A Trust Analysis Methodology for Pervasive Computing Systems

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

[1] University of Southampton, School of Electronics and Computer Science, Southampton SO17 1BJ, United Kingdom; {splp,mjb,mal}@ecs.soton.ac.uk
[2] QinetiQ ltd, WR14 3PS, Malvern, United Kingdom; cjmb@signal.qinetiq.com

Abstract. We present an analysis Trust Analysis Methodology for finding trust issues within pervasive computing systems. It is based on a systematic analysis of scenarios that describe the typical use of the pervasive system 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 in pervasive computing systems. The Trust Analysis Grid is then used to guide the design of the pervasive computing system.

Keywords: Trust Analysis Methodology, Pervasive Computing, Pervasive Scenario, Trust Analysis Grid.

1 Introduction

Pervasive computing [1] places the user at the centre of an environment populated by services accessible through devices embedded in physical objects. In contrast to the current mode of human-machine interaction, where all tasking is performed by a human being, pervasive computing seeks to soften the cognitive and physical burden for a human being within its environment, by enabling the system to reason about the user's situation. A plethora of new pervasive computing systems has been implemented in the recent years and their success is an important enabler for building electronic societies.

Agent systems have been shown to be a suitable paradigm for designing pervasive computing [2–4]. Collaborating agents, capable of describing, discovering and accessing services dynamically, can assume important information and service management roles in pervasive systems. To enable agents to reason about the capabilities of the services on offer, they may exploit the common understanding of terms, relations and services across communities promoted by the Semantic Web [5]. At the human-machine interface, agents may interact with humans to elicit and report information, and to provide an atmosphere of ambient intelligence [6].

The notion of trust has recently taken a central role in computing [7]. In systems like pervasive computing where the focus is on the user, technical features such as security are no longer sufficient to correctly design and implement

distributed systems. The subjective concept of trust not only enables users to better understand the paradigm of pervasive computing, but also opens new directions of research for solving existing problems, such as security [8], management of online communities [9] or e-Services lifecycle [10]. Despite much work tackling the issue of trust and some definitions of this concept spanning a wide range of domains [11], there is no clear and shared consensus on the definition of this concept, partly due to the fact that the definition depends on the context of use.

Agent technologies will be of particular importance in pervasive environments because they can embrace the subjective and uncertain aspects of trust. In some instances, agents within the system will effectively become extensions of the person, and must be given a capability to reason about trust in a way analogous to humans, in order to procure and offer pervasive information services seamlessly. To be able to determine how we can apply agent technologies that promote trust in pervasive computing, we need to fully understand the trust issues of significance within potential pervasive applications.

In this paper we describe a Trust Analysis Methodology for highlighting the key trust issues when designing pervasive systems. Section 2 outlines the various steps of the approach, from devising a scenario to analyzing it, and finally drawing conclusions regarding the system design. In Section 3, we give an example that illustrates the approach. We compare the Trust Analysis Methodology to related work in Section 4 and we finally conclude in Section 5.

2 The Trust Analysis Methodology

The goal of the Trust Analysis Methodology is to help in the design of the pervasive system by highlighting the trust issues inherent to the system. It is a guide rather than a model as it does not define rigorously exact terms but rather provides a means to discover trust issues. Nevertheless we can give the following definition of the term *trust* that is inspired from the literature and summarizes our whole approach.

Definition 1 (Trust). *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 **context** of trust has four elements:

- a group of three elements called *external context*:
  - the legal system (the law, legal entities, contractual agreements);
  - the social environment (non-legal entities, rules of communication and etiquette, culture, norms, social expectations);
  - the material environment (technologies, costs, limitations);
- and one element which constitutes the *internal context*:

- the moral state (intentions, prejudices and tendencies, beliefs other than trust, knowledge, past experiences and beliefs about other principals, emotions); a particular constituent of this element is the relative weight of the various constituents of the external context in the evaluation of the trust belief.

The **components** of the trust belief are:

– Data components
Source versus Interpretation, and Accuracy;
– System components
Audit trails, Authorisation, Identification, Personal Responsibility, Reliability and Availability;
– Subjective components
Reasoning, Usability and Harm.

Quantifiable constituents are combined to compute the value of a *trust metric*, which is the quantifiable part of the trust belief and is also dependent on the context.

Each part of this definition is explained in one of the five steps of the Trust Analysis Methodology. The methodology involves iteration over four steps, followed by a final fifth step. Each step is described below.

2.1 Step 1: The Scenario

Trust in distributed computing is often discussed in terms of abstract concepts or security features, and it is sometimes difficult to appreciate the impact of particular trust issues on the users of the system. Because of the human-centric nature of pervasive computing, it is critically important that trust is explored from the user's perspective. The Trust Analysis Methodology reflects this imperative by working on a scenario.

A *scenario* is a short, fictional narrative, set in the near future that describes people's daily lives, concentrating on their use of pervasive computing under examination. The scenario is user-focussed and usually avoids descriptions of how the technology works unless such descriptions clarify the users' interactions with the system. It is important that the scenario accurately reflects the way in which people would use the pervasive technologies to support them in their lives. Pervasive computing scenarios have already been designed in the context of ambient intelligence [6].

A scenario should not be too long and should focus on a specific set of features provided by the systems. This eases the writing of the scenario but may limit the scope of the methodology results, as a longer scenario made of smaller ones can introduce interactions between elements that were independent in the small scenarios. We consider the scenario as a living document that will evolve during the process of trust analysis to meet the needs of the users and of the system designer.

A key point in the acceptance of the pervasive computing paradigm is its applicability, which can be demonstrated by realistic scenarios. It is critically important that these scenarios are validated by subject matter experts, so that they plausibly depict people and processes within the application domain. A central principle for pervasive computing design is to fit the technology to the task, rather than the opposite. To this end, the scenario should be, if possible, validated by a person external to the trust analysis and the pervasive system design, so that her opinion is not biased towards the technical environment proposed.

To aid an understanding of trust issues, we have developed a number of plausible scenarios which contain use-cases that highlight the interactions between a user and her pervasive environment. We have also applied our methodology to scenarios borrowed from other pervasive systems. The scenarios form the foundations of our methodology and their development and analysis provide a valuable holistic view of trust that can guide the design of the pervasive system.

2.2 Step 2: The Trust Analysis

The second step and foundation of our Trust Analysis Methodology involves the *Trust Analysis Grid*. A sketch of a Trust Analysis Grid 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. Our previous work [12] used a different view on the Trust Analysis Grid, corresponding to a rotation by 90 degrees (or the rows and the columns are inverted) of the grid presented in this paper, which made the representation of the previous results in the form of a grid less intuitive because there is an unknown number of columns. The Trust Analysis Grid we use here is more suited to the study of a scenario as it enables the reviewer to follow the flow of narration.

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

Table 1. Trust Analysis Grid

Vignettes Since the scenario is 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 the scenario and constitutes a cohesive group with regard to the trust issues. In order to make the vignettes more readable, they can contain pieces of sentence that do not concern the trust analysis, for example when a sentence is incomplete. We will indicate the pieces of sentence of interest to the trust analysis by formatting them in *italic*. The various vignettes are examined in the Trust Analysis Grid in the order where they appear in the scenario.

Trust Issue Categories Following the trust analysis of several scenarios in a previous work [12], we classified the trust issues into eleven categories which are partitioned into three groups. Trust Issue Categories correspond to different facets of trust that complement each other and are denoted by labels that are defined in Table 2. We assume that the generalisations 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 scenario [12].

It is important to understand that each category denotes a kind of trust issue that is directly observed in a vignette, rather than being the consequence of such an observation. For example, the category *Source vs. Interpretation* generally follows from the category *Reasoning*, though this latter observation is not directly observable in the scenario at that point.

Trust Issues Groups The groups of Trust Issue Categories correspond to elements of trust at a higher level of abstraction. They are only used to organize the Trust Issue Categories according to their abstract similarities.

- **Subjective categories**
  Trust is inherently subjective in that it reflects the point of view of the trustor [14]. 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**
  These categories relate to the underlying components and services of the pervasive system used in the scenario. This system may involve a physical device, a computer program, or a more general socio-economic system.
- **Data categories**
  These two categories describe the properties of the data from the point of view of trust.

Values in the Grid Cells The first version of our methodology [12] simply checked each vignette against the eleven trust issue categories. This corresponded to putting an **X** mark in the grid cells to indicate that the corresponding trust issue occurred in the italic text in the vignette.

It was further found that this value format was not enough to describe accurately some trust issues. The grid cells of the Trust Analysis Grid can contain:

Data Group
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 is, the more confident the user will be that she can trust this particular part of the system.
System Group
Audit trails
An audit trail lists all the actions performed and the events occurring in the system. This information should not be modifiable, or at least a modification should be detected.
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 (this period may vary depending on the kind of actions).
Subjective Group
Personal Responsibility
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 pervasive systems as they can greatly impede the user. Little work has been done so far on this aspect of trust, as exemplified by Bottoni and al. [13].
Harm
This aspect goes hand in hand with trust, since trust is a belief, and it may be misleading and harm the user or the system. Loss of privacy, in the sense that personal data has been accessed against the will of its owner, or loss of financial assets are two examples.

Table 2. Trust Issue Categories and Groups

- an **X** mark to indicate that this particular trust issue applies in its general stance; 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* 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**); see Table 1 for an example;
- the name of a more precise issue; for example, the trust issue category *Harm* can be refined into **physical** or **financial**.

2.3 Step 3: Peer Review

In the third step, 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. This peer review may be the occasion to discover some missing trust issues and complement the reviewer's point of view.

In practice, the peer review is a very useful exercise as it forces the reviewers to explain their trust analysis, thus clarifying it. The peer review is typically done during a meeting where the reviewers go through their Trust Analysis Grids 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.

The peer review may also be the consequence of trust analysis made from the point of view of users of the system who have a different role. 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 pervasive system is consistent.

2.4 Step 4: Scenario Refinement

In the fourth step, the scenario is refined by adding new text and vignettes, or removing existing ones. The purpose of a scenario is to provide a framework which illustrates possible applications of the pervasive computing system, and to extract the most relevant trust issues. It is important that the scenario reflects 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.

The updated scenario is again validated by the domain experts who first validated it and another trust analysis is run by going back to step 2. This sequence, composed of the trust analysis (step 3) and the scenario refinement (step 4), is iterated until the reviewer and system designer believe that it covers

adequately, respectively, the trust requirements and the functionalities of the pervasive system.

2.5 Step 5: Guiding the Design of the Pervasive System

The four previous steps have already provided some insight into the trust issues underpinning the pervasive systems and are a means to virtually 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 our methodology consists in using the Trust Analysis Grid to draw some guidelines in order to help in the design of the pervasive systems under consideration. We describe below two possible approaches for the examination of the Trust Analysis Grid.

Identifying Significant Areas A simple visual examination of the Trust Analysis Grid can give the system designer an overview of where the significant areas are in the scenario. Because of its visual nature and the fact that its vertical dimension corresponds to the sequential flow of the scenario, the Trust Analysis Grid can, in a certain way, be considered as a map of the trust issues in the pervasive system under examination. The various *areas* of this map can give us some guidance on how to best design the system as we show below with three examples.

Firstly, we can decompose the Trust Analysis Grid into three areas corresponding to the three groups of Trust Issue Categories *Subjective*, *System* and *Data*. This general 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 pervasive computing may best practice his abilities. A *Data*-group system can be designed by an expert in data management and processing.

Secondly, we can also examine each column of the Trust Analysis Grid individually. The full columns indicate that the corresponding trust issue is predominating in the pervasive 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 few Trust Issue Categories which have the most values in the Trust Analysis Grid should correspond to an additional verification pass following the system design (in reverse order of how full the Trust Issue Categories are, so that the most full Trust Issue Category is verified last) that will check that these concerns are mitigated.

Thirdly, a row or a sequence of rows where a lot of 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 behaviour, 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 the Scenario Rather than use the previous informal guidelines, one can try to analyze the Trust Analysis Grid in a more systematic way to draw some more accurate conclusions. Though it is not easy because of the subjective nature of the trust issues that are represented in this grid, it can still shed an interesting light on it. As the purpose of our approach is to help in the design of the pervasive system, any means to understand how best to do this is beneficial to the system designer.

We first tried to understand how to introduce the technological elements into our approach. This was done by devising a Trust Analysis Grid of the various common technologies and techniques used in pervasive computing, see Table 3. We then have two Trust Analysis Grids, one corresponding to the scenario and the other 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 Trust Analysis Grid of the scenario.

This pattern matching technique differs from the previous heuristic method in that it relates the informal analyses of scenario and technologies, and provides a point of anchorage for a more formal approach. As more scenarios are analyzed against the Trust Analysis Grid in Table 3, this grid will be refined in order to better represent the trust issues of pervasive computing.

3 Illustration of the Methodology

3.1 Introduction

To illustrate the Trust Analysis Methodology, we describe here a scenario based on a Pervasive Theme Park. The scenario is presented in Section 3.2, then its Trust Analysis Grid is given in Section 3.3 and Section 3.4 finally outlines some comments on the system design.

This park is named *Vaughn Park* and provides a lot of fun rides for the whole family. It is fully equipped with pervasive computers, for example information kiosks with tactile screens can be very easily found in all parts of the park. The park is designed as a closed environment, meaning that no intruder can penetrate inside and that no customer or computer signal can get outside. Customers buy tickets equipped with location technology (e.g. wi-fi, RFID) at the entrance of the park and then indicate if they belong to a group, for example a family.

Vignette in the scenario	Trust Issue Categories										
	Data		System						Subjective		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
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			
Authorisation 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	

Legend of the column numbers

(1): Source vs. Interpretation	(2): Accuracy	(3): Audit Trail
(4): Authorisation	(5): Identification	(6): Availability
(7): Reliability	(8): Personal Responsibility	(9): Reasoning
(10): Usability	(11): Harm	

Table 3. Trust Analysis Grid of Technologies

3.2 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' virtual queue, but there is an estimated wait of over an hour until they will be able to ride. Their parents suggest that they try one of the pervasive games the park offers. While they play, the parents have a coffee at a café.

The information kiosk can tell that they are 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 the system thinks that the *Treasure Hunt* game is a good candidate for them. Indeed it is, because they choose to play the game.

The first clue is a simple one: "Can you find a big squirrel?" (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.) 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 is not 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.

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 do not really know what to look for, so they do not 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 nineteen to the dozen about their great day out.

3.3 Trust Analysis

Table 4 presents the Trust Analysis Grid resulting from the analysis of the scenario. The column numbers are the same as in Table 3. We illustrate in the following how the grid is filled on the example of the penultimate row. This table is the one obtained after steps 3 (peer review) and 4 (scenario refinement) of the approach, these steps not being shown here for the sake of simplicity.

This row corresponds to a vignette in the scenario where 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, and it 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. Hence the marks in columns (9) and (10). 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. Hence the marks in columns (7) and (8). The other categories do not apply to this piece of sentence of interest.

3.4 Guiding the Design

We first notice that the Trust Analysis Grid 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, what greatly simplifies the security requirements, and that the services provided are not data-intensive. This result 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 each Trust Issue Category individually, *Personal Responsibility* and *Reasoning* are those who are most filled. The first category

Vignette in the scenario	Trust Issue Categories										
	Data		System						Subjective		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
estimated wait of over an hour		time unit									
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. Trust Analysis Grid of the Scenario

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 will include in the system inference 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 prevents to concentrate the system capabilities into one particular part of the scenario.

4 Related Work

The requirement aspects from our Trust Analysis Methodology share some similarities with the approach taken by the TRUST-EC project [15]. This approach

lists and analyzes the common applications in e-Commerce. It then deduces from this analysis a list of requirements for trust and confidence in this domain. These requirements all have an equivalent Trust Analysis Category in our methodology, which adds the categories *Source vs. Interpretation*, *Reasoning* and *Usability*. This points to the fact that many of the trust methodologies tackle the problems from a technical point of view, rather than a human-centric one. Thus, many of the subjective facets of trust are evaded, as these objective concepts are more directly applicable to real-world applications.

The i^* framework [16] is 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. The i^* graphs relate elements that are very similar to the content of our Trust Analysis Grid and can be viewed as a scenario annotated with keywords that are interrelated. The i^* framework's expressive power and planar nature hinder a structured view of the trust issues in the system. Furthermore, the only example of requirements studied in detail for the i^* framework is privacy, while usability and security (which should be decomposed into more atomic properties as identification and authorisation) are only briefly mentioned.

Tan's trust matrix model [17] is a means to analyze trust-building services for electronic commerce. It proposes to represent a service in the form of a grid. The grid rows correspond to properties of the service grouped into three layers, each layer playing the same role as our Trust Issue Groups. Some of those properties correspond to our Trust Issue Categories (Information Integrity, Document Interpretation) while others correspond to a fine-grain decomposition of some Trust Issue Categories. The grid columns correspond a theoretical decomposition of the notion of trust into four reasons, themselves decomposed into two sources. The trust analysis in this framework is more suited to the examination of a particular service offered by a system, but it fails to capture the temporal dimension of a scenario and is thus not directly applicable to pervasive computing. It is also more precise in that it considers more trust issues, but those issues are more specific to the kind of services examined in this work.

Bændeland and Stølen's [18] propose to analyze user trust in a net-bank scenario. Their approach is closer to software engineering as they use the UML notation to model the system and reuses the CORAS risk analysis methodology. They define trust as a hierarchy of assets which must be protected from threats, vulnerabilities and incidents. Evaluating in detail the risks associated with the system under examination enables them to propose the solution to trust issues. Their analysis methodology is partitioned into five sub-processes, from establishing the context to treating the risks. Our methodology spans the first two sub-processes of their methodology. On the other hand, this approach does not explicit any definition of trust.

5 Conclusion

If pervasive computing is to be the next successful paradigm for computing, usability and 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 trusted. Agent systems are a relevant technology to implement this vision, because it captures many of the aspects of pervasive computing applications, such as mentalistic attitudes, social behaviour and users' mobility. Agents will ultimately behave as humans by proxy in autonomic pervasive computing.

Trust is a key notion in this paradigm. It supports both a better understanding of the system by the user and a better representation of the users' needs and concerns, since it is a human notion. We are investigating ways to create trusted pervasive systems and devised a Trust Analysis Methodology to help in the design and implementation of such a system.

Our approach is based on five steps. Firstly, pervasive computing scenarios are written to illustrate the use of the system and are validated by subject matter experts, ensuring a realistic representation of the system and of the trust concerns from the user. Secondly, the scenario is analyzed by filling in a Trust Analysis Grid that provides a means to discover the relevant trust issues. The trust analyses are then cross-checked between the various reviewers in step 3, and this leads to a possible refinement of the scenario in step 4. Finally, the fifth step consists in deducing from the Trust Analysis Grid a set of more or less formal guidelines for the system designer.

We are still applying the Trust Analysis Methodology to other pervasive computing scenarios in order to further extend, refine and stabilize it. In particular, we are considering extending the format of the cell values, including the possibility to give accurate measures (numbers), and add to the cell value a probability of occurrence of the particular trust issue. This extension of the cell value format could form the basis of a numerical analysis that could, for example, be bound to a risk analysis [18, 19].

As we wish to model formally the design of the pervasive system [20], we would also like to verify the Trust Analysis Grid against our formal models. The final models should integrate a model of the particular agent technologies used to implement the system and will enable to have a higher confidence that the trust analysis is correct than the one obtained with semi-formal methodologies.

6 Acknowledgments

This work is supported by the Next Wave Technologies and Markets programme of the United Kingdom's Department of Trade and Industry in the context of the T-SAS project (`http://www.trustedagents.co.uk`). The authors would like to thank Elisabeth Ball for her useful help.

References

1. Huang, A., Ling, B., Ponnekanti, S.: Pervasive Computing - What is it Good For? In: Proceedings of the ACM International Workshop on Data Engineering for Wireless and Mobile Access, Seatle, WA, USA (1999) 84–91
2. Falcone, R., Singh, M.P., Tan, Y.H.: Trust in Cyber-societies, Integrating the Human and Artificial Perspectives. LNAI 2246. Springer (2001)
3. Nixon, P., Terzis, S.: Trust Management (First International Conference iTrust. LNCS 2692. Springer (2003)
4. Dimitrakos, T., Martinelli, F.: Proceedings of the 1^{st} International Workshop on Formal Aspects in Security and Trust (FAST 2003). Istituto di Informatica e Telematica (2003)
5. James Hendler: Agents and the Semantic Web. IEEE Intelligent Systems **16** (2001) 30–37
6. K. Ducatel, M. Bogdanowicz, F. Scapolo, J. Leijten, and J-C. Burgelman: Scenarios for Ambient Intelligence in 2010. Technical report, Information Society Technologies, European Commission (2001)
7. IST Advisory Group: Trust, dependability, security and privacy for IST in FP6. Technical report, IST Advisory Group (2002) `ftp://ftp.cordis.lu/pub/ist/docs/istag_kk4402464encfull.pdf`.
8. Heather, J., Hill, D.: I'm Not Signing That! In Dimitrakos, T., Martinelli, F., eds.: Proceedings of the 1^{st} International Workshop on Formal Aspects in Security and Trust (FAST 2003), Pisa, Italy (2003) 71–81
9. Ishaya, T., Mundy, D.: Trust Development and Management in Virtual Communities. In Jensen, C., Poslad, S., Dimitrakos, T., eds.: Proceedings of the Second International Conference on Trust Management (iTrust 2003). LNCS 2995, Oxford, United Kingdom (2004) 266–276
10. Rindebäck, C., Gustavsson, R.: Why Trust is Hard - Challenges in e-mediated Services. In: Proceedings of the 7^{th} International Workshop on Trust in Agent Societies, New York, USA (2004)
11. McKnight, D.H., Chervany, N.L.: The Meanings of Trust. Technical Report 94–04, Carlson School of Management, University of Minnesota (1996) `http://misrc.umn.edu/wpaper/WorkingPapers/ 9604.pdf`.
12. Butler, M., Leuschel, M., Presti, S.L., Allsopp, D., Beautement, P., Booth, C., Cusack, M., Kirton, M.: Towards a Trust Analysis Framework for Pervasive Computing Scenarios. In: Proceedings of the 6^{th} International Workshop on Trust, Privacy, Deception and Fraud in Agent Systems, Melbourne, Australia (2003)
13. Bottoni, P., Costabile, M.F., Levialdi, S., Matera, M., Mussio, P.: Trusty Interaction in Visual Environments. In Emiliani, P.L., Stephanidis, C., eds.: Proceedings of the 6^{th} ERCIM Workshop "USER INTERFACES FOR ALL" (UI4ALL), Florence, Italy (2000) 263–277
14. Gambetta, D.: Can We Trust Trust? In Gambetta, D., ed.: Trust: Making and Breaking Cooperative Relations. Department of Sociology, University of Oxford (2000) 213–237 Electronic edition. `http://www.sociology.ox.ac.uk/papers/gambetta213-237.pdf`.
15. Jones, S., Morris, P.: TRUST-EC: Requirements for Trust and Confidence in E-Commerce: Report of the Workshop held in Luxembourg, April $8^{th} - 9^{th}$. Technical Report EUR 18749 EN, European Communities EUR Report (1999) Issue 2, `http://dsa-isis.jrc.it/TrustEC/D1.pdf`.

16. Yu, E., Cysneiros, L.M.: Designing for Privacy in a Multi-agent World. In Falcone, R., Barber, S., abd Munindar Singh, L.K., eds.: Trust, Reputation, and Security: Theories and Practice, Bologna, Italy (2002) 209–223
17. Tan, Y.H.: A Trust Matrix Model for Electronic Commerce. In Nixon, P., Terzis, S., eds.: Trust Management (First International Conference iTrust 2003), Crete, Greece (2003) 33–45
18. Bændeland, G., Stølen, K.: Using Risk Analysis to Assess User Trust - A Net-Bank Scenario. In Jensen, C., Poslad, S., Dimitrakos, T., eds.: Trust Management (Second International Conference iTrust 2004), Oxford, United Kingdom (2004) 146–160
19. Storey, N.: Safety-Critical Computer Systems. Addison-Wesley (1996)
20. Butler, M., Leuschel, M., Presti, S.L., Turner, P.: The Use of Formal Methods in the Analysis of Trust. In Jensen, C., Poslad, S., Dimitrakos, T., eds.: Trust Management (Second International Conference iTrust 2004), Oxford, United Kingdom (2004) 333–339