Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Exploiting Trust Information to Cope with Malicious Entities in Multi-Agent Systems

Athirai A. Irissappane
School of Computer Engineering
Nanyang Technological University, Singapore
athirai001@e.ntu.edu.sg

1 Research Problem

Our research is within the area of artificial intelligence and multi-agent systems. More specifically, we focus on evaluating trust relationships between the agents in multi-agent e-marketplaces and sensor networks and aim to address the following problems,

- how to identify a trustworthy (good quality) agent
- how to cope with dishonest advisors i.e., agents who provide misleading opinions about others

To explain, in multi-agent e-marketplaces, self-interested selling agents may act maliciously by not delivering products with the same quality as promised. It is thus important for buying agents to analyze their quality and determine which sellers to do business with, based on their previous experience with the sellers. However, in most e-markets, buyers often encounter sellers with which they have no previous experience. In such cases, they can query other buyers (called advisors) about the sellers. But, advisors may act dishonestly by providing misleading opinions (unfair ratings) to promote low quality sellers or demote sellers with high quality [Irissappane *et al.*, 2014a]. Hence, it is necessary to evaluate the quality of advisors' opinions to determine their reliability.

Similarly, in the (resource-constraint) Wireless Sensor Networks (WSNs), each sensor node (agent) needs to select a trustworthy next-hop neighbor to route packets, in order to successfully deliver them to the sink. As sensor nodes are deployed in remote and unattended environments, they are easily susceptible to malicious activities i.e., on routing packets, a neighboring sensor node can drop packets instead of forwarding them. Thus, it is necessary to evaluate the trustworthiness of the neighboring nodes (in an energy-efficient manner) before routing packets, by querying other sensor nodes for information, when necessary.

2 Progress to Date

Up to date we have proposed: 1) a POMDP [Kaelbling *et al.*, 1998] based approach (called the SALE POMDP) to optimally select trustworthy sellers in e-marketplaces; 2) a routing scheme (called the Secure Routing POMDP) to identify malicious nodes and select a trustworthy next-hop neighbor in wireless sensor networks.

2.1 The SALE POMDP

Existing trust models [Jøsang *et al.*, 2007] mainly focus on accurately estimating seller trustworthiness rather than optimally choosing a good seller to perform transaction; they simply query 'all' advisors about the sellers' trustworthiness and fail to reason 'when' it is really necessary to query advisors. We have proposed the *(S)eller & (A)dvisor se(LE)ction* (SALE) POMDP [Irissappane *et al.*, 2014b] to optimally select sellers by selectively querying advisors. SALE POMDP enables optimal trade-offs of the expected benefit and cost of obtaining more information (about the sellers and advisors), aiming to maximize the total utility of the buyer.

The SALE POMDP can be described in terms of states, actions, transitions, observations and rewards: Its state consists of the quality levels of each seller, each advisor and status of the buyer's transaction (*notstarted, successful,unsuccessful,gaveup,finished*). SALE POMDP actions include, querying advisor about a seller, querying advisor about another advisor, buying from a seller and not buying from any seller in the e-market. We also assume that for query actions, the state does not change. When taking buy action, the state transitions to *successful*, if the selected seller is trustworthy and *unsuccessful*, otherwise. When taking do not buy action, state transitions to *gaveup*. SALE POMDP receives observations based on the quality levels of the seller/advisor for the query actions taken. The observation probabilities are such that trustworthy advisors give more accurate and consistent answers than untrustworthy ones. There is a small cost for the ask actions. A reward is given for buying from a good seller, otherwise a penalty is levied. There is a penalty for not buying from any seller, when in fact there is a trustworthy seller in the e-marketplace. To resolve scalability issues that arise due large state spaces, we use factored representations (dynamic Bayesian networks with conditional probability tables), and use symbolic Perseus solver that can exploit such factored nature while solving the SALE POMDP.

We conduct extensive evaluation on the ART testbed which demonstrates that SALE POMDP balances the cost of obtaining and benefit of more information more effectively, leading to more earnings, than traditional trust models. Specifically, we consider single transaction settings, where the buyer is a new comer to the market, as well as multiple transaction settings, where the buyer gains experience by involving in many transactions. Results show that, in both cases, our approach leads to better earnings and reasonable accuracy in predicting seller trustworthiness than other trust models, even in the presence of malicious advisors.

2.2 The Secure Routing POMDP (SRP)

POMDP provides a principled approach for decision making under uncertainty and is an ideal choice for nodes in WSNs that need to choose a suitable next-hop neighbor to route packets with only limited information. Also, POMDPs can effectively model the exploration/exploitation tradeoff, i.e., whether to gather more information about sensor nodes to deal with malicious activities or decide to route packets with available information, in an energy-efficient manner. Unfortunately, the routing problem in WSNs is too large to be modeled as a single POMDP. Thus, we propose a *hierarchical* POMDP (called *Secure Routing* POMDP SRP [Irissappane *et al.*, 2015]) approach, exploiting the fact that this problem admits a natural decomposition in subtasks.

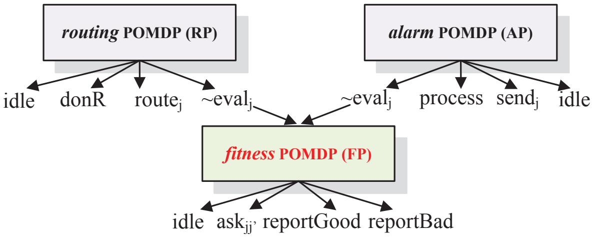

Figure 1: Secure Routing POMDP (SRP)

In particular (as shown in Fig 1), at the higher level of the hierarchy, we select a next-hop neighbor based on its trustworthiness, performed by the *routing* POMDP (RP). We also have the *alarm* POMDP (AP) to raise alarms about untrustworthy nodes. At the lower level, we evaluate the trustworthiness of nodes based on more detailed observations and trust propagation mechanisms, handled by the *fitness* POMDP (FP). The state of RP is composed of the quality levels of each neighboring sensor node. There is a reward for selecting a suitable node to route packets and a penalty, otherwise. For AP, the state is composed of the quality levels of neighbors and variables to represent nodes about which alarms were sent and received. In AP, there is a reward for sending or processing truthful alarms and a penalty for false alarms. For FP, apart from the quality levels of the neighbors, the detailed behavior of the nodes are maintained as state variables i.e., if a neighboring node has qualities $\in \{good, bad\}$, which are determined by several aspects, such as remaining energy, distance from the sink and selfishness in routing packets, the state space of FP contains all these aspects. In FP, there is a reward for correctly reporting the quality of a particular node and a penalty for incorrectly reporting about a node's quality.

The actions for the different POMDPs are shown in Fig 1 (for details refer [Irissappane *et al.*, 2015]). We can see that $eval_j$ (evaluate trustworthiness of node j) is an abstract action i.e., when it is called by RP/AP, control is transferred to FP. Specifically, every time a node receives a packet to route, an episode of RP is activated. This RP episode can trigger multiple episodes of FP. FP evaluates the trustworthiness of nodes and reports it to RP based on which RP selects a next-hop neighbor. The AP is activated after each packet is handled by RP. AP can also call FP to perform node evaluations.

We conduct experiments in a simulated as well as a real testbed. Results show that SRP can successfully route packets, even in the presence of various attacks, in an energy-efficient manner than other trust based routing schemes. Specifically, SRP can cope with network based attacks such as black-hole, on-off attacks, obtaining high packet delivery ratio (ratio of packets successfully delivered to the sink) and lower residual-energy (average remaining energy of the nodes). SRP can also deal with trust based attacks e.g., random, camouflage, etc. [Irissappane *et al.*, 2014a].

3 Future Work

The presented research opens up many directions of future work. Though we use factored representations to reduce the state space complexity of the SALE POMDP model, SALE POMDP cannot scale to more than a handful number of agents (say 10-12), as computing the optimal policy becomes intractable. Also, in our second work i.e., the SRP model, though we propose a hierarchical formulation to reduce the complexity of SRP, our experiments suggest that it still cannot scale to a large number (more than $6-7$) of sensor nodes. As future work, we plan to reduce the complexity in computing the POMDP policy for both models, using a distributed approach i.e., dividing the larger POMDP problem into smaller problems and solving each of them individually. We will also investigate using finite-state controllers to further improve the complexity in representing the POMDP policies.

Both SALE POMDP and SRP employ parameters i.e., transition, observation probabilities which are manually specified. While we experimentally show that SALE POMDP and SRP are robust against the choice of these parameters, an interesting direction of future work is to automatically optimize these (e.g., using Bayesian learning).

References

[Irissappane *et al.*, 2014a] Athirai A Irissappane, Siwei Jiang, and Jie Zhang. A biclustering-based approach to filter dishonest advisors in multi-criteria e-marketplaces. In *AAMAS*, 2014.

[Irissappane *et al.*, 2014b] Athirai A Irissappane, Frans A Oliehoek, and Jie Zhang. A POMDP based approach to optimally select sellers in electronic marketplaces. In *AAMAS*, 2014.

[Irissappane *et al.*, 2015] Athirai A Irissappane, Jie Zhang, Frans A Oliehoek, and Partha S Dutta. Secure routing in wireless sensor networks via POMDPs. In *IJCAI*, 2015.

[Jøsang *et al.*, 2007] Audun Jøsang, Roslan Ismail, and Colin Boyd. A survey of trust and reputation systems for online service provision. *Decision support systems*, 43(2):618–644, 2007.

[Kaelbling *et al.*, 1998] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.