Towards a Formal Model of Trust in e-Commerce

Theo Dimitrakos

IT Department, CLRC Rutherford Appleton Laboratory, OX11 0QX, UK
t.dimitrakos@rl.ac.uk

Abstract. In this paper we provide a working definition of trust that is relevant to e-commerce, summarise current trends in formalising trust, and discuss plans for developing an integrated modelling framework. We do not attempt to define a general model of trust, but rather to analyse trust relationships of particular interest in e-commerce. In doing so, we illustrate what is missing from existing formalisations, indicate which aspects of existing models we find particularly useful, and suggest a form of integration that seems to facilitate a solution.

1 Introduction

Trust underlies almost every aspect of human interaction. However, trust may take different forms in different contexts, and there is clear evidence of fundamental differences between trust in commerce in the physical world and trust in e-commerce. In the physical world, we derive much of our notions of trust in commerce from the tangible nature of the entities in our environment. Our trust relies on personal contact between human agents negotiating, the tangibility of the objects and services under negotiation, the expense and difficulty of physical fraudulence, such as setting up a shop, and the existence of a clearly defined legal framework as a reference to our negotiations. Personal contact in virtual communities is limited, the legal framework is vague, and the objects and services under negotiation are less tangible. Thereby, the traditional notions of trust in commerce need to be rethought, while new definitions and properties of trust in e-commerce have to be developed.

The emerging virtual communities require richer models of trust, in order to distinguish between the different types of trust, and then accommodate the implications of this classification within the context of a specific service. A major shortcoming of current solutions is that they fail to incorporate in their decision making evidence or opinions collected by an agent (human or software) through the agent's own experience of the system, or via communication with other agents who cohabit the system. This, not only makes the evolution of e-commerce systems harder, but it also impedes the ability of implemented systems to adapt to changes in trust and to set up new relationships. In order to be able to handle trust dynamics, future solutions will have to incorporate methods to simulate learning, reasoning and analysing transaction and environmental risks with respect to the agents' view of the system they inhabit. The needs for flexibility and scalability can be better addressed by separating the trust management framework from the purpose of the application. (Similar opinions are expressed in [17] and [5].)

2 A Working Definition of Trust in e-Commerce

To date, there is little consensus in the literature on what trust is, although its importance has been recognised. On the other hand, as it is elaborated in [5], many researchers assume an (unprovided) definition of trust and use the term in a very specific way related to authentication and authorisation or to paying for purchases. In [3] we survey various attempts to provide some definition of trust that is suitable for e-commerce. Some aspects of these definitions are common, other are complementary. For example, [5] emphasises that trust is a belief in the competence of an entity that depends on a specified context. While [14] lay stress on that the entity that manifests trust (which we will call the *"trustor"*) is the human - not the system. They also emphasise that trust is in part subjective. A somewhat similar view is expressed in [10] where entities are distinguished into those who have free will, called *passionate*, and those who don't, called *rational.* According to [14, 10] trustors are always *passionate* entities. The definition of [7] focuses on another important aspect of trust: in commerce, *trust is relative to a business relationship.* One entity may trust another entity for one specific business and not in general. Even in the context of the same business relationship one may be trusted for one transaction but not for another. This diversity of the purpose of trust is also mentioned in [10] but not incorporated into a definition. Finally none of the above has emphasised the fact that trust, in commerce, is inherently measurable and it exists and evolves in time.[1]

We identify their common and complementary aspects and incorporate them into the following wider definition of trust in e-commerce.

Definition 1 *Trust of a party A in a party B for a service X is the measurable belief of A in that B behaves dependably for a critical period within a specified context.*

Remarks: *(1)* A *"party"* can be an individual entity or a collective of humans or software agents, or a system. Obviously, the trusting subject has to be an entity which can form a belief. Such entities may be either humans or adequately programmed software agents. *(2)* The term *"service"* is used in a deliberately broad sense to include e-payments and other transactions, recommendations, issuing certificates, guaranteeing for someone in another service, etc. *(3) "Dependability"* is deliberately understood broadly to include *security, safety, reliability, timeliness, maintainability*; *(4)* The *"crtical period"* may be in the past (history), the duration of the service (from now and until end of service), future (a scheduled or forecasted critical time slot), or always. *(4) "Context"* includes the business context, the relevant agreements, the service history, the technology infrastructure and the legislative, regulatory systems that may apply.

Notably our definition differs from [14] and [10] with respect to the trusting subjects. Intelligent agents who negotiate over the Web can be either humans or programs and in both cases they need to manifest trusting intentions and establish trusting relationships. Indeed, intelligent software agents are adaptive autonomous programs which feature the ability to acquire knowledge and take decisions of their own. For adaptive programs are able to improve their competence at dealing with their goals over time, and autonomous programs may operate in an independent fashion without explicit guidance, deciding on how to relate stimuli from the environment to actions leading towards the satisfaction of their goals. Software agents are therefore influenced by their environment and alter their behaviour through learning and exercise. Within open distributed systems, such programs may appear to behave passionately as they continuously interact with their environment and respond to stimuli that we cannot predict.

We also note that distrust, accounting to what extent we can ignore claims of another party about her own or a third party's trustworthiness and their proclaimed actions or commitments, can be modelled as a measurable belief in that a party behaves in a *non-dependable* manner for a critical period within a specified context. Distrust is useful in order to revoke previously agreed trust, obstruct the propagation of trust, communicate the information that a party is "blacklisted" for a class of potential business transactions.

2.1 Properties of Trust and Distrust

Trust exists in relation to some specific service and with respect to a specific business context. Thus, the particular characteristics of trust may differ from business to business. Nevertheless, there are some common delimiters which indicate the existence of general principles that govern trust in e-commerce.

Proposition 2 (G1) *[dis]trust is a measurable belief.*(G2) *[dis]trust is directed.*(G3) *[dis]trust is relativised to a business transaction.*(G4) *[dis]trust exists in time.*(G5) *[dis]trust evolves in time.*(G6) *[dis]trust is reflexive, yet trust in oneself is measurable.*(G7) *[dis]trust is not necessarily symmetric.*(G8) *[dis]trust between collectives does not necessarily distribute to [dis]trust between their members.*(G9) *distrust is different than "trust not to".*

[1] In [5] it is mentioned that there are "levels of trust" which change in time but measurability is not treated as an inherent characteristic of trust.

Remarks:

G1 states that agent A may trust agent B more than A trusts agent C for the same business. The metric is based on evidence, experience and perception. The measurement can be quantitative (e.g. as a probability) or relative (e.g. by means of a partial order). There are some interesting arguments, mainly of a philosophical nature, for and against each of these alternative metrics. In practice, either type of metric may be preferable depending on the deployed trust management scheme.

G2 and *G7* state that different parties with different roles in a transaction may have different views on trust in each other or in third parties. To a certain extent, trust is subjective.

G3 states that trust has to refer to a particular business transaction. Different laws may govern different trusting relationships for different business transactions.

G4,G5 reflect the fact that trust depends on a sequence of events. Assume A trusts B for a business relationship that lasts for a limited period. During a business transaction, the more A realises she can depend on B the more A trusts B. On the other hand, A's trust in B may decrease if B proves to be less dependable than A expected. For example, if B is less competent than A expected, if some reliable source discredits B, if A finds out that B trusts one of A's competitors, etc. At the end of the transaction A notes B's performance and uses this and any other relevant information to re-evaluate her trust in B before she enters into a future business relationship with B. The fact that A trusted B in the past does not in itself guarantee that A will trust B in the future.

G6 supports the ability of an agent to delegate or offer a task to another agent in order to improve efficiency or reduce risk.

G8 distinguishes trust in a collective from trust in its members. On the assumption that A trusts a group of contractors to deliver (as a group) in a collaborative project, one cannot conclude that A trusts each member of the team to deliver in the project. A potentially bad performance of a member of the group can be overshadowed by potentially excelling performance of another.

G9 distinguishes between distrust in an agent's expected behaviour and trust in a complementary behaviour. A may distrust B to fly a plane, but this is different from A trusting B not to fly a plane. [Dis]trust is related to the potential of B's behaviour contributing towards a state that is [non-]dependable in A's view of the world. Since the complement of a [non-]dependable state may include states that are themselves [non-]dependable, "trust not to" and "distrust to" are not necessarily equivalent.

Propagation of Trust. Neither trust nor distrust is necessarily transitive. On the assumption that A trusts B to buy on credit and B trusts C to buy on credit, one cannot conclude that A trusts C to buy on credit. A may have some control over B's resources but no control over C's resources. Or perhaps a third party guarantees B's credit but no third party that A trusts has the authority to guarantee for C's credit.

As we elaborate in the sequel, at least **unintentional** transitivity within a locus may be endorsed depending on the specific context. For this purpose, and in order to avoid referring to some specific business context, we distinguish three special *roles* that entities mediating in a trust relationship may play. These roles are: *guarantors*, *intermediates*, and *advisers*. An entity may play more than one mediating role in a business relationship. For example, depending on the criticality of a transaction, certification authorities may play the roles of a guarantor and/or of an adviser.

Guarantor is a party offering a formal promise or assurance that all obligations of the parties she guarantees for will be fulfilled in the context of a transaction and will be of a specified quality and durability. Usually, the guarantors assist the establishment or facilitate the increase of trust in a specific transaction by underwriting (a part of) the risk associated with the transaction. Typical examples include credit card companies, companies issuing e-wallets and companies managing Internet payments. All parties in a business transaction have to exhibit sufficiently strong trust in each other or a guarantor in order to for the transaction to happen.

Intermediate is a party that intervenes between other parties in a business transaction and mediates so that they establish a business relationship with or without their knowledge. We distinguish the following four types of intermediate, depending on the information they provide with respect to the stakeholders in the services the mediate in. An intermediate who identifies the parties she is mediating between to each other is called *transparent intermediate.* A typical example is the `Bookworlds` book-club which advertised that when ordering on-line, if they don't have a certain book in their stock, they will provide it to their customers through the services of `bol.com`, a mainstream virtual bookstore. A trivial example is an entity that simply redirects to another entity. An intermediate who identifies the existence of the parties she is mediating between to each other but not their identity is called *translucent intermediate.* Typical examples include Internet retailers who advertise that their products will be delivered by a courier company, without specifying which. An intermediate who hides the existence of the parties she is mediating between from each other is called *overcast intermediate.* Typical examples are virtual enterprises, and ventures who provide a portfolio of e-services, some of which may be outsourced to unidentified strategic allies. An intermediate who is authorised to act as a substitute of another entity is called a *proxy intermediate.* A proxy is transparent but she does not reveal her own existence to the trustor.

Adviser is a party that offers recommendations about the credibility of another party. Advisers include the authorities maintaining blacklists for a community. Typical examples include, credit scoring authorities, reputation systems, systems that offer citations, etc.

Proposition 3 (P1) *[dis]trust is not transferred along an overcast intermediate.*(P2) *trust is transitively transferred (to commitment) along transparent intermediates.*(P3) *[dis]trust in all subcontractors of a transparent intermediate is transferred to an inclination to [dis]trust the intermediate.*(P4) *trust is transferred anonymously along translucent intermediates.*(P5) *trust in an adviser is transferred to the recommended parties.*(P6) *distrust in the recommended parties is transferred to (inclination to distrust) the adviser.*(P7) *distrust propagates through trust.*(P8) *distrust obstructs the propagation of trust.*

Remarks:

P1 states that overcast intermediates obstruct the propagation of trust by hiding the identity of those they mediate for. Assume A trusts an overcast intermediate T for a service X and T trusts another party B to subserve X. A is not aware that B subserves but A's measurement of trust in T may be influenced by T's trust in B to subserve, if T decides to communicates the measurement of her trust in B as the measurement of trust in herself (T) for the subservice in question.

P2 states that, if a party A trusts a transparent intermediate T for X and T offers to mediate to a party B whom she trusts for X then, if A accepts, A is committed to trust B for X. Note that distrust is not necessarily transferred along a transparent intermediate. On the assumption that A distrusts `Bookworlds` one cannot infer that A distrusts `bol.com`.

P3 states that, if a party A trusts all subcontractors of a transparent intermediate T for a service X, then A is inclined to trust T for this service;[2]

P4 states that, if a party A trusts a translucent intermediate T for X and T trusts a party B to subserve for X then A is committed to trust B to subserve X without being necessarily aware of B's identity. On the other hand, B's trust in T depends on the expectation that those who use the service will not be aware of her (B) identity.

P5 states that trust in a recommended party depends and assumes trust in the adviser. The measurement of trust however is balanced against other recommendations. Obviously, distrust is not transferred in this direction.

P6 is complementary to *P5* as it provides a means to question the integrity of an adviser on the basis of distrust in some of her recommendations.

P7,P8 state that trust expands business activity space while distrust prunes this expansion.

It should come as no surprise that by defining trust as a measurable belief one gives ground to conflicts. The following is a typical example. First, party A certainly trusts an advisor T_1 for a transaction X and T_1 certainly trusts (and recommends) party B for this transaction. Second,

[2] The measurement of trust of A in T depends on A's trust in the subcontractors as a collective and any other information A may have stored or collected about T.

party A also certainly trusts an advisor T_2 for the same transaction and T_2 certainly distrusts B for that service (and alarm A). Then A faces a conflict: A is willing to commit to trust B for X and also willing to commit to distrust B for the same transaction. Such conflicts naturally appear when reasoning in systems where one's perception and knowledge evolve in time, and manifest that belief and knowledge need to be monitored and reassessed. An important aspect of trust management (section 4) is concerned with providing strategies to overcome such conflicts.

3 Formalisations of Trust

There have been some attempts to formalise aspects of trust relationships within a logical framework. The majority of these attempts involve variants of first order logic or tailored modal logics with distinctive deontic elements. In general, one can expect the logic used for modelling trust to be rich enough to represent actions and interactions between distributed agents, temporal constraints and deontic statements relating to duty and obligation as ethical concepts. In [3] we review some indicative proposals and compare them with our recommendation in an attempt to elaborate why logic alone is not sufficient for modelling the complex trust relationships that appear in open distributed systems.

3.1 Modelling Patterns of Trusting Behaviour

A formal theory to reason about situations where an agent is given institutionalised power to ensure a certain state of affairs is presented in [9]. The focus is on a conditional relation $\Rightarrow_x$ called *"counts as"* which is used for encoding the idea that, within a given institution, the performance of an act in a given context by a designated agent "counts as" a way of establishing a particular institutional fact. A sentence of the form $A \Rightarrow_x B$ describes that according to an institution x establishing that the state of affairs described by A is, counts as a means of establishing that the state of affairs described by B is. This is extended in [8] to a formal theory about trust and deception. The theory is based on an integration of a modal action logic to specify the actions of an agent, a belief logic and a deontic logic to describe commitment. The action logic component is axiomatically similar to a relativised classical modal system of type ECT, with one further axiom schema asserting that logical truths fall outside anyone's agency. A family of modalities E_i is associated with this component, where $E_i\varphi$ denotes that *" agent i brings it about that φ"*. The belief component logic is axiomatically similar to a relativised KD45, reading expressions of the form $B_i\varphi$ as *"agent i believes that φ"*. The deontic component logic is axiomatically similar to a relativised normal modal system of type KD. Sentences of the form $O_x\varphi$ are read as *"the optimal functioning of system x requires the establishment of φ"*.

The resultant formalism is rich enough to accommodate and reason with statements like *"agent 1 believes that within x agent 2 makes φ happen"* $(((E_2\psi \Longrightarrow_x O_x\varphi) \land B_1E_2\psi) \rightarrow B_1\varphi)$. Such a formal language can be used to reason about deception and (dogmatic) trust in an entity. For example $(\neg B_2\varphi \land E_2B_1\varphi)$ denotes that *"agent 2 does not believe φ but gets agent 1 to believe it"*. Although the formalism is well suited for describing patterns of trusting behaviour, it relies on the existence of trust relations and cannot accommodate reasoning about procedures that establish trust between agents. There is also little or no support for capturing lack of trust or uncertainty in an agent's trust.

3.2 Modelling Subjective Opinions

The problem of assigning trust values in the presence of uncertainty is addressed though, in Jøsang's trust model which incorporates the concept of an *opinion* based on subjective logic [13]. The deployed metric has its origin in Dempster-Shafer *Theory of Evidence* [18]. The first step in applying this belief model is to define a set of possible situations which is called the frame of discernment. A frame of discernment delimits a set of possible states of a given system, exactly one of which is assumed to be true at any one time. The elementary states in the frame of discernment Θ are called atomic states because they do not contain substates. An agent's opinion is a representation of a belief and is modelled as a quadruple (b, d, u, a) where b measures belief, d measures disbelief, u measures uncertainty, and a measures relative atomicity, such that $b + d + u = 1$. Roughly, relative atomicity a normalises uncertainty by incorporating the percentage of the atomic states in Θ that are covered by a state $x \in 2^{\Theta}$ about which the opinion is formed. See [13] for further details. Assuming an opinion $\omega = (b, d, u, a)$, the probability expectation of ω is $b + a * u$ which is analogous to the pignistic probability described in [19]. Opinions can be strictly ordered by firstly sorting them in a total order according to probability expectation and then sorting remaining equals according to certainty. Opinions can be deterministically established if all available evidence can be analysed statistically [11].

Subjective logic is defined by integrating classical logic and probabilistic measurement. In addition to the usual connectives of conjunction, disjunction and negation, subjective logic has connectives for *recommendation* $\otimes$ (which forms the opinion of an agent A for a proposition φ based on B's recommendation about φ and A's opinion about this recommendation) and *consensus* $\oplus$ (which represents the opinion of an imaginary agent [A+B] about φ based on A and B's opinions about φ). If the opinions are dogmatic (i.e. $b = 1$ or $d = 1$) then conjunction, disjunction and negation in subjective logic are the same as in classical logic. The operator $\otimes$ is similar to the "discounting operator" of [18]. The operator $\oplus$ is an improvement to "Dempster's rule" of [18]. This operator is partial: as one would expect $\oplus$ reduces uncertainty, therefore a consensus of two incompatible dogmatic opinions cannot be defined. These operators are used in [12] in order to reason about decisions involved in authenticating public keys based on recommendations and certificates and illustrate that, in this case, trust in remote agents can be determined by embedding trust recommendations inside public key certificates.

3.3 Integration Scenarios

None of the above mentioned formalisations is rich enough to capture the mixture of measuring trust in the presence of uncertainty, on the one hand, and relating trust to commitment or expectation to exhibit a prescribed behaviour, on the other hand, that is prominent in our analysis (elaborated in section 2 of [3]). However, as we already elaborated, it seems plausible that an appropriate integration of the formalisms provided in [4] and [13] may provide the basis for formalising our model. A defect that will survive such an integration is that one cannot guarantee that users will assign values appropriately and therefore cannot ensure determining trust consistently. However, analogous problems of determining measures of probability are well-recognised in risk analysis in the process industry (e.g. safety) and finance. In principle, this is a matter of achieving the right means of abstracting information from reality into a mathematical model. We expect that by using logical reasoning along side risk analysis in an integrated trust management framework built around a rigorous system description, one can produce the right guidelines and metrics to make the logical models work.

4 Trust Management

Solutions to the shortcomings of existing trust management systems can be better addressed by separating the trust management framework from the purpose of the application. To achieve this, we need to develop methods to model trust management, i.e, to systematise the process by which control mechanisms and trust policies can be developed.

We define trust management as *the conception, evaluation and enforcement of trusting intentions.* Trust management aims to provide a coherent framework for determining the conditions under which a party A takes the risk to depend on a party B with respect to a service X even though negative consequences are possible. Increasing the levels of trust facilitates processes to become more efficient but also increases the risk of the exploitation of any vulnerabilities of a computer dependent system (including the agents and the infrastructure). One would consequently aim to *maximise trust while minimising risk.* Hence, trust management subsumes and relies on risk management: First, one may employ tailored risk analysis techniques in order to capture and measure trust from the environment. Second, risk management allows us to combine risk with trust in order to form a policy (understood as a description of trusting intentions).

Inclinations	– Situational Trust	
	– Beliefs:	- benevolence
		- honesty
		- competence
		- predictability
		- dispositional trust
	– System Trust	

Intentions	– Resource Access
	– Provision of Service
	– Certification-based Trust
	– Reputation-based Trust
	– Delegation Trust
	– Underwriting Trust
	– Infrastructure Trust

Behaviour	– Enactment Trust
	– Enablement Trust
	– Regulatory Trust
	– Recommendatory Trust

Fig. 1. Classifications of trust.

Third, risk management should enable the evaluation of the impact of a failure in trust and help device a countermeasure strategy.

In order to be able to apply risk management as means of assessing dependability and abstracting the relevant information from reality into formal models, one has to provide a conceptual classification of the different aspects of trust and the different ways they can influence behaviour. We adapt the solution of [17] which is based on the conceptual framework developed in [16]. Our adaptation incorporates the following concepts.

Dependable Behaviour describing the extent to which a party behaves dependably including the act and effects of trusting, i.e., the extent to which a party has decided to depend on other parties for a specified period within a specified context. It implies acceptance of risks (potential of negative consequences) and their effect.

Dependable Intentions describing the extent to which a party is willing to depend on other parties (including oneself) for a specified period, within a specified context and in relation to a specific service. Dependable intentions encompass policies (i.e. rules that can be used to change the behaviour of a system [2]), and meta-policies (i.e. policies about which policies can coexist in the system or what are permitted attribute values for a valid policy [2]). Meta-policies are particularly useful for *resolving conflicts* [15]: a conflict may arise, for example, if there are two policies according to one of which a party is willing to commit into a specific business relationship with another party, whilst according to the other the same business relationship with the same party has to be avoided.

Risk Management aims to control risk; it is the "total process of identifying, controlling and minimising the impact of uncertain events" [6]. It is about managing resources wisely, protecting clients from harm, and safeguarding assets. The risk analysis construct, in particular, is critical for achieving the right means of abstracting information from reality into a mathematical model and supporting the formation of policies. (See figure 1.) It allows one to identify threats, analyse and measure risks, assess treatments of risk and incorporate this information into the formal models. The importance of risk analysis as a means of abstraction from the real world into formal or mathematical models is recognised in the process industry and finance, and elegant models and techniques have been developed for risk management in these areas.

Trust Inclinations is a intentionally broad term, which we use in order to refer to the tendencies of an agent to a particular aspect, state, character or action. They include the following constructs in McKnight *et al.*'s conceptual framework: *situational trust* (i.e., the extent to which a party is willing to depend on an unspecified party in a specific role in a given circumstance); *beliefs* (i.e., a party's schema about the environment it inhabits); *system trust* (i.e., the extent to which a party believes that she can a depend on the known institutional structures and the underlying technology infrastructure).

The above-mentioned dimensions of trust relationships give rise to complementary classifications of trust. The first dimension gives rise to a conceptual classification focusing on how a party perceives trust and forms trust intentions. The second dimension gives rise to an operational classification of trust focusing on how the intention to trust is controlled and exercised. The third dimension gives gives rise to a classification of trust which focus on the roles of the stakeholders and the types of trust that appear in interaction between agents within a context. We further elaborate these classifications in section 4 of [3].

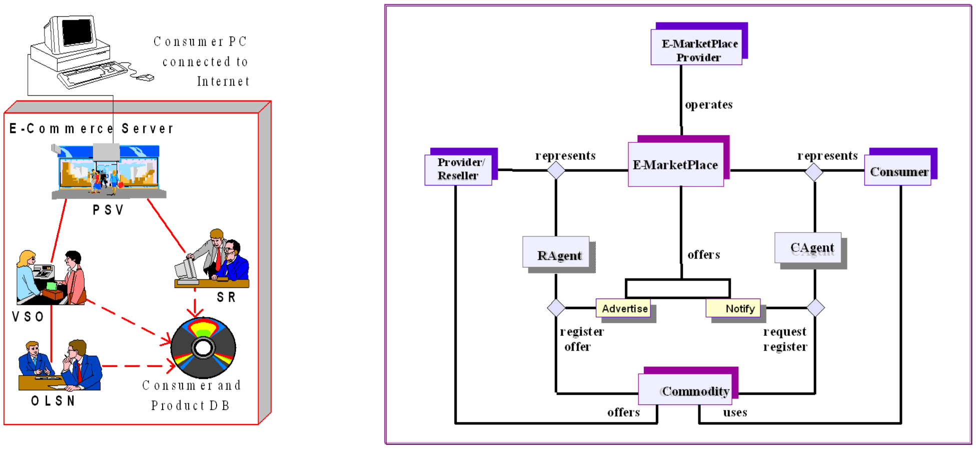

Fig. 2. An outline of the E-commerce platform (left) and Basic Business Relations and Entities underlying OLSN (right).

5 An e-Commerce Scenario

In section 5 of [3] we summarise one of the two major user trial scenarios of the European project CORAS [20, 1], a recent industry lead EC Framework V project involving 10 partners (3 commercial, 5 research institutes and 2 educational) from 4 European countries. CORAS is developing a framework for precise, unambiguous, and efficient risk analysis of security critical systems. This framework will be evaluated by means of major user trials in e-commerce and e-medicine.

The e-commerce user trial will use the Home Shopping Tool (HST) component of the ACTIVE platform. The ACTIVE platform was developed in the framework of the R&D project ACTIVE, (co-funded by the European Commission under the ESPRIT program EP 27046), which aimed to introduce a global Electronic Commerce platform that supports integrated retail services, providing an intelligent interface upon which the involved parties (retailers, suppliers and consumers) establish a tied and trusted relationship.

The e-commerce user trial will use parts of a Home Shopping Tool (HST) provided by Intracom S.A. (http://www.intracom.gr). HST provides to the consumers a personalised environment based on a set of advanced services and facilities that transform the shopping process to an entertainment experience, while at the same time acts on behalf of and for the interest of the consumer. It also supports separate modules for the support of negotiation sections and personalisation (personal WEB pages).

HST delivers a personalised, targeted marketing experience to the consumers through the realisation of a variety of services including *personalised shopping* (the consumer can specify the categories of products and bookmarks to products she would like to access as she enters the store), *catalogue information*, *shopping facilities* (such as quick shopping, shopping baskets and shopping lists), *product search*, *product recommendations*, *sales negotiation*, *e-payment*, and *user management facilities* including analysis and presentation of Point Of Sales (POS), product orders, and consumer profile data.

Notably, the consumers and suppliers are provided with an agent-based automated bargaining mechanism. This mechanism allows customers to find and negotiate products of their interest with various suppliers, and suppliers to promote their products and try to attract customers. The consumer can create an agent, and order him to find and negotiate the purchase of a product according to his preferred product attributes (e.g. product name, quantity, price and time). This agent will get involved in a negotiation process and will try to reach a mutual agreement according to the mandate given by her creator.

Furthermore, HST provides an open payment architecture that could incorporate most of the payment systems currently available, a Consumer Information Model (CIM) provides the necessary metadata for the description and update of all basic concepts and services. Finally, for each consumer, specific information is gathered for the purpose of behaviour analysis and can be made available to the platform operator.

These services are offered to the users with the help of the following software modules depicted in figure 2: The Virtual Shopping Operator (VPO), the Shopping Recommender (SR), the On-line Sales Negotiator (OLSN), the Personalised Store Visualiser.

VPO is responsible for the provision of basic purchasing services such as, electronic baskets, which hold the selected products, access to product information through electronic catalogues and Point of Sales (POS) type of services (i.e. secure order processing and payment).

SR provides recommendations to the users regarding products purchasing i.e. discounts, special offers, new products, contests and lotteries, etc. taking into account the consumer's profile, which is registered in the Consumer Information Model (CIM) database. In addition, recommendations provided, based on other consumers' comments.

OLSN finds, based on agent technologies, negotiate and purchase products on behalf of the shoppers and retailers, based on a set of user-specified constraints such as, desired price, a highest (or lowest) acceptable price, technical specification, a date by which to sell (or buy), etc. Launching of agents is optional depending on consumer or retailer requests. Some basic business relations and entities of OLSN are depicted in Figure 5.

PSV is responsible for the GUI and virtual store customisation. This provides a personalised view of the virtual store based on user profiles and on advanced presentation schemes.

In the course of the CORAS project, the E-commerce platform will be modelled in order to perform risk analysis on the security aspects of the platform. We will further analyse these models (in parallel to CORAS) and use them as indicative examples in order to highlight basic security and fairness aspects that appear in retail (B2C) e-commerce and relate them to trust.[3] During this analysis we will assess the effectiveness of, and further develop, the model of trust outlined in this paper. The results of the security risk analysis being conducted in CORAS will be used as input to our working model of trust management. However, note that the focus of CORAS is on security assurance rather than trust and, to that extent, our analysis may not directly relate to the assessment criteria that apply to CORAS.

6 Conclusion

In this paper we provided a working definition of trust in e-commerce and classify basic trust relationships underlie e-service provision. As indicative example, we refered to [3] where some typical e-commerce security aspects to trust on the basis of a "real-life", industrial size, e-commerce application. We also surveyed recent trends in formalising trust, we indicated what we think is missing from existing formalisations of trust, and we discussed how risk analysis and formal modelling can be combined to facilitate a solution. This discussion is timely in view of our plans for extending a modelling framework under development in order to incorporate trust elements in the development and deployment of e-commerce enabling technology. (See [3] for an elaborate presentation of these plans.) So far, our research has provided evidence that there are methods, formalisms and conceptual models which, if appropriately integrated, can bridge the gap between systems modelling, trust and risk management in e-commerce. However, there is still a long way to go. Effective solutions to such problems require interdisciplinary approaches, which provide a fertile ground for the application of many tools from cognitive sciences and economics in addition to computer science.

[3] Many fairness and security properties in such e-commerce platforms either appear as refinements or rely on more abstract trust properties. For example, being authenticated implies being trusted enough to do business with, either on the basis of positive experience and sufficiently low risk, or on the basis of credentials provided by an authority. In the latter case trust in the issuing authority is balanced against any potential negative information broadcasted from other authorities. Authorisation assumes trust in the identity of the authorised party or her suitability for the role she is requesting to play. The access rights depend on the level of trust and the risks associated with the actions available to the requester after authorisation is granted. Assuring accountability (including non-repudiation as a special case) implies trust in a party to be aware of its own history and to be ready to provide explanation for her history throughout the duration of a service.

Acknowledgement

Motivating and fruitful discussions with the CORAS partners and the participants of the CLRC/CORAS workshop on *Semi-formal Modelling, e-Risk and e-Trust*[4] contributed to the improvement of this paper.

References

1. CORAS - A Platform for Risk Analysis of Security Critical Systems (http://www.nr.no/coras). See also http://www.itd.clrc.ac.uk/Activity/CORAS
2. N. Damianou, N. Dulay, E. Lupu and M. Sloman. The Ponder Policy Specification Language Proc. Policy 2001: Workshop on Policies for Distributed Systems and Networks, Bristol, UK, 29-31 Jan. 2001, Springer-Verlag LNCS 1995, pp. 18-39
3. T. Dimitrakos. System Models, e-Risks and e-Trust - Towards bridging the gap? Tehcnical Report. CLRC ISE-ITD Feb. 2001. http://www.itd.clrc.ac.uk/Publications/1331/eTrustReportFeb01.pdf
4. B.S. Firozabadi and M. Sergot. Power and Permission in Security Systems, in 7th International Workshop In Security Protocols, 1999, Cambridge, UK. LNCS, Springer-Verlag.
5. T. Grandison and M. Sloman. A Survey of Trust in Internet Applications IEEE Communications Surveys and Tutorials, Fourth Quarter 2000,
6. Information technology - Security techniques - Guidelines for the management of IT Security (GMITS) - Part 1: Concepts and models for IT Security. ISO/IEC TR 13335-1: 1996.
7. S. Jones, TRUST-EC: requirements for Trust and Confidence in E-Commerce, European Commission, Joint Research Centre, 1999.
8. A.J.I. Jones and B.S. Firozabadi On the characterisation of a Trusting agent - Aspects of a Formal Approach. In Workshop on Deception, Trust and Fraud in Agent Societies, 2000.
9. A.J.I. Jones and M.J. Sergot. A Formal Characterisation of Institutionalised Power. In Journal of the IGPL, vol. 4, no. 3, 1996, pp.427-443.
10. A. Jøsang, The right type of trust for distributed systems. In Proceedings of the New Security Paradigms Workshop, ACM, 1996.
11. A. Jøsang, A subjective metric of authentication. In D.Gollmann, editor, Proceedings of the 5th European Symposium on Research in Computer Security (ESORICS'98), Springer-Verlag, 1998
12. A. Jøsang, Trust-based decision making for electronic transactions. In L.Yngstr-m and T.Svensson, editors,Proceedings of the Fourth Nordic Workshop on Secure IT Systems (NORDSEC'99), Stockholm University Report 99-005, 1999.
13. A. Jøsang, A logic for uncertain probabilities, International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, Vol 9. No. 3, June 2001, World Scientific Publishing Company.
14. A. Kini and J. Choobineh, Trust in Electronic Commerce: Definition and Theoretical Considerations, In Proceedings of the 31st Annual Hawaii International Conference on System Sciences, 1998, Hawaii,
15. E.C. Lupu and M. Sloman, Conflicts in Policy-Based Distributed Systems Management. IEEE Trans. on Software Engineering, 25(6): 852-869 Nov.1999.
16. D.H. McKnight and N.L. Chervany. What is Trust? A Conceptual Analysis and an Interdisciplinary Model. In Proceedings of the 2000 Americas Conference on Information Systems (AMCIS2000). AIS, Long Beach, CA, August 2000.
17. Dean Povey, Developing Electronic Trust Policies Using a Risk Management Model. In Rainer Baumgart (Ed.): Secure Networking - CQRE (Secure) '99, International Exhibition and Congress Dsseldorf, Germany, LNCS, Vol. 1740, Springer, 1999
18. G. Shafer. A Mathematical Theory of Evidence. Princeton University Press, 1976.
19. Ph. Smets and R. Kennes. The transferable belief model. Artificial Intelligence, 66:191-234, 1994.
20. K. Stølen. CORAS - A Framework for Risk Analysis of Security Critical Systems. To appear in the Proceedings of the International Conference on Dependable Systems and Networks, 2001

[4] Workshop hosted at Rutherford Appleton Laboratory on March 2001 in conjunction the 2nd CORAS meeting. See http://www.itd.clrc.ac.uk/Activity/CORAS+1087 for contributed technical talks and project presentations.