

Representing and Reasoning about Game Strategies

Dongmo Zhang¹ Michael Thielscher²

¹ University of Western Sydney, Australia
dongmo@scm.uws.edu.au

² University of New South Wales, Australia
mit@cse.unsw.edu.au

Technical Report
UNSW-CSE-TR-201405
February 2014

THE UNIVERSITY OF
NEW SOUTH WALES



School of Computer Science and Engineering
The University of New South Wales
Sydney 2052, Australia

Abstract

As a contribution to the challenge of building game-playing AI systems, we develop and analyse a formal language for representing and reasoning about strategies. Our logical language builds on the existing general Game Description Language (GDL) and extends it by a standard modality for linear time along with two dual connectives to express preferences when combining strategies. The semantics of the language is provided by a standard state-transition model. As such, problems that require reasoning about games can be solved by the standard methods for reasoning about actions and change. We also endow the language with a specific semantics by which strategy formulas are understood as move recommendations for a player. To illustrate how our formalism supports automated reasoning about strategies, we demonstrate two example methods of implementation: first, we formalise the semantic interpretation of our language in conjunction with game rules and strategy rules in the Situation Calculus; second, we show how the reasoning problem can be solved with Answer Set Programming.

1 Introduction

Strategic reasoning has been a major research theme in game theory. However, “much of game theory is about the question whether strategic equilibria exist”, as Johan van Benthem points out, “*but there are hardly any explicit languages for defining, comparing, or combining strategies*” [23]. The intrinsic difficulty of modelling strategic reasoning is that reasoning about strategies is not purely deductive but combines temporal reasoning, counterfactual reasoning, reasoning about actions and preferences, and multi-agent interaction. If any logic is used, such a logic must be able to support these reasoning mechanisms.

In recent years a number of specific logical formalisms have been proposed for specifying strategic behaviour of agents in multi-agent systems [14, 25, 15, 4, 12]. These frameworks offer modelling facilities and inference mechanisms for specifying strategy reasoning in multi-agent systems. At the same time, the conceptually simple, general *Game Description Language* (GDL) has been developed as a practical language for encoding the rules of arbitrary games so that they can be understood by general game-playing systems, whose task is to learn to play unknown games without human intervention [7, 21].

In this report, we explore the middle ground between pure game specification languages like GDL on the one hand, and existing expressive formalisms for strategic reasoning on the other hand. Our main contributions can be summarised as follows:

1. We show how a simple extension of GDL using a standard modality for linear time suffices to describe strategies in addition to the mere rules of a game, and we present a specific semantics by which formulas in this language can be understood as move recommendations for a player.
2. We enrich our language by two *preference* operators, respectively called prioritised disjunction and prioritised conjunction, and show how to use them to describe strategies that are complete (i.e., provide a move recommendation in every state) and deterministic (i.e., move recommendations are always unique).
3. We demonstrate two example methods of implementation: a formalisation of the semantic interpretation of our language in the Situation Calculus [17] and a translation of a subset of our language into Answer Set Programming [6].

These results are accompanied by a thorough mathematical analysis of the language and its semantics, in particular with regard to our novel preference operators. The advantage of our framework is to allow for concise but complete representations of games and strategic behaviour of agents so that the approach promises to be practically useful for the design of game-playing agents. We will use a simple game scenario as a running example to explain the basic concepts and to demonstrate how our language can be used to write strategies for game-playing agents. Our results lay the foundations for a variety of applications of automated reasoning about strategies in AI systems, including the following:

- A general game player can be fed not only with the mere rules of a new game but also with declarative descriptions of tailor-made strategies.

- General game players can use a high-level, declarative representation of their strategies as the basis for learning, maintaining, and reasoning about them.
- Players can also use the language and its inference mechanism to represent, revise and reason with their beliefs about their opponents' strategies.

The remainder of the report proceeds as follows. We will begin our technical exposition with the definition of a simple state transition semantics for games. In Section 3, we will then define the semantic concept of a strategy within this framework, and in Section 4 we will develop a formal game specification language based on GDL for describing both game rules and strategies in a concise manner. In Section 5, we will introduce the two preference connectives and show how to use these to combine strategies into more complex ones. In Section 6, we demonstrate how to verify if a strategy could bring out an expected outcome with the model-based approach. In Section 7, we will encode the semantic interpretation of our language in conjunction with game rules and strategy rules in the Situation Calculus and also illustrate how the reasoning problem can be solved with Answer Set Programming. We conclude with a discussion of related work.

2 State Transition Games

To set the stage for our work, we first define formally a general state transition model for games.

Definition 1 A state transition game G is a tuple $(N, W, \mathcal{A}, \bar{w}, t, l, u, g)$, where

1. N is a non-empty finite set of players;
2. W is a non-empty set of states (possible worlds);
3. $\mathcal{A} = \bigcup_{i \in N} A^i$, where A^i is a non-empty finite set of actions for player $i \in N$;
4. $\bar{w} \in W$, representing the initial state;
5. $l \subseteq W \times \mathcal{A}$ is a binary legality relation, describing what actions are allowed in which states;
6. $u : \mathcal{A} \times W \mapsto W$ is an update function, specifying the state transitions;
7. $g : N \mapsto 2^W$ is a goal function, specifying the winning states of each agent;
8. $t \subseteq W$, representing the terminal states.

To keep our formalism as simple as possible, we assume that all actions are performed asynchronously, each by a single player (although games need not be turn-taking). We also assume that different agents have different actions, that is, $A^i \cap A^j = \emptyset$ for any $i \neq j$, but of course two actions may have the same effects.

A sequence $w_0 \xrightarrow{a_0} w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{m-1}} w_m$ is called a *complete path* if

1. $w_0 = \bar{w}$, $w_m \in t$, and $w_j \in W$ for all $0 < j < m$;

2. $a_j \in \mathcal{A}$ for all $0 \leq j < m$;
3. $(w_j, a_j) \in l$ for all $0 \leq j < m$; and
4. $u(a_j, w_j) = w_{j+1}$ for all $0 \leq j < m$.

Any segment $w_k \xrightarrow{a_k} w_{k+1} \xrightarrow{a_{k+1}} \dots \xrightarrow{a_{l-1}} w_l$ of a complete path is called a *reachable path*. We let $\mathcal{P}(G)$ denote all reachable paths in a state transition game G . Note that a single state without action can be a reachable path (i.e., where $k = l$). We call such a singleton path $w \in W \cap \mathcal{P}(G)$ a *reachable state*.

A state-action pair (w, a) is called a *reachable legal move* (or simply a *move*) in G if there is a w' such that $w \xrightarrow{a} w'$ is a reachable path. The set of all such moves in G is denoted by $\Omega(G)$. More formally,

$$\Omega(G) = \{(w, a) : \text{there is } w' \in W \text{ such that } w \xrightarrow{a} w' \in \mathcal{P}(G)\} \quad (2.1)$$

Furthermore, the moves for player i are denoted by $\Omega^i(G)$:

$$\Omega^i(G) = \{(w, a) \in \Omega(G) : a \in A^i\} \quad (2.2)$$

For convenience, we let

$$l^i(w) = \{a : (w, a) \in \Omega^i(G)\} \quad (2.3)$$

To facilitate the presentation of our framework, we will use as our running example a simple two-player game, which we call *CrossDot* game.

Example 1 (CrossDot Game) Two players take turns in placing either a cross “×” (player 1) or a dot “.” (player 2) into an empty box in a line of $m \geq 2$:

$$\underbrace{\square \square \square \dots \square}_m$$

Each box can contain at most one object. The first player to successfully fill k ($1 < k \leq m$) consecutive boxes will end and win the game, where k is arbitrary but fixed. If all boxes have been filled without a winner, the game ends with a tie. We assume that the “×”-player goes first.

To describe the scenario in terms of the state transition game, let $N = \{1, 2\}$ be the players and

$$W = \{(t_1, t_2, x_1, x_2, \dots, x_m) : t_1, t_2 \in \{0, 1\} \ \& \ x_1, \dots, x_m \in \{\square, \boxtimes, \square\}\}$$

the set of possible states, where t_1, t_2 specify whose turn it is ($t_i = 1$ means player i 's turn) and x_1, \dots, x_m indicate the status of the boxes. The initial state is $\bar{w} = (1, 0, \square, \square, \dots, \square)$.

We write a_j^i to denote the action of player i marking the j th box. Let $A^i = \{a_j^i : i \in N \ \& \ 1 \leq j \leq m\}$. We refrain from explicitly listing the legality relation, the update function, and the terminal and goal states for the players as this is possible but considerably lengthy even for a very simple game like this; the syntactic *axiomatisation* of this game given in the following section will be much more concise and practical. Let us just pick some random examples:

For $m = 4$ and $k = 2$, the state $(1, 0, \square, \boxtimes, \boxtimes, \square)$ is a terminal state and also a goal state for player 1. Under our assumption that player 1 takes the first move, the state-action pair $((1, 0, \square, \square, \boxtimes, \square), a_2^1)$, say, satisfies the legality relation, whereas $((1, 0, \square, \square, \boxtimes, \square), a_2^2)$ does not because it is not player 2's turn at this stage. As an example of the update function we have $u(a_2^1, (1, 0, \square, \square, \boxtimes, \square)) = (0, 1, \square, \boxtimes, \boxtimes, \square)$.

For future reference, by $G_{CrossDot}^{m,k}$ we denote the instance of our game that consists of m boxes and has winning length k . ■

3 Strategies in State Transition Games

In a multi-agent game environment, agents strive to achieve their goals. How individual agents act is determined by their *strategies*: At a state of a game, the strategies that the agent applies determine which actions the agent will take. For example, a strategy of a player in the game of chess reflects which moves the player would play in certain positions that can be reached according to the standard chess rules¹. In terms of the state transition game, a strategy of a player can be defined as a relation between game states and (legal) actions the player will perform.

Definition 2 A strategy S of player i is a subset of $W \times A^i$ such that $S \subseteq \Omega^i(G)$.

Intuitively, a strategy of a player specifies which actions the player should take in which states. Note that our concept of strategy is significantly different from the ones in the context of alternating-time temporal logic (ATL) [1, 25, 26]. In ATL, a strategy is a function that maps each state (or a sequence of states) to an action. In other words, a strategy specifies which action has to do exactly in each state or each history of states. However, our concept of strategy can express a “rough idea” of what to do. A strategy may suggest no action, one action or several actions to do in each state. In the later sections, we will further develop a technology to combine and refine strategies in order to generate a strategy with desirable properties.

3.1 Properties of Strategies

We say that a strategy S is *valid* if $S \neq \emptyset$. A strategy S of player i is *complete* if for each reachable state $w \in W \cap \mathcal{P}(G)$, there is a move $(w, a^i) \in S$ unless $l^i(w) = \emptyset$. In other words, a complete strategy provides the player with a “complete” guideline that always provides the player with one or more suggestions how to act when it is his move. A strategy S is *deterministic* if for any $(w, a) \in S$ and $(w, a') \in S$, we have $a = a'$. A strategy is *functional* if it is complete and deterministic. An agent with a functional strategy knows precisely what to do in any reachable game state. For instance, if an agent $i \in N$ has a default action $a_0^i \in A^i$ that is always legal in any reachable state, then the following simple strategy for this player is complete and deterministic, hence functional:

$$S^i = \{(w, a_0^i) : w \in W \cap \mathcal{P}(G)\}$$

¹A player might apply different strategies in different situations of a game. A single strategy does not necessarily determine the moves for all possible legal positions of a game.

Example 2 Consider an instance of the CrossDot game $G_{CrossDot}^{4,2}$. The following is an example of strategies for player 1 that intuitively says “fill a box next to one you have marked before”:

$$S^1 = \{ ((1, 0, \boxtimes, \square, \square, \square), a_2^1), ((1, 0, \boxtimes, \square, \square, \square), a_2^1), \\ ((1, 0, \square, \boxtimes, \square, \square), a_1^1), ((1, 0, \square, \boxtimes, \square, \square), a_1^1), \\ ((1, 0, \square, \square, \boxtimes, \square), a_3^1), ((1, 0, \square, \square, \boxtimes, \square), a_2^1), \\ ((1, 0, \square, \square, \square, \boxtimes), a_3^1), ((1, 0, \square, \square, \square, \boxtimes), a_2^1), \\ ((1, 0, \square, \square, \square, \square), a_4^1), ((1, 0, \square, \square, \square, \square), a_3^1), \\ ((1, 0, \square, \square, \square, \square), a_4^1), ((1, 0, \square, \square, \square, \square), a_3^1) \}$$

It is easy to see that the strategy is valid but neither complete nor deterministic.

■

4 Strategy Representation and Semantics

The above example shows that directly representing a strategy in a state transition game requires to list every move that complies with the strategy. In the following, we develop a syntactical representation that allows to describe a strategy much more concisely.

4.1 Formal Game Specification Language

We first present a logic-based, general game description language with linear time.

Definition 3 Consider a propositional modal language \mathcal{L} with these components:

- a finite set Φ of propositional variables;
- a finite set N of agent symbols;
- a finite set A^i of action symbols for each $i \in N$;
- propositional connectives $\neg, \wedge, \vee, \supset$ and \equiv ;¹
- pseudo-function symbols $does(\cdot), legal(\cdot),$ and $wins(\cdot)$;
- modal operator \bigcirc ;
- special propositional symbols $init$ and $terminal$.

Formulas in \mathcal{L} are defined as follows:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \bigcirc\varphi \mid does(a) \mid legal(a) \mid wins(i) \mid init \mid terminal$$

where $p \in \Phi$, $i \in N$ and $a \in \mathcal{A} = \bigcup_{i \in N} A^i$.

¹Only \neg and \wedge are treated as primitives.

Note that we overload N and A^i as they occur in both syntax and semantics. They can be distinguished from the context. As in Definition 1, we assume that $A_i \cap A_j = \emptyset$ if $i \neq j$.

We call *does*, *legal* and *wins* pseudo-functions because formally each instance $does(a)$, $legal(a)$ or $wins(i)$ is taken to be an individual propositional symbol. This language is a direct adaptation of the general *Game Description Language* [7] and allows to describe games in a compact way.

EXAMPLE 1 (*continued*)² We use the propositional symbols p_j^i to represent the fact that box j is filled with i 's marker, where $i \in \{1, 2\}$ and $j \in \{1, \dots, m\}$. In addition, we use two specific propositional symbols $turn(1)$, $turn(2)$ to represent players' turns. Putting all the propositional symbols together, we have $\Phi_{CrossDot} = \{p_j^i : 1 \leq i \leq 2 \ \& \ 1 \leq j \leq m\} \cup \{turn(i) : i = 1 \text{ or } 2\}$. With this, we are able to describe the game rules in our logical language.

To begin with, the following rules specify the initial game state:

$$init \supset \neg p_j^i \quad \text{for all } i \in N \text{ and } j \leq m \quad (4.1)$$

$$init \supset turn(1) \wedge \neg turn(2) \quad (4.2)$$

The following statement defines the winning conditions: *Player i wins if there is j such that $j + k - 1 \leq m$ and $p_j^i \wedge \dots \wedge p_{j+k-1}^i$* , i.e.

$$wins(i) \equiv \bigvee_{j=1}^{m-k+1} \bigwedge_{l=j}^{j+k-1} p_l^i \quad (4.3)$$

With this, the condition for termination is:

$$terminal \equiv wins(1) \vee wins(2) \vee \bigwedge_{1 \leq j \leq m} (p_j^1 \vee p_j^2) \quad (4.4)$$

As before, let a_j^i denote the action of player i filling box j . Legality of the actions of each player i can be described thus:

$$\neg(p_j^1 \vee p_j^2) \wedge turn(i) \wedge \neg terminal \equiv legal(a_j^i) \quad (4.5)$$

The effects of the actions are given by³

$$p_j^i \vee does(a_j^i) \equiv \bigcirc p_j^i \quad (4.6)$$

$$turn(1) \supset \bigcirc \neg turn(1) \wedge \bigcirc turn(2) \quad (4.7)$$

$$turn(2) \supset \bigcirc \neg turn(2) \wedge \bigcirc turn(1) \quad (4.8)$$

Let $\Sigma_{CrossDot}^{m,k}$ be the set of axioms (4.1)–(4.8). ■

In the following, we interpret the language using our state transition model of Definition 1. A *valuation function* $v : W \mapsto 2^\Phi$ specifies which propositions are true at each state. Propositional formulas and their truth values can be defined accordingly.

²While all examples in this report are based on the *CrossDot* game, some examples are specifically numbered just for the purpose of cross-referencing.

³To avoid too much complexity, we follow [7] in not offering any solution to the frame problem here. Game descriptions simply include all necessary frame axioms.

Definition 4 Let $G = (N, W, \mathcal{A}, \bar{w}, t, l, u, g)$ be a state transition game and v a valuation function. We call the pair $M = (G, v)$ a state transition model. Let $\delta = w_0 \xrightarrow{a_0} w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{m-1}} w_m \in \mathcal{P}(G)$ be a reachable path in G and φ a formula. We say that δ satisfies φ under M (written $M, \delta \models \varphi$) according to the following definition:

$M, \delta \models p$	<i>iff</i> $p \in v(w_0)$ ($p \in \Phi$)
$M, \delta \models \text{does}(a)$	<i>iff</i> $a = a_0$
$M, \delta \models \text{init}$	<i>iff</i> $w_0 = \bar{w}$
$M, \delta \models \text{terminal}$	<i>iff</i> $w_0 \in t$
$M, \delta \models \text{legal}(a)$	<i>iff</i> $(w_0, a) \in l$
$M, \delta \models \text{wins}(i)$	<i>iff</i> $w_0 \in g(i)$
$M, \delta \models \neg\varphi$	<i>iff</i> $M, \delta \not\models \varphi$
$M, \delta \models \varphi_1 \wedge \varphi_2$	<i>iff</i> $M, \delta \models \varphi_1$ and $M, \delta \models \varphi_2$
$M, \delta \models \bigcirc\varphi$	<i>iff</i> $M, w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{m-1}} w_m \models \varphi$

It is worth clarifying that in the limit case $m = 0$ (i.e., $\delta = w_0$) we have that $M, \delta \models \text{does}(a)$ and $M, \delta \models \bigcirc\varphi$ hold for any a and φ .

A formula φ is *valid* in model M , denoted $M \models \varphi$, if it is satisfied by any reachable path in the game, that is, $M, \delta \models \varphi$ for all $\delta \in \mathcal{P}(G)$. Let Σ be a set of sentences in \mathcal{L} , then M is a *model* of Σ if $M \models \varphi$ for all $\varphi \in \Sigma$.

Observation 1 Consider the CrossDot game $G = G_{\text{CrossDot}}^{m,k}$ introduced in Example 1. Let v be a valuation function such that for each state $w = (t_1, t_2, x_1, \dots, x_m) \in W$, $v(w) = \{\text{turn}(i) : t_i = 1\} \cup \{p_j^1 : x_j = \boxtimes \ \& \ 1 \leq j \leq m\} \cup \{p_j^2 : x_j = \boxminus \ \& \ 1 \leq j \leq m\}$. Let $M = (G, v)$. Then M is a model of $\Sigma_{\text{CrossDot}}^{m,k}$ (see equations (4.1)–(4.8)).

Proof: Given any reachable path $\delta = w_0 \xrightarrow{a_0} w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{e-1}} w_e$, we only have to verify that each axiom in $\Sigma_{\text{CrossDot}}^{m,k}$ is satisfied by δ in M . This is straightforward. Consider for instance the proof that $M, \delta \models \text{init} \supset \neg p_j^i$ for any $i \in N$ and $j < m$: If $w_0 \neq \bar{w}$, this holds trivially because $M, \delta \models \neg \text{init}$. If $w_0 = \bar{w}$, then $M, \delta \models \text{init}$ and by the definition of v , $p_j^i \notin v(w_0)$, hence $M, \delta \models \neg p_j^i$ for any $i \in N$ and $j < m$. The other axioms can be verified similarly. ■

In order to develop a syntactical representation for strategies, we introduce the following specific concepts. Let $\delta = w_0 \xrightarrow{a_0} w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{m-1}} w_m$. We call δ a *reachable path starting with the move* (w_0, a_0) . The set of all the reachable paths in game G that starts with (w, a) is denoted by $(w, a)^\sim$. Given a state transition model $M = (G, v)$, for any move $(w, a) \in \Omega(G)$, a formula φ is *valid under move* (w, a) , denoted by $M \models_{(w,a)} \varphi$, if $M, \delta \models \varphi$ for all reachable path $\delta \in (w, a)^\sim$.

Example 3 Consider the same instance of the CrossDot game $G_{\text{CrossDot}}^{4,2}$ as in Example 2. Let M be the state transition model defined in Observation 1 for the case $m = 4$ and $k = 2$. Assume that $(w, a) = ((1, 0, \boxtimes, \boxminus, \square, \square), a_3^1)$. It is easy to verify the following:

$$\begin{aligned} M \models_{(w,a)} \text{legal}(a_4^1) \\ M \not\models_{(w,a)} \text{does}(a_3^2) \end{aligned}$$

$$\begin{aligned}
M &\models_{(w,a)} \neg(p_3^1 \vee p_3^2) \wedge \bigcirc p_3^1 \wedge \bigcirc(\text{legal}(a_4^2) \wedge \text{does}(a_4^2)) \\
M &\models_{(w,a)} \bigcirc \bigcirc \text{terminal}
\end{aligned}$$

■

4.2 Describing Strategies

We now turn to the syntactical representation of strategies using the language introduced above. For any state transition model $M = (G, v)$ and formula $\varphi \in \mathcal{L}$, let

$$S^i(\varphi) = \{(w, a) \in \Omega^i(G) : M \models_{(w,a)} \varphi\} \quad (4.9)$$

In other words, $S^i(\varphi)$ comprises all moves for player i under which φ is valid.

Definition 5 *Given a state transition model M , let S be a strategy of player i according to Definition 2. A formula φ in \mathcal{L} is a representation of S iff $S = S^i(\varphi)$.*

In the following, we will call “strategy” both a set of moves (i.e., a strategy in the sense of Definition 2) and its representation (using \mathcal{L}). They should be easy to distinguish from the context. Obviously, for any formula $\varphi \in \mathcal{L}$, $S^i(\varphi)$ is a strategy for player i . For instance, a tautology \top is a strategy that allows to take any reachable legal move. Another example is $\text{does}(a^i)$, where $a^i \in A^i$, which means —if read as strategy for player i — to take a^i only at any reachable state.

While the language we consider in this report is propositional and finite, a state transition game can have infinitely many worlds which cannot be distinguished by the propositions that hold in them. A strategy that assigns different actions to two indistinguishable worlds cannot be described in our language. However, every strategy that is *Markovian* in the following sense does have a representation in \mathcal{L} .

Definition 6 *Given a state transition model $M = (G, v)$, a strategy S is Markovian if for all $(w_1, a) \in S$ and $w_2 \in \mathcal{P}(G)$ such that $v(w_1) = v(w_2)$, $(w_2, a) \in S$.*

Proposition 1 *Given a state transition model $M = (G, v)$ of \mathcal{L} , any Markovian strategy S of a player has a representation in \mathcal{L} .*

Proof: For each $(w, a) \in S$, since the set Φ of propositional variables is finite, $v(w)$ and $\Phi \setminus v(w)$ are both finite. Thus the following is a well-formed propositional formula:

$$\varphi(w, a) = \left(\bigwedge_{p \in v(w)} p \right) \wedge \left(\bigwedge_{p \in \Phi \setminus v(w)} \neg p \right) \wedge \text{does}(a) \quad (4.10)$$

Again \mathcal{A} being finite implies the set $\{\varphi(w, a) : (w, a) \in S\}$ to be finite even though S may be infinite. Therefore the following is also a well-formed propositional formula:

$$\varphi = \bigvee_{(w,a) \in S} \varphi(w, a)$$

Now we show φ is a representation of S . Obviously, for any $(w, a) \in S$, $M \models_{(w,a)} \varphi$. For any $(w, a) \in \Omega^i(G)$, where i is the player under consideration, assume

that $M \models_{(w,a)} \varphi$. By the construction of φ , there is a move $(w', a') \in S$ such that

$$M \models_{(w,a)} \left(\bigwedge_{p \in v(w')} p \right) \wedge \left(\bigwedge_{p \in \Phi \setminus v(w')} \neg p \right) \wedge \text{does}(a')$$

It turns out that $v(w) = v(w')$. Since S is Markovian, we have $(w, a') \in S$. Note that $M \models_{(w,a)} \text{does}(a')$ implies $a = a'$ by Definition 4. We then have $(w, a) \in S$, as desired. ■

Note that Markovian strategies are history-independent. If we want to specify strategies in which move choices depend on the history of a game, then we need to extend our language with syntactic means to talk about past states. This can be done by adding the inverse of the operator \bigcirc , written \bigcirc^{-1} , to denote that some property holds in “the previous state”. We leave this extension for future work.

The reader is reminded that formulas in our language have been endowed with two different semantics. If it is used to represent a property, it has a truth value as normal propositional formula. If a formula is used to represent a strategy, then it no longer has a truth value but represents a set of moves for a player. Moreover, the same formula represents different strategies for different players. For instance, if we use the constant proposition $\top (=_{def} p \vee \neg p)$ to represent a strategy for a player, it means all moves that are available to that player at each state. Obviously different players have different available moves.

The following observation shows how our language can be used to describe a useful strategy for our running example in a compact and concise manner.

Observation 2 *Strategy S^1 from Example 2 is represented by*

$$\varphi = \bigvee_{1 < j \leq 4} (p_{j-1}^1 \wedge \neg p_j^1 \wedge \neg p_j^2 \wedge \text{does}(a_j^1)) \vee \bigvee_{1 \leq j < 4} (\neg p_j^1 \wedge \neg p_j^2 \wedge p_{j+1}^1 \wedge \text{does}(a_j^1))$$

Proof: First we show $S^1 \subseteq S^1(\varphi)$. This can be done by verifying $M \models_{(w,a)} \varphi$ one by one for each $(w, a) \in S^1$. For instance, consider $w = (1, 0, \boxtimes, \square, \square, \square,)$ and $a = a_2^1$. Then $M \models_{(w,a)} p_1^1 \wedge \neg p_2^1 \wedge \neg p_2^2$ and $M \models_{(w,a)} \text{does}(a_2^1)$. Thus $M \models_{(w,a)} \varphi$.

In order to show that $S^1(\varphi) \subseteq S^1$, we can prove that for each reachable legal move (w, a) of player 1, $(w, a) \notin S^1$ implies $M \not\models_{(w,a)} \varphi$. This can be done by enumerating all reachable legal moves in $\Omega^1(G_{CrossDot}^{4,2}) \setminus S^1$. For instance, let $w_0 = (1, 0, \boxtimes, \square, \square, \square)$ and $a_0 = a_3^1$. Obviously, (w_0, a_0) is a reachable legal move and $(w_0, a_0) \notin S^1$. To show $M \not\models_{(w_0, a_0)} \varphi$, consider the reachable path $\delta = w_0 \xrightarrow{a_0} w_1$, where $w_1 = (0, 1, \boxtimes, \square, \boxtimes, \square)$. It is easy to verify that $M, \delta \models \neg p_2^1 \wedge \neg p_4^1 \wedge \text{does}(a_3^1)$ while $M, \delta \models \neg \text{does}(a')$ for any $a' \neq a_3^1$. Thus we have $M, \delta \not\models \varphi$. We then yield that $M \not\models_{(w_0, a_0)} \varphi$. Other cases can be verified similarly. ■

The above observation shows that we can still use our logical sense to design a strategy despite the significant differences of semantics between a propositional formula and a strategy representation. When doing so it is important to keep it in mind that a propositional formula bears a very different meaning when understood as a strategy rather than a state description. In the following we

demonstrate the speciality of strategy representation with a few examples based on the CrossDot game $G = G_{CrossDot}^{4,2}$.

1. *One formula can represent different strategies for different players:*

$$\begin{aligned}
S^1(\bigcirc p_1^1) &= \{(1, 0, \square, x_2, x_3, x_4), a_1^1\} \in \mathcal{P}(G) : x_2, x_3, x_4 \in \{\square, \boxtimes, \square\}\} \\
&\cup \{(1, 0, \boxtimes, \square, x_3, x_4), a_2^1\} \in \mathcal{P}(G) : x_3, x_4 \in \{\square, \boxtimes, \square\}\} \\
&\cup \{(1, 0, \boxtimes, x_2, \square, x_4), a_3^1\} \in \mathcal{P}(G) : x_2, x_4 \in \{\square, \boxtimes, \square\}\} \\
&\cup \{(1, 0, \boxtimes, x_2, x_3, \square), a_4^1\} \in \mathcal{P}(G) : x_2, x_3 \in \{\square, \boxtimes, \square\}\} \\
S^2(\bigcirc p_1^1) &= \{(0, 1, \boxtimes, \square, x_3, x_4), a_2^2\} \in \mathcal{P}(G) : x_3, x_4 \in \{\square, \boxtimes, \square\}\} \\
&\cup \{(0, 1, \boxtimes, x_2, \square, x_4), a_3^2\} \in \mathcal{P}(G) : x_2, x_4 \in \{\square, \boxtimes, \square\}\} \\
&\cup \{(0, 1, \boxtimes, x_2, x_3, \square), a_4^2\} \in \mathcal{P}(G) : x_2, x_3 \in \{\square, \boxtimes, \square\}\}
\end{aligned}$$

In other words, if $\bigcirc p_1^1$ represents a strategy of player 1, it means that player 1 is to fill the first box with his marker if the box is currently empty or to fill any other empty box if the box has already been filled with a cross whenever it is his turn. However, if the formula $\bigcirc p_1^1$ represents a strategy of player 2, it means player 2 wishes the first box to be filled by player 1. In this case, he waits for the states to come and then do whatever is feasible to him.

2. *Forcing another player into a particular action:*

$$\begin{aligned}
S^1(\bigcirc \text{does}(a_1^2)) &= \{(1, 0, \square, \square, x_3, x_4), a_2^1\} \in \mathcal{P}(G) : x_3, x_4 \in \{\boxtimes, \square\}\} \\
&\cup \{(1, 0, \square, x_2, \square, x_4), a_3^1\} \in \mathcal{P}(G) : x_2, x_4 \in \{\boxtimes, \square\}\} \\
&\cup \{(1, 0, \square, x_2, x_3, \square), a_4^1\} \in \mathcal{P}(G) : x_2, x_3 \in \{\boxtimes, \square\}\}
\end{aligned}$$

In general, a player does not have control on the other player's actions. But a player may be able to create a situation in which the opponent has no choice other than the desired action. In the above example, player 1 is enforcing a situation in which player 2 has no other option but to perform a_1^2 when it becomes his turn.

3. *Thinking forward:*

$$\begin{aligned}
S^2((\bigcirc(\neg \text{does}(a_1^1)) \supset \bigcirc \bigcirc \text{does}(a_1^2)) &= \{(0, 1, \square, \square, \square, \boxtimes), a_2^2\}, (0, 1, \square, \square, \square, \boxtimes), a_3^2\}, \\
&\{(0, 1, \square, \square, \boxtimes, \square), a_2^2\}, (0, 1, \square, \square, \boxtimes, \square), a_4^2\}, \\
&\{(0, 1, \square, \boxtimes, \square, \square), a_3^2\}, (0, 1, \square, \boxtimes, \square, \square), a_4^2\}
\end{aligned}$$

This is a strategy for player 2 to try to find an action so that “he can fill box 1 in his next turn as long as player 1 will not fill it beforehand.” Note that $S^2(\bigcirc \bigcirc \text{does}(a_1^2)) = \emptyset$ because he can do nothing to guarantee that $\text{does}(a_1^2)$ is doable in his next turn.

4. *Try anything to win the game within one step:*

$$\begin{aligned}
S^1(\bigcirc \text{wins}(1)) &= \{(1, 0, \boxtimes, \square, \square, \square), a_2^1\}, (1, 0, \boxtimes, \square, \square, \square), a_2^1\}, \\
&\{(1, 0, \square, \boxtimes, \square, \square), a_1^1\}, (1, 0, \square, \boxtimes, \square, \square), a_1^1\}, \\
&\{(1, 0, \square, \boxtimes, \square, \square), a_3^1\}, (1, 0, \square, \square, \boxtimes, \square), a_2^1\}, \\
&\{(1, 0, \square, \square, \boxtimes, \square), a_2^1\}, (1, 0, \square, \square, \square, \boxtimes), a_3^1\}, \\
&\{(1, 0, \square, \square, \square, \boxtimes), a_3^1\}
\end{aligned}$$

This is a simple strategy for player 1 that aims to try any available action to win the game in one step.

We have seen from the above examples that even though a strategy rule is written in the syntax of (temporal) propositional formulas, its semantics is significantly different from propositional logic therefore we must be very cautious when we start to use the language to design game strategies.

To complete this section, we present a property of our strategy representation for later use.

Lemma 1 *Given a state transition model M , for each player i ,*

1. $S^i(\varphi_1 \wedge \varphi_2) = S^i(\varphi_1) \cap S^i(\varphi_2)$
2. $S^i(\varphi_1) \cup S^i(\varphi_2) \subseteq S^i(\varphi_1 \vee \varphi_2)$

Proof: To prove (1), assume that $(w, a) \in S^i(\varphi_1 \wedge \varphi_2)$. It follows that $a \in A^i$ and $M \models_{(w,a)} \varphi_1 \wedge \varphi_2$. Then for each reachable path $\delta \in (w, a)^{\rightsquigarrow}$, we have $M, \delta \models \varphi_1 \wedge \varphi_2$, which implies $M, \delta \models \varphi_1$ and $M, \delta \models \varphi_2$. It turns out that $M \models_{(w,a)} \varphi_1$ and $M \models_{(w,a)} \varphi_2$. We get $(w, a) \in S^i(\varphi_1)$ and $(w, a) \in S^i(\varphi_2)$. The other direction is similar.

To prove (2), assume that $(w, a) \in S^i(\varphi_1) \cup S^i(\varphi_2)$. Without loss of generality, suppose that $(w, a) \in S^i(\varphi_1)$, which implies $a \in A^i$ and $M \models_{(w,a)} \varphi_1$. Hence, $M, \delta \models \varphi_1$ for each reachable path $\delta \in (w, a)^{\rightsquigarrow}$. It follows that $M, \delta \models \varphi_1 \vee \varphi_2$. Thus we yield $M \models_{(w,a)} \varphi_1 \vee \varphi_2$. Given that $a \in A^i$, we conclude $(w, a) \in S^i(\varphi_1 \vee \varphi_2)$. ■

5 Strategy Composition

As mentioned in the introduction, the main motivation of this work is to introduce a formal logical language for defining, comparing and combining strategies. We have provided our formal definition of strategies in both syntactical and semantical levels in the previous sections. To facilitate the composition of strategies, in this section we extend our language by two specific connectives, called *prioritised disjunction* and *prioritised conjunction*, respectively.

5.1 Prioritised Disjunction and Conjunction

The idea behind these two new connectives is the following. The prioritised disjunction “ ∇ ” extends the choice of actions such that if a first strategy fails to apply then a second one offers more options, and if that fails too then a third strategy may offer more options still, and so on. Conversely, the prioritised conjunction “ Δ ” narrows down the choice of actions: if a first strategy allows too many options, then a second strategy may be used to constrain these options, a third strategy may narrow down the options even further, and so on—up to the point where the next strategy in line would lead to empty option.

Definition 7 *The set of strategy rules is the smallest set such that*

1. *a formula in \mathcal{L} is a strategy rule;*
2. *if r_1, \dots, r_m are strategy rules, then so is $r_1 \nabla r_2 \nabla \dots \nabla r_m$;*
3. *if r_1, \dots, r_m are strategy rules, then so is $r_1 \Delta r_2 \Delta \dots \Delta r_m$.*

Note that the new strategy connectives are introduced as macros rather than as additional connectives in the language \mathcal{L} . This is so because we do not want to allow the nesting of the strategy connectives with logical connectives (while nested strategy rules are allowed). For instance, $(\varphi_1 \Delta \psi_1) \nabla (\varphi_2 \Delta \psi_2)$ is a syntactically correct strategy rule while $\varphi \supset (\psi_1 \nabla \psi_2)$ is not.

5.2 Semantics of Strategy Rules

Given a strategy S , we let $S \upharpoonright_w = \{(w', a) \in S : w' = w\}$.

Definition 8 *Let r be a strategy rule. We define $S^i(r)$ recursively on the structure of r as follows:*

1. If $r = \varphi \in \mathcal{L}$, then $S^i(r) = S^i(\varphi)$.
2. If $r = r_1 \nabla r_2 \nabla \dots \nabla r_m$, then $(w, a) \in S^i(r)$ iff there exists k ($1 \leq k \leq m$) such that $(w, a) \in S^i(r_k)$ and $\bigcup_{j < k} (S^i(r_j) \upharpoonright_w) = \emptyset$.
3. If $r = r_1 \Delta r_2 \Delta \dots \Delta r_m$, then $(w, a) \in S^i(r)$ iff there exists k ($1 \leq k \leq m$) such that
 - (a) $(w, a) \in \bigcap_{j \leq k} S^i(r_j)$ and
 - (b) $k = m$ or $\bigcap_{j \leq k+1} (S^i(r_j) \upharpoonright_w) = \emptyset$.

Intuitively, $r_1 \nabla r_2 \nabla \dots \nabla r_m$ represents a strategy that combines the m strategies r_1, r_2, \dots, r_m in such a way that a strategy rule r_k becomes applicable only if none of the higher prioritised rules r_j ($j < k$) is applicable. $r_1 \Delta r_2 \Delta \dots \Delta r_m$ tries to apply as many strategy rules all together as possible but gives higher priority to the left rules than the right rules if conflicts occur.

The following observations show that the connective ∇ is indeed a kind of prioritised disjunction and the connective Δ is indeed a kind of prioritised conjunction.

Lemma 2 *Given a state transition model M , for each player i ,*

1. $S^i(r_1) \subseteq S^i(r_1 \nabla r_2)$
2. $S^i(r_1 \Delta r_2) \subseteq S^i(r_1)$
3. If $S^i(r_1) = \emptyset$, then $S^i(r_1 \nabla r_2) = S^i(r_2)$.
4. If $S^i(r_2) = \emptyset$, then $S^i(r_1 \Delta r_2) = S^i(r_1)$.
5. $S^i(\varphi_1 \nabla (\varphi_1 \wedge \varphi_2)) = S^i(\varphi_1)$
6. $S^i((\varphi_1 \wedge \varphi_2) \nabla \varphi_2) \subseteq S^i(\varphi_2)$
7. $S^i(\varphi_1 \Delta (\varphi_1 \vee \varphi_2)) = S^i(\varphi_1)$
8. $S^i(\varphi_2) \subseteq S^i((\varphi_1 \vee \varphi_2) \Delta \varphi_2)$

Proof: The proof of (1)-(4) is straightforward from Definition 8.

To show (5), let $(w, a) \in S^i(\varphi_1 \nabla (\varphi_1 \wedge \varphi_2))$. If $(w, a) \notin S^i(\varphi_1)$, then by Definition 8, $(w, a) \in S^i(\varphi_1 \wedge \varphi_2)$. By Lemma 1 (1), this implies $(w, a) \in S^i(\varphi_1)$ —a contradiction. Hence, $(w, a) \in S^i(\varphi_1)$. We have proved $S^i(\varphi_1 \nabla (\varphi_1 \wedge \varphi_2)) \subseteq S^i(\varphi_1)$. The other direction follows (1).

To show (6), let $(w, a) \in S^i((\varphi_1 \wedge \varphi_2) \nabla \varphi_2)$. Suppose that $(w, a) \notin S^i(\varphi_1 \wedge \varphi_2)$, then by Definition 8, we have $(w, a) \in S^i(\varphi_2)$, as desired. If, on the other hand, $(w, a) \in S^i(\varphi_1 \wedge \varphi_2)$, then by Lemma 1 (1) we also have $(w, a) \in S^i(\varphi_2)$.

To show (7), assume that $(w, a) \in S^i(\varphi_1)$. By Lemma 1 (2), $(w, a) \in S^i(\varphi_1 \vee \varphi_2)$. Hence, $(w, a) \in S^i(\varphi_1) \upharpoonright_w \cap S^i(\varphi_1 \vee \varphi_2) \upharpoonright_w$, which implies $(w, a) \in S^i(\varphi_1 \Delta (\varphi_1 \vee \varphi_2))$. The other direction follows (2).

To show (8), assume that $(w, a) \in S^i(\varphi_2)$. By Lemma 1 (2), $(w, a) \in S^i(\varphi_1 \vee \varphi_2)$. It follows that $(w, a) \in S^i(\varphi_1 \vee \varphi_2) \upharpoonright_w \cap S^i(\varphi_2) \upharpoonright_w$. By Definition 8, we have $(w, a) \in S^i((\varphi_1 \vee \varphi_2) \Delta \varphi_2)$. ■

Note that in the above lemma, only one direction of inclusion in items (6) and (8) holds. The other direction does not. For instance, assume that $S^i(\varphi_1) = \{(w, a)\}$ and $S^i(\varphi_2) = \{(w, a), (w, b)\}$. Then $S^i(\varphi_1 \wedge \varphi_2) = \{(w, a)\}$. It turns out that $S^i((\varphi_1 \wedge \varphi_2) \nabla \varphi_2) = \{(w, a)\}$. This shows that $S^i(\varphi_2) \not\subseteq S^i((\varphi_1 \wedge \varphi_2) \nabla \varphi_2)$. Furthermore, if $S^i(\varphi_1) = \{(w, a)\}$ and $S^i(\varphi_2) = \emptyset$. By Lemma 1, we have $(w, a) \in S^i(\varphi_1 \vee \varphi_2)$. Thus $(w, a) \in S^i((\varphi_1 \vee \varphi_2) \Delta \varphi_2) \neq \emptyset$, which means that the other direction of item (8) does not hold.

The following lemma shows that the prioritised disjunction can be reduced to binary connectives. Interestingly, the prioritised disjunction does not have such a nice property.

Lemma 3 *Given a state transition model M , for each player i ,*

$$S^i(r_1 \nabla r_2 \nabla \dots \nabla r_m) = S^i(r_1 \nabla (r_2 \nabla \dots \nabla r_m)) = S^i((r_1 \nabla \dots \nabla r_{m-1}) \nabla r_m)$$

Proof: We will prove the first equation, the other one will be quite similar. We consider two cases. If $S^i(r_1) \upharpoonright_w = \emptyset$ then

$$\begin{aligned} & (w, a) \in S^i(r_1 \nabla r_2 \nabla \dots \nabla r_m) \\ \text{iff } & (w, a) \in S^i(r_2 \nabla \dots \nabla r_m) \\ \text{iff } & (w, a) \in S^i(r_1 \nabla (r_2 \nabla \dots \nabla r_m)) \end{aligned}$$

On the other hand, if $S^i(r_1) \upharpoonright_w \neq \emptyset$ then

$$\begin{aligned} & (w, a) \in S^i(r_1 \nabla \dots \nabla r_m) \\ \text{iff } & (w, a) \in S^i(r_1) \\ \text{iff } & (w, a) \in S^i(r_1 \nabla (r_2 \nabla \dots \nabla r_m)) \end{aligned}$$

■

We remark that the prioritised conjunction cannot likewise be reduced to binary connectives. In general neither $S^i(r_1 \Delta r_2 \Delta \dots \Delta r_m) = S^i((r_1 \Delta \dots \Delta r_{m-1}) \Delta r_m)$ nor $S^i(r_1 \Delta r_2 \Delta \dots \Delta r_m) = S^i(r_1 \Delta (r_2 \Delta \dots \Delta r_m))$ is true. For example, let $S^i(r_1) = \{(w, a_1), (w, a_2)\}$; $S^i(r_2) = \emptyset$; and $S^i(r_3) = \{(w, a_1), (w, a_3)\}$. Then $S^i(r_1 \Delta r_2 \Delta r_3) = S^i(r_1) = \{(w, a_1), (w, a_2)\}$. But $S^i((r_1 \Delta r_2) \Delta r_3) = S^i(r_1 \Delta r_3) = \{(w, a_1)\}$. On the other hand, if $S^i(r_2) = \{(w, a_2), (w, a_3)\}$ and $S^i(r_1)$ and $S^i(r_3)$ are as before, then we get $S^i(r_1 \Delta r_2 \Delta r_3) = \{(w, a_2)\}$ while $S^i(r_1 \Delta (r_2 \Delta r_3)) = \{(w, a_1), (w, a_2)\}$.

5.3 Complete Strategies and Deterministic Strategies

When we build a software game player, we need to instruct it what to do in each possible situation. In other words, each player should be equipped with a functional strategy. In the following, we demonstrate how to generate a complete and/or deterministic strategy by using our prioritised connectives.

We say that a strategy rule is *consistent* for player i if it represents a valid strategy; a rule is *complete* for i if it represents a complete strategy for i ; and *deterministic* for i if it represents a deterministic strategy for i . Similarly, a strategy rule is *functional* if it is complete and deterministic. Note that all these concepts are player-specific.

The following theorems show a number of nice properties of the prioritised connectives, which give us a guideline for how to design a strategy with desired properties. The first result deals with consistency of strategy rules.

Theorem 1 *Given a state transition model M ,*

1. $r_1 \nabla \cdots \nabla r_m$ is consistent if and only if there is a k ($1 \leq k \leq m$) such that r_k is consistent.
2. $r_1 \Delta \cdots \Delta r_m$ is consistent if and only if r_1 is consistent.

Proof: To prove (1), let $r = r_1 \nabla \cdots \nabla r_m$. Assume that $S^i(r)$ is non-empty for player i , then there is $(w, a) \in S^i(r)$ such that $(w, a) \in S^i(r_k)$ for some k , which means that r_k is consistent for player i . Conversely, let k be the smallest number such that $S^i(r_k)$ is non-empty. This implies that $S^i(r_1 \nabla \cdots \nabla r_m) = S^i(r_k \nabla \cdots \nabla r_m)$. By the above two lemmas we have $S^i(r_k) \subseteq S^i(r_1 \nabla \cdots \nabla r_m)$, which means that r is consistent.

To prove (2), we also let $r = r_1 \Delta \cdots \Delta r_m$. Obviously if r_1 is inconsistent, so is r . Assume that r_1 is consistent. Then there is $(w, a) \in S^i(r_1)$. Let k be the biggest number such that $\bigcap_{j \leq k} (S^i(r_j) \upharpoonright_w) \neq \emptyset$. Obviously $k \geq 1$. By Definition 8 we have $\bigcap_{j \leq k} (S^i(r_j) \upharpoonright_w) \subseteq S^i(r) \upharpoonright_w$. Thus r is consistent. ■

The second result shows us how to generate a complete or deterministic strategy.

Theorem 2 *Given a state transition model M , for each player i*

1. If r_1 or r_2 is complete, so is $r_1 \nabla r_2$.
2. If r_1 is complete, so is $r_1 \Delta \cdots \Delta r_m$.
3. If r_1 and r_2 are deterministic, so is $r_1 \nabla r_2$.
4. If r_1 is deterministic, so is $r_1 \Delta \cdots \Delta r_m$.

Proof: (1) and (4) are straightforward from Definition (8).

To show (2), assume that $r_1 \Delta \cdots \Delta r_m$ is incomplete. Then there exists a reachable state $w \in W \cap \mathcal{P}(G)$ such that $S^i(r_1 \Delta \cdots \Delta r_m) \upharpoonright_w = \emptyset$. By Definition 8, this can happen only if $S^i(r_1) \upharpoonright_w = \emptyset$, which contradicts the assumption that r_1 is complete.

To prove (3), assume that $(w, a), (w, a') \in S^i(r_1 \nabla r_2)$. If $S^i(r_1) \upharpoonright_w \neq \emptyset$, we have $(w, a), (w, a') \in S^i(r_1)$ according to Definition 8. Since r_1 is deterministic, this implies $a = a'$. On the other hand, if $S^i(r_1) \upharpoonright_w = \emptyset$, then we have

$(w, a), (w, a') \in S^i(r_2) \upharpoonright_w$, which also implies $a = a'$ since r_2 is deterministic. ■

Statement (1) in the above theorem provides us with an easy way of generating a complete strategy: *create a trivial complete strategy first and then combine it with other strategies using the prioritised disjunction*. Note that creating a trivial complete strategy is rather easy: *let the agent do anything available*.¹ Statement (2) tells us that once we get a complete strategy, we can further refine the strategy targeting more specific properties, say deterministic thus functional, using the prioritised conjunction without losing its completeness. Statement (3) shows us another feasible way of generating a functional strategy: instead of create a complete strategy then refine it into a deterministic one, we can *devise a set of specific deterministic strategies first and then combine them into a complete one*. Example 4 demonstrates how the above mentioned can be applied to the CrossDot game. Before doing that, let's show another nice property of the prioritised connectives.

Theorem 3 *Given a state transition model M , for each player i*

1. *If r_1 is complete, then $S^i(r_1 \nabla \dots \nabla r_m) = S^i(r_1)$.*
2. *If r_1 is deterministic, then $S^i(r_1 \Delta \dots \Delta r_m) = S^i(r_1)$.*

Proof: To prove (1), we only have to show $S^i(r_1 \nabla \dots \nabla r_m) \subseteq S^i(r_1)$. Assume that $(w, a) \in S^i(r_1 \nabla \dots \nabla r_m)$. Because r_1 is complete, $S^i(r_1) \upharpoonright_w \neq \emptyset$. By Definition 8 we have $(w, a) \in S^i(r_1)$.

To show (2), we need to show that $S^i(r_1) \subseteq S^i(r_1 \Delta \dots \Delta r_m)$. If $(w, a) \in S^i(r_1)$, then $S^i(r_1) \upharpoonright_w = \{(w, a)\}$ since r_1 is deterministic. By Definition 8 it follows that $(w, a) \in S^i(r_1 \Delta \dots \Delta r_m)$. ■

Combining the two statements of this theorem, we can say that if r_1 is functional, then $S^i(r_1 \nabla \dots \nabla r_m) = S^i(r_1 \Delta \dots \Delta r_m) = S^i(r_1)$. This means that once a strategy is functional, neither extending nor refining it with strategies of lower priority has any effect. In fact, our goal of introducing the prioritised connectives is to facilitate the design of functional strategies. Once a functional strategy has been obtained, these connectives automatically stop working.

Example 4 Consider the CrossDot game scenario in Example 1. We define a few strategy rules for player i as follows:

- Fill a box next to a box that contains player i 's mark:

$$fill_next^i = \bigvee_{1 < j \leq m} (\neg p_j^1 \wedge \neg p_j^2 \wedge p_{j-1}^i \wedge does(a_j^i)) \vee \bigvee_{1 \leq j < m} (\neg p_j^1 \wedge \neg p_j^2 \wedge p_{j+1}^i \wedge does(a_j^i)) \quad (5.1)$$

- Fill an isolated box (i.e., whose immediate neighbours are empty):

$$fill_isolated^i = \bigvee_{1 < j < m} ((\neg p_{j-1}^1 \wedge \neg p_{j-1}^2) \wedge (\neg p_{j+1}^1 \wedge \neg p_{j+1}^2) \wedge (\neg p_j^1 \wedge \neg p_j^2) \wedge does(a_j^i)) \quad (5.2)$$

¹In this case, we might have to give the trivial strategy the lowest priority in the disjunction.

- Fill any empty box:

$$fill_any^i = \bigvee_{1 \leq j \leq m} (\neg p_j^1 \wedge \neg p_j^2 \wedge does(a_j^i)) \quad (5.3)$$

- Try $fill_next^i$ first. If this fails, try $fill_isolated^i$, then try $fill_any^i$:

$$combined^i = fill_next^i \nabla fill_isolated^i \nabla fill_any^i \quad (5.4)$$

Observation 3 Let $M = (G, v)$ be the state transition model defined in Observation 1. The strategy rule $combined^i$ is complete for player i .

Proof: We first prove that the strategy $fill_any^i$ is complete. Assume an arbitrary reachable state $w \in W \cap \mathcal{P}(G)$ such that $l^i(w) \neq \emptyset$ (recall formula (2.3) for the definition of l^i and the definition of completeness of a strategy in Section 3.1). Let $a_j^i \in l^i(w)$. It follows that $M \models_{(w, a_j^i)} legal(a_j^i)$. By (4.5), we know that $M \models_{(w, a_j^i)} \neg p_j^1 \wedge \neg p_j^2$, hence $M \models_{(w, a_j^i)} \neg p_j^1 \wedge \neg p_j^2 \wedge does(a_j^i)$, which implies $(w, a_j^i) \in S^i(fill_any^i)$. Therefore $fill_any^i$ is a complete strategy rule for player i . According to Theorem 2 (1), we know that the strategy rule $combined^i$ is also complete. \blacksquare

A complete strategy for a player gives the player a feasible option (if any) for any given situation. However, a software game player could still be unsure about what to do under a complete strategy if there is more than one option in the same situation. The prioritised conjunction provides us with a way to narrow down multiple options.

Example 5 For our running example game, let $c_t^i = \bigvee_{j=1}^t does(a_j^i)$ where $1 \leq t \leq m$ (with m being the overall number of boxes as usual), then c_t^i represents the strategy of i to place an object in any box between 1 and a given number t . Let

$$thoughtful^i = combined^i \triangle c_m^i \triangle \dots \triangle c_1^i \quad (5.5)$$

where $combined^i$ is the strategy rule defined in Example 4 (cf. (5.4)).

It is not hard to see that if $combined^i$ gives more than one boxes to fill, $thoughtful^i$ will choose the left most one.

Observation 4 Let $M = (G, v)$ be the state transition model defined in Observation 1. The strategy rule $thoughtful^i$ is functional for player i .

Proof: By Theorem 2 (2), $thoughtful^i$ is complete because $combined^i$ is complete. We prove that $thoughtful^i$ is deterministic. Assume that

$$(w, a_{j_1}^i), (w, a_{j_2}^i) \in S^i(thoughtful^i)$$

such that, without loss of generality, $j_1 \leq j_2$. Let t be the smallest number (between 1 and m) such that $H = S^i(combined^i) \upharpoonright_w \cap \bigcap_{j=t}^m S^i(c_j^i) \upharpoonright_w$ is not empty. Such a t exists because $S^i(combined^i) \upharpoonright_w \cap S^i(c_m^i) \upharpoonright_w \neq \emptyset$. By Definition (8), we have $(w, a_{j_1}^i) \in H$ and $(w, a_{j_2}^i) \in H$, hence $(w, a_{j_1}^i), (w, a_{j_2}^i) \in S^i(c_t^i) \upharpoonright_w$. Hence $j_1 \leq j_2 \leq t$. If $t = 1$, then $a_{j_1}^i = a_{j_2}^i = a_1^i$. Otherwise, i.e. if $t > 1$, then $S^i(combined^i) \upharpoonright_w \cap \bigcap_{j=t-1}^m S^i(c_j^i) \upharpoonright_w = \emptyset$, which implies that neither $(w, a_{j_1}^i)$ nor

$(w, a_{j_2}^i)$ belongs to $S^i(c_{t-1}^i) \upharpoonright_w$. Note that $S^i(c_{t-1}^i) \upharpoonright_w = \{(w, a_j^i) \in \mathcal{P}(G) : j \leq t-1\}$. Thus $t \leq j_1 \leq j_2$. By the above assumption, we have $j_1 = j_2 = t$. We have proved that $thoughtful^i$ is deterministic, thus it is functional. ■

To summarise, the prioritised connectives provide a natural way of refining a strategy. If a strategy is too restricted in that it can only be applied to few states, it can be extended using prioritised disjunction. If, on the other hand, a strategy is too generic in that it leaves too many options, it can be strengthened using prioritised conjunction. Once a strategy has been functional, further extension or refinement using the prioritised disjunction or conjunction take no effect as long as give the existing strategy the highest priority.

6 Reasoning About Strategies

Knowing how to write strategies for a game-playing agent, we now consider the question whether a strategy meets its goal, for instance, if it is guaranteed to lead to a winning state or to a desirable state given the strategies that the other players use. In this section, we will demonstrate by using our running example how to reason about strategies within our framework. We also demonstrate how to design a strategy using the prioritised connectives to meet desired properties.

6.1 Compliance with a strategy

In order to verify whether a strategy can bring about an expected result for a player in a game, we assume that the player complies with the strategy all the way through the game and observes the outcome of the game.

Let $M = (G, v)$ be a state transition model of a game G and S a strategy of the game for player i . We say that M (or G) *complies with S by player i* if for each reachable move (w, a) of player i in G , $(w, a) \in S$. In other words, player i follows the strategy S whenever he makes a move.¹

The following observation shows that for any CrossDot game when $k = 2$ and $m > 2$, Player 1 wins as long as he plays the strategy $thoughtful^i$ all the way through the game.

Observation 5 *Let M be a state transition model for the CrossDot game with $k = 2$ and $m > 2$. Assume that player i takes the first turn. If M complies with the strategy $S^i(thoughtful^i)$ by player i , then*

$$M \models terminal \supset wins(1)$$

In other words, player i wins as long as he takes the first turn and follows the strategy rule $thoughtful^i$.

Proof: Without loss of generality, we assume that $i = 1$. We have to show that $M, \delta \models terminal \supset wins(i)$ for any reachable path δ . Since the definitions

¹We may understand the concept in the following way. Given an arbitrary state transition model $M = (G, v)$ and a strategy S of player i in M , let M' be another state transition model that is exactly the same as M except for the legality of actions for player i in such a way that the legality relation l' in M' is $l \cap S$, where l is the legality relation in M . We then view M' as the reduction of M once player i complies with strategy S .

of *terminal* and *wins* do not contain *does*, *legal* and \bigcirc (cf. formulas (4.3) & (4.4)), we only need to show for any reachable terminal state w_e (which is a special case of a reachable path) $M, w_e \models \text{wins}(1)$. To this end, we assume a complete path $\delta = w_0 \xrightarrow{a_0} w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{e-1}} w_e$, where $w_0 = \bar{w}$ and $w_e \in t$. Since player 1 has the first turn and plays with strategy *thoughtful*¹, the action the player takes in the initial state must be a_2^1 , i.e., $a_0 = a_2^1$. If player 2 responds with action a_1^2 , i.e., $a_1 = a_1^2$, player 1 will then take action a_3^1 and win the game at state w_3 . If player 2 responds with any other action, player 1 will take a_1^1 and also will win the game. In any case, δ ends up with $w_e = w_3$ in which player 1 wins. Therefore $M, w_e \models \text{wins}(1)$. We conclude $M \models \text{terminal} \supset \text{wins}(1)$. ■

The strategy *thoughtful* seems like a “smart” strategy that can guarantee a winning state for the player who takes the first turn when $k = 2$. However it is less mighty when taken by the second player, in which case the strategy cannot even compete the following “trivial” strategy

$$\text{fill_leftmost}^i = c_m^i \Delta \dots \Delta c_1^i \quad (6.1)$$

where c_j^i was defined in Example 5.

Observation 6 *Let M be a state transition model for a CrossDot game with $k = 2$ and $m > 2$. If M complies with $S^1(\text{fill_leftmost}^1)$ by Player 1 and with $S^2(\text{thoughtful}^2)$ by Player 2, then*

$$M \models \text{terminal} \supset \text{wins}(1)$$

Proof: It is not hard to prove that fill_leftmost^1 is functional. Since thoughtful^2 is also functional, there is only one complete path in the game:

$$\begin{aligned} (1, 0, \square, \square, \square, \square, \dots, \square) &\xrightarrow{a_1^1} (0, 1, \boxtimes, \square, \square, \square, \dots, \square) \\ &\xrightarrow{a_3^2} (1, 0, \boxtimes, \square, \square, \square, \dots, \square) \xrightarrow{a_2^1} (0, 1, \boxtimes, \boxtimes, \square, \square, \dots, \square) \end{aligned}$$

which implies that player 1 takes the left most box, followed by player 2 fills an isolated box and finally player 1 fills the second box and wins. We can easily verify that $\text{terminal} \supset \text{wins}(1)$ is valid with any segment of the path in M . ■

The failure of *thoughtful* is not because it is not “smart” enough but because the *thoughtful*² strategy requires to fill an isolated box at the first move, which provides the first player with a chance to win.

6.2 Reasoning About Other Players’ Strategies

The examples of strategies we have shown up to now are all from a single player’s viewpoint, which is obviously not sufficient. We should also reason about other players’ strategies.

From Observation 6 we learnt that a player should check for existing threats before applying any “aggressive” strategy, like *thoughtful*. The following formula defines a *defence* strategy, which says that if my opponent can win by filling box j , then I should mark j first to prevent an immediate loss:

$$\begin{aligned} \text{defence}^i = & \\ & (\bigwedge_{j=1}^m (\bigcirc(\text{does}(a_j^{-i}) \wedge \bigcirc \text{wins}(-i)) \supset \text{does}(a_j^i))) \Delta c_m^i \Delta \dots \Delta c_1^i \end{aligned} \quad (6.2)$$

where $-i$ stands for the opponent of i . Note that $defence^i$ is deterministic but not complete.

The following observation shows that if the second player plays $defence^2 \nabla thoughtful^2$, which means to protect himself before attacking his opponent, then player 1 cannot win with the $fill_leftmost$ strategy (ref. equation 6.1).

Observation 7 *Let M be a state transition model for the CrossDot game with $k = 2$ and $m > 2$. If M complies with $S^1(fill_leftmost^1)$ by Player 1 and with $S^2(defence^2 \nabla thoughtful^2)$ by Player 2, then*

$$M \models terminal \supset \neg wins(1)$$

Proof: Let $cautious^2 = defence^2 \nabla thoughtful^2$. Since $thoughtful^2$ is complete, the strategy $cautious^2$ is complete. Both $thoughtful^2$ and $defence^2$ are deterministic, which implies that $cautious^2$ is deterministic according to Theorem 2(3). Therefore $cautious^2$ is functional. Note that $fill_leftmost^1$ is also functional. Thus there is only one complete path in the game. Assume that the complete path is $\delta = w_0 \xrightarrow{a_0} w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{e-1}} w_e$. The first two states in the path must be

$$w_0 = (1, 0, \square, \square, \square, \square, \dots, \square) \text{ and } w_1 = (0, 1, \boxtimes, \square, \square, \square, \dots, \square)$$

and the action taken by player 1 must be $a_0 = a_1^1$. To know what action player 2 takes, we calculate $S^2(defence^2)|_{w_1}$. To simplify the calculation, let

$$defe^2 = \bigwedge_{j=1}^m (\bigcirc(does(a_j^1) \wedge \bigcirc wins(1)) \supset does(a_j^2))$$

We will show $S^2(defe^2)|_{w_1} = \{(w_1, a_2^2)\}$, which implies $S^2(defence^2)|_{w_1} = S^2(defe^2)|_{w_1} = \{(w_1, a_2^2)\}$. To verify $(w_1, a_2^2) \in S^2(defe^2)|_{w_1}$, assume any reachable path $\delta' \in (w_1, a_2^2)^{\rightsquigarrow}$. We verify

$$M, \delta' \models \bigwedge_{j=1}^m (\bigcirc(does(a_j^1) \wedge \bigcirc wins(1)) \supset does(a_j^2))$$

This is obviously true because if $j = 2$ then $does(a_j^2)$ is satisfied; otherwise $\bigcirc(does(a_j^1) \wedge \bigcirc wins(1))$ is false for any δ' .

Next we show $(w_1, a_1) \notin S^2(defe^2)|_{w_1}$ provided $a_1 \neq a_2^2$. Let $\delta' = w_1 \xrightarrow{a_1} w_2 \xrightarrow{a_2^1} w_3$. It is easy to verify that

$$M, \delta' \models \bigcirc(does(a_2^1) \wedge \bigcirc wins(1)) \wedge \neg does(a_2^2)$$

which implies

$$M, \delta' \not\models \bigwedge_{j=1}^m (\bigcirc(does(a_j^1) \wedge \bigcirc wins(1)) \supset does(a_j^2))$$

Note that $\delta' \in (w_1, a_1)^{\rightsquigarrow}$. Therefore we have

$$M \not\models_{(w_1, a_1)} \bigwedge_{j=1}^m (\bigcirc(does(a_j^1) \wedge \bigcirc wins(1)) \supset does(a_j^2))$$

which implies $(w_1, a_1) \notin S^2(defe^2) \upharpoonright_{w_1}$. This completes the proof that $S^2(defe^2) \upharpoonright_{w_1} = \{(w_1, a_2^2)\}$. Hence, for the complete path δ we find that $a_1 = a_2^2$ and $w_2 = (1, 0, \boxtimes, \square, \square, \square, \dots, \square)$. The game continues in the same way until all boxes are filled without a winner. ■

The instances of the CrossDot game we considered above are limited to the simple case where $k = 2$. In the following, we proposed a solution to the game in the general setting where k can be any number larger than 2 (and $\leq m$). Consider the following strategies:

- Fill a box next to an opponent's box:

$$fill_o_next^i = \left(\bigvee_{1 \leq j < m} (\neg p_j^1 \wedge \neg p_j^2 \wedge p_{j+1}^{-i} \wedge does(a_j^i)) \right) \nabla \left(\bigvee_{1 < j \leq m} (\neg p_j^1 \wedge \neg p_j^2 \wedge p_{j-1}^{-i} \wedge does(a_j^i)) \right)$$

Note that priority is given to the left empty box if exists.

- Passive defence:

$$passive_defence^i = (defence^i \nabla fill_o_next^i \nabla fill_any^i) \Delta c_m^i \Delta \dots \Delta c_1^i \tag{6.3}$$

It is not hard to see that $passive_defence^i$ is functional. The following observation shows that such a passive defence strategy guarantees no loss for a player no matter whether it is taken by the first player or the second player, and no matter what the strategy of the other player, for any instance of the CrossDot game with $k > 2$ and $m > 2$.

Observation 8 *Let M be a state transition model for a CrossDot game with $k > 2$ and $m \geq k$. If M complies with $S^i(passive_defence^i)$ by player i , then*

$$M \models terminal \supset \neg win(-i)$$

where $-i$ represents the opponent player of i . In other words, a player never loses as long as he plays the $passive_defence$ strategy.

Proof: Obviously we only have to prove the claim for the case when $k = 3$ because if a strategy can help a player effectively blocking his opponent to own three consecutive boxes, he can also block his opponent to form any longer line of consecutive boxes. Also if we can prove the strategy to be effective for player 2, it is sufficient to show that it is also effective for player 1 because he can easily ignore his first move. Altogether this allows us to restrict the proof to $k = 3$ and $i = 2$.

Since $passive_defence^2$ is functional, to show $M \models terminal \supset \neg win(1)$, we only have to verify that for any complete path $\delta = w_0 \xrightarrow{a_0} w_1 \xrightarrow{a_1} \dots \xrightarrow{a_{e-1}} w_e$, we have $M, w_e \models \neg win(1)$. Recall that w_e is a special reachable path that contains only the terminal state w_e . If w_e is a winning state for player 2 or a tie state, we have $M, w_e \models \neg win(1)$.

Suppose by contradiction that w_e is a winning state of player 1. It then turns out e is an odd number and there are at least three consecutive crosses in some states along the complete path δ . In other words, there must be a cross

next to two other crosses. However, we will show by induction on the length of the path that this can never happen. More precisely, we claim that *in each state w_{2j} ($0 \leq j \leq (e-1)/2$), each box filled with a cross must be either adjacent to a box with a dot or at an end (left or right) of the line, meanwhile the box on its left must not be empty*. Note that if the claim is true in w_{e-1} , then player 1 cannot win in state w_e no matter which action he does in state w_{e-1} .

Obviously the claim holds when $j = 0$ because all boxes are empty in the initial state w_0 . Now we assume that the claim holds in state $w_{2(j-1)}$ where $j \leq (e-1)/2$. We show that the claim also holds at w_{2j} .

Suppose that the action $a_{2(j-1)}$ that player 1 took in state $w_{2(j-1)}$ was a_v^1 , i.e., $a_{2(j-1)} = a_v^1$. Obviously this action does not affect the status of any box other than the box itself and its neighbours, i.e., $v-1$, v and $v+1$ (if exists). Therefore we only consider the effect of player 1's action a_v^1 on these three possible boxes.

If box v is not at the right end, by the induction assumption, box $v+1$ can never be with a cross in state $w_{2(j-1)}$; otherwise box v were not empty thus a_v^1 were not doable. Therefore box $v+1$ (if exists) satisfies the claimed conditions trivially in state w_{2j} .

We then consider the status of boxes v and $v-1$ with three cases:

(i). If box $v-1$ (if exists) has been filled with a dot, both boxes $v-1$ and v satisfy the conditions of the claim, which remain true in state w_{2j} .

(ii). If box $v-1$ has been already filled with a cross before state $w_{2(j-1)}$, by induction assumption, box $v-1$ is either at the left end or next to a box with a dot on its left. In both situations, box $v-1$ satisfies the claimed conditions that remain true in state w_{2j} . After player 1 fills box v with a cross in state $w_{2(j-1)}$, player 2 must respond with action $a_{2j-1} = a_{v+1}^2$ by applying *defence*² in state w_{2j-1} to prevent an immediate loss unless the box v is at the right end or $v+1$ has already occupied by a dot (note that box $v+1$ can never be with a cross). In both cases, box v satisfies the claimed conditions in state w_{2j} .

(iii). If box $v-1$ was empty in state $w_{2(j-1)}$, by the construction of *passive_defence*ⁱ, player 2 must respond with action a_{v-1}^2 by applying strategy *fill_o_next*² after player 1 fills box v . In this case, both boxes $v-1$ and v satisfy the claim conditions in state w_{2j} . Note that *fill_o_next*² is not applicable to any box other than $v-1$ because the induction assumption guarantees that in state $w_{2(j-1)}$ any box with a cross does not have an empty box on its left and *fill_o_next*² gives priority to the boxes with left empty boxes.

We have verified the claim holds in each state w_{2j} ($0 \leq j \leq (e-1)/2$), including state w_{e-1} , which implies that the player 2 does not lose in state w_e . ■

Note that the *defence* strategy (cf. 6.2) plays a crucial role in the solution. Assume that a game comes to the following situation after a few moves:

$$\square, \boxtimes, \square, \square, \square, \boxtimes, \boxtimes, \square, \dots$$

and it is player 2's turn. If player 2 does not take strategy *defence*² but applies *fill_o_next*², he will take action a_3^2 instead of a_8^2 , which gives player 1 a chance to win:

$$\square, \boxtimes, \square, \square, \square, \boxtimes, \boxtimes, \boxtimes, \dots$$

Although the *passive_defence* strategy can effectively prevent from losing a game, it is hard to win a game because it does not encode any winning strategy. The reader is invited to extend the strategy with more aggressive rules, say *thoughtful*, so that if the other player is “not that smart”, it can have a chance to win.

7 Computing With Strategies

We now turn to the question of how to actually compute with strategy rules. Generally speaking, the conceptual simplicity of our language and the fact that it is not tied to a specific action formalism should make it easy to incorporate knowledge of strategies into various methods for the design and analysis of intelligent agents. To illustrate this, we will adopt here a very general calculus for reasoning about actions, the Situation Calculus (see, e.g., [17]), and show how our strategy representation can be easily integrated. Reasoning problems about strategies can take different forms, and we will specifically consider two of them: We will show how the calculus allows you to infer the possible outcomes of a game given information about the strategies of all players, and we will illustrate how players can reason about the strategies of opponents to infer their best course of action. We will also show how our Situation Calculus dialect forms the basis for an encoding of game rules and strategies into Answer Set Programs that allows to compute with a restricted class of strategies.

7.1 Example: Situation Calculus

The Situation Calculus is a formalism for reasoning about actions and change that is based on classical predicate logic with a few pre-defined language elements:

- constant s_0 , denoting the *initial situation*, and constructor $\text{DO}(\alpha, \sigma)$, denoting the situation resulting from doing action α in situation σ ;
- predicate $\text{HOLDS}(\varphi, \sigma)$, denoting that *fluent* φ (i.e., an atomic state feature) is true in situation σ ;
- predicate $\text{POSS}(\alpha, \sigma)$, denoting that action α is possible in situation σ .

For our purpose, we will extend the base language of the Situation Calculus by the two game-specific predicates $\text{WINS}(\iota, \sigma)$ and $\text{TERMINAL}(\sigma)$ meaning, respectively, that the game is won for player ι in situation σ and that σ is a terminal game position. With this, any axiom in our game specification language can be easily rewritten for the Situation Calculus along the line of [19]. To this end, let A and S be two distinct variables (standing for any action and situation, respectively), then a formula φ in our language \mathcal{L} can be rewritten into Situation

Calculus $\varphi^{SC}[A, S]$ inductively:

$$\begin{aligned}
p^{SC}[A, S] &:= \text{HOLDS}(p, S) \\
(\neg\varphi)^{SC}[A, S] &:= \neg\varphi^{SC}[A, S] \\
(\varphi \wedge \psi)^{SC}[A, S] &:= \varphi^{SC}[A, S] \wedge \psi^{SC}[A, S] \\
\text{does}(a)^{SC}[A, S] &:= A = a \\
\text{legal}(a)^{SC}[A, S] &:= \text{POSS}(a, S) \\
\text{wins}(i)^{SC}[A, S] &:= \text{WINS}(i, S) \\
(\bigcirc\varphi)^{SC}[A, S] &:= (\forall A') \varphi^{SC}[A', \text{Do}(A, S)] \\
\text{init}^{SC}[A, S] &:= S = s_0 \\
\text{terminal}^{SC}[A, S] &:= \text{TERMINAL}(S)
\end{aligned}$$

EXAMPLE 1 (*continued*) Recall the specification of the game in Example 1. Applying the construction from above to formulas (4.1)–(4.8), and performing some slight syntactic simplifications on the result, yields the following encoding of our example game in the Situation Calculus.

$$\begin{aligned}
&\neg\text{HOLDS}(p_j^i, s_0) \\
&\text{HOLDS}(\text{turn}(1), s_0) \wedge \neg\text{HOLDS}(\text{turn}(2), s_0) \\
&\text{WINS}(i, S) \equiv \bigvee_{j=1}^{m-k+1} \bigwedge_{l=j}^{j+k-1} \text{HOLDS}(p_l^i, S) \\
&\text{TERMINAL}(S) \equiv \text{WINS}(1, S) \vee \text{WINS}(2, S) \vee \\
&\quad (\bigwedge_{1 \leq j \leq m} (\text{HOLDS}(p_j^1, S) \vee \text{HOLDS}(p_j^2, S))) \\
&(\neg(\text{HOLDS}(p_j^1, S) \vee \text{HOLDS}(p_j^2, S)) \wedge \\
&\text{HOLDS}(\text{turn}(i), S) \wedge \neg\text{TERMINAL}(s)) \equiv \text{POSS}(a_j^i, S) \\
&\text{HOLDS}(p_j^i, S) \vee A = a_j^i \equiv \text{HOLDS}(p_j^i, \text{Do}(A, S)) \\
&\text{HOLDS}(\text{turn}(1), S) \supset \neg\text{HOLDS}(\text{turn}(1), \text{Do}(A, S)) \wedge \\
&\quad \text{HOLDS}(\text{turn}(2), \text{Do}(A, S)) \\
&\text{HOLDS}(\text{turn}(2), S) \supset \neg\text{HOLDS}(\text{turn}(2), \text{Do}(A, S)) \wedge \\
&\quad \text{HOLDS}(\text{turn}(1), \text{Do}(A, S))
\end{aligned}$$

7.2 Adding Strategy Rules

As a general logic-based formalism, our example calculus allows for a straightforward encoding of strategy rules with the help of a direct encoding of their interpretation according to (4.9) and Definition 8. Based on the rewriting rules

from above, for a strategy rule r over language \mathcal{L} the Situation Calculus encoding $r[A, S]$, where A and S are variables, is inductively obtained as follows.

$$\begin{aligned} \varphi[A, S] &:= \text{Poss}(A, S) \wedge \varphi^{SC}[A, S] \\ (r_1 \nabla r_2 \nabla \dots \nabla r_n)[A, S] &:= \\ & \quad r_1[A, S] \\ & \quad \vee \\ & \quad r_2[A, S] \wedge \neg(\exists A') r_1[A', S] \\ & \quad \vee \dots \vee \\ & \quad r_n[A, S] \wedge \neg(\exists A') (r_1[A', S] \vee \dots \vee r_{n-1}[A', S]) \\ (r_1 \Delta r_2 \Delta \dots \Delta r_n)[A, S] &:= \\ & \quad r_1[A, S] \wedge \neg(\exists A') (r_1[A', S] \wedge r_2[A', S]) \\ & \quad \vee \\ & \quad r_1[A, S] \wedge r_2[A, S] \wedge \neg(\exists A') (r_1[A', S] \wedge r_2[A', S] \wedge r_3[A', S]) \\ & \quad \vee \dots \vee \\ & \quad r_1[A, S] \wedge r_2[A, S] \wedge \dots \wedge r_n[A, S] \end{aligned}$$

7.3 Computing with Strategies: Examples

The encoding of strategy rules can be used in the Situation Calculus to define a special predicate, which we denote by $\text{STRAT}(\alpha, \sigma)$, whose intended meaning is that it is possible according to some player's strategy to take action α in situation σ . Consider, for example, a given set of complete strategy rules $\{r_1, \dots, r_n\}$ for each player, then this can be embedded into a Situation Calculus encoding of a game as follows:

$$\text{STRAT}(A, S) \equiv r_1[A, S] \vee \dots \vee r_n[A, S]$$

Such information about strategies can be used for a variety of purposes. Specifically, as we will briefly illustrate next, it can be used to infer possible outcomes of a game under a given set of strategy rules and to help a player to decide on a course of action by reasoning about opponents' strategies.

Inferring possible outcomes. Based on the predicate STRAT , the set of all possible playouts of a game according to players' strategies can be recursively defined as follows.

$$\begin{aligned} & \text{STRATEGIC}(s_0) \\ & \text{STRATEGIC}(S) \wedge \text{STRAT}(A, S) \supset \text{STRATEGIC}(\text{Do}(A, S)) \end{aligned}$$

This predicate determines all paths that are reachable if all players follow their given strategy rules. Hence, all possible playouts under these strategies can be determined by $\text{STRATEGIC}(S) \wedge \text{TERMINAL}(S)$.

Reasoning about opponents' strategies. Another way of using reasoning about strategies is for players to use knowledge or belief about their opponents' strategies in order to compute their own best course of actions in case the opponents do follow the assumed strategy rules. As an example, we will consider the encoding of a generalised form of Minimax evaluation in the Situation Calculus.

```

1 action(a(I,J)).
2
3 holds(turn(1),0).
4 legal(a(I,J),T) :- not holds(p(1,J),T), holds(turn(I),T),
5                     not holds(p(2,J),T), not terminal(T).
6 holds(p(I,J), T+1) :- holds(p(I,J),T).
7 holds(p(I,J), T+1) :- does(a(I,J)).
8 holds(turn(1),T+1) :- holds(turn(2),T).
9 holds(turn(2),T+1) :- holds(turn(1),T).
10
11 wins(I,T) :- holds(p(I,J),T), ..., holds(p(I,J+k),T).
12 terminal(T) :- wins(I,T).
13 terminal(T) :- not freecell(T).
14 freecell(T) :- not holds(p(1,J)), not holds(p(2,J)).
15
16 1 { does(A,T) : action(A) } 1.
17 :- does(A,T), not legal(A,T).
18
19 :- non_strategic(T).
20 non_strategic(T) :- does(A,T), not strat(A,T).
21
22 strat(A,T) :- fill_next(A,T).
23 strat(A,T) :- not exists_fill_next(T), fill_any(A,T).
24 fill_next(a(I,J),T) :- holds(p(I,J-1),T), legal(a(I,J),T),
25                       not holds(p(1,J),T), not holds(p(2,J),T).
26 fill_next(a(I,J),T) :- holds(p(I,J+1),T), legal(a(I,J),T),
27                       not holds(p(1,J),T), not holds(p(2,J),T).
28 exists_fill_next(T) :- fill_next(A,T).
29 fill_any(a(I,J),T) :- legal(a(I,J),T).

```

Figure 7.1: An ASP encoding of the game from Example 1, including a strategy rule. For the sake of brevity, we have omitted the domain definitions for variables I , J , and T (the latter ranging from 0 to a given time horizon).

To this end, let us take the stance of particular player i and assume that this player's belief about the opponents' strategy rules is encoded using predicate STRAT as above. Let us assume also that TURN(i, S) expresses the fact that it is this player i 's turn in situation S . We can then define recursively the notion of a win situation for i as follows:

$$\text{WINS}(i, S) \supset \text{WIN}(S) \quad (7.1)$$

$$\begin{aligned} \text{TURN}(i, S) \wedge (\exists A)(\text{POSS}(A, S) \wedge \text{WIN}(\text{DO}(A, S))) \\ \supset \text{WIN}(S) \end{aligned} \quad (7.2)$$

$$\begin{aligned} \neg \text{TURN}(i, S) \wedge (\forall A)(\text{STRAT}(A, S) \supset \text{WIN}(\text{DO}(A, S))) \\ \supset \text{WIN}(S) \end{aligned} \quad (7.3)$$

According to this definition, a situation is winnable for our player i if he can choose a course of action whenever it is his turn (7.2) such that if all other players choose any actions according to their strategy (7.3) then a situation that is won by definition will be reached (7.1).

7.4 Computing with Strategies in ASP

Under specific conditions, problems that require reasoning about games and strategies can be straightforwardly solved by Answer Set Programming (ASP), which provides a way of computing models for logic programs for which particularly efficient implementations have been developed in the recent past, such as [5] just to mention one. ASP has been used successfully in the past for reasoning

about actions and plan generation (see, e.g., [10]) as well as for endgame search in general game playing [20]. In this section, we build on these existing methods and show how ASP can be used to compute all possible outcomes of a game under given strategies for the players. We assume the reader to be familiar with basic notions and notations of ASP, as can be found in [6].

The standard use of ASP for computing with actions is to replace the branching time structure of the Situation Calculus by linear time. For deterministic games with complete specification of the initial state and a given time horizon, the game rules can be encoded as an ASP in such a way that each answer set corresponds to a reachable path and vice versa (see, e.g., [20]). For our running example game, Figure 7.1 (lines 1–18) constitute an ASP encoding that follows this principle. Specifically, the so-called *weight atom* in clause 16 requires each answer set to include one and only one action at each point in time, and the so-called *constraint* in clause 17 rules out any answer set in which the chosen action is not legal.

A given ASP that provides a linear-time encoding of a game can be extended by encodings of strategy rules for the players in such a way that the answer sets comply with the strategies. This provides a computational method for inferring the possible outcomes of a game under a given set of strategies.

Any strategy φ in our language that does not include the \bigcirc -operator can be encoded in this linear way according to the following inductive coding scheme. Let \mathbf{n} be a unique predicate name that stands for (the satisfaction of) φ , then:

φ	encoding
p	$\mathbf{n}(A, T) :- \text{holds}(p, T).$
$\neg\psi_1$	$\mathbf{n}(A, T) :- \text{not } \mathbf{n}_1(A, T).$
$\psi_1 \wedge \psi_2$	$\mathbf{n}(A, T) :- \mathbf{n}_1(A, T), \mathbf{n}_2(A, T).$
$\text{does}(a)$	$\mathbf{n}(A, T) :- A = a.$
$\text{legal}(a)$	$\mathbf{n}(A, T) :- \text{legal}(a, T).$
$\text{wins}(i)$	$\mathbf{n}(A, T) :- \text{wins}(i, T).$
init	$\mathbf{n}(A, T) :- T = 0.$
terminal	$\mathbf{n}(A, T) :- \text{terminal}(T).$
$\psi_1 \nabla \psi_2 \nabla \dots \nabla \psi_m$	$\mathbf{n}(A, T) :- \mathbf{n}_1(A, T).$ $\mathbf{n}(A, T) :- \mathbf{n}_2(A, T), \text{not } \mathbf{n}'(T).$ $\mathbf{n}(A, T) :- \mathbf{n}_3(A, T), \text{not } \mathbf{n}''(T).$ \dots $\mathbf{n}'(T) :- \mathbf{n}_1(A', T).$ $\mathbf{n}''(T) :- \mathbf{n}_2(A', T).$ \dots
$\psi_1 \Delta \psi_2 \Delta \dots \Delta \psi_m$	$\mathbf{n}(A, T) :- \mathbf{n}_1(A, T), \text{not } \mathbf{n}'(T).$ $\mathbf{n}(A, T) :- \mathbf{n}_1(A, T), \mathbf{n}_2(A, T), \text{not } \mathbf{n}''(T).$ \dots $\mathbf{n}'(T) :- \mathbf{n}_1(A', T), \mathbf{n}_2(A', T).$ $\mathbf{n}''(T) :- \mathbf{n}_1(A', T), \mathbf{n}_2(A', T), \mathbf{n}_3(A', T).$ \dots

Where necessary, this is accompanied by clauses with head \mathbf{n}_i to encode the sub-formulas ψ_i (for $i = 1, 2, \dots$), which are obtained inductively.

Clauses 19–29 in Figure 7.1 serve as an example of applying this encoding to constrains the answer sets for our *CrossDot* game to those where both

players i follow the simple strategy $(fill_next^i) \nabla (fill_any^i)$ (cf. Example 4): Constraint 19 in conjunction with clause 20 rejects all answer sets that do not comply with the strategy definition. Clauses 24–27 encode the first part of the strategy rule for both players (in a slightly more compact form than obtained by strictly applying the coding scheme from above), and clause 23 (in conjunction with clauses 28–29) encodes the second part.

7.5 Further extensions and computational complexity

The encoding scheme from above does not extend to strategies that include the \bigcirc -operator since their evaluation requires a counterfactual lookahead. Hence, they cannot be directly represented in an ASP based on a linear time structure. There are two conceivable ways to overcome this limitation.

1. If the axiomatisation of a game supports the definition of a regression operator similar to the one in the standard Situation Calculus [16], then any strategy of the form $\bigcirc\varphi$ can be regressed to a formula φ' that is logically equivalent under the game axioms and contains one less occurrence of the \bigcirc -operator. The repeated application of regression will yield a \bigcirc -free formula, which then can be encoded in the same way as above.
2. Alternatively, we can extend the linear time structure to allow for more than one sequence of actions in a single same answer set. A similar approach has been shown to be practically viable for proving epistemic properties in general games using ASPs [9].

A detailed formalisation and analysis of either solution goes beyond the scope of this report and is left for future work.

Up to now we have demonstrated with two example implementations how our formalism supports automated reasoning. At this stage, the efficiency of the implementations has not been analysed in detail. In fact, we can consider different aspects of efficiency with regard to both the description of strategies and the problem of reasoning about them.

The complexity of translating a strategy in our language to situation calculus is linear, and so is the translation to ASP for \bigcirc -free strategy rules. However, generally speaking it is unrealistic to expect high efficiency for a generic strategy reasoning mechanism because in theory, as an extension of GDL, our language can describe any finite game with perfect information. Verifying an arbitrary strategy such as, “try any possible action to win,” is equivalent to solving a game and therefore equally complex. It is well known that the complexity of finding a winning strategy for complex games like Japanese Go, which can be specified in GDL, is EXPTIME-complete [18].

Moreover, the computational complexity of solving a game sometimes is independent of the length of the game description and the strategies that are used to solve the game. Therefore, a complexity analysis for domain-independent strategy reasoning mechanisms can be meaningless¹. However, it is possible that certain restrictions on both game descriptions and strategy rules may lead to specific upper bounds for the computational complexity of solving these games. Restrictions on the number of lookahead steps can also affect the complexity of

¹For instance, the description of game rules for Chinese Go is quite similar to the one for Japanese Go but their complexity is significantly different.

reasoning. In addition, since our strategies represent some ideas of how to play a game well, it is possible that we express these ideas in our language and design specific algorithms to automatically generate possible moves that correspond to the ideas. The efficiency of these algorithms is crucial to the design of a game player. We leave these issues for future investigation.

8 Related Work

Modelling and specifying strategies is a fundamental research theme in game theory. Researchers in artificial intelligence have recently joined in the research but mostly focus on modelling of strategic reasoning with the help of logical approaches. A number of logical frameworks have been proposed in the literature for strategy representation and reasoning [13, 1, 8, 3, 12, 24]. Most of the frameworks were built on either *Coalition Logic* (CL), *Alternating-Time Temporal Logic* (ATL), or *Propositional Dynamic Logic* (PDL).

Both coalition logic [13] and alternating-time temporal logic [1] were developed to model strategic abilities of coalitions in multi-agent systems. The modality $\langle C \rangle \varphi$ (or $\langle C \rangle \bigcirc \varphi$) expresses that “a group of agents, C , has a joint strategy to bring about φ no matter what strategies the other agents choose”. In CL, a strategy of a player is simply an action available to the player. In ATL, a strategy is a function that maps a sequence of states to an action (in ATL). In both logics, strategies stay on the semantic level without syntactical representation.

A number of extensions of either CL or ATL aim to bring strategies to the syntactical level. van der Hoek et al. [25] proposed an extension of ATL, named CATL for *Counterfactual ATL*, with a variation to the coalition modality, $C_i(\sigma, \varphi)$, representing the counterfactual statement, “if agent i had committed to a strategy σ , then φ would hold”. Strategies in that framework can be explicitly represented on the syntactical level using dynamic logic-like modalities, even though program connectives of dynamic logic are not allowed to be used for combining strategies. Walther et al. [26] refine the work of CATL into an axiomatic logical system with a different semantics. However, strategies are still restricted to primitive forms, which means that the combination of strategies is not supported. Similar restrictions have also been applied in several other ATL- or CL-like logical frameworks for strategic reasoning, such as [3, 12, 11].

Another approach to strategy representation and reasoning is to treat a strategy as a program so that PDL-style program connectives can be used to combine strategies [22, 15, 24]. van Benthem proposed a logical framework, named Temporal Forcing Logic (TFL), with a modality $[\sigma, i]\varphi$, meaning that “player i applies strategy σ , against any play of the others, to force the game to a state in which φ holds”, where a strategy can be defined as any PDL program. Similar proposal can also be found in [15]. Such an “intuitive analogue to strategies” provides a close approximation to strategy representation; nevertheless, a strategy has essential differences from a program, which requires specific ways of composition and reasoning as we have shown in the previous sections.

We like to stress that our treatment of strategies is different from all of the abovementioned approaches in the following aspects. Firstly, we can use the same propositional formulas for different purposes. A propositional formula with the standard semantics of propositional modal logic can represent properties of

the game state and be used in domain-dependent axioms. But we also represent a strategy with the help of a propositional formula, by endowing the formula with a specific semantics. This makes strategy design much easier and efficient. Secondly, we view a strategy as a set of possible moves, i.e., a set of state-action pairs, rather than a function from a state (or a sequence of states) to an action. In this sense, our strategies represent “rough ideas”, which can then be combined and refined. Thirdly, instead of using PDL-style program connectives, we introduced two prioritised connectives for combining strategies, which, as we have seen, provides for a very natural and convenient design of strategies.

We want to mention that the idea of the prioritised disjunction was inspired by Brewka et al.’s [2] Qualitative Choice Logic (QCL). QCL contains a non-standard propositional connective $A \overrightarrow{\vee} B$ with the meaning, “*A if possible; but if A is impossible then at least B*”. We found that the semantics of the connective fits strategies very well. It is an interesting question for future work whether our prioritised conjunction can be integrated into QCL.

9 Conclusion

In this report we have introduced a logical language to describe, compose and combine strategies for game-playing agents. The language derives from the general Game Description Language (GDL) and extends it by a single temporal operator \bigcirc and two new prioritised connectives: ∇ and Δ . The basic components of GDL facilitate the representation of initial and terminal conditions, winning criteria and legality of actions (i.e., preconditions). The temporal operator allows us to describe the effects of actions. These form the basic language for describing game- and player-specific strategies. The newly introduced connectives allow us to combine simple strategies into more complicated and refined ones. When we use the language to describe a strategy, we endow it with a specific semantics so that we can compose a strategy in a logical way and the actual moves the strategy represents can be generated in an automatic way.

We have thoroughly analysed the properties of the language. In fact, the nice properties in particular of the new connectives give us great freedom in practice for strategy design: we can start with a strategy that formalises one specific idea. If it is too restricted, we can extend it with more generic ones using the prioritised disjunction, and if it is too generic, we can refine it with more specific strategies using the prioritised conjunction. We have shown also how strategies can be embedded into existing methods for the design and analysis of intelligent agents in order to solve problems that involve reasoning about strategies, including the computation of possible outcomes of games under given strategies.

Our current implementation for strategy reasoning combines reasoning about actions and change with an encoding of game rules and strategies using Answer Set Programming. It would be more efficient to develop a specific method of model checking based on the structure of our state transition model and the syntax of strategy composition. We leave this for future work.

Acknowledgements. This research was partially supported by the Australian Research Council through Discovery Project DP0988750. The second author is the recipient of an ARC Future Fellowship (project number FT 0991348). He is

also affiliated with the University of Western Sydney. We thank the anonymous reviewers of an earlier version of this report for their constructive comments and suggestions. We also thank Guifei Jiang for her valuable input.

References

- [1] Alur, R., T. A. Henzinger, and O. Kupferman: 2002, ‘Alternating-time temporal logic’. *J. ACM* **49**(5), 672–713.
- [2] Brewka, G., S. Benferhat, and D. L. Berre: 2004, ‘Qualitative choice logic’. *Artificial Intelligence* **157**(1-2), 203–237.
- [3] Chatterjee, K., T. A. Henzinger, and N. Piterman: 2007, ‘Strategy logic’. In: *CONCUR 2007–Concurrency Theory*. Springer, pp. 59–73.
- [4] Chatterjee, K., T. A. Henzinger, and N. Piterman: 2010, ‘Strategy logic’. *Information and Computation* **208**(6), 677 – 693.
- [5] Gebser, M., B. Kaufmann, R. Kaminski, M. Ostrowski, T. Schaub, and M. T. Schneider: 2011, ‘Potassco: The Potsdam Answer Set Solving Collection’. *AI Communications* **24**(2), 107–124.
- [6] Gelfond, M.: 2008, ‘Answer sets’. In: V. van Harmelen, F.; Lifschitz and B. Porter (eds.): *Handbook of Knowledge Representation*. Elsevier, pp. 285–31.
- [7] Genesereth, M., N. Love, and B. Pell: 2005, ‘General Game Playing: Overview of the AAAI Competition’. *AI Magazine* **26**(2), 62–72.
- [8] Goranko, V. and G. Van Drimmelen: 2006, ‘Complete axiomatization and decidability of alternating-time temporal logic’. *Theoretical Computer Science* **353**(1), 93–117.
- [9] Haufe, S. and M. Thielscher: 2012, ‘Automated Verification of Epistemic Properties for General Game Playing’. In: *Proceedings of the Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*. pp. 339–349, AAAI Press.
- [10] Lifschitz, V.: 2002, ‘Answer set programming and plan generation’. *Artificial Intelligence* **138**(1-2), 39–54.
- [11] Lorini, E.: 2010, ‘A Dynamic Logic of Agency II: Deterministic DLA, Coalition Logic, and Game Theory’. *Journal of Logic, Language and Information* **19**(3), 327–351.
- [12] Mogavero, F., A. Murano, and M. Y. Vardi: 2010, ‘Reasoning About Strategies’. In: *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2010*. pp. 133–144, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.
- [13] Pauly, M.: 2002, ‘A modal logic for coalitional power in games’. *Journal of logic and computation* **12**(1), 149–166.

- [14] Pauly, M. and R. Parikh: 2003, ‘Game Logic - An Overview’. *Studia Logica* **75**(2), 165–182.
- [15] Ramanujam, R. and S. E. Simon: 2008, ‘Dynamic Logic on Games with Structured Strategies.’. In: *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning (KR’08)*. pp. 49–58.
- [16] Reiter, R.: 1991, ‘The Frame Problem in the Situation Calculus: A Simple Solution (sometimes) and a Completeness Result for Goal Regression’. In: V. Lifschitz (ed.): *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*. Academic Press, pp. 359–380.
- [17] Reiter, R.: 2001, *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT Press.
- [18] Robson, J. M.: 1983, ‘The Complexity of Go’. In: *Proceedings of the IFIP 9th World Computer Congress on Information Processing*. pp. 413–417.
- [19] Schiffel, S. and M. Thielscher: 2011, ‘Reasoning About General Games Described in GDL-II’. In: *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI-11)*. pp. 846–851, AAAI Press.
- [20] Thielscher, M.: 2009, ‘Answer Set Programming for Single-Player Games in General Game Playing’. In: *Proceedings of the International Conference on Logic Programming (ICLP)*. pp. 327–341, Springer.
- [21] Thielscher, M.: 2011, ‘General Game Playing in AI Research and Education’. In: *Proceedings of the German Conference on Artificial Intelligence (KI)*. pp. 26–37, Springer.
- [22] van Benthem, J.: 2001, ‘Games in Dynamic-Epistemic Logic’. *Bulletin of Economic Research* **53**(4), 219–48.
- [23] van Benthem, J.: 2008, ‘In praise of strategies’. In: J. V. Eijck and R. Verbrugge (eds.): *Games, Actions, and Social Software*, ILLC scientific publications. Institute for Logic, Language and Computation (ILLC), University of Amsterdam.
- [24] van Benthem, J.: 2013, ‘Reasoning about Strategies’. In: *Computation, Logic, Games, and Quantum Foundations. The Many Facets of Samson Abramsky*. Springer, pp. 336–347.
- [25] van der Hoek, W., W. Jamroga, and M. Wooldridge: 2005, ‘A logic for strategic reasoning’. In: *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-05)*. pp. 157–164, ACM.
- [26] Walther, D., W. van der Hoek, and M. Wooldridge: 2007, ‘Alternating-time temporal logic with explicit strategies’. In: *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge*. New York, NY, USA, pp. 269–278, ACM.