

Dynamic Logic on Games with Structured Strategies

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 propositional dynamic logic whose programs are regular expressions over game - strategy pairs. At the atomic level, these are finite extensive form game trees with structured strategy specifications, whereby a player's strategy may depend on properties of the opponent's strategy. The advantage of imposing structure not merely on games or on strategies but on game - strategy pairs, is that we can speak of a composite game g followed by g' whereby if the opponent played a strategy s in g , the player responds with s' in g' to ensure a certain outcome. In the presence of iteration, a player has significant ability to strategise taking into account the explicit structure of games. We present a complete axiomatization of the logic and prove its decidability. The tools used combine techniques from PDL, CTL and game logics.

Overview

Strategies are the unsung heroes of game theory.
Johan van Benthem.

In one sense, game theory is all about strategic reasoning. Games are defined by sets of rules that specify what moves are available to each player, and according to her own preferences over the possible outcomes, every player plans her strategy. If the game is rich enough, the player has access to a wide range of strategies, and the choice of what strategy to employ in a game situation depends not only on the player's understanding of how the game can proceed from then on, but also based on his expectation of what strategies other players are following.

While this observation holds true of much of game playing, we find such reasoning hardly typical of analysis in game theory. In this respect game theory largely consists of reasoning *about* games rather than reasoning *in* games. It is assumed that the entire structure of the game is laid out in front of us, and we reason from above, predicting how rational players would play, and such predictions are summarised into assertions on existence of equilibria. This type of study mostly suffices to focus on existence of strategies for certain outcomes.

And yet, as Aumann and Dreze (2005) point out, this is not how game theory started. The seminal work of von Neu-

mann and Morgenstern envisaged game theory as constituting advice for players in game situations, so that strategies may be synthesized accordingly. While this was summarily achieved for two-person zero-sum games, advice functions for multi-player games with overlapping objectives have been hard to come by. Aumann and Dreze argue that such a prescriptive game theory must account for the beliefs and expectations each player has about strategies followed by other players. Clearly, in any such study, strategies cannot be viewed as unstructured atomic objects arbitrarily picked from a suitably large set, but accorded first class citizenship. That is, they are seen as composite objects, function determined by structure. This calls for a grammar of strategy construction, which in turn depends on the structure of the game in which the strategy is employed.

Strategies with unbounded memory constitute *global* reasoning at the level of the game arena, since, in principle, details about game structure and trajectories of plans can be coded up into them. However, bounded memory strategies can only act *locally*, but can exploit game structure effectively. The maxim, *Think globally, act locally*, is apt for structure sensitive strategizing.

There have been many logical studies in this direction. The work on alternating temporal logic (Alur, Henzinger, and Kupferman 1998) considers selective quantification over paths that are possible outcomes of games in which players and an environment alternate moves. The emphasis is on the existence of a strategy for a coalition of players to force an outcome. In (Harrenstein et al. 2003) and (van der Hoek, Jamroga, and Wooldridge 2005), logics are developed to describe equilibrium concepts and for strategic reasoning. (Chatterjee, Henzinger, and Piterman 2007) looks at a logic where quantification over strategy terms is part of the logical formalism and study its relationship with alternating temporal logic and other variants. All of the above mentioned logics have the common property that the game arena is taken to be fixed and a functional notion of strategy is adopted. Strategies are taken to be atomic objects whereby the logical structure present within the strategy is not taken into account for analysis.

The idea of taking into account the structure available within strategies and making assertions about a specific strategy leading to a specified outcome is, of course, not new. Van Benthem (2001; 2002) uses dynamic logic to de-

scribe games as well as strategies. When dealing with finite extensive form games, this approach of describing the complete strategy explicitly in a dynamic logic framework is appropriate, however the technique does not generalise satisfactorily to games on graphs.

On the other hand, propositional game logic (Parikh 1985), the seminal work on logical aspects of game theory, talks of existence of strategies, but builds composite structure into games. (Goranko 2003) looks at an algebraic characterisation of games and presents a complete axiomatization of identities of the basic game algebra. Pauly (2001) has built on this to provide interesting relationships between programs and games, and to describe coalitions to achieve desired goals. Goranko (2001) relates Pauly’s coalition logics with work done in alternating temporal logic. In this line of work, the game itself is structurally built from atomic objects. However, the reasoning done is about existence of strategies and not reasoning *with* strategies: the ability of a player to strategize in response to the opponent’s actions. (Ghosh 2008) presents a complete axiomatisation of a logic describing both games and strategies in a dynamic logic framework, but again the assertions are about atomic strategies.

In this paper, we make a small contribution to the logical study of games and strategies. We look at a framework where both games and strategies are structurally built and where strategizing by players is explicitly represented in the formulas of the logic. We suggest that considering game - strategy pairs is useful: 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 . 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 opponent employing π ; when g_2 starts from any of the leaf nodes of this subtree, play according to σ . We encode this as $(g_1, \pi); (g_2, \sigma)$, and see (g_2, σ) as a response to (g_1, π) . Thus the “programs” of this logic are game - strategy pairs of this kind.

We consider a propositional dynamic logic, the programs of which are regular expressions over atomic pairs of the form (g, σ) where g is a finite game tree in extensive form, and σ is a strategy specification, structured syntactically. The central syntactic device consists of interactive structure in strategies and algebraic structure not only on games but on game - strategy pairs. While the technical result is a complete axiomatization and the decidability of the satisfiability problem, we see this contribution as an advocacy of studying algebraic structure on strategies, induced by that on games.

Preliminaries

Game tree

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 $\Sigma = \Sigma_1 \cup \Sigma_2$.

Let $(S, \Longrightarrow, s_0)$ be a finite tree rooted at s_0 on the set of vertices S and $\Longrightarrow: (S \times \Sigma) \rightarrow S$. An *extensive*

form game tree is given by $T = (S, \Longrightarrow, s_0, \lambda)$ where S is the set of game positions and s_0 is the initial game position. For a game position $s \in S$, let $\vec{s} = \{s' \in S \mid s \xrightarrow{a} s' \text{ for some } a \in \Sigma\}$. A game position s is a leaf node (or terminal node) if $\vec{s} = \emptyset$, let S^{leaf} denote the set of all leaf nodes of T . The turn function $\lambda: S \rightarrow \{1, 2\}$ associates each game position with a player.

Technically we need player labelling only at the non-leaf nodes. However, for the sake of uniform presentation, we do not distinguish between leaf nodes and non-leaf nodes as far as player labelling is concerned.

Figure 1(a) shows an example game tree. Here nodes are labelled with the players and edges represents the actions. A *play* in T is a finite path $\rho: s_0 \xrightarrow{a_0} s_1 \cdots \xrightarrow{a_k} s_k$ where s_k is a leaf node.

Let $\bar{i} = 2$ when $i = 1$ and $\bar{i} = 1$ when $i = 2$. A *strategy* for player i , is a subtree of T where for each player i node, there is a unique outgoing edge and for player \bar{i} , every move is included. Figure 1(b) shows a strategy for player i in the game tree Figure 1(a). For $i \in \{1, 2\}$, let Ω^i denote the set of all strategies for player i in the game. For a tree T , let $frontier(T)$ denote the set of all leaf nodes of T .




Figure 1: Game and strategy.

The formulas of the logic refer to extensive form game trees. One convenient way of representing the tree is to specify it in the following syntax.

Syntax for game trees: Let $Nodes$ be a finite set. The finite game structure is specified using the syntax:

$$G := (i, x) \mid \Sigma_{a_m \in J} ((i, x), a_m, t_{a_m})$$

where $J \subseteq \Sigma_i$, $x \in Nodes$, $i \in \{1, 2\}$ and $t_{a_m} \in G$.

Given $g \in G$ we define the tree T_g generated by g inductively as follows.

- $g \equiv (i, x)$: $T_g = (S_g, \Longrightarrow_g, \lambda_g, s_{g,0})$ where $S_g = \{s_x\}$, $\lambda_g(s_x) = i$ and $s_{g,0} = s_x$.
- $g \equiv ((i, x), a_1, t_{a_1}) + \cdots + ((i, x), a_k, t_{a_k})$: Inductively we have trees T_1, \dots, T_k where for $j: 1 \leq j \leq k$, $T_j = (S_j, \Longrightarrow_j, \lambda_j, s_{j,0})$. Define $T_g = (S_g, \Longrightarrow_g, \lambda_g, s_{g,0})$ where
 - $S_g = \{s_x\} \cup S_{T_1} \cup \dots \cup S_{T_k}$ and $s_{g,0} = s_x$.
 - $\lambda_g(s_x) = i$ and for all j , for all $s \in S_{T_j}$, $\lambda_g(s) = \lambda_j(s)$.

The edge relation is the union of the edge relation on the individual tree along with the edges $s_x \xrightarrow{a_j} s_{j,0}$ for $j: 1 \leq j \leq k$.

Strategy specification

We give a syntax to specify strategies in a structured manner. Atomic strategy formulas specify, for a player, what conditions she tests for before making a move. We consider the case when these conditions are simply boolean formulas. Composite strategy specifications are built from atomic ones using connectives (without negation). We use an implication of the form: “if the opponent’s play conforms to a strategy π then play σ ”. This connective is crucial to capture the notion of players strategizing in response to opponents actions.

For a countable set of propositions P^i , let $\Psi(P^i)$ be the boolean formulas over P^i built using the following syntax:

$$\Psi(P^i) := p \in P^i \mid \neg\psi \mid \psi_1 \vee \psi_2.$$

For $i \in \{1, 2\}$, let $Strat^i(P^i)$ be the set of strategy specifications given by the following syntax:

$$Strat^i(P^i) := [\psi \mapsto a] \mid \sigma_1 + \sigma_2 \mid \sigma_1 \cdot \sigma_2 \mid \pi \Rightarrow \sigma$$

where $\pi \in Strat^{\bar{i}}(P^1 \cap P^2)$, $\psi \in \Psi(P^i)$ and $a \in \Sigma_i$.

The idea is to use the above constructs to specify properties of strategies. For instance the interpretation of a player i specification $[\psi \mapsto a]^i$ will be to choose move “ a ” for every i node where ψ holds. Consider the game given in Figure 1 (a). Suppose the proposition p holds at the root, then the strategy depicted in Figure 1 (b) conforms to the specification $[\psi \mapsto a]^1$.

The specification $\pi \Rightarrow \sigma$ says, at any node player i sticks to the specification given by σ if on the history of the play, all moves made by \bar{i} conform to π . In strategies, this captures the aspect of players actions being responses to the opponent’s moves. The opponent’s complete strategy may not be available, the player makes a choice taking into account the apparent behaviour of the opponent on the history of play.

Let $\Sigma_i = \{a_1, \dots, a_m\}$, we use the abbreviation $null^i \equiv [\top \mapsto a_1] + \dots + [\top \mapsto a_m]$. The intuitive meaning is, any strategy of player i conforms to $null^i$.

Semantics: Given a state u and a valuation $V : u \rightarrow 2^P$, the truth of a formula $\psi \in \Psi(P^i)$ is defined as follows:

- $u \models p$ iff $p \in V(u)$.
- $u \models \neg\psi$ iff $u \not\models \psi$.
- $u \models \psi_1 \vee \psi_2$ iff $u \models \psi_1$ or $u \models \psi_2$.

We consider game trees along with a valuation function $V : S \rightarrow 2^P$. Given a strategy μ of player i and a node $s \in \mu$, let $\rho_s : s_0 a_0 s_1 \dots s_m = s$ be the unique path in μ from the root node to s . For all $j : 0 \leq j < m$, let $out_{\rho_s}(s_j) = a_j$ and $out_{\rho_s}(s)$ be the unique outgoing edge in μ at s . For a strategy specification $\sigma \in Strat^i(P^i)$, we define when μ conforms to σ (denoted $\mu \models_i \sigma$) as follows:

- $\mu \models_i \sigma$ iff for all player i nodes $s \in \mu$, we have $\rho_s, s \models_i \sigma$

where we define $\rho_s, s_j \models_i \sigma$ for any s_j in ρ_s as,

- $\rho_s, s_j \models_i [\psi \mapsto a]^i$ iff $s_j \models \psi$ implies $out_{\rho_s}(s_j) = a$.
- $\rho_s, s_j \models_i \sigma_1 + \sigma_2$ iff $\rho_s, s_j \models_i \sigma_1$ or $\rho_s, s_j \models_i \sigma_2$.
- $\rho_s, s_j \models_i \sigma_1 \cdot \sigma_2$ iff $\rho_s, s_j \models_i \sigma_1$ and $\rho_s, s_j \models_i \sigma_2$.

- $\rho_s, s_j \models_i \pi \Rightarrow \sigma$ iff for all player \bar{i} nodes $s_k \in \rho_s$ such that $k \leq j$, if $\rho_s, s_k \models_{\bar{i}} \pi$ then $\rho_s, s_j \models_i \sigma$.

Above, $\pi \in Strat^{\bar{i}}(P^1 \cap P^2)$ and $\psi \in \Psi(P^i)$.

Reasoning about strategies

We present a logic to reason about strategies with respect to a single extensive form game tree g . Strategy specifications are employed in the formulas of the logic to partially specify strategies rather than giving a complete description.

Syntax: Let $g \in G$ be an extensive form game tree. The syntax of the logic is given by:

$$\Phi := p \in P \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid \langle\langle g, \sigma \rangle\rangle\gamma$$

where $i \in \{1, 2\}$, $\sigma \in Strat^i(P^i)$ and $\gamma \in \Psi(P)$.

The intuitive meaning of $\langle\langle g, \sigma \rangle\rangle\gamma$ is: in the game g , the player has a strategy conforming to the specification σ which ensures γ . Since we are considering a fixed game g , this implies that γ holds at all the leaf node of the appropriate strategy. The restriction of γ to boolean formulas over the set of propositions is due to this reason. Nesting of the modality $\langle\langle g, \sigma \rangle\rangle$ does not make sense for a fixed game. At a later stage we will look at composing games at which point γ can be taken to be any arbitrary formula.

Semantics: The model $M = (T_g, V)$ where $T_g = (S, \Longrightarrow, s_0, \lambda)$ is the extensive form game tree associated with g and V is the valuation function $V : S \rightarrow 2^P$.

The truth of a formula $\alpha \in \Phi$ in a model M and 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 $M, s \models \alpha_2$.
- $M, s \models \langle\langle g, \sigma \rangle\rangle\gamma$ iff $\exists \mu \in \Omega^i$ such that $\mu \models_i \sigma$ and for all $s' \in frontier(\mu)$, $M, s' \models \gamma$.

The formula $\langle\langle g, \sigma \rangle\rangle\gamma$ says that there exists a strategy for player i conforming to μ such that all the leaf nodes satisfy γ . The dual $[\langle\langle g, \sigma \rangle\rangle\gamma]$ says that for all strategies of player i conforming to σ , there exists a leaf node which satisfy γ .

Strategy comparison

Consider the formula $\langle\langle g, null^i \rangle\rangle\gamma$. The formula asserts that player i can ensure the reward γ no matter what player \bar{i} does. This makes no reference to *how* player i may achieve this objective, and thus is similar to assertions in most game logics. Now consider the formula $\langle\langle g, \sigma \rangle\rangle\gamma$. This says something stronger: that there exists a strategy μ satisfying σ for player i such that irrespective of what player \bar{i} plays, γ is guaranteed. Here, the mechanism μ used by player i to ensure γ is specified by the property σ .

The extensive form game tree g merely defines the rules of how the game progresses and terminates. However, to compare strategies of players, we need to specify the objectives. For $i \in \{1, 2\}$, let R_i be a finite set of rewards for player i , $\preceq^i \subseteq R_i \times R_i$, be a preference ordering on R_i and let $R = R_1 \times R_2$. Let the payoff function $payoff : S^{leaf} \rightarrow R$ associate each leaf node with a reward. For a leaf node s , and

$payoff(s) = (r_1, r_2)$, let $payoff(s)[i]$ denote the i 'th component of r , i.e. $payoff(s)[1] = r_1$ and $payoff(s)[2] = r_2$.

In order to refer to rewards of the players in formulas of the logic, we use special propositions to code them up. This is similar to the approach adopted in (Bonanno 2002). Without loss of generality assume that $r_1^1 \preceq^1 r_1^2 \preceq^1 \dots \preceq^1 r_1^l$. Let $\Theta_1 = \{\theta_1^1, \dots, \theta_1^l\}$ 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 $i \in \{1, 2\}$, $\{\theta_i^1, \dots, \theta_i^j\} \subseteq V(s)$ iff $payoff(s)[i] = r_i^j$.

The preference ordering on the rewards for each player is simply inherited from the implication available in the logic.

Coming to the notion of strategy comparison, we say that σ is better for player i than σ' if the following condition holds: irrespective of what player \bar{i} plays if there exists a strategy μ' satisfying σ' such that θ_i is guaranteed, then there also exists a strategy μ satisfying σ which guarantees θ_i . This can be expressed by the formula,

$$BT^i(\sigma, \sigma') \equiv \bigwedge_{\theta_i \in \Theta_i} (\langle (g, \sigma') \rangle \theta_i \supset \langle (g, \sigma) \rangle \theta_i)$$

Given a finite set of strategy specifications Υ^i for player i , we say that σ is the best strategy if the following holds:

$$Best^i(\sigma) \equiv \bigwedge_{\sigma' \in \Upsilon^i} BT^i(\sigma, \sigma')$$

Note that in the case of a finite extensive form game tree, we can code up the game positions uniquely using propositions. In this case, it is possible to represent a complete strategy in terms of a strategy specification. At each game position, it specifies a unique action. Suppose the number of player i game positions are k and the proposition p_i^1, \dots, p_i^k uniquely identifies all of these positions, then the specification representing a complete strategy would have the form $\sigma \equiv [p_i^1 \mapsto a_1] \dots [p_i^k \mapsto a_k]$. In this particular scenario, the notion of strategy comparison and best strategy reduces to the classical notions by taking the set Υ^i to be the set of all strategies for player i .

Composition of game - strategy pairs

In the previous section we looked at strategies being defined by their properties. Strategy specifications are structurally built and the reasoning performed was with respect to one fixed extensive form game tree. Instead of working with a single game, we can look at complex games arising out of composition of these atomic games. In this context, we argue that reasoning about game - strategy pairs and their composition is more useful than composing games and analysing strategies separately. Here we present a logic to reason about game - strategy pairs. Both strategy specification and game structure is embedded into the syntax of the logic.

The logic

The logic is a simple dynamic logic where we take regular expressions over game-strategy pairs as programs in the logic. The formulas of the logic can then be used to specify the result of a player following a particular strategy in a specified game enabled at a state.

Syntax: For $i \in \{1, 2\}$, let P^i be a countable set of propositions and $P = P^1 \cup P^2$. The syntax for the logic is given by:

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

where $\xi \in \Gamma$, the set Γ consists of game strategy pairs which is defined below. As a convention we use $\top \equiv p \vee \neg p$. We will also make use of the following abbreviation:

- Let $g^i = ((i, x), a, (j, y))$ and $g^{\bar{i}} = ((\bar{i}, x), a, (j, y))$,
 - $\langle a \rangle \alpha \equiv turn_i \supset \langle g^i, [\top \mapsto a]^i \rangle \alpha \wedge turn_{\bar{i}} \supset \langle g^{\bar{i}}, [\top \mapsto a]^{\bar{i}} \rangle \alpha$

From the semantics it will be clear that this gives the usual interpretation 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 .

In the syntax of the logic, ξ represents regular expressions over game-strategy pairs (g, σ) . The intuitive meaning of $\langle \xi \rangle \alpha$ being that in the game g the player has a strategy conforming to the specification σ which ensures α .

Game strategy pairs: Syntax for game strategy specification pair is given by:

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

where $g \in G$, $\sigma \in Strat^i(P^i)$.

The atomic construct (g, σ) as mentioned in the earlier section, specifies that in game g a strategy conforming to specification σ is employed. Game strategy pairs are then composed using standard dynamic logic connectives. $\xi_1 + \xi_2$ 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 strategy pairs of the form $(g_1, \sigma_1); (g_2, \sigma_2)$. This is where the extensive form game tree interpretation makes the main difference. Since the strategy specifications are intended to be partial, a pair (g, σ) gives rise to a set of possibilities and therefore composition over these trees need to be performed. ξ^* is the iteration of the ';' operator.

Model: The formulas of the logic express properties about game trees and strategies which are composed using tree regular expressions. These formulas are to be interpreted on game positions and they make assertions about the frontier of the game trees which results from the pruning performed as dictated by the strategy specification. Therefore the models of the logic are game trees, but this can potentially be an infinite set of finite game trees. Alternatively, we can think of these game trees as being obtained from unfoldings of a Kripke structure. As we will see later, the logic cannot distinguish between these two.

A model $M = (W, \longrightarrow, \lambda, V)$ where W is the set of states (or game positions), the relation $\longrightarrow \subseteq W \times \Sigma \times W$, player labelling $\lambda : W \rightarrow \{1, 2\}$ and $V : W \rightarrow 2^P$.

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

- $M, w \models p$ iff $p \in V(w)$.
- $M, w \models \neg\alpha$ iff $M, w \not\models \alpha$.
- $M, w \models \alpha_1 \vee \alpha_2$ iff $M, w \models \alpha_1$ or $M, w \models \alpha_2$.
- $M, w \models \langle \xi \rangle \alpha$ iff $\exists(w, X) \in R_\xi$ such that $\forall w' \in X$ we have $M, w' \models \alpha$.

In the semantics of $\langle \xi \rangle \alpha$, the state w can be thought of as the starting game position and X , the set of leaf nodes of the game. We require that the player has a strategy conforming to the specification to ensure that α holds in all of the leaf nodes.

For $\xi \in \Gamma$, we have $R_\xi \subseteq W \times 2^W$. To define the relation formally, let us first assume that R is defined for the atomic case, namely when $\xi = (g, \sigma)$. The semantics for composite game strategy pairs is given as follows:

- $R_{\xi_1; \xi_2} = \{(u, X) \mid \exists Y = \{v_1, \dots, v_k\}$ 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 $\bigcup_{j=1, \dots, k} X_j = X\}$.
- $R_{\xi_1 \cup \xi_2} = R_{\xi_1} \cup R_{\xi_2}$.
- $R_{\xi^*} = \bigcup_{n \geq 0} (R_\xi)^n$.

In the atomic case when $\xi = (g, \sigma)$ we want a pair (u, X) to be in R_ξ if the game g is enabled at state u and there is a strategy conforming to the specification σ such that X is the set of leaf nodes of the strategy. In order to make this precise, we will require the following notations and definitions.

Restriction on trees: For $w \in W$, let T_w denote the tree unfolding of M starting at w . Given a state w and $g \in G$, let $T_w = (S_M^w, \Longrightarrow_M, \lambda_M, s_w)$ and $T_g = (S_g, \Longrightarrow_g, \lambda_g, s_{g,0})$. 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 inductively as follows: $T_w \upharpoonright g = (S, \Longrightarrow, \lambda, s_0, f)$ where $f : S \rightarrow S_g$. Initially $S = \{s_w\}$, $\lambda(s_w) = \lambda_M(s_w)$, $s_0 = s_w$ and $f(s_w) = s_{g,0}$.

For any $s \in S$, let $f(s) = t \in S_g$. Let $\{a_1, \dots, a_k\}$ be the outgoing edges of t , i.e. for all $j : 1 \leq j \leq k$, $t \xrightarrow{a_j} t_j$. For each a_j , let $\{s_j^1, \dots, s_j^m\}$ be the nodes in S_M^w such that $s \xrightarrow{a_j}_M s_j^l$ for all $l : 1 \leq l \leq m$. Add nodes s_j^1, \dots, s_j^m to S and the edges $s \xrightarrow{a_j} s_j^l$ for all $l : 1 \leq l \leq m$. Also set $\lambda(s_j^l) = \lambda_M(s_j^l)$ and $f(s_j^l) = t_j$.

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

- $\forall s \in S, \lambda(s) = \lambda_g(f(s))$.

For a game tree T , let $\Omega^i(T)$ denote the set of strategies of player i on the game tree T and $frontier(T)$ denote the set of all leaf nodes of T .

Atomic game-strategy pair: For atomic game-strategy pair $\xi = (g, \sigma)$ we define R_ξ as follows:

Let g be the game with a single node $g = (i, x)$,

- $R_{(g, \sigma)} = \{(u, \{u\})\}$ if $enabled(g, u)$ holds, for all $i \in \{1, 2\}$, for all $\sigma \in Strat^i(P^i)$.

For $g = ((i, x), a_1, t_{a_1} + \dots + (i, x), a_k, t_{a_k})$

- $R_{(g, \sigma)} = \{(u, X) \mid enabled(g, u) \text{ and } \exists \mu \in \Omega^i(T_u \upharpoonright g) \text{ such that } \mu \models_i \sigma \text{ and } frontier(\mu) = X\}$.




Figure 2: Model

Example 1 Let the extensive form game g be the one given in Figure 1(a) and the Kripke structure M be as shown in Figure 2(a). For the node u of the structure the restriction $T_u \upharpoonright g$ is shown in Figure 2(b). This is the maximal subtree of T_u according to the structure dictated by g . For instance at node v_1 there are two x_1 labelled edges present in M and therefore both have to be included in $T_u \upharpoonright g$ as well.

Now consider the player 1 strategy specification $\sigma \equiv null^1$. At node u , the choice 'a' can ensure player 1 the states $\{w_1, v_2, w_3\}$ and the choice 'b' can ensure the states $\{w_4, w_5\}$. Therefore the relation $R_{(g, \sigma)} = \{(u, \{w_1, v_2, w_3\}), (u, \{w_4, w_5\}), (v_1, \{w_1, v_2, w_3\}), (v_2, \{w_4, w_5\})\}$.

Suppose $M, u \models p$ and consider the specification $\sigma \equiv [p \mapsto a]^1$. Since p holds at the root, player 1 is restricted to make the choice 'a' at u . Hence the relation in this case would be $R_{(g, \sigma)} = \{(u, \{w_1, v_2, w_3\}), (v_1, \{w_1, v_2, w_3\}), (v_2, \{w_4, w_5\})\}$.




Figure 3:

Example 2 To illustrate the logic, consider the games g_1 and g_2 given in Fig. 3. Let u be a state of the model where g_1 is enabled. Let g denote the game $g_1; g_2$, i.e. the game obtained by pasting g_2 at each of the leaf nodes of g_1 . We use the following notation:

- w^{a_1} : denotes the state reached after action a_1 .
- w^{a_1, b_1} : the state reached on following actions a_1 and a_2 .
- $w_{x_1, y_1}^{a_1, b_1}$: the state reached on the sequence of action $a_1 b_1 x_1 y_1$.

Let win_1 , win_2 and p be propositions whose valuations are given by $V(win_2) = \{w^{a_1, b_1}, w^{a_2, b_2}\}$, $V(win_1) = \{w^{a_1, b_1}, w^{a_2, b_2}\}$ and $V(p) = \{w^{a_1}\}$. Consider the following specifications:

- $\pi \equiv [p \mapsto b_1]^2 \cdot [\neg p \mapsto b_2]^2$.
- $\sigma \equiv [\top \mapsto x_1]^1$.

It is easy to see that $\langle\langle g_1, \pi \rangle\rangle win_2$ holds at u . Player 1 does not have a strategy in the composite game g to ensure win_1 . However, in the composite pair $\xi = (g_1, \pi); (g_2, \sigma)$, it is easy to see that $\langle\xi\rangle win_1$ holds. Assuming that in the game g_1 player 2 plays according to π then in g_2 by using a strategy which conforms to σ player 1 can ensure win_1 . In some sense this says that reasoning in the game g is different from reasoning in g_1 composed with g_2 . In the latter, the additional structural information is available which can be used for strategizing.

For simple game structures it is quite obvious that such reasoning can be done with a past modality. It is iteration which provides the actual expressive power. In the presence of iteration, the analysis asserts the fact that players can take into account the structure of the game and the opponent's strategy. In particular while strategizing, a player can make use of the fact that the opponent is using a bounded memory strategy and that with the type of strategy that is being used the opponent can be forced into a particular region of the game graph.

The above mentioned reasoning can also be thought of as players trying to attain certain local goals. If player 2 plays to achieve the local goal win_2 then player 1 can use this information and respond with a strategy in g_2 to achieve the objective win_1 . Players can then try to achieve their global objective by performing appropriate composition of the local objectives.

Even at the atomic level, the game structure can be quite complicated. At this level, strategy specifications enable reasoning about strategies satisfying certain invariant properties. Here strategizing in response to the opponent's action is captured by the construct $\pi \Rightarrow \sigma$.

Axiom system

We now present an axiomatization of the valid formulas of the logic. We will use the following notations:

For a set $A = \{a_1, \dots, a_k\} \subseteq \Sigma$, we will use the notation $\mathfrak{R}(i, x, A)$ to denote the game $((i, x), a_1, t_{a_1} + \dots + (i, x), a_k, t_{a_k})$.

For game g , we use the formula g^\vee to denote that the game structure g is enabled. This is defined as:

- For $g = (i, x)$, let $g^\vee \equiv turn_i$.
- For $g = \mathfrak{R}(i, x, A)$, let
 - $g^\vee \equiv turn_i \wedge (\bigwedge_{j=1, \dots, k} (\langle a_j \rangle \top \wedge [a_j] t_{a_j}^\vee))$.

The axiom schemes

(A1) Propositional axioms:

- (a) All the substitutional instances of tautologies of PC.
- (b) $turn_i \equiv \neg turn_{\bar{i}}$.

(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) $\langle \xi_1 \cup \xi_2 \rangle \alpha \equiv \langle \xi_1 \rangle \alpha \vee \langle \xi_2 \rangle \alpha$.
- (b) $\langle \xi_1; \xi_2 \rangle \alpha \equiv \langle \xi_1 \rangle \langle \xi_2 \rangle \alpha$.
- (c) $\langle \xi^* \rangle \alpha \equiv \alpha \vee \langle \xi \rangle \langle \xi^* \rangle \alpha$.

(A4) $\langle g, \sigma \rangle \alpha \equiv g^\vee \wedge push(g, \sigma, \alpha)$.

Inference rules

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

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

Axiom (A2a) does not hold for general game strategy pairs (i.e. $\xi \in \Gamma$). In particular $\langle \xi \rangle (\alpha_1 \vee \alpha_2) \supset \langle \xi \rangle \alpha_1 \vee \langle \xi \rangle \alpha_2$ is not valid. However (A2a) is sound since $\langle a \rangle$ asserts properties about a single edge.

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. $R_{\xi_1; \xi_2} = \{(u, X) | \exists Y \text{ such that } (u, Y) \in R_{\xi_1} \text{ and for all } v \in Y, (v, X) \in R_{\xi_2}\}$. It is easy to see that under this interpretation the formula $\langle \xi_1 \rangle \langle \xi_2 \rangle \alpha \supset \langle \xi_1; \xi_2 \rangle \alpha$ is not valid. A soundness argument for axiom (A3b) is given in the appendix.

Definition of push: For all $i \in \{1, 2\}$, $g \in G$, $\sigma \in Strat^i(P^i)$ and $\alpha \in \Phi$, we define $push(g, \sigma, \alpha)$ as follows. We have various cases depending on the structure of g .

The case when g is an atomic game, i.e. $g = (i, x)$, for all $i \in \{1, 2\}$ and $\sigma \in Strat^i(P^i)$ we have,

(C1) $push(g, \sigma, \alpha) \equiv \alpha$.

Suppose $g = \mathfrak{R}(i, x, A)$ for $A = \{a_1, \dots, a_k\}$. For each $a_m \in A$ let $g_{a_m} = ((i, x), a_m, (j_m, y_m))$, where (j_m, y_m) is the root of t_{a_m} .

For $\pi \equiv [\psi \mapsto a]^\bar{i}$, $\pi_1 + \pi_2, \pi_1 \cdot \pi_2 \in Strat^{\bar{i}}(P^{\bar{i}})$.

(C2) $push(g, \pi, \alpha) \equiv \bigwedge_{a_m \in A} [a_m] push(t_{a_m}, \pi, \alpha)$.

(C3) $push(g, \sigma \Rightarrow \pi, \alpha) \equiv$

$$\bigwedge_{a_m \in A} (\langle g_{a_m}, \sigma \rangle \top \supset [a_m] push(t_{a_m}, \sigma \Rightarrow \pi, \alpha))$$

$$\wedge \neg \langle g_{a_m}, \sigma \rangle \top \supset [a_m] push(t_{a_m}, null^{\bar{i}}, \alpha).$$

(C4) $push(g, [\psi \mapsto a]^i, \alpha) \equiv$
 $(\psi \supset \langle a \rangle push(t_a, [\psi \mapsto a]^i, \alpha))$
 $\wedge (\neg \psi \supset (\bigvee_{a_m \in A} \langle a_m \rangle push(t_{a_m}, [\psi \mapsto a]^i, \alpha))).$

(C5) $push(g, \sigma_1 \cdot \sigma_2, \alpha) \equiv$

$$\bigvee_{a_m \in A} (\langle g_{a_m}, \sigma_1 \rangle push(t_{a_m}, \sigma_1 \cdot \sigma_2, \alpha))$$

$$\wedge \langle g_{a_m}, \sigma_2 \rangle push(t_{a_m}, \sigma_1 \cdot \sigma_2, \alpha).$$

(C6) $push(g, \sigma_1 + \sigma_2, \alpha) \equiv$

$$\bigvee_{a_m \in A} (\langle g_{a_m}, \sigma_1 \rangle push(t_{a_m}, \sigma_1 + \sigma_2, \alpha))$$

$$\vee \langle g_{a_m}, \sigma_2 \rangle push(t_{a_m}, \sigma_1 + \sigma_2, \alpha).$$

$$(C7) \text{ push}(g, \pi \Rightarrow \sigma, \alpha) \equiv \bigvee_{a_m \in A} (\langle g_{a_m}, \sigma \rangle \text{push}(t_{a_m}, \pi \Rightarrow \sigma, \alpha)).$$

The soundness of axiom (A4) is shown in the appendix.

Completeness

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 $V(w) = \{p \in P \mid p \in w\}$ and $\lambda(w) = i$ iff $\text{turn}_i \in w$. The model $M = (W, \longrightarrow, \lambda, V)$ where $W = \mathcal{AT}(\alpha_0)$. Once the Kripke structure is defined, the game theoretic semantics given earlier defines the relation $R_{(g,\sigma)}$ on $W \times 2^W$ for $g \in T$ and a strategy specification σ .

However to show the completeness result, we need to also specify the relation between a pair (u, X) being in $R_{(g,\sigma)}$ and the consistency requirement on u and X . In other words, we need to define a new relation $R'_{(g,\sigma)}$ in terms of consistency of u and X and show that the following property holds:

$$(P1) (u, X) \in R'_{(g,\sigma)} \text{ iff } (u, X) \in R_{(g,\sigma)}.$$

The first attempt would be to say $(u, X) \in R'_{(g,\sigma)}$ iff $\widehat{u} \wedge \langle g, \sigma \rangle \widetilde{X}$ is consistent. But this definition need not satisfy (\Rightarrow) of (P1). The trouble is, in the game theoretic definition of $R_{(g,\sigma)}$, we require X to be the exact set of leaves of g for which player has a strategy conforming to σ . If the definition of R had instead been ‘‘upward closed’’, i.e. $(u, X) \in R_{(g,\sigma)}$ implies for any $Y \supseteq X$, $(u, Y) \in R_{(g,\sigma)}$, then this approach would work.

The second attempt would be to say $(u, X) \in R'_{(g,\sigma)}$ iff for all $w \in X$, we have $\widehat{u} \wedge \langle g, \sigma \rangle \widehat{w}$ is consistent. It is quite easy to see that this definition is also unsatisfactory. The closure of the formula is quite rich in the sense that the tree structure as dictated by the axioms are present in the closure. Therefore for individual atoms u and w , unless g is a single edge game, $\widehat{u} \wedge \langle g, \sigma \rangle \widehat{w}$ need not be consistent at all.

What we really need is the minimal set X such that $\widehat{u} \wedge \langle g, \sigma \rangle \widetilde{X}$ is consistent. For this set X , we have that the pair $(u, X) \in R_{(g,\sigma)}$. Lemma 1 given below formalises this fact.

Lemma 1 *For all $g \in G$, for all $i \in \{1, 2\}$ and $\sigma \in \text{Strat}^i(P^i)$, for all $X \subseteq W$ and for all $u \in W$ the following holds:*

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

A detailed proof can be found in the appendix. Item 1 follows from the axioms and the fact that $CL(\alpha_0)$ is rich

enough that it has the tree structure built into it as dictated by the axioms. For item 2, we basically need to show the following two things:

- The game g is enabled at u .
- The existence of a strategy μ on g which conforms to the specification σ such that the leaf nodes of μ is $X' \subseteq X$.

The strategy construction is similar to the technique used to build the witness tree in CTL for the $\forall \exists$ quantifier. The idea is to start at u and extend in stages, making sure that for a player i node the choice conforms to σ and for a player \bar{i} node all the branches are taken into account. Since the analysis is done over tree structures, it is evident at this point that the techniques used are different from the ones in dynamic logic.

Lemma 2 *For all $\xi \in \Gamma$, for all $X \subseteq W$ and $u \in W$, if $\widehat{u} \wedge \langle \xi \rangle \widetilde{X}$ is consistent then there exists $X' \subseteq X$ such that $(u, X') \in R_\xi$.*

Proof is given in the appendix.

Lemma 3 *For all $\langle \xi \rangle \alpha \in CL(\alpha_0)$, for all $u \in W$, $\widehat{u} \wedge \langle \xi \rangle \alpha$ is consistent iff there exists $(u, X) \in R_\xi$ such that $\forall w \in X$, $\alpha \in w$.*

Proof: (\Rightarrow) Follows from lemma 2 by considering the set $X_\alpha = \{w \in W \mid \alpha \in w\}$.

(\Leftarrow) Suppose $\exists (u, X) \in R_\xi$ such that $\forall w \in X$, $\alpha \in w$. We need to show that $\widehat{u} \wedge \langle \xi \rangle \alpha$ is consistent, this is done by induction on the structure of ξ .

- The case when $\xi \equiv (g, \sigma)$ follows easily from lemma 1 and $\xi \equiv \xi_1 \cup \xi_2$ follows from the induction hypothesis and axiom (A3a).
- $\xi \equiv \xi_1; \xi_2$: Since $(u, X) \in R_{\xi_1; \xi_2}$, there exists $Y = \{v_1, \dots, v_k\}$, there exists sets $X_1, \dots, X_k \subseteq X$ such that $\bigcup_{j=1, \dots, k} X_j = X$, for all $j : 1 \leq j \leq k$, $(v_j, X_j) \in R_{\xi_2}$ and $(u, Y) \in R_{\xi_1}$. By induction hypothesis, for all j , $\widehat{v}_j \wedge \langle \xi_2 \rangle \alpha$ is consistent. Since v_j is an atom and $\langle \xi_2 \rangle \alpha \in CL(\alpha_0)$, we get $\langle \xi_2 \rangle \alpha \in v_j$. Again by induction hypothesis we have $\widehat{u} \wedge \langle \xi_1 \rangle \langle \xi_2 \rangle \alpha$ is consistent. Hence from (A3b) we have $\widehat{u} \wedge \langle \xi_1; \xi_2 \rangle \alpha$ is consistent.

- $\xi \equiv \xi_1^*$: If $u \in X$ then $\vdash \widehat{u} \supset \widetilde{X}$. We have $\vdash \widetilde{X} \supset \alpha$ and hence we get $\widehat{u} \wedge \alpha$ is consistent. From axiom (A3c) we have $\widehat{u} \wedge \langle \xi_1^* \rangle \alpha$ is consistent.

Else we have $(u, X) \in R_{\xi_1; \xi_1^*}$. Let $Z_0 = X$ and $Z_{n+1} = Z_n \cup \{w \mid (w, Z') \in R_{\xi_1}, Z' \in Z_n\}$. Take the least m such that $u \in Z_m$. We have for all $w \in Z_{m-1}$, $\vdash \widehat{w} \supset \langle \xi_1^* \rangle \widetilde{X}'$ for some $X' \subseteq X$. We also have $(u, Z'_m) \in R_{\xi_1}$ for some $Z'_m = \{v_1, \dots, v_k\} \subseteq Z_m$. Let $X_1, \dots, X_k \subseteq X$ such that $\forall j : 1 \leq j \leq k$, we have $(v_j, X_j) \in R_{\xi_1^*}$ and $X' = \bigcup_{j=1, \dots, k} X_j$. By an argument similar to the previous case we can show that $\widehat{u} \wedge \langle \xi_1 \rangle \langle \xi_1^* \rangle \widetilde{X}'$ is consistent. Hence we get $\widehat{u} \wedge \langle \xi_1; \xi_1^* \rangle \alpha$ is consistent. Therefore from axiom (A3c) we have $\widehat{u} \wedge \langle \xi_1^* \rangle \alpha$ is consistent. \square

Theorem 4 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 3 by a routine inductive argument.

Decidability: Since the size of the action set $|\Sigma|$ 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 . From the completeness theorem we get that for a formula α_0 , if α_0 is satisfiable then it has a model of exponential size, i.e. $|W| = \mathcal{O}(2^{\alpha_0})$. For all game strategy pairs ξ occurring in α_0 , the relation R_ξ can be computed in time exponential in the size of the model. Therefore it follows that the logic is decidable in nondeterministic double exponential time.

Extensions

Concurrency operator: Concurrency as introduced in game logic (van Benthem, Ghosh, and Liu 2007) can be represented in our framework with the addition of the operator $\xi_1 \times \xi_2$ in the syntax of game strategy pairs. For instance, $(g_1, \sigma_1) \times (g_2, \sigma_2)$ would mean that the game g_1 is played with a strategy conforming to σ_1 and concurrently, the game g_2 is played with a strategy conforming to σ_2 . The semantics can be defined in the usual manner:

- $R_{\xi_1 \times \xi_2} = \{(u, X) \mid X = X_1 \cup X_2 \text{ such that } (u, X_1) \in R_{\xi_1} \text{ and } (u, X_2) \in R_{\xi_2}\}$.

It is easy to see that the completeness theorem also follows with the addition of the following axiom.

- $\langle \xi_1 \times \xi_2 \rangle \alpha \equiv \langle \xi_1 \rangle \alpha \wedge \langle \xi_2 \rangle \alpha$.

Test operator: The test operator as in dynamic logic can also be added into the syntax of game strategy pairs. For $\beta \in \Phi$, the interpretation of $\beta?$ $\in \Gamma$ would be to test whether β holds at the particular state and if yes, continue else fail. The semantics can be given as:

- $R_{\beta?} = \{(u, \{u\}) \mid M, u \models \beta\}$.

The test operator gives the ability of checking for certain conditions and then deciding which game to proceed with. This construct is particularly interesting in our framework, since unlike programs we have players in the game. For instance, let π denote the strategy specification of player 2 and σ the specification of player 1. The formula $(g_1, \pi); win_2?; (g_2, \sigma)$ says that in g_1 if player 2 by employing a strategy conforming to π can ensure win_2 then proceed with the game g_2 where player 1 plays σ . Note that if the test fails then g_2 is not played. This is in contrast to the tests performed in a strategy specification. In a specification if the test fails then the player is free to choose any action.

With the addition of the following axiom, the completeness theorem also goes through.

- $\langle \beta? \rangle \alpha \equiv \beta \supset \alpha$

Discussion

The logical interaction between strategy specifications and game structure is explicated by the axioms, but this is as

yet unsatisfactory. Ideally, this is best accomplished by an equational theory \vdash_E so that one rule suffices (in the presence of induction): from $\vdash_E (g_1, \sigma_1) = (g_2, \sigma_2)$ infer $\vdash \langle (g_1, \sigma_1) \rangle \alpha \equiv \langle (g_2, \sigma_2) \rangle \alpha$. We need further work on strategy structure as we have on Kleene algebras.

As we remarked at the beginning, we see this study as initial: one among the natural but missing game theoretic notions is that of players' ability to switch strategies: whereby a player plays strategy μ_1 till a particular objective is achieved and then switches to strategy μ_2 . While we can describe such changes at least at the atomic level, the *rationale* for switching is missing.

Shoham (2003) advocates incorporating elements of rationality and utility into programming languages. This makes eminent sense; we merely add a footnote that strategies provide the environments in which such programs (with goals and preferences) are to be interpreted.

Acknowledgements

We thank Sujata Ghosh for the various discussions and valuable comments. We also thank the anonymous referees for their valuable comments and suggestions.

Appendix

Soundness

Axiom (A3b): Suppose $\langle \xi_1; \xi_2 \rangle \alpha \supset \langle \xi_1 \rangle \langle \xi_2 \rangle \alpha$ is not valid. Then there exists M and u such that $M, u \models \langle \xi_1; \xi_2 \rangle \alpha$ and $M, u \not\models \langle \xi_1 \rangle \langle \xi_2 \rangle \alpha$. Since $M, u \models \langle \xi_1; \xi_2 \rangle \alpha$, from semantics we have there exists $(u, X) \in R_{\xi_1; \xi_2}$ such that $\forall w \in X$, $M, w \models \alpha$. From definition of R , $\exists Y = \{v_1, \dots, v_k\}$ 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 $\bigcup_{j=1, \dots, k} X_j = X$. Therefore we get $\forall v_k \in Y$, $M, v_k \models \langle \xi_2 \rangle \alpha$ and hence from semantics, $M, u \models \langle \xi_1 \rangle \langle \xi_2 \rangle \alpha$. This gives the required contradiction.

Suppose $\langle \xi_1 \rangle \langle \xi_2 \rangle \alpha \supset \langle \xi_1; \xi_2 \rangle \alpha$ is not valid. Then there exists M and u such that $M, u \models \langle \xi_1 \rangle \langle \xi_2 \rangle \alpha$ and $M, u \not\models \langle \xi_1; \xi_2 \rangle \alpha$. We have $M, u \models \langle \xi_1 \rangle \langle \xi_2 \rangle \alpha$ iff there exists $(u, Y) \in R_{\xi_1}$ such that $\forall v_k \in Y$, $M, v_k \models \langle \xi_2 \rangle \alpha$. $M, v_k \models \langle \xi_2 \rangle \alpha$ iff there exists $(v_k, X_k) \in R_{\xi_2}$ such that $\forall w_k \in X_k$, $M, w_k \models \alpha$. Let $X = \bigcup_k X_k$, from definition of R we get $(u, X) \in R_{\xi_1; \xi_2}$. Hence from semantics $M, u \models \langle \xi_1; \xi_2 \rangle \alpha$.

Axiom (A4): To show the soundness of axiom (A4), we need to consider the cases (C2) to (C7). Soundness of one direction (\supset) is easy to see. Let us consider the other direction (\subset). The root of g is an i node therefore any \bar{i} strategy should consider all moves enabled at the root. (A4) case (C2) says, for an \bar{i} specification which is not of the form $\sigma \Rightarrow \pi$, if at all enabled edges a_m , the subtree t_{a_m} satisfies $\langle t_{a_m}, \pi \rangle \alpha$ then $\langle g, \pi \rangle \alpha$ holds. Case (C4) has a player i specification. This says that if at the root node there is some choice a_j that player i can make conforming to the specification such that for the subtree $\langle t_{a_j}, [\psi \mapsto a]^i \rangle \alpha$ holds then the number of branches at the root is irrelevant and therefore $\langle g, [\psi \mapsto a]^i \rangle \alpha$ holds as well. For (C5) the important point to note is the fact if an edge $u \xrightarrow{a} w$ satisfies a specification

σ then all w' with $u \xrightarrow{a} w'$ satisfies σ . This is because satisfaction of σ depends only on u and the action a , it does not depend on the target node. Case (C6) and (C7) also follows quite easily.

The interesting case is when the root of g is an i node and when the specification is of the form $\sigma \Rightarrow \pi$, this is specified in (C3). For a strategy τ of player \bar{i} to satisfy $\sigma \Rightarrow \pi$ on g , it should make sure of the following:

- for each choice $a_m \in A$, if the choice conforms with σ then the strategy on t_{a_m} should satisfy π .
- for each choice $a_m \in A$, which does not conform with σ player \bar{i} is allowed to employ any strategy on the game t_{a_m} .

From the above observation, the soundness of (A4) case (C3) follows easily.

Detailed proofs

For a model M , a state $u \in W$ and a formula $\psi \in \Psi$, we use the notation $M, u \models \psi$ to mean $u \models \psi$. The following proposition is easy to show using a standard inductive argument.

Proposition 5 For all $i \in \{1, 2\}$, for all $\psi \in \Psi(P^i)$, for all $u \in W$ we have $M, u \models \psi$ iff $\psi \in u$.

Lemma 1. For all $g \in G$, for all $i \in \{1, 2\}$ and $\sigma \in \text{Strat}^i(P^i)$, for all $X \subseteq W$ and for all $u \in W$ the following holds:

1. if $(u, X) \in R_{(g, \sigma)}$ then $\hat{u} \wedge \langle g, \sigma \rangle \tilde{X}$ is consistent.
2. if $\hat{u} \wedge \langle g, \sigma \rangle \tilde{X}$ is consistent then there exists $X' \subseteq X$ such that $(u, X') \in R_{(g, \sigma)}$.

Proof: By induction on the structure of (g, σ) .

For atomic game $g = (i, x)$, from axiom (A4) case (C1) we get $\langle (i, x), \sigma \rangle \alpha \equiv \text{turn}_i \wedge \alpha$. The lemma follows from this quite easily. For the case when g is a single edge, i.e. $g = ((i, x), a, (j, y))$, it is easy to see that the lemma holds.

Let $g = \mathfrak{R}(i, x, A)$ for $A = \{a_1, \dots, a_k\}$.

$\sigma \equiv [\psi \mapsto a]^t$:

Suppose $(u, X) \in R_{(g, \sigma)}$, since $\text{enabled}(g, u)$ holds we have there exists sets Y_1, \dots, Y_k such that for all $j : 1 \leq j \leq k$, for all $w_j \in Y_j$ we have $u \xrightarrow{a_j} w_j$. Since u is an i node, any strategy of i will pick a unique edge at u . We have the following two cases:

- $M, u \models \psi$: From semantics, the strategy should choose a w_a such that $u \xrightarrow{a} w_a$ and $(w_a, X) \in R_{(t_a, \sigma)}$. By induction hypothesis, we have $\hat{w}_a \wedge \langle t_a, \sigma \rangle \tilde{X}$ is consistent. Hence $\hat{u} \wedge \langle a \rangle \langle t_a, \sigma \rangle \tilde{X}$ is consistent.
- $M, u \not\models \psi$: The strategy can choose any w_j such that $u \xrightarrow{a_j} w_j$ and $(w_j, X) \in R_{(t_j, \sigma)}$. By induction hypothesis, $\hat{w}_j \wedge \langle t_j, \sigma \rangle \tilde{X}$ is consistent. Hence $\hat{u} \wedge \langle a_j \rangle \langle t_j, \sigma \rangle \tilde{X}$ is consistent.

From axiom (A4) case (C4) we get $\hat{u} \wedge \langle g, \sigma \rangle \tilde{X}$ is consistent.

Suppose $\hat{u} \wedge \langle g, \sigma \rangle \tilde{X}$ is consistent. From axiom (A4) it follows that there exists sets Y_1, \dots, Y_k such that for all $j : 1 \leq j \leq k$, for all $w_j \in Y_j$ we have $u \xrightarrow{a_j} w_j$ and hence $\text{enabled}(g, u)$ holds. Let $X = \{v_1, \dots, v_m\}$. We have the following two cases:

- if $M, u \models \psi$: then from case (C4), $\hat{u} \wedge (\langle a \rangle \langle t_a, \sigma \rangle \tilde{X})$ is consistent. Hence we get there exists w_a such that $u \xrightarrow{a} w_a$ and $\hat{w}_a \wedge \langle t_a, \sigma \rangle \tilde{X}$ is consistent. By induction hypothesis there exists $X' \subseteq X$ such that $(w_a, X') \in R_{(t_a, \sigma)}$ and by definition of R we have $(u, X') \in R_{(g, \sigma)}$.
- if $M, u \not\models \psi$: then from case (C4), $\hat{u} \wedge \bigvee_{a_j \in A} \langle a_j \rangle \langle t_j, \sigma \rangle \tilde{X}$. Therefore there exists w_j such that $u \xrightarrow{a_j} w_j$ and $\hat{w}_j \wedge \langle t_j, \sigma \rangle \tilde{X}$ is consistent. By induction hypothesis there exists $X' \subseteq X$ such that $(w_j, X') \in R_{(t_j, \sigma)}$ and therefore we have $(u, X') \in R_{(g, \sigma)}$.

$\sigma \equiv [\psi \mapsto a]^{\bar{i}}, \pi_1 + \pi_2, \pi_1 \cdot \pi_2 \in \text{Strat}^{\bar{i}}(P^{\bar{i}})$:

Suppose $(u, X) \in R_{(g, \pi)}$, since $\text{enabled}(g, u)$ holds, we have there exists Y_1, \dots, Y_k such that for all $j : 1 \leq j \leq k$, for all $w_j \in Y_j$, we have $u \xrightarrow{a_j} w_j$. Since u is an i node, any strategy τ of \bar{i} conforming to π will have all the branches at u (by definition of strategy). Therefore we get for all w_j with $u \xrightarrow{a_j} w_j$, there exists $X_j \subseteq X$ such that $(w_j, X_j) \in R_{(t_j, \pi)}$ and $X = \bigcup_{j=1, \dots, k} X_j$. By induction hypothesis and the fact that $X_j \subseteq X$, we have $\hat{w}_j \wedge \langle t_j, \pi \rangle \tilde{X}$ is consistent. Hence from axiom (A4) case (C2), we conclude that $\hat{u} \wedge \langle g, \sigma \rangle \tilde{X}$ is consistent.

Suppose $\hat{u} \wedge \langle g, \pi \rangle \tilde{X}$ is consistent. From axiom (A4) we get that $\hat{u} \wedge g^\vee$ is consistent. This implies that there exists sets Y_1, \dots, Y_k such that for all $j : 1 \leq j \leq k$, for all $w_j \in Y_j$ we have $u \xrightarrow{a_j} w_j$ and hence $\text{enabled}(g, u)$ holds. From case (C2), we have $\hat{u} \wedge (\bigwedge_{a_j \in A} [a_j] \langle t_{a_j}, \pi \rangle \alpha)$ is consistent. Therefore for all j such that $u \xrightarrow{a_j} w_j$, we have $w_j \wedge \langle t_{a_j}, \pi \rangle \tilde{X}$ is consistent. By induction hypothesis there exists $X'_j \subseteq X$ such that $(w_j, X'_j) \in R_{(t_j, \pi)}$. Let $X' = \bigcup_{j=1, \dots, k} X'_j$, by definition of R we have $(u, X') \in R_{(g, \pi)}$.

The cases when $\sigma \equiv \sigma_1 \cdot \sigma_2, \sigma_1 + \sigma_2, \pi \Rightarrow \sigma_1$ follows easily from axiom (A4) cases (C5) and (C6). Since the root of g is an i node the case when $\sigma \equiv \pi \Rightarrow \sigma_1$, also follows from case (C7) and the induction hypothesis.

The interesting case is when the root of g is an i node and when the specification is $\sigma_1 \Rightarrow \pi$.

Let $g = \mathfrak{R}(i, x, A)$ where $A = \{a_1, \dots, a_k\}$ and $\sigma \equiv \sigma_1 \Rightarrow \pi$.

Suppose $(u, X) \in R_{g, \sigma}$ since $\text{enabled}(g, u)$ holds, its easy to show that $\hat{u} \wedge g^\vee$ is consistent. For a strategy τ of player \bar{i} to satisfy $\sigma_1 \Rightarrow \pi$ on g , it should make sure of the following:

- for each edge $a_j \in A$, if $u \xrightarrow{a_j} w_j$ conforms with σ_1 then the strategy on t_j should satisfy π .
- for each edge $a_j \in A$, if $u \xrightarrow{a_j} w_j$ does not conform with σ_1 then any strategy can be employed on the game t_j .

From the above observations and axiom (A4) case (C3), we get $\hat{u} \wedge \langle g, \sigma_1 \Rightarrow \pi \rangle \tilde{X}$ is consistent.

Part 2 of the lemma again follows from (C3) and a similar argument. \square

Lemma 2. For all $\xi \in \Gamma$, for all $X \subseteq W$ and $u \in W$, if $\hat{u} \wedge \langle \xi \rangle \tilde{X}$ is consistent then there exists $X' \subseteq X$ such that $(u, X') \in R_\xi$.

Proof: By induction on the structure of ξ .

- $\xi \equiv (g, \sigma)$: Suppose $\hat{u} \wedge \langle g, \sigma \rangle \tilde{X}$ is consistent. From lemma 1 item 2, it follows that there exists $X' \subseteq X$ such that $(u, X') \in R_\xi$.
- $\xi \equiv \xi_1 \cup \xi_2$: By axiom (A3a) we get $\hat{u} \wedge \langle \xi_1 \rangle \tilde{X}$ is consistent or $\hat{u} \wedge \langle \xi_2 \rangle \tilde{X}$ is consistent. By induction hypothesis there exists $X_1 \subseteq X$ such that $(u, X_1) \in R_{\xi_1}$ or there exists $X_2 \subseteq X$ such that $(u, X_2) \in R_{\xi_2}$. Hence we have $(u, X_1) \in R_{\xi_1 \cup \xi_2}$ or $(u, X_2) \in R_{\xi_1 \cup \xi_2}$.
- $\xi \equiv \xi_1; \xi_2$: By axiom (A3b), $\hat{u} \wedge \langle \xi_1 \rangle \langle \xi_2 \rangle \tilde{X}$ is consistent. Hence $\hat{u} \wedge \langle \xi_1 \rangle (\bigvee (\hat{w} \wedge \langle \xi_2 \rangle \tilde{X}))$ is consistent, where the join is taken over all $w \in Y = \{w \mid w \wedge \langle \xi_2 \rangle \tilde{X} \text{ is consistent}\}$. So $\hat{u} \wedge \langle \xi_1 \rangle \tilde{Y}$ is consistent. By induction hypothesis, there exists $Y' \subseteq Y$ such that $(u, Y') \in R_{\xi_1}$. We also have that for all $w \in Y$, $\hat{w} \wedge \langle \xi_2 \rangle \tilde{X}$ is consistent. Therefore we get for all $w_j \in Y' = \{w_1, \dots, w_k\}$, $\hat{w}_j \wedge \langle \xi_2 \rangle \tilde{X}$ is consistent. By induction hypothesis, there exists $X_j \subseteq X$ such that $(w_j, X_j) \in R_{\xi_2}$. Let $X' = \bigcup_{j=1, \dots, k} X_k \subseteq X$, we get $(u, X') \in R_{\xi_1; \xi_2}$.
- $\xi \equiv \xi_1^*$: Let Z be the least set containing X and closed under the condition: for all w , if $\hat{w} \wedge \langle \xi_1 \rangle \tilde{Z}$ is consistent, then $w \in Z$. By definition of Z and induction hypothesis, we get for all $w \in Z$, there exists $X_w \subseteq X$ such that $(w, X_w) \in R_{\xi_1^*}$. It is also easy to see that $\vdash \tilde{X} \supset \tilde{Z}$. Using standard techniques, it is also easy to show that $\vdash \langle \xi_1 \rangle \tilde{Z} \supset \tilde{Z}$.

Applying the induction rule (IND), we have $\vdash \langle \xi_1^* \rangle \tilde{Z} \supset \tilde{Z}$.

By assumption, $\hat{u} \wedge \langle \xi_1^* \rangle \tilde{X}$ is consistent. So $\hat{u} \wedge \langle \xi_1^* \rangle \tilde{Z}$ is consistent. Hence $\hat{u} \wedge \tilde{Z}$ is consistent and therefore $u \in Z$. Thus we have $(u, X') \in R_{\xi_1^*}$ for some $X' \subseteq X$. \square

References

- Alur, R.; Henzinger, T. A.; and Kupferman, O. 1998. Alternating-time temporal logic. *Lecture Notes in Computer Science* 1536:23–60.
- Aumann, R. J., and Dreze, J. H. 2005. When all is said and done, how should you play and what should you expect? http://www.ma.huji.ac.il/raumann/pdf/dp_387.pdf.

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

Chatterjee, K.; Henzinger, T. A.; and Piterman, N. 2007. Strategy logic. In *Proceedings of the 18th International Conference on Concurrency Theory (CONCUR)*, volume 4703 of *Lecture Notes in Computer Science*, 59–73. Springer.

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

Goranko, V. 2001. Coalition games and alternating temporal logics. *Proceedings of 8th conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII)* 259–272.

Goranko, V. 2003. The basic algebra of game equivalences. *Studia Logica* 75(2):221–238.

Harrenstein, P.; van der Hoek, W.; Meyer, J.-J.; and Witteven, C. 2003. A modal characterization of Nash equilibrium. *Fundamenta Informaticae* 57:2-4:281–321.

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

Pauly, M. 2001. *Logic for Social Software*. Ph.D. Dissertation, University of Amsterdam.

Shoham, Y. 2003. Rational programming. <http://robotics.stanford.edu/~shoham/wwwpapers/RatProg.pdf>.

van Benthem, J.; Ghosh, S.; and Liu, F. 2007. Modelling simultaneous games with concurrent dynamic logic. In *A Meeting of Minds, Proceedings of the Workshop on Logic, Rationality and Interaction*, 243–258.

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

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

van der Hoek, W.; Jamroga, W.; and Wooldridge, M. 2005. A logic for strategic reasoning. *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)* 157–164.