ECAI 2016
G.A. Kaminka et al. (Eds.)

doi:10.3233/978-1-61499-672-9-1636

Towards a Framework for Detecting Opportunism in Multi-Agent Systems

Jieting Luo[1] and **John-Jules Meyer**[2] and **Max Knobbout**[3]

1 Introduction

Consider a common social scenario. A seller sells a cup to a buyer and it is known by the seller beforehand that the cup is actually broken. The buyer buys the cup without knowing it is broken. Since the buyer's value gets demoted, the behavior performed by the seller is usually forbidden by social norms. Such a social behavior intentionally performed by the seller is first named opportunistic behavior (or opportunism) by economist Williamson [6]. It is a typical social behavior that is motivated by self-interest and takes advantage of knowledge asymmetry about the behavior to achieve own gains, regardless of the principles [3]. This definition implies that, given a social context, opportunistic behavior results in promoting agents' own value while demoting social value. Therefore, it is prohibited by norms in most societies. In the context of multi-agent systems, we constrain such a selfish behavior through setting enforcement norms, in the sense that agents receive a corresponding sanction when they violate the norm. On the one hand, it is important to detect it, as it has undesirable results for the participating agents. On the other hand, as opportunism is always in the form of cheating, deception and betrayal, meaning that the system does not know what the agent performs or even the motivation behind it (for example, in a distributed system), monitors can only observe agents' opportunistic behavior indirectly. Therefore, there has to be a monitoring mechanism that can detect the performance of opportunistic behavior in the system. This paper introduces the idea of using a logical framework based on the specification of actions to verify whether agents in the system perform opportunistic behavior.

2 Framework

Since monitors cannot observe the performance of opportunism directly, the action can only be identified through the information about the context where the action can be performed and the property change in the system, which is called *action specification* [5] or *action description* [2]. Usually an action can be specified through its precondition and its effect (postcondition): the precondition specifies the scenario where the action can be performed whereas the postcondition specifies the scenario resulting from performing the action. For example, the action, dropping a glass to the ground, can be specified as holding a glass as its precondition and the glass getting broken as its effect. Therefore, we assume that every action has a pair of the form $\langle \psi, \psi' \rangle$, where ψ is the precondition of action a and ψ' is the effect of performing action a in the context of ψ.

[1] Utrecht University, the Netherlands, email: J.Luo@uu.nl
[2] Utrecht University, the Netherlands, email: J.J.C.Meyer@uu.nl
[3] Delft University of Technology, the Netherlands, email: M.Knobbout@tudelft.nl

The models that we use are transition systems, which consist of agents, states, actions, transitions between states by actions and a valuation function mapping a state to a set of true propositions. When an action is performed in a certain state s, the system might progress to a different state in which different propositions might hold. Such a system is a very generic way of modelling a multi-agent system (see for example [7]). Since we have already introduced the notion of action specification $\langle \psi, \psi' \rangle$, all the possible state transitions are defined such that they go from a ψ-state to a ψ'-state. We also extend the standard framework with a monitor relation $\mathcal{M}$, which represents the indistinguishability of a monitor over different states. Sometimes we also use $s\langle a \rangle$ to denote the state resulting from the performance of action a in state s. The logical language we use in this paper is propositional logic extended with action modality for reasoning about dynamic worlds. The syntax and the semantics are defined in the same way as propositional logic except the formula $\langle a \rangle \varphi$. This formula holds if and only if φ is true after action a is performed in state s.

Similar to [1], we simply consider a norm as a subset of all the state transitions that is decided by designers of the system. In other words, if a norm is denoted as η, a state transition is an η-violation if and only if it is in the set η. We define norms of the form $\eta(\varphi, a)$, interpreted as it is forbidden to perform action a in a φ-state. This is the most common form in which the action and the context where the action is forbidden are explicitly represented, regardless of the effect that the action brings about. For example, it is forbidden to smoke in a non-smoking area. Of course, it is only a choice in this paper and more forms of norms are described and constructed based on our logical framework in the full paper [4].

3 Defining Opportunism

Before we propose our monitoring approach for opportunism, we should formally define opportunism from the perspective of the system so that the system knows what to detect for monitoring opportunism. In our previous paper [3], we emphasized opportunistic behavior is performed by intent rather than by accident. However, monitors cannot read agents' mental states, so for monitoring we assume that agents violate the norms always by intention from a pragmatic perspective. For example, we always assume that speeding is performed with intention. In this paper we remove all the references to the mental states from the formal definition of opportunism in our previous paper [3], assuming that the system can tell agents' value promotion/demotion causing by an action. In a sentence, from the perspective of the system, opportunistic behavior performed by an agent in a social context can be simply defined as a behavior that causes norm violations and promotes his own value. In this short paper, opportunism is denoted as $Opportunism(\eta, a)$, interpreted as action a is opportunistic behavior with respect to norm η. Opportunistic behavior results in promoting agents' own value, which can be interpreted as that opportunistic agents prefer the state that results from opportunistic behavior rather than the initial state. For having preferences over different states, we argue that agents always evaluate the truth value of specific propositions in those states based on their value systems. Based on this understanding, we define a function of the form $EvalRef(V_i, s, s')$ mapping a value system and two states to a proposition an agent refers to for specifying his preference over two states.

4 Monitoring Opportunism

In this paper, a monitor is considered as an external observer to evaluate a state transition with respect to a given norm. However, a monitor can only verify state properties instead of observing the performance of actions directly. Our approach to solve this problem is to check how things change in a given state transition and reason about the action taking place in between.

We first define a state monitor $m_{state}(\varphi)$, which can evaluate the validity of a given property in a given state. Because a monitor can be seen as an external observer that can observe agents' activities, we can define state monitors in this paper in a similar way as we define knowledge in epistemic logic, and correspondingly adopt S5 properties.

Definition 1 (State Monitors)**.** *Given a monitoring transition system $\mathcal{I}$, a value system set V, and a propositional formula φ, a state monitor m_{state} for φ over $\mathcal{I}$ is defined as follows: $\mathcal{I}, s \vDash m_{state}(\varphi)$ iff for all s' $s\mathcal{M}s'$ implies $\mathcal{I}, V, s' \vDash \varphi$.*

$m(\varphi)$ is read as φ "is detected" to be φ. As the $\mathcal{M}$-relation is reflexive, we have the validity $\vDash m_{state}(\varphi) \to \varphi$, meaning that what the state monitor detects is always considered to be true.

State monitors are the basic units in our monitoring mechanism. We can combine state monitors to check how things change in a given state transition and evaluate it with respect to a given set of norms. As we defined in Section 3, opportunistic behavior performed by an agent is a behavior that causes norm violations and promotes his own value. In other words, opportunism is monitored with respect to a norm and a value system of an agent. Based on this definition, we design a monitoring opportunism approach $m_{opp}((\varphi, a), \langle\psi, \psi'\rangle, a')$ with respect to norm $\eta(\varphi, a)$.

Definition 2. *Given a monitoring transition system $\mathcal{I}$, a value system set V, a norm $\eta(\varphi, a)$, and a pair $\langle\psi, \psi'\rangle$ of action a, in order to check action a' performed by agent i in state s is opportunistic behavior, we can combine monitors as follows:*

$$\mathcal{I}, s \vDash m_{opp}((\varphi, a), \langle\psi, \psi'\rangle, a') := m_{state}(\varphi \wedge \psi) \wedge \langle a' \rangle m_{state}(\psi')$$

where $\varphi \wedge \psi$ implies $\neg p$ and ψ' implies p, and $p = EvalRef(V_i, s, s\langle a' \rangle)$.

In order to check whether action a' is opportunistic behavior (violates norm $\eta(\varphi, a)$ and promotes own value), we verify if action a' is performed in φ-state. Besides, we check if action a' is the action that the norm explicitly states. Since the monitors cannot observe the performance of action a' directly, we only can identify action a' to be possibly action a with pair $\langle\psi, \psi'\rangle$ by checking if formulas ψ and ψ' are successively satisfied in the state transition by action a'. With this approach we have a candidate set of states for state s and a candidate set of states for state $s\langle a' \rangle$ and any two states from them satisfy the resulting property of function $EvalRef$, which means that given the partial information the execution of action a' in state s brings about p thus promoting agent i's value.

However, since the monitors can only verify state properties instead of observing the performance of the action directly, we cannot guarantee that an action that is detected to be opportunistic was indeed opportunistic, because there might exist more than one action that can be represented by pair $\langle\psi, \psi'\rangle$. That is, formula $\mathcal{I}, V, s \vDash m_{opp}((\varphi, a), \langle\psi, \psi'\rangle, a') \to Opportunism((\varphi, a), a')$ might not hold. Given this problem, we need to investigate in which case or with what requirement the action that is detected by the opportunism monitor is indeed opportunistic behavior. In order to guarantee that action a' that is detected to be opportunistic was indeed opportunistic, we should make sure that, within the actions available in φ-state, there exists only one action that can be represented with pair $\langle\psi, \psi'\rangle$. With this condition, action a' is indeed action a as norm $\eta(\varphi, a)$ indicates, so we can guarantee that action a' that is detected to be opportunistic was indeed opportunistic. This approach implies: in order to better monitor opportunistic behavior, we should appropriately find an action pair $\langle\psi, \psi'\rangle$ such that the possible actions in between can be strongly restricted and minimized. Assuming that the action pair we use is $\langle\top, \top\rangle$, the possibility that the opportunism monitor makes an error is extremely high, because every action that is available in a φ-state will be detected to be opportunistic behavior. However, sometimes it is difficult to find a unique pair $\langle\psi, \psi'\rangle$ for the action we monitor, especially when we cannot limit the available actions with the given context (a φ-state). So it is important to have more information not only about the action, but also about the context where the action performed and the system. All these issues will be elaborated and discussed in the full paper [4].

5 CONCLUSION

In this paper, we introduce the idea of verifying opportunism, which is a behavior that causes norm violation and promotes agents' own value. Our logical framework is developed based on the specification of actions. In particular, we investigated how to evaluate agents' actions to be opportunistic with respect to norms when those actions cannot be observed directly. Future work can investigate more formal properties to improve the effectiveness of our monitoring approach for opportunism: whenever an action is detected to be opportunistic, it was indeed opportunistic; whenever an action was opportunistic, it is indeed detected. Monitoring costs can be another interesting topic to be studied based on our monitoring approach.

REFERENCES

[1] Nils Bulling, Mehdi Dastani, and Max Knobbout, 'Monitoring norm violations in multi-agent systems', in *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pp. 491–498, (2013).

[2] J Fiadeiro and T Maibaum, 'Temporal reasoning over deontic specifications', *Journal of Logic and Computation*, **1**(3), 357–395, (1991).

[3] Jieting Luo and John-Jules Meyer, 'A formal account of opportunism based on the situation calculus', *AI & SOCIETY*, 1–16, (2016).

[4] Jieting Luo, John-Jules Meyer, and Max Knobbout, 'Monitoring opportunism in multi-agent systems', in *Proceedings of the Coordination, Organizations, Institutions and Norms workshop @ECAI2016*, p. to appear. Springer.

[5] Raymond Reiter, *Knowledge in action: logical foundations for specifying and implementing dynamical systems*, MIT press, 2001.

[6] Oliver E Williamson, 'Markets and hierarchies: analysis and antitrust implications: a study in the economics of internal organization', (1975).

[7] Michael Wooldridge, *An introduction to multiagent systems*, John Wiley & Sons, 2009.