Networking for Multi-agents: Beyond a Local View (Extended Abstract)

Sarah N. Lim Choi Keung and Nathan Griffiths
Department of Computer Science, University of Warwick
Coventry CV4 7AL, United Kingdom
{slck, nathan}@dcs.warwick.ac.uk

ABSTRACT

Decentralised multi-agent systems have empowered individual agents to actively select their course of action to achieve their goals. However, this leaves agents with a limited view of their environment, limiting awareness mainly to their interaction partners. We believe that through interactions and recommendation exchange, agents can build a network of their environment and expand their local view. Agents can extract emergent information from the relationships and interactions they observe. Analysing and understanding such information can improve an agent's efforts in minimising risk associated with the uncertainty of agent interactions, including collusion detection and the discovery of new providers.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents, Multiagent systems*

General Terms

Design, Experimentation, Reliability

Keywords

Agent Systems, Trust, Reputation, Relationships, Collusion

1. INTRODUCTION

Agents in open and dynamic multi-agent systems (MAS) interact with others to reach goals they cannot achieve alone. Agents may, however, be unreliable or dishonest, threatening the success of interactions. Thus, the selection of appropriate partners and an understanding of inter-agent relationships becomes crucial. Trust and reputation have been proposed as partial solutions to the problems caused by uncertainty in interactions. However, trust and reputation only provide a local view of the evaluating agent's environment. Decentralised systems lack the global view that provides complete information for decision making. In this respect, we develop a mechanism that enables agents to extend their environment view beyond a localised one. Agents build and maintain a network of agents, including providers and recommenders, and track how they relate and interact. From this extended view, emergent information about interactions and relationships can be discovered, and used to encourage successful interactions and long-term benefits.

Cite as: Networking for Multi-agents: Beyond a Local View, (Extended Abstract), Sarah N. Lim Choi Keung, Nathan Griffiths, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1315–1316

2. RELATED WORK

Social networks are effective for communicating information, for instance, consumers share product information by word-of-mouth through referrals, and influence purchasing choices [1]. Milgram demonstrates through the *Six Degrees of Separation* concept that referrals are effective in searching large social networks [5]. With dynamism comes the link prediction problem, which concerns the accurate prediction of edges in a network [3]. It has parallels with the discovery of emergent environment information through the agents' overlapping local views, giving a wider perspective of other agents' transactions and social links.

Trust and reputation models [2, 4, 6] propose using trust from direct interactions and reputation from third parties to predict future agent performance. Some models include a social network or agent neighbourhood. In ReGreT [6], a neighbourhood is a group of agents with some common knowledge and neighbourhood reputation reflects the behaviour of the whole group. FIRE uses neighbourhoods to search for relevant witnesses [2]. However, these models do not specify how the network is built and used to understand agent relationships. In our work, the use of social networks is different to multi-agent learning techniques. As building the network is not the main focus, we use a simpler mechanism for agents to acquire environment information, based mainly on their local interactions.

3. AGENT NETWORK BUILDING MODEL

Our proposed network model captures the dynamic behaviour of agents, their interactions and any emergent information. The model has three main components: data collection, network building, and analysis of interaction data.

3.1 Data Collection

The data collection component is based on our previous work on a model of trust and reputation [4]. Agents use trust with direct and indirect recommendations to inform their decision making. An evaluator records summaries of direct interactions and combines these with recommendations to make decisions about the trustworthiness of other agents. We supplement this by recording any recommendations received, and using these to build a network. Consider the representation of a customer agent, a_c, acting as an evaluator. Agent a_c records a partial history of provider interactions, $H_{i_s} = (\mathbb{P}\, i_s, count^+, count^-, ST, ST_c)$, where $i_s = (a_c, a_p, s, t)$ is a service interaction. The provider agent is a_p, s is the service performed at time t, $count^+$ and $count^-$ are the number of positive and negative interactions experienced by a_c respectively. ST is the situational trust in a_p and ST_c is the confidence in the situational trust value. The service s is defined as the service type and a set of dimensions, each defined as: $d = (d_{type}, d_e, d_a)$, where d_{type} is the dimension, d_e is the expected value, and d_a is the actual value following an interaction. The evaluator a_c also holds a history of the recommendations, $H_{i_r} = (\mathbb{P}\, i_r, count^+, count^-, RT, RT_c)$ where $\mathbb{P}\, i_r$ is the set of recommendations, RT is the recommendation trust in the witness and RT_c is the confidence in that trust. Recommendations are defined as $i_r = (a_c, a_t, a_r, s, t, r)$ where a_t is the target, a_r is the witness who gives recommendation r at time t, and s is the service recommended. Recommendations can be direct, $r^d = (s, a_r, count^+, count^-)$ or indirect, $r^i = (a_{r'}, r^d_{a_{r'}}) \vee (a_{r'}, r^i_{a_{r''}})$, where $a_{r'}$ is an indirect recommender and $r^d_{a_{r'}}$ is the direct recommendation of $a_{r'}$, and $r^i_{a_{r''}}$ is the indirect recommendation of the next witness in the recommender chain $a_{r''}$.

3.2 Network Building and Data Analysis

As an evaluator interacts with providers and witnesses it gathers information about interactions and relationships to build an agent network to better understand its environment. We consider three graph structures to represent an agent's environment: provider graph, witness graph, and a combined provider-witness graph. The nodes represent agents and the edges correspond to links between agents, including the strength of the link in terms of experience. For example, Algorithm 1 shows how part of the witness graph is constructed, r_μ is the currently processed recommendation. For a direct recommendation, an edge is created for each new recommender and the recommendation count is incremented. Indirect recommendations are updated recursively, with edges created or updated from the further recommender $a_{r''}$ in the chain to a closer one $a_{r'}$.

Algorithm 1 Indirect recommendations

for all indirect recommendations r^i **do**
  if $r^i.a_{r'} \notin \mathbb{P}\, a_{r'}$ **then**
    add edge($a_{r'}, a_c$)
  increment $count_{response}$
  repeat
    if $r^i.a_{r''} \notin \mathbb{P}\, a_r$ **then**
      add edge($a_{r''}, a_{r'}$)
  until $r_\mu = r^d$

Trust and reputation data, together with agent network details can give insight into other aspects of the environment. For instance, network information reinforces trust in the roles of witnesses to give accurate information. We focus on emergent information from the multi-agent network, relating to the open issue of collusion. Collusion detection strategies, typically rely on a global view to identify colluding agents. However, no such global view is available to individuals in a decentralised MAS. Despite the limitations of agents' local views, we believe that combining local views with recommendation information gives usable information for simple collusion detection.

Examples of simple collusion include: witness and target

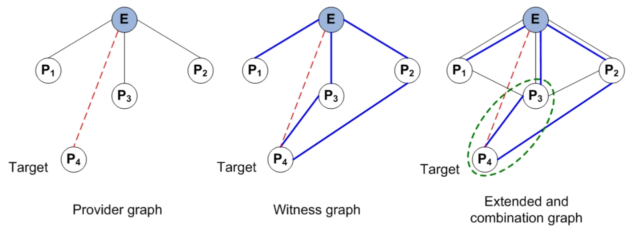

Figure 1: Collusion between target and witness

collusion, where the witness promotes the target, collusion among witnesses to manipulate a target's reputation, and provider collusion over price. Collusive behaviour is characterised by elements such as heavy agent interactions or similar responses to queries as witnesses. Witness and target collusion is depicted in Figure 1. Witness $P3$ and target $P4$ collude to promote $P4$ as a trustworthy provider to evaluator E. Algorithm 2 shows the collusion detection process after target a_β has just provided service s_β following recommendations. After some time period, the opinions of the witnesses about the target can be checked for consistency.

Algorithm 2 Witness and Target Collusion Detection

for all direct recommendations r^d **do**
  if $(r^d.a_t = a_r)$ AND $(r^d.s = s_\beta)$ **then**
    for all dimensions $d \in r^d.s$ **do**
      if $d_a < d_e$ **then**
        add a_r to $\mathbb{P}\, colluders$

4. ONGOING AND FUTURE WORK

A network of agents with the associated interactions and relationship data potentially provides the necessary clues to extract the information that is crucial to understanding the environment. In our ongoing work, we are investigating the detection of various types of collusion among agents, including collusive witnesses, and evaluating our work with an implementation of the network model. We also aim to devise algorithms to enable systematic analysis of agent networks and interaction history information.

5. REFERENCES

[1] J. J. Brown and P. H. Reingen. Social ties and word-of-mouth referral behavior. *Journal of Consumer Behaviour*, 14(3):350–362, 1987.

[2] T. D. Huynh, N. R. Jennings, and N. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Journal of Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006.

[3] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *Proceedings of the 12th International Conference on Information and Knowledge Management*, pages 556–559, 2003.

[4] S. N. Lim Choi Keung and N. Griffiths. Indirect recommendations for improved trust assessment. In *Proceedings of 11th International Workshop on Trust in Agent Societies*, pages 33–40. IFAAMAS, 2008.

[5] S. Milgram. The small world problem. *Psychology Today*, 2:60–67, 1967.

[6] J. Sabater. *Trust and Reputation in Agent Societies.* PhD thesis, Universitat Aùtonoma de Barcelona, 2003.