Logics, Logic Programming, as well as F-Logic as starting points for the development of a number of WSML language variants: WSML-Core, WSML-DL, WSML-Flight, WSML-Rule, and WSML-Full. The WSML variants differ in logical expressiveness and in the underlying language paradigms. They allow users to make the trade-off between the degree of expressiveness and the implied complexity on a per-application basis. The WSML has two alternative layerings:

- WSML-Core → WSML-DL → WSML-Full, and
- WSML-Core → WSML-Flight → WSML-Rule → WSML-Full.

For both layerings, WSML-Core and WSML-Full denote the least and most expressive layers respectively. The two layerings are disjoint to a certain extent in the sense that inter-operation between the Description Logic variant (WSML-DL) and the Logic Programming variants (WSML-Flight and WSML-Rule) is only possible through a common core (WSML-Core) or through a very expressive superset (WSML-Full).

The WSML, can be seen as a testing ground for the development of formal techniques for Web service description.

Web Service Modeling Execution Environment (WSMX)

The Web service Execution Environment (WSMX) [4] is an execution environment which enables discovery, selection, mediation, and invocation of Semantic Web services described according to the philosophy of WSMO. It is thus a WSMO reference implementation. The WSMX provides a tangible testbed for the WSMO in order to prove its viability as a means to achieve dynamic interoperability between Semantic Web services.

In short, WSMX functionality can be summarized as performing discovery, mediation, selection and invocation of Web services on receiving a user goal specified in WSML. The user goal is first matched against the formal descriptions of Web services registered with the WSMX. If successful, one or more service descriptions (ranked according to user preference) can be returned. The most appropriate service as selected by the user is then invoked and the result returned to user. Prior to the invocation step, the WSMX ensures that the data provided for the service invocation is in the format that the Web service expects. If necessary, a data mediation process is performed to ensure inter-operability between different entities. Presently, the WSMX architecture relies on a set of loosely-coupled main components that provide functionality for each step of the Web service usage process: discovery, selection, mediation and invocation.

Final Remarks

Semantic Web services, by combining recent Web-related trends, constitute one of the most promising research directions to improve the integration of applications within and across enterprise boundaries. In this context, the WSMO aims to provide the conceptual and technical means to realize Semantic Web services, improving the cost-effectiveness, scalability and robustness of current solutions. The WSML provides a formal syntax and semantics for the WSMO by offering variants based on different logics in order to provide several levels of logical expressiveness and thus allowing tradeoffs between expressivity and computability). Finally, the WSMX provides a reference implementation for the WSMO and the interoperation of Semantic Web services.

References

[1] http://www.wsmo.org/

[2] http://www.sdk-cluster.org/

[3] http://www.wsmo.org/wsml

[4] http://www.wsmx.org/

Trust: Challenges and Opportunities

Nathan Griffiths
University of Warwick
United Kingdom
nathan@dcs.warwick.ac.uk

Introduction

Trust is fundamental in distributed systems where individual components interact to achieve some overall objective. In small-scale or closed systems this trust can be implicit, imbued to the individual components and the system overall by its designers and implementers. In open or large-scale systems however, it is becoming increasingly common for trust to be explicitly represented and reasoned about by the components, or agents, in the system. In recent years trust has become a hot research topic, with numerous conferences and workshops attracting both academic researchers and industrial representatives from solutions providers in areas as diverse as telecoms, logistics and e-business. This article gives a brief overview of the alternative approaches to trust and attempts to identify some of the important research questions.

The current interest in trust creates an active environment for trust researchers and practitioners. However, it also raises some challenges. Trust research has parallels with agent research a decade ago — it is an exciting area of clear value, but there is a risk that debate about definitions and mechanisms might add confusion and delay widespread adoption. Just as there was (is?) no consensus definition of agents, there is similar debate over trust. For example, how does trust relate to reputation? Is trust an individual (experience-based) notion, or should it encompass others' (potentially subjective) views? Are trusted agents secure and reliable, or do they simply have "good" intentions? These questions are important, but it is equally important to ensure that confusion is avoided, and that as a community we have a clear overall view.

Psychological and Cognitive Approaches

Many trust models take a cognitive view of agents and trust, typically relying on folk psychology notions such as belief and desire. Agents trust others with respect to some activity or the performance of some task, and consider trustworthiness according to beliefs about such aspects as competence, disposition, willingness, dependence, and fulfilment [1]. The level of trust is determined by these beliefs, along with past experiences and possible recommendations from others. Some cognitive approaches also consider modelling the desires/motivations of other agents, and incorporate this into assessing trustworthiness. These approaches give a powerful mechanism for reasoning about interactions and the trust, power and dependence relationships between agents. However, it can be difficult to translate their richness into a practical implemented system, since the data structures and the functions needed to manipulate them are expensive to maintain.

Numerical Approaches

Numerical approaches are perhaps the most commonly used with numerous proposed mechanisms, in which agents represent the trustworthiness of others in numeric intervals, typically [-1, +1] or [0, 1]. The lower bound corresponds to complete distrust and the upper bound to blind trust. Agents either keep track of a numerical value that is updated by some function after each interaction, or are equipped with a function to transform a history of interactions into a numerical value. Some approaches decompose trust into separate values for different situations, or according to the different dimensions of an interaction (such as cost, quality, timeliness). Numerical methods tend to use a trust threshold and only when trust is above the threshold will cooperation take place. The main advantage of numerical approaches is their simplicity, since it is relatively inexpensive to incorporate trust into an agent's decision making. This simplicity, however, is also their disadvantage. Firstly, numerical approaches do not provide the richness of reasoning available with other techniques. Secondly, there is limited meaning to the values themselves, which encumbers the sharing of information between agents and external reasoning about the system.

> *Trust research has parallels with agent research a decade ago — it is an exciting area of clear value, but there is a risk that debate about definitions and mechanisms might add confusion and delay widespread adoption. Just as there was (is?) no consensus definition of agents, there is similar debate over trust.*

Probabilistic Methods

Probabilistic methods are a subset of numerical approaches in which trust is represented in the interval [0,1]. However, this number represents a probability and has a clearer semantics associated with it. There are many probabilistic approaches ranging from those based on simple objective and subjective probabilities, to those using more complex Bayesian probability distributions. Decisions are made in a similar manner to numerical approaches by use of trust thresholds and maximising the probability of success.

Reputation Systems

Many reputation-based approaches have been proposed, ranging from centralised systems that aggregate feedback (*à la* eBay), through decentralised feedback systems, to numerical and probabilistic approaches that are augmented with recommendations from other agents. Reputation systems generally use a combination of direct experience, recommendations, and knowledge of the social structure of the system to represent and reason about trust. The incorporation of reputation greatly enhances the richness of a trust model, but typically increases complexity, and opens up questions about issues such as collusion.

Certificates and Keys

Where trust is viewed as a mechanism for ensuring security, it tends to be achieved via protocols, certificates or keys. Some approaches define detailed interaction protocols that ensure detection of any deviation from expected behaviour, and define the actions or sanctions that should be taken in such cases. Other approaches use trusted third parties to provide verification and authentication, and to act as intermediaries in interactions. However, the most common security-oriented approach is to use certificates and keys: a certification authority issues a certificate verifying that an agent's public key is owned by that agent. These public keys can then be used to sign and encrypt data to ensure authentication and privacy. Certificates and keys provide a powerful mechanism for achieving security, but it can be difficult to combine them with more general social- or service-oriented approaches.

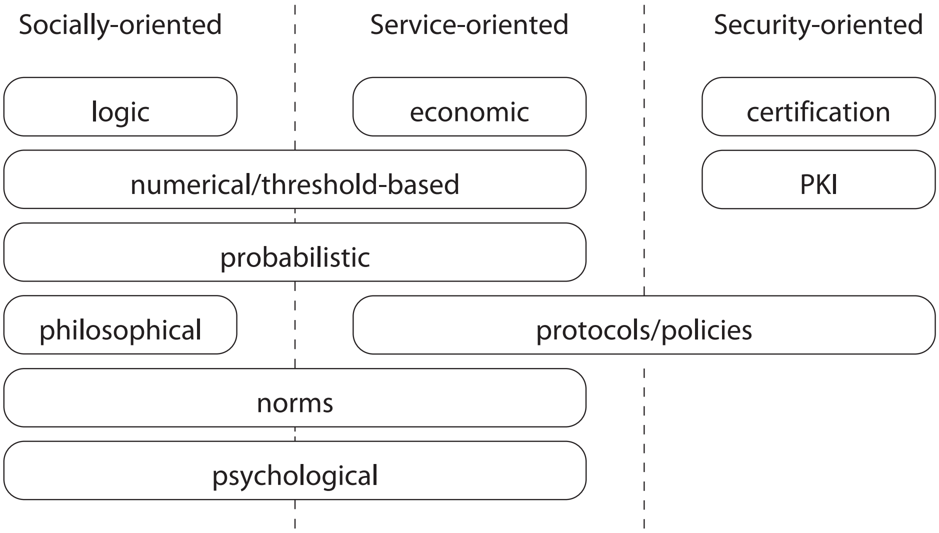

Figure 1. Selected techniques and mechanisms for trust.

Challenges for Trust Researchers

Computational trust is a young and active research area and numerous techniques have been proposed, as introduced above. However, there are a number of open research challenges. Since there is no overarching view of trust it is difficult for implementers to select a trust mechanism for a given environment. There is a need for researchers to frame their work within the context of the general trust landscape to enable simple comparisons between approaches. A trust taxonomy would begin to address this problem. Furthermore, trust should be seen as a fundamental component of any multi-agent system. Again, a taxonomy of trust would assist in enabling this, especially if aided by the provision of suitable software tools to assist in incorporating trust into agents. Finally, much trust research takes place in simulated semi-closed environments, and a number of issues must be addressed to enable trust mechanisms to be effective in real-world domains. We therefore propose three primary challenges for trust researchers.

A Trust Taxonomy A taxonomy of trust would be valuable, firstly to aid implementers in selecting appropriate trust techniques for a particular context and, secondly, to assist researchers in positioning their work, and comparing it to other approaches. The division of trust literature into socially-, service- and security-oriented trust might a starting point, but perhaps what should be aimed for is a pattern library of trust techniques. Ideally, an implementer should be able to select a trust approach easily, based on the characteristics of the domain.

Agent Platforms Trust is fundamental to multi-agent systems. However, many existing agent platforms do not incorporate trust by default. Trust should be viewed as a fundamental component of any agent platform, and implementers should be able to select an appropriate "pattern" for a given domain. Existing platforms need to be augmented with "trust wrappers", while new platforms should include placeholders for trust mechanisms by default.

Coping with the "Real-World" The real-world is a complex place, with interactions potentially failing in numerous ways for a variety of reasons, including malicious motivations of cooperative partners, competition between agents, and unavoidable environmental change. Current trust models are typically relatively simplistic in updating trust after interactions. If trust is to be effectively applied in real-world applications, then there must be a mechanism to distinguish between intentional and unintentional failure. Furthermore, as Marsh proposes, there must be a distinction between trust, mistrust and distrust, i.e. trust, misplaced and incorrect trust, and explicit distrust (c.f. information, misinformation and disinformation) [2].

Summary

Trust is a rich notion and an area of active research. In this article we have tried to give a flavour of the breadth of trust research. As in any young research area there are a number of open challenges; this article has presented a personal view in identifying three areas that the author sees as meriting immediate attention. Clearly, there are numerous other open questions, but the author hopes that a focus on these challenges will help facilitate more widespread use of trust in agent-based systems.

[1] C. Castelfranchi. Trust mediation in knowledge management and sharing. In C. Jensen, S. Poslad, and T. Dimitrakos, editors, Proceedings of the Second International Conference on Trust Management (iTrust 2004), pages 304–318, 2004.

[2] S. Marsh and M. R. Dibben. Trust, untrust, distrust and mistrust — an exploration of the dark(er) side. In P. Herrman, V. Issarny, and S. Shiu, editors, Proceedings of the Third International Conference on Trust Management (iTrust 2005), pages 17–33. Springer-Verlag, 2005.

[3] S. D. Ramchurn, D. Huynh, and N. R. Jennings. Trust in multi-agent systems. Knowledge Engineering Review, 19(1):1–25, 2004.

Diversity of Trust Research

There is a growing corpus of literature on trust, a detailed overview of which can be found in the further reading identified below. There is no overarching taxonomy of trust research, and the wide applicability of trust gives rise to a wide range of approaches. However, we can broadly divide these approaches into three areas.

Socially-Oriented Trust - Typically influenced by social science, psychology or philosophy, socially-oriented trust is viewed as a social notion for modelling and reasoning about the relationships between agents. Socially-oriented trust often considers issues such as the motivations of agents, and the power and dependence relationships between them.

Service-Oriented Trust - Taking a pragmatic view, service-oriented trust is a mechanism for achieving, maintaining, and reasoning about quality of services and interactions. Agents typically maintain their own trust information about others, possibly incorporating the recommendations of others, and use this to inform their decision making processes.

Security-Oriented Trust - Taking the view of trust as a mechanism for ensuring security, encompassing issues of authentication, authorisation, access control, privacy, etc. Security-oriented trust also includes work on "trusted computing", i.e. building trusted platforms to ensure privacy and security.

In each of these areas a range of techniques and mechanisms have been proposed, drawing on work in areas as diverse as logic and the social sciences. For any given situation there are generally several alternative candidate trust models/techniques for an implementer to choose from, as illustrated in Figure 1. (Note that this figure is not intended to be exhaustive, and many other techniques exist.)

Trust literature can be further divided according to whether it is concerned with *individual- or system-level trust*. In the former, individual agents model and reason about others, while in the latter agents are forced to be trustworthy by externally imposed regulatory protocols, and mechanisms [3]. Additionally, some trust models are *centralised* and have a single repository of information, while others are *decentralised* with individuals maintaining their own information. Finally, we can distinguish between models concerned with *direct trust* where agents trust others directly based on their experiences, and *recommendation* trust where trust is based on the recommendations of others. For each of these categories of trust there are a number of alternative approaches discussed in the literature, the most common of which are also briefly introduced.

Benchmark simulations for Multi-Agent Learning

Maarten van Someren (University of Amsterdam)

maarten@science.uva.nl

Evaluating methods and systems on publicly available datasets has proved to be a succesful methodology in Machine Learning and in Information Retrieval. Evaluating multi-agent learning systems requires simulation environments, even though this can be more difficult to achieve than for the other areas (e.g. the traffic light simulation, see http://sourceforge.net/projects/stoplicht).

This effort may become part of the new network of excellence KDUbiq which will start in December or January. If you are interested then please contact me by email (maarten@science.uva.nl).