

Prospectives for Modelling Trust in Information Security

Audun Jøsang *

Department of Telematics, NTNU, 7034 Trondheim, Norway
Information Security Research Centre, QUT, Brisbane Qld 4001, Australia **
Email: ajos@item.ntnu.no

Abstract. This paper describes trust in information security as a subjective human belief. On this background, four formal models for trust which have been proposed in the recent years are analysed with the purpose of determining their strong and weak sides. From this we try to define general criteria for the feasibility of modelling trust.

1 Introduction

A wide range of methods are used to achieve security in IT systems. Whether it is cryptographic algorithms and protocols, or high level security models and policies, they all have in common the implicit purpose of creating trust. This observation has led some researchers to propose models where trust is explicitly expressed.

Trust is a human cognitive phenomenon, and not for example a property of a system or of an agent. The purpose of modelling trust must therefore be to model how a human observer would assess the security of a system or the honesty of an agent. This would enable automatic appreciation of trust *as a human would do it*. The advantage is that trust can be assessed quickly, and situations which would have been too complex for the human mind can be efficiently analysed.

2 Definition of Trust

Trust is a very general concept which can be used in almost any context. For the purpose of IT security, it is desirable to give trust a more specific meaning. Our definition of trust is taken from [Jøs96].

A human would be trusted if he or she is believed to be benevolent, and distrusted if believed to be malicious. It may be true that there is a finite number of factors which determine whether a human will behave in a benevolent or malicious way, but in practice it is impossible to determine all these factors, and the behaviour of a human will therefore always be partly unpredictable. For all practical purposes, whatever the underlying mechanism may be, we will call the

* This research was supported by Norwegian Research Council Grant No.116417/410.

** This research was carried out while the author was visiting the ISRC at QUT.

mental process which decides between benevolent and malicious behaviour *the free will*, and we designate agents possessing this type of free will as *passionate*. We define trust in a passionate agent as *the belief that it will behave without malicious intent*.

Algorithms, protocols, software and hardware can hardly be characterised as passionate or having a free will, but they can still be trusted. We will call this type of agent *rational* as opposed to passionate. Because a rational entity has no free will, it is not expected to be malicious or benevolent. What exactly is being trusted is that it will resist any attempt of manipulation by malicious agents. We therefore define trust in a rational entity as *the belief that it will resist attacks from malicious agents*.

3 Models for Trust in Rational Entities

BAN-Logic and *Security Evaluation Criteria* are two examples of methods to create trust in rational entities. BAN-Logic has a narrow scope whereas security evaluation can be very general.

3.1 Trust Representation in BAN-Logic

The BAN-Logic[BAN89] can be used to verify the correctness of security protocols. A security protocol is a formal specification of a message sequence with the purpose of reaching some specific security goals. An authentication protocol will for example establish the authenticity of a message which otherwise can not be trusted.

A typical problem when specifying security protocols is that they often get too complex to be fully understood even by their designers. As a result, serious vulnerabilities can remain undetected, and this has led to the need to formally verify the original formal specification.

The general principle of BAN-Logic is to transform the steps of a security protocol and its initial assumptions into logical formulas which can be manipulated by a set of logical rules or postulates to infer the conclusion of the protocol. It can then be seen whether the conclusions correspond exactly to the specified purpose of the protocol. As an illustration, the principle of propagating trust to the authenticity of a document is in the BAN-Logic captured as a postulate:

$$\frac{P \models \overset{K}{\mapsto} Q, P \triangleleft \{X\}_{K^{-1}}}{P \models Q \sim X}$$

Here, $\models$ means “believes”, and $\overset{K}{\mapsto} Q$ that Q has K as public key, so that $P \models \overset{K}{\mapsto} Q$ means that P believes that Q has K as public key. Further, $\triangleleft$ means “sees”, and $\{X\}_{K^{-1}}$ represents the message X signed with the private key K^{-1} , so that $P \triangleleft \{X\}_{K^{-1}}$ means that P sees the digital signature of X produced by the private key K^{-1} . Finally $\sim$ means “once said” so that $P \models Q \sim X$ means that P believes that Q once said X . The whole postulate or inference rule can then

be translated into: *If P believes that K is Q 's public key, and P sees message X signed with Q 's corresponding private key K^{-1} , then P will believe that Q once said X .*

We will not describe the other postulates or principles of the logic in more detail, but simply point out that there is no proof for the validity of the postulates other than their purely intuitive credibility. Such are all logics, and their validity depends directly on the validity of their axioms and postulates. Most people would agree that the postulate mentioned above intuitively makes sense, and this is the only basis for its validity.

With BAN-Logic it is possible to prove that a particular protocol produces the intended beliefs based on the initial assumptions *when the protocol is executed as specified*. Although this is important it is not sufficient. In fact, there are examples of attacks which are not captured by BAN-Logic, because the protocol is executed in a non-standard way (see e.g. [Syv93]). Also, BAN-Logic is unable to prove that no other set of initial assumptions can produce the same resulting beliefs. That can only be done by checking all possible combinations of initial beliefs. Although the BAN-like logics have weaknesses, they generally have a good reputation and are widely used to verify security protocols.

3.2 Trust from Evaluation Assurance

Security evaluation is based on a set of evaluation criteria and accredited evaluation laboratories which perform the evaluation under supervision of a national authority. A successful evaluation leads to the determination of an assurance level which shall reflect to which degree the TOE¹ can be trusted.

The idea of formally determining trust in computer systems was first developed within the military community and later spread to the commercial sector. The earliest set of criteria, TCSEC[USD85], defined a grading of 7 assurance levels, and when other criteria followed (e.g. ITSEC[EC92] and CC[ISO96]), they also tended to use 7 assurance levels as an attempt to maintain compatibility with TCSEC. When considering the large amount and the complex nature of the evidence which an evaluator must consider, there is nothing which logically indicates that this can be translated into a discrete value such as an assurance level. For this reason, it must be recognised that the idea of determining an assurance level for systems is as much a managerial requirement as it is a natural characteristic of a system, and that the definition of 7 assurance levels is totally ad hoc.

On the other hand, a more rich and thereby complex classification of assurance could easily become useless because users would not be able to understand it. That is in reality the dilemma we are facing; the simpler the classification of assurance becomes, the less it is able to reflect the diverse aspects of security, but on the other hand, the richer the classification, the less useful it becomes.

A user's trust in a system in reality is based on a variety of evidence. At least four types of evidence can be mentioned: 1) evaluation assurance, 2) assessment

¹ TOE = Target of Evaluation

of operational environment and threats, 3) security advisory reports, and 4) security incidents and intrusion detection. All these and possibly other types of evidence should be considered together in order for the user or system operator to get a good picture of the system's security state.

Security evaluation alone is not meant to be sufficient as a basis for trusting systems. If properly conducted, it can give an indication of the maximum security level a system can provide in its operational environment.

4 Models for Trust in Passionate Agents

Defining models for trust in passionate agents is intuitively a much bigger challenge than for rational entities. We will first analyse a rather general model for determining trust in remote agents of a distributed system, and subsequently look at a model for the weakening of shared control schemes as a function of trust between individuals.

4.1 Trust Representation in the BBK-Scheme

A good model for how the security of remote agents can be estimated would be a powerful aid to deciding which entities it will be most advantageous to interact with. Because remote systems are operated by humans and human organisations we are here faced with the problem of determining trust in passionate agents.

One contribution to this debate was given by Beth *et al* [BBK94]. The model which we will call the BBK-Scheme consists of a method for extracting trust values based on experiences from the real world, and secondly a method for deriving new trust values from existing ones within a network of trust relationships.

The Model. The system consists of entities which communicate via links. Each entity has a unique identifier and may have a secret which can be used for authentication purposes. The entities can generate, read and modify any message on any link. Entities may have some computational power e.g. for the encryption and decryption of messages. Some entities are distinguished as *authentication servers* as they support the authentication of other entities.

To model degrees of trust, the notion of *numbers of positive/negative experiences* is used. An entity can assign a certain value to each task it entrusts to another entity. This value can be thought of as the potential loss in case the task is not fulfilled. Each transaction increments the number of positive or negative experiences by one. Some of the central results from [BBK94] are listed below:

- (a) A *Trust Classification* is defined according to [YKB93] where each type of trusted task corresponds to a trust class, meaning that an entity can be trusted to perform specific tasks like e.g. key generation or clock synchronisation, without necessarily being trusted for other tasks. Each estimated trust value in the BBK-Scheme thus corresponds to a specific trust class.

- (b) A distinction is made between *Direct Trust* which consists of trusting an entity to perform a specific task, and *Recommended Trust* which consists of trusting an entity to recommend another entity for a specific task.
- (c) Definition of the trust level ν based on the confidence level α according to the conceptual formula:

$$\nu = \int_{\alpha}^1 f_x(\theta) d\theta \quad 0 < \alpha < 1 \quad (1)$$

where θ is the trust and $f_x(\theta)$ can be called the *trust density function* of an entity with regard to trust class x .

- (d) Estimation of Direct Trust ν_d , based on p positive observations, is done according to the formula:

$$\nu_d = 1 - \alpha^p \quad (2)$$

- (e) Estimation of Recommended Trust ν_r , based on p positive and n negative experiences, is done according to the formula:

$$\nu_r = \begin{cases} 1 - \alpha^{p-n} & \text{if } p > n \\ 0 & \text{else} \end{cases} \quad (3)$$

- (f) Estimation of New Direct Trust ν_n , based on a sequence of Recommended Trust ν_r , and Direct Trust ν_d , is done using a special ring-dot product:

$$\begin{aligned} \nu_n &= \nu_r \odot \nu_d \\ &= 1 - (1 - \nu_d)^{\nu_r} \end{aligned} \quad (4)$$

- (g) Estimation of Combined Recommended Trust when there are m different recommended trust paths between the trusting and the trusted entity, according to the formula:

$$\nu_{com} = \frac{1}{m} \sum_{i=1}^m \nu_i \quad \nu_i \neq 0 \quad (5)$$

Points (d), (e), (f) and (g) from the list above will be discussed below.

Analysis of the BBK-Scheme. Although the trust density function $f_x(\theta)$ is not stated explicitly in [BBK94], it is suggested by the proof of the formula (2) for Direct Trust. $f_x(\theta)$ is in reality an estimation of an entity's "trustworthiness" based on positive (and negative) experiences using Bayesian calculus.

Consequence of a long history. The paper states that no negative experiences are accepted if an entity is to be trusted at all, and when that is the case, the estimated trust level according to (2) can only increase. As a consequence, a long history either implies almost absolute trust or none at all. This can be misused by a malicious agent simply by cooperating during a certain period in order to accumulate high trust from other agents, and then to defect for a transaction of sufficiently high value.

Adaption to changing trustworthiness. Recommended trust is treated differently from direct trust in that negative experiences are accepted according to (3), which makes it possible to have a relatively low trust after a long history. But if the entity undergoes changes, its behaviour pattern may change, and so may the average numbers of positive and negative observations. A long history of positive observations before the change takes place will have as consequence that new observations will hardly have any influence at all on the trust level. A long history of roughly equally shared positive and negative experiences on the other hand is equivalent to a short history, and the trust will be very easily influenced by new observations. The conclusion is that (3) is unsuitable in environments where the trustworthiness of entities is likely to change.

Recommended trust. The method to derive New Direct Trust from recommended trust goes as follows: *A* trusts *B* with trust level ν_r to recommend someone else for a specific trust class. *B* trusts *C* directly with trust value ν_d for that trust class, so he recommends *C* to *A*. According to the ring-dot product of formula (4), *A* can then derive a new direct trust level ν_n .

From formula (4), we see that $\nu_d \rightarrow 1 \Rightarrow \nu_n \rightarrow 1$, disregarding the value of ν_r . This is illustrated in Fig.1 where the New Direct Trust is given as a function of the Direct Trust and the Recommended Trust.




Fig. 1. New Direct Trust = Direct Trust $\odot$ Recommended Trust

The application of the ring-dot product would for example lead to cases like: *If you tell me that you trust NN by 100% and I only trust you by 1% to recommend me somebody, then I also trust NN by 100%.* This is rather counterintuitive and would make any system which depends on this scheme extremely vulnerable to malicious manipulation.

Combination of trust values. When there are several recommendation trust paths between the trusting and the trusted entity, each yielding a different value,

the different Recommended Trust values are combine into a single value by taking the arithmetic average according to (5). The reason why neither the maximum nor the minimum value should be selected is explained by the fear of being overrun by a single unjustified value, without any further explanation of what “unjustified” means in this case. The fear that a trust value could be false after all, really is an expression of mistrust in the scheme itself. The consequences are important to notice, because it is the realisation that we are often faced with uncertainty about uncertainties, i.e. second and higher order uncertainties.

4.2 Trust Representation in Shared Control Schemes

A shared control scheme is a method for enforcing the concurrence of a pre-designated number of participants before a secret piece of information can be recovered or a jointly controlled event initiated. Such schemes are useful in cases where no single individual can be trusted to perform a given task, but where a group of individuals can.

The simplest shared control schemes are unanimous consent schemes, in which an action can only be performed by n designated individuals all acting in concert. Somewhat more complex are the threshold schemes, in which any k members of a group of n individuals can perform the action, but no $k - 1$ can. Because such schemes are applied to “individuals” they relate to trust in passionate agents.

Simmons and Meadows [SM95] found that what they call *additional trust* between the participants can create unintended combinations of participants which are able to execute an action. For example, if A and B are required to act concurrently to execute a task, and A entrusts B to fill his role, then suddenly, B can execute the task alone. [SM95] proposes a model for studying the consequences of additional trust in shared control schemes.

The Model. The concurrency Γ is the set of subsets of the participants who are able to execute a task, and the notation $\Gamma = x_1 + x_2 + \dots + x_n$ designates that the x_i ’s are the members of Γ . Each element x of Γ is in turn denoted as $A_1 \dots A_m$ where the A_j ’s are the elements of x . Thus $\Gamma = AB + BC$ designates a particular concurrency in which either A and B or B and C must act in concert in order to execute the task. For the additional trust, the notation $A \rightarrow B$ designates that A entrusts B to perform his role in the execution of the task. Simmons *et al* present a method which will determine the minimum additional trust which is needed to reduce Γ to a given weaker concurrency.

For example, if we want to determine the exact sufficient additional trust which leads to the weakening of concurrency $\Gamma_1 = AB + BC$ to $\Gamma_2 = B + AC$, written as $(AB + BC) \rightarrow (B + AC)$, the method will identify the following equivalence:

$$(AB + BC) \rightarrow (B + AC) \Leftrightarrow ((A \rightarrow B) \sqcap (B \rightarrow AC)) \sqcup ((C \rightarrow B) \sqcap (B \rightarrow AC)) \quad (6)$$

where $\sqcup$ is the logical “or” and $\sqcap$ the logical “and”.

The paper also proposes a method to determine all weaker concurrencies as a function of a given concurrency Γ . For example, it can be used to determine the hierarchy of concurrencies weaker than Γ_1 caused by additional trust. As can be expected, the situation where the task can be executed by each individual acting alone represents the bottom of the hierarchy, as illustrated in Fig.2.




Fig. 2. Lattice diagram of concurrencies weaker than $\Gamma_1 = AB + BC$

Analysis of the Model. The “entrusting” process denoted by $A \rightarrow B$ is in reality a transfer of capability, so that the model gives a method for determining minimum concurrencies as a function of capability accumulation. This capability accumulation is caused by additional trust. However, the capability transfer can happen in a multitude of ways, for example through threat, blackmailing, force, conspiracy, accident or even legitimate capability allocation. The cause behind the capability transfer, whether it is additional trust or something else, is in fact secondary to the model. It is the capability accumulation itself, and not trust, which is the central input parameter to the model. In addition, it is doubtful how B would be able to entrust AC according to $B \rightarrow AC$ in (6) because it would require B to set up a micro shared control scheme.

The model presented in [SM95] can be useful for analysing consequences of capability accumulation. In a real case, it would also be interesting to know the likelihood of this to happen. An assessment of the different causes can indicate the likelihood of the concurrency weakening to occur, and it can then be determined whether the scheme can still be trusted in that case.

5 Emerging Paradigms

In order to find appropriate models for trust it is absolutely necessary to understand exactly how trust is created. For example, knowing (or believing) that it is practically impossible to produce a correct digital signature without knowing the corresponding private key makes the observer believe that a message was signed by somebody who knew the secret signature key. This process takes place in the

observers mind, and a great deal of knowledge and experience is necessary for it to happen. We will define the *Belief Producing Process* as *the mental process which creates belief based on evidence and reasoning power*.

New knowledge can dramatically alter the whole process, for example if it suddenly was discovered that factoring large numbers no longer was hard. Belief is a state of mind and the Belief Producing Process is constantly active and ready to review old beliefs in the light of new evidence and experience.

Ultimately, it is this Belief Producing Process which directly or indirectly is the target of any formal model of trust, and it is absolutely necessary that it is thoroughly understood in order for the model to be of any value. More specifically, we will require 1) that the process can be clearly and explicitly described and 2) that the description is credible. In order to achieve the transition from subjective to common trust, it is in addition required that the credibility be commonly agreed upon.

5.1 Trust by Control or by Faith

For the purpose of modelling trust it would be interesting to know whether different types of underlying evidence give rise to different types of trust. For example, imagine two system managers within the same organisation, each operating their own system and each respecting the other manager's skills. Each manager has first-hand evidence about his or her own system but only second-hand evidence about the other manager's system. Although each manager's trust in the general security of both systems may be rated equal, the basis for the trust is very different. Trusting the own system is based on control whereas trusting the other manager's system is more based on faith.

This example illustrates two fundamentally different types of basis for trust. In case of rational entities, trust is typically based on control, and can be modelled for example with BAN-Logic or Security Evaluation Criteria. Trust in passionate entities on the other hand is typically based on faith, simply because passionate agents are difficult to control. Control is usually preferable to faith because it is based on rational and explicit evidence. Faith on the other hand is more a fuzzy feeling and is often based on vague and irrational evidence. There is no clear line between control and faith, and often trust is based on a combination of both. For modelling purposes everything must be made explicit, and it therefore seems easier to model trust which tends to be based on control.

5.2 Recommending Trust

It is usually assumed that trust is transitive, i.e. if A trusts B , and B trusts C , then it should be possible for A to trust C if only B recommended C to A . However, trust is *not* necessarily transitive. For example if person B has made a good deal with a particular shopkeeper C at the market, and B then recommends the shopkeeper to person A , transitivity would require A 's final trust in the shopkeeper C to be something like:

$$t(A \rightarrow C) = t(A \rightarrow B) \cdot t(B \rightarrow C)$$

However, it may be that the shopkeeper does not like person *A* and therefore would cheat him. Anybody with this knowledge would also know that the formula does not hold. In many cases it is probably possible to recommend trust, but that does not mean that it is true in general. It would be useful to define criteria for when trust can be assumed to be transitive.

In our opinion, a recommendation to trust a third party, whether it is rational or passionate, can only be considered as *evidence* so that in principle there is no such thing as a “chain of trust”. A recommendation can only come from passionate agents directly or indirectly, so that a chain of recommendations necessarily would have to be based on passionate agents. To expect that trust itself can propagate across more than one link in the chain would be a farce.

We do not know any case where proxies or security certificates are being transferred through multiple agents in a network purely based on trust between pairs of agents. All such applications take place within a hierarchy or a security domain with an authority to enforce the security policy, and in that case it is more than just trust between pairs of agents. For trust to propagate along a chain of agents, it is not only depends on exchanged evidence such as certificates, but also on circumstantial evidence such as liability of agents in the chain and enforcement of policies. This circumstantial evidence introduces a degree of control, without which the trust would be purely based in faith.

When comparing hierarchic structures such as X.509 [ITU89] with anarchic structures such as PGP[Zim95], the most apparent difference is of course the architecture, but even more important is that the former usually will be anchored in professional or government organisations with certain liabilities, whereas the latter typically is based on informal relationships and undefined roles. Both can be excellent for propagating trust, but the basis behind the trust is very different, and one may be more suitable for a particular application than the other. A formal hierarchic structure seems easier to model than an informal anarchic structure because the former has a higher degree of control.

5.3 Combining Conflicting Evidence

The question of how conflicting evidence can be properly combined is quite intriguing. The Dempster-Shafer Theory[Sha76] proposes a framework which makes it possible to distinguish between uncertainty and ignorance by allowing commitment of belief to a proposition without requiring any commitment of belief to its negation. The amount of uncommitted belief represents the extent to which the perceiver feels ignorant. Classical probability theory depend on the sum of probabilities of possible events to be 1, and the advantage of the Dempster-Shafer Theory is that it can be applied also in situations where this is not the case.

The Dempster-Shafer Theory does contain seemingly counterintuitive aspects. We will use an example from [Coh86] where two experts, E1 and E2, assign belief to three hypothesis H1, H2 and H3 according to table 1.

The third column shows the resulting belief in the three hypothesis when the conflicting advice from the two experts are combined according to the Dempster's

	E1	E2	E1&E2
H1	0.99	0	0
H2	0.01	0.01	1.00
H3	0	0.99	0

Table 1. Combined belief assignments according to Dempster's rule

rule. All the belief has been assigned to H2 despite the fact that neither expert put more than 0.01 belief in it. This result would probably be rejected in a real situation, and a perceiver would instead make the conclusion that one or both of the experts are mistaken. The perceiver has to decide which of the experts should be trusted the most and he would have to assign belief or trust in the experts themselves, and this implies that doubt about the experts are created by their conflicting advice. By assigning belief to the experts themselves we take the step into second order belief reasoning of which the mathematical framework is even more complex and less explored than for first order.

5.4 Understanding Human Behaviour

Neoclassic economy is perhaps the science which in recent years has claimed the biggest success in modelling human behaviour. It is based on the assumption that human beings are “rational utility-maximising individuals”, or in other words rational and selfish and always seeking to maximise their material well-being. Neoclassical economists have extended economic methodology to what are usually regarded as noneconomic phenomena like politics, bureaucracy, racism, the family and fertility[Bec76]. It is then a small step to extend this approach to information security, as already suggested by e.g. [RJ96]. However, there are at least two pitfalls to be avoided when using this approach.

Firstly, economic theory tries to predict how markets will behave in general, while still allowing for some individuals to diverge from the norm. The same would be true when applied to information security, so that this approach will not be appropriate to determine trust in a particular person, but rather to a group or a population.

Secondly, the assumptions behind neoclassical theory can only be partly true, as expressed by many critics (e.g.[GS94]). This does however not undermine the basic structure of the neoclassical edifice. That is, people will act as self interested individuals often enough for the predictions to be useful, but human beings act for non-utilitarian ends in irrational ways sufficiently often that the neoclassical model provides an incomplete picture of human nature. Trust is essentially an aspect of human nature that does not follow strict rational choices. It is for example partly thanks to mutual trust that strategy game deadlocks such as in the prisoners dilemma are resolved by humans in real life.

On this background, it still uncertain to which degree game theory and rational choice can be useful in modelling trust in passionate agents.

6 Conclusion

It is important to keep in mind that trust is a human cognitive phenomenon. By analysing four trust models, we have have invoked some of the challenges for understanding how the human mind conceives trust based on evidence, and this must always be the reference when defining trust models and judging their correctness.

References

- [BAN89] Michael Burrows, Martín Abadi, and Roger Needham. A logic of authentication. Technical report, DEC Systems Research Center, February 1989. Research Report 39.
- [BBK94] T. Beth, M. Borchering, and B. Klein. Valuation of trust in open networks. In *ESORICS 94. Brighton, UK*, November 1994.
- [Bec76] Gary Becker. *The economic approach to human behavior*. University of Chicago Press, 1976.
- [Coh86] M.S. Cohen. An expert system framework for non-monotonic reasoning about probabilistic assumptions. In L.N. Kanal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*. North-Holland, 1986.
- [EC92] EC. *Information Technology Security Evaluation Criteria (ITSEC)*. The European Commission, 1992.
- [GS94] Donald P. Green and Ian Shapiro. *Pathologies of Rational Choice Theory: A Critique of Applications in Political Science*. Yale Univ. Press, 1994.
- [ISO96] ISO. *Evaluation Criteria for IT Security (Common Criteria), documents N-1401/1404*. ISO/IEC JTC1/SC 27, 1996.
- [ITU89] ITU. *X.509, The Directory - Authentication Framework*. International Telecommunications Union, 1989.
- [Jøs96] A. Jøsang. The right type of trust for distributed systems. In C. Meadows, editor, *Proc. of the 1996 New Security Paradigms Workshop*. ACM, 1996.
- [RJ96] Lars Rasmussen and Sverker Jansson. Simulated social control for secure internet commerce. In Catherine Meadows, editor, *Proceedings of the 1996 New Security Paradigms Workshop*. ACM, 1996.
- [Sha76] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [SM95] G.J. Simmons and C. Meadows. The role of trust in information integrity protocols. *Journal of Computer Security*, 3(1):71–84, 1995.
- [Syv93] Paul F. Syverson. On key distribution protocols for repeated authentication. *Operating Systems Review*, 27(4), October 1993.
- [USD85] USDoD. *Trusted Computer System Evaluation Criteria (TCSEC)*. US Department of Defence, 1985.
- [YKB93] R. Yahalom, B. Klein, and Th. Beth. Trust relationships in secure systems - a distributed authentication perspective. In *Proceedings of the 1993 IEEE Symp. on Research in Security and Privacy*, pages 150–164, 1993.
- [Zim95] P.R. Zimmermann. *The Official PGP User's Guide*. MIT Press, 1995.