Information and Computation ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Reasoning about graded strategy quantifiers

Vadim Malvone [a], Fabio Mogavero [b], Aniello Murano [a,*], Loredana Sorrentino [a]

[a] Università di Napoli Federico II, Department DIETI, Via Claudio, n. 21, 80125, Napoli, Italy
[b] University of Oxford, Department of Computer Science, Wolfson Building, Parks Road Oxford, OX1 3QD, UK

ARTICLE INFO

Article history:
Received 30 March 2016
Available online xxxx

Keywords:
Strategic reasoning
Strategy logic
Counting quantifiers

ABSTRACT

In this paper we introduce and study *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 x satisfying φ. Dually, via the universal construct $[\![x < g]\!] \varphi$, one can ensure that all but less than g strategies x satisfy φ.

Strategies in GSL are counted semantically. This means that strategies inducing the same outcome, even though looking different, are counted as one. While this interpretation is natural, it requires a suitable machinery to allow for such a counting, as we do. Precisely, we introduce a non-trivial equivalence relation over strategy profiles based on the strategic behavior they induce.

To give an evidence of GSL usability, we investigate some basic questions about the *Vanilla* GSL[1G] fragment, that is the vanilla restriction of the well-studied *One-Goal Strategy Logic* fragment of SL augmented with graded strategy quantifiers. We show that the model-checking problem for this logic is PTIME-COMPLETE. We also report on some positive results about the determinacy.

1. Introduction

Formal methods in system design are a renowned story of success. Breakthrough contributions in this field comprise *model checking* [1,2] and *temporal logics* such as LTL [3], CTL [1], CTL* [4], and the like. First applications of these methodologies involved *closed systems* [5] generally analyzing whether a Kripke structure, modeling the system, meets a temporal logic formula, specifying the desired behavior [6]. In the years several algorithms have been proposed in this setting and some implemented as tools [7]. Nevertheless these approaches turn to be useless when applied to *open systems* [5]. 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 [8,9]. In the last decade this methodology has been fruitfully extended in several directions (see [10–12] for some 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* [13]. One of the most important developments in this field is *Alternating-Time Temporal Logic*

---

[*] Corresponding author.
E-mail address: murano@na.infn.it (A. Murano).

http://dx.doi.org/10.1016/j.ic.2017.08.010

Please cite this article in press as: V. Malvone et al., Reasoning about graded strategy quantifiers, Inf. Comput. (2017), http://dx.doi.org/10.1016/j.ic.2017.08.010

(ATL*, for short), introduced by Alur, Henzinger, and Kupferman [13]. This logic allows to reason about strategies of agents having the satisfaction of temporal goals as the payoff criterion. Formally, it is obtained as a generalization of CTL*, in which the existential $\mathtt{E}$ and the universal $\mathtt{A}$ *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 its 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) [14,15], a more powerful formalism for strategic reasoning. As a key aspect, SL 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 [16,17], for an argumentation). This has attracted the interest of the formal verification community to *graded modalities*. These have been first studied in classic modal logic [18] and then exported to the field of *knowledge representation* to allow quantitative bounds on the set of individuals satisfying specific properties. Specifically, they are *counting quantifiers* in first-order logics [19], *number restrictions* in *description logics* [20–23] and *numerical constraints* in query languages [24].

First applications of graded modalities in formal verification concern closed systems. In [25], *graded* μCALCULUS has been introduced in order to express statements about a given number of immediately accessible worlds. Successively in [26–28, 16], 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 [16]. This work has been recently extended in [29] to address GCTL*, a graded extension of CTL*.

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 [30] and an extension of ATL with graded path modalities (GATL, for short) [31]. 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 different approach by formally introducing a 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. As in SL, we use the binding operator to associate these strategies to agents.

As far as the counting of strategies is concerned, 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 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 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 [32]. 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 [33], *i.e.*, if a strategy fails for any reason, it is possible to apply the others.

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.

Outline The sequel of the paper is structured as follows. Section 2 introduces GSL and provides some preliminaries. Section 3 introduces, by means of axioms, the equivalence relation used to count strategies. Section 4 shows how to transform a game from concurrent to turn-based. Section 5 and Section 6 address the determinacy and the model-checking problem for Vanilla GSL[1G]. Finally we conclude in Section 7 with some discussions and future work.

Please cite this article in press as: V. Malvone et al., Reasoning about graded strategy quantifiers, Inf. Comput. (2017), http://dx.doi.org/10.1016/j.ic.2017.08.010