Holistic Trust Design of E-Services

Stéphane Lo Presti[1], Michael Butler[2], Michael Leuschel[3] and Chris Booth[4]

[1] Royal Holloway, University of London, Information Security Group, Egham, Surrey TW20 0EX, United Kingdom; Phone: +44 (0)1784 414346; Fax: +44 (0)1784 430766; Stephane.Lo-Presti@rhul.ac.uk

[2] University of Southampton, School of Electronics and Computer Science, Southampton SO17 1BJ, United Kingdom; Phone: +44 (0)23 8059 2435; Fax: +44 (0)23 8059 3045; mjb@ecs.soton.ac.uk

[3] Institut für Informatik, Heinrich-Heine Universität Düsseldorf, Universitätsstr. 1, D-40225 Düsseldorf; Phone: +49 (0)211-8110711; Fax: +49 (0)211-81-10712; leuschel@cs.uni-duesseldorf.de

[4] QinetiQ ltd, WR14 3PS, Malvern, United Kingdom; Phone: +44 (0)1684 89 6400; Fax: +44 (0)1684 89 4389; cjmb@signal.qinetiq.com

Holistic Trust Design of E-Services

Abstract. As a central issue of modern e-services, trust has to be tackled early during the development phases. We present and compare in this chapter various works and methodologies that contribute to this aspect. A Holistic Trust Design Methodology that combines useful aspects encountered in the existing works is then described in detail. It is based on a systematic analysis of scenarios that describe the typical use of the e-service by using a Trust Analysis Grid. The Trust Analysis Grid is composed of eleven Trust Issue Categories that cover the various aspects of the concept of trust and is used to guide the design of the computing system by analyzing and refining the scenarios, and providing hints at the suitability of technologies for the scenario. We illustrate this methodology on several examples.

INTRODUCTION

Trust has recently been recognized as a crucial and central property of modern systems that provide e-services in a variety of contexts. Because failing to address this issue correctly may have a profound and costly impact on the e-service development, the issue of trust must be tackled early during the development, so as to identify and mitigate it as early as possible. This chapter covers methodologies that help to do so.

Trust is a human notion that goes beyond technical aspects of the system. It is important that it is not confused with other concepts, e.g. security, so that users understand and thus have confidence into the system. This aspect is reinforced by the rapid growth of e-services developed in, for example, pervasive computing (Huandget al., 1999) or multi-agent systems (Hanssens et al., 2002).

Trust defies traditional analysis in that it encompasses a wide range of other issues at a high level of abstraction, for example security, risk, social engineering or the law, in an ever-increasing complex arrangement. The recent literature on trust (see (Rindeback & Gustavsson, 2004), (Josang et al., 2005) or (Staab et al., 2004) for example) shows a number of ways with which trust can be dealt. But the literature lacks a holistic point of view that can help understand which techniques or technologies are best in various contexts and circumstances.

The design phase of the system development is the most appropriate time for analysis of trust in the system. This is the approach used to tackle more traditional issues like risk (Storey, 1996) and security (Anderson, 2001), and it has proven successful in improving the quality of systems. It can be seen as a process whose output is a set of requirements that must be addressed in the subsequent phases of the development.

Based on those two ideas of holistic design, trust is considered in this chapter as an evolving, contextual and composite belief that one principal (trustor) has that another principal (trustee) will perform certain actions with certain expected results, when not all information about those actions is available. The various elements of this definition will be detailed in the remainder of this chapter.

The first section presents current works on methodologies to help design trustworthy e-services. Then we present a methodology that builds upon the current understanding

of trust and improves on the existing trust design methodologies. It provides a holistic analysis framework to help design trustworthy e-services where the user is the focus of attention. This framework is applied to several realistic systems under development, including e-health and e-learning, in the next section.

EXISTING METHODOLOGICAL WORK TACKLING THE TRUST ISSUE

There is a huge corpus of work on the issue of trust, but few concentrate on this issue during the design of a system development and fewer propose methodologies to help the design process. We present here representative contributions to this topic and discuss their advantages and disadvantages as holistic trust design methodologies.

Typologies of Trust

There are many general works that analyze and decompose the notion of trust so as to provide general guidelines and understanding into this notion. Some of them suggest in particular the holistic nature of the concept of trust, by combining together the various aspects of trust presented in other works.

Though these works provide a necessary and fundamental insight into the notion of trust that any holistic trust design methodology should provide, and in particular a basic decomposition into its more elemental properties, they are not sufficient by themselves for devising an e-service design methodology as they provide no concrete help to the e-service designer.

TRUST-EC e-Commerce Requirements

The TRUST-EC project (Jones & Morris, 1999) lists the common applications in e-Commerce and analyzes how they consider trust, taken as a kind of reliance of system stakeholders that is more general than dependability. The list of non-functional requirements for trust for the e-Commerce business derived from this analysis is comprised of confidentiality, integrity, availability, identification, prevention, traceability, quality, risk management, and authentication. This work concludes on the possible necessity of *"a requirements process for e-business, based on the framework, which will assist developers in structuring both the process of eliciting trust requirements (by providing a checklist of issues to be discussed with different stakeholders), and the way in which such requirements are documented"* (Jones & Morris, 1999).

McKnight and Chervany's Conceptual Framework

McKnight and Chervany (McKnight & Chervany, 1996) analyze research papers on trust, taken from a wide range of domains: management, sociology, economics, politics, science, psychology. They summarize the various views on trust in a conceptual framework based on six constructs, which themselves are decomposed into more elements. The basic concepts of the generic definition of trust of McKnight and Chervany are trusting intention, trusting behavior, trusting beliefs, system trust, dispositional trust, and situational decision. These general concepts are refined into notions like feelings of security, vulnerability, honesty, situational normality, belief-in-person, or trust stance. This work is probably the one that provides the best view on the holistic nature of trust, in that it gathers information from domains whose problems are quite orthogonal, and it shows the intrinsically diverse and interwoven

structure of trust. On the other hand, it covers such a wide scope of trust issues that it is difficult to use it concretely.

Jøsang, Ismail and Boyd's Analysis of Trust-related Applications and Models

Jøsang, Ismail and Boyd (Josang et al., 2005) give a more technical overview of the research on trust, in particular by focusing on modern reputation systems. They describe five trust classes (borrowed from Grandison and Sloman's classification (Grandison & Sloman, 2000)) that describe the various layers of the concept of trust, namely provision trust, access trust, delegation trust, identity trust, and context trust. These classes are defined in the context of a trust purpose, which enables to instantiate the classes in a given situation. The semantics of trust considered in a given application is then described against two dimensions, the levels of subjectivity and specificity. Many examples of modern systems using reputation are presented, for example eBay (EBay, 2006), or Slashdot (Slashdot, 2006).

General Methodologies

This section presents methodologies to help in various aspects of the design of systems and e-services where trust is central. The following methodologies are based on ad-hoc concepts and methods but will show the possible approaches currently taken to design trustworthy systems.

Security Requirements Specification Method

Tan, Titkov and Poslad (Tan et al., 2002) devise an abstract security model that they apply to the case of an e-Banking service implemented as a multi-multi-agent system. This model is based on the concepts of asset, safeguard to protect the asset, threat to the asset and profile to specify policies and relationships between the previous concepts. A graphical notation is used to illustrate the model, where security domains are shown as groupings of concepts and profiles. The various assets are then described more precisely in the case of the e-Banking scenario, where the various services needed are modeled as assets.

That analysis enables to discuss the various service design issues in the e-Banking scenario. But first the methodology is aimed at Multi-Agent Systems, though it could still be used more generally at a conceptual level, and secondly, the analyzed issues are related to the more technical side of trust, namely security, and they require expert knowledge (authentication, security policies). The ideas underlying this approach are good, but not developed enough to make it a holistic trust design methodology.

Matrix Model

Tan's trust matrix model (Tan, 2003) is a means to analyze trust-building services between trading partners in e-commerce. It proposes to represent an e-service in the form of a grid. The grid rows correspond to properties of the service grouped into three layers. The grid columns correspond to a theoretical decomposition of the notion of trust into four reasons, namely social signs, personal experience, understanding, and communality. Each reason is divided into two sources, depending on whether they correspond to trust created by a party of the transaction or a control mechanism.

The trust analysis in this framework is suited to the examination of a particular service offered by a system. It is also quite precise in that it considers a lot of trust issues, but those issues are specific to the kind of services examined in this work, namely Business-to-Business first trade situation. This approach is very rich in terms of the holistic view it provides into trust, but weak in terms of design features as it is not aimed at this task.

Decompositional Model of Trust used in System Design

Yan and Cofta (Yan & Cofta, 2003) define trust domains as areas of mobile communication where the definition of trust, which is a set of statements and goals, is common between the various elements. Gaps between these trust domains, implied by the subjectivity of the trust definitions, are bridged with particular components that are responsible for ensuring trust at a level above the one of the domains. The methodology is based on a graphical representation of entities, domains and their interconnections, enabling a view of the system at a higher level of abstraction.

This methodology brings intuitiveness to system design and enables to treat the trust issue at a variable level of granularity, but it lacks accurateness to express more specific properties and localized problems, and it does not provide clear guidelines to help the system designer understand the various trust issues. It is in particular unclear what the application of the methodology would bring to examples other than the one presented.

Holistic Analyses of Trust Relationships

In the context of virtual communities, Ishaya and Mundy (Ishaya & Mundy, 2004) indicate that the potential barriers to trust development and management are fivefold: sociological, psychological, technological, legal, economic. They develop mechanisms to support the provision of trust, summarized by three tables that present relevant questions to ask about the elements of trust, the trust building process, and the security factors.

Similarly, Grimsley and Meehan (Grimsley et al., 2004) consider trust from the perspective of internet-mediated community relations and decompose it into various dimensions, which are not necessarily orthogonal. The community trust compact and an experience management matrix described in this work are based on three dimensions, corresponding to the notions of information, control and influence. These two conceptual tools are represented by tables guiding the design of the project under consideration.

These two works present a holistic view on trust mainly focused on the socio-psychological phenomenon and where the subjective aspects of trust are captured by many concepts and notions expressed in long questions. These textual approaches at design are not satisfactory as they add a layer of subjectivity and are prone to errors. Furthermore they are very difficult to apply to technical systems, as they mainly consider abstract properties of human systems.

Framework to derive Trust Assumptions

Haley et al. (Haley et al., 2004) apply principles of requirements engineering to security systems with the goal of deriving trust assumptions during the system-level analysis. A graphical notation is used to define domains. Interfaces, phenomena and

constraints are described in context and problem frame diagrams. Constraints explored in this work are security requirements that help identify threats and vulnerabilities at the system-level. The constraints are completed by trust assumptions that help satisfy the security requirements and that specify the system designer's point of view on trust.

This approach provides powerful features for e-service design, but it considers trust in the technical sense, i.e. implicit assumptions in the security model that need to be made explicit. The subjective, and thus holistic, nature of trust is not represented. Though apparently intuitive, the graphical notation can rapidly lead to cluttered diagrams, which impede on an effective design.

The *i** Graphical Design Framework

The *i** framework (Yu & Cysneiros, 2002; Yu & Liu, 2001) is a framework proposed to model non-functional requirements (privacy and security) in multi-agent systems. A composite graph is used to represent the relationships between actors of a system. Relationships are of four types: goal, task, resource and *softgoal* (goal that have no qualitative measure of satisfaction and are attached to graph edges). Trust relationships are expressed as softgoals. The SD (Strategic Dependency) and SR (Strategic Rationale) models enable to view the system at various levels of abstraction. An *i** graph looks like a scenario annotated with interrelated keywords.

The *i** framework is interesting as it focuses on a user-centered scenario describing the e-service under consideration in its context of daily use and uses intuitive graphical features to represent properties. But as illustrated in (Yu & Cysneiros, 2002) *i** graphs can very quickly become unreadable, and thus unusable to the e-service designer.

Methods with a Formal Background

This section covers some methodologies that use formal methods to design and specify the electronic system. Computing science formal methods stem from mathematics and aim to help design, develop, analyze and validate software so that it is correct, error-free and robust. The formal dimension of these methodologies provides a basis for automated reasoning and tool support, thus greatly easing the work of the designer and opening the door for systematic analysis.

TROPOS

TROPOS (Giorgini et al., 2004, 2005) is a graphical methodology for modeling and representing the various system dependencies with the aim of analyzing trust requirements in security properties of electronic systems. TROPOS extends the i* framework (see above) by defining new concepts (actor, resource) and relationships (delegation, negative authorization) with the goal of focusing on trust. The diagrammatical notation is formalized into the logical language Datalog so that the specification can be checked automatically with software. A CASE tool has been devised to draw TROPOS diagrams and checking its correctness.

TROPOS builds upon i* and extends it with more precise security notions. It removes the disadvantage of needing to manage the complexity of its graph thanks to the

formalization. On the other hand, this approach requires some knowledge of formal verification that may not be available in the design project.

CORAS

The CORAS methodology (Braendeland & Stolen, 2004; Vraalsen et al., 2005) implements risk assessment techniques using an extension of the UML semi-formal method. The expressiveness of UML profile is extended conceptually by introducing normative modalities borrowed from deontic logics so that legal risks can be better specified. It covers the risk dimension of trust, but does not explicit any definition of trust. It requires and enables the system designer to use the definition most suited for his system. Various tools are available to use the methodology (Vraalsen et al., 2005) and a fully-formal view of CORAS is under development.

The CORAS approach is close to software engineering as it uses the UML notation to model the system and it bases its analysis on a scenario. In these scenarios, a hierarchy of assets describes trust that must be protected from threats, vulnerabilities, and incidents. Evaluating in detail the risks associated with the system under examination enables to propose the solution to trust issues. The CORAS analysis methodology is partitioned into five sub-processes that are: establish the context; identify risks; analyze risks; evaluate risks; treat risks. This sequence is completed by a monitoring and review process that run in parallel and can restart the sequence of sub-processes.

Conclusion

A study of the existing works reveals that a decomposition of trust is central to any trust design methodology. General methodologies focus on a functional subset of properties, more well-known and less subjective, except for Holistic analyses of trust relationships (Grimsley et al., 2004; Ishaya & Mundy, 2004). But this limitation removes the holistic nature of trust, in that focusing on a sub property while eluding the others transforms the issue from trust to a different concept (e.g. security, risk).

Most methodologies adopt an ad-hoc approach, without building on traditional design methodologies. This makes more difficult and costly the tackling of the trust issue in e-service design, and thus constitutes an impediment to the adoption of such methodologies in e-service development. While methodologies using a graphical representation bring intuitiveness to this task, it is quickly limited because of graph cluttering, though this can be improved as in TROPOS (Giorgini et al., 2004, 2005) and CORAS (Vraalsen et al., 2005; Vraalsen et al., 2005) via tool-support. This later methodology is based on traditional design methodologies (risk analysis, UML) but mainly focuses on risk. Haley et al. (Haley et al., 2004) is also based on existing requirements engineering methods.

Textual notations are not limited by the graphical notation and can express complex notions in a human-understandable way, but it is only useful to the designer if guided as in the Matrix model as long descriptions are prone to misinterpretations and difficult to use. Furthermore, these notations can easily be incorporated into existing design methodologies, without the need to modify the design process.

A HOLISTIC TRUST DESIGN PHASE

As seen in the previous section, it is still quite difficult to find an existing methodology able help in designing e-services so as to make them globally trustworthy. It is particularly difficult in computing systems where the subjectivity induced by plain text English is in conflict with the operational nature of computing systems. This is made even more difficult in the context of emerging and changing technologies, where the level of abstraction is quite low and the focus is centered on the user.

We describe a holistic methodology to analyze trust during the design of a system that focuses on the system user and provides insight into the subjective nature of the system. This methodology can be seen as a holistic trust design phase that is added at the start of the system development. Many of the subjective facets of trust are captured by this methodology, as well as objective concepts that are more directly applicable to real-world applications. The holistic trust design phase is composed of five steps that are structured as shown in Figure 1. Each step of the holistic trust design is described in the following sections. We then illustrate the methodology on several examples.

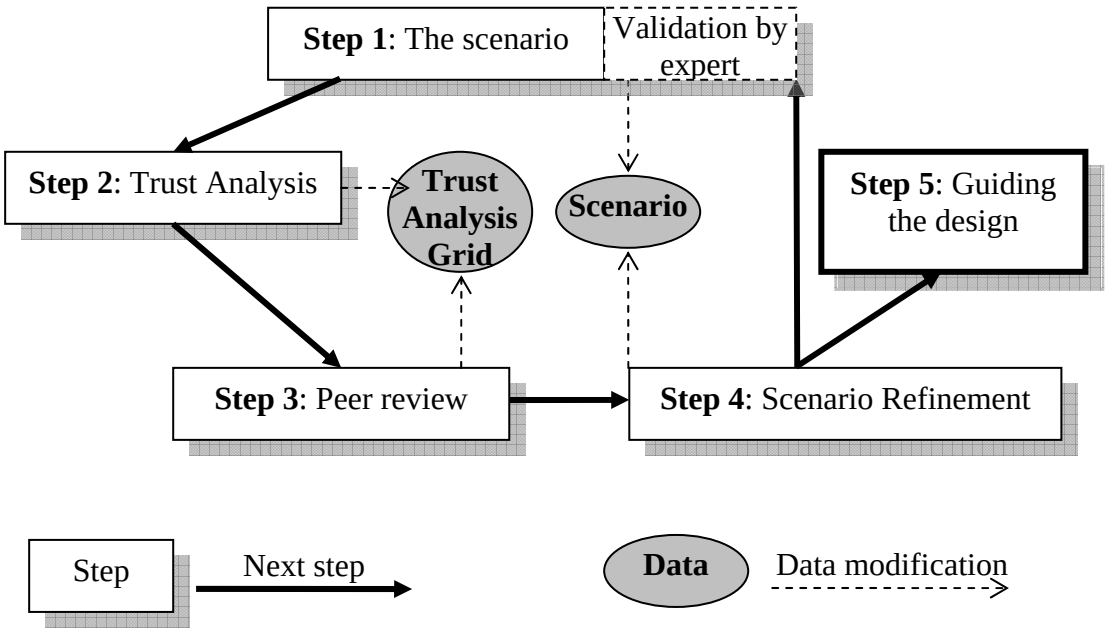

Figure 1: The holistic trust design

Step 1: The Scenario

Because of the human-centric nature of trust and modern e-services, it is critically important that trust is explored from the user's perspective, rather than in terms of abstract concepts or security features, so as to appreciate the impact of particular trust issues on the users of the system. The holistic trust design phase reflects this imperative by working on *scenarios*. The scenarios form the foundations of the methodology and their development and analysis provide a valuable holistic view of trust that can guide the design of the electronic system.

A *scenario* is a short, fictional narrative, set in the near future that describes people's daily lives, concentrating on their use of e-services under examination. The scenarios

are user-focused and usually avoid descriptions of how the technology works unless such descriptions clarify the users' interactions with the system.

It is important that the scenarios reflect the way in which people would use the e-services to support them in their daily lives, in order to fit the technology to the task rather than the opposite. It is critical that the scenarios are validated by subject matter experts, so that they plausibly depict people and processes within the application domain. This validation should be done, if possible, by a person external to the trust analysis and the system design, so that her opinion is not biased towards the technical environment proposed.

When writing scenarios, there is a trade-off between length and accuracy, but for the purpose of system development scenarios should focus on a specific set of features provided by the systems. The writing of scenarios is thus eased but this may limit the scope of the results, as a longer scenario can introduce interactions between elements that would be independent in smaller scenarios. Scenarios are living documents that will evolve during the process of trust analysis to meet the needs of the users and of the system designer. Their iterative development provides insight into the system and enables system designers to explore the various dimensions of their system.

Step 2: The Trust Analysis

The second step and foundation of the holistic trust design phase involves the *Trust Analysis Grid*. A sketch of a Trust Analysis Grid (TAG) is given in Table 1. The rows of the grid correspond to vignettes in the scenario. The columns of the grid correspond to categories of trust issues that will be checked against the vignettes.

	Trust Issue Categories										
	Data		**System**					**Subjective**			
Vignette in the scenario	**Source vs. Interpretation**	**Accuracy**	**Audit Trail**	**Authorization**	**Identification**	**Availability**	**Reliability**	**Personal Responsibility**	**Reasoning**	**Usability**	**Harm**
First vignette	X			XX			Y		YYY		physical

Table 1: Trust Analysis Grid

These trust issue categories are also grouped into *trust issues group* for the sake of convenience (see next sections). This choice of rows and columns is particularly suited to the study of scenarios as it enables the reviewer to follow the flow of narration vertically.

Vignettes

Since scenarios are written in a narrative style, only certain sentences and pieces of sentences are of interest for analyzing the trust issues. A vignette corresponds to one or several pieces of one or several sentences of a scenario and constitutes a cohesive group with regards to the trust issues. The pieces of sentence of interest to the trust analysis are indicated by formatting them in *italic*, the rest being pieces of the

sentence that do not concern the trust analysis but are displayed to ease the reading of the TAG. The various vignettes are examined in the order where they appear in the scenario.

Trust Issue Categories

Trust issue categories correspond to a specific property that represents one of the different facets of trust. They complement each other and are denoted by column labels in the TAG. These categories have been determined by successive analysis of various scenarios developed and a study of the state of the art on trust. We assume that the generalizations that we derived from the trust analysis are plausible because they have been derived from the user's interaction with the system represented in the plausible scenarios.

Each trust issue category denotes a facet of the notion of trust that is directly observed in a vignette, rather than being the consequence of such an observation. For example, the category *Source vs. Interpretation* generally may follow from the category *Reasoning*, though this latter observation is not directly observable in the scenario at that point. The eleven trust issue categories are:

- *Source vs. Interpretation*

An interpretation is data that has been obtained after the processing of other data (the source). The interpretation is generally less trusted than the source data itself.

- *Accuracy*

The level of detail of an information determines how precisely trust can be evaluated in the system. The higher the accuracy, the more confident users will be that they can trust this particular part of the system.

- *Audit trails*

An audit trail lists all the actions performed, who performed them or gave permission to perform them, and the events occurring in the system. This information should not be modifiable, or at least a modification should be detected and recorded along with the previous version.

- *Authorization*

Any agent accessing a piece of information or requesting a service must have the permission from the system to do so, which in turn may require that the user has authorized it (or not denied it).

- *Identification*

Identity is important to differentiate the participants and communicate with one of them. On the other hand, this identity may be limited (e.g. pseudonym) in certain contexts in order to provide privacy.

- *Reliability*

This property indicates that a service operates according to its specification. Similarly, the property can refer to the integrity of the data produced by the service.

- *Availability*

Availability corresponds to the temporal constraints on a service that ensure that the flow of action in the system is not stopped for a period of time longer than expected.

- *Personal Responsibility*

The system cannot check everything, but some things people do will affect the system’s trust. A person must remain responsible for the actions she performs, since they are not mediated by a trusted system. The property of accountability is important to put a significant level of trust in the system.

- *Reasoning*

Each participant manipulates the data to process it, in order to make decisions or answer a request. This process can weaken the trust another participant has in the system if this reasoning does not appear correct.

- *Usability*

This aspect of trust encompasses various elements, like the intrusiveness of the mechanisms used to interact with the user, or its usefulness. It is a crucial element of trust in computer systems as they can greatly impede the user. If a system is hard to use correctly, it may then be used incorrectly, and this will in turn reduce the trust.

- *Harm*

At the heart of trust is the notion of avoiding harm, since trust is a belief based on uncertain and approximate knowledge. It encompasses situations like loss of privacy (in the sense that personal data has been accessed against the will of its owner), breach of confidentiality, loss of financial assets, physical or emotional damage, and more generally risk.

Trust Issues Groups

The trust issue categories are grouped together into groups that correspond to properties at a higher level of abstraction. They are only used to organize the trust issue categories according to their abstract similarities. The three trust issue groups are:

- *Subjective categories*

Trust issue categories: *Personal Responsibility, Reasoning, Usability, Harm*
Trust is inherently subjective in that it reflects the point of view of the trustor. The subjective categories involve the agent's internal state and knowledge and express its beliefs. They also provide part of the context that is used to interpret trust relationships.

- *System categories*

Trust issue categories: *Audit Trail, Authorisation, Identification, Availability, Reliability*
These categories relate to the underlying components and services of the computing system used in the scenario. This system may involve a physical device, a computer program, or a more general socio-economic system.

- *Data categories*

Trust issue categories: *Source vs. Interpretation, Accuracy*
These two categories describe the properties of the data from the point of view of trust.

Grid Cells Values

The TAG is populated with values that can be of various forms, each providing slightly different means to represent the trust issues. The grid cells of the TAG can contain:

- An **X** mark

X indicates that this particular trust issue applies in its general stance in this vignette; the marks **XX** and **XXX** indicate values that are *more*, or respectively *much more*, important as those marked with an **X** on the same row; on the other hand, **X** cell values are not comparable between different rows.
If in a given row with four filled cells, one needs to relate two of them in terms of importance (for example **X** and **XX**) and also relate the two others, but independently

from the first two, then one can use different letters **X** and **Y** for the two pairs of values. The second one could be for example **Y** and **YYY**.

- The name of a more precise issue

It is sometimes necessary to indicate more precisely which aspect of the trust issue category is involved in a given vignette. This is done by putting a word as a cell value. For example, the trust issue category *Harm* can be refined into **physical** or **financial**.

- A signed number

A natural number represents the scale of the trust issue for a given vignette, 1 being the least important (but still present, as 0 is not used and instead the cell is empty) and increasing values indicating more important occurrence of the trust issue. The number is preceded with a negative sign (-) to represent the fact the contribution of this vignette to the trust issue is negative, i.e. it is an issue. On the contrary a positive sign (+) is used to denote the fact that the vignette addresses the particular trust issue. Note that values corresponding to the trust issue category *Harm* are always negative.

Colors are also used to represent our judgment about the trust issues, as they emphasize that these judgments are subjective. Two colors, a light and a dark, are used to represent visually the convention expressed by the number sign. It can be used in conjunction with the number value, or alone to give more visual information.

Step 3: Peer Review

In the third step of the holistic trust design phase, the initial examination of trust issues in step 2 undergoes peer review and cross-checking. Peer review supports the extraction of trust issues from the perspective of another potential user, who may have a different view on trust issues. It may be thus the occasion to discover some missing trust issues by complementing the reviewer's point of view.

In practice, the peer review is a very useful exercise as it forces the reviewers to explain and clarify their trust analysis. The peer review is typically done during a meeting where the reviewers go through their TAGs and compare them. Since trust is a subjective matter, they may argue on whether or not a particular trust issue arises at one point of the scenario. This disagreement may mean that a choice between contradicting requirements must be made by the system designer.

Disagreement occurring during the peer review may also be the consequence of trust analyses made from the point of view of users of the system who have a different roles, for example an end user and a system administrator. The trust analyses are not generally compatible due to contradictory requirements occurring between the roles, but the peer review ensures that the overall approach to analyzing the e-services is consistent.

Step 4: Scenarios Refinement

In the fourth step, scenarios are refined by adding new text and vignettes, or removing existing ones to address comments that have been made during the peer review. The purpose of the scenarios is to provide a framework which illustrates possible applications of the system, and to extract the most relevant trust issues. It is important that the scenarios reflect the trust concerns of all the stakeholders involved, and it

should be updated to represent different priorities. However, these concerns evolve as the trust analysis progresses and makes explicit the various trust issues.

When scenarios are updated, the sequence of steps is then executed again. The updated scenarios are validated by the domain experts who first validated it (step 1), another trust analysis is made (step 2), a peer review is organized (step 3) and the scenarios are possibly refined another time (step 4). This sequence is iterated until reviewers and system designers believe that the scenarios cover all the functionalities of the system and the trust analyses depict in a satisfying fashion the understanding of the system.

Step 5: Guiding the Design of the System

The four previous steps provided some insight into the trust issues underpinning the systems and are a means to explore the possible solutions provided by the system. In that sense, it follows the traditional design phase in software development based on use-cases. The last and final step of the holistic trust design phase consists in using the TAG to draw some guidelines in order to make design decisions.

Identifying Significant Areas

A simple visual examination of the TAG can give the system designer an overview of where the areas that are significant regarding trust are in the scenarios. Because of its visual nature and the fact that its vertical dimension corresponds to the sequential flow of a scenario, the TAG can be viewed as a map of the trust issues in the system under examination. The various *areas* of this map can give us some guidance on how to best design the system.

Firstly, we can decompose the TAG into three areas corresponding to the three groups of trust issue categories Subjective, System and Data. This abstract typology of trust indicates the kind of expertise that is required for designing the system. A Subjective-group system may require a system designer with knowledge of social science and/or the law, and human-computer interface. A System-group system corresponds to a system where the infrastructure plays a central role and where a technical experts in e-services may best practice his abilities. A Data-group system may need to be designed by an expert in data management and processing and/or privacy and data usage.

Secondly, we can also examine each column of the TAG individually. The columns that are full indicate that the corresponding trust issue is predominating in the system. This means that the system components proposed to solve this trust issue category in the design are given special attention and that enough resources are devoted to them. Ideally, the trust issue categories which have the most cell values in the TAG should require an additional verification pass following the system design. This verification should be made in reverse order, so that the most full trust issue category is verified last, and should carefully check that these concerns are mitigated.

Thirdly, a row or a sequence of rows where a lot of cell values are present probably indicates a crucial point in the scenario. This corresponds to a part of the system that is critical regarding trust and where additional attention must be paid. Another sub-scenario may be created to describe in more precise terms how the user interacts with

the system and the system behavior, and then a new trust analysis can be run. Following the system design, this point in the scenario must be verified thoroughly.

Matching Technologies against Scenarios

Rather than using the previous guidelines, one can try to analyze the TAG in a more systematic way to draw some more precise conclusions. Though it is not easy because of the subjective nature of the trust issues that are represented, it can still shed an interesting light on the design issues and in particular its technical feasibility. As the purpose of our approach is to help in the design of e-services, any means to understand how best to do this is beneficial to the e-service designer.

In order to introduce the technological elements into the holistic trust design, a TAG of the various common technologies and techniques used in modern e-services is used. An example of such a technology TAG is presented in Table 2. We then have two TAGs, one corresponding to the scenario and the other one to the technologies. The suitability of a particular technology at a given point (sequence of vignettes) in the scenario is given in terms of how close its pattern (a row of eleven cell values) matches the area corresponding to this point in the TAG of the scenario.

This pattern matching technique differs from the previous heuristic method in that it relates the informal analyses of scenarios and the technologies, and provides a point of anchorage for a more formal approach. As scenarios and its TAG are refined through the iterative sequence of steps of the holistic trust design, the technology TAG is completed to match the system's technological needs.

	Trust Issue Categories										
	Data		**System**					**Subjective**			
Technology	**Source vs. Interpretation**	**Accuracy**	**Audit Trail**	**Authorization**	**Identification**	**Availability**	**Reliability**	**Personal Responsibility**	**Reasoning**	**Usability**	**Harm**
Wireless Network			X	X	X	X	X	X		X	X
Grid Computing				X	X	X	X		X	X	X
Peer-to-Peer Network			X	X	X	X		X			X
Sensors	X	X		X		X		X		X	
Data Records	X		X		X	X		X	X		X
Network Traffic	X	X					X			X	
Audio and Video Data	X	X		X			X	X		X	
Speech Data	X	X					X		X	X	
Pads				X				X	X	X	
Location and Context		X							X		
HUDs										X	X
Personal Agents	X	X	X			X		X	X	X	
Service Agents		X	X		X	X	X		X		

Encryption							X				
Digital Signatures				X				X			
Authorization Mechanism			X		X	X	X	X			
Authentication		X		X	X		X			X	
Time Limited Leases			X	X	X	X	X	X		X	
Domain-based Security				X	X					X	

Table 2: Technology TAG

EXAMPLES

We describe here three examples of concrete applications of the holistic trust design phase that will help to understand it and demonstrate its usefulness. For the sake of brevity, some stories have been truncated and unnecessary elements were removed. The trust analysis is presented in the form of the trust analysis grid, using various cell value formats for the different scenarios. Step 3 of the holistic trust design phase is not discussed for the sake of conciseness.

E-healthcare

We present a scenario in the context of e-healthcare. In this futuristic but realistic scenario, various police and health workers collaborate on a crash scene. The scenario actors are performing their work and duties using computer devices and e-services to improve the patient treatment.

The Scenario

Neil is driving to work, when, suddenly, brake lights flare and Neil is jolted alert. There seems to be a wall of slowing cars and smoke is pouring from the wheels of the car in front. Neil's car was too close to avoid a collision. As the motorway grinds to a halt, it appears that three cars have crashed. Other motorists have managed to avoid the initial accident on both sides, but some have had minor collisions.

The emergency services already know much of the situation. As soon as the cars' airbags were triggered by the crash, the cars transmitted a distress call, including their location (given by the navigation systems) and the number of occupants (detected by simple pressure sensors in the seats). The first car's phone was too badly damaged to transmit its call, but for 999 calls it was able to piggy-back on the phone of the second car using short-range networking.

The emergency control room dispatches a small number of police, fire and ambulance vehicles immediately. The incoming calls from other motorists, and images from a traffic camera on a nearby bridge, seem to confirm the seriousness of the accident, and further vehicles are dispatched. The controller also sends an incident support vehicle to assist with clear-up. The information known so far is shared between all of the vehicles en-route. Information on traffic flow and speed is also shared between the vehicles to enable them to avoid blocked or slow routes. The dispatch and arrival of the vehicles is logged automatically to provide statistics on response times.

The traffic police are first on scene, and begin making the area safe as best they can. The video feed from their speed camera is available to the control room, but at low bandwidth. A still image is shared with the vehicles en-route though. The police confirm the number of vehicles involved, and the number of casualties. They quickly take a few evidential photos of the scene, and begin basic first responder treatment. These photos are shared with the en-route vehicles and the hospitals.

Neil is awake now. One of the policemen is trying to hand him over to the first paramedic on scene, but the policeman is told to keep holding Neil's head still while the paramedic triages the other casualties. He informs the ambulances and the control room of his findings by radio. The control room enters this information into the log for the incident, which is shared with the receiving hospitals.

The ambulances are arriving on scene and, after checking with the fire-fighters that the scene is secure, the paramedics continue treatment. They record their assessment and treatment onto normal paper report forms, but these are backed by smart clipboards that record and recognize the handwriting and ticked boxes, and can forward that information if required. Each patient is given an RFID tag, normally on a wristband, to enable the incident records to follow them around the system.

Neil seems to be relatively unhurt, but is immobilized with a cervical collar and board until spinal injury can be ruled out. The spinal board, sadly, doesn't have any sensors yet, but it does have an embedded RFID tag to identify which ambulance organization it belongs to. The fire-fighters are busy cutting up the car in front, and one of them is taking a few quick photos with his helmet camera for the incident support crew, to give them some idea of the scene.

Trust Analysis

We present below in Table 3 the TAG for the scenario described above. Due to the high number of actors and situations described, the TAG is quite big, but it could be split into smaller TAGs if necessary.

An interesting vignette in this scenario with regards to the technologies envisaged in pervasive systems is the *piggy-backing* of the first car's phone. In this situation, the *piggy-backing* implies that the phone call does not originate from the *source* and that some sort of *authorization* system enabled the first car's phone to use the resources of an unknown car around it. The *availability* of this technology is here paramount, but on the other hand could lead to abuses (*harm*) if the call serves other purposes than emergency. All other trust issue categories were not represented here.

Vignette in the scenario	**Trust Issue Categories**		
	Data	**System**	**Subjective**

	Source vs. Interpretation	**Accuracy**	**Audit Trail**	**Authorization**	**Identification**	**Availability**	**Reliability**	**Personal Responsibility**	**Reasoning**	**Usability**	**Harm**
three cars have *crashed*		X									X
motorists have *managed to avoid* the … accident	X		X					X			
minor collisions		X									
emergency services already know *much of* the situation		X						X			X
the cars *transmitted* a distress call			X			X	X				
piggy-back on the phone of the second car	X			X		X					X
emergency control room *dispatches…* vehicles			X						X		
the *seriousness* of the accident		X									
The information known so far is *shared*				X		X		X			
The dispatch and arrival of the vehicles is *logged*			X			X					
traffic police are *first* on scene			X					X			
making the area *safe*											X
The *video feed…* at *low bandwidth*	X	X	X			X					X
the *number* of vehicles involved, and the *number* of casualties		X	X				X				
evidential photos of the scene	X		X				X	X			
begin *basic first responder treatment*				X				X			
These photos are *shared*	X			X		X	X				
the paramedic *triages* the other casualties							X		X		
enters this information into the *log*			X			X					X
shared with the receiving hospitals	X			X		X	X				
the scene is *secure*											X
normal paper report forms	X									X	

smart clipboards that *record* and *recognize* the handwriting and ticked boxes, and can *forward* that information	X		X	X			X		X		
RFID tag					X					X	X
the incident records	X		X								
The spinal board, sadly, doesn't have any *sensors* …		X				X	X				
… but it does have an embedded *RFID tag*					X					X	X

Table 3: TAG of the E-Healthcare scenario

Guiding the Design

This long scenario was part of the bunch of scenarios that helped devise the design methodology and improve the trust analysis grid by going through several rounds of analysis. Though realistic and validated by a first aider, the system was not designed and developed further. The scenario is here to illustrate the first steps of the holistic trust design phase, but design and development are not part of the requirements.

The scenario and TAG enable to look at the system and services from two different viewpoints, thus separating functionalities from properties. For example the scenario made use of technologies not yet fully implemented (pervasive network, RFID), but providing the right functionalities and the TAG lead to hard requirements on these technologies (availability, harm).

Theme Park

This scenario is set in the context of a pervasive Theme Park, named Vaughn Park, which is fully equipped with pervasive computers that provide e-services to the customers. Park tickets are embedded with location technology (e.g. wi-fi, RFID that enable to provide context-dependent services. The focus of this scenario is a virtual queuing system where users can queue for ride and play a treasure hunt game that will guide them to the ride.

This scenario contains an hypothetical part (between brackets) that illustrate the expressive power of the holistic trust design, where various design branch can be explored by adding optional features and situations.

The Scenario

Janet and John are having a great time at Vaughn Park, but now that they have been on all the rides they wanted to, except for Hubris which has a long queue, they are beginning to get a little bored. They and their parents have joined Hubris' queue, but there is an estimated wait of over an hour until they'll be able to ride. Their parents suggest that they try one of the pervasive games the park offers.

The information kiosk can tell that they're waiting for Hubris, and it also knows that Janet and John have been on many of the rides that are likely to interest

them, so it thinks that the Treasure Hunt game is a good candidate for them. Indeed it is, so they choose to play the game.

The first clue is a simple one: "Can you find a big squirrel?" (If they weren't old enough to be reading yet, they could be given picture-only clues, but only if their parents played along with them.) Janet remembers that there is a squirrel on one of the Merry-go-rounds in the green area.

When they find the Merry-go-round they go up to an information kiosk. The kiosk knows they are playing Treasure Hunt, and that they are looking for a big squirrel. The one on the Merry-go-round isn't the one it had in mind, so it displays a message saying "Good try! But this one isn't big enough, can you find a bigger one? It's quite close!"

John notices a topiary cat on the other side of the Merry-go-round, and wonders whether there might be some more topiary nearby.

What a surprise! There is the squirrel, sculpted in the hedging. And neither of them had noticed it at all when they were on the Merry-go-round! The nearby kiosk congratulates them warmly, and asks them to find a big cleaning implement. They don't really know what to look for, so they don't move. The kiosk gives them a bigger clue: "who might use a cleaning implement, but not necessarily for cleaning?" That's it! Off they go, to the haunted house, which has a witch!

After successfully solving several more clues, the final clue leads them to Hubris, where they find their parents waiting. Their ride is great, and they go home afterwards talking very fast with a lot of excitement about the great day out.

Trust Analysis

Before we give the TAG in Table 4 below, we detail the analysis of one vignette of the scenario. This vignette is when the children do not understand what they should be looking for, and the kiosk gives them a clue. The piece of sentence of interest here is the *bigger clue*. It is a *usability* feature, as it will make the Treasure Hunt game more usable to the children. It also requires the system to *reason* about the situation, e.g. detect that the children are waiting for a clue because they are not standing in front of the kiosk. The input of this process is the activity of the children, which is their *personal responsibility*, while the output is a clue, which should *reliably* help the children. The other categories do not apply to this piece of sentence of interest.

	Trust Issue Categories										
	Data		**System**					**Subjective**			
Vignette in the scenario	**Source vs. Interpretation**	**Accuracy**	**Audit Trail**	**Authorization**	**Identification**	**Availability**	**Reliability**	**Personal Responsibility**	**Reasoning**	**Usability**	**Harm**

estimated wait of over an hour		tim e									
The information kiosk *can tell…*									X		
... that *they* are waiting					X	X		X			
it also *knows* that Janet and John have been			X						XX		
rides that are *likely* to interest them		XX						X	XX		
So the system *thinks* that			X						XX		
they *choose* to play the game			X						X		
If they were not old enough to be reading yet, they could be given picture-only clues, but only if their parents played along with them								XX	X	XX	X
The kiosk *knows* they are playing Treasure Hunt					X			X	X		
the one it had *in mind* ...				X					X		
... so it *displays* a message								X		X	X
The nearby kiosk *congratulates them warmly*					X			X	X	X	
They *don't really know* what to look for								X			X
The kiosk gives them *a bigger clue*							X	X	X	X	
the final clue leads them to Hubris				X				X	X	X	

Table 4: TAG of the Theme Park scenario

Guiding the Design

We first notice that the TAG is mostly filled with trust issues from the group *Subjective*. This is explained by the fact that the Pervasive Theme Park is a closed environment and this greatly simplifies the security requirements. Furthermore, the services provided are not data-intensive. This indicates that the application described in the scenario is a quite subtle application, what corresponds to intuition that it is user-friendly, and that the emphasis should be put on the perception of the system by the user during the design.

From the point of view of individual trust issue categories, *Personal Responsibility* and *Reasoning* are the ones that are most filled. The first category corresponds to the fact that the user has the total freedom to walk around the Pervasive Theme Park during the game and she is responsible for her actions when looking for the clues,

while the system is not interacting with her. The second category underlines the fact that the game corresponds to a hunt and must be adapted to the way the user performs it. To make the application more trustworthy, the user expects the system to act with her in a way consistent with the status of her hunt. The system designer shall include inference and pro-active capabilities of good quality.

Finally, no row distinguishes itself from the others, notably due to the overall importance and continuous presence of trust issues from the *Subjective* group. This may be explained by the fact that the application is focused on the user whose mobility avoids concentrating the system capabilities into one particular part of the scenario.

E-learning

This last example describes uses of an e-learning application named ShowNTalk. This commercial application allows pupils to create and show presentations on PDAs and is aimed at improving the reading, writing, presentational and IT skills of primary school pupils. The scenario served as a basis for the implementation of its first prototype, while the TAG helped in identifying the trust issues and improve the application.

The Scenario

Liz, a primary school teacher, has decided to assign her class of pupils an assignment to do on their PDA's. The children will prepare a multimedia presentation on the different types of cloud formations, working individually on their PDA's.

After creating an introductory slide about clouds on his PDA Sam notices his geography book has a chapter on cloud formations. Sam uses the PDA to take a picture of a page in a book which has a picture of a cumulus cloud and some text describing it. He uses the voice annotation feature to record himself reading aloud the text on the page. He later finds a page in the same book about stratus clouds but there's no picture so he uses the stylus on the PDA to input the text from the book and adds a voice recording of the text. When looking outside Sam spots a cirrus cloud in the sky and goes outside to use the camera function to take a picture of the cloud and add some text using the stylus.

Later in the day the pupils present their work to their teacher using the slideshow option in the software. The software automatically scrolls through the slides, showing the pictures and text whilst playing any audio files present with each. Liz asks each pupil to upload their work to the server.

After her pupils have gone home Liz reviews some of the presentations starting with Sam's. She checks Sam's previous work on the server and notices a comment saying that his reading needed some work. From his latest presentation she notices that his reading of the text on many of the slides is much better so she adds a comment to the presentation stored on the server.

Trust Analysis

The following TAG in Table 5 uses the color yellow (light) for issues that need to be addressed and green (dark) for those that are addressed.

Vignette in the scenario	Trust Issue Categories										
	Data		System					Subjective			
	Source vs. Interpretation	Accuracy	Audit Trail	Authorization	Identification	Availability	Reliability	Personal Responsibility	Reasoning	Usability	Harm
Liz gives assignment to pupils											
Sam takes picture of page in book											
Sam records himself reading text from book											
Sam copies text from book into PDA using stylus											
Sam takes pictures of cirrus cloud											
Sam adds text to slide about cirrus cloud											
Sam uses slideshow options											
Sam uploads work to server											
Liz reviews some presentations											
Liz checks Sam's previous work on server											
Liz notices a comment											
Liz adds a comment											

Table 5: TAG of the E-Learning scenario

Guiding the Design

The trust analysis reveals various strengths and weaknesses of the system with regard to trust. On the positive side, the trust issue categories *Availability*, *Reliability* and *Usability* are well addressed. This shows that the application depicted in the scenario is user-friendly, as these trust issue categories are among the ones that are perceived directly by the user.

The TAG shows that the ShowNTalk application has several limitations in the areas of *Identification* and *Authorization*, two security properties. For example, there is no way to record which pupil is using a given PDA, or whether the pupil is entitled to synchronize his work with the school's server. The trust analysis shows clearly that addressing these two trust issue categories should reduce greatly the number of trust

issues in the system (at least from the point of view of this scenario, which is representative of the system use). This analysis result led the development team to introduce a login system on the PDA and an access control list on the server. The prototype tests showed that this design improved the trustworthiness of the prototype and that this was overall a good design decision as it was focused on a specific part of the system and used the

The introduction of this security component in the system affects the *Usability* issue, which need to be addressed in the next version of the application, notably by improving the GUI and considering HCI aspects of the application. This is a more traditional aspect of the design that can now be considered separately from the holistic trust design.

FUTURE TRENDS

As trust and e-services get more pervasive in computing systems, reliance on them will increase too and understanding the various aspects composing trust will be as important as mastering how these aspects are composed. As each aspect of trust gets more and more diverse and complex (Josang et al., 2005), design will be particularly important in order to ensure the trustworthiness of whole systems. This trend should follow the general trend in the software industry to try to address significant issues as early as possible during the development process.

The subjective side of trust encompasses at the moment the facets of trust that are understood the least, for example usability, see Bottoni et al. (Bottoni et al., 2000) for a rare example, and the law. These topics will get more attention as trust is better understood by researchers and developers, and this trend will increase as more application will use trust and users will have to consider this issue more often.

While independent aspects of trust are better understood (dynamics, risk, etc.), the study of their combination will become critical. Standardization efforts will help in introducing some coherence in the field, as well as activities proposing to compare the various propositions like the Agent Reputation and Trust Testbed (ART Testbed, 2006).

Security is a growing concern nowadays, as vulnerabilities become more difficult to find and threats cause increasing damages. In many ways, trust is tied to security, though this relationship is not well understood yet. Anderson (Anderson, 2001) highlights the need for security usability, while the 14th International Workshop on Security Protocols ("Fourteenth International Workshop on Security Protocols", 2006) propose to “put the human back in the protocol" because he is “directly aware of the security requirement, or has access to the correct system state at the outer level of abstraction.” This later aspect is reminiscent of the holistic nature of trust.

CONCLUSION

If modern e-services are to be successful, compelling applications will not be enough to make it enter people's daily life. More efforts are needed to both find a suitable way to implement it and to make it trustworthy. Trust is a key notion in these systems. It supports both a better understanding of the system by the user and a better representation of the user's needs and concerns, since it is a concept inspired by a human notion.

Much in the recent years proposed ways to address part of the trust issue, but few tried to tackle the problem from a holistic point of view, where the various parts are put in the perspective of real systems. The holistic nature of trust is difficult to capture, as well as its subjective nature. Firstly the issue has to be tackled early during the development cycle so as to think about the system as a whole rather than concentrate on its parts. Secondly scenarios, or use-cases, are intuitive and powerful means to explore the system ideas by putting them in the context of its use. Lastly, guidelines should be provided on the various facets of trust issues, rather than imposing a particular definition of trust. An important point here is that they should guide the system designers at an adequate level of abstraction.

Many methodologies exist, but most do not address these requirements as they tackle specific aspects of trust or are difficult to apply to the design of concrete e-services. They put on the system designer the responsibility of bringing together the various parts that can be gathered from the different methodologies.

On the contrary, the previous requirements are addressed by the holistic trust design presented in this chapter, as it is a process preceding the system design, it centers on scenarios and the Trust Analysis Grid guides the designer during the analysis of the system. The holistic trust design phase is based on 5 steps and revolves around the scenarios and the TAG by iterating the steps until a stable design of the system is achieved from the point of view of trust. By focusing on domain expert-validated scenarios, the design phase stays close to the user concerns that are crucial to achieve trustworthiness. By providing an abstract decomposition of trust that covers a wide range of topics, it helps the system designer to identify the trust issues and to try to mitigate them in subsequent round of trust analysis.

REREFERENCES

Anderson, R. (2001). *A Guide to Building Dependable Distributed Systems*: John Wiley and Sons.

ART Testbed. (2006). Retrieved March 25, 2006 from http://www.lips.utexas.edu/art-testbed.

Bottoni, P., Costabile, M. F., Levialdi, S., Matera, M., & Mussio, P. (2000). Trusty Interaction in Visual Environments. In *Proceedings of the 6th ERCIM Workshop "USER INTERFACES FOR ALL" (UI4ALL)* (pp. 263-277). Florence, Italy.

Braendeland, G., & Stolen, K. (2004). Using Risk Analysis to Assess User Trust. In *Proceedings of the Second International iTrust Conference* (pp. 146-160). Oxford, UK: Springer LNCS 2995.

EBay. (2006). Retrieved March 25, 2006 from http://www.ebay.com.

Fourteenth International Workshop on Security Protocols. (2006). March 27-29 2006, Retrieved March 25, 2006 from http://homepages.feis.herts.ac.uk/~strrjh/SP2006

Giorgini, P., Massaci, F., Mylopoulos, J., & Zannone, N. (2004). Requirements Engineering Meets Trust Management. In *Proceedings of the Second International iTrust Conference* (pp. 176-190). Oxford, UK: Springer LNCS 2995.

Giorgini, P., Massaci, F., Mylopoulos, J., & Zannone, N. (2005). Modelling social and Individual Trust in Requirements Engineering Methodologies. In *Proceedings of the Third International iTrust Conference* (pp. 161-176). Paris, France: Springer LNCS 3477.

Grandison, T., & Sloman, M. (2000). A Survey of Trust in Internet Applications. *IEEE Communications Surveys and Tutorials, 3*(4), 2-16.

Grimsley, M., Meehan, A., & Tan, A. (2004). Managing Internet-Mediated Community Trust Relations. In *Proceedings of the Second International iTrust Conference* (pp. 277-290). Oxford, UK: Springer LNCS 2995.

Haley, C., Laney, R., Moffett, J., & Nuseibeh, B. (2004). Picking Battles: The Impact of Trust Assumptions on the Elaboration of Security Requirements. In *Proceedings of the Second International iTrust Conference* (pp. 347-354). Oxford. UK: Springer LNCS 2995.

Hanssens, N., Kulkarni, A., Tuchinda, R., & Horton, T. (2002). Building Agent-Based Intelligent Workspaces. In *Proceedings of the International Conference on Internet Computing, IC'2002* (Vol. 3, pp. 675-681). Las Vegas, NV, USA: CSREA Press.

Huandg, C., Ling, B. C., & Ponnekanti, S. (1999). Pervasive Computing – What is it Good for? In *Proceedings of the ACM International Workshop on Data Engineering for Wireless and Mobile Access* (pp. 84-91). Seattle, USA.

Ishaya, T., & Mundy, D. (2004). Trust Development and Management in Virtual Communities. In *Proceedings of the Second International iTrust Conference* (pp. 266–276). Oxford, UK: Springer LNCS 2995.

Jones, S., & Morris, P. (1999). TRUST-EC: Requirements for Trust and Confidence in Commerce: Report of the Workshop held in Luxembourg, April 8th-9th. *European Communities EUR Report, 2*.

Josang, A., Ismail, R., & Boyd, C. (2005). A Survey of Trust and Reputation Systems for Online Service Provision. *Decision Support Systems, (to appear).*

McKnight, D. H., & Chervany, N. L. (1996). *The Meanings of Trust.* (No. 96-04, Technical Report MISRC Working Paper Series): University of Minnesota, Management Information Systems Research Center.

Rindeback, C., & Gustavsson, R. (2004). Why Trust is Hard - Challenges in e-mediated Services. In *Proceedings of the 7th International Workshop on Trust in Agent Societies* (pp. 180-199). New York, USA: Springer LNCS 3577.

Slashdot. (2006). Retrieved March 25, 2006 from http://slashdot.org.

Staab, S., Bhargava, B., Lilien, L., Rosenthal, A., Winslett, M., Sloman, M., et al. (2004). The Pudding of Trust. *IEEE Intelligent Systems Journal*, *19*(5), 74-88.

Storey, N. (1996). *Safety-Critical Computer Systems*: Addison-Wesley.

Tan, J., Titkov, L., & Poslad, S. (2002). Securing Agent-Based e-Banking Services. In *Trust, Reputation, and Security: Theories and Practice (AAMAS 2002 International Workshop)* (pp. 148-162): Springer LNAI 2631

Tan, Y. H. (2003). A Trust Matrix Model for Electronic Commerce. In *Proceedings of the First International iTrust Conference* (pp. 33-45). Crete, Greece: Springer LNCS 2692.

Vraalsen, F., Braber, F., Lund, M., & Stolen, K. (2005). the CORAS Tool for Security Risk Analysis. In *Proceedings of the Third International iTrust Conference* (pp. 402-405). Paris, France: Springer LNCS 3477.

Vraalsen, F., Lund, M., Mahler, T., Parent, X., & Stolen, K. (2005). Specifying Legal Risk Scenarios Using the CORAS Threat Modelling Language. In *Proceedings of the Third International iTrust Conference* (pp. 45-60). Paris, France: Springer LNCS 3477.

Yan, Z., & Cofta, P. (2003). Methodology to Bridge Different Domains of Trust in Mobile Communications. In *Proceedings of the First International iTrust Conference* (pp. 211-224). Crete, Greece: Springer LNCS 2692.

Yu, E., & Cysneiros, L. (2002). Designing for Privacy in a Multi-Agent World. In *Trust, Reputation, and Security: Theories and Practice (AAMAS 2002 International Workshop)* (pp. 209-223): Springer LNAI 2631.

Yu, E., & Liu, L. (2001). Using the i* Strategic Actors Framework. In *Trust in Cyber-societies, Integrating the Human and Artificial Perspectives* (pp. 175-194): Springer LNAI 2246