

Trust Evaluation Through Relationship Analysis

Ronald Ashri¹, Sarvapali D. Ramchurn¹, Jordi Sabater², Michael Luck¹, Nicholas R. Jennings¹

¹School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK

²LABSS, Institute of Cognitive Science and Technology, Viale Marx 15, 00137 Roma, Italy

¹{ra,sdr,mml,nrj}@ecs.soton.ac.uk, ²jsabater@iia.csic.es

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems, Coherence and coordination.*

General Terms

Design, Reliability, Management.

Keywords

Trust, Reputation, Relationships, Multi-Agent Systems.

ABSTRACT

Current mechanisms for evaluating the trustworthiness of an agent within an electronic marketplace depend either on using a history of interactions or on recommendations from other agents. In the first case, these requirements limit what an agent with no prior interaction history can do. In the second case, they transform the problem into one of trusting the recommending agent. However, these mechanisms do not consider the *relationships* between agents that arise through *interactions* (such as buying or selling) or through overarching *organisational* structures (such as hierarchical or flat), which can also aid in evaluating trustworthiness. In response, this paper outlines a method that enables agents to evaluate the trustworthiness of their counterparts, based solely on an analysis of such relationships. Specifically, relationships are identified using a generic technique in conjunction with a basic model for agent-based marketplaces. They are then interpreted through a trust model that enables the inference of trust valuations based on the different types of relationships. In this way, we provide a further component for a trust evaluation model that addresses some of the limitations of existing work.

1. INTRODUCTION

Agents generally interact by making commitments to, or contracts with, one another to carry out particular tasks. However, in most realistic environments there is no guarantee that a contracted agent will actually enact its commitments (because it may defect

to gain higher utility, or because there is uncertainty about whether the task can be achieved). In such situations, computational models of trust (here defined as the positive expectation that an interaction partner will act benignly and cooperatively in situations in which defecting would prove more profitable to itself [4]) have an important role to play. First, they can help determine the most reliable interaction partner (i.e. those in which the agent has the highest trust) and second, they can influence the interaction process itself (e.g., an agent's negotiation stance may vary according to the opponent's trust level).

However, when an agent first enters an environment, it has no history of interactions (with the other agents in that environment) that it can analyse to decide who to trust. In such circumstances, current research suggests two possible solutions [9]. It could interact with all agents and then derive trust measures from the history of interactions (as in [12, 10]). Alternatively, it could request reputation information acquired from an existing social network [12], in which reputation is understood as a third party's estimate of trustworthiness. However, there are a number of problems in each of these alternatives. Firstly, if the agent interacts with each agent, it inevitably risks making losses if the counterparts it interacts with are not trustworthy. Secondly, if the agent relies on reputation information, then it cannot be sure that the agents providing the information are doing so truthfully. In both cases, others may be unreliable because they have conflicting interests with the agent (e.g., if they compete in the same market) or because they can collude to exploit the agent (e.g. if some agents know each other and all share their gains from exploiting an agent). In consequence, if an agent could take into account relationships such as competition or situations such as collusion it could produce more robust valuations. However, existing mechanisms do not adequately consider the relationships between agents that arise through the *interactions* between them, which may lead to competitive relationships, or through overarching *organisational* structures, which may lead them to collude. Furthermore, when relationships *are* taken into account in a limited manner, such as in [12], the information used is implicitly assumed and no mechanisms are provided to enable the agent to *discover* that information dynamically, nor to react to changing relationships between counterparts. An agent also needs to be able to *identify* and *interpret* such relationships in a changing environment.

Against this background, in this paper we address just this need by developing a method for identifying such relationships between agents in an electronic marketplace and then using this information to enhance trust valuations. We advance the state of the art in the following ways. First, we develop a process for agents to dynamically identify relationships between agents in an electronic marketplace. Second, we identify the general types of relationships

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.

that should be considered with regards to trust and discuss the types of reasoning such information can enable. Finally, we make use of an ontology-based framework to analyse relationships, providing a realistic application of semantic web technologies.

The following section provides an overview of our approach. Section 3 describes the relationship identification process used and how that is mapped to the specific context of an electronic marketplace. Section 4 introduces the most relevant relationship types with regard to trust and Section 5 discusses how knowledge of such relationship can affect trust valuations as well as how the work here can directly be used by existing trust models. Conclusions and further work are given in Section 6.

2. OVERVIEW OF APPROACH

Our overarching aim is to improve the derivation of trust values for agents by taking into account information about the relationships agents can infer its counterparts have with each other and itself. Below we provide an overview of this process, while the rest of the paper expands on each individual step.

1. **Relationship Identification** First, we must *identify* relationships between agents. Keep in mind that it would be unrealistic to expect to know exactly what relationships exist between various agents, rather what we attempt to do is identify what relationships *may* exist given what information is available to us. In order to do so we make use of a generic relationship identification process [1] that uses a model of agent interaction with the environment to infer when agents may be related given their individual capabilities. We then map this generic process to a more specific *Agent-Based Market Model* (ABMM) that is based on a typical e-commerce example of an electronic marketplace. The generic process provides us with the necessary conceptual grounding, while the ABMM provides us with a domain-specific view (in the case an e-commerce marketplace).
2. **Relationship Characterisation** Having identified the possible relationships, within the context of the ABMM, we then distinguish the types of relationships that are most relevant with regards to trust. These types of relationships will provide us with a set of *patterns* that agents can use in trust evaluation.
3. **Relationship Interpretation** Using these relationship patterns, and their interpretation through additional information about the specific context in which an agent operates we then discuss how such relationships can be interpreted to derive trust valuations.

3. RELATIONSHIP IDENTIFICATION

3.1 Relationship Identification Model

The ability to identify the different types of relationships in an environment can be used to enable us to reason about the possible underlying motives of agents and, as a result, derive trust valuations. For example, consider a situation in which agent α wants to sell us a product, and α and β belong to the same organisation. Then, if we ask β for a rating of α 's product quality we should not place much credence on the reply, since an ulterior motive, relating to the overall gain of the organisation that α and β belong to, could bias the reply.

However, in order for the identification process to be widely applicable we require a principled approach that can be made part of




Figure 1: Region of influence affects viewable environment

an agent program. In this section, we briefly present such a relationship identification process, which we will subsequently map to a specific domain, which for the purposes of this paper is an electronic marketplace.

This process is based on a model of interaction between agents and their environment that makes no assumptions about any internal agent components, since they cannot be observed. The focus, therefore, is on the *interface* between individual agents and their environment, through the capabilities of agents. The notions that underpin this model are based on the SMART framework [6], and are discussed in more detail in [1, 2], so they are only briefly described below. We first present the underlying agent model provided by SMART and then explain how we use it to create a model of interaction with the environment, and by consequence other agents.

Agents. For the purposes of relationship analysis, an agent is considered an entity described by a set of *attributes*. Attributes are simply describable features of the *environment*, and are the only characteristics that are manifest. Agents are able to perform *actions*, which can change the environment by adding or removing attributes. Agents also pursue *goals*, which are desirable environment states described by non-empty sets of attributes. Agents are denoted by the set Ag noted as $\alpha, \beta, \dots \in Ag$.¹

Agent Perception and Action. Agent actions are divided into those that retrieve the values of attributes, representing the agent's *sensor capabilities*, and those that attempt to *change* attribute values of the environment, representing the agent's *actuator capabilities*. We note the set of actions as A .

Viewable Environment and Region of Influence. Given that agents interact with the environment through actuators and sensors, and that the environment as a whole is defined through a set of attributes, we can intuitively think of actuators and sensors as defining *regions of the environment*, or subsets of the entire set of attributes that make up the environment. The attributes that an agent's actuators can *manipulate* define a *Region of Influence (RoI)*, while the attributes that an agent's sensors can *view* define a *Viewable Environment (VE)*. We note an agent α having a certain region of influence as RoI_α and an agent β having a certain viewable environment as VE_β .

The VE and the RoI of an agent provide us with a model that relates an agent and its individual capabilities to the environment. In order to identify relationships between agents we need to look at how their VEs and $RoIs$ overlap. The different ways in which these overlaps occur plays a role in determining the possible relationships between them. In Figure 1, these concepts are illustrated by using an ellipse to represent the VE and a pentagon shape for the RoI . We

¹In the interest of space we have decided to adopt a simplified notation for the purposes of this paper. A more formal elaboration of SMART and the generic relationship model, can be seen in [1, 2], where Z is used as the formal language [14].

use this notation throughout when illustrating different situations.

Goals. Knowledge of an agent's goals, in addition to its *RoI* and *VE*, can provide better information about its possible relationships with other agents, since it will identify where within an entire *RoI* (or outside its *RoI*) an agent is most likely to attempt to operate. We note the set of goals as G and an agent α having a particular goal as $g_\alpha \in G_\alpha$ where $G_\alpha \subseteq G$.

In the broadest sense, agents can have only two types of goals. On the one hand, they may want to *effect some change* in the environment, which implies changing attributes of the environment, while, on the other hand, they may just want *some information about the environment* which does not lead to any direct changes in the environment. Distinguishing between these two types of goals is important since the latter can only be achieved directly by an agent if that goal is in the *RoI* of the agent, while the former can only be achieved if the goal is in the *VE* of the agent.

We distinguish between these two types of goals by using the same terminology as the dMARS system, as their approach is suitable to our context of use and it is compatible with the SMART framework through the formalisation in [5]. Essentially, a *query* goal is one for which an agent tries to elicit some information, either from its internal beliefs or from the environment. As such, it can be satisfied if it falls within an agent's *VE*. Conversely, an *achievement* goal may require that the agent performs certain actions in order to change the environment, if the environment is not already in the desired goal state. Thus, an *achievement* goal can be satisfied if it lies within an agent's *RoI*. We note query goals as $q(g_\alpha)$ and achievement goals as $a(g_\beta)$.

3.2 Agent-Based Market Model

The ABMM aims to capture most of the features of a typical e-commerce scenario by which sellers and buyers trade in an online market.² The model is defined using the OWL (Web Ontology Language) [8] standard, which is a natural choice, since it enables us to check the consistency of the model and reason about it through the underlying description logics [3], using widely available tools.

The information types and the relationships between them are illustrated in Figure 2, in which large arrowheads represent inheritance relationships, while smaller arrowheads represent property references. In this model, an *Agent* is considered to be any entity that *requires* or *sells* a number of *Products* (e.g. memory chips or computer processors), and can be affiliated to an *Organisation*. We distinguish between an *AtomicProduct* (e.g. computer chips), for which no further division of the product into components takes place, and a *CompositeProduct* (e.g. desktop computer), which comprises several atomic products. This distinction allows us to better represent the situation in which an agent requires a composite product whose atomic components must be sourced from a number of seller agents. Now, an agent *buys from* or *sells in* a *Market*, which has a number of resources of type *Product*. A *Market* is *regimented by* an *Institution*. An *Institution* is an entity that regiments the roles and relationships of the interacting agents, and determines the rules of encounter that prescribe what an agent can do at what point in time [7].

Similarly to a *Product*, a *Market* can be a *SingleMarket* or a *CompositeMarket*. A composite market is one in which the goods traded are inter-related (e.g. buyers and sellers of particular car parts and second-hand cars). The composite market may also be different from a single market in that it is regimented by more than one institution.

²We believe these features of the model are necessary rather than sufficient ones. The model can be easily adapted to cope with more features if required as will be discussed later on.

Once an agent has information about its counterparts as described and related by the ABMM model, it can begin the process of identifying how they are related. For example, if two agents sell the same products in a market, we can assume that they are competitors. If they sell complementary products in a composite market and belong to the same organisation, then we could assume that their opinions of each other may be biased.

3.3 Mapping the ABMM to SMART model

In order to identify what types of relationship can exist between agents, based on the ABMM model, and to ensure that we do this through a principled approach we map the ABMM to the generic identification model. We then make use of the principles of the generic identification model to determine what relationships may exist within agents in this context.

Agent Attributes, Actions In the specific case of market agents, the attributes include information such as the organisation an agent belongs to, the market in which it operates, the available products, and so forth. The basic actions are the ability to buy or sell products. With regard to the ABMM, sensor capabilities are those that allow the agent to perceive other agents in a market and identify relevant information such as available products. The most relevant actuator capabilities are those that allow it to sell or buy a product.

Query and Achievement Goals. A query goal must be instantiated when an agent wishes to buy something from the market, since it must query the market to identify the agents that are able to sell a product, while an achievement goal is instantiated when an agent wishes to sell something in the market. Note that following a query goal an actual transaction must take place between two agents, that will require an achievement goal on the part of the buying agent as well as the selling one. However, we do not attempt to identify such goals since we assume that the market only places buyers and sellers in contact, but information about an actual transaction remains private to the pair of agents involved in the transaction.

Viewable Environment and Region of Influence. With respect to the ABMM, we assume that the *VE* of an agent defines a region of the market that an agent is able to view. This means that an agent can instantiate query goals within its *VE* to identify sellers that sell the desired product as defined by the query goal. Furthermore, if two agents belong to the same organisation, we could decide, depending on the nature of the organisation, to represent the *VE* of each individual agent as the sum of the *VEs* of each member of the organisation, reflecting the assumption that agents would share information. Now, the *RoI* of an agent represents the products that an agent can sell in a market. An agent can instantiate achieve goals within its *RoI* signifying that it is willing to sell a product within a specific market. Therefore, for a buyer and seller to be matched the buyer must have a query goal within its *VE* that overlaps with an achieve goal within the *RoI* of a seller.

In order to clarify these notions we illustrate them through an example. Returning to Figure 1, we show a situation in which α 's *RoI* overlaps with β 's *VE*, and both agents' *VEs* overlap. Given this information, we can infer that α and β operate in a common market where the *VEs* overlap. Furthermore, α sells a product in that market where its *RoI* overlaps with the common *VEs*.

Now, assume that β has the goal to buy a product from A . β instantiates a query goal and uses its sensory capabilities to identify the price and other relevant attributes of the product. In addition, assume that β 's *RoI* represents the sale of a product that β constructs using, in part, the products bought from α . As a result, β now becomes *dependent* on α making that product available at an appropriate price. Thus, whenever α performs an action that affects that product in some way, it will eventually *influence* β 's actions,

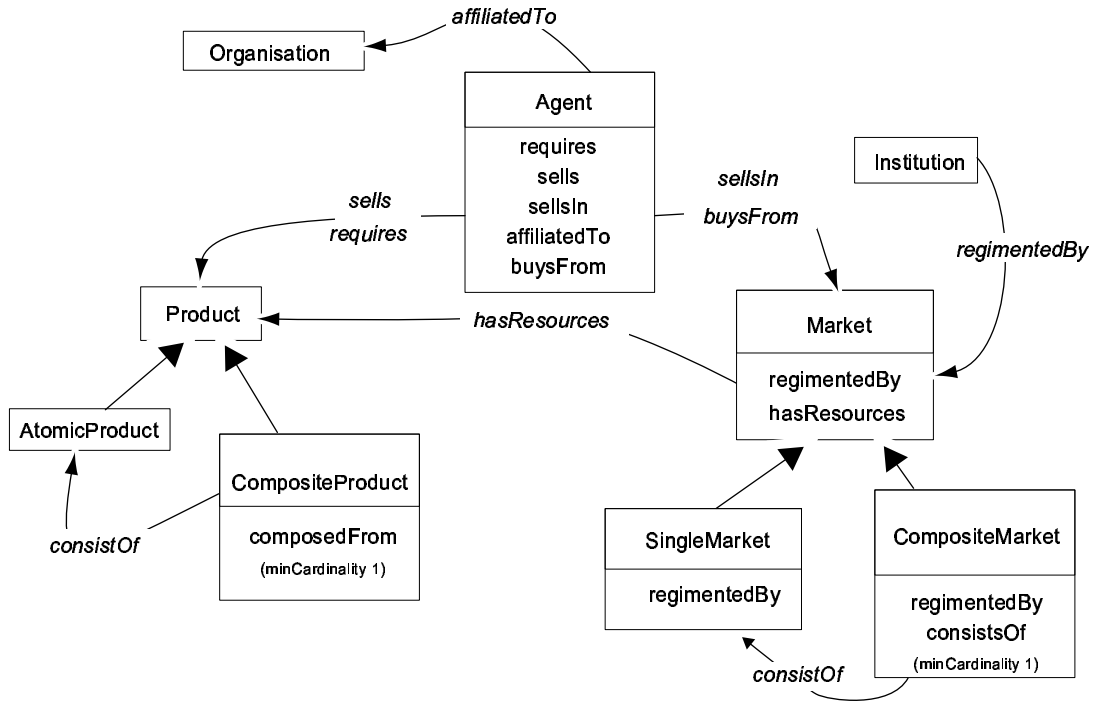


Figure 2: Agent-Based Market Model

since β must now react to the changes when producing its own product.

In the next section, we describe how this model can be used to interpret relationships and describe some of the more significant ones with regards to trust.

4. RELATIONSHIP CHARACTERISATION

Given the conceptual grounding provided by the generic identification model and its mapping to the ABMM we can now identify a variety of different types of relationships between agents. However, as there is a large number of possible relationships and combinations of relationships, as well as a large number of possible interpretations of those relationships, it is necessary to minimise the amount of decision making needed in order to interpret the significance of the relationship. To this end, we define some *basic* types of relationships that are clearly relevant to trust valuations and can be combined to describe more complex types.

The relationship types, of the set R , defined here build on results from an existing trust model³ [12], and represent the types we consider most salient with regards to trust. Each type is represented by a *pattern* representing a specific configuration of *VEs*, *RoIs* and goals.

Furthermore, in addition to the type of relationship pattern (which is identified as either existing or not) we must also consider the *context* in which the relationship is developing. The context, noted as C , deals with issues that are not directly captured by the pattern alone. Issues such as the abundance of a product, the number of sellers of the product, and the amount being bought, define what

³Note here that our aim is not to define a new trust model. Rather we focus on showing how existing trust models can use relationship analysis techniques to enhance their efficiency. In this context [12] provide a good starting point since they make a clear link to relationships in their model.

we term the *intensity* $I : C \times \mathcal{R} \rightarrow [0, 1]$ of the relationships \mathcal{R} . The instantiation of an intensity calculation function will depend on the type of application. For example we could define the intensity of a relationship as dependant of the ratio of the amount of goods that each agent trades in a market or the number of seller of a specific type of product.

Note that in order to make this analysis we need to consider a specific context (in our case that of e-markets). The same configuration of *VEs* and *RoIs* in a different context, or based on a different ABMM, would possibly have a completely different interpretation in terms of relationships.

4.1 Trade

A *Trade* relationship exists when an agent is able to buy a product from another agent within the same market. This is the most basic type of relationship and just signifies that if one agent wished to buy something from another, this is possible because they can both interact in the same market. We formally specify a trade relationship between agents α and β as $Trade(a(g_\beta^y))_{RoI_\beta \subseteq VE_\alpha}$ for an agent β that can achieve a goal that is within α 's *VE*. The goal here is to sell a product y that α may or may not wish to acquire.

4.2 Dependency

When an agent α selling goods in a market that β can view and buy from *and*, at the same time, β has the goal (represented by a square) to buy the goods α is selling in that market, we have a *Dependency* relationship (illustrated in Figure 3(a)). The intensity of this relationship depends on several factors: the number of sellers in that market that provide the same product, the abundance of these products, etc. Notice that these factors can ultimately determine who depends on whom. For instance, in a market where there are many sellers providing the same product, and very few buyers interested in that product, the roles in a dependency relation like the one described above are interchanged. Then the seller depends

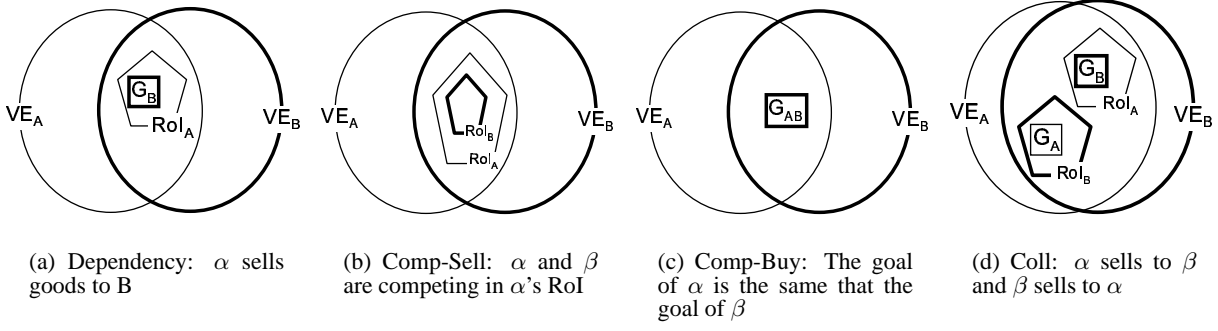


Figure 3: Key Relationship Patterns

on the buyer. We specify a dependency relationships in terms of goals in the following way: $Dep(q(g_\alpha^y), a(g_\beta^y))_{RoI_\beta \subseteq VE_\alpha}$, where y is the product β is selling to α (i.e. α wants to achieve the goal of having y), and β 's region of influence is within α 's viewable environment as for trade relationships.

4.3 Competition

In the configuration of Figure 3(b), termed a *Comp-Sell* configuration, the RoIs of agents α and β intersect. This implies that both agents are selling the same goods in the same market. This reflects a competition in that area of influence. The intensity of this competitive relation is determined by several factors such as market share, profit, and cost of goods. This relationship is noted as $Comp(a(g_\alpha^y), a(g_\beta^y))_{RoI_\alpha \subseteq RoI_\beta}$, where y is the product each β and α desire to sell in the environment. Obviously, the agents may also intend to buy products and compete in doing so (figure 3(c)). This can be easily represented by swapping $a(g^y)$ for $q(g^y)$ as in $Comp(q(g_\alpha^y), q(g_\beta^y))_{RoI_\alpha \subseteq RoI_\beta}$.

The configuration of Figure 3(c) also reflects a competitive relationships, termed (*Comp-Buy*). In this case, α and β have the same goal, indicating that they want to buy the same products. The intensity of this competitive relation is based on factors similar to the dependence relation presented in Figure 3(a): the number of sellers in that market that provide the products required by α and β , and the abundance of these products in general.

4.4 Collaboration

Figure 3(d) shows a configuration in which α has a goal in the *RoI* of β and β has a goal in the *RoI* of α . This means that α is selling goods to β and, at the same time, β is selling (different) goods to α . This configuration, called a *Coll* configuration, is a composition of two Trade-Dep configurations (see Figure 3(a)). The relationships generated by the two Trade-Dep configurations are a trade and dependence relation between α and β , and a trade and dependence relation between β and α . If α depends on β and β depends on α , we say there is a *collaboration* between α and β . A collaborative relationship between agents α and β is represented as $Coll(Dep(\alpha, \beta), Dep(\beta, \alpha))_{RoI_\alpha, VE_\alpha \subseteq RoI_\beta, VE_\beta}$.

4.5 Tripartite relationships

In this section we provide an indication of how relationships between more agents can be considered, if at least one more agent is added to the analysis.

In the majority of cases, the resulting configurations can be decomposed in terms of the configurations above. One of the exceptions to this is the configuration showed in Figure 4, in which α is




Figure 4: α and β competing to sell to C while α is selling to β

selling a product that β requires, while at the same time α and β are competing in order to sell γ a product. Thus, there are three *Dependency* configurations and a *Comp-Sell* configuration. However in this configuration there is a special situation, since while is a competitive relation between α and β at the same time β depends on α . This situation gives a privileged position to α with respect to β , which has to be considered when analysing trust.

5. RELATIONSHIP INTERPRETATION

Having introduced some of the most salient relationships that can be directly inferred using the ABMM and the relationship identification process, we discuss here how we can make use of them to derive trust valuations. Given our initial definition of trust (in Section 1), we argue that an agent should *distrust* its counterpart whenever the latter has an opportunity to defect, as can be inferred from the relationships with the counterparts and the counterpart's relationships with others. We deal with the trust valuations in two parts: (i) where one agent tries to infer the trustworthiness of its counterpart (bipartite relationships) and (ii) where one agent tries to infer the trustworthiness of its counterpart and both or one of them is related to other agents (multipartite).

5.1 The Trust Model

Here, we describe the trust model (based on [12, 10]) to illustrate how relationships can be directly factored in to agents' analysis of their counterparts' trustworthiness. In this respect, we capture trust of an agent α in β as the combination of α 's confidence in β (based on direct interactions) with the reputation β has in the society of agents (based on reports of direct interactions β has had with other agents). We will note the trust function of an agent as $T : Ag \times Ag \rightarrow [0, 1]$. We also define the confidence an agent has in its opponent as $Conf : Ag \times Ag \rightarrow [0, 1]$ and the reputation of an agent as $Rep : Ag \times 2^{Ag} \rightarrow [0, 1]$. Thus, we allow

the confidence of an agent to be calculated according to the evaluation of direct interactions with it, which in the simplest case is an average of all efficiency reports $\eta^\beta \in [0, 1]$ perceived by α about β : $Conf(\alpha, \beta) = \frac{((n-1) \cdot Conf(\alpha, \beta) + \eta)}{n}$, where $n \in \mathcal{Z}$ is the number of interactions the agent α has been engaged in with β . The confidence model of an agent is usually bootstrapped according to some belief about the reliability of other agents that can be gathered from the relationships that the agents are aware of (as will be seen in sections 5.2 and 5.3). We will note this starting value as s_{conf} . As more reports are obtained (when the agents interact), the confidence values will tend to approximate the expected efficiency of the opponent. On the other hand, the reputation of β is calculated from a set of reports η_γ^β where γ represents the agent providing the information and β the agent that is evaluated. In the simplest case, $Rep(\beta, Ag) = \sum_{\gamma \in Ag} w_\gamma \cdot \eta_\gamma^\beta$ where $\sum w_\gamma = 1$. The values of w_γ may be calculated in a number of ways. For example, in the ReGrE system, the values for w_γ are chosen according to the relationships agents have with each other while in [13] it is shown how these values are determined according to the truthfulness of the agent γ providing a report. Given this, the trust of α in β is calculated as $T(\alpha, \beta) = \kappa \cdot Conf(\alpha, \beta) + (1 - \kappa) \cdot Rep(\beta, Ag)$, where $\kappa \in [0, 1]$. Usually we could expect κ to increase as the number of interactions between α and β (that is n), increases. Given these definitions, in the following sections, we present a conceptual framework within which relationships and combinations of these can be used to select both the starting value of confidence models and weights of reports used in evaluating reputation and trust in general.

5.2 Bipartite Relationships

The reasoning that knowledge of bipartite relationships enables with regard to trust is described below for each type of relationship identified above.

1. *Dependency* – if β is dependent on α , then α may have an opportunity to exploit β if β has no other choice than α as an interaction partner. In the case where the intensity of dependence is high (e.g. in terms of amount of goods traded and percentage of total costs to β), β 's trust in α should be the lowest possible (and conversely if the dependence is low). This can be modelled in the confidence function by setting the initial value of confidence to 0, with $R = Dep(q(g_\beta^y), a(g_\alpha^y))_{RoI_\beta \subseteq VE_\alpha}$ as follows:

$$\text{If } R \text{ Then } s_{conf} = 1 - k \cdot I(C, R)$$

where $k \in [0, 1]$ is a constant, C the context and $I(C, R)$ the intensity of the relationship. Here, we set the initial trust to an arbitrarily low value if the intensity of the relationship is high according to the context and a high value if the intensity is low.

2. *Comp-Sell* and *Comp-Buy* – these competitive relationships obviously do not favour trust between agents since it is in their interest to undermine each other in all possible ways. In such cases it becomes more important to have a more refined understanding of how such agents are related to other agents within the environment and how the information they transmit should be interpreted, as we will see in the next section.
3. *Coll* – in this case, both agents gain by not defecting during their interactions since they both depend on each other to achieve their goals. Depending on how intense this relationship is, we would expect these agents to trust each other

highly if they are strongly interdependent, and not place much trust in each other if they are not equally dependent on each other (e.g. if α depends on B more than β depends on α , β could defect on α and α will not be able to compensate such defections by defecting on β). In these cases, we specify a high starting confidence value and also allow agents to modify the weights they place on the reports provided by their collaborators about other agents in the reputation model. Hence, for

$$R = Coll(Dep(\alpha, \beta), Dep(\beta, \alpha))_{RoI_\alpha, VE_\alpha \subseteq RoI_\beta, VE_\beta} \text{ and a context } C,$$

$$\text{If } R \text{ Then } w_\beta = k \cdot I(C, R) \text{ and } s_{conf} = k \cdot I(C, R)$$

where w_β is the weight given by α to β 's report about other agents in the reputation model and $k \in [0, 1]$ is a constant. Thus, the weight and starting confidence values are directly proportional to the intensity of the collaboration.

As can be seen, when each of the above relationships is detected, a particular rule will fire. However, it may happen that two or more of these relationships may happen at the same time. In such cases, a number of solutions are possible. One possible way of deciding the values for weights and s_{conf} is to use fuzzy sets with the intensity $I(C, R)$ as the tuning mechanism to combine the outputs of each rule. Another way would be to use case-based reasoning to select (over repeated interactions), the most significant rule(s).

5.3 Multipartite Relationships

We can now discuss how the intensity of relationships and their combinations (i.e. where agents can be related to more than one other agent) can lead to more informed decisions given the environment these agents interact in. Some examples are given below.

1. α depends on β in *Dependency* while β and γ are in *Coll*. In this case β may have an incentive to misrepresent the reliability of γ to α . This may happen because β could gain from a more profitable γ , and would therefore provide unrealistically high ratings for γ . Alternatively, β may wish to continue to hold γ captive in their collaboration, and provide unrealistically low ratings so that γ is not able to become more independent. Given the intensity of these relationships, the credibility of β 's reports will be decreased or increased to different degrees. We represent this as follows. Given a relationship $R_1 = Comp(a(g_\gamma^y), a(g_\beta^y))$ and $R_2 = Dep(q(g_\alpha^x), a(g_\beta^x))$, and a context C , the rule is:

$$\text{If } R_1 \text{ and } R_2 \text{ Then } w_\beta^\gamma = (1 - k) \cdot I(C, R_1) \times I(C, R_2) \quad (1)$$

where k is a constant and w_β^γ is the weight chosen by α to weight β 's reports about γ in its reputation model. However, if roles between α and β were reversed, and β depended on α in *Dependency*, then the converse of the above reasoning may apply and α might give more value to β 's reports on γ . This means that for a relationship $R_3 = Dep(q(g_\beta^x), a(g_\alpha^x))$, then:

$$\text{If } R_1 \text{ and } R_3 \text{ Then } w_\beta^\gamma = k \cdot I(C, R_1) \times I(C, R_3)$$

where $k \in [0, 1]$ is a constant.

2. α is in *Comp-Sell* or *Comp-Buy* with β and β is in *Coll* with γ . In this case, α will obviously distrust β 's reports about γ since β could gain from giving false reports about γ to α , as discussed above. To this, we apply a similar rule as in equation 1.

3. α is in *Coll* with β and β is in *Coll* with γ and α is in *Coll* with γ . All the agents should trust each other fully. This situation may arise if all the agents form part of the same organisation or form a cartel. The latter form of collaboration might be very profitable to the agents α , β and γ , but might affect the performance of the system within which they operate with other agents. In such circumstances, it is up to the institution to ensure good behaviour. So, formally, for an agent α and given:

$$R_1 = \text{Coll}(\text{Dep}(\alpha, \beta), \text{Dep}(\beta, \alpha))_{\text{RoI}_{\alpha, \text{V E}_{\alpha} \subseteq \text{RoI}_{\beta, \text{V E}_{\beta}}}}$$

$$R_2 = \text{Coll}(\text{Dep}(\alpha, \gamma), \text{Dep}(\gamma, \alpha))_{\text{RoI}_{\alpha, \text{V E}_{\alpha} \subseteq \text{RoI}_{\gamma, \text{V E}_{\gamma}}}}$$

$$R_3 = \text{Coll}(\text{Dep}(\gamma, \beta), \text{Dep}(\beta, \gamma))_{\text{RoI}_{\gamma, \text{V E}_{\gamma} \subseteq \text{RoI}_{\beta, \text{V E}_{\beta}}}}$$

then:

If R_1 and R_2 and R_3

Then $w_{\beta}^{\gamma} = w_{\gamma}^{\beta} = k_1 \cdot I(C, R_1) \times I(C, R_2)$

and $w_{\alpha}^{\beta} = w_{\beta}^{\alpha} = k_2 \cdot I(C, R_2) \cdot I(C, R_3)$

and $w_{\alpha}^{\gamma} = w_{\gamma}^{\alpha} = k_3 \cdot I(C, R_1) \cdot I(C, R_3)$

where $k_1, k_2, k_3 \in [0, 1]$ are constants chosen by agents α, β , and γ respectively, and the weights w_b^a are chosen by $c \neq a, b$, where $a, b, c \in \{\alpha, \beta, \gamma\}$. As can be seen, the weights given to reports by two agents to each other are the same since we use the same constant (k_1, k_2, k_3) to scale the intensity of relationships. This assumes that the importance given to each agent's report is the same since they may be considered equally trustworthy or efficient. However, some situations may arise where this is not the case and we aim to study these in future work.

One particular model that studies how such combinations of relationships influence trust valuations of agents is the ReGreT model [11]. ReGreT uses fuzzy sets to capture the intensity of relationships that might exist between agents or groups of agents in order to elicit quantitative measures of trustworthiness of agents or the ratings they may provide. For example, in ReGreT if a witness α is in a 'strong' *Coll* relationship with β , then agent γ who is strongly dependent on α in *Dependency* might value ratings of α about β very 'low', where 'strong' and 'low' represent fuzzy sets characterising the intensity of relationships.

Similarly, and also by means of fuzzy rules, the ReGreT system uses relationships together with direct experiences in order to assign trust values to other agents. In other words, it uses prejudice as a mechanism for evaluation. For example, if an agent α usually offers good quality resources and agent β has a 'strong' *Coll* relation with it, the system will assign to agent B a high trust value associated to the quality of the products (see [11] for more details).

In this sense, trust and reputation systems like ReGreT that rely on relationships to improve the computation of reputation and trust values can take advantage of the work presented in this paper.

6. CONCLUSIONS AND FURTHER WORK

In this paper, we have presented a novel process for identifying relationships between agents in an electronic marketplace and discussed how this information can be used to reason about the trustworthiness of agents. By doing this, we address a shortcoming of existing trust models, since they typically do not consider such relationships and where they attempt to incorporate them in a trust model they provide no mechanisms for *identifying* them automatically and do not enable sophisticated reasoning about the complex range of scenarios that may arise. Furthermore, we have discussed how the work presented here can be used directly within existing trust models, such as ReGreT, to supplement its existing approach. By combining existing models with this work we further the development of robust trust valuation models in a constructive and immediate manner.

In the future, we aim to deal with more complex combinations of relationships (i.e. more than 3 agents in all relationships) and explore ways of analysing such combinations of relationships using learning or case-based reasoning tools.

Acknowledgment

Jordi Sabater enjoys a Sixth Framework Programme Marie Curie Intra-European fellowship, contract No. MEIF-CT-2003-500573. This research was also funded by the EPSRC Semantic Firewall project (ref. GR/S45744/01).

7. REFERENCES

- [1] R. Ashri, M. Luck, and M. d'Inverno. On Identifying and Managing Relationships in Multi-Agent Systems. In George Gottlob and Toby Walsh, editors, *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 743–748. Morgan Kaufmann Publishers, 2003.
- [2] R. Ashri, M. Luck, and M. d'Inverno. Identifying opportunities and constraints for goal achievement through relationship analysis. In N. R. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Proc. 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1452–1453. ACM, 2004.
- [3] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [4] P. Dasgupta. Trust as a commodity. In D. Gambetta, editor, *Trust: Making and Breaking Cooperative Relations*, pages 49–72. Blackwell, 1998.
- [5] M. d'Inverno, D. Kinny, M. Luck, and M. Wooldridge. A Formal Specification of dMARS. In M. P. Singh, A. S. Rao, and M. Wooldridge, editors, *Intelligent Agents IV*, volume 1365 of *LNCS*, pages 155–176. Springer, 1996.
- [6] M. d'Inverno and M. Luck. *Understanding Agent Systems*. Springer, 2nd edition, 2004.
- [7] M. Esteva, D. de la Cruz, and C. Sierra. ISLANDER: an electronic institutions editor. In *The First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1045–1052. ACM Press, 2002.
- [8] D. L. McGuinness and F. van Harmelen. OWL Web Ontology Language: Overview. <http://www.w3.org/TR/2003/PR-owl-features-20031215/>.
- [9] S. D. Ramchurn, D. Huynh, and N. R. Jennings. Devising a trust model using confidence and reputation. *The Knowledge Engineering Review*, 19(1):1–25, 2004.
- [10] S. D. Ramchurn, C. Sierra, L. Godo, and N. R. Jennings. Devising a trust model using confidence and reputation. *Applied Artificial Intelligence*, 18(9-10):833–852, 2004.
- [11] J. Sabater. *Trust and reputation for agent societies*. PhD thesis, Universitat Autònoma de Barcelona (UAB), 2003.
- [12] J. Sabater and C. Sierra. REGRET: a reputation model for gregarious societies. In C. Castelfranchi and L. Johnson, editors, *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 475–482, 2002.
- [13] M. Schillo, M. Rovatsos, and P. Funk. Using trust for detecting deceitful agents in artificial societies. *Applied Artificial Intelligence*, 14(8):825–848, 2000.
- [14] J.M. Spivey. *The Z Notation*. Prentice Hall, 2nd edition, 1992.