Divide and Conquer: The role of trust and assurance in the design of secure socio-technical systems

Ivan Flechais
Oxford University
Computing Laboratory
Wolfson Building
Parks Road
UK – Oxford OX1 3QD

Jens Riegelsberger
Department of Computer Science
University College London
Gower Street
UK – London WC1E 6BT

M. Angela Sasse
Department of Computer Science
University College London
Gower Street
UK – London WC1E 6BT

ivan.flechais@comlab.ox.ac.uk, {j.riegelsberger,a.sasse}@cs.ucl.ac.uk

In order to be effective, secure systems need to be both correct (i.e. effective when used as intended) and dependable (i.e. actually being used as intended). In this paper we look at the problems of achieving dependable security, given that all secure systems involve people and they can be unreliable. We present a brief look at the approach the Human-Computer Interactions in Security (HCISec) community propose to address this. We then discuss the nature and role of trust in secure systems, highlighting the social problems that can be the consequence of security policies and technologies that do not take trust relationships into account, and the benefits that ensue for systems that are designed with these factors in mind. We conclude by presenting some design principles that provide ways of making people more likely to behave in a secure manner thanks to a better understanding of the factors and motivations involved.

1. Introduction

Security deals with ensuring the deterrence, avoidance, prevention, detection and reaction to events in a system that are undesirable to the owner of that system.

That is to say that security is concerned with building countermeasures that effectively mitigate, handle, or completely avoid threats to the assets of a system.

In order to effectively counter threats, however, there are two different aspects that must be addressed:

- Correctness: the designed countermeasures will neutralise the threat if working as intended
- Dependability: the degree to which designed countermeasures are working as intended

"A computer is secure if you can depend on it and its software to behave as you expect" [8]. Although this definition is open to debate in that it implies that security exists in the reader's expectations of computer and software behaviour, it is useful in underlining the importance of dependability in computer security. It has been argued that an emerging sentiment in security research is *"correctness" is not the issue; "de-*

	Category	Description	Example
Technical countermeasure	Prevention	Stop attacks from happening	Firewalls, access control, etc.
	Detection	Notice and identify an attack	Intrusion detection systems, Automatic terrorist profiling
	Reaction	Stop or mitigate an attack	Automated response mechanisms linked to intrusion detection systems
	Deterrence	Discourage abuse	The visibility of technical countermeasures. E.g. CCTV
Social countermeasure	Prevention	Stop attacks from happening	Don't share passwords, lock your screen, have security guards on the gate
	Detection	Notice and identify an attack	Sysadmins, alert users, audit checking
	Reaction	Stop or mitigate an attack	Sysadmins or emergency response teams
	Deterrence	Discourage abuse	Prosecution, financial ruin, job loss, prison

Table 1. Technical and Social countermeasures

pendability" is' [4]. The point is that the ability to know how the system is going to behave is now being recognised as very important, in addition to building a system that actually counters threats.

A secure system is part of a wider socio-technical system whose goal is the achievement of a production task [1, 5, 22, 25]. As such it involves both human and technical components that work together in achieving the goal of securing that system. Dependability is therefore measured by the degree to which this socio-technical system behaves in the way it is expected to. Technical systems and policies are well defined and on the whole far easier to predict (even though composing different systems can have unexpected emergent behaviours). As the effectiveness of social engineering attacks [14] or the varying individual adherence to organisational security policies demonstrates [5], the behaviour of the social side of a secure system is much less dependable.

One answer to increasing the dependability of the security in a system is to limit the role that people have in that secure system, using technical countermeasures to deal with the unpredictability of people. This already seems to be happening in areas of technical detection mechanisms, for example with facial recognition for the identification of terrorists. Table 1 shows that there are technical and social countermeasures for every dimension of security (prevention, detection, reaction and deterrence). While the number of technical solutions is increasing, they cannot be expected to fully replace social countermeasures.

Technical countermeasures are very good at conducting repetitive security tasks, such as access control, virus checking or integrity checking. They start to become less

effective when less well-defined tasks are required, such as anomaly detection (i.e. intrusion detection), facial recognition or deterrence. Social countermeasures, although lacking the meticulous nature of technical countermeasures, are very flexible and can be extremely effective countermeasures – especially since deterrence tends to be a social mechanisms (for example prosecution and jail) It becomes evident that a good secure system takes advantage of the strengths of both types of countermeasures, whilst minimising their disadvantages.

People are essential in the operation of a secure system and can themselves be very effective security countermeasures. The problem is that people can be unreliable, whether for malicious reasons or not. Much of this revolves around the notion of trust; trust in people to behave in a secure manner, and trust between people as a means of fostering a good working culture. Throughout this paper, the notion of trust between two actors differs from the traditional security definition of trust, where a *"trusted system or component is one whose failure can break the security policy"* [3]. The consequence of the traditional definition is that the human and affective aspects of trust are overlooked in favour of the end result of a trust relationship. Instead of looking at trust from the point of view of the damage that can be caused by someone breaking that trust, we take the view that trust is *"an attitude of positive expectation that one's vulnerabilities will not be exploited"* (see section 3).

In the next section we will present a brief overview of some of the work and approaches the Human-Computer Interactions in Security (HCISec) community has taken to address the problems of people and security. This will be followed by a look at how trust relationships are built between two actors (whether human or not) and what key factors foster well-placed trust. We will then present a discussion as to how trust and security factors can work together in order to make people more dependable in their application of security policies, and therefore less likely to become unwitting victims of social engineering attacks. We conclude by introducing a number of design principles that may favour well-placed trust between organisations and their employees, as well as fostering a trusting environment within the organisation.

2. People and Security

Kahn [10], cited by Anderson [2], *"attributes the Russian disasters of World War 1 to the fact that their soldiers found the more sophisticated army cipher systems too hard to use, and reverted to using simple systems which the Germans could solve without great difficulty"*. This statement seems to expound the notion that good security is hard to use.

Bruce Schneier [23], however, makes the point that *"security is only as good as its weakest link, and people are the weakest link in the chain"*, indicating that good security has to acknowledge the weaknesses of people. Other authors [1, 9, 14, 15, 26] also argue that secure systems are broken through human issues, such as bad security configuration. They state that ease of use is necessary in order to get people to behave securely, and that good security is not necessarily hard to use [1].

As a consequence of this, the whole field of HCISec is largely focussing on building better tools [6, 24] and improving the user interfaces to these tools [9, 26]. We believe that this will undoubtedly result in improvements to the usability of security and as a consequence also improve its propensity to be used, however we also believe the user interface is only the tip of the iceberg when getting people to behave securely.

When Saltzer and Schroeder identified in 1975 the need for *'psychological acceptability'* [20] in secure systems, they were referring to the need for better interfaces. This notion effectively extends beyond user interfaces because a secure system is psychologically acceptable if the *user cost* (i.e. the sum total of the psychological, cognitive and physical load required of a user in a given task) of using it is not excessive compared to the *user benefits* (i.e. the incentives and advantages of engaging in a given task). This goes beyond the security user interface, and affects the user in the wider context of system use.

For example, in most organisations security is a secondary 'enabling task', compared to the primary 'production task' [21]. One of the costs of applying security is how much it will interfere with production tasks. Sample benefits of applying security might be the avoidance of sanctions, or the peer acceptance of a particular "security conscious" group. In organisations that are only concerned with productivity, and do not concern themselves with either enforcement or reward for secure behaviour, the cost of applying security is high compared to the benefits, and therefore the human element of the security is less likely to behave as intended. From the point of view of the organisation, this can be seen as a breech of trust: the organisation expects its employees to behave in a specific manner, and they do not.

3. Trust

Where does trust fit into the security picture? The term 'trust' is frequently used and misused in security literature – for example when referring to trusted paths and trust chains. Trust is not an absolute, objective or necessarily transferable phenomenon. In fact, trust can be defined as *"an attitude of positive expectation that one's vulnerabilities will not be exploited"* [12, 13, 18, 19] which is subjective, emotional and possibly non-rational – something that the security literature frequently ignores with it's own uses of trust, which tend to follow the lines that something is trusted when they mean that it is dependable.

A useful starting point when looking at the role trust plays in security is to identify which factors influence an actor's decision on engaging in a trust relationship. Actors can refer to people, but also to organisations, institutions, job roles (such as bank clerks, couriers or policemen), etc. The term trust relationship is only meaningful in a situation that is risky, that is to say where the actors stand to lose something or the outcome of the situation is uncertain. For example this can be seen in online transactions, where the customer pays the vendor in the expectation that the vendor will send the desired goods. The customer cannot ensure compliance from the vendor and has to trust that they will keep their side of the bargain. The vendor, on the other hand,

has the option to default on sending the goods and a number of factors can influence this decision.

The following factors can be seen in Figure 1, which is the result of the analysis of the mechanics of trust between a *trustor* (i.e. trusting actor) and a *trustee* (i.e. trusted actor) by [18]. It consist of a number of factors that affect how trust is signalled, how these signals are understood and how they affect a given trust relation. The main factors consist of:

- Intrinsic properties
  - Motivation, Ability, Internalised Norms and Benevolence
- Contextual properties
  - Temporal, Social and Institutional embeddedness.

Both contextual and intrinsic properties play a role in the establishment of a trust relationship. Intrinsic properties refer to factors that are internal to the trustor and trustee, such as the propensity to take risks, benefits of engaging in a trust relationship and personal cost of breaking trust. Contextual properties refer to factors that exist outside both actors, such as law enforcement, expectations of future interactions or reputation.

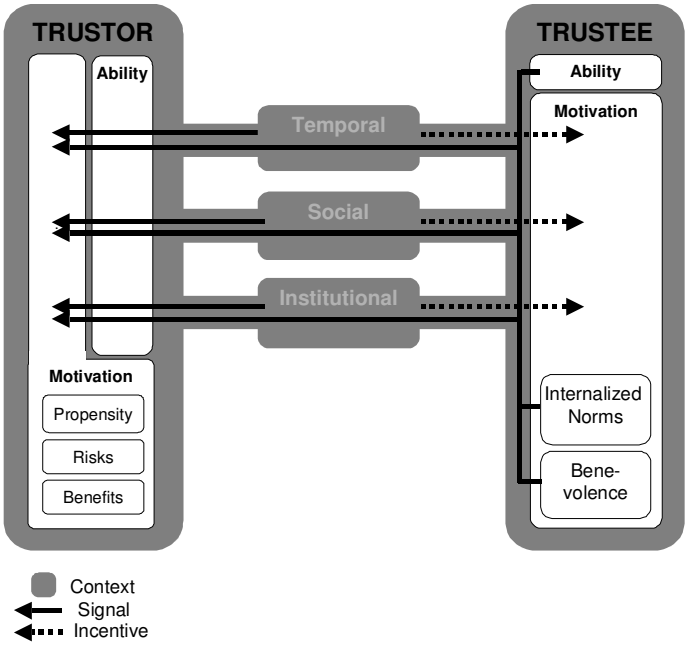

Fig. 1. Model of Trust

Intrinsic Properties

Trustor: Motivation
Motivation refers to an actor's incentive for engaging in a trust relationship. It is affected by factors such as propensity to trust, perception of risk, benefits of engaging in the relationship and the availability of other options that may achieve similar results. These are subjective characteristics that vary betweens actors. Propensity to trust relates to the trustor's inclination to be trusting. The perception of the risk of engaging in a trust relationship refers to the potential for loss, not only financial, but for example the psychological cost of having been naïve. The benefits capture what the actor stands to gain from a successful trust relationship, such as financial profit, a reduction in cognitive effort, time saving, etc. The availability of other options also affects the likelihood of a trust relationship forming should these other options be perceived to be less risky, more beneficial, or less needing of trust. Finally the signals about the situational and intrinsic properties affecting the trust relationship play a significant role in the motivation to engage in the relationship.

Trustor: Ability
For the trustor, this refers to the individual knowledge and understanding of the signals and situations which affect the formation of a trust relationship. For example, a trustor's assessment of the risk inherent in a given trust relationship is affected by both past experience, as well as new knowledge. An employee may be happy leaving new cleaning staff in their office alone, but may feel differently should a theft occur.

Trustee: Ability
The trustee's ability to actually achieve a given task. The trustee may be willing but unable to actually perform in the manner expected. A good example of this can be seen in the number of password resets that are necessary as a consequence of people forgetting them [22]. People may be willing to have different passwords for different systems, but the limitations of human memory make it exceedingly difficult for anyone to remember a large number of complicated passwords without writing them down – which contravenes the security policy.

Trustee: Motivation: Internalised Norms
Actors have been observed to behave in a trustworthy manner despite not having any external incentive to do so. Partly this can be put down to habit, but also to internalised norms that affect that actor. These norms can induce the actor to behave in an untrustworthy manner, for selfish actors motivated only by immediate gain. They can also induce the actor to behave in a trustworthy manner, a trait that can be referred to as *integrity*:*"… the trustee adheres to a set of principles that the trustor finds acceptable."* [12]. For example, some people based on their upbringing, are more diligent in following the rules.

Trustee: Motivation: Benevolence
"Human behaviour in romantic relationships is an example of trustworthy action motivated by strong feelings of benevolence. In such relationships the well-being of the other forms part of one's own gratification. Benevolence – albeit to a lesser degree – also applies to relationships between work colleagues or friends." [18] In relationships of benevolence, actors do not expect immediate or equal returns, but this sentiment only evolves over time, and after a number of successful trust exchanges. This factor can be crucial in breaking security policy. Social engineering relies on the target's benevolence and willingness to be helpful in order to break policies.

Contextual Properties

Temporal embeddedness
This is the notion that two actors' decision to engage in a trust relationship is affected by their expectation of future interactions. This is one of the reasons why, for example, many disgruntled employees are willing to vandalise and cause damage to systems they have access to, since they have no expectations of future interactions with their employers.

Social embeddedness
Social embeddedness represents the group interactions and ensuing reputation that an actor gains from his behaviour towards members of that group. The incentive to behave in a trustworthy manner is no longer linked to future interactions with a single actor, but to future interactions with actors likely to hear of their reputation. Within organisations, this can be a powerful motivator for a newcomer to conform to the existing security behaviours (or lack thereof) inside their immediate peer group.

Institutional embeddedness
This refers to the organisations (e.g. employee's company, ethics committees, or consumer rights groups) or institutions (e.g. law) that have the power to sanction untrustworthy behaviour or behaviour that is below expectations. Here the significant factor for engaging in a trust relationship is governed by the type, strictness and severity of punishment. This type of sanctioning is governed by strict rules defining what the institutions' expectations are in a given situation. An example of this could be the threat of being excluded from a professional group as a result of objectionable behaviour under a given code of practice.

This type of deterrence is currently the most widely used means of *assuring* compliance to the security policy.

Assurance vs. Trust

Assurance consists of the *contextual* factors that organisations can put in place to ensure a specific outcome to a trust exchange. These are currently mainly restricted to detecting and sanctioning undesirable behaviour. *Trust* on the other hand is based on

an understanding of the *intrinsic* properties that pertain to a given actor. High-security endeavours seek to identify whether an actor is intrinsically trustworthy by conducting background checks. These are intended to determine whether the actor has any past evidence of law-breaking, indicative of individuals whose integrity may be less than satisfactory.

As stated by [18], organisations are more productive if they have *social capital* [16] – i.e. trust that is based on shared informal norms that promote cooperation [7]. Some authors claim that reported failures of systems to yield the expected productivity gains in organisations [11] partially stems from a reduction in opportunities to build social capital [17].

Security policies are generally designed to encourage actions that can be readily interpreted as untrusting. For example, refusing to share a password with a colleague, locking your computer screen or checking the credentials of a technician are all signs of distrust in any usual setting and are considered to be basic security practices.

The design of current secure systems rarely considers the need for – or existence of – trust between the different operators running the system. The attitude that prevails in system design is that the operator of the system must and will perform a task, with little to no thought going into how this will affect him in the wider context of his organisation. Ignoring this issue can lead to damage to the formation of social capital in the organisation, or even provide a means of forming social capital through the breaking of security practices, i.e. employees bonding together in the knowledge that everyone is breaking the rules with them.

For example, it may be that for confidentiality purposes a medical data provider has specified a policy that separates different kinds of medical data (i.e. general health, sexual health, aids, cancer, etc.) and restricts access to these. There are cases where a particular organisation has a number of different projects utilising these resources, all needing different access privileges so as not to be given unnecessary information. This can put that particular organisation in a position where the different projects combined have access to all areas of information from the provider, yet the policy is designed to prevent some projects from accessing parts of this information. So, although the organisation as a whole is trusted with the totality of the confidential medical information, the policy requirements result in particular individuals being allowed more access than others. Lacking any justification for this measure, this can be interpreted as a lack of trust of an organisation in its employees, as well as resulting in increasing the administration costs of maintaining the access control.

Is it possible to foster a secure environment – i.e. trust employees, whilst at the same time not undermining the trust relationships that have to build in order to build social capital?

4. Analysing Trust and Security

Where Trust and Security meet

The role of security in an organisation is to dependably handle the threats to the assets of that organisation. In order to do this, both technical and social countermeasures are necessary, and ensuring that these countermeasures are actually applied is of equal importance. In an ideal world, from the organisation's point of view, there should be no need to trust any of the employees within the organisation, because the rules and procedures in place would be sufficiently reliable as to avoid any risk of employees acting undesirably.

In the case of well-defined, repetitive and predictable tasks, this lends itself well to creating and enforcing a policy that compels employees to apply security, whilst preventing them from abusing the system. A good example of this can be seen in the banking sector which has evolved a vast number of procedures, both technical and social, to prevent employees from stealing money. The disadvantage of using these types of *contextual* measures is that it takes away a lot of flexibility, and makes organisations slow to respond to new situations.

In areas where job requirements are vague, or there is a specific need for flexibility, these kinds of rigid policies cannot be made to work because they are too complex, constraining or costly. In these cases, the only available option is to choose, encourage and trust employees to behave in a secure manner, rather than enforce it.

Breaking Trust

As we have seen in section 3, security policies can require operators to behave in an untrusting manner. For example a trustor requires a trustee to divulge his password in order to allow the trustor to finish some urgent work. In this exchange the *motivation* for the trustor to engage in this trust relationship is that he has a high potential *benefit* – i.e. finishing the urgent work – and the other *options* are more inconvenient and time-consuming – getting his password reset or reissued.

The trustee has the *ability* to divulge his password, and may feel *benevolent* which may influence him to choose to divulge in order to help a colleague. On the other hand the trustee may have a degree of *integrity* that prevents him from behaving in such a way as to disobey the security policy. The motivation to refuse or accept to share the password is also affected by external factors.

Expectations of *future involvement* may tip the balance in favour of breaking the security, since it is very likely that the trustee will interact again with the trustor as they are colleagues. The trustor may also be a part of a larger group of colleagues and in cases where security is not important to this s*ocial group*, they might give the trustee a bad reputation or affect the relationships between him and the group should he decide to refuse. The final factor is the degree to which the *organisation* detects and punishes transgressions and rewards good behaviour.

Should the trustee refuse to violate the security policy, unless the trustor understands and agrees with the motivation to do so, he may feel untrusted and untrust-

worthy, which can create tensions between the two actors, and will definitely hinder the creation of social capital. In this case three trust signals can positively influence the adherence to the security policy without harming the trust relationship between both actors:

- Providing an alternative to initiating the trust exchange in the first place. This can be done by giving the trustor an easy way of accessing the systems he needs, for example by reissuing his password in a timely manner or providing a limited access based on an anonymous login and monitoring the activity of the trustor. This is the kind of approach that HCISec is trying to achieve by making it easier to use secure systems.
- Having a security conscious culture within the peer group. Both actors can relate to this, even if they do not overtly recognise it (i.e. if everyone is careful with their passwords and refuses to divulge them, then the peer pressure to behave in the same way is significant).
- Ensuring that the detection and punishment for breaking the security rules are effective. Very stringent enforcement of security policies will result in adherence to the policy. This is a very straightforward means of preventing rule breaking because it is easily understood by both trustor and trustee, who have a lot more to lose than gain.

Coming back to the banking example, the stringent security measures in place do not create tensions amongst staff because it is well understood that the detection and punishment for a transgression is taken very seriously. This in turn can foster an environment where no one breaks the rules, thereby reinforcing the motivation to avoid transgression. Stringent enforcement can only happen in the areas where the expectations are as well defined as the punishments. As we have seen above, the disadvantage of this approach is that it stifles flexibility and this makes it inappropriate for a significant number of jobs that require security.

Middle Ground

As seen in section 3, there are two extremes of security:

- Assurance: complete control over what employees must and can't do, together with stringent enforcement.
- Trust: no control over what employees can do, and only trust and encouragement for them to behave in a secure manner.

The problems start to occur when trying to secure a system that exists in the middle ground of being able to support well-defined security policies, whilst still requiring a degree of flexibility. In cases like this where the security policy in place is either not well-defined (in order to maintain flexibility) it is essential that the enforcement of that policy be both strictly specified and applied to everyone in the organisation. In addition, it is essential to foster an environment which encourages employees to behave in a trustworthy manner. Table 2 describes a set of design principles that make use of the trust warranting factors identified in section 3.

Design principle	Description	Relevant Property
Simplifying security	Make the task of behaving securely easier through better tools and simpler interfaces but also through simple policy rules – exceptions to the rules can be sources of confusion and abuse.	Ability Motivation (other options)
Promoting a security culture.	A security culture should be encouraged by ensuring that the security policy is neither excessive (i.e. for every countermeasure there is a corresponding threat) nor unfair (i.e. the boss is allowed to avoid security measures). In addition to this checks should be regularly made to ensure the policy is in use, transgressions are detected and punished according to a published code of conduct, and secure behaviour is rewarded.	Ability Social embeddedness Organisational embeddedness Motivation (avoiding punishment, benefiting from reward)
Group membership Group identity	Specifically grouping people into security groups, together with their own responsibilities and rewards can make security a more immediate concern for employees. By making the groups smaller expectations of future interactions are greatly increased, thus harnessing that particular factor.	Ability Temporal, Social and Institutional embeddedness Benevolence
Educating employees about security	By providing employees with training as to what is expected and required and what are the threats.	Ability Motivation (Perception of Risk) Benevolence

Table 2. Principles for fostering dependable behaviour from the social elements of a secure system

5. Conclusion

Getting a secure system to behave dependably is a complex task. Assurance mechanisms can achieve a degree of success, but in most real-world situations, organisations either cannot afford the costs of maintaining such a stringent system, or need to be flexible. This means that these systems have to rely on people behaving in a secure manner. We have looked at the field of trust and identified a number of fac-

tors that affect an individual's propensity to behave in a trustworthy manner. We are convinced that these factors can be applied to improving the dependability of an individual's security behaviour, and have presented a number of design principles aimed at addressing this.

6. References

[1] Adams, A. & Sasse, M. A. *Users Are Not The Enemy.* Communications of the ACM 1999. Vol. 42, No. 12 December

[2] Anderson, R. *Why Cryptosystems Fail.* ACM Conf.Computer and Communication Security CCS'93 1993. pp 215-227.

[3] Anderson, R. *Security Engineering: A Guide to Building Dependable Distributed Systems*. 2001. Wiley.

[4] Baker, D. *Fortresses built upon sand.* Proceedings of the New Security Paradigms Workshop 1996.

[5] Brostoff, S. & Sasse, M. A. *Safe and Sound: a safety-critical approach to security design.* New Security Paradigms Workshop 2001.

[6] Dourish, P. & Redmiles, D. *An Approach to Usability Security Based on Event Monitoring and Visualization.* Proc.New Security Paradigms Workshop (Virginia Beach, VA) 2002.

[7] Fukuyama, F. *Social Capital and the Civil Society.* 2nd Conference on Second Generation Reforms 1999. Washington, DC: IMF.

[8] Garfinkel, S. & Spafford, G. *Practical UNIX and Internet Security*. 1996. O'Reilly.

[9] Ka-Ping, Y. *User Interaction Design for Secure Systems.* 2002. http://zesty.ca/sid

[10] Kahn, D. *The Codebreakers*. 1967. Macmillan.

[11] Landauer, T. K. *The Trouble with Computers: Usefulness, Usability, and Productivity*. 1996. Cambridge, MA: MIT Press.

[12] Mayer, R. C., Davis, J. H., & Schoorman, F. D. *An Integrative Model of Organizational Trust.* Academy of Management Review 1995. 20(3) , pp 709-734.

[13] McAllister, D. J. *Affect- and Cognitation-based Trust as Foundations for Interpersonal Cooperation in Organizations.* Academy of Management Journal 1995. 38(1) , pp 24-59.

[14] Mitnick, K. D. & Simon, W. L. *The Art of Deception: Controlling the Human Element of Security*. 2003. John Wiley & Sons Inc.

[15] Poulsen, K. *Mitnick to lawmakers: People, phones and weakest links.* 2000. http://www.politechbot.com/p-00969.html

[16] Putnam, R. D. *Bowling Alone: The Collapse and Revival of American Community.* 2000. New York: Simon & Schuster.

[17] Resnick, P. *Beyond Bowling Together: SocioTechnical Capital*. HCI in the New Millenium 2002. pp 242-272. Boston, MA: Addison-Wesley.

[18] Riegelsberger, J., Sasse, M. A., & McCarthy, J. *The Mechanics of Trust: A Framework for Research and Design.* International Journal of Human Computer Studies 2004. 62(3), pp 381-422.

[19] Rousseau, D. M., Sitkin, S. B., Burt, R. S., & Camerer, C. *Not so different after all: A cross-discipline view of trust.* Academy of Management Review 1998. 23(3) , pp 393-404.

[20] Saltzer, J. H. & Schroeder, M. D. *The protection of information in computer systems.* IEEE 1975.

[21] Sasse, M. A. *Computer Security: Anatomy of a Usability Disaster, and a Plan for Recovery.* CHI 2003 2003.

[22] Sasse, M. A., Brostoff, S., & Weirich, D. *Transforming the 'weakest link': a human-computer interaction approach to usable and effective security.* BT Technical Journal 2001. 19 , pp 122-131.

[23] Schneier, B. *Secrets and Lies*. 2000. John Wiley & Sons.

[24] Smetters, D. K. & Grinter, R. E. *Moving from the design of usable security technologies to the design of useful secure applications.* New Security Paradigms Workshop.September 23-26, 2002, Virginia Beach, VA 2002.

[25] Weirich, D. & Sasse, M. A. *Pretty Good Persuasion: A first step towards effective password security in the real world.* New Security Paradigms Workshop 2001.

[26] Whitten, A. & Tygar, J. D. *Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0.* Proceedings of the 8th USENIX Security Symposium, August 1999, Washington 1999.