

SPECIFYING MULTIAGENT ENVIRONMENTS IN THE GAME DESCRIPTION LANGUAGE

Stephan Schiffel and Michael Thielscher

Department of Computer Science, Dresden University of Technology, 01062 Dresden, Germany

Keywords: Knowledge Representation, Multiagent Systems, Game Playing.

Abstract: The Game Description Language (GDL) has been developed for the purpose of formalizing game rules. It serves as the input language for general game players, which are systems that learn to play previously unknown games without human intervention. In this paper, we show that GDL can be readily used as a specification language for a large class of multiagent environments. The resulting specifications are declarative, compact, and easy to understand and maintain. At the same time they can be fully automatically understood and used by autonomous agents who intend to participate in these environments. Our main result is a formal characterization of the class of multiagent environments that can be described in GDL.

1 INTRODUCTION

A novel and challenging research problem for Artificial Intelligence, General Game Playing is concerned with the development of systems that learn to play a previously unknown game solely on the basis of the rules. The Game Description Language (GDL) (Love et al., 2006) has been developed to formalize the rules of any finite, information-symmetric n -player game in such a way that the description can be automatically processed by a general game player (Genesereth et al., 2006). As a declarative language, GDL supports specifications that are modular and easy to develop, understand, and maintain. While the basic semantics for GDL is grounded in standard logic, the language uses several pre-defined predicates as keywords, whose intended meaning is only informally described in (Love et al., 2006).

In this paper, we show that GDL can be readily used as a specification language for a large class of multiagent environments. This allows for formalizing the physics and laws that govern an arbitrary domain in such a way that agents can automatically understand the rules and thus know how to participate in this environment. There is a variety of potential applications for machine processable descriptions of multiagent environments: the rules of an e-marketplace can be made accessible to agents, the interface of interactive Internet platforms for software agents can be formally described, and agent competitions can be run without revealing detailed problem specifications

in advance. In each of these cases, an autonomous agent—or a team of agents—can learn how to participate in a new or modified environment without the need to be (re-)programmed for each specific case. Because GDL uses a decidable subset of logic programming, autonomous agents require just a simple, standard reasoning module to be able to understand and effectively process a given set of rules. Moreover, if an agent environment is specified in GDL, successful general game playing systems such as (Kuhlmann et al., 2006; Clune, 2007; Schiffel and Thielscher, 2007; Finnsson and Björnsson, 2008) can be readily employed as intelligent agents for these environments.

The main result in this paper is the definition of a formal class of multiagent environments which can be expressed in GDL and, conversely, which can be used to provide a semantics for any GDL game description. As a by-product we thus obtain a formal semantics for the special, pre-defined keywords in GDL.

The rest of the paper is organized as follows. In Section 2, we formally define the class of deterministic, synchronous multiagent environments. In Section 3, we show how these can be axiomatically described in GDL, and in the section that follows, we present the converse result by showing how all GDL games can be interpreted as a deterministic, asynchronous multiagent environment. We conclude in Section 5.




Figure 1: A simple multiagent domain: two “guard” agents Ag_1 and Ag_2 shall cooperatively try to catch Ag_3 , whose goal in turn is to escape via one of the three exits at locations $(1,1)$, $(5,1)$, and $(1,5)$. All agents act synchronously and can move horizontally or vertically to an adjacent position. Ag_3 is caught when it ends up in the same location as Ag_1 or Ag_2 , or when it crosses path with one of them in a simultaneous move.

2 MULTIAGENT ENVIRONMENTS

As the running example in this paper, we will consider the multiagent domain depicted in Figure 1. As a discrete environment it can be formally described as a finite state transition system. However, even though it is obviously just a toy-size example, it has a considerably large state space, rendering an explicit encoding difficult—and practically impossible for even slightly larger environments. Fortunately it is possible to exploit the fact that any natural and realistic multiagent environment has an internal structure, which allows one to describe its dynamics with the help of symbols that represent individual components. Our example domain, for instance, can be formally described using the following symbolic expressions: Ag_1, Ag_2, Ag_3 , representing the three agents; $At(r, x, y)$, where $r \in \{Ag_1, Ag_2, Ag_3\}$ and $x, y \in \{1, \dots, 5\}$, representing the position of each agent; and $Move(d)$, where $d \in \{North, South, East, West\}$, along with $Stay$ and $Exit$, representing the possible actions.

Based on a suitable collection of symbols, multiagent domains can be formally described as follows.

Definition 1 Let Σ be a countable set of ground (i.e., variable-free) symbolic expressions. A (discrete, synchronous, deterministic) *multiagent environment* is a structure

$$(R, s_1, t, l, u, g)$$

where

- $R \subseteq \Sigma$ finite (the agents, or *roles*);

- $s_1 \subseteq \Sigma$ finite (the *initial state*);
- $t \subseteq 2^\Sigma$ (the *terminal states*);
- $l \subseteq R \times \Sigma \times 2^\Sigma$ (the action preconditions, or *legality relation*);
- $u : (R \mapsto \Sigma) \times 2^\Sigma \mapsto 2^\Sigma$ (the transition function, or *update function*);
- $g \subseteq R \times \mathbb{N} \times 2^\Sigma$ (the utility, or *goal relation*).

Here, 2^Σ denotes the set of all *finite* subsets of Σ , and for any $r \in R$ and $S \in 2^\Sigma$, $l(r, a, S)$ holds for finitely many $a \in \Sigma$. \square

This definition deserves some explanation. For the sake of simplicity, the symbolic expressions are not categorized—there is no formal distinction between symbols for objects, state components, actions, etc. For practical purposes, it is important that states are finitely representable; hence, while possibly infinitely many symbols give rise to infinitely many states, a state itself is an element of the set of all *finite* subsets of the given symbols. The legality relation $l(r, a, S)$ defines a to be a legal action for agent r in state S . Again for the sake of practical usability, it is assumed that every agent in every state has only finitely many possible actions. The update function takes an action for each agent and (synchronously) applies the joint actions to a current state, resulting in the updated state. For the sake of simplicity, we take natural numbers $n \in \mathbb{N}$ as the utility of a state S for agent r in the goal relation $g(r, n, S)$.

For illustration, consider a formalization of the multiagent environment of Figure 1 using the symbols introduced above.

- $R = \{Ag_1, Ag_2, Ag_3\}$;
- $s_1 = \{At(Ag_1, 1, 1), At(Ag_2, 5, 1), At(Ag_3, 5, 5)\}$;
- t contains all states

$$\{At(Ag_1, x_1, y_1), At(Ag_2, x_2, y_2)\}$$

(that is, where Ag_3 has escaped) along with all states

$$\{At(Ag_1, x_1, y_1), At(Ag_2, x_2, y_2), At(Ag_3, x_3, y_3)\}$$

in which $x_1 = x_3 \wedge y_1 = y_3$ or $x_2 = x_3 \wedge y_2 = y_3$ (that is, where Ag_3 has been caught);

- l is defined as follows: each agent can always *Stay*; in every non-terminal state each agent can *Move(d)* in any direction d unless this would lead outside the physical environment; Ag_3 can *Exit* from any of the locations $(1,1)$, $(5,1)$, or $(1,5)$, provided it has not been caught.
- u is defined as follows: actions *Stay*, *Exit*, and *Move(d)* have the expected effects on the individual locations of the agents, with the exception

that when the paths of Ag_3 and either of Ag_1 or Ag_2 (or both) cross in a simultaneous move, then Ag_3 ends up (caught) in the same location as Ag_1 or Ag_2 , respectively. For illegal actions or states that are not reachable, u may be arbitrarily defined.

- g shall be defined as true for $n = 100$ and $r = Ag_3$ in terminal states in which this agent has escaped; conversely, g holds for $n = 100$ and both Ag_1 and Ag_2 in terminal states in which Ag_3 got caught. In all other states the goal relation gives value 0 for all three agents.

We have thus obtained a formal, symbolic description of the example multiagent environment. However, this specification is not yet amenable to automatic processing by an autonomous agent, because it uses natural language to describe some of the components. If this were to be translated into an explicit enumeration of the transition function, this would again yield too large a description to be of any practical use. In the following section, we show how the Game Description Language can be readily used to provide a fully axiomatic, compact description of arbitrary multiagent environments; a description that on the one hand is declarative and easy to understand and maintain by humans, and on the other hand can be fully automatically processed by artificial, autonomous agent systems.

3 AXIOMATIZING MULTIAGENT ENVIRONMENTS AS GAME DESCRIPTIONS

The Game Description Language (GDL) has been developed to formalize the rules of any finite game with complete information in such a way that the description can be automatically processed by a general game player. In this section, we first recapitulate the GDL syntax from (Love et al., 2006) and then show how the multiagent environments defined in the preceding section can be formally described in this language.

3.1 General GDL Syntax

GDL is based on the standard syntax of logic programs, including negation. A logic program is a set of clauses according to the following definition (see, e.g., (Lloyd, 1987)).

Definition 2

- A *term* is either a variable, or a function symbol applied to terms as arguments (a *constant* is a function symbol with no argument);
- An *atom* is a predicate symbol applied to terms as arguments;
- A *literal* is an atom or its negation;
- A *clause* is an implication $h \Leftarrow b_1 \wedge \dots \wedge b_n$ where *head* h is an atom and *body* $b_1 \wedge \dots \wedge b_n$ a conjunction of literals ($n \geq 0$).

□

We adopt the Prolog convention according to which variables are denoted by uppercase letters and predicate and function symbols start with a lowercase letter. (The interested reader may take a peek at Figure 2 at this point to see some example clauses, which in fact constitute a complete GDL axiomatization of our running example domain.) GDL imposes some general restrictions on a set of clauses, with the intention to ensure finite derivability.

Definition 3 The *dependency graph* for a set G of clauses is a directed, labeled graph whose nodes are the predicate symbols that occur in G and where there is a *positive edge* $p \xrightarrow{+} q$ if G contains a clause $p(\vec{s}) \Leftarrow \dots \wedge q(\vec{t}) \wedge \dots$, and a *negative edge* $p \xrightarrow{-} q$ if G contains a clause $p(\vec{s}) \Leftarrow \dots \wedge \neg q(\vec{t}) \wedge \dots$.

To constitute a *valid GDL specification*, a set of clauses G and its dependency graph Γ must satisfy the following.

1. There are no cycles involving a negative edge in Γ (this is also known as being *stratified* (Apt et al., 1987; van Gelder, 1989));
2. Each variable in a clause occurs in at least one positive atom in the body (this is also known as being *allowed* (Lloyd and Topor, 1986));
3. If p and q occur in a cycle in Γ and G contains a clause

$$p(\vec{s}) \Leftarrow b_1(\vec{t}_1) \wedge \dots \wedge q(v_1, \dots, v_k) \wedge \dots \wedge b_n(\vec{t}_n)$$

then for every $i \in \{1, \dots, k\}$,

- v_i is variable-free, or
- v_i is one of s_1, \dots, s_m ($:= \vec{s}$), or
- v_i occurs in some \vec{t}_j ($1 \leq j \leq n$) such that b_j does not occur in a cycle with p in Γ .

□

Stratified logic programs are known to admit a specific *standard model*; we refer to (Apt et al., 1987) for details and just mention the following properties:

1. To obtain the standard model, clauses with variables are replaced by their (possibly infinitely many) ground instances.

Table 1: GDL keywords.

<code>role(R)</code>	R is a player
<code>init(P)</code>	P holds in the initial position
<code>true(P)</code>	P holds in the current position
<code>legal(R,M)</code>	player R has legal move M
<code>does(R,M)</code>	player R does move M
<code>next(P)</code>	P holds in the next position
<code>terminal</code>	the current position is terminal
<code>goal(R,N)</code>	player R gets goal value N in the current position

2. Clauses are interpreted as reverse implications.
3. The standard model is minimal while interpreting negation as non-derivability (the “negation-as-failure” principle (Clark, 1978));

The second and third restriction in Definition 3 essentially guarantee that a logic program entails a *finite* number of ground atoms via its standard model. This is necessary to enable agents to make effective use of a set of game rules.

3.2 GDL Keywords

As a tailor-made specification language, GDL uses a few pre-defined predicate symbols. These are shown in Table 1 together with their informal meaning. A further, standard predicate is `distinct(X,Y)` to express (syntactic) inequality of two terms.¹

GDL imposes additional restrictions on the use of these keywords.

Definition 4 A *valid GDL specification* is a set of clauses G that, in addition to the restrictions in Definition 3, satisfies the following conditions.

- `role` only appears in the head of clauses that have an empty body;
- `init` only appears as head of clauses and is not connected, in the dependency graph for G , to any of `true`, `legal`, `does`, `,`, `goal`;
- `true` only appears in the body of clauses;
- `does` only appears in the body of clauses and is not connected, in the dependency graph for G , to any of `legal`, `terminal`, `goal`;
- `nly` appears as head of clauses.

□

¹The semantics of this predicate is given by tacitly assuming the addition of the clause

$$\text{distinct}(s,t) \Leftarrow$$

for every pair s,t of syntactically different ground terms.

According to the informal semantics given in (Love et al., 2006), a GDL specification G is to be understood as follows. The derivable instances of `role(R)` define the players. The initial state is composed of the derivable instances of `init(P)`. In order to determine the legal moves of a player in any given state, this state has to be encoded first, using the keyword `true`. More precisely, let $S = \{p_1, \dots, p_n\}$ be a state (e.g., the derivable instances of `init(P)` at the beginning), then G is extended by the clauses

$$\begin{aligned} \text{true}(p_1) &\Leftarrow \\ \dots & \\ \text{true}(p_n) &\Leftarrow \end{aligned} \quad (1)$$

Those instances of `legal(R,A)` which are derivable from this extended program define all legal actions A for player R in state S . In the same way, the clauses for `terminal` and `goal(R,N)` define termination and goalhood (of value N for player R) *relative* to the encoding of a given state. Determining a state transition, finally, requires the encoding of the current state along with clauses representing a joint move. Specifically, if players r_1, \dots, r_n make moves a_1, \dots, a_n , then

$$\begin{aligned} \text{does}(r_1, a_1) &\Leftarrow \\ \dots & \\ \text{does}(r_n, a_n) &\Leftarrow \end{aligned} \quad (2)$$

must be added to G , and then the derivable instances of `next(P)` compose the updated state.

3.3 Multiagent Environments in GDL

GDL provides all necessary features for declarative, formal descriptions of arbitrary multiagent environments as defined in Section 2. Of course there are many possible ways in which any specific environment can be axiomatized. We therefore define two sets of GDL clauses as *logically equivalent* if for any finite set of ground clauses (1) and (2) added, the two standard models of the two resulting logic programs agree on the interpretation of all GDL keywords. Before we can show how to formalize multiagent environments in GDL, we need the following syntactic definitions.

For any finite subset $S = \{p_1, \dots, p_n\} \subseteq \Sigma$ of a set of ground terms, the following conjunction axiomatizes S as the current state:

$$S^{\text{true}} \stackrel{\text{def}}{=} \text{true}(p_1) \wedge \dots \wedge \text{true}(p_n) \wedge \neg p_S \quad (3)$$

Here, p_S is an auxiliary predicate, one for every finite $S \subseteq \Sigma$, whose purpose is to ensure that the conjunction does not hold for states that are strict supersets of S :

$$p_S \Leftarrow \text{true}(X) \wedge \text{distinct}(X, p_1) \wedge \dots \wedge \text{distinct}(X, p_n)$$

Hence, p_S is true for any state in which at least one state component X is true that differs syntactically from any of p_1, \dots, p_n , that is, the elements of S . This ensures that the conjunction defined in (3) characterizes state S exactly.

Furthermore, for any function $A : \{r_1, \dots, r_n\} \mapsto \Sigma$, where $r_1, \dots, r_n \in \Sigma$, the following conjunction axiomatizes A as a joint action:

$$A^{\text{does}} \stackrel{\text{def}}{=} \text{does}(r_1, A(r_1)) \wedge \dots \wedge \text{does}(r_n, A(r_n))$$

We are now ready to show how GDL can be used to axiomatize multiagent domains.

Definition 5 Let $E = (R, s_1, t, l, u, g)$ be a multiagent environment based on ground symbolic expressions Σ , then any valid set of GDL clauses is an *axiomatic description* of E if it is logically equivalent to the following.

- $\text{role}(r) \Leftarrow$ for each $r \in R$;
- $\text{init}(p) \Leftarrow$ for each $p \in s_1$;
- $\text{terminal} \Leftarrow S^{\text{true}}$ for each $S \in t$;
- $\text{legal}(r, a) \Leftarrow S^{\text{true}}$ for each $(r, a, S) \in l$;
- $() \Leftarrow A^{\text{does}} \wedge S^{\text{true}}$ for each $p \in u(A, S)$ and $A : R \mapsto \Sigma, S \subseteq \Sigma$;
- $\text{goal}(r, n) \Leftarrow S^{\text{true}}$ for each $(r, n, S) \in g$.

□

It is important to realize that this direct axiomatization, where all relations and functions are encoded explicitly, is used solely to define the intended semantics. In practice, of course, a domain can be described in a much more compact manner, using variables, logical equivalence, and possibly auxiliary predicates. As an example, Figure 2 depicts a complete GDL specification of the multiagent environment introduced in Section 2. It is not too difficult to verify that this is a valid set of clauses according to Definition 3 and 4 and that it is indeed a correct axiomatic description of this domain according to Definition 5.

The specification of the Game Description Language in (Love et al., 2006) lacks a fully formal definition of the intended meaning of a specification. This

is why there are no formal grounds on which it could actually be proved that Definition 5 yields a correct description of a multiagent environment. In fact, we can and will use our formal concept of a multiagent domain to provide just this precise semantics for GDL in terms of a transition system.

4 A MULTIAGENT SEMANTICS FOR GDL

In the preceding section, we have shown how GDL provides a declarative, compact language to formally describe a large class of multiagent environments in a machine processable fashion. In this section, we show how the abstract model of a multiagent environment can in turn be used to provide a formal semantics for GDL in terms of a transition system. In this way we make precise what is only informally described in (Love et al., 2006).

Any valid game description G in GDL contains a finite set of function symbols, including constants, which implicitly determines a (usually infinite) set of ground terms. This set constitutes the symbol base Σ in the transition-based semantics for G . The syntactic restrictions in GDL ensure finite derivability, so that each state, the set of roles, etc. are all finite subsets of Σ . The following definition of the semantics of a GDL description is straightforwardly obtained by reversing the mapping from a multiagent environment into GDL (cf. Definition 5).

Definition 6 Let G be a valid GDL specification, whose signature determines the set of ground terms Σ . The *semantics* of G is the multiagent environment (R, s_1, t, l, u, g) where²

- $R = \{r \in \Sigma : G \models \text{role}(r)\}$;
- $s_1 = \{p \in \Sigma : G \models \text{init}(p)\}$;
- $t = \{S \in 2^\Sigma : G \cup S^{\text{true}} \models \text{terminal}\}$;
- $l = \{(r, a, S) : G \cup S^{\text{true}} \models \text{legal}(r, a, S)\}$, where $r \in R, a \in \Sigma$, and $S \in 2^\Sigma$;
- $u(A, S) = \{p \in \Sigma : G \cup A^{\text{does}} \cup S^{\text{true}} \models ()\}$, for all $A : (R \mapsto \Sigma)$ and $S \in 2^\Sigma$;
- $g = \{(r, n, S) : G \cup S^{\text{true}} \models \text{goal}(r, n, S)\}$, where $r \in R, n \in \mathbb{N}$, and $S \in 2^\Sigma$.

□

This definition provides a formal semantics for GDL in terms of abstract multiagent environments. Finite derivability in valid GDL specifications implies that

²Below, entailment (\models) is via the standard model of a set of clauses.

```

role(ag1) ←
role(ag2) ←
role(ag3) ←

init(at(ag1,1,1)) ←
init(at(ag2,5,1)) ←
init(at(ag3,1,5)) ←

terminal ← true(at(ag1,X,Y)) ∧ true(at(ag3,X,Y))
terminal ← true(at(ag2,X,Y)) ∧ true(at(ag3,X,Y))
terminal ← ¬remain

remain ← true(at(ag3,X,Y))

legal(R,stay) ← true(at(R,X,Y))
legal(ag3,exit) ← ¬terminal ∧ true(at(ag3,1,1))
legal(ag3,exit) ← ¬terminal ∧ true(at(ag3,5,1))
legal(ag3,exit) ← ¬terminal ∧ true(at(ag3,1,5))
legal(R,move(D)) ← ¬terminal ∧ true(at(R,U,V)) ∧ adjacent(U,V,D,X,Y)

adjacent(X,Y1,north,X,Y2) ← co(X) ∧ succ(Y1,Y2)
adjacent(X,Y1,south,X,Y2) ← co(X) ∧ succ(Y2,Y1)
adjacent(X1,Y, east,X2,Y) ← co(Y) ∧ succ(X1,X2)
adjacent(X1,Y, west,X2,Y) ← co(Y) ∧ succ(X2,X1)

co(1) ←      co(2) ←      co(3) ←      co(4) ←      co(5) ←
succ(1,2) ←  succ(2,3) ←  succ(3,4) ←  succ(4,5) ←

next(at(R,X,Y)) ← does(R,stay) ∧ true(at(R,X,Y))
next(at(R,X,Y)) ← does(R,move(D)) ∧ true(at(R,U,V)) ∧ adjacent(U,V,D,X,Y) ∧ ¬capture(R)
next(at(ag3,X,Y)) ← true(at(ag3,X,Y)) ∧ capture(ag3)

capture(ag3) ← true(at(ag3,X,Y)) ∧ true(at(R,U,V)) ∧ does(ag3,move(D1)) ∧
               does(R,move(D2)) ∧ adjacent(X,Y,D1,U,V) ∧ adjacent(U,V,D2,X,Y)

goal(R,0) ← role(R) ∧ ¬terminal
goal(R,0) ← role(R) ∧ distinct(R,ag3) ∧ terminal ∧ ¬remain
goal(R,100) ← role(R) ∧ distinct(R,ag3) ∧ terminal ∧ true(at(ag3,X,Y))
goal(ag3,0) ← terminal ∧ true(at(ag3,X,Y))
goal(ag3,100) ← terminal ∧ ¬remain

```

Figure 2: A complete, formal description of the multiagent environment of Figure 1.

the entailment relation is decidable, which in turn ensures that the definition of the semantics is effective.

In the preceding section we have seen that one and the same multiagent environment can be axiomatically described in many different ways. With the help of Definition 6 it is now easy to verify that two logically equivalent GDL descriptions (as defined in Section 3.3) describe exactly the same environment.

Proposition 7 *The semantics of two logically equivalent, valid GDL descriptions coincide.*

Proof: By definition, two logically equivalent GDL descriptions agree on the interpretation of all GDL keywords for all finite additions of clauses (1) and (2). It is easy to see, then, that the various components of their semantics according to Definition 6 must be identical. ■

Based on this result it is also straightforward to

prove that Definition 6 indeed provides the complement to the encoding of a multiagent environment in GDL.

Proposition 8 *Let E be a multiagent environment and G any axiomatic description thereof, then the semantics of G is E .*

Proof: Consider the generic encoding of E given in Definition 5. It is easy to verify that the standard model for this set of clauses, augmented by any finite set of facts about relations `true` and `does` (cf. clauses (1) and (2), respectively, in Section 3.2), determines a semantics (R, s_1, t, l, u, g) via Definition 6 which equals E . The claim follows from Proposition 7 and the fact that any GDL encoding for E is logically equivalent to the generic clauses given in Definition 5. ■

5 DISCUSSION

We have shown how the Game Description Language, developed in the context of General Game Playing, can be readily used as a declarative language to provide compact and machine processable specifications of a large class of multiagent environments. This can be applied to formalize the rules, for example, of an e-marketplace, of publicly accessible agent platforms on the Internet, of problem domains used in agent competitions, etc. By automatically processing these specifications, autonomous agents can fully automatically learn how to participate in a new or modified environment without the need to be (re-)programmed. Moreover, successful off-the-shelf general game playing systems can be readily employed as intelligent agents for these environments.

It is interesting to note that GDL has been originally developed as problem specification language for a competition (Genesereth et al., 2006), much like the Planning Domain Description Language (PDDL) (McDermott, 2000), which today is a quasi standard for the specification of planning domains. GDL can be viewed as a generalization of PDDL to domains with multiple agents, because solving a planning problem can be understood as playing a single-player game. Indeed, most features of current versions of PDDL can be expressed in GDL, though with one notable exception: sensing actions are not included in the current version of GDL. Although a GDL specification leaves agents with uncertainty about how the world evolves (an agent can decide on its own actions but not on those of all other agents), the language has been written for games without information asymmetry. An important research issue for the near future is to extend the Game Description Language so as to support descriptions of games with asymmetric information and sensing actions, which is a typical feature of card games, for instance. This would then provide a suitable formalization language for an even larger class of multiagent environments than considered in this paper.

In the second part of the paper, we have used the concept of a multiagent environment to provide a formal, transition-based semantics for GDL. With this we have made precise what is only informally described in (Love et al., 2006). Our semantics for GDL in terms of multiagent environments is related to an existing formal characterization of GDL by a game structure (van der Hoek et al., 2007). The main difference of the latter in comparison to our work are:

- It is restricted to propositional GDL;
- It puts further restrictions on GDL, such as not allowing predicate `init` to occur in clause with

non-empty bodies;

- It uses an inductive definition of the set of all states in order to obtain only those which are reachable from the initial state. Since it is possible to give valid GDL specifications of games that do not terminate, this definition would be undecidable in the general setting.³

These restrictions have been imposed because the focus in (van der Hoek et al., 2007) lies on the use of Temporal Logic for the purpose of verifying properties of games, such as termination or winnability. In contrast to this, the semantics given in the present paper covers full GDL.

ACKNOWLEDGEMENTS

We are grateful to the anonymous reviewers of this paper for helpful suggestions. This research was partially supported by *Deutsche Forschungsgemeinschaft* under Contract TH 541/16-1.

REFERENCES

- Apt, K., Blair, H. A., and Walker, A. (1987). Towards a theory of declarative knowledge. In Minker, J., editor, *Foundations of Deductive Databases and Logic Programming*, chapter 2, pages 89–148. Morgan Kaufmann.
- Clark, K. (1978). Negation as failure. In Gallaire, H. and Minker, J., editors, *Logic and Data Bases*, pages 293–322. Plenum Press.
- Clune, J. (2007). Heuristic evaluation functions for general game playing. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 1134–1139, Vancouver. AAAI Press.
- Finnsson, H. and Björnsson, Y. (2008). Simulation-based approach to general game playing. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 259–264, Chicago. AAAI Press.

³It is worth noting that this does not contradict the finite derivability property of valid GDL specifications, which just implies that all *local* reasoning problems are decidable. More specifically, given a particular state it is decidable whether an action is possible, and given a joint action it is also decidable what properties hold in the updated state, etc. On the other hand, GDL is expressive enough to describe any Turing machine as a “game” using clauses like

$$\begin{aligned} \text{init}(\text{head}(0)) &\Leftarrow \\ \text{next}(\text{head}(\text{succ}(\bar{x}))) &\Leftarrow \text{true}(\text{head}(\bar{x})) \wedge \\ &\quad \text{does}(\text{tm}, \text{move_forward}) \end{aligned}$$

Hence, reachability of states is generally undecidable in GDL.

- Genesereth, M., Love, N., and Pell, B. (2006). General game playing. *AI Magazine*, 26(2):73–84.
- Kuhlmann, G., Dresner, K., and Stone, P. (2006). Automatic heuristic construction in a complete general game player. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 1457–1462, Boston. AAAI Press.
- Lloyd, J. (1987). *Foundations of Logic Programming*. Series Symbolic Computation. Springer, second, extended edition.
- Lloyd, J. and Topor, R. (1986). A basis for deductive database systems II. *Journal of Logic Programming*, 3(1):55–67.
- Love, N., Hinrichs, T., Haley, D., Schkufza, E., and Genesereth, M. (2006). General Game Playing: Game Description Language Specification. Technical Report LG-2006-01, Stanford Logic Group, Computer Science Department, Stanford University, 353 Serra Mall, Stanford, CA 94305. Available at: games.stanford.edu.
- McDermott, D. (2000). The 1998 AI planning systems competition. *AI Magazine*, 21(2):35–55.
- Schiffel, S. and Thielscher, M. (2007). Fluxplayer: A successful general game player. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 1191–1196, Vancouver. AAAI Press.
- van der Hoek, W., Ruan, J., and Wooldridge, M. (2007). Strategy logics and the game description language. In *Proceedings of the Workshop on Logic, Rationality and Interaction*, Beijing, China.
- van Gelder, A. (1989). The Alternating Fixpoint of Logic Programs with Negation. In *Proceedings of the 8th Symposium on Principles of Database Systems*, pages 1–10. ACM SIGACT-SIGMOD.

