Dynamic logic on normal form games

R. Ramanujam and Sunil Simon

The Institute of Mathematical Sciences
C.I.T. Campus, Chennai 600 113, India.
E-mail: {jam,sunils}@imsc.res.in

Abstract. We consider a dynamic logic of game composition, where atomic games are in normal form. We suggest that it is useful to consider a modality indexed by game - play pairs. We show how we can reason not only about notions like strategy comparison, dominated strategies and equilibrium in such a framework, but also strategic response, whereby the choice of a player depends on plays observed in the past. This makes for a significant difference in the presence of unbounded iteration. We present a complete axiomatization of the logic and prove its decidability.

1 Overview

The central innovation introduced by game theory is its strategic dimension. A player's environment is not neutral, and she expects that other players will try to outguess her plans. Reasoning about such expectations and strategizing one's own response accordingly constitutes the main logical challenge of game theory.

Games are defined by sets of rules that specify what moves are available to each player, and every player plans her strategy according to her own preferences over the possible outcomes. In an extensive form game, the moves of players are explicitly presented and therefore strategies are not abstract atomic entities, but have a certain structure associated with them. The choice of which strategy to employ depends not only on the game structure but also on her expectation of what strategy other players choose. Thus at any game position, the past as well as the possible futures and players' expectations determine strategies.

In contrast, strategies are presented in an abstract way in a normal form game and the reasoning in such a game is driven by outcome specifications. Thus normal form games can be viewed as extensive form games abstracted into a tree of depth one, where edges are labelled by a tuple of strategies, one for each player, and thus strategies are atomic. Therefore there is no past and future that strategies refer to, and we only speak of notions like rational response, dominant strategies, equilibrium and so on. However, when we consider repeated games, or games composed of smaller games, the notion of strategic response of a player to other players' moves becomes relevant, pretty much in the same way as it is used in extensive form games. History information, as well as epistemic attitudes of players become relevant.

There have been several logical studies from this viewpoint. Notable among these is the work on alternating temporal logic (ATL) [AHK02] which considers selective quantification over paths that are possible outcomes of games in which players and an environment alternate moves. An ATL model is a concurrent game structure which consists of a single game arena whose edges correspond to concurrent moves of the players. Moves in the arena can therefore be thought of as a strategy profile of an appropriate normal form game. Thus each game position of the arena is associated with a single normal form game. The formulas of ATL make assertions about the tree unfolding of this arena. The emphasis is on the existence of a strategy for a coalition of players to force an outcome. Since the game tree encodes the past information, the logic itself can be extended with past modalities as well as knowledge modalities in order to reason about the history information and epistemic conditions used in strategizing by players ([JvdH04],[vdHW02]). Extensions of ATL where strategies are allowed to be named and referred to in the formulas of the logic are proposed in ([vdHJW05],[WvdHW07]). ([Ago06], [Bor07]) extends ATL with the ability to specify actions of players in the formulas.

The running thread in this line of work is the notion of strategies as being atomic, lacking structure. In principle, since concurrent actions (strategy profiles) can be named, their labels can be used to refer to strategies, and hence we can speak of b being a response to a in the past, to achieve α. However, the tree models carry only temporal information, and strategies lack syntactic structure, and this is reflected in reasoning. Thus ATL-based logics can be seen as analogous to temporal logics (for games), as opposed to dynamic and process logics.

When games are themselves structured, strategic response reflects such structure as well. For games of bounded length, an action labelled modal logic reflects game and strategy structure well, but when we consider unbounded play as arising from unbounded repetition of games, the situation is different. This is the spirit in which game logic [Par85] was proposed and underlying framework of coalition logic [Pau01]. The strategies used by a player in such a composite game would depend on not just the outcome specification but also what strategy was used, especially by opponents, in the past. The history information can then be analysed by taking into account the underlying structure of the composite game.

We suggest that in reasoning about structured games, it is useful for the strategies of players to also reflect the structure. Thus rather than reasoning about the strategies in the composed game, one should look at strategies in the atomic game and compose such atomic game strategy pairs.

Suppose that we have a 2-player 2-stage game g_1 followed by g_2. Consider player 1 strategizing at the end of g_1, when g_2 is about to start; her planning depends not only how g_2 is structured, but also how her opponent had played in g_1, and the outcomes that resulted in g_1 for both of them. Thus her strategizing in the composite game $g_1; g_2$ is best described as follows: consider g_1 in extensive form as a tree, and the subtree obtained by the set of plays η_1; when g_2 starts from any of the leaf nodes of this subtree, consider the play η_2. We encode this as $(g_1, \eta_1); (g_2, \eta_2)$, and see (g_2, η_2) as a response to (g_1, η_1). Thus the "programs" of this logic are game - play pairs of this kind.

For extensive form games, this was done in [RS08], where we look at strategic reasoning done in extensive form games by making explicit use of the structure of the game tree. It defines a propositional dynamic logic where programs are regular expressions over game strategy pairs. This gives the ability to reason about the strategic response of players based on what happened in the past. A complete axiomatization of the dynamic logic is presented and the decidability of the logic is also shown. This paper proves similar results for composition of normal form games.

We consider composition of game play pairs in normal form games, corresponding to the fact that the reasoning performed in single stage is mostly outcome based. If we restrict the reasoning to bounded repetition of games or to multistage games where the number of stages are bounded, then we do not need to look at composition of game play pairs. It is the presence of unbounded iteration of games which makes it necessary to introduce a dynamic structure on game play pairs. We therefore study a dynamic logic where programs consist of regular expressions over game play pairs in normal form games. While the main technical result is a complete axiomatization for the logic, the central objective of the paper is to highlight the logical differences between composition of normal form games and that of extensive form games, in terms of the reasoning involved.

We wish to emphasize that what we study here is really a dynamic logic of tree composition. When we consider only bounded games, the logic is subsumed by the ATL frameworks, but the class of games with unbounded iteration studied in Section 5 is our main object of study. However, rather than presenting it all at one go, we discuss strategic response for bounded games before considering repetition.

In the case of extensive form games, the idea of taking into account the structure available within strategies and making assertions about a specific strategy leading to a specified outcome is developed in [vB01,vB02], where van Benthem uses dynamic logic to describe games as well as strategies. [Gho08] presents a complete axiomatisation of a logic describing both games and strategies in a dynamic logic framework where assertions are made about atomic strategies. The techniques developed in [Gho08] can be easily transferred to normal form games. Our point of departure from this line of work is in talking about the strategic response of players in the logical framework.

2 Preliminaries

Normal form games

Let $N = \{1, 2\}$ be the set of players, Σ_i for $i \in \{1, 2\}$ be a finite set of action symbols which represent moves of players and $\Sigma = \Sigma_1 \times \Sigma_2$. For each player i, let R_i be the finite set of rewards, $\preceq^i \subseteq R_i \times R_i$ be a preference ordering on R_i and let $R = R_1 \times R_2$.

Normal form (or strategic form) games are one shot games where the strategies of players corresponds to choosing an action from the action set. A strategy profile is simply a pair of actions, one for each player. A play of the game corresponds to each player choosing an action simultaneously without knowledge of the action picked by the other player. Thus a strategy profile constitutes a play in the game. Each play is associated with a pair of rewards for the players, the outcome of the play.

Suppose $|\Sigma_1| = m$ and $|\Sigma_2| = k$, then a strategic form game can be represented as an $m \times k$ matrix A where the actions of player 1 constitute the rows of the matrix and that of player 2 the columns. The matrix entries specify the outcome of the play for each player, i.e. elements from R. An example game is given in Fig. 1. Here $\Sigma_1 = \{b, c\}$ and $\Sigma_2 = \{x, y\}$. The action profile (b, x) where player 1 chooses to play b and player 2 chooses x, results in the reward r_1^1 for player 1 and r_2^1 for player 2.

	x	y
b	(r_1^1, r_2^1)	(r_1^2, r_2^2)
c	(r_1^3, r_2^3)	(r_1^4, r_2^4)

Fig. 1. Matrix game

Let $\bar{\imath} = 2$ when $i = 1$ and $\bar{\imath} = 1$ when $i = 2$. Unless specified, we will use the convention that $i = 1$ and $\bar{\imath} = 2$. We use b, c to denote the actions of player i, x, y to denote the actions of player $\bar{\imath}$ and a to denote the strategy profile.

Strategy comparison and equilibrium

For player i, the ordering $\preceq^i$ on the rewards R_i induces an ordering on the strategy profiles as: $(b, x) \preceq^i (c, y)$ iff $A(b, x)[i] \preceq^i A(c, y)[i]$. Having defined the preference ordering on strategy profiles, the various game theoretic notions of interest include:

- Weak domination: We say a strategy b of player i weakly dominates a strategy c if for all $x \in \Sigma_{\bar{\imath}}$, $(c, x) \preceq^i (b, x)$.
- Best response: Given strategies b and x of player i and $\bar{\imath}$ respectively, we say that b is the best response for x iff for all $c \in \Sigma_i$, $(c, x) \preceq^i (b, x)$. A similar definition can be given for the best response of player $\bar{\imath}$.
- Equilibrium: A strategy profiles (b, x) constitutes a Nash equilibrium iff b is the best response for x and x is the best response for b.

3 Reasoning in strategic form games

As opposed to extensive form games where the game structure is explicit, normal form games are specified by the set of abstract strategies and the outcomes. In this scenario a player cannot strategize based on the past moves of the opponent. Strategizing would rather be based on his expectation of what strategies the opponent will choose, along with the outcomes which can be ensured by the player. In this section we look at how to logically reason about such abstract strategies with respect to the outcomes.

For a logical analysis, it is convenient to view the normal form game as a tree of depth one, where the edges are labelled by pairs of actions, one for each player. Formally $g = (S, \longrightarrow, s_0, \lambda)$ where S is the set of states, s_0 is the root of the tree. The transition function $\longrightarrow: s_0 \times \Sigma \to S$ is a partial function also called the move function. The reward function $\lambda : S \to (R_1 \times R_2)$. For a node $s \in S$, let $\vec{s} = \{a \in \Sigma \mid \exists s' \in S \text{ where } s \xrightarrow{a} s'\}$ and $\Sigma^g = \{a \in \Sigma \mid \exists s, s' \in S \text{ where } s \xrightarrow{a} s'\}$. Thus for a game g, the set Σ^g constitutes all the strategy profiles of g.

The game tree corresponding to the strategic form game in Fig. 1 is shown in Fig. 2(i). A play is simply an edge in the tree, this corresponds to both the players picking an action. A strategy for player i is the subtree of g where for player i a unique action is chosen and for player $\bar{\imath}$ all the actions are taken into account. A strategy for player 1 in the game given in Fig. 2(i) where he picks action "b", is shown in Fig. 2(ii). For the rest of the paper, we will use the tree representation for strategic games. For a pair $a = (b, x)$ and $j \in \{1, 2\}$, we denote by $a[j]$ the j^{th} component of a.

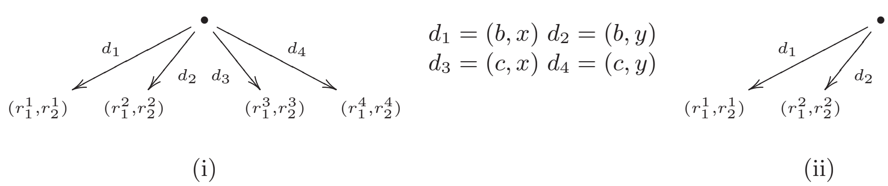

Fig. 2. Normal form game

The logic

Syntax: Let P be a countable set of propositions and g be a strategic form game (in the tree representation). The syntax of the logic is given by:

$$\Phi := p \in P \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid \langle g, \eta\rangle^{\forall}\alpha$$

where $\eta \subseteq \Sigma^g$.

For the moment we let the semantic tree structure g be part of the syntax of the formulas. In section 4 we show how game trees can be specified in the logic in a syntactic manner. For a game g (specified as a tree of depth one), $\eta \subseteq \Sigma^g$ represents a set of plays in g. The intuitive meaning of the construct $\langle g, \eta \rangle^{\forall} \alpha$ is to say that the formula α holds at all nodes which results from the plays of g specified by η. Note that strategies of a particular player can be easily represented in η. This can be done by fixing a single action for the player and considering all plays in g where this action is fixed.

Semantics: The model $M = (g, V)$ where $g = (S, \longrightarrow, s_0, \lambda)$ is a normal form game and $V : S \to 2^P$ is a valuation function. To be able to perform strategic reasoning in the logic, we need to refer to rewards of the players in the formula. This is taken care of by using special propositions to code them up, in the spirit of the approach taken in [Bon02]. The preference ordering is then simply inherited by the implication available in the logic. Formally, let $R_1 = \{r_1^1, \ldots, r_1^l\}$ be the set of rewards for player 1. Without loss of generality we assume that $r_1^1 \preceq^1 r_1^2 \preceq^1 \cdots \preceq^1 r_1^l$. Let $\Theta_1 = \{\theta_1^1, \ldots, \theta_1^l\}$ be a set of special propositions used to encode the rewards in the logic, i.e. θ_1^j corresponds to the reward r_1^j. Likewise for player 2, corresponding to the set R_2, we have a set of propositions Θ_2. The valuation function satisfies the condition:

- For all states s, for all $i \in \{1, 2\}$, $\{\theta_i^1, \ldots, \theta_i^j\} \subseteq V(s)$ iff $\lambda(s)[i] = r_i^j$.

The truth of a formula $\alpha \in \Phi$ in the model M at a position s (denoted $M, s \models \alpha$) is defined as follows:

- $M, s \models p$ iff $p \in V(s)$.
- $M, s \models \neg\alpha$ iff $M, s \not\models \alpha$.
- $M, s \models \alpha_1 \vee \alpha_2$ iff $M, s \models \alpha_1$ or $M, s \models \alpha_2$.
- $M, s \models \langle g, \eta \rangle^{\forall} \alpha$ iff s is not a leaf node and $\forall s' \in \mathit{tail}(g, \eta)$, $M, s' \models \alpha$.

where for game g and $\eta \in 2^{\Sigma}$, $\mathit{tail}(g, \eta) = \{s' \mid s_0 \xrightarrow{a} s' \text{ and } a \in \eta\}$.

When $\langle g, \eta \rangle^{\forall} \alpha$ is asserted at the root node s_0 of the game tree g, we get the following interpretation: $\langle g, \eta \rangle^{\forall} \alpha$ holds iff α holds at all leaf nodes resulting from plays specified by η. Since we are working with a single tree of depth one, interpreting $\langle g, \eta \rangle^{\forall} \alpha$ at the leaf nodes does not make sense. The dual modality $[g, \eta]^{\exists} \alpha$, would say that there exists a play of g specified in η such that α holds at the leaf node of the play.

Strategy comparison in the logic: We show that the various strategizing notions discussed in the earlier section can be expressed in the logic. For a game g, let $\Sigma^g = \{a_1, \ldots, a_k\}$ be the strategy profiles occurring in g. For $i \in \{1, 2\}$, let $\Sigma_i^g = \{a_1[i], \ldots, a_k[i]\}$ and for $b \in \Sigma_i^g$, let $\Sigma_g(b) = \{a \in \Sigma^g \mid a[i] = b \text{ and } a[\overline{\imath}] \in \Sigma_{\overline{\imath}}^g\}$. $\Sigma_g(b)$ thus consists of all the strategy profiles where player i's strategy is fixed to b. Consider the formula:

$$\mathit{ensures}^i(g, \gamma) \equiv \bigvee_{b \in \Sigma_i^g} \langle g, \Sigma_g(b) \rangle^{\forall} \gamma.$$

$ensures^i(g,\gamma)$ says that given that the opponent chooses an action from the set $\Sigma^g_{\bar{\imath}}$, there is a strategy for player i to achieve γ no matter what choice player $\bar{\imath}$ makes. In the case of $\gamma \in R_i$, this corresponds to the rewards that player i can ensure. If player i expects that $\bar{\imath}$ will choose only actions from the set $\Sigma' \subseteq \Sigma^g_{\bar{\imath}}$, then the restriction of $ensures^i(g,\gamma)$ to Σ' specifies what player i can ensure in terms of his expectation. A player during the phase of strategizing might take into consideration what he can ensure given his expectation about the strategies of the opponent. The related concept of weakly dominating strategies can be defined as follows:

$$DOM^i(b,b') \equiv \bigwedge_{x\in\Sigma^g_{\bar{\imath}}} \bigwedge_{\theta_i\in\Theta_i} \Big(\langle g,(b',x)\rangle^{\forall}\theta_i \supset \langle g,(b,x)\rangle^{\forall}\theta_i\Big).$$

This says that whatever reward that can be ensured using the strategy b' can also be ensured with the strategy b. In other words, this says that for player i, the strategy b weakly dominates b'.

Given a strategy x of player $\bar{\imath}$ we can express the fact that the strategy b is better than b' for player i using the formula

$$Better^i_x(b,b') \equiv \textstyle\bigwedge_{\theta_i\in\Theta_i}(\langle g,(b',x)\rangle^{\forall}\theta_i \supset \langle g,(b,x)\rangle^{\forall}\theta_i)$$

We can express b is the best response of player i for x as $BR^i_x(b) \equiv \bigwedge_{b'\in\Sigma^g_i} Better^i_x(b,b')$.

Having defined best response, the fact that a strategy profile (b,x) constitutes an equilibrium can be expressed as: $EQ(b,x) \equiv BR^i_x(b) \wedge BR^{\bar{\imath}}_b(x)$.

	d^2	c^2
d^1	(P^1,P^2)	(T^1,S^2)
c^1	(S^1,T^2)	(R^1,R^2)

Fig. 3. Prisoner's Dilemma

Example: Consider the prisoner's dilemma game given in Fig. 3. Let the actions c and d correspond to cooperate and defect respectively. The preference ordering over the rewards for $i \in \{1,2\}$ is given by $S^i \preceq^i P^i \preceq^i R^i \preceq^i T^i$. Let the propositions representing the rewards be $\{\theta^S_i, \theta^P_i, \theta^R_i, \theta^T_i\}$. Consider the formulas:

- $\alpha_1 \equiv \langle g,(c^1,d^2)\rangle^{\forall}\theta^S_1 \supset \langle g,(d^1,d^2)\rangle^{\forall}\theta^S_1$.
- $\alpha_2 \equiv \langle g,(c^1,c^2)\rangle^{\forall}\theta^R_1 \supset \langle g,(d^1,c^2)\rangle^{\forall}\theta^R_1$.

The formula α_1 holds since we have $S^1 \preceq^1 P^1$ and α_2 holds since $R^1 \preceq^1 T^1$. The formula $\alpha_1 \wedge \alpha_2$ states that irrespective of the move made by player 2, it is better for player 1 to choose d^1. In other words, "defect" is a dominant strategy for player 1 in this game.

α_1 says that the strategy d^1 is better than c^1 for player 1 against the strategy d^2 of player 2. Since there are only two strategies available for player 1, we get that d^1 is the best response for d^2. A similar reasoning with respect to player 2 shows that d^2 is the best response for d^1. From which we get that the strategy profile (d^1, d^2) constitutes an equilibrium profile.

It can be seen, that quite a lot of reasoning that is done in the case of normal form games can be captured by considering game play pairs. The game play pairs in effect, provides us the power of reasoning about restrictions of the full game tree and the ability to compare various such restrictions in terms of the outcomes they guarantee. A player can thus make use of notions like dominant strategy, guaranteed outcome, best response and so on to come up with an appropriate plan of action for the game. The important strategizing notion which is missing in this approach is that of strategic response of a player to the opponent's action. To capture this aspect we need to move over to a model where instead of working with a fixed normal form game, we have a finite set of games and where composition of these games can be performed.

4 Strategic response

For the sake of clarity, in subsequent sections, we concentrate on the structure of the game g with respect to the moves of the players and disregard the rewards associated in the game structure. In section 7, after the logic is presented in full generality, we mention the changes required to take care of the rewards present in the game.

Since formulas of the logic refer to the normal form game trees, we first present a syntax for representing such trees.

Syntax for strategic form games: Let *Nodes* denote a finite set of nodes, the strategic form game tree is specified using the syntax:

$$\mathrm{G} := \Sigma_{a_m \in J}(x, a_m, y_m).$$

where $x, y_m \in \mathit{Nodes},\ J = J_1 \times J_2$ for $J_2 \subseteq \Sigma_1$ and $J_2 \subseteq \Sigma_2$.

The game tree T_g generated by the game $g \in \mathrm{G}$ is defined as follows. Let $g = (x, a_1, y_1) + \ldots + (x, a_k, y_k),\ T_g = (S_g, \Longrightarrow_g, s_{g,0})$ where

- $S_g = \{s_x, s_{y_1}, \ldots, s_{y_k}\}$ and $s_{g,0} = s_x$.
- For $1 \leq j \leq k$ we have $s_x \overset{a_j}{\Longrightarrow}_g s_{y_j}$.

Syntax: In addition to the set of propositions, let $\mathcal{G} \subseteq \mathrm{G}$ be a finite set of games. The syntax of the logic is very similar to what was presented earlier:

$$\Phi := p \in P \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid \langle g, \eta \rangle^{\forall} \alpha$$

where $g \in \mathcal{G}$ and $\eta \subseteq \Sigma^g$.

Models: Formulas of the logic express properties about normal form game trees and plays in the game. Since the modality $\langle g, \eta \rangle^{\forall}$ can be nested, we are in effect talking about finite trees which are generated by composing individual game trees. However there can be an infinite set of finite game trees. One way of giving a finite presentation is to think of the tree being obtained by unfolding of a Kripke structure. As we will see later, the logic cannot distinguish between these two. A model $M = (W, \longrightarrow, V)$ where W is the set of states (or game positions), the relation $\longrightarrow \subseteq W \times \Sigma \times W$ and $V : W \to 2^P$ is a valuation function.

Semantics: The truth of a formula $\alpha \in \Phi$ in a model M and a state u is defined as in the earlier case. The only difference is in the interpretation of $\langle g, \eta \rangle^{\forall} \alpha$ which is given as:

- $M, u \models \langle g, \eta \rangle^{\forall} \alpha$ iff $enabled(g, u)$ and for all $w \in tail(T_u \upharpoonright g, \eta)$, $M, w \models \alpha$.

Intuitively $enabled(g, u)$ says that the g structure can be embedded at state u of the model, with respect to compatibility with the action labels. $tail(T_u \upharpoonright g, \eta)$ is the set of nodes of the resulting embedded tree when restricted to plays in η. $M, w \models \langle g, \eta \rangle^{\forall} \alpha$ says that firstly g can be embedded at u and if X is the set of all states resulting from the plays specified in η, then the formula α holds in all $w \in X$. The dual $[g, \eta]^{\exists} \alpha$ says: if g can be embedded at the state u then there exists a state w resulting from the plays specified in η such that α holds at w. Formally the tree embedding and the restriction operation is defined below.

Restriction on trees: For $w \in W$, let T_w denote the tree unfolding of M starting at w. Given a state w and $g \in \mathcal{G}$, let $T_w = (S_M^w, \Longrightarrow_M, s_w)$ and $T_g = (S_g, \Longrightarrow_g, s_{g,0})$. The restriction of T_w with respect to the game g (denoted $T_w \upharpoonright g$) is the subtree of T_w which is generated by the structure specified by T_g. The restriction is defined as follows: $T_w \upharpoonright g = (S, \Longrightarrow, s_0, f)$ where $f : S \to S_g$. Initially $S = \{s_w\}$, $s_0 = s_w$ and $f(s_w) = s_{g,0}$.

Let $\{a_1, \ldots, a_k\}$ be the outgoing edges of $s_{g,0}$, i.e. for all $j : 1 \leq j \leq k$, $s_{g,0} \overset{a_j}{\Longrightarrow}_g t_j$. For each a_j, let $\{s_j^1, \ldots, s_j^m\}$ be the nodes in S_M^w such that $s_w \overset{a_j}{\Longrightarrow}_M s_j^l$ for all $l : 1 \leq l \leq m$. Add nodes $s_j^1, \ldots, s_j^m$ to S and the edges $s_0 \overset{a_j}{\Longrightarrow} s_j^l$ for all $l : 1 \leq l \leq m$. Also set $f(s_j^l) = t_j$.

We say that a game g is enabled at w (denoted $enabled(g, w)$) if the tree $T_w \upharpoonright g = (S, \Longrightarrow, s_0, f)$ has the following property:

- $\vec{s_0} = f(\vec{s_0})$.

As an illustration of the restriction operation, consider the game g shown in Fig. 2(i) (disregarding the payoff labels). Let the Kripke structure M be as given in Fig. 4(i). For the node u of the Kripke structure, the restriction $T_u \upharpoonright g$ is shown in Fig. 4(ii)

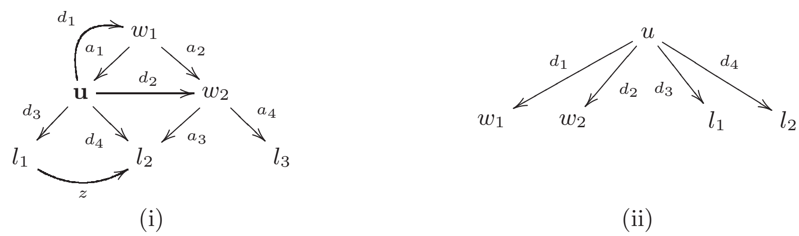

Fig. 4. Restriction

Example: Strategic response of players can easily be expressed in the logic. For instance, in the prisoner's dilemma game, the "tit-for-tat" strategy for player 1 would be to copy the action of player 2 in the earlier stage. This can be represented as: $\langle g,(c^1,d^2)\rangle^{\forall}\langle g,(d^1,d^2)\rangle^{\forall}\alpha \wedge \langle g,(c^1,c^2)\rangle^{\forall}\langle g,(c^1,c^2)\rangle^{\forall}\alpha$.

The above formalism enables reasoning of bounded levels of strategic response by players. The next step would be to look at unbounded iteration or composition of games. This cannot be achieved by the nesting of modalities and therefore the dynamic structure needs to be brought in at the level of the game play pairs.

5 Unbounded game composition

The syntax of the game play pair is enriched as follows:

$$\Gamma := (g,\eta) \mid \xi_1;\xi_2 \mid \xi_1 \cup \xi_2 \mid \xi^* \mid \beta?$$

where $g \in \mathrm{G}$, $\eta \subseteq \Sigma^g$ and $\beta \in \Phi$.

Here we allow g to be any normal form game tree in G (syntax for trees given in section 4). The atomic game play pair (g,η) would have the same interpretation as before. $\xi_1 \cup \xi_2$ would mean playing ξ_1 or ξ_2. Sequencing in our setting does not mean the usual relational composition of games. Rather, it is the composition of game play pairs of the form $(g_1,\eta_1);(g_2,\eta_2)$. A pair (g,η) gives rise to a tree and therefore composition over these trees need to be performed. ξ^* is the iteration of the ';' operator and $\beta?$ tests whether the formula β holds at the current state.

The syntax of the formulas of the logic is given by:

$$\Phi := p \in P \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid \langle\xi\rangle^{\forall}\alpha$$

where $\xi \in \Gamma$.

Models for the logic are Kripke structures as in the earlier case and the semantics remains the same except for the construct $\langle\xi\rangle^{\forall}\alpha$ which is interpreted as:

- $M,u \models \langle\xi\rangle^{\forall}\alpha$ iff $\exists(u,X) \in R_\xi$ such that $\forall w \in X$ we have $M,w \models \alpha$.

For $\xi \in \Gamma$, we have $R_\xi \subseteq W \times 2^W$. The definition of R in the atomic case is same as the interpretation of game play pair used earlier. i.e.:

- $R_{(g,\sigma)} = \{(u, X) \mid \mathit{enabled}(g, u) \text{ and } X = \mathit{tail}(T_u \upharpoonright g, \eta)\}$.

The semantics for composite game strategy pairs is given as follows:

- $R_{\xi_1;\xi_2} = \{(u, X) \mid \exists Y = \{v_1, \ldots, v_k\}$ such that $(u, Y) \in R_{\xi_1}$ and $\forall v_j \in Y$ there exists $X_j \subseteq X$ such that $(v_j, X_j) \in R_{\xi_2}$ and $\bigcup_{j=1,\ldots,k} X_j = X\}$.
- $R_{\xi_1 \cup \xi_2} = R_{\xi_1} \cup R_{\xi_2}$.
- $R_{\xi^*} = \bigcup_{n \geq 0} (R_\xi)^n$.
- $R_{\beta?} = \{(u, \{u\}) \mid M, u \models \beta\}$.

The formulas of the logic can not only make assertions about strategies of players but also about the game structure itself. Thus states of the Kripke structure can be viewed as being associated with a set of atomic normal form games. The restriction operation identifies the specific game under consideration, which in turn is determined by the assertions made by formulas of the logic. Consider the following formula:

- $\langle (g, \eta_2); (g', \eta_1) \rangle^\forall \mathit{win}_1$ where η_2 is a strategy for player 2 in game g and η_1 a strategy of player 1 in g'.

This says that assuming in game g, player 2 plays according to strategy η_2 then in g', player 1 can follow η_1 and ensure win_1. Note that this is not same as saying player 1 can ensure win_1 in the composed game $g = g; g'$. The fact that player 2 employed strategy η_2 in game g is used in strategizing by player 1. However, this specification involves only bounded level of strategic response and can thus be expressed in an ATL like framework extended with the appropriate action modalities and past operators. Consider a construct of the form:

- $((g_1, \eta_1); ((g_2, \eta_2) \cup (g_3, \eta_3)))^*; \mathit{win}_2?; (g, \eta)$
  where η_1, η_2 and η_3 are player 2 strategies in games g_1, g_2 and g_3 respectively and η is a player 1 strategy in game g.

This says that if player 2 can ensure win_2 by iterating the structure g_1 followed by g_2 or g_3 and employing strategies η_1 followed by η_2 or η_3 then player 1 plays according to η in game g. Here not only does player 1 assert that player 2 can ensure win_2 but also makes assertions about the specific game structure that is enabled and the atomic strategies that player 2 employs. Iteration performed here does not correspond to the assertion that a property holds through out the history. To express such properties, one needs to shift from the ATL setting to a dynamic logic framework.

6 Axiom system

We now present an axiomatization of the valid formulas of the logic. We will find the following notations and abbreviations useful.

For $a \in \Sigma$, let g_a denote the normal form game with a unique strategy profile a, we define $\langle a \rangle \alpha$ as:

– $\langle a\rangle\alpha \equiv \langle g_a, \{a\}\rangle^{\forall}\top \wedge [g_a, \{a\}]^{\exists}\alpha$.

From the semantics it is easy to see that for $a \in \Sigma$, this gives the usual semantics for $\langle a\rangle\alpha$, i.e. $\langle a\rangle\alpha$ holds at a state u iff there is a state w such that $u \xrightarrow{a} w$ and α holds at w.

For a game $g = (x, a_1, y_1) + \ldots + (x, a_k, y_k)$, the formula $g^{\checkmark}$ denotes that the game structure g is enabled. This is defined as:

– $g^{\checkmark} \equiv \bigwedge_{j=1,\ldots,k} \langle a_j\rangle\top$.

The axiom schemes

(A1) Propositional axioms:
 (a) All the substitutional instances of tautologies of PC.
(A2) Axiom for single edge games:
 (a) $\langle a\rangle(\alpha_1 \vee \alpha_2) \equiv \langle a\rangle\alpha_1 \vee \langle a\rangle\alpha_2$.
(A3) Dynamic logic axioms:
 (a) $\langle \xi_1 \cup \xi_2\rangle^{\forall}\alpha \equiv \langle \xi_1\rangle^{\forall}\alpha \vee \langle \xi_2\rangle^{\forall}\alpha$.
 (b) $\langle \xi_1; \xi_2\rangle^{\forall}\alpha \equiv \langle \xi_1\rangle^{\forall}\langle \xi_2\rangle^{\forall}\alpha$.
 (c) $\langle \xi^*\rangle^{\forall}\alpha \equiv \alpha \vee \langle \xi\rangle^{\forall}\langle \xi^*\rangle^{\forall}\alpha$.
 (d) $\langle \beta?\rangle^{\forall}\alpha \equiv \beta \supset \alpha$.

For $g = (x, a_1, y_1) + \ldots + (x, a_n, y_n)$ and $\eta \subseteq \Sigma^g$,

(A4) $\langle g, \eta\rangle^{\forall}\alpha \equiv g^{\checkmark} \wedge (\bigwedge_{a\in\eta}[a]\alpha)$.

Inference rules

$$(MP)\ \frac{\alpha, \quad \alpha \supset \beta}{\beta} \qquad (NG)\ \frac{\alpha}{[a]\alpha}$$

$$(IND)\ \frac{\langle \xi\rangle^{\forall}\alpha \supset \alpha}{\langle \xi^*\rangle^{\forall}\alpha \supset \alpha}$$

Since the relation R is synthesised over tree structures, the interpretation of sequential composition is quite different from the standard one. Consider the usual relation composition semantics for $R_{\xi_1;\xi_2}$, i.e. $R_{\xi_1;\xi_2} = \{(u, X) | \exists Y$ such that $(u, Y) \in R_{\xi_1}$ and for all $v \in Y$, $(v, X) \in R_{\xi_2}\}$. It is easy to see that under this interpretation the formula $\langle \xi_1\rangle^{\forall}\langle \xi_2\rangle^{\forall}\alpha \supset \langle \xi_1; \xi_2\rangle^{\forall}\alpha$ is not valid.

7 Completeness

Here we present an overview of the completeness proof for the logic. Details and the full proof can be found in [RS08].

To show completeness, we prove that every consistent formula is satisfiable. Let α_0 be a consistent formula, and $CL(\alpha_0)$ denote the subformula closure of α. Let $\mathcal{AT}(\alpha_0)$ be the set of all maximal consistent subsets of $CL(\alpha_0)$, referred to as atoms. We use u, w to range over the set of atoms. Each $u \in \mathcal{AT}$ is a finite set of formulas, we denote the conjunction of all formulas in u by $\widehat{u}$. For a nonempty subset $X \subseteq \mathcal{AT}$, we denote by $\widetilde{X}$ the disjunction of all $\widehat{u}, u \in X$. Define a transition relation on $\mathcal{AT}(\alpha_0)$ as follows: $u \xrightarrow{a} w$ iff $\widehat{u} \wedge \langle a \rangle \widehat{w}$ is consistent. The valuation V is defined as $V(w) = \{p \in P \mid p \in w\}$. The model $M = (W, \longrightarrow, V)$ where $W = \mathcal{AT}(\alpha_0)$. Once the Kripke structure is defined, the semantics given earlier defines the relation $R_{(g,\eta)}$ on $W \times 2^W$ for $g \in \mathrm{G}$.

However for the completeness theorem, we need to also specify the relation between a pair (u, X) being in $R_{(g,\eta)}$ and the consistency requirement on u and X. This is done in the following lemma:

Lemma 7.1. *For all $g \in \mathrm{G}$, for all $i \in \{1, 2\}$ and for all $\eta \subseteq \Sigma^g$, for all $X \subseteq W$ and for all $u \in W$ the following holds:*

1. *if $(u, X) \in R_{(g,\eta)}$ then $\widehat{u} \wedge \langle g, \eta \rangle^{\forall} \widetilde{X}$ is consistent.*
2. *if $\widehat{u} \wedge \langle g, \eta \rangle^{\forall} \widetilde{X}$ is consistent then there exists $X' \subseteq X$ such that $(u, X') \in R_{(g,\eta)}$.*

Using techniques developed in propositional dynamic logic, the following two lemmas can be shown.

Lemma 7.2. *For all $\xi \in \Gamma$, for all $X \subseteq W$ and $u \in W$, if $\widehat{u} \wedge \langle \xi \rangle^{\forall} \widetilde{X}$ is consistent then there exists $X' \subseteq X$ such that $(u, X') \in R_\xi$.*

Lemma 7.3. *For all $\langle \xi \rangle^{\forall} \alpha \in CL(\alpha_0)$, for all $u \in W$, $\widehat{u} \wedge \langle \xi \rangle^{\forall} \alpha$ is consistent iff there exists $(u, X) \in R_\xi$ such that $\forall w \in X$, $\alpha \in w$.*

Theorem 7.1. *For all $\beta \in CL(\alpha_0)$, for all $u \in W$, $M, u \models \beta$ iff $\beta \in u$.*

The theorem follows from lemma 7.3 by a routine inductive argument.

Decidability: Since $|\Sigma|$ is constant, the size of $CL(\alpha_0)$ is linear in $|\alpha_0|$. Atoms are maximal consistent subsets of $CL(\alpha_0)$, hence $|\mathcal{AT}(\alpha_0)|$ is exponential in the size of α_0. It follows from the completeness theorem that given a formula α_0, if α_0 is satisfiable then it has a model of exponential size. For all $\xi \in \Gamma$ occurring in α_0, the relation R_ξ can be computed in time exponential in the size of the model. Therefore we get that the logic is decidable in nondeterministic double exponential time.

Adding rewards to the game structure: The syntax of game trees presented in section 4 can be easily modified to include the payoff (reward) information for the game. Each node "y_j" needs to be replaced with a tuple of the form $r_j = (r_1^j, r_2^j)$ where $r_j \in R$. Models are Kripke structures $M = (W, \longrightarrow, V, \lambda)$ where $\lambda : W \to R$. For a game g the generated tree $T_g = (S_g, \Longrightarrow_g, \lambda_g, s_{g,0})$. The tree restriction $T_w \upharpoonright g$ (presented in section 4) is therefore a structure of the form $(S, \Longrightarrow, \lambda, s_0, f)$ where $\lambda : S \to R$. The condition for a game g being enabled at a state w needs to capture the rewards of the game as well and therefore needs to be modified as follows:

– $\forall s \in S \setminus \{s_0\}$, $\lambda(s) = \lambda_g(f(s))$.
– $\vec{s_0} = f(\vec{s_0})$.

As mentioned in section 3, let Θ_i be the finite set of special propositions coding up the rewards of players R_i for $i \in \{1, 2\}$. For a game $g = (x, a_1, (r_1^1, r_2^1)) + \ldots + (x, a_k, (r_1^k, r_2^k))$, the enabling of g can be represented in the logic as:

– $g^{\checkmark} \equiv \bigwedge_{j=1,\ldots,k}(\langle a_j \rangle \top \wedge [a_j](\theta_1^j \wedge \theta_2^j))$.

In the axiom scheme (section 6), the following two axioms are added along with the propositional axioms to capture the ordering of the rewards.

– $(\bigvee_{\theta_i \in \Theta_i} \theta_i)$ for $i \in \{1, 2\}$.
– $\bigwedge_{\theta_i^j \in \Theta_i}(\theta_i^j \supset \bigwedge_{k=1,\ldots,j} \theta_i^k)$ for $i \in \{1, 2\}$.

It is easy to check that with the above mentioned modification, the completeness theorem follows.

8 Discussion

By considering game play pairs, we are able to reason about restrictions of the game tree and thereby express game theoretic notions like a player's best response for an opponents strategy and equilibrium. In contrast, the approach taken in [RS08] is closer to the style of game logics: the reasoning is about what a player can ensure by following a certain strategy specification where all possible strategies of the opponent is taken into account. However, at the compositional level, the axiom system remains the same. This shows that the framework being considered is quite general, and is not dependent on the exact game representation. For a specific representation under consideration, once the axioms for the atomic case are presented appropriately, the theory lifts quite neatly.

This paper deals with games of perfect information, since at the end of each stage, all the players know the strategy profile along with the outcomes. It also operates within the framework of foundations for game theory in modal logics. In this sense, it does not try to offer new models for game theory but explicate the reasoning involved. It is worth noting that almost all the analysis performed in reasoning about games, including the related works mentioned earlier, are based on games of perfect information. Coming up with logical formalisms and extending the techniques to reason in games of imperfect information is a challenging task.

To come up with prescriptive mechanisms which provides advice to players on how to play, it is essential to be able to represent a player's expectations about the behaviour of the opponent. The expectations need not necessarily be represented in a probabilistic manner. Introducing expectations of players is particularly interesting in the framework of unbounded game composition as it allows players to learn from the past information, revise their expectations and accordingly make use of it to generate sophisticated plans. Enriching the framework to be able to represent expectations of players is left as future work.

Acknowledgements

We thank the anonymous referees for their helpful comments and suggestions.

References

[Ago06] T. Agotnes. Action and knowledge in alternating time temporal logic. *Synthese*, 149(2):377–409, 2006.

[AHK02] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.

[Bon02] G. Bonanno. Modal logic and game theory: Two alternative approaches. *Risk Decision and Policy*, 7:309–324, December 2002.

[Bor07] S. Borgo. Coalitions in action logic. In *Proceedings IJCAI'07*, pages 1822–1827, 2007.

[Gho08] S. Ghosh. Strategies made explicit in dynamic game logic. In *Logic and the Foundations of Game and Decision Theory*, 2008.

[JvdH04] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2-3):185–219, 2004.

[Par85] R. Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–140, 1985.

[Pau01] M. Pauly. *Logic for Social Software*. PhD thesis, University of Amsterdam, October 2001.

[RS08] R. Ramanujam and S. Simon. Dynamic logic on games with structured strategies. In *The Principles of Knowledge Representation and Reasoning (to appear)*, 2008. `http://www.imsc.res.in/~sunils/papers/pdf/rs-kr08.pdf`.

[vB01] J. van Benthem. Games in dynamic epistemic logic. *Bulletin of Economic Research*, 53(4):219–248, 2001.

[vB02] J. van Benthem. Extensive games as process models. *Journal of Logic Language and Information*, 11:289–313, 2002.

[vdHJW05] W. van der Hoek, W. Jamroga, and M. Wooldridge. A logic for strategic reasoning. *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 157–164, 2005.

[vdHW02] W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems*, pages 1167–1174, 2002.

[WvdHW07] D. Walther, W. van der Hoek, and M. Wooldridge. Alternating-time temporal logic with explicit strategies. In *Theoretical Aspects of Rationality and Knowledge*, pages 269–278, 2007.