

Strategy Games: A Renewed Framework

Fabio Mogavero
Università degli Studi di Napoli
Federico II
fm@fabiomogavero.com

Aniello Murano
Università degli Studi di Napoli
Federico II
murano@na.infn.it

Luigi Sauro
Università degli Studi di Napoli
Federico II
luigi.sauro74@gmail.com

ABSTRACT

Game theory is a useful framework in multi-agent system verification. In this context, an important contribution is given by modal logics for strategic ability, such as ATL^* , SL, and the like, which allow to describe the interaction among players for the achievement of specific goals. However, most of the attempts carried out so far have focused on specific forms of games in which strategic-reasoning aspects are glued with agent temporal goals.

In this paper, we propose a revisit of logics for strategic reasoning by following the general guidelines of game-theory, where the definition of solution concepts abstracts away the underlying cost-benefit analysis. Specifically, we introduce a game-theoretic framework consisting of three entities: the *arena*, shaping the rules of the game, the *extension* describing the property of interest of each possible play, and the *schema*, representing the agent roles and their interaction. In particular, to describe the latter, we use a variant of SL, by casting away its LTL temporal core. The new framework allows to gain in flexibility, modularity, and technical simplification, as well as, to grasp new features and better complexity results about logics previously studied. To give an evidence of this, by rephrasing SL in our framework, we improve known results about some of its fragments.

Categories and Subject Descriptors

I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods—*Modal logic, Temporal logic, Representation languages*; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Theory, Verification

Keywords

Game models, Strategic logics, Formal verification.

1. INTRODUCTION

In the multi-agent system design and verification, logics for the strategic reasoning play a fundamental role [2,3,5,8,13,17].

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

These logics allow to check open system reliability by rephrasing the decision question as an interaction among different players in a game-strategic setting, which may compete or collaborate for the achievement of specific *goals*. A pioneering logic in this matter is *Alternating-Time Temporal Logic* (ATL^* , for short) [2]. This logic makes use of strategic modalities of the form “ $\langle\langle A \rangle\rangle$ ” and “ $[[A]]$ ”, for a set A of *agents*. Their formulas are interpreted over *concurrent game structures* (CGS, for short) [2], which are transition systems modeling the interaction among agents (a.k.a. *players*) and whose states are labeled with a set of atomic proposition true in that state. Given a CGS \mathcal{G} and a set A of agents, an ATL^* formula $\langle\langle A \rangle\rangle\psi$ is satisfied at a state s of \mathcal{G} if there is a set of strategies for agents in A such that, no matter which strategies are executed by agents not in A , the resulting outcome of the interaction in \mathcal{G} satisfies the temporal goal ψ at s . In particular, ψ is build upon the syntax rules of the classical linear-time temporal-logic LTL [15]. Since its introduction, ATL^* has been intensively studied and its model checking has been proved to be 2EXPTIME-complete [2]. Also, several practical applications of this logic have been investigated through the implementation of useful verification tools [6].

The benefits of ATL^* brought to multi-agent system verification have opened to a broader modeling scenario and several variants of this logic have been successively introduced and studied, in order to balance expressiveness with respect to complexity of the related decision questions [1,3,7,10,11,13,19]. Among the others, a recent story of success is *Strategy Logic* (SL, for short) [13]. This logic has the benefit to include all main logics for strategic reasoning previously introduced and to be powerful enough to express fundamental concepts of games, such as *Nash equilibrium*, which instead cannot be expressed in ATL^* , unless one considers substantial extensions [16].

A key aspect of SL is that, differently from ATL^* , strategies are not treated implicitly through its modalities, but rather used as first-order objects that can be existentially and universally quantified. Specifically, SL uses the existential $\langle\langle x \rangle\rangle$ and the universal $[[x]]$ strategic modalities, which can be read as “*there exists a strategy x* ” and “*for all strategies x* ”, respectively. In SL, strategies represent general conditional plans that at each step prescribe an action on the base of the previous history. Thus, strategies are not intrinsically glued to a specific agent. Then, to use a strategy, an explicit *binding* operator (a, x) allows to bind an agent a to the strategy associated with a variable x . On the other hand, SL uses CGSs as underlying structures as well as temporal goals are given in terms of LTL rules, as it is for ATL^* .

Although ATL^* and SL count several interesting applications in computer science, they only use a restricted power offered by game theory. This is mainly due to the rigidity of the syntax used to define these logics, which overall come out as an evolution of model checking and temporal logics, originally conceived for the verification of closed systems. Indeed, from one side we use CGS as a model, which is a Kripke structure further equipped with a description of player moves. Hence, we inherit an intrinsic labeling of the states with atomic propositions. From the other side, we extend path modalities with game modalities, but we limit ourself to use predicates over paths, *i.e.*, game evolutions, mainly expressed in terms of LTL formulas. Additionally, this part is glued inside the syntax of the logic itself. In game theory speaking, player rules, goals, and verification properties are independent entities, usually given and treated separately. Conversely, having them all glued in one object leads to the loss of several important features, as well as, it limits the applicability of the resulting model. Yet in terms of game theory, combining these entities is as bizarre as letting a solution concept to depend from the specific form of agents utility functions. Consider for example Nash equilibria. Using a somewhat standard notation, a strategy profile ξ is a Nash equilibrium if, for each other strategy profile ξ' consisting in a unilateral deviation of an agent i from ξ , $u_i(\xi') \leq u_i(\xi)$, where u_i is the utility function of the agent i . Notice that, since the utility functions occur as generic functions, the definition of Nash equilibria does not rely on how they are determined: different cost-benefit analysis methods can be used as a black box in the pure strategic reasoning.

In this paper, we propose a revisit of classical logics for strategic reasoning by following the general guidelines of the game-theory where the definition of solution concepts (*i.e.*, the pure strategic reasoning) abstracts away the underlying cost-benefit analysis (together with the temporal constraints used to determinate it). Specifically, we introduce a game-theoretic framework consisting of three entities: the *arena*, shaping the rules of the game, the *extension* describing the property of interest of each possible play, and the *schema*, representing properties of agent interactions (*e.g.* Nash equilibrium). In particular, to describe the latter, we use a variant of SL, by casting away its LTL temporal core.

We show that the introduced framework has several advantages. First, it allows to reason about the game model in a modular way. More specifically, we can investigate the complexity of a decision problem by looking at the exact contribution of each component of the strategic reasoning. As often it happens in multi-agent system verification, system and specification parameters may vary significantly one from another and the best methodology to be used to address a specific decision problem is a trade-off. Second, one can embed different formalisms to describe temporal goals. In particular, we can use very expressive extensions of the LTL syntax rules by using the linear $\mu\text{CALCULUS}$ [18], or vice versa use a very restricted form of LTL such as the common fragment of LTL and CTL. As another example, one can inject goals directly represented by means of automata [20].

As a positive side effect, the proposed game framework let us to re-address the main features of SL and obtain new and better complexity results regarding the model-checking problem of some of its fragments, once the related extension is represented by means of classic LTL formulas. Specifically, three fragments of SL have been recently introduced

and deeply investigated: $\text{SL}[\text{IG}]$, $\text{SL}[\text{CG}]$, and $\text{SL}[\text{DG}]$ [10–12]. The latter two logics, which include the former, are at the moment the maximal fragments of SL known to have an elementary model-checking and precisely 2EXPTIME-COMplete . Here, we give new game-based decision procedures for the model-checking problem of all these fragments, which allow to improve their overall complexities. Notably, we use as decision tool of the game framework, the framework itself.

Due to space limit, most of the proofs are just sketched and reported in the full version of the article.

2. STRATEGY GAMES

In this section, we introduce a reformulation of the game-theoretic framework through which we model strategic reasonings. Differently from other approaches, where a unique structure is defined, here we separate the basic aspects of a game in three components: arena, extension, and schema. Roughly speaking, an arena describes the game board and its moves, *i.e.*, the physical world where agents act. An extension specifies object-level conditions over game evolutions, *i.e.*, how to determine the winner of a play. Finally, a schema abstractly represents strategic properties of the game, like determinacy, existence of Nash equilibria, and so on.

2.1 Arenas

We now formally define the concept of arena.

DEFINITION 2.1. *Arena.* - A multi-agent concurrent arena is a tuple $\mathcal{A} \triangleq \langle \text{Ag}, \text{Ac}, \text{St}, \text{tr} \rangle \in \text{Ar}(\text{Ag})$, where Ag is the finite set of agents, a.k.a. players, Ac is the set of actions, a.k.a. moves, St is the non-empty sets of states, a.k.a. positions, and $\text{Ar}(\text{Ag})$ is the class of all arenas over Ag . Assume $\text{Dc} \triangleq \text{Ag} \rightarrow \text{Ac}$ to be the set of decisions, *i.e.*, partial functions describing the choices of an action by some agent. Then, $\text{tr} : \text{Dc} \rightarrow (\text{St} \rightarrow \text{St})$ denotes the transition function mapping every decision $\delta \in \text{Dc}$ to a partial function $\text{tr}(\delta) \subseteq \text{St} \times \text{St}$ representing a deterministic graph over the states.

Intuitively, an arena can be seen as a *labeled transition graph*, where labels are possibly incomplete agent decisions, which determine the transition to be executed at each step of a play in dependence of the choices made by the agents in the relative state. Incomplete decisions allow us to represent any kind of legal move in a state, where some agents or a particular combination of actions may not be active. Note that, since $\text{tr}(\delta)$ is a partial function, we are implicitly assuming that the arena is deterministic, *i.e.*, a decision $\delta \in \text{Dc}$ may be not applicable in a state $s \in \text{St}$, but if it is, then it leads to the unique successor $\text{tr}(\delta)(s)$.

An arena \mathcal{A} naturally induces a graph $\mathcal{G}(\mathcal{A}) \triangleq \langle \text{St}, \text{Ed} \rangle$, whose vertexes are represented by the states and the edge relation $\text{Ed} \triangleq \bigcup_{\delta \in \text{Dc}} \text{tr}(\delta)$ is obtained by rubbing out all labels on the transitions. A path $\pi \in \text{Pth}$ in \mathcal{A} is simply a path in $\mathcal{G}(\mathcal{A})$. Similarly, the *order* $|\mathcal{A}| \triangleq |\mathcal{G}(\mathcal{A})|$ (resp., *size* $\|\mathcal{A}\| \triangleq \|\mathcal{G}(\mathcal{A})\|$) of \mathcal{A} is the order (resp., size) of its induced graph. Note that there could be states where no transitions are available, *i.e.*, $\text{dom}(\text{Ed}) \subset \text{St}$. If this is the case, the states in $\text{St} \setminus \text{dom}(\text{Ed})$ are called *sink-states*.

As usual in the study of extensive-form games, finite paths also describe the possible evolutions of a play up to a certain point. For this reason, they are called in the game-theoretic jargon *histories* and the corresponding set is denoted by $\text{Hst} \triangleq \{\rho \in \text{Pth} : |\rho| < \omega\}$.

A *strategy* is a partial function $\sigma \in \text{Str} \triangleq \text{Hst} \rightarrow \text{Ac}$ prescribing, whenever defined, which action has to be performed for a certain history of the current outcome. In addition, a *profile* is a function $\xi \in \text{Prf} \triangleq \text{Ag} \rightarrow \text{Str}$ assigning to every agent the adopted strategy. Roughly speaking, a strategy is a generic conditional plan to achieve one or more goals, which specifies “*what to do*” but not “*who will do it*”, while a profile specifies a unique behavior for each agent by associating her with a strategy. Note that a given strategy can be used by more than one agent at the same time.

For a profile $\xi \in \text{Prf}$, to identify which action an agent $a \in \text{Ag}$ has chosen to perform on a history $\rho \in \text{Hst}$, we first extract the corresponding strategy $\xi(a)$ and then we determinate the action $\xi(a)(\rho)$, whenever defined. To identify, instead, the whole decision on ρ , we apply the flipping operator to ξ . We get so a function $\hat{\xi} : \text{Hst} \rightarrow \text{Dc}$ such that $\hat{\xi}(\rho)(a) = \xi(a)(\rho)$, which maps each history to the planned decision.

A path $\pi \in \text{Pth}$ is *coherent w.r.t.* a profile $\xi \in \text{Prf}$ (ξ -*coherent*, for short) iff, for all $i \in [1, |\pi|]$, it holds that $(\pi)_i = \text{tr}(\hat{\xi}((\pi)_{<i}))((\pi)_{i-1})$, *i.e.*, $(\pi)_i$ is the unique successor of $(\pi)_{i-1}$ identified by the agent decision $\hat{\xi}((\pi)_{<i})$ described in the profile ξ on the history $(\pi)_{<i}$. Moreover, π is a ξ -*play* iff it is a maximal ξ -coherent path, *i.e.*, it is not a proper prefix of any another ξ -coherent path. Clearly, an infinite path is ξ -coherent iff it is a ξ -play. Given a state $s \in \text{St}$, the determinism of the arena ensures that there exists exactly one ξ -play π starting in s , *i.e.*, $\text{fst}(\pi) = s$. Such a play is called (ξ, s) -*play*. For this reason, we use the *play function* $\text{play} : \text{Prf} \times \text{St} \rightarrow \text{Pth}$ to identify, for each profile $\xi \in \text{Prf}$ and state $s \in \text{St}$, the corresponding (ξ, s) -*play* $\text{play}(\xi, s)$. Observe that finite ξ -plays do not necessarily end in a sink-state, since it may be that it cannot be extended in a way that is coherent *w.r.t.* the profile ξ (below we return on this point).

Arenas describe generic mathematical structures, where the basilar game-theoretic notions of history, strategy, profile, and play can be defined. However, in several contexts, some constraints rule out how the function tr maps partial decisions to transitions. For example, some definitions of *bounded* concurrent arenas require that, in a given state, each agent is associated with a finite non-empty set of actions, whose possible combinations are all applicable in that state. Another case is that of classic *turn-based* arenas, where there is at most one active agent for each state. Both these examples represent particular cases of our general definition of arena, since they can be embedded in this formalism by forcing the transition function to satisfy some further condition. In order to illustrate such embeddings, by means of the functions $\text{dc} : \text{St} \rightarrow 2^{\text{Dc}}$ and $\text{ag} : \text{St} \rightarrow 2^{\text{Ag}}$ defined as follows, we preliminary introduce the set of decisions and agents that trigger some transition in a given state $s \in \text{St}$:

$$\text{dc}(s) \triangleq \{\delta \in \text{Dc} : s \in \text{dom}(\text{tr}(\delta))\};$$

$$\text{ag}(s) \triangleq \{a \in \text{Ag} : \exists \delta \in \text{dc}(s) . a \in \text{dom}(\delta)\}.$$

Thanks to this notation, a restriction to bounded concurrent arenas can be provided by saying that, for each state $s \in \text{St}$, there is a left-total relation $R \subseteq \text{Ag} \times \text{Ac}$ with $|\text{rng}(R)| < \omega$ such that $\text{dc}(s) = \{\delta \in \text{Dc} : \delta \subseteq R\}$. Similarly, a restriction to turn-based arenas is given by $|\text{ag}(s)| \leq 1$.

As a running example, used to illustrate the introduced definitions, consider the arena \mathcal{A}_S depicted in Figure 1. It represents a model of a simple *scheduler system* in which two *processes*, P_1 and P_2 , can require the access to a *shared resource*, like a processor, and an *arbiter* A is used to solve




Figure 1: Scheduler Arena \mathcal{A}_S .

all conflicts that may arise when contending requests are made. In particular, the arbiter can preempt a process owning the resource to allow the other one to access to it. The processes have two actions to interact with the system: r for request and f for free/release. The former is used to request the resource from the system, when this is not yet owned, while the latter releases it, when this is not needed anymore. In addition, a process can simply avoid to make any action, in case it wants to remain in the current state. The arbiter, from its side, has two actions to decide which process has to receive the resource: 1 for P_1 and 2 for P_2 . The whole scheduler system can reside in the following four states: I , 1, 2, and A . The idle state I indicates that both processes do not own the resource, while i , with $i \in \{1, 2\}$, that process P_i is using it. Finally, the arbitrage state A represents the situation in which an action from the arbiter is required in order to solve a conflict between contending requests. For readability reasons, a decision is graphically represented by an arrow \mapsto with a sequence of agents on left hand side and the sequence of corresponding actions on right hand side. Formally, \mathcal{A}_S is the arena $\langle \text{Ag}, \text{Ac}, \text{St}, \text{tr} \rangle$, where all components but the transition function are set as follows: $\text{Ag} = \{A, P_1, P_2\}$, $\text{Ac} = \{r, f, 1, 2\}$, and $\text{St} = \{I, 1, 2, A\}$. In addition, we have that $\text{ag}(s) = \{P_1, P_2\}$, for all $s \in \{I, 1, 2\}$, and $\text{ag}(A) = \{A\}$. Moreover, $\text{dc}(I) = \{\emptyset, P_1 \mapsto r, P_2 \mapsto r, P_1 P_2 \mapsto rr\}$, $\text{dc}(A) = \{A \mapsto 1, A \mapsto 2\}$, and $\text{dc}(i) = \{\emptyset, P_i \mapsto f, P_{3-i} \mapsto r, P_i P_{3-i} \mapsto fr\}$, for all $i \in \{1, 2\}$.

2.2 Extensions

Clearly, a game, besides its dynamics, consists also of private objectives or a global goal, which typically require to check some properties over paths. For instance, relevant conditions in chess are whether the game ends up in a checkmate or a stalemate, which determine a winner or a draw, respectively. More in general, path properties may represent fairness, reachability, safety, or other conditions, usually expressed by means of some kind of temporal language, like LTL [15] or the linear $\mu\text{CALCULUS}$ [18], which combine atomic propositions over states with suitable temporal operators. As game theorists usually do, our aim is to separate the strategic reasoning from the specification of such properties and represent them as generic Boolean conditions over paths.

DEFINITION 2.2. *Extension.* - An extension is a tuple $\mathcal{E} \triangleq \langle \mathcal{A}, \text{Pr}, \text{pr} \rangle \in \text{Ex}(\text{Ag}, \text{Pr})$, where $\mathcal{A} \in \text{Ar}(\text{Ag})$ is the underlying arena, Pr is the finite non-empty set of predicates, $\text{pr} : \text{Pr} \rightarrow 2^{\text{Pth}}$ is the predicate function mapping each

predicate $p \in \text{Pr}$ to the set of paths $\text{pr}(p)$ satisfying it, and $\text{Ex}(\text{Ag}, \text{Pr})$ is the class of all extensions over Ag and Pr .

Intuitively, an extension classifies paths by means of a set of monadic predicates that describe relevant temporal properties of plays for the specific domain under analysis.

An extension is *Borelian* (resp., *regular*) iff, for all predicates $p \in \text{Pr}$, the induced language $\text{pr}(p) \subseteq \text{St}^\infty$ of finite and infinite words over the states is Borelian (resp., regular) [14].

Continuing with the scheduler example, consider the regular extension \mathcal{E}_S consisting of the scheduler arena \mathcal{A}_S and two predicates p_1 and p_2 expressing that, whenever a conflict arises, every process P_k will have at some point in the future access to the resource. Formally, for each path $\pi \in \text{Pth}$, it holds that $\pi \in \text{pr}(p_k)$ iff, for all $i \in [0, |\pi|[$ with $(\pi)_i = \mathbf{A}$, there exists $j \in]i, |\pi|[$ such that $(\pi)_j = \mathbf{k}$.

2.3 Schemas

Coming back to the chess analysis, note that we have defined the chessboard together with the possible legal moves (the arena). Also, we have introduced the relevant properties, such as checkmate or stalemate, which have to be checked during the game (the extension). However, we have not yet completely grasped the notion of game. Indeed, an initial configuration needs to be specified and, more important, it is necessary to indicate which behavior we expect from the players. To do this, we introduce the concept of schema. It abstractly describes the possible roles and the allowed interactions of agents, by means of a relation between an extension, with an associated initial state, and some elements describing which behaviors that extension satisfies.

DEFINITION 2.3. *Schema.* - A schema over sets of agents Ag and predicates Pr is a tuple $\mathcal{S} \triangleq \langle \text{Tr}, \models \rangle \in \text{Sc}(\text{Ag}, \text{Pr})$, where Tr is the non-empty set of targets, $\models \subseteq \text{Ex}(\text{Ag}, \text{Pr}) \times_{\mathcal{E}} \text{St}_{\mathcal{E}} \times \text{Tr}$ is the schema relation describing which targets $t \in \text{Tr}$ can be achieved on an extension $\mathcal{E} \in \text{Ex}(\text{Ag}, \text{Pr})$ starting from a given initial state $s \in \text{St}_{\mathcal{E}}$, in symbols $\mathcal{E}, s \models t$, and $\text{Sc}(\text{Ag}, \text{Pr})$ is the class of all schemas over Ag and Pr .

In the scheduler example, an interesting target t is to avoid starvation for a process, due to the selfish behavior of the other one. To precisely state t , one can either define an ad-hoc schema relation or introduce a suitable formal language that allows to describe an entire class of possible schemas. Clearly, we follow the second approach from now on.

2.4 Games

By summing up the definitions of arena, extension, and schema, we now formalize the concept of game.

DEFINITION 2.4. *Game.* - Let $\mathcal{S} \in \text{Sc}(\text{Ag}, \text{Pr})$ be a schema over the sets of agents Ag and predicates Pr . Then, a game w.r.t. \mathcal{S} is a tuple $\mathcal{G} \triangleq \langle \mathcal{E}, s, t \rangle \in \text{Gm}(\mathcal{S})$, where $\mathcal{E} \in \text{Ex}(\text{Ag}, \text{Pr})$ is the underlying extension, $s \in \text{St}_{\mathcal{E}}$ is the designated initial state, $t \in \text{Tr}$ is the prescribed target, and $\text{Gm}(\mathcal{S})$ is the class of all games over \mathcal{S} .

A game $\mathcal{G} \triangleq \langle \mathcal{E}, s, t \rangle$ is *fulfilled* iff $\mathcal{E}, s \models t$. Intuitively, this means that all agents can achieve the target t on the predicate arena \mathcal{E} starting from s . The *fulfillment problem* is to decide whether a game is fulfilled. Observe that this problem is an abstract generalization of the model-checking problem of a language against a model, which we here respectively represent via a schema and an extension.

3. STRATEGY LOGICS

By analogies with object-oriented languages, a schema can be seen as an interface which allows to describe the notion of fulfillment of a game in a very general way. Following the same analogy, a schema has to be implemented in order to make the game concrete. It is for this reason that we have introduced schemas in our framework. Indeed one can change an arena or an extension without changing the schema. Clearly, the more powerful the language used to formalize the schema is the wider the class of strategic notions we can grasp is as well. Hence, we borrow the strategic modalities from SL, instead of the less powerful coalition modalities of ATL^* , as a language for the targets, and its semantics as the corresponding schema relation.

3.1 Syntax

Strategic reasoning primary concerns to analyze different counterfactual hypothesis such as “what if two agents cooperate...”, “what if another agent tries to obstruct them...”, or “what if an agent deviates from a given behavior...”. From a logical point of view, this means to consider different patterns of quantification over strategies, as well as, to bind such strategies to the agents. For this reason, we introduce two *strategy quantifiers*, the existential $\langle\langle x \rangle\rangle$ and the universal $\llbracket x \rrbracket$, and an *agent binding* (a, x) , where a is an agent and x belongs to a countable set Vr of variables. Intuitively, these operators can be respectively read as “there exists a strategy x ”, “for all strategies x ”, and “bind agent a to the strategy associated with x ”. Clearly, such a bag of logical quantifiers would be completely pointless without the possibility to express properties over paths. Hence, the set Pr of predicates is part of our syntax together with standard Boolean operators.

It remains to discuss an important issue regarding which strategies can be assigned to the agents. As mentioned in the previous section, given a profile ξ and an initial state s , the play $\pi = \text{play}(\xi, s)$ can be a finite path even if its last state $\text{lst}(\pi)$ is not a sink-state. Formally, this happens when the decision $\hat{\xi}(\pi)$ is not in $\text{dc}(\text{lst}(\pi))$ even if $\text{dc}(\text{lst}(\pi)) \neq \emptyset$. To give an intuitive example, a chess match could be trivially blocked in its initial configuration just because the white player decides not to perform any move. Clearly, it is a matter of the specific application domain whether the agents are free to adopt any strategy, even the blocking ones, or they are forced to trigger a transition whenever it is possible. To be as general as possible, we do not implicitly restrict the strategies that can be assigned to the agents, but we introduce a further constant \odot to express the fact that we are restricting to no blocking assignments.

DEFINITION 3.1. *Syntax.* - SL formulas are defined by means of the following context-free grammar, where $a \in \text{Ag}$, $p \in \text{Pr}$, and $x \in \text{Vr}$:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle\langle x \rangle\rangle\varphi \mid \llbracket x \rrbracket\varphi \mid (a, x)\varphi \mid \odot.$$

Usually, to provide the semantics of a predicative logic, it is necessary to define the concepts of free and bound *placeholders*. For example, first order logic has the variables as unique type of placeholder. In SL, instead, since strategies can be associated to both agents and variables, we use the set of *free agents/variables* $\text{free}(\varphi)$ as the subset of $\text{Ag} \cup \text{Vr}$ containing (i) all agents a for which there is no binding (a, x) before the occurrence of a predicate p and (ii) all variables x for which there is a binding (a, x) but no quantification $\langle\langle x \rangle\rangle$

or $\llbracket x \rrbracket$. In case $\text{free}(\varphi) \cap \text{Ag} = \emptyset$ (resp., $\text{free}(\varphi) \cap \text{Vr} = \emptyset$), the formula φ is named *agent-closed* (resp., *variable-closed*). A *sentence* is a both agent- and variable-closed formula.

3.2 Semantics

Similarly as in first order logic, the interpretation of a formula makes use of an assignment function which associates placeholders to some elements of the domain. In particular, an *assignment* is a (possibly partial) function $\chi \in \text{Asg} \triangleq (\text{Vr} \cup \text{Ag}) \rightarrow \text{Str}$ mapping variables and agents to strategies. An assignment χ is *complete* iff it is defined on all agents, i.e., $\text{Ag} \subseteq \text{dom}(\chi)$. In this case, it directly identifies the profile $\chi|_{\text{Ag}}$ given by the restriction of χ to Ag . In addition, $\chi[e \mapsto \sigma]$, with $e \in \text{Vr} \cup \text{Ag}$ and $\sigma \in \text{Str}$, is the assignment defined on $\text{dom}(\chi[e \mapsto \sigma]) \triangleq \text{dom}(\chi) \cup \{e\}$ which differs from χ only in the fact that e is associated with σ . Formally, $\chi[e \mapsto \sigma](e) = \sigma$ and $\chi[e \mapsto \sigma](e') = \chi(e')$, for all $e' \in \text{dom}(\chi) \setminus \{e\}$. A path $\pi \in \text{Pth}$ is *coherent w.r.t.* an assignment $\chi \in \text{Asg}$ (χ -*coherent*, for short) in case the restriction of χ on agents can be extended into a profile in respect of which π is a coherent path. Formally, π is χ -coherent iff there exists a profile $\xi \in \text{Prf}$ with $\xi|_{(\text{dom}(\chi) \cap \text{Ag})} = \chi|_{\text{Ag}}$ such that π is ξ -coherent. Finally, an assignment χ is *non-blocking* in a state $s \in \text{St}$ in case that for each χ -coherent history ρ starting from s and not ending in a sink-state, χ in ρ agrees with one of the triggering decisions of $\text{lst}(\rho)$. Formally, an assignment χ is *non-blocking* in s iff for all χ -coherent histories $\rho \in \text{Hst}$ such that $\text{fst}(\rho) = s$ and $\text{dc}(\text{lst}(\rho)) \neq \emptyset$, there exists a decision $\delta \in \text{dc}(\text{lst}(\rho))$ such that $\delta(a) = \chi(a)(\rho)$, for all agents $a \in \text{dom}(\chi) \cap \text{Ag}$. We are ready to define the semantics of our logics.

DEFINITION 3.2. Semantics. - *Let \mathcal{E} be an extension. Then, for all SL formulas φ , assignments $\chi \in \text{Asg}$ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$, and states $s \in \text{St}$, the modeling relation $\mathcal{E}, \chi, s \models \varphi$ is inductively defined as follows.*

- 1) $\mathcal{E}, \chi, s \models p$ iff $\text{play}(\chi|_{\text{Ag}}, s) \in \text{pr}(p)$.
- 2) *Boolean operators are interpreted as usual.*
- 3) $\mathcal{E}, \chi, s \models \langle\langle x \rangle\rangle \varphi$ iff there exists a strategy $\sigma \in \text{Str}$ such that $\mathcal{E}, \chi[x \mapsto \sigma], s \models \varphi$.
- 4) $\mathcal{E}, \chi, s \models \llbracket x \rrbracket \varphi$ iff for all strategies $\sigma \in \text{Str}$, $\mathcal{E}, \chi[x \mapsto \sigma], s \models \varphi$.
- 5) $\mathcal{E}, \chi, s \models (a, x) \varphi$ iff $\mathcal{E}, \chi[a \mapsto \chi(x)], s \models \varphi$.
- 6) $\mathcal{E}, \chi, s \models \odot$ iff χ is non-blocking from s .

Note that the satisfaction of a sentence φ does not depend on assignments, hence we omit them and write $\mathcal{E}, s \models \varphi$.

In what follows, $\langle\langle a, x \rangle\rangle \varphi$ and $\llbracket a, x \rrbracket \varphi$ are abbreviations for $(a, x)(\odot \wedge \varphi)$ and $(a, x)(\odot \rightarrow \varphi)$, respectively. The formula $\langle\langle a, x \rangle\rangle \varphi$ is used to ensure non-blocking assignments when the variable x is under the scope of an existential quantification (intuitively, “ a adopts x , which needs to be non-blocking and to satisfy φ ”). Similarly, $\llbracket a, x \rrbracket \varphi$ forces non-blocking assignments in case x is under the scope of a universal quantification (intuitively, “ a adopts x , which needs to satisfy φ in case it is non-blocking”).

In the scheduler example, the ability to avoid starvations can be represented by the formula $\varphi = \langle\langle x \rangle\rangle \langle\langle y_1 \rangle\rangle \langle\langle y_2 \rangle\rangle \llbracket z \rrbracket \phi$, where $\phi = \langle\langle \mathbf{A}, x \rangle\rangle (\langle\langle \mathbf{P}_1, y_1 \rangle\rangle (\langle\langle \mathbf{P}_2, z \rangle\rangle p_1 \wedge \langle\langle \mathbf{A}, x \rangle\rangle \langle\langle \mathbf{P}_2, y_2 \rangle\rangle (\langle\langle \mathbf{P}_1, z \rangle\rangle p_2$ with p_1 and p_2 being the same predicates as before. It is easy to see that for all $s \in \{\mathbf{I}, \mathbf{1}, \mathbf{2}, \mathbf{A}\}$, it holds that $\mathcal{E}_S, s \models \varphi$.

In particular, the arbiter strategy consists in alternating the access to the resource between the two processes, while the processes have to request the resource at least twice. Note that φ requires a unique strategy for the arbiter in order to coordinate with both the processes independently. Therefore, in terms of ATL coalition modalities, φ is weaker than $\langle\langle \{\mathbf{A}\} \rangle\rangle (p_1 \wedge p_2)$, but stronger than $\langle\langle \{\mathbf{A}, \mathbf{P}_1\} \rangle\rangle p_1 \wedge \langle\langle \{\mathbf{A}, \mathbf{P}_2\} \rangle\rangle p_2$. Actually, φ cannot be expressed in ATL^* .

Once the sets of agents and predicates are fixed, SL induces a schema \mathcal{S}_{SL} , where the set of targets is SL itself, i.e., its sentences, and the schema relation is given by the definition of the semantics. A *strategy game* is a game w.r.t. \mathcal{S}_{SL} . The set of all strategy games is denoted for short by $\text{SG} \triangleq \text{Gm}(\mathcal{S}_{\text{SL}})$. We now introduce some restrictions of SG. First, we consider the case in which the numbers $|\text{Ag}|$ of agents and $|\text{Vr}|$ of variables that are used to write a formula are fixed to the a priori values $n, m \in [1, \omega[$. We name these two fragments $\text{SG}[n_{\text{AG}}]$ and $\text{SG}[m_{\text{VAR}}]$, respectively. Also, by $\text{SG}[\text{TB}]$ we refer to games whose arena is turn-based. Finally, $\text{SG}[\Sigma_k]$ (resp., $\text{SG}[\Pi_k]$) denotes those games where the target has a quantification prefix of k alternation starting with an existential (resp., universal) quantifier.

3.3 Binding Fragments

We now consider three fragments of SL that allow to formalize interesting game properties not expressible in ATL^* . For example, they can express that an agent can join two or more different coalitions without producing mutual conflicts or that a winning condition is weaker than another one.

In what follows a *quantification prefix* over a set $V \subseteq \text{Vr}$ of variables is a finite word $\wp \in \{\langle\langle x \rangle\rangle, \llbracket x \rrbracket : x \in V\}^{|\text{V}|}$ of length $|\text{V}|$ such that each variable $x \in V$ occurs just once in \wp . By $\text{Qn}(V)$ we indicate the set of quantification prefixes over V , whereas $\langle\langle \wp \rangle\rangle$ (resp. $\llbracket \wp \rrbracket$) denote the set of variables occurring *existentially* (resp. *universally*) quantified in \wp . A *binding prefix* over V is a word $\flat \in \{(a, x), \langle\langle a, x \rangle\rangle, \llbracket a, x \rrbracket : a \in \text{Ag} \wedge x \in V\}^{|\text{Ag}|}$ such that each agent in Ag occurs exactly once in \flat . By Bn we indicate the set of all binding prefixes.

DEFINITION 3.3. Binding Fragments. - *SL[BG] formulas are defined by the following context-free grammar:*

$$\begin{aligned} \varphi &:= \wp \phi \mid \neg \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi, \\ \phi &:= \flat \psi \mid \neg \phi \mid \phi \wedge \phi \mid \phi \vee \phi, \\ \psi &:= p \mid \neg \psi \mid \psi \wedge \psi \mid \psi \vee \psi. \end{aligned}$$

where $\wp \in \text{Qn}(\text{free}(\psi))$, $\flat \in \text{Bn}$, and $p \in \text{Pr}$. Moreover, the fragments obtained by replacing the second row with $\phi := \flat \psi$, $\phi := \flat \psi \mid \phi \wedge \phi$, or $\phi := \flat \psi \mid \phi \vee \phi$, are called $\text{SL}[1\text{G}]$, $\text{SL}[\text{CG}]$, $\text{SL}[\text{DG}]$, respectively.

In what follows, we also consider some constraints on the number of agents/variables or the type of arenas and predicates. More specifically, TB indicates that we restrict to turn-based arenas, and k_{VAR} (resp. k_{AG}) bounds the maximal number of variables (resp. agents) in a formula to k . Complexity classes restrict the satisfaction problem of the predicates occurring in a formula. Finally, a slash / is used to consider the union of two fragments. For example, $\text{SG}[2\text{CG}, 3_{\text{AG}}, \Sigma_1]$ is the fragment of $\text{SL}[\text{CG}]$ with two goals, at most 3 agents, and NP-TIME-HARD predicates, whereas $\text{SG}[1\text{G}, \text{TB}, 2_{\text{AG}}/\text{VAR}]$ is the fragment of $\text{SG}[1\text{G}]$ restricted to turn-based arenas and with at most two agents or variables.

Two fragments of sentences L_1 and L_2 can be compared in terms of their expressiveness. Formally, two sentences φ_1

and φ_2 are equivalent just in the case that, for all extensions \mathcal{E} and $s \in \text{St}$, it holds that $\mathcal{E}, s \models \varphi_1$ iff $\mathcal{E}, s \models \varphi_2$. Then, we say that L_2 is *at least as expressive as* L_1 , in symbols $L_1 \preceq L_2$, if every sentence $\varphi_1 \in L_1$ is equivalent to some sentence $\varphi_2 \in L_2$. If $L_1 \preceq L_2$, but $L_2 \not\preceq L_1$, then L_2 is *more expressive than* L_1 , in symbols $L_1 \prec L_2$.

Clearly, we have that $\text{SL}[\text{BG}]$ is at least as expressive as both $\text{SL}[\text{CG}]$ and $\text{SL}[\text{DG}]$, which in turn are at least as expressive as $\text{SL}[\text{1G}]$. However, since $\text{SL}[\text{CG}]$ and $\text{SL}[\text{DG}]$ are not closed under negation, while this is the case for $\text{SL}[\text{BG}]$ and $\text{SL}[\text{1G}]$, we have that the two inclusions are necessarily strict.

THEOREM 3.1. Expressiveness Hierarchy. - *The following expressiveness relations hold: $\text{SL}[\text{1G}] \prec \text{SL}[\text{CG}]$, $\text{SL}[\text{1G}] \prec \text{SL}[\text{DG}]$, $\text{SL}[\text{CG}] \prec \text{SL}[\text{BG}]$ and $\text{SL}[\text{DG}] \prec \text{SL}[\text{BG}]$.*

Since our framework trivially subsumes classic reachability games [4], we immediately obtain an hardness result *w.r.t.* the size of the arena.

THEOREM 3.2. Arena Lower Bound. - *The fulfilling problem for $\text{SG}[\text{1G}, \text{TB}, \text{2AG}, \text{2VAR}]$ is PTIME-HARD in the size of the arena.*

We can also obtain hardness results *w.r.t.* the size of schemas. In particular, due to the first-order power of SL , we can reduce the truth problem of a QBF sentence to the fulfilling of a game on a fixed arena. Notably, the predicates used in the reduction can be evaluated in constant-time.

THEOREM 3.3. Target Lower Bound. - *The fulfilling problems for $\text{SG}[\text{CG}, \text{3AG}, \Sigma_k]$ and $\text{SG}[\text{DG}, \text{3AG}, \Pi_k]$ are Σ_{k+1}^{P} -HARD and Π_{k+1}^{P} -HARD, respectively.*

PROOF (SKETCH). We only sketch the proof for $\text{SG}[\text{CG}]$, since $\text{SG}[\text{DG}]$ is just the dual case.

Let $\alpha = \varphi \bigwedge_{i=1}^n (l_i^1 \vee l_i^2 \vee l_i^3)$ be a Σ_k QBF sentence, where, *w.l.o.g.*, the matrix is in 3CNF. Consider an arena having as agents $\text{Ag} = \{1, 2, 3\}$, actions $\text{Ac} = \{0, 1\}$, and states $\text{St} = \{\text{I}\} \cup \text{D}$, where $\text{D} = \{\delta \in \text{Dc} : \text{dom}(\delta) = \text{Ag}\}$ contains the possible eight total decisions, which are seen as sink-states. In addition, if $\delta \in \text{D}$ then $\text{tr}(\delta)(\text{I}) = \{\delta\}$ else $\text{tr}(\delta)(\text{I}) = \emptyset$.

We construct an $\text{SG}[\text{CG}]$ game $\mathfrak{G} = (\mathcal{E}, \text{I}, \varphi)$ that is fulfilled iff α is true. The extension is built on the n predicates p_1, \dots, p_n , whose valuations are set as follows: $\text{pr}(p_i) \triangleq \{\text{I} \cdot \delta : \delta \models l_i^1 \vee l_i^2 \vee l_i^3\}$. With $\delta \models l_i^1 \vee l_i^2 \vee l_i^3$ we mean that the formula $l_i^1 \vee l_i^2 \vee l_i^3$ is evaluated to true when assigning the values $\delta(j)$ to the variables x_i^j of the literal l_i^j . It only remains to define the target¹: $\varphi = \varphi \bigwedge_{i=1}^n b_i p_i$, where $b_i = (1, x_i^1)(2, x_i^2)(3, x_i^3)$. At this point, it is easy to see that $\mathcal{E}, \text{I} \models \varphi$ iff α is true. Hence, the thesis is derived from the fact that the truth problem for Σ_k QBF sentences is Σ_{k+1}^{P} -HARD. \square

4. GAME-TYPE CONVERSIONS

In this section, we describe two game-type conversions that can be used to directly solve strategy games, provided that the predicates of the extensions are decidable. Furthermore, they substantially improve and simplify the already known decision procedures for classic $\text{SL}[\text{1G}]$, $\text{SL}[\text{CG}]$, and $\text{SL}[\text{DG}]$. In particular, *w.r.t.* [10, 12], we can avoid the computational overhead resulting from the use of an automata-theoretic approach. This fact makes even clearer that, in our framework, the concept of game is considered, at the same time, as the object of study and the technical tool to analyze it.

¹Here we are making an abuse of notation since the quantification prefix should be modified to be an SL one.

4.1 From $\text{SG}[\text{1G}]$ to $\text{SG}[\text{1G}, \text{TB}, \text{2AG}/\text{VAR}]$

First we recall that an important property used as a foundation aspect of all known elementary decision procedures for the fragments of SL is the *behavioral semantics*. Roughly speaking, a strategy logic has this property whenever the choices of an existential quantified strategy does not depend on how the other universal strategies will behave in the future or in a counterfactual play [12]. The original proof of equivalence between the classic and behavioral semantics of $\text{SL}[\text{1G}]$ [10] is based on a quite complex reduction from the verification problem of the behavioral modeling relation $\mathcal{E}, s \models_B \varphi b \psi$ to the solution of a classic form of two-player turn-based games. In doing so, the size of the arena is exponential in the number of actions of the original arena and doubly exponential in the number of universal variable from which an existential one depends.

Here, we propose a new reduction from generic $\text{SG}[\text{1G}]$ to two-agent turn-based $\text{SG}[\text{1G}]$. Differently from previous reductions, the quantifications over actions are not treated in a one-shot fashion, but rather are emulated by finite games between the two players, the first choosing the value of the existential variables and the second the universal ones. This results in an arena just polynomial in the number of actions and exponential in the dependences between the variables.

THEOREM 4.1. One-Goal Reduction. - *For a $\text{SG}[\text{1G}, k\text{VAR}]$ Borelian \mathfrak{G} of order n , there is a $\text{SG}[\text{1G}, \text{TB}, \text{2AG}/\text{VAR}]$ Borelian \mathfrak{G}^* of order $n \cdot 2^{O(k)}$ such that \mathfrak{G} is fulfilled iff \mathfrak{G}^* is.*

PROOF (SKETCH). Let $\mathfrak{G} = (\mathcal{E}, s, \varphi)$, with target $\varphi = \varphi b \psi$, extension $\mathcal{E} = \langle \mathcal{A}, \text{Pr}, \text{pr} \rangle$, and arena $\mathcal{A} = \langle \text{Ag}, \text{Ac}, \text{St}, \text{tr} \rangle$, we show how to construct $\mathfrak{G}^* = \langle \mathcal{E}^*, \varphi^*, t^* \rangle$, with extension $\mathcal{E}^* = \langle \mathcal{A}^*, \text{Pr}^*, \text{pr}^* \rangle$ and arena $\mathcal{A}^* = \langle \text{Ag}^*, \text{Ac}^*, \text{St}^*, \text{tr}^* \rangle$.

We start with \mathcal{A}^* . It has two agents, \exists and \forall , the former trying to prove that $\mathcal{E}, s \models \varphi$, while the latter the opposite. To achieve their task, for each state in \mathcal{A} , they give an evaluation to the existential and universal variables, respectively, by choosing an appropriate action. Following this idea, we set $\text{Ag}^* \triangleq \{\exists, \forall\}$ and $\text{Ac}^* \triangleq \text{Ac}$. The state space has to maintain an information about the position in \mathcal{A} together with the index of the variable that has still to be evaluated and the values already associated to the previous variables. To do this, we set $\text{St}^* \triangleq \text{St} \times [0, |\varphi|] \times_i (\text{Vr}(\varphi_{<i}) \rightarrow \text{Ac})$. Observe that, when the game is in a state $(s, |\varphi|, \zeta)$, all quantifications are already resolved and it is time to evaluate the corresponding decision by composing ζ with the binding b .

Before proceeding with the definition of the transition function, it is helpful to identify which are the active agents for each possible state (s, i, ζ) . When $i < |\varphi|$ points to an existential variable in φ , *i.e.*, $\text{type}(\varphi_i) \triangleq \exists$, the unique owner of the state is \exists . Similarly, if $\text{type}(\varphi_i) \triangleq \forall$, the active agent is \forall . In the case $i = |\varphi|$, instead, there are no more choices to do, so, the related state is deterministic, *i.e.*, it has no active agents. Formally, for all $(s, i, \zeta) \in \text{St}^*$, we have that if $i < |\varphi|$ then $\text{ag}((s, i, \zeta)) \triangleq \{\text{type}(\varphi_i)\}$ else $\text{ag}((s, i, \zeta)) \triangleq \emptyset$.

The transition function is defined as follows. For each state (s, i, ζ) with $i < |\varphi|$ and decision $\text{type}(\varphi_i) \mapsto c$, we simply need to increase the counter i and associate the variable $\text{vr}(\varphi_i)$ of φ_i with action c . Formally, we set $\text{tr}^*(\text{type}(\varphi_i) \mapsto c)((s, i, \zeta)) \triangleq (s, i + 1, \zeta[\text{vr}(\varphi_i) \mapsto c])$. For a state $(s, |\pi|, \zeta)$, instead, we just define a transition to the state $(s', 0, \emptyset)$, where s' is the successor of s in the arena \mathcal{A} following the decision $\zeta \circ b$, whenever active. Formally, we have $\text{tr}^*(\emptyset) \triangleq \{(s, |\varphi|, \zeta) \mapsto (\text{tr}(\zeta \circ b)(s), 0, \emptyset) : \zeta \circ b \in \text{dc}(s)\}$.

Now, we define the extension \mathcal{E}^* . The predicates and their path valuations are simply inherited from the original extension \mathcal{E} , *i.e.*, $\text{Pr}^* \triangleq \text{Pr}$ and $\pi^* \in \text{pr}^*(p)$ iff $\pi \in \text{pr}(p)$, where $\pi \in \text{Pth}$ is the unique path in \mathcal{A} such that $\pi_i^* = (\pi_i, 0, \emptyset)$. Intuitively, a path π^* satisfies a predicate p^* iff its projection π on the states of \mathcal{A} does the same.

Finally, the initial state is $s^* \triangleq (s, 0, \emptyset)$ and the target is $\varphi^* \triangleq \langle\langle x \rangle\rangle \langle\langle y \rangle\rangle (\exists x) (\forall y) \psi$.

At this point, by means of Martin's determinacy theorem for Borelian games [9], it is not hard to show that \mathcal{D} is equivalent to \mathcal{D}^* , *i.e.*, \mathcal{D} is fulfilled iff \mathcal{D}^* is. \square




Figure 2: A full arena \mathcal{A} .

As an example, consider the $\text{SL}[1\text{G}]$ game $\mathcal{D} \triangleq \langle \mathcal{E}, s, \varphi \rangle$, whose underlying arena \mathcal{A} is the one depicted in Figure 2, having in s_0 all decisions active and s_1, s_2 as sink-states. The extension \mathcal{E} contains a unique predicate p having the path valuation $\text{pr}(p) = \{s_0 \cdot s_1\}$. Moreover, the target φ is $\wp b p$, where $\wp = \langle\langle x \rangle\rangle \langle\langle y \rangle\rangle$ and $b = (\alpha, x)(\beta, y)$. It is easy to see that \mathcal{D} is fulfilled, by means of the copy-cat strategy for β over α .




Figure 3: $\text{SL}[1\text{G}]$ reduction of \mathcal{A} .

By applying the above construction, we obtain the game \mathcal{D}^* , whose arena is reported in Figure 3. The box state belongs to \forall , while the diamond ones to \exists . Circle and pentagonal states do not belong to any agent and, in particular, the latter are sink-state. Note that, for the sake of space, we omit the state component in all but the pentagonal nodes. Similarly, the two sink-states do not report the $0, \emptyset$ part. The path valuation of the predicate p is $\text{pr}^*(p) = \{(s_0, 0, \emptyset) \cdot (s_0, 1, \zeta_1) \cdot (s_0, 2, \zeta_2) \cdot (s_1, 0, \emptyset) : \zeta_1 = \zeta_2 = \emptyset \vee \zeta_1 = x \mapsto 0 \wedge \zeta_2 = xy \mapsto 00\}$. Finally, the initial state is $s^* = (s_0, 0, \emptyset)$ and the target is $\varphi^* \triangleq \langle\langle y \rangle\rangle \langle\langle x \rangle\rangle (\exists x) (\forall y) \psi$. It can be seen that \mathcal{D}^* is fulfilled, as well. Indeed, \exists can choose \emptyset on $(s_0, 1, \emptyset)$ and $\exists \mapsto 0$ on $(s_0, 1, x \mapsto 0)$, forcing the play to reach $(s_1, 0, \emptyset)$.

4.2 From $\text{SG}[\text{CG}/\text{DG}, k\text{VAR}]$ to $\text{SG}[1\text{G}, (k+1)\text{AG}/\text{VAR}]$

Recently, it has been proved that the two $\text{SL}[1\text{G}]$ extensions $\text{SL}[\text{CG}]$ and $\text{SL}[\text{DG}]$ enjoy the behavioral semantics as well [12]. Actually, in the hierarchy obtained by restricting the combination of goals, they represent the maximal fragments of SL satisfying such a property. Indeed, there are $\text{SL}[\text{BG}]$ satisfiable sentences that are not behaviorally satisfiable [10].

As in the case of $\text{SL}[1\text{G}]$, the proof for these fragments was a direct reduction to a two-player turn-based game. However, its structure was much more complex as it further required to keep memory in the states of the information about the current position of each binding. This fact results in a state space whose cardinality was $O(|\text{St}|^h)$, where h is the number of bindings used in the related sentence.

Here, we propose a completely different approach. We reduce the fulfilling of a conjunctive/disjunctive game to the

same problem of a concurrent one-goal game. In particular, the construction simulates the conjunction (*resp.*, disjunction) of bindings by means of a dedicated new agent, whose relative strategy is universally (*resp.*, existentially) quantified. By doing so, we can avoid to keep memory of all positions at the same time, obtaining a state space whose size is $O(|\text{St}| \cdot 2^h)$.

THEOREM 4.2. Conjunctive/Disjunctive-Goal Reduction.
- For each $\text{SG}[\text{hCG}/\text{DG}, k\text{VAR}] \mathcal{D}$ of order n , there is an $\text{SG}[1\text{G}, (k+1)\text{AG}/\text{VAR}] \mathcal{D}^*$ of order $O(n \cdot 2^h)$ such that \mathcal{D} is fulfilled iff \mathcal{D}^* is.

PROOF (SKETCH). Given an $\text{SG}[\text{CG}]$ (*resp.*, $\text{SG}[\text{DG}]$) $\mathcal{D} = \langle \mathcal{E}, s, \varphi \rangle$, with target $\varphi = \wp \bigwedge_{b \in B} b\psi_b$ (*resp.*, $\varphi = \wp \bigvee_{b \in B} b\psi_b$), extension $\mathcal{E} = \langle \mathcal{A}, \text{Pr}, \text{pr} \rangle$, and arena $\mathcal{A} = \langle \text{Ag}, \text{Ac}, \text{St}, \text{tr} \rangle$, we show how to construct $\mathcal{D}^* = \langle \mathcal{E}^*, \varphi^*, t^* \rangle$, with extension $\mathcal{E}^* = \langle \mathcal{A}^*, \text{Pr}^*, \text{pr}^* \rangle$ and arena $\mathcal{A}^* = \langle \text{Ag}^*, \text{Ac}^*, \text{St}^*, \text{tr}^* \rangle$.

We start with \mathcal{A}^* . Its agents are the free variables of $\phi = \otimes_{b \in B} b\psi_b$, with $\otimes \in \{\wedge, \vee\}$, *i.e.*, those quantified in \wp , extended with the fresh agent \otimes , whose role is to simulate the corresponding Boolean combination of goals: $\text{Ag}^* \triangleq \text{free}(\phi) \cup \{\otimes\}$. The actions are the ones in \mathcal{A} augmented with a fresh one for each binding prefix in B . The latter lets the agent \otimes to choose which goal to verify: $\text{Ac}^* \triangleq \text{Ac} \cup B$. The state space is just the Cartesian product of St with the power-set of the binding prefixes set B , as we have to keep memory on which goals are still active in a given state: $\text{St}^* \triangleq \text{St} \times (2^B \setminus \{\emptyset\})$.

Before defining the transition function, it is useful to determine which decisions are active in a given state. To do this, note that we need agent \otimes to be active in every state $(s, X) \in \text{St}^*$ and let him to choose between the binding prefixes in X that have still to be verified from s onward. Moreover, the remaining part of an active decision has to ensure that, once it is composed with the selected binding prefix, it returns an active decision of \mathcal{A} in s . Formally, we set $\text{dc}((s, X)) \triangleq \{\delta^* \in \text{Dc}^* : \otimes \in \text{dom}(\delta^*) \wedge \delta^*(\otimes) \in X \wedge \delta^* \circ \delta^*(\otimes) \in \text{dc}(s)\}$.

We now set the transition function tr^* . Given an active decision $\delta^* \in \text{dc}((s, X))$ on a state $(s, X) \in \text{St}^*$, the successor state (s', X') is set as follows. The component s' is the successor of s following the decision $\delta^* \circ \delta^*(\otimes)$ obtained by assigning the value of the variables to the agents of \mathcal{A} by means of the binding prefix $\delta^*(\otimes)$ chosen by \otimes . The component X' is the subset of X containing all binding prefixes that, once selected, lead from s to s' . Formally, we have that $\text{tr}^*(\delta^*)((s, X)) \triangleq (\text{tr}(\delta^* \circ \delta^*(\otimes))(s), \{b \in X : \text{tr}(\delta^* \circ \delta^*(\otimes))(s) = \text{tr}(\delta^* \circ b)(s)\})$.

Now, we define the extension \mathcal{E}^* . The paths $\pi^* \in \text{Pth}^*$ in \mathcal{A}^* for a goal $b\psi_b$ are those having b as element of the second component of all its states, *i.e.*, $b \in X_i$, for all $i \in [0, |\pi^*|]$, where $(\pi^*)_i = (s_i, X_i)$. We denote the set of all these paths by $P_b \subseteq \text{Pth}^*$. We consider a unique predicate p^* representing the paths that satisfy the formulas ψ_b , for each binding prefix b of interest. This means that, we first set $\text{Pr}^* \triangleq \{p^*\}$. Then, for all $\pi^* \in \text{Pth}^*$, we define the predicate function pr^* as follows, where $\pi \in \text{Pth}$ is obtained from π^* by projecting out the second component of all its states. If $\otimes = \wedge$, then $\pi^* \in \text{pr}^*(p^*)$ iff, for each $b \in B$ with $\pi^* \in P_b$, it holds that $\pi \in \text{pr}(\psi_b)$. If $\otimes = \vee$, instead, $\pi^* \in \text{pr}^*(p^*)$ iff there exists $b \in B$ with $\pi^* \in P_b$ such that $\pi \in \text{pr}(\psi_b)$.

The initial state s^* is set as (s, B) . Finally, we define the target φ^* . The quantification prefix \wp^* is the juxtaposition of the prefix \wp with either a universal or an existential quantifier over the fresh variable \otimes , depending on whether \otimes is a conjunction or a disjunction. As binding prefix b^* , we use the one that associates with all agents the variables having the same

names. In particular, the type of each binding depends on the type of quantification. Hence, the target is $\varphi^* = \varphi^*b^*p^*$. \square

As an example of the above reduction, consider the SL[CG] game $\mathcal{D} = \langle \mathcal{E}, s_0, \varphi \rangle$, whose underlying arena \mathcal{A} is again the one depicted in Figure 2. In the extension $\mathcal{E} = \langle \mathcal{A}, \text{Pr}, \text{pr} \rangle$, there are two predicates p_1 and p_2 such that $\text{pr}(p_i) = \{s_0 \cdot s_i\}$, for all $i \in \{1, 2\}$. The target φ is the sentence $\llbracket x \rrbracket \langle \langle y \rangle \rangle \bigwedge_{i=1}^2 b_i p_i$, where $b_1 = (\alpha, y)(\beta, y)$ and $b_2 = (\alpha, x)(\beta, y)$. It is easy to see that \mathcal{D} is fulfilled.

Now, by applying the construction described above, we obtain an SL[1G] game \mathcal{D}^* whose arena, represented in Figure 4, has $(s_0, \{b_1, b_2\})$ as the unique state with active agents. The unique predicate p^* is associated with the following path valuation:

$\text{pr}^*(p^*) = \{(s_0, \{b_1, b_2\}) \cdot (s_1, \{b_1\}), (s_0, \{b_1, b_2\}) \cdot (s_2, \{b_2\})\}$. The initial state of \mathcal{D}^* is $(s_0, \{b_1, b_2\})$ and the target φ^* is the sentence $\varphi^*b^*p^*$, where $\varphi^* = \llbracket x \rrbracket \langle \langle y \rangle \rangle \llbracket \otimes \rrbracket$ and $b^* = (\alpha, x)(\beta, y)(\otimes, \otimes)$. It is not hard to see that also \mathcal{D}^* is fulfilled, as well.




Figure 4: SL[CG/DG] reduction of \mathcal{A} .

5. DISCUSSION

Several logics for strategic reasoning such as ATL* and SL borrows LTL operators to specify temporal properties over the evolutions of a game. This seemed a natural choice, since LTL allows to specify a widely range of conditions such as fairness, reachability, safety and so on. Nevertheless, some contexts may require more expressive languages than LTL or, conversely, if someone is interested in a specific game, say chess, it could be profitable to implement an ad-hoc procedure to check checkmates or stalemates. Summarily, LTL is not the all-inclusive solution to the verification of temporal properties.

For this reason, we revisit the classical logic of strategic reasonings and propose a new game framework which focuses on the “pure” strategic components, leaving the problem to check temporal properties to an external oracle. This, on the one hand, enables to isolate the contribution of the solely strategic modalities to the model-checking problem. In particular, we show new lower bounds which make very minimal assumptions on the complexity of the temporal part. As a side effect, the proposed framework has allowed to readdress and substantially improve previous complexity results for the SL fragments SL[1G], SL[CG], and SL[DG]. In particular for SL[1G], the exponential dependence on the number of actions occurring in formula decreases to a polynomial one and the double exponential dependence on the number of dependences between variables to a single exponential. For SL[CG] and SL[DG], instead, we have a polynomial but significant improvement in the order of the game we need to solve, since we pass from a $O(|St|^h)$ to $O(|St| \cdot 2^h)$, where h is the number of agent bindings.

Acknowledgments

We thank M. Benerecetti for useful discussions on the reduction from conjunctive/disjunctive-goal to one-goal games.

This paper is partially supported by the FP7 EU project 600958-SHERPA and the Italian projects IndAM “Logiche di Gioco Estese”, PON OR.C.HE.S.T.R.A., and POR Embedded System Cup B25B09090100007.

6. REFERENCES

- [1] T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-Time Temporal Logics with Irrevocable Strategies. In *TARK’07*, pages 15–24, 2007.
- [2] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *JACM*, 49(5):672–713, 2002.
- [3] K. Chatterjee, T.A. Henzinger, and N. Piterman. Strategy Logic. *IC*, 208(6):677–693, 2010.
- [4] N. Immerman. Number of Quantifiers is Better Than Number of Tape Cells. *JCSS*, 22(3):384–406, 1981.
- [5] W. Jamroga and W. van der Hoek. Agents that Know How to Play. *FI*, 63(2-3):185–219, 2004.
- [6] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: A Model Checker for the Verification of Multi-Agent Systems. In *CAV’09*, LNCS 5643, pages 682–688. Springer, 2009.
- [7] A.D.C. Lopes, F. Laroussinie, and N. Markey. ATL with Strategy Contexts: Expressiveness and Model Checking. In *FSTTCS’10*, LIPIcs 8, pages 120–132. Leibniz-Zentrum fuer Informatik, 2010.
- [8] E. Lorini. A Dynamic Logic of Agency II: Deterministic DLA, Coalition Logic, and Game Theory. *JLLI*, 19(3):327–351, 2010.
- [9] A.D. Martin. Borel Determinacy. *AM*, 102(2):363–371, 1975.
- [10] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. Reasoning About Strategies: On the Model-Checking Problem. Technical report, arXiv, 2011.
- [11] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. What Makes ATL* Decidable? A Decidable Fragment of Strategy Logic. In *CONCUR’12*, LNCS 7454, pages 193–208. Springer, 2012.
- [12] F. Mogavero, A. Murano, and L. Sauro. On the Boundary of Behavioral Strategies. In *LICS’13*, pages 263–272. IEEE Computer Society, 2013.
- [13] F. Mogavero, A. Murano, and M.Y. Vardi. Reasoning About Strategies. In *FSTTCS’10*, LIPIcs 8, pages 133–144. Leibniz-Zentrum fuer Informatik, 2010.
- [14] D. Perrin and J. Pin. *Infinite Words*. Pure and Applied Mathematics. Elsevier, 2004.
- [15] A. Pnueli. The Temporal Logic of Programs. In *FOCS’77*, pages 46–57. IEEE Computer Society, 1977.
- [16] W. van der Hoek, W. Jamroga, and M. Wooldridge. A Logic for Strategic Reasoning. In *AAMAS’05*, pages 157–164. Association for Computing Machinery, 2005.
- [17] J. van Eijck. PDL as a Multi-Agent Strategy Logic. In *TARK’13*, pages 206–215, 2013.
- [18] M.Y. Vardi. A Temporal Fixpoint Calculus. In *POPL’88*, pages 250–259. Association for Computing Machinery, 1988.
- [19] D. Walther, W. van der Hoek, and M. Wooldridge. Alternating-Time Temporal Logic with Explicit Strategies. In *TARK’07*, pages 269–278, 2007.
- [20] P. Wolper. Temporal Logic Can Be More Expressive. *IC*, 56(1-2):72–99, 1983.