Multi-Agent Action Modeling through Action Sequences and Perspective Fluents

Chitta Baral, Gregory Gelfond, Enrico Pontelli and Tran Cao Son

Abstract

Actions in a multi-agent setting have complex characteristics. They may not only affect the real world, but also affect the knowledge and beliefs of agents in the world. In many cases, the effect on the beliefs or knowledge of an agent is not due to that agent actively doing some actions, but could be simply the result of that agent's perspective in terms of where it is looking. In dynamic epistemic logic (DEL), such multi-agent actions are expressed as complex constructs or as Kripke model type structures. This paper uses the multi-agent action language $m\mathcal{A}+$ to show how one can take advantage of some of the *perspective fluents* of the world to model complex actions, in the sense of DEL, as simple action sequences. The paper describes several plan modules using such actions. Such plan modules will be helpful in planning for belief and knowledge goals in a multi-agent setting, as planning from scratch would often be prohibitively time consuming.

Introduction

Reasoning about actions and their effects in a multi-agent setting involves reasoning about agents' knowledge and belief about the world as well as about each other's knowledge and belief. Except very recently (Baral et al. 2012; 2013; Bolander and Andersen 2011), this direction has been pursued more in the literature of dynamic epistemic logic (DEL) (van Ditmarsch, van der Hoek, and Kooi 2008; Baltag and Moss 2004; van Benthem, van Eijck, and Kooi 2006; Baral et al. 2012). In these works, actions such as *public announcements* and *semi-private observations* are complex objects. For example, the complex action involving two agents a and b, where a senses the value of p, and it is common knowledge between a and b that b observes a sensing, is expressed in one of the DEL as $L_{ab}(L_a?p \cup L_a?\neg p)$. Alternatively, in the *Concealed Coin Domain* of (Baltag and Moss 2004), a semi-private observation such as *"agent A peeks into the box to sense which face of the coin is up, with agent B watching him, and C unaware of what has transpired"* is encoded by an update model, a graph-like general, but complex, representation of multi-agent actions.

In the graphical approach, there is no notion of a general action of *observing the value of a fluent*, and hence a slightly different occurrence, such as one where both B and C are aware of A's actions, necessitates a completely different representation. Thus, occurrences of the same *elementary action* which differ only with respect to how the agents perceive them are treated as separate classes of actions. As a result, there is a need to specify numerous different action models for each possible occurrence of an elementary action. For example, Fig. 1 shows the update models for two kinds of announcements: *"A publicly announces that the coin is facing heads up,"* and *"A privately tells B that the coin lies heads up,"* in a domain with agents A, B, and C.

In this paper, our contention is that, while the representation of actions as update models (Baltag and Moss 2004) is quite general, *when possible*, it is advantageous to express it as an action sequence, as is done in the traditional reasoning about actions and planning community. For example, the action with respect to the coins can be simply expressed, from the perspective of agent A, by the sequence: *distract(A,C), signal(A,B), peek(A)*. The intuition behind this action sequence is: first, agent A distracts C, thus causing C to look away. Then, A signals B to make sure B is watching A. When A peeks at the coin, B (who is watching) knows about it and C (who is looking away) will be unaware of it. In our view, B's watching of A and C's looking away from A are properties of the world (albeit of a special kind, which we call *perspective fluents*) and thus one need not make them part of a complex action model as in (Baltag and Moss 2004). A reason for having sequences of simple actions as an alternative to update models, is that many existing planning algorithms in simpler domains are designed to analyze and produce sequences of simple actions. Thus, having the same structure of plans will make it easier to enhance existing planning algorithms to the multi-agent domain with knowledge and belief goals.

Another concern about the language of update models is that it often ignores the *executability of actions*: consider, for example, the case where an agent wishes to perform a public announcement but only a subset of the agents is actually listening. How exactly will an agent execute the action of public announcement? A realistic approach would be for him/her to first make sure that everyone is paying attention (e.g., not listening to music using an earphone) and then announce it loud enough for everyone to hear. This can be achieved by a *sequence of actions* consisting of first signaling each of the agents[1] and then speaking loudly.

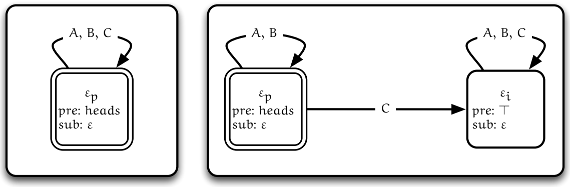

Figure 1: Action model for public (left) and semi-private announcement (right)

To address the above concerns, we investigate the use of an alternative notion of what an action is in multi-agent setting with stress on their executability. In addition, we introduce a notion of perspective fluents and identify some basic actions that mainly manipulate such fluents. We show how many of the known multi-agent actions can be modeled as a sequence of simple executable actions. We also develop a new multi-agent domain, the "Escapee Domain" and introduce several multi-agent action primitives such as group formation and dissolution, and discuss how they can be simulated using action sequences and perspective fluents.

Background

We use the multi-agent action language $m\mathcal{A}+$ (Baral et al. 2012) extended with the notion of *perspective* fluents which were not mentioned there. Theories of $m\mathcal{A}+$ are defined over a multi-agent domain, $\mathcal{D}$, with a signature $\Sigma = (\mathcal{AG}, \mathcal{F}, \mathcal{A})$ where $\mathcal{AG}$, $\mathcal{F}$, and $\mathcal{A}$, are finite, disjoint, nonempty sets of symbols respectively defining the *names of the agents within* $\mathcal{D}$, the *properties of* $\mathcal{D}$ (or *fluents*), and the *elementary actions* which the agents may perform.

The language $\mathcal{L}$ for representing and reasoning about the knowledge and beliefs of the agents is comprised of formulae built from fluents, propositional connectives, and a set of modal operators $\mathbf{B}_i$, with one such operator per agent in $\mathcal{AG}$. Two types of formulae, *fluent formulae* (those that built over $\mathcal{F}$) and *belief formulae* are considered, where a belief formula is *(i)* a fluent formula, or *(ii)* $\mathbf{B}_i\psi$ where ψ is a belief formula, or *(iii)* $\psi \vee \phi$, $\psi \wedge \phi$, or $\neg\psi$, where ψ and ϕ are belief formulae. In addition, given a formula φ and $\emptyset \nsubseteq S \subseteq \mathcal{AG}$, $\mathbf{C}_S\varphi$ represent the common knowledge among S. A *Kripke structure* is a tuple $\langle \Omega, \pi, \{R\}_{\alpha \in Agents} \rangle$, where Ω is a set of state symbols, π is a function associating an interpretation of $\mathcal{F}$ to each element of Ω, and $R_\alpha \subseteq \Omega \times \Omega$ for $\alpha \in Agents$. The satisfaction relation between belief formulae and a Kripke world (M, ω) is defined in the usual way.

Direct effects of actions are represented using laws which are statements of the form:

$$a \textbf{ causes } \lambda \textbf{ if } \phi \tag{1}$$

$$a \textbf{ determines } f \tag{2}$$

$$a \textbf{ communicates } \varphi \tag{3}$$

[1]This can be parallelized and we do not discount such genuine parallelization.

where a is an action, f is a fluent, λ is a fluent literal, ϕ is a conjunction of fluent literals, and φ is a belief formula. (1) (called *ontic action*) is read as "performing the action a in a state which satisfies ϕ causes λ to be true." (2) states that "if an agent performs the *sensing* action a, he will learn the value of the fluent f." (3) indicates that "if an agent performs the *communication* action a, the truth value of φ will be announced." Only truthful announcements are allowed.

(1)–(3) only describe the *direct effects* of their respective actions. In general, an agent's actions may indirectly affect the knowledge/beliefs of his fellows, depending on the *frames of reference* that the agents have w.r.t the action occurrence. In $m\mathcal{A}+$, for any given action occurrence, the agents are divided into three groups:[2] those who are fully aware of both the action occurrence and its effects; those who are aware of the occurrence but not of the full consequences of the action; and those who are oblivious as to what has transpired. Frames of reference are dynamic in nature and are described by statements of the form:

$$X \textbf{ observes } a \textbf{ if } \phi \tag{4}$$

$$X \textbf{ aware of } a \textbf{ if } \phi \tag{5}$$

where X is a set of agent names, a is an action, and ϕ is a belief formula. (4), called *observation axioms*, define the set of agents who are fully aware of both the action occurrence and its effects. (5), called *awareness axioms*, define the set of agents who are aware of the occurrence, but only partially of its effects. By default, it is assumed that all other agents within the domain are oblivious. Fluents appearing in ϕ are called *perspective fluents*.

In $m\mathcal{A}+$, the only assumptions made regarding the frames of reference of the agents are that those who are fully aware of an action occurrence and its effects, as well as those who are aware only of the occurrence, know the frames of reference of all of the agents within the domain.

An agent's ability to perform an action is determined by the current state of the domain. This is described by *executability conditions*, which are statements of the form:

$$\textbf{executable } a \textbf{ if } \phi \tag{6}$$

where a is an action and ϕ is a belief formula. Statements of this form are read as: "action a is executable in any state of the domain which satisfies ϕ." In $m\mathcal{A}+$, it is assumed that a unique executability condition governs each action (if any).

Initial state axioms are used to describe the initial state and are statements of the form:

$$\textbf{initially } \varphi \tag{7}$$

where φ is a belief formula.

A *domain specification (or domain)* $\mathcal{D}$ in $m\mathcal{A}+$ is a collection of statements of the form (1)–(6). Consistency of $\mathcal{D}$ is defined in the usual way. An *action theory* is a pair, $(\mathcal{D}, I)$, where $\mathcal{D}$ is a consistent domain specification and I is a set of initial state axioms.

The **semantics** of an $m\mathcal{A}+$ action theory $(\mathcal{D}, I)$ is defined by a transition function $\Phi_{\mathcal{D}}$ which maps pairs of actions and

[2]This can be further generalized.

states into sets of states. $m\mathcal{A}+$ defines a state as a *Kripke world* which is a pair (M, ω), where M is a Kripke structure and $\omega \in \Omega$ (denoting the "real" state of the world). The function $\Phi_{\mathcal{D}}$ is defined separately for each type of actions and takes into consideration the frame of reference of the agents with respect to that action occurrence. The *frames of reference* with respect to an occurrence of an action a, in the state σ, partition the agents of the domain into three sets as follows: (i) the set of agents who are *fully aware* of a, denoted by $f(\sigma, a)$, is $\{\alpha \in \mathcal{AG} \mid [\alpha \textbf{ observes } a \textbf{ if } \phi] \in \mathcal{D} \wedge (M, \omega) \models \phi\}$; (ii) the set of agents who are *partially aware* of a, denoted by $p(\sigma, a)$, is $\{\alpha \in \mathcal{AG} \mid [\alpha \textbf{ aware of } a \textbf{ if } \phi] \in \mathcal{D} \wedge (M, \omega) \models \phi\}$; and (iii) the set of agents who are *oblivious* of a, denoted by $o(\sigma, a)$, is $\mathcal{AG} \setminus (f(\sigma, a) \cup p(\sigma, a))$.

$\Phi_{\mathcal{D}}$ guarantees that: (a) fully aware agents would have the correct information about the successor state; (b) partially aware agents know about the action occurrence and can reason about its effects on the beliefs of fully aware agents but might not know the effects of the actions; and (c) oblivious agents believe that nothing changes. Furthermore, (a)-(c) are common knowledge amongst agents in the same partition.

The definition of $\Phi_{\mathcal{D}}$ makes use of the notion of a update model in (Baltag and Moss 2004; van Benthem, van Eijck, and Kooi 2006). It is shown that $\Phi_{\mathcal{D}}$ coincides with the semantics for elementary actions using update models (Baral et al. 2012). The technical details and the properties of $\Phi_{\mathcal{D}}$ can be found in the aforementioned paper on $m\mathcal{A}+$. The function $\Phi_{\mathcal{D}}$ is then extended to define $\Phi^*_{\mathcal{D}}$ that allows for reasoning about the effects of a sequence of actions in the usual way (see, e.g., (Gelfond and Lifschitz 1998)).

Definition 1 *An action description* $\Delta = (\mathcal{D}, I)$ entails *the query* φ **after** Π, *denoted by* $\Delta \models \varphi$ **after** Π, *if:*

- $\Phi^*_{\mathcal{D}}(\Pi, I_0) \neq \{\bot\}$, *and*
- $\sigma = (M, \omega) \models \varphi$ *for each* $\sigma \in \Phi^*_{\mathcal{D}}(\Pi, I_0)$

where I_0 *is the initial belief state of* Δ.

Modeling Simple Multi-Agent Actions

In this section, we start with presenting and axiomatizing the "Strongbox Domain," a variation of the "Concealed Coin Domain" from (Baltag and Moss 2004; Baral and Gelfond 2011), and show how to simulate the following classes of actions: private and semi-private observation; and semi-private announcements by sequences of elementary actions.

Strongbox Domain: Three agents A, B and C are together in a room containing a strongbox in which there is a coin. This fact is common knowledge amongst the agents, as is the fact that none of them knows which face of the coin is showing. To this basic description, we add the following: (i) initially the agents are paying attention to their surroundings and this fact is common knowledge amongst them; (ii) an agent may distract or signal one of his fellows rendering them respectively inattentive/attentive; (iii) agents who are paying attention are aware of the actions that their fellows may perform; (iv) an agent may peek into the strongbox to learn which face of the coin is showing; and (v) an agent may communicate a formula to a fellow agent.

The domain can be encoded using the following domain signature: $\mathcal{AG} = \{A, B, C\}$; $\mathcal{F} = \{heads, attentive(\alpha)\}$; and $\mathcal{A} = \{peek(\alpha), tell(\alpha_1, \alpha_2, heads), signal(\alpha_1, \alpha_2), distract(\alpha_1, \alpha_2)\}$, where α, α_1 and α_2 are variables over $\mathcal{AG}$. $attentive(\alpha)$ is a *perspective fluent* whose truth determines whether or not α is actively paying attention to the environment as stated in (i); this fluent can be manipulated by a new kind of ontic actions, called $signal$ and $distract$ (in (ii)) whose direct effects are described by the laws:

$$signal(\alpha_1, \alpha_2) \textbf{ causes } attentive(\alpha_2) \quad (8)$$

$$distract(\alpha_1, \alpha_2) \textbf{ causes } \neg attentive(\alpha_2) \quad (9)$$

Item (iii) describes how the perspective fluents affect the frames of reference that the agents may have with respect to various action occurrences. This information can be encoded using the following axioms:

$$\{\alpha_1, \alpha_2\} \textbf{ observes } signal(\alpha_1, \alpha_2) \quad (10)$$

$$\{\alpha\} \textbf{ observes } signal(\alpha_1, \alpha_2) \textbf{ if } attentive(\alpha) \quad (11)$$

$$\{\alpha_1, \alpha_2\} \textbf{ observes } distract(\alpha_1, \alpha_2) \quad (12)$$

$$\{\alpha\} \textbf{ observes } distract(\alpha_1, \alpha_2) \textbf{ if } attentive(\alpha) \quad (13)$$

Peeking into the strongbox is a *sensing action* and is described by the sensing axiom:

$$peek(\alpha) \textbf{ determines } heads \quad (14)$$

The agent performing the action is fully aware of the action occurrence. Any agents who are paying attention to their environment are simply observers of the action. For example, if A peeks into the box, and B and C are both paying attention, then B and C will be aware that A peeked. Agents who are not attentive however will be oblivious as to the actions of their fellows. This is represented in $m\mathcal{A}+$ as follows:

$$\{\alpha\} \textbf{ observes } peek(\alpha) \quad (15)$$

$$\{\alpha_2\} \textbf{ aware of } peek(\alpha_1) \textbf{ if } attentive(\alpha_2) \quad (16)$$

Communication between agents is described in a similar fashion, with the caveat that only attentive agents may communicate:

$$tell(\alpha_1, \alpha_2, heads) \textbf{ communicates } heads \quad (17)$$

$$\{\alpha_1, \alpha_2\} \textbf{ observes } tell(\alpha_1, \alpha_2, heads) \quad (18)$$

$$\{\alpha_3\} \textbf{ aware of } tell(\alpha_1, \alpha_2, heads) \textbf{ if} \quad (19)$$

$$\textbf{executable } tell(\alpha_1, \alpha_2, heads) \textbf{ if} \quad (20)$$
$$attentive(\alpha_1) \wedge attentive(\alpha_2)$$

The Initial State: According to scenario, we know that: A, B and C are together in a room containing a strongbox in which there is a coin. This fact is common knowledge amongst the agents, as is the fact that none of them knows which face of the coin is showing. Furthermore, let us suppose that the coin is actually showing heads, and that it is common knowledge that all of the agents are paying attention to their environment. This information could be encoded by the following initial state axioms:

$$\textbf{initially } heads \quad (21)$$

$$\textbf{initially } \mathbf{C}_{\{A,B,C\}} attentive(\alpha) \; (\alpha \in \{A, B, C\}) \quad (22)$$

$$\textbf{initially } \mathbf{C}_{\{A,B,C\}} \neg \mathbf{B}_\alpha heads \wedge \neg \mathbf{B}_\alpha \neg heads \quad (23)$$

For the rest of this section, let σ_0 be the initial belief state of the action description $(\mathcal{D}, I)$ where $\mathcal{D}$ is the domain description in $m\mathcal{A}+$ of the Strongbox Domain and I is the set of initial state axioms (21)–(23).

Private Observation: Suppose that agent A wishes to perform the following complex action in the sense of (Baltag and Moss 2004): "A peeks into the box and sees that the coin is facing heads up, while both agents B and C are oblivious as to what has transpired." In $m\mathcal{A}+$, such a *private observation* would be modeled by the following plan, Π_0:

$$[distract(A,B), distract(A,C), peek(A)]$$

Let φ_0 be the following belief formula:

$$\neg attentive(B) \wedge \neg attentive(C) \wedge attentive(A) \wedge$$
$$\mathbf{B}_A heads \wedge \mathbf{B}_A(\neg attentive(B) \wedge \neg attentive(C)) \wedge$$
$$\mathbf{C}_{\{B,C\}}(\neg \mathbf{B}_A heads \wedge \neg \mathbf{B}_A \neg heads)$$

Proposition 1 states the consequences of performing the plan, Π_0, in state, σ_0, satisfy our intuition — that A knows which face of the coin is showing, and that the beliefs of B and C are unchanged.

Proposition 1 $(\mathcal{D}, \sigma_0) \models \varphi_0$ **after** Π_0.

Semi-Private Observation Now suppose that instead of his private observation, agent A performs the following complex action: "A peeks into the box and sees that the coin is facing heads up, with agent B watching him and C oblivious of what has transpired." Such a *semi-private observation* would correspond to the following plan in $m\mathcal{A}+$, Π_1:

$$[distract(A,C), peek(A)]$$

Let φ_1 be the following belief formula:

$$\neg attentive(C) \wedge attentive(A) \wedge attentive(B) \wedge$$
$$\mathbf{B}_A heads \wedge \mathbf{B}_A(attentive(B) \wedge \neg attentive(C)) \wedge$$
$$\mathbf{C}_{\{A,B\}}(\mathbf{B}_A heads \vee \mathbf{B}_A \neg heads) \wedge$$
$$\mathbf{C}_{\{A,B\}}(\neg \mathbf{B}_B heads \wedge \neg \mathbf{B}_B \neg heads) \wedge$$
$$\mathbf{C}_{\{A,B\}}(\neg \mathbf{B}_C heads \wedge \neg \mathbf{B}_C \neg heads) \wedge$$
$$\mathbf{C}_{\{A,B\}}\mathbf{B}_C(\mathbf{C}_{\{A,B,C\}}(\neg \mathbf{B}_\alpha heads \wedge \neg \mathbf{B}_\alpha \neg heads))$$

For a semi-private observation our intuition, as captured by φ_1, is more complex than for private observations. In this case, our intuition informs us that A has now learned that the value of $heads$ is true, and that B is aware that A has learned the value of fluent. The beliefs of the oblivious agent C have not changed. As with private observation, Prop. 2 states that the consequences of performing Π_1 in σ_0 match our intuition.

Proposition 2 $(\mathcal{D}, \sigma_0) \models \varphi_1$ **after** Π_1.

Semi-Private Announcements: Now suppose that after he peeked into the strongbox, A wishes to tell C what he saw, and to have B aware of this fact. In other words: "A tells C that he saw heads, with B watching them communicate." If we assume that this action takes place after the execution of Π_1 then such a *semi-private announcement* would correspond to the following plan, Π_2:

$$[signal(A,C), tell(A,C,heads)]$$

Let Π_3 be the plan obtained from appending Π_2 to Π_1: $[distract(A,C), peek(A), signal(A,C), tell(A,C,heads)]$

Let $\varphi_2 = \mathbf{C}_{\{A,B,C\}} attentive(\alpha) \wedge \mathbf{C}_{\{A,C\}} heads \wedge$
$$\mathbf{C}_{\{A,B,C\}}\mathbf{B}_B(\mathbf{C}_{\{A,C\}} heads \vee \mathbf{C}_{\{A,C\}} \neg heads)$$

Proposition 3 states the consequences of an occurrence of the semi-private communication between agents A and C, which occurs after A's semi-private observation of the value of the fluent $heads$, comports with φ_2.

Proposition 3 $(\mathcal{D}, \sigma_0) \models \varphi_2$ **after** Π_3.

Modeling Complex Multi-Agent Actions

The techniques discussed in the previous section for representing basic multi-agent interactions can be readily applied towards more advanced domains. In this section we illustrate how this approach may be used to represent and reason about a more complex multi-agent domain known as the "Escapee Domain". In the context of the "Escapee Domain", we show how the following classes of actions may be simulated by action sequences: private announcements; private ontic actions; group formation; collaborative action; and finally group dissolution.

Escapee Domain: Suppose that agent A is held captive by a hostile agent B. In order to escape, A must open his cell without B's knowledge. Fortunately, agent C is a double agent in B's organization and may release A from his cell. C does not want to break his cover however, so he may release A only if B is not watching. Once A has been released, he must work together with C to subdue agent B, and then make his escape. A will only work with C if he believes that C is an ally.

A simple refinement of this domain gives us the following signature: $\mathcal{AG} = \{A, B, C\}$; $\mathcal{F} = \{free(\alpha), bound(\alpha), captor(\alpha_1, \alpha_2), united(\alpha_1, \alpha_2), attentive(\alpha), allies(\alpha_1, \alpha_2)\}$; and $\mathcal{A} = \{escape(\alpha), release(\alpha_1, \alpha_2), subdue(\alpha_1, \alpha_2, \alpha_3)$ $unite(\alpha_1, \alpha_2), disband(\alpha_1, \alpha_2), signal(\alpha_1, \alpha_2),$ $distract(\alpha_1, \alpha_2), tell(\alpha_1, \alpha_2, \varphi)\}$
where $\alpha, \alpha_i \in \mathcal{AG}$ and φ is of the form $allies(\alpha_1, \alpha_2)$.

As was the case with the Strongbox Domain, $attentive$ is a perspective fluent. In addition, the fluent $united$, is also a perspective fluent which will be used to model collaborative actions. These fluents may be manipulated directly by the agents via the *perspective altering actions* $signal$/$distract$ and $unite$/$disband$, respectively. The actions $signal$ and $distract$ are represented in a similar manner to the one presented in the formalization of the Strongbox Domain.

In general, agents may $unite$ in order to act together. In the context of the Escapee Domain, an agent must be unbound before he may unite with another agent to collaboratively perform some action. In addition, an agent will only unite with someone whom he believes is an ally. Once they are done collaborating, they may $disband$.

$$unite(\alpha_1, \alpha_2) \ \mathbf{causes} \ united(\alpha_1, \alpha_2) \quad (24)$$

$$\mathbf{executable} \ unite(\alpha_1, \alpha_2) \ \mathbf{if} \quad (25)$$
$$\neg bound(\alpha_1) \wedge \neg bound(\alpha_2) \wedge \mathbf{B}_{\alpha_1} allies(\alpha_1, \alpha_2)$$

$$disband(\alpha_1, \alpha_2) \ \mathbf{causes} \ \neg united(\alpha_1, \alpha_2) \quad (26)$$

The observation axioms governing the frames of reference of the agents with respect to occurrences of the actions $unite$ and $disband$ follow the same pattern as those for the actions $signal$ and $distract$:

$$\{\alpha_1, \alpha_2\} \textbf{ observes } unite(\alpha_1, \alpha_2) \tag{27}$$

$$\{\alpha_3\} \textbf{ observes } unite(\alpha_1, \alpha_2) \textbf{ if } attentive(\alpha_3) \tag{28}$$

$$\{\alpha_1, \alpha_2\} \textbf{ observes } disband(\alpha_1, \alpha_2) \tag{29}$$

$$\{\alpha_3\} \textbf{ observes } disband(\alpha_1, \alpha_2) \textbf{ if } attentive(\alpha_3) \tag{30}$$

A single agent may $release$ another agent causing him to no longer be bound. A pair of agents working together may $subdue$ an agent, causing him to be bound.

$$release(\alpha_1, \alpha_2) \textbf{ causes } \neg bound(\alpha_2) \tag{31}$$

$$subdue(\alpha_1, \alpha_2, \alpha_3) \textbf{ causes } bound(\alpha_3) \tag{32}$$

$$\textbf{executable } subdue(\alpha_1, \alpha_2, \alpha_3) \textbf{ if } united(\alpha_1, \alpha_2) \vee united(\alpha_2, \alpha_1) \tag{33}$$

The observation axioms concerning the action $release$ are quite straightforward. As before, the agents directly involved in the action occurrence will be aware of it, as will any attentive agents. The same holds true for the action $subdue$.

$$\{\alpha_1, \alpha_2\} \textbf{ observes } release(\alpha_1, \alpha_2) \tag{34}$$

$$\{\alpha_3, \} \textbf{ observes } release(\alpha_1, \alpha_2) \textbf{ if } attentive(\alpha_3) \tag{35}$$

$$\{\alpha_1, \alpha_2, \alpha_3\} \textbf{ observes } subdue(\alpha_1, \alpha_2, \alpha_3) \tag{36}$$

The representation of the action $escape$ is fairly straightforward. Intuitively, once an agent has escaped, he is free. From the domain description, we know that an agent (in this case A), may only escape once his captor has been subdued.

$$escape(\alpha) \textbf{ causes } free(\alpha) \tag{37}$$

$$\textbf{executable } escape(\alpha_1) \textbf{ if } captor(\alpha_2, \alpha_1) \wedge bound(\alpha_2) \wedge (\neg united(\alpha_1, \alpha_3) \vee \neg united(\alpha_3, \alpha_1)) \tag{38}$$

The agent who escapes will be aware of this action occurrence, as will any attentive agents.

$$\{\alpha\} \textbf{ observes } escape(\alpha) \tag{39}$$

$$\{\alpha_2\} \textbf{ observes } escape(\alpha_1) \textbf{ if } attentive(\alpha_2) \tag{40}$$

Lastly, an agent may tell another agent some facts about the domain. The action $tell$ is a communication action, and is represented by the following communication axiom:

$$tell(\alpha_1, \alpha_2, \varphi) \textbf{ communicates } \varphi \tag{41}$$

where φ is of the form $allies(\alpha_1, \alpha_2)$. Agents may eavesdrop however, and therefore in the Escapee Domain, communication must be done with caution. For this domain, we assume that attentive agents are fully aware of what is said between their fellows.[3] This assumption is described by the following observation axioms:

$$\{\alpha_1, \alpha_2\} \textbf{ observes } tell(\alpha_1, \alpha_2, \varphi) \tag{42}$$

$$\{\alpha_3, \} \textbf{ observes } tell(\alpha_1, \alpha_2, \varphi) \textbf{ if } attentive(\alpha_3) \tag{43}$$

[3]Other domains, may call for different assumptions. For example, in another less "paranoid" domain, agents who are attentive may be considered only *partially aware* of occurrences of communication actions.

The Initial State: Let us assume that, initially, it is common knowledge amongst the agents that: A is bound, and that all of the agents are attentive. Furthermore, let us assume that A does not know/believe that C is an ally, and that this fact is not known to B. This information could be encoded by the following initial state axioms:

$$\textbf{initially } \mathbf{C}_{\{A,B,C\}} attentive(\alpha) \wedge bound(A) \tag{44}$$

$$\textbf{initially } \mathbf{C}_{\{A,B,C\}}(\neg \mathbf{B}_A allies(A, C) \wedge \neg \mathbf{B}_B allies(A, C)) \tag{45}$$

$$\textbf{initially } \mathbf{B}_C allies(A, C) \tag{46}$$

where $\alpha \in \{A, B, C\}$. For the rest of this section, let σ_0 be a belief state of the action description consisting of the Escapee domain and the initial state axioms (44)–(46).

Private Announcements: Agent A and agent C are allies. Before they can collaborate towards his escapee, C must confirm his identity to agent A. In the approach taken by (Baltag and Moss 2004; van Benthem, van Eijck, and Kooi 2006), this would correspond to the following action: "C tells A that they are allies, with agent B oblivious of what has transpired." In, $m\mathcal{A}+$, a *private announcement* such as this would be achieved by the following plan, Π_4:

$$[distract(C, B), tell(C, A, allies(A, C))]$$

The first action, $distract(C, B)$, is an ontic action which alters the value of the perspective fluent $attentive(B)$. As a consequence, B is no longer paying attention to his environment, and is therefore oblivious of any subsequence action occurrences. Once this has been done, C can freely inform A that they are allies, without fear of being overheard by B.

Let φ_3 be the following belief formula:

$$\mathbf{C}_{\{A,C\}}(\neg attentive(B)) \wedge \mathbf{C}_{\{A,C\}}(allies(A, C)) \wedge \mathbf{C}_{\{A,C\}}(\neg \mathbf{B}_B allies(A, C)) \wedge \neg \mathbf{B}_B allies(A, C)$$

Proposition 4 states that the successor belief state achieved after the execution of Π_4 matches our intuition about the effects of a private announcement – namely that the agents involved know what has been communicated while the beliefs of oblivious agents are unchanged.

Proposition 4 $(\mathcal{D}, \sigma_0) \models \varphi_3$ **after** Π_4.

Private World Altering Actions: Alternatively, before revealing himself to A, C may seek to secretly free him from his confinement. In other words: "C releases A without the knowledge of B." A *private world altering action* such as this could be accomplished by the following plan Π_5:

$$[distract(C, B), release(C, A)]$$

Let φ_4 be the following belief formula:

$$\mathbf{C}_{\{A,C\}}(\neg attentive(B)) \wedge \mathbf{C}_{\{A,C\}}(\neg bound(A)) \wedge \mathbf{C}_{\{A,B,C\}} \mathbf{B}_B bound(A)$$

Alternatively, let Π_6 be the following plan:

$$[distract(C, B), tell(C, A, allies(A, C)), release(C, A)]$$

which corresponds to agent C releasing A after he secretly informs him that they are allies.

Propositions 5 and 6 describe the belief state resulting from the execution of Π_5 and Π_6, respectively; these match our intuition concerning the effects of private world altering actions—namely, that the agents who are aware of the occurrence are aware of its effects, while the beliefs of oblivious agents remain unchanged.

Proposition 5 $(\mathcal{D}, \sigma_0) \models \varphi_4$ **after** Π_5

Proposition 6 $(\mathcal{D}, \sigma_0) \models \varphi_4 \wedge \varphi_3$ **after** Π_6

Collaboration: In general, agents often need to cooperate in order to achieve their goals. In the Escapee Domain, agents A and C need to work together in order to subdue B. In order to collaborate, agents need to unite, thereby forming a group. Once a group has been formed, the collaborative activity is performed, and then the group is dissolved. This process can itself be modeled as a sequence of elementary actions. We begin by examining group formation.

Group Formation: Group formation may be done publicly or privately, once again depending on the frames of reference of the respective agents. In this particular domain, we are interested in *private group formation* of the form "agents A and C get together to collaborate without the knowledge of agent B." The action $unite$ in this domain is interesting, in that is has a complex executability condition, requiring that the formula $\neg bound(\alpha_1) \wedge \neg bound(\alpha_2) \wedge \mathbf{B}_{\alpha_1} allies(\alpha_1, \alpha_2)$ be true in the state in which the action occurrence takes place. In addition, in order for the escape to be successful, this must be done in secret. Observe that group formation is a form of world changing (ontic) actions, in that it only affects the values of perspective fluents.

Let Π_7 be the plan:

$$[distract(C, B), tell(C, A, allies(A, C)), \\ release(C, A), unite(A, C)]$$

Π_7 corresponds to agents A and C uniting together after B has been distracted, and C has revealed himself to A; and let

$$\varphi_5 = \mathbf{C}_{\{A,C\}} united(A, C) \wedge \mathbf{C}_{\{A,B\}} \mathbf{B}_C bound(A)$$

Proposition 7 $(\mathcal{D}, \sigma_0) \models \varphi_5$ **after** Π_7

It should be clear that *public group formation* could be modeled by performing the singleton action $unite(\alpha_1, \alpha_2)$ in a belief state where all other agents are attentive.

Collaborative Action: Once agents have united together, they may perform *collaborative actions*. In this domain, $subdue(A, C, B)$ is an action which requires the collaboration on the part of agents A and C. If the action occurs in a state which satisfies $\neg attentive(B)$, the plan consisting of the single action $[subdue(A, C, B)]$ models *private collaboration*, otherwise it represents the *public collaboration* of agents A and C. In order for A and C to subdue B, they must work together in secret. Let Π_8 be the following sequence of elementary actions:

$$[distract(C, B), tell(C, A, allies(A, C)), \\ release(C, A), unite(A, C), subdue(A, C, B)]$$

$subdue(A, C, B)$ is executable after the execution of:

$$[distract(C, B), tell(C, A, allies(A, C)), \\ release(C, A), unite(A, C)]$$

In addition:

Proposition 8 $(\mathcal{D}, \sigma_0) \models \mathbf{C}_{\{A,B,C\}}(bound(B))$ **after** Π_8

Group Dissolution: After the collaborative action is done, agents may disband their group, and operate separately. As for previous actions, this may be *private* or *public* depending on the frames of reference of the agents who are not involved in the action occurrence. Dissolution, like its counterpart of group formation is essentially a world changing, affecting the value of the perspective fluent $united$.

Let Π_9 be the plan obtained by appending the action $disband(A, C)$ to Π_8:

$$[distract(C, B), tell(C, A, allies(A, C)), release(C, A), \\ unite(A, C), subdue(A, C, B), disband(A, C)]$$

Proposition 9 $(\mathcal{D}, \sigma_0) \models \mathbf{C}_{\{A,C\}} \neg united(A, C)$ **after** Π_9

Once the group has been dissolved, agent A may again perform a *private world changing action* and escape. Propositions 7, 8 and 9, all state the effects of the collaborative action between A and C, and match our intuition regarding their private, collaboration on a world altering action.

Conclusions and Final Remarks

In this paper, we used a high-level action language, $m\mathcal{A}+$, for representing and reasoning about dynamic multi-agent domains to show how various multi-agent actions can be expressed as sequence of simple actions in the presence of perspective fluents and perspective actions. We used the Strongbox and developed a new Escapee Domain. In particular, we showed how previously complex actions, such as the private announcement action of "agent C informs agent A that $allies(A, C)$ is true, while agent B is unaware of what has transpired" could be easily represented in a more elaboration tolerant fashion as a sequence of elementary actions involving the manipulation of *perspective fluents* on the part of the agents. We then showed how various multi-agent interactions (e.g., group formation/dissolution; collaborative action) could be represented in a similar manner. We note that the use of perspective fluents in this paper divides the agents in a domain into three disjoint categories while, in practice, the number of categories could be higher. In this sense, $m\mathcal{A}+$ is not as expressive as the well-known action models formalization. We conjecture that by allowing the perspective fluents to be nested, this issue will be address and plan to address it in the near future.

It is our belief, that the introduction of such fluents and actions, not only makes the task of knowledge representation simpler, but also makes the language more readily applicable towards the planning and plan execution problems for agents in multi-agent environments. In particular, the simple canonical "public announcement" action in multi-agent systems is not so trivial to execute when agents may be inattentive. In that case, first they have to be made attentive and only then something can be publicly announced. This amounts to a sequence of elementary actions. Since most existing planning algorithms are about finding action sequences, our view of multi-agent actions (as sequence of elementary actions) in this paper makes them more amenable to planning than the use of action/update models.

References

Baltag, A., and Moss, L. S. 2004. Logics for Epistemic Programs. *Synthese* 139(2):165–224.

Baral, C., and Gelfond, G. 2011. On Representing Actions in Multi-Agent Domains. In Balduccini, M., and Son, T. C., eds., *Proceedings of the Symposium on Constructive Mathematics*, 213–232. Springer.

Baral, C.; Gelfond, G.; Son, T. C.; and Pontelli, E. 2012. An Action Language for Reasoning about Beliefs in Multi-Agent Domains. In *Proceedings of the 14th International Workshop on Non-Monotonic Reasoning*.

Baral, C.; Gelfond, G.; Son, T. C.; and Pontelli, E. 2013. Reasoning about the Beliefs of Agents in Multi-Agent Domains in the Presence of State Constraints: The Action Language mAL. In Leite, J. a.; Son, T. C.; Torroni, P.; Torre, L.; and Woltran, S., eds., *Computational Logic in Multi-Agent Systems*, volume 8143 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 290–306.

Bolander, T., and Andersen, M. 2011. Epistemic Planning for Single and Multi-Agent Systems. *Journal of Applied Non-Classical Logics* 21(1).

Gelfond, M., and Lifschitz, V. 1998. Action Languages. *Electronic Transactions on AI* 3.

van Benthem, J.; van Eijck, J.; and Kooi, B. 2006. Logics of communication and change. *Information and Computation* 204(11):1620–1662.

van Ditmarsch, H.; van der Hoek, W.; and Kooi, B. 2008. *Dynamic Epistemic Logic*. Springer.