

The Epistemic Logic Behind the Game Description Language

Ji Ruan and Michael Thielscher

School of Computer Science and Engineering
The University of New South Wales, Australia
{jiruan,mit}@cse.unsw.edu.au

Abstract

A *general* game player automatically learns to play arbitrary new games solely by being told their rules. For this purpose games are specified in the game description language GDL, a variant of Datalog with function symbols and a few known keywords. In its latest version GDL allows to describe nondeterministic games with any number of players who may have imperfect, asymmetric information. We analyse the epistemic structure and expressiveness of this language in terms of epistemic modal logic and present two main results: (1) The operational semantics of GDL entails that the situation at any stage of a game can be characterised by a multi-agent epistemic (i.e., S5-) model; (2) GDL is sufficiently expressive to model any situation that can be described by a (finite) multi-agent epistemic model.

Introduction

General game playing aims at building systems that automatically learn to play arbitrary new games solely by being told their rules (Pitrat 1971; Pell 1993). The Game Description Language (GDL) is a special-purpose, rule-based language for defining games (Love et al. 2006). GDL is used in the AAAI General Game Playing Competition, where participants are provided with previously unknown games specified in this language, and are required to dynamically and autonomously determine how best to play these games (Genesereth, Love, and Pell 2005). A recent extension to GDL allows to describe games with randomness and imperfect information (Thielscher 2010). This opens the door to nondeterministic games like Poker, in which players have imperfect, incomplete, and asymmetric information.

The game description language is a variant of Datalog with function symbols and a few known keywords. By applying a standard semantics for logic programs, a game description G can be interpreted by a state transition system. The execution model underlying GDL then induces a *game model* for G , which determines all possible ways in which the game may develop and what information the players acquire as the game proceeds (Love et al. 2006; Thielscher 2010). However, an open question has been to what extent this game model, including its implicit epistemic

structure due to imperfect and asymmetric information, satisfies standard properties of epistemic logic, and how expressive it is compared to this logic. The latter is particularly interesting because at first glance GDL seems to be constrained by the fact that all players have complete knowledge of the game rules and in particular the initial position.

In this paper we analyse the epistemic structure and expressiveness of GDL in terms of epistemic logic. Seminal work in this area is (Hintikka 1962), and since then many philosophers have been interested in further developing the notions of knowledge and belief using a possible world semantics. In the late 1980s these approaches were picked up and further developed by computer scientists, cf. (Halpern and Vardi 1986; Fagin et al. 1995). This development was originally motivated by the need to reason about communication protocols, where one is typically interested in what knowledge different parties to a protocol have before, during and after a run (i.e., an execution sequence) of the protocol. Apart from computer science, there is much interest in the dynamics of knowledge and belief in areas as diverse as artificial intelligence (Moore 1990), multi-agent systems (Rao and Georgeff 1991), and game theory (Aumann and Brandenburger 1995).

Here, we present, and formally prove, two main results:

1. The game model for any (syntactically valid) GDL game entails that at any round of the game the situation that arises can be characterised by a multi-agent S5-model.
2. Given an arbitrary (yet finite) epistemic model it is possible to construct a GDL game description which produces the situation described by this model.

This is complemented by an analysis of entailment of epistemic formulas in GDL and a study of how existing systems for Automated Theorem Proving and Epistemic Model Checking can be combined into a proof system to systematically verify epistemic properties of GDL descriptions.

The remainder of the paper proceeds as follows. The next section recapitulates both GDL and epistemic logic. The third section analyses the epistemic logic behind GDL and shows how the situations during a game can always be characterised by a standard epistemic model that entails the exact same formulas. The fourth section provides the construction of a GDL game for any given epistemic model. We conclude with a short discussion on model checking in our framework.

role(?r)	?r is a player
init(?f)	?f holds in the initial position
true(?f)	?f holds in the current position
legal(?r, ?m)	?r can do move ?m
does(?r, ?m)	player ?r does move ?m
next(?f)	?f holds in the next position
terminal	the current position is terminal
goal(?r, ?v)	goal value for role ?r is ?v
sees(?r, ?p)	?r perceives ?p in the next position
random	the random player

Table 1: GDL-II keywords: the top eight comprise standard GDL while the last two have been added in GDL-II.²

Preliminaries

Describing Games with GDL

General Game Playing requires a formal language for describing the rules of arbitrary games. A complete game description consists of the names of the players, a specification of the initial position, the legal moves and how they affect the position, along with the terminating and winning criteria. The emphasis of the game description language GDL is on high-level, declarative game rules that are easy to understand and maintain. At the same time, GDL has a precise semantics and is fully machine-processable. Moreover, background knowledge is not required—a set of rules is all a player needs to know in order to be able to play a hitherto unknown game.

A variant of Datalog with function symbols, the game description language uses a few known *keywords*; cf. Table 1. Original GDL is suitable for describing finite, synchronous, and deterministic n -player games with complete information about the game state (Love et al. 2006).³ The extended game description language GDL-II allows to specify arbitrary games with randomness and imperfect/incomplete information (Thielscher 2010; 2011). Valid game descriptions must satisfy certain syntactic restrictions, which ensure that all deductions “ \vdash ” used in Definition 1 below are finite and decidable; we refer to (Love et al. 2006) for details.

We need two abbreviations: Let $S = \{f_1, \dots, f_k\}$ be a *state*, that is, a finite set of ground terms (containing the position features that hold in S), then

$$S^{\text{true}} \stackrel{\text{def}}{=} \{\text{true}(f_1), \dots, \text{true}(f_k)\}$$

Furthermore, let $M = \langle m_1, \dots, m_n \rangle$ be a *joint move*, that is, a move (m_i) for each player (r_i), then

$$M^{\text{does}} \stackrel{\text{def}}{=} \{\text{does}(r_1, m_1), \dots, \text{does}(r_n, m_n)\}$$

Definition 1. (Thielscher 2010) Let G be a valid GDL-II specification whose signature determines the set of ground terms Σ . Let 2^Σ be the set of finite subsets of Σ . The semantics of G is given by the following state transition system.

²The keywords are accompanied by the auxiliary, pre-defined predicate *distinct* (X, Y), meaning the syntactic inequality of the two arguments (Love et al. 2006).

³Synchronous means that all players move simultaneously. In this setting, turn-taking games are modelled by allowing players only one legal move, without effect, if it is not their turn.

- $R = \{r \in \Sigma : G \vdash \text{role}(r)\}$ (the roles);
- $s_0 = \{f \in \Sigma : G \vdash \text{init}(f)\}$ (the initial position);
- $t = \{S \in 2^\Sigma : G \cup S^{\text{true}} \vdash \text{terminal}\}$ (the terminal positions);
- $l = \{(r, m, S) : G \cup S^{\text{true}} \vdash \text{legal}(r, m)\}$, for all $r \in R$, $m \in \Sigma$, and $S \in 2^\Sigma$ (the legal moves);
- $u(M, S) = \{f \in \Sigma : G \cup M^{\text{does}} \cup S^{\text{true}} \vdash \text{next}(f)\}$, for all joint moves M and $S \in 2^\Sigma$ (the update function);
- $\mathcal{I} = \{(r, M, S, p) : G \cup M^{\text{does}} \cup S^{\text{true}} \vdash \text{sees}(r, p)\}$, for all $r \in R \setminus \{\text{random}\}$, $M \in \Sigma^{|R|-1}$, $S \in 2^\Sigma$, and $p \in \Sigma$ (the information relation, determining the players’ percepts);
- $g = \{(r, v, S) : G \cup S^{\text{true}} \vdash \text{goal}(r, v)\}$, for all $r \in R \setminus \{\text{random}\}$, $v \in \{0, \dots, 100\}$ and $S \in 2^\Sigma$ (the goal relation).

Different runs of a game can be described by *developments*, which are sequences of states and moves by each player up to a certain round, and a player *cannot distinguish* two developments if he takes the same moves and gets the same percepts in the two (Thielscher 2010).

Definition 2. Let $\langle R, s_0, t, l, u, \mathcal{I}, g \rangle$ be the semantics of a GDL-II description G according to Definition 1, then a development δ is a sequence

$$\langle s_0, M_1, s_1, \dots, s_{d-1}, M_d, s_d \rangle$$

where $d \geq 0$ and for all $i \in \{1, \dots, d\}$, M_i is a joint move such that each move is legal in s_{i-1} , and states are updated thus: $s_i = u(M_i, s_{i-1})$. The length of a development δ , denoted as $\text{len}(\delta)$, is the number of states in δ . By $M_i(j)$ we denote agent j ’s move in the joint move M_i .

Consider two developments $\delta = \langle s_0, M_1, s_1, \dots \rangle$ and $\delta' = \langle s_0, M'_1, s'_1, \dots \rangle$. The player $j \in R \setminus \{\text{random}\}$ cannot distinguish δ from δ' (written as $\delta \sim_j \delta'$) if, and only if,

1. $\text{len}(\delta) = \text{len}(\delta')$ and
2. for all $i \in \{1, \dots, \text{len}(\delta) - 1\}$:
 - (a) $\{p : (j, M_i, s_{i-1}, p) \in \mathcal{I}\} = \{p : (j, M'_i, s'_{i-1}, p) \in \mathcal{I}\}$
 - (b) $M_i(j) = M'_i(j)$.

Modal Epistemic Logic

In order to analyse the epistemic logic behind GDL-II and its semantics, we recapitulate basic notions from standard Modal Epistemic Logic (Fagin et al. 1995).

Definition 3. (Language) A basic Modal Epistemic Logic Language for epistemic formulas is given by the following Backus-Naur Form:

$$\phi := P \mid \neg\phi \mid \phi \wedge \psi \mid K_i\phi \mid C_B\phi$$

where P is an atomic proposition, i an agent, and B a non-empty set of agents. $\top, \perp, \vee, \rightarrow$ are defined as usual.

Intuitively, $K_i\phi$ means agent i knows ϕ , and $C_B\phi$ means that ϕ is common knowledge among the agents in B ; for example, “agent k knows that agent j knows P ” can be expressed as K_iK_jP . To give precise meanings to this language, we need multi-agent epistemic models.

```

1 role(generalA). role(generalB). role(random).
2 succ(0,1). succ(1,2). ... succ(8,9).
3 time(3am). time(9pm).
4 init(round(0)).
5 gets_message(?g,?m) <= role(?g), distinct(?g,random),
6                               does(?g1,send(?m)),
7                               does(random,pass), distinct(?g,?g1).
8 gets_new_message(?g) <= gets_message(?g,?m).
9 has_a_message(?g) <= true(message(?g,?m)).
10 legal(random,noop) <= true(round(0)).
11 legal(random,pass) <= not true(round(0)).
12 legal(random,stop) <= not true(round(0)).
13 legal(generalA,settime(?t)) <= true(round(0)), time(?t).
14 legal(generalB,noop) <= true(round(0)).
15 legal(generalB,noop) <= true(control(generalB)),
16                               not has_a_message(generalB).
17 legal(generalA,noop) <= true(control(generalB)).

18 legal(generalB,noop) <= true(control(generalA)).
19 legal(?g,send(ack(?m))) <= true(control(?g)),
20                               true(message(?g,?m)).
21 sees(?g,?m) <= gets_message(?g,?m).
22 next(message(generalA,?t))
23                               <= does(generalA,settime(?t)).
24 next(attack_time(?t)) <= does(generalA,settime(?t)).
25 next(attack_time(?t)) <= true(attack_time(?t)).
26 next(message(?g,?m)) <= gets_message(?g,?m).
27 next(message(?g,?m)) <= true(message(?g,?m)),
28                               not gets_new_message(?g).
29 next(control(generalA)) <= true(round(0)).
30 next(control(generalA)) <= true(control(generalB)).
31 next(control(generalB)) <= true(control(generalA)).
32 next(round(?n)) <= true(round(?m)), succ(?m,?n).
33
34 terminal <= true(round(9)).

```

Figure 1: A GDL-II description of the Two Generals' Coordinated Attack Game: G_{ca} .

Definition 4. A multi-agent epistemic model E is a structure $\langle W, \{\sim_i: i \in Ag\}, V \rangle$, where W is a set of possible worlds, Ag is a set of agents, each $\sim_i \subseteq W \times W$ is an equivalence relation (called the accessibility relation)⁴ for agent i , and $V: W \mapsto 2^{Atoms}$ is a valuation function that assigns each world a set of atomic propositions.

Definition 5. Given an epistemic model E and an epistemic formula ϕ , the entailment relation \models is defined as follows:

- $E, w \models P$ iff $P \in V(w)$;
- $E, w \models \neg\phi$ iff $E, w \not\models \phi$;
- $E, w \models \phi \wedge \psi$ iff $E, w \models \phi$ and $E, w \models \psi$;
- $E, w \models K_i\phi$ iff for all w' , if $w \sim_i w'$ then $E, w' \models \phi$;
- $E, w \models C_B\phi$ iff for all w' , if $w \sim_B w'$ then $E, w' \models \phi$.

where \sim_B is the transitive and reflexive closure of $\cup_{i \in B} \sim_i$.

Common knowledge $C_B\phi$ is an appropriate concept when it comes to analysing the knowledge of players in GDL-II games, in particular since the game rules themselves are always common knowledge.

From GDL-II to Epistemic Models

Before we present our results in all technical detail, we introduce a running example adopted from (Fagin et al. 1995).

Example 1. (Coordinated Attack Problem) A valley separates two hills. Two armies, each on its own hill and led by General A and B, respectively, are preparing to attack their common enemy in the valley. The two generals must have their armies attack the valley at the same time in order to succeed. The only way for the two generals to communicate is by sending messengers through the valley. Unfortunately, there is a chance that any given messenger sent through the

⁴Note that in general the accessibility relation \sim_i does not have to be an equivalence relation, but since epistemic models in this paper are meant to represent the knowledge of agents, rather than belief or other modalities, we restricted \sim_i to be an equivalence relation; hence all epistemic models are *S5-models* (Fagin et al. 1995).

valley will be stopped by the enemy, in which case the message is lost but the content is not leaked. The problem is to come up with algorithms that the generals can use, including sending messages and processing received messages, that can allow them to correctly agree upon a time to attack.

It was proved that such a coordinated attack is impossible (Fagin et al. 1995). We show that GDL-II is flexible enough to specify problems like this that involve complex epistemic situations and arguments. Specifically, we will use the game semantics to show why a coordinated attack is not possible.

Figure 1 describes a version of this problem in GDL-II. Generals A and B are modelled as two roles, and the enemy is modelled by the ‘random’ role. For the sake of simplicity, General A starts by choosing between just two possible attack times: ‘3am’ or ‘9pm’, and then sends his choice as a message to B. Subsequently, each general takes control in turn, and if he receives a message m then he sends an acknowledgement $ack(m)$ back to the other general, otherwise he does *noop*; simultaneously, the ‘random’ role always chooses randomly either *pass*, which allows the message to go to the other general, or *stop*, which intercepts the message. For the sake of simplicity, we assume that the game terminates at round 9 and ignore the goal values.

The semantics of a game description according to Definition 1 determines a state transition system. We can then use the operational semantics implicit in Definition 2 to introduce the special concept of *epistemic game models* for GDL-II.

Definition 6. (GDL-II Epistemic Game Model) Consider a GDL-II description G with semantics $\langle R, s_0, t, l, u, \mathcal{I}, g \rangle$. The epistemic game model of G , denoted by $E(G)$, is a structure $\langle W, Ag, \{\sim_i: i \in Ag\}, V \rangle$ where

- W is the set of developments of G ;
- Ag is the set of roles $R \setminus \{random\}$;
- $\sim_i \subseteq W \times W$ is the accessibility relation for agent $i \in Ag$ given by $(\delta, \delta') \in \sim_i$ (also written as $\delta \sim_i \delta'$) iff role i cannot distinguish development δ from δ' ;

- $V : W \rightarrow 2^\Sigma$ is an interpretation function which associates with each development δ the set of ground terms in Σ that are true in the last state of δ .

For example, from the game description G_{ca} for the Coordinated Attack Problem, we derive a game model $E(G_{ca})$, partly depicted in Figure 2, in two steps. The first step is to use the game semantics for GDL-II to determine all states that are reachable from the initial state. A joint move is depicted as (a, b, c) , where a, b, c are the moves of, respectively, General A, General B, and ‘random’. For instance, there are two possible joint moves at s_0 , $M_1 = (\text{settime}(3am), \text{noop}, \text{noop})$ and $M_2 = (\text{settime}(9pm), \text{noop}, \text{noop})$, which transit s_0 to s_1 and s_2 , respectively. From s_1 there are again two possible joint moves which result, respectively, in s_{11} where B receives A’s message, and in s_{12} where B receives nothing. Accordingly at state s_{11} , it is legal for B to send an acknowledgement, after which s_{11} transits to either of two possible states, s_{111} and s_{112} . This process is reiterated until a terminal state is reached.

The second step in constructing the game model for G_{ca} is to collect all developments and to determine the individual accessibility relations. For example, consider the two developments $\delta_1 = \langle s_0, M_1, s_1 \rangle$ and $\delta_2 = \langle s_0, M_2, s_2 \rangle$. It is easy to check that $\delta_1 \not\sim_A \delta_2$ since General A moves differently in M_1 and M_2 . On the other hand, $\delta_1 \sim_B \delta_2$ since General B takes the same move in M_1 and M_2 and perceives nothing in both cases.

Based on our concept of an epistemic game model for GDL-II, we can define how to interpret formulas in the basic epistemic language over such models in a similar fashion as Definition 5.

Definition 7. Given an epistemic game model $E(G)$, a development δ , and an epistemic formula ϕ , the entailment relation \models is defined as follows:

- $E(G), \delta \models P$ iff $P \in V(\text{last}(\delta))$;
- $E(G), \delta \models \neg\phi$ iff $E(G), \delta \not\models \phi$;
- $E(G), \delta \models \phi \wedge \psi$ iff $E(G), \delta \models \phi$ and $E(G), \delta \models \psi$;
- $E(G), \delta \models K_i\phi$ iff for all δ' , if $\delta \sim_i \delta'$ then $E(G), \delta' \models \phi$;
- $E(G), \delta \models C_B\phi$ iff for all δ' , if $\delta \sim_B \delta'$ then $E(G), \delta' \models \phi$.

where $\text{last}(\delta)$ is the last state of development δ , and \sim_B is the transitive and reflexive closure of $\cup_{i \in B} \sim_i$.

Coming back to our running example, a simple and elegant argument can be given now on why a coordinated attack is never possible. First, using the epistemic language of Definition 3 we can express knowledge conditions such as:

- K_AP for “General A knows that P ,” where P is an atomic expression, e.g. $\text{attack_time}(3am)$;
- $\neg K_B K_AP$ for “General B does not know if General A knows P ,”
- $C_{\{A,B\}}P$ for “ P is common knowledge to both generals.”

Let $P = \text{attack_time}(3am)$ and $\delta_1, \delta_{11}, \delta_{111}$ be the left-most developments with length 1, 2, and 3 in Figure 2, then we can verify each of the following: $E(G_{ca}), \delta_1 \models K_AP \wedge \neg K_B K_AP$; $E(G_{ca}), \delta_{11} \models K_B K_AP \wedge \neg K_A K_B K_AP$; and

$E(G_{ca}), \delta_{111} \models K_A K_B K_A P \wedge \neg K_B K_A K_B K_A P$. This implies that $E(G_{ca}), \delta \models \neg C_{\{A,B\}}P$ for each development $\delta = \delta_1, \delta_{11}, \delta_{111}$. That is to say, the attack time is not common knowledge among A and B even after the successful delivery of all messages during three rounds. We can generalise this to developments of arbitrary length. Given that this common knowledge is a prerequisite for a coordinated attack, it follows that the latter can never be accomplished.

In general, it is easy to show that the epistemic game model we constructed for GDL-II is equivalent to the standard concept of models and entailment in Epistemic Modal Logic. Specifically, we can pick an arbitrary game at any round and build an epistemic model for this situation such that the truth of epistemic formulas is preserved.

Proposition 1. Given an arbitrary GDL-II description G and any round of playing $k \geq 0$ (with round 0 corresponding to the initial state), we can derive a finite epistemic model E such that this round of the game is characterised by E , which is to say, $E, \delta \models \phi$ if, and only if, $E(G), \delta \models \phi$.

Proof. Let $E(G) = \langle W, Ag, \{\sim_i : i \in Ag\}, V \rangle$ be constructed from G according to Definition 6, and assume that game play is at round k . Based on $E(G)$, we construct a finite epistemic model $E = \langle W', \{\sim_i : i \in Ag'\}, V' \rangle$ as follows:

1. W' is the set of any game development $\delta \in W$ with $\text{len}(\delta) = k + 1$;
2. Ag' is the same set of agents as Ag ;
3. \sim'_i is the equivalence relation \sim_i restricted on the new domain W' , i.e., $\sim'_i = \sim_i \cap (W' \times W')$;
4. V is a valuation function with $P \in V'(\delta)$ iff $P \in V(\delta)$ for any atomic proposition P and $\delta \in W'$.

We show by induction on the structure of formula ϕ that for all $\delta \in W'$: $E, \delta \models \phi$ iff $E(G), \delta \models \phi$. The propositional cases follow from the fact that the valuation does not change. For the case of $\phi := K_i\psi$, by definition, we have that $E, \delta \models K_i\psi$ iff for all δ' , if $\delta \sim'_i \delta'$ then $E, \delta' \models \psi$. If two developments δ, δ' have different lengths, then any agent can distinguish them, so if $\delta \sim'_i \delta'$, then $\text{len}(\delta) = \text{len}(\delta') = k + 1$, which means that $\delta' \in W'$ as well. So by induction, for all δ' , if $\delta \sim'_i \delta'$ then $E, \delta' \models \phi$ iff for all δ' , if $\delta \sim_i \delta'$ then $E(G), \delta' \models \phi$; therefore $E, \delta \models K_i\psi$ iff $E(G), \delta \models K_i\psi$. For the case of $\phi := C_B\psi$, the reasoning is similar since the developments in the transitive and reflexive closure of $\cup_{i \in B} \sim_i$ also are all of the same length $k + 1$. \square

As a corollary we can show, say, that the round number is a common knowledge for all agents; in our example game:

$$E(G_{ca}), \delta \models \bigwedge_k (\text{round}(k) \rightarrow C_{\{A,B\}}\text{round}(k)).$$

From Epistemic Models to GDL-II

We now look at the other direction and show that for any given finite multi-agent epistemic model E we can construct a valid GDL-II description for a game with a development that leads to E . As a matter of fact, a (very abstract) game can always be constructed where a single move suffices to bring about an arbitrary given epistemic model.

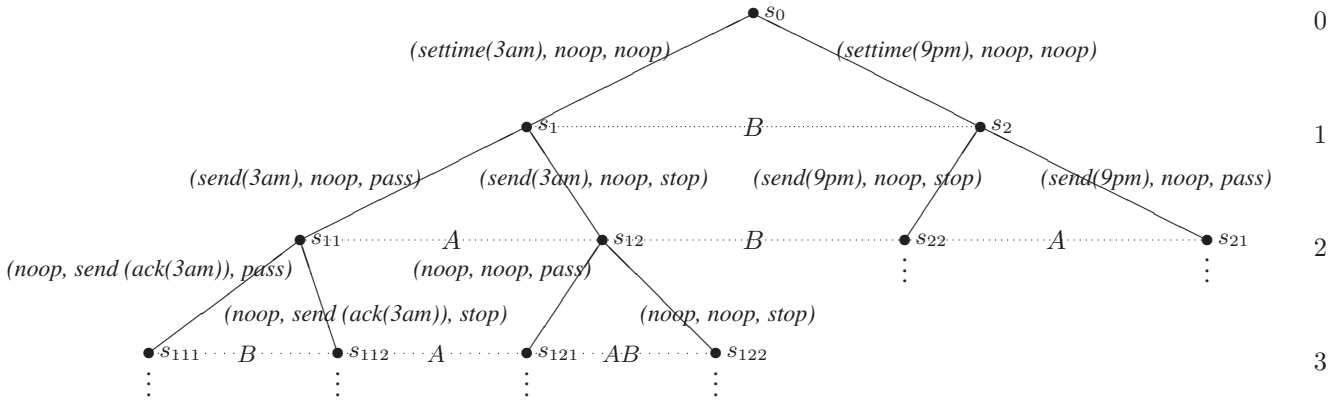


Figure 2: Epistemic game model $E(G_{ca})$ for the Two Generals' Coordinated Attack Game.

Theorem 1. For an arbitrary finite multi-agent epistemic model $E = \langle W, \{\sim_i: i \in Ag\}, V \rangle$ a GDL-II game description G can be constructed such that E can emerge after one step of play in G , which is to say, E is isomorphic to a sub-model of $E(G)$ for the situation after the first move.

Proof. Let $W = \{w_1, \dots, w_k\}$ and $Ag = \{1, \dots, n\}$, then game G can be constructed as follows:

```

1 role(1). ... role(n). role(random).
2
3 world(w1). ... world(wk).
4 legal(random, select(?w)) <= world(?w).
5 legal(?r, noop) <= role(?r), distinct(?r, random).
6
7 val(w1, P1). ... val(wk, Pm).
8 next(?P) <= does(random, select(?w)), val(?w, ?P).
9
10 equiv(1, wa, wa). equiv(1, wa, wb). ... equiv(n, wx, wy).
11 sees(?r, class(?w2)) <= does(random, select(?w1)),
12    equiv(?r, ?w1, ?w2).

```

The game has $n + 1$ roles, namely, the n agents plus the standard ‘random’ role (line 1). Initially, ‘random’ has a legal move $select(w)$ for any world $w \in W$ (line 3–4) while all other players can only do $noop$. The move $select(w)$ results in a state in which all atomic propositions hold that are true in world w (line 8). This rule uses an explicit enumeration of all pairs (w, P) such that $P \in V(w)$ (line 7). Furthermore, in order to arrive at the desired epistemic structure, the players get to see *all* worlds in their equivalence class $\{w' : (w, w') \in \sim_i\}$ (line 11–12). This rule uses an enumeration of all triples (i, w_a, w_b) with $w_a \sim_i w_b$ (line 10).⁵

We show that G indeed gives E according to the semantics in Definition 1. The initial state is $s_0 = \{\}$, and then $G \cup s_0^{\text{true}}$ entails $legal(random, select(w_j))$ for all $j \in [1..k]$, and $legal(1, noop), \dots, legal(n, noop)$. In other words, each agent in Ag can only do $noop$, while ‘random’ may select an arbitrary world from E . Define joint move $M^j := (noop, \dots, noop, select(w_j))$ and consider, then, successor states $s_x = u(M^x, s_0)$ and $s_y = u(M^y, s_0)$, corresponding to the two developments $\delta_x = \langle s_0, M^x, s_x \rangle$

⁵We omit definitions for terminal and goal.

and $\delta_y = \langle s_0, M^y, s_y \rangle$, respectively. If $w_x \sim_i w_y$, then agent i gets to see both $class(w_x)$ and $class(w_y)$ in both states s_x and s_y , in which case the agent cannot distinguish δ_x from δ_y because also his action is the same in both M_1 and M_2 . On the other hand, if $w_x \not\sim_i w_y$ then agent i can distinguish the two developments based on his percepts. Altogether this process gives us an epistemic game model $E(G)$, from which we obtain a standard epistemic model $E' = \langle W', \{\sim'_i: i \in Ag\}, V' \rangle$ as follows: W' is the set of all developments of length 2 from $E(G)$, while \sim'_i and V' are restrictions of, respectively, \sim_i and V on W' .

Now E and E' are isomorphic: Each world $w_j \in W$ corresponds to the state $s_j = u(M^j, s_0)$ and hence to the development $\delta_j \in W'$. Moreover, $(w_x, w_y) \in \sim_i$ iff $(\delta_x, \delta_y) \in \sim'_i$ for agent i . Finally, for all atomic proposition P we have that $P \in V(w_j)$ iff $P \in V'(\delta_j)$. \square

Model Checking

With an epistemic framework for GDL-II, we are now able to reason about epistemic properties of games. For example, given that agents may have only partial observation ability, it is easy to construct games in which agents do not have sufficient information to derive their legal moves; this may render a game unfair or even unplayable. A desirable property of a game may be to avoid such a situation, and we can use the epistemic structure to verify that a GDL-II description obeys this property. We express this property in the basic epistemic language by this formula for agent i :

$$\phi_i = \bigwedge_m (legal(i, m) \rightarrow K_i legal(i, m))$$

To check such properties systematically amounts to a model checking problem: given a GDL-II description G , a round number k , and an epistemic formula ϕ , verify that $E(G), \delta \models \phi$ for all δ with $len(\delta) \leq k$. In the following we sketch a model checking method using two existing tools, the Answer Set Programming system POTASSCO (Gebser et al. 2011) and the epistemic model checker DEMO (van Eijck 2007).

The first step is to compute all reachable states in game G up to round k . For this, we utilise the method in (Schif-

fel and Thielscher 2009): based on GDL-II description G , we can first derive negation-free clauses \mathcal{D} that define the domains of the features (predicate $fdom$) and moves (predicate $mdom$) in G ; e.g.,

```
1 fdom(round(0)). fdom(control(generalA)). ...
2 mdom(noop). mdom(settime(3am)). ...
```

From this we can construct an Answer Set Program with a one-to-one correspondence between its answer sets and the developments of given length k . These answer sets also indicate what each player gets to see after each move. Using all these reachable states and moves, we then derive an epistemic model E that characterises the game in round k according to Proposition 1.

The second step is to specify the epistemic model E and epistemic formula ϕ in the model checker DEMO. In particular, an epistemic model is represented as

```
Mo [state] [(state, formula)] [(Agent, state, state)]
```

where Mo is an indicator of the model; $[state]$ is a list of states (here, ‘developments’); $[(state, formula)]$ is a valuation function with $formula$ being atomic; and finally $[(Agent, state, state)]$ represents the accessibility relations of the agents. A formula is encoded as

```
Form = Top | Prop Prop | Neg Form | Conj [Form]
      | Disj [Form] | K Agent Form | CK [Agent] Form
```

The syntax of DEMO does not allow for implications, but these can be encoded equivalently using the conjunctive normal form. For example, the implication $legal(i, m) \rightarrow K_i legal(i, m)$ is logically equivalent to the disjunction $\neg legal(i, m) \vee K_i legal(i, m)$, which can then be encoded as $Disj [Neg legal(i, m), K i legal(i, m)]$.

Conclusion

We analysed the epistemic structure and expressiveness of GDL-II in terms of epistemic modal logic and presented two results: (1) The operational semantics of GDL entails that the situation in any round of a game can be characterised by a multi-agent epistemic model; (2) GDL is sufficiently expressive to model any situation that can be described by a (finite) multi-agent epistemic model. We also sketched a model checking method for systematically verifying whether a given GDL-II description satisfies important epistemic properties. Further investigations will include the computational complexity of this model checking problem, and an extension of the basic epistemic language.

In an accompanying paper we show how GDL-II can be formally translated into the Situation Calculus as a first-order axiomatisation that allows players to reason about their percepts and what they know about the legality and effects of moves based on the game description (Schiffel and Thielscher 2011). Other related work describes the use of model checking to verify properties of general games (Ruan, van der Hoek, and Wooldridge 2009), but this is restricted to original GDL and hence to games where players can maintain complete state information. There is of course a large body of work on the epistemic structure of imperfect-information games, but ours is the first application of this

line of research to formally analyse the epistemic structure behind the general Game Description Language GDL-II.

Acknowledgements. We thank Jérôme Lang for suggesting the example for the paper. The second author is the recipient of an Australian Research Council Future Fellowship (project FT 0991348).

References

- Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.
- Aumann, R., and Brandenburger, A. 1995. Epistemic conditions for nash equilibrium. *Econometrica* 63:1161–1180.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning About Knowledge*. MIT Press.
- Gebser, M.; Kaminski, R.; Kaufmann, B.; Ostrowski, M.; Schaub, T.; Schneider, M. 2011. Potassco: The Potsdam answer set solving collection. *AI Communications* 24. (System available at potassco.sourceforge.net).
- Genesereth, M.; Love, N.; and Pell, B. 2005. General game playing: Overview of the AAI competition. *AI Magazine* 26(2):62–72.
- Halpern, J., and Vardi, M. 1986. The complexity of reasoning about knowledge and time. In *Proc. 18th ACM Symposium on Theory of Computing*, 304–315.
- Hintikka, J. 1962. *Knowledge and Belief*. Cornell University Press.
- Love, N.; Hinrichs, T.; Haley, D.; Schkufza, E.; and Genesereth, M. 2006. General Game Playing: Game Description Language Specification. Tech. Report LG-2006-01, Stanford University. (Available at games.stanford.edu).
- Moore, R. 1990. A formal theory of knowledge and action. In *Readings in Planning*. Morgan Kaufmann. 480–519.
- Pell, B. 1993. *Strategy Generation and Evaluation for Meta-Game Playing*. Ph.D. Dissertation, Cambridge.
- Pitrat, J. 1971. A general game playing program. In Findler, N., and Meltzer, B., eds., *Artificial Intelligence and Heuristic Programming*, 125–155.
- Rao, A., and Georgeff, M. 1991. Modeling rational agents within a BDI-architecture. In *Proc. KR*, 473–484.
- Ruan, J.; van der Hoek, W.; and Wooldridge, M. 2009. Verification of games in the game description language. *Journal of Logic and Computation* 19(6):1127–1156.
- Schiffel, S., and Thielscher, M. 2011. Reasoning about general games described in GDL-II, 2011. In *Proc. AAI (this volume)*.
- Schiffel, S., and Thielscher, M. 2009. Automated theorem proving for general game playing. In *Proc. IJCAI*, 911–916.
- Thielscher, M. 2010. A general game description language for incomplete information games. In *Proc. AAI*, 994–999.
- Thielscher, M. 2011. The general game playing description language is universal. In *Proc. IJCAI*.
- van Eijck, J. 2007. DEMO—a demo of epistemic modelling. In *Interactive Logic*. Texts in Logic and Games 1, 305–363.