

Axioms for composite 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 logic for reasoning about composite strategies in games, where players' strategies are like programs, composed structurally. These depend not only on conditions that hold at game positions but also on properties of other players' strategies. We present an axiomatization for the logic and prove its completeness.

1 Summary

Extensive form turn-based games are trees whose nodes are game positions and branches represent moves of players. With each node is associated a player whose turn it is to move at that game position. A player's strategy is then simply a subtree which contains a unique successor for every node where it is this player's turn to make a move, and contains all successors (from the game tree) for nodes where other players make moves. Thus a strategy is an advice function that tells a player what move to play when the game reaches any specific position. In two-player win/loss games, analysis of the game amounts to seeing if either player has a winning strategy from any starting position, and if possible, synthesize such a winning strategy.

In multi-player games where the outcomes are not merely winning and losing, the situation is less clear. Every player has a preference for certain outcomes and hence cooperation as well as conflict become strategically relevant. Moreover, each player has some expectations (and assumptions) about strategies adopted by other players, and fashions her response appropriately. In such situations, game theory tries to explain what *rational* players would do.

In so-called **small** (normal form) games, where the game consists of a small fixed number of moves (often one move chosen independently by each player), strategies have little structure, and prediction of stable behaviour (equilibrium strategy profiles) is possible. However, this not only becomes difficult in games with richer structure and long sequences of moves, it is also less clear how to postulate behaviour of rational players. Moreover, if we look to game theory not only for existence of equilibria but also *advice* to players on how to play, the structure of strategies followed by players becomes relevant.

Even in games of perfect information, if the game structure is sufficiently rich, we need to re-examine the notion of strategy as a function that determines a player's move in every game position. Typically, the game position is itself only **partially known**, in terms of properties that the player can test for. Viewed in this light, strategies are like **programs**, built up systematically from atomic decisions like *if b then a* where b is a condition checked by the player to hold (at some game position) and a is a move available to the player at that position.

There is another dimension to strategies, namely that of responses to other players' moves. The notion of each player independently deciding on a strategy needs to be re-examined as well. A player's chosen strategy depends on the player's perception of apparent strategies followed by other players. Even when opponents' moves are visible, an opponent's strategy is not known completely as a function. Therefore the player's strategy is necessarily partial as well.

The central idea of this paper is to suggest that it helps to study strategies given by their properties. Hence, assumptions about strategies can be partial, and these assumptions can in turn be structurally built into the specification of other strategies. This leads us to proposing a logical structure for strategies, where we can reason with assertions of the form “(partial) strategy σ ensures the (intermediate) condition α ”.

This allows us to look for *induction principles* which can be articulated in the logic. For instance, we can look at what conditions must be maintained locally (by one move) to influence an outcome eventually. Moreover, we can compare strategies in terms of what conditions they can enforce.

The main contributions of this paper are:

- We consider non-zero-sum games over finite graphs, and consider best response strategies (rather than winning strategies).
- The reasoning carried out works explicitly with the structure of strategies rather than existence of strategies.
- We present a logic with structured strategy specifications and formulas describe how strategies ensure outcomes.
- We present an axiom system for the logic and prove that it is complete.

1.1 Other work

Recently, the advent of computational tasks on the world-wide web and related security requirements have thrown up many game theoretic situations. For example, signing contracts on the web requires interaction between principals who do not know each other and typically distrust each other. Protocols of this kind which involve *selfish agents* can be easily viewed as strategic games of imperfect information. These are complex interactive processes which critically involve players reasoning about each others' strategies to decide on how to act.

Game logics are situated in this context, employing modal logics (in the style of logics of programs) to study logical structure present in games. Parikh's work on propositional game logic ([Par85]) initiated the study of game structure using algebraic properties. Pauly ([Pau01]) has built on this to provide interesting relationships between programs and games, and to describe coalitions to achieve desired goals. Bonnano ([Bon91]) suggested obtaining game theoretic solution concepts as characteristic formulas in modal logic. van Benthem ([vB01]) uses dynamic logic to describe games as well as strategies.

On the other hand, the work on Alternating Temporal Logic ([AHK98]) considers selective quantification over paths that are possible outcomes of games in which players and an environment alternate moves. Here, we talk of the existence of a strategy for a coalition of players to force an outcome. [Gor01] draws parallels between these two lines of work, that of Pauly's coalition logics and alternating temporal logic. It is to be noted that in these logics, the reasoning is about existence of strategies, and the strategies themselves do not figure in formulas.

In the work of [HvdHMW03] and [vdHJW05], van der Hoek and co-authors develop logics for strategic reasoning and equilibrium concepts and this line of work is closest to ours in spirit. Our point of departure is in bringing logical structure into strategies rather than treating strategies as atomic. In particular, the strategy specifications we use are partial (in the sense that a player may assume that an opponent plays a whenever p holds, without knowing under what conditions the opponent strategy picks another move b), allowing for more generality in reasoning.

2 Game Arenas

We begin with a description of game models on which formulas of the logic will be interpreted. We use the graphical model for extensive form turn-based multiplayer games, where at most one player gets to move at each game position.

Game Arena

Let $N = \{1, 2, \dots, n\}$ be a non-empty finite set of players and $\Sigma = \{a_1, a_2, \dots, a_m\}$ be a finite set of action symbols, which represent *moves* of players. A **game arena** is a finite graph $\mathcal{G} = (W, \longrightarrow, w_0, \chi)$ where W is the set of nodes which represents the *game positions*, $\longrightarrow: (W \times \Sigma) \rightarrow W$ is a function also called the move function, w_0 is the initial node of the game.

Let the set of successors of $w \in W$ be defined as $\vec{w} = \{w' \in W \mid w \xrightarrow{a} w' \text{ for some } a \in \Sigma\}$. A node w is said to be *terminal* if $\vec{w} = \emptyset$. $\chi : W \rightarrow N$ assigns to each node w in W the player who “owns” w : that is, if $\chi(w) = k$ and w is not terminal then player k has to pick a move at w .

In an arena defined as above, the play of a game can be viewed as placing a token on w_0 . If player k owns the game position w_0 i.e. $\chi(w_0) = k$ and she picks an action ‘ a ’ which is enabled for her at w_0 , then the new game position moves the token to w' where $w_0 \xrightarrow{a} w'$. A play in the arena is simply a sequence of such moves. Formally, a play in \mathcal{G} is a finite path $\rho = w_0 \xrightarrow{a_1} w_1 \xrightarrow{a_2} \dots \xrightarrow{a_k} w_k$ where w_k is terminal, or it is an infinite path $\rho = w_0 \xrightarrow{a_1} w_1 \xrightarrow{a_2} \dots$, and where $\forall i : w_i \xrightarrow{a_i} w_{i+1}$ holds. Let *Plays* denote the set of all plays in the arena.

With a game arena $\mathcal{G} = (W, \longrightarrow, w_0, \chi)$, we can associate its *tree unfolding* also referred to as the *extensive form game tree* $\mathcal{T} = (S, \Rightarrow, s_0, \lambda)$ where (S, \Rightarrow) is a countably infinite tree rooted at s_0 with edges labelled by Σ and $\lambda : S \rightarrow W$ such that:

- $\lambda(s_0) = w_0$.
- For all $s, s' \in S$, if $s \xrightarrow{a} s'$ then $\lambda(s) \xrightarrow{a} \lambda(s')$.
- If $\lambda(s) = w$ and $w \xrightarrow{a} w'$ then there exists $s' \in S$ such that $s \xrightarrow{a} s'$ and $\lambda(s') = w'$.

Given the tree unfolding of a game arena \mathcal{T} , a node s in it, we can define the *restriction* of \mathcal{T} to s , denoted \mathcal{T}_s to be the subtree obtained by retaining only the unique path from root s_0 to s and the subtree rooted at s .

Games and Winning Conditions

Let \mathcal{G} be an arena as defined above. The arena merely defines the rules about how the game progresses and terminates. More interesting are winning conditions, which specify the game outcomes. Naturally, each player may have some preferences relating to the outcomes.

What we need is a finitely specified condition describing each player’s preference on outcomes. We choose to use *regular* conditions: since the graph is finite, in every infinite play, some positions are visited infinitely often. Rather than specify preferences on plays, we do so on the set of positions through which the play *cycles*. Let C_G denote the set of all cycles in the game arena G and T_G denote the set of all terminal nodes in the game arena. For each $k \in N$, let $\preceq^k \subseteq ((C_G \cup T_G) \times (C_G \cup T_G))$ be a complete reflexive, transitive binary relation denoting the preference relation of player k . Clearly, we have induced orderings: for $k \in N$, let $\preceq^k \subseteq (Plays \times Plays)$, defined in the obvious manner.

Then a game is defined as the pair $G = (\mathcal{G}, (\preceq^i)_{i \in N})$.

Strategies

For simplicity we will restrict ourselves to two player games, i.e. $N = \{1, 2\}$. It is easy to extend the notions introduced here to the general case where we have n players.

Let the game graph be represented by $\mathcal{G} = (W^1, W^2, \longrightarrow, s_0)$ where W^1 is the set of positions of player 1, W^2 that of player 2. Let $W = W^1 \cup W^2$.

Let \mathcal{T} be the tree unfolding of the arena and s_1 a node in it. A *strategy* for player 1 at node s_1 is given by: $\mu = (S_\mu^1, S_\mu^2, \Rightarrow_\mu, s_1)$ is a subtree of \mathcal{T}_{s_1} which contains the unique path from root s_0 to s_1 in \mathcal{T} and is the least subtree satisfying the following properties:

- $s_1 \in S_\mu^1$, where $\chi(\lambda(s_1)) = 1$.
- For every s in the subtree of \mathcal{T}_G rooted at s_1 ,
 - if $s \in S_\mu^1$ then for some $a \in \Sigma$, for each s' such that $s \xrightarrow{a} s'$, we have $s \xrightarrow{a}_\mu s'$.
 - if $s \in S_\mu^2$, then for every $b \in \Sigma$, for each s' such that $s \xrightarrow{b} s'$, we have $s \xrightarrow{b}_\mu s'$.

Let Ω_i denote the set of all strategies of Player i in G , for $i = 1, 2$. A strategy profile $\langle \mu, \tau \rangle$ defines a unique play ρ_μ^τ in the game \mathcal{G} .

3 The logic

We now present a logic for reasoning about composite strategies. The syntax of the logic is presented in two layers, that of **strategy specification** and **game formulas**.

Atomic strategy formulas specify, for a player, what conditions she tests for before making a move. Since these are intended to be bounded memory strategies, the conditions are stated as **past time** formulas of a simple tense logic. Composite strategy specifications are built from atomic ones using connectives (without negation). We crucially use an implication of the form: “if the opponent’s play conforms to a strategy π then play σ ”.

Game formulas describe the game arena in a standard modal logic, and in addition specify the result of a player following a particular strategy at a game position, to choose a specific move a , to *ensure* an intermediate outcome α . Using these formulas one can specify how a strategy helps to eventually *win* an outcome α .

Before we describe the logic and give its semantics, some preliminaries will be useful. Below, for any countable set X , let $Past(X)$ be a set of formulas given by the following syntax:

$$\psi \in Past(X) := x \in X \mid \neg\psi \mid \psi_1 \vee \psi_2 \mid \diamond\psi.$$

Such past formulas can be given meaning over finite sequences. Given any sequence $\xi = t_0 t_1 \cdots t_m$, $V : \{t_0, \dots, t_m\} \rightarrow 2^X$, and k such that $0 \leq k \leq m$, the truth of a past formula $\psi \in Past(X)$ at k , denoted $\xi, k \models \psi$ can be defined as follows:

- $\xi, k \models p$ iff $p \in V(t_k)$.
- $\xi, k \models \neg\psi$ iff $\xi, k \not\models \psi$.
- $\xi, k \models \psi_1 \vee \psi_2$ iff $\xi, k \models \psi_1$ or $\xi, k \models \psi_2$.
- $\xi, k \models \diamond\psi$ iff there exists a $j : 0 \leq j \leq k$ such that $\xi, j \models \psi$.

Strategy specifications

For simplicity of presentation, we stick with two player games, where the players are Player 1 and Player 2. Let $\bar{i} = 2$ when $i = 1$ and $\bar{i} = 1$ when $i = 2$.

Let $P^i = \{p_0^i, p_1^i, \dots\}$ be a countable set of proposition symbols where $\tau_i \in P_i$, for $i \in \{1, 2\}$. Let $P = P^1 \cup P^2 \cup \{leaf\}$. τ_1 and τ_2 are intended to specify, at a game position, which player’s turn it is to move. *leaf* specifies whether the position is a terminal node.

Further, the logic is parametrized by the finite alphabet set $\Sigma = \{a_1, a_2, \dots, a_m\}$ of players’ moves and we only consider game arenas over Σ .

Let $Strat^i(P^i)$, for $i = 1, 2$ be the set of strategy specifications given by the following syntax:

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

where $\pi \in Strat^{\bar{i}}(P^1 \cap P^2)$, $\psi \in Past(P^i)$ and $a_k \in \Sigma$.

The idea is to use the above constructs to specify properties of strategies. For instance the interpretation of a player i specification $[p \mapsto a]^i$ will be to choose move “ a ” for every i node where p holds. $\pi \Rightarrow \sigma$ would say, 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} conforms to π .

For a game tree \mathcal{T} , a node s and a strategy specification $\sigma \in Strat^i(P^i)$, we define $\mathcal{T} \upharpoonright \sigma = (S_\sigma, \Longrightarrow_\sigma, s_0)$ to be the least subtree of \mathcal{T}_s which contains the unique path from s_0 to s and satisfies the following property.

- For every s' in S_σ ,
 - if s' is an i node then for all s'' with $s' \xrightarrow{a} s''$ and $a \in \sigma(s')$, we have $s' \xrightarrow{a}_\sigma s''$.
 - if s' is an \bar{i} node then for all s'' with $s' \xrightarrow{a} s''$, we have $s' \xrightarrow{a}_\sigma s''$.

Given a game tree \mathcal{T} and a node s in it, let $\rho_{s_0}^s : s_0 \xrightarrow{a_1} s_1 \cdots \xrightarrow{a_m} s_m = s$ denote the unique path from s_0 to s . For a strategy specification $\sigma \in Strat^i(P^i)$ and a node s we define $\sigma(s)$ as follows:

- $[\psi \mapsto a]^i(s) = \begin{cases} \{a\} & \text{if } s \in W^i \text{ and } \rho_{s_0}^s, m \models \psi \\ \Sigma & \text{otherwise} \end{cases}$
- $(\sigma_1 + \sigma_2)(s) = \sigma_1(s) \cup \sigma_2(s)$.
- $(\sigma_1 \cdot \sigma_2)(s) = \sigma_1(s) \cap \sigma_2(s)$.
- $(\pi \Rightarrow \sigma)(s) = \begin{cases} \sigma(s) & \text{if } \forall j : 0 \leq j < m, a_j \in \pi(s_j) \\ \Sigma & \text{otherwise} \end{cases}$

We say that a path $\rho_s^{s'} : s = s_1 \xrightarrow{a_1} s_2 \cdots \xrightarrow{a_{m-1}} s_m = s'$ in \mathcal{T} conforms to σ if $\forall j : 1 \leq j < m, a_j \in \sigma(s_j)$. When the path constitutes a proper play, i.e. when $s = s_0$, we say that the play conforms to σ .

Syntax

The syntax of the logic is given by:

$$\Pi ::= p \in P \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid \langle a \rangle \alpha \mid \langle \bar{a} \rangle \alpha \mid \Box \alpha \mid (\sigma)_i : c \mid \sigma \rightsquigarrow_i \beta$$

where $c \in \Sigma$, $\sigma \in \text{Strat}^i(P^i)$, $\beta \in \text{Past}(P^i)$. The derived connectives \wedge , \supset and $[a]\alpha$ are defined as usual. Let $\Diamond\alpha = \neg\Box\neg\alpha$, $\langle N \rangle \alpha = \bigvee_{a \in \Sigma} \langle a \rangle \alpha$, $[N]\alpha = \neg\langle N \rangle \neg\alpha$, $\langle P \rangle \alpha = \bigvee_{a \in \Sigma} \langle \bar{a} \rangle \alpha$ and $[P]\alpha = \neg\langle P \rangle \neg\alpha$.

The formula $(\sigma)_i : c$ asserts, at any game position, that the strategy specification σ for player i suggests that the move c can be played at that position. The formula $\sigma \rightsquigarrow_i \beta$ says that from this position, there is a way of following the strategy σ for player i so as to ensure the outcome β . These two modalities constitute the main constructs of our logic.

Semantics

The models for the logic are extensive form game trees along with a valuation function. A model $M = (\mathcal{T}, V)$ where $\mathcal{T} = (S^1, S^2, \longrightarrow, s_0)$ is a game tree as defined in section 2, and $V : S \rightarrow 2^P$ is the valuation function, such that:

- For $i \in \{1, 2\}$, $\tau_i \in V(s)$ iff $s \in S^i$.
- $leaf \in V(s)$ iff $\text{moves}(s) = \phi$.

where for any node s , $\text{moves}(s) = \{a \mid s \xrightarrow{a} s'\}$.

The truth of a formula $\alpha \in \Pi$ in a model M and position s (denoted $M, s \models \alpha$) is defined by induction on the structure of α , as usual. Let $\rho_{s_0}^s$ be $s_0 \xrightarrow{a_0} s_1 \cdots \xrightarrow{a_{m-1}} s_m = s$.

- $M, s \models p$ 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 $M, s \models \alpha_2$.
- $M, s \models \langle a \rangle \alpha$ iff there exists $s' \in W$ such that $s \xrightarrow{a} s'$ and $M, s' \models \alpha$.
- $M, s \models \langle \bar{a} \rangle \alpha$ iff $a = a_{m-1}$ and $M, s_{m-1} \models \alpha$.
- $M, s \models \Diamond\alpha$ iff there exists $j : 0 \leq j \leq m$ such that $M, s_j \models \alpha$.
- $M, s \models (\sigma)_i : c$ iff $c \in \sigma(s)$.
- $M, s \models \sigma \rightsquigarrow_i \beta$ iff
 - for all s' in $\mathcal{T}_s \upharpoonright \sigma$, such that $s \Longrightarrow^* s'$, we have $M, s' \models \beta \wedge (\tau_i \supset \text{enabled}_\sigma)$.

where $\text{enabled}_\sigma \equiv \bigvee_{a \in \Sigma} (\langle a \rangle \text{True} \wedge (\sigma)_i : a)$.

The notions of satisfiability and validity can be defined in the standard way. A formula α is satisfiable iff there exists a model M such that $M, s_0 \models \alpha$. A formula α is said to be valid iff for all models M , we have $M, s_0 \models \alpha$.




Fig. 1.

4 Example

Consider the game shown in Fig. 1. Players alternate moves with 1 starting at s_0 . There are two cycles $C_1 : s_5 \rightarrow s_6 \rightarrow s_7 \rightarrow s_8 \rightarrow s_5$, $C_2 : s_1 \rightarrow s_2 \rightarrow s_3 \rightarrow s_4 \rightarrow s_1$ and two terminal nodes t_1 and t_2 . Let the preference ordering of player 1 be $t_2 \preceq^1 t_1 \preceq^1 C_2 \preceq^1 C_1$. As far as player 2 is concerned $t_1 \preceq^2 C_1$ and he is indifferent between C_2 and t_2 . However, he prefers C_2 or t_2 over $\{C_1, t_1\}$. Equilibrium reasoning will advise player 1 to move “a” since at s_7 it is irrational for 2 to move x as it will result in 2’s worst outcome. However the utility difference between C_1 and t_1 for 2 might be negligible compared to the incentive of staying in the “left” path. Therefore 2 might decide to punish 1 for moving b when 1 knew that $\{C_2, t_2\}$ was equally preferred by 2. Even though t_1 is the worst outcome, at s_7 player 2 can play x to implement the punishment. Let $V(p_j) = \{s_3, s_7\}$, $V(p_{init}) = \{s_0\}$, $V(p_{good}) = \{s_0, s_1, s_2, s_3, s_4\}$ and $V(p_{punish}) = \{s_0, s_5, s_6, s_7, t_1\}$. The local objective of 2 will be to remain on the good path or to implement the punishment. Player 2 strategy specification can be written as

$$\pi \equiv ([p_{init} \mapsto b]^1 \Rightarrow [p_j \mapsto x]^2) \cdot ([p_{init} \mapsto a]^1 \Rightarrow [p_j \mapsto y]^2).$$

We get that $\pi \rightsquigarrow_2 (p_{good} \vee p_{punish})$. Player 1 if he knows 2’s strategy will be tempted to play “a” at s_0 by which the play will end up in C_2 .

5 Axiom system

We now present our axiomatization of the valid formulas of the logic. Before we present the axiomatization, we will find some abbreviations useful:

- $root = \neg \langle P \rangle True$ defines the root node to be one that has no predecessors.
- $\delta_i^\sigma(a) = \tau_i \wedge (\sigma)_i : a$ denotes that move “a” is enabled by σ at an i node.
- $inv_i^\sigma(a, \beta) = (\tau_i \wedge (\sigma)_i : a) \supset [a](\sigma \rightsquigarrow_i \beta)$ denotes the fact that after an “a” move by player i which conforms to σ , $\sigma \rightsquigarrow_i \beta$ continues to hold.
- $inv_{\bar{i}}^\sigma(\beta) = \tau_{\bar{i}} \supset [N](\sigma \rightsquigarrow_{\bar{i}} \beta)$ says that after any move of \bar{i} , $\sigma \rightsquigarrow_{\bar{i}} \beta$ continues to hold.
- $conf_\pi = \Box(\langle \bar{a} \rangle \tau_{\bar{i}} \supset \langle \bar{a} \rangle (\pi)_{\bar{i}} : a)$ denotes that all opponent moves in the past conform to π .

The axiom schemes

(A0) All the substitutional instances of the tautologies of PC

(A1) (a) $[a](\alpha_1 \supset \alpha_2) \supset ([a]\alpha_1 \supset [a]\alpha_2)$

- (b) $[\bar{a}](\alpha_1 \supset \alpha_2) \supset ([\bar{a}]\alpha_1 \supset [\bar{a}]\alpha_2)$
- (A2) (a) $\langle a \rangle \alpha \supset [a]\alpha$
 (b) $\langle \bar{a} \rangle \alpha \supset [\bar{a}]\alpha$
 (c) $\langle \bar{a} \rangle True \supset \neg \langle \bar{b} \rangle True$ for all $b \neq a$
- (A3) (a) $\alpha \supset [a]\langle \bar{a} \rangle \alpha$
 (b) $\alpha \supset [\bar{a}]\langle a \rangle \alpha$
- (A4) (a) $\diamond root$
 (b) $\Box \alpha \equiv (\alpha \wedge [P]\Box \alpha)$
- (A5) (a) $([\psi \mapsto a^i]_i : a \text{ for all } a \in \Sigma)$
 (b) $\tau_i \wedge ([\psi \mapsto a^i]_i : c \equiv \neg \psi \text{ for all } a \neq c)$
- (A6) (a) $(\sigma_1 + \sigma_2)_i : c \equiv \sigma_1 : c \vee \sigma_2 : c$
 (b) $(\sigma_1 \cdot \sigma_2)_i : c \equiv \sigma_1 : c \wedge \sigma_2 : c$
 (c) $(\pi \Rightarrow \sigma)_i : c \equiv conf_\pi \supset (\sigma)_i : c$
- (A7) $\sigma \rightsquigarrow_i \beta \supset (\beta \wedge inv_i^\sigma(a, \beta) \wedge inv_{\bar{i}}^\sigma(\beta) \wedge (\neg leaf \supset enabled_\sigma))$

Inference rules

$$\begin{array}{l}
 (MP) \frac{\alpha, \alpha \supset \beta}{\beta} \quad (NG) \frac{\alpha}{[a]\alpha} \quad (NG-) \frac{\alpha}{[\bar{a}]\alpha} \\
 (Ind-past) \frac{\alpha \supset [P]\alpha}{\alpha \supset \Box \alpha} \\
 (Ind \rightsquigarrow) \frac{\alpha \wedge \delta_i^\sigma(a) \supset [a]\alpha, \alpha \wedge \tau_{\bar{i}} \supset [N]\alpha, \alpha \wedge \neg leaf \supset enabled_\sigma, \alpha \supset \beta}{\alpha \supset \sigma \rightsquigarrow_i \beta}
 \end{array}$$

The axioms are mostly standard. After the Kripke axioms for the $\langle a \rangle$ modalities, we have axioms that ensure determinacy of both $\langle a \rangle$ and $\langle \bar{a} \rangle$ modalities, and an axiom to assert the uniqueness of the latter. We then have axioms that relate the previous and next modalities with each other, as well as to assert that the past modality steps through the $\langle \bar{a} \rangle$ modality. An axiom asserts the existence of the root in the past. The rest of the axioms describe the semantics of strategy specifications.

The rule *Ind-past* is standard, while *Ind \rightsquigarrow* illustrates the new kind of reasoning in the logic. It says that to infer that the formula $\sigma \rightsquigarrow_i \beta$ holds in all reachable states, β must hold at the asserted state and

- for a player i node after every move which conforms to σ , β continues to hold.
- for a player \bar{i} node after every enabled move, β continues to hold.
- player i does not get stuck by playing σ .

To see the soundness of (A7), suppose it is not valid. Then there exists a node s such that $M, s \models \sigma \rightsquigarrow_i \beta$ and one of the following holds:

- $M, s \not\models \beta$: in which case from semantics we get that $M, s \not\models \sigma \rightsquigarrow_i \beta$ which is a contradiction.
- $M, s \not\models inv_i^\sigma(a, \beta)$: In this case, we have $s \in W^i$, $M, s \models (\sigma)_i : a$ and $M, s' \not\models \sigma \rightsquigarrow_i \beta$ where $s \xrightarrow{a} s'$. This implies that there is a path $\rho_s^{s'k}$ which conforms to σ and either $M, s_k \not\models \beta$ or $moves(s_k) \cap \sigma(s_k) = \phi$. But since $s \xrightarrow{a} s'$, we have $\rho_s^{s'k}$ conforms to σ as well. From which it follows that $M, s \not\models \sigma \rightsquigarrow_i \beta$ which is a contradiction.
- $M, s \not\models inv_{\bar{i}}^\sigma(\beta)$: We have a similar argument as above.
- $M, s \not\models \neg leaf \supset enabled_\sigma$: This means that $M, s \models \neg leaf$ and $M, s \not\models enabled_\sigma$. Therefore $moves(s) \cap \sigma(s) = \phi$ and by semantics we have $M, s \not\models \sigma \rightsquigarrow_i \beta$ which is a contradiction.

To show that the induction rule preserves validity, suppose that the premise is valid and the conclusion is not. Then for some node s we have $M, s \models \alpha$ and $M, s \not\models \sigma \rightsquigarrow_i \beta$. i.e. there is a path $\rho_s^{s'k}$ which conforms to σ such that $M, s_k \not\models \beta$ or s_k is a non-leaf node and $\sigma(s_k) \cap moves(s_k) = \phi$. Consider the shortest such path. Since the premise is assumed to be valid, we get for all s' in $\rho_s^{s'k-1}$, $M, s' \models \alpha$. Assume that $\alpha \supset \beta$ is valid.

If $M, s_k \not\models \beta$ then we have two cases to consider

1. $s_{k-1} \in W^i$, in which case we have $M, s_{k-1} \models \alpha \wedge \delta_i^\sigma(a_{k-1})$ and $M, s_k \not\models \alpha$. From which we get that $M, s_{k-1} \not\models (\alpha \wedge \delta_i^\sigma(a_{k-1})) \supset [a_{k-1}]\alpha$.
2. $s_{k-1} \in W^{\bar{i}}$. We have $M, s_{k-1} \models \alpha \wedge \tau_{\bar{i}}$ and $M, s_{k-1} \not\models [a_{k-1}]\alpha$ from which we get $M, s_{k-1} \not\models (\alpha \wedge \tau_{\bar{i}}) \supset [N]\alpha$.

Either case we get a contradiction to the validity of a premise. If s_k is a non-leaf node and $\sigma(s_k) \cap \text{moves}(s_k) = \emptyset$ then we have $M, s_k \models \alpha \wedge \neg \text{leaf}$ and $M, s_k \not\models \text{enabled}_\sigma$. Therefore $M, s_k \not\models (\alpha \wedge \neg \text{leaf}) \supset \text{enabled}_\sigma$.

6 Completeness

To show completeness, we prove that every consistent formula is satisfiable. Let α_0 be a consistent formula, and let W denote the set of all maximal consistent sets (MCS). We use w, w' to range over MCS's. Since α_0 is consistent, there exists an MCS w_0 such that $\alpha_0 \in w_0$.

Define a transition relation on MCS's as follows: $w \xrightarrow{a} w'$ iff $\{\langle a \rangle \alpha \mid \alpha \in w'\} \subseteq w$.

We will find it useful to work not only with MCSs, but also with sets of subformulas of α_0 . For a formula α let $CL(\alpha)$ denote the subformula closure of α . Let \mathcal{AT} denote the set of all maximal consistent subsets of $CL(\alpha_0)$, referred to as atoms. Each $t \in \mathcal{AT}$ is a finite set of formulas, we denote the conjunction of all formulas in t by \hat{t} . For a nonempty subset $X \subseteq \mathcal{AT}$, we denote by \tilde{X} the disjunction of all $\hat{t}, t \in X$. Define a transition relation on \mathcal{AT} as follows: $t \xrightarrow{a} t'$ iff $\hat{t} \wedge \langle a \rangle \hat{t}'$ is consistent. Call an atom t a *root atom* if there does not exist any atom t' such that $t' \xrightarrow{a} t$ for some a .

Note that $t_0 = w_0 \cap CL(\alpha_0) \in \mathcal{AT}$.

Proposition 6.1. *There exist $t_1, \dots, t_k \in \mathcal{AT}$ and $a_1, \dots, a_k \in \Sigma$ ($k \geq 0$) such that $t_k \xrightarrow{a_k} t_{k-1} \dots \xrightarrow{a_1} t_0$, where t_k is a root atom.*

Proof. Consider the least set R containing t_0 and closed under the following condition: if $t_1 \in R$ and for some $a \in \Sigma$ there exists t_2 such that $t_2 \xrightarrow{a} t_1$, then $t_2 \in R$. Now, if there exists an atom $t' \in R$ such that t' is a root then we are done. Suppose not, then we have $\vdash \tilde{R} \supset \neg \text{root}$. But then we can show that $\vdash \tilde{R} \supset [P]\tilde{R}$. By rule *Ind-past* and above we get $\vdash \tilde{R} \supset \Box \neg \text{root}$. But then $t_0 \in R$ and hence $\vdash \hat{t}_0 \supset \tilde{R}$ and therefore we get $\vdash \hat{t}_0 \supset \Box \neg \text{root}$, contradicting axiom (A4a).

Above, we have additional properties: for any formula $\diamond \alpha \in t_k$, we also have $\alpha \in t_k$. Further, for all $j \in \{0, \dots, k\}$, if $\diamond \alpha \in t_j$, then there exists i such that $k \geq i \geq j$ and $\alpha \in t_i$. Both these properties are ensured by axiom (A4b).

Hence it is easy to see that there exist MCS's $w_1, \dots, w_k \in W$ and $a_1, \dots, a_k \in \Sigma$ ($k \geq 0$) such that $w_k \xrightarrow{a_k} w_{k-1} \dots \xrightarrow{a_1} w_0$, where $w_j \cap CL(\alpha_0) = t_j$. Now this path defines a (finite) tree $T_0 = (S_0, \Longrightarrow_0, s_0)$ rooted at s_0 , where $S_0 = \{s_0, s_1, \dots, s_k\}$, and for all $j \in \{0, \dots, k\}$, s_j is labelled by the MCS w_{k-j} . The relation \Longrightarrow_0 is defined in the obvious manner. From now we will simply say $\alpha \in s$ where s is the tree node, to mean that $\alpha \in w$ where w is the MCS associated with node s .

Inductively assume that we have a tree $T_k = (S_k, \Longrightarrow_k, s_0)$ such that the past formulas at every node have “witnesses” as above. Pick a node $s \in S_k$ such that $\langle a \rangle \text{True} \in s$ but there is no $s' \in S_k$ such that $s \xrightarrow{a} s'$. Now, if w is the MCS associated with node s , there exists an MCS w' such that $w \xrightarrow{a} w'$. Pick a new node $s' \notin S_k$ and define $T_{k+1} = S_k \cup \{s'\}$ and $\Longrightarrow_{k+1} = \Longrightarrow_k \cup \{(s, a, s')\}$, where w' is the MCS associated with s' . It is easy to see that every node in T_{k+1} has witnesses for past formulas as well.

Now consider $T = (S, \Longrightarrow, s_0)$ defined by: $S = \bigcup_{k \geq 0} S_k$ and $\Longrightarrow = \bigcup_{k \geq 0} \Longrightarrow_k$. Define the model $M = (T, V)$

where $V(s) = w \cap P$, where w is the MCS associated with s .

Lemma 6.1. *For any $s \in S$, we have the following properties.*

1. if $[a]\alpha \in s$ and $s \xrightarrow{a} s'$ then $\alpha \in s'$.
2. if $\langle a \rangle \alpha \in s$ then there exists s' such that $s \xrightarrow{a} s'$ and $\alpha \in s'$.

3. if $[\bar{a}]\alpha \in s$ and $s' \xrightarrow{a} s$ then $\alpha \in s'$.
4. if $\langle \bar{a} \rangle \alpha \in s$ then there exists s' such that $s' \xrightarrow{a} s$ and $\alpha \in s'$.
5. if $\exists \alpha \in s$ and $s' \xRightarrow{*} s$ then $\alpha \in s'$.
6. if $\diamond \alpha \in s$ then there exists s' such that $s' \xRightarrow{*} s$ and $\alpha \in s'$.

Lemma 6.2. For all $\psi \in \text{Past}(P)$, for all $s \in S$, $\psi \in s$ iff $\rho_s, s \models \psi$.

Lemma 6.3. For all i , for all $\sigma \in \text{Strat}^i(P^i)$, for all $c \in \Sigma$, for all $s \in S$, $(\sigma)_i : c \in s$ iff $c \in \sigma(s)$.

Proof. The proof is by induction on the structure of σ . The nontrivial cases are as follows:

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

(\Rightarrow) Suppose $([\psi \mapsto a])_i : c \in s$. If $c = a$ then the claim holds trivially. If $c \neq a$ then from (A5a) we get that $\neg \psi \in s$, from lemma 6.2 $\rho_s, s \not\models \psi$. Therefore by definition we have $[\psi \mapsto a]^i(s) = \Sigma$ and $c \in \sigma(s)$.

(\Leftarrow) Conversely, suppose $([\psi \mapsto a])_i : c \notin s$. From (A5a) we have $a \neq c$. From (A5b) we get $\psi \in s$. By lemma 6.2 $\rho_s, s \models \psi$. Therefore $c \notin \sigma(s)$ by definition.

$\sigma \equiv \sigma_1 + \sigma_2$:

$(\sigma_1 + \sigma_2)_i : c \in s$ iff by (A6a) and property of MCS, $(\sigma_1)_i \in s$ or $(\sigma_2)_i \in s$ iff by induction hypothesis, $c \in \sigma_1(s)$ or $c \in \sigma_2(s)$ iff $c \in \sigma_1(s) \cup \sigma_2(s)$ iff by definition $c \in (\sigma_1 + \sigma_2)(s)$.

$\sigma \equiv \sigma_1 \cdot \sigma_2$: A similar argument as above.

$\sigma \equiv \pi \Rightarrow \sigma'$: Let $\rho_{s_0}^s : s_0 \xrightarrow{a_0} \dots \xrightarrow{a_{k-1}} s_k = s$ be the unique path from the root to s .

(\Rightarrow) Suppose $(\pi \Rightarrow \sigma')_i : c \in s$. To show $c \in (\pi \Rightarrow \sigma')(s)$. Suffices to show that $\rho_{s_0}^s$ conforms to π implies $c \in \sigma'(s)$. From (A6c) we have $\text{conf}_\pi \supset (\sigma')_i : c \in s$. Rewriting this we get $\diamond(\langle \bar{a} \rangle \tau_{\bar{i}} \wedge [\bar{a}](\neg(\pi)_{\bar{i}} : a)) \vee (\sigma')_i : c \in s$. We have two cases,

- if $(\sigma')_i : c \in s$ then by induction hypothesis we get $c \in \sigma'(s)$. Therefore by definition $c \in (\pi \Rightarrow \sigma)_i(s)$.
- otherwise we have $\diamond(\langle \bar{a} \rangle \tau_{\bar{i}} \wedge [\bar{a}](\neg(\pi)_{\bar{i}} : a)) \in s$. From lemma 6.1(6), there exists $s_l \in \rho_s$ such that $\langle \bar{a} \rangle \tau_{\bar{i}} \wedge [\bar{a}](\neg(\pi)_{\bar{i}} : a) \in s_l$. By lemma 6.1(4) there exists $s_{l-1} \in \rho_s \cap W^{\bar{i}}$ such that $s_{l-1} \xrightarrow{a} s_l$. From lemma 6.1(3), $\neg(\pi)_{\bar{i}} : a \in s_{l-1}$. Since s_{l-1} is an MCS, we have $(\pi)_{\bar{i}} : a \notin s_{l-1}$. By induction hypothesis, $a \notin \pi(s_{l-1})$, therefore we have $\rho_{s_0}^s$ does not conform to π .

(\Leftarrow) Conversely, using (A6c) and a similar argument it can be shown that if $(\pi \Rightarrow \sigma')_i : c \notin s$ then $c \notin (\pi \Rightarrow \sigma')(s)$.

Theorem 6.1. For all $\alpha \in \Pi$, for all $s \in S$, $\alpha \in s$ iff $M, s \models \alpha$.

Proof. The proof is by induction on the structure of α .

$\alpha \equiv (\sigma)_i : c$.

From lemma 6.3 we have $(\sigma)_i : c \in s$ iff $c \in \sigma(s)$ iff by semantics $M, s \models (\sigma)_i : c$.

$\alpha \equiv \sigma \rightsquigarrow_i \beta$.

(\Rightarrow) We show the following:

1. If $\sigma \rightsquigarrow_i \beta \in s$ and there exists a transition $s \xrightarrow{a} s'$ such that $a \in \sigma(s)$, then $\{\beta, \sigma \rightsquigarrow_i \beta\} \subseteq s'$.
Suppose $\sigma \rightsquigarrow_i \beta \in s$, from (A7) we have $\beta \in s$. We have two cases to consider.
 - $s \in W^i$: We have $\tau_i \in s$. Since $a \in \sigma(s)$, by lemma 6.3 we have $(\sigma)_i : a \in s$. From (A7) we get $[a](\sigma \rightsquigarrow_i \beta) \in s$. By lemma 6.1(1) we have $\sigma \rightsquigarrow_i \beta \in s'$.
 - $s \in W^{\bar{i}}$: We have $\tau_{\bar{i}} \in s$. From (A7) we get $[N](\sigma \rightsquigarrow_i \beta) \in s$, since s is an MCS we have for every $a \in \Sigma$, $[a](\sigma \rightsquigarrow_i \beta) \in s$. By lemma 6.1(1) we have $\sigma \rightsquigarrow_i \beta \in s'$.

By applying (A7) at s' we get $\beta \in s'$.

2. If $\sigma \rightsquigarrow_i \beta \in s$ and s is a non-leaf node, then $\exists s'$ such that $s \xrightarrow{a} s'$ and $a \in \sigma(s)$.

Suppose s is a non-leaf node. From (A7), $\bigvee_{a \in \Sigma} (\langle a \rangle \text{True} \wedge (\sigma)_i : a) \in s$. Since s is an MCS, there exists an

a such that $\langle a \rangle \text{True} \wedge (\sigma)_i : a \in s$. By lemma 6.1(2), there exists an s' such that $s \xrightarrow{a} s'$ and by lemma 6.3 $a \in \sigma(s)$.

(1) ensures that whenever $\sigma \rightsquigarrow_i \beta \in s$ and there exists a path $\rho_s^{s_k}$ which conforms to σ , then we have $\{\beta, \sigma \rightsquigarrow_i \beta\} \subseteq s_k$. Since $\beta \in \text{Past}(P)$, by lemma 6.2 we have $M, s_k \models \beta$. (2) ensures that for all paths $\rho_s^{s_k}$ which conforms to σ , if s_k is a non-leaf node, then $\text{moves}(s) \cap \sigma(s) \neq \phi$. Therefore we get $M, s \models \sigma \rightsquigarrow_i \beta$. (\Leftarrow) Conversely suppose $\sigma \rightsquigarrow_i \beta \notin s$, to show $M, s \not\models \sigma \rightsquigarrow_i \beta$. Suffices to show that there exists a path $\rho_s^{s_k}$ that conforms to σ such that $M, s_k \not\models \beta$ or s_k is a non-leaf node and $\text{moves}(s_k) \cap \sigma(s_k) = \phi$.

Lemma 6.4. *For all $t \in \mathcal{AT}$, $\sigma \rightsquigarrow_i \beta \notin t$ implies there exists a path $\rho_t^{t_k} : t = t_1 \xrightarrow{a_1}_{\mathcal{AT}} t_2 \dots \xrightarrow{a_{k-1}}_{\mathcal{AT}} t_k$ which conforms to σ such that one of the following conditions hold.*

- $\beta \notin t_k$.
- t_k is a non-leaf node and $\text{moves}(t_k) \cap \sigma(t_k) = \phi$.

We have $t = s \cap \text{CL}(\sigma \rightsquigarrow_i \beta)$ is an atom. By lemma 6.4, there exists a path in the atom graph $t = t_1 \xrightarrow{a_1}_{\mathcal{AT}} t_2 \dots \xrightarrow{a_k}_{\mathcal{AT}} t_k$ such that $\beta \notin t_k$ or t_k is a non-leaf node and $\text{moves}(t_k) \cap \sigma(t_k) = \phi$. t_1 can be extended to the MCS s . Let $t'_2 = t_2 \cup \{\alpha \mid [a_1]\alpha \in s\}$. Its easy to check that t'_2 is consistent. Consider any MCS s_2 extending t'_2 , we have $s \xrightarrow{a_1} s_2$. Continuing in this manner we get a path in $s = s_1 \xrightarrow{a_1} s_2 \dots \xrightarrow{a_{k-1}} s_k$ in M which conforms to σ where either $\beta \notin s_k$ or s_k is a non-leaf node and $\text{moves}(s_k) \cap \sigma(s) = \phi$.

7 Discussion

We have defined a logic for reasoning about composite strategies in games. We have presented an axiomatization for the logic and shown its completeness.

We again remark that the presentation has been given for two-player games only for easy readability. It can be checked that all the definitions and arguments given here can be appropriately generalized for n -player games.

While our emphasis in the paper has been on advocating syntactically constructed strategies, we make no claims to having the “right” set of connectives for building them. This will have to be decided by experience, gained by specifying several kinds of strategies which turn out to be of use in reasoning about games.

We believe that a framework of this sort will prove useful in reasoning about multi-stage and repeated games, where strategy revision based on learning other players’ strategies (perhaps partially) plays an important role.

References

- [AHK98] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Lecture Notes in Computer Science*, 1536:23–60, 1998.
- [Bon91] G. Bonanno. The logic of rational play in games of perfect information. *Economics and Philosophy*, 7:37–65, 1991.
- [Gor01] V. Goranko. Coalition games and alternating temporal logics. *Proceedings of 8th conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 259–272, 2001.
- [HvdHMW03] Paul Harrenstein, Wiebe van der Hoek, John-Jules Meyer, and Cees Witteven. A modal characterization of nash equilibrium. *Fundamenta Informaticae*, 57:2–4:281–321, 2003.
- [Par85] Rohit Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–140, 1985.
- [Pau01] Marc Pauly. *Logic for Social Software*. PhD thesis, University of Amsterdam, October 2001.
- [vB01] Johan van Benthem. Games in dynamic epistemic logic. *Bulletin of Economic Research*, 53(4):219–248, 2001.
- [vdHJW05] Wiebe van der Hoek, Wojtek Jamroga, and Michael Wooldridge. A logic for strategic reasoning. *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)*, pages 157–164, 2005.