


Logics for Multi-Agent Systems Verification



Giuseppe Perelli

Università degli studi di Napoli “Federico II”
Dipartimento di Matematica e Applicazioni “R. Caccioppoli”

Dottorato in **Scienze Computazionali e Informatiche**
Ciclo XXVII

A thesis submitted in fulfillment of the degree of
Doctor in Compute Science

Submission: March 31, 2015
Defense: Napoli, to be added
Revised version: March 25, 2015

© Copyright 2015
by
Giuseppe Perelli

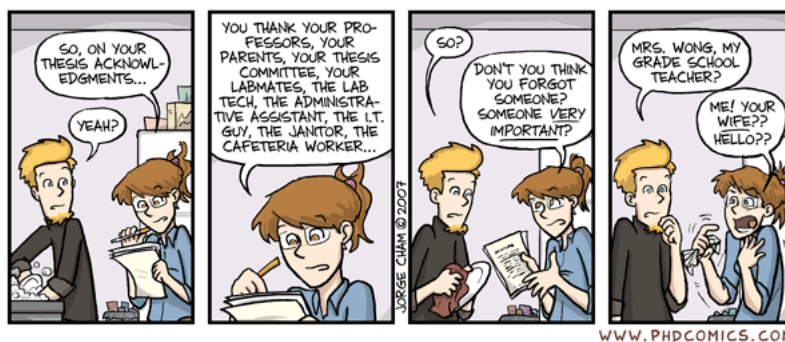
Supervisor: Prof. Dr. Aniello Murano

This paper intentionally left blank

“A problem is a chance for you to do your best.”
Duke Ellington

Acknowledgments

The approaching end of a period easily evokes a final summary. In the case of a period ending with a thesis, such summary becomes mandatory. Otherwise, you risk the readers to complain for missing this fundamental chapter. It is indeed not a coincidence that the acknowledgements chapter is the most read and carefully revised by many many reviewers, the most severe you can ever meet. If missing a citation in the references of a paper sometime makes the difference between acceptance and rejection, missing a mention in the acknowledgements of a thesis always makes the difference between friendship and enmity.



First of all, I wish to say thanks to my **family**, for their constant presence, patience, and for all the moral support they never denied to me in these years. For having let me go alone unconditionally with my attitudes and ambitions, giving me the possibility to undertake an excellent path of knowledge and reach the goal of a PhD.

A big thank to my supervisor, Prof. **Aniello Murano**, for having perfectly interpreted the role of *academic parent*. Never ending source of good advices, he has been so good in wisely dosing the rations of encouragements, congrats, and scoldings, which has built a deep relation of trust and friendship along these years. With great emotion, I say thanks also to my *academic grandparent* (as he likes to define himself!), Prof. **Moshe Vardi**. The period working at Rice University under his guidance has been of big personal and professional growth. In every meeting with him, I am always amazed by his wise visions of research and life in general. The way he takes care of all his students, with sympathy and affection, is priceless! To complete the picture of my academic family, I would say thanks to my elder brother **Fabio Mogavero**. Every scientific conversation with him is always inspired and deep. The born understanding between us when we do research is extraordinary. Still today I am stunned by our ability to reason in harmony on all the problems we meet working together. Moreover, I would say thanks to all the professors and researchers I have had the pleasure to work with during these years: **Orna Kupferman**, **Angelo Montanari**, **Adriano Peron**, and **Luigi Sauro**.

For what the friends regards, I cannot avoid to mention all the ones that have been actively part of my PhD life, in particular in Houston. My officemate **Kuldeep**, for the philosophical conversations, the black humor jokes, the biscuits, and lunches full of sriracha. As well as for having taught me some useful hindi word ... Friends from the research group **Dror**, **Lucas**, **Miguel**, **Morteza**, and **Sumit**, for the fruitful collaboration and for all the suggestions. All the

Duncan Hall friends **Alina, Dragos, Ryan, Sailesh, Sarah, and Ted**, which whom I shared my spare time inside and outside campus.

A special mention to my Italian friends **Roberto, Stefania** and **Andrea**, for having let me find a bit of Naples and Italy thousands of kilometers far away from home. Our casual meeting in a supermarket on the other side of world has made us feel immediately tuned together. I still remember my crazy walk of about 7 miles to reach them the first time we hanged out. It is like I already knew they deserve my efforts! I will never forget all the dinner invitations, the bike (and walk!) adventures, and all the problems they have helped me to fix. A special remark for the two weeks in which **Roberta** joined us. After coming back to Naples, we started sharing a coffee from time to time, and it is always a lot of fun. Last but not least thanks to **Rishi**, for the great sense of friendship and generosity they make him unique, and for having taught me the best hindi words to have success with Indian girls!

As some reader would have noticed, it seems I forgot to mention someone: the most severe reviewers of this chapter. This is clearly not true. In old Latin we say *dulcis in fundo* (sweetness in the end), and I would never forget the “women of my life”⁰: **Loredana** and **Sonali**. I met Loredana, three years ago. We started together the PhD adventure and our “colleagueship” immediately turned to be friendship. Our complicity has been very fundamental to effort both the office time and the several schools and conferences we participated together. Having such a special company, especially in the hard moments, would be a big fortune for everyone. Sonali has filled my days in Houston with her sweetness and happiness. Her genuine affection towards me makes me in a good mood just with a simple smile. Due to the long distance, we made an agreement: to meet every at most 3 years. I will do the best to respect this agreement, in order to be present in all the most important moments of our lives.

⁰I can still afford to have more than one!

Contents

Introduction	iii
1 Strategy Logic	1
1.1 Syntax and Semantics	1
1.2 Model-Checking Hardness	15
1.3 Satisfiability	19
2 Strategy Logic Fragments	26
2.1 Syntax	26
2.2 Behavioral property	30
2.3 Model Checking	41
2.4 Satisfiability	51
3 Synthesis	61
3.1 Rational synthesis	61
3.2 Qualitative Rational Synthesis	65
3.3 Quantitative Rational Synthesis	66
4 Pushdown ATL*	69
4.1 Pushdown Games	69
4.2 ATL* over pushdown systems	71
4.3 Model Checking	72

Introduction

In system design, *model checking* is a well-established formal method that allows to automatically check for global system correctness [CE81, QS81, CGP02]. In such a framework, in order to check whether a system satisfies a required property, we describe its structure in a mathematical model (such as *Kripke structures* [Kri63] or *labeled transition systems* [Kel76]), specify the property with a formula of a temporal logic (such as LTL [Pnu77], CTL [CE81], or CTL* [EH86]), and check formally that the model satisfies the formula. In the last decade, interest has arisen in analyzing the behavior of individual components or sets of them in systems with several entities. This interest has started in reactive systems, which are systems that interact continuously with their environments. In *module checking* [KVV01], the system is modeled as a module that interacts with its environment and correctness means that a desired property holds with respect to all such interactions (see also [FMP08] and [ALM⁺13, JM14] for recent works in this field).

Starting from the study of module checking, researchers have looked for logics focusing on the strategic behavior of agents in multi-agent systems [ÅGJ07, AHK02, Pau02, JvdH04, WvdHW07, BJ14]. One of the most important developments in this field is *Alternating-Time Temporal Logic* (ATL*, for short), introduced by Alur, Henzinger, and Kupferman [AHK02]. ATL* allows reasoning about strategies of agents with temporal goals. Formally, it is obtained as a generalization of CTL* in which the path quantifiers, *there exists* “E” and *for all* “A”, are replaced with *strategic modalities* of the form “ $\langle\langle A \rangle\rangle$ ” and “ $[[A]]$ ”, where A is a set of *agents* (a.k.a. *players*). Strategic modalities over agent sets are used to express 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 games between a coalition and its complement.

ATL* formulas are interpreted over *concurrent game structures* (CGS, for short) [AHK02], which model interacting processes. Given a CGS \mathcal{G} and a set A of agents, the ATL* formula $\langle\langle A \rangle\rangle\psi$ is satisfied at a state s of \mathcal{G} if there is a set of strategies for agents in A such that, no matter which strategies are executed by agents not in A, the resulting outcome of the interaction in \mathcal{G} satisfies ψ at s . Thus, ATL* can express properties related to the interaction among components, while CTL* can only express property of the global system. As an example, consider the property “processes α and β cooperate to ensure that a system (having more than two processes) never enters a failure state”. This can be expressed by the ATL* formula $\langle\langle \{\alpha, \beta\} \rangle\rangle G\neg fail$, where G is the classical LTL temporal operator “*globally*”. CTL*, in contrast, cannot express this property [AHK02]. Indeed, it can only assert whether the set of all agents may or may not prevent the system from entering a failure state. It turns out that both the model-checking and the satisfiability problems are elementarily decidable and, precisely, 2EXPTIME-COMplete [AHK02, Sch08].

Despite its powerful expressiveness, ATL* suffers from a strong limitation, due to the fact that strategies are only treated implicitly, through modalities that refer to games between competing coalitions. To overcome this problem, Chatterjee, Henzinger, and Piterman introduced *Strategy Logic* (CHP-SL, for short) [CHP07], a logic that treats strategies in *two-player*

turn-based games as explicit *first-order objects*. In CHP-SL, the ATL* formula $\langle\langle\{\alpha\}\rangle\rangle\psi$, for a system modeled by a CGS with agents α and β , becomes $\exists x.\forall y.\psi(x, y)$, i.e., “there exists a player- α strategy x such that for all player- β strategies y , the unique infinite path resulting from the two players following the strategies x and y satisfies the property ψ ”. The explicit treatment of strategies in this logic allows to state many properties not expressible in ATL*. In particular, it is shown in [CHP07] that ATL*, in the restricted case of two-agent turn-based games, corresponds to a proper one-alternation fragment of CHP-SL. The authors of that work have also shown that the model-checking problem for CHP-SL is decidable, although only a non-elementary algorithm for it, both in the size of system and formula, has been provided, leaving as an open question whether an algorithm with a better complexity exists or not. The complementary question about the decidability of the satisfiability problem for CHP-SL was also left open and it is not addressed in other papers apart the preliminary work [MMV10a].

While the basic idea exploited in [CHP07] to quantify over strategies and then to commit agents explicitly to certain of these strategies turns to be very powerful and useful [FKL10], CHP-SL still presents severe limitations. Among the others, it needs to be extended to the more general concurrent multi-agent setting. Also, the specific syntax considered there allows only a weak kind of strategy commitment. For example, CHP-SL does not allow different players to share the same strategy, suggesting that strategies have yet to become first-class objects in this logic. Moreover, an agent cannot change his strategy during a play without forcing the other to do the same.

These considerations, as well as all questions left open about decision problems, led us to introduce and investigate a new *Strategy Logic*, denoted SL, as a more general framework than CHP-SL, for explicit reasoning about strategies in *multi-agent concurrent games*. Syntactically, SL extends LTL by means of two *strategy quantifiers*, the existential $\langle\langle x \rangle\rangle$ and the universal $\llbracket x \rrbracket$, as well as *agent binding* (a, x) , where a is an agent and x a variable. Intuitively, these elements can be respectively read as “*there exists a strategy x* ”, “*for all strategies x* ”, and “*bind agent a to the strategy associated with x* ”. For example, in a CGS with three agents α, β, γ , the previous ATL* formula $\langle\langle\{\alpha, \beta\}\rangle\rangle G\neg fail$ can be translated in the SL formula $\langle\langle x \rangle\rangle\langle\langle y \rangle\rangle\llbracket z \rrbracket(\alpha, x)(\beta, y)(\gamma, z)(G\neg fail)$. The variables x and y are used to select two strategies for the agents α and β , respectively, while z is used to select one for the agent γ such that their composition, after the binding, results in a play where *fail* is never met. Note that we can also require, by means of an appropriate choice of agent bindings, that agents α and β share the same strategy, using the formula $\langle\langle x \rangle\rangle\llbracket z \rrbracket(\alpha, x)(\beta, x)(\gamma, z)(G\neg fail)$. Furthermore, we may vary the structure of the game by changing the way the quantifiers alternate, as in the formula $\langle\langle x \rangle\rangle\llbracket z \rrbracket\langle\langle y \rangle\rangle(\alpha, x)(\beta, y)(\alpha, z)(G\neg fail)$. In this case, x remains uniform w.r.t. z , but y becomes dependent on it. Finally, we can change the strategy that one agent uses during the play without changing those of the other agents, by simply using nested bindings, as in the formula $\langle\langle x \rangle\rangle\langle\langle y \rangle\rangle\llbracket z \rrbracket\langle\langle w \rangle\rangle(\alpha, x)(\beta, y)(\gamma, z)(G(\gamma, w)G\neg fail)$. The last examples intuitively show that SL is an extension of both ATL* and CHP-SL. It is worth noting that the pattern of modal quantifications over strategies and binding to agents can be extended to other linear-time temporal logics than LTL, such as the linear μ CALCULUS [Var88]. In fact, the use of LTL here is only a matter of simplicity in presenting our framework, and changing the embedded temporal logic only involves few side-changes in proofs and procedures.

Introduction

As one of the main results in this thesis, we show that the model-checking problem is non-elementarily decidable. To gain this, we use an *automata-theoretic approach* [KVV00]. Precisely, we reduce the decision problem for our logic to the emptiness problem of a suitable *alternating parity tree automaton*, which is an *alternating tree automaton* (see [GTW02], for a survey) along with a *parity acceptance condition* [MS95]. Due to the operations of projection required by the elimination of quantifications on strategies, which induce at any step an exponential blow-up, the overall size of the required automaton is non-elementary in the size of the formula, while it is only polynomial in the size of the model. Thus, together with the complexity of the automata-nonemptiness calculation, we obtain that the model checking problem is in PTIME, w.r.t. the size of the model, and NONELEMENTARY, w.r.t. the size of the specification. Hence, the algorithm we propose is computationally not harder than the best one known for CHP-SL and, notably, a non-elementary data complexity improvement, (i.e., with respect to the size of the model). The latter is remarkable as in formal verification often the specification is very small (or even constant) with respect to the system. Moreover, we prove that our problem has a non-elementary lower bound. Specifically, it is k -EXPSpace-HARD in the alternation number k of quantifications in the specification. In addition to the model-checking problem for SL, we address the satisfiability problem, showing its undecidability [MMV10a]. The proof is based on the reduction from the *recurrent domino problem* proposed for the first time by Wang [Wan61]. A key role for this result is played by the unbounded model property of SL, which is proved in [MMPV12] and reported in this Thesis.

The contrast between the high complexity of the model-checking satisfiability problems for our logic and the elementary ones for ATL^* has spurred us to investigate syntactic fragments of SL, strictly subsuming ATL^* , with a better complexity. In particular, by means of these sublogics, we want to understand why SL is computationally more difficult than ATL^* . To this aim, we introduce *Nested-Goal*, *Boolean-Goal*, and *One-Goal Strategy Logic*, respectively denoted by $SL_{[NG]}$, $SL_{[BG]}$, and $SL_{[1G]}$. These fragments encompass formulas in a special prenex normal form having nested temporal goals, Boolean combinations of goals, and a single goal at a time, respectively. For goal we mean an SL formula of the type $\mathfrak{b}\psi$, where \mathfrak{b} is a binding prefix of the form $(\alpha_1, x_1), \dots, (\alpha_n, x_n)$ containing all the involved agents and ψ is an agent-full formula. With more detail, the idea behind $SL_{[NG]}$ is that, when in ψ there is a quantification over a variable, then there are quantifications of all free variables contained in the inner subformulas. So, a subgoal of ψ that has a variable quantified in ψ itself cannot use other variables quantified out of this formula. Thus, goals can be only nested or combined with Boolean and temporal operators. $SL_{[BG]}$ and $SL_{[1G]}$ further restrict the use of goals. In particular, in $SL_{[1G]}$, each temporal formula ψ is prefixed by a quantification-binding prefix $\wp\mathfrak{b}$ that quantifies over a tuple of strategies and binds them to all agents.

As main results about these fragments, we prove that the model-checking problem for $SL_{[1G]}$ is 2EXPTIME-COMPLETE, thus not harder than the one for ATL^* . On the contrary, for $SL_{[NG]}$, it is NONELEMENTARY-HARD and thus we enforce the corresponding result for SL. Finally, by observing that $SL_{[NG]}$ includes $SL_{[BG]}$, we have that the model-checking problem for the latter is decidable with a 2EXPTIME lower-bound, given by $SL_{[1G]}$ model-checking complexity. The problem of determine its precise complexity is left open here.

Introduction

For what the satisfiability regards, we prove that the problem remains undecidable for $SL_{[BG]}$, since the reduction provided for the entire SL already holds in this fragment. On the contrary, the problem is $2EXPTIME$ -COMPLETE for the smallest fragment $SL_{[1G]}$, thus not harder than ATL^* .

To achieve all positive results about $SL_{[1G]}$, we introduce a fundamental property of the semantics of this logic, called *behavioral*¹, which allows us to strongly simplify the reasoning about strategies by reducing it to a set of reasonings about agent's choices (formally, agent's *actions*). This intrinsic characteristic of $SL_{[1G]}$, which unfortunately is not shared by the other fragments, asserts that, in a determined history of the play, the value of an existential quantified strategy depends only on the values of strategies, from which the first depends, on the same history. This means that, to choose an existential strategy, we do not need to know the entire structure of universal strategies, as for the whole SL, but only their values on the histories of interest. In other words, by means of the behavioral property, a one-shot second-order quantification over strategies can be eliminated in favor of a progressive first-order quantification over agent's choices. Technically, to describe the behavioral property, we make use of the machinery of *Skolem dependence function*, which defines a Skolemization procedure for SL, inspired by the one in first order logic (see [Hod01]).

By means of behavioral, we can modify the SL model-checking procedure via alternating tree automata in such a way that we avoid the projection operations by using a dedicated automaton that makes an action quantification for each node of the tree model. Consequently, the resulting automaton is only exponential in the size of the formula, independently from its alternation number. Thus, together with the complexity of the automata-nonemptiness calculation, we get that the model-checking procedure for $SL_{[1G]}$ is $2EXPTIME$. As it is explained in Chapter 2, the behavioral property is fundamental to prove also the bounded-tree model property. Clearly, the behavioral property also holds for ATL^* , as it is included in $SL_{[1G]}$. In particular, although it has not been explicitly stated, this property is crucial for most of the results achieved in literature about ATL^* by means of automata (see [Sch08], as an example). Moreover, we believe that our proof techniques are of independent interest and applicable to other logics as well (see for example [MP14]).

The formal representation in SL can be exploited to address several problems. For instance, we can investigate the problem of *Rational Synthesis*. *Synthesis* is the automated construction of a system from its specification. The basic idea is simple and appealing: instead of developing a system and verifying that it adheres to its specification, we would like to have an automated procedure that, given a specification, constructs a system that is correct by construction. The first formulation of synthesis goes back to Church [Chu63]; the modern approach to synthesis was initiated by Pnueli and Rosner, who introduced LTL (linear temporal logic) synthesis [PR89]. The LTL *synthesis problem* receives as input a specification in LTL and outputs a reactive system modeled by a finite-state transducer satisfying the given specification — if such exists. It is important to distinguish between input signals, assigned by the environment, and output signals, assigned by the system. A system should be able to cope with all values of the input signals, while setting the output signals to desired values [PR89].

¹We use this term as it has a direct correspondence with the “behavioral” concept used in game theory [Mye97, MMS13, MMS14].

Introduction

Therefore, the quantification structure on input and output signals is different. Input signals are universally quantified while output signals are existentially quantified.

Modern systems often interact with other systems. For example, the clients interacting with a server are by themselves distinct entities (which we call agents) and are many times implemented by systems. In the traditional approach to synthesis, the way in which the environment is composed of its underlying agents is abstracted. In particular, the agents can be seen as if their only objective is to conspire to fail the system. Hence the term “hostile environment” that is traditionally used in the context of synthesis. In real life, however, many times agents have goals of their own, other than to fail the system. The approach taken in the field of algorithmic game theory [NRTV07] is to assume that agents interacting with a computational system are *rational*, i.e., agents act to achieve their own goals. Assuming agents rationality is a restriction on the agents behavior and is therefore equivalent to softening the universal quantification on the environment.² Thus, the following question arises: can system synthesizers capitalize on the rationality and goals of agents interacting with the system?

In [FKL10], Fisman et al. positively answered this question by introducing and studying *rational synthesis*. The input to the rational-synthesis problem consists of LTL formulas specifying the objectives of the system and the agents that constitute the environment, and a solution concept, e.g., dominant strategies, Nash Equilibria, and the like. The atomic propositions over which the objectives are defined are partitioned among the system and the agents, so that each of them controls a subset of the propositions. The desired output is a strategy profile such that the objective of the system is satisfied in the computation that is the outcome of the profile, and the agents that constitute the environment have no incentive to deviate from the strategies suggested to them (formally, the profile is an equilibrium with respect to the solution concept). Fisman et al. showed that there are specifications that cannot be realized in a hostile environment but are realizable in a rational environment. Moreover, the rational-synthesis problem for LTL and common solution concepts used in game theory can be solved in 2EXPTIME, thus its complexity coincides with that of usual synthesis.

In this thesis, we present the following three contributions on Rational Synthesis, as they have been introduced in [KPV14]. First, we suggest an alternative definition to rational synthesis, in which the agents are rational but not cooperative. Second, we study a richer specification formalism, where the objectives of the system and the agents are not Boolean but quantitative. Third, we show that all these variants of the rational synthesis problems can be reduced to the model checking in fragments of *Strategy Logic* [MMPV14]. Before we describe our contributions in more detail, let us highlight a different way to consider rational synthesis and our contribution here. *Mechanism design* is a field in game theory and economics studying the design of games whose outcome (assuming agents rationality) achieves some goal [NR01, NRTV07]. The outcome of traditional games depends on the final position of the game. In contrast, the systems we reason about maintain an *on-going interaction* with their environment, and we reason about their behavior by referring not to their final state (in fact, we consider non-terminating systems, with no final state) but rather

²Early work on synthesis has realized that the universal quantification on the behaviors of the environment is often too restrictive. The way to address this point, however, has been by adding assumptions on the environment, which can be part of the specification (cf., [CHJ08]).

Introduction

to the *language* of computations that they generate. Rational synthesis can be viewed as a variant of mechanism design in which the game is induced by the objective of the system, and the objectives of both the system and the agents refer to their on-going interaction and are specified by temporal-logic formulas. Our contributions here correspond to the classic setting assumed in mechanism design: the agents need not be cooperative, and the outcome is not Boolean.

We argue that the definition of rational synthesis in [FKL10] is *cooperative*, in the sense that the agents that constitute the environment are assumed to follow the strategy profile suggested to them (as long as it is in an equilibrium). Here, we consider also a *non-cooperative* setting, in which the agents that constitute the environment may follow any strategy profile that is in an equilibrium, and not necessarily the one suggested to them by the synthesis algorithm. In many scenarios, the cooperative setting is indeed too optimistic, as the system cannot assume that the environment, even if it is rational, would follow a suggested strategy, rather than a strategy that is as good for it. Moreover, sometimes there is no way to communicate with the environment and suggest a strategy for it. From a technical point of view, we show that the non-cooperative setting requires reasoning about *all* possible equilibria, yet, despite this more sophisticated reasoning, it stays 2EXPTIME-COMplete. We achieve the upper bound by reducing rational synthesis to the model-checking problem for SL. While the model-checking problem for strategy logic is in general non-elementary, we show that it is possible to express rational synthesis in $SL[NG, 1-ALT]$ ³, which leads to the desired complexity. It is important to observe the following difference between the cooperative and the non-cooperative settings. In the cooperative one, we synthesize strategies for all agents, with the assumption that the agent that corresponds to the system always follows his suggested strategy and the agents that constitute the environment decide in a rational manner whether to follow their strategies. On the other hand, in the non-cooperative setting, we synthesize a strategy only for the agent that corresponds to the system, and we assume that the agents that constitute the environment are rational, thus the suggested strategy has to win against all rational behaviors of the environment.

Our second contribution addresses a weakness of the classical synthesis problem, a weakness that is more apparent in the rational setting. In classical synthesis, the specification is Boolean and describes the expected behavior of the system. In many applications, systems can satisfy their specifications at different levels of quality. Thus, synthesis with respect to Boolean specifications does not address designers needs. This latter problem is a real obstacle, as designers would be willing to give up manual design only after being convinced that the automatic procedure that replaces it generates systems of comparable quality. In the last years we see a lot of effort on developing formalisms that would enable the specification of such quality measures [BCHJ09, ABK13].

Classical applications of game theory consider games with quantitative payoffs. In the Boolean setting, we assumed that the payoff of an agent is, say, 1, if its objective is satisfied, and is 0 otherwise. In particular, this means that agents whose objectives are not satisfied have no incentive to follow any strategy, even if the profile satisfies the solution concept. In real-life, rational objectives are rarely Boolean. Thus, even beyond our goal of synthesizing

³By $SL[k-ALT]$ we denote the fragment of SL in which only a k number of alternation is allowed.

Introduction

systems of high quality, the extension of the synthesis problem to the rational setting calls also for an extension to a quantitative setting. Unfortunately, the full quantitative setting is undecidable already in the context of model checking [Hen10]. In [FKL10], Fisman et al. extended cooperative rational synthesis to objectives in the multi-valued logic LLTL, where specifications take truth values from a finite lattice.

We introduce here a new quantitative specification formalism, termed *Objective LTL*, (OLTL, for short). We first define the logic, and then study its rational synthesis. Essentially, an OLTL specification is a pair $\theta = \langle \Psi, f \rangle$, where $\Psi = \langle \psi_1, \psi_2, \dots, \psi_m \rangle$ is a tuple of LTL formulas and $f : \{0, 1\}^m \rightarrow \mathbb{Z}$ is a *reward function*, mapping Boolean vectors of length m to an integer. A computation π then maps θ to a reward in the expected way, according to the subset of formulas that are satisfied in π . In the rational synthesis problem for OLTL, the input consists of OLTL specifications for the system and the other agents, and the goal of the system is to maximize its reward with respect to environments that are in an equilibrium. Again, we distinguish between a cooperative and a non-cooperative setting. Note that the notion of an equilibrium in the quantitative setting is much more interesting, as it means that all agents in the environment cannot expect to increase their payoffs. We show that the quantitative setting is not more complex than the non-quantitative one, thus quantitative rational synthesis is complete for 2EXPTIME in both the cooperative and non-cooperative settings.

We conclude this thesis by presenting a result that aims to combine strategic specifications with pushdown systems. Starting from the works on module checking, two significant directions have been taken in open-system verification. One concerns extending the framework to more sophisticated systems while maintaining the dichotomy system-environment states in modeling. In this context, worthy of mention is the work on *pushdown module checking* [BMP10]. This has the merit of having handled the verification of infinite-state open systems and, thanks to the fact that the infinite number of states is induced by a recursive structure of finite size, the problem turns out to be decidable and precisely 3EXPTIME-COMPLETE for specifications in CTL*. Another direction has instead completely redesigned the module checking approach in order to handle the more involved scenario of multi-agent (concurrent) systems. This is, for instance, the above mentioned case of ATL*. However, the two approaches are incomparable as in module checking it is possible to use nondeterministic strategies.

Despite the undoubted utility of considering, from one hand, infinite-state open-system models induced by finite-size recursive structures and, from the other hand, multi-agent specifications, to the best of our knowledge no work has been devoted to the combination of the two.

In this thesis, we consider multi-agent pushdown systems and address the related model checking problem for specifications expressed in ATL*. To this aim, we first introduce *pushdown game structures* to properly model the infinite-state multi-agent system and formalize the model checking question. Then, by means of an automata-theoretic approach, we provide a 3EXPTIME solution to the addressed problem. Precisely, we construct a doubly-exponential size *pushdown parity tree automaton* that collects all execution trees satisfying the ATL* formula. Then by using the fact that the emptiness of this automaton can be checked in exponential time [KPV02], we get the desired result. We also provide a 2EXPSpace lower-bound

by showing a reduction from the model checking problem for one-counter systems against quantitative ATL* specifications [Ves14].

Related works. We now discuss some work related to the results included in the thesis.

Logics for strategic reasoning representation Several works have focused on extensions of ATL and ATL* to incorporate more powerful strategic constructs. Among them, we recall *Alternating-Time μ CALCULUS* ($A\mu$ CALCULUS, for short) [AHK02], *Game Logic* (GL, for short) [AHK02], *Quantified Decision Modality μ CALCULUS* ($QD\mu$ CALCULUS, for short) [Pin07], *Coordination Logic* (CL, for short) [FS10], (*ATL with plausibility* (ATL^+ , for short) [BJD08], (*ATL with Irrevocable strategies* (IATL, for short) [ÅGJ07], (*Memo-ryful ATL** ($mATL^*$, for short) [MMV10b], *Basic Strategy-Interaction Logic* (BSIL, for short) [WHY11] *Temporal Cooperation Logic* (TCL, for short) [HSW13], *Alternating-time Temporal Logic with Explicit Actions* (ATLEA, for short) [HLW13] and some extensions of ATL* considered in [BLLM09]. $A\mu$ CALCULUS and $QD\mu$ CALCULUS are intrinsically different from SL (as well as from CHP-SL and ATL*) as they are obtained by extending the propositional μ -calculus [Koz83] with strategic modalities. CL is similar to $QD\mu$ CALCULUS but with LTL temporal operators instead of explicit fixpoint constructors. GL is strictly included in CHP-SL, in the case of two-player turn-based games, but it does not use any explicit treatment of strategies, neither it does the extensions of ATL* introduced in [BLLM09], which consider restrictions on the memory for strategy quantifiers. ATL^+ enables to express rationality assumptions of intelligent agents in ATL. In IATL, the semantics of the logic ATL is changed in a way that, in the evaluation of the goal, agents can be forced to keep the strategy they have chosen in the past in order to reach the state where a goal is evaluated. $mATL^*$ enriches ATL* by giving the ability to agents to “relent” and change their goals and strategies depending on the history of the play. BSIL allows to specify behaviors of a system that can cooperate with several strategies of the environment for different requirements. TCL extends ATL by allowing successive definitions of agent strategies, with the aim of using the collaborative power of groups of agents to enforce different temporal objectives. ATLEA introduces explicit actions in the logic ATL to reason about abilities of agents under commitments to play precise actions. Thus, all above logics are different from SL, which we recall it aims to be a minimal but powerful logic to reason about strategic behavior in multi-agent systems.

At roughly the same time we have conceived Strategy Logic, another generalization of ATL*, named *ATL* with Strategy Contexts*, which results to be very expressive but a proper sublogic of SL, has been considered in [DLM10] (see also [DLM12, LM13] for more recent works). In this logic, a quantification over strategies does not reset the strategies previously quantified but allows to maintain them in a particular context in order to be reused. This makes the logic much more expressive than ATL*. On the other hand, as it does not allow agents to share the same strategy, it is not comparable with the fragments we have considered in this paper.

Recently, several extensions of SL have been also investigated. *Updating Strategy Logic* (U-SL, for short) has been considered in [CBC13] where, in addition to SL, an agent can

Introduction

refine its own strategies by means of an "unbinder" operator, which explicitly deletes the binding of a strategy to an agent. In [Bel14], an epistemic extension of SL with modal operators for individual knowledge has been considered, showing that the complexity of model checking for this logic is not worse than the one for (non-epistemic) SL. Last but not least, in [ČLMM14] a BDD-based model checker for the verification of systems against specifications expressed in the epistemic extension of SL, named MCMAS-SLK, has been introduced.

Synthesis The understanding that synthesis corresponds to a game in which the objective of each player is to satisfy his specification calls for a mutual adoption of ideas between formal methods and game theory. In *rational synthesis*, introduced in [FKL10], synthesis is defined in a way that takes into an account the rationality of the agents that constitute the environment and involves an assumption that an agent cooperates with a strategy in which his objective is satisfied. Here, we extend the idea and consider also *non-cooperative* rational synthesis, in which agents need not cooperate with suggested strategies and may prefer different strategies that are at least as beneficial for them.

Many variants of the classical synthesis problem has been studied. It is interesting to examine the combination of the rational setting with the different variants. To start, the cooperative and non-cooperative settings can be combined into a framework in which one team of agents is competing with another team of agents, where each team is internally cooperative, but the two teams are non-cooperative. Furthermore, we plan to study rational synthesis with *incomplete information* [KV99], where agents can view only a subset of the signals that other agents output, and rational *stochastic* synthesis [CMJ04], which models the unpredictability of nature and involves stochastic agents that assign values to their output signals by means of a distribution function. Beyond a formulation of the richer settings, one needs a corresponding extension of strategy logic and its decision problems.

As discussed above, classical applications of game theory consider games with quantitative payoffs. We added a quantitative layer to LTL by introducing Objective-LTL and studying its rational synthesis. In recent years, researchers have developed more refined quantitative temporal logics, which enable a formal reasoning about the quality of systems. In particular, we plan to study rational synthesis for the multi-valued logics LTL[F] [ABK13], which enables a prioritization of different satisfaction possibilities, and LTL[D] [ABK14], in which discounting is used in order to reduce the satisfaction value of specifications whose eventualities are delayed. The rational synthesis problem for these logics induce a game with much richer, possibly infinitely many, profiles, making the search for a stable solution much more challenging.

Pushdown systems In recent years, model checking of pushdown systems has received a lot of attention, largely due to the ability of these systems to capture the flow of procedure of calls and returns in programs [ABE⁺05]. The work in this area started with Muller and Schupp, who showed that the monadic second-order theory of graphs induced by pushdown systems is decidable [MS85]. Walukiewicz in [Wal96] showed that the model checking for pushdown systems with respect to *modal μ -calculus* is EXPTIME-COMplete. The problem remains

EXPTIME-COMPLETE also for CTL and LTL, while it becomes 2EXPTIME-COMPLETE for CTL* [Wal00, BEM97]. In [BMP10], open pushdown systems along with the module checking paradigm have been considered. This setting has been investigated under several restrictions including the imperfect-information case [ALM⁺13].

Strategy Logic

This preliminary chapter is organized as follows. In Section 1.1 we introduce the syntax and the semantics of *Strategy Logic*, together with the auxiliary definition of the underlying framework, namely *concurrent game structure*. Then, we discuss its expressiveness by showing examples, most of them borrowed by some classic problem in Game Theory. In Section 1.2 we discuss the model-checking problem for this logic, by showing a reduction from the satisfiability problem for QPTL, which is known to be NONELEMENTARY-COMPLETE. Finally, in Section 1.3 we show that the satisfiability problem for SL is undecidable, by means of a reduction from the *recurrent domino problem*, introduced in [Wan61].

1.1 Syntax and Semantics

Strategy Logic [MMV10a] (SL, for short) is an extension of the classic linear-time temporal logic LTL [Pnu77] along with the concepts of strategy quantifications and agent binding, which formalism allows to express strategic plans over temporal goals. The main distinctive feature of this formalism *w.r.t.* other logics with the same aim resides in the decoupling strategy instantiations, done through the quantifications, from their applications, by means of bindings. Consequently, the logic is not simply propositional but predicative, since we treat strategies as a first order concept via the use of agents and variables as explicit syntactic elements. This fact let us to write Boolean combinations and nesting of complex predicates, representing each one a different temporal goal, linked together by some common strategic choices.

1.1.1 Underlying framework

As semantic framework for SL, we use the *graph-based model* for *multi-player games* named *concurrent game structure* [AHK02], which is a generalization of *Kripke structures* [Kri63] and *labeled transition systems* [Kel76]. It allows to model *multi-agent systems* viewed as extensive form games, in which players perform *concurrent actions* to trigger different transitions over the graph.

Definition 1.1.1 (Concurrent Game Structures) A concurrent game structure (CGS, for short) is a tuple $\mathcal{G} \triangleq \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_0 \rangle$, where AP and Ag are finite non-empty sets of atomic propositions and agents, Ac and St are enumerable non-empty sets of actions and states, $s_0 \in \text{St}$ is a designated initial state, and $\text{ap} : \text{St} \rightarrow 2^{\text{AP}}$ is a labeling function that maps each state to the set of atomic propositions true in that state. Let $\text{Dc} \triangleq \text{Ac}^{\text{Ag}}$ be the set of decisions, i.e., functions from Ag to Ac representing the choices of an action for each

1.1. Syntax and Semantics

agent.¹ Then, $\text{tr} : \text{St} \times \text{Dc} \rightarrow \text{St}$ is a transition function mapping a pair of a state and a decision to a state.

To get familiar with the concept of CGS, we present here some running example of simple concurrent games.

We start by modeling the *paper, rock, and scissor* game.

Example 1.1.1 (Paper, Rock, and Scissor) Consider the classic two-player concurrent game paper, rock, and scissor (PRS, for short) as represented in Figure 1.1, where a play continues until one of the participants catches the move of the other. Vertexes are states of the game and labels on edges represent decisions of agents or sets of them, where the symbol $*$ is used in place of every possible action. In this specific case, since there are only two agents, the pair of symbols $**$ indicates the whole set Dc of decisions. The agents “Alice” and “Bob” in $\text{Ag} \triangleq \{A, B\}$ have as possible actions those in the set $\text{Ac} \triangleq \{P, R, S\}$, which stand for “paper”, “rock”, and “scissor”, respectively. During the play, the game can stay in one of the three states in $\text{St} \triangleq \{s_i, s_A, s_B\}$, which represent, respectively, the waiting moment, named idle, and the two winner positions. The latter ones are labeled accordingly with one of the atomic propositions in $\text{AP} \triangleq \{w_A, w_B\}$, in order to represent who is the winner between A and B. The catch of one action over another is described by the relation $C \triangleq \{(P, R), (R, S), (S, P)\} \subseteq \text{Ac} \times \text{Ac}$. We can now define the CGS $\mathcal{G}_{PRS} \triangleq \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_i \rangle$ for the PRS game. The labeling is given by $L(s_i) \triangleq \emptyset$, $L(s_A) \triangleq \{w_A\}$, and $L(s_B) \triangleq \{w_B\}$. Moreover, the transition function is defined as follows, where $D_A \triangleq \{dc \in \text{Dc}_{\mathcal{G}_{PRS}} : (dc(A), dc(B)) \in C\}$ and $D_B \triangleq \{dc \in \text{Dc}_{\mathcal{G}_{PRS}} : (dc(B), dc(A)) \in C\}$ are the sets of winning decisions for the two agents: if $s = s_i$ and $dc \in D_A$ then $\text{tr}(s, dc) \triangleq s_A$, else if $s = s_i$ and $dc \in D_B$ then $\text{tr}(s, dc) \triangleq s_B$, otherwise $\text{tr}(s, dc) \triangleq s$. Note that, when none of the two agents catches the action of the other, i.e., the used decision is in $D_i \triangleq \text{Dc}_{\mathcal{G}_{PRS}} \setminus (D_A \cup D_B)$, the play remains in the idle state to allow another try, otherwise it is stuck in a winning position forever.




Figure 1.1: The CGS \mathcal{G}_{PRS} .

We now describe a non-classic qualitative version of the well-known prisoner’s dilemma.

Example 1.1.2 (Prisoner’s Dilemma) In the prisoner’s dilemma (PD, for short), two accomplices are interrogated in separated rooms by the police, which offers them the same agreement. If one defects, i.e., testifies for the prosecution against the other, while the other cooperates, i.e., remains silent, the defector goes free and the silent accomplice goes to jail. If both cooperate, they remain free, but will be surely interrogated in the next future waiting




Figure 1.2: The CGS \mathcal{G}_{PD} .

¹In the following, we use both $X \rightarrow Y$ and Y^X to denote the set of functions from the domain X to the codomain Y .

for a defection. On the other hand, if they both defect, both go to jail. Note that no one of the two prisoners knows about the choice made by the other. This tricky situation can be modeled by the CGS $\mathcal{G}_{PD} \triangleq \langle AP, Ag, Ac, St, tr, ap, s_i \rangle$ depicted in Figure 1.2 on the preceding page, where the agents “Accomplice-1” and “Accomplice-2” in $Ag \triangleq \{A_1, A_2\}$ can chose an action in $Ac \triangleq \{C, D\}$, which stand for “cooperation” and “defection”, respectively. There are four states in $St \triangleq \{s_i, s_{A_1}, s_{A_2}, s_j\}$. In the idle state s_i the agents are waiting for the interrogation, while s_j represents the jail for both of them. Instead the remaining states s_{A_1} and s_{A_2} indicate the situations in which only one of the agents becomes definitely free. To characterize the different meaning of these states, we use the atomic propositions in $AP \triangleq \{f_{A_1}, f_{A_2}\}$, which denote who is “free”, by defining the following labeling: $L(s_i) \triangleq \{f_{A_1}, f_{A_2}\}$, $L(s_{A_1}) \triangleq \{f_{A_1}\}$, $L(s_{A_2}) \triangleq \{f_{A_2}\}$, and $L(s_j) \triangleq \emptyset$. The transition function tr can be easily deduced by the figure.

We can analyze also an extended version of the well-known prisoner’s dilemma [OR94] in which also the actions of the police are taken into account.

Example 1.1.3 (Prisoners and Police’s Dilemma)

In the prisoner and police’s dilemma (PPD, for short), apart from the classic agreement of the PD, the two accomplices also know that they can try to gain a better sentence by the judge if one spontaneously defects without being interrogated by police, since he is considered a “good willing man”. In this case, indeed, if the other cooperates, the defector becomes definitely free, while the other goes to jail with the possibility to eventually be released. It is important, however, that no both of them defect, otherwise the police can subtle act as they were interrogated. Moreover, differently from the PD, they are not free during the time in which they can be interrogated. This complex situation, can be modeled by the CGS

$\mathcal{G}_{PPD} \triangleq \langle AP, Ag, Ac, St, tr, ap, s_0 \rangle$ depicted in Figure 1.3, where there are three agents in $Ag \triangleq \{A_1, A_2, P\}$, with P being the police, and all of them can execute the two abstract actions in $Ac \triangleq \{0, 1\}$. For the accomplices, 0 and 1 have the meaning of “cooperate” and “defect”. For the police, on the contrary, they mean “wait” and “interrogate” in all the states but those in which one of the accomplices can eventually be released, where the meaning is “release” and “maintain”, instead. The set of states for the game is given by $\{s_i, s_{A_1}, s_{A_2}, s_j, s_{A_1j}, s_{A_2j}, s_{A_1A_2}\}$. The idle state s_i the two prisoners are waiting to be interrogated by police. They can even decide to defect before interrogation. The states s_{A_1} and s_{A_2} denotes the situation in which only one prisoner becomes definitely free. Moreover, the states s_{A_1j} and s_{A_2j} indicate when one of the prisoner is free while the other in the jail is waiting for his release. Finally, $s_{A_1A_2}$ denotes the state in which both have gained definitely the freedom. To represent the different meaning of these states, we use the atomic propositions f_{A_i} to denote that the prisoner A_i is free. Both the labeling function ap and the transition

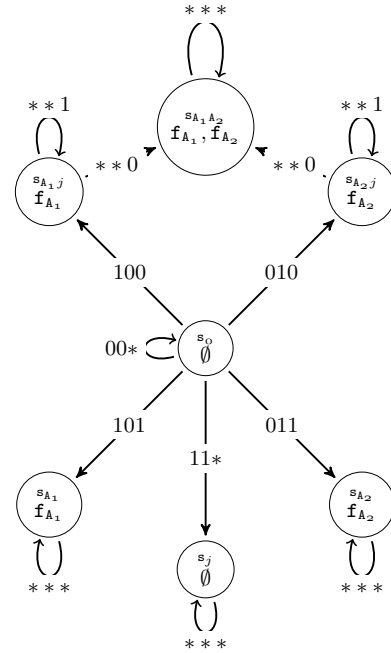


Figure 1.3: The CGS \mathcal{G}_{PPD} .

1.1. Syntax and Semantics

function tr can be extracted from the figure.

In addition to PPD, we model a very simple *preemptive scheduling protocol* for the access of processes to a shared resource.

Example 1.1.4 (Preemptive Scheduling) Consider the following preemptive scheduling protocol (PS, for short) describing the access rules of two processes to a shared resource in a preemptive way. When the resource is free and only one process asks for it, this process receives directly the grant. Instead, if there is a competition of requests, it is the scheduler that, in a nondeterministic way, determines who can access to the resource. Finally, in case one process owns the resource while the other asks for it, the scheduler can choose to make or not a preemption. These rules are formalized in the CGS $\mathcal{G}_{PS} \triangleq \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_0 \rangle$ of Figure 1.4, where the agents “Process-1”, “Process-2” and “Scheduler” in $\text{Ag} \triangleq \{P_1, P_2, S\}$ can choose between the two abstract actions in $\text{Ac} \triangleq \{0, 1\}$. The processes use the actions 0 and 1 to send or not the request to the scheduler, while this uses them in order to decide who can have the access to the resource in a situation of competition. There are five states $\text{St} \triangleq \{s_i, s_1, s_2, s_{1,2}, s'_1, s'_2\}$ in which the protocol can reside: the idle state s_i in which the resource is free; the three states s_1 , s_2 , and $s_{1,2}$ in which P_1 , P_2 or both are requesting the resource; the two states s'_1 and s'_2 in which the resource has been finally granted to P_1 and P_2 , respectively. To represent all information associated, we use the atomic propositions in $\text{AP} \triangleq \{r_1, r_2, g_1, g_2\}$, where r_i represents the request of P_i , while g_i the fact that the resource has been granted to P_i .

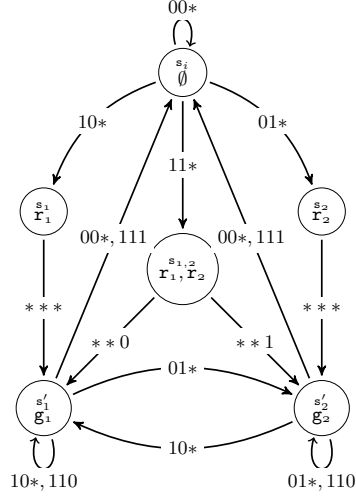


Figure 1.4: The CGS \mathcal{G}_{PS} .

As last example in which it is possible to discuss about the ability of an agent to be safe with respect to the behavior of malicious agents, we consider a generalization of the well known robber-vs-cops game (see [AF84] for a survey).

Example 1.1.5 (Romeo and Juliet Game) In the Romeo and Juliet game, Romeo, being able to move inside Juliet’s house, aims to reach her balcony, trying to avoid meeting both the Montagues and Capulets families, who are able to move concurrently with him in the house. Graphically, the house is divided in rooms and each room has doors on all internal walls. The topology of the house depends on the status of these doors that can be closed or open.




Figure 1.5: Romeo and Juliet game.

Additionally, Romeo is helped by the accomplice Shakespeare that controls the doors of the rooms. Specifically, at each step of the game, Shakespeare can switch the status of at most two doors at a time (i.e., from open to closed or viceversa). Hence, he can modify the shape of the house during a play. We consider Romeo meeting the families in two cases: if they are in the same room in a certain phase of the play (Romeo is captured) or if they are

in adjacent rooms with the connecting door open (Romeo is shot). In each step, the agents Romeo, Montagues family, and Capulets family can independently decide to stay in the room or, in the case a connection door to another room is open, move to this one. A sample game given in Figure 1.5 on the facing page is made by nine rooms, with Room 9 being the balcony, Romeo in Room 1, Montagues family in Room 3 and Capulets family in Room 7, respectively, and all the doors open as initial state of the game. We denote with $R = [1, 9]$ the set of rooms and $D = \{d = (r_1, r_2) \in R \times R : r_1 \text{ is adjacent to } r_2\}$ the set of doors.

We can model this situation with a CGS $\mathcal{G}_{RJ} = \langle AP, Ag, Ac, St, tr, ap, s_0 \rangle$ with $AP \triangleq \{c, r\}$ standing for Captured or Shot (c) and Reached (r); $Ag \triangleq \{R, M, C, S\}$, standing for Romeo (R), Montagues (M), Capulets (C), Shakespeare (S), respectively; $Ac \triangleq M \cup C$ where $M \triangleq \{\text{stay, up, down, left, right}\}$ is the set of moves for Romeo and Families and $C = \{d \in 2^D : |d| \leq 2\}$ be the set of possible switching of doors for Shakespeare, in which the empty set stands for no switching; $St \triangleq R \times R^2 \times \{0, 1\}^D$ be the set of states in which the first coordinate stands for the position of Romeo, the second and third coordinates stand for the positions of Montagues and Capulets, respectively, and the fourth coordinate is a function describing the status open or closed of each door in the house. The initial state is given by $(1, 3, 7, f_0)$ where $f_0 : D \rightarrow \{0, 1\}$ is the function constant to 0, i.e., all the doors are open. The labeling function is given as follows: for each state $s = (n, m_1, m_2, f) \in St$, $r \in L(s)$ iff $n = 9$,² $c \in L(s)$ iff there exists $i \in \{1, 2\}$ such that either $n = m_i$ or $(n, m_i) \in D$ and $f((n, m_i)) = 0$.³ The transition function can be easily deduced by the figure. For the sake of simplicity, all the states labeled with r only have a loop in the game.⁴

1.1.2 Syntax

Strategy Logic (SL, for short) syntactically extends LTL by means of two *strategy quantifiers*, the existential $\langle\langle x \rangle\rangle$ and the universal $\llbracket x \rrbracket$, and the *agent binding* (a, x) , where a is an agent and x a variable. Intuitively, these new elements can be read as “there exists a strategy x ”, “for all strategies x ”, and “bind agent a to the strategy associated with the variable x ”, respectively. The formal syntax of SL follows.

Definition 1.1.2 (SL Syntax) SL formulas are built inductively from the sets of atomic propositions AP , variables $\forall r$, and agents Ag , by using the following grammar, where $p \in AP$, $x \in \forall r$, and $a \in Ag$:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \forall r\varphi \mid \varphi \cup \varphi \mid \varphi \mathbf{R}\varphi \mid \langle\langle x \rangle\rangle\varphi \mid \llbracket x \rrbracket\varphi \mid (a, x)\varphi.$$

SL denotes the infinite set of formulas generated by the above rules.

Observe that, by construction, LTL is a proper syntactic fragment of SL, i.e., $LTL \subset SL$. In order to abbreviate the writing of formulas, we use the boolean values true \mathbf{t} and false \mathbf{f} and

²Romeo has reached the balcony.

³Romeo is in the same room with a family or they are close with the door open.

⁴For the sake of simplicity, we assume that each agent, at each state of the game, has a specific subset of possible choices of actions. For instance, if Romeo stays in the room 4 and the door (4, 5) is closed, than only the actions up, down, and stay are available to him. It is easy to see that such assumption can be overcome in a suitable, but even tricky, CGS that is somehow equivalent.

1.1. Syntax and Semantics

the well-known temporal operators future $F\varphi \triangleq \tau U\varphi$ and globally $G\varphi \triangleq f R\varphi$. Moreover, we use the italic letters x, y, z, \dots , possibly with indexes, as meta-variables on the variables $\mathbf{x}, \mathbf{y}, \mathbf{z}, \dots$ in V_I .

A first classic notion related to the syntax of SL is that of *subformula*, *i.e.*, a syntactic expression that is part of an a priori given formula. By $\text{sub}(\varphi)$ we formally denote the set of subformulas of an SL formula φ . For instance, consider $\varphi = \langle\langle \mathbf{x} \rangle\rangle(\alpha, \mathbf{x})(Fp)$. Then, it is immediate to see that $\text{sub}(\varphi) = \{\varphi, (\alpha, \mathbf{x})(Fp), (Fp), p, \tau\}$.

Usually, in predicative logics, we need the concepts of *free* and *bounded* placeholders, to correctly define the meaning of a formula. In SL, we have two different kind of placeholders: variables and agents. The former is used in the strategy quantifications, the latter to commit an agent, by means of bindings, to adhere to a strategy. Consequently, we need to differentiate the sets of free variables and free agents of an SL formula φ . The first contains the variables that are not in a scope of a quantification. The second, instead, contains the agents for which there is no related binding in the scope of a temporal operator. A formula without any free variable (*resp.*, agent) is named *variable-closed* (*resp.*, *agent-closed*). A formula that is both variable- and agent-closed, is named *sentence*. For a given SL formula φ , by $\text{free}(\varphi)$ we denote the set of both free variables and agents occurring in φ . The formal definition of $\text{free}(\cdot)$ has been given in [MMPV11, MMPV14]. Just as an example, let $\varphi = \langle\langle \mathbf{x} \rangle\rangle(\alpha, \mathbf{x})(\beta, \mathbf{y})(Fp)$ be a formula on the agents $\text{Ag} = \{\alpha, \beta, \gamma\}$. Then, we have $\text{free}(\varphi) = \{\gamma, \mathbf{y}\}$, since γ is an agent without any binding after Fp and \mathbf{y} has no quantification at all. Consider also the formulas $(\alpha, \mathbf{z})\varphi$ and $(\gamma, \mathbf{z})\varphi$, where the subformula φ is the same as above. Then, we have $\text{free}((\alpha, \mathbf{z})\varphi) = \text{free}(\varphi)$ and $\text{free}((\gamma, \mathbf{z})\varphi) = \{\mathbf{y}, \mathbf{z}\}$, since α is not free in φ but γ is, *i.e.*, $\alpha \notin \text{free}(\varphi)$ and $\gamma \in \text{free}(\varphi)$. So, $(\gamma, \mathbf{z})\varphi$ is agent-closed while $(\alpha, \mathbf{z})\varphi$ is not.

We now introduce other two general restrictions in which the numbers $|\text{Ag}|$ of agents and $|V_I|$ of variables that are used to write a formula are fixed to the a priori values $n, m \in [1, \omega]$, respectively. Moreover, we can also forbid the sharing of variables, *i.e.*, each variable is binded to one agent only, so, we cannot force two agents to use the same strategy. We name these three fragments $\text{SL}[n\text{-AG}]$, $\text{SL}[m\text{-VAR}]$, and $\text{SL}[FVS]$, respectively.

To make practice with the syntax of SL, we now describe few examples of some game-theoretic properties, which cannot be expressed neither in ATL^* nor in CHP-SL . We clarify this point later in the paper.

The first we introduce is the well-known concept of *Nash Equilibrium* in concurrent infinite games with Boolean payoffs.

Example 1.1.6 (Nash Equilibrium) *Consider the n agents $\alpha_1, \dots, \alpha_n$ of a game, each of them having, respectively, a possibly different temporal goal described by one of the LTL formulas ψ_1, \dots, ψ_n . Then, we can express the existence of a strategy profile $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ that is a Nash equilibrium (NE, for short) for $\alpha_1, \dots, \alpha_n$ w.r.t. ψ_1, \dots, ψ_n by using the SL sentence $\varphi_{NE} \triangleq \langle\langle \mathbf{x}_1 \rangle\rangle \cdots \langle\langle \mathbf{x}_n \rangle\rangle(\alpha_1, \mathbf{x}_1) \cdots (\alpha_n, \mathbf{x}_n) \psi_{NE}$, where $\psi_{NE} \triangleq \bigwedge_{i=1}^n (\langle\langle \mathbf{y} \rangle\rangle(\alpha_i, \mathbf{y})\psi_i) \rightarrow \psi_i$ is a variable-closed formula. Informally, this asserts that every agent α_i has \mathbf{x}_i as one of the best strategy w.r.t. the goal ψ_i , once all the other strategies of the remaining agents α_j , with $j \neq i$, have been fixed to \mathbf{x}_j . Note that here we are only considering equilibria under deterministic strategies.*

Chapter 1. Strategy Logic

Note that the syntactic feature of SL that allows us to represent Nash Equilibria is the binding construct. Indeed, by means of a suitable usage of it, we can compare, in a Boolean way, the outcomes of two strategy profiles in which only one agent has changed his choice. This cannot be done in ATL^* , since the use of an agent modality necessarily re-quantify the strategies associated with all the agents. Therefore, we cannot change in ATL^* just a strategy of a single selected agent.

In game theory, it is important to recall that an equilibrium is not always stable. Indeed, there are games like the PD of Example 1.1.2 on page 2 having Nash equilibria that are instable. One of the simplest concepts of stability that is possible to think is called *stability profile*.

Example 1.1.7 (Stability Profile) *Think about the same situation of the above example on NE. Then, a stability profile (SP, for short) is a strategy profile (x_1, \dots, x_n) for $\alpha_1, \dots, \alpha_n$ w.r.t. ψ_1, \dots, ψ_n such that there is no agent α_i that can choose a different strategy from x_i without changing its own payoff and penalizing the payoff of another agent α_j , with $j \neq i$. To represent the existence of such a profile, we can use the SL sentence $\varphi_{SP} \triangleq \langle\langle x_1 \rangle\rangle(\alpha_1, x_1) \cdots \langle\langle x_n \rangle\rangle(\alpha_n, x_n) \psi_{SP}$, where $\psi_{SP} \triangleq \bigwedge_{i,j=1, i \neq j}^n \psi_j \rightarrow \llbracket y \rrbracket((\psi_i \leftrightarrow (\alpha_i, y)\psi_i) \rightarrow (\alpha_i, y)\psi_j)$. Informally, with the ψ_{SP} subformula, we assert that, if α_j is able to achieve his goal ψ_j , all strategies y of α_i that left unchanged the payoff related to ψ_i , also let α_j to maintain his achieved goal. At this point, it is very easy to ensure the existence of an NE that is also an SP, by using the SL sentence $\varphi_{SNE} \triangleq \langle\langle x_1 \rangle\rangle(\alpha_1, x_1) \cdots \langle\langle x_n \rangle\rangle(\alpha_n, x_n) \psi_{SP} \wedge \psi_{NE}$.*

In a game in which not all agents are peers, we can have one or more of them that may vary the payoff of the others, without having a personal aim, *i.e.*, without looking for the maximization of their own payoffs. Such situations can usually arise when we have games with arbiters or similar characters, like supervisors or government authorities, that have to be fair, *i.e.*, they have to lay down an *equity governance*.

Example 1.1.8 (Equity Governance) *Let us think to a game similar to the one described in the previous example, in which there is also a supervisor agent β , which does not have a specific goal. By the way, the peers want him to be fair w.r.t. their own goals, *i.e.*, the supervisor has to use a strategy that does not have to prefer one agent over another. This concept is called equity governance (EG, for short). In order to formalize it, we can use the SL sentence $\varphi_{EG} \triangleq \llbracket x_1 \rrbracket \cdots \llbracket x_n \rrbracket(\alpha_1, x_1) \cdots (\alpha_n, x_n) \langle\langle y \rangle\rangle(\beta, y) \psi_{EG}$, where $\psi_{EG} \triangleq \bigwedge_{i,j=1, i < j}^n (\langle\langle z_1 \rangle\rangle(\beta, z_1)\psi_i \wedge (\langle\langle z_2 \rangle\rangle(\beta, z_2)\psi_j) \rightarrow (\psi_i \leftrightarrow \psi_j))$. Informally, the ψ_{EG} subformula asserts that, if there are two strategies z_1 and z_2 for β that allow α_i and α_j to achieve their own goals ψ_i and ψ_j , separately, then the unique strategy y previously chosen by the supervisor has to ensure the achievement of either both the goals or none of them. Note that the sentence φ_{EG} ensure the existence of an EG strategy y for β , in dependence of the strategies x_1, \dots, x_n chosen by the peers. To verify the existence of a uniform EG, we may use the SL sentence $\varphi_{UEG} \triangleq \langle\langle y \rangle\rangle(\beta, y) \psi'_{EG}$, with $\psi'_{EG} \triangleq \llbracket x_1 \rrbracket \cdots \llbracket x_n \rrbracket(\alpha_1, x_1) \cdots (\alpha_n, x_n) \psi_{EG}$, whose difference w.r.t. φ_{EG} resides only in the alternation of quantifiers. Finally, to verify the existence of a uniform EG that allows also the existence of an NE for the peers, we can use the SL sentence $\varphi_{UEG+NE} \triangleq \langle\langle y \rangle\rangle(\beta, y) (\psi'_{EG} \wedge \varphi_{NE})$.*

1.1. Syntax and Semantics

Usually, the equity of a supervisor does not ensure that the whole game can advance, *i.e.*, that the peers can achieve their respective goals. Indeed, there are games like the zero-sum ones in which the agents have opposite goals that cannot be achieved at the same time. However, there are different kind of games, as the PPD or the PS of Examples 1.1.3 on page 3 and 1.1.4 on page 4, in which a supervisor can try to help all peers in their intent, by applying an *advancement governance*.

Example 1.1.9 (Advancement Governance) *Consider the game described in the previous example of EG. Here, we want to consider an advancement governance (AG, for short) for the supervisor, *i.e.*, a strategy for β that allows the peers to achieve their own goals, if they have the will and possibility to do so. Formally, this concept can be expressed by using the SL sentence $\varphi_{AG} \triangleq \llbracket \mathbf{x}_1 \rrbracket \cdot \llbracket \mathbf{x}_n \rrbracket (\alpha_1, \mathbf{x}_1) \cdots (\alpha_n, \mathbf{x}_n) \langle \mathbf{y} \rangle (\beta, \mathbf{y}) \psi_{AG}$, where $\psi_{AG} \triangleq (\bigwedge_{i=1}^n (\langle \mathbf{z} \rangle (\beta, \mathbf{z}) \psi_i) \rightarrow \psi_i)$. Intuitively, the ψ_{AG} subformula expresses the fact that, if β has a strategy \mathbf{z} able to force a goal ψ_i , once the strategies of the peers $\alpha_1, \dots, \alpha_n$ have been fixed, then his a priori choice \mathbf{y} w.r.t. the goals has to force ψ_i as well. As in the case of EG, we can have an uniform version of AG, by using the SL sentence $\varphi_{UAG} \triangleq \langle \mathbf{y} \rangle (\beta, \mathbf{y}) \psi'_{AG}$, where $\psi'_{AG} \triangleq \llbracket \mathbf{x}_1 \rrbracket \cdots \llbracket \mathbf{x}_n \rrbracket (\alpha_1, \mathbf{x}_1) \cdots (\alpha_n, \mathbf{x}_n) \psi_{AG}$.*

Differently from the previous examples, one can consider the case in which the authority agent has his own goal to be satisfied, given the other agents to be in a certain equilibrium. In the context of system design [PR89], *rational synthesis* [FKL10] is a recent improvement of the classical reactive one. In this setting, the adversarial environment is not a monolithic block, but a set of agent components, each of them having their own goal. In the next example, we show that the most typical instances of a rational synthesis problem can be represented in SL.

Example 1.1.10 (Rational Synthesis) *Consider a solution concept that is representable in SL by means of a suitable formula ψ_{SC} , *e.g.*, NE, and a temporal goal ψ_β for the system agent. Here, we look for a rational synthesis solution for the players, *i.e.*, a strategy profile (y, x_1, \dots, x_n) such that, if β acts according to y , then ψ_β is satisfied and (x_1, \dots, x_n) is in an equilibrium according to the solution concept considered, *i.e.*, ψ_{SC} is satisfied. Formally, this concept can be expressed by using the SL sentence $\varphi_{RS} = \langle \mathbf{y} \rangle \langle \mathbf{x}_1 \rangle \cdots \langle \mathbf{x}_n \rangle (\beta, \mathbf{y}) (a_1, \mathbf{x}_1) \cdots (a_n, \mathbf{x}_n) (\psi_0 \wedge \psi_{SC})$. As an example, ψ_{SC} can be the formula ψ_{NE} of Example 1.1.6 on page 6. In this case, we obtain the rational synthesis problem for NE.*

1.1.3 Basic notions

Before continuing with the formal description of SL, we need to introduce some basic notions related to CGSs, such as those of *track*, *path*, *strategy*, and the like. They are the natural analogous of the ones for Kripke structures and, more generally, for game arenas. All these notions have been already introduced in [MMV10b]. However, for the sake of completeness, as well as for their importance in the definition of SL semantics, we fully report them in this section.

We start with the notions of *track* and *path*. Intuitively, tracks and paths of a CGS are legal sequences of reachable states that can be respectively seen as partial and complete

descriptions of possible outcomes of the game modeled by the structure itself. Formally, a *track* (resp., *path*) in a CGS \mathcal{G} is a finite (resp., an infinite) sequence of states $\rho \in \text{St}^*$ (resp., $\pi \in \text{St}^\omega$) such that, for all $i \in [0, |\rho| - 1]$ (resp., $i \in \mathbb{N}$), there exists a decision $\delta \in \text{Dc}$ such that $(\rho)_{i+1} = \text{tr}((\rho)_i, \delta)$ (resp., $(\pi)_{i+1} = \text{tr}((\pi)_i, \delta)$)⁵. A track ρ is *non-trivial* if it has non-zero length, i.e., $|\rho| > 0$ that is $\rho \neq \epsilon$ ⁶. The set $\text{Trk} \subseteq \text{St}^+$ (resp., $\text{Pth} \subseteq \text{St}^\omega$) contains all non-trivial tracks (resp., paths). Moreover, $\text{Trk}(s) \triangleq \{\rho \in \text{Trk} : \text{fst}(\rho) = s\}$ (resp., $\text{Pth}(s) \triangleq \{\pi \in \text{Pth} : \text{fst}(\pi) = s\}$) indicates the subsets of tracks (resp., paths) starting at a state $s \in \text{St}$ ⁷.

As an example, consider the CGS \mathcal{G}_{PS} in Figure 1.4 on page 4. Then $\rho = s_i \cdot s_1 \cdot s_1' \cdot s_1' \cdot s_i$ and $\pi = (s_i \cdot s_{1,2} \cdot s_1' \cdot s_i \cdot s_{1,2} \cdot s_2')^\omega$ are a track and a path, respectively. Moreover, we have that $\text{Trk}_{\mathcal{G}_{PS}} = \text{St}_{\mathcal{G}_{PS}}^*$ and $\text{Pth}_{\mathcal{G}_{PS}} = \text{St}_{\mathcal{G}_{PS}}^\omega$.

At this point, we can define the concept of *strategy*. Intuitively, a strategy is a scheme for an agent that contains all choices of actions as in dependence of the history of the current outcome. However, observe that here there is not an a priori connection between a strategy and an agent, since the same strategy can be used by more than one agent at the same time. Formally, a *strategy* in a CGS \mathcal{G} is a partial function $f : \text{Trk} \rightarrow \text{Ac}$ that maps each non-trivial track in its domain to an action. For a state $s \in \text{St}$, a strategy f is said *s-total* if it is defined on all tracks starting in s , i.e., $\text{dom}(f) = \text{Trk}(s)$. The set $\text{Str} \triangleq \text{Trk} \rightarrow \text{Ac}$ (resp., $\text{Str}(s) \triangleq \text{Trk}(s) \rightarrow \text{Ac}$) contains all (resp., *s-total*) strategies.

An example of strategy in the CGS \mathcal{G}_{PS} is given by the function $f_1 \in \text{Str}$ assigning the action 0 to all the tracks in which the state s_i occurs an even number of times and the action 1, otherwise. Another example of strategy is the function $f_2 \in \text{Str}$ assigning the action 1 on all possible tracks of the CGS.

We now introduce the notion of *assignment*. Intuitively, an assignment gives a valuation of variables with strategies, where the latter are used to determine the behavior of agents in the game. With more detail, as in the case of first order logic, we use this concept as a technical tool to quantify over strategies associated with variables, independently of agents to which they are related to. So, assignments are used precisely as a way to define a correspondence between variables and agents via strategies.

Definition 1.1.3 (Assignments) *An assignment in a CGS \mathcal{G} is a partial function $\chi : \text{Vr} \cup \text{Ag} \rightarrow \text{Str}$ mapping variables and agents in its domain to a strategy. An assignment χ is complete if it is defined on all agents, i.e., $\text{Ag} \subseteq \text{dom}(\chi)$. For a state $s \in \text{St}$, it is said that χ is *s-total* if all strategies $\chi(l)$ are *s-total*, for $l \in \text{dom}(\chi)$. The set $\text{Asg} \triangleq \text{Vr} \cup \text{Ag} \rightarrow \text{Str}$ (resp., $\text{Asg}(s) \triangleq \text{Vr} \cup \text{Ag} \rightarrow \text{Str}(s)$) contains all (resp., *s-total*) assignments. Moreover, $\text{Asg}(X) \triangleq X \rightarrow \text{Str}$ (resp., $\text{Asg}(X, s) \triangleq X \rightarrow \text{Str}(s)$) indicates the subset of *X-defined* (resp., *s-total*) assignments, i.e., (resp., *s-total*) assignments defined on the set $X \subseteq \text{Vr} \cup \text{Ag}$.*

As an example of assignment, consider the function $\chi_1 \in \text{Asg}$ in the CGS \mathcal{G}_{PS} , with $\text{dom}(\chi_1) = \{P_1, x\}$, such that $\chi_1(P_1) = f_1$ and $\chi_1(x) = f_2$. As another example, consider the assignment $\chi_2 \in \text{Asg}$ in the same CGS, with $\text{dom}(\chi_2) = \text{Ag}$, such that $\chi_2(S) = f_1$ and $\chi_2(P_1) = \chi_2(P_2) = f_2$. Note that χ_2 is complete, while χ_1 is not.

⁵The notation $(w)_i \in \Sigma$ indicates the *element* of index $i \in [0, |w|]$ of a non-empty sequence $w \in \Sigma^\infty$.

⁶The Greek letter ϵ stands for the *empty sequence*.

⁷By $\text{fst}(w) \triangleq (w)_0$ it is denoted the *first element* of a non-empty sequence $w \in \Sigma^\infty$.

1.1. Syntax and Semantics

Given an assignment $\chi \in \text{Asg}$, an agent or variable $l \in \text{Vr} \cap \text{Ag}$, and a strategy $f \in \text{Str}$, we need to describe the *redefinition* of χ , *i.e.*, a new assignment equal to the first one on all elements of its domain but l , on which it assumes the value f . Formally, with $\chi[l \mapsto f] \in \text{Asg}$ we denote the new assignment defined on $\text{dom}(\chi[l \mapsto f]) \triangleq \text{dom}(\chi) \cup \{l\}$ that returns f on l and is equal to χ on the remaining part of its domain, *i.e.*, $\chi[l \mapsto f](l) \triangleq f$ and $\chi[l \mapsto f](l') \triangleq \chi(l')$, for all $l' \in \text{dom}(\chi) \setminus \{l\}$. Intuitively, if we have to add or update a strategy that needs to be bound by an agent or variable, we can simply take the old assignment and redefine it by using the above notation. It is worth observing that, if χ and f are s -total then $\chi[l \mapsto f]$ is s -total, as well.

Now, we can formalize the concept of *play* in a game. Intuitively, a play is the unique outcome of the game determined by all agent strategies participating to it.

Definition 1.1.4 (Plays) *A path $\pi \in \text{Pth}(s)$ starting at a state $s \in \text{St}$ is a play w.r.t. a complete s -total assignment $\chi \in \text{Asg}(s)$ ((χ, s) -play, for short) if, for all $i \in \mathbb{N}$, it holds that $(\pi)_{i+1} = \text{tr}((\pi)_i, \delta)$, where $\delta(a) \triangleq \chi(a)((\pi)_{\leq i})$, for each $a \in \text{Ag}$ ⁸. The partial function $\text{play} : \text{Asg} \times \text{St} \rightarrow \text{Pth}$, with $\text{dom}(\text{play}) \triangleq \{(\chi, s) : \text{Ag} \subseteq \text{dom}(\chi) \wedge \chi \in \text{Asg}(s) \wedge s \in \text{St}\}$, returns the (χ, s) -play $\text{play}(\chi, s) \in \text{Pth}(s)$, for all pairs (χ, s) in its domain.*

As last example, consider again the CGS \mathcal{G}_{PS} and the complete s_0 -total assignment χ_2 defined above. Then, we have that $\text{play}(\chi_2, s_0) = (s_i \cdot s_{1,2} \cdot s_1' \cdot s_i \cdot s_{1,2} \cdot s_2')^\omega$.

Finally, we give the definition of global translation of a complete assignment associated with a state, which is used to capture, at a certain step of the play, what is the current state and its updated assignment.

Definition 1.1.5 (Global Translation) *For a given state $s \in \text{St}$ and a complete s -total assignment $\chi \in \text{Asg}(s)$, the i -th global translation of (χ, s) , with $i \in \mathbb{N}$, is the pair of a complete assignment and a state $(\chi, s)^i \triangleq ((\chi)_{(\pi)_{\leq i}}, (\pi)_i)$, where $\pi = \text{play}(\chi, s)$ and $(\chi)_{(\pi)_{\leq i}}$ denotes the $(\pi)_i$ -total assignment, with $\text{dom}((\chi)_{(\pi)_{\leq i}}) = \text{dom}(\chi)$, such that, for all $l \in \text{dom}(\chi)$, $\text{dom}((\chi)_{(\pi)_{\leq i}}(l)) = \{\rho \in \text{Trk}((\pi)_i) : (\pi)_{\leq i} \cdot \rho \in \text{dom}(\chi(l))\}$ and $(\chi)_{(\pi)_{\leq i}}(l)(\rho) = \chi(l)((\pi)_{\leq i} \cdot \rho)$, for all $\rho \in \text{dom}((\chi)_{(\pi)_{\leq i}}(l))$.*

1.1.4 Semantics

As already reported at the beginning of this section, just like ATL^* and differently from CHP-SL , the semantics of SL is defined *w.r.t.* concurrent game structures. For an SL formula φ , a CGS \mathcal{G} , one of its states s , and an s -total assignment χ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$, we write $\mathcal{G}, \chi, s \models \varphi$ to indicate that the formula φ holds at s in \mathcal{G} under χ . The semantics of SL formulas involving the atomic propositions, the Boolean connectives \neg , \wedge , and \vee , as well as the temporal operators X , U , and R is defined as usual in LTL . The novel part resides in the formalization of the meaning of strategy quantifications $\langle\langle x \rangle\rangle$ and $\llbracket x \rrbracket$ and agent binding (a, x) .

Definition 1.1.6 (SL Semantics) *Given a CGS \mathcal{G} , for all SL formulas φ , states $s \in \text{St}$, and s -total assignments $\chi \in \text{Asg}(s)$ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$, the modeling relation $\mathcal{G}, \chi, s \models \varphi$ is inductively defined as follows.*

⁸The notation $(w)_{\leq i} \in \Sigma^*$ indicates the *prefix* up to index $i \in [0, |w|]$ of a non-empty sequence $w \in \Sigma^\infty$.

1. $\mathcal{G}, \chi, s \models p$ if $p \in \text{ap}(s)$, with $p \in \text{AP}$.
2. For all formulas φ, φ_1 , and φ_2 , it holds that:
 - (a) $\mathcal{G}, \chi, s \models \neg\varphi$ if not $\mathcal{G}, \chi, s \models \varphi$, that is $\mathcal{G}, \chi, s \not\models \varphi$;
 - (b) $\mathcal{G}, \chi, s \models \varphi_1 \wedge \varphi_2$ if $\mathcal{G}, \chi, s \models \varphi_1$ and $\mathcal{G}, \chi, s \models \varphi_2$;
 - (c) $\mathcal{G}, \chi, s \models \varphi_1 \vee \varphi_2$ if $\mathcal{G}, \chi, s \models \varphi_1$ or $\mathcal{G}, \chi, s \models \varphi_2$.
3. For a variable $x \in \text{Vr}$ and a formula φ , it holds that:
 - (a) $\mathcal{G}, \chi, s \models \langle\langle x \rangle\rangle\varphi$ if there is an s -total strategy $f \in \text{Str}(s)$ such that $\mathcal{G}, \chi[x \mapsto f], s \models \varphi$;
 - (b) $\mathcal{G}, \chi, s \models \llbracket x \rrbracket\varphi$ if, for all s -total strategies $f \in \text{Str}(s)$, it holds that $\mathcal{G}, \chi[x \mapsto f], s \models \varphi$.
4. For an agent $a \in \text{Ag}$, a variable $x \in \text{Vr}$, and a formula φ , it holds that $\mathcal{G}, \chi, s \models (a, x)\varphi$ if $\mathcal{G}, \chi[a \mapsto \chi(x)], s \models \varphi$.
5. Finally, if the assignment χ is complete, for all formulas φ, φ_1 , and φ_2 , it holds that:
 - (a) $\mathcal{G}, \chi, s \models \mathbf{X}\varphi$ if $\mathcal{G}, (\chi, s)^1 \models \varphi$;
 - (b) $\mathcal{G}, \chi, s \models \varphi_1 \mathbf{U} \varphi_2$ if there is an index $i \in \mathbb{N}$ with $k \leq i$ such that $\mathcal{G}, (\chi, s)^i \models \varphi_2$ and, for all indexes $j \in \mathbb{N}$ with $k \leq j < i$, it holds that $\mathcal{G}, (\chi, s)^j \models \varphi_1$;
 - (c) $\mathcal{G}, \chi, s \models \varphi_1 \mathbf{R} \varphi_2$ if, for all indexes $i \in \mathbb{N}$ with $k \leq i$, it holds that $\mathcal{G}, (\chi, s)^i \models \varphi_2$ or there is an index $j \in \mathbb{N}$ with $k \leq j < i$ such that $\mathcal{G}, (\chi, s)^j \models \varphi_1$.

Intuitively, at Items 3a and 3b, respectively, we evaluate the existential $\langle\langle x \rangle\rangle$ and universal $\llbracket x \rrbracket$ quantifiers over strategies, by associating them to the variable x . Moreover, at Item 4, by means of an agent binding (a, x) , we commit the agent a to a strategy associated with the variable x . It is evident that, due to Items 5a, 5b, and 5c, the LTL semantics is simply embedded into the SL one.

In order to complete the description of the semantics, we now give the classic notions of *model* and *satisfiability* of an SL sentence. We say that a CGS \mathcal{G} is a *model* of an SL sentence φ , in symbols $\mathcal{G} \models \varphi$, if $\mathcal{G}, \emptyset, s_0 \models \varphi$.⁹ In general, we also say that \mathcal{G} is a *model* for φ on $s \in \text{St}$, in symbols $\mathcal{G}, s \models \varphi$, if $\mathcal{G}, \emptyset, s \models \varphi$. An SL sentence φ is *satisfiable* if there is a model for it.

It remains to formalize the concepts of *implication* and *equivalence* between SL formulas, which are useful to describe transformations preserving the meaning of a specification. Given two SL formulas φ_1 and φ_2 , with $\text{free}(\varphi_1) = \text{free}(\varphi_2)$, we say that φ_1 *implies* φ_2 , in symbols $\varphi_1 \Rightarrow \varphi_2$, if, for all CGSs \mathcal{G} , states $s \in \text{St}$, and $\text{free}(\varphi_1)$ -defined s -total assignments $\chi \in \text{Asg}(\text{free}(\varphi_1), s)$, it holds that if $\mathcal{G}, \chi, s \models \varphi_1$ then $\mathcal{G}, \chi, s \models \varphi_2$. Accordingly, we say that φ_1 is *equivalent* to φ_2 , in symbols $\varphi_1 \equiv \varphi_2$, if both $\varphi_1 \Rightarrow \varphi_2$ and $\varphi_2 \Rightarrow \varphi_1$ hold.

In the rest of the paper, especially when we describe a decision procedure, we may consider formulas in *existential normal form* (*enf*, for short) and *positive normal form* (*pnf*, for short),

⁹The symbol \emptyset stands for the empty function.

1.1. Syntax and Semantics

i.e., formulas in which only existential quantifiers appear or, respectively, the negation is applied solely to atomic propositions. In fact, it is to this aim that we have considered in the syntax of SL both the Boolean connectives \wedge and \vee , the temporal operators U , and R , and the strategy quantifiers $\langle\langle \cdot \rangle\rangle$ and $\llbracket \cdot \rrbracket$. Indeed, all formulas can be linearly translated in *enf* and *pnf* by using De Morgan's laws together with the following equivalences, which directly follow from the semantics of the logic: $\neg X\varphi \equiv X\neg\varphi$, $\neg(\varphi_1 U \varphi_2) \equiv (\neg\varphi_1)R(\neg\varphi_2)$, $\neg\langle\langle x \rangle\rangle\varphi \equiv \llbracket x \rrbracket\neg\varphi$, and $\neg(a, x)\varphi \equiv (a, x)\neg\varphi$.

At this point, in order to better understand the meaning of the SL semantics, we discuss some examples of formulas interpreted over the CGSs previously described.

We start by explaining how a strategy can be shared by different agents.

Example 1.1.11 (Shared Variable) Consider the SL sentence $\varphi = \langle\langle x \rangle\rangle \llbracket y \rrbracket \langle\langle z \rangle\rangle ((\alpha, x)(\beta, y)(Xp) \wedge (\alpha, y)(\beta, z)(Xq))$. It is immediate to note that both agents α and β use the strategy associated with y to achieve

simultaneously the LTL temporal goals Xp and Xq . A model for φ is given by the CGS $\mathcal{G}_{SV} \triangleq \langle \{p, q\}, \{\alpha, \beta\}, \{0, 1\}, \{s_0, s_1, s_2, s_3\}, L, tr, s_0 \rangle$, where $L(s_0) \triangleq \emptyset$, $L(s_1) \triangleq \{p\}$, $L(s_2) \triangleq \{p, q\}$, $L(s_3) \triangleq \{q\}$, $tr(s_0, (0, 0)) \triangleq s_1$, $tr(s_0, (0, 1)) \triangleq s_2$, $tr(s_0, (1, 0)) \triangleq s_3$, and all the remaining transitions (with any decision) go to s_0 . In Figure 1.6, we report a graphical representation of the structure. Clearly, $\mathcal{G}_{SV} \models \varphi$ by letting, on s_0 , the variables x to chose action 0 (the goal $(\alpha, x)(\beta, y)(Xp)$ is satisfied for any choice of y , since we can move from s_0 to either s_1 or s_2 , both labeled with p) and z to choose action 1 when y has action 0 and, vice versa, 0 when y has 1 (in both cases, the goal $(\alpha, y)(\beta, z)(Xq)$ is satisfied, since one can move from s_0 to either s_2 or s_3 , both labeled with q).




Figure 1.6: The CGS \mathcal{G}_{SV} .

We now discuss an application of the concepts of Nash equilibrium and stability profile to both the prisoner's dilemma and the paper, rock, and scissor game.

Example 1.1.12 (Equilibrium Profiles) Let us first consider the CGS \mathcal{G}_{PD} of the prisoner's dilemma described in the Example 1.1.2 on page 2. Intuitively, each of the two accomplices A_1 and A_2 want to avoid the prison. These goals can be represented by the LTL formulas $\psi_{A_1} \triangleq Gf_{A_1}$ and $\psi_{A_2} \triangleq Gf_{A_2}$, respectively. The existence of a Nash equilibrium in \mathcal{G}_{PD} for the two accomplices w.r.t. the above goals can be written as $\phi_{NE} \triangleq \langle\langle x_1 \rangle\rangle (A_1, x_1) \langle\langle x_2 \rangle\rangle (A_2, x_2) \psi_{NE}$, where $\psi_{NE} \triangleq ((\llbracket y \rrbracket (A_1, y) \psi_{A_1}) \rightarrow \psi_{A_1}) \wedge ((\llbracket y \rrbracket (A_2, y) \psi_{A_2}) \rightarrow \psi_{A_2})$, which results to be an instantiation of the general sentence φ_{NE} of Example 1.1.6 on page 6. In the same way, the existence of a stable Nash equilibrium can be represented with the sentence $\phi_{SNE} \triangleq \langle\langle x_1 \rangle\rangle (A_1, x_1) \langle\langle x_2 \rangle\rangle (A_2, x_2) \psi_{NE} \wedge \psi_{SP}$, where $\psi_{SP} \triangleq (\psi_1 \rightarrow \llbracket y \rrbracket ((\psi_2 \leftrightarrow (A_2, y) \psi_2) \rightarrow (A_2, y) \psi_1)) \wedge (\psi_2 \rightarrow \llbracket y \rrbracket ((\psi_1 \leftrightarrow (A_1, y) \psi_1) \rightarrow (A_1, y) \psi_2))$, which is a particular case of the sentence φ_{SNE} of Example 1.1.7 on page 7. Now, it is easy to see that $\mathcal{G}_{PD} \models \phi_{SNE}$ and, so, $\mathcal{G}_{PD} \models \phi_{NE}$. Indeed, an assignment $\chi \in \text{Asg}_{\mathcal{G}_{PD}}(Ag, s_i)$, for which $\chi(A_1)(s_i) = \chi(A_2)(s_i) = D$, is a stable equilibrium profile, *i.e.*, it is such that $\mathcal{G}_{PD}, \chi, s_i \models \psi_{NE} \wedge \psi_{SP}$. This is due to the fact that, if an agent A_k , for $k \in \{1, 2\}$, choses another strategy $f \in \text{Str}_{\mathcal{G}_{PD}}(s_i)$, he is still unable to achieve his goal ψ_k , *i.e.*, $\mathcal{G}_{PD}, \chi^{A_k \mapsto f}, s_i \not\models \psi_k$, so, he

cannot improve his payoff. Moreover, this equilibrium is stable, since the payoff of an agent cannot be made worse by the changing of the strategy of the other agent. However, it is interesting to note that there are instable equilibria too. One of these is represented by the assignment $\chi' \in \text{Asg}_{\mathcal{G}_{PD}}(\text{Ag}, \mathbf{s}_i)$, for which $\chi'(\mathbf{A}_1)(\mathbf{s}_i^j) = \chi'(\mathbf{A}_2)(\mathbf{s}_i^j) = \mathbf{C}$, for all $j \in \mathbb{N}$. Indeed, we have that $\mathcal{G}_{PD}, \chi', \mathbf{s}_i \models \psi_{NE}$, since $\mathcal{G}_{PD}, \chi', \mathbf{s}_i \models \psi_1$ and $\mathcal{G}_{PD}, \chi', \mathbf{s}_i \models \psi_2$, but $\mathcal{G}_{PD}, \chi', \mathbf{s}_i \not\models \psi_{SP}$. The latter property holds because, if one of the agents \mathbf{A}_k , for $k \in \{1, 2\}$, choses a different strategy $f' \in \text{Str}_{\mathcal{G}_{PD}}(\mathbf{s}_i)$ for which there is a $j \in \mathbb{N}$ such that $f'(\mathbf{s}_i^j) = \mathbf{D}$, he cannot improve his payoff but makes surely worse the payoff of the other agent, i.e., $\mathcal{G}_{PD}, \chi'[\mathbf{A}_k \mapsto f'], \mathbf{s}_i \models \psi_k$ but $\mathcal{G}_{PD}, \chi'[\mathbf{A}_k \mapsto f'], \mathbf{s}_i \not\models \psi_{3-k}$. Finally, consider the CGS \mathcal{G}_{PRS} of the paper, rock, and scissor game described in the Example 1.1.1 on page 2 together with the associated formula for the Nash equilibrium $\phi_{NE} \triangleq \langle\langle \mathbf{x}_1 \rangle\rangle(\mathbf{A}, \mathbf{x}_1) \langle\langle \mathbf{x}_2 \rangle\rangle(\mathbf{B}, \mathbf{x}_2) \psi_{NE}$, where $\psi_{NE} \triangleq (\langle\langle \mathbf{y} \rangle\rangle(\mathbf{A}, \mathbf{y}) \psi_{\mathbf{A}} \rightarrow \psi_{\mathbf{A}}) \wedge (\langle\langle \mathbf{y} \rangle\rangle(\mathbf{B}, \mathbf{y}) \psi_{\mathbf{B}} \rightarrow \psi_{\mathbf{B}})$ with $\psi_{\mathbf{A}} \triangleq \mathbf{Fw}_{\mathbf{A}}$ and $\psi_{\mathbf{B}} \triangleq \mathbf{Fw}_{\mathbf{B}}$ representing the LTL temporal goals for Alice and Bob, respectively. Then, it is not hard to see that $\mathcal{G}_{PRS} \not\models \phi_{NE}$, i.e., there are no Nash equilibria in this game, since there is necessarily an agent that can improve his/her payoff by changing his/her strategy.

Example 1.1.13 (Law and Order) Consider the CGS \mathcal{G}_{PPD} given in Example 1.1.3 on page 3. It is easy to see that Police can ensure at least one prisoner to be definitely in jail. Indeed the formula $\varphi_1 = \langle\langle \mathbf{y} \rangle\rangle \langle\langle \mathbf{x}_1 \rangle\rangle \langle\langle \mathbf{x}_2 \rangle\rangle (\mathbf{P}, \mathbf{y}) (\mathbf{A}_1, \mathbf{x}_1) (\mathbf{A}_2, \mathbf{x}_2) ((\mathbf{FG} \neg \mathbf{f}_{\mathbf{A}_1}) \vee (\mathbf{FG} \neg \mathbf{f}_{\mathbf{A}_2}))$ is satisfied over \mathcal{G}_{PPD} . A way to see this is to consider the strategy f for \mathbf{P} given by $f(\rho) = 0$, for all $\rho \in \text{Trk}$, which allows agent \mathbf{P} to always avoid the state $\mathbf{s}_{\mathbf{A}_1, \mathbf{A}_2}$. On the other hand, the formula $\varphi_2 = \langle\langle \mathbf{y} \rangle\rangle \langle\langle \mathbf{x}_1 \rangle\rangle \langle\langle \mathbf{x}_2 \rangle\rangle (\mathbf{P}, \mathbf{y}) (\mathbf{A}_1, \mathbf{x}_1) (\mathbf{A}_2, \mathbf{x}_2) (\mathbf{FG}(\neg \mathbf{f}_{\mathbf{A}_1} \wedge \neg \mathbf{f}_{\mathbf{A}_2}))$ is not satisfied over \mathcal{G}_{PPD} . Indeed, if agents \mathbf{A}_1 and \mathbf{A}_2 use strategies f_1 and f_2 , respectively, such that $f_2(\mathbf{s}_i) = 1 - f_1(\mathbf{s}_i)$, we have that, whatever agent \mathbf{P} does, at least one of them gains freedom.

Example 1.1.14 (Fair Scheduler) Consider the CGS \mathcal{G}_{PS} of Example 1.1.4 on page 4 and suppose the Scheduler wants to ensure that, whatever process makes a request, the resource is eventually granted to it. We can represent this specification by means of the formula $\varphi = \langle\langle \mathbf{y} \rangle\rangle \langle\langle \mathbf{x}_1 \rangle\rangle \langle\langle \mathbf{x}_2 \rangle\rangle (\mathbf{S}, \mathbf{y}) (\mathbf{P}_1, \mathbf{x}_1) (\mathbf{P}_2, \mathbf{x}_2) (\mathbf{G}((\mathbf{r}_1 \rightarrow \mathbf{Fg}_1) \wedge (\mathbf{r}_2 \rightarrow \mathbf{Fg}_2)))$. It is easy to see that $\mathcal{G}_{PS} \models \varphi$. Indeed, consider the strategy f for \mathbf{S} defined as follows. For all tracks of the form $\rho \cdot \mathbf{s}'_i$, we set a possible preemption, i.e., $f(\rho \cdot \mathbf{s}'_i) = 1$. For the tracks of the form $\rho \cdot \mathbf{s}_{1,2}$ we set the action as prescribed in the sequel: (i) if there is no occurrence of \mathbf{s}'_1 and \mathbf{s}'_2 or the last occurrence is \mathbf{s}'_1 , then we release the resource to \mathbf{P}_1 , i.e., $f(\rho \cdot \mathbf{s}_{1,2}) = 0$; (ii) if the last occurrence in ρ between \mathbf{s}'_1 and \mathbf{s}'_2 is \mathbf{s}'_1 , then then we release the resource to \mathbf{P}_2 , i.e., $f(\rho \cdot \mathbf{s}_{1,2}) = 1$.

In the next example, we show the expressive power of alternation of quantifications in SL by pointing out an example based on the ‘‘Romeo and Juliet’’ game, shown in Example 1.1.5 on page 4.

Example 1.1.15 (Quantification modalities.) First consider the CGS \mathcal{G}_{RJ} described in Example 1.1.5 on page 4. Romeo \mathbf{R} , helped by Shakespeare \mathbf{S} , wants to reach the




Figure 1.7: Winning position.

1.1. Syntax and Semantics

balcony of his loved Juliet while never meeting the Montagues family \mathbb{M} nor the Capulets family \mathbb{C} . According to the CGS, this aim can be easily expressed with the LTL formula $\psi = \mathbb{F}\mathbf{r} \wedge \mathbb{G}\neg\mathbf{c}$. The game can be played with several quantification modalities, each of them having a different meaning. For example, we can consider the formula $\varphi_1 = \llbracket \mathbf{z} \rrbracket \llbracket \mathbf{w} \rrbracket \langle \langle \mathbf{x} \rangle \rangle \langle \langle \mathbf{y} \rangle \rangle \mathfrak{b}\psi$, where $\mathfrak{b} = (\mathbb{R}, \mathbf{x})(\mathbb{S}, \mathbf{y})(\mathbb{M}, \mathbf{z})(\mathbb{C}, \mathbf{w})$. This means that both \mathbb{R} and \mathbb{S} have full visibility about the moves of the families in the house. Conversely, we can consider the formula $\varphi_2 = \langle \langle \mathbf{x} \rangle \rangle \langle \langle \mathbf{y} \rangle \rangle \llbracket \mathbf{z} \rrbracket \llbracket \mathbf{w} \rrbracket \mathfrak{b}\psi$, which means that both \mathbb{R} and \mathbb{S} have to play a strategy that is winning no matter what the families do. Moreover, one can select a modality in which the alternation of the quantification is greater than one. For instance, in the formula $\varphi_3 = \langle \langle \mathbf{x} \rangle \rangle \llbracket \mathbf{z} \rrbracket \llbracket \mathbf{w} \rrbracket \langle \langle \mathbf{y} \rangle \rangle \mathfrak{b}\psi$, Romeo has to play independently on the families while Shakespeare can select a strategy according to what the families are doing in the house. In other word, a team of agents has visibility on the adversary team that differs agent by agent, and depends on the order in which their strategies are quantified in the formula. It is worth noting that, while φ_1 and φ_2 can be represented also in ATL^* , φ_3 does not have an equivalent ATL^* formula, due to its alternation number of quantifications equal to 2.

It is easy to see that the position in Figure 1.7 on the previous page, with rooms 1 and 9 completely isolated and families standing in the remaining sector of the house composed by rooms from 2 to 8, is winning for Romeo and Shakespeare. In fact, a trivial strategy for them is given by Romeo remaining in the room 1 until Shakespeare closes the doors needed to block the families in two rooms (possibly the same) of the house, then opening a path toward the balcony for Romeo. Note that such a path always exists, since there are three different non intersecting paths from room 1 to room 9 (i.e., $1 - 2 - 3 - 6 - 9$, $1 - 4 - 7 - 8 - 9$, and $1 - 5 - 9$) and the families can occupy at most two of them.

Moreover, it is easy to see that the winning position in Figure 1.7 on the preceding page cannot be reached if Shakespeare does not have full visibility on families' moves. Indeed, starting from the initial position, Shakespeare is forced to close doors (1, 2) and (1, 4), in order to prevent Romeo to be shot at the second step, since Montagues and Capulets can move in 2 and 4, respectively. But, by closing those doors, Montagues (resp., Capulets) can move in 6 (resp., 8), being able to reach the balcony room in the next step of the game then holding there indefinitely. This makes impossible to eventually satisfy \mathbf{r} without satisfying \mathbf{c} , as well. Hence, we have that $\mathcal{G}_{RJ} \not\models \varphi_2$.

On the other hand, if Shakespeare knows where the families are moving from the initial state, then he can close the nearest doors and prevent them to reach the key rooms, i.e., 1 and 9, in the next steps. Indeed, if Montagues and Capulets are going to move in 2 and 4, then Shakespeare closes the doors (1, 2) and (1, 4), otherwise, if Montagues and Capulets are going to move in 6 and 8, then Shakespeare closes the doors (6, 9) and (7, 9). Again, the "mixed" strategies of Montagues and Capulets given by moving in 2 and 8 or 4 and 6 can be neutralized by a suitable choice for Shakespeare. Moreover, it is easy to see that, after a well played initial step, the the winning position of Figure 1.7 on the previous page is easily reachable by Romeo and Shakespeare, as the latter can close the doors (1, 5) and (5, 9) at the second step and then close the remaining doors in order to isolate the rooms 1 and 9.

Thanks to this reasoning, we have that $\mathcal{G}_{RJ} \models \varphi_3$.

Finally, we discuss another examples that is useful to point out the power of forcing two or more agents to share a strategy. Hex is a two-player game, red vs blue, in which each player in turn places a stone of his color on a single empty hexagonal cell of the rhomboidal playing board having opposite sides equally colored, either red or blue. The goal of each player is to be the first to form a path connecting the opposing sides of the board marked by his color. It is easy to prove that the stealing-strategy argument does not lead to a winning strategy in Hex, i.e., if the player that moves second copies the moves of the opponent, he surely loses the play. It is possible to formalize this fact in SL as in the following example.

Example 1.1.16 (Stealing-Strategy argument in Hex) *First model Hex with a CGS \mathcal{G}_H whose states represent every possible configuration reached during a play between Player “r” red and “b” blue. Then, verify the negation of the stealing-strategy argument by checking $\mathcal{G}_H \models \langle\langle x \rangle\rangle(r, x)(b, x)(\text{Fcnc}_x)$. Intuitively, this sentence says that agent x has a strategy that, once it is copied (bound) by b it allows the former to win, i.e., to be the first to connect the related red edges (Fcnc_x).*

It should be noted that the syntax of SL does not allow to express manipulation of strategies. For example, in the Hex game, we cannot express, for Player b , the effects of a strategy obtained by a manipulation of the one chosen by Player r . As a further example, consider a CGS having $A_c = \{0, 1\}$. Then, in SL it is not possible to describe a comparison between the outcomes that a player may obtain by replacing one of his strategies f with the associated flipping f' , i.e., $f'(\rho) = 1 - f(\rho)$, for all $\rho \in \text{dom}(f)$.

1.2 Model-Checking Hardness

In this section, we show the non-elementary lower bound for the model-checking problem of SL. Precisely, we prove that, for sentences having alternation number k , this problem is k -EXSPACE-HARD. To this aim, in Subsection 1.2.1, we first recall syntax and semantics of *Quantified propositional temporal logic* [Sis83]. Then, in Subsection 1.2.2, we give a reduction from the satisfiability problem for this logic to the model-checking problem for SL.

1.2.1 Quantified propositional temporal logic

Quantified Propositional Temporal Logic (QPTL, for short) syntactically extends the old-style temporal logic with the *future* F and *global* G operators by means of two *proposition quantifiers*, the existential $\exists q$. and the universal $\forall q$., where q is an atomic proposition. Intuitively, these elements can be respectively read as “*there exists an evaluation of q* ” and “*for all evaluations of q* ”. The formal syntax of QPTL follows.

Definition 1.2.1 (QPTL Syntax) *QPTL formulas are built inductively from the sets of atomic propositions AP, by using the following grammar, where $p \in \text{AP}$:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid X\varphi \mid F\varphi \mid G\varphi \mid \exists p.\varphi \mid \forall p.\varphi.$$

QPTL denotes the infinite set of formulas generated by the above grammar.

1.2. Model-Checking Hardness

Similarly to SL, we use the concepts of subformula, free atomic proposition, sentence, and alternation number, together with the QPTL syntactic fragment of bounded alternation $\text{QPTL}[k\text{-ALT}]$, with $k \in \mathbb{N}$.

In order to define the semantics of QPTL, we have first to introduce the concepts of truth evaluations used to interpret the meaning of atomic propositions at the passing of time. A *temporal truth evaluation* is a function $\text{tte} : \mathbb{N} \rightarrow \{\mathbf{f}, \mathbf{t}\}$ that maps each natural number to a Boolean value. Moreover, a *propositional truth evaluation* is a partial function $\text{pte} : \text{AP} \rightarrow \text{TTE}$ mapping every atomic proposition in its domain to a temporal truth evaluation. The sets $\text{TTE} \triangleq \mathbb{N} \rightarrow \{\mathbf{f}, \mathbf{t}\}$ and $\text{PTE} \triangleq \text{AP} \rightarrow \text{TTE}$ contain, respectively, all temporal and propositional truth evaluations.


At this point, we have the tool to define the interpretation of QPTL formulas. For a propositional truth evaluation pte with $\text{free}(\varphi) \subseteq \text{dom}(\text{pte})$ and a number k , we write $\text{pte}, k \models \varphi$ to indicate that the formula φ holds at the k -th position of the pte .

Definition 1.2.2 (QPTL Semantics) *For all QPTL formulas φ , propositional truth evaluation $\text{pte} \in \text{PTE}$ with $\text{free}(\varphi) \subseteq \text{dom}(\text{pte})$, and numbers $k \in \mathbb{N}$, the modeling relation $\text{pte}, k \models \varphi$ is inductively defined as follows.*

1. $\text{pte}, k \models p$ iff $\text{pte}(p)(k) = \mathbf{t}$, with $p \in \text{AP}$.
2. For all formulas φ , φ_1 , and φ_2 , it holds that:
 - (a) $\text{pte}, k \models \neg\varphi$ iff not $\text{pte}, k \models \varphi$, that is $\text{pte}, k \not\models \varphi$;
 - (b) $\text{pte}, k \models \varphi_1 \wedge \varphi_2$ iff $\text{pte}, k \models \varphi_1$ and $\text{pte}, k \models \varphi_2$;
 - (c) $\text{pte}, k \models \varphi_1 \vee \varphi_2$ iff $\text{pte}, k \models \varphi_1$ or $\text{pte}, k \models \varphi_2$;
 - (d) $\text{pte}, k \models X\varphi$ iff $\text{pte}, k + 1 \models \varphi$;
 - (e) $\text{pte}, k \models F\varphi$ iff there is an index $i \in \mathbb{N}$ with $k \leq i$ such that $\text{pte}, i \models \varphi$;
 - (f) $\text{pte}, k \models G\varphi$ iff, for all indexes $i \in \mathbb{N}$ with $k \leq i$, it holds that $\text{pte}, i \models \varphi$.
3. For an atomic proposition $q \in \text{AP}$ and a formula φ , it holds that:
 - (a) $\text{pte}, k \models \exists q.\varphi$ iff there exists a temporal truth evaluation $\text{tte} \in \text{TTE}$ such that $\text{pte}_{q \rightarrow \text{tte}}, k \models \varphi$;
 - (b) $\text{pte}, k \models \forall q.\varphi$ iff for all temporal truth evaluations $\text{tte} \in \text{TTE}$ it holds that $\text{pte}_{q \rightarrow \text{tte}}, k \models \varphi$.

Obviously, a QPTL sentence φ is *satisfiable* if $\emptyset, 0 \models \varphi$. Observe that the described semantics is slightly different but completely equivalent to that proposed and used in [SVW87] to prove the non-elementary hardness result for the satisfiability problem.

1.2.2 Non-elementary lower-bound



We can show how the solution of QPTL satisfiability problem can be reduced to that of the model-checking problem for SL, over a constant size CGS with a unique atomic proposition.

Figure 1.8: The CGS \mathcal{G}_{Rdc} .

In order to do this, we first prove the following auxiliary lemma, which actually represents the main step of the above mentioned reduction.

Lemma 1.2.1 (QPTL Reduction) *There is a one-agent CGS \mathcal{G}_{Rdc} such that, for each QPTL[k -ALT] sentence φ , with $k \in \mathbb{N}$, there exists an SL variable-closed formula $\bar{\varphi}$ such that φ is satisfiable iff $\mathcal{G}_{Rdc}, \chi, s_0 \models \bar{\varphi}$, for all complete assignments $\chi \in \text{Asg}(\text{Ag}, s_0)$.*

Proof. Consider the one-agent CGS $\mathcal{G}_{Rdc} \triangleq \langle \{p\}, \{\alpha\}, \{f, t\}, \{s_0, s_1\}, L, \text{tr}, s_0 \rangle$ depicted in Figure 1.8, where the two actions are the Boolean values false and true and where the labeling and transition functions L and tr are set as follows: $L(s_0) \triangleq \emptyset$, $L(s_1) \triangleq \{p\}$, and $\text{tr}(s, dc) = s_0$ iff $dc(\alpha) = f$, for all $s \in \text{St}$ and $dc \in \text{Dc}$. Moreover, consider the transformation function $\bar{\cdot} : \text{QPTL} \rightarrow \text{SL}$ inductively defined as follows:

- $\bar{q} \triangleq (\alpha, \mathbf{x}_q)\mathbf{X}p$, for $q \in \text{AP}$;
- $\overline{\exists q. \varphi} \triangleq \langle \langle \mathbf{x}_q \rangle \rangle \bar{\varphi}$;
- $\overline{\forall q. \varphi} \triangleq \llbracket \mathbf{x}_q \rrbracket \bar{\varphi}$;
- $\overline{\text{Op} \varphi} \triangleq \text{Op} \bar{\varphi}$, where $\text{Op} \in \{\neg, \mathbf{X}, \mathbf{F}, \mathbf{G}\}$;
- $\overline{\varphi_1 \text{Op} \varphi_2} \triangleq \bar{\varphi}_1 \text{Op} \bar{\varphi}_2$, where $\text{Op} \in \{\wedge, \vee\}$.

It is not hard to see that a QPTL formula φ is a sentence iff $\bar{\varphi}$ is variable-closed. Furthermore, we have that $\text{alt}(\bar{\varphi}) = \text{alt}(\varphi)$.

At this point, it remains to prove that, a QPTL sentence φ is satisfiable iff $\mathcal{G}_{Rdc}, \chi, s_0 \models \bar{\varphi}$, for all total assignments $\chi \in \text{Asg}(\{\alpha\}, s_0)$. To do this by induction on the structure of φ , we actually show a stronger result asserting that, for all subformulas $\psi \in \text{sub}(\varphi)$, propositional truth evaluations $\text{pte} \in \text{PTE}$, and $i \in \mathbb{N}$, it holds that $\text{pte}, i \models \psi$ iff $\mathcal{G}_{Rdc}, (\chi, s_0)^i \models \bar{\psi}$, for each total assignment $\chi \in \text{Asg}(\{\alpha\} \cup \{\mathbf{x}_q \in \text{Vr} : q \in \text{free}(\psi)\}, s_0)$ such that $\chi(\mathbf{x}_q)((\pi)_{\leq n}) = \text{pte}(q)(n)$, where $\pi \triangleq \text{play}(\chi, s_0)$, for all $q \in \text{free}(\psi)$ and $n \in [i, \omega]$.

Here, we only show the base case of atomic propositions and the two inductive cases regarding the proposition quantifiers. The remaining cases of Boolean connectives and temporal operators are straightforward and left to the reader as a simple exercise.

- $\psi = q$.

By Item 1 of Definition 1.2.2 on the preceding page of QPTL semantics, we have that $\text{pte}, i \models q$ iff $\text{pte}(q)(i) = \mathbf{t}$. Thus, due to the above constraint on the assignment, it follows that $\text{pte}, i \models q$ iff $\chi(\mathbf{x}_q)((\pi)_{\leq i}) = \mathbf{t}$. Now, by applying Items 4 and 5a of Definition 1.1.6 on page 10 of SL semantics, we have that $\mathcal{G}_{Rdc}, (\chi, s_0)^i \models (\alpha, \mathbf{x}_q)\mathbf{X}p$ iff $\mathcal{G}_{Rdc}, (\chi'[\alpha \mapsto \chi'(\mathbf{x}_q)], s')^1 \models p$, where $(\chi', s') = (\chi, s_0)^i$. At this point, due to the particular structure of the CGS \mathcal{G}_{Rdc} , we have that $\mathcal{G}_{Rdc}, (\chi'[\alpha \mapsto \chi'(\mathbf{x}_q)], s')^1 \models p$ iff $(\pi')_1 = s_1$, where $\pi' \triangleq \text{play}(\chi'[\alpha \mapsto \chi'(\mathbf{x}_q)], s')$, which in turn is equivalent to $\chi'(\mathbf{x}_q)((\pi')_{\leq 0}) = \mathbf{t}$. So, $\mathcal{G}_{Rdc}, (\chi, s_0)^i \models (\alpha, \mathbf{x}_q)\mathbf{X}p$ iff $\chi'(\mathbf{x}_q)((\pi')_{\leq 0}) = \mathbf{t}$. Now,

1.2. Model-Checking Hardness

by observing that $(\pi')_{\leq 0} = (\pi)_i$ and using the above definition of χ' , we obtain that $\chi'(\mathbf{x}_q)((\pi')_{\leq 0}) = \chi(\mathbf{x}_q)((\pi)_{\leq i})$. Hence, $\text{pte}, i \models q$ iff $\text{pte}(q)(i) = \chi(\mathbf{x}_q)((\pi)_{\leq i}) = \mathbf{t} = \chi'(\mathbf{x}_q)((\pi')_{\leq 0})$ iff $\mathcal{G}_{Rdc}, (\chi, \mathbf{s}_0)^i \models (\alpha, \mathbf{x}_q)\text{Xp}$.

- $\psi = \exists q.\psi'$.

[Only if]. If $\text{pte}, i \models \exists q.\psi'$, by Item 3a of Definition 1.2.2 on page 16, there exists a temporal truth evaluation $\text{tte} \in \text{TTE}$ such that $\text{pte}_{q \rightarrow \text{tte}}, i \models \psi'$. Now, consider a strategy $f \in \text{Str}(\mathbf{s}_0)$ such that $f((\pi)_{\leq n}) = \text{tte}(n)$, for all $n \in [i, \omega[$. Then, it is evident that $\chi^{\mathbf{x}_{q \rightarrow f}}(\mathbf{x}_{q'})((\pi)_{\leq n}) = \text{pte}_{q \rightarrow \text{tte}}(q')(n)$, for all $q' \in \text{free}(\psi)$ and $n \in [i, \omega[$. So, by the inductive hypothesis, it follows that $\mathcal{G}_{Rdc}, (\chi^{\mathbf{x}_{q \rightarrow f}}, \mathbf{s}_0)^i \models \overline{\psi'}$. Thus, we have that $\mathcal{G}_{Rdc}, (\chi, \mathbf{s}_0)^i \models \langle\langle \mathbf{x}_q \rangle\rangle \overline{\psi'}$.

[If]. If $\mathcal{G}_{Rdc}, (\chi, \mathbf{s}_0)^i \models \langle\langle \mathbf{x}_q \rangle\rangle \overline{\psi'}$, there exists a strategy $f \in \text{Str}(\mathbf{s}_0)$ such that $\mathcal{G}_{Rdc}, (\chi^{\mathbf{x}_{q \rightarrow f}}, \mathbf{s}_0)^i \models \overline{\psi'}$. Now, consider a temporal truth evaluation $\text{tte} \in \text{TTE}$ such that $\text{tte}(n) = f((\pi)_{\leq n})$, for all $n \in [i, \omega[$. Then, it is evident that $\chi^{\mathbf{x}_{q \rightarrow f}}(\mathbf{x}_{q'})((\pi)_{\leq n}) = \text{pte}_{q \rightarrow \text{tte}}(q')(n)$, for all $q' \in \text{free}(\psi)$ and $n \in [i, \omega[$. So, by the inductive hypothesis, it follows that $\text{pte}_{q \rightarrow \text{tte}}, i \models \psi'$. Thus, by Item 3a of Definition 1.2.2 on page 16, we have that $\text{pte}, i \models \exists q.\psi'$.

- $\psi = \forall q.\psi'$.

[Only if]. For each strategy $f \in \text{Str}(\mathbf{s}_0)$, consider a temporal truth evaluation $\text{tte} \in \text{TTE}$ such that $\text{tte}(n) = f((\pi)_{\leq n})$, for all $n \in [i, \omega[$. It is evident that $\chi^{\mathbf{x}_{q \rightarrow f}}(\mathbf{x}_{q'})((\pi)_{\leq n}) = \text{pte}_{q \rightarrow \text{tte}}(q')(n)$, for all $q' \in \text{free}(\psi)$ and $n \in [i, \omega[$. Now, since $\text{pte}, i \models \forall q.\psi'$, by Item 3b of Definition 1.2.2 on page 16, it follows that $\text{pte}_{q \rightarrow \text{tte}}, i \models \psi'$. So, by the inductive hypothesis, for each strategy $f \in \text{Str}(\mathbf{s}_0)$, it holds that $\mathcal{G}_{Rdc}, (\chi^{\mathbf{x}_{q \rightarrow f}}, \mathbf{s}_0)^i \models \overline{\psi'}$. Thus, we have that $\mathcal{G}_{Rdc}, (\chi, \mathbf{s}_0)^i \models \llbracket \mathbf{x}_q \rrbracket \overline{\psi'}$.

[If]. For each temporal truth evaluation $\text{tte} \in \text{TTE}$, consider a strategy $f \in \text{Str}(\mathbf{s}_0)$ such that $f((\pi)_{\leq n}) = \text{tte}(n)$, for all $n \in [i, \omega[$. It is evident that $\chi^{\mathbf{x}_{q \rightarrow f}}(\mathbf{x}_{q'})((\pi)_{\leq n}) = \text{pte}_{q \rightarrow \text{tte}}(q')(n)$, for all $q' \in \text{free}(\psi)$ and $n \in [i, \omega[$. Now, since $\mathcal{G}_{Rdc}, (\chi, \mathbf{s}_0)^i \models \llbracket \mathbf{x}_q \rrbracket \overline{\psi'}$, it follows that $\mathcal{G}_{Rdc}, (\chi^{\mathbf{x}_{q \rightarrow f}}, \mathbf{s}_0)^i \models \overline{\psi'}$. So, by the inductive hypothesis, for each temporal truth evaluation $\text{tte} \in \text{TTE}$, it holds that $\text{pte}_{q \rightarrow \text{tte}}, i \models \psi'$. Thus, by Item 3b of Definition 1.2.2 on page 16, we have that $\text{pte}, i \models \forall q.\psi'$.

Thus, we are done with the proof. \square

Now, we can show the full reduction that allows us to state the existence of a non-elementary lower-bound for the model-checking problem of SL.

Theorem 1.2.1 (SL Model-Checking Hardness) *The model-checking problem for $\text{SL}[k\text{-ALT}]$ is $k\text{-EXPSpace-HARD}$.*

Proof. Let φ be a QPTL $[k\text{-ALT}]$ sentence, $\overline{\varphi}$ the related SL $[k\text{-ALT}]$ variable-closed formula, and \mathcal{G}_{Rdc} the CGS of Lemma 1.2.1 on the previous page of QPTL reduction. Then, by applying the previous mentioned lemma, it is easy to see that φ is satisfiable iff $\mathcal{G}_{Rdc} \models \llbracket \mathbf{x} \rrbracket (\alpha, \mathbf{x}) \overline{\varphi}$ iff $\mathcal{G}_{Rdc} \models \langle\langle \mathbf{x} \rangle\rangle (\alpha, \mathbf{x}) \overline{\varphi}$. Thus, the satisfiability problem for QPTL can be reduced to the model-checking problem for SL. Now, since the satisfiability problem for QPTL $[k\text{-ALT}]$ is $k\text{-EXPSpace-HARD}$ [SVW87], we have that the model-checking problem for SL $[k\text{-ALT}]$ is $k\text{-EXPSpace-HARD}$ as well. \square

1.3 Satisfiability

In Section 1.2 it has been shown that the model-checking for SL is NONELEMENTARY-HARD. Here, we prove that the satisfiability problem is even harder, *i.e.*, undecidable. To do this, we first introduce a sentence that is satisfiable only on unbounded models. Then, by using this result, we prove the undecidability result through a reduction of the classic domino problem [Wan61].

1.3.1 Unbounded models

We now show that SL does not enjoy the bounded-tree model property. Informally, a modal logic satisfies the bounded-tree model property if, whenever a formula is satisfiable, it is so on a model in which all states have the number of successors bounded by an a priori fixed constant. Clearly, if a logic invariant under unwinding enjoys the finite model property, it enjoys the bounded-tree model property as well. The other direction may not hold, instead.

To prove this result for SL, we introduce, in the following definition, the sentence φ^{ord} to be used as a counterexample.

Definition 1.3.1 (Ordering Sentence) Let $x_1 < x_2 \triangleq \langle\langle y \rangle\rangle \varphi(x_1, x_2, y)$ be an agent-closed formula, named partial order, on the sets $AP = \{p\}$ and $Ag = \{\alpha, \beta\}$, where $\varphi(x_1, x_2, y) \triangleq ((\alpha, x_1)(\beta, y)(Xp)) \wedge ((\alpha, x_2)(\beta, y)(X\neg p))$. Then, the ordering sentence $\varphi^{ord} \triangleq \varphi^{unb} \wedge \varphi^{trn}$ is the conjunction of the following two sentences, called unboundedness and transitivity strategy requirements:

1. $\varphi^{unb} \triangleq \llbracket x_1 \rrbracket \langle\langle x_2 \rangle\rangle x_1 < x_2$;
2. $\varphi^{trn} \triangleq \llbracket x_1 \rrbracket \llbracket x_2 \rrbracket \llbracket x_3 \rrbracket (x_1 < x_2 \wedge x_2 < x_3) \rightarrow x_1 < x_3$.

Intuitively, φ^{unb} asserts that, for each strategy in x_1 , there is a different strategy in x_2 that is in relation of $<$ w.r.t. the first one, *i.e.*, $<$ has no upper bound, due to the fact that, by the definition of $\varphi(x_1, x_2, y)$, it is not reflexive. Moreover, φ^{trn} ensures that the relation $<$ is transitive too. Consequently, φ^{ord} induces a strict partial pre-order on the strategies.

Obviously, in order to be useful, the sentence φ^{ord} needs to be satisfiable, as reported in the following lemma.

Lemma 1.3.1 (Ordering Satisfiability) The sentence φ^{ord} is satisfiable.

Proof. To prove that φ^{ord} is satisfiable, consider the unbounded CGS \mathcal{G}^* , where (i) $AP \triangleq \{p\}$, (ii) $Ag \triangleq \{\alpha, \beta\}$, (iii) $Ac_{\mathcal{G}^*} \triangleq \mathbb{N}$, (iv) $St_{\mathcal{G}^*} \triangleq \{s_0, s_1, s_2\}$, (v) $s_{0\mathcal{G}^*} = s_0$, (vi) $L_{\mathcal{G}^*}(s_0) = L_{\mathcal{G}^*}(s_2) \triangleq \emptyset$ and $L_{\mathcal{G}^*}(s_1) \triangleq \{p\}$, and (vii) $tr_{\mathcal{G}^*}$ is such that if $\delta \in P \triangleq \{\delta \in Dc_{\mathcal{G}^*} : \delta(\alpha) \leq \delta(\beta)\}$ then $tr_{\mathcal{G}^*}(s_0, \delta) = s_1$ else $tr_{\mathcal{G}^*}(s_0, \delta) = s_2$, and $tr_{\mathcal{G}^*}(s, \delta) = s$, for all $s \in \{s_1, s_2\}$ and $\delta \in Dc_{\mathcal{G}^*}$ (see Figure 1.9).




Figure 1.9: The CGS \mathcal{G}^* model of φ^{ord} .

1.3. Satisfiability

Now, it is easy to see that $\mathcal{G}^* \models \varphi^{unb}$, since for every strategy f_{x_1} for x_1 , consisting of picking a natural number $n = f_{x_1}(s_0)$ as an action at the initial state, we can reply with a strategy f_{x_2} for x_2 having $f_{x_2}(s_0) > n$ and a strategy f_y for y having $f_y(s_0) = n$. Formally, we have that $\mathcal{G}^*, \chi, s_0 \models \varphi(x_1, x_2, y)$ iff $\chi(x_1)(s_0) \leq \chi(y)(s_0) < \chi(x_2)(s_0)$, for all assignments $\chi \in \text{Asg}_{\mathcal{G}^*}(\{x_1, x_2, y\}, s_0)$.

By a similar reasoning, we can see that $\mathcal{G}^* \models \varphi^{trn}$. Indeed, consider three strategies f_{x_1} , f_{x_2} , and f_{x_3} for the variables x_1 , x_2 , and x_3 , respectively, which correspond to picking three natural numbers $n_1 = f_{x_1}(s_0)$, $n_2 = f_{x_2}(s_0)$, and $n_3 = f_{x_3}(s_0)$. Now, if $\mathcal{G}^*, \chi, s_0 \models x_1 < x_2$ and $\mathcal{G}^*, \chi, s_0 \models x_2 < x_3$, for an assignments $\chi \in \text{Asg}_{\mathcal{G}^*}(\{x_1, x_2, x_3\}, s_0)$ where $\chi(x_1) = f_{x_1}$, $\chi(x_2) = f_{x_2}$, and $\chi(x_3) = f_{x_3}$, we have that $n_1 < n_2$ and $n_2 < n_3$. Consequently, $n_1 < n_3$. Hence, by using a strategy f_y for y with $f_y(s_0) = f_{x_1}(s_0)$, we have $\mathcal{G}^*, \chi_{y \mapsto f_y}, s_0 \models \varphi(x_1, x_3, y)$ and thus $\mathcal{G}^*, \chi, s_0 \models x_1 < x_3$. \square

Next lemmas report two important properties of the sentence φ^{ord} , for the negative statements we want to show. Namely, they state that, in order to be satisfied, φ^{ord} must require the existence of strict partial order relations on strategies and actions that do not admit any maximal element. From this, as stated in Theorem 1.3.1 on the next page, we directly derive that φ^{ord} needs an infinite chain of actions to be satisfied, *i.e.*, it cannot have a bounded model.

Lemma 1.3.2 (Strategy Order) *Let \mathcal{G} be a model of φ^{ord} . Moreover, let $r^< \subseteq \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}) \times \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ be a relation between $s_{0\mathcal{G}}$ -total strategies of \mathcal{G} such that $r^<(f_1, f_2)$ holds iff $\mathcal{G}, \emptyset[x_1 \mapsto f_1][x_2 \mapsto f_2], s_{0\mathcal{G}} \models x_1 < x_2$, for all strategies $f_1, f_2 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$. Then, $r^<$ is a strict partial order without maximal element.*

Proof. The proof derives from the fact that $r^<$ satisfies the following properties:

1. *Irreflexivity:* $\forall f \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}). \neg r^<(f, f)$;
2. *Unboundedness:* $\forall f_1 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}) \exists f_2 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}). r^<(f_1, f_2)$;
3. *Transitivity:* $\forall f_1, f_2, f_3 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}). (r^<(f_1, f_2) \wedge r^<(f_2, f_3)) \rightarrow r^<(f_1, f_3)$.

Indeed, Items (ii) and (iii) are directly derived from the strategy unboundedness and transitivity requirements. The proof of Item (i) derives, instead, from the following reasoning. By contradiction, suppose that $r^<$ is not a strict order, *i.e.*, there is a strategy $f \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ for which $r^<(f, f)$ holds. This means that, at the initial state $s_{0\mathcal{G}}$ of \mathcal{G} , there exists an assignment $\chi \in \text{Asg}_{\mathcal{G}}(\{x_1, x_2, y\}, s_{0\mathcal{G}})$ for which $\mathcal{G}, \chi, s_{0\mathcal{G}} \models \varphi(x_1, x_2, y)$, where $\chi(x_1) = \chi(x_2) = f$. The last fact implies the existence of a successor of $s_{0\mathcal{G}}$ in which both p and $\neg p$ hold, which is clearly impossible. \square

Lemma 1.3.3 (Action Order) *Let \mathcal{G} be a model of φ^{ord} . Moreover, let $s^< \subseteq \text{Ac}_{\mathcal{G}} \times \text{Ac}_{\mathcal{G}}$ be a relation between actions of \mathcal{G} such that $s^<(c_1, c_2)$ holds iff, for all strategies $f_1, f_2 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ with $c_1 = f_1(s_{0\mathcal{G}})$ and $c_2 = f_2(s_{0\mathcal{G}})$, it holds that $r^<(f_1, f_2)$, where $c_1, c_2 \in \text{Ac}_{\mathcal{G}}$. Then, $s^<$ is a strict partial order without maximal element.*

Proof. The irreflexivity and transitivity of $s^<$ are directly derived from the fact that, by Lemma 1.3.2 on the facing page, $r^<$ is irreflexive and transitive too. The proof of the unboundedness property derives, instead, from the following reasoning. As first thing, observe that, since the formula $x_1 < x_2$ relies on Xp and $X\neg p$ as the only temporal operators, it holds that $r^<(f_1, f_2)$ implies $r^<(f'_1, f'_2)$, for all strategies $f_1, f_2, f'_1, f'_2 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ such that $f_1(s_{0\mathcal{G}}) = f'_1(s_{0\mathcal{G}})$ and $f_2(s_{0\mathcal{G}}) = f'_2(s_{0\mathcal{G}})$. Now, suppose by contradiction that $s^<$ does not satisfy the unboundedness property, *i.e.*, there is an action $c \in \text{Ac}_{\mathcal{G}}$ such that, for all actions $c' \in \text{Ac}_{\mathcal{G}}$, it results that $s^<(c, c')$ does not hold. Then, by the definition of $s^<$ and the previous observation, we derive the existence of a strategy $f \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ with $f(s_{0\mathcal{G}}) = c$ such that $r^<(f, f')$ does not hold, for any strategy $f' \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$, which is clearly impossible. \square

Now, we have all tools to prove that SL lacks of the bounded-tree model property, which hold, instead, for several commonly used multi-agent logics, such as ATL*.

Theorem 1.3.1 (SL Unbounded Model Property) *SL does not enjoy the bounded model property.*

Proof. To prove the statement, we show that the sentence φ^{ord} of Definition 1.3.1 on page 19 cannot be satisfied on a bounded CGS. Consider a CGS \mathcal{G} such that $\mathcal{G} \models \varphi^{ord}$. The existence of such a model is ensured by Lemma 1.3.1 on page 19. Now, consider the strict partial order without maximal element between actions $s^<$ described in Lemma 1.3.3 on the preceding page. By a classical result on first order logic model theory [EF95], the relation $s^<$ cannot be defined on a finite set. Hence, $|\text{Ac}| = \infty$. \square

1.3.2 Undecidable satisfiability

We finally show the undecidability of the satisfiability problem for SL through a reduction from the *recurrent domino problem*.

The *domino problem*, proposed for the first time by Wang [Wan61], consists of placing a given number of tile types on an infinite grid, satisfying a predetermined set of constraints on adjacent tiles. One of its standard versions asks for a compatible tiling of the whole plane $\mathbb{N} \times \mathbb{N}$. The *recurrent domino problem* further requires the existence of a distinguished tile type that occurs infinitely often in the first row of the grid. This problem was proved to be highly undecidable by Harel, and in particular, Σ_1^1 -COMPLETE [Har84]. The formal definition follows.

Definition 1.3.2 (Recurrent Domino System) *An $\mathbb{N} \times \mathbb{N}$ recurrent domino system $\mathcal{D} = \langle D, H, V, t_0 \rangle$ consists of a finite non-empty set D of domino types, two horizontal and vertical matching relations $H, V \subseteq D \times D$, and a distinguished tile type $t_0 \in D$. The recurrent domino problem asks for an admissible tiling of $\mathbb{N} \times \mathbb{N}$, which is a solution mapping $\partial : \mathbb{N} \times \mathbb{N} \rightarrow D$ such that, for all $x, y \in \mathbb{N}$, it holds that (i) $(\partial(x, y), \partial(x + 1, y)) \in H$, (ii) $(\partial(x, y), \partial(x, y + 1)) \in V$, and (iii) $|\{x \in \mathbb{N} : \partial(x, 0) = t_0\}| = \omega$.*

1.3. Satisfiability

Grid specification Consider the sentence $\varphi^{grd} \triangleq \bigwedge_{a \in \text{Ag}} \varphi_a^{ord}$, where $\varphi_a^{ord} = \varphi_a^{unb} \wedge \varphi_a^{trn}$ are the *order sentences* and φ_a^{unb} and φ_a^{trn} are the *unboundedness* and *transitivity* strategy requirements for agents α and β defined, similarly to Definition 1.3.1 on page 19, as follows:

1. $\varphi_a^{unb} \triangleq \llbracket z_1 \rrbracket \langle\langle z_2 \rangle\rangle z_1 <_a z_2$;
2. $\varphi_a^{trn} \triangleq \llbracket z_1 \rrbracket \llbracket z_2 \rrbracket \llbracket z_3 \rrbracket (z_1 <_a z_2 \wedge z_2 <_a z_3) \rightarrow z_1 <_a z_3$;

where $x_1 <_\alpha x_2 \triangleq \langle\langle y \rangle\rangle \varphi_\alpha(x_1, x_2, y)$ and $y_1 <_\beta y_2 \triangleq \langle\langle x \rangle\rangle \varphi_\beta(y_1, y_2, x)$ are the two *partial order* formulas on strategies of α and β , respectively, with $\varphi_\alpha(x_1, x_2, y) \triangleq (\beta, y)((\alpha, x_1)(Xp) \wedge (\alpha, x_2)(X\neg p))$ and $\varphi_\beta(y_1, y_2, x) \triangleq (\alpha, x)((\beta, y_1)(X\neg p) \wedge (\beta, y_2)(Xp))$. Intuitively, $<_\alpha$ and $<_\beta$ correspond to the horizontal and vertical ordering of the positions in the grid, respectively.

It is easy to show that φ^{grd} is satisfiable, by using the same candidate model \mathcal{G}^* (see Figure 1.9 on page 19) and a proof argument similar to that proposed in Lemma 1.3.1 on page 19 for the simpler order sentence.

Lemma 1.3.4 (Grid Ordering Satisfiability) *The sentence φ^{grd} is satisfiable.*

Proof. Let \mathcal{G}^* be the model described in Lemma 1.3.1 on page 19. On one hand, by using the same lemma, it is evident that $\mathcal{G}^* \models \varphi_\alpha^{ord}$. On the other hand, in order to prove that $\mathcal{G}^* \models \varphi_\beta^{ord}$, first observe that $\mathcal{G}^*, \chi, s_{0\mathcal{G}^*} \models \varphi_\beta(y_1, y_2, x)$ iff $\chi(y_1)(s_{0\mathcal{G}^*}) < \chi(x)(s_{0\mathcal{G}^*}) \leq \chi(y_2)(s_{0\mathcal{G}^*})$, for all assignments $\chi \in \text{Asg}_{\mathcal{G}^*}(\{y_1, y_2, x\}, s_{0\mathcal{G}^*})$. At this point, the thesis follows by a reasoning similar to the one proposed in the lemma. \square

Consider now a model \mathcal{G} of φ^{grd} and, for all agents $a \in \text{Ag}$, the relation $r_a^< \subseteq \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}) \times \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ between $s_{0\mathcal{G}}$ -total strategies defined as follows: $r_a^<(f_1, f_2)$ holds iff $\mathcal{G}, \emptyset[z_1 \mapsto f_1][z_2 \mapsto f_2], s_{0\mathcal{G}} \models z_1 <_a z_2$, for all strategies $f_1, f_2 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$. By using a proof similar to that of Lemma 1.3.2 on page 20, it is possible to see that $r_a^<$ is a *strict partial order without maximal element* on $\text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$.

Now, to apply the desired reduction, we need to transform $r_a^<$ into a total order over strategies, by using the following two lemmas.

Lemma 1.3.5 (Strategy Equivalence) *Let $r_a^{\equiv} \subseteq \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}) \times \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$, with $a \in \text{Ag}$, be the relation between strategies such that $r_a^{\equiv}(f_1, f_2)$ holds iff neither $r_a^<(f_1, f_2)$ nor $r_a^<(f_2, f_1)$ holds, for all $f_1, f_2 \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$. Then r_a^{\equiv} is an equivalence relation.*

Proof. It is immediate to see that the relation r_a^{\equiv} is reflexive, since $r_a^<$ is not reflexive. Moreover, it is symmetric by definition. Finally, due to the definition of the partial order formula $<_a$, it is also transitive and, thus, r_a^{\equiv} is an *equivalence relation*. Indeed, if both $r_a^{\equiv}(f_1, f_2)$ and $r_a^{\equiv}(f_2, f_3)$ hold, we have that either $\mathcal{G}, \chi, s_{0\mathcal{G}} \models (\beta, y)((\alpha, x_1)(Xp) \wedge (\alpha, x_2)(Xp))$ or $\mathcal{G}, \chi, s_{0\mathcal{G}} \models (\beta, y)((\alpha, x_1)(X\neg p) \wedge (\alpha, x_2)(X\neg p))$ and either $\mathcal{G}, \chi, s_{0\mathcal{G}} \models (\beta, y)((\alpha, x_2)(Xp) \wedge (\alpha, x_3)(Xp))$ or $\mathcal{G}, \chi, s_{0\mathcal{G}} \models (\beta, y)((\alpha, x_2)(X\neg p) \wedge (\alpha, x_3)(X\neg p))$, for all assignments $\chi \in \text{Asg}(\mathcal{G}, \{x_1, x_2, x_3, y\}, s_{0\mathcal{G}})$ such that $\chi(x_1) = f_1$, $\chi(x_2) = f_2$, and $\chi(x_3) = f_3$. Hence, we also have that either $\mathcal{G}, \chi, s_{0\mathcal{G}} \models (\beta, y)((\alpha, x_1)(Xp) \wedge (\alpha, x_3)(Xp))$ or $\mathcal{G}, \chi, s_{0\mathcal{G}} \models (\beta, y)((\alpha, x_1)(X\neg p) \wedge (\alpha, x_3)(X\neg p))$. Thus, $r_a^{\equiv}(f_1, f_3)$ holds, too. The same reasoning applies to r_β^{\equiv} . \square

Let $\text{Str}_a^{\equiv} \triangleq (\text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})/r_a^{\equiv})$ be the quotient set of $\text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ w.r.t. r_a^{\equiv} , for $a \in \text{Ag}$, i.e., the set of the related equivalence classes over $s_{0\mathcal{G}}$ -total strategies. Then, the following holds.

Lemma 1.3.6 (Strategy Total Order) *Let $s_a^< \subseteq \text{Str}_a^{\equiv} \times \text{Str}_a^{\equiv}$, with $a \in \text{Ag}$, be the relation between classes of strategies such that $s_a^<(F_1, F_2)$ holds iff $r_a^<(f_1, f_2)$ holds, for all $f_1 \in F_1$, $f_2 \in F_2$, and $F_1, F_2 \in \text{Str}_a^{\equiv}$. Then $s_a^<$ is a strict total order with minimal element but no maximal element.*

Proof. The fact that $s_a^<$ is a *strict partial order without maximal element* derives directly from the same property of $r_a^<$. Indeed, due to the definition of the partial order formula $<_a$, if $r_a^{\equiv}(f', f'')$ and $r_a^<(f', f)$ (resp., $r_a^<(f, f')$) hold, we obtain that $r_a^<(f'', f)$ (resp., $r_a^<(f, f'')$) holds too. Hence, if there are $f_1 \in F_1$ and $f_2 \in F_2$ such that $r_a^<(f_1, f_2)$ holds, we directly obtain that $s_a^<(F_1, F_2)$ holds as well, for all $F_1, F_2 \in \text{Str}_a^{\equiv}$ and $a \in \text{Ag}$.

Moreover, $s_a^<$ is total, since r_a^{\equiv} is an equivalence relation that cluster together all strategies of the agent a that are not in relation *w.r.t.* either $r_a^<$ or its inverse $(r_a^<)^{-1}$. Indeed, suppose by contradiction that there are two different classes $F_1, F_2 \in \text{Str}_a^{\equiv}$ such that neither $s_a^<(F_1, F_2)$ nor $s_a^<(F_2, F_1)$ holds. This means that, for all $f_1 \in F_1$ and $f_2 \in F_2$, neither $r_a^<(f_1, f_2)$ nor $r_a^<(f_2, f_1)$ holds and, so, $r_a^{\equiv}(f_1, f_2)$. But, this contradicts the fact that F_1 and F_2 are different equivalence classes.

Finally, it is important to note that in Str_a^{\equiv} there is also a minimal element *w.r.t.* $s_a^<$. Indeed, for a strategy $f \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$ for α (resp., for β) that forces the play to reach only nodes labeled with p (resp., $\neg p$) as successor of $s_{0\mathcal{G}}$, independently from the strategy of β (resp., α), the relation $r_a^<(f', f)$ (resp., $r_a^<(f', f)$) cannot hold, for any $f' \in \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}})$. \square

By a classical result on first-order model theory [EF95], the relation $s_a^<$ cannot be defined on a finite set. Hence, $|\text{Str}_a^{\equiv}| = \omega$, for all $a \in \text{Ag}$. Now, let $s_a^< \subseteq \text{Str}_a^{\equiv} \times \text{Str}_a^{\equiv}$ be the *successor* relation on Str_a^{\equiv} compatible with the strict total order $s_a^<$, *i.e.*, such that $s_a^<(F_1, F_2)$ holds iff (i) $s_a^<(F_1, F_2)$ holds and (ii) there is no $F_3 \in \text{Str}_a^{\equiv}$ for which both $s_a^<(F_1, F_3)$ and $s_a^<(F_3, F_2)$ hold, for all $F_1, F_2 \in \text{Str}_a^{\equiv}$. Then, we can represent the two sets of classes Str_a^{\equiv} and Str_a^{\equiv} , respectively, as the infinite ordered lists $\{F_0^\alpha, F_1^\alpha, \dots\}$ and $\{F_0^\beta, F_1^\beta, \dots\}$ such that $s_a^<(F_i^\alpha, F_{i+1}^\alpha)$ holds, for all indexes $i \in \mathbb{N}$. Note that F_0^α is the class of minimal strategies *w.r.t.* the relation $s_a^<$.

At this point, we have all the machinery to build an embedding of the plane $\mathbb{N} \times \mathbb{N}$ into a model \mathcal{G} of φ^{grd} . Formally, we consider the *bijjective map* $\aleph : \mathbb{N} \times \mathbb{N} \rightarrow \text{Str}_a^{\equiv} \times \text{Str}_b^{\equiv}$ such that $\aleph(i, j) = (F_i^\alpha, F_j^\beta)$, for all $i, j \in \mathbb{N}$.

Compatible tiling Given the grid structure built on the model \mathcal{G} of φ^{grd} through the bijective map \aleph , we can express that a tiling of the grid is admissible by making use of the formula $z_1 \prec_a z_2 \triangleq z_1 <_a z_2 \wedge \neg \langle\langle z_3 \rangle\rangle z_1 <_a z_3 \wedge z_3 <_a z_2$ corresponding to the successor relation $s_a^<$, for all $a \in \text{Ag}$. Indeed, it is not hard to see that $\mathcal{G}, \chi, s_{0\mathcal{G}} \models z_1 \prec_a z_2$ iff $\chi(z_1) \in F_i^\alpha$ and $\chi(z_2) \in F_{i+1}^\alpha$, for all indexes $i \in \mathbb{N}$ and assignments $\chi \in \text{Asg}_{\mathcal{G}}(\{z_1, z_2\}, s_{0\mathcal{G}})$. The idea here is to associate with each domino type $t \in \mathcal{D}$ a corresponding atomic proposition $t \in \text{AP}$ and to express the horizontal and vertical matching conditions via suitable object labeling. In particular, we can express that the tiling is locally compatible, the horizontal neighborhoods of a tile satisfy the H or V requirements, respectively. All these constraints can be formulated through the following three agent-closed formulas:

1. $\varphi^{t,loc}(x, y) \triangleq (\alpha, x)(\beta, y)(\mathbf{X}(t \wedge \bigwedge_{t' \neq t} \neg t'))$;

1.3. Satisfiability

$$2. \varphi^{t,hor}(x, y) \triangleq \bigvee_{(t,t') \in H} \llbracket x' \rrbracket (x \prec_\alpha x' \rightarrow (\alpha, x')(\beta, y)(\mathbf{X}t'));$$

$$3. \varphi^{t,ver}(x, y) \triangleq \bigvee_{(t,t') \in V} \llbracket y' \rrbracket (y \prec_\beta y' \rightarrow (\alpha, x)(\beta, y')(\mathbf{X}t')).$$

Informally, we have the following: $\varphi^{t,loc}(x, y)$ asserts that t is the only domino type labeling the successors of the root of the model \mathcal{G} that can be reached using the strategies related to the variables x and y ; $\varphi^{t,hor}(x, y)$ asserts that the tile t' labeling the successors of the root reachable through the strategies x' and y is compatible with t w.r.t. the horizontal requirement H , for all strategies x' that immediately follow that related to x w.r.t. the order $r_\alpha^<$; $\varphi^{t,ver}(x, y)$ asserts that the tile t' labeling the successors of the root reachable through the strategies x and y' is compatible with t w.r.t. the vertical requirement V , for all strategies y' that immediately follow that related to y w.r.t. the order $r_\beta^<$.

Finally, to express that the whole grid has an admissible tiling, we use the sentence $\varphi^{til} \triangleq \llbracket x \rrbracket \llbracket y \rrbracket \bigvee_{t \in \mathcal{D}} \varphi^{t,loc}(x, y) \wedge \varphi^{t,hor}(x, y) \wedge \varphi^{t,ver}(x, y)$ that asserts the existence of a domino type t satisfying the three conditions mentioned above, for every point identified by the strategies x and y .

Recurrent tile As last task, we impose that the grid embedded into \mathcal{G} has the distinguished domino type t_0 occurring infinitely often in its first row. To do this, we describe two formulas that determine if a row or a column is the first one w.r.t. the orders $s_\alpha^<$ and $s_\beta^<$, respectively. Formally, we use $0_a(z) \triangleq \neg \langle \langle z' \rangle \rangle z' <_a z$, for $a \in \text{Ag}$. One can easily prove that $\mathcal{G}, \chi, s_{0\mathcal{G}} \models 0_a(z)$ iff $\chi(z) \in F_0^a$, for all assignments $\chi \in \text{Asg}_{\mathcal{G}}(\{z\}, s_{0\mathcal{G}})$. Now, the infinite occurrence requirement on t_0 can be expressed with the following sentence: $\varphi^{rec} \triangleq \llbracket x \rrbracket \llbracket y \rrbracket (0_\beta(y) \wedge (0_\alpha(x) \vee (\alpha, x)(\beta, y)(\mathbf{X}t_0))) \rightarrow \langle \langle x' \rangle \rangle x <_\alpha x' \wedge (\alpha, x')(\beta, y)(\mathbf{X}t_0)$. Informally, φ^{rec} asserts that, when we are on the first row identified by the variable y and at a column pointed by x such that it is the first column or the node of the “intersection” between x and y is labeled by t_0 , we have that there exists a greater column identified by x' such that its “intersection” with y is labeled by t_0 as well.

Construction correctness At this point, we have all tools to formally prove the correctness of the undecidability reduction, by showing the equivalence between the satisfiability of the sentence φ^{dom} and finding a solution of the recurrent tiling problem.

Theorem 1.3.2 (Satisfiability) *The satisfiability problem for SL is highly undecidable. In particular, it is Σ_1^1 -HARD.*

Proof. For the direct reduction, assume that there exists a solution mapping $\partial : \mathbb{N} \times \mathbb{N} \rightarrow \mathcal{D}$ for the given recurrent domino system \mathcal{D} . Then, we can build a finite CGS \mathcal{G}_∂^* similar to the one used in Lemma 1.3.1 on page 19, which satisfies the sentence φ^{dom} :




Figure 1.10: Part of the CGS \mathcal{G}_∂^* model of φ^{dom} , where $\partial(0, 0) = \mathbf{t}_1$, $\partial(0, 1) = \mathbf{t}_2$, $\partial(1, 0) = \mathbf{t}_3$, and $\partial(1, 1) = \mathbf{t}_4$.

- (i) $\text{Ac}_{\mathcal{G}_\delta^*} \triangleq \mathbb{N}$;
- (ii) there are $2 \cdot |\text{D}| + 1$ different states $\text{St}_{\mathcal{G}_\delta^*} \triangleq \{s_0\} \cup (\{p, \neg p\} \times \text{D})$ such that $\text{L}_{\mathcal{G}_\delta^*}(s_0) \triangleq \emptyset$, $\text{L}_{\mathcal{G}_\delta^*}((p, t)) \triangleq \{p, t\}$, and $\text{L}_{\mathcal{G}_\delta^*}((\neg p, t)) \triangleq \{t\}$, for all $t \in \text{D}$;
- (iii) each state $(z, t) \in \{p, \neg p\} \times \text{D}$ has only self loops $\text{tr}_{\mathcal{G}_\delta^*}((z, t), \delta) \triangleq (z, t)$ and the initial state $s_{0\mathcal{G}_\delta^*} \triangleq s_0$ is connected to (z, t) through the decision δ , *i.e.*, $\text{tr}_{\mathcal{G}_\delta^*}(s_0, \delta) \triangleq (z, t)$, iff
 - (a) $t = \partial(\delta(\alpha), \delta(\beta))$ and
 - (b) $z = p$ iff $\delta(\alpha) \leq \delta(\beta)$, for all $\delta \in \text{Dc}_{\mathcal{G}_\delta^*}$.

By a simple case analysis on the subformulas of φ^{dom} , it is possible to see that $\mathcal{G}_\delta^* \models \varphi^{\text{dom}}$.

Conversely, let \mathcal{G} be a model of the sentence φ^{dom} and $\aleph : \mathbb{N} \times \mathbb{N} \rightarrow \text{Str}_\alpha^{\equiv} \times \text{Str}_\beta^{\equiv}$ the related bijective map built for the grid specification task. As first thing, we have to prove the existence of a coloring function $\partial : \text{Str}_\alpha^{\equiv} \times \text{Str}_\beta^{\equiv} \rightarrow \text{D}$ such that, for all pairs of classes of strategies $(F^\alpha, F^\beta) \in \text{Str}_\alpha^{\equiv} \times \text{Str}_\beta^{\equiv}$ and assignments $\chi \in \text{Asg}_{\mathcal{G}}(\{\alpha, \beta\}, s_{0\mathcal{G}})$ with $\chi(\alpha) \in F^\alpha$ and $\chi(\beta) \in F^\beta$, it holds that $\mathcal{G}, \chi, s_{0\mathcal{G}} \models \text{X}\partial(F^\alpha, F^\beta)$. Then, it remains to note that the solution mapping $\partial = \partial \circ \aleph$ built as a composition of the bijective map \aleph and the coloring function ∂ is an admissible tiling of the plane $\mathbb{N} \times \mathbb{N}$.

Due to the $\varphi^{\text{t,loc}}$ formula in the sentence φ^{til} , we have that, for all assignments $\chi \in \text{Asg}_{\mathcal{G}}(\{\alpha, \beta\}, s_{0\mathcal{G}})$, there exists just one domino type $t \in \text{D}$ satisfying the property $\mathcal{G}, \chi, s_{0\mathcal{G}} \models \text{X}t$. Let $\widehat{\partial} : \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}) \times \text{Str}_{\mathcal{G}}(s_{0\mathcal{G}}) \rightarrow \text{D}$ be the function that returns such a type, for all pairs of strategies of α and β , *i.e.*, such that $\mathcal{G}, \chi, s_{0\mathcal{G}} \models \text{X}\widehat{\partial}(\chi(\alpha), \chi(\beta))$, for all assignments $\chi \in \text{Asg}_{\mathcal{G}}(\{\alpha, \beta\}, s_{0\mathcal{G}})$. Now, it is not hard to see that, due to the formulas $\varphi^{\text{t,hor}}$ and $\varphi^{\text{t,ver}}$ in the sentence φ^{til} , it holds (i) $(\widehat{\partial}(f_\alpha, f_\beta), \widehat{\partial}(f'_\alpha, f_\beta)) \in H$ and (ii) $(\widehat{\partial}(f_\alpha, f_\beta), \widehat{\partial}(f_\alpha, f'_\beta)) \in V$, for all $f_\alpha \in F_i^\alpha$, $f'_\alpha \in F_{i+1}^\alpha$, $f_\beta \in F_j^\beta$, $f'_\beta \in F_{j+1}^\beta$, and $i, j \in \mathbb{N}$. Moreover, the guess of the tile type t' adjacent to t is uniform *w.r.t.* the choice of the successor strategy. Indeed, the disjunctions $\bigvee_{(t,t') \in H}$ and $\bigvee_{(t,t') \in V}$ precede the universal quantifications $\llbracket x' \rrbracket$ and $\llbracket y' \rrbracket$ in the formulas $\varphi^{\text{t,hor}}$ and $\varphi^{\text{t,ver}}$, respectively. Thus, we have that, for all $f'_\alpha, f''_\alpha \in F_i^\alpha$ and $f'_\beta, f''_\beta \in F_j^\beta$ with $i, j \in \mathbb{N}$ and $i + j > 0$, it holds that $\widehat{\partial}(f'_\alpha, f'_\beta) = \widehat{\partial}(f''_\alpha, f''_\beta)$. Note that this fact is not necessarily true for strategies belonging to the minimal classes F_0^α and F_0^β , since the sentence φ^{dom} does not contain a relative requirement. However, every domino type $\widehat{\partial}(f_\alpha, f_\beta)$, with $f_\alpha \in F_0^\alpha$ and $f_\beta \in F_0^\beta$, can be used to label the origin of the plane $\mathbb{N} \times \mathbb{N}$ in order to obtain an admissible tiling. So, we can consider a function ∂ , defined as follows: (i) $\partial(F_0^\alpha, F_0^\beta) \in \{\widehat{\partial}(f_\alpha, f_\beta) : f_\alpha \in F_0^\alpha \wedge f_\beta \in F_0^\beta\}$; (ii) $\partial(F_i^\alpha, F_j^\beta) = \widehat{\partial}(f_\alpha, f_\beta)$, for all $f_\alpha \in F_i^\alpha$, $f_\beta \in F_j^\beta$, and $i, j \in \mathbb{N}$ with $i + j > 0$.

Clearly, (i) $(\partial(F_i^\alpha, F_j^\beta), \partial(F_{i+1}^\alpha, F_j^\beta)) \in H$, (ii) $(\partial(F_i^\alpha, F_j^\beta), \partial(F_i^\alpha, F_{j+1}^\beta)) \in V$, and (iii) $|\{i : \partial(F_i^\alpha, F_0^\beta) = t_0\}| = \omega$, for all $i, j \in \mathbb{N}$. So, $\partial = \partial \circ \aleph$ is an admissible tiling. \square

Strategy Logic Fragments

Since model checking for SL is non-elementary hard and satisfiability is even undecidable, while the same problems for ATL^* are only 2EXPTIME-COMplete , a question that arises naturally is whether there are proper fragments of SL of practical interest strictly subsuming ATL^* yet still having the same elementary complexities.

In this chapter, we answer positively to this question and go even further. More precisely, we reveal a fundamental property that, if satisfied, allows to retain a 2EXPTIME-COMplete model-checking problem. This is the property of *behavioral* quantification. Standard quantification over strategies introduces a somewhat unintuitive interpretation in the semantics of games, since it instantiates full strategies, which includes also future actions, whereas players see only actions played in the past. In more details, quantification of the form $\forall\exists$ means that the existentially quantified strategy is chosen in response to the universally quantified strategy, which prescribes actions in response to histories that may never arise in the actual game. On the other hand, in classic games player reacts only to current and past actions of other players, so the response of the existentially quantified strategy to a given history should depend only on the response of the universally quantified strategy to this history. To formally introduce the concept of behavioral quantification, we use the notion of *Skolem dependence function* as a machinery (see also [Kol85] for more on this subject). Furthermore, we show how the behavioral property can be used to apply a 2EXPTIME procedure to solve the model-checking problem. Finally, we prove the bounded model property, which is fundamental to provide a decidable procedure for the satisfiability problem.

The chapter is organized as follows. In Section 2.1, we describe three syntactic fragments of SL, named $\text{SL}[\text{NG}]$, $\text{SL}[\text{BG}]$, and $\text{SL}[\text{1G}]$, having the peculiarity to use strategy quantifications grouped in atomic blocks. Then, in Section 2.2, we define the notion of Skolem dependence function, which is used, to introduce the concept of behavioral and then the prove a fundamental result, which is at the base of our elementary model-checking procedure for $\text{SL}[\text{1G}]$. In Section 2.3, we provide an automata-based 2EXPTIME procedure to solve the model-checking problem for $\text{SL}[\text{1G}]$. Finally, in Section 2.4 we first prove the *bounded-tree model property* for $\text{SL}[\text{1G}]$ and then, by means of this, we provide a 2EXPTIME procedure for the satisfiability problem of this fragment.

2.1 Syntax

In order to formalize the syntactic fragments of SL we want to investigate, we first need to define the concepts of *quantification* and *binding prefixes*. A *quantification prefix* over a set $V \subseteq V_r$ of variables is a finite word $\varphi \in \{\langle\langle x \rangle\rangle, \llbracket x \rrbracket : x \in V\}^{|V|}$ of length $|V|$ such that each

variable $x \in V$ occurs just once in \wp , i.e., there is exactly one index $i \in [0, |V|]$ such that $(\wp)_i \in \{\langle\langle x \rangle\rangle, \llbracket x \rrbracket\}$.

A *binding prefix* over a set of variables $V \subseteq \text{Vr}$ is a finite word $b \in \{(a, x) : a \in \text{Ag} \wedge x \in V\}^{|\text{Ag}|}$ of length $|\text{Ag}|$ such that each agent $a \in \text{Ag}$ occurs just once in b , i.e., there is exactly one index $i \in [0, |\text{Ag}|]$ for which $(b)_i \in \{(a, x) : x \in V\}$.

Finally, $\text{Qnt}(V) \subseteq \{\langle\langle x \rangle\rangle, \llbracket x \rrbracket : x \in V\}^{|\text{V}|}$ and $\text{Bnd}(V) \subseteq \{(a, x) : a \in \text{Ag} \wedge x \in V\}^{|\text{Ag}|}$ denote, respectively, the sets of all quantification and binding prefixes over variables in V .

We now have all tools to define the syntactic fragments we want to analyze, which we name *Nested-Goal*, *Boolean-Goal*, and *One-Goal Strategy Logic*, respectively ($\text{SL}[\text{NG}]$, $\text{SL}[\text{BG}]$, and $\text{SL}[\text{1G}]$, for short). For a *goal* we mean an SL agent-closed formula of the kind $b\varphi$, with $\text{Ag} \subseteq \text{free}(\varphi)$, being $b \in \text{Bnd}(\text{Vr})$ a binding prefix. The idea behind $\text{SL}[\text{NG}]$ is that, when there is a quantification over a variable used in a goal, we are forced to quantify over all free variables of the inner subformula containing the goal itself, by using a quantification prefix. In this way, the subformula is build only by nesting and Boolean combinations of goals. In addition, with $\text{SL}[\text{BG}]$ we avoid nested goals sharing the variables of a same quantification prefix, but allow their Boolean combinations. Finally, $\text{SL}[\text{1G}]$ forces the use of a different quantification prefix for each single goal in the formula. The formal syntax of $\text{SL}[\text{NG}]$, $\text{SL}[\text{BG}]$, and $\text{SL}[\text{1G}]$ follows.

Definition 2.1.1 ($\text{SL}[\text{NG}]$, $\text{SL}[\text{BG}]$, and $\text{SL}[\text{1G}]$ **Syntax**) *$\text{SL}[\text{NG}]$ formulas are built inductively from the sets of atomic propositions AP , quantification prefixes $\text{Qnt}(V)$ for any $V \subseteq \text{Vr}$, and binding prefixes $\text{Bnd}(\text{Vr})$, by using the following grammar, with $p \in \text{AP}$, $\wp \in \cup_{V \subseteq \text{Vr}} \text{Qnt}(V)$, and $b \in \text{Bnd}(\text{Vr})$:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \varphi \mathbf{R}\varphi \mid \wp\varphi \mid b\varphi,$$

where in the formation rule $\wp\varphi$ it is ensured that φ is agent-closed and $\wp \in \text{Qnt}(\text{free}(\varphi))$. In addition, $\text{SL}[\text{BG}]$ formulas are determined by splitting the above syntactic class in two different parts, of which the second is dedicated to build the Boolean combinations of goals avoiding their nesting:

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \varphi \mathbf{R}\varphi \mid \wp\psi, \\ \psi &::= b\varphi \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi, \end{aligned}$$

where in the formation rule $\wp\psi$ it is ensured that $\wp \in \text{Qnt}(\text{free}(\psi))$.

Finally, the simpler $\text{SL}[\text{1G}]$ formulas are obtained by forcing each goal to be coupled with a quantification prefix:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \varphi \mathbf{R}\varphi \mid \wp b\varphi,$$

where in the formation rule $\wp b\varphi$ it is ensured that $\wp \in \text{Qnt}(\text{free}(b\varphi))$.

$\text{SL} \supset \text{SL}[\text{NG}] \supset \text{SL}[\text{BG}] \supset \text{SL}[\text{1G}]$ denotes the syntactic chain of infinite sets of formulas generated by the respective grammars with the associated constraints on free variables of goals.

Intuitively, in $\text{SL}[\text{NG}]$, $\text{SL}[\text{BG}]$, and $\text{SL}[\text{1G}]$, we force the writing of formulas to use atomic blocks of quantifications and bindings, where the related free variables are strictly coupled with those that are effectively quantified in the prefix just before the binding. In a nutshell, we can only write formulas by using sentences of the form $\wp\psi$ belonging to a kind of *prenex normal form* in which the quantifications contained into the *matrix* ψ only belong to the prefixes \wp' of some inner subsentence $\wp'\psi' \in \text{snt}(\wp\psi)$.

We want to remark that the nesting of goals in SL formulas is much more expressive than the nesting of quantification in ATL^* . Indeed, in ATL^* , we can make only a nesting of quantifications that are in some sense complete w.r.t. players, which implies, at every nesting, a total replacement of the strategy profile for all players. This is due to the fact that the quantifier operator over a set of agents A in ATL^* includes implicitly the quantification of strategies for agents not in A , i.e., each quantifier ranges over the whole set of agents, and the binding is immediately applied to the agents. Instead, in the case of SL, the quantification is referred to a single strategy and the binding with an agent can be postponed. This power of nesting allows to embed very expressive logics such as QPTL, as it is shown in Section 1.2.1

An $\text{SL}[\text{NG}]$ sentence ϕ is *principal* if it is of the form $\phi = \wp\psi$, where ψ is agent-closed and $\wp \in \text{Qnt}(\text{free}(\psi))$. By $\text{psnt}(\varphi) \subseteq \text{snt}(\varphi)$ we denote the set of all principal subsentences of the formula φ .

To start to practice with the above fragments, consider again the sentence φ of Example 1.1.11 on page 12. It is easy to see that it actually belongs to $\text{SL}[\text{BG}, 2\text{-AG}, 3\text{-VAR}, 2\text{-ALT}]$, and so, to $\text{SL}[\text{NG}]$, but not to $\text{SL}[\text{1G}]$, since it is of the form $\wp(b_1Xp \wedge b_2Xq)$, where the quantification prefix is $\wp = \langle\langle x \rangle\rangle[\langle\langle y \rangle\rangle]\langle\langle z \rangle\rangle$ and the binding prefixes of the two goals are $b_1 = (\alpha, x)(\beta, y)$ and $b_2 = (\alpha, y)(\beta, z)$.

Along the paper, sometimes we assert that a given formula φ belongs to an SL syntactic fragment also if its syntax does not precisely correspond to what is described by the relative grammar. We do this in order to make easier the reading and interpretation of the formula φ itself and only in the case that it is simple to translate it into an equivalent formula that effectively belongs to the intended logic, by means of a simple generalization of classic rules used to put a formula of first order logic in the prenex normal form. For example, consider the sentence φ_{NE} of Example 1.1.6 on page 6 representing the existence of a Nash equilibrium. This formula is considered to belong to $\text{SL}[\text{BG}, n\text{-AG}, 2n\text{-VAR}, \text{FVS}, 1\text{-ALT}]$, since it can be easily translated in the form $\phi_{NE} = \wp \bigwedge_{i=1}^n b_i\psi_i \rightarrow b\psi$, where $\wp = \langle\langle x_1 \rangle\rangle \cdots \langle\langle x_n \rangle\rangle[\langle\langle y_1 \rangle\rangle] \cdots [\langle\langle y_n \rangle\rangle]$, $b = (\alpha_1, x_1) \cdots (\alpha_n, x_n)$, $b_i = (\alpha_1, x_1) \cdots (\alpha_{i-1}, x_{i-1})(\alpha_i, y_i)(\alpha_{i+1}, x_{i+1}) \cdots (\alpha_n, x_n)$, and $\text{free}(\psi_i) = \text{Ag}$. As another example, consider the sentence φ_{SP} of Example 1.1.7 on page 7 representing the existence of a stability profile. Also this formula is considered to belong to $\text{SL}[\text{BG}, n\text{-AG}, 2n\text{-VAR}, \text{FVS}, 1\text{-ALT}]$, since it is equivalent to $\phi_{SP} = \wp \bigwedge_{i,j=1, i \neq j}^n b\psi_j \rightarrow ((b\psi_i \leftrightarrow b_i\psi_i) \rightarrow b_i\psi_j)$. Furthermore, it is immediate to see that, Example 1.1.15 on page 13, the formulas φ_1 and φ_2 are in $\text{SL}[\text{1G}, 4\text{-AG}, 4\text{-VAR}, \text{FVS}, 1\text{-ALT}]$, while φ_3 is in $\text{SL}[\text{1G}, 4\text{-AG}, 4\text{-VAR}, \text{FVS}, 2\text{-ALT}]$. Note that both ϕ_{NE} and ϕ_{SP} are principal sentences.

Now, it is interesting to observe that CTL^* and ATL^* are exactly equivalent to $\text{SL}[\text{1G}, \text{FVS}, 0\text{-ALT}]$ and $\text{SL}[\text{1G}, \text{FVS}, 1\text{-ALT}]$, respectively. Moreover, $\text{GL}[\text{AHK02}]$ is the very simple fragment of $\text{SL}[\text{BG}, \text{FVS}, 1\text{-ALT}]$ that forces all goals in a formula to have a common part containing all variables quantified before the unique possible alternation of

the quantification prefix. Finally, we have that CHP-SL is the $\text{SL}_{[\text{BG}, 2\text{-AG}, \text{FVS}]}$ fragment interpreted over turn-based CGS.

From now on, we refer to CTL^* , ATL^* , and CHP-SL formulas by denoting them with their translation in SL.

Remark 2.1.1 (SL_[NG] Model-Checking Hardness) *It is well-known that the non-elementary hardness result for the satisfiability problem of QPTL [SVW87] already holds for formulas in prenex normal form. Now, it is not hard to see that the transformation described in Lemma 1.2.1 on page 17 of QPTL reduction puts QPTL_[k-ALT] sentences φ in prenex normal form into SL_[NG, 1-AG, k-ALT] variable-closed formulas $\bar{\varphi} = \wp\psi$. Moreover, the derived SL_[1-AG, k-ALT] sentence $\langle\langle \mathbf{x} \rangle\rangle(\alpha, \mathbf{x})\wp\psi$ used in Theorem 1.2.1 on page 18 of SL model-checking hardness is equivalent to the SL_[NG, 1-AG, k-ALT] principal sentence $\langle\langle \mathbf{x} \rangle\rangle\wp(\alpha, \mathbf{x})\psi$, since \mathbf{x} is not used in the quantification prefix \wp . Thus, the hardness result for the model-checking problem holds for SL_[NG, 1-AG, k-ALT] as well. However, it is important to observe that, unfortunately, it is not known if such a hardness result holds for SL_[BG] and, in particular, for CHP-SL.*

Remark 2.1.2 (SL_[BG] Satisfiability) *In Section 1.3.2 on page 21, we prove the undecidability of the satisfiability problem for SL. Now, it is not hard to see that the formula φ^{dom} used to reduce the domino problem in Theorem 1.3.2 on page 24 actually lies in the SL_[BG] fragment. Hence, the satisfiability for this logic is undecidable too. On the other hand, later in the thesis, we prove that the same problem for SL_[1G] is 2EXPTIME-COMPLETE, thus not harder than the one for ATL^* .*

Differently from SL_[1G], in fragments like SL_[BG] we are able to associate different strategies with the same agent. This feature allows to represent game properties such as Nash Equilibria, Equilibrium Profiles, and the like. This syntactic flexibility, however, is not the feature that forces these fragments to have non-elementary complexities for the model-checking problem. Indeed, as it has been recently pointed out in [MMS13], the same problem for some of the SL fragments preserving this feature is 2EXPTIME-COMPLETE. Thus, the non-elementary hardness in the model-checking problem for SL_[NG] is due to the ability to nest the binding of strategies by means of temporal operators. A deeper study of these aspects is out of the scope of this paper.

At this point, we prove that ATL^* is strictly less expressive than SL_[1G] and, consequently, than SL_[BG] and SL_[NG]. To do this, we show the existence of two structures that result to be equivalent only *w.r.t.* sentences having alternation number bounded by 1. It can be interesting to note that, we use an ad-hoc technique based on a brute-force check to verify that all ATL^* formulas cannot distinguish between the two structures. A possible future line of research is to study variants of the Ehrenfeucht-Fraïssé game [EF95, Hod93] for SL, which would allow to determine in a more direct way whether two structures are equivalent or not *w.r.t.* a particular SL fragment.

Theorem 2.1.1 (SL_[1G] vs ATL^* Expressiveness) *There exists an SL_[1G] sentence having no ATL^* equivalent sentence.*




Figure 2.1: Alternation-2 non-equivalent CGS s.

Proof. Consider the two CGS s $\mathcal{G}_1 \triangleq \langle \{p\}, \{\alpha, \beta, \gamma\}, \{0, 1\}, \{s_0, s_1, s_2\}, L, tr_1, s_0 \rangle$ and $\mathcal{G}_2 \triangleq \langle \{p\}, \{\alpha, \beta, \gamma\}, \{0, 1, 2\}, \{s_0, s_1, s_2\}, L, tr_2, s_0 \rangle$ depicted in Figure 2.1, where $L(s_0) = L(s_2) \triangleq \emptyset$, $L(s_1) \triangleq \{p\}$, $D_1 \triangleq \{00^*, 11^*\}$, and $D_2 \triangleq \{00^*, 11^*, 12^*, 200, 202, 211\}$. Moreover, consider the $SL[1G]$ sentence $\varphi^* \triangleq \wp^* b^* Xp$, where $\wp^* \triangleq \llbracket x \rrbracket \langle y \rangle \llbracket z \rrbracket$ and $b^* \triangleq (\alpha, x)(\beta, y)(\gamma, z)$. Then, it is easy to see that $\mathcal{G}_1 \models \varphi^*$ but $\mathcal{G}_2 \not\models \varphi^*$. Indeed, $\mathcal{G}_1, \chi_1, s_0 \models b^* Xp$, for all $\chi_1 \in \text{Asg}_{\mathcal{G}_1}(\{x, y, z\}, s_0)$ such that $\chi_1(y)(s_0) = \chi_1(x)(s_0)$, and $\mathcal{G}_2, \chi_2, s_0 \models b^* X\neg p$, for all $\chi_2 \in \text{Asg}_{\mathcal{G}_2}(\{x, y, z\}, s_0)$ such that $\chi_2(x)(s_0) = 2$ and $\chi_2(z)(s_0) = (\chi_2(y)(s_0) + 1) \bmod 3$.

Now, due to the particular structure of the CGS s \mathcal{G}_i under exam, with $i \in \{1, 2\}$, for each path $\pi \in \text{Pth}_{\mathcal{G}_i}(s_0)$, we have that either $L((\pi)_j) = \{p\}$ or $L((\pi)_j) = \emptyset$, for all $j \in [1, \omega[$, i.e., apart from the initial state, the path is completely labeled either with $\{p\}$ or with \emptyset . Thus, it is easy to see that, for each ATL^* formula $\wp b \psi$, there is a literal $l_\psi \in \{p, \neg p\}$ such that $\mathcal{G}_i \models \wp b \psi$ iff $\mathcal{G}_i \models \wp b X l_\psi$, for all $i \in \{1, 2\}$. W.l.o.g., we can suppose that $b = b^*$, since we are always able to uniformly rename the variables of the quantification and binding prefixes without changing the meaning of the sentence.

At this point, it is easy to see that there exists an index $k \in \{1, 2, 3\}$ for which it holds that either $\wp_k b^* X l_\psi \Rightarrow \wp b^* X l_\psi$ or $\wp b^* X l_\psi \Rightarrow \overline{\wp} b^* X l_\psi$, where $\wp_1 \triangleq \llbracket x \rrbracket \llbracket z \rrbracket \langle y \rangle$, $\wp_2 \triangleq \langle x \rangle \langle y \rangle \llbracket z \rrbracket$, and $\wp_3 \triangleq \llbracket y \rrbracket \llbracket z \rrbracket \langle x \rangle$. Thus, to prove that every ATL^* formula cannot distinguish between \mathcal{G}_1 and \mathcal{G}_2 , we can simply show that the sentences $\wp_k b^* X l$, with $k \in \{1, 2, 3\}$ and $l \in \{p, \neg p\}$, do the same. In fact, it holds that $\mathcal{G}_i \models \wp_k b^* X l$, for all $i \in \{1, 2\}$, $k \in \{1, 2, 3\}$, and $l \in \{p, \neg p\}$. Hence, the thesis holds. The check of the latter fact is trivial and left to the reader as an exercise. \square

2.2 Behavioral property

2.2.1 Skolem Dependence Functions

We now introduce the concept of Skolem dependence function (Sdf, for short) of a quantification and show how any quantification prefix contained into an SL formula can be represented by an adequate choice of a Skolem dependence function over strategies. The main idea here is inspired by what Skolem proposed for the first order logic in order to eliminate all existential quantifications over variables, by substituting them with second order existential quantifications over functions, whose choice is uniform w.r.t. the universal variables.

First, we introduce some notation regarding quantification prefixes. Let $\wp \in \text{Qnt}(V)$ be a quantification prefix over a set $V(\wp) \triangleq V \subseteq V_r$ of variables. By $\langle\langle \wp \rangle\rangle \triangleq \{x \in V(\wp) : \exists i \in [0, |\wp|[. (\wp)_i = \langle\langle x \rangle\rangle\}$ and $\llbracket \wp \rrbracket \triangleq V(\wp) \setminus \langle\langle \wp \rangle\rangle$ we denote the sets of *existential*

and *universal variables* quantified in \wp , respectively. For two variables $x, y \in V(\wp)$, we say that x *precedes* y in \wp , in symbols $x <_{\wp} y$, if x occurs before y in \wp , i.e., there are two indexes $i, j \in [0, |\wp|[,$ with $i < j$, such that $(\wp)_i \in \{\langle\langle x \rangle\rangle, \llbracket x \rrbracket\}$ and $(\wp)_j \in \{\langle\langle y \rangle\rangle, \llbracket y \rrbracket\}$. Moreover, we say that y is *functional dependent* on x , in symbols $x \rightsquigarrow_{\wp} y$, if $x \in \llbracket \wp \rrbracket$, $y \in \langle\langle \wp \rangle\rangle$, and $x <_{\wp} y$, i.e., y is existentially quantified after that x is universally quantified, so, there may be a dependence between a value chosen by x and that chosen by y . This definition induces the set $\text{Dep}(\wp) \triangleq \{(x, y) \in V(\wp) \times V(\wp) : x \rightsquigarrow_{\wp} y\}$ of *dependence pairs* and its derived version $\text{Dep}(\wp, y) \triangleq \{x \in V(\wp) : x \rightsquigarrow_{\wp} y\}$ containing all variables from which y depends. Finally, we use $\bar{\wp} \in \text{Qnt}(V(\wp))$ to indicate the quantification derived from \wp by *dualizing* each quantifier contained in it, i.e., for all indexes $i \in [0, |\wp|[,$ it holds that $(\bar{\wp})_i = \langle\langle x \rangle\rangle$ iff $(\wp)_i = \llbracket x \rrbracket$, with $x \in V(\wp)$. It is evident that $\langle\langle \bar{\wp} \rangle\rangle = \llbracket \wp \rrbracket$ and $\llbracket \bar{\wp} \rrbracket = \langle\langle \wp \rangle\rangle$. As an example, let $\wp = \llbracket x \rrbracket \langle\langle y \rangle\rangle \langle\langle z \rangle\rangle \llbracket w \rrbracket \langle\langle v \rangle\rangle$. Then, we have $\langle\langle \wp \rangle\rangle = \{y, z, v\}$, $\llbracket \wp \rrbracket = \{x, w\}$, $\text{Dep}(\wp, x) = \text{Dep}(\wp, w) = \emptyset$, $\text{Dep}(\wp, y) = \text{Dep}(\wp, z) = \{x\}$, $\text{Dep}(\wp, v) = \{x, w\}$, and $\bar{\wp} = \langle\langle x \rangle\rangle \llbracket y \rrbracket \llbracket z \rrbracket \langle\langle w \rangle\rangle \llbracket v \rrbracket$.

Finally, we define the notion of *valuation* of variables over a generic set D , called *domain*, i.e., a partial function $v : \text{Vr} \rightarrow D$ mapping every variable in its domain to an element in D . By $\text{Val}_D(V) \triangleq V \rightarrow D$ we denote the set of all valuation functions over D defined on $V \subseteq \text{Vr}$.

At this point, we give a general high-level semantics for the quantification prefixes by means of the following main definition of *Skolem dependence function*.

Definition 2.2.1 (Skolem Dependence Function) *Let $\wp \in \text{Qnt}(V)$ be a quantification prefix over a set $V \subseteq \text{Vr}$ of variables, and D a set. Then, a Skolem dependence function for \wp over D is a function $\theta : \text{Val}_D(\llbracket \wp \rrbracket) \rightarrow \text{Val}_D(V)$ satisfying the following properties:*

1. $\theta(v) \upharpoonright_{\llbracket \wp \rrbracket} = v$, for all $v \in \text{Val}_D(\llbracket \wp \rrbracket)$; ¹
2. $\theta(v_1)(x) = \theta(v_2)(x)$, for all $v_1, v_2 \in \text{Val}_D(\llbracket \wp \rrbracket)$ and $x \in \langle\langle \wp \rangle\rangle$ such that $v_1 \upharpoonright_{\text{Dep}(\wp, x)} = v_2 \upharpoonright_{\text{Dep}(\wp, x)}$.

$\text{SDF}_D(\wp)$ denotes the set of all Skolem dependence functions for \wp over D .

Intuitively, Item 1 asserts that θ takes the same values of its argument w.r.t. the universal variables in \wp and Item 2 ensures that the value of θ w.r.t. an existential variable x in \wp does not depend on variables not in $\text{Dep}(\wp, x)$. To get a better insight into this definition, a Skolem dependence function θ for \wp can be considered as a set of classical *Skolem functions* that, given a value for each variable in $V(\wp)$ that is universally quantified in \wp , returns a possible value for all the existential variables in \wp , in a way that is consistent w.r.t. the order of quantifications. Observe that, each $\theta \in \text{SDF}_D(\wp)$ is injective, so, $|\text{rng}(\theta)| = |\text{dom}(\theta)| = |D|^{|\llbracket \wp \rrbracket|}$. Moreover, $|\text{SDF}_D(\wp)| = \prod_{x \in \langle\langle \wp \rangle\rangle} |D|^{|\text{Dep}(\wp, x)|}$. As an example, let $D = \{0, 1\}$ and $\wp = \llbracket x \rrbracket \langle\langle y \rangle\rangle \llbracket z \rrbracket \in \text{Qnt}(V)$ be a quantification prefix over $V = \{x, y, z\}$. Then, we have that $|\text{SDF}_D(\wp)| = 4$ and $|\text{SDF}_D(\bar{\wp})| = 8$. Moreover, the Skolem dependence functions $\theta_i \in \text{SDF}_D(\wp)$ with $i \in [0, 3]$ and $\bar{\theta}_i \in \text{SDF}_D(\bar{\wp})$ with $i \in [0, 7]$, for a particular fixed order, are such that $\theta_0(v)(y) = 0$, $\theta_1(v)(y) = v(x)$, $\theta_2(v)(y) = 1 - v(x)$, and $\theta_3(v)(y) = 1$,

¹By $g \upharpoonright_Z : (X \cap Z) \rightarrow Y$ we denote the *restriction* of a function $g : X \rightarrow Y$ to the elements in the set Z .

for all $\bar{v} \in \text{Val}_D(\llbracket \wp \rrbracket)$, and $\bar{\theta}_i(\bar{v})(x) = 0$ with $i \in [0, 3]$, $\bar{\theta}_i(\bar{v})(x) = 1$ with $i \in [4, 7]$, $\bar{\theta}_0(\bar{v})(z) = \bar{\theta}_4(\bar{v})(z) = 0$, $\bar{\theta}_1(\bar{v})(z) = \bar{\theta}_5(\bar{v})(z) = \bar{v}(y)$, $\bar{\theta}_2(\bar{v})(z) = \bar{\theta}_6(\bar{v})(z) = 1 - \bar{v}(y)$, and $\bar{\theta}_3(\bar{v})(z) = \bar{\theta}_7(\bar{v})(z) = 1$, for all $\bar{v} \in \text{Val}_D(\llbracket \wp \rrbracket)$.

We now prove the following fundamental theorem that describes how to eliminate the strategy quantifications of an SL formula via a choice of a suitable Skolem dependence function over strategies. This procedure can be seen as the equivalent of *Skolemization* in first order logic (see [Hod93], for more details).

Theorem 2.2.1 (SL Strategy Quantification) *Let \mathcal{G} be a CGS and $\varphi = \wp\psi$ an SL formula, being $\wp \in \text{Qnt}(V)$ a quantification prefix over a set $V \subseteq \text{free}(\psi) \cap \text{Vr}$ of variables. Then, for all assignments $\chi \in \text{Asg}(\text{free}(\varphi), s_0)$, the following holds: $\mathcal{G}, \chi, s_0 \models \varphi$ iff there exists a Skolem dependence function $\theta \in \text{SDF}_{\text{Str}(s_0)}(\wp)$ such that $\mathcal{G}, \chi \uplus \theta(\chi'), s_0 \models \psi$, for all $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$.²*

Proof. The proof proceeds by induction on the length of the quantification prefix \wp . For the base case $|\wp| = 0$, the thesis immediately follows, since $\llbracket \wp \rrbracket = \emptyset$ and, consequently, both $\text{SDF}_{\text{Str}(s_0)}(\wp)$ and $\text{Asg}(\llbracket \wp \rrbracket, s_0)$ contain the empty function only (we are assuming, by convention, that $\wp(\emptyset) \triangleq \emptyset$).

We now prove, separately, the two directions of the inductive case.

[Only if]. Suppose that $\mathcal{G}, \chi, s_0 \models \varphi$, where $\wp = \text{Qn} \cdot \wp'$. Then, two possible cases arise: either $\text{Qn} = \langle\langle x \rangle\rangle$ or $\text{Qn} = \llbracket x \rrbracket$.

- $\text{Qn} = \langle\langle x \rangle\rangle$.

By Item 3a of Definition 1.1.6 on page 10 of SL semantics, there is a strategy $f \in \text{Str}(s_0)$ such that $\mathcal{G}, \chi^{x \mapsto f}, s_0 \models \wp'\psi$. Note that $\llbracket \wp \rrbracket = \llbracket \wp' \rrbracket$. By the inductive hypothesis, we have that there exists a Skolem dependence function $\theta \in \text{SDF}_{\text{Str}(s_0)}(\wp')$ such that $\mathcal{G}, \chi^{x \mapsto f} \uplus \theta(\chi'), s_0 \models \psi$, for all $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. Now, consider the function $\theta' : \text{Asg}(\llbracket \wp \rrbracket, s_0) \rightarrow \text{Asg}(V, s_0)$ defined by $\theta'(\chi') \triangleq \theta(\chi')[x \mapsto f]$, for all $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$. It is easy to check that θ' is a Skolem dependence function for \wp over $\text{Str}(s_0)$, i.e., $\theta' \in \text{SDF}_{\text{Str}(s_0)}(\wp)$. Moreover, $\chi^{x \mapsto f} \uplus \theta(\chi') = \chi \uplus \theta(\chi')[x \mapsto f] = \chi \uplus \theta'(\chi')$, for $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$. Hence, $\mathcal{G}, \chi \uplus \theta'(\chi'), s_0 \models \psi$, for all $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$.

- $\text{Qn} = \llbracket x \rrbracket$.

By Item 3b of Definition 1.1.6 on page 10, we have that, for all strategies $f \in \text{Str}(s_0)$, it holds that $\mathcal{G}, \chi^{x \mapsto f}, s_0 \models \wp'\psi$. Note that $\llbracket \wp \rrbracket = \llbracket \wp' \rrbracket \cup \{x\}$. By the inductive hypothesis, we derive that, for each $f \in \text{Str}(s_0)$, there exists a Skolem dependence function $\theta^f \in \text{SDF}_{\text{Str}(s_0)}(\wp')$ such that $\mathcal{G}, \chi^{x \mapsto f} \uplus \theta^f(\chi'), s_0 \models \psi$, for all $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. Now, consider the function $\theta' : \text{Asg}(\llbracket \wp \rrbracket, s_0) \rightarrow \text{Asg}(V, s_0)$ defined by $\theta'(\chi') \triangleq \theta^{\chi'(x)}(\chi' \upharpoonright_{\llbracket \wp' \rrbracket})[x \mapsto \chi'(x)]$, for all $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$. It is evident that θ' is a Skolem dependence function for \wp over $\text{Str}(s_0)$, i.e., $\theta' \in \text{SDF}_{\text{Str}(s_0)}(\wp)$. Moreover, $\chi^{x \mapsto f} \uplus \theta^f(\chi') = \chi \uplus \theta^f(\chi')[x \mapsto f] = \chi \uplus \theta'(\chi'[x \mapsto f])$, for $f \in \text{Str}(s_0)$ and $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. Hence, $\mathcal{G}, \chi \uplus \theta'(\chi'), s_0 \models \psi$, for all $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$.

²By $g_1 \uplus g_2 : (X_1 \cup X_2) \rightarrow (Y_1 \cup Y_2)$ we denote the operation of *union* of two functions $g_1 : X_1 \rightarrow Y_1$ and $g_2 : X_2 \rightarrow Y_2$ defined on disjoint domains, i.e., $X_1 \cap X_2 = \emptyset$.

[If]. Suppose that there exists a Skolem dependence function $\theta \in \text{SDF}_{\text{Str}(s_0)}(\wp)$ such that $\mathcal{G}, \chi \uplus \theta(\chi'), s_0 \models \psi$, for all $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$, where $\wp = \text{Qn} \cdot \wp'$. Then, two possible cases arise: either $\text{Qn} = \langle\langle x \rangle\rangle$ or $\text{Qn} = \llbracket x \rrbracket$.

- $\text{Qn} = \langle\langle x \rangle\rangle$.

There is a strategy $f \in \text{Str}(s_0)$ such that $f = \theta(\chi')(x)$, for all $\chi' \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$. Note that $\llbracket \wp \rrbracket = \llbracket \wp' \rrbracket$. Consider the function $\theta' : \text{Asg}(\llbracket \wp' \rrbracket, s_0) \rightarrow \text{Asg}(V \setminus \{x\}, s_0)$ defined by $\theta'(\chi') \triangleq \theta(\chi') \upharpoonright_{(V \setminus \{x\})}$, for all $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. It is easy to check that θ' is a Skolem dependence function for \wp' over $\text{Str}(s_0)$, i.e., $\theta' \in \text{SDF}_{\text{Str}(s_0)}(\wp')$. Moreover, $\chi \uplus \theta(\chi') = \chi \uplus \theta'(\chi')[x \mapsto f] = \chi^{x \mapsto f} \uplus \theta'(\chi')$, for $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. Then, it is evident that $\mathcal{G}, \chi^{x \mapsto f} \uplus \theta'(\chi'), s_0 \models \psi$, for all $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. By the inductive hypothesis, we derive that $\mathcal{G}, \chi^{x \mapsto f}, s_0 \models \wp' \psi$, which means that $\mathcal{G}, \chi, s_0 \models \wp$, by Item 3a of Definition 1.1.6 on page 10 of SL semantics.

- $\text{Qn} = \llbracket x \rrbracket$.

First note that $\llbracket \wp \rrbracket = \llbracket \wp' \rrbracket \cup \{x\}$. Also, consider the functions $\theta^{f'} : \text{Asg}(\llbracket \wp' \rrbracket, s_0) \rightarrow \text{Asg}(V \setminus \{x\}, s_0)$ defined by $\theta^{f'}(\chi') \triangleq \theta(\chi'[x \mapsto f]) \upharpoonright_{(V \setminus \{x\})}$, for each $f \in \text{Str}(s_0)$ and $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. It is easy to see that every $\theta^{f'}$ is a Skolem dependence function for \wp' over $\text{Str}(s_0)$, i.e., $\theta^{f'} \in \text{SDF}_{\text{Str}(s_0)}(\wp')$. Moreover, $\chi \uplus \theta(\chi') = \chi \uplus \theta^{f'}(\chi') \upharpoonright_{\llbracket \wp' \rrbracket} [x \mapsto \chi'(x)] = \chi^{x \mapsto \chi'(x)} \uplus \theta^{f'}(\chi') \upharpoonright_{\llbracket \wp' \rrbracket}$, for $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$.

Then, it is evident that $\mathcal{G}, \chi^{x \mapsto f} \uplus \theta^{f'}(\chi'), s_0 \models \psi$, for all $f \in \text{Str}(s_0)$ and $\chi' \in \text{Asg}(\llbracket \wp' \rrbracket, s_0)$. By the inductive hypothesis, we derive that $\mathcal{G}, \chi^{x \mapsto f}, s_0 \models \wp' \psi$, for all $f \in \text{Str}(s_0)$, which means that $\mathcal{G}, \chi, s_0 \models \wp$, by Item 3b of Definition 1.1.6 on page 10.

Thus, the thesis of the theorem holds. \square

As an immediate consequence of the previous result, we derive the following corollary that restricts to SL sentences.

Corollary 2.2.1 (SL Strategy Quantification) *Let \mathcal{G} be a CGS and $\wp = \wp \psi$ an SL sentence, where ψ is agent-closed and $\wp \in \text{Qnt}(\text{free}(\psi))$. Then, $\mathcal{G} \models \wp$ iff there exists a Skolem dependence function $\theta \in \text{SDF}_{\text{Str}(s_0)}(\wp)$ such that $\mathcal{G}, \theta(\chi), s_0 \models \psi$, for all $\chi \in \text{Asg}(\llbracket \wp \rrbracket, s_0)$.*

2.2.2 Behavioral quantifications

We now have all tools we need to introduce the property of behavioral for a particular class of Skolem dependence functions. Intuitively, a Skolem dependence function over functions from a set T to a set D is behavioral if it can be split into a set of Skolem dependence functions over D, one for each element of T. This idea allows us to enormously simplify the reasoning about strategy quantifications, since we can reduce them to a set of quantifications over actions, one for each track in their domains. This means that, under certain conditions, we can transform a Skolem dependence function $\theta \in \text{SDF}_{\text{Str}(s)}(\wp)$ over strategies in a function $\tilde{\theta} : \text{Trk}(s) \rightarrow \text{SDF}_{\text{Ac}}(\wp)$ that associates with each track a Skolem dependence function over actions.

To formally develop the above idea, we have first to introduce the generic concept of adjoint function and state an auxiliary lemma.

Definition 2.2.2 (Adjoint Functions) Let D , T , U , and V be four sets, and $m : (T \rightarrow D)^U \rightarrow (T \rightarrow D)^V$ and $\tilde{m} : T \rightarrow (D^U \rightarrow D^V)$ two functions. Then, \tilde{m} is the adjoint of m if $\tilde{m}(t)(\widehat{g}(t))(x) = m(g)(x)(t)$, for all $g \in (T \rightarrow D)^U$, $x \in V$, and $t \in T$ ³

Intuitively, \tilde{m} is the adjoint of m if the dependence from the set T in both domain and codomain of the latter can be extracted and put as a common factor of the functor given by the former. This means also that, for every pair of functions $g_1, g_2 \in (T \rightarrow D)^U$ such that $\widehat{g}_1(t) = \widehat{g}_2(t)$ for some $t \in T$, it holds that $m(g_1)(x)(t) = m(g_2)(x)(t)$, for all $x \in V$. It is immediate to observe that if a function has an adjoint then this adjoint is unique. At the same way, if one has an adjoint function then it is possible to determine the original function without any ambiguity. Thus, it is established a one-to-one correspondence between functions admitting an adjoint and the adjoint itself.

Next lemma formally states the property briefly described above, i.e., that each Skolem dependence function over a set $T \rightarrow D$, admitting an adjoint function, can be represented as a function, with T as domain, which returns Skolem dependence functions over D as values.

Lemma 2.2.1 (Adjoint Skolem Dependence Functions) Let $\wp \in \text{Qnt}(V)$ be a quantification prefix over a set $V \subseteq \text{Vr}$ of variables, D and T two sets, and $\theta : \text{Val}_{T \rightarrow D}(\llbracket \wp \rrbracket) \rightarrow \text{Val}_{T \rightarrow D}(V)$ and $\tilde{\theta} : T \rightarrow (\text{Val}_D(\llbracket \wp \rrbracket) \rightarrow \text{Val}_D(V))$ two functions such that $\tilde{\theta}$ is the adjoint of θ . Then, $\theta \in \text{SDF}_{T \rightarrow D}(\wp)$ iff, for all $t \in T$, it holds that $\tilde{\theta}(t) \in \text{SDF}_D(\wp)$.

We now define the formal meaning of the behavioral of a Skolem dependence function over functions.

Definition 2.2.3 (Behavioral Skolem Dependence Functions) Let $\wp \in \text{Qnt}(V)$ be a quantification prefix over a set $V \subseteq \text{Vr}$ of variables, D and T two sets, and $\theta \in \text{SDF}_{T \rightarrow D}(\wp)$ a Skolem dependence function for \wp over $T \rightarrow D$. Then, θ is behavioral if it admits an adjoint function. $\text{BSDF}_{T \rightarrow D}(\wp)$ denotes the set of all behavioral Skolem dependence functions for \wp over $T \rightarrow D$.

It is important to observe that, unfortunately, there are Skolem dependence functions that are not behavioral. To easily understand why this is actually the case, it is enough to count both the number of Skolem dependence functions $\text{SDF}_{T \rightarrow D}(\wp)$ and those of adjoint functions $T \rightarrow \text{SDF}_D(\wp)$, where $|D| > 1$, $|T| > 1$, and \wp is such that there is an $x \in \llbracket \wp \rrbracket$ for which $\text{Dep}(\wp, x) \neq \emptyset$. Indeed, it holds that $|\text{SDF}_{T \rightarrow D}(\wp)| = \prod_{x \in \llbracket \wp \rrbracket} |D|^{|T| \cdot |D|^{|T| \cdot |\text{Dep}(\wp, x)|}} > \prod_{x \in \llbracket \wp \rrbracket} |D|^{|T| \cdot |D|^{\text{Dep}(\wp, x)|}} = |T \rightarrow \text{SDF}_D(\wp)|$. So, there are much more Skolem dependence functions, a number double exponential in $|T|$, than possible adjoint functions, whose number is only exponential in this value. Furthermore, observe that the simple set $\text{Qnt}^{\exists^* \forall^*}(V) \triangleq \{\wp \in \text{Qnt}(V) : \exists i \in [0, |\wp|] \cdot \llbracket (\wp)_{<i} \rrbracket = \emptyset \wedge \llbracket (\wp)_{\geq i} \rrbracket = \emptyset\}$, for $V \subseteq \text{Vr}$, is the maximal class of quantification prefixes that admits only behavioral Skolem dependence functions over $T \rightarrow D$, i.e., it is such that each $\theta \in \text{SDF}_{T \rightarrow D}(\wp)$ is behavioral, for all $\wp \in \text{Qnt}^{\exists^* \forall^*}(V)$. This is due to the fact that there are no functional dependences between variables, i.e., for each $x \in \llbracket \wp \rrbracket$, it holds that $\text{Dep}(\wp, x) = \emptyset$.

³By $\widehat{g} : Y \rightarrow X \rightarrow Z$ we denote the operation of *flipping* of a function $g : X \rightarrow Y \rightarrow Z$.

Finally, we introduce an important semantics for syntactic fragments of SL, which is based on the concept of behavioral Skolem dependence function over strategies, and we refer to the related satisfiability concept as *behavioral satisfiability*, in symbols \models_B . Intuitively, such a semantics has the peculiarity that a strategy used in an existential quantification, in order to satisfy a formula, is only chosen between those that are behavioral w.r.t. the universal quantifications. In this way, when we have to decide what is its value c on a given track ρ , we do it only in dependence of the values on the same track of the strategies so far quantified, but not on their whole structure, as it is the case for the classic semantics, instead. This means that c does not depend on the values of the other strategies on tracks ρ' that extend ρ , i.e., it does not depend on future choices made on ρ' . In addition, we have that c does not depend on values of parallel tracks ρ' that only share a prefix with ρ , i.e., it is independent on choices made on the possibly alternative futures ρ' . The behavioral semantics of SL[NG] formulas involving atomic propositions, Boolean connectives, temporal operators, and agent bindings is defined as for the classic one, where the modeling relation \models is substituted with \models_B , and we omit to report it here. In the following definition, we only describe the part concerning the quantification prefixes.

Definition 2.2.4 (SL[NG] Behavioral Semantics) *Let \mathcal{G} be a CGS, $s \in \text{St}$ one of its states, and $\varphi\psi$ an SL[NG] formula, where ψ is agent-closed and $\varphi \in \text{Qnt}(\text{free}(\psi))$. Then $\mathcal{G}, \emptyset, s \models_B \varphi\psi$ if there is a behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}(s)}(\varphi)$ for φ over $\text{Str}(s)$ such that $\mathcal{G}, \theta(\chi), s \models_B \psi$, for all $\chi \in \text{Asg}(\llbracket\varphi\rrbracket, s)$.*

It is immediate to see a strong similarity between the statement of Corollary 2.2.1 on page 33 of SL strategy quantification and the previous definition. The only crucial difference resides in the choice of the kind of Skolem dependence function. Moreover, observe that, differently from the classic semantics, the quantifications in the prefix are not treated individually but as an atomic block. This is due to the necessity of having a strict correlation between the point-wise structure of the quantified strategies.

Remark 2.2.1 (SL Behavioral Semantics) *It can be interesting to know that we do not define a behavioral semantics for the whole SL, since we are not able, at the moment, to easily use the concept of behavioral Skolem dependence function, when the quantifications are not necessarily grouped in prefixes, i.e., when the formula is not in prenex normal form. In fact, this may represent a challenging problem, whose solution is left to future works.*

Due to the new semantics of SL[NG], we have to redefine the related concepts of model and satisfiability, in order to differentiate between the classic relation \models and the behavioral one \models_B . Indeed, as we show later, there are sentences that are satisfiable but not behavioral satisfiable. We say that a CGS \mathcal{G} is a *behavioral model* of an SL[NG] sentence φ , in symbols $\mathcal{G} \models_B \varphi$, if $\mathcal{G}, \emptyset, s_0 \models_B \varphi$. In general, we also say that \mathcal{G} is a *behavioral model* for φ on $s \in \text{St}$, in symbols $\mathcal{G}, s \models_B \varphi$, if $\mathcal{G}, \emptyset, s \models_B \varphi$. An SL[NG] sentence φ is *behavioral satisfiable* if there is an behavioral model for it.

We have to modify the concepts of implication and equivalence, as well. Indeed, also in this case we can have pairs of equivalent formulas that are not behavioral equivalent, and vice versa. Thus, we have to be careful when we use natural transformation between

formulas, since it can be the case that they preserve the meaning only under the classic semantics. An example of this problem can arise when one want to put a formula in *pnf*. Given two $\text{SL}_{[\text{NG}]}$ formulas φ_1 and φ_2 with $\text{free}(\varphi_1) = \text{free}(\varphi_2)$, we say that φ_1 *behaviorally implies* φ_2 , in symbols $\varphi_1 \Rightarrow_{\text{B}} \varphi_2$, if, for all CGS \mathcal{G} , states $s \in \text{St}$, and $\text{free}(\varphi_1)$ -defined s -total assignments $\chi \in \text{Asg}(\text{free}(\varphi_1), s)$, it holds that if $\mathcal{G}, \chi, s \models_{\text{B}} \varphi_1$ then $\mathcal{G}, \chi, s \models_{\text{B}} \varphi_2$. Accordingly, we say that φ_1 is *behaviorally equivalent* to φ_2 , in symbols $\varphi_1 \equiv_{\text{B}} \varphi_2$, if both $\varphi_1 \Rightarrow_{\text{B}} \varphi_2$ and $\varphi_2 \Rightarrow_{\text{B}} \varphi_1$ hold.

2.2.3 Behavioral and non-behavioral semantics

Finally, we show that the introduced concept of behavioral satisfiability is relevant to the context of our logic, as its applicability represents a demarcation line between “easy” and “hard” fragments of SL. Moreover, we believe that it is because of this fundamental property that several well-known temporal logics are so robustly decidable [Var96].

Remark 2.2.2 ($\text{SL}_{[\text{NG}, 0\text{-ALT}]}$ **Behavioral**) *It is interesting to observe that, for every CGS \mathcal{G} and $\text{SL}_{[\text{NG}, 0\text{-ALT}]}$ sentence φ , it holds that $\mathcal{G} \models \varphi$ iff $\mathcal{G} \models_{\text{B}} \varphi$. This is an immediate consequence of the fact that all quantification prefixes φ used in φ belong to $\text{Qnt}^{\exists^* \forall^*}(\text{V})$, for a given set $\text{V} \subseteq \text{Vr}$ of variables. Thus, as already mentioned, the related Skolem dependence functions on strategies $\theta \in \text{SDF}_{\text{Str}(s_0)}(\varphi)$ are necessarily behavioral.*

By Corollary 2.2.1 on page 33 of SL strategy quantification, it is easy to see that the following coherence property about the behavioral of the $\text{SL}_{[\text{NG}]}$ satisfiability holds. Intuitively, it asserts that every behaviorally satisfiable sentence in *pnf* is satisfiable too.

Theorem 2.2.2 ($\text{SL}_{[\text{NG}]}$ **Behavioral Coherence**) *Let \mathcal{G} be a CGS, $s \in \text{St}$ one of its states, φ an $\text{SL}_{[\text{NG}]}$ formula in *pnf*, and $\chi \in \text{Asg}(s)$ an s -total assignment with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$. Then, it holds that $\mathcal{G}, \chi, s \models_{\text{B}} \varphi$ implies $\mathcal{G}, \chi, s \models \varphi$.*

Proof. The proof proceeds by induction on the structure of the formula. For the sake of succinctness, we only show the crucial case of principal subsentences $\phi \in \text{psnt}(\varphi)$, i.e., when ϕ is of the form $\varphi\psi$, where $\varphi \in \text{Qnt}(\text{free}(\psi))$ is a quantification prefix, and ψ is an agent-closed formula.

Suppose that $\mathcal{G}, \vartheta, s \models_{\text{B}} \varphi\psi$. Then, by Definition 2.2.4 on the previous page of $\text{SL}_{[\text{NG}]}$ behavioral semantics, there is a behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}(s)}(\varphi)$ such that, for all assignments $\chi \in \text{Asg}(\llbracket \varphi \rrbracket, s)$, it holds that $\mathcal{G}, \theta(\chi), s \models_{\text{B}} \psi$. Now, by the inductive hypothesis, there is a Skolem dependence function $\theta \in \text{SDF}_{\text{Str}(s)}(\varphi)$ such that, for all assignments $\chi \in \text{Asg}(\llbracket \varphi \rrbracket, s)$, it holds that $\mathcal{G}, \theta(\chi), s \models \psi$. Hence, by Corollary 2.2.1 on page 33 of SL strategy quantification, we have that $\mathcal{G}, \vartheta, s \models \varphi\psi$. \square

However, it is worth noting that the converse property may not hold, as we show in the next theorem, i.e., there are sentences in *pnf* that are satisfiable but not behavioral satisfiable. Note that the following results already holds for CHP-SL.

Theorem 2.2.3 ($\text{SL}_{[\text{BG}]}$ **Non-Behavioral**) *There exists a satisfiable $\text{SL}_{[\text{BG}, 1\text{-AG}, 2\text{-VAR}, 1\text{-ALT}]}$ sentence in *pnf* that is not behaviorally satisfiable.*

Proof. Consider the $\text{SL}[\text{BG}, 1\text{-AG}, 2\text{-VAR}, 1\text{-ALT}]$ sentence $\varphi \triangleq \varphi_1 \wedge \varphi_2$ in *pnf* where $\varphi_1 \triangleq \wp(\psi_1 \wedge \psi_2)$, with $\wp \triangleq \llbracket \mathbf{x} \rrbracket \langle \langle \mathbf{y} \rangle \rangle$, $\psi_1 \triangleq (\alpha, \mathbf{x})\mathbf{Xp} \leftrightarrow (\alpha, \mathbf{y})\mathbf{X}\neg\mathbf{p}$, and $\psi_2 \triangleq (\alpha, \mathbf{x})\mathbf{XXp} \leftrightarrow (\alpha, \mathbf{y})\mathbf{XXp}$, and $\varphi_2 \triangleq \llbracket \mathbf{x} \rrbracket (\alpha, \mathbf{x})\mathbf{X}(\langle \langle \mathbf{x} \rangle \rangle (\alpha, \mathbf{x})\mathbf{Xp} \wedge (\langle \langle \mathbf{x} \rangle \rangle (\alpha, \mathbf{x})\mathbf{X}\neg\mathbf{p}))$. Moreover, note that the $\text{SL}[1\text{G}, 1\text{-AG}, 1\text{-VAR}, 0\text{-ALT}]$ sentence φ_2 is equivalent to the CTL formula $\text{AX}((\text{EXp}) \wedge (\text{EX}\neg\mathbf{p}))$. Then, it is easy to see that the turn-based CGS \mathcal{G}_{Rdc} of Figure 1.8 on page 17 satisfies φ . Indeed, $\mathcal{G}_{Rdc}, \theta(\chi), \mathbf{s}_0 \models \psi_1 \wedge \psi_2$, for all assignments $\chi \in \text{Asg}(\{\mathbf{x}\}, \mathbf{s}_0)$, where the non-behavioral Skolem dependence function $\theta \in \text{SDF}_{\text{Str}(\mathbf{s}_0)}(\wp)$ is such that $\theta(\chi)(\mathbf{y})(\mathbf{s}_0) = \neg\chi(\mathbf{x})(\mathbf{s}_0)$ and $\theta(\chi)(\mathbf{y})(\mathbf{s}_0 \cdot s_i) = \chi(\mathbf{x})(\mathbf{s}_0 \cdot s_{1-i})$, for all $i \in \{0, 1\}$.

Now, let \mathcal{G} be a generic CGS. If $\mathcal{G} \not\models \varphi$, by Theorem 2.2.2 on the facing page of $\text{SL}[\text{NG}]$ behavioral coherence, it holds that $\mathcal{G} \not\models_{\text{B}} \varphi$. Otherwise, we have that $\mathcal{G} \models \varphi$ and, in particular, $\mathcal{G} \models \varphi_1$, which means that $\mathcal{G} \models \wp(\psi_1 \wedge \psi_2)$. At this point, to prove that $\mathcal{G} \not\models_{\text{B}} \varphi$, we show that, for all behavioral Skolem dependence functions $\theta \in \text{BSDF}_{\text{Str}(\mathbf{s}_0)}(\wp)$, there exists an assignment $\chi \in \text{Asg}(\{\mathbf{x}\}, \mathbf{s}_0)$ such that $\mathcal{G}, \theta(\chi), \mathbf{s}_0 \not\models_{\text{B}} \psi_1 \wedge \psi_2$. To do this, let us fix a behavioral Skolem dependence function θ and an assignment χ . Also, assume $s_1 \triangleq \text{tr}(\mathbf{s}_0, \emptyset[\alpha \mapsto \chi(\mathbf{x})(\mathbf{s}_0)])$ and $s_2 \triangleq \text{tr}(\mathbf{s}_0, \emptyset[\alpha \mapsto \theta(\chi)(\mathbf{y})(\mathbf{s}_0)])$. Now, we distinguish between two cases.

- $\mathbf{p} \in \text{L}(s_1)$ iff $\mathbf{p} \in \text{L}(s_2)$. In this case, we can easily observe that $\mathcal{G}, \theta(\chi), \mathbf{s}_0 \not\models \psi_1$ and consequently, by Theorem 2.2.2 on the preceding page, it holds that $\mathcal{G}, \theta(\chi), \mathbf{s}_0 \not\models_{\text{B}} \psi_1 \wedge \psi_2$. So, we are done.
- $\mathbf{p} \in \text{L}(s_1)$ iff $\mathbf{p} \notin \text{L}(s_2)$. If $\mathcal{G}, \theta(\chi), \mathbf{s}_0 \not\models \psi_2$ then, by Theorem 2.2.2 on the facing page, it holds that $\mathcal{G}, \theta(\chi), \mathbf{s}_0 \not\models_{\text{B}} \psi_1 \wedge \psi_2$. So, we are done. Otherwise, let $s_3 \triangleq \text{tr}(s_1, \emptyset[\alpha \mapsto \chi(\mathbf{x})(\mathbf{s}_0 \cdot s_1)])$ and $s_4 \triangleq \text{tr}(s_2, \emptyset[\alpha \mapsto \theta(\chi)(\mathbf{y})(\mathbf{s}_0 \cdot s_2)])$. Then, it holds that $\mathbf{p} \in \text{L}(s_3)$ iff $\mathbf{p} \in \text{L}(s_4)$. Now, consider a new assignment $\chi' \in \text{Asg}(\{\mathbf{x}\}, \mathbf{s}_0)$ such that $\chi'(x)(\mathbf{s}_0 \cdot s_2) = \chi(x)(\mathbf{s}_0 \cdot s_2)$ and $\mathbf{p} \in \text{L}(s_3')$ iff $\mathbf{p} \notin \text{L}(s_4)$, where $s_3' \triangleq \text{tr}(s_1, \emptyset[\alpha \mapsto \chi'(\mathbf{x})(\mathbf{s}_0 \cdot s_1)])$. Observe that the existence of such an assignment, with particular reference to the second condition, is ensured by the fact that $\mathcal{G} \models \varphi_2$. At this point, due to the behavioral of the Skolem dependence function θ , we have that $\theta(\chi')(\mathbf{y})(\mathbf{s}_0 \cdot s_2) = \theta(\chi)(\mathbf{y})(\mathbf{s}_0 \cdot s_2)$. Consequently, it holds that $s_4 = \text{tr}(s_2, \emptyset[\alpha \mapsto \theta(\chi')(\mathbf{y})(\mathbf{s}_0 \cdot s_2)])$. Thus, $\mathcal{G}, \theta(\chi'), \mathbf{s}_0 \not\models \psi_2$, which implies, by Theorem 2.2.2 on the preceding page, that $\mathcal{G}, \theta(\chi'), \mathbf{s}_0 \not\models_{\text{B}} \psi_1 \wedge \psi_2$. So, we are done.

Thus, the thesis of the theorem holds. \square

It is interesting to note that, at the moment, we do not know if such a result can be extended to the simpler GL fragment.

Remark 2.2.3 (Kinds of Non-Behavioral) *It is worth remarking that the kind of non-behavioral of the sentence φ shown in the above theorem can be called essential, i.e., it cannot be eliminated, due to the fact that φ is satisfiable but not behavioral satisfiable. However, there are different sentences, such as the conjunct φ_1 in φ , having both models on which they are behaviorally satisfiable and models, like the CGS \mathcal{G}_{Rdc} , on which they are only non-behaviorally satisfiable. Such a kind of non-behavioral can be called non-essential, since it can be eliminated by a suitable choice of the underlying model. Note that a similar reasoning can be done for the dual concept of behavioral, which we call essential if all models satisfying a given sentence behaviorally satisfy it as well.*

Before continuing, we want to show the reason why we have redefined the concepts of implication and equivalence in the context of behavioral semantics. Consider the $\text{SL}[\text{BG}, 1\text{-AG}, 2\text{-VAR}, 1\text{-ALT}]$ sentence φ_1 used in Theorem 2.2.3 on page 36 of $\text{SL}[\text{BG}]$ non-behavioral result. It is not hard to see that it is equivalent to the $\text{SL}[1\text{G}, 1\text{-AG}, 1\text{-VAR}, 0\text{-ALT}]$ $\varphi' \triangleq (\langle\langle \mathbf{x} \rangle\rangle(\alpha, \mathbf{x})\psi_1 \leftrightarrow \langle\langle \mathbf{x} \rangle\rangle(\alpha, \mathbf{x})\psi_2) \wedge (\langle\langle \mathbf{x} \rangle\rangle(\alpha, \mathbf{x})\psi_3 \leftrightarrow \langle\langle \mathbf{x} \rangle\rangle(\alpha, \mathbf{x})\psi_4)$, where $\psi_1 \triangleq \mathbf{X}(\mathbf{p} \wedge \mathbf{Xp})$, $\psi_2 \triangleq \mathbf{X}(\neg\mathbf{p} \wedge \mathbf{Xp})$, $\psi_3 \triangleq \mathbf{X}(\mathbf{p} \wedge \mathbf{X}\neg\mathbf{p})$, and $\psi_4 \triangleq \mathbf{X}(\neg\mathbf{p} \wedge \mathbf{X}\neg\mathbf{p})$. Note that φ' is in turn equivalent to the CTL^* formula $(\mathbf{E}\psi_1 \leftrightarrow \mathbf{E}\psi_2) \wedge (\mathbf{E}\psi_3 \leftrightarrow \mathbf{E}\psi_4)$. However, φ_1 and φ' are not behaviorally equivalent, since we have that $\mathcal{G}_{Rdc} \not\models_{\text{B}} \varphi_1$ but $\mathcal{G}_{Rdc} \models_{\text{B}} \varphi'$, where \mathcal{G}_{Rdc} is the CGS of Figure 1.8 on page 17.

At this point, we can proceed with the proof of the behavioral satisfiability for $\text{SL}[1\text{G}]$. It is important to note that there is no gap, in our knowledge, between the logics that are behaviorally satisfiable and those that are not, since the fragment $\text{SL}[\text{BG}, 1\text{-AG}, 2\text{-VAR}, 1\text{-ALT}]$ used in the previous theorem cannot be further reduced, due to the fact that otherwise it collapses into $\text{SL}[1\text{G}]$. Before starting, we have to describe some notation regarding classic two-player games on infinite words [PP04], which are used here as a technical tool. Note that we introduce the names of scheme and match in place of the more usual strategy and play, in order to avoid confusion between the concepts related to a CGS and those related to the tool.

A *two-player arena* (TPA, for short) is a tuple $\mathcal{A} \triangleq \langle \mathbf{N}_e, \mathbf{N}_o, \text{Ed}, n_o \rangle$, where \mathbf{N}_e and \mathbf{N}_o are non-empty non-intersecting sets of *nodes* for player *even* and *odd*, respectively, $\text{Ed} \triangleq E_e \cup E_o$, with $E_e \subseteq \mathbf{N}_e \times \mathbf{N}_o$ and $E_o \subseteq \mathbf{N}_o \times \mathbf{N}_e$, is the *edge relation* between nodes, and $n_o \in \mathbf{N}_o$ is a designated *initial node*.

An *even position* in \mathcal{A} is a finite non-empty sequence of nodes $v \in \mathbf{N}_e^+$ such that $(v)_0 = n_o$ and, for all $i \in [0, |v| - 1[$, there exists a node $n \in \mathbf{N}_o$ for which $((v)_i, n) \in E_e$ and $(n, (v)_{i+1}) \in E_o$ hold. In addition, an *odd position* in \mathcal{A} is a finite non-empty sequence of nodes $v = v' \cdot n \in \mathbf{N}_e^+ \cdot \mathbf{N}_o$, with $n \in \mathbf{N}_o$, such that v' is an even position and $(\text{lst}(v'), n) \in E_e$. By Pos_e and Pos_o we denote, respectively, the sets of even and odd positions.

An *even* (resp., *odd*) *scheme* in \mathcal{A} is a function $s_e : \text{Pos}_e \rightarrow \mathbf{N}_o$ (resp., $s_o : \text{Pos}_o \rightarrow \mathbf{N}_e$) that maps each even (resp., odd) position to an odd (resp., even) node in a way that is compatible with the edge relation E_e (resp., E_o), i.e., for all $v \in \text{Pos}_e$ (resp., $v \in \text{Pos}_o$), it holds that $(\text{lst}(v), s_e(v)) \in E_e$ (resp., $(\text{lst}(v), s_o(v)) \in E_o$). By Sch_e (resp., Sch_o) we indicate the sets of even (resp., odd) schemes.

A *match* in \mathcal{A} is an infinite sequence of nodes $\varpi \in \mathbf{N}_e^\omega$ such that $(\varpi)_0 = n_o$ and, for all $i \in \mathbb{N}$, there exists a node $n \in \mathbf{N}_o$ such that $((\varpi)_i, n) \in E_e$ and $(n, (\varpi)_{i+1}) \in E_o$. By Mtc we denote the set of all matches. A *match map* $\text{mtc} : \text{Sch}_e \times \text{Sch}_o \rightarrow \text{Mtc}$ is a function that, given two schemes $s_e \in \text{Sch}_e$ and $s_o \in \text{Sch}_o$, returns the unique match $\varpi = \text{mtc}(s_e, s_o)$ such that, for all $i \in \mathbb{N}$, it holds that $(\varpi)_{i+1} = s_o((\varpi)_{\leq i} \cdot s_e((\varpi)_{\leq i}))$.

A *two-player game* (TPG, for short) is a tuple $\mathcal{H} \triangleq \langle \mathcal{A}, \text{Wn} \rangle$, where \mathcal{A} is a TPA and $\text{Wn} \subseteq \text{Mtc}$. On one hand, we say that player even wins \mathcal{H} if there exists an even scheme $s_e \in \text{Sch}_e$ such that, for all odd schemes $s_o \in \text{Sch}_o$, it holds that $\text{mtc}(s_e, s_o) \in \text{Wn}$. On the other hand, we say that player odd wins \mathcal{H} if there exists an odd scheme $s_o \in \text{Sch}_o$ such that, for all even schemes $s_e \in \text{Sch}_e$, it holds that $\text{mtc}(s_e, s_o) \notin \text{Wn}$.

In the following, for a given binding prefix $b \in \text{Bn}(V)$ with $V \subseteq \text{Vr}$, we denote by $\text{bnd}_b : \text{Ag} \rightarrow V$ the function associating with each agent the related variable in b , i.e., for all $a \in \text{Ag}$, there is $i \in [0, |b|[$ such that $(b)_i = (a, \text{bnd}_b(a))$.

As first step towards the proof of the behavioral of $SL[1G]$, we have to give a construction of a two-player game, based on an a priori chosen CGS, in which the players are explicitly viewed one as a Skolem dependence function and the other as a valuation, both over actions. This construction results to be a deep technical evolution of the proof method used for the dualization of alternating automata on infinite objects [MS87].

Definition 2.2.5 (Dependence-vs-Valuation Game) *Let \mathcal{G} be a CGS, $s \in \text{St}$ one of its states, $P \subseteq \text{Pth}(s)$ a set of paths, $\wp \in \text{Qnt}(V)$ a quantification prefix over a set $V \subseteq \text{Vr}$ of variables, and $\flat \in \text{Bn}(V)$ a binding. Then, the dependence-vs-valuation game for \mathcal{G} in s over P w.r.t. \wp and \flat is the TPG $\mathcal{H}(\mathcal{G}, s, P, \wp, \flat) \triangleq \langle \mathcal{A}(\mathcal{G}, s, \wp, \flat), P \rangle$, where the TPA $\mathcal{A}(\mathcal{G}, s, \wp, \flat) \triangleq \langle \text{St}, \text{St} \times \text{SDF}_{\text{Ac}}(\wp), \text{Ed}, s \rangle$ has the edge relations defined as follows:*

- $E_e \triangleq \{(t, (t, \theta)) : t \in \text{St} \wedge \theta \in \text{SDF}_{\text{Ac}}(\wp)\};$
- $E_o \triangleq \{((t, \theta), \text{tr}(t, \theta(\mathbf{v}) \circ \text{bnd}_{\flat})) : t \in \text{St} \wedge \theta \in \text{SDF}_{\text{Ac}}(\wp) \wedge \mathbf{v} \in \text{Val}_{\text{Ac}}(\llbracket \wp \rrbracket)\}^4.$

In the next lemma we state a fundamental relationship between dependence-vs-valuation games and their duals. Basically, we prove that if a player wins the game then the opposite player can win the dual game, and vice versa. This represents one of the two crucial steps in our behavioral proof.

Lemma 2.2.2 (Dependence-vs-Valuation Duality) *Let \mathcal{G} be a CGS, $s \in \text{St}$ one of its states, $P \subseteq \text{Pth}(s)$ a set of paths, $\wp \in \text{Qnt}(V)$ a quantification prefix over a set $V \subseteq \text{Vr}$ of variables, and $\flat \in \text{Bn}(V)$ a binding. Then, player even wins the TPG $\mathcal{H}(\mathcal{G}, s, P, \wp, \flat)$ iff player odd wins the dual TPG $\mathcal{H}(\mathcal{G}, s, \text{Pth}(s) \setminus P, \bar{\wp}, \bar{\flat})$.*

Now, we are going to give the definition of the important concept of *encasement*. Informally, an encasement is a particular subset of paths in a given CGS that “works to encase” a behavioral Skolem dependence function on strategies, in the sense that it contains all plays obtainable by complete assignments derived from the evaluation of the above mentioned Skolem dependence function. In our context, this concept is used to summarize all relevant information needed to verify the behavioral satisfiability of a sentence.

Definition 2.2.6 (Encasement) *Let \mathcal{G} be a CGS, $s \in \text{St}$ one of its states, $P \subseteq \text{Pth}(s)$ a set of paths, $\wp \in \text{Qnt}(V)$ a quantification prefix over a set $V \subseteq \text{Vr}$ of variables, and $\flat \in \text{Bn}(V)$ a binding. Then, P is an encasement w.r.t. \wp and \flat if there exists a behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}(s)}(\wp)$ such that, for all assignments $\chi \in \text{Asg}(\llbracket \wp \rrbracket, s)$, it holds that $\text{play}(\theta(\chi) \circ \text{bnd}_{\flat}, s) \in P$.*

In the next lemma, we give the second of the two crucial steps in our behavioral proof. In particular, we are able to show a one-to-one relationship between the winning in the dependence-vs-valuation game of player even and the verification of the encasement property of the associated winning set. Moreover, in the case that the latter is a Borelian set, by using Martin’s Determinacy Theorem [Mar75], we obtain a complete characterization of the winning concept by means of that of encasements.

⁴By $g_2 \circ g_1 : X \rightarrow Z$ we denote the operation of *composition* of two functions $g_1 : X \rightarrow Y_1$ and $g_2 : Y_2 \rightarrow Z$ with $Y_1 \subseteq Y_2$.

Lemma 2.2.3 (Encasement Characterization) *Let \mathcal{G} be a CGS, $s \in \text{St}$ one of its states, $P \subseteq \text{Pth}(s)$ a set of paths, $\wp \in \text{Qnt}(V)$ a quantification prefix over a set $V \subseteq \text{Vr}$ of variables, and $b \in \text{Bn}(V)$ a binding. Then, the following hold:*

- (i) *player even wins $\mathcal{H}(\mathcal{G}, s, P, \wp, b)$ iff P is an encasement w.r.t. \wp and b ;*
- (ii) *if player odd wins $\mathcal{H}(\mathcal{G}, s, P, \wp, b)$ then P is not an encasement w.r.t. \wp and b ;*
- (iii) *if P is a Borelian set and it is not an encasement w.r.t. \wp and b then player odd wins $\mathcal{H}(\mathcal{G}, s, P, \wp, b)$.*

Finally, we have all technical tools useful to prove the behavioral of the satisfiability for $\text{SL}[1G]$. Intuitively, we describe a bidirectional reduction of the problem of interest to the verification of the winning in the dependence-vs-valuation game. The idea behind this construction resides in the strong similarity between the statement of Corollary 2.2.1 on page 33 of SL strategy quantification and the definition of the winning condition in a two-player game. Indeed, on one hand, we say that a sentence is satisfiable iff “there exists a Skolem dependence function such that, for all all assignments, it holds that ...”. On the other hand, we say that player even wins a game iff “there exists an even scheme such that, for all odd schemes, it holds that ...”. In particular, for the $\text{SL}[1G]$ fragment, we can resolve the gap between these two formulations, by using the concept of behavioral quantification.

Theorem 2.2.4 (SL[1G] Behavioral) *Let \mathcal{G} be a CGS, φ an SL[1G] formula, $s \in \text{St}$ a state, and $\chi \in \text{Asg}(s)$ an s -total assignment with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$. Then, it holds that $\mathcal{G}, \chi, s \models \varphi$ iff $\mathcal{G}, \chi, s \models_{\text{B}} \varphi$.*

Proof. The proof proceeds by induction on the structure of the formula. For the sake of succinctness, we only show the most important inductive case of principal subsentences $\phi \in \text{psnt}(\varphi)$, i.e., when ϕ is of the form $\wp b \psi$, where $\wp \in \text{Qnt}(V)$ and $b \in \text{Bnd}(V)$ are, respectively, a quantification and binding prefix over a set $V \subseteq \text{Vr}$ of variables, and ψ is a variable-closed formula.

[If]. The proof of this direction is practically the same of the one used in Theorem 2.2.2 on page 36 of $\text{SL}[\text{NG}]$ behavioral coherence. So, we omit to report it here.

[Only if]. Assume that $\mathcal{G}, \emptyset, s \models \wp b \psi$. Then, it is easy to see that, for all behavioral Skolem dependence functions $\theta \in \text{BSDF}_{\text{Str}(s)}(\overline{\wp})$, there is an assignment $\chi \in \text{Asg}(\llbracket \overline{\wp} \rrbracket, s)$ such that $\mathcal{G}, \theta(\chi) \circ \text{bnd}_b, s \models \psi$. Indeed, suppose by contradiction that there exists a behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}(s)}(\overline{\wp})$ such that, for all assignments $\chi \in \text{Asg}(\llbracket \overline{\wp} \rrbracket, s)$, it holds that $\mathcal{G}, \theta(\chi) \circ \text{bnd}_b, s \not\models \psi$, i.e., $\mathcal{G}, \theta(\chi) \circ \text{bnd}_b, s \models \neg \psi$, and so $\mathcal{G}, \theta(\chi), s \models b \neg \psi$. Then, by Corollary 2.2.1 on page 33 of SL strategy quantification, we have that $\mathcal{G}, \emptyset, s \models \overline{\wp} b \neg \psi$, i.e., $\mathcal{G}, \emptyset, s \models \neg \wp b \psi$, and so $\mathcal{G}, \emptyset, s \not\models \wp b \psi$, which is impossible.

Now, let $P \triangleq \{\text{play}(\chi, s) \in \text{Pth}(s) : \chi \in \text{Asg}(\text{Ag}, s) \wedge \mathcal{G}, \chi, s \not\models \psi\}$. Then, it is evident that, for all behavioral Skolem dependence functions $\theta \in \text{BSDF}_{\text{Str}(s)}(\overline{\wp})$, there is an assignment $\chi \in \text{Asg}(\llbracket \overline{\wp} \rrbracket, s)$ such that $\text{play}(\theta(\chi) \circ \text{bnd}_b, s) \notin P$.

At this point, by Definition 2.2.6 on the previous page of encasement, it is clear that P is not an encasement w.r.t. $\overline{\wp}$ and b . Moreover, since ψ describes a regular language, the derived set P is Borelian [PP04]. Consequently, by Item iii of Lemma 2.2.3 of encasement characterization,

we have that player odd wins the TPG $\mathcal{H}(\mathcal{G}, s, P, \bar{\varphi}, b)$. Thus, by Lemma 2.2.2 on page 39 of dependence-vs-valuation duality, player even wins the dual TPG $\mathcal{H}(\mathcal{G}, s, \text{Pth}(s) \setminus P, \varphi, b)$. Hence, by Item i of Lemma 2.2.3 on the facing page, we have that $\text{Pth}(s) \setminus P$ is an encasement w.r.t. φ and b . Finally, again by Definition 2.2.6 on page 39, there exists an behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}(s)}(\varphi)$ such that, for all assignments $\chi \in \text{Asg}(\llbracket \varphi \rrbracket, s)$, it holds that $\text{play}(\theta(\chi) \circ \text{bnd}_b, s) \in \text{Pth}(s) \setminus P$.

Now, it is immediate to observe that $\text{Pth}(s) \setminus P = \{\text{play}(\chi, s) \in \text{Pth}(s) : \chi \in \text{Asg}(\text{Ag}, s) \wedge \mathcal{G}, \chi, s \models \psi\}$. So, by the inductive hypothesis, we have that $\text{Pth}(s) \setminus P = \{\text{play}(\chi, s) \in \text{Pth}(s) : \chi \in \text{Asg}(\text{Ag}, s) \wedge \mathcal{G}, \chi, s \models_{\text{B}} \psi\}$, from which we derive that there exists a behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}(s)}(\varphi)$ such that, for all assignments $\chi \in \text{Asg}(\llbracket \varphi \rrbracket, s)$, it holds that $\mathcal{G}, \theta(\chi) \circ \text{bnd}_b, s \models_{\text{B}} \psi$. Consequently, by Definition 2.2.4 on page 35 of $\text{SL}_{[\text{NG}]}$ behavioral semantics, we have that $\mathcal{G}, \emptyset, s \models_{\text{B}} \varphi$. \square

As an immediate consequence of the previous theorem, we derive the following fundamental corollary, restricted to $\text{SL}_{[1\text{G}]}$ sentences.

Corollary 2.2.2 (SL_[1G] Behavioral) *Let \mathcal{G} be a CGS and φ an SL_[1G] sentence. Then, $\mathcal{G} \models \varphi$ iff $\mathcal{G} \models_{\text{B}} \varphi$.*

It is worth to observe that the behavioral property for $\text{SL}_{[1\text{G}]}$ is a crucial difference w.r.t. $\text{SL}_{[\text{BG}]}$, which allows us to obtain a behavioral decision procedure for the related model-checking problem, as described in the last part of the next section.

2.3 Model Checking

In this section, we study the model-checking problem for SL and show that, in general, it is non-elementarily decidable, while, in the particular case of $\text{SL}_{[1\text{G}]}$ sentences, it is just 2EXPTIME-COMplete, as for ATL^* . For the algorithmic procedures, we follow an *automata-theoretic approach* [KVV00], reducing the decision problem for the logics to the emptiness problem of an automaton. In particular, we use a bottom-up technique through which we recursively label each state of the CGS of interest by all principal subsentences of the specification that are satisfied on it, starting from the innermost subsentences and terminating with the sentence under exam. In this way, at a given step of the recursion, since the satisfaction of all subsentences of the given principal sentence has already been determined, we can assume that the matrix of the latter is only composed by Boolean combinations and nesting of goals whose temporal part is simply LTL. The procedure we propose here extends that used for ATL^* in [AHK02] by means of a richer structure of the automata involved in.

The rest of this section is organized as follows. In Subsection 2.3.1, we recall the definition of alternating parity tree automata. Then, in Subsection 2.3.2, we build an automaton accepting a tree encoding of a CGS iff this satisfies the formula of interest, which is used to prove the main result about SL and $\text{SL}_{[\text{NG}]}$ model checking. Finally, in Subsection 2.3.3, we refine the previous result to obtain an elementary decision procedure for $\text{SL}_{[1\text{G}]}$.

2.3.1 Alternating tree automata

Nondeterministic tree automata are a generalization to infinite trees of the classical *nondeterministic word automata* on infinite words (see [Tho90], for an introduction). *Alternating tree automata* are a further generalization of nondeterministic tree automata [MS87]. Intuitively, on visiting a node of the input tree, while the latter sends exactly one copy of itself to each of the successors of the node, the former can send several own copies to the same successor. Here we use, in particular, *alternating parity tree automata*, which are alternating tree automata along with a *parity acceptance condition* (see [GTW02], for a survey).

We now give the formal definition of alternating tree automata.

Definition 2.3.1 (Alternating Tree Automata) *An alternating tree automaton (ATA, for short) is a tuple $\mathcal{A} \triangleq \langle \Sigma, \Delta, Q, \delta, q_0, \aleph \rangle$, where Σ , Dir , and Q are, respectively, non-empty finite sets of input symbols, directions, and states, $q_0 \in Q$ is an initial state, \aleph is an acceptance condition to be defined later, and $\delta : Q \times \Sigma \rightarrow \mathbb{B}^+(\text{Dir} \times Q)$ is an alternating transition function that maps each pair of states and input symbols to a positive Boolean combination on the set of propositions of the form $(d, q) \in \text{Dir} \times Q$, a.k.a. moves.*

On one side, a *nondeterministic tree automaton* (NTA, for short) is a special case of ATA in which each conjunction in the transition function δ has exactly one move (d, q) associated with each direction d . This means that, for all states $q \in Q$ and symbols $\sigma \in \Sigma$, we have that $\delta(q, \sigma)$ is equivalent to a Boolean formula of the form $\bigvee_i \bigwedge_{d \in \text{Dir}} (d, q_{i,d})$. On the other side, a *universal tree automaton* (UTA, for short) is a special case of ATA in which all the Boolean combinations that appear in δ are conjunctions of moves. Thus, we have that $\delta(q, \sigma) = \bigwedge_i (d_i, q_i)$, for all states $q \in Q$ and symbols $\sigma \in \Sigma$.

The semantics of the ATAs is given through the following concept of run.

Definition 2.3.2 (ATA Run) *A run of an ATA $\mathcal{A} = \langle \Sigma, \Delta, Q, \delta, q_0, \aleph \rangle$ on a Σ -labeled Dir-tree $\mathcal{T} = \langle T, \nu \rangle$ is a $(\text{Dir} \times Q)$ -tree R such that, for all nodes $x \in R$, where $x = \prod_{i=1}^n (d_i, q_i)$ and $y \triangleq \prod_{i=1}^n d_i$ with $n \in [0, \omega]$, it holds that (i) $y \in T$ and (ii), there is a set of moves $S \subseteq \text{Dir} \times Q$ with $S \models \delta(q_n, \nu(y))$ such that $x \cdot (d, q) \in R$, for all $(d, q) \in S$.*

In the following, we consider ATAs along with the *parity acceptance condition* (APT, for short) $\aleph \triangleq (F_1, \dots, F_k) \in (2^Q)^+$ with $F_1 \subseteq \dots \subseteq F_k = Q$ (see [KVV00], for more). The number k of sets in the tuple \aleph is called the *index* of the automaton. We also consider ATAs with the *co-Büchi acceptance condition* (ACT, for short) that is the special parity condition with index 2.

Let R be a run of an ATA \mathcal{A} on a tree \mathcal{T} and w one of its branches. Then, by $\text{inf}(w) \triangleq \{q \in Q : |\{i \in \mathbb{N} : \exists d \in \text{Dir}. (w)_i = (d, q)\}| = \omega\}$ we denote the set of states that occur infinitely often as the second component of the letters along the branch w . Moreover, we say that w satisfies the parity acceptance condition $\aleph = (F_1, \dots, F_k)$ if the least index $i \in [1, k]$ for which $\text{inf}(w) \cap F_i \neq \emptyset$ is even.

At this point, we can define the concept of language accepted by an ATA.

Definition 2.3.3 (ATA Acceptance) *An ATA $\mathcal{A} = \langle \Sigma, \Delta, Q, \delta, q_0, \aleph \rangle$ accepts a Σ -labeled Dir-tree \mathcal{T} iff there exists a run R of \mathcal{A} on \mathcal{T} such that all its infinite branches satisfy the acceptance condition \aleph .*

By $L(\mathcal{A})$ we denote the language accepted by the ATA \mathcal{A} , i.e., the set of trees \mathcal{T} accepted by \mathcal{A} . Moreover, \mathcal{A} is said to be *empty* if $L(\mathcal{A}) = \emptyset$. The *emptiness problem* for \mathcal{A} is to decide whether $L(\mathcal{A}) = \emptyset$.

We finally show a simple but useful result about the APT direction projection. To do this, we first need to introduce an extra notation. Let $f \in \mathcal{B}(P)$ be a Boolean formula on a set of propositions P . Then, by $f^{p/q|p \in P'}$ we denote the formula in which all occurrences of the propositions $p \in P' \subseteq P$ in f are replaced by the proposition q belonging to a possibly different set.

Theorem 2.3.1 (APT Direction Projection) *Let $\mathcal{A} \triangleq \langle \Sigma \times \text{Dir}, \text{Dir}, Q, \delta, q_0, \aleph \rangle$ be an APT over a set of m directions with n states and index k . Moreover, let $d_0 \in \text{Dir}$ be a distinguished direction. Then, there exists an NPT $\mathcal{N}^{d_0} \triangleq \langle \Sigma, \text{Dir}, Q', \delta', q'_0, \aleph' \rangle$ with $m \cdot 2^{O(k \cdot n \cdot \log n)}$ states and index $O(k \cdot n \cdot \log n)$ such that, for all Σ -labeled Dir-tree $\mathcal{T} \triangleq \langle T, \nu \rangle$, it holds that $\mathcal{T} \in L(\mathcal{N}^{d_0})$ iff $\mathcal{T}' \in L(\mathcal{A})$, where \mathcal{T}' is the $(\Sigma \times \text{Dir})$ -labeled Dir-tree $\langle T, \nu' \rangle$ such that $\nu'(t) \triangleq (\nu(t), \text{lst}(d_0 \cdot t))$, for all $t \in T$.*

Proof. As first step, we use the well-known nondeterminization procedure for APTs [MS95] in order to transform the APT \mathcal{A} into an equivalent NPT $\mathcal{N} = \langle \Sigma \times \text{Dir}, \text{Dir}, Q'', \delta'', q''_0, \aleph'' \rangle$ with $2^{O(k \cdot n \cdot \log n)}$ states and index $k' = O(k \cdot n \cdot \log n)$. Then, we transform the latter into the new NPT $\mathcal{N}^{d_0} \triangleq \langle \Sigma, \text{Dir}, Q', \delta', q'_0, \aleph' \rangle$ with $m \cdot 2^{O(k \cdot n \cdot \log n)}$ states and same index k' , where $Q' \triangleq Q'' \times \text{Dir}$, $q'_0 \triangleq (q''_0, d_0)$, $\aleph' \triangleq (F_1 \times \text{Dir}, \dots, F_{k'} \times \text{Dir})$ with $\aleph'' \triangleq (F_1, \dots, F_{k'})$, and $\delta'((q, d), \sigma) \triangleq \delta''(q, (\sigma, d))[(d', q') / (d', (q', d')) \mid (d', q') \in \text{Dir} \times Q'']$, for all $(q, d) \in Q'$ and $\sigma \in \Sigma$. Now, it easy to see that \mathcal{N}^{d_0} satisfies the declared statement. \square

2.3.2 SL Model Checking

A first step towards our construction of an algorithmic procedure for the solution of the SL model-checking problem is to define, for each possible formula φ , an alternating parity tree automaton $\mathcal{A}_\varphi^{\mathcal{G}}$ that recognizes a tree encoding \mathcal{T} of a CGS \mathcal{G} , containing the information on an assignment χ on the free variables/agents of φ , iff \mathcal{G} is a model of φ under χ . The high-level idea at the base of this construction is an evolution and merging of those behind the translations of QPTL and LTL, respectively, into nondeterministic [SVW87] and alternating [MSS88] Büchi automata.

To proceed with the formal description of the model-checking procedure, we have to introduce a concept of encoding for the assignments of a CGS.

Definition 2.3.4 (Assignment-State Encoding) *Let \mathcal{G} be a CGS, $s \in \text{St}_{\mathcal{G}}$ one of its states, and $\chi \in \text{Asg}_{\mathcal{G}}(\mathbb{V}, s)$ an assignment defined on the set $\mathbb{V} \subseteq \text{Vr} \cup \text{Ag}$. Then, a $(\text{Val}_{\text{Ac}_{\mathcal{G}}}(\mathbb{V}) \times \text{St}_{\mathcal{G}})$ -labeled $\text{St}_{\mathcal{G}}$ -tree $\mathcal{T} \triangleq \langle T, u \rangle$, where $T \triangleq \{\rho_{\geq 1} : \rho \in \text{Trk}_{\mathcal{G}}(s)\}$, is an assignment-state encoding for χ if it holds that $u(t) \triangleq (\widehat{\chi}(s \cdot t), \text{lst}(s \cdot t))$, for all $t \in T$.*

Observe that there is a unique assignment-state encoding for each given assignment.

In the next lemma, we prove the existence of an APT for each CGS and SL formula that is able to recognize all the assignment-state encodings of an a priori given assignment, made the assumption that the formula is satisfied on the CGS under this assignment.

Lemma 2.3.1 (SL Formula Automaton) *Let \mathcal{G} be a CGS and φ an SL formula. Then, there exists an APT $\mathcal{A}_\varphi^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_\varphi, \delta_\varphi, q_{0\varphi}, \aleph_\varphi \rangle$ such that, for all states $s \in \text{St}_{\mathcal{G}}$ and assignments $\chi \in \text{Asg}_{\mathcal{G}}(\text{free}(\varphi), s)$, it holds that $\mathcal{G}, \chi, s \models \varphi$ iff $\mathcal{T} \in \text{L}(\mathcal{A}_\varphi^{\mathcal{G}})$, where \mathcal{T} is the assignment-state encoding for χ .*

Proof. The construction of the APT $\mathcal{A}_\varphi^{\mathcal{G}}$ is done recursively on the structure of the formula φ , which w.l.o.g. is supposed to be in *enf*, by using a variation of the transformation, via alternating tree automata, of the S1S and SkS logics into nondeterministic Büchi word and tree automata recognizing all models of the formula of interest [Büc62, Rab69].

The detailed construction of $\mathcal{A}_\varphi^{\mathcal{G}}$, by a case analysis on φ , follows.

- If $\varphi \in \text{AP}$, the automaton has to verify if the atomic proposition is locally satisfied or not. To do this, we set $\mathcal{A}_\varphi^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\emptyset) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \{\varphi\}, \delta_\varphi, \varphi, (\{\varphi\}) \rangle$, where $\delta_\varphi(\varphi, (v, s)) \triangleq \mathbf{t}$, if $\varphi \in \text{L}_{\mathcal{G}}(s)$, and $\delta_\varphi(\varphi, (v, s)) \triangleq \mathbf{f}$, otherwise. Intuitively, $\mathcal{A}_\varphi^{\mathcal{G}}$ only verifies that the state s in the labeling of the root of the assignment-state encoding of the empty assignment \emptyset satisfies φ .
- The boolean case $\varphi = \neg\varphi'$ is treated in the classical way, by simply dualizing the automaton $\mathcal{A}_{\varphi'}^{\mathcal{G}} = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi')) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi'}, \delta_{\varphi'}, q_{0\varphi'}, \aleph_{\varphi'} \rangle$ [MS87].
- The boolean cases $\varphi = \varphi_1 \text{Op} \varphi_2$, with $\text{Op} \in \{\wedge, \vee\}$, are treated in a way that is similar to the classical one, by simply merging the two automata $\mathcal{A}_{\varphi_1}^{\mathcal{G}} = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi_1)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi_1}, \delta_{\varphi_1}, q_{0\varphi_1}, \aleph_{\varphi_1} \rangle$ and $\mathcal{A}_{\varphi_2}^{\mathcal{G}} = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi_2)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi_2}, \delta_{\varphi_2}, q_{0\varphi_2}, \aleph_{\varphi_2} \rangle$ into the automaton $\mathcal{A}_\varphi^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_\varphi, \delta_\varphi, q_{0\varphi}, \aleph_\varphi \rangle$, where the following hold:

- $\mathcal{Q}_\varphi \triangleq \{q_{0\varphi}\} \cup \mathcal{Q}_{\varphi_1} \cup \mathcal{Q}_{\varphi_2}$, with $q_{0\varphi} \notin \mathcal{Q}_{\varphi_1} \cup \mathcal{Q}_{\varphi_2}$;
- $\delta_\varphi(q_{0\varphi}, (v, s)) \triangleq \text{Op} \delta_{\varphi_1}(q_{0\varphi_1}, (v_{|\text{free}(\varphi_1)}, s)) \text{Op} \delta_{\varphi_2}(q_{0\varphi_2}, (v_{|\text{free}(\varphi_2)}, s))$, for all $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}$;
- $\delta_\varphi(q, (v, s)) \triangleq \delta_{\varphi_1}(q, (v_{|\text{free}(\varphi_1)}, s))$, if $q \in \mathcal{Q}_{\varphi_1}$, and $\delta_\varphi(q, (v, s)) \triangleq \delta_{\varphi_2}(q, (v_{|\text{free}(\varphi_2)}, s))$, otherwise, for all $q \in \mathcal{Q}_{\varphi_1} \cup \mathcal{Q}_{\varphi_2}$ and $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}$;
- $\aleph_\varphi \triangleq (F_{1\varphi}, \dots, F_{k\varphi})$, where (i) $\aleph_{\varphi_1} \triangleq (F_{1\varphi_1}, \dots, F_{k_1\varphi_1})$ and $\aleph_{\varphi_2} \triangleq (F_{1\varphi_2}, \dots, F_{k_2\varphi_2})$, (ii) $h = \min\{k_1, k_2\}$ and $k = \max\{k_1, k_2\}$, (iii) $F_{i\varphi} \triangleq F_{i\varphi_1} \cup F_{i\varphi_2}$, for $i \in [1, h]$, (iv) $F_{i\varphi} \triangleq F_{i\varphi_j}$, for $i \in [h+1, k-1]$ with $k_j = k$, and (v) $F_{k\varphi} \triangleq \mathcal{Q}_\varphi$.

- The case $\varphi = \text{X}\varphi'$ is solved by running the automaton $\mathcal{A}_{\varphi'}^{\mathcal{G}} = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi')) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi'}, \delta_{\varphi'}, q_{0\varphi'}, \aleph_{\varphi'} \rangle$ on the successor node of the root of the assignment-state encoding in the direction individuated by the assignment itself. To do this, we use the automaton $\mathcal{A}_\varphi^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_\varphi, \delta_\varphi, q_{0\varphi}, \aleph_\varphi \rangle$, where the following hold:

- $\mathcal{Q}_\varphi \triangleq \{q_{0\varphi}\} \cup \mathcal{Q}_{\varphi'}$, with $q_{0\varphi} \notin \mathcal{Q}_{\varphi'}$;
- $\delta_\varphi(q_{0\varphi}, (v, s)) \triangleq (\text{tr}_{\mathcal{G}}(s, v_{|\text{Ag}}), q_{0\varphi'})$, for all $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}$;

- $\delta_\varphi(q, (v, s)) \triangleq \delta_{\varphi'}(q, (v_{|\text{free}(\varphi')}, s))$, for all $q \in Q_{\varphi'}$ and $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_G$;
- $\aleph_\varphi \triangleq (F_{1\varphi'}, \dots, F_{k\varphi'} \cup \{q_{0\varphi}\})$, where $\aleph_{\varphi'} \triangleq (F_{1\varphi'}, \dots, F_{k\varphi'})$.

- To handle the case $\varphi = \varphi_1 U \varphi_2$, we use the automaton $\mathcal{A}_\varphi^G \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_G, \text{St}_G, Q_\varphi, \delta_\varphi, q_{0\varphi}, \aleph_\varphi \rangle$ that verifies the truth of the until operator using its one-step unfolding equivalence $\varphi_1 U \varphi_2 \equiv \varphi_2 \vee \varphi_1 \wedge X\varphi_1 U \varphi_2$, by appropriately running the two automata $\mathcal{A}_{\varphi_1}^G = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi_1)) \times \text{St}_G, \text{St}_G, Q_{\varphi_1}, \delta_{\varphi_1}, q_{0\varphi_1}, \aleph_{\varphi_1} \rangle$ and $\mathcal{A}_{\varphi_2}^G = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi_2)) \times \text{St}_G, \text{St}_G, Q_{\varphi_2}, \delta_{\varphi_2}, q_{0\varphi_2}, \aleph_{\varphi_2} \rangle$ for the inner formulas φ_1 and φ_2 . The definitions of \mathcal{A}_φ^G components follows:

- $Q_\varphi \triangleq \{q_{0\varphi}\} \cup Q_{\varphi_1} \cup Q_{\varphi_2}$, with $q_{0\varphi} \notin Q_{\varphi_1} \cup Q_{\varphi_2}$;
- $\delta_\varphi(q_{0\varphi}, (v, s)) \triangleq \delta_{\varphi_2}(q_{0\varphi_2}, (v_{|\text{free}(\varphi_2)}, s)) \vee \delta_{\varphi_1}(q_{0\varphi_1}, (v_{|\text{free}(\varphi_1)}, s)) \wedge (\text{tr}_G(s, v_{|\text{Ag}}, q_{0\varphi}))$, for all $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_G$;
- $\delta_\varphi(q, (v, s)) \triangleq \delta_{\varphi_1}(q, (v_{|\text{free}(\varphi_1)}, s))$, if $q \in Q_{\varphi_1}$, and $\delta_\varphi(q, (v, s)) \triangleq \delta_{\varphi_2}(q, (v_{|\text{free}(\varphi_2)}, s))$, otherwise, for all $q \in Q_{\varphi_1} \cup Q_{\varphi_2}$ and $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_G$;
- $\aleph_\varphi \triangleq (F_{1\varphi}, \dots, F_{k\varphi})$, where (i) $\aleph_{\varphi_1} \triangleq (F_{1\varphi_1}, \dots, F_{k_1\varphi_1})$ and $\aleph_{\varphi_2} \triangleq (F_{1\varphi_2}, \dots, F_{k_2\varphi_2})$, (ii) $h = \min\{k_1, k_2\}$ and $k = \max\{k_1, k_2\}$, (iii) $F_{i\varphi} \triangleq \{q_{0\varphi}\} \cup F_{i\varphi_1} \cup F_{i\varphi_2}$, for $i \in [1, h]$, (iv) $F_{i\varphi} \triangleq \{q_{0\varphi}\} \cup F_{i\varphi_j}$, for $i \in [h+1, k-1]$ with $k_j = k$, and (v) $F_{k\varphi} \triangleq Q_\varphi$.

It is important to observe that the initial state $q_{0\varphi}$ is included in all sets of the parity acceptance condition, in particular in $F_{1\varphi}$, in order to avoid its regeneration for an infinite number of times.

- To handle the case $\varphi = \varphi_1 R \varphi_2$, we use the automaton $\mathcal{A}_\varphi^G \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_G, \text{St}_G, Q_\varphi, \delta_\varphi, q_{0\varphi}, \aleph_\varphi \rangle$ that verifies the truth of the release operator using its one-step unfolding equivalence $\varphi_1 R \varphi_2 \equiv \varphi_2 \wedge (\varphi_1 \vee X\varphi_1 R \varphi_2)$, by appropriately running the two automata $\mathcal{A}_{\varphi_1}^G = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi_1)) \times \text{St}_G, \text{St}_G, Q_{\varphi_1}, \delta_{\varphi_1}, q_{0\varphi_1}, \aleph_{\varphi_1} \rangle$ and $\mathcal{A}_{\varphi_2}^G = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi_2)) \times \text{St}_G, \text{St}_G, Q_{\varphi_2}, \delta_{\varphi_2}, q_{0\varphi_2}, \aleph_{\varphi_2} \rangle$ for the inner formulas φ_1 and φ_2 . The definitions of \mathcal{A}_φ^G components follows:

- $Q_\varphi \triangleq \{q_{0\varphi}\} \cup Q_{\varphi_1} \cup Q_{\varphi_2}$, with $q_{0\varphi} \notin Q_{\varphi_1} \cup Q_{\varphi_2}$;
- $\delta_\varphi(q_{0\varphi}, (v, s)) \triangleq \delta_{\varphi_2}(q_{0\varphi_2}, (v_{|\text{free}(\varphi_2)}, s)) \wedge (\delta_{\varphi_1}(q_{0\varphi_1}, (v_{|\text{free}(\varphi_1)}, s)) \vee (\text{tr}_G(s, v_{|\text{Ag}}, q_{0\varphi}))$, for all $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_G$;
- $\delta_\varphi(q, (v, s)) \triangleq \delta_{\varphi_1}(q, (v_{|\text{free}(\varphi_1)}, s))$, if $q \in Q_{\varphi_1}$, and $\delta_\varphi(q, (v, s)) \triangleq \delta_{\varphi_2}(q, (v_{|\text{free}(\varphi_2)}, s))$, otherwise, for all $q \in Q_{\varphi_1} \cup Q_{\varphi_2}$ and $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_G$;
- $\aleph_\varphi \triangleq (F_{1\varphi}, \dots, F_{k\varphi})$, where (i) $\aleph_{\varphi_1} \triangleq (F_{1\varphi_1}, \dots, F_{k_1\varphi_1})$ and $\aleph_{\varphi_2} \triangleq (F_{1\varphi_2}, \dots, F_{k_2\varphi_2})$, (ii) $h = \min\{k_1, k_2\}$ and $k = \max\{k_1, k_2\}$, (iii) $F_{1\varphi} \triangleq F_{1\varphi_1} \cup F_{1\varphi_2}$, (iv) $F_{i\varphi} \triangleq \{q_{0\varphi}\} \cup F_{i\varphi_1} \cup F_{i\varphi_2}$, for $i \in [2, h]$, (iv) $F_{i\varphi} \triangleq \{q_{0\varphi}\} \cup F_{i\varphi_j}$, for $i \in [h+1, k-1]$ with $k_j = k$, and (v) $F_{k\varphi} \triangleq Q_\varphi$.

It is important to observe that, differently from the case of the until operator, the initial state $q_{0\varphi}$ is included in all sets of the parity acceptance condition but $F_{1\varphi}$, in order to allow its regeneration for an infinite number of time.

- The case $\varphi = (a, x)\varphi'$ is solved by simply transforming the transition function of the automaton $\mathcal{A}_{\varphi'}^{\mathcal{G}} = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi')) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi'}, \delta_{\varphi'}, q_{0\varphi'}, \aleph_{\varphi'} \rangle$, by setting the value of the valuations in input w.r.t. the agent a to the value of the same valuation w.r.t. the variable x . The definitions of the transition function for $\mathcal{A}_{\varphi}^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi'}, \delta_{\varphi}, q_{0\varphi'}, \aleph_{\varphi'} \rangle$ follows: $\delta_{\varphi}(q, (v, s)) \triangleq \delta_{\varphi'}(q, (v', s))$, where $v' = v_{a \rightarrow v(x) \upharpoonright \text{free}(\varphi')}$, if $a \in \text{free}(\varphi')$, and $v' = v$, otherwise, for all $q \in \mathcal{Q}_{\varphi'}$ and $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}$.
- To handle the case $\varphi = \langle\langle x \rangle\rangle\varphi'$, assuming that $x \in \text{free}(\varphi')$, we use the operation of existential projection for nondeterministic tree automata. To do this, we have first to nondeterminize the APT $\mathcal{A}_{\varphi'}^{\mathcal{G}}$, by applying the classic transformation [MS95]. In this way, we obtain an equivalent NPT $\mathcal{N}_{\varphi'}^{\mathcal{G}} = \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi')) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi'}, \delta_{\varphi'}, q_{0\varphi'}, \aleph_{\varphi'} \rangle$. Now, we make the projection, by defining the new NPT $\mathcal{A}_{\varphi}^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\varphi'}, \delta_{\varphi}, q_{0\varphi'}, \aleph_{\varphi'} \rangle$ where $\delta_{\varphi}(q, (v, s)) \triangleq \bigvee_{c \in \text{AcG}} \delta_{\varphi'}(q, (v_{x \rightarrow c}, s))$, for all $q \in \mathcal{Q}_{\varphi'}$ and $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(\varphi)) \times \text{St}_{\mathcal{G}}$.

At this point, it only remains to prove that, for all states $s \in \text{St}_{\mathcal{G}}$ and assignments $\chi \in \text{Asg}_{\mathcal{G}}(\text{free}(\varphi), s)$, it holds that $\mathcal{G}, \chi, s \models \varphi$ iff $\mathcal{T} \in \text{L}(\mathcal{A}_{\varphi}^{\mathcal{G}})$, where \mathcal{T} is the assignment-state encoding for χ . The proof can be developed by a simple induction on the structure of the formula φ and is left to the reader as a simple exercise. \square

We now have the tools to describe the recursive model-checking procedure on nested subsentences for SL and its fragments under the general semantics.

To proceed, we have first to prove the following theorem that represents the core of our automata-theoretic approach.

Theorem 2.3.2 (SL Sentence Automaton) *Let \mathcal{G} be a CGS, $s \in \text{St}_{\mathcal{G}}$ one of its states, and φ an SL sentence. Then, there exists an NPT $\mathcal{N}_{\varphi}^{\mathcal{G}, s}$ such that $\mathcal{G}, \emptyset, s \models \varphi$ iff $\text{L}(\mathcal{N}_{\varphi}^{\mathcal{G}, s}) \neq \emptyset$.*

Proof. To construct the NPT $\mathcal{N}_{\varphi}^{\mathcal{G}, s}$ we apply Theorem 2.3.1 on page 43 of APT direction projection with distinguished direction s to the APT $\mathcal{A}_{\varphi}^{\mathcal{G}}$ derived by Lemma 2.3.1 on page 44 of SL formula automaton. In this way, we can ensure that the state labeling of nodes of the assignment-state encoding is coherent with the node itself. Observe that, since φ is a sentence, we have that $\text{free}(\varphi) = \emptyset$, and so, the unique assignment-state encoding of interest is that related to the empty assignment \emptyset .

[Only if]. Suppose that $\mathcal{G}, \emptyset, s \models \varphi$. Then, by Lemma 2.3.1 on page 44, we have that $\mathcal{T} \in \text{L}(\mathcal{A}_{\varphi}^{\mathcal{G}})$, where \mathcal{T} is the dependence-state encoding for \emptyset . Hence, by Theorem 2.3.1 on page 43, it holds that $\text{L}(\mathcal{N}_{\varphi}^{\mathcal{G}, s}) \neq \emptyset$.

[If]. Suppose that $\text{L}(\mathcal{N}_{\varphi}^{\mathcal{G}, s}) \neq \emptyset$. Then, by Theorem 2.3.1 on page 43, there exists an $(\{\emptyset\} \times \text{St}_{\mathcal{G}})$ -labeled $\text{St}_{\mathcal{G}}$ -tree \mathcal{T} such that $\mathcal{T} \in \text{L}(\mathcal{A}_{\varphi}^{\mathcal{G}})$. Now, it is immediate to see that \mathcal{T} is the assignment-state encoding for \emptyset . Hence, by Lemma 2.3.1 on page 44, we have that $\mathcal{G}, \emptyset, s \models \varphi$. \square

Before continuing, we define the length $\text{lng}(\varphi)$ of an SL formula φ as the number $|\text{sub}(\varphi)|$ of its subformulas. We also introduce a generalization of the Knuth's double arrow notation in order to represent a tower of exponentials: $a \uparrow\uparrow_b 0 \triangleq b$ and $a \uparrow\uparrow_b (c + 1) \triangleq a^{\uparrow\uparrow_b c}$, for all $a, b, c \in \mathbb{N}$.

At this point, we prove the main theorem about the non-elementary complexity of SL model-checking problem.

Theorem 2.3.3 (SL Model Checking) *The model-checking problem for SL is PTIME-COMplete w.r.t. the size of the model and NONELEMENTARY w.r.t. the size of the specification.*

Proof. By Theorem 2.3.2 on the preceding page of SL sentence automaton, to verify that $\mathcal{G}, \emptyset, s \models \varphi$, we simply calculate the emptiness of the NPT $\mathcal{N}_{\varphi}^{\mathcal{G}, s}$ having $|\text{St}_{\mathcal{G}}| \cdot (2 \uparrow\uparrow_m m)$ states and index $2 \uparrow\uparrow_m m$, where $m = O(\text{lng}(\varphi) \cdot \log \text{lng}(\varphi))$. It is well-known that the emptiness problem for such a kind of automaton with n states and index h is solvable in time $O(n^h)$ [KV98]. Thus, we get that the time complexity of checking whether $\mathcal{G}, \emptyset, s \models \varphi$ is $|\text{St}_{\mathcal{G}}|^{2 \uparrow\uparrow_m m}$. Hence, the membership of the model-checking problem for SL in PTIME w.r.t. the size of the model and NONELEMENTARY w.r.t. the size of the specification directly follows. Finally, by getting the relative lower bound on the model from the same problem for ATL* [AHK02], the claim is proved. \square

Finally, we show a refinement of the previous result, when we consider sentences of the SL_[NG] fragment.

Theorem 2.3.4 (SL_[NG] Model Checking) *The model-checking problem for SL_[NG] is PTIME-COMplete w.r.t. the size of the model and $(k + 1)$ -EXPTIME w.r.t. the maximum alternation k of the specification.*

Proof. By Theorem 2.3.2 on the facing page of SL sentence automaton, to verify that $\mathcal{G}, \emptyset, s \models \wp\psi$, where $\wp\psi$ is an SL_[NG] principal sentence without proper subsentences, we can simply calculate the emptiness of the NPT $\mathcal{N}_{\wp\psi}^{\mathcal{G}, s}$ having $|\text{St}_{\mathcal{G}}| \cdot (2 \uparrow\uparrow_m k)$ states and index $2 \uparrow\uparrow_m k$, where $m = O(\text{lng}(\psi) \cdot \log \text{lng}(\psi))$ and $k = \text{alt}(\wp\psi)$. Thus, we get that the time complexity of checking whether $\mathcal{G}, \emptyset, s \models \wp\psi$ is $|\text{St}_{\mathcal{G}}|^{2 \uparrow\uparrow_m k}$. At this point, since we have to do this verification for each possible state $s \in \text{St}_{\mathcal{G}}$ and principal subsentence $\wp\psi \in \text{psnt}(\varphi)$ of the whole SL_[NG] specification φ , we derive that the bottom-up model-checking procedure requires time $|\text{St}_{\mathcal{G}}|^{2 \uparrow\uparrow_{\text{lng}(\varphi)} k}$, where $k = \max\{\text{alt}(\wp\psi) : \wp\psi \in \text{psnt}(\varphi)\}$. Hence, the membership of the model-checking problem for SL in PTIME w.r.t. the size of the model and $(k + 1)$ -EXPTIME w.r.t. the maximum alternation k of the specification directly follows. Finally, by getting the relative lower bound on the model from the same problem for ATL* [AHK02], the thesis is proved. \square

2.3.3 SL_[1G] Model Checking

We now show how the concept of behavioral of Skolem dependence functions over strategies can be used to enormously reduce the complexity of the model-checking procedure for the SL_[1G] fragment. The idea behind our approach is to avoid the use of projections used to

handle the strategy quantifications, by reducing them to action quantifications, which can be managed locally on each state of the model without a tower of exponential blow-ups.

To start with the description of the ad-hoc procedure for $\text{SL}[1\text{G}]$, we first prove the existence of a UCT for each CGS and $\text{SL}[1\text{G}]$ goal $b\psi$ that recognizes all the assignment-state encodings of an a priori given assignment, made the assumption that the goal is satisfied on the CGS under this assignment.

Lemma 2.3.2 (SL[1G] Goal Automaton) *Let \mathcal{G} be a CGS and $b\psi$ an SL[1G] goal without principal subsentences. Then, there exists an UCT $\mathcal{U}_{b\psi}^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(b\psi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{b\psi}, \delta_{b\psi}, q_{ob\psi}, \aleph_{b\psi} \rangle$ such that, for all states $s \in \text{St}_{\mathcal{G}}$ and assignments $\chi \in \text{Asg}_{\mathcal{G}}(\text{free}(b\psi), s)$, it holds that $\mathcal{G}, \chi, s \models b\psi$ iff $\mathcal{T} \in \text{L}(\mathcal{U}_{b\psi}^{\mathcal{G}})$, where \mathcal{T} is the assignment-state encoding for χ .*

Proof. A first step in the construction of the UCT $\mathcal{U}_{b\psi}^{\mathcal{G}}$ is to consider the UCW $\mathcal{U}_{\psi} \triangleq \langle 2^{\text{AP}}, \mathcal{Q}_{\psi}, \delta_{\psi}, \mathcal{Q}_{ob\psi}, \aleph_{\psi} \rangle$ obtained by dualizing the NBW resulting from the application of the classic Vardi-Wolper construction to the LTL formula $\neg\psi$ [VW86]. Observe that $\text{L}(\mathcal{U}_{\psi}) = \text{L}(\psi)$, i.e., \mathcal{U}_{ψ} recognizes all infinite words on the alphabet 2^{AP} that satisfy the LTL formula ψ . Then, define the components of $\mathcal{U}_{b\psi}^{\mathcal{G}} \triangleq \langle \text{Val}_{\text{AcG}}(\text{free}(b\psi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{b\psi}, \delta_{b\psi}, q_{ob\psi}, \aleph_{b\psi} \rangle$ as follows:

- $\mathcal{Q}_{b\psi} \triangleq \{q_{ob\psi}\} \cup \mathcal{Q}_{\psi}$, with $q_{ob\psi} \notin \mathcal{Q}_{\psi}$;
- $\delta_{b\psi}(q_{ob\psi}, (v, s)) \triangleq \bigwedge_{q \in \mathcal{Q}_{ob\psi}} \delta_{b\psi}(q, (v, s))$, for all $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(b\psi)) \times \text{St}_{\mathcal{G}}$;
- $\delta_{b\psi}(q, (v, s)) \triangleq \bigwedge_{q' \in \delta_{\psi}(q, \text{L}_{\mathcal{G}}(s))} (\text{tr}_{\mathcal{G}}(s, v \circ \text{bnd}_b), q')$, for all $q \in \mathcal{Q}_{\psi}$ and $(v, s) \in \text{Val}_{\text{AcG}}(\text{free}(b\psi)) \times \text{St}_{\mathcal{G}}$;
- $\aleph_{b\psi} \triangleq \aleph_{\psi}$.

Intuitively, the UCT $\mathcal{U}_{b\psi}^{\mathcal{G}}$ simply runs the UCW \mathcal{U}_{ψ} on the branch of the encoding individuated by the assignment in input. Thus, it is easy to see that, for all states $s \in \text{St}_{\mathcal{G}}$ and assignments $\chi \in \text{Asg}_{\mathcal{G}}(\text{free}(b\psi), s)$, it holds that $\mathcal{G}, \chi, s \models b\psi$ iff $\mathcal{T} \in \text{L}(\mathcal{U}_{b\psi}^{\mathcal{G}})$, where \mathcal{T} is the assignment-state encoding for χ . \square

Now, to describe our modified technique, we introduce a new concept of encoding regarding the behavioral Skolem dependence functions over strategies.

Definition 2.3.5 (Behavioral Dependence-State Encoding) *Let \mathcal{G} be a CGS, $s \in \text{St}_{\mathcal{G}}$ one of its states, and $\theta \in \text{BSDF}_{\text{Str}_{\mathcal{G}}(s)}(\wp)$ a behavioral Skolem dependence function over strategies for a quantification prefix $\wp \in \text{Qnt}(\text{V})$ over the set $\text{V} \subseteq \text{Vr}$. Then, a $(\text{SDF}_{\text{AcG}}(\wp) \times \text{St}_{\mathcal{G}})$ -labeled $\text{St}_{\mathcal{G}}$ -tree $\mathcal{T} \triangleq \langle \text{T}, \text{u} \rangle$, where $\text{T} \triangleq \{\rho_{\geq 1} : \rho \in \text{Trk}_{\mathcal{G}}(s)\}$, is a behavioral dependence-state encoding for θ if it holds that $\text{u}(t) \triangleq (\theta(s \cdot t), \text{lst}(s \cdot t))$, for all $t \in \text{T}$.*

Observe that there exists a unique behavioral dependence-state encoding for each behavioral Skolem dependence function over strategies.

In the next lemma, we show how to handle locally the strategy quantifications on each state of the model, by simply using a quantification over actions, which is modeled by the choice of an action Skolem dependence function. Intuitively, we guess in the labeling what is the right part of the Skolem dependence function over strategies for each node of the tree and then verify that, for all assignments of universal variables, the corresponding complete assignment satisfies the inner formula.

Lemma 2.3.3 (SL[1G] Sentence Automaton) *Let \mathcal{G} be a CGS and $\wp b\psi$ an SL[1G] principal sentence without principal subsentences. Then, there exists a UCT $\mathcal{U}_{\wp b\psi}^{\mathcal{G}} \triangleq \langle \text{SDF}_{\text{AcG}}(\wp) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\wp b\psi}, \delta_{\wp b\psi}, q_{0\wp b\psi}, \aleph_{\wp b\psi} \rangle$ such that, for all states $s \in \text{St}_{\mathcal{G}}$ and behavioral Skolem dependence functions over strategies $\theta \in \text{BSDF}_{\text{Str}_{\mathcal{G}}(s)}(\wp)$, it holds that $\mathcal{G}, \theta(\chi), s \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)$, iff $\mathcal{T} \in \text{L}(\mathcal{U}_{\wp b\psi}^{\mathcal{G}})$, where \mathcal{T} is the behavioral dependence-state encoding for θ .*

Proof. By Lemma 2.3.2 on the preceding page of SL[1G] goal automaton, there is an UCT $\mathcal{U}_{b\psi}^{\mathcal{G}} = \langle \text{Val}_{\text{AcG}}(\text{free}(b\psi)) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{b\psi}, \delta_{b\psi}, q_{0b\psi}, \aleph_{b\psi} \rangle$ such that, for all states $s \in \text{St}_{\mathcal{G}}$ and assignments $\chi \in \text{Asg}_{\mathcal{G}}(\text{free}(b\psi), s)$, it holds that $\mathcal{G}, \chi, s \models b\psi$ iff $\mathcal{T} \in \text{L}(\mathcal{U}_{b\psi}^{\mathcal{G}})$, where \mathcal{T} is the assignment-state encoding for χ .

Now, transform $\mathcal{U}_{b\psi}^{\mathcal{G}}$ into the new UCT $\mathcal{U}_{\wp b\psi}^{\mathcal{G}} \triangleq \langle \text{SDF}_{\text{AcG}}(\wp) \times \text{St}_{\mathcal{G}}, \text{St}_{\mathcal{G}}, \mathcal{Q}_{\wp b\psi}, \delta_{\wp b\psi}, q_{0\wp b\psi}, \aleph_{\wp b\psi} \rangle$, with $\mathcal{Q}_{\wp b\psi} \triangleq \mathcal{Q}_{b\psi}$, $q_{0\wp b\psi} \triangleq q_{0b\psi}$, and $\aleph_{\wp b\psi} \triangleq \aleph_{b\psi}$, which is used to handle the quantification prefix \wp atomically, where the transition function is defined as follows: $\delta_{\wp b\psi}(q, (\theta, s)) \triangleq \bigwedge_{v \in \text{Val}_{\text{AcG}}(\llbracket \wp \rrbracket)} \delta_{b\psi}(q, (\theta(v), s))$, for all $q \in \mathcal{Q}_{\wp b\psi}$ and $(\theta, s) \in \text{SDF}_{\text{AcG}}(\wp) \times \text{St}_{\mathcal{G}}$. Intuitively, $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$ reads an action Skolem dependence function θ on each node of the input tree \mathcal{T} labeled with a state s of \mathcal{G} and simulates the execution of the transition function $\delta_{b\psi}(q, (v, s))$ of $\mathcal{U}_{b\psi}^{\mathcal{G}}$, for each possible valuation $v = \theta(v')$ on $\text{free}(b\psi)$ obtained from θ by a universal valuation $v' \in \text{Val}_{\text{AcG}}(\llbracket \wp \rrbracket)$. It is important to observe that we cannot move the component set $\text{SDF}_{\text{AcG}}(\wp)$ from the input alphabet to the states of $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$, by making a related guessing of the Skolem dependence function θ in the transition function, since we have to ensure that all states in a given node of the tree \mathcal{T} , i.e., in each track of the original model \mathcal{G} , make the same choice for θ .

Finally, it remains to prove that, for all states $s \in \text{St}_{\mathcal{G}}$ and behavioral Skolem dependence function over strategies $\theta \in \text{BSDF}_{\text{Str}_{\mathcal{G}}(s)}(\wp)$, it holds that $\mathcal{G}, \theta(\chi), s \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)$, iff $\mathcal{T} \in \text{L}(\mathcal{U}_{\wp b\psi}^{\mathcal{G}})$, where \mathcal{T} is the behavioral dependence-state encoding for θ .

[Only if]. Suppose that $\mathcal{G}, \theta(\chi), s \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)$. Since ψ does not contain principal subsentences, we have that $\mathcal{G}, \theta(\chi), s \models b\psi$. So, due to the property of $\mathcal{U}_{b\psi}^{\mathcal{G}}$, it follows that there exists an assignment-state encoding $\mathcal{T}_{\chi} \in \text{L}(\mathcal{U}_{b\psi}^{\mathcal{G}})$, which implies the existence of an $(\text{St}_{\mathcal{G}} \times \mathcal{Q}_{b\psi})$ -tree R_{χ} that is an accepting run for $\mathcal{U}_{b\psi}^{\mathcal{G}}$ on \mathcal{T}_{χ} . At this point, let $R \triangleq \bigcup_{\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)} R_{\chi}$ be the union of all runs. Then, due to the particular definition of the transition function of $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$, it is not hard to see that R is an accepting run for $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$ on \mathcal{T} . Hence, $\mathcal{T} \in \text{L}(\mathcal{U}_{\wp b\psi}^{\mathcal{G}})$.

[If]. Suppose that $\mathcal{T} \in \text{L}(\mathcal{U}_{\wp b\psi}^{\mathcal{G}})$. Then, there exists an $(\text{St}_{\mathcal{G}} \times \mathcal{Q}_{\wp b\psi})$ -tree R that is an accepting run for $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$ on \mathcal{T} . Now, for each $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)$, let R_{χ} be the run for $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$ on the assignment-state encoding \mathcal{T}_{χ} for $\theta(\chi)$. Due to the particular definition of the transition function of $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$, it is easy to see that $R_{\chi} \subseteq R$. Thus, since R is accepting, we have that R_{χ} is accepting as well. So, $\mathcal{T}_{\chi} \in \text{L}(\mathcal{U}_{b\psi}^{\mathcal{G}})$. At this point, due to the property of $\mathcal{U}_{b\psi}^{\mathcal{G}}$, it follows that $\mathcal{G}, \theta(\chi), s \models b\psi$. Now, since ψ does not contain principal subsentences, we have that $\mathcal{G}, \theta(\chi), s \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)$. \square

At this point, we can prove the following theorem that is at the base of the elementary model-checking procedure for SL[1G].

Theorem 2.3.5 (SL[1G] Sentence Automaton) *Let \mathcal{G} be a CGS, $s \in \text{St}_{\mathcal{G}}$ one of its states, and $\wp b\psi$ an SL[1G] principal sentence without principal subsentences. Then, there exists an NPT $\mathcal{N}_{\wp b\psi}^{\mathcal{G},s}$ such that $\mathcal{G}, \emptyset, s \models \wp b\psi$ iff $L(\mathcal{N}_{\wp b\psi}^{\mathcal{G},s}) \neq \emptyset$.*

Proof. As in the general case of SL sentence automaton, we have to ensure that the state labeling of nodes of the behavioral dependence-state encoding is coherent with the node itself. To do this, we apply Theorem 2.3.1 on page 43 of APT direction projection with distinguished direction s to the UPT $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$ derived by Lemma 2.3.3 on the previous page of the SL[1G] sentence automaton, thus obtaining the required NPT $\mathcal{N}_{\wp b\psi}^{\mathcal{G},s}$.

[Only if]. Suppose that $\mathcal{G}, \emptyset, s \models \wp b\psi$. By Corollary 2.2.2 on page 41 of SL[1G] behavioral, it means that $\mathcal{G}, \emptyset, s \models_{\text{B}} \wp b\psi$. Then, by Definition 2.2.4 on page 35 of SL[NG] behavioral semantics, there exists a behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}_{\mathcal{G}}(s)}(\wp)$ such that $\mathcal{G}, \theta(\chi), s \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)$. Thus, by Lemma 2.3.3 on the preceding page, we have that $\mathcal{T} \in L(\mathcal{U}_{\wp b\psi}^{\mathcal{G}})$, where \mathcal{T} is the behavioral dependence-state encoding for θ . Hence, by Theorem 2.3.1 on page 43, it holds that $L(\mathcal{N}_{\wp b\psi}^{\mathcal{G},s}) \neq \emptyset$.

[If]. Suppose that $L(\mathcal{N}_{\wp b\psi}^{\mathcal{G},s}) \neq \emptyset$. Then, by Theorem 2.3.1 on page 43, there exists an $(\text{SDF}_{\text{Ac}_{\mathcal{G}}}(\wp) \times \text{St}_{\mathcal{G}})$ -labeled $\text{St}_{\mathcal{G}}$ -tree \mathcal{T} such that $\mathcal{T} \in L(\mathcal{U}_{\wp b\psi}^{\mathcal{G}})$. Now, it is immediate to see that there is a behavioral Skolem dependence function $\theta \in \text{BSDF}_{\text{Str}_{\mathcal{G}}(s)}(\wp)$ for which \mathcal{T} is the behavioral dependence-state encoding. Thus, by Lemma 2.3.3 on the previous page, we have that $\mathcal{G}, \theta(\chi), s \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s)$. By Definition 2.2.4 on page 35 of SL[NG] behavioral semantics, it holds that $\mathcal{G}, \emptyset, s \models_{\text{B}} \wp b\psi$. Hence, by Corollary 2.2.2 on page 41 of SL[1G] behavioral, it means that $\mathcal{G}, \emptyset, s \models \wp b\psi$. \square

Finally, we show in the next fundamental theorem the precise complexity of the model-checking for SL[1G].

Theorem 2.3.6 (SL[1G] Model Checking) *The model-checking problem for SL[1G] is PTIME-COMplete w.r.t. the size of the model and 2EXPTIME-COMplete w.r.t. the size of the specification.*

Proof. By Theorem 2.3.5 of SL[1G] sentence automaton, to verify that $\mathcal{G}, \emptyset, s \models \wp b\psi$, we simply calculate the emptiness of the NPT $\mathcal{N}_{\wp b\psi}^{\mathcal{G},s}$. This automaton is obtained by the operation of direction projection on the UCT $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$, which is in turn derived by the UCT $\mathcal{U}_{b\psi}^{\mathcal{G}}$. Now, it is easy to see that the number of states of $\mathcal{U}_{b\psi}^{\mathcal{G}}$, and consequently of $\mathcal{U}_{\wp b\psi}^{\mathcal{G}}$, is $2^{\text{O}(\text{lg}(\psi))}$. So, $\mathcal{N}_{\wp b\psi}^{\mathcal{G},s}$ has $|\text{St}_{\mathcal{G}}| \cdot 2^{2^{\text{O}(\text{lg}(\psi))}}$ states and index $2^{\text{O}(\text{lg}(\psi))}$.

The emptiness problem for such a kind of automaton with n states and index h is solvable in time $\text{O}(n^h)$ [KV98]. Thus, we get that the time complexity of checking whether $\mathcal{G}, \emptyset, s \models \wp b\psi$ is $|\text{St}_{\mathcal{G}}|^{2^{\text{O}(\text{lg}(\psi))}}$. At this point, since we have to do this verification for each possible state $s \in \text{St}_{\mathcal{G}}$ and principal subsentence $\wp b\psi \in \text{psnt}(\wp)$ of the whole SL[1G] specification \wp , we derive that the whole bottom-up model-checking procedure requires time $|\text{St}_{\mathcal{G}}|^{2^{\text{O}(\text{lg}(\wp))}}$. Hence, the membership of the model-checking problem for SL[1G] in PTIME w.r.t. the size of the model and 2EXPTIME w.r.t. the size of the specification directly follows. Finally the thesis is proved, by getting the relative lower bounds from the same problem for ATL^* [AHK02]. \square

2.4 Satisfiability

As shown in Section 1.3.2, SL does not have the bounded model property and, as a trivial consequence, its satisfiability problem is undecidable. By means of Remark 2.1.2 on page 29, we derive that also the fragment $\text{SL}_{[\text{BG}]}$ has an undecidable satisfiability problem. On the contrary, it is well-known that the satisfiability problem for ATL^* is 2EXPTIME-COMplete [Sch08]. This gap in complexity between SL and ATL^* gives naturally rise to the question of which are the inherent properties of ATL^* that make the problem decidable. Here we answer such question by analyzing the syntactic fragment $\text{SL}_{[1G]}$ introduced in Section 2.1. We show that this fragment retains all positive properties of ATL^* , such as the decision-tree model property and the bounded model property, which allows us to show that its satisfiability problem is 2EXPTIME-COMplete . A fundamental feature used as a tool to prove the announced properties is again the *behavioral satisfiability*.

2.4.1 Tree-model property

The satisfiability procedure we propose in this thesis is based on the use of alternating tree automata. Consequently, we need to establish a kind of *tree model property*, which is based on a special sub-class of CGSs, namely, the *concurrent game-trees* (CGT, for short), whose structure of the underlying graph is a tree.

Definition 2.4.1 (Concurrent Game Trees) A concurrent game tree (CGT, for short) is a CGS $\mathcal{T} \triangleq \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_0 \rangle$, where (i) $\text{St} \subseteq \text{Dir}^*$ is a Dir-tree for a given set Dir of directions and (ii) if $t \cdot e \in \text{St}$ then there is a decision $\delta \in \text{Dc}$ such that $\text{tr}(t, \delta) = t \cdot e$, for all $t \in \text{St}$ and $e \in \text{Dir}$. Furthermore, \mathcal{T} is a decision tree (DT, for short) if (i) $\text{St} = \text{Dc}^*$ and (ii) if $t \cdot \delta \in \text{St}$ then $\text{tr}(t, \delta) = t \cdot \delta$, for all $t \in \text{St}$ and $\delta \in \text{Dc}$.

Intuitively, CGTs are CGSs having a transition relation with a tree shape and DTs have, in addition, the states that uniquely determine the history of the computation leading to them. Observe that, for each non trivial track ρ (resp., path π), there exists a unique finite (resp., infinite) sequence of decisions $\delta_0 \cdot \dots \cdot \delta_{|\rho|-2} \in \text{Dc}^*$ (resp., $\delta_0 \cdot \delta_1 \cdot \dots \in \text{Dc}^\omega$) such that $\rho_{(i+1)} = \text{tr}(\rho_{(i)}, \delta_i)$ (resp., $\pi_{(i+1)} = \text{tr}(\pi_{(i)}, \delta_i)$), for all $i \in [0, |\rho| - 1[$ (resp., $i \in \mathbb{N}$).

We now define a generalization for CGSs of the classic concept of *unwinding* of labeled transition systems, namely the *decision-unwinding* (see Figure 2.2 and Figure 2.3 on the next page, for an example), that allows to show that $\text{SL}_{[1G]}$ enjoys the decision-tree model property.




Figure 2.2: A CGS.
 $\text{ap}'(t) = \text{ap}(\text{unw}(t))$, for all $t \in \text{Dc}^*$ and $\delta \in \text{Dc}$.

Definition 2.4.2 (Decision-Unwinding) Let $\mathcal{G} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_0 \rangle$ be a CGS. Then, the decision-unwinding of \mathcal{G} is the DT $\mathcal{G}_{\text{DU}} \triangleq \langle \text{AP}, \text{Ag}, \text{Ac}, \text{Dc}^*, \text{ap}', \text{tr}', \varepsilon \rangle$ for which there is a surjective function $\text{unw} : \text{Dc}^* \rightarrow \text{St}$ such that (i) $\text{unw}(\varepsilon) = s_0$, (ii) $\text{unw}(\text{tr}'(t, \delta)) = \text{tr}(\text{unw}(t), \delta)$, and (iii)

Observe that, due to its construction, each CGS \mathcal{G} has a unique associated decision-unwinding \mathcal{G}_{DU} .

We are now able to prove that $\text{SL}[1\text{G}]$ satisfies the decision-tree model property.




Figure 2.3: Part of the decision-unwinding of the CGS in Figure 2.2.

Theorem 2.4.1 (SL[1G] Decision-Tree Model Property) *Let φ be a satisfiable SL[1G] sentence. Then, there exists a DT \mathcal{T} such that $\mathcal{T} \models \varphi$.*

Proof. The proof proceeds by structural induction on the sentence SL[1G]. For the Boolean combination of principal sentences, the induction is trivial. For the case of a principal sentence φ of the form $\wp b \psi$, by Theorem 2.2.4 on page 40, we derive that there exists a behavioral Skolem map $\theta \in \text{BSDF}_{\text{Str}_{\mathcal{G}}(s_0)}(\wp)$ such that $\mathcal{G} \models_{\theta} \wp b \psi$. Furthermore, there exists the adjoint function to tracks $\tilde{\theta}$ of θ . Now, consider the decision unwinding $\mathcal{T} = \mathcal{G}_{DU}$ of \mathcal{G} and the lifting $\Gamma : \text{Trk}_{\mathcal{T}} \rightarrow \text{Trk}_{\mathcal{G}}$ of the unwinding function unw such that $\Gamma(\rho) = \text{unw}(\rho_0) \cdot \dots \cdot \text{unw}(\rho_{|\rho|-1})$, for all $\rho \in \text{Trk}_{\mathcal{T}}$. At this point, consider the function $\tilde{\theta}'$ such that $\tilde{\theta}'(\rho') \triangleq \tilde{\theta}(\Gamma(\rho'))$, for all $\rho' \in \text{Trk}_{\mathcal{T}}$. Clearly, since $\tilde{\theta}$ is a Skolem map over actions, so $\tilde{\theta}'$ is as well. Then, consider the Skolem map θ' for which the function $\tilde{\theta}'$ is its adjoint. By induction on the nesting of principal subsentences in $\varphi = \wp b \psi$, we now prove that $\mathcal{T} \models_{\theta'} \wp b \psi$. As base case, *i.e.*, when ψ is an LTL formula, consider an assignment $\chi' \in \text{Asg}_{\mathcal{T}}(\llbracket \wp \rrbracket, \varepsilon)$ and the induced play $\pi' = \text{play}(\theta'(\chi') \circ b, \varepsilon)$ over \mathcal{T} . Moreover, consider an assignment $\chi \in \text{Asg}_{\mathcal{G}}(\llbracket \wp \rrbracket, s_0)$ such that, for all placeholders $l \in \text{dom}(\chi)$ and tracks $\rho' \in \text{Trk}_{\mathcal{T}}$, it holds that $\chi(l)(\Gamma(\rho')) = \chi'(l)(\rho')$. From the satisfiability of φ on \mathcal{G} , we derive that $\pi \models \psi$, where $\pi = \text{play}(\theta(\chi) \circ b, s_0)$. Indeed, assume for a while that $(\pi)_{\leq k} = \Gamma((\pi')_{\leq k})$. Then, from the definition of the labeling in the decision-tree unwinding, it easily follows that $\text{ap}((\pi)_k) = \text{ap}'((\pi')_k)$, and, so we derive $\pi' \models \psi$, from $\pi \models \psi$. Consequently, this holds for all $\chi' \in \text{Asg}_{\text{Str}_{\mathcal{T}}(\varepsilon)}(\llbracket \wp \rrbracket)$ and, so, we have that $\mathcal{T} \models_{\theta'} \wp b \psi$. The inductive case, *i.e.*, when ψ contains some subsentence, easily follows by considering the principal subsentences as fresh atomic propositions.

It remains to prove, by induction on k , that $(\pi)_{\leq k} = \Gamma((\pi')_{\leq k})$, for all $k \in \mathbb{N}$. As base case, we have that $(\pi)_0 = s_0 = \Gamma(\varepsilon) = (\pi')_0$. As inductive case, assume that $(\pi)_{\leq k} = \Gamma((\pi')_{\leq k})$. Then, in particular, we have that $(\pi)_k = \Gamma((\pi')_k) = \text{unw}((\pi')_k)$. By definition of play, it holds that $(\pi)_{k+1} = \text{tr}((\pi)_k, (\tilde{\theta}((\pi)_{\leq k}))(\hat{\chi}) \circ b)$, which is, by inductive hypothesis, equal to $\text{tr}(\text{unw}((\pi')_k), (\tilde{\theta}(\Gamma((\pi)_{\leq k}))) (\hat{\chi})) \circ b$. Now, by the definition of $\tilde{\theta}'$ and χ , we obtain $\text{tr}(\text{unw}((\pi')_k), (\tilde{\theta}(\Gamma((\pi)_{\leq k}))) (\hat{\chi})) \circ b = \text{tr}(\Gamma((\pi')_k), (\tilde{\theta}'((\pi')_{\leq k})) (\hat{\chi}')) \circ b$. Finally, by the definition of unw and Γ , we have that $\text{tr}(\Gamma((\pi')_k), (\tilde{\theta}'((\pi')_{\leq k})) (\hat{\chi}')) \circ b = \Gamma((\pi')_{k+1})$. \square

2.4.2 Bounded model property

In order to prove the bounded-tree model property for $\text{SL}[1G]$, we first need to introduce the new concept of *disjoint satisfiability*, which regards the verification of different instances of the same subsentence of the original specification. Intuitively, it asserts that either these instances can be checked on disjoint subtrees of the tree model or, if two instances use part of the same subtree, they are forced to use the same dependence map as well. This notion is a reformulation of the notion of explicit model introduced for ATL^* in [Sch08]. This intrinsic characteristic of $\text{SL}[1G]$ is fundamental for the building of a unique automaton that checks the truth of all subsentences, by simply merging their respective automata, without using a projection operation to eliminate their own alphabets, which otherwise may be in conflict. In this way, we are also able to avoid an exponential blow-up. A deeper discussion on this point is reported later in the paper.

Definition 2.4.3 (Disjoint Satisfiability) *Let \mathcal{T} be a DT and $\varphi = \wp b \psi$ be a $\text{SL}[1G]$ principal sentence. Moreover, let $S \triangleq \{s \in \text{St}_{\mathcal{T}} : \mathcal{T}, s \models \varphi\}$. Then, \mathcal{T} satisfies φ disjointly over S if there exist two functions $\text{head} : S \rightarrow \text{SM}_{\text{Ac}}(\wp)$ and $\text{body} : \text{Trk}(\varepsilon) \rightarrow \text{SM}_{\text{Ac}}(\wp)$ such that, for all $s \in S$ and $\chi \in \text{Asg}_{\text{Str}(s)}(\llbracket \varphi \rrbracket)$ it holds that $\mathcal{T}, \theta(\chi), s \models b\psi$, where the behavioral Sdf $\theta \in \text{BSDF}_{\text{Str}(s)}$ is defined, by means of its adjoint, as follows:*

- (i) $\tilde{\theta}(s) \triangleq \text{head}(s)$;
- (ii) $\tilde{\theta}(\rho) \triangleq \text{body}(\rho' \cdot \rho)$, for all $\rho \in \text{Trk}(s)$ with $|\rho| > 1$, where $\rho' \in \text{Trk}(\varepsilon)$ is the unique track such that $\rho' \cdot \rho \in \text{Trk}(\varepsilon)$ ⁵.

The disjoint satisfiability holds for all $\text{SL}[1G]$ formulas. To prove this fact, we first introduce the preliminary definition of *twin decision-tree*. Intuitively, in such a kind of tree, each action is flanked by a twin one having the same purpose of the original. This allows to satisfy two sentences requiring the same actions in a given state in two different branches of the tree itself, which is what the disjoint satisfiability precisely requires.

Definition 2.4.4 (Twin Decision Tree) *Let $\mathcal{T} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, \varepsilon \rangle$ be a DT. Then, the twin decision tree of \mathcal{T} is the DT $\mathcal{T}' \triangleq \langle \text{AP}, \text{Ag}, \text{Ac}', \text{St}', \text{tr}', \text{ap}', \varepsilon' \rangle$ with $\text{Ac}' = \text{Ac} \times \{\text{new}, \text{cont}\}$ and $\varepsilon' = (\varepsilon, \text{new})$. The labeling and the transition function are defined by means of a set of projection functions introduced below:*

- the function $\text{prj}_{\text{Ac}} : \text{Ac}' \rightarrow \text{Ac}$ returns the first component of the action in Ac' , i.e., $\text{prj}_{\text{Ac}}((c, \iota)) = c$, for all $(c, \iota) \in \text{Ac}'$;
- the function $\text{prj}_{\text{Dc}} : \text{Dc}' \rightarrow \text{Dc}$ projects out the flags on all the actions in the decision, returning a corresponding decision in \mathcal{T} , i.e., $\text{prj}_{\text{Dc}}(\delta')(a) = \text{prj}_{\text{Ac}}(\delta'(a))$, for all $\delta' \in \text{Dc}'$ and $a \in \text{Ag}$;
- the function $\text{prj}_{\text{St}} : \text{St}' \rightarrow \text{St}$ returns the corresponding state in \mathcal{T} , according to the projection made on the decisions, i.e., $\text{prj}_{\text{St}}(\varepsilon') = \varepsilon$ and $\text{prj}_{\text{St}}(s' \cdot \delta') = \text{prj}_{\text{St}}(s') \cdot \text{prj}_{\text{Dc}}(\delta')$, for all $s' \in \text{St}'$ and $\delta' \in \text{Dc}'$;

⁵Existence and uniqueness of ρ' is guaranteed by the fact that \mathcal{T} is a DT.

- analogously, the function $\text{prj}_{\text{Trk}} : \text{Trk}' \rightarrow \text{Trk}$, returns the concatenation of the projected states, i.e., $\text{prj}_{\text{Trk}}(s') = \text{prj}_{\text{St}}(s')$, for all $s' \in \text{St}'$, and $\text{prj}_{\text{Trk}}(\rho' \cdot s') = \text{prj}_{\text{Trk}}(\rho') \cdot \text{prj}_{\text{St}}(s')$, for all $\rho' \in \text{Trk}'$ and $s' \in \text{St}'$.

Then, $\text{ap}'(s') \triangleq \text{ap}(\text{prj}_{\text{St}}(s'))$.

Observe that $\text{prj}_{\text{St}}(\text{tr}'(s', \delta')) = \text{tr}(\text{prj}_{\text{St}}(s'), \text{prj}_{\text{Dc}}(\delta'))$, for all $s' \in \text{St}'$ and $\delta' \in \text{Dc}'$. We can now prove the disjoint satisfiability property for $\text{SL}[1G]$.

Theorem 2.4.2 (Disjoint Satisfiability) *Let $\varphi = \wp b \psi$ be an $\text{SL}[1G]$ principal sentence and $\mathcal{T} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, \varepsilon \rangle$ a DT. Moreover, let $S \triangleq \{s \in \text{St} : \mathcal{T}, s \models \varphi\}$. Then the twin decision tree \mathcal{T}' of \mathcal{T} disjointly satisfies φ over $S' \triangleq \{s' \in \text{St}' : \text{prj}(s') \in S\}$.*

Proof.[Proof idea] Starting from the fact that $\mathcal{T}, s \models \varphi$, for all $s \in S$, by means of Theorem 2.2.4 on page 40, we derive the existence of a behavioral Sdf θ_s . Such a θ_s is used to define a behavioral Sdf θ_s' in \mathcal{T}' in which the existential agents suitably select either *new* or *cont* as second component, in order to guarantee the satisfaction of different instances over different branches of the twin decision tree. Indeed, it allows to properly define the two functions head and body and, consequently, the behavioral Sdf θ' for which we finally prove that $\mathcal{T}', s' \models \varphi$, for all $s' \in S'$. Since θ' has been built from the head and body functions, the disjoint satisfiability is immediately derived. \square

Proof. Let $s \in S$ be one of the states on which φ is satisfied. Since $\mathcal{T}, s \models \varphi$, by Theorem 2.2.4 on page 40, we have that there exists $\theta_s \in \text{BSDF}_{\text{Str}(s)}(\varphi)$ such that $\mathcal{T}, \theta_s(\chi), s \models b\psi$, for all states s -total assignments $\chi \in \text{Asg}_{\text{Str}_{\mathcal{T}}(s)}(\llbracket \varphi \rrbracket)$. Then, consider the adjoint function $\tilde{\theta}_s' : \text{Trk}' \rightarrow \text{SM}_{\text{Ac}'}(\varphi)$ defined, for all states $s' \in \text{St}'$, decisions $\delta' \in \text{Dc}'$, and tracks $\rho' \in \text{Trk}'$ as follows:

- $\tilde{\theta}_s'(s')(\delta')(x) = (\tilde{\theta}_s(\text{prj}_{\text{St}}(s'))(\text{prj}_{\text{Dc}}(\delta'))(x), \text{new})$, if $\text{prj}_{\text{St}}(s') = s$;
- $\tilde{\theta}_s'(\rho')(\delta')(x) = (\tilde{\theta}_s(\text{prj}_{\text{Trk}}(\rho'))(\text{prj}_{\text{Dc}}(\delta'))(x), \text{cont})$, otherwise.

At this point, we assume the function head : $\text{St}' \rightarrow \text{SM}_{\text{Ac}'}(\varphi)$ to be defined as follows: $\text{head}(s') \triangleq \tilde{\theta}_s'(s')$. Moreover, we set the function body : $\text{Trk}'(\varepsilon) \rightarrow \text{SM}_{\text{Ac}'}(\varphi)$ in such a way that it agrees with $\tilde{\theta}_s'$ on all tracks $\rho' = s_0' \cdot \dots \cdot s_n'$ for which there is an index $i \in \{0, \dots, n\}$ such that, for all agents $a \in \text{Ag}$ with $b(a) \in \llbracket \varphi \rrbracket$, it holds that:

- $\text{lst}((\rho')_i)(a) = (c_a, \text{new})$, for some $c_a \in \text{Ac}$, and
- $\text{lst}((\rho')_j)(a) = (c_a, \text{cont})$, for all $j \in \{i + 1, \dots, n\}$ and for some $c_a \in \text{Ac}$.

Note that the tracks of this form are such that the players bound to an existentially quantified variable have selected an action flagged by *new* on the i -th step of the game and then keep playing with the *cont* flag. Intuitively, they are starting the verification of a subsentence right in the i -th state of the track, by keeping it separated from the verification of the other subsentences, which are addressed with the *cont* flag.

For all the other tracks ρ' , instead, the value of $\text{body}(\rho')$ may be arbitrary.

Now, consider the behavioral Sdf $\theta' \in \text{BSDF}_{\text{Str}}(\wp)$ defined by means of the functions head and body as prescribed by Definition 2.4.3 on page 53. It remains to prove that $\mathcal{T}', s' \models_{\theta'} \wp b \psi$, for all $s' \in S'$. We proceed by induction on the nesting of the principal subsentences of φ . As base case, assume that such nesting is 0. This means that ψ is an LTL formula. Now, let $\chi' \in \text{Asg}_{\text{Str}(s')}(\llbracket \wp \rrbracket)$. By construction, it is not hard to see that there exists an assignment $\chi \in \text{Asg}_{\text{Str}(\text{prj}(s'))}(\llbracket \wp \rrbracket)$ for which the play $\pi' \triangleq \text{play}'(\theta'(\chi') \circ b, s')$ satisfies the equality $\text{prj}_{\text{Pth}}(\pi') = \text{play}(\theta(\chi) \circ b, \text{prj}(s')) = \pi$ ⁶. Thus, since $\mathcal{T}, s \models_{\theta} \wp b \psi$, we have that $\pi \models \psi$. Moreover, it holds that $\text{ap}'((\pi')_i) = \text{ap}((\pi)_i)$, for all $i \in \mathbb{N}$, which implies that $\pi' \models \psi$. Consequently, we can conclude that $\mathcal{T}', s' \models_{\theta'} \wp b \psi$. The inductive case, easily follows by considering the inner principal subsentences as fresh atomic propositions. \square

We now have all tools to prove the bounded model property of $\text{SL}[1G]$.

Theorem 2.4.3 (Bounded Model Property of $\text{SL}[1G]$) *Let φ be a $\text{SL}[1G]$ sentence and \mathcal{T} be a DT such that $\mathcal{T} \models \varphi$. Then, there exists a bounded DT \mathcal{T}' such that $\mathcal{T}' \models \varphi$.*

The proof makes use of some instruments and formalisms for *First-Order Logic* (FOL, for short) that are introduced in [MP14]. For the sake of completeness, here we give an informal discussion of such object. A *language signature* is a tuple $\mathcal{L} = \langle \text{Ar}, \text{Rl}, \text{ar} \rangle$ in which Ar and Rl are two finite non-empty sets of *arguments* and *relations*, respectively, and $\text{ar} : \text{Rl} \rightarrow 2^{\text{Ar}} \setminus \{\emptyset\}$ is a function mapping each relation in Rl to its non-empty set of arguments. Language signatures are used to reformulate FOL syntax in terms of *binding forms*, which are a way to associate variables to relations by means of bindings. The interpretation of FOL formulas is given on *relational structures*, which are tuples $\mathcal{R} = \langle \text{Dm}, \text{rl} \rangle$ with Dm being a non-empty domain and where $\text{rl}(r) \subseteq \text{ar}(r) \rightarrow \text{Dm}$ is a set of functions, representing the tuples on which the relation $r \in \text{Rl}$ is interpreted as true.

Proof.[Proof Idea] The key idea used to prove the theorem is based on the *finite model property* of the *One-Binding* fragment of FOL ($\text{FOL}[1B]$, for short), proved in [MP14], which allows to define a bounded-tree model \mathcal{T}' , which preserves the satisfiability of φ . In particular, for each state s^* of a tree \mathcal{T} satisfying φ , we build a first-order structure and a $\text{FOL}[1B]$ formula η_{s^*} that characterizes the topology of the successors of s_* in \mathcal{T} . Then, since $\text{FOL}[1B]$ enjoys the finite model property, we are able to build a finite first-order structure for η_{s^*} from which we can build the bounded model \mathcal{T}' of φ . Such a construction is based on both the disjoint and behavioral satisfiability of $\text{SL}[1G]$. For each state s^* in \mathcal{T} , we consider a set given by pairs of subsentences η of φ and states s , on which it holds that $\mathcal{T}, s \models_{\theta} \eta$, where such satisfaction is forced to pass through s^* for at least one universal assignment fed to θ . This means that at least one play used to satisfy η passes through s^* . By the disjoint satisfiability, we have to cope with at most two Sdfs for each subsentence η , those given by head_{η} and body_{η} , implying that the total number of Sdfs to take into account, for all s^* , is finite. From that, we define a related $\text{FOL}[1B]$ sentence η_{s^*} , having a model derived from the topology of the successors of s^* whose elements are constituted by the actions of \mathcal{T} . Now, by applying the finite model property to η_{s^*} , we derive the existence of a model for the formula η_{s^*} with a finite domain Ac'_{η, s^*} . Exploiting the finite model built for all states s^* , we are able to define the labeling of \mathcal{T}' and a behavioral Sdf θ' in such a way that $\mathcal{T}', \varepsilon \models_{\theta'} \varphi$. \square

⁶By prj_{Pth} we are denoting the natural lifting of the function prj_{Trk} to paths.

Proof.[Proof of Theorem 2.4.3 on the preceding page] We give the proof for the case of $\varphi = \wp b \psi$, since the Boolean combination of principal sentences easily follows from this one. Given a tree-model \mathcal{T} for φ , derived by the tree-model property of SL[1G] of Theorem 2.4.1 on page 52, for each state $s^* \in \text{St}$, consider the set $\Phi_{s^*} \subseteq \text{St} \times \text{psnt}(\varphi)$ of states of \mathcal{T} and principal subsentences of φ such that $(s, \eta) \in \Phi_{s^*}$ iff (i) $\mathcal{T}, s \models \eta$ and (ii) there exists an assignment $\chi \in \text{Asg}_{\text{Str}(s)}(\llbracket \wp \eta \rrbracket)$ such that $s^* = (\text{play}(\theta_{(s,\eta)}^{s^*})(\chi)) \circ b_{\eta}, s)_n$, for some $n \in \mathbb{N}$, where the behavioral Sdf $\theta_{(s,\eta)}^{s^*}$ is defined by means of its adjoint, which is in its turn built from the functions head_{η} and body_{η} , given by Theorem 2.4.2 on page 54, applied on $\eta = \wp b \psi_{\eta}$. Observe that, for a fixed η , if $s_1, s_2 \in \text{St}$, with $s_1 \neq s^*$ and $s_2 \neq s^*$, $\widetilde{\theta_{(s,\eta)}^{s^*}}(\rho_{s_1}) = \text{head}_{\eta}(\rho'_{s_1} \cdot \rho_{s_1}) = \text{head}_{\eta}(\rho'_{s_2} \cdot \rho_{s_2}) = \widetilde{\theta_{(s,\eta)}^{s^*}}(\rho_{s_2})$, where ρ_{s_1} and ρ_{s_2} are the unique tracks ending in s^* and starting in s_1 and s_2 , respectively, while ρ'_{s_1} and ρ'_{s_2} are the unique tracks such that $\rho'_{s_1} \cdot \rho_{s_1}$ and $\rho'_{s_2} \cdot \rho_{s_2}$ start from ε . Now, for a given Sdf over Actions $\tilde{\theta} \in \text{SM}_{\text{Ac}}(\wp)$ and a given state $s \in \text{St}$, define the set $\text{Succ}_{\tilde{\theta}, b}(s) \triangleq \{s' \in \text{St} : \exists v \in \text{Ac}^{\llbracket \wp \rrbracket}. \text{tr}(s, \vartheta(v) \circ b) = s'\}$, where $\vartheta \in \text{Ac}^{\llbracket \wp \rrbracket} \rightarrow \text{Ac}^{\text{Tr}(\wp)}$ is a Sdf for \wp over actions. Intuitively, the set $\text{Succ}_{\tilde{\theta}, b}(s)$ defines the set of states that can be reached in one step from s by prescribing the agents that are bound by b to an existential variable to move according to the Sdf ϑ .

At this point, consider the language signature $\mathcal{L} = \langle \text{Ar}, \text{Rl}, \text{ar} \rangle = \langle \text{Ag}, \text{AP}, \text{ar} \rangle$ with $\text{ar}(p) = \text{Ag}$, for all $p \in \text{AP}$, where each atomic proposition is viewed as a relation having the agents as arguments and, so, the decisions as elements of its interpretation. Moreover, for all sets $P \subseteq \text{AP}$, let $\text{mask}^P = \bigwedge_{p \in P} p \wedge \bigwedge_{q \in \text{AP} \setminus P} \neg q$ be the FOL[1B] formula asserting that only the relations in P hold. Finally, for all $(s, \eta) \in \Phi_{s^*}$, consider the FOL[1B] sentence $\eta_s^* = \wp \eta b_{\eta} \bigvee_{s \in \text{Succ}_{\tilde{\theta}_{(s,\eta)}^{s^*}}(\text{trk}(s^*), b_{\eta})} \text{mask}^{\text{ap}(s)}$ ⁷. Clearly, by definition, each η_s^* is satisfied by the relational structure $\mathcal{R}_{s^*} = \langle \text{Ac}, \text{rl}_{s^*} \rangle$ with $\text{rl}_{s^*}(p) \triangleq \{\delta \in \text{Dc} : p \in \text{ap}(\text{tr}(s^*, \delta))\}$, where a relation p is interpreted as true on all decisions that allow s^* to reach a state in which p holds. Indeed, for each partial valuation $v \in \text{Ac}^{\llbracket \wp \rrbracket}$, it holds that either $s' = \text{tr}(s, \text{head}_{\eta}(s)(v)) \in \text{Succ}_{\text{head}_{\eta}(s), b_{\eta}}(s)$ or $s'' = \text{tr}(s, \text{body}_{\eta}(\text{trk}(s))(v)) \in \text{Succ}_{\text{body}_{\eta}(\text{trk}(s)), b_{\eta}}(s)$, which implies that either $\mathcal{R}_{s^*}, \text{head}_{\eta}(s)(v) \models \text{mask}^{\text{ap}(s')}$ or $\mathcal{R}_{s^*}, \text{body}_{\eta}(\text{trk}(s))(v) \models \text{mask}^{\text{ap}(s')}$. Hence, we have that $\mathcal{R}_{s^*} \models \eta_s^*$ and, since this is true for all $(s, \eta) \in \Phi_{s^*}$, we derive that $\mathcal{R}_{s^*} \models \bigwedge_{(s,\eta) \in \Phi_{s^*}} \eta_s^*$.

At this point, from the finite model property of FOL[1B], we derive that there exists a finite relational structure $\mathcal{R}'_{s^*} = \langle \text{Dm}'_{s^*}, \text{rl}'_{s^*} \rangle$ such that $\mathcal{R}'_{s^*} \models \bigwedge_{(s,\eta) \in \Phi_{s^*}} \eta_s^*$. Moreover, we define $\widetilde{\theta'_{(s,\eta)}^{s^*}}$ to be such that $\mathcal{R}'_{s^*} \models \widetilde{\theta'_{(s,\eta)}^{s^*}} \eta_s^*$. Observe that, in the proof of finite model property for FOL[1B] [MP14], the bound on Dm'_{s^*} only depends on the quantification and binding prefixes given in the formulas η_s^* , which all occur in φ . Thus, the size of Dm'_{s^*} does not depend on η and, *w.l.o.g.*, we can assume that $\text{Dm}'_{s^*} = \text{Ac}'$ for all $s^* \in \text{St}$. Moreover, again from the finite model property proof of FOL[1B], there exists a function $M_{s^*} : \text{Dc}' \rightarrow \text{Dc}$, with $\text{Dc}' = \text{Ac}'^{\text{Ag}}$, such that, for all $(s, \eta) \in \Phi_{s^*}$ and $\delta' \in \text{rng}(\widetilde{\theta'_{(s,\eta)}^{s^*}})$, we have that $M_{s^*}(\delta') \in \text{rng}(\widetilde{\theta_{(s,\eta)}^{s^*}})$, where $\widetilde{\theta_{(s,\eta)}^{s^*}}$ is the Sdf used to satisfy η_s^* on \mathcal{R} . At this point, we define a DT \mathcal{T}' having Ac' as set of

⁷By $\text{trk}(s)$ we are denoting the unique track starting from ε and ending in s .

actions. In order to define the labeling function ap' , first consider the mapping $\Gamma : \text{St}' \rightarrow \text{St}$ recursively defined as follows:

- $\Gamma(\varepsilon) = \varepsilon$;
- $\Gamma(s' \cdot \delta') = \Gamma(s') \cdot M_{\Gamma(s')}(\delta')$.

By means of Γ , define $\text{ap}'(s') \triangleq \text{ap}(\Gamma(s'))$ for all $s' \in \text{St}'$. It remains to prove that $\mathcal{T}' \models \varphi$. We do this by using a Sdf $\theta' \in \text{SM}_{\text{Str}'(\varepsilon)}(\wp)$ defined from the adjoint $\tilde{\theta}'$ introduced in the following.

Let ρ' be a track in \mathcal{T}' and consider $s' = \text{lst}(\rho')$. If $(\varepsilon, \varphi) \in \Phi_{\Gamma(s')}$, define $\tilde{\theta}'(\rho') = \theta'_{(s, \eta)}^{s^*}$. For all other tracks, the value of $\tilde{\theta}'$ may be arbitrary. We now show that $\mathcal{T}, \emptyset, \varepsilon \models_{\theta'} \varphi$, by induction on the nesting of principal subsentences. As base case, suppose that φ has nesting 0. This implies that it is of the form $\wp b \psi$ with ψ being an LTL formula. Then, consider a universal assignment $\chi' \in \text{Asg}_{\text{Str}'(\varepsilon)}(\llbracket \wp \rrbracket)$ and then the total assignment $\theta'(\chi')$. This determines a play $\pi' = \text{play}(\theta'(\chi'), \varepsilon)$. Now, consider a universal assignment $\chi \in \text{Asg}_{\text{Str}(\varepsilon)}(\llbracket \wp \rrbracket)$ such that, for all $x \in \llbracket \wp \rrbracket$ and $\rho' \in \text{Trk}'(\varepsilon)$, it holds that $\chi'(x)(\rho') = \chi(x)(\Gamma(\rho'))$, where Γ is the lifting over tracks of the mapping over states defined above, *i.e.*, by $\Gamma(\rho')$ is the track in \mathcal{T} obtained from ρ' by mapping each state s' in ρ' into $s = \Gamma(s')$. It holds that $\pi = \text{play}(\theta(\chi) \circ b, \varepsilon)$ is such that, for all $i \in \mathbb{N}$, $(\pi)_{\leq i} = \Gamma((\pi')_{\leq i})$. Indeed, by induction on i , as base case, we have that $(\pi)_{\leq 0} = \varepsilon = \Gamma(\varepsilon) = (\pi')_{\leq 0}$. As inductive case, suppose that $(\pi)_{\leq i} = \Gamma((\pi')_{\leq i})$. Then, $(\pi)_{i+1} = \text{tr}((\pi)_i, \tilde{\theta}((\pi)_{\leq i}) \circ b) = \text{tr}(\Gamma((\pi')_i), \tilde{\theta}(\Gamma((\pi')_{\leq i})) \circ b) = \Gamma(\text{tr}((\pi')_i, \tilde{\theta}'((\pi')_{\leq i}) \circ b)) = \Gamma((\pi')_{i+1})$. Thus, according to the definition of ap' , we have that $\text{ap}'((\pi')_i) = \text{ap}((\pi)_i)$, for all $i \in \mathbb{N}$. Since $\pi \models \psi$, we derive that $\pi' \models \psi$. This holds for all possible universal assignments $\chi' \in \text{Asg}_{\text{Str}'(\varepsilon)}(\llbracket \wp \rrbracket)$. Hence, it holds that $\mathcal{T}' \models_{\theta'} \wp b \psi$. The inductive case follows by considering the principal subsentences of φ as fresh atomic propositions. \square

2.4.3 Satisfiability procedure

We finally solve the satisfiability problem for $\text{SL}[1G]$ and show that it is 2EXPTIME-COMplete , as for ATL^* . The algorithmic procedure is based on an automata-theoretic approach, which reduces the decision problem for the logic to the emptiness problem of a suitable universal Co-Büchi tree automaton (UCT, for short) [GTW02]. From an high-level point of view, the automaton construction seems similar to what was proposed in literature for CTL^* [KVW00] and ATL^* [Sch08]. However, our technique is completely new, since it is based on the novel notions of behavioral semantics and disjoint satisfiability.

Principal sentences. To proceed, we first have to introduce the concept of encoding for an assignment and the labeling of a DT.

Definition 2.4.5 (Assignment-Labeling Encoding) *Let \mathcal{T} be a DT, $t \in \text{St}_{\mathcal{T}}$ one of its states, and $\chi \in \text{Asg}_{\mathcal{T}}(\mathbb{V}, t)$ a t -total assignment defined on the set $\mathbb{V} \subseteq \text{Vr}$. A $(\text{Val}_{\text{Ac}_{\mathcal{T}}}(\mathbb{V}) \times 2^{\text{AP}})$ -labeled $\text{Dc}_{\mathcal{T}}$ -tree $\mathcal{T}' \triangleq \langle \text{St}_{\mathcal{T}}, u \rangle$ is an assignment-labeling encoding for χ on \mathcal{T} if $u(\text{lst}((\rho)_{\geq 1})) = (\hat{\chi}(\rho), \text{ap}_{\mathcal{T}}(\text{lst}(\rho)))$, for all $\rho \in \text{Trk}_{\mathcal{T}}(t)$ ⁸.*

⁸Note that $\text{lst}(\varepsilon) = \varepsilon$.

Observe that there is a unique assignment-labeling encoding for each assignment over a given DT.

Now, we prove the existence of a UCT $\mathcal{U}_{b\psi}^{\text{Ac}}$ for each SL[1G] goal $b\psi$ having no principal subsentences. The $\mathcal{U}_{b\psi}^{\text{Ac}}$ recognizes all the assignment-labeling encodings \mathcal{T}' of an a priori given assignment χ over a generic DT \mathcal{T} , whenever the goal is satisfied on \mathcal{T} under χ . Intuitively, we start with a UCW, recognizing all infinite words on the alphabet 2^{AP} that satisfy the LTL formula ψ , obtained by a simple variation of the Vardi-Wolper construction [VW86]. Then, we run it on the encoding tree \mathcal{T}' by following the directions identified by the assignment in its labeling.

Lemma 2.4.1 (SL[1G] Goal Automaton) *Let $b\psi$ an SL[1G] goal without principal subsentences and Ac a finite set of actions. Then, there exists an UCT $\mathcal{U}_{b\psi}^{\text{Ac}} \triangleq \langle \text{Val}_{\text{Ac}}(\text{free}(b\psi)) \times 2^{\text{AP}}, \text{Dc}, \text{Q}_{b\psi}, \delta_{b\psi}, q_{0b\psi}, \aleph_{b\psi} \rangle$ such that, for all DT $s \mathcal{T}$ with $\text{Ac}_{\mathcal{T}} = \text{Ac}$, states $t \in \text{St}_{\mathcal{T}}$, and t -total assignments $\chi \in \text{Asg}_{\mathcal{T}}(\text{free}(b\psi), t)$, it holds that $\mathcal{T}, \chi, t \models b\psi$ iff $\mathcal{T}' \in \text{L}(\mathcal{U}_{b\psi}^{\text{Ac}})$, where \mathcal{T}' is the assignment-labeling encoding for χ on \mathcal{T} .*

Proof. A first step in the construction of the UCT $\mathcal{U}_{b\psi}^{\text{Ac}}$, is to consider the UCW $\mathcal{U}_{\psi} \triangleq \langle 2^{\text{AP}}, \text{Q}_{\psi}, \delta_{\psi}, \text{Q}_{0\psi}, \aleph_{\psi} \rangle$ obtained by dualizing the NBW resulting from the application of the classic Vardi-Wolper construction to the LTL formula $\neg\psi$ [VW86]. Observe that $\text{L}(\mathcal{U}_{\psi}) = \text{L}(\psi)$, i.e., this automaton recognizes all infinite words on the alphabet 2^{AP} that satisfy the LTL formula ψ . Then, define the components of $\mathcal{U}_{b\psi}^{\text{Ac}} \triangleq \langle \text{Val}_{\text{Ac}}(\text{free}(b\psi)) \times 2^{\text{AP}}, \text{Dc}, \text{Q}_{b\psi}, \delta_{b\psi}, q_{0b\psi}, \aleph_{b\psi} \rangle$, as follows:

- $\text{Q}_{b\psi} \triangleq \{q_{0b\psi}\} \cup \text{Q}_{\psi}$, with $q_{0b\psi} \notin \text{Q}_{\psi}$;
- $\delta_{b\psi}(q_{0b\psi}, (v, \sigma)) \triangleq \bigwedge_{q \in \text{Q}_{0\psi}} \delta_{b\psi}(q, (v, \sigma))$, for all $(v, \sigma) \in \text{Val}_{\text{Ac}}(\text{free}(b\psi)) \times 2^{\text{AP}}$;
- $\delta_{b\psi}(q, (v, \sigma)) \triangleq \bigwedge_{q' \in \delta_{\psi}(q, \sigma)} (v \circ b, q')$, for all $q \in \text{Q}_{\psi}$ and $(v, \sigma) \in \text{Val}_{\text{Ac}}(\text{free}(b\psi)) \times 2^{\text{AP}}$;
- $\aleph_{b\psi} \triangleq \aleph_{\psi}$.

Intuitively, the UCT $\mathcal{U}_{b\psi}^{\text{Ac}}$ simply runs the UCW \mathcal{U}_{ψ} on the branch of the encoding individuated by the assignment in input. Thus, it is easy to see that, for all states $t \in \text{St}_{\mathcal{T}}$ and t -total assignments $\chi \in \text{Asg}_{\mathcal{T}}(\text{free}(b\psi), t)$, it holds that $\mathcal{T}, \chi, t \models b\psi$ iff $\mathcal{T}' \in \text{L}(\mathcal{U}_{b\psi}^{\text{Ac}})$, where \mathcal{T}' is the assignment-labeling encoding for χ on \mathcal{T} . \square

We now introduce a new concept of encoding regarding the behavioral dependence maps over strategies.

Definition 2.4.6 (Behavioral Dependence-Labeling Encoding) *Let \mathcal{T} be a DT, $t \in \text{St}_{\mathcal{T}}$ one of its states, and $\theta \in \text{BSDF}_{\text{Str}_{\mathcal{T}}(t)}(\wp)$ a behavioral dependence map over strategies for a quantification prefix $\wp \in \text{Qnt}(V)$ over the set $V \subseteq \text{Vr}$. A $(\text{SM}_{\text{Ac}_{\mathcal{T}}}(\wp) \times 2^{\text{AP}})$ -labeled Dir-tree $\mathcal{T}' \triangleq \langle \text{St}_{\mathcal{T}}, u \rangle$ is a behavioral dependence-labeling encoding for θ on \mathcal{T} if $u(\text{lst}((\rho)_{\geq 1})) = (\tilde{\theta}(\rho), \text{ap}_{\mathcal{T}}(\text{lst}(\rho)))$, for all $\rho \in \text{Trk}_{\mathcal{T}}(t)$.*

Observe that also in this case there exists a unique behavioral dependence-labeling encoding for each behavioral dependence map over strategies.

Finally, in the next lemma, we show how to locally handle the strategy quantifications on each state of the model, by simply using a quantification over actions modeled by the choice of an action dependence map. Intuitively, we guess in the labeling what is the right part of the dependence map over strategies for each node of the tree and then verify that, for all assignments of universal variables, the corresponding complete assignment satisfies the inner formula.

Lemma 2.4.2 (SL[1G] Sentence Automaton) *Let $\wp b\psi$ be an SL[1G] principal sentence without principal subsentences and Ac a finite set of actions. Then, there exists an UCT $\mathcal{U}_{\wp b\psi}^{\text{Ac}} \triangleq \langle \text{SM}_{\text{Ac}}(\wp) \times 2^{\text{AP}}, \text{Dc}, \mathcal{Q}_{\wp b\psi}, \delta_{\wp b\psi}, q_{0\wp b\psi}, \aleph_{\wp b\psi} \rangle$ such that, for all DT $s \mathcal{T}$ with $\text{Ac}_{\mathcal{T}} = \text{Ac}$, states $t \in \text{St}_{\mathcal{T}}$, and behavioral dependence maps over strategies $\theta \in \text{BSDF}_{\text{St}_{\mathcal{T}}(t)}(\wp)$, it holds that $\mathcal{T}, \theta(\chi), t \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{T}}(\llbracket \wp \rrbracket, t)$, iff $\mathcal{T}' \in \text{L}(\mathcal{U}_{\wp b\psi}^{\text{Ac}})$, where \mathcal{T}' is the behavioral dependence-labeling encoding for θ on \mathcal{T} .*

Proof. By Lemma 2.4.1 on the preceding page of SL[1G] goal automaton, there is an UCT $\mathcal{U}_{b\psi}^{\text{Ac}} \triangleq \langle \text{Val}_{\text{Ac}}(\text{free}(b\psi)) \times 2^{\text{AP}}, \text{Dc}, \mathcal{Q}_{b\psi}, \delta_{b\psi}, q_{0b\psi}, \aleph_{b\psi} \rangle$ such that, for all DTs \mathcal{T} with $\text{Ac}_{\mathcal{T}} = \text{Ac}$, states $t \in \text{St}_{\mathcal{T}}$, and assignments $\chi \in \text{Asg}_{\mathcal{T}}(\text{free}(b\psi), t)$, it holds that $\mathcal{T}, \chi, t \models b\psi$ iff $\mathcal{T}' \in \text{L}(\mathcal{U}_{b\psi}^{\text{Ac}})$, where \mathcal{T}' is the assignment-labeling encoding for χ on \mathcal{T} .

Now, transform $\mathcal{U}_{b\psi}^{\text{Ac}}$ into the new UCT $\mathcal{U}_{\wp b\psi}^{\text{Ac}} \triangleq \langle \text{SM}_{\text{Ac}}(\wp) \times 2^{\text{AP}}, \text{Dc}, \mathcal{Q}_{\wp b\psi}, \delta_{\wp b\psi}, q_{0\wp b\psi}, \aleph_{\wp b\psi} \rangle$, with $\mathcal{Q}_{\wp b\psi} \triangleq \mathcal{Q}_{b\psi}$, $q_{0\wp b\psi} \triangleq q_{0b\psi}$, and $\aleph_{\wp b\psi} \triangleq \aleph_{b\psi}$, which is used to handle the quantification prefix \wp atomically, where the transition function is defined as follows: $\delta_{\wp b\psi}(q, (\theta, \sigma)) \triangleq \bigwedge_{v \in \text{Val}_{\text{Ac}}(\llbracket \wp \rrbracket)} \delta_{b\psi}(q, (\theta(v), \sigma))$, for all $q \in \mathcal{Q}_{\wp b\psi}$ and $(\theta, \sigma) \in \text{SM}_{\text{Ac}}(\wp) \times 2^{\text{AP}}$. Intuitively, $\mathcal{U}_{\wp b\psi}^{\text{Ac}}$ reads an action dependence map θ on each node of the input tree \mathcal{T}' labeled with a set of atomic propositions σ and simulates the execution of the transition function $\delta_{b\psi}(q, (v, \sigma))$ of $\mathcal{U}_{b\psi}^{\text{Ac}}$, for each possible valuation $v = \theta(v')$ on $\text{free}(b\psi)$ obtained from θ via a universal valuation $v' \in \text{Val}_{\text{Ac}}(\llbracket \wp \rrbracket)$. It is worth observing that we cannot move the component set $\text{SM}_{\text{Ac}}(\wp)$ from the input alphabet to the states of $\mathcal{U}_{\wp b\psi}^{\text{Ac}}$ by making a related guessing of the dependence map θ in the transition function, since the automaton is universal and we have to ensure that all states in a given node of the tree \mathcal{T}' , i.e., in each track of the original model \mathcal{T} , make the same choice for θ .

Finally, it remains to prove that, for all states $t \in \text{St}_{\mathcal{T}}$ and behavioral dependence maps over strategies $\theta \in \text{BSDF}_{\text{St}_{\mathcal{T}}(t)}(\wp)$, it holds that $\mathcal{T}, \theta(\chi), t \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{T}}(\llbracket \wp \rrbracket, t)$, iff $\mathcal{T}' \in \text{L}(\mathcal{U}_{\wp b\psi}^{\text{Ac}})$, where \mathcal{T}' is the behavioral dependence-labeling encoding for θ on \mathcal{T} .

[Only if]. Suppose that $\mathcal{T}, \theta(\chi), t \models_{\text{B}} b\psi$, for all $\chi \in \text{Asg}_{\mathcal{T}}(\llbracket \wp \rrbracket, t)$. Since ψ does not contain principal subsentences, we have that $\mathcal{T}, \theta(\chi), t \models b\psi$. So, due to the property of $\mathcal{U}_{b\psi}^{\text{Ac}}$, it follows that there exists an assignment-labeling encoding $\mathcal{T}'_{\chi} \in \text{L}(\mathcal{U}_{b\psi}^{\text{Ac}})$, which implies the existence of a $(\text{Dc} \times \mathcal{Q}_{b\psi})$ -tree R_{χ} that is an accepting run for $\mathcal{U}_{b\psi}^{\text{Ac}}$ on \mathcal{T}'_{χ} . At this point, let $R \triangleq \bigcup_{\chi \in \text{Asg}_{\mathcal{T}}(\llbracket \wp \rrbracket, t)} R_{\chi}$ be the union of all runs. Then, due to the particular definition of the transition function of $\mathcal{U}_{\wp b\psi}^{\text{Ac}}$, it is not hard to see that R is an accepting run for $\mathcal{U}_{\wp b\psi}^{\text{Ac}}$ on \mathcal{T}' defined as above. Hence, $\mathcal{T}' \in \text{L}(\mathcal{U}_{\wp b\psi}^{\text{Ac}})$.

[If]. Suppose that $\mathcal{T}' \in \text{L}(\mathcal{U}_{\wp b\psi}^{\text{Ac}})$. Then, there exists a $(\text{Dc} \times \mathcal{Q}_{\wp b\psi})$ -tree R that is an accepting run for $\mathcal{U}_{\wp b\psi}^{\text{Ac}}$ on \mathcal{T}' . Now, for each $\chi \in \text{Asg}_{\mathcal{T}}(\llbracket \wp \rrbracket, t)$, let R_{χ} be the run for $\mathcal{U}_{\wp b\psi}^{\text{Ac}}$ on the assignment-state encoding \mathcal{T}'_{χ} for $\theta(\chi)$ on \mathcal{T} . Due to the particular definition of the

transition function of $\mathcal{U}_{\wp b \psi}^{\text{Ac}}$, it is not hard to see that $R_\chi \subseteq R$. Thus, since R is accepting, we have that R_χ is accepting as well. So, $\mathcal{T}'_\chi \in L(\mathcal{U}_{b \psi}^{\text{Ac}})$. At this point, due to the property of $\mathcal{U}_{b \psi}^{\text{Ac}}$, it follows that $\mathcal{T}, \theta(\chi), t \models b \psi$. Since ψ does not contain principal subsentences, we have that $\mathcal{T}, \theta(\chi), t \models_{\text{B}} b \psi$, for all $\chi \in \text{Asg}_{\mathcal{T}}(\llbracket \wp \rrbracket, t)$. \square

Full sentences. By summing up all previous results, we are now able to solve the satisfiability problem for the full SL[1G] fragment.

To construct the automaton for a given SL[1G] sentence φ , we first consider all UCT $\mathcal{U}_\phi^{\text{Ac}}$, for an assigned bounded set Ac , previously described for the principal sentences $\phi \in \text{psnt}(\varphi)$, in which the inner subsentences are considered as atomic propositions. Then, thanks to the disjoint satisfiability property of Definition 2.4.3 on page 53, we can merge them into a unique UCT \mathcal{U}_φ that supplies the dependence map labeling of internal components $\mathcal{U}_\phi^{\text{Ac}}$, by using the two functions head and body contained into its labeling. Moreover, observe that the final automaton runs on a b -bounded decision tree, where b is obtained from Theorem 2.4.3 on page 55 on the bounded-tree model property.

Theorem 2.4.4 (SL[1G] Automaton) *Let φ be an SL[1G] sentence. Then, there exists an UCT \mathcal{U}_φ such that φ is satisfiable iff $L(\mathcal{U}_\varphi) \neq \emptyset$.*

Finally, by a simple calculation of the size of \mathcal{U}_φ and the complexity of the related emptiness problem, we state in the next theorem the precise computational complexity of the satisfiability problem for SL[1G].

Theorem 2.4.5 (SL[1G] Satisfiability) *The satisfiability problem for SL[1G] is 2EXPTIME-COMplete.*

Proof. By Theorem 2.4.4 of SL[1G] automaton, to verify whether an SL[1G] sentence φ is satisfiable we can calculate the emptiness of the UPT \mathcal{U}_φ . This automaton is obtained by merging all UCTs $\mathcal{U}_\phi^{\text{Ac}}$, with $\phi = \wp b \psi \in \text{psnt}(\varphi)$, which in turn are based on the UCTs $\mathcal{U}_{b \psi}^{\text{Ac}}$ that embed the UCws \mathcal{U}_ψ . By a simple calculation, it is easy to see that \mathcal{U}_φ has $2^{O(|\varphi|)}$ states.

Now, by using a well-known nondeterminization procedure for APTs [MS95], we obtain an equivalent NPT \mathcal{N}_φ with $2^{2^{O(|\varphi|)}}$ states and index $2^{O(|\varphi|)}$.

The emptiness problem for such a kind of automaton with n states and index h is solvable in time $O(n^h)$. Thus, we get that the time complexity of checking whether φ is satisfiable is $2^{2^{O(|\varphi|)}}$. Hence, the membership of the satisfiability problem for SL[1G] in 2EXPTIME directly follows. Finally the thesis is proved, by getting the relative lower bound from the same problem for CTL* [VS85]. \square

Synthesis

The chapter is organized as follows. In Section 3.1 we provide the basic notions and formal definitions of rational synthesis. Then, in Section 3.2 we introduce and solve the problem of *qualitative rational synthesis*, showing that it is in 2EXPTIME-COMplete. Finally, in Section 3.3 we address also a quantitative variant of rational synthesis, proving that it is in 2EXPTIME-COMplete, as well as the qualitative case.

3.1 Rational synthesis

We define two variants of rational synthesis. The first, *cooperative rational synthesis*, was introduced in [FKL10]. The second, *non-cooperative rational synthesis*, was introduced in [KPV14].

We work with the following model: the world consists of a *system* and an environment composed of k agents: $\alpha_1, \dots, \alpha_k$. For uniformity, we refer to the system as Agent α_0 . We assume that Agent α_i controls a set X_i of propositions, and the different sets are pairwise disjoint. At each point in time, each agent sets his propositions to certain values. Let $X = \bigcup_{0 \leq i \leq k} X_i$, and $X_{-i} = X \setminus X_i$. Each agent α_i (including the system) has an objective φ_i , specified as an LTL formula over X .

This setting induces the CGS $\mathcal{G} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_0 \rangle$ defined as follows. The set of agents $\text{Ag} = \{\alpha_0, \alpha_1, \dots, \alpha_k\}$ consists of the system and the agents that constitute the environment. The actions of Agent α_i are the possible assignments to its variables. Thus, $\text{Ac}_i = 2^{X_i}$. We use Ac and Ac_{-i} to denote the sets 2^X and $2^{X_{-i}}$, respectively. The nodes of the game record the current assignment to the variables. Hence, $\text{St} = \text{Ac}$, and for all $s \in \text{St}$ and $\langle c_0, \dots, c_k \rangle \in \text{Ac}_0 \times \text{Ac}_1 \times \dots \times \text{Ac}_k$, we have $\delta(s, c_0, \dots, c_k) = \langle c_0, \dots, c_k \rangle$.

A strategy for the system is a function $f_0 : \text{Trk} \rightarrow \text{Ac}_0$. In the standard synthesis problem, we say that f_0 realizes φ_0 if, no matter which strategies the agents composing the environment follow, all the paths in which the system follows f_0 satisfy φ_0 . In rational synthesis, on instead, we assume that the agents that constitute the environment are rational, which soften the universal quantification on the behavior of the environment.

Recall that the rational-synthesis problem gets a solution concept as a parameter. As discussed in the introduction, the fact that a strategy profile is a solution with respect to the concept guarantees that it is not worthwhile for the agents constituting the environment to deviate from the strategies assigned to them. Several solution concepts are studied and motivated in game theory. Here, we focus on the concepts of *dominant strategy* and *Nash equilibrium*, defined below.

The common setting in game theory is that the objective for each agent is to maximize his *payoff* – a real number that is a function of the outcome of the game. We use $\text{payoff}_i :$




Figure 3.1: A game.

$\text{Path} \rightarrow \mathbb{R}$ to denote the payoff function of Agent α_i . That is, payoff_i assigns to each possible path π a real number $\text{payoff}_i(\pi)$ expressing the payoff of α_i on π . For a strategy profile P , we use $\text{payoff}_i(P)$ to abbreviate $\text{payoff}_i(\text{play}(P, s_0))$. In the case of an LTL goal ψ_i , we define $\text{payoff}_i(\pi) = 1$ if $\pi \models \psi_i$ and $\text{payoff}_i(\pi) = 0$, otherwise.

The simplest and most appealing solution concept is dominant-strategies solution [OR94]. A *dominant strategy* is a strategy that an agent can never lose by adhering to, regardless of the strategies of the other agents. Therefore, if there is a profile of strategies $P = \langle f_0, \dots, f_k \rangle$ in which all strategies f_i are dominant, then no agent has an incentive to deviate from the strategy assigned to him in P . Formally, P is a *dominant strategy profile* if for every $1 \leq i \leq k$ and for every (other) profile P' , we have that $\text{payoff}_i(P') \leq \text{payoff}_i(P'[i \leftarrow f_i])$.

As an example, consider the game in Figure 3.1, played by three agents, Alice, Bob, and Charlie, whose actions are $\{a_1, a_2\}$, $\{b_1, b_2\}$, and $\{c_1, c_2\}$, respectively. The arrows are labeled with the possible action of the agents. Each agent wants to visit a state marked with his initial letter, infinitely often. In this game, the strategy for Bob of always choosing b_2 on his node 2 is dominant, while all the possible strategies for Charlie are dominant. On the other hand, Alice has no dominant strategies, since her goal essentially depends on the strategies adopted by the other agents. In several games, it can happen that agents have not any dominant strategy. For this reason, one would consider also other kind of solution concepts.

Another well known solution concept is Nash equilibrium [OR94]. A strategy profile is a *Nash equilibrium* if no agent has an incentive to deviate from his strategy in P provided that the other agents adhere to the strategies assigned to them in P . Formally, P is a *Nash equilibrium profile* if for every $1 \leq i \leq k$ and for every (other) strategy f'_i for agent α_i , we have that $\text{payoff}_i(P[i \leftarrow f'_i]) \leq \text{payoff}_i(P)$. An important advantage of Nash equilibrium is that it is more likely to exist than an equilibrium of dominant strategies [OR94]¹. A weakness of Nash equilibrium is that it is not nearly as stable as a dominant-strategy solution: if one of the other agents deviates from his assigned strategy, nothing is guaranteed.

For the case of repeated-turn games like infinite games, a suitable refinement of Nash Equilibria is the *Subgame perfect Nash-equilibrium* [Sel75] (SPE, for short). A strategy profile $P = \langle f_0, \dots, f_k \rangle$ is an SPE if for every possible history of the game, no agent α_i has an incentive to deviate from her strategy f_i , assuming that the other agents follow their strategies in P . Intuitively, an SPE requires the existence of a Nash Equilibrium for each subgame starting from a randomly generated finite path of the original one. In [FKL10], the authors have studied cooperative rational synthesis also for the solution concept of SPE. To

¹In particular, all k -agent turn-based games with ω -regular objectives have Nash equilibrium [CMJ04].

do this, the synthesis algorithm in [FKL10] was extended to consider all possible histories of the game. In SL such a path property can be expressed combining strategy quantifiers with temporal operators. Indeed, the formula $\varphi = \llbracket \vec{x} \rrbracket \mathbf{b}(\vec{x}) \mathbf{G} \psi(\vec{y})$, with $\text{free}(\varphi) = \vec{y}$, states that, for all possible profile strategies the agents can follow, the game always reaches a position in which the formula $\psi(\vec{y})$ holds. Thus, for all possible paths that can be generated by agents, the property holds. By replacing $\psi(\vec{y})$ with the above formula, we then obtain a formula that represents SPEs. Hence, the cooperative and non-cooperative synthesis problems can be asserted in SL also for SPE, and our results hold also for this solution concept.

In rational synthesis, we control the strategy of the system and assume that the agents that constitute the environment are rational. Consider a strategy profile $P = \langle f_0, \dots, f_k \rangle$ and a solution concept γ (that is, γ is either “dominant strategies” or “Nash equilibrium”). We say that P is *correct* if $\text{play}(P)$ satisfies φ_0 . We say that P is in a f_0 -fixed γ -equilibrium if the agents composing the environment have no incentive to deviate from their strategies according to the solution concept γ , assuming that the system continues to follow f_0 . Thus, P is a f_0 -fixed dominant-strategy equilibrium if for every $1 \leq i \leq k$ and for every (other) profile P' in which Agent 0 follows f_0 , we have that $\text{payoff}_i(P') \leq \text{payoff}_i(P'[i \leftarrow f_i])$. Note that for the case of Nash equilibrium, adding the requirement that P is f_0 -fixed does not change the definition of an equilibrium.

In the context of objectives in LTL, we assume the following simple payoffs. If the objective φ_i of Agent α_i holds, then his payoff is 1, and if φ_i does not hold, then the payoff of Agent i is 0. Accordingly, $P = \langle f_0, \dots, f_k \rangle$ is in a dominant-strategy equilibrium if for every $1 \leq i \leq k$ and profile $P' = \langle f'_0, \dots, f'_k \rangle$ with $f'_0 = f_0$, if $\text{play}(P') \models \varphi_i$, then $\text{play}(P'[i \leftarrow f_i]) \models \varphi_i$. Also, P is in a Nash-equilibrium if for every $1 \leq i \leq k$ and strategy f'_i , if $\text{play}(P[i \leftarrow f'_i]) \models \varphi_i$, then $\text{play}(P) \models \varphi_i$.

Definition 3.1.1 (Rational synthesis) *The input to the rational-strategy problem is a set X of atomic propositions, partitioned into X_0, \dots, X_k , LTL formulas $\varphi_0, \dots, \varphi_k$, describing the objectives of the system and the agents composing the environment, and a solution concept γ . We distinguish between two variants of the problem:*

1. *In Cooperative rational synthesis [FKL10], the desired output is a strategy profile P such that $\text{play}(P)$ satisfies φ_0 and P is a f_0 -fixed γ -equilibrium.*
2. *In Non-cooperative rational synthesis, the desired output is a strategy f_0 for the system such that for every strategy profile P that includes f_0 and is a f_0 -fixed γ -equilibrium, we have that $\text{play}(P)$ satisfies φ_0 .*

Thus, in the cooperative variant of [FKL10], we assume that once we suggest to the agents in the environment strategies that are in a γ -equilibrium, they will adhere to the suggested strategies. In the non-cooperative variant we introduce here, the agents may follow any strategy profile that is in a γ -equilibrium, and thus we require the outcome of all these profiles to satisfy φ_0 . It is shown in [FKL10] that the cooperative rational synthesis problem is 2EXPTIME-COMPLETE.

Note that the input to the rational synthesis problem may not have a solution, so when we solve the rational-synthesis problem, we first solve the *rational-realizability* problem,

which asks if a solution exists. As with classical synthesis, the fact that SL model-checking algorithms can be easily modified to return a regular witness for the involved strategies in case an existentially quantified strategy exists, makes the realizability and synthesis problems strongly related.

Example 3.1.1 Consider a file-sharing network with the system and an environment consisting of two agents. The system controls the signal d_1 and d_2 (Agent α_1 and α_2 can download, respectively) and it makes sure that an agent can download only when the other agent uploads. The system's objective is that both agents will upload infinitely often. Agent α_1 controls the signal u_1 (Agent α_1 uploads), and similarly for Agent α_2 and u_2 . The goal of both agents is to download infinitely often.

Formally, the set of atomic propositions is $X = \{d_1, d_2, u_1, u_2\}$, partitioned into $X_0 = \{d_1, d_2\}$, $X_1 = \{u_1\}$, and $X_2 = \{u_2\}$. The objectives of the system and the environment are as follows.

- $\varphi_0 = G(\neg u_1 \rightarrow \neg d_2) \wedge G(\neg u_2 \rightarrow \neg d_1) \wedge GFu_1 \wedge GFu_2$,
- $\varphi_1 = GFd_1$,
- $\varphi_2 = GFd_2$.

First, note that in standard synthesis, φ_0 is not realizable, as a hostile environment needs not upload. In the cooperative setting, the system can suggest to both agents the following TIT FOR TAT strategy: upload at the first time step, and from that point onward upload iff the other agent uploads. The system itself follows a strategy f_0 according to which it enables downloads whenever possible (that is, d_2 is valid whenever Agent α_1 uploads, and d_1 is valid whenever Agent α_2 uploads). It is not hard to see that the above three strategies are all dominant. Indeed, all the three objectives are satisfied. Thus, the triple of strategies is a solution for the cooperative setting, for both solution concepts.

What about the non-cooperative setting? Consider the above strategy f_0 of the system, and consider strategies for the agents that never upload. The tuple of the three strategies is in a f_0 -fixed Nash equilibrium.

This ensures strategies for the environment to be dominant. Indeed, if Agent α_2 changes her strategy, φ_1 is still satisfied and vice-versa. Indeed, as long as Agent α_2 sticks to his strategy, Agent α_1 has no incentive to change his strategy, and similarly for Agent α_2 . Thus, f_0 is not a solution to the non-cooperative rational synthesis problem for the solution concept of Nash equilibrium. On the other hand, we claim that f_0 is a solution to the non-cooperative rational synthesis problem for the solution concept of dominant strategies. For showing this, we argue that if f_1 and f_2 are dominant strategies for α_1 and α_2 , then φ_0 is satisfied in the path that is the outcome of the profile $P = \langle f_0, f_1, f_2 \rangle$. To see this, consider such a path $\pi = \text{play}(f_0, f_1, f_2)$. We necessarily have that $\pi \models \varphi_1 \wedge \varphi_2$. Indeed, otherwise f_1 and f_2 would not be dominant, as we know that with the strategies described above, α_1 and α_2 can satisfy their objectives. Now, since $\pi \models \varphi_1 \wedge \varphi_2$, we know that u_2 and u_1 hold infinitely often in π . Also, it is not hard to see that the formulas $G(\neg u_1 \rightarrow \neg d_2)$ and $G(\neg u_2 \rightarrow \neg d_1)$ are always satisfied in the context of f_0 , no matter how the other agents behave. It follows that $\pi \models \varphi_0$, thus f_0 is a solution of the non-cooperative rational synthesis problem for dominant strategies.

3.2 Qualitative Rational Synthesis

In this section we study cooperative and non-cooperative rational synthesis and show that they can be reduced to the model-checking problem for $SL[NG]$. The cooperative and non-cooperative rational synthesis problems for several solution concepts can be stated in $SL[NG]$.

We first show how to state that a given strategy profile $\vec{y} = (y_0, \dots, y_k)$ is in a y_0 -fixed γ -equilibrium. For $\alpha_i \in Ag$, let φ_i be the objective of Agent α_i . Moreover, for a given strategy profile \vec{y} , by $b(\vec{y}) = (\alpha, y_0) \dots (\alpha, y_n)$ we denote the binding prefix assigning each strategy to the corresponding agent, and by $b(\vec{y}_{-i}, z_i) = (\alpha, y_0) \dots (\alpha, y_{i-1})(\alpha, z_i)(\alpha, y_{i+1}) \dots (\alpha, y_n)$ we denote the binding prefix assigning the each strategy to the corresponding agent but α , to which it assigns the strategy z_i . For a solution concept γ and a strategy profile $\vec{y} = (y_0, \dots, y_k)$, the formula $\varphi^\gamma(\vec{y})$, expressing that the profile \vec{y} is a y_0 -fixed γ -equilibrium, is defined as follows.

- For the solution concept of dominant strategies, we define:

$$\varphi^\gamma(\vec{y}) := \llbracket \vec{z} \rrbracket \bigwedge_{i=1}^k (b(\vec{z})\varphi_i \rightarrow b(\vec{z}_{-i}, y_i))\varphi_i.$$

- For the solution concept of Nash equilibrium, we define:

$$\varphi^\gamma(\vec{y}) := \llbracket \vec{z} \rrbracket \bigwedge_{i=1}^k (b(\vec{y}_{-i}, z_i)\varphi_i \rightarrow b(\vec{y})\varphi_i).$$

- For the solution concept of Subgame Perfect Equilibrium, we define:

$$\varphi^\gamma(\vec{y}) := \llbracket \vec{x} \rrbracket b(\vec{x}_{-0}, y_0)F \bigwedge_{i=1}^k \llbracket z_i \rrbracket b(\vec{y}_{-i}, z_i)\varphi_i \rightarrow b(\vec{y})\varphi_i.$$

We can now state the existence of a solution to the cooperative and non-cooperative rational-synthesis problem, respectively, with input $\varphi_0, \dots, \varphi_k$ by the closed formulas:

1. $\varphi_{cRS}^\gamma := \langle\langle y_0 \rangle\rangle \langle\langle y_1 \rangle\rangle \dots \langle\langle y_k \rangle\rangle (\varphi^\gamma(\vec{y}) \wedge \varphi_0)$;
2. $\varphi_{noncRS}^\gamma := \langle\langle y_0 \rangle\rangle \llbracket y_1 \rrbracket \dots \llbracket y_k \rrbracket (\varphi^\gamma(\vec{y}) \rightarrow \varphi_0)$.

Indeed, the formula 1 specifies the existence of a strategy profile $P = \langle f_0, \dots, f_k \rangle$ that is f_0 -fixed γ -equilibrium and such that the outcome satisfies φ_0 . On the other hand, the formula 2 specifies the existence of a strategy f_0 for the system such that the outcome of all profiles that are in a f_0 -fixed γ -equilibrium satisfy φ_0 .

As shown above, all the solution concepts we are taking into account can be specified in $SL[NG]$ with formulas whose length is polynomial in the number of the agents and in which the alternation depth of the quantification is 1. Hence we can apply the known complexity results for $SL[NG]$:

Theorem 3.2.1 (Cooperative and non-cooperative rational-synthesis complexity) *The cooperative and non-cooperative rational-synthesis problems in the qualitative setting are 2EXPTIME-complete.*

Proof. Consider an input $\varphi_0, \dots, \varphi_k, X$, and γ to the cooperative or non-cooperative rational-synthesis problem. As explained in Section 3.1, the input induces a game \mathcal{G} with nodes in 2^X . As detailed above, there is a solution to the cooperative (resp., non-cooperative)

problem iff the SL[NG] formula φ_{cRS}^γ (resp., φ_{noncRS}^γ) is satisfied in \mathcal{G} . The upper bound then follows from the fact that the model checking problem for SL[NG] formulas of alternation depth 1 is in 2EXPTIME in the size of the formula (Cf., Theorem 2.3.4 on page 47). Moreover, the model-checking algorithm can return finite-state transducers that model strategies that are existentially quantified.

For the lower bound, it is easy to see that the classical LTL synthesis problem is a special case of the cooperative and non-cooperative rational synthesis problem. Indeed, $\varphi(I, O)$ is realizable against a hostile environment iff the solution to the non-cooperative rational synthesis problem for a system that has an objective φ and controls I and an environment that consists of a single agent that controls O and has an objective *True*, is positive. \square

3.3 Quantitative Rational Synthesis

As discussed in this chapter introduction, a weakness of classical synthesis algorithms is the fact that specifications are Boolean and miss a reference to the quality of the satisfaction. Applications of game theory consider games with quantitative payoffs. Thus, even more than the classical setting, the rational setting calls for an extension of the synthesis problem to a quantitative setting. In this section we introduce *Objective LTL*, a quantitative extension of LTL, and study an extension of the rational-synthesis problem for specifications in Objective LTL. As opposed to other multi-valued specification formalisms used in the context of synthesis of high-quality systems [BCHJ09, ABK13], Objective LTL uses the syntax and semantics of LTL and only augments the specification with a reward function that enables a succinct and convenient prioritization of sub-specifications.

Objective LTL (OLTL, for short) is an extension of LTL in which specifications consist of sets of LTL formulas weighted by functions. Formally, an OLTL *specification* over a set X of atomic propositions is a pair $\theta = \langle \Psi, f \rangle$, where $\Psi = \langle \psi_1, \psi_2, \dots, \psi_m \rangle$ is a tuple of LTL formulas over X and $f : \{0, 1\}^m \rightarrow \mathbb{Z}$ is a *reward function*, mapping Boolean vectors of length m to integers. We assume that f is given by a polynomial. We use $|\Psi|$ to denote $\sum_{i=1}^m |\psi_i|$. For a computation $\pi \in (2^X)^\omega$, the *signature of Ψ in π* , denoted $\text{sig}(\Psi, \pi)$, is a vector in $\{0, 1\}^m$ indicating which formulas in Ψ are satisfied in π . Thus, $\text{sig}(\Psi, \pi) = \langle v_1, v_2, \dots, v_m \rangle$ is such that for all $1 \leq i \leq m$, we have that $v_i = 1$ if $\pi \models \psi_i$ and $v_i = 0$ if $\pi \not\models \psi_i$. The *value of θ in π* , denoted $\text{val}(\theta, \pi)$ is then $f(\text{sig}(\Psi, \pi))$. Thus, the interpretation of an OLTL specification is quantitative. Intuitively, $\text{val}(\theta, \pi)$ indicates the level of satisfaction of the LTL formulas in Ψ in π , as determined by the priorities induced by f . We note that the input of weighted LTL formulas studied in [CY98] is a special case of Objective LTL.

Example 3.3.1 Consider a system with m buffers, of capacities c_1, \dots, c_m . Let full_i , for $1 \leq i \leq m$, indicate that buffer i is full. The OLTL specification $\theta = \langle \Psi, f \rangle$, with $\Psi = \langle \text{Ffull}_1, \text{Ffull}_2, \dots, \text{Ffull}_m \rangle$ and $f(v) = c_1 \cdot v_1 + \dots + c_m \cdot v_m$ enables us to give different satisfaction values to the objective of filling a buffer. Note that $\text{val}(\langle \Psi, f \rangle, \pi)$ in a computation π is the sum of capacities of all filled buffers.

In the quantitative setting, the objective of Agent α is given by means of an OLTL specification $\theta_i = \langle \Psi_i, f_i \rangle$, specifications describe the payoffs to the agents, and the objective

of each agent (including the system) is to maximize his payoff. For a strategy profile P , the payoff for Agent α in P is simply $\text{val}(\theta_\alpha, \text{play}(P))$.

In the quantitative setting, rational synthesis is an optimization problem. Here, in order to easily solve it, we provide a decision version by making use of a threshold. It is clear that the optimization version can be solved by searching for the best threshold in the decision one.

Definition 3.3.1 (Quantitative rational synthesis) *The input to the quantitative rational-strategy problem is a set X of atomic propositions, partitioned into X_0, \dots, X_k , OLTL specifications $\theta_0, \theta_1, \dots, \theta_k$, with $\theta_i = \langle \Psi_i, f_i \rangle$, and a solution concept γ . We distinguish between two variants of the problem:*

1. *In cooperative quantitative rational synthesis, the desired output for a given threshold $t \in \mathbb{N}$ is a strategy profile P such that $\text{payoff}_0(P) \geq t$ and P is in a f_0 -fixed γ -equilibrium.*
2. *In non-cooperative quantitative rational synthesis, the desired output for a given threshold $t \in \mathbb{N}$ is a strategy f_0 for the system such that, for each strategy profile P including f_0 and being in a f_0 -fixed γ -equilibrium, we have that $\text{payoff}_0(P) \geq t$.*

Now, we introduce some auxiliary formula that helps us to formulate also the quantitative rational-synthesis problem in $\text{SL}[\text{NG}]$.

For a tuple $\Psi = \langle \psi_1, \dots, \psi_m \rangle$ of LTL formulas and a signature $v = \{v_1, \dots, v_m\} \in \{0, 1\}^m$, let $\text{mask}(\Psi, v)$ be an LTL formula that characterizes computations π for which $\text{sig}(\Psi, \pi) = v$. Thus, $\text{mask}(\Psi, v) = (\bigwedge_{i:v_i=0} \neg\psi_i) \wedge (\bigwedge_{i:v_i=1} \psi_i)$.

We adjust the SL formulas $\Phi^\gamma(\vec{y})$ described in Section 3.2 on page 65 to the quantitative setting. Recall that $\Phi^\gamma(\vec{y})$ holds iff the strategy profile assigned to \vec{y} is in a f_0 -fixed γ -equilibrium. There, the formula is a conjunction over all agents in $\{\alpha, \dots, \alpha\}$, stating that Agent α does not have an incentive to change his strategy. In our quantitative setting, this means that the payoff of Agent α in an alternative profile is not bigger than his payoff in \vec{y} . For two strategy profiles, assigned to \vec{y} and \vec{y}' , an SL formula that states that Agent α has no incentive that the profile would change from \vec{y} to \vec{y}' can state that the signature of Ψ in $\text{play}(\vec{y}')$ results in a payoff to Agent α that is smaller than his current payoff. Formally, we have that:

$$\Phi_i^{eq}(\vec{y}, \vec{y}') = \bigvee_{v \in \{1, \dots, m_i\}: f_i(v) \leq \text{payoff}_i(\vec{y})} b(\vec{y}') \text{mask}(\Psi_i, v).$$

We can now adjust $\Phi^\gamma(\vec{y})$ for all the cases of solution concepts we are taking into account.

- For the solution concept of dominant strategies, we define:

$$\Phi^\gamma(\vec{y}) := \bigwedge_{i \in \{1, \dots, n\}} \llbracket \vec{z} \rrbracket \Phi_i^{eq}(\vec{y}, (\vec{z}[\alpha \leftarrow y_0]));$$

- For the solution concept of Nash equilibrium, we define:

$$\Phi^\gamma(\vec{y}) := \bigwedge_{i \in \{1, \dots, n\}} \llbracket \vec{z} \rrbracket \Phi_i^{eq}(\vec{y}, (\vec{y}[\alpha \leftarrow z_i]));$$

- For the solution concept of Subgame Perfect Equilibrium, we define:

$$\varphi^\gamma(\vec{y}) := \llbracket \vec{x} \rrbracket b(\vec{x}[\alpha \leftarrow y_0]) \text{F} \bigwedge_{i \in \{1, \dots, n\}} \llbracket \vec{z} \rrbracket \Phi_i^{eq}(\vec{y}, (\vec{z}[\alpha \leftarrow y_0])).$$

Once we adjust $\Phi^\gamma(\vec{y})$ to the quantitative setting, we can use the same SL formula used in the non-quantitative setting to state the existence of a solution to the rational synthesis problem. We have the following:

- $\Phi_{RS}^\gamma := \langle\langle y_0 \rangle\rangle \langle\langle y_1 \rangle\rangle \dots \langle\langle y_k \rangle\rangle (\Phi^\gamma(\vec{y}) \wedge \varphi_0)$;
- $\Phi_{nonRS}^\gamma := \langle\langle y_0 \rangle\rangle \llbracket y_1 \rrbracket \dots \llbracket y_k \rrbracket (\Phi^\gamma(\vec{y}) \rightarrow \varphi_0)$.

Theorem 3.3.1 *The cooperative and non-cooperative quantitative rational-synthesis problems are 2EXPTIME-COMPLETE.*

Proof. We can reduce the problems to the model-checking problem of the SL formulas Φ_{RS}^γ and Φ_{nonRS}^γ , respectively. We should, however, take care when analyzing the complexity of the procedure, as the formulas $\Phi_i^{eq}(\vec{y}, \vec{y}')$, which participate in Φ_{RS}^γ and Φ_{nonRS}^γ involve a disjunction over vectors in $\{0, 1\}^{m_i}$, resulting in Φ_{nonRS}^γ of an exponential length.

While the above prevents us from using the doubly exponential known bound on SL model checking for formulas of alternation depth 1 as is, it is not difficult to observe that the run time of the model-checking algorithm in [MMPV14], when applied to Φ_{nonRS}^γ , is only doubly exponential. The reason is the fact that the inner exponent paid in the algorithm for SL model checking is due to the blow-up in the translation of the formula to a nondeterministic Büchi automaton over words (NBW, for short). In this translation, the exponentially many disjuncts are dominated by the exponential translation of the innermost LTL to NBW. Thus, the running time of the algorithm is doubly exponential, and it can return the witnessing strategies.

Hardness in 2EXPTIME follows easily from hardness in the non-quantitative setting. □

Pushdown ATL[★]

The chapter is organized as follows. In Section 4.1 we introduce the definition of PGS, where we also provide an example that helps clarifying the setting. There, we also show how a PGS can be embedded into a (infinite-state) CGS. In Section 4.2 we recall the syntax and the semantics of ATL[★] over CGSs and then, exploiting the results in the previous section, we define the model-checking problem of ATL[★] over PGSs. Finally, in Section 4.3 we show that this problem can be solved in 3EXPTIME by means of an automata-theoretic approach. There, we also show a 2EXPSpace-HARD lower-bound.

4.1 Pushdown Games

Classically, ATL[★] formulas are interpreted over *Concurrent Game Structures* [AHK02]. Here, we instead interpret ATL[★] formulas over a new semantic framework, which we call *Pushdown Game Structure*. Intuitively, this new formalism provides a concurrent game structure in which a stack is added and the labeling and transition functions depend on its content. In this section, we also show that every pushdown game structure can be transformed into a suitable concurrent game structure, so providing the required interpretation of ATL[★] formulas over the former. However, note that the latter requires a infinite number of states, used to represent all the possible configurations the pushdown system can enter. We start with the definition of pushdown game structures.

Definition 4.1.1 (Pushdown Game Structure) A Pushdown Game Structure (PGS, for short) is a tuple $\mathcal{P} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{Loc}, \Gamma, \text{tr}, \text{ap}, l_0 \rangle$, where AP, Ag, Ac, Loc and Γ are finite sets of atomic propositions, agents, actions, locations, and stack alphabet, respectively, $l_0 \in \text{Loc}$ is an initial location, and $\text{ap} : \text{Loc} \times \Gamma_{\perp} \rightarrow 2^{\text{AP}}$ is a labeling function, where $\Gamma_{\perp} = \Gamma \cup \{\perp\}$ and \perp is the special bottom stack symbol not contained in Γ . Let $\text{Dc} \triangleq \text{Ac}^{\text{Ag}}$ be the set of decisions, i.e., functions from Ag to Ac representing the action choices for each agent. Then, $\text{tr} : \text{Loc} \times \Gamma_{\perp} \times \text{Dc} \rightarrow \text{Loc} \times \Gamma_{\perp}^*$ is a transition function mapping a location, a stack symbol, and a decision to a location and a word in the stack alphabet.

A pair $s = (l, \alpha) \in \text{St} \triangleq \text{Loc} \times (\Gamma^* \cdot \{\perp\})$ is called *state* or *configuration*. We write $\text{top}(\alpha)$ for the left most symbol of α and call it the *top* of the stack content α . The PGS moves according to the transition function. This means that, if it is in the location l , the top of the stack content is γ , and the agents make a decision δ , then $\text{tr}(l, \gamma, \delta) = (l', \alpha)$ means that the execution moves to the location l' and the symbol γ is replaced with α on the top of the stack content. We assume that, if \perp is popped, then it is pushed right back, and that is the only case in which it is pushed. This means that \perp is always on the bottom of the stack and nowhere else. The stack containing only the symbol \perp is said to be *empty*.

As the stack has no a priori bound on its size, the set St is assumed to be possibly infinite. Saying this, it turns out that PGSs are *infinite-state multi-agent systems*.

The notion of labeling and transition can be lifted to states, as follows. For a state $s = (l, \alpha)$, we define $\text{ap}(s) = \text{ap}(l, \text{top}(\alpha))$. Moreover, for a decision $\delta \in \text{Dc}$, we define $\text{tr}((l, \gamma \cdot \alpha), \delta) = (l', \beta \cdot \alpha)$, with $(l', \beta) = \text{tr}(l, \gamma, \delta)$.

Note that for a classical *pop* we write the empty word ε on the stack. To make a classical *push* one has to first put back the read top symbol and then push the required word. The transition function also allows to perform in one step a *pop-push* operation that replaces the top stack symbol with the required word. To get familiar with PGSs, we give an example.

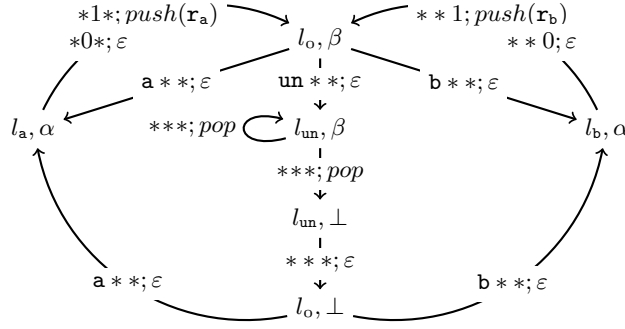


Figure 4.1: A Pushdown system scheduler.

Example 4.1.1 (Pushdown scheduler) Take a system consisting of two processes a and b that may access to a common resource via the respective requests r_a and r_b and a scheduler s that can grant in a LIFO order the processes requests, all memorized into a stack. As model we use a PGS $\mathcal{P} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{Loc}, \Gamma, \text{tr}, \text{ap}, l_o \rangle$, with $\text{AP} = \{r_a, r_b, g_a, g_b, a, b, e\}$, $\text{Ag} = \{s, a, b\}$, $\text{Ac} = \{a, b, \text{un}, 0, 1\}$, $\text{Loc} = \{l_o, l_a, l_b, l_{\text{un}}\}$, and $\Gamma = \{r_a, r_b\}$. The scheduler s controls the location l_o by means of the actions a , b , and un , standing for “ a can make a request”, “ b can make a request”, and “the system can unload the stack requests”, respectively. Accordingly, they lead to the locations l_a , l_b , and l_{un} . On l_a and l_b , the agents a and b , respectively, can either make a request via action 1 or skip it with action 0. In the former case, the request is recorded into the stack by writing the symbol r_x , for $x \in \{a, b\}$; otherwise, in the latter case there is no operation over the stack. Finally, the location l_{un} triggers the granting phase by emptying the stack. During this phase, neither a nor b can make any further request. This can be seen as a legitimate constraint by thinking how classical synchronizing and backup systems are designed.

The labeling function, for all $\gamma \in \Gamma_{\perp}$ and $x \in \{a, b\}$, is defined as follows: $\text{ap}(l_o, \perp) = \emptyset$, $\text{ap}(l_o, r_x) = \{x\}$, $\text{ap}(l_x, \gamma) = \{x\}$, $\text{ap}(l_{\text{un}}, r_x) = \{g_x\}$, and $\text{ap}(l_{\text{un}}, \perp) = \{e\}$. Intuitively, propositions a and b means that agents a and b are authorized to make a request, respectively. The proposition r_x , instead, occurs when the corresponding request has been just made by agent x . On the other hand, the proposition g_x occurs when the request r_x has been just granted. Finally, e indicates that the unloading phase is terminated and so the stack is empty.

The transition function tr is described directly in Figure 4.1. The labeling of the edges have the following meaning. First, note that it is composed of two parts separated by a semi-column. The left part represents the decision of the agents, given in the order $s < a < b$. The right part represents the stack operation. As an example, the label $*1*; \text{push}(r_a)$ says that agent a is making a request r_a and the symbol r_a is pushed on the stack, where the symbol $*$ denotes any possible action for the other agents. The nodes represent all possible states. Note that, for the locations l_a and l_b we have collapsed the two possible configurations

with $\beta = \perp$ and $\beta \neq \perp$ since the transition over them does not depend on the stack content. Finally, observe that the stack is unbounded and so an execution might generate an infinite number of distinguished states. Also, observe that the stack is fundamental to keep track of the order in which the requests appear.

Clearly, a PGS $\mathcal{P} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{Loc}, \Gamma, \text{tr}, \text{ap}, l_o \rangle$ can be suitably turned into a CGS $\mathcal{G}_{\mathcal{P}} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_o \rangle$, where $\text{St} = \text{Loc} \times \Gamma_{\perp}^*$, $s_o = (l_o, \perp)$, and the functions ap and tr are the lifting on states of the corresponding functions in \mathcal{P} . Intuitively, the states of \mathcal{G} are used to implicitly represent both the current location and store the stack content. Despite this, it is important to observe that, while a PGS has a finite number of control locations, the corresponding CGS necessarily has an infinite number of control states, as the number of different stack contents is unbounded.

4.2 ATL* over pushdown systems

In this section, we recall the syntax of ATL* and introduce its semantics over PGS via its representation in terms of CGS (with infinite states). We start with the definition of ATL* syntax.

Definition 4.2.1 (ATL* Syntax) *ATL* formulas are built inductively from the set of atomic propositions AP and agents Ag, by using the following grammar, where $p \in \text{AP}$ and $A \subseteq \text{Ag}$:*

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid X\varphi \mid \varphi U\varphi \mid \langle\langle A \rangle\rangle\varphi.$$

A *sentence* is a Boolean combination of ATL* formulas of the form $\langle\langle A \rangle\rangle\psi$. Intuitively, $\langle\langle A \rangle\rangle\psi$ means that each agent in A has a strategy such that, whatever the other agents do, the resulting play satisfies ψ .

We now provide some examples of ATL* formulas that will be useful in the sequel. Precisely, we consider $\varphi_1 = \langle\langle \{s\} \rangle\rangle(\text{GF}a \wedge \text{GF}b \wedge \text{GF}e)$, $\varphi_2 = \langle\langle \emptyset \rangle\rangle(\text{GF}a \wedge \text{GF}b \wedge \text{GF}e)$ and $\varphi_3 = \langle\langle \emptyset \rangle\rangle((r_a \wedge X(\neg eUr_b)) \rightarrow (\text{F}g_b \rightarrow X\text{F}g_a))$ over the sets AP and Ag given in Example 4.1.1 on the preceding page. The formula φ_1 states that agent s has a way to let propositions a , b , and e to occur infinitely often. The formula φ_2 states that, no matters how the agents behave, propositions a , b , and e occur infinitely often. Finally, the formula φ_3 states that, whenever a request r_b occurs after an r_a one in the same loading phase, then, if r_b is eventually granted, then r_a is granted later on as well.

We now provide the semantics of ATL*.

Definition 4.2.2 (ATL* Semantics) *For a CGS $\mathcal{G} = \langle \text{AP}, \text{Ag}, \text{Ac}, \text{St}, \text{tr}, \text{ap}, s_o \rangle$ and a path $\pi \in \text{Pth}$, the modeling relation $\mathcal{G}, \pi \models \varphi$ is inductively defined as follows:*

- *The atomic and boolean cases are defined as usual;*
- $\mathcal{G}, \pi \models \langle\langle A \rangle\rangle\varphi$ *if there is an assignment χ_A such that $\mathcal{G}, \pi' \models \varphi$, for all $\pi' \in \text{play}(\chi_A, (\pi)_0)$;*
- $\mathcal{G}, \pi \models X\varphi$ *if $\mathcal{G}, (\pi)_{\geq 1} \models \varphi$;*

- $\mathcal{G}, \pi \models \varphi_1 \cup \varphi_2$ if there exists $j \in \mathbb{N}$ such that $\mathcal{G}, (\pi)_{\geq i} \models \varphi_1$, for all $i < j$, and $\mathcal{G}, (\pi)_{\geq j} \models \varphi_2$.

For a sentence φ and two paths π_1, π_2 with $(\pi_1)_0 = (\pi_2)_0$, it holds that $\mathcal{G}, \pi_1 \models \varphi$ iff $\mathcal{G}, \pi_2 \models \varphi$. Indeed, according to the semantics of the existential quantification $\langle\langle A \rangle\rangle\psi$, the only element of the path to take into account is the first one. For this reason, for a sentence φ we write $\mathcal{G}, s \models \varphi$ if $\mathcal{G}, \pi \models \varphi$ for some $\pi \in \text{Pth}(s)$. Finally, we say that \mathcal{G} satisfies φ , and write $\mathcal{G} \models \varphi$, if $\mathcal{G}, s_0 \models \varphi$.

To get familiar with the semantics, consider the PGS \mathcal{P} given in Example 4.1.1 on page 70 and the formulas φ_1, φ_2 , and φ_3 given above. It is easy to see that $\mathcal{G}_{\mathcal{P}} \models \varphi_1$. Indeed, the strategy f_1 that allows the scheduler to pick infinitely often the actions **a**, **b**, and **un**, makes all the generated paths to satisfy $\text{GFa} \wedge \text{GFb} \wedge \text{GFe}$. On the other hand, it is easy to see that $\mathcal{G}_{\mathcal{P}} \not\models \varphi_2$. Indeed, the strategy f_2 for **s** such that $f_2(\rho) = \mathbf{a}$, for all $\rho \in \text{Trk}$, never makes **b** and **e** to occur in the generated paths. Finally, the formula φ_3 simply states that, whatever the three agents in the system do, the unloading phases always grant the requests in the reverse order they occur.

Definition 4.2.3 (Model-checking) For a given PGS \mathcal{P} and an ATL* formula φ , the model-checking problem is to decide whether $\mathcal{G}_{\mathcal{P}} \models \varphi$.

4.3 Model Checking

In this section, we provide a 3EXPTIME upper-bound and a 2EXPSpace lower-bound for the model-checking problem of ATL* over PGS.

4.3.1 Upper-bound Complexity

For the upper bound we use an automata-theoretic approach. We start with some notation and the definition of nondeterministic pushdown automata. See [KPV02, KVV00] for more.

For a given set $D \subseteq \mathbb{N}$, a D -tree T is a prefix closed subset of D^* , i.e., a set in which, if $x \cdot d \in T$ then $x \in T$. The elements of T are called *nodes* and the empty word ε is called the *root* of T . For $x \in T$, the set of *children* of x is $\text{children}(T, x) \triangleq \{x \cdot i \in T : i \in D\}$. For $x, y \in T$, we write $x \prec y$ to mean that x is a proper prefix of y , i.e., there exists $z \in D^*$ such that $x \cdot z = y$. For $x \in T$, a path in T from x is a minimal set $\pi \subseteq T$ such that $x \in \pi$ and for each $y \in \pi$ such that $\text{children}(T, y) \neq \emptyset$, there exist exactly one node in $\text{children}(T, y)$ belonging to π . For an alphabet Σ , a Σ -labeled tree is a pair $\mathcal{T} = \langle T, V \rangle$ where T is a tree and $V : T \rightarrow \Sigma$ maps each node in T to an element in Σ . In the following, we mainly consider Σ to be the set power set 2^{AP} of atomic propositions AP.

A *Nondeterministic Pushdown Tree Automata* (PD-NTA, for short), over Σ -labeled trees, is a tuple $\mathcal{A} = \langle \Sigma, \Gamma, Q, q_0, \alpha_0, \delta, F \rangle$, where Σ and Γ are *finite* input and stack alphabet sets, Q is a finite set of states, q_0 is an initial state, $\alpha_0 \in \Gamma^* \cdot \perp$ is an initial stack content, $\delta : Q \times \Sigma \times \Gamma_{\perp} \rightarrow 2^{2^{(Q \times \Gamma^*)}}$ is a transition function such that, for all $(q, \sigma, \gamma) \in Q \times \Sigma \times \Gamma_{\perp}$, $\delta(q, \sigma, \gamma)$ is a *finite* set, and F is an acceptance condition over Q .

When the automaton is in a state q , reading an input node x labeled with $\sigma \in \Sigma$, and the stack contains a word $\gamma \cdot \alpha$ in Γ_{\perp}^* , it chooses, for some $k \in \mathbb{N}$, a tuple $\langle (q_1, \beta_1), \dots, (q_k, \beta_k) \rangle \in \delta(q, \sigma, \gamma)$ and splits in k copies such that, for each $1 \leq i \leq k$, the i -th copy is in the state q_i and the stack content is updated by removing γ and pushing β_i . Then, it reads the node $x \cdot i$.

A run of a PD-NTA on a Σ -labeled tree $\mathcal{T} = \langle T, V \rangle$ is a $(Q \times \Gamma_{\perp}^*)$ -labeled tree $\langle T, r \rangle$ such that $r(\varepsilon) = (q_0, \alpha_0)$ and for each $x \in T$ with $r(x) = (q, \gamma \cdot \alpha)$, there is a tuple $\langle (q_1, \beta_1), \dots, (q_k, \beta_k) \rangle \in \delta(q, V(x), \gamma)$ for some $k \in \mathbb{N}$, such that, for all $1 \leq i \leq k$, $r(x \cdot i) = (q_i, \beta_i \cdot \alpha)$ if $\gamma \neq \perp$, and $r(x \cdot i) = (q_i, \beta_i \cdot \perp)$, otherwise.

Given a path π starting from ε , by $\text{inf}_r(\pi)$ we denote the subset of states q such that there are infinitely many $x \in \pi$ such that $r(x) \in \{q\} \times \Gamma_{\perp}^*$. A path satisfies a *parity* condition $F = \{F_1, \dots, F_k\}$, with $F_i \subseteq F_{i+1}$, for all $i < k$, and $F_k = Q$, if the minimum index $1 \leq i \leq k$ such that $\text{inf}_r(\pi) \cap F_i \neq \emptyset$ is even. A run $\langle T, r \rangle$ is *accepting* if every path satisfies the acceptance condition. The PD-NTA \mathcal{A} accepts an input tree $\langle T, V \rangle$ iff there is an accepting run of \mathcal{A} over it. The language of \mathcal{A} , denoted by $\mathcal{L}(\mathcal{A})$, contains all the trees accepted by \mathcal{A} . The emptiness problem for PD-NTA is to decide, for a given \mathcal{A} , whether $\mathcal{L}(\mathcal{A}) = \emptyset$. In [KPV02] it is reported the following.

Theorem 4.3.1 *The emptiness problem for a parity PD-NTA is EXPTIME-COMPLETE.*

In several branching-time temporal-logic verification settings, the automata-theoretic approach has been fruitfully applied. Very close to our case are the procedures deployed for model checking pushdown systems over CTL^{*} specifications [BEM97, BMP10, Boz06] and finite-state CGSs over ATL^{*} specifications [AHK02]. The former is a top-down procedure that first builds an automaton accepting all the trees that satisfy the formula and then checks for the membership problem of the tree unwinding of the pushdown model. Precisely, to get a tight complexity, it starts with a single-exponential alternating¹ parity tree automaton and the membership problem results in a special alternating pushdown tree automaton named *one-letter*, with no blow-up in size, whose emptiness can be checked in exponential-time² resulting in an overall doubly-exponential time solution. The procedure for ATL^{*}, instead, uses a doubly-exponential bottom-up approach based on the idea of labeling each state of the structure with subformulas true in that state. In our setting we can neither proceed with the membership problem nor use a bottom-up procedure. Indeed, because of ATL^{*}, we need to consider not just the unwinding of the model but the tree execution induced by the player existentially quantified in the formula. Moreover, because of the possible infinite number of configurations induced by the PGS, a bottom-up procedure could never terminate. For this reason, we use a top-down approach that constructs a doubly exponential PD-NTA that simultaneously checks whether a tree is an execution of the structure and a model of the formula. As far as we know, this is the first top-down automata-theoretic approach exploited for ATL^{*}. Some details about this automata construction are reported in the following.

Theorem 4.3.2 *The model-checking problem for ATL^{*} on PGS can be solved in 3EXPTIME.*

¹Automata having as transition relation a positive Boolean combination of states and directions [KVV00].

²Recall that in general the emptiness check for alternating pushdown automata is undecidable [KPV02].

Proof. We give an intuition behind the automata construction by providing some details on how to extend the one introduced in [BEM97] used to solve the model-checking problem for branching-time specifications over pushdown systems. The mentioned approach starts with a tree automaton accepting all tree models of a formula φ , namely the formula automaton \mathcal{A}_φ , over which one can build a PD-NTA $\mathcal{A}_{\mathcal{P},\varphi}$ accepting the unwinding of the pushdown structure \mathcal{P} iff it is contained in the language of the formula automaton \mathcal{A}_φ . To handle ATL^* , one can start with a doubly-exponential parity tree automaton as an adaptation of the one provided in [Sch08]³. Moreover, in order to correctly evaluate the formula over a PGS we need not just to consider the unwinding of the structure but rather the execution trees induced by the formula and precisely from the players existentially quantified in it. This results in selecting at each node subsets of children upon the choices of the players. As the number of these subsets is linear in the number of the decisions of the structure, the overall size of the PD-NTA we construct remains doubly-exponential. Thus, from Theorem 4.3.1 on the previous page we derive a 3EXPTIME procedure. \square

4.3.2 A Lower-bound Complexity

In this section, the model-checking problem for ATL^* over PGSs is shown to be 2EXPSpace-HARD by means of a reduction from the model-checking problem for *quantitative* ATL^* over *one-counter games* [Ves14]. We first give the definition of one-counter games.

Definition 4.3.1 (One-counter Game) *A one-counter game is a tuple $\mathcal{M} = \langle \text{AP}, \text{Ag}, \text{Loc}, (\text{Loc}_a)_{a \in \text{Ag}}, l_0, R, \text{L} \rangle$ where AP is a finite set of atomic propositions, Ag is a finite set of agents, Loc is a finite set of locations, $(\text{Loc}_a)_{a \in \text{Ag}}$ is a partition of Loc , assigning to each agent its own set of locations, $R \subseteq \text{Loc} \times \{-1, 0, 1\} \times \text{Loc}$ is a transition relation, $\text{L} : \text{St} \rightarrow 2^{\text{AP}}$ is a labeling function, and l_0 is an initial location.*

Intuitively, a computation in \mathcal{M} starts from the initial location l_0 with a counter set to the value 0. Then, at each location $l \in \text{Loc}_a$, the owner agent a selects a successor l' and updates the counter according to the transition relation R , *i.e.*, if a selects the triple $(l, \iota, l') \in R$, then the computation moves to the location l' and the current value v of the counter is updated to $v + \iota$. As additional feature, if the counter value is 0, then the transitions providing a decreasing of the counter are disabled. Formally, a computation in \mathcal{M} is an infinite sequence of pairs $\eta = (l_0, v_0) \cdot (l_1, v_1) \dots \in (\text{Loc} \times \mathbb{N})^\omega$ such that, for all $i \in \mathbb{N}$, we have that $(l_i, v_{i+1} - v_i, l_{i+1}) \in R$. For convenience, by $\text{loc}(\eta) \in \text{Loc}^\omega$ and $\text{cnt}(\eta) \in \mathbb{N}^\omega$, we denote the projections of η on locations and counter values, respectively.

One-counter games can be used to interpret *quantitative* ATL^* (QATL^* , for short), which is obtained from ATL^* , by extending the syntax as in the following.

Definition 4.3.2 (Quantitative ATL^*) *Formulas of QATL^* are built inductively from the set of atomic propositions AP and agents Ag , by using the following grammar, where $p \in \text{AP}$, r is a variable ranging on \mathbb{N} , $\bowtie \in \{\leq, <, =, >, \geq, \equiv_k\}$, $c \in \mathbb{N}$, and $\text{A} \subseteq \text{Ag}$:*

³In [Sch08] it is given a single-exponential alternating automaton that can be easily translated into a non-deterministic one with a single exponential-time blowup.

$$\varphi := p \mid r \bowtie c \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathbf{X}\varphi \mid \varphi\mathbf{U}\varphi \mid \langle\langle A \rangle\rangle\varphi.$$

As in ATL^{*}, a QATL^{*} sentence is a Boolean combination of formulas of the form $\langle\langle A \rangle\rangle\psi$. Formulas in QATL^{*} are interpreted over one-counter games and computations. For each one-counter game \mathcal{M} and a computation η , and two natural numbers $c, k \in \mathbb{N}$, we say that $\mathcal{M}, \eta \models r \bowtie c$ if $(\text{cnt}(\eta))_0 \bowtie c$, for all $\bowtie \in \{\leq, <, =, >, \geq, \equiv_k\}$, while for the atomic proposition, Boolean, temporal, and agent quantification cases, the semantics is defined as in ATL^{*}. Also in the case of QATL^{*} sentences, the satisfaction on a given computation η depends only on its first configuration. Then, the relation $\mathcal{M}, (l, i) \models \varphi$ is well-defined and, so, the model-checking problem consists on checking whether $\mathcal{M}, (l_0, 0) \models \varphi$.

In [Ves14], it has been shown that the model-checking problem for QATL^{*} over one-counter games, with counter coded in unary, is 2EXPSpace-Complete and PSPACE-Complete for QATL^{*} formulas of bounded size. We now show that such a problem can be reduced to the model-checking problem of ATL^{*} over PGSs. Indeed, consider the following construction.

Construction 4.3.1 *Let $\mathcal{M} = \langle \text{AP}, \text{Ag}, \text{Loc}, (\text{Loc}_a)_{a \in \text{Ag}}, s_0, R, \mathbf{L} \rangle$ be a one-counter model and φ an QATL^{*} formula with n_1, \dots, n_k being the natural numbers for which there is an occurrence of $r \equiv_{n_i} c_i$ in φ , for some $c_i \in \mathbb{N}$. Then the derived PGS is given by $\mathcal{P}_{\mathcal{M}, \varphi} = \langle \text{AP}', \text{Ag}', \text{Ac}', \text{Loc}', \Gamma, \text{tr}, \text{ap}, s_0 \rangle$ where: $\text{AP}' = \text{AP} \cup \{o, \text{zero}, \sharp, \text{mod}_{n_1}, \dots, \text{mod}_{n_k}, \text{mod}_{n_1}^0, \dots, \text{mod}_{n_1}^{n_1-1}, \dots, \text{mod}_{n_k}^0, \dots, \text{mod}_{n_k}^{n_k-1}\}$, $\text{Ag}' = \text{Ag} \cup \{\text{cnt}\}$, $\text{Ac}' = \text{Loc} \times \{-1, 0, +1\} \cup \{o, \sharp, \text{mod}_{n_1}, \dots, \text{mod}_{n_k}\}$, $\text{Loc}' = \text{Loc} \cup \{l_\sharp^i, l_0^{n_1}, \dots, l_{n_1-1}^{n_1}, \dots, l_0^{n_k}, \dots, l_{n_k-1}^{n_k}\}$, $\Gamma = \{1\}$, ap is such that: $\text{ap}(l, \gamma) = \mathbf{L}(l) \cup \{o\}$, for all $l \in \text{St}$, $\text{ap}(l_\sharp, 1) = \{\sharp\}$ and $\text{ap}(l_\sharp, \perp) = \{\sharp, \text{zero}\}$, and $\text{ap}(l_j^{n_i}, 1) = \{\text{mod}_{n_i}\}$ and $\text{ap}(l_j^{n_i}, \perp) = \{\text{mod}_{n_i}, \text{mod}_{n_i}^j\}$. The transition function is given as follows. For each $(l, \iota, l') \in R$, with $l \in \text{Loc}_a$, we have that: $\text{tr}(l, \gamma, \delta[\text{cnt} \mapsto o][a \mapsto (l', \iota)]) = (l', \varepsilon)$ if $\iota = -1$, $\text{tr}(l, \gamma, \delta[\text{cnt} \mapsto o][a \mapsto (l', \iota)]) = (l', 1)$ if $\iota = 0$, $\text{tr}(l, \gamma, \delta[\text{cnt} \mapsto o][a \mapsto (l', \iota)]) = (l', 1 \cdot 1)$ if $\iota = +1$. $\text{tr}(l, \gamma, \delta[\text{cnt} \mapsto \sharp]) = (l_\sharp, \varepsilon)$, $\text{tr}(l, \gamma, \delta[\text{cnt} \mapsto \text{mod}_{n_i}]) = (l_1^{n_i}, \varepsilon)$, $\text{tr}(l_\sharp, \gamma, \delta) = (l_\sharp, \varepsilon)$, $\text{tr}(l_i^{n_j}, 1, \delta) = (l_{(i+1) \bmod n_j}^{n_j}, \varepsilon)$, $\text{tr}(l_i^{n_j}, \perp, \delta) = (l_i^{n_j}, \varepsilon)$.*

Intuitively, the PGS $\mathcal{P}_{\mathcal{M}, \varphi}$ emulates the counter in \mathcal{M} with a one-symbol stack, in the usual way. When the transition in \mathcal{M} asks for an increasing of the counter, then a the symbol 1 is pushed on the top of the stack. On the other hand, when the transition in \mathcal{M} asks for a decreasing of the counter, then, the symbol in the top of the stack is popped out. Note that the original model never decreases the counter set to the value 0. This means that $\mathcal{P}_{\mathcal{M}, \varphi}$ never performs a pop when the stack is empty. Finally, if the transition in \mathcal{M} does not modify the counter, then $\mathcal{P}_{\mathcal{M}, \varphi}$ does the same. For what the labeling regards, for a location $l \in \text{Loc}$ we extend the labeling given in \mathcal{M} with an auxiliary atomic proposition o , to keep track in the execution that $\mathcal{P}_{\mathcal{M}, \varphi}$ is emulating the one generated in \mathcal{M} . For what the additional locations concerns, we label configurations having l_\sharp as location with the atomic proposition \sharp , to keep track that the execution is performing a check on the size of the stack. Moreover, we label (l, \perp) also with the atomic proposition zero , in order to remember that the stack has been emptied. For what the locations of the form $l_i^{n_j}$ regards, we label all the corresponding configurations with the atomic proposition mod_{n_j} to remember that the

execution is performing a check modulo n_j of the stack, and the one with the symbol \perp also with $\text{mod}_i^{n_j}$ to remember that the stack has been emptied in m steps, with $m \equiv_{n_j} i$. Moreover, in order to emulate the checks of the QATL* formula, we add some auxiliary location. For example, we add a location $l_{\#}$ from which the game loops indefinitely and performs a pop of the stack up to empty it, whatever is the decision made by the agents. Agent *cnt* has the power to enter the location $l_{\#}$, from each original location $l \in \text{Loc}$, at any execution point.




Figure 4.2: The transition to the counting location.

To better understand this point, assume the execution to be in a location l of the PGS, as depicted in Figure 4.2. In the case the agent *cnt* picks the action o , then the execution moves to a location and makes a stack operation according to the choices of the rest of the agents, emulating the execution in \mathcal{M} . On the contrary, in the case the agent *cnt* picks the action $\#$, then the execution moves to the location $l_{\#}$. Observe that, once a transition to $l_{\#}$ is performed, the number of steps (including such transition) in which the stack gets empty is exactly the length of its content.

Thus, for each configuration $s = (l, \alpha)$ the fact that $|\alpha| \leq c$, for a given $c \in \mathbb{N}$, is equivalent to the fact that $\mathcal{P}, s \models \langle\langle \{\text{cnt}\} \rangle\rangle (\text{X}\# \wedge \text{X}^c \text{zero})$ ⁴. Indeed, with the temporal property $\text{X}\#$ we are forcing the agent *cnt* to move on $l_{\#}$. Then, the stack is emptied in exactly $|\alpha|$ steps, so less than c , which implies that the temporal specification $\text{X}^c \text{zero}$ is satisfied. From this, we can derive that $\mathcal{P}, s \models \langle\langle \{\text{cnt}\} \rangle\rangle (\text{X}\# \wedge \text{X}^c \text{zero})$ iff $\mathcal{M}, (l, |\alpha|) \models r \leq c$. By exploiting the equivalences $(r < c) \leftrightarrow (r \leq c - 1)$, $(r > c) \leftrightarrow \neg(r \leq c)$, $(r \geq c) \leftrightarrow \neg(r > c)$, and $(r = c) \leftrightarrow ((r \leq c) \wedge (r \geq c))$, we can check any ordering of the counter by checking the satisfaction of a suitable ATL* formula in $\mathcal{P}_{\mathcal{M}, \varphi}$.

We can address the problem also for the formula $r \equiv_n c$. Indeed, we commit the agent *cnt* to pick a suitable action mod_n and so the PGS to enter a substructure suitably defined for this kind of check. As an example, consider the situation depicted in Figure 4.3 for \equiv_2 . In this case, when the agent *cnt* pick its action mod_2 , then the execution moves to l_2^1 and, from that point onward, switches from l_2^1 to l_2^0 and viceversa, popping the top of the stack at each step. As soon as the stack is emptied, the execution loops in the location reached when the last popping has been performed. Also in this case, it is not hard to see that the execution starts looping on the location l_2^i such that $i = |\alpha| \text{mod} 2$. This implies that, for each configuration $s = (l, \alpha)$, the fact that $|\alpha| \equiv_n c$, for some $n, c \in \mathbb{N}$, is equivalent to check whether $\mathcal{P}_{\mathcal{M}, \varphi}, (l, \alpha) \models \langle\langle \{\text{cnt}\} \rangle\rangle (\text{Xmod}_n \wedge \text{Fmod}_n^c)$. Indeed, the temporal specification Xmod_n is used to force the agent *cnt* to pick the action mod_n , while the formula




Figure 4.3: The transition to the counting modulo 2 location.

⁴By X^c we are denoting c occurrences of the operator X .

Fmod_n^c is satisfied iff the stack is emptied in a number m of steps such that $m \equiv_n c$.

Finally, consider a QATL^{*} formula of the form $\varphi = \langle\langle A \rangle\rangle\psi$. Interpreting φ over a one-counter game \mathcal{M} is equivalent to interpret the formula $\langle\langle A \cup \{\text{cnt}\} \rangle\rangle \text{XGo} \wedge \psi$ on the derived PGS \mathcal{P} . Indeed, with the temporal specification XGo we are essentially forcing the agent cnt to always pick its action o . This implies that the other agents in A can apply the same strategies they use in \mathcal{M} in order to satisfy the formula ψ also in \mathcal{P} .

By applying the above reasoning, for a QATL^{*} formula φ we define recursively the following transformation $\tau(\varphi)$: $\tau(\mathbf{p}) = \mathbf{p}$; $\tau(\mathbf{r} \leq \mathbf{c}) = \langle\langle \{\text{cnt}\} \rangle\rangle (\text{X}\# \wedge \text{X}^c \text{zero})$; $\tau(\mathbf{r} \equiv_n \mathbf{c}) = \langle\langle \{\text{cnt}\} \rangle\rangle (\text{Xmod}_n \wedge \text{Fmod}_n^c)$; $\tau(\neg\psi) = \neg\tau(\psi)$; $\tau(\psi_1 \wedge \psi_2) = \tau(\psi_1) \wedge \tau(\psi_2)$; $\tau(\text{X}\psi) = \text{X}\tau(\psi)$; $\tau(\psi_1 \text{U}\psi_2) = \tau(\psi_1) \text{U}\tau(\psi_2)$; $\tau(\langle\langle A \rangle\rangle\psi) = \langle\langle A \cup \{\text{cnt}\} \rangle\rangle (\text{XGo} \wedge \tau(\psi))$.

It is not hard to prove the following theorem.

Theorem 4.3.3 (Model-checking Reduction) *For a given one-counter game \mathcal{M} and a QATL^{*} formula φ , it holds that $\mathcal{M} \models \varphi$ iff $\mathcal{P}_{\mathcal{M},\varphi} \models \tau(\varphi)$*

Note that, for counters coded in unary, we have that the construction of the PGS reported above is of size polynomial in the size of the one-counter game. From this and from Theorem 4.3.3 we derive the following.

Corollary 4.3.1 (Model-checking Lower-bound) *The model-checking problem for ATL^{*} over PGS is in 2EXPSpace-HARD.*

Bibliography

- [ABE⁺05] R. Alur, M. Benedikt, K. Etessami, P. Godefroid, T. W. Reps, and M. Yannakakis. Analysis of Recursive State Machines. *ACM Trans. Program. Lang. Syst.*, 27(4):786–818, 2005.
- [ABK13] S. Almagor, U. Boker, and O. Kupferman. Formalizing and Reasoning about Quality. In *ICALP'13*, volume 7966 of *LNCS*, pages 15–27, 2013.
- [ABK14] S. Almagor, U. Boker, and O. Kupferman. Discounting in LTL. In *TACAS'14*, volume 8413 of *LNCS*, pages 424–439. Springer, 2014.
- [AF84] M. Aigner and M. Fromme. A Game of Cops and Robbers. *Discrete Applied Mathematics*, 8(1):1 – 12, 1984.
- [ÅGJ07] T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-Time Temporal Logics with Irrevocable Strategies. In *Theoretical Aspects of Rationality and Knowledge'07*, pages 15–24, 2007.
- [AHK02] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *Journal of the ACM*, 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.*, 223:1–17, 2013.
- [BCHJ09] R. Bloem, K. Chatterjee, T. Henzinger, and B. Jobstmann. Better Quality in Synthesis through Quantitative Objectives. In *CAV'09*, volume 5643 of *LNCS*, pages 140–156. Springer, 2009.
- [Bel14] Francesco Belardinelli. Reasoning about Knowledge and Strategies: Epistemic Strategy Logic. In *SR*, pages 27–33, 2014.
- [BEM97] A. Bouajjani, J. Esparza, and O. Maler. Reachability Analysis of Pushdown Automata: Application to Model-Checking. In *CONCUR'97*, pages 135–150, 1997.
- [BJ14] N. Bulling and W. Jamroga. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.
- [BJD08] Nils Bulling, Wojciech Jamroga, and Jürgen Dix. Reasoning about temporal properties of rational play. *Ann. Math. Artif. Intell.*, 53(1-4):51–114, 2008.
- [BLLM09] T. Brihaye, A. Da Costa Lopes, F. Laroussinie, and N. Markey. ATL with Strategy Contexts and Bounded Memory. In *Logical Foundations of Computer Science'09*, LNCS 5407, pages 92–106. Springer, 2009.

- [BMP10] L. Bozzelli, A. Murano, and A. Peron. Pushdown Module Checking. *Formal Methods in System Design*, 36(1):65–95, 2010.
- [Boz06] L. Bozzelli. Complexity Results on Branching-Time Pushdown Model Checking. In *VMCAI'06*, pages 65–79, 2006.
- [Büc62] J.R. Büchi. On a Decision Method in Restricted Second-Order Arithmetic. In *International Congress on Logic, Methodology, and Philosophy of Science'62*, pages 1–11. Stanford University Press, 1962.
- [Bul14] N. Bulling. A Survey of Multi-Agent Decision Making. *KI*, 28(3):147–158, 2014.
- [CBC13] C. Charetton, J. Brunel, and D. Chemouil. Towards an Updatable Strategy Logic. In *Proceedings 1st International Workshop on Strategic Reasoning, SR 2013, Rome, Italy, March 16-17, 2013.*, pages 91–98, 2013.
- [CE81] E.M. Clarke and E.A. Emerson. Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic. In *Logic of Programs'81*, LNCS 131, pages 52–71. Springer, 1981.
- [CGP02] E.M. Clarke, O. Grumberg, and D.A. Peled. *Model Checking*. MIT Press, 2002.
- [CHJ08] K. Chatterjee, T. Henzinger, and B. Jobstmann. Environment Assumptions for Synthesis. In *CONCUR'08*, volume 5201 of *LNCS*, pages 147–161. Springer, 2008.
- [CHP07] K. Chatterjee, T.A. Henzinger, and N. Piterman. Strategy Logic. In *Concurrency Theory'07*, LNCS 4703, pages 59–73. Springer, 2007.
- [Chu63] A. Church. Logic, Arithmetics, and Automata. In *International Congress of Mathematicians'62*, pages 23–35, 1963.
- [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.
- [CMJ04] K. Chatterjee, R. Majumdar, and M. Jurdzinski. On Nash Equilibria in Stochastic Games. In *CSL'04*, volume 3210 of *LNCS*, pages 26–40. Springer, 2004.
- [CY98] C. Courcoubetis and M. Yannakakis. Markov Decision Processes and Regular Events. *IEEE Transaction on automatic control*, 43(10):1399 – 1418, 1998.
- [DLM10] A. Da Costa, F. Laroussinie, and N. Markey. ATL with Strategy Contexts: Expressiveness and Model Checking. In *Foundations of Software Technology and Theoretical Computer Science'10*, LIPIcs 8, pages 120–132, 2010.

- [DLM12] A. Da Costa, F. Laroussinie, and N. Markey. Quantified CTL: Expressiveness and model checking - (extended abstract). In *CONCUR*, volume 7454 of *Lecture Notes in Computer Science*, pages 177–192. Springer, 2012.
- [EF95] H.D. Ebbinghaus and J. Flum. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer, 1995.
- [EH86] E.A. Emerson and J.Y. Halpern. “Sometimes” and “Not Never” Revisited: On Branching Versus Linear Time. *Journal of the ACM*, 33(1):151–178, 1986.
- [FKL10] D. Fisman, O. Kupferman, and Y. Lustig. Rational Synthesis. In *TACAS*, pages 190–204, 2010.
- [FMP08] A. Ferrante, A. Murano, and M. Parente. Enriched μ -Calculi Module Checking. *Logical Methods in Computer Science*, 4(3), 2008.
- [FS10] B. Finkbeiner and S. Schewe. Coordination Logic. In *Computer Science Logic’10*, LNCS 6247, pages 305–319. Springer, 2010.
- [GTW02] E. Grädel, W. Thomas, and T. Wilke. *Automata, Logics, and Infinite Games: A Guide to Current Research*. LNCS 2500. Springer, 2002.
- [Har84] D. Harel. A Simple Highly Undecidable Domino Problem. In *Logic and Computation Conference’84*. North-Holland, 1984.
- [Hen10] T.A. Henzinger. From Boolean to Quantitative Notions of Correctness. In *POPL’10*, pages 157–158. ACM, 2010.
- [HLW13] A. Herzig, E. Lorini, and D. Walther. Reasoning about Actions Meets Strategic Logics. In *Logic, Rationality, and Interaction*, pages 162–175. Springer, 2013.
- [Hod93] W. Hodges. *Model Theory*. Encyclopedia of Mathematics and its Applications. Cambridge, 1993.
- [Hod01] W. Hodges. *Elementary Predicate Logic.*, volume 1 of *Handbook of philosophical logic*. Springer, 2001.
- [HSW13] C. Huang, S. Schewe, and F. Wang. Model-Checking Iterated Games. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 154–168. Springer, 2013.
- [JM14] W. Jamroga and A. Murano. On module checking and strategies. In *AAMAS’14*, pages 701–708, 2014.
- [JvdH04] W. Jamroga and W. van der Hoek. Agents that Know How to Play. *Fundamenta Informaticae*, 63(2-3):185–219, 2004.
- [Kel76] R.M. Keller. Formal Verification of Parallel Programs. *Communication of the ACM*, 19(7):371–384, 1976.

- [Kol85] F. Kolaitis. Game Quantification. In J Barwise and S. Feferman, editors, *Handbook of Model-Theoretic Logics*, pages 365–421. SpringerVerlag, 1985.
- [Koz83] D. Kozen. Results on the Propositional μ Calculus. *Theoretical Computer Science*, 27(3):333–354, 1983.
- [KPV02] O. Kupferman, N. Piterman, and M. Y. Vardi. Pushdown Specifications. In *LPAR'02*, pages 262–277, 2002.
- [KPV14] O. Kupferman, G. Perelli, and M. Y. Vardi. Synthesis with Rational Environments. In *EUMAS'14*, 2014. To appear.
- [Kri63] S.A. Kripke. Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, 16:83–94, 1963.
- [KV98] O. Kupferman and M.Y. Vardi. Weak Alternating Automata and Tree Automata Emptiness. In *Symposium on Theory of Computing'98*, pages 224–233, 1998.
- [KV99] O. Kupferman and M. Y. Vardi. Church's problem revisited. *Bulletin of Symbolic Logic*, 5(2):245–263, 1999.
- [KVV00] O. Kupferman, M.Y. Vardi, and P. Wolper. An Automata Theoretic Approach to Branching-Time Model Checking. *Journal of the ACM*, 47(2):312–360, 2000.
- [KVV01] O. Kupferman, M.Y. Vardi, and P. Wolper. Module Checking. *Information and Computation*, 164(2):322–344, 2001.
- [LM13] F. Laroussinie and N. Markey. Satisfiability of ATL with Strategy Contexts. In *GandALF*, volume 119 of *EPTCS*, pages 208–223, 2013.
- [LR06] A. Lomuscio and F. Raimondi. MCMAS: A Model Checker for Multi-agent Systems. In *Tools and Algorithms for the Construction and Analysis of Systems'06*, LNCS 3920, pages 450–454. Springer, 2006.
- [Mar75] A.D. Martin. Borel Determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
- [MMPV11] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. Reasoning About Strategies: On the Model-Checking Problem. Technical report, arXiv, 2011.
- [MMPV12] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. What Makes ATL* Decidable? A Decidable Fragment of Strategy Logic. In *Concurrency Theory'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. *ACM Trans. Comput. Log.*, 15(4):34, 2014.

- [MMS13] F. Mogavero, A. Murano, and L. Sauro. On the Boundary of Behavioral Strategies. In *Logic in Computer Science'13*, pages 263–272. IEEE Computer Society, 2013.
- [MMS14] F. Mogavero, A. Murano, and L. Sauro. A Behavioral Hierarchy of Strategy Logic. In *CLIMA XV*, LNCS 8624, pages 869–876. Springer, 2014.
- [MMV10a] F. Mogavero, A. Murano, and M.Y. Vardi. Reasoning About Strategies. In *Foundations of Software Technology and Theoretical Computer Science'10*, LIPIcs 8, pages 133–144. Leibniz-Zentrum fuer Informatik, 2010.
- [MMV10b] F. Mogavero, A. Murano, and M.Y. Vardi. Relentful Strategic Reasoning in Alternating-Time Temporal Logic. In *Logic for Programming Artificial Intelligence and Reasoning'10*, LNAI 6355, pages 371–387. Springer, 2010.
- [MP14] F. Mogavero and G. Perelli. On the Remarkable Features of Binding Forms. *CoRR*, abs/1404.1531, 2014.
- [MS85] D. E. Muller and P. E. Schupp. The Theory of Ends, Pushdown Automata, and Second-Order Logic. *Theor. Comput. Sci.*, 37:51–75, 1985.
- [MS87] D.E. Muller and P.E. Schupp. Alternating Automata on Infinite Trees. *Theoretical Computer Science*, 54(2-3):267–276, 1987.
- [MS95] D.E. Muller and P.E. Schupp. Simulating Alternating Tree Automata by Nondeterministic Automata: New Results and New Proofs of Theorems of Rabin, McNaughton, and Safra. *Theoretical Computer Science*, 141(1-2):69–107, 1995.
- [MSS88] D.E. Muller, A. Saoudi, and P.E. Schupp. Weak Alternating Automata Give a Simple Explanation of Why Most Temporal and Dynamic Logics are Decidable in Exponential Time. In *Logic in Computer Science'88*, pages 422–427. IEEE Computer Society, 1988.
- [Mye97] R.B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.
- [NR01] N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior*, 35(1-2):166–196, 2001.
- [NRTV07] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [OR94] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [Pau02] M. Pauly. A Modal Logic for Coalitional Power in Games. *Journal of Logic and Computation*, 12(1):149–166, 2002.

- [Pin07] S. Pinchinat. A Generic Constructive Solution for Concurrent Games with Expressive Constraints on Strategies. In *Automated Technology for Verification and Analysis'07*, LNCS 4762, pages 253–267. Springer, 2007.
- [Pnu77] A. Pnueli. The Temporal Logic of Programs. In *Foundation of Computer Science'77*, pages 46–57. IEEE Computer Society, 1977.
- [PP04] D. Perrin and J. Pin. *Infinite Words*. Pure and Applied Mathematics. Elsevier, 2004.
- [PR89] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190, 1989.
- [QS81] J.P. Queille and J. Sifakis. Specification and Verification of Concurrent Programs in Cesar. In *Symposium on Programming'81*, LNCS 137, pages 337–351. Springer, 1981.
- [Rab69] M.O. Rabin. Decidability of Second-Order Theories and Automata on Infinite Trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.
- [Sch08] S. Schewe. ATL* Satisfiability is 2ExpTime-Complete. In *International Colloquium on Automata, Languages, and Programming'08*, LNCS 5126, pages 373–385. Springer, 2008.
- [Sel75] R. Selten. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory*, 4(1):25–55, 1975.
- [Sis83] A.P. Sistla. *Theoretical Issues in the Design and Certification of Distributed Systems*. PhD thesis, Harvard University, Cambridge, MA, USA, 1983.
- [SVW87] A.P. Sistla, M.Y. Vardi, and P. Wolper. The Complementation Problem for Büchi Automata with Applications to Temporal Logic. *Theoretical Computer Science*, 49:217–237, 1987.
- [Tho90] W. Thomas. Automata on Infinite Objects. In *Handbook of Theoretical Computer Science (vol. B)*, pages 133–191. MIT Press, 1990.
- [Var88] M.Y. Vardi. A Temporal Fixpoint Calculus. In *Principles of Programming Languages'88*, pages 250–259. Association for Computing Machinery, 1988.
- [Var96] M.Y. Vardi. Why is Modal Logic So Robustly Decidable? In *Descriptive Complexity and Finite Models'96*, pages 149–184. American Mathematical Society, 1996.
- [Ves14] Steen Vester. Model-checking Quantitative Alternating-time Temporal Logic on One-counter Game Models. Technical report, ArXiv, 2014.
- [VS85] M.Y. Vardi and L.J. Stockmeyer. Improved Upper and Lower Bounds for Modal Logics of Programs: Preliminary Report. In *Symposium on Theory of Computing'85*, pages 240–251, 1985.

- [VW86] M.Y. Vardi and P. Wolper. An Automata-Theoretic Approach to Automatic Program Verification. In *Logic in Computer Science '86*, pages 332–344. IEEE Computer Society, 1986.
- [Wal96] I. Walukiewicz. Pushdown Processes: Games and Model Checking. In *CAV'96*, pages 62–74, 1996.
- [Wal00] I. Walukiewicz. Model Checking CTL Properties of Pushdown Systems. In *FSTTCS'00*, pages 127–138, 2000.
- [Wan61] H. Wang. Proving Theorems by Pattern Recognition II. *Bell System Technical Journal*, 40:1–41, 1961.
- [WHY11] F. Wang, C. Huang, and F. Yu. A Temporal Logic for the Interaction of Strategies. In *CONCUR 2011*, pages 466–481. Springer, 2011.
- [WvdHW07] Dirk Walther, Wiebe van der Hoek, and Michael Wooldridge. Alternating-Time Temporal Logic with Explicit Strategies. In *TARK*, pages 269–278, 2007.