2009 International Conference on Advanced Information Networking and Applications Workshops

A Property based Framework for Trust and Reputation in Mobile Computing

Kevin Jones, Helge Janicke, Antonio Cau
De Montfort University
Leicester
UK
{kij, heljanic, acau}@dmu.ac.uk

Abstract

In this paper we reinforce the benefits of a trust based approach in mobile multi-agent systems. Whilst we do not specify a trust model we provide a framework for property based trust, defining the observations which underpin many pre-existing trust models. This framework maintains subjectivity which should be the basis of any trust based system yet enables entities to communicate and collaborate trust information regardless of the trust model used for deliberation. We also introduce 'Trust Communities' as a representation of property based trust information thus enabling the collaboration of reputation information.

I. Introduction

In recent years there has grown a desire to understand the social aspects of trust and apply these within the sphere of computer science. Elements of this exist in HCI (Human Computer Interaction) [1] in determining the requirements of a user, such that they trust the system. However, the introduction of distributed service-oriented computing architectures [2], [3] has given rise to a new adaption of trust thus, an agent (human or computational) is able to reason about the uncertainty of using a particular service. The service oriented architecture is central to the idea of web services [4], [5], pervasive computing [6], grid computing [7] and multi-agent systems [8]. This understanding highlights the importance and application of trust based systems in a number of application domains.

Mobile software agents are autonomous entities capable of migrating from one host to another, continuing their execution within their new environment. As the execution is performed within a potentially untrusted environment, the requirement for security is critical.

There is a research history of using trusted-third parties, tamper-proof hardware and cryptographic techniques as solutions towards ensuring a more secure environment within which agents can operate. We propose a framework that uses trust for deliberation and decision making based on knowledge of previous behaviour. As such we do not purport to eradicate malicious behaviour by hosts but rather discourage it based upon reputation models and endower agents with the ability, based around a notion of trust, such that it can avoid exposure to potentially malicious entities.

The use of trust in order to achieve this minimises and standardises the complex calculations and interactions required to make an informed decision about the reliability of an entity. Agents cooperate in order to increase the data used in trust deliberation. Whilst the notion of using trust models is not new, we will take the view that in an open system where trusted hardware is not guaranteed, the assumption that all entities will use the same trust model is too strong.

In the case where multiple trust models and weighting factors are used within the same system it becomes difficult to collaborate given no standardisation of communicated trust. In this paper we look at the use of properties enabling agents with different and subjective trust models to compute trust whilst maintaining collaboration and reputation information.

We also provide the notion of communities using reputation information so that the available trust data for decision making is increased. Subjectivity at this level is not with the information itself as it is a collective opinion however, the encapsulation of this information into the trust model and therefore subsequent decision making is done by individual agents thus, remains subjective.

In the remainder of this paper we provide a suitable background to our work in Section II by first looking at related work. In Section III we present the notion of property (observations) as the underpinning of trust. This

978-0-7695-3639-2/09 $25.00 © 2009 IEEE
DOI 10.1109/WAINA.2009.204

is extended in Section IV utilising collaboration in order to incorporate indirect trust information and continued in Section V with the introduction of trust communities and reputation. Finally we conclude in Section VI.

II. Related Work

Trust is often considered as a human, social behaviour that is innate to most and forms part of every day life [9], [10]. No surprise then that early work on trust surrounded the study of interactions between humans and its importance within such interaction [11]. As trust appears to be innate this leads to many varied perceptions over what constitutes trust. Indeed, there is no common consensus as to a single definition of trust but the importance of trust is stressed by Luhmann [12] in the suggestion that trust 'is a basic fact of human life' and that it 'arises from our inherent inability to handle complexity'. A minimal understanding is provided by Guinnane [13] in that trust is a 'three-part relationship involving at least two actors and one act'.

Deutsch [9] talks of 'probability' and suggests an unknown element to an action, given the results of the outcome are contingent on the behaviour of another. The definition is further explored in the thesis of Teacy [14] who continues to explain what he believes to be the 5 key points of this definition:

1) Trust occurs between pairs of entities - trust is the assessment of one entity (trustee), from the perspective of another entity (trustor).
2) Trust relates to a particular action - Although sometimes we talk generally about our trust in an individual, a high level of trust in someone to perform one type of action does not imply a high level of trust in them to perform another.
3) Trust is a subjective probability - Trust is subjective, because it is assessed from the unique perspective of the trustor. It is dependent both on the individual set of evidence available to the trustor and her relationship with the trustee.
4) Trust is defined to exist before the respective action can be monitored - Trust is a prior belief about an entity's actions. It is an assessment made in a context of uncertainty. Once the trustor knows the outcome of an action, she no longer needs to assess trust in relation to that outcome.
5) Trust is situated in a context in which it affects the trustor's own action - By this, we mean that our interest is limited only to those actions of a trustee that have relevance to the trustor.

In order to compute trust, there must first be a formal model by which trust is measured, calculated, and communicated. A number of these models exist including work by Marsh [15] who provides a heuristic formalism to the problem of trust and cooperation which is extended to provide temporal considerations. Carbone et. al. [16], [17] by contrast opt for a policy approach to trust management using a lattice of *high, medium, low, ask, and unknown* in addition to information ordering and trust ordering techniques. This technique however, does not consider the temporal nature or trust or the autonomy of dynamic trust management approach of agent systems.

McDonald and Yasinsac [18], [19] also model the relationships between entities and are even inclusive of principles such as *code developer, application owner* and *host manager*. Such principles are then associated with a trust relationship and enforced through security mechanisms based on the CIA triad for security and through tamper-proof hardware.

More recent work includes that of Ramchurn et. al. [20] in the formalisation of trust as confidence and reputation for the selection of interaction partners and that of Carbone et. al [17] providing a calculus for trust management. Abdul-Rahman and Hailes [21] model trust relationships and provide a protocol by which trust recommendations can be communicated.

Perhaps the closest related work is that of Teacy et. al. [22] with the TRAVOS model and of Derbas et. al. [23] with TRUMMAR model. These both assume an open mobile environment without the presence of trusted third parties. Trust is then propagated through their system as a series of collaborations.

With each of these models however, the collaboration and communication of trust assumes that all entities adopt them as a single model and indeed, utilise them in exactly the same way w.r.t weighting measurements. This for us appears too ridged a restriction and search for a method by which these models are still adopted but are interoperable and used more subjectively.

Both TRUMMAR and TRAVOS deal specifically with the problems of mobile agent security and adopt a trust based approach. In their approach however, the trust models used are fixed such that agents can deliberate over trust dependant upon the number of observations available describing previous behaviour of an entity. Thus, subjectivity in this manner directly relates subjective opinion with the observations available.

The adoption of property based trust however, enables agents to be truly subjective not only based on the different observations available but also with the use of completely different trust models and weighting factors. For us this is intuitive with the mapping of human concepts of trust into agent systems.

III. Property Based Trust

Collaborative entities such as agents interact with and observe their environment and others within it. Observations may be rudimentary such as the identification of entities within the environment or more complex such as the outcome of an interaction. It is these elements of observations that we denote to be properties. An observation consists of many properties which change over time and/or in response to an action, these changes are measurable.

Agents provide services and encapsulate series of actions as behaviours, these enable an agent to act within the environment. The deliberation of an agent determines which behaviours are executed. As behaviours are not necessarily transparent, especially when performing a black-box style service not all actions of that behaviour are visible. In this instance observations are made over the entire behaviour as opposed to every action comprising that behaviour.

We believe that to discuss trust in a meaningful manner, we must first consider carefully the observations that are made as the basis for the computation of trust values. Whilst many existing trust models assume observations we believe that it is the properties of these observations that are implicitly used.

We also consider interactions to be temporal, thus a property can change over time as part of an interaction. A behavioural observation is therefore the measurement of one or more pre- and post-observations w.r.t a property over the course of a given time period. Segregation of such a time period spans the duration of a behaviour or potentially an entire interaction.

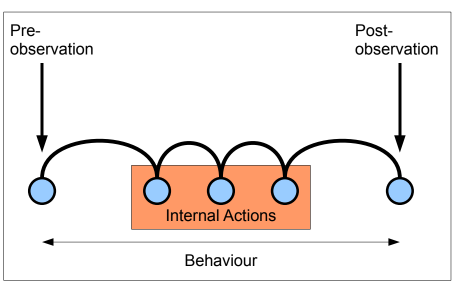

Fig. 1. Observation points of a behaviour

The alternative to this is to classify observations for each attribute change. With this approach however, the sequence of interactions considered within a behaviour is not easily transparent. Figure 1 shows the pre- and post-observation points from the perspective of an observer with the view of a service consumer therefore, a number of actions performed by the provider as part of the behaviour (Internal Actions) are not visible.

We predefine properties such that a given behaviour is observed in accordance with one or more properties of that behaviour. Properties are (semi-) subjective such that each agent computes trust from its own selection of properties. We use semi-subjective as in practise there is a finite number of observations that can be made, the subjectivity evolves from selection.

Examples of measurable properties within an observation include, time, interaction partners (service provider and service consumer), response time, liveliness, security mechanisms, etc. At this abstraction level we do not specify the number of properties an observation may encapsulate as this remains with system developers. There are however, a number of standard properties that must be observed for identification and processing purposes including; descriptor, observer ID, observed ID, and timestamps.

If we are to map a human concept of trust into autonomous agents we must denote trust to be highly subjective and action dependant thus, multiple behaviours are never comparatively equal. Time, observer, or observed i.e. those involved in the interaction are variant. This is however, advantageous as in composing observations from a series of observations these differences are considered.

What is needed therefore is a mechanism to manage and compute trust from these observations whilst maintaining the subjectivity with which it is so closely linked.

IV. Trust Collaboration

Agents are autonomous entities, which deliberate for themselves over their behaviour in relation to the achievement of a goal. They are able to sense the environment and communicate with others to form collaborations. Within such collaborations, agents interact providing each other with services. As stated, agents make observations over properties of these interactions.

To use trust in decision making, observations of a number of previous interactions are required. It is fair to assume that a single entity may have limited interactions with others in the system such that gathering enough observations to make useful computations of trust is not possible. Therefore, direct observations alone do not provide enough information about the environment and the entities within it.

To improve on this there is a need for agents to collaborate w.r.t trust information. Such an approach is not uncommon and indeed much work on trust utilises the notions of Direct (self observed) trust, Indirect (recommended) trust, and Reputation based information. This quickly increases an agent's knowledge of the system and each layer of trust information can be weighted accordingly in the calculation of a trust value.

As we see trust (and the computation of trust values) as subjective, the communication of trust becomes difficult. The assumption that all entities within the system are utilising the same trust model or indeed, are using the same model in the same way to deliberate behaviour does not hold. This is intuitive with the way we as humans quantify trust, some are optimistic, some pessimistic, and each have different concepts of what constitutes trustworthy and untrustworthy.

It is not beneficial to discuss the communication of the computed trust value itself (although we do not rule this out if it is predetermined that the same trust model, weighting mechanisms, and deliberation approach is used between agents). Collaboration of trust information is therefore the communication of observations either de-facto or as an aggregation of a number of observations. Such indirect observations classified as indirect trust are then assimilated into the available observations upon which to base a trust decision. Again, this is intuitive to the manner in which humans utilise trust such that we ask our associates their opinion of others.

This concept is also beneficial in gathering reputation information as this is the opinion of many different entities within the system about a single entity. Thus, it is a set of observations provided by different entities. In this instance, using the observations approach patterns of behaviours becomes visible. A real world example of such a concept would be user feedback on auctioning or retail websites such that multiple users leave feedback based upon a number of criteria.

We also introduce a further descriptor of trust information utilising Trust Communities in which groups of associated entities are collated thus, enabling the introduction of trust between groups of entities as opposed to individual entities.

V. Trust Communities

We introduce the notion of trust communities in order to further extend the information available to entities that use trust for decision making and to reduce the workload required to achieve this. Communities are formed when a number of entities share the same property value, these communities can then be composed into further communities if these properties can be composed. This can be seen in Figure 2 such that entities form a community and thus, are classified together in relation to a property. These properties are those which we can observe and therefore, used as the basis for trust information.

We distinguish two types of trust community:

- **Perceived Community:** is calculated internally by an entity in order to categorise the trust information it possesses. This provides a mechanism by which to quickly look up the community of an entity when necessary rather than recalculating information
- **Reputation Community:** operates at a system level such that all entities operate with a single communities view thus, enabling community level trust whereby an entity becomes more or less trustworthy dependant upon its associates w.r.t a property.

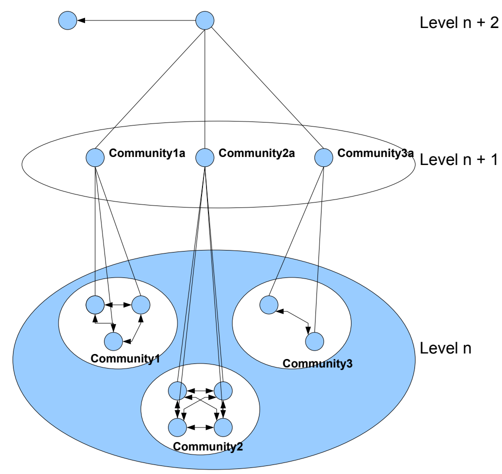

Fig. 2. Trust Community View

We foresee that the perceived community view provides a mechanism by which to place entities in sets dependant upon observations regarding their Quality of Service (QoS). This adds to the understanding of trust as it enables any entity to make similarities between those sharing properties. An example of this would be to use response time as the property by which communities are composed. Entities are placed in communities given their previously observed response time of $< 0.4m/s$, $< 0.6m/s$ or $> 0.6m/s$.

We utilise basic set theory to compose communities such that it is possible to denote communities as being composed using 'OR' $\vee$ (seen as a rectangle in Figure 3) or 'AND' $\wedge$ (seen as oval). To do so is more intuitive when considering the composition of multiple communities based upon different properties. This can be seen in Figure 3 whereby membership to the communities is based upon property x, property y, and property z. These communities are composed such that it is possible to describe a community consisting of entities in any of the compositional communities (in the figure as $x \vee y$) or only entities that meet properties of both communities (in the figure shown as $(x \vee y) \wedge z$).

As perceived communities are updated dynamically as observations and recommendations are introduced there is a partial updating of agents knowledge and re-calculation of the entire trust model is not required. Using a representation of other entities in communities in this manner also enables a mobile agent to migrate with merely the representation of it knowledge rather that the cumbersome and extensive observations that it describes thus, an agent can still make trust decisions without direct access to its observations.

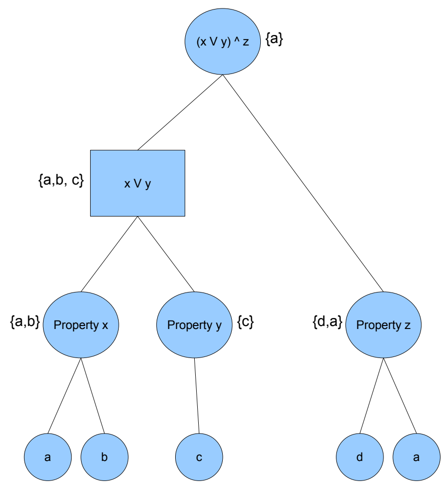

Fig. 3. Community for Response Time QoS

By contrast *reputation communities* are more centralised and are shared throughout the system. They are based upon more static (although not completely static) properties such as owner, developer, role, location, trust model, etc. They also encompass the observations of many entities. Those within a community share not only a similar property but are bound by the rules of that community thus, to be part of a community of Doctors for example the requirement for participation in communities of PhD or MD can be used. As such it becomes possible to compute trust on a researcher entity based upon membership to the PhD community.

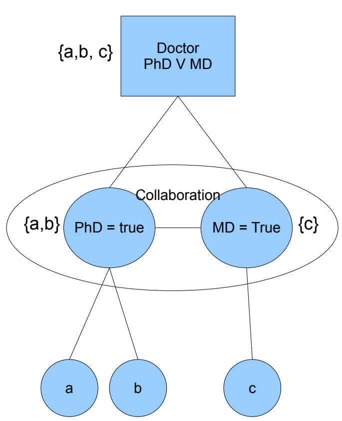

Fig. 4. Reputation Community for Doctorate

Reputation communities are self governing such that agreement must be met to enable admission or revocation of entities to the community. In addition one entity acts as a broker for the community thus, is elected to represent the community as a whole. From this it is possible to implement community level trust whereby communities communicate their reputation information to others via the broker.

We can see this is a very powerful method by which to improve trust based decision making and is intuitive to the way humans look to increase their trust in others using trust by association and certification, the pre-requisites for communities.

To take a rather current real-world example; a community composed of entities meeting the property classification of a bank operate and collaborate with each other. There is a trust between banks within the community as they can loan money to each other. Banks have customers (borrowers and savers) whom form a community themselves, firstly as a community of customers of one specific bank and composed as customers of all banks. Banks trust their borrowing customers that they will pay back what is owed and their savings customers that they will keep savings in their account. Savings customers trust their bank to keep their money safe and borrowing customers have trust in their banks to keep repayments at a reasonable rate. This is therefore a bi-directional trust relationship between the two communities.

The interesting trust relationships within this scenario have lead to the breakdown of trust between banks themselves in the community and between banks and borrowing customers. Therefore a reduction of trust in the complete banks community occurs leading to savings customers trusting less and a cyclic spiral of trust occurs.

If we are to introduce an additional community known as a government, with the responsibility of increasing the level of trust in other community level relationships. The objective for this is to enable the entities within the government community to perform their task (and as a side effect increase the level of trust of customers' community in it).

The notion here is that as humans we don't always consider every entity separately but rather as a community. We feel that this is a useful addition to the work on trust and the introduction of trust information between groups of entities.

As with perceived communities the calculation of trust remains with an individual agent and remains subjective as we continue to work with observation information and not specific trust models. However, communities provide additional information by which to make the computation.

VI. Conclusions and Future Work

We have presented a framework using observation properties as the underpinning of existing trust models and therefore increasing collaboration and communication of trust within the system given an environment where entities compute trust using different models or weighting measures within these models. It is these observations, or aggregated observations that are communicated between entities in order to provide recommended trust information. Thus, we avoid the requirement for translation between trust models and maintain the subjectivity of decision making with the agent itself.

We also present an extension to this approach using trust communities which we believe provides greater scope of information to the agent upon which to base its trust decision. In the case of mobile agents the community representation also enables an agent resident upon a remote host and without access to large observation data stores to make trust decisions based upon its last known information as shown in the community representation.

For reputation information the use of reputation communities also enables agents to act as a collective in respect to trust. This is a powerful notion as it not only enables entities to collaboratively decide who is considered part of their community but also to collaborate as a community with others.

We are already utilising the notion of observation properties in a prototype for architecture comparisons and have found communication of observations thus, collaboration of trust to be feasible. We do wish to compare this solution with the effectiveness of a system incorporating multiple agents and differing trust models whereby collaboration can only occur between agents with the same model.

Further future work includes the incorporate communities into our existing architecture implementation for trust deliberation and provide empirical evidence as to the effects the additional information has on the selection process of an agent in avoiding malicious behaviour towards it.

References

[1] S. van Mulken, E. André, and J. Müller, "An empirical study on the trustworthiness of life-like interface agents," in *HCI (2)*, H.-J. Bullinger and J. Ziegler, Eds. Lawrence Erlbaum, 1999, pp. 152–156.

[2] M. N. Huhns, M. P. Singh, M. H. Burstein, K. S. Decker, E. H. Durfee, T. W. Finin, L. Gasser, H. J. Goradia, N. R. Jennings, K. Lakkaraju, H. Nakashima, H. V. D. Parunak, J. S. Rosenschein, A. Ruvinsky, G. Sukthankar, S. Swarup, K. P. Sycara, M. Tambe, T. Wagner, and R. L. Z. Gutierrez, "Research directions for service-oriented multiagent systems," *IEEE Internet Computing 9*, vol. 6, pp. 65–70, 2005.

[3] T. Eymann, F. Klgl, W. Lamersdorf, M. Klusch, and M. N. Huhns, "Multiagent system technologies," in *In proc. Third German Conference, MATES 2005, Koblenz, Germany*. Springer, September 2005.

[4] W. T. Tsai, M. Malek, Y. Chen, and F. Bastani, "Perspectives on service-oriented computing and service-oriented system engineering," in *SOSE '06: Proceedings of the Second IEEE International Symposium on Service-Oriented System Engineering*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 3–10.

[5] S. A. McIlraith, T. C. Son, and H. Zeng, "Semantic web services," *IEEE Intelligent Systems*, vol. 16, no. 2, pp. 46–53, 2001.

[6] L. T. Yang, J. Ma, M. Takizawa, and T. K. Shih, Eds., *Proceedings of the 2005 International Conference on Pervasive Systems and Computing, PSC 2005, Las Vegas, Nevada, June 27-30, 2005*. CSREA Press, 2005.

[7] I. T. Foster, "A new era in computing: Moving services onto grid," in *ISPDC*. IEEE Computer Society, 2005, p. 3.

[8] N. R. Jennings, "On agent-based software engineering," *Artif. Intell.*, vol. 117, no. 2, pp. 277–296, 2000.

[9] M. Deutsch, "Behavior and heredity: Statement by the society for the psychological study of social issues," *American Psychologist*, vol. 28, pp. 620–621, 1973.

[10] A. Baier, "Trust and antitrust," *Ethics*, vol. 96, p. 231260, 1985.

[11] D. Hume, *Enquiries Concerning The Human Understanding and Concerning The Principles of Morals (1737)*. Oxford University Press 2nd Ed, 1957.

[12] N. Luhmann, "Trust and power," Chichester: John Wiley, 1979.

[13] T. W. Guinnane, "Trust: A concept too many," Working Papers 907, Economic Growth Center, Yale University, 2005.

[14] W. T. L. Teacy, "Agent-based trust and reputation in the context of inaccurate information sources," Ph.D. dissertation, University of Southampton, 2006.

[15] S. Marsh and M. R. Dibben, "Trust, untrust, distrust and mistrust - an exploration of the dark(er) side," in *iTrust*, ser. Lecture Notes in Computer Science, P. Herrmann, V. Issarny, and S. Shiu, Eds., vol. 3477. Springer, 2005, pp. 17–33.

[16] M. Carbone, M. Nielsen, and V. Sassone, "A formal model for trust in dynamic networks," in *SEFM*. IEEE Computer Society, 2003, pp. 54–.

[17] ——, "A calculus for trust management," in *FSTTCS*, ser. Lecture Notes in Computer Science, K. Lodaya and M. Mahajan, Eds., vol. 3328. Springer, 2004, pp. 161–173.

[18] J. McDonald, A. Yasinsac, and W. Thompson, "Trust in mobile agent systems," Florida State University, Tech. Rep., 2005.

[19] J. T. McDonald and A. Yasinsac, "Application security models for mobile agent systems," *Electr. Notes Theor. Comput. Sci.*, vol. 157, no. 3, pp. 43–59, 2006.

[20] S. D. Ramchurn, N. R. Jennings, C. Sierra, and L. Godo, "Devising a trust model for multi-agent interactions using confidence and reputation," *Applied Artificial Intelligence*, vol. 18, no. 9-10, pp. 833–852, 2004.

[21] A. Abdul-Rahman and S. Hailes, "Supporting trust in virtual communities," in *HICSS*, 2000.

[22] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck, "Travos: Trust and reputation in the context of inaccurate information sources," *Autonomous Agents and Multi-Agent Systems*, vol. 12, no. 2, pp. 183–198, 2006.

[23] G. Derbas, A. I. Kayssi, H. Artail, and A. Chehab, "Trummar - a trust model for mobile agent systems based on reputation," in *ICPS*. IEEE Computer Society, 2004, pp. 113–120.