

Extended Graded Modalities in Strategy Logic

Benjamin Aminof

Vadim Malvone, Aniello Murano, Sasha Rubin

Technische Universität Wien, Austria

Università degli Studi di Napoli Federico II, Italy

benj@forsyte.tuwien.ac.at

{malvone,murano,rubin}@unina.it

Strategy Logic (SL) is a logical formalism for strategic reasoning in multi-agent systems. Its main feature is that it has variables for strategies that are associated to specific agents with a binding operator. We introduce Graded Strategy Logic (GRADED_NSL), an extension of SL by graded quantifiers over tuples of strategy variables, i.e., “there exist at least g different tuples (x_1, \dots, x_n) of strategies” where g is a cardinal from the set $\mathbb{N} \cup \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$. We prove that the model-checking problem of GRADED_NSL is decidable. We then turn to the complexity of fragments of GRADED_NSL. When the g ’s are restricted to finite cardinals, written GRADED_NSL, the complexity of model-checking is no harder than for SL, i.e., it is non-elementary in the quantifier rank. We illustrate our formalism by showing how to count the number of different strategy profiles that are Nash equilibria (NE), or subgame-perfect equilibria (SPE). By analyzing the structure of the specific formulas involved, we conclude that the important problems of checking for the existence of a unique NE or SPE can both be solved in 2EXPTIME, which is not harder than merely checking for the existence of such equilibria.

1 Introduction

Strategy Logic (SL) is a powerful formalism for reasoning about strategies in multi-agent systems [37, 38]. Strategies tell an agent what to do — they are functions that prescribe an action based on the history. The key idea in SL is to treat strategies as first-order object [18]. A strategy x can be quantified existentially $\langle\langle x \rangle\rangle$ (read: there exists a strategy x) and, dually, universally $[[x]]$ (read: for all strategies x). Strategies are not intrinsically glued to specific agents: the *binding* operator (α, x) allows one to bind an agent α to the strategy x . SL is built over the structure of the linear-time temporal-logic LTL [43] and strictly subsumes several other well-known logics for the strategic reasoning including ATL^{*} [3] and the like [1, 14, 24, 27, 28, 34, 49, 50].

We extend SL by replacing the quantification $\langle\langle x \rangle\rangle$ and $[[x]]$ over strategy variables with *graded quantification over tuples of strategy variables*: $\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}$ (read $\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}$ as “there exist at least g different tuples (x_1, \dots, x_n) of strategies”) and its dual $[[x_1, \dots, x_n]]^{< g}$, where $g \in \mathbb{N} \cup \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$. Here, two tuples are different if they are different in some component, and two strategies are different if they disagree on some history. That is, we count strategies syntactically in a way similar as it is usually done in graded extensions of modal and description logics [7, 8, 13, 15, 21, 22, 25, 31, 45].

We address the model-checking problem for GRADED_NSL and, by means of an automata-theoretic approach, we prove that it is decidable. Our algorithm is not elementary, i.e., it is not bounded by any tower of exponentials. Let GRADED_NSL denote the fragment in which all grades are finite (i.e., $g \in \mathbb{N}$). We studied this fragment in [4] and report the results here. The model-checking problem has the same complexity as SL. That is, model checking GRADED_NSL formulas with a nesting depth $k > 0$ of blocks

of quantifiers¹ is in $(k + 1)\text{EXPTIME}$, and that for the special case where the formula starts with a block of quantifiers, it is in $k\text{EXPTIME}$. We also study natural fragments of GRADED_NSL, in line with the fragments of SL. In particular, the Nested-Goal fragment of SL requires that bindings and quantifications appear in exhaustive blocks (see Section 2 for details), and can express Nash Equilibrium (NE) [37]. Similarly, we define Nested-Goal GRADED_NSL, a syntactic fragment of GRADED_NSL. We show that Nested-Goal GRADED_NSL has the same model-checking complexity as Nested-Goal SL, i.e., non-elementary in the alternation number of the quantifiers appearing in the formula.²

Since many natural formulas have a small number of quantifiers, and even smaller nesting depth of blocks of quantifiers, the complexity of the model-checking problem is not as bad as it seems. Several solution concepts can be expressed as SL formulas with a small number of quantifiers [9, 18, 26, 29, 30, 37]. We illustrate this by expressing *uniqueness* of various solution concepts, including winning-strategies in two-player zero-sum games and equilibria in multi-player non zero-sum games. It is important to observe that, when dealing with NE, to establish whether the game admits a *unique* equilibrium is a very challenging and important question [2, 19, 41]. This problem has an impact on the predictive power of NE since, in case there are multiple equilibria, the outcome of the game cannot be pinned down [20, 42, 51]. Before our works involving Strategy Logic, there was no uniform framework to deal with this problem. So far, this question has mainly been addressed algorithmically and for very restrictive game topologies [40]. By means of GRADED_NSL, in this paper we show how to address and efficiently solve this problem and in a general way. In particular, thanks to the small nesting depth of the blocks of quantifiers of the formulas involved, we show that checking the uniqueness of winning strategies, NE, or subgame-perfect equilibria (SPE) can be solved in 2EXPTIME , which is no worse than checking their existence.

Related work. The work closest to ours is [4] that also considers a graded extension of Strategy Logic, aimed at counting strategies in solution concepts. Precisely, strategies in [4] are counted syntactically, as we do here, but graded quantifiers numbers are limited to finite cardinals. This explains the words “extended graded modalities” in the title of our work. In this paper, for the sake of completeness, we have decided to report in summary the main results from [4] as an introductory material. Also, most of the definitions we provide extend those reported in [4] and we refer to that paper for additional comments and observations.

Another work close to ours is [35]. It also introduces a graded extension of SL, called GSL. However, in contrast with our work and [4] (and by extending an idea of counting paths in [5, 10]), GSL provides a quite intricate way of counting strategies: it gives a semantic definition of equivalence of strategies, and the graded modalities count non-equivalent strategies. While this approach is sound, it heavily complicates the model-checking problem. Indeed, only a very weak fragment of GSL has been solved in [35] by exploiting an *ad hoc* solution that does not seem to be scalable to (all of) GSL. Such a fragment is orthogonal to ATL^* and, as far as we know, not enough powerful to represent NE neither SPE.

Outline. The sequel of the paper is organized as follows. In Section 2 we introduce GRADED_NSL and provide some preliminary related concepts. In Section 3 we recall the notion of objective and illustrate how to express some solution concepts and their uniqueness. In Section 4 we address the model-checking problem for GRADED_NSL and its fragments. We conclude with Section 5 in which we have a discussion and suggestions for future work.

¹A *block* of quantifiers is a maximally-consecutive sequence of quantifiers of the same type, i.e., either all existential, or all universal.

²The *alternation number* of a Nested-Goal formula is, roughly speaking, the maximum number of existential/universal quantifier switches in a consecutive sequence of quantifiers.

2 Graded Strategy Logic

In this section we introduce Graded Strategy Logic, which we call GRADED_{SL} for short.

2.1 Models

Sentences of GRADED_{SL} are interpreted over *concurrent game structures*, just as for ATL and SL [3, 37].

Definition 2.1 A concurrent game structure (CGS) is a tuple $\mathcal{G} \triangleq \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, s_I, \text{ap}, \text{tr} \rangle$, where

- AP is a finite set atomic propositions,
- Ag is a finite set of agents,
- Ac is a finite set of actions,
- St is a finite set of states,
- $s_I \in \text{St}$ is the initial state,
- $\text{ap} : \text{St} \rightarrow 2^{\text{AP}}$ is the labeling function mapping each state to the set of atomic propositions true in that state, and
- let $\text{Dc} \triangleq \text{Ag} \rightarrow \text{Ac}$ be the set of decisions, i.e., functions describing the choice of an action by every agent. Then, $\text{tr} : \text{Dc} \rightarrow (\text{St} \rightarrow \text{St})$, a transition function, maps every decision $\delta \in \text{Dc}$ to a function $\text{tr}(\delta) : \text{St} \rightarrow \text{St}$.

We will usually take the set Ag of agents to be $\{\alpha_1, \dots, \alpha_n\}$. A *path* (from s) is a finite or infinite non-empty sequence of states $s_1 s_2 \dots$ such that $s = s_1$ and for every i there exists a decision δ with $\text{tr}(\delta)(s_i) = s_{i+1}$. The set of paths starting with s is denoted $\text{Pth}(s)$. The set of finite paths from s , called the *histories* (from s), is denoted $\text{Hst}(s)$. A *strategy* (from s) is a function $\sigma \in \text{Str}(s) \triangleq \text{Hst}(s) \rightarrow \text{Ac}$ that prescribes which action has to be performed given a history. We write $\text{Pth}, \text{Hst}, \text{Str}$ for the set of all paths, histories, and strategies (no matter where they start). We use the standard notion of equality between strategies, [33], i.e., $\sigma_1 = \sigma_2$ iff for all $\rho \in \text{Hst}$, $\sigma_1(\rho) = \sigma_2(\rho)$. This extends to equality between two n -tuples of strategies in the natural way, i.e., coordinate-wise.

2.2 Syntax

GRADED_{SL} extends SL by replacing the singleton strategy quantifiers $\langle\langle x \rangle\rangle$ and $\llbracket x \rrbracket$ with the graded (tupled) quantifiers $\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}$ and $\llbracket x_1, \dots, x_n \rrbracket^{< g}$, respectively, where each x_i belongs to a countable set of variables Vr and $g \in \mathbb{N} \cup \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$ is called the *grade* of the quantifier. Intuitively, these are read as “there exist at least g tuples of strategies (x_1, \dots, x_n) ” and “all but less than g many tuples of strategies”, respectively. The syntax (α, x) denotes a *binding* of the agent α to the strategy x .

Definition 2.2 GRADED_{SL} formulas are built inductively by means of the following grammar, where $p \in \text{AP}$, $\alpha \in \text{Ag}$, $x, x_1, \dots, x_n \in \text{Vr}$ such that $x_i \neq x_j$ for $i \neq j$, $n \in \mathbb{N}$, and $g \in \mathbb{N} \cup \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$:

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \times\varphi \mid \varphi \cup \varphi \mid \langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}\varphi \mid (\alpha, x)\varphi.$$

Notation. Whenever we write $\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}$ we mean that $x_i \neq x_j$ for $i \neq j$, i.e., the variables in a tuple are distinct. Shorthands are derived as usual. Specifically, $\text{true} \triangleq p \vee \neg p$, $\text{false} \triangleq \neg \text{true}$, $F\varphi \triangleq \text{true} \cup \varphi$, and $G\varphi \triangleq \neg F\neg\varphi$. Also, we have that $\llbracket x_1, \dots, x_n \rrbracket^{< g}\varphi \triangleq \neg\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}\neg\varphi$.

In order to define the semantics, we first define the concept of *free placeholders* in a formula, which refer to agents and variables. Intuitively, an agent or variable is free in φ if it does not have a strategy

associated with it (either by quantification or binding) but one is required in order to ascertain if φ is true or not. The set of *free agents* and *free variables* of a GRADED_{SL} formula φ is given by the function $\text{free} : \text{GRADED}_{\text{SL}} \rightarrow 2^{\text{Ag} \cup \text{Vr}}$ defined as follows:

- $\text{free}(p) \triangleq \emptyset$, where $p \in \text{AP}$;
- $\text{free}(\neg\varphi) \triangleq \text{free}(\varphi)$;
- $\text{free}(\varphi_1 \vee \varphi_2) \triangleq \text{free}(\varphi_1) \cup \text{free}(\varphi_2)$;
- $\text{free}(\mathbf{X}\varphi) \triangleq \text{Ag} \cup \text{free}(\varphi)$;
- $\text{free}(\varphi_1 \cup \varphi_2) \triangleq \text{Ag} \cup \text{free}(\varphi_1) \cup \text{free}(\varphi_2)$;
- $\text{free}(\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g} \varphi) \triangleq \text{free}(\varphi) \setminus \{x_1, \dots, x_n\}$;
- $\text{free}((\alpha, x)\varphi) \triangleq \text{free}(\varphi)$, if $\alpha \notin \text{free}(\varphi)$, where $\alpha \in \text{Ag}$ and $x \in \text{Vr}$;
- $\text{free}((\alpha, x)\varphi) \triangleq (\text{free}(\varphi) \setminus \{\alpha\}) \cup \{x\}$, if $\alpha \in \text{free}(\varphi)$, where $\alpha \in \text{Ag}$ and $x \in \text{Vr}$.

A formula φ without free agents (resp., variables), i.e., with $\text{free}(\varphi) \cap \text{Ag} = \emptyset$ (resp., $\text{free}(\varphi) \cap \text{Vr} = \emptyset$), is called *agent-closed* (resp., *variable-closed*). If φ is both agent- and variable-closed, it is called a *sentence*.

SL has a few natural syntactic fragments, the most powerful of which is called Nested-Goal SL. Similarly, we define *Nested-Goal* GRADED_{SL} (abbreviated GRADED_{SL}[NG]), as a syntactic fragment of GRADED_{SL}. As in SL[NG], in GRADED_{SL}[NG] we require that bindings and quantifications appear in exhaustive blocks. I.e., whenever there is a quantification over a variable in a formula ψ it is part of a consecutive sequence of quantifiers that covers all of the free variables that appear in ψ , and whenever an agent is bound to a strategy then it is part of a consecutive sequence of bindings of all agents to strategies. Finally, formulas with free agents are not allowed. To formalize GRADED_{SL}[NG] we first introduce some notions. A *quantification prefix* over a finite set $V \subseteq \text{Vr}$ of variables is a sequence $\wp \in \{\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}, \llbracket x_1, \dots, x_n \rrbracket^{<g} : n \in \mathbb{N}, x_1, \dots, x_n \in V \wedge g \in \mathbb{N} \cup \{\aleph_0, \aleph_1, 2^{\aleph_0}\}\}^*$ such that each $x \in V$ occurs exactly once in \wp . A *binding prefix* is a sequence $\flat \in \{(a, x) : \alpha \in \text{Ag} \wedge x \in \text{Vr}\}^*$ such that each $\alpha \in \text{Ag}$ occurs exactly once in \flat . We denote the set of binding prefixes by Bn , and the set of quantification prefixes over V by $\text{Qn}(V)$.

Definition 2.3 GRADED_{SL}[NG] formulas are built inductively using the following grammar, with $p \in \text{AP}$, $\wp \in \text{Qn}(V)$ ($V \subseteq \text{Vr}$), and $\flat \in \text{Bn}$:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \cup \varphi \mid \wp\varphi \mid \flat\varphi,$$

where in the rule $\wp\varphi$ we require that φ is agent-closed and $\wp \in \text{Qn}(\text{free}(\varphi))$.

Formulas of GRADED_{SL}[NG] can be classified according to their *alternation number*, i.e., the maximum number of quantifier switches in a quantification prefix. Formally:

Definition 2.4 The alternation number of a GRADED_{SL}[NG] formula is given by:

- $\text{alt}(p) \triangleq 0$, where $p \in \text{AP}$;
- $\text{alt}(\text{OP}\varphi) \triangleq \text{alt}(\varphi)$, where $\text{OP} \in \{\neg, \mathbf{X}, \flat\}$;
- $\text{alt}(\varphi_1 \text{OP} \varphi_2) \triangleq \max(\text{alt}(\varphi_1), \text{alt}(\varphi_2))$ where $\text{OP} \in \{\vee, \cup\}$;
- $\text{alt}(\wp\varphi) \triangleq \max(\text{alt}(\varphi), \text{alt}(\wp))$ where \wp is a quantification prefix;
- $\text{alt}(\flat) \triangleq \sum_{i=1}^{|\flat|-1} \text{switch}(\flat_i, \flat_{i+1})$, where $\text{switch}(Q, Q') = 0$ if Q and Q' are either both universal or both existential quantifiers, and 1 otherwise.

The *quantifier rank* of φ is the maximum nesting of quantifiers in φ , e.g., $\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g}(\alpha_1, x_1) \dots (\alpha_n, x_n) \wedge_{i=1}^n (\langle\langle y \rangle\rangle(\alpha_i, y) \psi_i) \rightarrow \psi_i$ has quantifier rank 2 if each ψ_i is quantifier free. Moreover, a *quantifier-block* of φ is a maximally-consecutive sequence of quantifiers in φ of the same type (i.e., either all existential, or all universal). The *quantifier-block rank* of φ is exactly like the quantifier rank except that a quantifier block of j quantifiers contributes 1 instead of j to the count.

We conclude this subsection by introducing *One-Goal GRADED SL*, written $\text{GRADED SL}[1G]$. The importance of this fragment in SL stems from the fact that it strictly includes ATL^* while maintaining the same complexity for both the model checking and the satisfiability problems, i.e. 2EXPTIME-COMplete [37]. However, it is commonly believed that Nash Equilibrium cannot be expressed in this fragment. The definition of $\text{GRADED SL}[1G]$ follows.

Definition 2.5 $\text{GRADED SL}[1G]$ formulas are built inductively using the following grammar, with $p \in \text{AP}$, $\wp \in \text{Qn}(\text{V})$ ($\text{V} \subseteq \text{Vr}$), and $b \in \text{Bn}$:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid X\varphi \mid \varphi \cup \varphi \mid \wp b \varphi,$$

where \wp is a quantification prefix over $\text{free}(b\varphi)$.

Finally, we subscript with \mathbb{N} and write $\text{GRADED}_{\mathbb{N}}\text{SL}$, $\text{GRADED}_{\mathbb{N}}\text{SL}[1G]$, and $\text{GRADED}_{\mathbb{N}}\text{SL}[\text{NG}]$ for the fragments in which all grades are from the set \mathbb{N} .

2.3 Semantics

As for SL, the interpretation of a GRADED SL formula requires a valuation of its free placeholders.

Definition 2.6 An assignment (from s) is a function $\chi \in \text{Asg}(s) \triangleq (\text{Vr} \cup \text{Ag}) \rightarrow \text{Str}(s)$ mapping variables and agents to strategies.

We denote by $\chi[e \mapsto \sigma]$, with $e \in \text{Vr} \cup \text{Ag}$ and $\sigma \in \text{Str}(s)$, the assignment that differs from χ only in the fact that e maps to σ . Extend this definition to tuples: for $\bar{e} = (e_1, \dots, e_n)$ with $e_i \neq e_j$ for $i \neq j$, define $\chi[\bar{e} \mapsto \bar{\sigma}]$ to be the assignment that differs from χ only in the fact that e_i maps to σ_i (for each i).

Since an assignment ensures that all free variables are associated with strategies, it induces a play.

Definition 2.7 For an assignment $\chi \in \text{Asg}(s)$ the (χ, s) -play denotes the path $\pi \in \text{Pth}(s)$ such that for all $i \in \mathbb{N}$, it holds that $\pi_{i+1} = \text{tr}(\text{dc})(\pi_i)$, where $\text{dc}(\alpha) \triangleq \chi(\alpha)(\pi_{\leq i})$ for $\alpha \in \text{Ag}$. The function $\text{play} : \text{Asg} \times \text{St} \rightarrow \text{Pth}$, with $\text{dom}(\text{play}) \triangleq \{(\chi, s) : \chi \in \text{Asg}(s)\}$, maps (χ, s) to the (χ, s) -play $\text{play}(\chi, s) \in \text{Pth}(s)$.

The notation $\pi_{\leq i}$ (resp. $\pi_{< i}$) denotes the prefix of the sequence π of length i (resp. $i - 1$). Similarly, the notation π_i denotes the i th symbol of π . Thus, $\text{play}(\chi, s)_i$ is the i th state on the play determined by χ from s .

The following definition of χ_i says how to interpret an assignment χ starting from a point i along the play, i.e., for each placeholder e , take the action the strategy $\chi(e)$ would do if it were given the prefix of the play up to i followed by the current history.

Definition 2.8 For $\chi \in \text{Asg}(s)$ and $i \in \mathbb{N}$, writing $\rho \triangleq \text{play}(\chi, s)_{\leq i}$ (the prefix of the play up to i) and $t \triangleq \text{play}(\chi, s)_i$ (the last state of ρ) define $\chi_i \in \text{Asg}(t)$ to be the assignment from t that maps $e \in \text{Vr} \cup \text{Ag}$ to the strategy that maps $h \in \text{Hst}(t)$ to the action $\chi(e)(\rho_{< i} \cdot h)$.

The semantics of GRADED SL mimics the one for SL as given in [37]. Given a CGS \mathcal{G} , for all states $s \in \text{St}$ and assignments $\chi \in \text{Asg}(s)$, we now define the relation $\mathcal{G}, \chi, s \models \varphi$, read φ holds at s in \mathcal{G} under χ .

Definition 2.9 Fix a CGS \mathcal{G} . For all states $s \in \text{St}$ and assignments $\chi \in \text{Asg}(s)$, the relation $\mathcal{G}, \chi, s \models \varphi$ is defined inductively on the structure of φ :

- $\mathcal{G}, \chi, s \models p$ iff $p \in \text{ap}(s)$;
- $\mathcal{G}, \chi, s \models \neg\varphi$ iff $\mathcal{G}, \chi, s \not\models \varphi$;
- $\mathcal{G}, \chi, s \models \varphi_1 \vee \varphi_2$ iff $\mathcal{G}, \chi, s \models \varphi_1$ or $\mathcal{G}, \chi, s \models \varphi_2$;
- $\mathcal{G}, \chi, s \models \mathsf{X}\varphi$ iff $\mathcal{G}, \chi_1, \text{play}(\chi, s)_1 \models \varphi$;
- $\mathcal{G}, \chi, s \models \varphi_1 \cup \varphi_2$ iff there is an index $i \in \mathbb{N}$ such that $\mathcal{G}, \chi_i, \text{play}(\chi, s)_i \models \varphi_2$ and, for all indexes $j \in \mathbb{N}$ with $j < i$, it holds that $\mathcal{G}, \chi_j, \text{play}(\chi, s)_j \models \varphi_1$;
- $\mathcal{G}, \chi, s \models (\alpha, x)\varphi$ iff $\mathcal{G}, \chi[\alpha \mapsto \chi(x)], s \models \varphi$;
- $\mathcal{G}, \chi, s \models \langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g} \varphi$ iff there exist g many n -tuples of strategies $\bar{\sigma}_i$ ($0 \leq i < g$) such that:
 - $\bar{\sigma}_i \neq \bar{\sigma}_j$ for $i \neq j$, and
 - $\mathcal{G}, \chi[\bar{x} \mapsto \bar{\sigma}_i], s \models \varphi$ for $0 \leq i < g$.

Intuitively, $\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq g} \varphi$ expresses that the number of distinct tuples of strategies that satisfy φ is at least g .

As usual, if χ and χ' agree on $\text{free}(\varphi)$, then $\mathcal{G}, \chi, s \models \varphi$ if and only if $\mathcal{G}, \chi', s \models \varphi$, i.e., the truth of φ does not depend on the values the assignment takes on placeholders that are not free. Thus, for a sentence φ , we write $\mathcal{G} \models \varphi$ to mean that $\mathcal{G}, \chi, s_I \models \varphi$ for some (equivalently, for all) assignments χ , and where s_I is the initial state of \mathcal{G} .

3 Illustrating GRADED_{SL}: uniqueness of solutions

In game theory, players have objectives that are summarized in a payoff function that maps plays to real numbers. In order to specify such payoffs with formulas, we follow a formalization from [30] called *objective LTL*. We then discuss appropriate solution concepts, and show how to express these in GRADED_{SL}.

Let \mathcal{G} be a CGS with n agents. Let $m \in \mathbb{N}$ and fix, for each agent $\alpha_i \in \text{Ag}$, an *objective* tuple $S_i \triangleq \langle f_i, \varphi_i^1, \dots, \varphi_i^m \rangle$, where $f_i : \{0, 1\}^m \rightarrow \mathbb{Z}$, and each φ_i^j is an LTL formula over AP. If π is a play, then agent α_i receives payoff $f_i(\bar{h}) \in \mathbb{N}$ where \bar{h}_j , the j th bit of \bar{h} , is 1 if and only if $\pi \models \varphi_i^j$. We assume agents are trying to maximize their payoffs.

In case each $f_i : \{0, 1\}^m \rightarrow \{-1, 1\}$ we say that the game is *win/lose*. If $\sum_{1 \leq i \leq n} f_i(\bar{h}) = 0$ for all π , then \mathcal{G} is a *zero-sum game*, otherwise it is a *non zero-sum game*.

Two player, Win/Lose, Zero-sum games. In these types of games the main solution concept is the *winning strategy*. In general, a strategy is winning for agent α_i if and only if for all strategies of adversarial agents the resulting induced play has payoff 1 for agent i . In [36] the authors describe a two-player game named “Cop and the Robber”, played in a maze, in which the objective of the Robber is to reach an exit (and thus the objective of the Cop is to ensure the Robber never reaches the exit). The authors describe two closely related mazes in which the Robber has, respectively, exactly one and exactly two winning strategies. Both these properties can be easily expressed by GSL. For instance, the Robber has a single LTL objective $\varphi \triangleq F \text{exit}$, and the following formula of GRADED_NSL expresses that the Robber has exactly two winning strategies:

$$\langle\langle x \rangle\rangle^{\geq 2} \llbracket y \rrbracket (\text{Robber}, x)(\text{Cop}, y)\varphi \wedge \neg \langle\langle x \rangle\rangle^{\geq 3} \llbracket y \rrbracket (\text{Robber}, x)(\text{Cop}, y)\varphi$$

Non zero-sum games. The central solution concept in these type of games is the Nash Equilibrium. A tuple of strategies, one for each player, is called a *strategy profile*. A strategy profile is a *Nash equilibrium* (NE) if no agent can increase his payoff by unilaterally choosing a different strategy. A game may have zero, one, or many NE. Suppose $\phi_{NE}(\bar{x})$ expresses that the profile $\bar{x} \triangleq (x_1 \dots x_n)$ is a NE. Then the following formula expresses that there is a unique NE:

$$\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq 1} \phi_{NE}(\bar{x}) \wedge \neg \langle\langle x_1, \dots, x_n \rangle\rangle^{\geq 2} \phi_{NE}(\bar{x}). \quad (1)$$

We now describe ϕ_{NE} , first in the win/lose setting, and then in the general case.

First, consider a win/lose setting in which the objective of agent $\alpha_i \in \text{Ag}$ is a single LTL formula φ_i and the payoff is 1 if the formula is true and -1 otherwise. The following formula expresses that $\bar{x} \triangleq (x_1 \dots x_n)$ is a Nash Equilibrium:

$$[[y_1]] \dots [[y_n]] \bigwedge_{i=1}^n (b_i \varphi_i \rightarrow b \varphi_i)$$

where $b = (\alpha_1, x_1) \dots (\alpha_n, x_n)$, and $b_i = (\alpha_1, x_1) \dots (\alpha_{i-1}, x_{i-1}) (\alpha_i, y_i) (\alpha_{i+1}, x_{i+1}) \dots (\alpha_n, x_n)$. Although this is not in the nested-goal fragment, (the quantifiers don't bind the x s), plugging it into Formula (1) results in a formula of $\text{GRADED}_{\mathbb{N}}\text{SL}[\text{NG}]$.

Next, consider the case that each agent α_i has a general objective tuple $S_i \triangleq \langle f_i, \varphi_i^1, \dots, \varphi_i^m \rangle$. Given a vector $\bar{h} \in \{0, 1\}^m$, let $gd_i(\bar{h}) \triangleq \{\bar{h}' \in \{0, 1\}^m \mid f_i(\bar{h}') \geq f_i(\bar{h})\}$ be the set of vectors \bar{h}' for which the payoff for agent α_i is not worse than for \bar{h} . Also, let $\eta_i^{\bar{h}}$ be the formula obtained by taking a conjunction of the formulas $\varphi_i^1, \dots, \varphi_i^m$ or their negations according to \bar{h} , i.e., by taking φ_a^j if the j th bit in \bar{h} is 1, and otherwise taking $\neg \varphi_a^j$. Formally, $\eta_i^{\bar{h}} \triangleq \bigwedge_{j \in \{1 \leq j \leq m \mid \bar{h}_j = 1\}} \varphi_i^j \wedge \bigwedge_{j \in \{1 \leq j \leq m \mid \bar{h}_j = 0\}} \neg \varphi_i^j$. Observe that the following formula says that $\bar{x} \triangleq (x_1 \dots x_n)$ is a Nash Equilibrium:

$$[[y_1]] \dots [[y_n]] \bigwedge_{i=1}^n \bigwedge_{\bar{h} \in \{0, 1\}^m} (b_i \eta_i^{\bar{h}} \rightarrow \bigvee_{\bar{h}' \in gd_i(\bar{h})} b \eta_i^{\bar{h}'})$$

Again, plugging it into Formula (1) results in a formula of $\text{GRADED}_{\mathbb{N}}\text{SL}[\text{NG}]$ expressing there is a unique NE.

It has been argued (in [30, 48]) that NE may be implausible when used for sequential games (of which iterated one shot games are central examples), and that a more robust notion is subgame-perfect equilibrium [46]. Given a game \mathcal{G} , a strategy profile is a *subgame-perfect equilibrium* (SPE) if starting at any reachable subgame, the profile is a NE. The following formula expresses that $\bar{x} \triangleq (\alpha_1, x_1) \dots (\alpha_n, x_n)$ is an SPE:

$$\phi_{SPE}(\bar{x}) \triangleq [[z_1, \dots, z_n]] (\alpha_1, z_1) \dots (\alpha_n, z_n) \text{G} \phi_{NE}(\bar{x})$$

Using graded modalities, we can thus express there is a unique SPE using the following $\text{GRADED}_{\mathbb{N}}\text{SL}[\text{NG}]$ formula:

$$\langle\langle x_1, \dots, x_n \rangle\rangle^{\geq 1} \phi_{SPE}(\bar{x}) \wedge \neg \langle\langle x_1, \dots, x_n \rangle\rangle^{\geq 2} \phi_{SPE}(\bar{x}).$$

4 Model-checking GRADED_{SL}

In this section we study the model-checking problem for GRADED_{SL} and show that it is decidable with a time-complexity that is non-elementary (i.e., not bounded by any fixed tower of exponentials). However, it

is elementary if the number of blocks of quantifiers is fixed. For the algorithmic procedures, we follow an *automata-theoretic approach* [32], reducing the decision problem for the logic to the emptiness problem of an automaton. The procedure we propose here extends that used for SL in [37]. The only case that is different is the new graded quantifier over tuples of strategies.

We start with the central notions of automata theory, and then show how to convert a GRADED SL sentence φ into an automaton that accepts exactly the (tree encodings) of the concurrent game structures that satisfy φ . This is used to prove the main result about GRADED SL model checking.

Automata Theory. A Σ -labeled Υ -tree T is a pair $\langle T, \nu \rangle$ where $T \subseteq \Upsilon^+$ is prefix-closed (i.e., if $t \in T$ and $s \in \Upsilon^+$ is a prefix of t then also $s \in T$), and $\nu : T \rightarrow \Sigma$ is a labeling function. Note that every word $w \in \Upsilon^+ \cup \Upsilon^\omega$ with the property that every prefix of w is in T , can be thought of as a path in T . Infinite paths are called *branches*.

Nondeterministic tree automata (NTA) are a generalization to infinite trees of the classical automata on words [47]. *Alternating tree automata* (ATA) are a further generalization of nondeterministic tree automata [23]. Intuitively, on visiting a node of the input tree, while an NTA sends exactly one copy of itself to each of the successors of the node, an ATA can send several copies to the same successor. We use the parity acceptance condition [32].

For a set X , let $B^+(X)$ be the set of positive Boolean formulas over X , including the constants **true** and **false**. A set $Y \subseteq X$ satisfies a formula $\theta \in B^+(X)$, written $Y \models \theta$, if assigning **true** to elements in Y and **false** to elements in $X \setminus Y$ makes θ true.

Definition 4.1 An *Alternating Parity Tree-Automaton* (APT) is a tuple $\mathcal{A} \triangleq \langle \Sigma, \Delta, Q, \delta, q_0, \mathcal{F} \rangle$, where

- Σ is the input alphabet,
- Δ is a set of directions,
- Q is a finite set of states,
- $q_0 \in Q$ is an initial state,
- $\delta : Q \times \Sigma \rightarrow B^+(\Delta \times Q)$ is an alternating transition function, and
- \mathcal{F} , an acceptance condition, is of the form $(F_1, \dots, F_k) \in (2^Q)^+$ where $F_1 \subseteq F_2 \dots \subseteq F_k = Q$.

The set $\Delta \times Q$ is called the set of *moves*. An NTA is an ATA in which each conjunction in the transition function δ has exactly one move (d, q) associated with each direction d .

An *input tree* for an APT is a Σ -labeled Δ -tree $T = \langle T, \nu \rangle$. A *run* of an APT on an input tree $T = \langle T, \nu \rangle$ is a $(\Delta \times Q)$ -tree R such that, for all nodes $x \in R$, where $x = (d_1, q_1) \dots (d_n, q_n)$ (for some $n \in \mathbb{N}$), it holds that (i) $y \triangleq (d_1, \dots, d_n) \in T$ and (ii) there is a set of moves $S \subseteq \Delta \times Q$ with $S \models \delta(q_n, \nu(y))$ such that $x \cdot (d, q) \in R$ for all $(d, q) \in S$.

The acceptance condition allows us to say when a run is successful. Let R be a run of an APT \mathcal{A} on an input tree T and $u \in (\Delta \times Q)^\omega$ one of its branches. Let $\text{inf}(u) \subseteq Q$ denote the set of states that occur in infinitely many moves of u . Say that u satisfies the *parity acceptance condition* $\mathcal{F} = (F_1, \dots, F_k)$ if the least index $i \in [1, k]$ for which $\text{inf}(u) \cap F_i \neq \emptyset$ is even. A run is *successful* if all its branches satisfy the parity acceptance condition \mathcal{F} . An APT *accepts* an input tree T iff there exists a successful run R of \mathcal{A} on T .

The *language* $L(\mathcal{A})$ of the APT \mathcal{A} is the set of trees T accepted by \mathcal{A} . Two automata are *equivalent* if they have the same language. The *emptiness problem* for alternating parity tree-automata is to decide, given \mathcal{A} , whether $L(\mathcal{A}) = \emptyset$. The *universality problem* is to decide whether \mathcal{A} accepts all trees.

4.1 From Logic to Automata

Following an automata-theoretic approach, we reduce the model-checking problem of GRADED_{SL} to the emptiness problem for alternating parity tree automata [37]. The main step is to translate every GRADED_{SL} formula φ (i.e., φ may have free placeholders), concurrent-game structure \mathcal{G} , and state s , into an APT that accepts a tree if and only if the tree encodes an assignment χ such that $\mathcal{G}, \chi, s \models \varphi$.

We first describe the encoding, following [37]. Informally, the CGS \mathcal{G} is encoded by its “tree-unwinding starting from s ” whose nodes represent histories, i.e., the St-labeled St-tree $\mathbb{T} \triangleq \langle \text{Hst}(s), u \rangle$ such that $u(h)$ is the last symbol of h . Then, every strategy $\chi(e)$ with $e \in \text{free}(\varphi)$ is encoded as an Ac-labeled tree over the unwinding. The unwinding and these strategies $\chi(e)$ are viewed as a single $(\text{VAL} \times \text{St})$ -labeled tree where $\text{VAL} \triangleq \text{free}(\varphi) \rightarrow \text{Ac}$.

Definition 4.2 *The encoding of χ (w.r.t. φ, \mathcal{G}, s) is the $(\text{VAL} \times \text{St})$ -labeled St-tree $\mathbb{T} \triangleq \langle \mathbb{T}, u \rangle$ such that \mathbb{T} is the set of histories h of \mathcal{G} starting with s and $u(h) \triangleq (f, q)$ where q is the last symbol in h and $f : \text{free}(\varphi) \rightarrow \text{Ac}$ is defined by $f(e) \triangleq \chi(e)(h)$ for all $e \in \text{free}(\varphi)$.³*

The following lemma is proved by induction on the structure of the formula φ , as in [37]. The idea for handling the new case, i.e., the graded quantifier $\langle \langle x_1, \dots, x_n \rangle \rangle^{\geq g} \psi$, is to build an APT that is a projection of an APT that itself checks that each of the g tuples of strategies satisfies ψ and that each pair of g tuples is distinct.

Lemma 4.1 *For every GRADED_{SL} formula φ , CGS \mathcal{G} , and state $s \in \text{St}$, there exists an APT \mathcal{A}_φ such that for all assignments χ , if \mathbb{T} is the encoding of χ (w.r.t. φ, \mathcal{G}, s), then $\mathcal{G}, \chi, s \models \varphi$ iff $\mathbb{T} \in \text{L}(\mathcal{A}_\varphi)$.*

Proof 4.1 *As in [37] we induct on the structure of the formula φ to construct the corresponding automaton \mathcal{A}_φ . The Boolean operations are easily dealt with using the fact that disjunction corresponds to non-determinism, and negation corresponds to dualizing the automaton. Note (\dagger) that thus also conjunction is dealt with due to De Morgan’s laws. The temporal operators are dealt with by following the unique play (determined by the given assignment) and verifying the required subformulas, e.g., for $\text{X} \psi$ the automaton, after taking one step along the play, launches a copy of the automaton for ψ . All of these operations incur a linear blowup in the size of the automaton. The only case that differs from [37] is the quantification, i.e., we need to handle the case that $\varphi = \langle \langle x_1, \dots, x_n \rangle \rangle^{\geq g} \psi$. Recall that $\mathcal{G}, \chi, s \models \langle \langle x_1, \dots, x_n \rangle \rangle^{\geq g} \psi$ iff there exists g many tuples $\bar{\sigma}_i$ of strategies such that: $\bar{\sigma}_a \neq \bar{\sigma}_b$ for $a \neq b$, and $\mathcal{G}, \chi[\bar{x} \mapsto \bar{\sigma}_i], s \models \psi$ for $0 \leq i < g$.*

There are two cases.

Case $g \in \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$.

Write MSOL for monadic second-order logic in the signature of trees. We use the following two results.

1. *For every MSOL formula $\alpha(Y)$ there exists an APT accepting the set of trees Y such that $\alpha(Y)$ holds. Conversely, for every APT there is an MSOL formula $\alpha(Y)$ that holds on those Y that are accepted by the APT, see [47].*
2. *For every MSOL formula $\alpha(\bar{X}, Y)$ and $\kappa \in \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$ there exists an MSOL formula $\beta(\bar{X})$ equivalent to “there exist κ many trees Y such that $\alpha(\bar{X}, Y)$ ” [6].*

Thus, consider $\langle \langle x_1, \dots, x_n \rangle \rangle^{\geq g} \psi$. By induction, there is an APT \mathcal{D} for ψ . Compile \mathcal{D} into an MSOL formula α , and then produce an MSOL formula β that holds iff “there exist κ many tuples of trees such that α ” (recall that a tuple of trees is coded as a single tree). Finally, convert β back into an APT.

³In case $\text{free}(\varphi) = \emptyset$, then f is the (unique) empty function. In this case, the encoding of every χ may be viewed as the tree-unwinding from s .

Note that the blowup in the translations (MSOL to APT, and closure under “there exists κ many trees”) is non-elementary.

Case $g \in \mathbb{N}$.

We show how to build an NPT for φ that mimics this definition: it will be a projection of an APT \mathcal{D}_ψ , which itself is the intersection of two automata, one checking that each of the g tuples of strategies satisfies ψ , and the other checking that each pair of the g tuples of strategies is distinct.

In more detail, introduce a set of fresh variables $X \triangleq \{x_i^j \in \text{Vr} : i \leq n, j \leq g\}$, and consider the formulas ψ^j (for $j \leq g$) formed from ψ by renaming x_i (for $i \leq n$) to x_i^j . Define $\psi' \triangleq \bigwedge_{j \leq g} \psi^j$. Note that, by induction, each ψ^j has a corresponding APT, and thus, using the conjunction-case (\dagger) above, there is an APT \mathcal{B} for ψ' . Note that the input alphabet for \mathcal{B} is $(\text{free}(\psi') \rightarrow \text{Ac}) \times \text{St}$ and that $X \subseteq \text{free}(\psi')$.

On the other hand, let \mathcal{C} be an APT with input alphabet $(\text{free}(\psi') \rightarrow \text{Ac}) \times \text{St}$ that accepts a tree $T = \langle T, \nu \rangle$ if and only if for every $a \neq b \leq g$ there exists $i \leq n$ and $h \in T$ such that $\nu(h) = (f, q)$ and $f(x_i^a) \neq f(x_i^b)$.

Form the APT \mathcal{D}_ψ for the intersection of \mathcal{B} and \mathcal{C} .

Now, using the classic transformation [39], we remove alternation from the APT \mathcal{D}_ψ to get an equivalent NPT \mathcal{N} (note that this step costs an exponential). Finally, use the fact that NPTs are closed under projection (with no blowup) to get an NPT for the language $\text{proj}_X(L(\mathcal{N}))$ of trees that encode assignments χ satisfying φ .

For completeness we recall this last step. If L is a language of Σ -labeled trees with $\Sigma \triangleq A \rightarrow B$, and $X \subset A$, then the X -projection of L , written $\text{proj}_X(L)$, is the language of Σ' -labeled trees with $\Sigma' \triangleq A \setminus X \rightarrow B$ such that $T \triangleq \langle T, \nu \rangle \in \text{proj}_X(L)$ if and only if there exists an X -labeled tree $\langle T, w \rangle$ such that the language L contains the tree $\langle T, u \rangle$ where $u : T \rightarrow (A \rightarrow B)$ maps $t \in T$ to $\nu(t) \cup w(t)$. Now, if \mathcal{N} is an NPT with input alphabet $\Sigma \triangleq A \rightarrow B$, and if $X \subset A$, then there is an NPT with input alphabet $\Sigma' \triangleq A \setminus X \rightarrow B$ with language $\text{proj}_X(L(\mathcal{N}))$.

The proof that the construction is correct is immediate. \square

We make some remarks about this lemma. First, all the cases in the induction incur a linear blowup except for the quantification case. For the quantification case, if $g \in \mathbb{N}$ then the translation incurs an exponential blowup (the translation from an APT to an NPT results in an exponentially larger automaton [32]); in case $g \in \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$ the blowup is non-elementary. In case all grades are from \mathbb{N} , we can say a little more. Note that a block of k identical quantifiers only costs, in the worst case, a single exponential (and not k many exponentials) because we can extend the proof above to deal with a block of quantifiers at once. Thus, we get that the size of the APT for φ is a tower of exponentials whose height is the quantifier-block rank of φ .

Theorem 4.1 *The model-checking problem for GRADEDSL is decidable. Regarding complexity:*

1. *The complexity is not bounded by any fixed tower of exponentials.*
2. *The complexity is PTIME-COMPLETE w.r.t. the size of the model.*
3. *If all grades are restricted to be in \mathbb{N} , then the complexity is $(k+1)\text{EXPTIME}$ if $k \geq 1$ is the quantifier-block rank of φ . Moreover, if φ is the form $\wp\psi$, where \wp is a quantifier-block, and ψ is of quantifier-block rank $k-1$, then the complexity is $k\text{EXPTIME}$.*

Proof 4.2 *The lower-bounds already holds for SL [37]. For decidability, use Lemma 4.1 to transform the CGS and φ into an APT and test its emptiness. For the upper bound in item 2., use the fact that the membership problem for APT is in PTIME. For the upper bound in item 3., proceeds as follows. The complexity of checking emptiness (or indeed, universality) of an APT is in EXPTIME [32]. As discussed*

after Lemma 4.1, for the case that all grades are in \mathbb{N} , the size of the APT is a tower of exponentials whose height is the quantifier-block rank of φ . This gives the $(k + 1)\text{EXPTIME}$ upper bound.

Theorem 4.2 *The model-checking problem for $\text{GRADED}_{\mathbb{N}}\text{SL}_{[\text{NG}]}$ is PTIME-COMplete w.r.t. the size of the model, and $(k + 1)\text{-EXPTIME}$ when restricted to formulas of maximum alternation number k and grades in \mathbb{N} .*

Proof 4.3 *The lower bound already holds for $\text{SL}_{[\text{NG}]}$ [37], and the PTIME upper bound is inherited from Theorem 4.1. The upper bound for $\text{GRADED}_{\mathbb{N}}\text{SL}_{[\text{NG}]}$ is obtained by following the same reasoning for $\text{SL}_{[\text{NG}]}$ of the singleton existential quantifier [37] but using the automaton construction as in Theorem 4.1.*

Directly from the statements reported above, we get the following results:

Theorem 4.3 *Checking the uniqueness of NE, and checking the uniqueness of SPE, can be done in 2EXPTIME .*

We conclude this section with the complexity of the model checking problem for $\text{GRADED}_{\mathbb{N}}\text{SL}_{[1\text{G}]}$. Also in this case one can derive the lower bound from the one holding for the corresponding sub-logic in SL (i.e., $\text{SL}_{[1\text{G}]}$) and the upper bound by using the same algorithm for $\text{SL}_{[1\text{G}]}$ but plugging a (yet no more complex) different automata construction for the new existential quantifier modality. Indeed the model checking problem for $\text{GRADED}_{\mathbb{N}}\text{SL}_{[1\text{G}]}$ is 2EXPTIME-COMplete . It is worth recalling that $\text{SL}_{[1\text{G}]}$ strictly subsumes ATL^* [37]. It is quite immediate to see that this also holds in the graded setting (note that ATL^* already allows quantifying over tuples of agents' (bound) strategies). As the model checking for ATL^* is already 2EXPTIME-HARD , we get that also for the graded extension for this logic, which we name GATL^* , the model checking problem is 2EXPTIME-COMplete . The model checking results for both GATL^* and $\text{GRADED}_{\mathbb{N}}\text{SL}_{[1\text{G}]}$ are reported in the following theorem.

Theorem 4.4 *The model-checking problem for GATL^* and $\text{GRADED}_{\mathbb{N}}\text{SL}_{[1\text{G}]}$ is PTIME-COMplete w.r.t. the size of the model and 2EXPTIME-COMplete in the size of formula.*

5 Conclusion

We introduced $\text{GRADED}_{\mathbb{N}}\text{SL}$, a logic for strategic reasoning that has the ability to count tuples of strategies, i.e., partial strategy-profiles. The logic is elegant, simple, and powerful: indeed, we can use it to solve the unique NE problem for LTL objectives in 2EXPTIME , and thus at the same complexity that is required to merely decide if a NE exists.

The positive results presented in this paper open several directions for future work:

1. Study the relative expressive power of fragments of GSL .
2. In order to express uniqueness of solution concepts, we added grades to the strategy quantifiers. One may, instead, study the expressive power of SL with the atomic expression $x = y$ (saying that strategies x and y are the same), just as one does for first-order logic with and without equality. Note that equality cannot replace the infinite grades. However, if one restricts to finite grades, are there advantages of doing it one way rather than the other?
3. We showed how to compile quantified formulas into automata. Other generalized quantifiers have been studied in the context of automata and monadic second-order logic [11, 12, 44], including the boundedness quantifier and the modulo-counting quantifiers. Thus, one might try understand which generalized quantifiers are useful for strategic reasoning.
4. One might implement $\text{GRADED}_{\mathbb{N}}\text{SL}$ and its model-checking procedure in a formal verification tool such as SLK-MCMAS [16, 17].

Acknowledgments

We thank Michael Wooldridge for suggesting uniqueness of Nash Equilibria as an application of graded strategy logic. Benjamin Aminof is supported by the Austrian National Research Network S11403-N23 (RiSE) of the Austrian Science Fund (FWF) and by the Vienna Science and Technology Fund (WWTF) through grant ICT12-059. Sasha Rubin is a Marie Curie fellow of the Istituto Nazionale di Alta Matematica. Aniello Murano is partially supported by the GNCS 2016 project: Logica, Automi e Giochi per Sistemi Auto-adattivi.

References

- [1] T. Ågotnes, V. Goranko & W. Jamroga (2007): *Alternating-time temporal logics with irrevocable strategies*. In: *TARK-2007*, pp. 15–24, doi:10.1145/1324249.1324256.
- [2] E. Altman, H. Kameda & Y. Hosokawa (2002): *Nash Equilibria in Load Balancing in Distributed Computer Systems*. *IGTR* 4(2), pp. 91–100, doi:10.1142/S0219198902000574.
- [3] R. Alur, T.A. Henzinger & O. Kupferman (2002): *Alternating-Time Temporal Logic*. *JACM* 49(5), pp. 672–713, doi:10.1145/585265.585270.
- [4] B. Aminof, V. Malvone, A. Murano & S. Rubin (2016): *Graded Strategy Logic: Reasoning about Uniqueness of Nash Equilibria*. In: *AAMAS 2016, IFAAMAS*, pp. 698–706.
- [5] B. Aminof, A. Murano & S. Rubin (2015): *On CTL* with Graded Path Modalities*. In: *LPAR-20 2016*, pp. 281–296, doi:10.1007/978-3-662-48899-7_20.
- [6] V. Bárány, L. Kaiser & A. M. Rabinovich (2010): *Expressing Cardinality Quantifiers in Monadic Second-Order Logic over Trees*. *Fundam. Inform.* 100(1-4), pp. 1–17, doi:10.3233/FI-2010-260.
- [7] E. Bárcenas, P. Genevès, N. Layaïda & A. Schmitt (2011): *Query Reasoning on Trees with Types, Interleaving, and Counting*. In: *IJCAI 2011*, pp. 718–723, doi:10.5591/978-1-57735-516-8/IJCAI11-127.
- [8] E. Bárcenas & J. Lavelle (2014): *Global Numerical Constraints on Trees*. *Logical Methods in Computer Science* 10(2), doi:10.2168/LMCS-10(2:10)2014.
- [9] F. Belardinelli (2015): *A Logic of Knowledge and Strategies with Imperfect Information*. In: *LAMAS 15*.
- [10] A. Bianco, F. Mogavero & A. Murano (2012): *Graded Computation Tree Logic*. *TOCL* 13(3), pp. 25:1–53, doi:10.1145/2287718.2287725.
- [11] M. Bojańczyk (2004): *A Bounding Quantifier*. In: *CSL 2004, LNCS 3210* 3210, Springer, pp. 41–55, doi:10.1007/978-3-540-30124-0_7.
- [12] M. Bojanczyk & S. Torunczyk (2012): *Weak MSO+U over infinite trees*. In Christoph Dürr & Thomas Wilke, editors: *STACS 2012, LIPIcs 14*, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, pp. 648–660, doi:10.4230/LIPIcs.STACS.2012.648.
- [13] P.A. Bonatti, C. Lutz, A. Murano & M.Y. Vardi (2008): *The Complexity of Enriched muCalculi*. *LMCS* 4(3), pp. 1–27, doi:10.2168/LMCS-4(3:11)2008.
- [14] T. Brihaye, A. Da Costa Lopes, F. Laroussinie & N. Markey (2009): *ATL with Strategy Contexts and Bounded Memory*. In: *LFCS 2009*, pp. 92–106, doi:10.1007/978-3-540-92687-0_7.
- [15] D. Calvanese, G. De Giacomo & M. Lenzerini (1999): *Reasoning in Expressive Description Logics with Fixpoints based on Automata on Infinite Trees*. In: *IJCAI 99*, pp. 84–89.
- [16] P. Čermák, A. Lomuscio, F. Mogavero & A. Murano (2014): *MCMAS-SLK: A Model Checker for the Verification of Strategy Logic Specifications*. In: *CAV’14, LNCS 8559*, Springer, pp. 524–531, doi:10.1007/978-3-319-08867-9_34.
- [17] P. Čermák, A. Lomuscio & A. Murano (2015): *Verifying and Synthesising Multi-Agent Systems against One-Goal Strategy Logic Specifications*. In: *AAAI 2015*, pp. 2038–2044.

- [18] K. Chatterjee, T.A. Henzinger & N. Piterman (2010): *Strategy Logic*. IC 208(6), pp. 677–693, doi:10.1016/j.ic.2009.07.004.
- [19] R. Cornes, R. Hartley & T. Sandler (1999): *An Elementary Proof via Contraction*. *Journal of Public Economic Theory* 1(4), pp. 499–509, doi:10.1111/1097-3923.00023.
- [20] J. Bramel D. Simchi-Levi, X. Chen (2013): *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management*. Science and Business Media, Springer.
- [21] S. Dal-Zilio & D. Lugiez (2003): *XML Schema, Tree Logic and Sheaves Automata*. In: *RTA 2003*, pp. 246–263, doi:10.1007/3-540-44881-0_18.
- [22] S. Demri & D. Lugiez (2010): *Complexity of modal logics with Presburger constraints*. *J. Applied Logic* 8(3), pp. 233–252, doi:10.1016/j.jal.2010.03.001.
- [23] E. A. Emerson & C. S. Jutla (1991): *Tree Automata, Mu-Calculus and Determinacy (Extended Abstract)*. In: *ASFCS 1991*, pp. 368–377, doi:10.1109/SFCS.1991.185392.
- [24] A. Ferrante, A. Murano & M. Parente (2008): *Enriched Mu-Calculi Module Checking*. *LMCS* 4(3), pp. 1–21, doi:10.2168/LMCS-4(3:1)2008.
- [25] Erich Grädel (1999): *On The Restraining Power of Guards*. *J. Symb. Log.* 64(4), pp. 1719–1742, doi:10.2307/2586808.
- [26] J. Gutierrez, P. Harrenstein & M. Wooldridge (2014): *Reasoning about Equilibria in Game-Like Concurrent Systems*. In: *KR 2014, AACL*.
- [27] A. Herzig, E. Lorini & D. Walther (2013): *Reasoning about Actions Meets Strategic Logics*. In: *LORI 2013*, pp. 162–175, doi:10.1007/978-3-642-40948-6_13.
- [28] W. van der Hoek, W. Jamroga & M. Wooldridge (2005): *A logic for strategic reasoning*. In: *AAMAS 2005*, pp. 157–164, doi:10.1145/1082473.1082497.
- [29] O. Kupferman, G. Perelli & M. Vardi (2016): *Synthesis with rational environments*. *Annals of Mathematics and Artificial Intelligence*, pp. 1–18, doi:10.1007/s10472-016-9508-8.
- [30] O. Kupferman, G. Perelli & M. Y. Vardi (2014): *Synthesis with Rational Environments*. In: *EUMAS 2014*, pp. 219–235, doi:10.1007/978-3-319-17130-2_15.
- [31] O. Kupferman, U. Sattler & M.Y. Vardi (2002): *The Complexity of the Graded muCalculus*. In: *CADE’02, LNCS 2392*, Springer, pp. 423–437, doi:10.1007/3-540-45620-1_34.
- [32] O. Kupferman, M.Y. Vardi & P. Wolper (2000): *An Automata Theoretic Approach to Branching-Time Model Checking*. *JACM* 47(2), pp. 312–360, doi:10.1145/333979.333987.
- [33] K. Leyton-Brown & Y. Shoham (2008): *Essentials of Game Theory: A Concise, Multidisciplinary Introduction (Synthesis Lectures on Artificial Intelligence and Machine Learning)*. M&C, doi:10.2200/S00108ED1V01Y200802AIM003.
- [34] A. Da Costa Lopes, F. Laroussinie & N. Markey (2010): *ATL with Strategy Contexts: Expressiveness and Model Checking*. In: *FSTTCS 2010*, pp. 120–132, doi:10.4230/LIPIcs.FSTTCS.2010.120.
- [35] V. Malvone, F. Mogavero, A. Murano & L. Sorrentino (2015): *On the Counting of Strategies*. In: *TIME 2015*, pp. 170–179, doi:10.1109/TIME.2015.19.
- [36] V. Malvone, A. Murano & L. Sorrentino: *Games with additional winning strategies*. In: *CILC, 2015*, CEUR.
- [37] F. Mogavero, A. Murano, G. Perelli & M.Y. Vardi (2014): *Reasoning About Strategies: On the Model-Checking Problem*. *TOCL* 15(4), pp. 34:1–42, doi:10.1145/2631917.
- [38] F. Mogavero, A. Murano & M.Y. Vardi (2010): *Reasoning About Strategies*. In: *FSTTCS’10, LIPIcs 8*, Leibniz-Zentrum fuer Informatik, pp. 133–144, doi:10.4230/LIPIcs.FSTTCS.2010.133.
- [39] D. E. Muller & P. E. Schupp (1995): *Simulating Alternating Tree Automata by Nondeterministic Automata: New Results and New Proofs of the Theorems of Rabin, McNaughton and Safra*. *Theor. Comput. Sci.* 141(1&2), pp. 69–107, doi:10.1016/0304-3975(94)00214-4.

- [40] A. Orda, R. Rom & N. Shimkin (1993): *Competitive routing in multiuser communication networks*. *IEEE/ACM Trans. Netw.* 1(5), pp. 510–521, doi:10.1109/90.251910.
- [41] G.P. Papavassilopoulos & J. B. Cruz (1979): *On the Uniqueness of Nash Strategies for a Class of Analytic Differential Games*. *Journal of Optimization Theory and Applications* 27(2), pp. 309–314, doi:10.1007/BF00933234.
- [42] L. Pavel (2012): *Game Theory for Control of Optical Networks*. Science and Business Media, Springer, doi:10.1007/978-0-8176-8322-1.
- [43] A. Pnueli (1977): *The Temporal Logic of Programs*. In: *FOCS'77*, IEEE Computer Society, pp. 46–57, doi:10.1109/SFCS.1977.32.
- [44] S. Rubin (2008): *Automata Presenting Structures: A Survey of the Finite String Case*. *Bulletin of Symbolic Logic* 14(2), pp. 169–209, doi:10.2178/bsl/1208442827.
- [45] H. Seidl, T. Schwentick & A. Muscholl (2008): *Counting in trees*. In: *Logic and Automata: History and Perspectives*, pp. 575–612.
- [46] R. Selten (1965): *Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragertragheit*. *Zeitschrift für die gesamte Staatswissenschaft* 121, pp. 301–324.
- [47] W. Thomas (1990): *Automata on infinite objects*. *Handbook of theoretical computer science, Volume B*, pp. 133–191.
- [48] M. Ummels (2006): *Rational Behaviour and Strategy Construction in Infinite Multiplayer Games*. In: *FSTTCS*, pp. 212–223, doi:10.1007/11944836_21.
- [49] D. Walther, W. van der Hoek & M. Wooldridge (2007): *Alternating-time temporal logic with explicit strategies*. In: *TARK 2007*, pp. 269–278, doi:10.1145/1324249.1324285.
- [50] F. Wang, C. Huang & F. Yu (2011): *A Temporal Logic for the Interaction of Strategies*. In: *CONCUR 2011, LNCS 6901*, Springer, pp. 466–481, doi:10.1007/978-3-642-23217-6_31.
- [51] Y. Zhang & M. Guizani (2011): *Game Theory for Wireless Communications and Networking*. CRC Press.