

Strategy Composition in Dynamic Games with Simultaneous Moves

Sujata Ghosh¹, Neethi Konar¹ and R. Ramanujam²

¹Indian Statistical Institute, Chennai, India

²Institute of Mathematical Sciences, Chennai, India

Keywords: Dynamic Games, Simultaneous Moves, Repeated Normal Form Games, Top-down Strategizing, Strategy Logic.

Abstract: Sometimes, in dynamic games, it is useful to reason not only about the existence of strategies for players, but also about what these strategies are, and how players select and construct them. We study dynamic games with simultaneous moves, repeated normal form games and show that this reasoning can be carried out by considering a single game, and studying composition of “local” strategies. We study a propositional modal logic in which such reasoning is carried out, and present a sound and complete axiomatization of the valid formulas.

1 INTRODUCTION

Game theory provides a normative view of what rational players should do in decision making scenarios taking into account their beliefs and expectations regarding other players’ (decision-makers’) behaviour. In this sense, the reasoning involved is more *about* games, rather than *within* the games themselves. The entire game structure is known to us, and we can predict how rational players would play, so that we can ask whether rational play could / would result in equilibrium strategy profiles, whereby players would not deviate from such strategic choices. Strategies of players in a game are assumed to be complete plans of actions that prescribe a unique move at each playable position in the game.

When games are finite, and strategies are complete plans, each player has only finitely many strategies to choose from and the study of *normal form games* simply abstracts strategies and sets of choices and studies the effect of each player making a choice simultaneously. Probabilistic (or *mixed* strategies) become the focus of such a study, and equilibrium theory then assigns probabilities to expected modes of player behaviour.

Even if we retain the structure of games, and study them as trees of possible sequences of player moves, as in *extensive form games*, a backward induction procedure (BI procedure (Osborne and Rubinstein, 1994)) can be employed to effectively compute optimal strategies for players, leading to predictions of stable play by rational players. Questions of how

players may arrive at selecting such strategies and playing them, and their expectations of other players symmetrically choosing such strategies are (rightly) glossed over.

And, when the moves are simultaneous, for example, in the case of infinitely *repeated normal form games*, even though there exists an equilibrium strategy (the grim strategy) (Rasmusen, 2007), Folk theorem tells us that under certain reasonable conditions, the claim that a particular behaviour arises is meaningless in such a game.

1.1 Strategies as Partial Plans

For dynamic games with simultaneous moves, one can think of both individual and group strategies, and the BI-procedure works bottom up on the game tree, and assumes players who have the computational and reasoning ability required to work it out (and hence each player can assume that other rational players do likewise). If the game is finite but consists of a tree of large size, such an assumption is untenable. Strategizing during play (rather than about the entire game tree) is meaningful in such situations and we need to consider players who are decisive and active agents but limited in their computational and reasoning ability.¹ Unlike the BI-procedure, strategizing during

¹The notion of players whose rationality is also limited in some way is interesting but more complex to formalize; for our considerations perfectly rational but resource bounded players suffice.

play follows the flow of time and hence works top down. Hence, unless a player has access to the entire subtree issuing at a node, she cannot compute optimal strategies, however well she is assured of their existence. It is in fact for this reason that although the determinacy of chess was established by Zermelo (Zermelo, 1913) the game remains fascinating to play as well as study even today.

Indeed, resource limited players working top down are forced to strategize *locally*, by selecting what part of the past history they choose to carry in their memory, and how much they can look ahead in their analysis. In combinatorial games, complexity considerations dictate such economizing in strategy selection. Predicting rational play by resource limited players is then quite interesting.

When game situations involve *uncertainty*, as inevitably happens in the case of games with large structure or large number of players, such top down strategizing is further necessitated by players having only a partial view of not only the past and the future but also the present as well. Once again, we are led to the notion of a strategy as something different from a complete plan, something analogous to a heuristic, whose applicability is dictated by local observations of game situations, for achieving local outcomes, based on expectations of other players' locally observed behaviour. The notion of locality in this description is imprecise, and pinning it down becomes an interesting challenge for a formal theory.

As an example, consider a heuristic in chess such as *pawn promotion*. This is generic advice to any player in any chess game, but it is local in the sense that it fulfils only a short term goal, it is not an advice for winning the game. A more interesting example is the heuristic employed by the computer Deep Blue against Gary Kasparov (on February 10, 1996) threatening Kasparov's queen with a knight (in response to Kasparov's 11th move). The move famously slowed down Kasparov for 27 minutes, and was later hailed as an important strategy.² The point is that such strategizing involves more than "look-ahead".

1.2 Strategy Selection

The foregoing discussion motivates a formal study of strategies in extensive form games, where we go beyond looking for existence of strategies for players to ensure desired outcomes, but take into account strategy structure and strategy selection as well. This work was initiated in (Ramanujam and Simon, 2008b), and we extend it to incorporate extensive

² https://en.wikipedia.org/wiki/Deep_Blue_versus_Kasparov,_1996,_Game_1

form games with simultaneous moves, so that they too can be studied through 'top-down strategizing'. Relating strategy structure with game structure was independently taken up in (Ghosh, 2008) and (Ramanujam and Simon, 2008a), see (Ghosh and Ramanujam, 2012) for a survey of such work.

When we reason about the existence of strategies for players to ensure desired outcomes, we can simply formalize the collection of strategies as a set and a rational player can be depended upon to pick the right one from the set and play it. In this case, reasoning about strategies amounts to assigning names to strategies and we can speak of player A playing a strategy σ to ensure outcome α from game position s . Such a player, in effect, chooses between outcomes and we reason about players' consideration of choices by other players.

In the case of a resource limited player who relies on partial and local plans, choice of strategies can be seen as composition of local plans to make 'more global' plans, and in this sense we can ascribe structure to strategies. As an example, consider a chess player who strategizes locally to capture either a knight or a bishop, and makes further conditional plans based on the success of either attempt, while at the same time formulating backup plans to counter unforeseen disasters along the way. In such a situation, each player can be seen as composing partial strategies, and exercising selections at each stage of the composition. Such strategy structure would then include hypothesizing about partial strategies of other players (as witnessed by their moves) as well. This mutual recursion in strategy structure and selection has been explicated in the logical study of (Ramanujam and Simon, 2008b) and in this paper we extend the approach to a class of games with simultaneous moves.

1.3 Related Work

When an extensive form game is presented as a finite or infinite tree, strategies constitute selective quantification over paths. When every edge of the tree corresponds to a normal form game, we obtain a *concurrent game structure*. The temporal evolution of such structures is studied in the pioneering work on Alternating time temporal logic (ATL) (Alur et al., 2002). In this logic, one reasons about groups of players having a strategy to force certain outcomes. Since the game played at one node of the tree is essentially different from that at any other node in the tree, strategizing is local at any node in the basic framework. These can be construed as subgames and hence ATL can be also seen as a logic of game composition. Further,

named strategies can be introduced as in extensions of ATL with explicit treatment of strategies (such as in (Chatterjee et al., 2007; van der Hoek et al., 2005; Walther et al., 2007; Ågotnes, 2006)). However, they principally reason about the existence of functional strategies in both normal form and extensive form games. For a detailed survey, see (Bulling et al., 2015).

Strategy composition arises from a different perspective. The point of departure here is in working with the heuristic notion of strategies as partial plans, and studying compositional structure in strategies. In this sense the contrast of ATL to this work is akin to that of temporal logics to process logics (which incorporate dynamic logic into temporal reasoning).

The *stit* frameworks of (Horty, 2001) work with notions such as “agent sees to it that” a particular condition holds, and automatically refers to agents having strategies to achieve certain goals. Extensive form game versions of *stit* have been discussed in (Broersen, 2009; Broersen, 2010), where strategies are considered as sets of histories of the game. Each such history gives a full play of the game, and hence, only *total* strategies are taken into account. See (Broersen and Herzig, 2015) for a detailed survey. We note here that such reasoning may be relevant for strategy logics with more detailed agency. Strategies as move recommendations of players based on game description language are considered in (Zhang and Thielscher, 2015).

Our work, which is an extension of the work done in (Ramanujam and Simon, 2008b), is closest in spirit to logical studies on games such as (Bentham, 2002; Bentham, 2007), (Harrenstein et al., 2003), (Bonanno, 2001). However, rather than formalizing the notion of backward induction and its epistemic enrichments, we study top down reasoning in the same basic framework as in these logics.

An aspect that comes out of our studies of compositional strategy structures is the fact that we are able to model players’ responses to other players *while playing a game*, even in the case of games with simultaneous moves, which occurs naturally in *repeated normal form games*. Thus one can study the different strategic responses of the players leading to different outcomes. Simple games or strategies are combined to form complicated structures which provide a way to describe actual plays of a game - how a player with limited resources, without knowing how a game might proceed in the future, can actually go about playing the game.

1.4 Contributions

Thus the main contributions of this paper are the following:

- We propose a logical language in which strategies are partial plans, have compositional structure and we reason about agents employing such strategies to achieve desired outcomes. The focus is on games with simultaneous moves (also known as concurrent games structures).
- We present a complete Hilbert-style axiomatization of this logic. A decision procedure can be extracted with some more work, as demonstrated in (Ramanujam and Simon, 2008b).

The syntax for composing strategies presented here is not proposed to be definitive, but merely illustrative. When we build libraries of strategies employed in games that programs modelling player agents make use of, we will have a more realistic understanding of compositional structure in strategies.

We note here that plays in many popular dynamic board games with simultaneous moves (e.g. RoboRally) are actually based on heuristic strategizing and local plans, and a library of strategies will aid the game developers and the general game playing crowd towards providing a better game-playing and strategy-developing experience. In what follows we will use the repeated normal form game, Iterated Prisoner’s Dilemma (Rasmusen, 2007) to explicate the concepts introduced.

2 PRELIMINARIES

We start with defining dynamic games with simultaneous moves, the basic underlying structure for this work, and what is meant by player-strategies in such games. Let $N = \{1, 2, \dots, n\}$ be a non-empty set of players. For each player $i \in N$, we associate a finite set Γ^i , the set of symbols which constitute player i ’s actions. Let $\tilde{\Gamma} = \prod_{i \in N} \Gamma^i$, the cartesian product of the Γ^i ’s over all $i \in N$. Throughout the text we will denote the elements of $\tilde{\Gamma}$ by γ . For each $i \in N$, let γ^{-i} denote the tuple $(\gamma^1, \dots, \gamma^{i-1}, \gamma^{i+1}, \dots, \gamma^n)$, and $\tilde{\Gamma}^{-i}$ denote the set of all such γ^{-i} ’s.

In case of Iterated Prisoner’s Dilemma (*IPD*), we have that $N = \{1, 2\}$. For each player $i \in N$, we associate a finite set $\Gamma^i = \{c_i(\text{confess}), d_i(\text{defect})\}$, and $\tilde{\Gamma} = \prod_{i \in N} \Gamma^i$.

Game Arena with Simultaneous Moves. A game arena with simultaneous moves is a tuple $G = (W, \rightarrow, w_0, \chi)$ such that W is the set of game positions, w_0 is the initial game position,

$\chi(w) = \chi^1(w) \times \dots \times \chi^n(w)$, for any $w \in W$, with $\chi^i : W \rightarrow 2^{\Gamma^i} \setminus \{\emptyset\}$ giving the set of possible actions that an agent can take when she is at a certain game position, and the move function $\rightarrow : W \times \tilde{\Gamma} \rightarrow W$ satisfies that for all $w, v \in W$, if $w \xrightarrow{\gamma} v$, then $\gamma[i] \in \chi^i(w)$.

Extensive Form Game Tree. Given a game arena $G = (W, \rightarrow, w_0, \chi)$, we can associate its tree unfolding, also referred to as the extensive form game tree, $T_G = (S, \Rightarrow, s_0, X, \lambda)$, where (S, \Rightarrow) is a tree rooted at s_0 with edges labelled by members of $\tilde{\Gamma}$, $X : S \rightarrow 2^{\tilde{\Gamma}}$, and $\lambda : S \rightarrow W$ is such that:

- $\lambda(s_0) = w_0$,
- $\forall s, s' \in S$, if $s \xrightarrow{\gamma} s'$, then $\lambda(s) \xrightarrow{\gamma} \lambda(s')$,
- if $\lambda(s) = w$ and $w \xrightarrow{\gamma} w'$, then $\exists s' \in S$ s.t. $s \xrightarrow{\gamma} s'$ and $\lambda(s') = w'$.
- $X(s) = \chi(\lambda(s))$

We have $X^i(s) = \chi^i(\lambda(s))$. We define $moves(s) = \{\gamma \in \tilde{\Gamma} : \exists s' \in S, \text{ such that } s \xrightarrow{\gamma} s'\}$, and $moves(s)_i$ is the set of i th projections of the members of $moves(s)$.

When we are given the tree unfolding T_G of a game arena G and a node s in it, we define the restriction of T_G to s , denoted by T_s to be the subtree obtained by retaining the unique path from root s_0 to s and the tree rooted at s .

Let us come back to our example. The game arena in case *IPD* is a tuple $G = (W, \rightarrow, w_0, \chi)$ such that $W = \{w\}$, $w_0 = w$ is the initial and only game position, $\chi(w) = \chi^1(w) \times \chi^2(w)$, with $\chi^i : w \mapsto \Gamma^i$ giving the set of possible actions that an agent can take when she is at a certain game position. The arena consists of a single game state and four loops labelled by (c_1, c_2) , (c_1, d_2) , (d_1, c_2) and (d_1, d_2) . The tree unfolding of this arena gives us the *IPD*. Let us now define strategies in the following.

Strategies. Let a game be represented by $G = (W, \rightarrow, w_0, \chi)$. Let T_G be the tree unfolding of the game arena G and s be a node in it. A strategy for player i at node s , μ^i is given by $T_s \upharpoonright \mu^i = (S_{\mu^i}, \Rightarrow_{\mu^i}, s, X_{\mu^i})$ which is the subtree of T_s containing the unique path from root s_0 to s and is the least subtree satisfying the following property: $\forall s \in S_{\mu^i}, \exists$ a unique $\gamma^i \in \Gamma^i$ such that,

$\forall \gamma^{-i} \in \tilde{\Gamma}^{-i}, \forall s'$ with $s \xrightarrow{(\gamma^i, \gamma^{-i})} s', s \xrightarrow{(\gamma^i, \gamma^{-i})} \mu^i s'$. X_{μ^i} is the restriction of X to S_{μ^i} . The idea is that we pick a single action for player i and all possible actions for other players and consider those tuples of moves corresponding to each game position in the subtree rooted at s .

For example, the tuples of actions (c_1, c_2) and (c_1, d_2) at each node of the *IPD* tree constitute a strategy tree

for player 1. Let Ω^i denote the set of all strategies of player i in G . Given a game tree T_G and a node s in it, let $\rho_{s_0}^s : s_0 \xrightarrow{\gamma_1} s_1 \dots \xrightarrow{\gamma_m} s_m = s$ denote the unique path from s_0 to s . In what follows we restrict the number n of elements in N to 2 for convenience. The ensuing discussion and results would follow similarly (with minor modifications) for any arbitrary $n \geq 2$.

3 STRATEGY LOGIC

We now propose a strategy logic (*SL*) to reason about such strategies and their compositions and to describe what these strategies can ensure in the positions where they are enabled. We give the syntax in two levels, as in (Ramanujam and Simon, 2008b), the first level consists of the *strategy specification language*, and the second level provides a syntax for reasoning in games with simultaneous moves.

3.1 Strategy Specifications

We first provide a specification language to describe such strategies. These strategies can be given as advices from the point of view of an outsider advising the players how to play in such games, and also can be considered from the players' perspectives at game positions, given the facts that hold in such positions. A propositional syntax with certain past formulas is used to describe observations at game positions.

Observation Syntax. Let \mathbb{P} be a countable set of propositions. Then we define,

$$\phi \in \Phi : p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid \diamond\phi$$

where $p \in \mathbb{P}$. Here, $\diamond\phi$ denotes that ϕ has happened sometime in the past and can be evaluated in terms of finite sequences.

Specification Syntax. For $i = 1, 2$, the set of strategy specifications, $Strat^i(\mathbb{P})$ is given by,

$$\sigma \in Strat^i(\mathbb{P}) : [\phi \mapsto \gamma^i]^i \mid \sigma_1 + \sigma_2 \mid \sigma_1 \cdot \sigma_2 \mid \pi \Rightarrow_i \sigma,$$

where, $\phi \in \Phi$, and $\pi \in Strat^{\bar{i}}(\mathbb{P})$. Here $\bar{i} = 2(1)$ if $i = 1(2)$.

The main idea is to use the above constructs to specify properties of strategies as well as to combine them to describe a play of the game, say. For instance the interpretation of a player i 's specification $[\phi \mapsto \gamma^i]^i$ where $\phi \in \Phi$, is to choose move γ^i at every game position where ϕ holds. At positions where ϕ does not hold, the strategy is allowed to choose any enabled move. The constructs '+' and ' \cdot ' correspond

to ‘or’ and ‘and’ respectively. The strategy specification $\sigma_1 + \sigma_2$ says that the strategy of player i conforms to the specification σ_1 or σ_2 . The construct $\sigma_1 \cdot \sigma_2$ says that the strategy conforms to specifications σ_1 and σ_2 . The strategy specification $\pi \Rightarrow_i \sigma$ says that if player \bar{i} 's strategy conforms to the specification π in the history of the game, then play σ .

Semantics. Let $V : S \rightarrow 2^{\mathbb{P}}$ be a valuation. Then a model is given by $M = (T_G, V)$, where T_G is an extensive form game tree (defined in Section 2) and V is a valuation. Let M be a model and s be a node in it. Let $\rho_{s_0}^s : s_0 \xrightarrow{\gamma_1} s_1 \dots \xrightarrow{\gamma_m} s_m = s$ denote the unique path from s_0 to s . We define for all $k \in \{0, \dots, m\}$,

- $\rho_{s_0}^s, k \models p$ iff $p \in V(s_k)$,
- $\rho_{s_0}^s, k \models \neg\phi$ iff $\rho_{s_0}^s, k \not\models \phi$,
- $\rho_{s_0}^s, k \models \phi_1 \vee \phi_2$ iff $\rho_{s_0}^s, k \models \phi_1$ or $\rho_{s_0}^s, k \models \phi_2$,
- $\rho_{s_0}^s, k \models \diamond\phi$ iff there exists a $j : 0 \leq j < k$ such that $\rho_{s_0}^s, j \models \phi$.

For any $\sigma^i \in \text{Strat}^i(\mathbb{P})$, we can define $\sigma^i(s)$ (the set of player i actions at node s conforming to the strategy specification σ^i) as follows:

$$([\phi \mapsto \gamma^i](s)) = \begin{cases} \{\gamma^i\} & \text{if } \rho_{s_0}^s, m \models \phi \text{ and } \gamma^i \in X^i(s), \\ X^i(s) & \text{otherwise;} \end{cases}$$

$$(\sigma_1^i + \sigma_2^i)(s) = \sigma_1^i(s) \cup \sigma_2^i(s),$$

$$(\sigma_1^i \cdot \sigma_2^i)(s) = \sigma_1^i(s) \cap \sigma_2^i(s).$$

$$(\pi \Rightarrow_i \sigma)(s) = \begin{cases} \sigma(s) & \text{if } \forall j : 0 \leq j < m, \gamma_j^{\bar{i}} \in \pi(s_j) \\ X^i(s) & \text{otherwise.} \end{cases}$$

Here, $X^i(s)$ denotes the set of moves enabled for a player i at s , $\gamma_j^{\bar{i}}$ denotes an action of player \bar{i} , and $\pi(s_j)$ denotes the set of actions enabled at s_j for player \bar{i} by the strategy specification π .

Given a game tree T_G and a node s in it, and a strategy specification $\sigma \in \text{Strat}^i(\mathbb{P})$, we define $T_s \upharpoonright \sigma = (S_\sigma, \Rightarrow_\sigma, s, X_\sigma)$ to be the subtree of T_s containing the unique path from root s_0 to s and is the least subtree satisfying the following property: $\forall s \in S_\sigma, \exists \gamma^i \in \Gamma^i$ such that, $\forall \gamma^{\bar{i}} \in \Gamma^{\bar{i}}, \forall s', s \xrightarrow{\gamma^i} s'$, where $\gamma = (\gamma^i, \gamma^{\bar{i}})$ iff $\gamma^i \in \sigma(s)$. X_σ is the restriction of X to S_σ .

In *IPD* one can consider strategies like *always cooperate* or *always defect*, which can be easily described in the syntax as follows:

- always cooperate - $[\top \mapsto c]^i$
- always defect - $[\top \mapsto d]^i$

Consider the following non-trivial strategy:

- Start by choosing d (defect)
- As long as the opponent cooperates, cooperate, or as long as the opponent defects, defect

This strategy can be represented as follows: $[root \mapsto d]^i.([\top \mapsto c]^{\bar{i}} \Rightarrow_i [\top \mapsto c]^i) + ([\top \mapsto d]^{\bar{i}} \Rightarrow_i [\top \mapsto d]^i)$. Here, $root$ can be considered as an atomic formula, true only at the root node. A commonly known strategy is given in the following. If both players choose this strategy in the infinite *IPD*, then one can find a simple perfect equilibrium (Rasmusen, 2007).

The Grim Strategy

- Start by choosing d
- Continue to choose d unless some player has chosen c , in which case, choose c forever.

We will show how this Strategy Logic, *SL* can describe such strategies. Let us now provide a language to reason with such strategies.

3.2 Logic

We are now ready to propose a logic, *SL* for describing strategic reasoning in games with simultaneous moves.

SL syntax The syntax of *SL* is defined as follows:

$$\alpha \in \Theta : p \in \mathbb{P} \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid \langle \gamma \rangle \alpha \mid \overline{\langle \gamma \rangle} \alpha \mid \diamond\alpha \mid \sigma^i : \gamma^i \mid \sigma^i \hookrightarrow_i \phi$$

where, $\gamma^i \in \Gamma^i$, $\sigma^i \in \text{Strat}^i(\mathbb{P})$, and $\phi \in \Phi$.

The formula $\langle \gamma \rangle \alpha$ says that after a γ move α holds, and $\overline{\langle \gamma \rangle} \alpha$ says that α holds before the γ move was taken. The formula $\sigma^i : \gamma^i$ asserts, at any game position, that the strategy specification σ^i for player i suggests that the move γ^i can be played at that position. The formula $\sigma^i \hookrightarrow_i \phi$ says that from this position, there is a way of following the strategy σ^i for player i so as to ensure the outcome ϕ . These two modalities constitute the main constructs of our logic.

The connectives \wedge, \supset and the formulas $[\gamma]\alpha, \overline{[\gamma]}\alpha$ are defined as usual, $\exists\alpha = \neg\forall\neg\alpha$, $\langle N \rangle \alpha = \bigvee_{\gamma \in \Gamma} \langle \gamma \rangle \alpha$, $[N]\alpha = \neg\langle N \rangle \neg\alpha$, $\langle P \rangle \alpha = \bigvee_{\gamma \in \Gamma} \overline{\langle \gamma \rangle} \alpha$, $[P]\alpha = \neg\langle P \rangle \neg\alpha$, $root = \neg\langle P \rangle \top$, $enabled_{\gamma^1} = \bigvee_{\gamma^2 \in \Gamma^2} \langle (\gamma^1, \gamma^2) \rangle \top$, $enabled_{\gamma^2} = \bigvee_{\gamma^1 \in \Gamma^1} \langle (\gamma^1, \gamma^2) \rangle \top$.

SL semantics Let M be a model as described in Section 3.1 and s be a node in it. Let $\rho_{s_0}^s$ be $s_0 \xrightarrow{\gamma_1} s_1 \dots \xrightarrow{\gamma_m} s_m = s$, as earlier. The truth definition of the *SL* formulas is given 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, $M, s \models \alpha_2$,
- $M, s \models \langle \gamma \rangle \alpha$ iff $\exists s' \in S$ s.t. $s \xrightarrow{\gamma} s'$ and $M, s' \models \alpha$,
- $M, s \models \overline{\langle \gamma \rangle} \alpha$ iff $m > 0, \gamma = \gamma_m$ and $M, s_{m-1} \models \alpha$,
- $M, s \models \diamond \alpha$ iff $\exists j : 0 \leq j < m$ s.t. $M, s_j \models \alpha$,
- $M, s \models \sigma^i ; \gamma^i$ iff $\gamma^i \in \sigma^i(s)$,
- $M, s \vdash \sigma^i \hookrightarrow_i \phi$ iff for all $s' \in T_s \downarrow \sigma^i$ with $s \Rightarrow^* s'$ we have $M, s' \models \phi \wedge \text{enabled}_{\sigma^i}$, where $\text{enabled}_{\sigma^i} = \bigvee_{\gamma \in \Gamma} (\langle \gamma \rangle \top \wedge \sigma^i ; \gamma^i)$.

The Grim Strategy (followed by player 1, say) described earlier can be represented as follows: ($\text{root} \supset ((d_1, c_2) \top \vee (d_1, d_2) \top) \wedge ((P) \diamond ((\top \mapsto c_2)^2 + [\top \mapsto d_2]^2) : 2c_2) \supset ((c_1, c_2) \top \vee (c_1, d_2) \top)$). We now provide the main technical result of the paper, that is a sound and complete axiomatization of SL .

4 A COMPLETE AXIOMATIZATION

The following axioms and rules provide a sound and complete axiomatization for SL . A proof sketch is provided below.

- (A₁) tautologies in classical propositional logic.
- (A₂) (a) $[\gamma](\alpha_1 \supset \alpha_2) \supset ([\gamma]\alpha_1 \supset [\gamma]\alpha_2)$
(b) $\overline{[\gamma]}(\alpha_1 \supset \alpha_2) \supset (\overline{[\gamma]}\alpha_1 \supset \overline{[\gamma]}\alpha_2)$
- (A₃) (a) $\langle \gamma \rangle \alpha \supset \overline{[\gamma]}\alpha$
(b) $\overline{\langle \gamma \rangle} \alpha \supset [\gamma]\alpha$
(c) $\langle \gamma \rangle \top \supset \neg \overline{\langle \gamma \rangle} \top$, for all $\gamma' \neq \gamma$
- (A₄) (a) $\alpha \supset [\gamma]\overline{\langle \gamma \rangle} \alpha$
(b) $\alpha \supset \overline{[\gamma]}\langle \gamma \rangle \alpha$
- (A₅) (a) $\diamond \text{root}$
(b) $\Box \alpha \equiv (\alpha \wedge [P] \Box \alpha)$
- (A₆) (a) $\text{enabled}_{\gamma^i} \equiv [\phi \mapsto \gamma^i]^i ; \gamma^i, \gamma^i \in \Gamma^i$
(b) $[\phi \mapsto \gamma^i]^i ; \gamma^i \equiv \neg \phi \wedge \text{enabled}_{\gamma^i}, \gamma^i \neq \gamma^i$
- (A₇) (a) $(\sigma_1^i + \sigma_2^i) ; \gamma^i \equiv \sigma_1^i ; \gamma^i \vee \sigma_2^i ; \gamma^i$
(b) $(\sigma_1^i \cdot \sigma_2^i) ; \gamma^i \equiv \sigma_1^i ; \gamma^i \wedge \sigma_2^i ; \gamma^i$
(c) $(\pi \mapsto_i \sigma) ; \gamma^i \equiv ((\Box(\langle \gamma \rangle \top \supset \overline{\langle \gamma \rangle}(\pi ; \gamma^i)) \supset (\sigma ; \gamma^i)) \wedge (\neg \Box(\langle \gamma \rangle \top \supset \overline{\langle \gamma \rangle}(\pi ; \gamma^i)) \supset \text{enabled}_{\gamma^i}))$
- (A₈) $\sigma^i \hookrightarrow_i \phi \supset (\phi \wedge \boxtimes \wedge \text{enabled}_{\sigma^i})$, where \boxtimes is $\sigma^i ; \gamma^i \supset [\gamma^i](\sigma^i \hookrightarrow_i \phi)$.
- (A₉) (a) $\bigvee_{\gamma^i \in \Gamma^i} \text{enabled}_{\gamma^i}$
(b) $\bigvee_{\gamma^i \in \Gamma^i} \text{enabled}_{\gamma^i}$

Inference Rules

$$\text{M.P. } \frac{\alpha, \alpha \supset \beta}{\beta}$$

$$\text{G}_1 \frac{\alpha}{[\gamma]\alpha} \quad \text{G}_2 \frac{\alpha}{\overline{[\gamma]}\alpha}$$

$$\text{P } \frac{\alpha \supset [P]\alpha}{\alpha \supset \Box \alpha}$$

$$\text{I}_1 \frac{\alpha \wedge \sigma^1 ; \gamma^1 \supset \bigvee_{\gamma^2 \in \Gamma^2} [\gamma^1, \gamma^2]\alpha, \alpha \supset \text{enabled}_{\sigma^1}. \alpha \supset \phi}{\alpha \supset \sigma^1 \hookrightarrow_1 \phi}$$

$$\text{I}_2 \frac{\alpha \wedge \sigma^2 ; \gamma^2 \supset \bigvee_{\gamma^1 \in \Gamma^1} [\gamma^1, \gamma^2]\alpha, \alpha \supset \text{enabled}_{\sigma^2}. \alpha \supset \phi}{\alpha \supset \sigma^2 \hookrightarrow_2 \phi}$$

The axioms (A₂) and rules G₁ and G₂ show that the modalities $[\gamma]$ and $\overline{[\gamma]}$ are normal modalities. Axioms (A₃) account for uniquely labelled moves. Axioms (A₄) show that the modalities $[\gamma]$ and $\overline{[\gamma]}$ are converses of each other. Axioms (A₅) and rule P take care of the past modality. Axioms (A₆) deal with the atomic strategy advices, whereas axioms (A₇) deal with the complex ones. For example, axiom (A₇(c)) says that the response specification suggests the move γ^i iff whenever the opponent strategy specification π suggests γ^i in the history of the game, player i strategy σ suggests γ^i , and if not, γ^i has to be enabled at the game position. Finally, the (A₈) axiom (if ϕ is ensured under the strategy σ then ϕ holds, σ is enabled and if σ suggests γ^i at the game position, then after every γ move containing γ^i , ϕ is ensured under the strategy σ) and the I₁ and I₂ rules give the necessary and sufficient conditions for ensuring ϕ under the strategy σ^i , and (A₉) ensures infiniteness of each branch of the game tree. Note that we are assuming game trees without leaf nodes. This is not a matter of principle but rather of convenience. We can incorporate finite branches to our game trees, thereby having leaf nodes and change the axiomatization accordingly. The completeness proof will work.

The validities of these axioms and rules can be checked routinely. We now provide a proof sketch for the completeness.

4.1 Completeness

To prove completeness, we have to prove that every consistent formula is satisfiable. We mention all the intermediate propositions and lemmas to prove the result without providing any detailed proof due to lack of space.

Let α_0 be a consistent formula and let W denote the set of 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: $w \xrightarrow{(\gamma^1, \gamma^2)} w'$ iff $\{\langle \gamma^1, \gamma^2 \rangle \alpha \mid \alpha \in w'\} \subseteq w$. For a formula α , let $cl(\alpha)$ denote the subformula closure of α . In addition to the usual downward closure, we also require that $\diamond root \in cl(\alpha)$ and $\sigma \hookrightarrow_i \phi \in cl(\alpha)$ implies that $\phi, enabled_\sigma \in cl(\alpha)$.

Let AT denote the set of all maximal consistent subsets of $cl(\alpha)$, referred to as atoms. Each $t \in AT$ is a finite set of formulas, we denote the conjunction of all formulas in t by \hat{t} . For a nonempty subset X of AT , we denote by \tilde{X} the disjunction of all \hat{t} , $t \in X$. Define a transition relation on AT as follows: $t \xrightarrow{(\gamma^1, \gamma^2)} t'$ iff $\hat{t} \wedge \langle \gamma^1, \gamma^2 \rangle \hat{t}'$ is consistent. Call an atom t a root atom if there does not exist any atom t' such that $t' \xrightarrow{(\gamma^1, \gamma^2)} t$ for some (γ^1, γ^2) in $\tilde{\Gamma}$. Note that $t_0 = w_0 \cap cl(\alpha_0) \in AT$.

Proposition 4.1 *There exist $t_1, \dots, t_k \in AT$ and $(\gamma_1^1, \gamma_1^2), \dots, (\gamma_k^1, \gamma_k^2) \in \tilde{\Gamma}$ with $(k \geq 1)$ such that $t_k \xrightarrow{(\gamma_k^1, \gamma_k^2)} t_{k-1} \dots \xrightarrow{(\gamma_1^1, \gamma_1^2)} t_0$, where t_k is a root atom.*

Lemma 4.2 *For every $t_1, t_2 \in AT$, the following are equivalent:*

- 1) $\hat{t}_1 \wedge \langle \gamma^1, \gamma^2 \rangle \hat{t}_2$ is consistent.
- 2) $\langle \gamma^1, \gamma^2 \rangle \hat{t}_1 \wedge \hat{t}_2$ is consistent.

Lemma 4.3 *Consider the path $t_k \xrightarrow{(\gamma_k^1, \gamma_k^2)} t_{k-1} \dots \xrightarrow{(\gamma_1^1, \gamma_1^2)} t_0$ where t_k is a root atom. Then:*

- (i) For all $j \in \{0, \dots, k-1\}$, if $[\overline{\gamma^1, \gamma^2}] \alpha \in t_j$ and $t_{j+1} \xrightarrow{(\gamma^1, \gamma^2)} t_j$, then $\alpha \in t_{j+1}$,
- (ii) For every $j \in \{0, \dots, k-1\}$, if $\langle \overline{\gamma^1, \gamma^2} \rangle \alpha \in t_j$ and $t_{j+1} \xrightarrow{(\gamma^1, \gamma^2)} t_j$, then $\langle \overline{\gamma^1, \gamma^2} \rangle \alpha \in t_{j+1}$, and $\alpha \in t_{j+1}$, and
- (iii) For all $j \in \{0, \dots, k-1\}$, if $\diamond \alpha \in t_j$, then there exists i such that $j \leq i \leq k$ and $\alpha \in t_i$.

Thus there exist *mcs*'s w_1, w_2, \dots, w_k , and $(\gamma_1^1, \gamma_1^2), \dots, (\gamma_k^1, \gamma_k^2) \in \tilde{\Gamma}$ such that $w_k \xrightarrow{(\gamma_k^1, \gamma_k^2)} w_{k-1} \dots \xrightarrow{(\gamma_1^1, \gamma_1^2)} w_0$, where $w_j \cap cl(\alpha_0) = t_j$. This path defines a tree $T_0 = (S_0, \Rightarrow_0, s_0)$ rooted at s_0 , where $S_0 = \{s_0, \dots, s_k\}$ and for all $j \in \{0, \dots, k\}$, s_j is labelled by the *mcs* w_{k-j} . The relation is \Rightarrow_0 is defined as usual. From now on we will denote $\alpha \in s$ whenever $\alpha \in w$, w being the *mcs* associated with s , where s is a tree node.

For $k \geq 0$, we can inductively define $T_k = (S_k, \Rightarrow_k, s_0)$ such that the past formulas at every node will have witnesses as ensured by lemma 4.3(iii). Given any T_k , one can extend it to $T_{k+1} = (s_{k+1}, \Rightarrow_{k+1}, s_{k+1})$ by adding new nodes corresponding to formulas like

$\langle \gamma^1, \gamma^2 \rangle \top \in s$, for some $s \in S_k$, for which there exists no $s' \in S_k$ such that $s \xrightarrow{(\gamma_1^1, \gamma_1^2)} s'$. The new node would correspond to the *mcs* w associated with s and some other w' , where $w \xrightarrow{(\gamma^1, \gamma^2)} w'$.

Let $T = (S, \Rightarrow, s_0, X)$ where $S = \bigcup_{k \geq 0} S_k$, $\Rightarrow = \bigcup_{k \geq 0} \Rightarrow_k$, and $X : S \rightarrow 2^{\tilde{\Gamma}}$ is given as follows: $(\gamma^1, \gamma^2) \in X(s)$ iff $enabled_{\gamma^1}, enabled_{\gamma^2} \in w$, where w is the *mcs* associated with s . It follows that $X^i(s) = \{\gamma^i : enabled_{\gamma^i} \in w\}$. We define the model $M = (T, V)$ where $V(s) = w \cap \mathbb{P}$, where w is the *mcs* associated with s . The following lemma proves some important properties which are used in the later proofs.

Lemma 4.4 *For every $s \in S$, we have:*

- (i) if $[\overline{\gamma^1, \gamma^2}] \alpha \in s$ and $s \xrightarrow{(\gamma^1, \gamma^2)} s'$, then $\alpha \in s'$;
- (ii) if $\langle \overline{\gamma^1, \gamma^2} \rangle \alpha \in s$, then there exists s' s.t. $s \xrightarrow{(\gamma^1, \gamma^2)} s'$ and $\alpha \in s'$;
- (iii) if $[\overline{\gamma^1, \gamma^2}] \alpha \in s$ and $s' \xrightarrow{(\gamma^1, \gamma^2)} s$, then $\alpha \in s'$;
- (iv) if $\langle \overline{\gamma^1, \gamma^2} \rangle \alpha \in s$, then there exists s' such that $s' \xrightarrow{(\gamma^1, \gamma^2)} s$ and $\alpha \in s'$;
- (v) if $\exists \alpha \in s$ and $s' \Rightarrow^* s$, then $\alpha \in s'$;
- (vi) if $\diamond \alpha \in s$, then there exists s' such that $s' \Rightarrow^* s$ and $\alpha \in s'$.

The following result takes care of boolean formulas.

Proposition 4.5 *For all $\phi \in \Phi$, for all $s \in S$, $\phi \in s$ iff $\rho_s, s \models \phi$.*

The following lemmas and propositions take care of the strategy constructs.

Lemma 4.6 *For all i , for all $\sigma^i \in Strat^i(\mathbb{P})$, for all $\gamma^i \in \Gamma^i$, for all $s \in S$, $\sigma^i : \gamma^i \in s$ iff $\gamma^i \in \sigma^i(s)$.*

Lemma 4.7 *For all $t \in AT$, $\sigma \hookrightarrow_i \phi \notin t$ implies that there exists a path $\rho_t^k : t = t_1 \xrightarrow{(\gamma_1^1, \gamma_1^2)} t_2 \dots \xrightarrow{(\gamma_{k-1}^1, \gamma_{k-1}^2)} t_k$ which conforms to σ such that one of the following holds:*

- (i) $\phi \notin t_k$
- (ii) $moves(t_k)_i \cap \sigma(t_k) = \emptyset$

Proposition 4.8 *For all $s \in S$, $\sigma \hookrightarrow_i \phi \in s$ iff $M, s \models \sigma \hookrightarrow_i \phi$.*

Finally, with all the required lemmas and propositions in hand, we can prove the following:

Theorem 4.9 *For all $\alpha \in \Theta$, for all $s \in S$, $\alpha \in s$ iff $M, s \models \alpha$.*

Consequently, the logic SL is weakly complete. The completeness proof suggests the following automata theoretic decision procedure. For every strategy specification we can construct an advice automaton with output, and given a formula ϕ , we can construct a tree

automaton whose states are the consistent subsets of subformulas of ϕ . These two automata are run in parallel. Since the number of strategy specifications is linear in the size of ϕ , the size of the tree automaton is doubly exponential in the size of ϕ . Emptiness checking for the automaton is polynomial, thus yielding us a decision procedure that runs in double exponential time. This is a modification of the procedure presented in (Ramanujam and Simon, 2008b).

5 FUTURE WORK

In this work we present a sound and complete axiomatization of a Strategy Logic, *SL* which is used to model strategic reasoning in dynamic games with simultaneous moves, in particular, in infinite repeated normal form games. It would be interesting to see the precise connection with the framework developed in (Ramanujam and Simon, 2008b), in particular, whether the strategy logic for turn-based games can be embedded in the logic proposed here. Dually, we could also consider the game tree as obtained by composition from a collection of “small” game trees constituting subgames and strategies as complete plans on them to ensure local outcomes. We may investigate such ideas in dynamic games with simultaneous moves, based on the work in (Ramanujam and Simon, 2008a; Ghosh, 2008).

REFERENCES

- Ågotnes, T. (2006). Action and knowledge in alternating time temporal logic. *Synthese*, 149(2):377–409.
- Alur, R., Henzinger, T. A., and Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, 49:672–713.
- Benthem, J. v. (2002). Extensive games as process models. *Journal of Logic Language and Information*, 11:289–313.
- Benthem, J. v. (2007). In praise of strategies. In Eijck, J. v. and Verbrugge, R., editors, *Foundations of Social Software*, Studies in Logic, pages 283–317. College Publications.
- Bonanno, G. (2001). Branching time logic, perfect information games and backward induction. *Games and Economic Behavior*, 36(1):57–73.
- Broersen, J. and Herzig, A. (2015). Using STIT theory to talk about strategies. In van Benthem, J., Ghosh, S., and Verbrugge, R., editors, *Models of Strategic Reasoning: Logics, Games and Communities*, FoLLI-LNAI State-of-the-Art Survey, LNCS 8972, pages 137–173. Springer.
- Broersen, J. M. (2009). A stit-logic for extensive form group strategies. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*, pages 484–487, Washington, DC, USA. IEEE Computer Society.
- Broersen, J. M. (2010). CTL.STIT: enhancing ATL to express important multi-agent system verification properties. In *Proceedings of the ninth international joint conference on Autonomous agents and multiagent systems (AAMAS 2010)*, Toronto, Canada, pages 683–690, New York, NY, USA. ACM.
- Bulling, N., Goranko, V., and Jamroga, W. (2015). Logics for reasoning about strategic abilities in multi-player games. In van Benthem, J., Ghosh, S., and Verbrugge, R., editors, *Models of Strategic Reasoning: Logics, Games and Communities*, FoLLI-LNAI State-of-the-Art Survey, LNCS 8972, pages 93–136. Springer.
- Chatterjee, K., Henzinger, T. A., and Piterman, N. (2007). Strategy logic. In *Proceedings of the 18th International Conference on Concurrency Theory*, volume 4703 of *Lecture Notes in Computer Science*, pages 59–73. Springer.
- Ghosh, S. (2008). Strategies made explicit in dynamic game logic. In *Proceedings of the Workshop on Logic and Intelligent Interaction, ESSLLI 2008*, pages 74–81.
- Ghosh, S. and Ramanujam, R. (2012). Strategies in games: A logic-automata study. In Bezanishvili, N. and Goranko, V., editors, *Lectures on Logic and Computation - ESSLLI 2010, ESSLLI 2011, Sel. Lecture Notes*, pages 110–159. Springer.
- Harrenstein, P., van der Hoek, W., Meyer, J., and Witteven, C. (2003). A modal characterization of Nash equilibrium. *Fundamenta Informaticae*, 57(2-4):281–321.
- Horty, J. (2001). *Agency and Deontic Logic*. Oxford University Press.
- Osborne, M. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press, Cambridge, MA.
- Ramanujam, R. and Simon, S. (2008a). Dynamic logic on games with structured strategies. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR-08)*, pages 49–58. AAAI Press.
- Ramanujam, R. and Simon, S. (2008b). A logical structure for strategies. In *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, volume 3 of *Texts in Logic and Games*, pages 183–208. Amsterdam University Press.
- Rasmusen, E. (2007). *Games and Information*. Blackwell Publishing, Oxford, UK, 4 edition.
- 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*, pages 157–164.
- Walther, D., van der Hoek, W., and Wooldridge, M. (2007). Alternating-time temporal logic with explicit strategies. In *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2007)*, pages 269–278.
- Zermelo, E. (1913). Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In *Proceedings of the Fifth Congress Mathematicians*, pages 501–504. Cambridge University Press.
- Zhang, D. and Thielscher, M. (2015). Representing and reasoning about game strategies. *Journal of Philosophical Logic*, 44(2):203–236.