp. authors) and agent statements (resp. tweets) correspond to sensors and sensor data. Statements (resp. tweets) made by trusted agents (resp. authors) are usually trustworthy (believed), and data obtained from trusted sensors are usually trustworthy. If various statements (resp. tweets) made by an agent (resp. author) can be corroborated for validity, the trust in the agent (resp. author) is enhanced. Similarly, if various sensor observations made by a sensor can be corroborated for validity, then the trust in the sensor is enhanced.

Overall, we advocate a balanced, iterative approach to trust that marries both theory and practice. On the theoretical side, we investigate models of trust to analyze and specify the nature of trust and trust computation. On the practical side, we propose to uncover aspects that provide a basis for trust formation and techniques to extract trust information from concrete social/sensor networks and interactions. We also cited simple use-cases for integrating social networks and sensor networks. We expect the development, specification and standardization of formal models of trust, and techniques and tools to extract trust information from social media and sensor web to provide the basis for applying semantic web technologies to trust management on the large scale, benefitting quality, interoperability and self-configuration.

REFERENCES

- [1] D. Artz and Y. Gil, "A Survey of Trust in Computer Science and the Semantic Web," *Journal of Web Semantics*,. Volume 5, Issue 2. pp. 58–71, 2007.
- [2] V. G. Bintzios, T. G. Papaioannou, and G. D. Stamoulis, "An effective approach for accurate estimation of trust of distant information sources in the semantic Web," Second International Workshop on Security, Privacy and Trust in Pervasive and Ubiquitous Computing, 6 pages, 2006.
- [3] Breslin et al, "Integrating Social Networks and Sensor Networks," W3C Workshop on the Future of Social Networking, January 2009.
- [4] S. Brin and L. Page, "The Anatomy of a Large-Scale Hypertextual Web Search Engine," *WWW7 / Computer Networks* 30(1-7), pp. 107–117, 1998.
- [5] J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler, "The Semantic Web Trust Layer," Developer's Day Talk at the 13th international conference on World Wide Web. 2004.
- [6] J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler, "Named graphs, provenance and trust," Proceedings of the 14th international conference on World Wide Web. pp. 613–622. 2005.

- [7] M. C. Fernandez-Gago, R. Roman, and J. Lopez, "A Survey on the Applicability of Trust Management Systems for Wireless Sensor Networks," Security, Privacy and Trust in Pervasive and Ubiquitous Computing, 2007. SECPerU 2007. Third International Workshop on Volume , Issue , 19-19 July 2007 Page(s):25 - 30
- [8] S. Ganeriwal, L. Balzano, M. B. Srivastava, "Reputation-based Framework for High Integrity Sensor Networks," *ACM Transactions on Sensor Networks (TOSN)*, Vol. 4 , No. 3. pp: 1–37. 2008.
- [9] J. Golbeck, *Computing and applying trust in web-based social networks*. Doctoral Dissertation, University of Maryland, College Park, 2005. 2005.
- [10] J. Golbeck and J. Hendler, "Inferring Trust Relationships in Web-based Social Network," *ACM Transactions on Internet Technology*. Vol. 6, Issue 4. pp. 497 - 529. 2006.
- [11] R. Guha, Ravi Kumar, Prabhakar Raghavan, Andrew Tomkins, "Propagation of Trust and Distrust," *International World Wide Web Conference (WWW2004)*, pp. 403–412, 2004.
- [12] O. Hartig, "Querying Trust in RDF Data with tSPARQL," *Proceedings of the 6th European Semantic Web Conference (ESWC'09)*, LNCS 5554. pp. 5–20. 2009.
- [13] C. Henson, Active Perception. (Informal Draft)
- [14] J. Huang, and M. S. Fox, "An ontology of trust: formal semantics and transitivity," *Proceedings of the 8th international Conference on Electronic Commerce (ICEC'06)*. pp. 259–270. 2006.
- [15] A. Josang, R. Hayward, and S. Pope, "Trust network analysis with subjective logic," *Proceedings of the 29th Australasian Computer Science Conference (ACSC-06)*, pp. 85–94, 2006.
- [16] A. Josang, R. Ismail, and C. Boyd, "A survey of trust and reputation systems for online service provision," *Decision Support Systems*, Vol. 43 No. 2, pp. 618–644, 2007.
- [17] Y. Katz and J. Golbeck, "Social Network-based Trust in Prioritized Default Logic," *Proceedings of The Twenty-First National Conference on Artificial Intelligence (AAAI-06)*, 2006.
- [18] U. Kuter and J. Golbeck, "Semantic Web Service Composition in Social Environments," *Proceedings of International Semantic Web Conference*. ISWC-2009, LNCS 5823. pp. 344–358. 2009.
- [19] H. Ma, I. King, and M. R. Lyu, "Learning to recommend with social trust ensemble," *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, 2009.
- [20] P. Massa and C. Hayes, "Page-reRank: Using Trusted Links to Re-rank Authority," *Web Intelligence Conference*, pp. 614–617, 2005.
- [21] P. Massa and P. Avesani, "Controversial Users demand Local Trust Metrics: an Experimental Study on Epinions.com Community," *Proceedings of The Twentieth National Conference on Artificial Intelligence (AAAI-05)*, pp. 121–126, 2005.
- [22] P. Massa and P. Avesani, "Trust-aware recommender systems," *Proceedings of the 2007 ACM Conference On Recommender Systems*, pp. 17–24, 2007.
- [23] M. Richardson, R. Agrawal and P. Domingos, "Trust Management for the Semantic Web," *Proceedings of the Second International Semantic Web Conference*, ISWC-2003, LNCS 2870, Springer. pp. 351–368, 2003.
- [24] S. Schenk, "On the Semantics of Trust and Caching in the Semantic Web," *Proceedings of the Seventh International Semantic Web Conference*. ISWC 2008, LNCS 5318, Springer. pp. 533–549, 2008.
- [25] K. Thirunarayan and R. Verma, "A Framework for Trust and Distrust Networks," *Proceedings of Web 2.0 Trust Workshop (W2Trust)*, June 2008.
- [26] K. Thirunarayan, D. K. Althuru, C. A. Henson, and A. P. Sheth, "A Local Qualitative Approach to Referral and Functional Trust," *Proceedings of the The 4th Indian International Conference on Artificial Intelligence (IICAI-09)*, December 2009.
- [27] L. Viljanen, "Towards an Ontology of Trust," *Proceedings of TrustBus 2005*. LNCS 3592, Springer. pp. 175-184. 2005.
- [28] Y. Wang, Y. Hori, and K. Sakurai, "Characterizing Economic and Social Properties of Trust and Reputation Systems in P2P Environment," *Journal of Computer Science and Technology*. pp. 129–140. 2008.