

Trust-Based Security Protocols

Henry Hexmoor, Sandeep Bhattaram, and Seth Wilson
 Department of Computer Science and Computer Engineering
 University of Arkansas, Fayetteville, Arkansas 72701
 {hexmoor, sbhatta, slw08}@uark.edu

Abstract –

An agent-based model of interpersonal trust is used to develop information sharing security protocols in the spirit of soft security. We demonstrate that availability information can be balanced against security breaches. Our protocols are simple, malleable, and scalable.

Index Terms – Security protocols, soft security, agents

I. INTRODUCTION

IN large law enforcement agencies, information is gathered and compiled by many individuals from independent and disparate sources. In government agencies like the FBI and CIA, it is not necessarily the information from any one source that is critical to the assessment of potential crimes or threats, but rather the totality of the information that is significant. Therefore, in order to build the intelligence and knowledge required for such agencies to avert or tackle any unlawful practices, active information sharing among the individuals and, in general, the agencies is one of the most vital activities.

One of the major concerns of information sharing is the security. Information that is obtained by an unauthorized individual, whether internal or external, may be harmful. In terms of information security, there are two elements for which concern must be given: confidentiality and integrity. Agencies implement policies that regulate and direct the sharing of information in light of these two concerns.

Many psychological and social factors need to be taken into account before assuming policies that characterize secure information sharing in and amongst agencies. Herein, we use the terms organization and agency interchangeably. These factors include time, the mobility of the modern workforce, its affects on long-term relationships, and demonstrating concern for the employee.

Most organizational policies take little notice of the impersonal nature of the work environment and communications, and mutuality in trust and respect [1]. These codified policies, thus, are ad hoc, difficult to follow, not malleable, and inadequately treat risks involved in information sharing. These problems can complicate secure information sharing to a point at which the organization ceases to share information efficiently and effectively. We have therefore embarked on developing trust-based security policies in information sharing that consider trust as the key parameter and an undeclared psychological understanding in building the organizational social capital.

This paper is organized as follows: Section II discusses the literature and other work related to our scope of interest. In section III, we present the motivation behind the choice of a trust-based approach to the problem. Section IV introduces our trust-based model and trust-based information sharing algorithm. Results and analysis of our simulations are put forth in section V. A discussion of the implications of our model ensues in section VI. In section VII, we make some concluding remarks and intentions for the future of this work.

II. RELATED WORK

A. Information Sharing

Information sharing is a critical activity for almost every institution. Herein we have mentioned its importance to the various arms of governmental justice institutions that engage in ongoing information gathering and dissemination. This topic has also been addressed by Phillips, Jr., Ting, and Demurjian in [2]. In this work, they consider the issues surrounding information sharing in dynamic coalitions pursuant to some crisis. They recognize that efficient and effective information exchange is crucial in such situations. Unlike our work, however, they combat issues of security (confidentiality and integrity) by considering the use of hard security measures such as role based, discretionary and mandatory access control, as well as cryptography.

In [3], Goecks and Cosley introduce their proprietary information sharing system called NuggetMine. This system provides facilities for the building, maintenance, and utilization of a repository of information called a *mine*. The authors stress a similar point therein concerning the importance of the assimilation of pieces of information: individually, they may seem insignificant, but collectively they may stimulate enlightenment. This work, however, does not address issues of security.

Another important influence to our work is a work in which the authors consider information dissemination in the context of a wireless sensor network [4]. Individual sensors, which represent can represent a social organization [5], in such a network aggregate data in order to provide a comprehensive and multi-dimensional view of the surrounding environment. In addition to the support provided again by the common theme of information sharing, one of the most important contributions of this work to our own is its analysis of the various dissemination algorithms. Our model adopts a variant of the classic flooding data dissemination algorithm discussed in section IV.

B. Trust and Trust-based Security

Many different definitions of trust have been espoused by various researchers in other literature. This inconsistency has led to confusion about the notion of trust. Before we move further into this discussion, we introduce the general definition of trust upon which our work is based: Agent x 's *trust* in agent y is agent x 's estimation of the probability that agent y will preserve agent x 's welfare with regard to the action to be performed [6]. We assume in our model that trust between individuals is a useful piece of information that is readily available [6], [7].

Rasmusson and Jansson in [8] propose a school of thought focusing on *soft security* methodologies in open systems that bring about a *secure social control*. In [9], a trust model based on distributed recommendations is proposed as a solution to issues of security in on-line transactions. The concept of social control as a soft security mechanism in a distributed system introduced in [8], [9] was particularly influential to our work. Additionally, the decentralized nature of trust management mirrored our intentions to support a more dynamic environment.

For the Semantic Web, trust has become an important tool for gauging the reliability of information. Authors in [8] and [10] identify the shortcomings of traditional hard security measures such as digital signatures and certificates – confirming the source of a document says little about the trustworthiness of its contents. The work of Gil and Ratnakar in [11] echoes the concern of trustworthiness in terms of the reliability and credibility of sources of information on the Web. Their trust rating system, TRELIS, is influential in spirit, but fails to suit the need for decentralized trust management.

III. MOTIVATION

A security policy based on trust has several important advantages that address the shortcomings of traditional policies. Trust is a universally understood social phenomenon and commodity [12]. It affords “naturalness” to the policy and focuses on the fact that building a larger connected human workforce begins with the strength of the relationships at the individual level. Organizational units may therefore more easily adapt to security policies based on trust.

Another important advantage is that trust-based policies mimic the fluidity of human relationships. These relationships predictably evolve, strengthening or weakening according to the established history of experiences between individuals [13].

Trust is also effective at preserving the security of information. Unlike hard security measures such as firewalls and access control, trust-based security does not offer iron-clad guarantees. Instead, we speak of levels and forms of assurance. These might be in terms of probability or possibility (i.e. Fuzzy logic).

As a social notion, trust produces a social influence among the members of an organization. These social influences encourage secure information sharing and deter breaches of

security. Collective behaviors of groups of individuals are directed using social control techniques based on social influence inference inherent in social structures such as organizations [13].

Under a trust-based security policy, sensitive information will not be shared with those insufficiently trusted, although there is a possibility that such information may be obtained by an unauthorized individual through some covert channel. Such a breach of security results in a breach of trust, which is quickly and easily remedied. Perpetually untrustworthy individuals are eliminated from the organization, and in turn the overall security of the organization is maintained.

IV. A WORKING ALGORITHM FOR TRUST-BASED INFORMATION SHARING

In our model, organizations and their constituent agents are considered social entities. Given this societal orientation, the agents we establish base their relationships and interactions upon the human resource of trust. Herein, we introduce the terminology related to the trust-based model and discuss the trust-based information sharing algorithm.

Interaction trust relation is an adjacency matrix that contains trust values between agents of an organization. The trust values, which are assumed to be known a priori, are quantified in a non-discrete manner as a real value in the range 0.0 to 1.0. A trust value of 1.0 indicates agent x 's complete trust in agent y . If there exists no relation between agent x and agent y , agent x either has lack of trust in agent y or agent x is ignorant or cannot make a trust-related judgment about agent y . We make a simplifying assumption by interpreting the aforementioned cases as the same and assigning a trust value of 0.0 to them. An agent is cognizant of its trust in its neighboring agents. However, the neighboring agents are not apprised of the trust the agent has in each of them. Interaction trust relations are assumed to satisfy the reflexivity property. However, they do not hold the symmetry and transitivity properties.

Trust-based information sharing is the exchange of information that is owned by an agent among agents in the organization in succession based on trust with an assumption that the information might satisfy at least one agent's requirement, including its originator. A *cycle* is measured as the time taken by an agent to receive or generate a fact and share it with her neighbors. Trust-based information sharing occurs continuously in each cycle and for every agent. Each agent in an organization produces a unique fact from the fact warehouse, which is a set of known and static facts. Each generated fact has a unique identifier consisting of two fields, the current cycle number and the unique initial of the initiator agent. Each agent in an organization also has a fact requirement that is selected from the fact warehouse. An agent's required fact, which is assumed to be a constant in all the cycles, and the fact produced by her in each cycle can be the same. An agent might require the fact produced by any agent (including herself). Also, an agent is not cognizant of other agents' requirements. These are few simplifying

assumptions that affect information sharing and its availability.

In trust-based information sharing an agent is *satisfied* if and only if the fact required by her is the same as the fact she has access to. An agent's *accessible facts* constitute the fact generated by her in the current cycle and the facts that she receives from her neighbors in the previous cycle, if any. In each cycle, every agent produces a unique fact and then checks for her fact requirement in the facts that are accessible to her in the current cycle. An agent is not willing to receive any other fact for the current cycle if any fact among the ones she received satisfies her fact requirement; otherwise she is willing to receive facts for the current cycle. Each of the facts accessible to an agent that do not satisfy her fact requirement is, in turn, shared with her trusted neighbors that are not satisfied with their accessible facts. The *trusted neighbors* of an agent are the set of agents whom she trusts with a trust value greater than *trust threshold*, which is a user-defined static minimal trust value that guards information sharing. The recipient agents check these facts for their respective requirements in the next cycle. If their requirements are also not met, facts continue to be shared with their trusted neighbors. Such sharing sustains until the fact's *expiration interval*, a user-defined system level metric representing the number cycles for which a fact can be shared after its creation, is completed. At the end of every cycle the initiator agents monitor the current recipients of each of the unexpired facts they generated by querying the fact's fact pedigree. The *fact pedigree* of a fact is a document that is accessible only to the initiator agent. It contains the signatures of all agents who have received that fact. The initiator agents determine if their trust in the recipient agents warrants their receiving the fact and classifies them into *intended receivers*, i.e., agents that are not satisfied with the received fact but are trusted by the initiator, and *unintended receivers*, i.e., agents that are not satisfied with the received fact and are not trusted by the initiator. Thus, the initiator agents can easily determine the covert channels that might lead to unintended receivers, and correspondingly apply trust policy models in order to eliminate them.

Trust policy models define the various combinations of trust update policies that regulate and update the interaction trust relations in the fact pedigree that involves an unintended receiver. The overarching goal we considered for the design of these policies is to simultaneously maximize the information sharing and minimize the unintended receivers. Trust policy models have been devised such that varying degrees of restriction can be applied on the paths leading to unintended receivers.

The first trust update policy updates the trust relation between every consecutive pair of agents in a fact path starting from the unintended receiver back toward the initiator agent. Every agent's trust in the succeeding agent is reduced exponentially based on the relative role played by each agent in propagating the fact finally to the unintended receiver. The *degree of responsibility* for every agent in the agent fact path is the depth of the agent (with the fact originator as the root

element) in the fact pedigree. The trust update value at every agent is proportional to the exponential of the degree of responsibility of the succeeding agent with user-defined *trust decrement factor* as the base.

The second trust update policy is a variant of the first. For every trust relation in the fact pedigree that led to an unintended receiver, trust is updated with the same user-defined trust decrement factor. This policy model is more restrictive than the first, as we hold every agent in the agent fact path equally responsible for the transmission of the fact to the unintended receiver rather than proportionally.

A third trust policy model formulated is one that updates the trust value between the initiator agent and every agent in the path, if a trust relation exists, by the user-defined trust decrement factor. By reducing the initiator's trust value in her neighbors that are on the fact pedigree leading to an unintended receiver, this policy model impedes the first step in the fact transmission (i.e. to her neighbors) that may lead to unintended receivers. This property makes it the most restrictive of the three.

The main focus of and the drive for this research, as described earlier, is to maximize the information sharing among a group of agents while safeguarding its security in terms of confidentiality and integrity. *Information availability*, IA, is the degree to which information is freely available when shared among a group of agents. IA, a system level metric, is the sum of number of satisfied agents and the number of intended receivers expressed as a percentage of the total facts shared. Note that the sets intended receivers and satisfied agents are disjoint. *Security Measure*, SM, a system level metric, is a measure of the number of unintended receivers expressed as a percentage of the total facts shared. *Total facts shared* is the cumulative of the count of intended and unintended receivers. IA, SM, and total facts shared are measured for every cycle in the algorithm. The aim of the devised algorithms is to maximize IA and minimize SM, which satisfies our intended objective. The cycle number at which IA and SM converge to their ideal values, i.e. 100% and 0% respectively, termed *IA saturation cycle*, is also measured in the algorithm given in Fig. 1. In the algorithm, pseudo-share addresses the issue of agents checking for their required fact among the received facts, sent in the current cycle, in the next cycle as mentioned earlier in this section.

```

1 initialize agents, fact warehouse, fact requirements,
  interaction trust relations, trust threshold, trust
  update, expiration interval, and total number of cycles
2 for each cycle {
3   every agent generates a fact
4   every agent checks for fact requirement from her
  accessible facts
5   pseudo share the received facts each agent does
  not require
6   share the generated fact each agent does not
  require
7   share the pseudo share
8   each initiator agent checks for current unintended
  or intended receivers and applies trust update policy if
  there are any unintended receivers

```

Fig. 1. An outline of the Trust-Based Information Sharing Algorithm.

V. SIMULATIONS

In our simulation we create a user-defined number of agents randomly. The trust values connecting two agents are discrete values randomly chosen between 0.0 and 1.0, inclusive. The interaction trust relation among these agents is also generated randomly and can contain any topology (i.e. cycles or non-connected graphs or agents can have reciprocity in trust, etc). Thus, the random trust relation acts as a generalization for any network of agents. The fact warehouse is randomly generated from a group of numbers between 0 and a pre-defined upper bound that is greater than zero. The fact warehouse acts as a template and is never empty. The fact requirement for each agent and the fact generated in each cycle are chosen randomly from the range of facts in the fact warehouse. It is assumed that agents abide by trust threshold conditions, and that there are no unauthorized intruders, illegal manipulation of the content of the fact, fact losses during information-sharing, or delays in the simulations. Each simulation is replicated ten times and the results obtained are averaged. The simulations, developed and implemented in Java, and their analysis are presented in the following subsections.

A. IA and SM over Cycles

IA, SM as defined in section IV have been observed over the total number of cycles of the simulation and the results of one such simulation are given in Fig. 2. Fig. 3 represents IA, SM, results for the simulation without the application of any trust policy model.

From the results shown in Fig. 2, it is easily observed that when a trust update policy is applied, IA and SM reach their expected levels as the number of cycles increase. Saturation of IA, after the sixth cycle, suggests that information is freely shared between agents via the trusted fact paths. SM value, after the sixth cycle in Fig. 2, at level zero suggests that all covert channels cease to exist. At cycle 1, IA and SM are at their desired levels, but as the initial cycles are completed, the SM increases as more information is shared among agents that include unintended receivers. These unintended receivers are quickly eliminated subsequently by the application of trust update policies which results in the fall of SM to zero. Comparing these results with the ideal case as given in Fig. 4, we find a variation in the IA and SM in the initial cycles until respective desired levels are obtained. IA saturation cycle in Fig. 2 is 6.

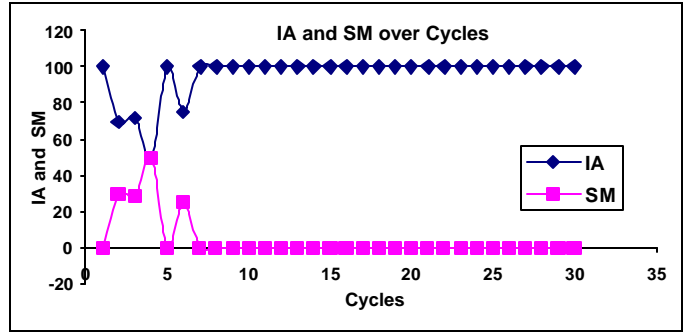


Fig. 2. IA and SM over Cycles




Fig. 3. IA and SM over Cycles without trust update policies

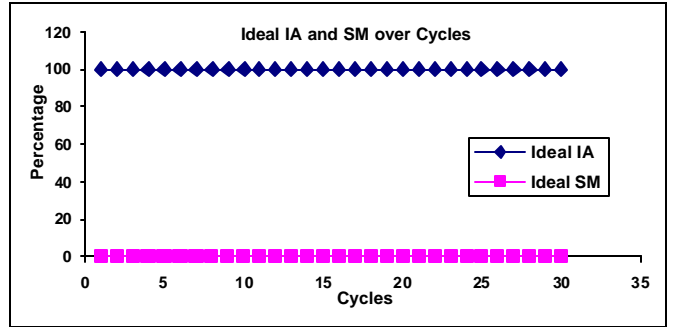


Fig. 4. Ideal IA and SM over Cycles

B. IA Saturation cycle over trust policy models

From the analysis of the results in section V-A, we find that the model proposed produces results that differ from the ideal case. Simulations in this section are performed to reduce the value of IA saturation cycle. The results of the simulation in Fig. 5 suggest that the value of IA saturation cycle depends on the degree of restriction of the trust policy model that is applied on the fact pedigree that leads to an unintended receiver.

The three types of trust policy models discussed in section IV can be ordered according to an increasing degree of restriction. For simplicity, the first, second, and third trust policy models are assigned numerical values of 1, 2 and 3 respectively, denoting the increasing degrees of restriction. The more the restriction a trust policy model applies, the greater the rate at which covert channels will be eliminated, and the lesser would be the value of IA saturation cycle.

Various combinations of trust policy models can be applied to the fact pedigree in the algorithm. The first and third, and the second and third trust policy models can be combined to

achieve greater restriction. They are assigned numerical values of 4 and 5 respectively. From (1) and Fig. 5 we find that as the degree of restriction of the trust policy model increases, the cycle at which IA saturates decreases, i.e., the model reaches a stable state and converges to the ideal nature of dissemination very quickly. Therefore,

$$IA \text{ saturation cycle} \propto 1 / \text{Degree of restriction of TPM} \quad (1)$$




Fig. 5. IA Saturation Cycle over trust policy models

C. Trust Threshold over Total Facts Shared

The trust threshold restricts the communication between two agents. In real-life organizations this can be mapped to any policy that restricts the occurrence of an action. Higher values of trust threshold indicate greater restrictions on performing the action of sharing the fact. This, however, reduces the total facts shared. Thus, the greater the trust threshold, the lesser the total facts shared.

A trade-off needs to be established between the trust threshold and the desirable total facts shared. This relationship is formulated in (2) and is depicted in Fig. 6. Here the total facts shared is the average of ten simulation values of the summation of the total facts shared at each cycle.

$$\text{Trust threshold} \propto 1 / \text{Total facts shared} \quad (2)$$

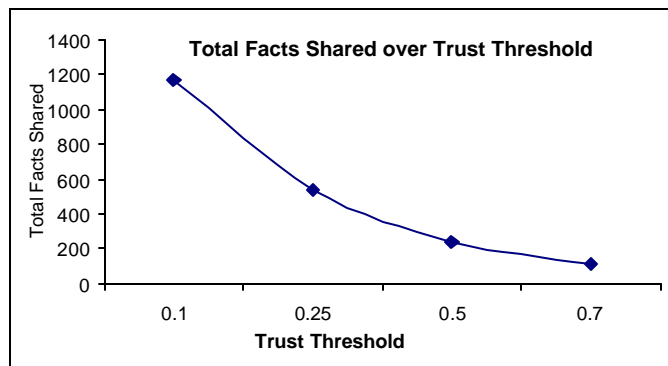


Fig. 6. Total Facts Shared over Trust Threshold

VI. DISCUSSIONS

The trust-based information-sharing algorithm is a variant of the classic flooding data dissemination algorithm. In classic flooding, an agent that needs to share information makes copies of the information and transmits one to each of

her neighbors. Also, whenever an agent receives a fact from another agent, she shares it with her neighbors excepting the sender. In our proposed model, an agent that receives a fact can transmit it back to the sender as long as the sender is trusted.

Two deficiencies of the classic flooding algorithm, implosion and overlap, are circumvented in the trust-based information sharing algorithm. Implosion occurs when an agent is the recipient of the same fact from two different sources via different paths. Overlap occurs when the agents generate facts from overlapping sources of information, which can sometimes result in same information [4]. Thus, the same fact is duplicated at the recipient in both the cases.

First, in our model, each fact generated has a unique identity. Thus, if an agent receives the same fact that is produced from two different sources, or if the same agent produces the same fact in two different cycles, and if an agent receives them from same or different sources, in effect she has received two different facts. Second, every agent in the system, as mentioned above, is concerned with observing the paths taken by each of her generated facts. Thus, even though the same agent receives the same fact from different sources, they have traversed different channels of agents to reach the current recipient. This supports every initiator's purpose of monitoring the possible conduits of information sharing among the agents. Also note that an agent can receive multiple copies of the same fact via the same path in the same cycle. In this case, the agent prefers the freshest fact, while the others continue to be shared. Thus, the implosion and overlap problems of classic flooding algorithm are overcome.

Trust among dominance levels in an organization produces security class structures isomorphic to the partially ordered sets of the organizational structure itself. This allows our policies to replicate the same hierarchical nature of established military security models such as Bell Lapadula, as well as the heterogeneity of ubiquitous organizational security models like Role Based Access Control.
... (in progress).

VII. FUTURE WORK

Through simulation we have observed that our goal of maximizing information availability while minimizing security breaches involves finding a strict balance of the formalized parameters of our model (trust policy model, trust threshold, total facts shared, etc). Additionally, we must extend our observations beyond IA saturation. We are currently developing a *rehabilitation* process by which agents who are no longer trusted and effectively eliminated from the network are allowed to re-enter. With this in place, we anticipate the discovery of the optimal range of the various parameters that arrive at the ideal.

We also consider a socio-organizational dynamic as a beneficial addition to our model. There are two polar concepts involved therein: the overall organizational goal and individual goals and aspirations of its employees. In this vein, we envision a system of *pride* and *respect* as a means by which these goals might coexist [14]. An agent-employee

who refrains from excessive selfish behavior will garner a level of respect from her parent organization. At the same time that agent's pride (or *allegiance*) in the organization increases. Inherently, the agent seeks this respect, and so, therefore, this is her motivation.

Presently our model only considers intra-organizational information sharing. A logical extension of this model is to an inter-organizational level. In the inter-organizational model we consider the additional trust relationships between two organizations. We propose a means of inter-organizational information sharing via a special agent called a *boundary spanner*. The boundary spanner is an organizational role donned by the most reliable agents. Implied in this special arrangement is the consideration of two distinctly identified yet related forms of trust: interpersonal and inter-organizational. We are interested in exploring their relationship, specifically how the interpersonal trust of constituent boundary spanners affects the inter-organizational trust of their parent-organizations.

VIII. CONCLUSION

The trust-based information sharing model presented in this work makes several contributions to research in the field of distributed security. Our contention that this model sufficiently and effectively guarantees the high availability of information while limiting security breaches has been supported by the results of the simulations. These simulations have demonstrated that the use of a soft security mechanism such as trust is as effective as hard security mechanisms. The inherent advantage of our model is in its simplicity, malleability, and scalability. Our model could be easily adapted to any open system for information sharing, like e-institutions and security agencies. We therefore conclude that trust-based security policies, such as the one presented herein, are beneficial and have practical application in real-world settings.

REFERENCES

- [1] Ralph C. Bledsoe, et al., "The 21st Century Federal Manager: A Study of Changing Roles and Competencies", a report by a panel of the Human Resources Management Consortium, National Academy of Public Administration, July 2002.
- [2] Phillips, Jr., C. E., Ting, T.C., and Demurjian, S.A. (2002), "Information Sharing and Security in Dynamic Coalitions", in *Proceedings of the Seventh ACM Symposium on Access Control Models and Technologies*: 87-96. Monterey, California.
- [3] Goecks, J. and Cosley, D. (2002), "NuggetMine: Intelligent Groupware for Opportunistically Sharing Information Nuggets", in *Proceedings of the 7th Annual International Conference on Intelligent User Interfaces*: 87-94. San Francisco, California.
- [4] Rabiner-Heinzelman, W., Kulik, J., and Balakrishnan, H. (1999), "Adaptive Protocols for Information Dissemination in Wireless Sensor Networks", in *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking*: 174-185. Seattle, Washington.
- [5] V. Lesser, C.L. Ortiz, M.Tambe, "Distributed Sensor Networks: A Multiagent Perspective", *Kluwer*, 2003.
- [6] G. Beavers and H. Hexmoor, 2003. "Understanding Agent Trust", in *Proceedings of The International Conference on Artificial Intelligence (IC-AI 2003)*: 769-775.
- [7] Golbeck, Jennifer, Bijan Parsia, James Hendler, "Trust Networks on the Semantic Web", *Proceedings of Cooperative Intelligent Agents 2003*, Helsinki, Finland.
- [8] Lars Rasmussen, Sverker Jansson, "Simulated social control for secure Internet commerce", *Proceedings of the 1996 workshop on New Security Paradigms*: 18-25, September 17-20, 1996, Lake Arrowhead, California, United States
- [9] Abdul-Rahman, A. and Hailes, S. (1998), "A Distributed Trust Model", in *Proceedings of the 1997 workshop on New Security Paradigms*: 48-60. Langdale, Cumbria, United Kingdom.
- [10] Golbeck, J., Parsia, B., and Hendler, J. (2003), "Trust Networks on the Semantic Web", in *Proceedings of Cooperative Intelligent Agents*, Helsinki, Finland.
- [11] Gil, Y. and Ratnakar, V. (2002), "Trusting Information Sources One Citizen at a Time", in *Proceedings of the First International Semantic Web Conference*, Sardinia, Italy.
- [12] Sydow, J. (1998), "Understanding the Constitution of Interorganizational Trust", in Lane, C. and Bachmann, R. (eds.), *Trust Within and Between Organizations: Conceptual Issues and Empirical Applications*. New York: Oxford University Press 1998.
- [13] Hardin, R. (2002), *Trust & Trustworthiness*. New York: Russel Sage Foundation.
- [14] Tyler T., "Why do people rely on others? Social identity and social aspects of trust", In: Cook K, ed. *Trust in Society*, vol. II. New York: Russell Sage Foundation; 2001:285-307.