On the Counting of Strategies

Vadim Malvone, Fabio Mogavero, Aniello Murano, Loredana Sorrentino
Università degli Studi di Napoli Federico II

***Abstract*—In game theory, a classic qualitative question is to check whether a designated set of players has a winning strategy. In several safety-critical applications, however, it is important to ensure that some redundant strategies also exist, to be possibly used in case of some fault.**

In this paper, we introduce *Graded Strategy Logic* (GSL), an extension of *Strategy Logic* (SL) with *graded quantifiers*. SL is a powerful formalism that allows to describe useful game concepts in multi-agent settings by explicitly quantifying over strategies treated as first-order citizens. In GSL, by means of the existential construct $\langle\langle x \geq g \rangle\rangle \varphi$ one can enforce that there exist at least g strategies satisfying φ. Dually, via the universal construct $[[x < g]]\varphi$ one can ensure that all but less than g strategies satisfy φ.

As different strategies may induce the same outcome, although looking different, they need to be counted as one. While this interpretation is natural, it heavily complicates the definition and thus the reasoning about GSL. In order to accomplish this specific way of counting, we formally introduce a suitable equivalence relation over profiles based on the strategic behavior they induce.

To give evidence of GSL usability, we investigate basic questions of one of its *vanilla* fragment, namely GSL[1G]. In particular, we report on positive results about the determinacy of games and the related model-checking problem, which we show to be PTIME-COMPLETE.

I. Introduction

Formal methods in system design are a renowned story of success. Breakthrough contributions in this field comprise *temporal logics*, such as LTL [Pnu77], CTL [CE81], or CTL* [EH86], and *model checking* [CE81], [QS81]. First applications of these methodologies involved *closed systems* [HP85] generally analyzing whether a Kripke structure, modeling the system, meets a temporal logic formula, specifying the desired behavior [CGP02]. In the years several algorithms have been proposed in this setting and some implemented as tools [BBF$^+$10]. Nevertheless these approaches turn to be useless when applied to *open systems* [HP85]. The latter are characterized, in the simplest situation, by an ongoing interaction with an external environment on which the whole system behavior deeply relies. To be able to deal with the unpredictability of the environment, extensions of the basic verification techniques have come out. A first attempt worth of note is *module checking* where a Kripke structure is replaced by a specific two-player arena. Module checking has been first introduced in [KV96], [KVW01]. In the last decade this methodology has been fruitfully extended in several directions (see [ALM$^+$13], [JM14], [JM15] for a list of related works).

Starting from the study of module checking, researchers have looked for logics focusing on the *strategic behavior* of players in *multi-agent systems* [AHK02]. One of the most important developments in this field is *Alternating-Time Temporal Logic* (ATL*, for short), introduced by Alur, Henzinger, and Kupferman [AHK02]. This logic allows to reason about strategies of agents having the satisfaction of temporal goals as payoff criterion. Formally, it is obtained as a generalization of CTL*, in which the existential E and the universal A *path quantifiers* are replaced with *strategic modalities* of the form $\langle\langle \mathrm{A} \rangle\rangle$ and $[[\mathrm{A}]]$, where A is a set of *agents*. Strategic modalities over agent teams are used to describe cooperation and competition among them in order to achieve certain goals. In particular, these modalities express selective quantifications over those paths that are the result of infinite interaction between a coalition and its complement.

Despite its expressiveness, ATL* suffers from the strong limitation that strategies are treated only implicitly in the semantics of such modalities. This restriction makes the logic less suited to formalize several important solution concepts, such as *Nash Equilibrium*. These considerations led to the introduction of *Strategy Logic* (SL, for short) [CHP07], [MMV10], a more powerful formalism for strategic reasoning. As a key aspect, this logic treats strategies as *first-order objects* that can be determined by means of the existential $\langle\langle x \rangle\rangle$ and universal $[[x]]$ quantifiers, which can be respectively read as *"there exists a strategy x"* and *"for all strategies x"*. Remarkably, a strategy in SL is a generic conditional plan that at each step prescribes an action on the base of the history of the play. Such a plan is not intrinsically glued to a specific agent but an explicit binding operator (a, x) allows to link an agent a to the strategy associated with a variable x.

A common aspect about all logics mentioned above is that quantifications are either existential or universal. *Per contra*, there are several real scenarios in which "more precise" quantifications are crucially needed (see [BMM12], [MMS15], for an argumentation). This has attracted the interest of the formal verification community to *graded modalities*. They have been first studied in classic modal logic [Fin72] and then exported to the field of *knowledge representation* to allow quantitative bounds on the set of individuals satisfying a certain property. In particular, they are considered as *counting quantifiers* in first-order logics [GOR97] and *number restrictions* in *description logics* [HB91].

First applications of graded modalities in formal verification concern closed systems. In [KSV02], *graded* μCALCULUS has been introduced in order to express statements about a given number of immediately accessible worlds. Successively in [FNP09], [BMM09], [BMM10], [BMM12], the notion of graded modalities have been extended to deal with number of paths. Among the others graded CTL (GCTL, for short) has been introduced with a suitable axiomatization of counting [BMM12].

In open systems verification, we are aware of just two orthogonal approaches in which graded modalities have been investigated, but in a very restricted form: module checking for graded μCALCULUS [FMP08] and an extension of ATL with graded path modalities (GATL, for short) [FNP10]. In particular, the former involves a counting of one-step moves among two agents, the latter allows for a more restricted counting on the histories of the game, but in a multi-player setting. Both approaches suffer of several limitations. First, not surprisingly, they cannot express powerful game reasoning due to the limitation of the underlying logic. Second, it is based on a very rigid and restricted counting of strategies.

In this paper, we take a completely different approach by formally introducing a solid machinery to count strategies in a multi-agent setting and use it upon the powerful framework of SL. Precisely, we introduce and study *Graded Strategy Logic* (GSL) which extends SL with the existential $\langle\langle x \geq g \rangle\rangle \varphi$ and universal $[[x < g]] \varphi$ graded strategy quantifiers. They allow to express that there are *at least* g or *all but less than* g strategies x satisfying φ, respectively. Then, by using the classical binding operator of SL, it is possible to associate these strategies to specific agents.

As far as the counting of strategies concerns, one of the main difficulties resides on the fact that some strategies, although looking different, produce the same outcome and therefore have to be counted as one. To overcome this problem while preserving a correct counting over paths for the underlining logic SL, we formally introduce a suitable equivalence relation over profiles based on the strategic behavior they induce. This is by its own an important contribution of this paper.

To show the applicability of GSL we positively investigate basic game-theoretic and verification questions over a powerful fragment of GSL. Recall that model checking is non-elementary-complete for SL and this has spurred researchers to investigate fragments of the logic for practical applications. Here, we concentrate on the *vanilla* version of the SL[1G] fragment of SL. We recall that SL[1G] was introduced in [MMPV12]. As for ATL, vanilla SL[1G] (for the first time introduced here) requires that two successive temporal operators in a formula are always interleaved by a strategy quantifier. We prove that the model-checking problem for this logic is PTIME-COMPLETE. We also show positive results about the determinacy of turn-based games.

GSL can have useful applications in several multi-agent game scenarios. For example, in safety-critical systems, it may be worth knowing whether a controller agent has a redundant winning strategy to play in case of some fault. Having more than a strategy may increase the chances for a success [ATO$^+$09]. Such a redundancy can easily be expressed in GSL by requiring that at least two different strategies exist for the achievement of the safety goal. The universal graded strategy quantifier may turn useful to grade the "security" of a system. For example, one can check whether preventing the use of at most k strategies, the remaining ones are all winning. In a network this may correspond to prevent some attacks while leaving the communication open.

Due to the lack of space, all proofs are omitted. We also refer to [MMPV14] for an introduction to SL and SL[1G].

Outline: The sequel of the paper is structured as follows. In Section II, we introduce GSL and provide some preliminary related concepts. In Section III, we describe a customized equivalence relation to count strategies by means of several axioms, one for each operator of GSL[1G]. In Section IV, we address the determinacy and the model-checking problem for the vanilla GSL[1G] fragment of GSL. Finally we conclude in Section V by giving some discussion and future works.

II. GRADED STRATEGY LOGIC

In this section, we introduce GSL. As stated in the Introduction, GSL extends SL to allow for reasoning about the number of winning strategies for the agents. We also recall that SL extends LTL with two strategy quantifiers and a binding operator to associate agents to strategies.

A. Model

Similarly to SL, as semantic framework we use a *game structure* [AHK02] in which the system is modeled as a game where players perform actions chosen as a function on the history of the play.

Definition II.1 (Game Structure). *A* game structure *is a tuple* $\mathcal{G} \triangleq \langle \mathrm{AP}, \mathrm{Ag}, \mathrm{Ac}, \mathrm{St}, \mathsf{tr}, \mathsf{ap}, s_I \rangle$, *where* AP *is a set of* atomic propositions, Ag, Ac, *and* St *are finite non-empty sets of* agents, actions *and* states, *respectively,* $s_I \in \mathrm{St}$ *is an* initial state, *and* $\mathsf{ap} : \mathrm{St} \to 2^{\mathrm{AP}}$ *is a* labeling function *mapping each state to the set of atomic propositions true in that state. Let* $\mathrm{Dc} \triangleq \mathrm{Ag} \rightharpoonup \mathrm{Ac}$ *be the set of* decisions, i.e., *partial functions describing the choices of an action by some agent. Then,* $\mathsf{tr} : \mathrm{Dc} \to (\mathrm{St} \rightharpoonup \mathrm{St})$ *denotes the* transition function *mapping every decision* $\delta \in \mathrm{Dc}$ *to a partial function* $\mathsf{tr}(\delta) \subseteq \mathrm{St} \times \mathrm{St}$ *representing a deterministic graph over the states.*

A game structure $\mathcal{G}$ naturally induces a graph $\langle \mathrm{St}, Ed \rangle$ with $Ed = \bigcup_{\delta \in \mathrm{Dc}} \mathsf{tr}(\delta)$, where the infinite paths starting at the initial state s_I represent all possible *plays* (whose set is denoted by Pth) and its finite paths are called *histories* (whose set is denoted by Hst). A *strategy* is a function $\sigma \in \mathrm{Str} \triangleq \mathrm{Hst} \to \mathrm{Ac}$ prescribing which action has to be performed given a certain history. We say that $\sigma \in \mathrm{Str}(\mathrm{A}) \subseteq \mathrm{Str}$ is A-*coherent w.r.t.* a set of agents A if $\sigma(\rho \cdot s) \in \mathrm{Ac}$, for all histories $\rho \cdot s \in \mathrm{Hst}$ and agents $a \in \mathrm{A}$. We assume, *w.l.o.g.* that, in each state, all agents can always take an action. Hence, we assume no *end states*.

As a running example, consider the game structure $\mathcal{G}_S$ depicted in Figure 1. It models a *scheduler system* that comprises three agents, including two *processes*, $\mathtt{P_1}$ and $\mathtt{P_2}$, willing to access a *shared resource* (such as a processor), and an *arbiter* $\mathtt{A}$ used to solve conflicts arisen under contending requests. The processes can use four actions: $\mathtt{i}$ for idle, which means that the process does not want to change the current situation in which the entire system resides, $\mathtt{r}$ for (resource) request, used to ask the resource, when this is not yet owned, $\mathtt{f}$ for free (a resource), used to release the resource, when this

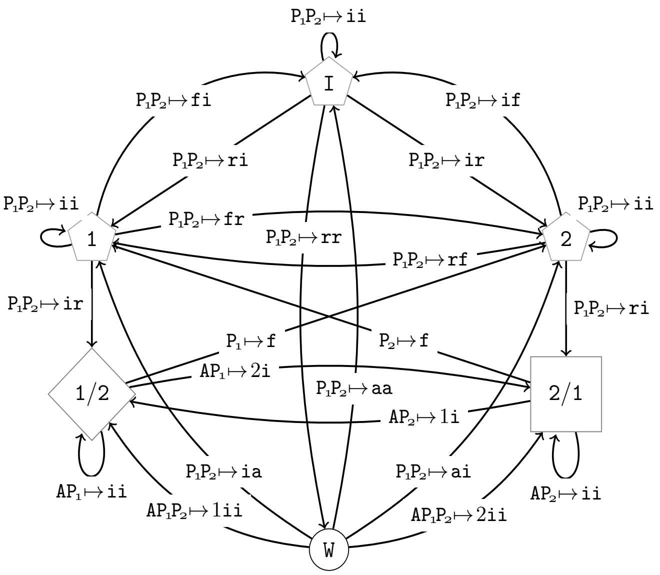

Figure 1. A scheduler system $\mathcal{G}_S$.

is yet owned, and a for abandon (a pending request), that is asserted by a process that, although has asked for the resource, did not obtain it and so it decides to relinquish the request. The system can reside in the states I, 1, 2, 1/2, 2/1 and W. The first three are ruled by the processes, the last by all the agents, and 1/2 (resp, 2/1) by $\mathtt{P}_1$ (resp., $\mathtt{P}_2$) and A. The idle state I indicates that none of the processes owns the resource, while a state $k \in \{1, 2\}$ asserts that process $\mathtt{P}_k$ is using it. The state 1/2 (resp. 2/1) indicates that the process $\mathtt{P}_1$ (resp., $\mathtt{P}_2$) has the resource, while its competitor requires it. Finally, the waiting state W represents the case in which an action from the arbiter is required in order to solve a conflict. To denote who is the owner of the resource, we label 1 and 1/2 (resp., 2 and 2/1) with the atomic proposition $\mathtt{r}_1$ (resp., $\mathtt{r}_2$). A decision is graphically represented by $\vec{\mathtt{a}} \mapsto \vec{\mathtt{c}}$, where $\vec{\mathtt{a}}$ is a sequence of agents and $\vec{\mathtt{c}}$ is a sequence of corresponding actions. For example $\mathtt{P}_1\mathtt{P}_2 \rightarrow \mathtt{ir}$ indicates that agents $\mathtt{P}_1$ and $\mathtt{P}_2$ take actions i and r, respectively.

B. Syntax

GSL extends SL by replacing its universal and existential *strategy quantifiers* $\langle\!\langle x \rangle\!\rangle$ and $[\![x]\!]$, where x belongs to a countable set of variables Vr, with their graded version $\langle\!\langle x \geq g \rangle\!\rangle$ and $[\![x < g]\!]$, in which the finite number $g \in \mathbb{N}$ denotes the corresponding *degree*. Intuitively, these quantifiers are read as *"there exist at least g strategies"* and *"all but less than g strategies"*, respectively.

Definition II.2 (GSL Syntax)**.** GSL formulas *are built inductively by means of the following context-free grammar, where* $a \in \text{Ag}$, $x \in \text{Vr}$, *and* $g \in \mathbb{N}$*:*

$$\varphi := \text{LTL}(\varphi) \mid \langle\!\langle x \geq g \rangle\!\rangle\varphi \mid [\![x < g]\!]\varphi \mid (a, x)\varphi.$$

By LTL(φ) we mean the set of LTL formulas. As usual, to provide the semantics of a predicative logic, it is necessary to define the concept of free and bound *placeholders* of a formula. As for SL, since strategies can be associated to both agents and variables, we need the set of *free agents/variables* $\mathsf{free}(\varphi)$ 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 temporal operator and *(ii)* all variables x for which there is a binding (a, x) but no quantification $\langle\!\langle x \geq g \rangle\!\rangle$ or $[\![x < g]\!]$. A detailed definition can be found in [MMPV14]. In case $\mathsf{free}(\varphi) = \emptyset$ the formula φ is named *sentence*. Since a variable x may be bound to more than one agent at the time, we also need the subset $\mathsf{shr}(\varphi, x)$ of Ag containing those agents for which a binding (a, x) occurs in φ.

For complexity reasons, we restrict to the *One-Goal* fragment of GSL (GSL[1G], for short), which is the graded extension of SL[1G] [MMPV14]. To formalize its syntax, we first introduce some notions. A *quantification prefix* over a set $\text{V} \subseteq \text{Vr}$ of variables is a word $\wp \in \{\langle\!\langle x \geq g \rangle\!\rangle, [\![x < g]\!] : x \in \text{V} \wedge g \in \mathbb{N}\}^{|\text{V}|}$ of length $|\text{V}|$ such that each $x \in \text{V}$ occurs just once in $\wp$. A *binding prefix* over $\text{A} \subseteq \text{Ag}$ is a word $\flat \in \{(a, x) : a \in \text{A} \wedge x \in \text{Vr}\}^{|\text{A}|}$ such that each $a \in \text{A}$ occurs exactly once in $\flat$. GSL[1G] restricts GSL by forcing, after a quantification prefix, a single goal to occur *i.e.*, a formula of the kind $\flat\psi$, where $\flat$ is a binding prefix on all the agent Ag. We now give the syntax of GSL[1G].

Definition II.3 (GSL[1G] Syntax)**.** GSL[1G] formulas *are built inductively through the following grammar:*

$$\varphi := \text{LTL}(\varphi) \mid \wp\flat\varphi,$$

with $\wp$ *quantification prefix over* $\mathsf{free}(\flat\varphi)$ *and* $\flat\varphi$ *a goal.*

An example of a GSL[1G] property, in the context of the scheduler system, is given by the sentence $\varphi = \wp\flat\psi$, with $\wp = \langle\!\langle \mathtt{x} \geq k \rangle\!\rangle [\![\mathtt{y}_1 < g_1]\!] [\![\mathtt{y}_2 < g_2]\!]$, $\flat = (\mathtt{A}, \mathtt{x})(\mathtt{P}_1, \mathtt{y}_1)(\mathtt{P}_2, \mathtt{y}_2)$, and $\psi = \mathtt{F}(\mathtt{r}_1 \vee \mathtt{r}_2)$. It states the existence of at least k strategies for the arbiter A ensuring that one of the two processes $\mathtt{P}_1$ and $\mathtt{P}_2$ receives the resource, once less than g_1 and g_2 strategies can be avoided by them, respectively.

C. Semantics

As for SL, the interpretation of a GSL formula requires a valuation for its free placeholders. This is done via *assignments*, *i.e.*, partial functions $\chi \in \text{Asg} \triangleq (\text{Vr} \cup \text{Ag}) \rightharpoonup \text{Str}$ mapping variables/agents to strategies. By $\text{Asg}(\text{X}) \subseteq \text{Asg}$ we denote the set of assignments over $\text{X} \subseteq \text{Vr} \cup \text{Ag}$.

An assignment χ is *complete* if it is defined on all agents in Ag, *i.e.*, $\chi(a) \in \text{Str}(\{a\})$, for all $a \in \text{Ag} \subseteq \mathsf{dom}(\chi)$. In this case, it directly identifies the profile $\chi_{\restriction \text{Ag}}$ given by the restriction of χ to Ag. In addition, $\chi[e \mapsto \sigma]$, with $e \in \text{Vr} \cup \text{Ag}$ and $\sigma \in \text{Str}$, denotes the assignment defined on $\mathsf{dom}(\chi[e \mapsto \sigma]) \triangleq \mathsf{dom}(\chi) \cup \{e\}$ that differs from χ only on 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 $e' \in \mathsf{dom}(\chi) \setminus \{e\}$. Finally, for a formula φ, we say that χ is φ*-coherent* iff *(i)* $\mathsf{free}(\varphi) \subseteq \mathsf{dom}(\chi)$, *(ii)* $\chi(a) \in \text{Str}(\{a\})$, for all $a \in \mathsf{dom}(\chi) \cap \mathrm{Ag}$, and *(iii)* $\chi(x) \in \mathrm{Str}(\mathsf{shr}(\varphi, x))$, for all $x \in \mathsf{dom}(\chi) \cap \mathrm{Vr}$.

We now define the semantics of a GSL formula φ *w.r.t.* a game structure $\mathcal{G}$ and a φ-coherent assignment χ. In particular, we write $\mathcal{G}, \chi \models \varphi$ to indicate that φ holds in $\mathcal{G}$ under χ. The semantics of LTL formulas and agent bindings are defined as in SL. The definition of graded strategy quantifiers, instead, makes use of a family of equivalence relations $\equiv_{\mathcal{G}}^{\varphi}$ on assignments that depend on the structure $\mathcal{G}$ and the formula φ under examination. This equivalence is used to reasonably count the number of strategies that satisfy a formula *w.r.t.* an *a priori* fixed criterion. Observe that we use a relation on assignments instead of a more direct one on strategies, since the classification may also depend on the context determined by the strategies previously quantified. In Section III, we will come back on the properties the equivalence has to satisfy in order to be used in the semantics of GSL.

Definition II.4 (GSL Semantics)**.** *Let $\mathcal{G}$ be a Game Structure and φ a* GSL *formula. For all φ-coherent assignments $\chi \in$* Asg*, the relation $\mathcal{G}, \chi \models \varphi$ is inductively defined as follows.*

1) *All* LTL *operators are interpreted as usual.*
2) *For each $x \in \mathrm{Vr}$, $g \in \mathbb{N}$, and $\varphi \in$* GSL*, it holds that:*
   a) $\mathcal{G}, \chi \models \langle\langle x \geq g \rangle\rangle \varphi$ *iff* $|(\{\chi[x \mapsto \sigma] : \sigma \in \varphi[\mathcal{G}, \chi](x)\}/{\equiv_{\mathcal{G}}^{\varphi}})| \geq g$;
   b) $\mathcal{G}, \chi \models [\![x < g]\!] \varphi$ *iff* $|(\{\chi[x \mapsto \sigma] : \sigma \in \neg\varphi[\mathcal{G}, \chi](x)\}/{\equiv_{\mathcal{G}}^{\neg\varphi}})| < g$;

   where $\eta[\mathcal{G}, \chi](x) \triangleq \{\sigma \in \mathrm{Str}(\mathsf{shr}(\eta, x)) : \mathcal{G}, \chi[x \mapsto \sigma] \models \eta\}$ is the set of $\mathsf{shr}(\eta, x)$-coherent strategies that, being assigned to x in χ, satisfy η.
3) *For each $a \in$* Ag*, $x \in$* Vr*, and $\varphi \in$* GSL*, it holds that $\mathcal{G}, \chi \models (a, x)\varphi$ iff $\mathcal{G}, \chi[a \mapsto \chi(x)] \models \varphi$.*

Intuitively, the existential quantifier $\langle\langle x \geq g \rangle\rangle \varphi$ allows us to count the number of equivalence classes *w.r.t.* $\equiv_{\mathcal{G}}^{\varphi}$ over the set of assignments $\{\chi[x \mapsto \sigma] : \sigma \in \varphi[\mathcal{G}, \chi](x)\}$ that, extending χ, satisfy φ. The universal quantifier $[\![x < g]\!] \varphi$ is the dual of $\langle\langle x \geq g \rangle\rangle \varphi$ and counts how many classes *w.r.t.* $\equiv_{\mathcal{G}}^{\neg\varphi}$ there are over the assignments $\{\chi[x \mapsto \sigma] : \sigma \in \neg\varphi[\mathcal{G}, \chi](x)\}$ that, extending χ, do not satisfy φ. It is worth noting that all GSL formulas with degree 1 are SL formulas. Also, the verification of a sentence φ does not depend on assignments, so, we just write $\mathcal{G} \models \varphi$.

Consider again the sentence $\varphi = \langle\langle \mathtt{x} \geq k \rangle\rangle [\![\mathtt{y}_1 < g_1]\!] [\![\mathtt{y}_2 < g_2]\!] (\mathtt{A}, \mathtt{x})(\mathtt{P}_1, \mathtt{y}_1)(\mathtt{P}_2, \mathtt{y}_2)\mathtt{F}(\mathtt{r}_1 \vee \mathtt{r}_2)$ of the scheduler example. Once a reasonable equivalence relation on assignments is fixed (see Section III), one can see that $\mathcal{G}_S \models \varphi$ with $\mathtt{k} \geq 0$ and $(g_1, g_2) = (1, 2)$ but $\mathcal{G}_S \not\models \varphi$ with $(k, g_1, g_2) = (1, 1, 1)$. Indeed, if the processes use the same strategy, they may force the play to be in $(\mathtt{I}^+ \cdot \mathtt{W})^* \cdot \mathtt{I}^\omega + (\mathtt{I}^+ \cdot \mathtt{W})^\omega$, so they either avoid to do a request or relinquish a request that is not immediately served. Consequently, to satisfy φ, we need to verify the property against all but one strategy of $\mathtt{P}_2$, *i.e.*, the one used by $\mathtt{P}_1$. Under these assumptions, we can see that the arbiter $\mathtt{A}$ has an infinite number of different strategies by suitably choosing the actions on all histories ending in the state $\mathtt{W}$.

Before continuing, we show how graded ATL [FNP10] can be translated to GSL[1G]. In [FNP10], the authors introduce two different semantics for their logic, called *off-line* and *on-line*. Under the first one, over a game structure with agents α and $\overline{\alpha}$, the graded ATL formula $\langle\langle \alpha \rangle\rangle^g \psi$ is equivalent to the GSL[1G] sentence $\langle\langle \mathtt{x} \geq g \rangle\rangle [\![\overline{\mathtt{x}} < 1]\!] (\alpha, \mathtt{x})(\overline{\alpha}, \overline{\mathtt{x}})\psi$. Under the second semantics, instead, it is equivalent to the sentence $[\![\overline{\mathtt{x}} < 1]\!] \langle\langle \mathtt{x} \geq g \rangle\rangle (\alpha, \mathtt{x})(\overline{\alpha}, \overline{\mathtt{x}})\psi$. Note that the counting over strategies in graded ATL is limited to existential agents and, so, the SL[1G] formula $[\![\overline{\mathtt{x}} < 2]\!] \langle\langle \mathtt{y} \geq 1 \rangle\rangle (\alpha, \mathtt{x})(\overline{\alpha}, \mathtt{y})\psi$ does not have any ATL equivalent. Moreover, the criteria used for the strategy classification is strictly coupled with the temporal operators $\mathtt{X}\varphi$, $\varphi_1 \mathtt{U} \varphi_2$, and $\mathtt{G}\varphi$ along the syntax, and we do not see how this can be extended to the whole LTL, unless one uses the approach proposed in [BMM12].

III. Strategy Equivalence

Our definition of GSL semantics makes use of an arbitrary family of equivalence relation on assignments. This choice introduces flexibility in its description, since one can come up with different logics by opportunely choosing different equivalences.

In this section, we focus on a particular relation whose key feature is to classify as equivalent all assignments that reflect the same "*strategic reasoning*", although they may have completely different structures. Just to get an intuition about what we mean, consider two assignments χ_1 and χ_2 and the corresponding involved strategies associated with the agents $\mathtt{a}_1$ and $\mathtt{a}_2$. Assume now that, for each $i \in \{1, 2\}$, the homologous strategies $\chi_1(\mathtt{a}_i)$ and $\chi_2(\mathtt{a}_i)$ only differ on histories never met by a play because of a specific combination of their actions. Clearly, χ_1 and χ_2 induce the same agent behaviors, which means to reflect the same strategic reasoning. Therefore, it is natural to set them as equivalent, as we do. In addition, take two equivalent formulas. We have that either two assignments are equivalent for both formulas or for none of them. Furthermore, if two assignments do not satisfy the same formulas, they are not equivalent.

In the sequel, in order to illustrate the introduced concepts, we analyze subformulas of the previously described sentence $\langle\langle \mathtt{x} \geq k \rangle\rangle [\![\mathtt{y}_1 < 1]\!] [\![\mathtt{y}_2 < 2]\!] (\mathtt{A}, \mathtt{x})(\mathtt{P}_1, \mathtt{y}_1)(\mathtt{P}_2, \mathtt{y}_2)\mathtt{F}(\mathtt{r}_1 \vee \mathtt{r}_2)$, together with their negations, over the game structure $\mathcal{G}_S$ of Figure 1.

A. Elementary Requirements

Logics usually admit syntactic redundancy. For example, in LTL we have $\neg\mathtt{X}(\mathtt{p} \wedge \mathtt{q}) \equiv \mathtt{X}\neg(\mathtt{p} \wedge \mathtt{q}) \equiv \mathtt{X}(\neg\mathtt{p} \vee \neg\mathtt{q})$. Also, the semantics is normally closed under substitution. Yet for LTL, this means that $\neg\mathtt{X}(\mathtt{p} \wedge \mathtt{q})$ can be replaced with $\mathtt{X}\neg(\mathtt{p} \wedge \mathtt{q})$ or $\mathtt{X}(\neg\mathtt{p} \vee \neg\mathtt{q})$, without changing the meaning of a formula. GSL should not be an exception. To ensure this, we require the invariance of the equivalence relation on assignments *w.r.t.* the syntax of the involved formulas.

Definition III.1 (Syntax Independence)**.** *An equivalence relation on assignments $\equiv_{\mathcal{G}}^{\cdot}$ is* syntax independent *if, for any pair* *of equivalent formulas* φ_1 *and* φ_2 *and* $(\mathsf{free}(\varphi_1) \cup \mathsf{free}(\varphi_2))$*-coherent assignments* $\chi_1, \chi_2 \in \mathrm{Asg}$*, we have that* $\chi_1 \equiv_{\mathcal{G}}^{\varphi_1} \chi_2$ *iff* $\chi_1 \equiv_{\mathcal{G}}^{\varphi_2} \chi_2$.

As declared above, our aim is to classify as equivalent *w.r.t.* a formula φ all assignments that induce the same strategic reasoning. Therefore, we cannot distinguish them *w.r.t.* the satisfiability of φ itself.

Definition III.2 (Semantic Consistency)**.** *An equivalence relation on assignments* $\equiv_{\mathcal{G}}^{\cdot}$ *is* semantically consistent *if, for any formula* φ *and* φ*-coherent assignments* $\chi_1, \chi_2 \in \mathrm{Asg}$*, we have that if* $\chi_1 \equiv_{\mathcal{G}}^{\varphi} \chi_2$ *then either* $\mathcal{G}, \chi_1 \models \varphi$ *and* $\mathcal{G}, \chi_2 \models \varphi$ *or* $\mathcal{G}, \chi_1 \not\models \varphi$ *and* $\mathcal{G}, \chi_2 \not\models \varphi$*.*

B. Play Requirement

We now deal with the equivalence relation for the basic case of temporal properties. Before disclosing the formalization, we give an intuition on how to evaluate the equivalence of two complete assignments χ_1 and χ_2 *w.r.t.* their agreement on the verification of a generic LTL property ψ. Let π_1 and π_2 with $\pi_1 \neq \pi_2$ be the plays satisfying ψ induced by χ_1 and χ_2, respectively. Also, consider their maximal common prefix $\rho = \mathsf{prf}(\pi_1, \pi_2) \in \mathrm{Hst}$. If ρ can be extended to a play in such a way that ψ does not hold, we are sure that the reasons why both the assignments satisfy the property are different, as they reside in the parts where the two plays diverge. Consequently, we can assume χ_1 and χ_2 to be non-equivalent *w.r.t.* ψ. Conversely, if all infinite extensions of ρ necessarily satisfy ψ, we may affirm that this is already a witness of the verification of the property by the two plays and, so, by the two assignments. Hence, we can assume χ_1 and χ_2 to be equivalent *w.r.t.* ψ.

In the following, we often make use of the concept of witness of an LTL formula ψ as the set $\mathrm{W}_{\psi} \triangleq \{\rho \in \mathrm{Hst} : \forall \pi \in \mathrm{Pth} \,.\, \rho < \pi \Rightarrow \pi \models \psi\}$ containing all histories that cannot be extended to a play violating the property.

Definition III.3 (Play Consistency)**.** *An equivalence relation on assignments* $\equiv_{\mathcal{G}}^{\cdot}$ *is* play consistent *if, for any* LTL *formula* ψ *and* ψ*-coherent assignments* $\chi_1, \chi_2 \in \mathrm{Asg}$*, we have that* $\chi_1 \equiv_{\mathcal{G}}^{\psi} \chi_2$ *iff either* $\pi_1 = \pi_2$ *or* $\mathsf{prf}(\pi_1, \pi_2) \in \mathrm{W}_{\psi}$*, where* $\pi_1 = \mathsf{play}(\chi_{1 \restriction \mathrm{Ag}})$ *and* $\pi_2 = \mathsf{play}(\chi_{2 \restriction \mathrm{Ag}})$ *are the plays induced by* χ_1 *and* χ_2*, respectively, and* $\mathrm{W}_{\psi} \subseteq \mathrm{Hst}$ *is the witness set of* ψ*.*

To see how to apply the above definition, consider the formula $\psi = \mathtt{F}(\mathtt{r}_1 \vee \mathtt{r}_2)$ and let W_{ψ} be the corresponding witness set, whose minimal histories can be represented by the regular expression $\mathtt{I}^{+} \cdot (1+2) + (\mathtt{I}^{+} \cdot \mathtt{W})^{+} \cdot (1+2+1/2+2/1)$. Moreover, let $\chi_1, \chi_2, \chi_3 \in \mathrm{Asg}(\{\mathtt{A}, \mathtt{P}_1, \mathtt{P}_2\})$ be three complete assignments on which we want to check the play consistency. We assume that each χ_i associates a strategy $\chi_i(\mathtt{a}) = \sigma_i^{\mathtt{a}}$ with the agent $\mathtt{a} \in \{\mathtt{A}, \mathtt{P}_1, \mathtt{P}_2\}$ as defined in the following, where $\rho, \rho_{\mathtt{s}} \in \mathrm{Hst}$ with $\mathsf{lst}(\rho) \neq \mathtt{I}$ and $\rho_{\mathtt{s}} \cdot \mathtt{s} \in \mathrm{Hst}$: for the arbiter $\mathtt{A}$, we set $\sigma_{1/2}^{\mathtt{A}}(\rho_{\mathtt{W}} \cdot \mathtt{W}) \triangleq 2$, $\sigma_{1/2/3}^{\mathtt{A}}(\rho_{1/2} \cdot 1/2) = \sigma_{2}^{\mathtt{A}}(\rho_{2/1} \cdot 2/1) \triangleq \mathtt{i}$, and $\sigma_{3}^{\mathtt{A}}(\rho_{\mathtt{W}} \cdot \mathtt{W}) = \sigma_{1/3}^{\mathtt{A}}(\rho_{2/1} \cdot 2/1) \triangleq 1$; for the processes, instead, we set $\sigma_{1/2/3}^{\mathtt{P}_1}(\rho) = \sigma_{1/2/3}^{\mathtt{P}_2}(\rho) \triangleq \mathtt{i}$, $\sigma_{1/2}^{\mathtt{P}_1}(\rho_{\mathtt{I}} \cdot \mathtt{I}) = \sigma_{1/2/3}^{\mathtt{P}_2}(\rho_{\mathtt{I}} \cdot \mathtt{I}) \triangleq \mathtt{r}$, and $\sigma_{3}^{\mathtt{P}_1}(\rho_{\mathtt{I}} \cdot \mathtt{I}) \triangleq \mathtt{i}$. Now, one can see that $\chi_1 \equiv_{\mathcal{G}}^{\psi} \chi_2$, but $\chi_1 \not\equiv_{\mathcal{G}}^{\psi} \chi_3$.

Indeed, χ_1, χ_2, and χ_3 induce the plays $\pi_1 = \mathtt{I} \cdot \mathtt{W} \cdot 2/1 \cdot 1/2^{\omega}$, $\pi_2 = \mathtt{I} \cdot \mathtt{W} \cdot 2/1^{\omega}$, and $\pi_3 = \mathtt{I} \cdot 2^{\omega}$, respectively, where $\rho_{12} = \mathsf{prf}(\pi_1, \pi_2) = \mathtt{I} \cdot \mathtt{W} \cdot 2/1$ and $\rho_{13} = \mathsf{prf}(\pi_1, \pi_3) = \mathtt{I}$ are the corresponding common prefixes. Thus, ρ_{12} belongs to the witness W_{ψ}, while ρ_{13} does not.

As another example, consider the formula $\overline{\psi} = \mathtt{G}(\neg \mathtt{r}_1 \wedge \neg \mathtt{r}_2)$, which is equivalent to the negation of the previous one, and observe that its witness set $\mathrm{W}_{\overline{\psi}}$ is empty. Moreover, let $\chi_1, \chi_2, \chi_3 \in \mathrm{Asg}(\{\mathtt{A}, \mathtt{P}_1, \mathtt{P}_2\})$ be the three complete assignments we want to analyze. The strategies for the arbiter $\mathtt{A}$ are defined as above, while those of the processes follows: $\overline{\sigma}_{1/2/3}^{\mathtt{P}_i}(\rho) \triangleq \mathtt{i}$, $\overline{\sigma}_{1/2}^{\mathtt{P}_i}(\rho_{\mathtt{I}} \cdot \mathtt{I}) \triangleq \mathtt{r}$, $\overline{\sigma}_{1/2}^{\mathtt{P}_i}(\rho_{\mathtt{W}} \cdot \mathtt{W}) \triangleq \mathtt{a}$, and $\overline{\sigma}_{3}^{\mathtt{P}_i}(\rho_{\mathtt{I}} \cdot \mathtt{I}) = \overline{\sigma}_{3}^{\mathtt{P}_i}(\rho_{\mathtt{W}} \cdot \mathtt{W}) \triangleq \mathtt{i}$, for all $i \in \{1, 2\}$ and $\rho, \rho_{\mathtt{s}} \in \mathrm{Hst}$ with $\mathsf{lst}(\rho) \notin \{\mathtt{I}, \mathtt{W}\}$ and $\rho_{\mathtt{s}} \cdot \mathtt{s} \in \mathrm{Hst}$. Now, one can see that $\chi_1 \equiv_{\mathcal{G}}^{\overline{\psi}} \chi_2$, but $\chi_1 \not\equiv_{\mathcal{G}}^{\overline{\psi}} \chi_3$. Indeed, χ_1 and χ_2 induce the same play $(\mathtt{I} \cdot \mathtt{W})^{\omega}$, while χ_3 runs along $\mathtt{I}^{\omega}$. Thus, χ_1 and χ_2 are equivalent, but χ_1 and χ_3 are not.

C. Strategy Requirements

The semantics of a binding construct $\varphi = (a, x)\eta$ involves a redefinition of the underlying assignment χ, as it asserts that φ holds under χ once the inner part η is satisfied by associating the agent a to the strategy $\chi(x)$. Thus, the equivalence of two assignments χ_1 and χ_2 *w.r.t.* φ necessarily depends on that of their extensions on a *w.r.t.* η.

Definition III.4 (Binding Consistency)**.** *An equivalence relation on assignments* $\equiv_{\mathcal{G}}^{\cdot}$ *is* binding consistent *if, for a formula* $\varphi = (\mathtt{a}, \mathtt{x})\eta$ *and* φ*-coherent assignments* $\chi_1, \chi_2 \in \mathrm{Asg}$*, we have that* $\chi_1 \equiv_{\mathcal{G}}^{\varphi} \chi_2$ *iff* $\chi_1[\mathtt{a} \mapsto \chi_1(\mathtt{x})] \equiv_{\mathcal{G}}^{\eta} \chi_2[\mathtt{a} \mapsto \chi_2(\mathtt{x})]$*.*

To get familiar with the above concept, consider the formula $\flat\psi$, where $\flat \triangleq (\mathtt{A}, \mathtt{x})(\mathtt{P}_1, \mathtt{y}_1)(\mathtt{P}_2, \mathtt{y}_2)$, and let $\chi_1, \chi_2, \chi_3 \in \mathrm{Asg}(\{\mathtt{x}, \mathtt{y}_1, \mathtt{y}_2\})$ be the assignments assuming as values the strategies $\chi_i(\mathtt{x}) \triangleq \sigma_i^{\mathtt{A}}$ and $\chi_i(\mathtt{y}_j) \triangleq \sigma_i^{\mathtt{P}_j}$ previously defined, where $i \in \{1, 2, 3\}$ and $j \in \{1, 2\}$. Then, by definition, it is immediate to see that $\chi_1 \equiv_{\mathcal{G}}^{\flat\psi} \chi_2$, but $\chi_1 \not\equiv_{\mathcal{G}}^{\flat\psi} \chi_3$.

Before continuing with the analysis of the equivalence, it is important to make an observation about the dual nature of the existential and universal quantifiers *w.r.t.* the counting of strategies. We do this by exploiting the classic game-semantics metaphor originally proposed for first-order logic by Lorenzen and Hintikka, where the choice of an existential variable is done by a player called $\exists$ and that of the universal ones by its opponent $\forall$. Consider a sentence $\langle\langle x_1 \geq g_1 \rangle\rangle [[x_2 < g_2]]\eta$, having $\langle\langle \mathrm{y}_1 \geq h_1 \rangle\rangle \eta_1$ and $[[\mathrm{y}_2 < h_2]]\eta_2$ as two subformulas in η. When player $\exists$ tries to choose h_1 different strategies $\mathtt{y}_1$ to satisfy η_1, it also has to maximize the number of strategies x_1 by verifying $[[x_2 < g_2]]\eta$ to be sure that the constraint $\geq g_1$ of the first quantification is not violated. At the same time, player $\forall$ tries to do the opposite while choosing h_2 different strategies y_2 not satisfying η_2, *i.e.*, it needs to maximize the number of strategies x_2 falsifying η in order to violate the constraint $< g_2$ of the second quantifier.

With this observation in mind, we now treat the equivalence for the existential quantifier. Two assignments χ_1 and χ_2 are equivalent *w.r.t.* a formula $\varphi = \langle\!\langle x \geq g \rangle\!\rangle \eta$ if player $\exists$ is not able to find a strategy σ among those satisfying η, to associate with the variable x, that allows the corresponding extensions of χ_1 and χ_2 on x to induce different behaviors *w.r.t.* η. In other words, $\exists$ cannot distinguish between the two assignments, as they behave the same independently of the way they are extended.

Definition III.5 (Existential Consistency)**.** *An equivalence relation on assignments $\equiv^{\cdot}_{\mathcal{G}}$ is* existentially consistent *if, for any formula $\varphi = \langle\!\langle \mathtt{x} \geq g \rangle\!\rangle \eta$ and φ-coherent assignments $\chi_1, \chi_2 \in \mathrm{Asg}$, we have that $\chi_1 \equiv^{\varphi}_{\mathcal{G}} \chi_2$ iff, for each strategy $\sigma \in \eta[\mathcal{G}, \chi_1](\mathtt{x}) \cup \eta[\mathcal{G}, \chi_2](\mathtt{x})$, it holds that $\chi_1[\mathtt{x} \mapsto \sigma] \equiv^{\eta}_{\mathcal{G}} \chi_2[\mathtt{x} \mapsto \sigma]$.*

To clarify the above definition, consider the formula $\varphi = \langle\!\langle \mathtt{y}_2 \geq 2 \rangle\!\rangle \flat \overline{\psi}$ and let $\chi_1, \chi_2, \chi_3 \in \mathrm{Asg}(\{\mathtt{x}, \mathtt{y}_1\})$ be the three assignments having as values the strategies $\chi_i(\mathtt{x}) \triangleq \overline{\sigma}^{\mathtt{A}}_i$ and $\chi_i(\mathtt{y}_1) \triangleq \overline{\sigma}^{\mathtt{P}_1}_i$ previously defined, where $i \in \{1, 2, 3\}$. By a matter of calculation, one can see that $\chi_1 \equiv^{\varphi}_{\mathcal{G}} \chi_2$, but $\chi_1 \,/\!\!\equiv^{\varphi}_{\mathcal{G}} \chi_3$. By definition, $\chi_1 \equiv^{\varphi}_{\mathcal{G}} \chi_2$ iff, for each strategy $\overline{\sigma} \in (\flat \overline{\psi})[\mathcal{G}, \chi_1](\mathtt{y}_2) \cup (\flat \overline{\psi})[\mathcal{G}, \chi_2](\mathtt{y}_2)$, it holds that $\chi_1[\mathtt{y}_2 \mapsto \overline{\sigma}] \equiv^{\flat \overline{\psi}}_{\mathcal{G}} \chi_2[\mathtt{y}_2 \mapsto \overline{\sigma}]$. Now, observe that the strategy $\overline{\sigma}^{\mathtt{P}_2}_1$ introduced above is the unique one that allows χ_1 and χ_2 to satisfy $\flat \overline{\psi}$ once extended on $\mathtt{y}_2$. At this point, we can easily show that $\chi_1[\mathtt{y}_2 \mapsto \overline{\sigma}^{\mathtt{P}_2}_1] \equiv^{\flat \overline{\psi}}_{\mathcal{G}} \chi_2[\mathtt{y}_2 \mapsto \overline{\sigma}^{\mathtt{P}_2}_1]$, as the derived complete assignments $\chi_1[\mathtt{y}_2 \mapsto \overline{\sigma}^{\mathtt{P}_2}_1] \circ \flat$ and $\chi_2[\mathtt{y}_2 \mapsto \overline{\sigma}^{\mathtt{P}_2}_1] \circ \flat$ induce the same play $(\mathtt{I} \cdot \mathtt{W})^{\omega}$. The non-equivalence of χ_1 and χ_3 easily follows from the fact that $\overline{\sigma}^{\mathtt{P}_2}_1 \not\in (\flat \overline{\psi})[\mathcal{G}, \chi_3](\mathtt{y}_2)$, as $\chi_3[\mathtt{y}_2 \mapsto \overline{\sigma}^{\mathtt{P}_2}_1] \circ \flat$ induces the play $\mathtt{I} \cdot 2^{\omega}$ that does not satisfy $\overline{\psi}$. Thus, $\chi_1[\mathtt{y}_2 \mapsto \overline{\sigma}^{\mathtt{P}_2}_1] \not\equiv^{\flat \overline{\psi}}_{\mathcal{G}} \chi_3[\mathtt{y}_2 \mapsto \overline{\sigma}^{\mathtt{P}_2}_1]$.

We conclude with the equivalence for the universal quantifier. Two assignments χ_1 and χ_2 are equivalent *w.r.t.* a formula $\varphi = [\![x < g]\!] \eta$ if, for each index $i \in \{1, 2\}$ and strategy σ_i player $\forall$ chooses among those satisfying η under χ_i, there is a strategy σ_{3-i} this player can choose among those satisfying η under χ_{3-i} such that, once the two strategies are associated with the variable x, they make the corresponding extensions of assignments equivalent *w.r.t.* η. This means that the parts of the game structure that are reachable under χ_1 and χ_2 contain exactly the same information *w.r.t.* the verification of the inner formula. In other words, $\forall$ cannot distinguish between the two assignments, as the induced subtrees of possible plays are practically the same.

Definition III.6 (Universal Consistency)**.** *An equivalence relation on assignments $\equiv^{\cdot}_{\mathcal{G}}$is* universally consistent *if, for any formula $\varphi = [\![\mathtt{x} < g]\!] \eta$ and φ-coherent assignments $\chi_1, \chi_2 \in$ Asg, we have that $\chi_1 \equiv^{\varphi}_{\mathcal{G}} \chi_2$ iff, for all $i \in \{1, 2\}$ and strategy $\sigma_i \in \eta[\mathcal{G}, \chi_i](\mathtt{x})$, there is a strategy $\sigma_{3-i} \in \eta[\mathcal{G}, \chi_{3-i}](\mathtt{x})$ such that $\chi_1[\mathtt{x} \mapsto \sigma_1] \equiv^{\eta}_{\mathcal{G}} \chi_2[\mathtt{x} \mapsto \sigma_2]$.*

Finally, to better understand the above definition, consider the formula $\varphi = [\![\mathtt{y}_1 < 1]\!] \eta$, where $\eta = [\![\mathtt{y}_2 < 2]\!] \flat \psi$, and let $\chi_1, \chi_2, \chi_3 \in \mathrm{Asg}(\{\mathtt{x}\})$ be the three assignments having as values the strategies $\chi_i(\mathtt{x}) \triangleq \sigma^{\mathtt{A}}_i$ previously defined, where $i \in \{1, 2, 3\}$. One can see that $\chi_1 \equiv^{\varphi}_{\mathcal{G}} \chi_2$, but $\chi_1 \not\equiv^{\varphi}_{\mathcal{G}} \chi_3$.

First, observe that $\eta[\mathcal{G}, \chi_1](\mathtt{y}_1) = \eta[\mathcal{G}, \chi_2](\mathtt{y}_1) = \mathrm{Str}$. Indeed, for all strategies $\sigma \in \mathrm{Str}$, we have that $\mathcal{G}, \chi_1[\mathtt{y}_1 \mapsto \sigma] \models \eta$ and $\mathcal{G}, \chi_2[\mathtt{y}_1 \mapsto \sigma] \models \eta$, since $\mathcal{G}, \chi_1[\mathtt{y}_1 \mapsto \sigma, \mathtt{y}_2 \mapsto \sigma'] \models \flat \psi$ and $\mathcal{G}, \chi_2[\mathtt{y}_1 \mapsto \sigma, \mathtt{y}_2 \mapsto \sigma'] \models \flat \psi$, for all $\sigma' \in \mathrm{Str}$ such that $\sigma \neq \sigma'$. This is due to the fact that the plays π_1 and π_2 induced by the two complete assignments $\chi_1[\mathtt{y}_1 \mapsto \sigma, \mathtt{y}_2 \mapsto \sigma'] \circ \flat$ and $\chi_2[\mathtt{y}_1 \mapsto \sigma, \mathtt{y}_2 \mapsto \sigma'] \circ \flat$ differ from $(\mathtt{I}^{+} \cdot \mathtt{W})^{*} \cdot \mathtt{I}^{\omega}$ and $(\mathtt{I}^{+} \cdot \mathtt{W})^{\omega}$, as the strategies of the two processes are different. Also, they share a common prefix $\rho = \mathsf{prf}(\pi_1, \pi_2)$ belonging to W_{ψ}, since the strategies of the arbiter only differ on the histories ending in the state $2/1$. We can now show that χ_1 and χ_2 are equivalent, by applying the above definition in which we assume that $\sigma_i = \sigma_{3-i}$.

To prove that χ_1 and χ_3 are non-equivalent, we show that there is a strategy $\sigma \in \eta[\mathcal{G}, \chi_1](\mathtt{y}_1)$ for χ_1 such that, for all strategies $\sigma' \in \eta[\mathcal{G}, \chi_3](\mathtt{y}_1)$ for χ_3, it holds that $\chi_1[\mathtt{y}_1 \mapsto \sigma] \not\equiv^{\eta}_{\mathcal{G}} \chi_3[\mathtt{y}_1 \mapsto \sigma']$. As before, observe that $\eta[\mathcal{G}, \chi_1](\mathtt{y}_1) = \eta[\mathcal{G}, \chi_3](\mathtt{y}_1) = \mathrm{Str}$ and choose $\sigma \in \mathrm{Str}$ as the strategy $\sigma^{\mathtt{P}_1}_1$ previously defined. At this point, one can easily see that all plays compatible with $\chi_1[\mathtt{y}_1 \mapsto \sigma] \circ \flat$ pass through either $\mathtt{I} \cdot 1$ or $\mathtt{I} \cdot \mathtt{W} \cdot 2/1$, while a play compatible with $\chi_3 \circ \flat$ cannot pass through the latter history. Thus, the non-equivalence of the two assignments immediately follows.

IV. Determinacy and Model Checking

In this section, we address the determinacy and the model checking problems for GSL[1G] over game structures. In particular, we provide procedures for the *vanilla* fragment of the logic in which all temporal properties are used as in ATL. Technically, for the model checking question, we make use of a technique that extends the one introduced in [MMS14b], which allows to reduce to a simplified equivalent version of the problem, over turn-based structures. For the matter of clarity, in the sequel we restrict to the case of structures involving only two players. Hence, we investigate the mentioned problems for a fragment of GSL that we name GSL[1G, 2AG].

A. Determinacy

Recall that determinacy has been first proved for classic Borel turn-based two-player games in [Mar75]. However, the proof used there does not directly apply to our graded setting. To give evidence of the differences between the two frameworks, observe that in SL[1G, 2AG] sentences like $\langle\!\langle x \rangle\!\rangle [\![\overline{x}]\!] \eta$ imply $[\![\overline{x}]\!] \langle\!\langle x \rangle\!\rangle \eta$, while in GSL[1G, 2AG] the corresponding implication $\langle\!\langle x \geq i \rangle\!\rangle [\![\overline{x} < j]\!] \eta \Rightarrow [\![\overline{x} < j]\!] \langle\!\langle x \geq i \rangle\!\rangle \eta$ does not hold. The determinacy property we are interested in is exactly the converse direction, *i.e.*, $[\![\overline{x} < j]\!] \langle\!\langle x \geq i \rangle\!\rangle \eta \Rightarrow \langle\!\langle x \geq i \rangle\!\rangle [\![\overline{x} < j]\!] \eta$. In particular, we extend the Gale-Stewart Theorem [PP04], by exploiting a deep generalization of the technique used in [FNP10]. The idea consists of a fixed-point calculation over the number of winning strategies an agent can select against all but a fixed number of those of its opponent. Regarding this approach, we want to remind that the simpler counting considered in [FNP10] is restricted to existential quantifications.

Construction IV.1 (Grading Function)**.** *Consider a two-agent turn-based game structure $\mathcal{G}$ with $\mathrm{Ag} = \{\alpha, \overline{\alpha}\}$. Moreover, let ψ be an* LTL *formula, where $\mathrm{W}_{\psi}, \mathrm{W}_{\neg\psi} \subseteq \mathrm{Hst}$ denotes the witness sets for ψ and $\neg\psi$, respectively. It is immediate to see that, in case $s_I \in \mathrm{W}_{\psi}$ (resp., $s_I \in \mathrm{W}_{\neg\psi}$), all strategy profiles are equivalent* w.r.t. *the temporal property ψ (resp., $\neg\psi$). If $s_I \in \mathrm{X} \triangleq \mathrm{Hst} \setminus (\mathrm{W}_{\psi} \cup \mathrm{W}_{\neg\psi})$, instead, we need to introduce a* grading function $\mathsf{G}^{\alpha}_{\psi} : \mathrm{X} \to \Gamma$, *where $\Gamma \triangleq \mathbb{N} \to (\mathbb{N} \cup \{\omega\})$, that allows to determine how many different strategies the agent α (resp., $\overline{\alpha}$) owns* w.r.t. *ψ (resp., $\neg\psi$). Informally, $\mathsf{G}^{\alpha}_{\psi}(\rho)(j)$ represents the number of winning strategies player α can put up against all but at most j strategies of its adversary $\overline{\alpha}$, once the current play has already reached the history $\rho \in \mathrm{X}$.*

Before continuing, observe that α sometimes has the possibility to commit a suicide, i.e.*, to choose a strategy leading directly to a history in $\mathrm{W}_{\neg\psi}$, with the hope to win the game by collapsing all strategies of its opponent into a unique class. The set of histories enabling this possibility is defined as follows: $\mathrm{S} \triangleq \{\rho \in \mathrm{X} : \exists \rho' \in \mathrm{W}_{\neg\psi} . \rho < \rho' \wedge \forall \rho'' \in \mathrm{Hst} . \rho \leq \rho'' < \rho' \Rightarrow \rho'' \in \mathrm{Hst}_{\alpha}\}$, where $\mathrm{Hst}_{\alpha} = \{\forall \rho \in \mathrm{Hst} : \mathsf{ag}(\mathsf{lst}(\rho)) = \{\alpha\}\}$ is the set of histories ending in a state controlled by α. Intuitively, this agent can autonomously extend a history $\rho \in \mathrm{S}$ into one $\rho' \in \mathrm{W}_{\neg\psi}$ that is surely loosing, independently of the behavior of $\overline{\alpha}$. Note that there may be several suicide strategies, but all of them are equivalent* w.r.t. *the property ψ. Also, against them, all counter strategies of $\overline{\alpha}$ are equivalent as well.*

At this point, to define the function $\mathsf{G}^{\alpha}_{\psi}$, we introduce the auxiliary functor $\mathsf{F}^{\alpha}_{\psi} : (\mathrm{X} \to \Gamma) \to (\mathrm{X} \to \Gamma)$, whose least fixpoint represents a function returning the maximum number of different strategies α can use against all but a precise fixed number of counter strategies of $\overline{\alpha}$. Formally, we have that:

$$\mathsf{F}^{\alpha}_{\psi}(\mathsf{f})(\rho)(j) \triangleq$$
$$\begin{cases} \sum_{\rho' \in \mathsf{suc}(\rho) \cap \mathrm{X}} \mathsf{f}(\rho')(0) + |\mathsf{suc}(\rho) \cap \mathrm{W}_{\psi}|, & \textit{if } \rho \in \mathrm{Hst}_{\alpha} \textit{ and } j = 0; \\ \sum_{\rho' \in \mathsf{suc}(\rho) \cap \mathrm{X}} \mathsf{f}(\rho')(j), & \textit{if } \rho \in \mathrm{Hst}_{\alpha} \textit{ and } j > 0; \\ \sum_{\mathsf{c} \in \mathrm{C}(\rho)(j)} \prod_{\rho' \in \mathsf{dom}(\mathsf{c})} \mathsf{f}(\rho')(\mathsf{c}(\rho')), & \textit{otherwise}; \end{cases}$$

where $\mathsf{suc}(\rho) = \{\rho' \in \mathrm{Hst} : \exists s \in \mathrm{St}.\ \rho s = \rho'\}$ and $\mathsf{C}(\rho)(i) \subseteq (\mathsf{suc}(\rho) \cap \mathrm{Z}) \rightharpoonup \mathbb{N}$ contains all partial functions $\mathsf{c} \in \mathsf{C}(\rho)(i)$ for which α owns a suicide strategy on the histories not in their domains, i.e., *$(\mathsf{suc}(\rho) \cap \mathrm{Z}) \setminus \mathsf{dom}(\mathsf{c}) \subseteq \mathrm{S}$, and the sum of all values assumed by c plus the number of successor histories that are neither surely winning nor contained in the domain of c equals to i,* i.e., *$i = \sum_{\rho' \in \mathsf{dom}(\mathsf{c})} \mathsf{c}(\rho') + |\mathsf{suc}(\rho) \setminus (\mathrm{X} \cup \mathsf{dom}(\mathsf{c}))|$.*

Intuitively, the first item of the definition simply asserts that the number of strategies $\mathsf{F}(\mathsf{f})(\rho)(0)$ that agent α has on the α-history ρ, without excluding any counter strategy of its adversary, is obtainable as the sum of the $\mathsf{f}(\rho')(0)$ strategies on the successor histories $\rho' \in \mathrm{X}$ plus a single strategy for each successor history that is surely winning. Similarly, the second item takes into account the case in which we can avoid exactly j counter strategies. The last item, instead, computes the number of strategies for α on the $\overline{\alpha}$-histories. In particular, through the set $\mathsf{C}(\rho)(j)$, it first determines in how many ways it is possible to split the number j of counter strategies to avoid among all successor histories of ρ. Then, for each of these splittings, it calculates the product of the corresponding numbers $\mathsf{f}(\rho')(\mathsf{c}(\rho'))$ of strategies for α.

We are finally able to define the grading function $\mathsf{G}^{\alpha}_{\psi}$ by means of the least fixpoint $\mathsf{f}^{\star} = \mathsf{F}^{\alpha}_{\psi}(\mathsf{f}^{\star})$ of the functor $\mathsf{F}^{\alpha}_{\psi}$ as follows: $\mathsf{G}^{\alpha}_{\psi}(\rho)(j) \triangleq \sum_{h=0}^{j} \mathsf{f}^{\star}(\rho)(h) + [\rho \in \mathrm{S} \wedge j \geq 1]$. Intuitively, $\mathsf{G}^{\alpha}_{\psi}(\rho)(j)$ is the sum of the numbers $\mathsf{f}^{\star}(\rho)(h)$ of winning strategies the agent α can exploit against all but exactly h strategies of its adversary $\overline{\alpha}$, for each $h \in [0, j]$. Moreover, if $\rho \in \mathrm{S}$, we need to add to this counting the suicide strategy that α can use once $\overline{\alpha}$ avoids to apply his unique counter strategy. This is formalized through the standard notation $[\eth]$ [GKP94] that is evaluated to 1, if the condition $\eth$ is true, and to 0, otherwise.

Thanks to the above construction, one can compute the maximum number of strategies that a player has at its disposal against all but a fixed number of strategies of the opponent. Next lemma, whose statement can be constructively proved by transfinite induction on the recursions of the functor $\mathsf{F}^{\alpha}_{\psi}$, precisely describes this fact. Indeed, we show how the satisfiability of a GSL[1G, 2AG] sentence $\langle\langle x \geq i \rangle\rangle [[\overline{x} \leq j]] (\alpha, x)(\overline{\alpha}, \overline{x})\psi$ can be decided via the computation of the associated grading function $\mathsf{G}^{\alpha}_{\psi}$, where by $[[\overline{x} \leq j]]\varphi$ we mean $[[\overline{x} < j + 1]]\varphi$.

Lemma IV.1 (Grading Function)**.** *Let $\mathcal{G}$ be a two-agent turn-based game structure, where $\mathrm{Ag} = \{\alpha, \overline{\alpha}\}$, and $\varphi = \langle\langle x \geq i \rangle\rangle [[\overline{x} \leq j]] (\alpha, x)(\overline{\alpha}, \overline{x})\psi$ a* GSL[1G, 2AG] *sentence. Moreover, let $\mathsf{G}^{\alpha}_{\psi}$ be the grading function and $\mathrm{W}_{\psi}, \mathrm{W}_{\neg\psi}, \mathrm{X} \subseteq \mathrm{Hst}$ the sets of histories obtained in Construction IV.1. Then, $\mathcal{G} \models \varphi$ iff one of the following three conditions hold:* (i) *$i \leq 1$, $j \geq 0$, and $s_I \in \mathrm{W}_{\psi}$;* (ii) *$i \leq 1$, $j \geq 1$, and $s_I \in \mathrm{W}_{\neg\psi}$;* (iii) *$i \leq \mathsf{G}^{\alpha}_{\psi}(s_I)(j)$ and $s_I \in \mathrm{X}$.*

Again by transfinite induction on its recursive structure, we can prove a quite natural but fundamental property of the grading function, *i.e.*, its duality in the form described in the next lemma. To give an intuition, assume that agent $\overline{\alpha}$ has at most j strategies to satisfy the temporal property $\neg\psi$ against all but at most i strategies of its adversary α. Then, it can be shown that the latter has more than i strategies to satisfy ψ against all but at most j strategies of the former.

Lemma IV.2 (Grading Duality)**.** *Let $\mathsf{G}^{\alpha}_{\psi}$ and $\mathsf{G}^{\overline{\alpha}}_{\neg\psi}$ be the grading functions and $\mathrm{X} \subseteq \mathrm{Hst}$ the set of histories obtainable by Construction IV.1. For all histories $\rho \in \mathrm{X}$ and indexes $i, j \in \mathbb{N}$, it holds that if $\mathsf{G}^{\overline{\alpha}}_{\neg\psi}(\rho)(i) \leq j$ then $i < \mathsf{G}^{\alpha}_{\psi}(\rho)(j)$.*

Summing up the above two results, we can easily prove that, on turn-based game structures, GSL[1G, 2AG] is determined. Indeed, suppose that $s_I \in \mathrm{X}$ and $\mathcal{G} \models [[\overline{x} \leq j]] \langle\langle x \geq i \rangle\rangle \flat\psi$, where $\flat = (\alpha, x)(\overline{\alpha}, \overline{x})$ (the case with $s_I \in \mathrm{W}_{\psi}$ immediately follows from classic Martin's Determinacy Theorem [Mar75], [Mar85]). Obviously, $\mathcal{G}$ does not satisfy the negation of this sentence, *i.e.*, $\mathcal{G} \not\models \langle\langle \overline{x} \geq j+1 \rangle\rangle [[x \leq i-1]] \flat\neg\psi$. Consequently, by Lemma IV.1, we have that $\mathsf{G}^{\overline{\alpha}}_{\neg\psi}(s_I)(i-1) \leq j$. Hence, by Lemma IV.2, it follows that $i \leq \mathsf{G}^{\alpha}_{\psi}(s_I)(j)$. Finally, again by Lemma IV.1, we obtain that $\mathcal{G} \models \langle\!\langle x \geq i \rangle\!\rangle [\![\overline{x} \leq j]\!] \flat\psi$, as required by the definition of determinacy.

Theorem IV.1 (Determinacy)**.** GSL[1G, 2AG] *on turn-based game structures is determined.*

B. Model Checking

We finally describe a solution of the model-checking problem for the above mentioned fragment of GSL[1G, 2AG], which only admits *simple temporal properties*, *i.e.*, $\varphi_1 \mathtt{U} \varphi_2$, $\varphi_1 \mathtt{R} \varphi_2$, and $\mathtt{X}\varphi$, where φ_1, φ_2, and φ are sentences. This fragment, called Vanilla GSL[1G, 2AG], is in relation with GSL[1G, 2AG], as CTL and ATL are for CTL* and ATL*, respectively.

The idea here is to exploit the characterization of the grading function stated in Lemma IV.1 in order to verify whether a game structure $\mathcal{G}$ satisfies a sentence $\varphi = \langle\!\langle \mathtt{x} \geq i \rangle\!\rangle [\![\overline{\mathtt{x}} \leq j]\!] (\alpha, \mathtt{x})(\overline{\alpha}, \overline{\mathtt{x}})\psi$, while avoiding the naive infinite calculation of least fixpoints $\mathsf{F}^{\alpha}_{\psi}$.

Fortunately, due to the simplicity of the temporal property ψ, we have that the four sets W_{ψ}, $\mathrm{W}_{\neg\psi}$, X, and S previously introduced are memoryless, *i.e.*, if a history belongs to them, every other history ending in the same state is also a member of these sets. Therefore, we can focus only on states by defining $\mathrm{W}_{\psi} \triangleq \{s \in \mathrm{St} : \mathcal{G}, s \models \mathtt{A}\psi\}$, $\mathrm{W}_{\neg\psi} \triangleq \{s \in \mathrm{St} : \mathcal{G}, s \models \mathtt{A}\neg\psi\}$, $\mathrm{X} \triangleq \mathrm{St} \setminus (\mathrm{W}_{\psi} \cup \mathrm{W}_{\neg\psi})$, and $\mathrm{S} \triangleq \{s \in \mathrm{St} : \mathcal{G}, s \models \mathtt{E}(\alpha \mathtt{U} \mathtt{A}\neg\psi)\}$ via very simple CTL properties. Observe that we use α and $\overline{\alpha}$ as labeling of a state to recognize its owner. Intuitively, W_{ψ} and $\mathrm{W}_{\neg\psi}$ contain the states from which agents α and $\overline{\alpha}$ can ensure, independently from the adversary, the properties ψ and $\neg\psi$, respectively. The set X, instead, contains the states on which we have still to determine the number of strategies at disposal of the two agents. Finally, S maintains the suicide states, *i.e.*, those states from which α can commit suicide by autonomously reaching $\mathrm{W}_{\neg\psi}$. In addition, since at most j strategies of $\overline{\alpha}$ can be avoided while reasoning on the sentence φ, we need just to deal with functions in the set $\Gamma \triangleq [0, j] \to (\mathbb{N} \cup \{\omega\})$ instead of $\Gamma \triangleq \mathbb{N} \to (\mathbb{N} \cup \{\omega\})$. Consequently, the functor $\mathsf{F}^{\alpha}_{\psi} : (\mathrm{X} \to \Gamma) \to (\mathrm{X} \to \Gamma)$ can be redefined as follows:

$$\mathsf{F}(\mathsf{f})(s)(h) \triangleq$$
$$\begin{cases} \sum_{s' \in \mathsf{suc}(s) \cap \mathrm{X}} \mathsf{f}(s')(0) + |\mathsf{suc}(s) \cap \mathrm{W}_{\psi}|, & \text{if } s \in \mathrm{St}_{\alpha} \text{ and } h = 0; \\ \sum_{s' \in \mathsf{suc}(s) \cap \mathrm{X}} \mathsf{f}(s')(h), & \text{if } s \in \mathrm{St}_{\alpha} \text{ and } h > 0; \\ \sum_{\mathsf{c} \in \mathrm{C}(s)(h)} \prod_{s' \in \mathsf{dom}(\mathsf{c})} \mathsf{f}(s')(\mathsf{c}(s')), & \text{otherwise;} \end{cases}$$

where $\mathsf{suc}(s) = \{s' \in \mathrm{St} : (s, s') \in Ed\}$ and $\mathsf{C}(s)(i) \subseteq (\mathsf{suc}(s) \cap \mathrm{Z}) \rightharpoonup \mathbb{N}$ contains all partial functions $\mathsf{c} \in \mathsf{C}(s)(i)$ for which α owns a suicide strategy on the states not in their domains, *i.e.*, $(\mathsf{suc}(s) \cap \mathrm{Z}) \setminus \mathsf{dom}(\mathsf{c}) \subseteq \mathrm{S}$, and the sum of all values assumed by c plus the number of successors that are neither surely winning nor contained in the domain of c equals to i, *i.e.*, $i = \sum_{s' \in \mathsf{dom}(\mathsf{c})} \mathsf{c}(s') + |\mathsf{suc}(s) \setminus (\mathrm{X} \cup \mathsf{dom}(\mathsf{c}))|$. Similarly, the grading function $\mathsf{G}^{\alpha}_{\psi} : \mathrm{X} \to \Gamma$ becomes $\mathsf{G}^{\alpha}_{\psi}(s)(h) \triangleq \sum_{l=0}^{h} \mathsf{f}^{\star}(s)(l) + [s \in \mathrm{S} \wedge h \geq 1]$, where $\mathsf{f}^{\star}$ is the least fixpoint of $\mathsf{F}^{\alpha}_{\psi}$.

Unfortunately, these redefinitions are not enough by their on to ensure that the fixpoint calculation can be done in a finite, possibly small, number of iterations of the functor. This is due to two facts: the functions in Γ have an infinite codomain and the game structure $\mathcal{G}$ has cycles inside. In order to solve such a problem, we make use of the following observation. Suppose that agent α has at least one strategy on one of its states $s \in \mathrm{St}_{\alpha}$ that is also part of a cycle in which no state of its opponent $\overline{\alpha}$ is adjacent to the set $\mathrm{W}_{\neg\psi}$. Then, α can use this cycle from s to construct an infinite number of nonequivalent strategies, by simply pumping-up the number of times he decides to traverse it before following the previously found strategy. Therefore, in this case, we avoid to compute the infinite number of iterations required to reach the fixpoint, by directly assuming ω as value. Formally, we introduce the functor $\mathsf{I} : (\mathrm{X} \to \Gamma) \to (\mathrm{X} \to \Gamma)$ defined as follows, where $\mathrm{L} \subseteq \mathrm{St}_{\alpha}$ denotes the set of α-states belonging to a cycle of the above kind: $\mathsf{I}(\mathsf{f})(s)(h) = \omega$, if $s \in \mathrm{L}$ and $\mathsf{f}(s)(h) > 0$, and $\mathsf{I}(\mathsf{f})(s)(h) = \mathsf{f}(s)(h)$, otherwise, for all $s \in \mathrm{St}$ and $h \in [0, j]$. By induction on the ordering and topology of the strong connected components of the underlying game structure, we can prove that $\mathsf{f}^{\star} = (\mathsf{I} \circ \mathsf{F}^{\alpha}_{\psi})(\mathsf{f}^{\star})$ iff $\mathsf{f}^{\star} = \mathsf{F}^{\alpha}_{\psi}(\mathsf{f}^{\star})$, *i.e.*, the functor obtained by composing I and $\mathsf{F}^{\alpha}_{\psi}$ have exactly the same fixpoint of $\mathsf{F}^{\alpha}_{\psi}$. Moreover, $\mathsf{f}^{\star} = (\mathsf{I} \circ \mathsf{F}^{\alpha}_{\psi})^{n}(\mathsf{f}_{\mathrm{o}})$ where $j \cdot |\mathcal{G}| \leq n$ and f_{o} is the zero function, *i.e.*, $\mathsf{f}_{\mathrm{o}}(s)(h) = 0$, for all $s \in \mathrm{St}$ and $h \in [0, j]$. Hence, we can compute $\mathsf{f}^{\star}$ in a number of iterations of $\mathsf{I} \circ \mathsf{F}^{\alpha}_{\psi}$ that is linear in both the degree j and the size of $\mathcal{G}$. Finally, observe that the computation of the set L can be done in quadratic time by using a classic Büchi procedure.

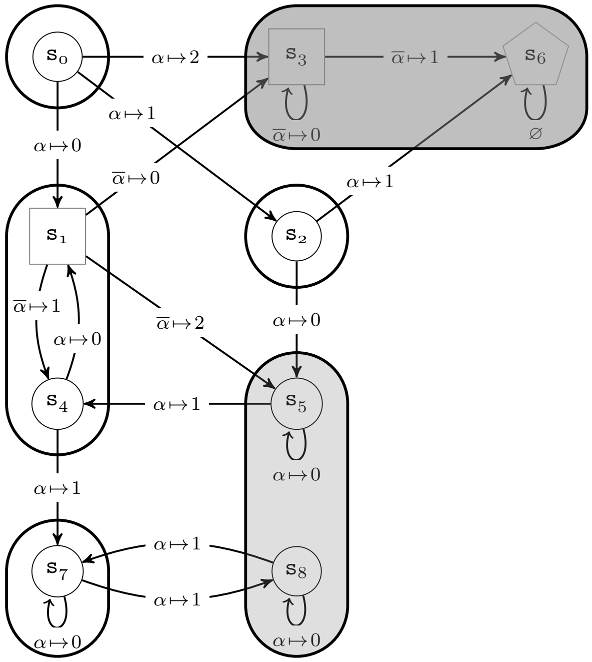

Figure 2. A two-player turn-based game structure.

As an example of an application of the model-checking procedure, consider the two-agent turn-based game structure $\mathcal{G}$ depicted in Figure 2, with the circle states ruled by α, the square ones by its opponent $\overline{\alpha}$, and where $\mathtt{s}_5$ and $\mathtt{s}_8$ are labeled by the atomic proposition p. Also, consider the vanilla GSL[1G, 2AG] sentence $\varphi = \langle\!\langle \mathtt{x} \geq i \rangle\!\rangle [\![\overline{\mathtt{x}} \leq j]\!] (\alpha, \mathtt{x})(\overline{\alpha}, \overline{\mathtt{x}})\mathtt{F}\mathtt{p}$. First, we need to compute the five preliminary sets of states $\mathrm{W}_{\mathtt{Fp}} = \{\mathsf{s}_5, \mathsf{s}_8\}$ (the light-gray area), $\mathrm{W}_{\neg\mathtt{Fp}} = \{\mathsf{s}_3, \mathsf{s}_6\}$ (the dark-gray area), $\mathrm{X} = \{\mathsf{s}_0, \mathsf{s}_1, \mathsf{s}_2, \mathsf{s}_4, \mathsf{s}_7\}$ (the white area partitioned into strong-connected components), $\mathrm{S} = \{\mathsf{s}_0, \mathsf{s}_2\}$, and $\mathrm{L} = \{\mathsf{s}_7\}$. Now, we can evaluate the fixpoint $\mathsf{f}^\star$ of the functor $\mathsf{I} \circ \mathsf{F}^{\alpha}_{\psi}$ that can be obtained, due to the topology of $\mathcal{G}$, after seven iterations, *i.e.*, $\mathsf{f}^\star = (\mathsf{I} \circ \mathsf{F}^{\alpha}_{\psi})^7(\mathsf{f}_0)$. Indeed, at the first one, the values on the states s_2 and s_7 are stabilized to $\mathsf{f}^\star(\mathsf{s}_2)(0) = 1$, $\mathsf{f}^\star(\mathsf{s}_7)(0) = \omega$, and $\mathsf{f}^\star(\mathsf{s}_2)(h) = \mathsf{f}^\star(\mathsf{s}_7)(h) = 0$, for all $h \in [1, j]$. After six iterations, we obtain $\mathsf{f}^\star(\mathsf{s}_1)(0) = 0$, $\mathsf{f}^\star(\mathsf{s}_1)(h) = \omega$, for all $h \in [1, j]$, and $\mathsf{f}^\star(\mathsf{s}_4)(h) = \omega$, for all $h \in [0, j]$. By computing the last iteration, we derive $\mathsf{f}^\star(\mathsf{s}_0)(0) = 1$ and $\mathsf{f}^\star(\mathsf{s}_0)(h) = \omega$, for all $h \in [1, j]$. Note that 7 is exactly the sum $1 + 5 + 1$ of iterations that the components of the longest chain $\{\mathsf{s}_7\} < \{\mathsf{s}_1, \mathsf{s}_4\} < \{\mathsf{s}_0\}$ need in order to stabilize the values on their states. Finally, we can verify whether $\mathcal{G} \models \varphi$, by computing the grading function $\mathsf{G}^{\alpha}_{\mathtt{Fp}}$ at s_0, whose values are $\mathsf{G}^{\alpha}_{\mathtt{Fp}}(s_0)(0) = 1$ and $\mathsf{G}^{\alpha}_{\mathtt{Fp}}(s_0)(h) = \omega$, for all $h \in [1, j]$. Thus, $\mathcal{G} \models \varphi$ iff $i = 1$ or $j > 0$.

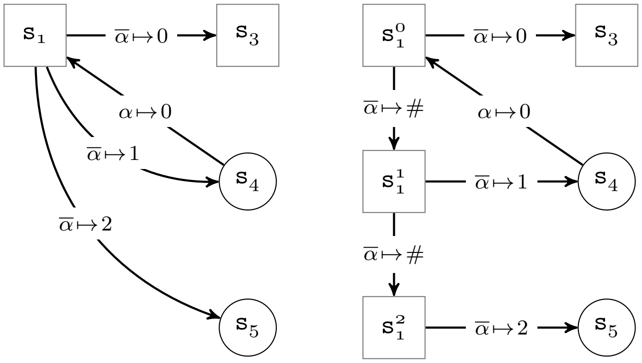

s_1
$\overline{\alpha} \mapsto 0$
s_3
$\alpha \mapsto 0$
$\overline{\alpha} \mapsto 1$
s_4
$\overline{\alpha} \mapsto 2$
s_5
s_1^0
$\overline{\alpha} \mapsto 0$
s_3
$\overline{\alpha} \mapsto \#$
$\alpha \mapsto 0$
s_1^1
$\overline{\alpha} \mapsto 1$
s_4
$\overline{\alpha} \mapsto \#$
s_1^2
$\overline{\alpha} \mapsto 2$
s_5

Figure 3. Degree transformation.

In order to obtain a PTIME procedure, we have also to ensure that each evaluation of the composed functor $\mathsf{I} \circ \mathsf{F}^{\alpha}_{\psi}$ can be computed in PTIME *w.r.t.* the above mentioned parameters. Actually, the whole I and the first two items of $\mathsf{F}^{\alpha}_{\psi}$ can easily be calculated in linear time. The third item, instead, may require a sum of an exponential number of elements. Indeed, due to all possible ways a degree j can split among the successors of a state s, the corresponding set $\mathrm{C}(s)(j)$ may contain an exponential number of functions. To avoid this, by exploiting a technique similar to the one proposed in [BMM10], [BMM12], we linearly transform a game structure into an equivalent one where all states ruled by $\overline{\alpha}$ have degree at most 2. In this way, the cardinality of $\mathrm{C}(s)(j)$ is bounded by j. For example, consider the left part of Figure 3 representing the substructure of the previous game structure $\mathcal{G}$ induced by the state s_1 together with its three successors. It is not hard to see that we can replace it, in $\mathcal{G}$, by the binary graph at its right, without changing the number of strategies that the two agents have at their disposal.

Theorem IV.2 (Model Checking). *The model-checking problem for* Vanilla GSL[1G] *is* PTIME *in both the size of the game structure and the sentence. Moreover it is* PTIME-HARD w.r.t. *both the data and combined complexity.*

Observe that the PTIME hardness *w.r.t.* the size of game is simply derived from the fact that classic reachability games [Imm81] are subsumed. Instead, the hardness *w.r.t.* the combined complexity is immediately obtained from the fact that GSL[1G] subsumes CTL [Sch02].

V. DISCUSSION

In multi-agent systems general questions to be investigated are: *"is there a winning strategy?"* or *"is the game surely winning?"* (*i.e.*, no matter which strategy the agent can play). In the years, several logics suitable for the strategic reasoning have been introduced and, by means of existential and universal modalities, this kind of questions has been addressed [AHK02]. However, these logics are not able to address quantitative aspects such as *"what is the number of winning strategies an agent can play?"* or, in general, to determine the *success rate of a game* [MMS15]. These questions are critical in dealing with solution concepts [Mye91] and in open-system verification [FMP08].

In this paper, we have introduced and studied GSL, an extension of Strategy Logic with graded modalities. The use of a powerful formalism such as Strategy Logic ensures the ability of dealing with very intricate game scenarios [MMV10]. The obvious drawback of this is a considerable amount of work on solving any related question [MMPV12]. One of the main difficulties we have faced in GSL has been the definition of the right methodology to count strategies. To this aim, we have introduced a suitable equivalence relation over strategy profiles based on the strategic behavior they induce and studied its robustness. Also, we have provided arguments and some examples along the paper to give evidence of the usefulness of GSL and the suitability of the proposed counting.

In order to provide results of practical use, we have investigated basic questions over a restricted fragment of GSL. Precisely we have considered the case in which the graded modalities are applied to the *vanilla* restriction of the one-goal fragment of SL [MMPV12]. The resulting logic, named *Vanilla* SL[1G], has been investigated in the turn-based setting. We have obtained positive results about determinacy and showed that the related model-checking problem is PTIME-COMPLETE.

The framework and the results presented in this paper open for several future work questions. First, it would be worth investigating the extension of existing formal verification tools such as MCMAS [LR06] with our results. We recall that MCMAS, originally developed for the verification for multi-agent models with respect to specification given in ATL [LR06], has been recently extended to handle Strategy Logic specifications [ČLMM14], [ČLM15]. Under our formalism it is possible to check, in a single evaluation process, that more than one strategy gives a fault and possibly correct all these errors. This in a way similar as the verification tool NuSMV has been extended to deal with *graded*-CTLverification [FNP10].

Another research direction regards investigating the graded extension of other formalism for the strategic reasoning such as ATL *with context* [BLLM09], [LLM10], as well as, for the sake of completeness, to determine the complexity of the model checking problem with respect to other fragments of Strategy Logic [MMS13], [MMS14a].

Finally, it would be really interesting to address the satisfiability for GSL[1G] too, by generalizing the solution procedure developed for SL[1G] [MMPV12]. However, we want to observe that, the technical tools described in this article are not powerful enough to solve this problem, since this also needs a bounded-width tree model property. So, further work is still required. Moreover, the procedure exploited for graded CTL [BMM09], [BMM10], [BMM12] cannot easily be applied to GSL[1G], due to the fact that the binary-tree unraveling used there would modify the way the strategies are valuated as equivalent.

Acknowledgment

This paper is partially supported by the FP7 EU project 600958-SHERPA.

References

[AHK02] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *JACM*, 49(5):672–713, 2002.

[ALM^{+}13] B. Aminof, A. Legay, A. Murano, O. Serre, and M. Y. Vardi. Pushdown module checking with imperfect information. *Inf. Comput.*, 223(1):1–17, 2013.

[ATO^{+}09] T. Antal, A. Traulsen, H. Ohtsuki, C.E. Tarnita, and M.A. Nowak. Mutation-Selection Equilibrium in Games with Multiple Strategies. *JTB*, 258(4):614–622, 2009.

[BBF^{+}10] B. Bérard, M. Bidoit, A. Finkel, F. Laroussinie, A. Petit, L. Petrucci, and P. Schnoebelen. *Systems and Software Verification: Model-Checking Techniques and Tools.* Springer, 2010.

[BLLM09] T. Brihaye, A. Da Costa Lopes, F. Laroussinie, and N. Markey. ATL with Strategy Contexts and Bounded Memory. In *LFCS'09*, LNCS 5407, pages 92–106. Springer, 2009.

[BMM09] A. Bianco, F. Mogavero, and A. Murano. Graded Computation Tree Logic. In *LICS'09*, pages 342–351. IEEE Computer Society, 2009.

[BMM10] A. Bianco, F. Mogavero, and A. Murano. Graded Computation Tree Logic with Binary Coding. In *CSL'10*, LNCS 6247, pages 125–139. Springer, 2010.

[BMM12] A. Bianco, F. Mogavero, and A. Murano. Graded Computation Tree Logic. *TOCL*, 13(3):25:1–53, 2012.

[CE81] E.M. Clarke and E.A. Emerson. Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic. In *LP'81*, LNCS 131, pages 52–71. Springer, 1981.

[CGP02] E.M. Clarke, O. Grumberg, and D.A. Peled. *Model Checking.* MIT Press, 2002.

[CHP07] K. Chatterjee, T.A. Henzinger, and N. Piterman. Strategy Logic. In *CONCUR'07*, LNCS 4703, pages 59–73. Springer, 2007.

[ČLM15] P. Čermák, A. Lomuscio, and A. Murano. Verifying and synthesising multi-agent systems against one-goal strategy logic specifications. In *AAAI'15*, pages 2038–2044. AAAI Press, 2015.

[ČLMM14] P. Čermák, A. Lomuscio, F. Mogavero, and A. Murano. MCMAS-SLK: A Model Checker for the Verification of Strategy Logic Specifications. In *CAV'14*, LNCS 8559, pages 524–531. Springer, 2014.

[EH86] E.A. Emerson and J.Y. Halpern. "Sometimes" and "Not Never" Revisited: On Branching Versus Linear Time. *JACM*, 33(1):151–178, 1986.

[Fin72] K. Fine. In So Many Possible Worlds. *NDJFL*, 13:516–520, 1972.

[FMP08] A. Ferrante, A. Murano, and M. Parente. Enriched Mu-Calculi Module Checking. *LMCS*, 4(3):1–21, 2008.

[FNP09] A. Ferrante, M. Napoli, and M. Parente. Model Checking for Graded CTL. *FI*, 96(3):323–339, 2009.

[FNP10] M. Faella, M. Napoli, and M. Parente. Graded Alternating-Time Temporal Logic. *FI*, 105(1-2):189–210, 2010.

[GKP94] R.L. Graham, D.E. Knuth, and O. Patashnik. *Concrete Mathematics - A Foundation for Computer Science (2nd ed.).* Addison-Wesley, 1994.

[GOR97] E. Grädel, M. Otto, and E. Rosen. Two-Variable Logic with Counting is Decidable. In *LICS'97*, pages 306–317. IEEE Computer Society, 1997.

[HB91] B. Hollunder and F. Baader. Qualifying Number Restrictions in Concept Languages. In *KR'91*, pages 335–346. Morgan Kaufmann, 1991.

[HP85] D. Harel and A. Pnueli. *On the Development of Reactive Systems.* Springer, 1985.

[Imm81] N. Immerman. Number of Quantifiers is Better Than Number of Tape Cells. *JCSS*, 22(3):384–406, 1981.

[JM14] W. Jamroga and A. Murano. On Module Checking and Strategies. In *AAMAS'14*, pages 701–708. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[JM15] W. Jamroga and A. Murano. Module checking of strategic ability. In *AAMAS'15*, pages 227–235. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[KSV02] O. Kupferman, U. Sattler, and M.Y. Vardi. The Complexity of the Graded mu-Calculus. In *CADE'02*, LNCS 2392, pages 423–437. Springer, 2002.

[KV96] O. Kupferman and M.Y. Vardi. Module checking. In *CAV'96*, LNCS 1102, pages 75–86. Springer, 1996.

[KVW01] O. Kupferman, M.Y. Vardi, and P. Wolper. Module Checking. *IC*, 164(2):322–344, 2001.

[LLM10] 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.

[LR06] A. Lomuscio and F. Raimondi. Model Checking Knowledge, Strategies, and Games in Multi-Agent Systems. In *AAMAS'06*, pages 161–168. International Foundation for Autonomous Agents and Multiagent Systems, 2006.

[Mar75] A.D. Martin. Borel Determinacy. *AM*, 102(2):363–371, 1975.

[Mar85] A.D. Martin. A Purely Inductive Proof of Borel Determinacy. In *SPM'82*, Recursion Theory., pages 303–308. American Mathematical Society and Association for Symbolic Logic, 1985.

[MMPV12] 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.

[MMPV14] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. Reasoning About Strategies: On the Model-Checking Problem. *TOCL*, 15(4):34:1–42, 2014.

[MMS13] F. Mogavero, A. Murano, and L. Sauro. On the Boundary of Behavioral Strategies. In *LICS'13*, pages 263–272. IEEE Computer Society, 2013.

[MMS14a] F. Mogavero, A. Murano, and L. Sauro. A Behavioral Hierarchy of Strategy Logic. In *CLIMA'14*, LNCS 8624, pages 148–165. Springer, 2014.

[MMS14b] F. Mogavero, A. Murano, and L. Sauro. Strategy Games: A Renewed Framework. In *AAMAS'14*, pages 869–876. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[MMS15] V. Malvone, A. Murano, and L. Sorrentino. Games with additional winning strategies. In *CILC'15*, CEUR Workshop Proceedings. To appear, pages 1–6, 2015.

[MMV10] F. Mogavero, A. Murano, and M.Y. Vardi. Reasoning About Strategies. In *FSTTCS'10*, LIPIcs 8, pages 133–144. Leibniz-Zentrum fuer Informatik, 2010.

[Mye91] R.B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, 1991.

[Pnu77] A. Pnueli. The Temporal Logic of Programs. In *FOCS'77*, pages 46–57. IEEE Computer Society, 1977.

[PP04] D. Perrin and J. Pin. *Infinite Words.* Pure and Applied Mathematics. Elsevier, 2004.

[QS81] J.P. Queille and J. Sifakis. Specification and Verification of Concurrent Programs in Cesar. In *SP'81*, LNCS 137, pages 337–351. Springer, 1981.

[Sch02] P. Schnoebelen. The complexity of temporal logic model checking. In *AIML'02*, pages 393–436, 2002.