Users and Trust in Cyberspace

Pekka Nikander and Kristiina Karvonen

Helsinki University of Technology
{Pekka.Nikander, Kristiina.Karvonen}@hut.fi

Abstract. The underlying belief and knowledge models assumed by various kinds of authentication protocols have been studied for well over 10 years now. On the other hand, the related question of the generic trust assumptions, which underlie the settings where the protocols are run, has received less attention. Furthermore, the notion of trust, as it is typically defined, has more been based on the formal model used than the real user requirements posed by the application context and the actual people using the system.

In this paper, we approach that problem from the users' point of view. We briefly describe what are the psychological bases on which typical people build their trust assumptions on, and consider how these are reflected in a typical e-commerce setting today. Given this background, we proceed to contemplate how the systems could be made more trustworthy by explicitly representing the trust assumptions and requirements, and how these digital expressions of trust could be instrumented to and integrated with actual authentication protocols. Thus, our aim is to broaden the view from a protocol centric approach towards considering the actual users, and to provide some initial requirements for future operating systems and user interface design.

1 Introduction

The majority of computer system users are relatively ignorant about the security, or non-security, of the systems they use. In fact, if asked, they tend to claim that they do not care [1]. However, if the same people are asked to explain how they adapt their behaviour according to the situation at hand, it rapidly becomes apparent that much of their behaviour is based on the *perceived* sense of security or insecurity. [2]

Still today, most people seem to consider computer systems almost godly; the computers are considered so complex that they could be understood, and it is generally believed and explained that computers themselves are unable to err, that is, whenever a computer seems to make a mistake, the fault is presumed to lie on the user. Furthermore, the computers seem to demand ungodly amounts of sacrifice in the form of time spent in getting them to produce the output desired (and getting them work in the first place). Besides, the publicity received by the various kinds of attacks against Internet based systems, along with the general opinion starting to consider personal computers unreliable[1], is slowly deepening the situation. Thus, a typical unconscious attitude to-

[1] Hereby we want to thank the parent organization of the workshop organizer for educating the general public's opinion, especially in the area of reliability of personal computer operating system and office effectiveness tools, to consider computers unreliable and untrustworthy.

wards a computer system may well resemble the attitude paid towards an austere god that requires blind faith and dedication.

When we consider this typical user attitude with the complexities involved in designing, implementing and verifying actual security protocols, we get an initial impression on the obstacle to be tackled. Fortunately, the work spend on analysing the underlying belief models of authentication protocols (e.g. [3], [4], [5], [6]), in modelling trust in technical sense (e.g. [7], [8], [9], [10]), and in providing infrastructures for expressing authorization and trust in open systems (e.g. [11], [12], [13], [14]), together with more user centric work (e.g. [1], [2], [15], [16]), allows us to outline a map of the problem area. Basically, we surmise that combining explicit, key oriented authorization with operating system level security and basic user interface design, could provide a sound technical basis for expressing the technical trustworthiness of a system in a way understandable to the average user. This, combined with enough of market reputation and supporting legislation, might be able to convert the Internet from its current insecure state into one where people could base their actions on reasonable security assumptions even when they do not possess deep technical knowledge about computer systems or information security.

In this paper, we attempt to give a glimpse to this possible solution approach. We start with the user centric view, considering how the aspects affecting the average user in making trust decisions, continuing to contemplate how some of these aspects could be expressed in digital form, and concluding with ideas how these digital expressions could be used in conjunction with actual security protocols.

1.1 About the Nature of Trust

The term “trust” is used in the literature to denote several different but related phenomena. Basically, a distinction can be made between two basic meanings. In computer security literature in general, the term is used to denote that something *must* be trusted (e.g. Trusted Computing Base, TCB). That is, something trusted is something that the users are *necessarily* dependent on. If a trusted component breaks, the security of all of the system breaks. On the other hand, in this paper, as well as elsewhere in more psychologically oriented literature, the term is used to denote that something *can* be trusted. That is, something trusted is something that the users *feel comfortable with* to be dependent on. If a trusted component breaks, the users feel betrayed. Probably some harm is done to the security of the system, but that is less relevant to this discussion.

This distinction should be kept in mind when reading this paper. In this paper, the term trust is used to denote the psychological attitude of being willing to be dependent on something or somebody. For example, if Alice trusts Bob (in some respect), Alice is willing to delegate control to Bob (over the issues covered), thereby making herself more dependent on Bob’s honesty and goodwill. Thus, whenever Alice expresses her trust, she, in fact, announces her willingness to trade a piece of her personal control for simplifying the situation. The usual reason for deciding to trust is the desire to make (future) decisions simpler. An existing trust relationship allows the user to proceed in her pursues more easily, without needing to contemplate whether the procedure is safe or not. In this respect, there is no distinction whether something is trusted because of a necessary need or due to a decision based on emotions and consideration.

The rest of this paper is organized as follows. In Sect. 2 we briefly describe some of the most relevant psychological aspects of trust, concentrating on how trust is created and lost in the current cyberspace. Next, in Sect. 3, we outline a proposal of how some of the user centric trust forming aspects could be represented in an non-forgeable digital form, while in Sect. 4 we consider how this information could benefit the integration of protocol level and operating system level security. Finally, Sect. 5 provides a brief discussion of some of the aspects involved.

2 Users and Trust

What does it mean to trust someone, or something? The concept of *trust* seems to imply lack of sufficient amount of knowledge [17], meaning that there is at least some amount of uncertainty involved [8][18][19]. On the other hand, trusting reduces the complexity of a situation. When we *decide* to trust rather than suspect — this is what it means when we talk of *a leap of trust* — the number of issues we have to consider is reduced, thereby simplifying the process of making decisions. Trusting also describes *an attitude* towards future expectations, as well as introduces the presence of implied risk in a given situation [19].

2.1 Technical vs. Psychological Trust

In the technical sense, there exists a number of reasonable well defined definitions for trust, e.g., [7], [8], [9], [10], [13], [20], and [21]. Thus, the concept of trust in a technical sense is rather well-defined, at least in comparison with the psychological definition of trust, which for the most part still remains unresolved and under discussion. In general, the technical approaches tend to consider trust as a more or less binary concept (with the exception of at least [8]); there either is trust or there is not. However, typically a distinction is made between various *types of trust*, e.g., distinguishing recommendations from "direct trust" (whatever the latter means).

A leap of trust is needed, because there is not conclusive amount of information available. This would, in fact, be the description of most real-world user situations. In this sense, Audun Jøsang's approach [8] seems to have some connections to the psychological sense of trust.

Understanding the real-world trust is crucial to understanding the actual security of any transactions on-line — maybe even more so than creating the technological solutions for these transactions. Users are often considered to be the weakest link in the security of on-line transactions, and rightly so; what else could they be, when they are not provided with sufficient amount of information and/or support on security-prone situations by the system and its user interface design? How could the users be expected to be able to make rational choices of whether an operation is secure and trustworthy or not, if they are not given the right information? This point is well expressed in the following quote by Eric Ketelaar, in his demand for trustworthy information [22].

> *"Why do we demand more of the quality of food or a car than we demand of that other essential: information? Reliability and authenticity determine the*

credibility and the usefulness of information. These concepts, developed in different cultures and at different times, are essential for our information society in its dependence on trust in information. In the creation and distribution of digital information, conditions should be met to ensure the reliability and authenticity of the information."

Trust can also be viewed as a historically emergent property of human interaction that is tied to a specific form of social organization. This means that modern forms of trust are rooted in the rights, obligations, and liberties of citizenship. Throughout history, people have always tried to ensure the authenticity of a document by several means: a seal, a special mark, witnesses, placing the document in safe-keeping with a public official, etc. Modern electronic systems also have these safeguards. They use passwords, cryptography, electronic sealing, digital signature, etc. Rules are needed about form, communication, and storage of information. [23]

2.2 The Untrustworthiness of the Web

Jacob Nielsen has described the current state of the Web as one of untrustworthiness, where "customers are traded like sheep" [24]. This has also been confirmed by other user studies, e.g., [16]. In practice, this mean that e-business has not taken the customers' need for security into any consideration at all. According to Nielsen, this has to change, however, if one wants to establish any decent business on the Net. *A culture of trust* must be promoted whenever there is a need to create a functioning network in a virtual world, and that's what the electronic marketplace essentially is. Mutual trust is always needed for good-quality relationships, be they between two people, a group of people, or between a user and an on-line service.

It is interesting, however, that according to another study by Hoffman et. al. [16], it was found that the negative perceptions of security and privacy increased along with the level of on-line proficiency — the more fluent the users were with using on-line service, the more conscious they seemed to be of the lurking risks of on-line transactions. In another study [25], the likely on-line consumer was described as someone with a "wired lifestyle": having years' of experience on the Internet, receiving a lot of e-mail, and searching for product information on the Internet, to name a few. In our studies, the behaviour of the users was exactly the opposite — the most educated and experienced users were most against and doubtful of on-line services that included transactions of money or private information [1].

2.3 Transferring real-world trust to on-line systems

Is real-world trust transferrable to the digital world? It seems that the answer is "yes". Trusting a bank stays more or less the same regardless of the media (there is, however, also some reports on studies that suggest exactly the opposite, e.g., [16]). More important than the place where the service is situated is the existing brand reputation and other users' opinions about the service provider. These elements create the *sense of place* that guides the social interactions, perception of privacy, and the nature of all transactions conducted on-line [2].

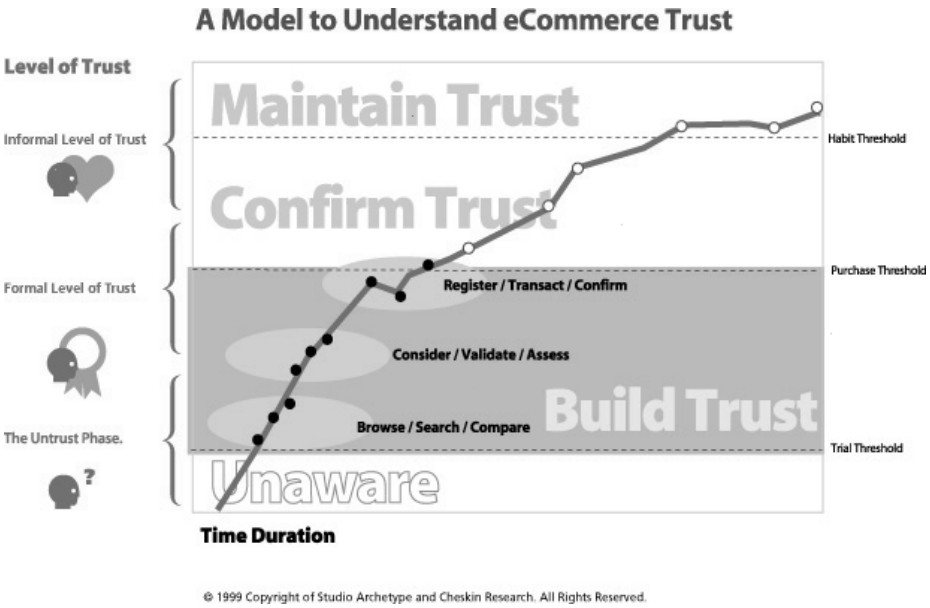

Fig. 1. An example of accumulation of trust as a function of time [15]

To start trusting is slow: trust is formed gradually, it takes quite a lot of time and repeated good experiences [15]. On-line trust can be described in terms of a human relationship. The initial stage is that of interest and suspicion; there has to be a motivation, a need, to get interested in the service, but this curiosity is stamped with distrust and suspicious cautiousness in the beginning of this flirtation with a new on-line service.

2.4 Recommendations, rumours and hear-say

To become better acquainted with the object of interest, additional information is gathered of the service through various media: the mass media, but especially from other people, i.e., friends, colleagues, experts etc. We are not alone in this world but surrounded by others, some friends, and some enemies. Listening to rumours play a big role in gaining information: finding appropriate knowledge is difficult and time-consuming, and users are not really motivated to find out about the technical details of the security of the service to begin with — all they want to know is whether it is safe to use the service or not, and not why this is so [2]. Also, users may often not have any way to judge whether the information gained is trustworthy [1]. Friends are trusted, and it should not be too surprising that the information they provide often forms the basis for decision-making also when starting the use of an on-line service. Trust in people is transformed into on-line trust [1].

2.5 Imposing Laws on On-line Behaviour

On-line trust depends on many factors, including consumer rights, freedom of expression, and social equity [26]. Trusting an on-line service provided by a well-known

bank is to a great extent based on users' knowledge (or assumptions) about the laws binding all the business operations of the bank. In most studies about computer security, the users report on finding legislative intervention by the state desirable and necessary for promoting on-line trust (e.g. [27], [16]). Furthermore, behind this trust in the legality of the bank is, to put it bluntly, the sub-conscious trust in the basic structures of the society to remain stable, the trust in the status-quo instead of anarchy — in short, the trust and belief in the good-doing nature of a social contract between men, in the Rousseaun sense [28].

Another kind of “social contract” is also suggested to be executed in the cooperative relationships built on the Net: According to the study by Hoffman et. al.[16], over 72% of Web users would have been willing to provide the service provider with personal information if only the sites would provide the customers with a statement about how this information would be used. Still, it seems that users do not consider information about themselves as merchandise, to be sold to the highest-bidding offer: in the Hoffman et. al. study, most users were found not to be interested in selling their personal information.

2.6 Trusted Portals

Most user studies investigating into perceived trustworthiness of on-line services have focused on evaluating the services of single companies that offer their services directly to the customer. In our user studies, this approach formed the starting point for the inquiry. We took the study one step further, however, by introducing the users to the idea of a *trusted portal*, i.e, a third party taking care of the on-line monetary transactions, and acting as a trusted third party between the Web merchant and the on-line customer. This third party was a party that has for centuries been trusted to handle our money properly — a well-known, long-since established bank.

Third parties acting as mediators is an idea repeatedly expressed in many studies (see, for example, Hoffman et.al.), and it is also more or less the same as is behind the idea of Certification Authorities (CA) that, in form of seals of approval, or trustmarks, would guarantee the safety of the Web merchant, such as the TRUSTe [29].

In our study, the bank acted as the host of the trusted portal. Thus, in that case the real life trust placed on the back was more or less completely transferred to the trusted portal in cyberspace.

2.7 Losing Trust

While gaining the trust of on-line users may be slow and painstaking from the entrepreneur's point of view, losing trust happens quickly. A single violation of trust may destroy the achievements of trust over a period of months or even years [24], [15]. And, once broken, the recovery of lost trust is difficult, if not impossible. One cause for losing trust may be an initial misunderstanding of how the system works. To avoid these misunderstandings, the service should be meaningful to the user: The more intelligible the service and the system behind it are to the user, the less likely she is to misinterpret them and the more willing and motivated she will be to put an effort to learn to use the service in a secure and educated way, i.e., to participate [30]. Motivation can

also be provided by personalised privacy tools; at present, a user with particular privacy needs and policy often lacks the means to fulfil them [31]. Users interested in their privacy often have to also conclude that the privacy information on most sites is confusing, incomplete, and inconsistent [32], even if the users would show an interest towards this kind of information. Privacy has also often be balanced against other, competing interests, both personal and others' [26].

3 Representing Trust in Digital Form

As we already mentioned in Sect. 2.1, a central problem lies in providing the users with information about the *real* security and trustworthiness of on-line operations so that they can make rational choices. Here, the terms *real security* and *real trustworthiness* necessarily refer to the social context of the user. That is, the user should be able to compare the security level of the system with the risks involved in the intended operation. To the average user, both the security level and the involved risks are at best vague, if comprehensible at all. Thus, the basis of the supposedly rational choices are based on the context, including individual trust decisions made by others, explicit recommendations, perceived brand reputation and other users' opinions, together with the quality of the relationships between the decision maker and the other individuals.

Today, the social context used in making decisions about the trustworthiness of services in the cyberspace cannot securely rely on the cyberspace itself [2]. That is, if I want to create an opinion of mine about the security level of a particular web server, I tend to prefer information received in real world, e.g., from my colleagues at the coffee table and through the public media.

In cyberspace, trust can be expressed. There is already a number of various techniques that attempt to express real life trust in various kinds of digital format. PGP is an example of such a system, where explicit real world trust is transferred in the digital form, and where the digital expressions of trust can be used infer the trustworthiness of previously unknown email addresses. Unfortunately, the current PGP approach is relatively rigid, and inherently bound to a single application, i.e., providing keys for securing email. [33]

Correspondingly, the PolicyMaker [11] and related approaches attempt to provide a more flexible platform for expressing authorization in a digital form. However, on that branch of security research, the focus has almost completely been on decentralizing access control systems. But, as we have argued ([11], [13]), these system can be, and should be, extended to handle also other forms of trusted information. That is, even the expressions of authorization information may be considered as a form of trust expressions, and the same kinds of certificates can be used to express many other forms of trust.

Thus, we propose that some form of authorization certificates, or rather *trust certificates*, is used for expressing *trust decisions* and *recommendations* made by the users. The same kind of certificates can also be used to represent the quality of relationships between individuals, allowing me to consider whose recommendations and opinions I trust and in which sense, and also to publish these considerations of mine. A suitable format for these kinds of trust certificates might be signed XML documents, allowing

basically any XML DTD to be used to express opinions. It is noteworthy that the usage of XML might even allow to express security policies in a digital, secure format. [34]

Certificates could also be used to express attitudes about brand names, and for associating specific networked servers with specific brand names. This would require, of course, that each brand name owner would publish or certify their own public key in some secure enough way, e.g., by publishing certificate fingerprints periodically in a newspaper. Given this kind of arrangement, users could also express their opinions about specific brands by referring to the public keys of those brand names.

Already as such, these kinds of techniques could be used to create digital counterparts of our real life social networks, and to express our opinions in a digital format. However, since the certificates are in machine readable form, and since XML documents can be relatively easily parsed and handled programmatically, it would make much more sense to integrate the handling of these kinds of trust expressions directly to the future operating systems and user interfaces. That is the topic of the next section.

4 Binding Trust to Operating Systems and Protocol Runs

The purpose of the explicit utterance of trust, in the form of certificates, is to promote *a culture of trust* (which we called for in Sect. 2.2), and to create *a secure sense of place*, allowing the users to conduct their tasks with a feeling of security that is based on real security measures. As discussed, essential elements in these are good-quality relationships, explicit brand reputation, and other users' opinions about service provides, among other things. All of these, along with basic recommendations and expressions of trust, can be represented in the form of digitally signed documents, i.e., certificates.

In order to be real useful, the handling of these kinds of trust expressions should be integrated to the trusted computing base (TCB) of the used computing system. That is, the security mechanisms of the underlying operating system should be extended to understand where, when, and for what purpose, trust is needed when conducting transactions over the network. In practice, this means that the operating system takes responsibility for securing the network connections, and whenever running an authentication protocol in order to open a new connection, takes care of evaluating the trust requirements of the requesting application together with the credentials of the server and client programs.

To put this in slightly more concrete terms, we might consider a multi-user operating system running TCP/IP protocol stack and using the IPSEC security protocols. In such a setting, the operating system would issue a security policy on all connection requests, allowing only such connections to be opened whose trust assumptions and security credentials match. The security policy would be based on the trust expressions the user has earlier stated, augmented with on-line user interaction when needed.

On opened connections, the trust assumptions and credentials would be separately bound to each IPSEC security association (SA), allowing SA sharing whenever the needs of a new connection match with ones provided by an existing SA. [13]

As another example, we have considered how Java/Jini based ad hoc communities could be secured with SPKI certificates, and how simple application specific trust rela-

tionships could be represented in that kind of information. [35]

One area still requiring considerably more study is the relationship of these kinds of security measures, enforced by the operating system, and the user interface. It seems that something similar to the trusted path is needed.

5 Discussion

In order to be really useful, quite a lot still needs to be done. First, it is not at all clear how the various kinds of trust relationships and their expressions could be turned into certificates or other kinds of signed documents. Second, the actual user expectations and their probable reactions to various kinds of automated trust evaluation mechanisms should be evaluated. Third, even the concept of trustworthiness needs more clarification, both in the formal sense and especially in a language understandable to the average user. Furthermore, it seems inevitable that some new legislation is also needed.

For example, considering an on-line service trustworthy means, among other things, considering the information provided for the service and all the conducted transactions to remain private. This, then, means that the information will not be available to others, and will not be used out of context, for example. But defining privacy is not an easy task. Privacy is a basic human requirement we have a fundamental right to, but this does not reduce its unambiguity. What is regarded as private varies across organisations, cultures and even individuals [36].

Good example of this are the findings of a study at the AT&T Labs-Research [37] on Net users' attitudes towards privacy, where it was concluded that users could be divided into at least three groups according to their privacy assessments. These included 1) privacy marginalists, who showed little or no interest in privacy matters, 2) privacy pragmatists, who were concerned about their privacy but were ready to trust the services if there was some sign of existing privacy protection, and 3) privacy fundamentalists, who were extremely concerned about their privacy and very suspicious of the online services. All these different groups seem to require different user interface designs, emphasizing different aspects of the underlying systems security.

Thus, as the users' expectations vary quite a lot, the mechanisms are not quite there yet, and it is unclear how the implementation of such mechanisms would effect the design and structure of operating systems and user interfaces, this work is in the very beginning at best. However, we wish that these contemplations would lead to new ideas and points of view, preferably eventually leading to an internet that is more secure, in practice, than the current one.

References

1. Kristiina Karvonen, "Creating Trust", in *Proceedings of the Fourth Nordic Workshop on Secure IT Systems (Nordsec '99)*, November 1-2, 1999, Kista, Sweden, pp. 21-36

2. Anne Adams and M. Angela Sasse, “Users are not the Enemy”, *Communications of the ACM*, Vol. 42, No. 12, December 1999, pp. 41-46
3. Martin Abadi, Mark R. Tutle, “A Semantics for a logic of authentication”, in *Proceedings of the 10th ACM Symposium on Principles of Distributed Computing*, pp. 201-216, ACM Press, Aug. 1991.
4. Michael Burrows, Martin Abadi, and Roger Needham, “A logic of authentication”, *ACM Transactions on Computer Systems*, 8:1, pp 18-36, Feb. 1990.
5. Paul Syverson and Paul C. van Oorschot, “On unifying some cryptographic protocol logics”, in *Proc. 1994 IEEE Computer Society Symposium on Research in Security and Privacy*, pp. 14-28, May 1994.
6. Pekka Nikander, *Modelling of Cryptographic Protocols*, Licenciate’s Thesis, Helsinki University of Technology, December 1997.
7. Thomas Beth, Malte Borcherding, and Birgit Klein, “Valuation of trust in open networks”, in *Proceedings of Computer Security--ESORICS’94*, Brighton, UK, 2-9 Nov. 1994.
8. A. Jøsang, *Modelling Trust in Information Society*, Ph.D. Thesis, Department of Telematics, Norwegian University of Science and Technology, Trondheim, Norway, 1998.
9. Raphael Yahalom, Birgit Klein, Thomas Beth, “Trust relationships in secure systems: a distributed authentication perspective”, in *Proc. 1993 IEEE Computer Society Symposium on Research in Security and Privacy*, pp. 150-164, IEEE Computer Society Press, May 1993.
10. Raphael Yahalom, Birgit Klein, Thomas Beth, “Trust-based navigation in distributed systems”, *Computing Systems*, 7:1, pp. 45-73, Winter 1994.
11. Matt Blaze, Joan Feigmenbaum, and Jack Lacy, “Decentralized trust management”, in *Proc. 1996 IEEE Computer Society Symposium on Research in Security and Privacy*, Oakland, CA, May 1996.
12. Ilari Lehti, and Pekka Nikander, “Certifying trust,” in *Proceedings of the Practice and Theory in Public Key Cryptography (PKC) ‘98*, Yokohama, Japan, Springer-Verlag, February 1998.
13. Pekka Nikander, *An Architecture for Authorization and Delegation in Distributed Object-Oriented Agent Systems*, Ph. D. Thesis, Helsinki University of Technology, March 1999.
14. G. U. Wilhelm, S. Staamann, L. Buttyán, “On the Problem of Trust in Mobile Agent Systems”, in P*roceedings of the 1998 Network And Distributed System Security Symposium*, March 11-13, 1998, San Diego, California, Internet Society, 1998.
15. *ECommerce Trust Study*, Cheskin Research and Studio Arhetype/Sapient, January 1999, `http://www.studioarchetype.com/cheskin/`
16. Donna L. Hoffman, Thomas P. Novak, and Marcos Peralta, “Building Consumer Trust On-line”, *Communications of the ACM*, April 1999, Vol. 42, No. 4, pp. 80-85
17. Lucas Cardholm, “Building Trust in an Electronic Environment”, in *Proceedings of the Fourth Nordic Workshop on Secure IT Systems (Nordsec '99)*, November 1-2, 1999, Kista, Sweden, pp. 5-20
18. A. Jøsang, “Trust-based decision making for electronic transactions,” in L.Yngström and T.Svensson (Eds.) *Proceedings of the Fourth Nordic Workshop on*

Secure IT Systems (NORDSEC'99), Stockholm, Sweden, Stockholm University Report 99-005, 1999.
19. M. Mühlfelder, U. Klein, S. Simon and H. Luczak, "Teams without Trust? Investigations in the Influence of Video-Mediated Communication on the Origin of Trust among Cooperating Persons", in *Behaviour & Information Technology,* Vol. 18, No. 5, 1999, pp. 349-360
20. Ronald Fagin and Joseph Y. Halpern, "I'm ok if you're ok: on the notion of trusting communication", *Journal of Philosophical Logic*, 17:4, pp. 329-354, Nov. 1988.
21. Gustavus J. Simmons and Catherine A. Meadows, "The role of trust in information integrity protocols", *Journal of Computer Security*, 3:2, 1994.
22. Eric Ketelaar, "Can We Trust Information?", in *International Information & Library Review,* Academic Press Limited, 1997, 29, pp. 333-338
23. A. B. Seligman, *The Problem of Trust*, Princeton University Press, New Jersey, 1997.
24. Jacob Nielsen, "Trust or Bust: Communicating Trustworthiness in Web Design", *Alertbox*, March 7, 1999, at `http://www.useit.com /alertbox/990307.htm`
25. Steven Bellman, Gerald L. Lohse, and Eric J. Johnson, "Predictors of On-line Buying Behaviour", *Communications of the ACM*, Vol. 42, No. 12, December 1999, pp. 32-38
26. Roger Clarke, "Internet Privacy Concerns Confirm the Case for Intervention", *Communications of the ACM*, Vol. 42, No. 2, February 1999, pp. 60-67
27. Kristiina Karvonen, "Enhancing Trust On-line", in *Proceedings of the Second International Workshop on Philosophy of Design and Information Technology (PhDIT '99)*, December 15-17, 1999, St.Ferréol, Toulouse, France, pp. 57-64
28. Jean-Jacques Rousseau, Maurice Cranston (Translator), *The Social Contract*, Reprint edition (September 1987), Penguin Books, USA.
29. Paola Benassi, "TRUSTe: An On-line Privacy Seal Program", *Communications of the ACM*, Vol. 42, No. 2, February 1999, pp. 56-59
30. Elena Rocco, "Trust Breaks Down in Electronic Contexts but Can Be Repaired by Some Initial Face-to-Face Contact", in *Proceedings of CHI '98*, April 18-23, 1998, Los Angeles, CA.
31. Tessa Lau, Oren Etzioni, and Daniel S. Weld, "Privacy Interfaces for Information Management", *Communications of the ACM*, Vol. 42, No. 10, October 1999, pp. 89-94
32. *Surfer Beware III: Privacy Policies without Privacy Protection*, Electronic Privacy Information Center (`www.epic.org`), December 1999, `http://www.epic.org/reports/surfer-beware3.htm`
33. Alma Whitten and J.D. Tygar, "Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0," in *Proceedings of the 8th USENIX Security Symposium*, August 1999.
34. Juha Pääjärvi, "XML Encoding of SPKI Certificates", work in progress, Internet draft `draft-paajarvi-xml-spki-cert-00.txt`, March 2000.
35. Pasi Eronen, Johannes Lehtinen, Jukka Zitting, and Pekka Nikander, "Extending Jini with Decentralized Trust Management", in *Proceedings of OpenArch'2000*, Tel Aviv, Israel.

36. Anne Adams and M. Angela Sasse, “Privacy Issues in Ubiquitous Multimedia Environments: Wake Sleeping Dogs, or Let Them Lie?”, Proceedings of Interact '99, IFIP TC.13 International Conference on Human-Computer Interaction, 30th August - 3rd September, 1999, Edinburgh, UK, pp. 214-221
37. I. F. Cranor, J. Reagle and M. S. Ackerman, “Beyond Concern: Understanding Net Users' Attitudes about On-line Privacy”, *AT&T Labs-Research Technical Report TR 99.4.3,* `http://www.research.att.com/library/trs/TRs/99/99.4/`
38. Ross J. Anderson, “Liability, trust and security standards”, in *Proceedings of the 1994 Cambridge Workshop on Security Protocols*, University of Cambridge, UK, Springer-Verlag 1994.
39. Gustavus J. Simmons, “An introduction to the mathematics of trust in security protocols”, in *Proc. Computer Security Foundations Workshop IV,* pp. 121-127, Franconia, N.H., 15–17 June, EEE Computer Society Press, Los Alamitos, CA, 1993.