

DEPARTMENT OF COMPUTING
IMPERIAL COLLEGE LONDON

MENG INDIVIDUAL PROJECT

A Model Checker for Strategy Logic

Author:
Petr ČERMÁK

Supervisor:
Prof. Alessio R. LOMUSCIO

June 2014

Abstract

There is a gap between game theory and practical model checking. Although logics for expressing game-theoretic concepts like Nash equilibria, notably Strategy Logic (SL), have been put forward recently, there are currently no tools supporting them. This report documents the design of the first practical model checking algorithms for two fragments of SL and their symbolic implementation using binary decision diagrams. We explain how we incorporated both algorithms in the existing model checker MCMAS and added support for strategy synthesis. We evaluate the performance of the algorithms on several scalable scenarios and compare it with the original tool. Our results demonstrate that SL model checking is feasible in practice despite its high computational complexity.

Acknowledgements

I would like to thank:

- Professor Alessio Lomuscio for his invaluable guidance throughout the whole project,
- Professor Aniello Murano and Dr. Fabio Mogavero for their feedback on the algorithms and their help with finding scalable scenarios for evaluation,
- the Department of Computing for funding my trips to the *2nd International Workshop on Strategic Reasoning* and the *26th International Conference on Computer Aided Verification*,
- my family and friends, who have always been there for me, and
- my girlfriend Klára for her love and support.

Contents

1	Introduction	11
1.1	Objectives	12
1.2	Challenges	13
1.3	Contributions	13
1.4	Published Work	14
2	Background	15
2.1	Frameworks for Modelling Systems	15
2.1.1	Kripke Models	15
2.1.2	Concurrent Game Structures	16
2.1.3	Interpreted Systems	17
2.2	Specification Languages	18
2.2.1	Linear Temporal Logic	18
2.2.2	Computation Tree Logic	20
2.2.3	Full Branching Time Logic	22
2.2.4	Alternating-Time Temporal Logic	24
2.2.5	Strategy Logic	26
2.2.6	Epistemic modalities	29
2.2.7	Model Checking Complexity	31
2.2.8	Summary	31
2.3	Verification Methods	32
2.3.1	Binary Decision Diagrams	32
2.3.2	Symbolic Model Checking	34
2.3.3	Summary	36
2.4	Automata and Games	36
2.4.1	ω -Automata	36
2.4.2	Symbolically Represented ω -Automata	38
2.4.3	Translating LTL Formulas to ω -automata	39
2.4.4	Determinisation	42
2.4.5	Games	49
2.4.6	Summary	56
2.5	Existing Tools	56
2.5.1	MCK	56
2.5.2	MCMAS	56
2.5.3	Mocha	57
2.5.4	NuSMV	57
2.5.5	PRISM	57
2.5.6	VerICS	58
2.6	Summary	59

3	Fragment Selection	61
3.1	Imperfect Recall	62
3.2	Complete Information	62
3.3	Toy Model	63
3.3.1	Formal Definition	64
3.3.2	Symbolic Implementation	64
3.4	Summary	66
4	Epistemic Strategy Logic	67
4.1	Logic	67
4.1.1	Syntax	67
4.1.2	Basic Concepts	68
4.1.3	Semantics	68
4.1.4	Comparison with Strategy Logic	70
4.1.5	Limitations	71
4.2	Model Checking	72
4.2.1	Complexity	72
4.2.2	Algorithm	73
4.2.3	Strategy Synthesis	76
4.2.4	Symbolic Implementation	79
4.3	Summary	83
5	One-Goal Strategy Logic	85
5.1	Logic	85
5.2	Model Checking	86
5.2.1	Algorithm	87
5.2.2	Complexity	97
5.2.3	Strategy Synthesis	98
5.2.4	Symbolic Implementation	102
5.2.5	Separate Determinisation	107
5.3	Summary	109
6	Implementation	111
6.1	Existing Tool	111
6.1.1	Functionality	111
6.1.2	Usage	112
6.1.3	Architecture	115
6.2	Epistemic Strategy Logic Extension	117
6.2.1	Functionality	118
6.2.2	Usage	118
6.2.3	Architecture	120
6.3	One-Goal Strategy Logic Extension	123
6.3.1	Functionality	123
6.3.2	Usage	124
6.3.3	Architecture	130
6.4	Experimental Results	133
6.4.1	Dining Cryptographers	133
6.4.2	Cake Cutting	134
6.4.3	Scheduler	140
6.4.4	Nim	143
6.4.5	Analysis	145
6.5	Summary	147

7	Evaluation	149
7.1	Theory	149
7.1.1	Strengths	150
7.1.2	Weaknesses	150
7.2	Implementation	151
7.2.1	Strengths	151
7.2.2	Weaknesses	152
7.3	Fragment Comparison	153
7.4	Summary	155
8	Conclusions	157
8.1	Summary of Work	157
8.2	Future Work	158
	Bibliography	161
A	Toy Model ISPL Code	167

Chapter 1

Introduction

Computer systems have an ever-increasing presence in our everyday lives. We rely on them in almost all activities both at work and during our free time. It will probably not be long until we meet driverless cars on the road and receive deliveries from unmanned quadcopters [19,24]. It seems that more and more decisions in the future will be made by computers without any human intervention. Yet computers are infamous for their frequent faults. Each of us has witnessed freezing applications countless times. While the usual type of damage caused by a computer bug today is the loss of several hours of unsaved work, autonomous system failures in the future might result in thousands of deaths in the blink of an eye. The six accidents caused by massive radiation overdoses by the Therac-25 medical electron accelerator between 1985 and 1987 remind us that such scenarios are not science fiction [59]. It should be clear that safety-critical systems need to be verified before being deployed in order to avoid loss of life and money. Fortunately, formal verification is now finally becoming the norm in many areas including processor design [53].

Model-checking is a popular approach to systems verification [62]. It entails translating the system and the property to be verified to a model \mathcal{M} and a formula φ respectively and then checking that the model satisfies the formula, i.e. that $\mathcal{M} \models \varphi$ holds. This approach can be applied in many areas including reactive systems, which never terminate, by modelling them as multi-agent systems, in which agents with local states interact with each other via synchronous actions. Various frameworks for modelling such systems are discussed in Section 2.1.

Several temporal logics for specifying system properties have been proposed, including Linear Temporal Logic (LTL) [80], Computation Tree Logic (CTL) [33], and Full Branching Time Logic (CTL*) [37]. More recently, Alternating-Time Logic (ATL) and Full Alternating-Time Logic (ATL*) [18] were introduced. ATL and ATL* are often more suitable for multi-agent systems since their operators refer to properties achievable by given groups of agents. Intuitively, an ATL formula $\langle\langle\{\alpha, \beta\}\rangle\rangle\varphi$ expresses that “agents α and β can together enforce φ ”. A detailed discussion of these logics can be found in Section 2.2. All of them except ATL*, which has a high model checking complexity, are now supported by multiple toolkits (see Section 2.5).

However, neither ATL nor ATL* syntax refers to strategies directly. Hence, the logics cannot express properties such as “agents α and β can enforce φ while sharing the same strategy”. A new formalism, *Strategy Logic* (SL), which treats strategies explicitly via first-order quantifiers $\langle\langle x \rangle\rangle$, $\llbracket x \rrbracket$ and agent binding (a, x) , has thus been introduced in [76]. SL strictly subsumes all logics in the ATL* hierarchy and can express game-theoretic concepts like Nash Equilibria. The aforementioned property, inexpressible in ATL*, can be written in SL as $\langle\langle x \rangle\rangle(a, x)(\beta, x)\varphi$. However, the price for the large expressiveness of SL includes NONELEMENTARYSPACE-HARD model checking complexity [72], undecidability of satisfiability [77], and non-behavioural strategies [75]. We are not aware of any existing model checker that would support SL.

In order to avoid the problems associated with full SL, several syntactic fragments have been proposed [72]. The least expressive subset, *One-Goal Strategy Logic* (SL[1G]), has the same model checking complexity as ATL*, namely 2EXPTIME-COMPLETE with respect to the size of the formula, while still strictly subsuming it. Moreover, SL[1G] strategies are behavioural and can therefore be synthesised [75]. We have designed a novel model checking algorithm for SL[1G] by reducing the problem to solving a two-player parity game. Our algorithm, presented in Chapter 5, admits an efficient symbolic implemen-

tation and supports general strategy synthesis. Moreover, we prove that it is correct and has optimal worst-case time complexity.

We introduce and investigate another variant of SL, *Epistemic Strategy Logic*, or Strategy Logic with Knowledge (SLK). Unlike SL and SL[1G], we define SLK on strategies with imperfect recall (i.e. memoryless strategies where an agent’s action is determined only by its current local state) and supports epistemic modalities expressing agents’ knowledge. We show that the theoretical complexity of SLK model checking is PSPACE with respect to both the size of the model and the formula. We then provide an efficient model checking algorithm for SLK which can be implemented symbolically and supports witness/counterexample strategy synthesis. We also prove that it is correct. Both the logic and the algorithm are discussed in Chapter 4.

MCMAS is a BDD-based model checker for the verification of multi-agent systems developed at Imperial College London [63]. The specification languages it currently supports include CTL and ATL with epistemic and deontic modalities. MCMAS and other existing model checking tools are discussed in Section 2.5. Various verification methods including BDDs are explained in Section 2.3.

We implemented the novel model checking algorithms for both SLK and SL[1G] as extensions for MCMAS and thereby extended the set of logics it supports. Both extensions support witness/counterexample strategy synthesis and the SL[1G] extension also supports general strategy synthesis. An overview of the existing MCMAS architecture as well as the usage and implementation of our algorithms is provided in Chapter 6. The experimental results we obtained using our extensions are presented in Section 6.4.

1.1 Objectives

Our goal was to *implement the first model checker for SL*. We believed that it would be a very important milestone for the development of SL. It would demonstrate the potential of the logic, provide a solid baseline for future research, and promote the synergy of systems verification and game theory. In order to achieve this, our objectives were to:

1. **Research SL and its fragments.** SL is a relatively new formalism for specifying properties in systems verification introduced in 2010 [76]. The first aim of this project was therefore to research SL as well as its syntactic fragments in order to have a good understanding of its expressiveness, complexity, decidability, and other properties. In other words, it was necessary to obtain basic knowledge of the research area to be able to set challenging yet achievable objectives for the project.
2. **Design an efficient model checking algorithm.** The main aim of this project was to come up with a model checking algorithm for SL or some of its variants. While theoretical automata-based model checking procedures have been proposed [72], it is important to design novel algorithms which use efficient data structures (e.g. binary decision diagrams discussed in Subsection 2.3.1) in order to reduce the effects of the high computational complexity of the problem.
3. **Implement the algorithm as a tool.** The next objective was to create a tool implementing the model checking algorithm we designed. The implementation of the algorithm had to be robust, efficient, and user-friendly. Moreover, the source code had to have a good design and appropriate documentation as we intended to release it as open source.

A standalone tool would have to support the whole verification framework including a modelling language parser, state space encoding, reachability calculation, etc. In order to avoid “reinventing the wheel”, we could extend an existing tool instead (several existing tools are presented in Section 2.5). This would give us more time to add more functionality and/or optimise the performance of the algorithm.

4. **Support strategy synthesis.** In addition to model checking, i.e. determining whether a given specification is true or false, we also wanted to synthesise the corresponding strategies for the agents. These strategies would provide us with a deeper insight into why a particular formula holds or does not hold. More importantly, strategy synthesis would make our tool much more useful as it would allow users to underspecify the agents’ behaviour and have it automatically generated.

1.2 Challenges

Overall, the project was very challenging due to its novelty. The work entailed transforming high-level theoretical concepts into low-level programming code with almost no prior experience in the field. The biggest challenges encountered throughout the project were:

1. **No existing algorithms.** While there has been a lot of focus on the theoretical aspects of SL in the past few years [72, 73, 75, 76], no practical model checking algorithm has been proposed or implemented for either SL or any of its fragments so far. In fact, we are not aware of any existing tool that would support ATL^* , which was introduced in 2002 [18] and is strictly subsumed even by the least expressive fragment of SL. Therefore, it was very difficult to find any starting points for our research into model checking algorithms.
2. **Undecidability.** SL is a very powerful formalism for expressing complex properties of multi-agent systems. While this is desirable from the user’s perspective, it makes designing and implementing efficient algorithms for handling the logic very difficult. Moreover, certain problems are impossible to solve due to undecidability. Unfortunately, this was the case when we considered adding incomplete information to the original SL, as described in Chapter 3. In order to achieve decidability, we had to consider less expressive variants of SL instead.
3. **Complexity.** The high model checking complexity of SL has a severe impact on performance. Firstly, the number of temporal operators within a formula we can handle is limited since both SLK and SL[1G] have high complexity with respect to the size of the formula, namely PSPACE and 2EXPTIME-COMPLETE. However, it turned out that that the PSPACE complexity of SLK with respect to the size of the model is an even bigger problem, rendering SLK model checking feasible only for very small state spaces in certain scenarios (see Section 6.4).
4. **Existing codebase.** MCMAS is one of the leading model checkers (see Subsection 2.5.2). However, despite being developed as open source, it is very badly designed from a software engineering point of view. The problems include a complete lack of testing, violation of OOP principles, frequent usage of global variables, and inconsistent code style. The tightly coupled code design made it very difficult for us to test our new functionality. Unfortunately, the lack of testing makes refactoring existing code very frustrating. In our opinion, the tool has outlived its original purpose and should be rewritten from scratch.

1.3 Contributions

We have designed and implemented model checking algorithms for two fragments of SL. We believe this is an important contribution to the research area, as our tool is the first model checker to support SL or any of its subsets. Throughout this project, we have developed:

1. **New SL fragment.** We introduced a new fragment of SL called *Epistemic Strategy Logic* or Strategy Logic with Knowledge (SLK). Unlike SL, SLK is defined with respect to imperfect recall semantics (i.e. agents have no memory of the past) and supports epistemic modalities expressing agents’ knowledge (see Subsection 2.2.6). Hence, SLK specifications can represent complex game-theoretic concepts like Nash equilibria under incomplete information.

We give the syntax and semantics of SLK. We also discuss its limitations and prove that the SLK model checking problem belongs to the PSPACE complexity class with respect to both the size of the model and the formula (see Subsection 4.2.1).

Section 4.1 provides a detailed description of SLK.

2. **Model checking algorithm for SLK.** We have designed a novel model checking algorithm for SLK which has worst-case exponential time complexity both in the size of the model and the formula. We prove its correctness and provide an efficient symbolic encoding of it using binary decision diagrams (see Subsections 2.3.1 and 2.3.2). In addition, we explain how the algorithm (and its symbolic encoding) can be used for the purposes of witness/counterexample strategy synthesis.

Section 4.2 provides a detailed description of the algorithm and its symbolic representation.

3. **Model checking algorithm for SL[1G].** We describe a novel algorithm which reduces the model checking problem for SL[1G] to the problem of solving two-player parity games. Moreover, we prove its correctness, show that it is optimal, and explain how it can be used to synthesise strategies for arbitrary SL[1G] formulas. Again, we also provide an efficient symbolic encoding of the algorithm.

We believe that this is the *biggest achievement of the project* as SL[1G] subsumes ATL*, which (as already mentioned in Section 1.2) was introduced in 2002 [18] but, as far as we know, has no practical implementation yet. Since ATL* has the same model checking complexity as SL[1G], our algorithm is also optimal for it.

Chapter 5 provides a detailed description of SL[1G], the algorithm, and its symbolic representation.

4. **Implementation of both algorithms.** We have extended the functionality of the MCMAS model checker (discussed in Subsection 2.5.2) with support for both SLK and SL[1G]. We also implemented strategy synthesis for both fragments. The tool will be released as open source¹ so that it could be further developed in the future.

Chapter 6 provides a detailed description of the existing MCMAS architecture and our extensions.

5. **Experimental evaluation.** We evaluated the performance of our SLK and SL[1G] MCMAS extensions on several scalable scenarios. Moreover, we demonstrated the expressiveness of the fragments as well as the newly implemented functionality by automatically synthesising strategies that enforce properties which were not supported by the original tool (e.g. we generate a strategy for a scheduler which ensures the existence of Nash equilibria, see Subsection 6.4.2).

The scenarios and the experimental results of our implementation are discussed in Section 6.4.

A detailed evaluation of the project can be found in Chapter 7.

1.4 Published Work

While working on the project, we published a tool paper on SLK implementation in MCMAS [27]. The paper was a joint work with my supervisor, Professor Alessio Lomuscio, and Dr. Fabio Mogavero and Professor Aniello Murano from the Computer Science Division (Sezione di Informatica) of the Department of Physical Science (Dipartimento di Scienze Fisiche) of the University of Naples “Federico II” (Università degli Studi di Napoli “Federico II”), Italy. The paper was accepted at the *26th International Conference on Computer Aided Verification 2014*. It summarises the procedure described in Section 4.2, its implementation provided in Section 6.2, and the experimental results presented in Section 6.4.

Given its novelty, we intend to publish a paper on our model checking algorithm for SL[1G] in the summer (again as a joint work).

¹The SLK extension, called MCMAS-SLK, has already been released under GNU General Public License (GPL) [5] and is available at <http://vas.doc.ic.ac.uk/software/tools/>.

Chapter 2

Background

In this chapter we describe concepts and methods from the area of software verification which the rest of this report is based on. We also present a list of existing tools related to our work. This chapter can be skipped at first reading and used as a reference in the subsequent chapters.

The ultimate goal of our project is to develop a tool that can verify SL specifications on multi-agent systems. Formal verification techniques typically comprise three components [49]:

1. a *framework for modelling systems*, typically a description language of some sort;
2. a *specification language* for describing the properties to be verified;
3. a *verification method* to establish whether the description of a system satisfies the specification.

Let us consider a system S and a property P to be verified (on the system). There are two main approaches to verification [49]:

- In *proof-based verification*, the system description is a set of formulas Γ_S and the property to be verified is another formula φ_P . The verification method consists of trying to prove $\Gamma_S \vdash \varphi_P$. This usually requires guidance and expertise from the user.
- In *model-based verification*, the system description is a model \mathcal{M}_S and the property to be verified is a formula ϕ_P . The verification method consists of computing whether \mathcal{M}_S satisfies ϕ_P ($\mathcal{M}_S \models \phi_P$). This can usually be done automatically for finite models.

In this project we focus purely on a model-based verification method called *model checking*, which refers to the process of determining whether a formula ϕ_P holds in a concrete model \mathcal{M}_S [49]:

$$\mathcal{M}_S \stackrel{?}{\models} \phi_P$$

A *model checker* is a tool (i.e. a computer program) which performs this calculation automatically.

We will now describe various approaches to each of the three components of formal verification using model checking: how the system S is represented as a model \mathcal{M}_S , how the property P is expressed as a formula ϕ_P , and how the answer ($\mathcal{M}_S \models \phi_P$, or $\mathcal{M}_S \not\models \phi_P$) is determined.

2.1 Frameworks for Modelling Systems

Given a system S , we aim to represent it using a model \mathcal{M}_S . We discuss here three formalisms of increasing complexity: Kripke models, concurrent game structures, and interpreted systems.

2.1.1 Kripke Models

Kripke semantics are arguably the most popular modal semantics nowadays [48, 55]. The basic idea of Kripke semantics is that a formula is true at some worlds, and false at others [62]. The semantics are based on two concepts: Kripke frames and Kripke models. Intuitively, a *Kripke frame* represents a set of worlds together with transitions between them.

Definition 2.1 (Kripke Frames). Let W be a non-empty set of worlds and $R \subseteq W \times W$ a binary *accessibility relation* on it. Then $\mathcal{F} \triangleq (W, R)$ is a *Kripke frame*.

For example, if we wanted to model a pedestrian traffic light as a Kripke frame $\mathcal{F}_{\text{TL}} = (W_{\text{TL}}, R_{\text{TL}})$, it would have two possible states $W_{\text{TL}} = \{\text{red}, \text{green}\}$. Assuming that the light changes at every temporal step, the accessibility relation of \mathcal{F}_{TL} would be $R_{\text{TL}} = \{(\text{red}, \text{green}), (\text{green}, \text{red})\}$. In order to be able to specify properties of the system, we need to fix a set of atomic propositions and specify the set of worlds where each of them holds. This is done by adding an assignment function to the Kripke frame. The result is a *Kripke model*:

Definition 2.2 (Kripke Models). Let $\mathcal{F} = (W, R)$ be a Kripke frame, AP a finite non-empty set of atomic propositions, and $h : AP \rightarrow 2^W$ an *assignment* into \mathcal{F} . Then $\mathcal{M} \triangleq (\mathcal{F}, h)$ is a *Kripke model*.

Suppose we wanted to express a property regarding the current colour of the traffic light. We would introduce two atoms $AP \triangleq \{\text{green}, \text{red}\}$ and use an assignment h_{TL} defined as $h_{\text{TL}}(\text{green}) \triangleq \{\text{green}\}$ and $h_{\text{TL}}(\text{red}) \triangleq \{\text{red}\}$. The resulting Kripke Model would be $\mathcal{M}_{\text{TL}} = (\mathcal{F}_{\text{TL}}, h_{\text{TL}})$.

The Kripke frame and model structures described above have only one accessibility relation R . However, sometimes we need to express multiple modalities within one model (e.g. temporal evolution and agents' knowledge, see Subsection 2.2.6). We can achieve this by augmenting the tuples with the corresponding accessibility relations as follows:

Definition 2.3 (Kripke Frames and Models with Multiple Modalities). Let W be a non-empty set of worlds, $R_1, \dots, R_n \subseteq W \times W$ binary *accessibility relations*, AP a finite non-empty set of atomic propositions, and $h : AP \rightarrow 2^W$ an *assignment*. Then $\mathcal{F}^+ = (W, R_1, \dots, R_n)$ is a *Kripke frame with multiple modalities* and $\mathcal{M}^+ = (\mathcal{F}^+, h)$ a *Kripke model with multiple modalities*.

Although Kripke frames/models have no concept of local states, which are typically used to represent agents' knowledge (see Subsection 2.1.3), it is possible to express epistemic accessibility (see Subsection 2.2.6 for an introduction to epistemic modalities) explicitly using a family of binary relations $\sim_i \subseteq W \times W$ on the set of worlds W for all agents $i \in \text{Agt}$. Intuitively, $w_1 \sim_i w_2$ holds iff agent $i \in \text{Agt}$ cannot distinguish between worlds $w_1, w_2 \in W$. We can then define a Kripke frame and model with epistemic modalities as:

$$\mathcal{F}_{\text{K}} \triangleq \left(W, R, (\sim_i)_{i \in \text{Agt}} \right) \qquad \mathcal{M}_{\text{K}} \triangleq (\mathcal{F}_{\text{K}}, h)$$

Note that even with multiple modalities, Kripke models are still strictly less expressive than concurrent game structures and interpreted systems because they lack the concept of agents' actions.

2.1.2 Concurrent Game Structures

Another framework for modelling systems are concurrent game structures [18], which generalise Kripke semantics by introducing the concept of actions. The temporal evolution of a concurrent game structure depends on the *concurrent actions* of all agents.

Definition 2.4 (Concurrent Game Structures). Let Agt and AP be non-empty sets of agents and atomic propositions respectively. Then a *concurrent game structure* is a tuple $\mathcal{G} \triangleq \langle AP, \text{Agt}, \text{Act}, \text{St}, \lambda, \tau, s_0 \rangle$ where:

- Act is a finite non-empty set of *actions*. $Dc \triangleq \text{Agt} \rightarrow \text{Act}$ is the set of *decisions* which map each agent to its choice of an action.
- St is a finite non-empty set of *states*.
- $\lambda : \text{St} \rightarrow 2^{AP}$ is a *labelling function* that maps each state to the set of atomic propositions true in that state.
- $\tau : \text{St} \times Dc \rightarrow \text{St}$ is a *transition function*. It maps the current state and a decision to the next state.
- $s_0 \in \text{St}$ is an *initial state*.

As we mentioned already, the main difference between Kripke models and concurrent game structures are actions, which enable us to reason about agents' strategies. Apart from them, the frameworks are equivalent¹: They both have a set of states (W, St) , a transition relation/function (R, τ) , and a labelling function² (h, λ) . Similarly to Kripke models, epistemic accessibility can be represented explicitly by a family of binary relations $\sim_i \subseteq St \times St$ for each agent $i \in Agt$.

2.1.3 Interpreted Systems

Interpreted systems are an extension of Kripke semantics for multi-agent systems [65] in which the global states of the whole system are composed of the internal states of all agents.

Definition 2.5 (Interpreted Systems). Let $\Sigma = \{1, \dots, n\}$ be a set of *players*, E a special agent representing the *environment*, $Agt = \Sigma \cup \{E\}$ the set of all *agents*, and AP a finite non-empty set of *atomic propositions*. Then an *interpreted system* is a tuple $\mathcal{I} \triangleq \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ where:

- L_i is a finite non-empty set of *internal states* of an agent $i \in Agt$. A tuple $g = (l_1, \dots, l_n, l_E) \in L_1 \times \dots \times L_n \times L_E$ is called a *global state*. Furthermore, $L_{\sigma E} \triangleq L_\sigma \times L_E$ is the set of *local states* of a player $\sigma \in \Sigma$. The set of local states of the environment is equal to its set of internal states, i.e. $L_{EE} \triangleq L_E$. $l_i(g)$ and $l_{iE}(g)$ denote the internal and local state of agent $i \in Agt$ in a global state $g \in G$ respectively.
- Act_i is a finite non-empty set of *actions* that an agent $i \in Agt$ may perform. $Act \triangleq Act_1 \times \dots \times Act_n \times Act_E$ denotes the set of *joint actions* of all agents. $a_i(a)$ denotes the action of agent $i \in Agt$ in the joint action $a \in Act$. $Act_A \triangleq \bigcap_{i \in A} Act_i$ is the set of *shared actions* of agents $A \subseteq Agt$ with $A \neq \emptyset$.
- $P_i : L_{iE} \rightarrow 2^{Act_i}$ is the *protocol* of an agent $i \in Agt$. For every local state $l_{iE} \in L_{iE}$ of the agent, $P_i(l_{iE}) \neq \emptyset$ is the set of actions available to agent i . The *global protocol* $P : G \rightarrow 2^{Act}$ is defined as $P(g) \triangleq \left\{ a \in Act \mid \bigwedge_{i \in Agt} a_i(a) \in P_i(l_{iE}(g)) \right\}$ for all global states $g \in G$.
- $t_i : L_{iE} \times Act \rightarrow L_i$ is the *evolution function* of the internal state of an agent $i \in Agt$.
- $I \subseteq L_1 \times \dots \times L_n \times L_E$ is a finite non-empty set of *initial* global states. $G \subseteq L_1 \times \dots \times L_n \times L_E$ is the set of *reachable* global states obtained by considering all the possible evolutions of the system from I .
- $h : AP \rightarrow 2^G$ is a *valuation function* that maps each propositional variable to the set of global reachable states in which it is true.

The evolution of the system is described by a *global evolution function* $t : G \times Act \rightarrow G$. It is defined as $t(g, a) = g'$ iff $\forall i \in Agt. t_i(l_{iE}(g), a) = l_i(g')$.

Intuitively, the environment agent E represents the part of the model that all agents can “see”. That is why its internal states $l_E \in L_E$ are included in the other agents' local states $(l_i, l_E) \in L_{iE}$. Note that the distinction between internal and local states of agents is often not emphasised in the literature. Internal states are sometimes referred to as *private local states* [65].

Interpreted systems are better suited for epistemic modalities expressing agents' knowledge (see Subsection 2.2.6) because they distinguish between local and global states (unlike Kripke models and concurrent game structures). The relationship between interpreted systems and concurrent game structures is discussed in [65].

We will consider interpreted systems exclusively in the rest of this report because they best represent multi-agent systems (for the reasons outlined above), whose verification our project focuses on.

¹The explicit initial state s_0 does not increase the expressiveness of concurrent game structures. We could simulate it in Kripke semantics by introducing a new atom $init \in AP$ and setting $h(init) \triangleq \{s_0\}$.

²While the signatures of $h : AP \rightarrow 2^W$ and $\lambda : St \rightarrow 2^{AP}$ differ, they are equivalent (assuming $St = W$) since $h(p) = \lambda^{-1}(p) = \{s \in St \mid p \in \lambda(s)\}$ and $\lambda(s) = h^{-1}(s) = \{p \in AP \mid s \in h(p)\}$ where $p \in AP$ and $s \in St$.

2.2 Specification Languages

Given a property P , we aim to describe it as a formula φ_P . We discuss multiple logics of increasing expressiveness for specifying φ_P . Unfortunately, higher expressiveness is usually accompanied by worse model checking complexity as we shall see in Subsection 2.2.7.

In order to compare the expressiveness of the logics described in this section, we will express the following sample properties using each of them (if possible):

1. A socket is opened and it will stay opened in the next two states.
2. If a device is currently connected, it *can* be disconnected in the future.
3. A task will be finished at some point in the future (no matter what happens).
4. It is possible to permanently delete a file (at some point in the future).
5. If the light is infinitely often red, it will also be infinitely often green (in the same evolution).
6. If a car *can* turn infinitely often left, it *can* also turn infinitely often right (possibly in a different evolution).
7. A firewall and an antivirus can (together) ensure that a computer will never be hacked.
8. If both players play the same strategies, neither of them will ever win.

2.2.1 Linear Temporal Logic

Linear Temporal Logic (LTL) [80] is a temporal logic which models time as a sequence of states, extending infinitely into the future [49]. These possible sequences of states are referred to as *paths*.

Definition 2.6 (Paths and Tracks). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system. A *path* $\pi \in G^\omega$ in \mathcal{I} is an infinite sequence of global states $g_0, g_1, g_2, \dots \in G$ such that, $\forall i \geq 0 \exists a \in Act. g_{i+1} = t(g_i, a)$. We write $\pi_{\geq i}$ for the suffix of π starting at g_i , and $\pi(i) = g_i$. $\mathbf{path}(g)$ denotes the set of all paths starting at a global state $g \in G$. $\mathbf{Pth} \triangleq \bigcup_{g \in G} \mathbf{path}(g)$ is the set of all paths in \mathcal{I} .

A *track* $\tau \in G^+$ in \mathcal{I} is a non-empty finite sequence of global states $g_0, g_1, \dots, g_{n-1} \in G$ such that, $\forall i \in [0..n-2] \exists a \in Act. g_{i+1} = t(g_i, a)$. The *length* of τ is $|\tau| \triangleq n$. $\mathbf{last}(\tau) \triangleq g_{n-1}$ refers to the *last global state* on τ . $\mathbf{track}(g)$ denotes the set of all tracks starting at a global state $g \in G$. $\mathbf{Trk} \triangleq \bigcup_{g \in G} \mathbf{track}(g)$ is the set of all tracks in \mathcal{I} .

Intuitively, a path represents a possible evolution of an interpreted system from a given global state $g \in G$ (not necessarily in I). In order to express properties over the possible evolutions of a system, LTL augments propositional logic with temporal connectives³ X , F , G , U , W , and R .

Definition 2.7 (LTL Syntax). LTL *formulas* are built inductively from the set of atomic propositions AP , by using the following grammar, where $p \in AP$:

$$\varphi ::= p \mid \top \mid \perp \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi \mid \varphi \leftrightarrow \varphi \mid X\varphi \mid F\varphi \mid G\varphi \mid \varphi U \varphi \mid \varphi W \varphi \mid \varphi R \varphi$$

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

The syntax allows for the construction of ambiguous formulas like $XaUb$, which can be interpreted as either $(Xa)Ub$, or $X(aUb)$. In order to resolve such ambiguities, the following binding priorities are used [49]: Unary operators (\neg , X , F , G) bind most tightly. The (decreasing) order of the other operators is U , R , W , \wedge , \vee , \rightarrow , and \leftrightarrow . Hence, the original formula is interpreted as $(Xa)Ub$ because X binds more tightly than U . These rules allow us to use fewer brackets. Note that we do not specify the associativity of binary operators. Thus, the formula $aUbUc$ is not well-formed, i.e. it needs brackets to disambiguate between $aU(bUc)$ and $(aUb)Uc$. We shall use similar rules for other logics.

The semantics of LTL are traditionally defined on Kripke models (see Definition 2.2). We will however define them on interpreted systems (see Definition 2.5) here for consistency. Note that, in order to save space, we shall assume throughout the rest of this report that AP is a finite set of propositional atoms.

³ G , F , and X are sometimes written symbolically as \square , \diamond , and \circ respectively.

Definition 2.8 (LTL Semantics). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system, and $\pi \in Pth$ a path in \mathcal{I} . We define $\mathcal{I}, \pi \models_{LTL} \varphi$ by induction on $\varphi \in LTL$:

- $\mathcal{I}, \pi \models_{LTL} \top$;
- $\mathcal{I}, \pi \models_{LTL} p$ iff $\pi(0) \in h(p)$ for $p \in AP$;
- $\mathcal{I}, \pi \models_{LTL} \neg\varphi$ iff $\mathcal{I}, \pi \not\models_{LTL} \varphi$;
- $\mathcal{I}, \pi \models_{LTL} \varphi_1 \wedge \varphi_2$ iff $\mathcal{I}, \pi \models_{LTL} \varphi_1$ and $\mathcal{I}, \pi \models_{LTL} \varphi_2$;
- $\mathcal{I}, \pi \models_{LTL} \varphi_1 \vee \varphi_2$ iff $\mathcal{I}, \pi \models_{LTL} \varphi_1$ or $\mathcal{I}, \pi \models_{LTL} \varphi_2$;
- $\mathcal{I}, \pi \models_{LTL} \varphi_1 \rightarrow \varphi_2$ iff $\mathcal{I}, \pi \not\models_{LTL} \varphi_1$ or $\mathcal{I}, \pi \models_{LTL} \varphi_2$;
- $\mathcal{I}, \pi \models_{LTL} \varphi_1 \leftrightarrow \varphi_2$ iff either $\mathcal{I}, \pi \not\models_{LTL} \varphi_1$ and $\mathcal{I}, \pi \not\models_{LTL} \varphi_2$, or $\mathcal{I}, \pi \models_{LTL} \varphi_1$ and $\mathcal{I}, \pi \models_{LTL} \varphi_2$;
- $\mathcal{I}, \pi \models_{LTL} X\varphi$ iff $\mathcal{I}, \pi_{\geq 1} \models_{LTL} \varphi$;
- $\mathcal{I}, \pi \models_{LTL} F\varphi$ iff there exists some $i \geq 0$ such that $\mathcal{I}, \pi_{\geq i} \models_{LTL} \varphi$;
- $\mathcal{I}, \pi \models_{LTL} G\varphi$ iff for all $i \geq 0$ we have $\mathcal{I}, \pi_{\geq i} \models_{LTL} \varphi$;
- $\mathcal{I}, \pi \models_{LTL} \varphi_1 U \varphi_2$ iff there exists some $i \geq 0$ such that $\mathcal{I}, \pi_{\geq i} \models_{LTL} \varphi_2$ and for all numbers $0 \leq j < i$ we have $\mathcal{I}, \pi_{\geq j} \models_{LTL} \varphi_1$;
- $\mathcal{I}, \pi \models_{LTL} \varphi_1 W \varphi_2$ iff either (i) there exists some $i \geq 0$ such that $\mathcal{I}, \pi_{\geq i} \models_{LTL} \varphi_2$ and for all $0 \leq j < i$ we have $\mathcal{I}, \pi_{\geq j} \models_{LTL} \varphi_1$, or (ii) for all $k \geq 0$ we have $\mathcal{I}, \pi_{\geq k} \models_{LTL} \varphi_1$;
- $\mathcal{I}, \pi \models_{LTL} \varphi_1 R \varphi_2$ iff either (i) there exists some $i \geq 0$ such that $\mathcal{I}, \pi_{\geq i} \models_{LTL} \varphi_1$ and for all $0 \leq j \leq i$ we have $\mathcal{I}, \pi_{\geq j} \models_{LTL} \varphi_2$, or (ii) for all $k \geq 0$, we have $\mathcal{I}, \pi_{\geq k} \models_{LTL} \varphi_2$.

A formula $\varphi \in LTL$ is true at a global state $g \in G$ in an interpreted system \mathcal{I} (written as $\mathcal{I}, g \models_{LTL} \varphi$) iff it is true on all paths $\pi \in path(g)$ starting at g (i.e. $\mathcal{I}, \pi \models_{LTL} \varphi$).

The intuitive meanings of the temporal operators X, F, G, U, W, and R are “neXt state”, “some Future state”, “all future states (Globally)”, “Until”, “Weak-until”, and “Release” respectively [49]. For example, the LTL formula $a U b$ expresses the property that a holds *until* b holds. There also exist past temporal connectives Y, S, O, and H which are opposite to X, U, F, and G respectively [49]. Their intuitive meanings are “Yesterday”, “Since”, “Once”, and “Historically”. However, adding these operators does *not* increase the expressiveness of LTL.

From this point onwards, we will omit the operators \perp , \rightarrow , and \leftrightarrow in all logics for the sake of conciseness. We can always express them using the following propositional equivalences:

$$\perp \equiv \neg\top \qquad \varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2 \qquad \varphi_1 \leftrightarrow \varphi_2 \equiv (\varphi_1 \wedge \varphi_2) \vee (\neg\varphi_1 \wedge \neg\varphi_2)$$

In fact, the set $\{\top, \neg, \wedge, X, U\}$ forms an *adequate set of connectives*⁴ for LTL [49], i.e. all the other operators (\vee, F, G, W , and R) can be expressed using equivalences:

$$\begin{aligned} \varphi_1 \vee \varphi_2 &\equiv \neg(\neg\varphi_1 \wedge \neg\varphi_2) & F\varphi &\equiv \top U \varphi & G\varphi &\equiv \neg F \neg\varphi \equiv \neg(\top U \neg\varphi) \\ \varphi_1 W \varphi_2 &\equiv \varphi_1 U \varphi_2 \vee G\varphi_1 & \varphi_1 U \varphi_2 &\equiv \varphi_1 U \varphi_2 \vee \neg(\top U \neg\varphi_1) & \varphi_1 R \varphi_2 &\equiv \neg(\neg\varphi_1 U \neg\varphi_2) \end{aligned}$$

Model checking of an LTL formula φ is typically performed by constructing a non-deterministic Büchi automaton $\mathfrak{A}_{\neg\varphi} = (S_{\mathfrak{A}}, 2^{AP}, I_{\mathfrak{A}}, R_{\mathfrak{A}}, \mathcal{F})$ equivalent to $\neg\varphi$ (see Section 2.4) and then checking if no fair path exists in the product of the interpreted system and the automaton. More precisely, $\mathcal{I}, g \not\models_{LTL} \varphi$ iff there exists an initial state $s_I \in I_{\mathfrak{A}}$ of the automaton and a path $\pi \in path((g, s_I))$ in $\mathcal{I} \times \mathfrak{A}_{\neg\varphi}$ starting at (g, s_I) such that \mathcal{F} is visited infinitely often along π . Please refer to [34] for more details about how this is performed in practice by converting LTL model checking to CTL model checking with fairness

⁴Note that an adequate set is *not* necessarily unique for a given logic. For example, we could replace \wedge with \vee and/or U with R and the result would still be an adequate set for LTL.

conditions using the tableau method. Another popular approach to LTL model checking is *bounded model checking* [32], which transforms the model checking problem into a propositional satisfiability problem.

We shall now demonstrate the expressive power of LTL. The sample properties listed at the beginning of this section can be expressed in LTL as follows:

1. $open \wedge X open \wedge XX open$.
2. This property cannot be expressed directly in LTL because it requires existential quantification (while LTL formulas are implicitly universally quantified over *all* paths). However, it is possible to express its *converse* as $connected \wedge G connected$. If the converse does *not* hold in a state, then the property is satisfied. Note that this shows that $\mathcal{I}, g \models_{LTL} \neg\varphi$ is in general *not equivalent* to $\mathcal{I}, g \not\models_{LTL} \varphi$ for a global state $g \in G$. Compare this with paths $\pi \in Pth$ for which $\mathcal{I}, \pi \models_{LTL} \neg\varphi$ is defined as $\mathcal{I}, \pi \not\models_{LTL} \varphi$.
3. $F finished$.
- 4, 6. Neither the properties nor their converses can be expressed in LTL because they combine both existential and universal quantification⁵.
- 5 $GF red \rightarrow GF green$.
- 7–8. These properties cannot be expressed in LTL because it cannot refer to agents' strategies.

We can see that LTL can only express properties that should hold on *all* paths. While it is possible to assert the existence of a path (see item 2 above), existential and universal path quantification cannot be combined within a single LTL formula (see items 4 and 6). Certain important properties therefore cannot be expressed in LTL. This shortcoming is addressed by Computation Tree Logic.

2.2.2 Computation Tree Logic

Computation Tree Logic (CTL) [33] is a temporal logic which allows us to quantify over paths explicitly [49]. Compared to LTL, CTL adds quantifiers A and E expressing “for all paths” and “there exists a path” respectively. Again, CTL extends the syntax of propositional logic with path and temporal connectives:

Definition 2.9 (CTL Syntax). CTL *formulas* are built inductively from the set of atomic propositions AP , by using the following grammar, where $p \in AP$:

$$\varphi ::= p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid AX\varphi \mid EX\varphi \mid AF\varphi \mid EF\varphi \mid AG\varphi \mid EG\varphi \mid A[\varphi U \varphi] \mid E[\varphi U \varphi]$$

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

Observe that every path quantifier is coupled with a temporal operator and vice versa. The binding priorities of the operators are the same as for LTL (e.g. $EF a \wedge b$ is interpreted as $(EF a) \wedge b$). Similarly to LTL, we define CTL semantics on interpreted systems (see Definition 2.5) for consistency:

Definition 2.10 (CTL Semantics). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system, and $g \in G$ a global state. We define $\mathcal{I}, g \models_{CTL} \varphi$ by induction on $\varphi \in LTL$:

- $\mathcal{I}, g \models_{CTL} \top$;
- $\mathcal{I}, g \models_{CTL} p$ iff $g \in h(p)$ for $p \in AP$;
- $\mathcal{I}, g \models_{CTL} \neg\varphi$ iff $\mathcal{I}, g \not\models_{CTL} \varphi$;
- $\mathcal{I}, g \models_{CTL} \varphi_1 \wedge \varphi_2$ iff $\mathcal{I}, g \models_{CTL} \varphi_1$ and $\mathcal{I}, g \models_{CTL} \varphi_2$;
- $\mathcal{I}, g \models_{CTL} \varphi_1 \vee \varphi_2$ iff $\mathcal{I}, g \models_{CTL} \varphi_1$ or $\mathcal{I}, g \models_{CTL} \varphi_2$;

⁵This might not be immediately obvious in property 6. Recall that $\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$. Hence the property can be interpreted as: on all paths the car does not turn infinitely often left or there is a path on which the car turns right infinitely often.

- $\mathcal{I}, g \models_{\text{CTL}} \text{AX } \varphi$ iff for all paths $\pi \in \text{path}(g)$, we have $\mathcal{I}, \pi(1) \models_{\text{CTL}} \varphi$;
- $\mathcal{I}, g \models_{\text{CTL}} \text{EX } \varphi$ iff there exists a path $\pi \in \text{path}(g)$ such that $\mathcal{I}, \pi(1) \models_{\text{CTL}} \varphi$;
- $\mathcal{I}, g \models_{\text{CTL}} \text{AF } \varphi$ iff for all paths $\pi \in \text{path}(g)$ there is some $i \geq 0$ such that $\mathcal{I}, \pi(i) \models_{\text{CTL}} \varphi$;
- $\mathcal{I}, g \models_{\text{CTL}} \text{EF } \varphi$ iff there exist a path $\pi \in \text{path}(g)$ and $i \geq 0$ such that $\mathcal{I}, \pi(i) \models_{\text{CTL}} \varphi$;
- $\mathcal{I}, g \models_{\text{CTL}} \text{AG } \varphi$ iff for all paths $\pi \in \text{path}(g)$ and $i \geq 0$ we have $\mathcal{I}, \pi(i) \models_{\text{CTL}} \varphi$;
- $\mathcal{I}, g \models_{\text{CTL}} \text{EG } \varphi$ iff there exists a path $\pi \in \text{path}(g)$ such that for all $i \geq 0$ we have $\mathcal{I}, \pi(i) \models_{\text{CTL}} \varphi$;
- $\mathcal{I}, g \models_{\text{CTL}} \text{A}[\varphi_1 \text{ U } \varphi_2]$ iff for all paths $\pi \in \text{path}(g)$ there exists $i \geq 0$ such that $\mathcal{I}, \pi(i) \models_{\text{CTL}} \varphi_2$, and for all $0 \leq j < i$, $\mathcal{I}, \pi(j) \models_{\text{CTL}} \varphi_1$.
- $\mathcal{I}, g \models_{\text{CTL}} \text{E}[\varphi_1 \text{ U } \varphi_2]$ iff there exists a path $\pi \in \text{path}(g)$ and $i \geq 0$ such that $\mathcal{I}, \pi(i) \models_{\text{CTL}} \varphi_2$, and for all $0 \leq j < i$, $\mathcal{I}, \pi(j) \models_{\text{CTL}} \varphi_1$.

The intuitive meanings of the temporal operators are the same as for LTL (see Subsection 2.2.1). For example, $\text{EG } p$ means that there exists an infinite path (starting in the current global state) such that p holds at every single state of the path. One possible adequate set of connectives for CTL is $\{\top, \neg, \wedge, \text{EX}, \text{AF}, \text{EU}\}$ [49].

Model checking of a CTL formula φ is typically⁶ performed by recursively calculating the sets of states $\|\psi\|_{\mathcal{I}} \triangleq \{g \in G \mid \mathcal{I}, g \models_{\text{CTL}} \psi\}$ of the interpreted system \mathcal{I} at which subformulas ψ of φ hold in a bottom-up manner. For example, assuming that we have already calculated $\|\psi\|_{\mathcal{I}}$, $\|\text{EX } \psi\|_{\mathcal{I}}$ is simply the set of global states of \mathcal{I} which have a possible transition into $\|\psi\|_{\mathcal{I}}$, which we denote $\text{pre}_{\exists}(\|\psi\|_{\mathcal{I}})$.

Definition 2.11 (Predecessors). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system and $X \subseteq G$ a set of global states. The *existential* and *universal predecessors* of X are defined as $\text{pre}_{\exists}(X) \triangleq \{g \in G \mid \exists a \in P(g). t(g, a) \in X\}$ and $\text{pre}_{\forall}(X) \triangleq \{g \in G \mid \forall a \in P(g). t(g, a) \in X\}$ respectively.

Once we have calculated $\|\varphi\|_{\mathcal{I}}$, determining whether the formula φ holds at a given global state $g \in G$ (i.e. $\mathcal{I}, g \models_{\text{CTL}} \varphi$) is equivalent to checking whether $g \in \|\varphi\|_{\mathcal{I}}$. The recursive model checking algorithm $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\cdot)$ for CTL, which calculates $\|\varphi\|_{\mathcal{I}}$, is defined as follows [49]:

Definition 2.12 (CTL Model Checking Algorithm). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system and $\varphi \in \text{CTL}$ a CTL formula. The function $\text{SAT}_{\text{CTL}}^{\mathcal{I}} : \text{CTL} \rightarrow 2^G$ is defined inductively as follows:

- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\top) \triangleq G$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(p) \triangleq h(p)$ for $p \in AP$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\neg\varphi) \triangleq G \setminus \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi)$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_1 \wedge \varphi_2) \triangleq \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_1) \cap \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_2)$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_1 \vee \varphi_2) \triangleq \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_1) \cup \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_2)$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{AX } \varphi) \triangleq \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\neg\text{EX } \neg\varphi)$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{EX } \varphi) \triangleq \text{pre}_{\exists}(\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi))$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{AF } \varphi) \triangleq \text{lfp}_X[\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi) \cup \text{pre}_{\forall}(X)]$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{EF } \varphi) \triangleq \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{E}[\top \text{ U } \varphi])$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{AG } \varphi) \triangleq \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\neg\text{EF } \neg\varphi)$;

⁶Alternative approaches exist. For example, a model checking algorithm for ACTL (a universal fragment of CTL where all path quantifiers are A) based on *bounded model checking* is presented in [87].

- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{EG } \varphi) \triangleq \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\neg \text{AF } \neg \varphi)$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{A}[\varphi_1 \text{ U } \varphi_2]) \triangleq \text{SAT}_{\text{CTL}}^{\mathcal{I}}(\neg (\text{E}[\neg \varphi_2 \text{ U } (\neg \varphi_1 \wedge \neg \varphi_2)] \vee \text{EG } \neg \varphi_2))$;
- $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\text{E}[\varphi_1 \text{ U } \varphi_2]) \triangleq \text{lfp}_X [\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_2) \cup (\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi_1) \cap \text{pre}_{\exists}(X))]$;

where lfp denotes the least fixed point.

Given a CTL formula φ , the algorithm calculates the set of global states of \mathcal{I} at which it holds. This fact is stated formally in the following proposition.

Proposition 2.1. Let \mathcal{I} be an interpreted system and $\varphi \in \text{CTL}$ an arbitrary CTL formula. Then we have $\text{SAT}_{\text{CTL}}^{\mathcal{I}}(\varphi) = \|\varphi\|_{\mathcal{I}} \triangleq \{g \in G \mid \mathcal{I}, g \models_{\text{CTL}} \varphi\}$.

Similarly to LTL, we now demonstrate the expressive power of CTL. The sample properties listed at the beginning of this section can be expressed in CTL as follows:

1. $\text{open} \wedge \text{AX open} \wedge \text{AXAX open}$.
2. $\text{connected} \rightarrow \text{EF } \neg \text{connected}$.
3. AF finished .
4. EF AG deleted .

6–8. These properties cannot be expressed in CTL because it does not refer to paths (properties 5 and 6⁷) and agents' strategies (properties 7 and 8).

Unlike LTL, CTL can express properties which combine existential and universal path quantifiers (property 4). However, it is limited by the fact that every temporal operator must be coupled with a path quantifier, which prevents CTL from expressing path specifications like property 5. Intuitively, we would like to express it as $\text{A}[\text{GF red} \rightarrow \text{GF green}]$, which is exactly what Full Branching Time Logic allows us to do.

2.2.3 Full Branching Time Logic

Full Branching Time Logic (CTL*) [37] is a superset of both LTL and CTL [49]. Moreover, it is strictly more expressive than the union of LTL and CTL, i.e. there exist CTL* formulas which cannot be expressed in either of the two logics. CTL* supports CTL path quantifiers A and E but relaxes the restrictions on their coupling with temporal operators:

Definition 2.13 (CTL* Syntax). CTL* *formulas* are built inductively from the set of atomic propositions AP . State formulas φ and path formulas ψ are defined by the following mutually recursive grammar, where $p \in AP$:

$$\begin{aligned} \varphi &::= p \mid \top \mid \neg \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \text{A}[\psi] \mid \text{E}[\psi] \\ \psi &::= \varphi \mid \neg \psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \text{X}\psi \mid \text{F}\psi \mid \text{G}\psi \mid \psi \text{ U } \psi \end{aligned}$$

CTL* denotes the infinite set of formulas generated by the above rules.

CTL* semantics are very similar to LTL and CTL semantics with only a few adjustments. We will also define them on interpreted systems (see Definition 2.5) for consistency:

Definition 2.14 (CTL* Semantics). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system, $g \in G$ a global state, and $\pi \in \text{Pth}$ a path in \mathcal{I} . We define $\mathcal{I}, g \models_{\text{CTL}^*} \varphi$ and $\mathcal{I}, \pi \models_{\text{CTL}^*} \psi$ by induction on state formulas φ and path formulas ψ :

⁷It would appear that the CTL formula $\text{EGEF } \varphi$ means that there exists a path along which φ holds infinitely often. Unfortunately, that is not the case. Consider a simple interpreted system with three global states g_1, g_2, g_3 , four possible transitions $g_1 \rightarrow g_1, g_1 \rightarrow g_2, g_2 \rightarrow g_3, \text{ and } g_3 \rightarrow g_3$, and an assignment $h(p) \triangleq \{g_2\}$. Then we have $\mathcal{I}, g_2 \models_{\text{CTL}} p$, so $\mathcal{I}, g_1 \models_{\text{CTL}} \text{EF } p$ (g_1 is an existential predecessor of g_2) and $\mathcal{I}, g_1 \models_{\text{CTL}} \text{EGEF } p$ (cycle in g_1). However, there is no path starting in g_1 along which p would be infinitely often true. A similar relationship between the CTL formula $\text{AGAF } p$ and the LTL formula $\text{GF } p$ is shown in [49].

- $\mathcal{I}, g \models_{\text{CTL}^*} A[\psi]$ iff for all paths $\pi \in \text{path}(g)$ starting at g we have $\mathcal{I}, \pi \models_{\text{CTL}^*} \psi$;
- $\mathcal{I}, g \models_{\text{CTL}^*} E[\psi]$ iff there exists a path $\pi \in \text{path}(g)$ starting at g such that $\mathcal{I}, \pi \models_{\text{CTL}^*} \psi$;
- $\mathcal{I}, g \models_{\text{CTL}^*} \varphi$ is defined in the same way as $\mathcal{I}, g \models_{\text{CTL}} \varphi$ for other CTL^* state formulas;
- $\mathcal{I}, \pi \models_{\text{CTL}^*} \varphi$ iff $\mathcal{I}, \pi(0) \models_{\text{CTL}^*} \varphi$;
- $\mathcal{I}, \pi \models_{\text{CTL}^*} \psi$ is defined in the same way as $\mathcal{I}, \pi \models_{\text{LTL}} \psi$ for other CTL^* path formulas.

Since LTL and CTL are subsets of CTL^* , any LTL/CTL formula can be transformed into an equivalent CTL^* formula:

- Let $\psi \in \text{LTL}$ be an LTL formula. It is equivalent to the CTL^* formula $A[\psi]$. This can be seen from the semantics of LTL:

$$\begin{aligned} \mathcal{I}, g \models_{\text{LTL}} \psi & \text{ iff } \forall \pi \in \text{path}(g). \mathcal{I}, \pi \models_{\text{LTL}} \psi \\ & \text{ iff } \forall \pi \in \text{path}(g). \mathcal{I}, \pi \models_{\text{CTL}^*} \psi \\ & \text{ iff } \mathcal{I}, g \models_{\text{CTL}^*} A[\psi] \end{aligned}$$

where the second equivalence follows from the fact that ψ contains no path quantifiers.

- CTL is a subset of CTL^* in which (i) we do not allow Boolean combinations of path formulas, and (ii) each temporal connective must be preceded by a path quantifier. An alternative definition of CTL syntax can thus be obtained by restricting path formulas in CTL^* syntax (Definition 2.13):

$$\psi ::= X\varphi \mid F\varphi \mid G\varphi \mid \varphi U \varphi$$

Hence, any CTL formula is also an equivalent CTL^* formula.

CTL^* model checking can be performed in the following bottom-up manner⁸: In order to calculate $\|A[\varphi]\|$, we replace each subformula $A[\psi]$ of φ with a new atom $p \in AP$ such that $h(p) \triangleq \|A[\psi]\|_{\mathcal{I}}$, which we calculate recursively. φ is now an LTL formula, so we can use the method involving automata outlined in Subsection 2.2.1 (or any other method for model checking LTL) to calculate $\|A[\varphi]\|_{\mathcal{I}}$. Please refer to [39] for more details about this reduction from CTL^* to LTL model checking.

Let us now investigate the expressive power of CTL^* . The sample properties listed at the beginning of this section can be expressed in CTL^* as follows:

1. $A[\text{open} \wedge X \text{open} \wedge XX \text{open}]$.
2. $\text{connected} \rightarrow E[F \neg \text{connected}]$.
3. $A[F \text{finished}]$.
4. $E[FA[G \text{deleted}]]$.
5. $A[G F \text{red} \rightarrow G F \text{green}]$.
6. $E[G F \text{left}] \rightarrow E[G F \text{right}]$.

7–8. These properties cannot be expressed in CTL^* because it does not refer to agents' strategies.

We can see that CTL^* can express all properties supported by LTL (1, 3, 5) and CTL (1–4). Moreover, it can express property 6, which could not be represented as a LTL or CTL formula. The only specifications (on our list) which CTL^* cannot express are those that involve reasoning about agents' ability to enforce certain properties on the system. This issue is addressed by Alternating-Time Temporal Logic.

⁸To make things simpler, we only consider the universal path quantifier A. The existential quantifier E can be equivalently expressed as $\neg A \neg$.

2.2.4 Alternating-Time Temporal Logic

Alternating-Time Temporal Logic (ATL) [18] is an extension of CTL for multi-agent systems. Instead of universal and existential path quantification, ATL allows for a more fine-grained control by parametrising temporal operators with sets of agents. This makes ATL more suitable for reasoning about open systems (such as multi-agent systems), which interact with their environment [18]. To do this, ATL introduces the $\langle\langle A \rangle\rangle$ quantifier, which ranges over the set of paths that can be enforced by the agents $A \subseteq \text{Agt}$ (regardless of what action the other agents $\text{Agt} \setminus A$ carry out) and replaces the CTL path quantifiers A and E.

Definition 2.15 (ATL Syntax). ATL *formulas* are built inductively from the set of atomic propositions AP and agents Agt , by using the following grammar, where $p \in AP$ and $A \subseteq \text{Agt}$:

$$\varphi ::= p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle\langle A \rangle\rangle X \varphi \mid \langle\langle A \rangle\rangle F \varphi \mid \langle\langle A \rangle\rangle G \varphi \mid \langle\langle A \rangle\rangle [\varphi \cup \varphi]$$

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

Unlike LTL and CTL semantics, ATL semantics are traditionally defined on concurrent game structures (see Definition 2.4). However, we will define them here on interpreted systems (see Definition 2.5) like we did for the previous formalisms for consistency. We first need to define the concept of a *strategy*.

Definition 2.16 (Individual Strategies). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system and $i \in \text{Agt}$ an agent. Then a partial function $f_i : \text{Trk} \rightarrow \text{Act}_i$ such that for all $\tau \in \text{dom}(f_i)$, $f_i(\tau) \in P_i(l_{iE}(\text{last}(\tau)))$, is an (*individual*) *strategy* for agent i which maps tracks to actions. $\text{Str}_i \subseteq \text{Trk} \rightarrow \text{Act}_i$ denotes the set of all strategies for agent i in \mathcal{I} . The set of all strategies is defined as $\text{Str} \triangleq \bigcup_{i \in \text{Agt}} \text{Str}_i$.

Let $A \subseteq \text{Agt}$ be a set of agents. Then an *agent assignment* is a function $\alpha : A \rightarrow \text{Str}$ which maps each agent $i \in A$ to a strategy $\alpha(i) \in \text{Str}_i$. The set of all agent assignment for agents A is denoted as AAsg_A . The set of all agent assignments is defined as $\text{AAsg} \triangleq \bigcup_{A \subseteq \text{Agt}} \text{AAsg}_A$.

Intuitively, an agent assignment maps agents to strategies, which in turn map histories of the system to the next actions of the agents. Given an agent assignment $\alpha \in \text{AAsg}_A$ for a group of agents $A \subseteq \text{Agt}$, we want to find the set of paths that the agents can enforce (if they follow α) regardless of the other agents' behaviour. This is referred to as the set of *outcomes*.

Definition 2.17 (Outcomes). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system, $g \in G$ a global state, $A \subseteq \text{Agt}$ a set of agents, and $\alpha \in \text{AAsg}_A$ an agent assignment for A . The set of *outcomes* $\text{out}(g, \alpha)$ contains all paths $\pi \in \text{Pth}$ such that (i) $\pi(0) = g$, and (ii) for all $k \geq 0$, there exists a joint action $a \in P(\pi(k))$, such that $t(\pi(k), a) = \pi(k+1)$ and for all $i \in A$, if $\pi_{\leq k} \in \text{dom}(\alpha(i))$, then $a_i = \alpha(i)(\pi_{\leq k})$.

Finally, we are ready to define ATL semantics. Intuitively, a formula $\langle\langle A \rangle\rangle \varphi$ is true iff the agents $A \subseteq \text{Agt}$ can enforce φ .

Definition 2.18 (ATL Semantics). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system, $g \in G$ a global state, and $A \subseteq \text{Agt}$ a set of agents. We define $\mathcal{I}, g \models_{\text{ATL}} \varphi$ by induction on $\varphi \in \text{ATL}$:

- $\mathcal{I}, g \models_{\text{ATL}} \top$;
- $\mathcal{I}, g \models_{\text{ATL}} p$ iff $g \in h(p)$ for $p \in AP$;
- $\mathcal{I}, g \models_{\text{ATL}} \neg\varphi$ iff $\mathcal{I}, g \not\models_{\text{ATL}} \varphi$;
- $\mathcal{I}, g \models_{\text{ATL}} \varphi_1 \wedge \varphi_2$ iff $\mathcal{I}, g \models_{\text{ATL}} \varphi_1$ and $\mathcal{I}, g \models_{\text{ATL}} \varphi_2$;
- $\mathcal{I}, g \models_{\text{ATL}} \varphi_1 \vee \varphi_2$ iff $\mathcal{I}, g \models_{\text{ATL}} \varphi_1$ or $\mathcal{I}, g \models_{\text{ATL}} \varphi_2$;
- $\mathcal{I}, g \models_{\text{ATL}} \langle\langle A \rangle\rangle X \varphi$ iff there exists an agent assignment $\alpha \in \text{AAsg}_A$, such that for all paths $\pi \in \text{out}(g, \alpha)$ we have $\mathcal{I}, \pi(1) \models_{\text{ATL}} \varphi$;

- $\mathcal{I}, g \models_{\text{ATL}} \langle\langle A \rangle\rangle \text{F} \varphi$ iff there exists an agent assignment $\alpha \in \text{AAsg}_A$, such that for all paths $\pi \in \text{out}(s, \alpha)$ there exists $i \geq 0$ such that $\mathcal{I}, \pi(i) \models_{\text{ATL}} \varphi$;
- $\mathcal{I}, g \models_{\text{ATL}} \langle\langle A \rangle\rangle \text{G} \varphi$ iff there exists an agent assignment $\alpha \in \text{AAsg}_A$, such that for all paths $\pi \in \text{out}(s, \alpha)$ and $i \geq 0$ we have $\mathcal{I}, \pi(i) \models_{\text{ATL}} \varphi$;
- $\mathcal{I}, g \models_{\text{ATL}} \langle\langle A \rangle\rangle [\varphi_1 \text{U} \varphi_2]$ iff there exists an agent assignment $\alpha \in \text{AAsg}$, such that for all paths $\pi \in \text{out}(s, \alpha)$ there exists $i \geq 0$ such that $\mathcal{I}, \pi(i) \models_{\text{ATL}} \varphi_2$ and for all $0 \leq j < i$ we have $\mathcal{I}, \pi(j) \models_{\text{ATL}} \varphi_1$.

The meaning of the temporal operators is the same as in LTL and CTL. For example, the ATL formula $\langle\langle \{a, b\} \rangle\rangle \text{G} p$ means that there are strategies for agents a and b such that no matter what the other agents do, p will be always true. We can now see the relationship between CTL and ATL: Given a CTL formula $\varphi \in \text{CTL}$, an equivalent ATL formula $\varphi' \in \text{ATL}$ can be obtained by replacing every subformula of the form $\text{E}[\psi]$ and $\text{A}[\psi]$ with $\langle\langle \text{Agt} \rangle\rangle \psi$ and $\langle\langle \emptyset \rangle\rangle \psi$ respectively [18].

Similarly to CTL* (see Subsection 2.2.3), which extends CTL by relaxing its syntax constraints (see Subsection 2.2.2), *Full Alternating-Time Logic* (ATL*) extends ATL by relaxing its syntax constraints, i.e. (i) ATL* syntax distinguishes between state and path formulas and (ii) quantifiers ($\langle\langle A \rangle\rangle$) and temporal operators (X, F, G, U) need not be coupled any more. ATL* semantics are modified analogously (see Definition 2.14). The result is a strictly more expressive logic [18].

The model checking algorithm for ATL is very similar to the one for CTL (see Definition 2.12). The main difference is that the predecessor function must take into account the quantified group of agents. For example, if we have already calculated the set of global states $\|\varphi\|_{\mathcal{I}}$ of an interpreted system \mathcal{I} where the ATL formula φ holds, the set of global states $\|\langle\langle A \rangle\rangle \text{F} \varphi\|_{\mathcal{I}}$ where $\langle\langle A \rangle\rangle \text{F} \varphi$ holds ($A \subseteq \text{Agt}$ is a set of agents) is equal to $\text{lfp}_X [\|\varphi\|_{\mathcal{I}} \cup \text{pre}_A(X)]$ where:

$$\text{pre}_A(X) \triangleq \left\{ g \in G \mid \exists a \in (\text{Act}_i)_{i \in A} \forall a' \in (\text{Act}_i)_{i \in \text{Agt} \setminus A} . t(g, (a, a')) \in X \right\}$$

Please refer to [65] for a complete model checking algorithm for ATL. Model checking ATL* is much more difficult. Although a theoretical model checking algorithm for ATL* is provided in [18] (in the form of a constructive proof of ATL* model checking complexity), we are not aware of any existing tool that would support it. If this is true, then we have created the first model checker for ATL* as part of this project (see Section 1.3).

We shall now demonstrate the expressive power of ATL. The sample properties listed at the beginning of this section can be expressed in ATL as follows:

1. $\text{open} \wedge \langle\langle \emptyset \rangle\rangle \text{X} \text{open} \wedge \langle\langle \emptyset \rangle\rangle \text{X} \langle\langle \emptyset \rangle\rangle \text{X} \text{open}$.
2. $\text{connected} \rightarrow \langle\langle \text{Agt} \rangle\rangle \text{F} \neg \text{connected}$.
3. $\langle\langle \emptyset \rangle\rangle \text{F} \text{finished}$.
4. $\langle\langle \text{Agt} \rangle\rangle \text{F} \langle\langle \emptyset \rangle\rangle \text{G} \text{deleted}$.
5. This property cannot be expressed in ATL because it does not refer to paths. It can however be expressed in ATL* as $\langle\langle \emptyset \rangle\rangle (\text{G F red} \rightarrow \text{G F green})$.
6. This property cannot be expressed in ATL because it does not refer to paths. It can however be expressed in ATL* as $[\langle\langle \text{Agt} \rangle\rangle \text{G F left}] \rightarrow [\langle\langle \text{Agt} \rangle\rangle \text{G F right}]$.
7. $\langle\langle \{\text{firewall, antivirus}\} \rangle\rangle \text{G} \neg \text{hacked}$.
8. This property cannot be expressed in ATL/ATL* because it does not refer to strategies directly.

ATL and ATL* allow us to express almost all properties on our list. The only exception is property 8, which cannot be expressed in either ATL or ATL* because it refers to agents' strategies directly, while both logics treat strategies only implicitly (via the $\langle\langle A \rangle\rangle$ quantifier). This limitation of ATL and ATL* is addressed by Strategy Logic.

2.2.5 Strategy Logic

Strategy Logic (SL) is a new formalism introduced in [76], which strictly subsumes ATL*. Instead of quantifying over agents in formulas, SL quantifies over strategies explicitly, allowing us to express properties like “agents a and b share the same strategy” or even game-theoretic concepts like Nash equilibria. To do this, it augments the syntax of LTL (see Definition 2.7) with three operators: universal strategy quantifier $\llbracket x \rrbracket \varphi$, existential strategy quantifier $\langle\langle x \rangle\rangle \varphi$, and agent binding $(a, x)\varphi$.

Definition 2.19 (SL Syntax). SL formulas are built inductively from the set of atomic propositions AP , strategy variables Var , and agents Agt , by using the following grammar, where $p \in AP$, $x \in Var$, and $i \in Agt$:

$$\varphi ::= p \mid \top \mid \neg \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{X} \varphi \mid \mathbf{F} \varphi \mid \mathbf{G} \varphi \mid \varphi \mathbf{U} \varphi \mid \langle\langle x \rangle\rangle \varphi \mid \llbracket x \rrbracket \varphi \mid (i, x)\varphi$$

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

We define $\text{vars}(\varphi) \subseteq Var$ to be the set of strategy variables quantified in an SL formula φ (e.g. $\text{vars}(\langle\langle x \rangle\rangle \llbracket y \rrbracket (a, x)(b, y)\mathbf{X}\langle\langle z \rangle\rangle (b, z)\mathbf{F}p) = \{x, y, z\}$). Similarly to first-order languages, we define the set of *free agents/variables* of a formula which contains (i) all agents for which there is no binding after the occurrence of a temporal operator and (ii) all variables for which there is a binding but no quantifications.

Definition 2.20 (Free Agents/Variables). The set of *free agents/variables* of an SL formula is given by the function $\text{free} : SL \rightarrow 2^{Agt \cup Var}$ defined as follows:

1. $\text{free}(\top) \triangleq \emptyset$;
2. $\text{free}(p) \triangleq \emptyset$, where $p \in AP$;
3. $\text{free}(\neg \varphi) \triangleq \text{free}(\varphi)$;
4. $\text{free}(\varphi_1 \wedge \varphi_2) \triangleq \text{free}(\varphi_1) \cup \text{free}(\varphi_2)$;
5. $\text{free}(\text{Op} \varphi) \triangleq Agt \cup \text{free}(\varphi)$, where $\text{Op} \in \{\mathbf{X}, \mathbf{F}, \mathbf{G}\}$;
6. $\text{free}(\varphi_1 \mathbf{U} \varphi_2) \triangleq Agt \cup \text{free}(\varphi_1) \cup \text{free}(\varphi_2)$;
7. $\text{free}(\mathbf{Qn} \varphi) \triangleq \text{free}(\varphi) \setminus \{x\}$, where $\mathbf{Qn} \in \{\llbracket x \rrbracket, \langle\langle x \rangle\rangle : x \in Var\}$;
8. $\text{free}((i, x)\varphi) \triangleq \text{free}(\varphi)$, if $i \notin \text{free}(\varphi)$, where $i \in Agt$ and $x \in Var$;
9. $\text{free}((i, x)\varphi) \triangleq (\text{free}(\varphi) \setminus \{i\}) \cup \{x\}$, if $i \in \text{free}(\varphi)$, where $i \in Agt$ and $x \in Var$.

A formula $\varphi \in SL$ is *agent-closed* (resp. *variable-closed*) iff $\text{free}(\varphi) \cap Agt = \emptyset$ (resp. $\text{free}(\varphi) \cap Var = \emptyset$). A formula $\varphi \in SL$ is a *sentence* iff it is both agent-closed and variable-closed.

Consider an interpreted system with agents $Agt = \{a, b, c\}$ and the SL formula $\varphi = \langle\langle x \rangle\rangle (a, x)(b, y)\mathbf{F}p$ [72]. We have $\text{free}(\varphi) = \{c, y\}$ because agent c is not bound to any variable after $\mathbf{F}p$ and variable y is not quantified. We also have $\text{free}((c, z)\varphi) = \{y, z\}$ and $\text{free}((a, z)\varphi) = \text{free}(\varphi)$ because $c \in \text{free}(\varphi)$ and $a \notin \text{free}(\varphi)$. Hence, $(c, z)\varphi$ is agent-closed while $(a, z)\varphi$ is not.

Before defining the basic concepts necessary for SL semantics, we briefly explain an important attribute of SL formulas called the *alternation number*, which we need in order to be able to describe SL model checking complexity (see Table 2.1). It refers to the maximum number of quantifier switches $\langle\langle \cdot \rangle\rangle \llbracket \cdot \rrbracket$, $\llbracket \cdot \rrbracket \langle\langle \cdot \rangle\rangle$, $\langle\langle \cdot \rangle\rangle \neg \langle\langle \cdot \rangle\rangle$ or $\llbracket \cdot \rrbracket \neg \llbracket \cdot \rrbracket$ that bind a variable in a subformula that is not a sentence. To determine the alternation number of an SL formula, we first replace all its subsentences with atoms and then count the number of quantifier switches. For example, consider an interpreted system with agents $Agt = \{a, b\}$ and the SL sentence $\varphi \triangleq \llbracket x \rrbracket \langle\langle y \rangle\rangle (a, x)(b, y)\mathbf{F} \varphi'$ with $\varphi' \triangleq \llbracket x \rrbracket \langle\langle y \rangle\rangle (a, x)(b, y)\mathbf{X}p$ [72]. The alternation number of φ is 1 because φ' is a sentence and there is one quantifier switch in $\llbracket x \rrbracket \langle\langle y \rangle\rangle (a, x)(b, y)\mathbf{F} \varphi'$ (the subsentence φ' is replaced with an atom $p_{\varphi'}$). However, if we replaced φ' with $\varphi'' \triangleq \llbracket x \rrbracket (a, x)\mathbf{X}p$ in φ , the alternation number of φ would be 2 because φ'' is not a sentence. Please refer to [72] for more details about this concept.

We shall now define several basic concepts relevant to SL semantics. Since we allow strategies to be *shared* among agents, we first need to slightly modify strategies (see Definition 2.16):

Definition 2.21 (Shared Strategies). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system and $A \subseteq Agt$ a non-empty set of agents. Then a partial function $f_A : Trk \rightarrow Act_A$, such that for all agents $i \in A$ and tracks $\tau \in \text{dom}(f_A)$, $f_A(\tau) \in P_i(l_{iE}(\text{last}(\tau)))$, is a *shared strategy* for agents A which maps tracks to actions. $SStr_A \subseteq Trk \rightarrow Act_A$ denotes the set of all shared strategies for agents A . The set of all shared strategies is defined as $SStr \triangleq \bigcup_{A \subseteq Agt}^{A \neq \emptyset} SStr_A$.

Observe that an individual strategy $f_i \in Str_i$ for an agent $i \in Agt$ (Definition 2.16) is also a shared strategy (Definition 2.21), i.e. $f_i \in SStr_{\{i\}}$. So shared strategies are a generalisation of individual strategies, i.e. $Str \subseteq SStr$. We shall see that many SL concepts are generalisations of those of ATL.

In order to determine the set of strategies over which a variable can quantify, we need to know which agents are bound to it. We thus introduce the set of *sharing agents*. Intuitively, it refers to the set of agents that share a variable within a formula.

Definition 2.22 (Sharing Agents). Let \mathcal{I} be an interpreted system. The function $\text{sharing} : SL \times Var \rightarrow 2^{Agt}$ returns the set of agents *sharing* a variable in an SL formula. It is defined inductively as follows:

1. $\text{sharing}(\top, x) \triangleq \emptyset$;
2. $\text{sharing}(p, x) \triangleq \emptyset$, where $p \in AP$;
3. $\text{sharing}(\neg\varphi, x) \triangleq \text{sharing}(\varphi, x)$;
4. $\text{sharing}(\varphi_1 \text{ Op } \varphi_2, x) \triangleq \text{sharing}(\varphi_1, x) \cup \text{sharing}(\varphi_2, x)$, where $\text{Op} \in \{\wedge, \vee\}$;
5. $\text{sharing}(\text{Op } \varphi, x) \triangleq \text{sharing}(\varphi, x)$, where $\text{Op} \in \{X, F, G\}$;
6. $\text{sharing}(\varphi_1 \cup \varphi_2, x) \triangleq \text{sharing}(\varphi_1, x) \cup \text{sharing}(\varphi_2, x)$;
7. $\text{sharing}(\text{Qn } \varphi, x) \triangleq \text{sharing}(\varphi, x)$, where $\text{Qn} \in \{\llbracket y \rrbracket, \langle\langle y \rangle\rangle : y \in Var\}$;
8. $\text{sharing}((i, x)\varphi, x) \triangleq \{i\} \cup \text{sharing}(\varphi, x)$;
9. $\text{sharing}((i, y)\varphi, x) \triangleq \text{sharing}(\varphi, x)$, where $y \in Var \setminus \{x\}$;

Note that for conciseness we assume throughout this report that every variable is quantified at most once in a given formula. This can be easily ensured by renaming variables which are not free in the formula (e.g. $(a, x)\langle\langle x \rangle\rangle(b, x)Xp \equiv (a, x)\langle\langle y \rangle\rangle(a, y)Xp \not\equiv (a, y)\langle\langle x \rangle\rangle(a, x)Xp$).

When we define the modelling relation for SL, we will somehow need to keep track of the strategies assigned to variables and agents. This information is encapsulated in an *assignment*.

Definition 2.23 (Assignments). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system. An *assignment* is a partial function $\chi : Var \cup Agt \rightarrow SStr$ which maps variables and agents in its domain to shared strategies. The set $Asg \subseteq Var \cup Agt \rightarrow SStr$ contains all possible assignments.

An assignment $\chi \in Asg$ is *complete* iff it is defined on all agents, i.e. $Agt \subseteq \text{dom}(\chi)$. The set $CAsg \subseteq Asg$ contains all complete assignments.

Observe that an agent assignment $\alpha \in AAsg_A$ for a set of agents $A \subseteq Agt$ (Definition 2.16) is also an assignment (Definition 2.23), i.e. $\alpha \subseteq Asg$. Hence, SL assignments are again a generalisation of ATL agent assignments, i.e. $AAsg \subseteq Asg$. In addition, we also have $AAsg_{Agt} \subseteq CAsg$ because $\text{dom}(\alpha) = Agt$ for all $\alpha \in AAsg_{Agt}$.

Assume that the history of the system has been $\rho \in Trk$. Given some strategy $f \in SStr$, we would like to determine the equivalent strategy $(f)_\rho \in SStr$ for the new history after ρ . Informally, we want to cut off the past. This is achieved by *translating* the strategy.

Definition 2.24 (Strategy Translation). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system, $f \in SStr_A$ a shared strategy for a set of agents $A \subseteq Agt$, and $\rho \in \text{dom}(f)$ a track in its domain. Then $(f)_\rho \in SStr_A$ denotes the *translation* of f along ρ with $\text{dom}((f)_\rho) \triangleq \{\rho' \in \text{track}(\text{last}(\rho)) \mid \rho \cdot \rho'_{\geq 1} \in \text{dom}(f)\}$ such that $(f)_\rho(\rho') \triangleq f(\rho \cdot \rho'_{\geq 1})$ for all $\rho' \in \text{dom}((f)_\rho)$.

For example, let $f \in SStr_A$ be a shared strategy for agents $A \subseteq Agt$ and $\rho = [g_0, g_1, g_2]$ and $\rho' = [g_2, g_3, g_4]$ two tracks such that $\rho \cdot \rho'_{\geq 1} = [g_0, g_1, g_2, g_3, g_4] \in \text{dom}(f)$. Then we have $f([g_0, g_1, g_2, g_3, g_4]) = (f)_\rho([g_2, g_3, g_4])$. Effectively, we have cut off the past history $[g_0, g_1]$ and treat g_2 as the starting state. This notion can be easily extended to assignment translation.

Definition 2.25 (Assignment Translation). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system, $\chi \in Asg$ an assignment, and $\rho \in Trk$ a track. Then $(\chi)_\rho \in Asg$ denotes the *translation* of χ along ρ such that $\text{dom}((\chi)_\rho) \triangleq \text{dom}(\chi)$ and $(\chi)_\rho(l) \triangleq (\chi(l))_\rho$ for all $l \in \text{dom}(\chi)$.

We can now define the notion of a *play*. Intuitively, a play is the unique evolution of an interpreted system determined by the current global state and agents' strategies (encapsulated in a complete assignment).

Definition 2.26 (Plays). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system. A path $\pi \in \text{path}(g)$ starting at a global state $g \in G$ is a *play* w.r.t. a complete assignment $\chi \in CAsg$ ((χ, g) -play, for short) if, for all $i \geq 0$, it holds that $\pi(i+1) = t(\pi(i), \langle \chi(a)(\pi_{\leq i}) : a \in Agt \rangle)$. The function $\text{play} : CAsg \times G \rightarrow Pth$ returns the (χ, g) -play $\text{play}(\chi, g) \in Pth(g)$, for all pairs (χ, g) in its domain.

Once again, SL plays are very similar to ATL outcomes (see Definition 2.17). The difference between the two concepts is that $\text{play}(\chi, g)$ is based on a complete assignment and therefore is unique (since the complete assignment determines the next action of every agent). In contrast, $\text{out}(g, \alpha)$ is a set of paths. The two concepts are essentially equivalent if we require that $\alpha \in AAsg_{Agt}$ since $AAsg_{Agt} \subseteq CAsg$.

Informally, given a complete assignment, which contains the strategies of all agents, and the current global state, we know exactly how the whole system will evolve. This idea is formalised in the following definition.

Definition 2.27 (Global Translation). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system, $\chi \in CAsg$ a complete assignment, and $g \in G$ a global state. Then the *i -th global translation* of (χ, g) , with $i \geq 0$ is defined as $(\chi, g)^i \triangleq ((\chi)_{\pi_{\leq i}}, \pi(i))$, where $\pi \triangleq \text{play}(\chi, g)$.

We need to define one more concept before we give the SL semantics. When an agent binding (i, x) is encountered, we must update the assignment so that the strategy assigned to agent $i \in Agt$ is the same as the one for the variable $x \in Var$. We generalise this concept of *redefinition* to an arbitrary partial function.

Definition 2.28 (Redefinition). Let A, B be two arbitrary sets, $f : A \rightarrow B$ a partial function, and $a \in A, b \in B$ elements of the sets. Then, $f[a \mapsto b] : A \rightarrow B$ denotes a new partial function defined on $\text{dom}(f[a \mapsto b]) \triangleq \text{dom}(f) \cup \{a\}$ such that $f[a \mapsto b](a) \triangleq b$ and $f[a \mapsto b](a') \triangleq f(a')$ for all $a' \in \text{dom}(f) \setminus \{a\}$.

We are now finally ready to define SL semantics. Intuitively, a formula $\langle\langle x \rangle\rangle \varphi$ is true iff there exists a strategy for x such that φ is true. Conversely, a formula $\llbracket x \rrbracket \varphi$ is true iff for all strategies assigned to x we have φ . $(a, x)\varphi$ simply binds the strategy assigned to variable x to agent a .

Definition 2.29 (SL Semantics). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system, $g \in G$ a global state, $\varphi \in SL$ an SL formula, and $\chi \in Asg$ an assignment with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$. The modelling relation $\mathcal{I}, \chi, g \models_{SL} \varphi$ is inductively defined as follows:

1. $\mathcal{I}, \chi, g \models_{SL} \top$.
2. $\mathcal{I}, \chi, g \models_{SL} p$ iff $g \in h(p)$, with $p \in AP$.
3. For all formulas $\varphi, \varphi_1, \varphi_2 \in SL$, it holds that:
 - (a) $\mathcal{I}, \chi, g \models_{SL} \neg \varphi$ iff $\mathcal{I}, \chi, g \not\models_{SL} \varphi$;
 - (b) $\mathcal{I}, \chi, g \models_{SL} \varphi_1 \wedge \varphi_2$ iff $\mathcal{I}, \chi, g \models_{SL} \varphi_1$ and $\mathcal{I}, \chi, g \models_{SL} \varphi_2$;
 - (c) $\mathcal{I}, \chi, g \models_{SL} \varphi_1 \vee \varphi_2$ iff $\mathcal{I}, \chi, g \models_{SL} \varphi_1$ or $\mathcal{I}, \chi, g \models_{SL} \varphi_2$.

4. For all variables $x \in Var$ and formulas $\varphi \in SL$, it holds that:
 - (a) $\mathcal{I}, \chi, g \models_{SL} \langle\langle x \rangle\rangle \varphi$ iff there exists a strategy $f \in SStr_{\text{sharing}(\varphi, x)}$ such that $\mathcal{I}, \chi[x \mapsto f], g \models_{SL} \varphi$;
 - (b) $\mathcal{I}, \chi, g \models_{SL} \llbracket x \rrbracket \varphi$ iff for all strategies $f \in SStr_{\text{sharing}(\varphi, x)}$ it holds that $\mathcal{I}, \chi[x \mapsto f], g \models_{SL} \varphi$.
5. For all agents $i \in Agt$, variables $x \in Var$, and formulas $\varphi \in SL$, it holds that $\mathcal{I}, \chi, g \models_{SL} (i, x)\varphi$ iff $\mathcal{I}, \chi[i \mapsto \chi(x)], g \models_{SL} \varphi$.
6. For all formulas $\varphi, \varphi_1, \varphi_2 \in SL$, it holds that:
 - (a) $\mathcal{I}, \chi, g \models_{SL} X\varphi$ iff $\mathcal{I}, (\chi, g)^1 \models_{SL} \varphi$;
 - (b) $\mathcal{I}, \chi, g \models_{SL} F\varphi$ iff there is an index $i \geq 0$ such that $\mathcal{I}, (\chi, g)^i \models_{SL} \varphi$;
 - (c) $\mathcal{I}, \chi, g \models_{SL} G\varphi$ iff for all indices $i \geq 0$ it holds that $\mathcal{I}, (\chi, g)^i \models_{SL} \varphi$;
 - (d) $\mathcal{I}, \chi, g \models_{SL} \varphi_1 \cup \varphi_2$ iff there is an index $i \geq 0$ such that $\mathcal{I}, (\chi, g)^i \models_{SL} \varphi_2$ and, for all indices $0 \leq j < i$ it holds that $\mathcal{I}, (\chi, g)^j \models_{SL} \varphi_1$.

To our best knowledge, there are no existing model checkers for SL. The aim of this project is to develop the first such tool (see Section 1.1). The *model checking problem* for SL is formally defined as:

Given an interpreted system \mathcal{I} , a global state $g \in G$, an SL formula $\varphi \in SL$, and an assignment $\chi \in Asg$ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$, determine whether $\mathcal{I}, \chi, g \models_{SL} \varphi$.

In addition to model checking, we are also interested in the problem of strategy synthesis. Informally, we want to construct strategies which make a particular SL formula true. More formally, the *strategy synthesis problem* is defined as:

Given an interpreted system \mathcal{I} , a global state $g \in G$, and an SL formula $\varphi \in SL$, find an assignment $\chi \in Asg$ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$ such that $\mathcal{I}, \chi, g \models_{SL} \varphi$.

Let us now investigate the expressive power of SL. The sample properties listed at the beginning of this section can be expressed in SL as follows:

1. $\wp_{\forall} (\text{open} \wedge X \text{open} \wedge XX \text{open})$.
2. $\text{connected} \rightarrow \wp_{\exists} F \neg \text{connected}$.
3. $\wp_{\forall} F \text{finished}$.
4. $\wp_{\exists} F \wp_{\forall} G \text{deleted}$.
5. $\wp_{\forall} (GF \text{red} \rightarrow GF \text{green})$.
6. $[\wp_{\exists} GF \text{left}] \rightarrow [\wp_{\exists} GF \text{right}]$.
7. $\langle\langle x_{\text{firewall}} \rangle\rangle \langle\langle x_{\text{antivirus}} \rangle\rangle (\llbracket x_i \rrbracket)_{i \in Agt \setminus \{\text{firewall}, \text{antivirus}\}} ((i, x_i))_{i \in Agt} G \neg \text{hacked}$.
8. $\llbracket x \rrbracket (\text{player}_1, x) (\text{player}_2, x) (\llbracket x_i \rrbracket (i, x_i))_{i \in Agt \setminus \{\text{player}_1, \text{player}_2\}} G (\neg \text{win}_1 \wedge \neg \text{win}_2)$.

where $\wp_{\exists} \triangleq (\langle\langle x_i \rangle\rangle (i, x_i))_{i \in Agt}$ and $\wp_{\forall} \triangleq (\llbracket x_i \rrbracket (i, x_i))_{i \in Agt}$ are two prefixes. This demonstrates that SL is more expressive than any of the previous formalisms including ATL* because it quantifies and binds strategies explicitly. However, there are other types of specifications which SL does not support. These include statements about agents' knowledge, which can be expressed using epistemic modalities.

2.2.6 Epistemic modalities

All formalisms defined so far deal only with time. However, there are other modalities (e.g. knowledge, necessity, belief, etc.) which can be expressed by extending one of the logics described in the previous subsections. In this subsection we discuss epistemic modalities expressing agents' *knowledge*.

Modelling agents' knowledge in multi-agent systems is very useful because it allows us to reason about the decision making process of individual agents as well as the interactions within a group of agents [42]. Syntactically, the knowledge of (groups of) agents is represented using epistemic connectives with the following informal meanings:

- $K_i \varphi$ – agent $i \in \text{Agt}$ knows that φ is true (individual knowledge);
- $E_A \varphi$ – each agent in $A \subseteq \text{Agt}$ knows that φ is true (group knowledge);
- $D_A \varphi$ – all agents $A \subseteq \text{Agt}$ (together) know that φ is true (distributed knowledge);
- $C_A \varphi$ – it is common knowledge⁹ among agents $A \subseteq \text{Agt}$ that φ is true (common knowledge).

Any of the logics described in the previous subsections can be augmented with the epistemic connectives. For instance, if we extend ATL (see Subsection 2.2.4) with the new operators, we obtain *Alternating-Time Temporal Logic with Knowledge* (ATLK).

Definition 2.30 (ATLK Syntax). ATLK *formulas* are built inductively from the set of propositional atoms AP and agents Agt , by using the following grammar, where $p \in AP$, $i \in \text{Agt}$, and $A \subseteq \text{Agt}$:

$$\varphi ::= p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle\langle A \rangle\rangle X\varphi \mid \langle\langle A \rangle\rangle F\varphi \mid \langle\langle A \rangle\rangle G\varphi \mid \langle\langle A \rangle\rangle[\varphi U \varphi] \mid K_i \varphi \mid E_A \varphi \mid D_A \varphi \mid C_A \varphi$$

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

The semantics of agents' knowledge is traditionally based on *epistemic accessibility* [61]. Intuitively, two states are *epistemically accessible* if they are indistinguishable by an agent. This is formally represented using a binary epistemic accessibility relation on states of the system. In interpreted systems (see Definition 2.5), the individual epistemic accessibility relation $\sim_i \subseteq G \times G$ of an agent $i \in \text{Agt}$ is naturally induced by its local states: two global states $g_1, g_2 \in G$ are indistinguishable by the agent iff their local states are the same, i.e. $l_{iE}(g_1) = l_{iE}(g_2)$. A formal definition follows.

Definition 2.31 (Epistemic Accessibilities). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system. Let $i \in \text{Agt}$ be an agent and $A \subseteq \text{Agt}$ a set of agents. Then the *epistemic accessibility relations* on G are defined as:

- The *individual* epistemic accessibility relation \sim_i of the agent i is defined by $g \sim_i g'$ if and only if $l_{iE}(g) = l_{iE}(g')$, i.e. the local states of agent i in global states g and g' are the same.
- The *group* epistemic accessibility relation \sim_A^E of the set of agents A is defined as $\sim_A^E \triangleq \bigcup_{j \in A} \sim_j$, i.e. the local states of at least one agent in A in global states g and g' are the same.
- The *distributed* epistemic accessibility relation \sim_A^D of the set of agents A is defined as $\sim_A^D \triangleq \bigcap_{j \in A} \sim_j$, i.e. the local states of all agents in A in global states g and g' are the same.
- The *common* epistemic accessibility relation \sim_A^C of the set of agents A is defined as $\sim_A^C \triangleq (\sim_A^E)^+$, i.e. transitive closure of group accessibility relation.

We then ascribe knowledge to agents in the following sense [42]: We say that an agent *knows* a fact iff it is true at all the worlds he considers possible, i.e. in all states that he cannot distinguish from the current one. This interpretation of knowledge is reflected in the semantics of the epistemic connectives K_i , E_A , D_A , and C_A .

Definition 2.32 (ATLK Semantics). Let $\mathcal{I} = \langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \rangle$ be an interpreted system, $g \in G$ a global state, $i \in \text{Agt}$ an agent, and $A \subseteq \text{Agt}$ a set of agents. We define $\mathcal{I}, g \models_{\text{ATLK}} \varphi$ by induction on $\varphi \in \text{ATLK}$:

- $\mathcal{I}, g \models_{\text{ATLK}} K_i \varphi$ iff for all states $g' \in G$ such that $g \sim_i g'$ we have $\mathcal{I}, g' \models_{\text{ATLK}} \varphi$;
- $\mathcal{I}, g \models_{\text{ATLK}} E_A \varphi$ iff for all states $g' \in G$ such that $g \sim_A^E g'$ we have $\mathcal{I}, g' \models_{\text{ATLK}} \varphi$;
- $\mathcal{I}, g \models_{\text{ATLK}} D_A \varphi$ iff for all states $g' \in G$ such that $g \sim_A^D g'$ we have $\mathcal{I}, g' \models_{\text{ATLK}} \varphi$;
- $\mathcal{I}, g \models_{\text{ATLK}} C_A \varphi$ iff for all states $g' \in G$ such that $g \sim_A^C g'$ we have $\mathcal{I}, g' \models_{\text{ATLK}} \varphi$;
- $\mathcal{I}, g \models_{\text{ATLK}} \varphi$ is defined in the same way as $\mathcal{I}, g \models_{\text{ATL}} \varphi$ for other ATLK formulas.

⁹ $C_A \varphi = \bigwedge_{i=1}^{\infty} E_A^i \varphi = E_A \varphi \wedge E_A E_A \varphi \wedge \dots$ i.e. each agent in A knows that φ is true, each agent in A knows that each agent in A knows that φ is true, ...

Logic	Complexity	Complexity w.r.t. $ \varphi $	Complexity w.r.t. $ \mathcal{I} $
LTL	$2^{O(\varphi)}O(\mathcal{I})$	PSPACE-COMPLETE	NLOGSPACE-COMPLETE
CTL	$O(\varphi \times \mathcal{I})$	LOGSPACE	NLOGSPACE-COMPLETE
CTL*	$2^{O(\varphi)}O(\mathcal{I})$	PSPACE-COMPLETE	NLOGSPACE-COMPLETE
ATL	$O(\varphi \times \mathcal{I})$	P-COMPLETE	P-COMPLETE
ATL*	$ \mathcal{I} ^{2^{O(\varphi)}}$	2EXPTIME-COMPLETE	P-COMPLETE
SL	—	NONELEMENTARY	P-COMPLETE

Table 2.1: Model checking complexities of various logics [18, 38, 58, 72, 84]. $|\varphi|$ and $|\mathcal{I}|$ denote the size of the formula and the model respectively. The precise lower and upper bounds on SL complexity with respect to the size of the formula are k -EXPSpace-HARD and $(k + 1)$ -EXPTIME respectively where k is the alternation number of φ [72].

2.2.7 Model Checking Complexity


In this subsection we discuss the computational complexity of the model checking problem. Table 2.1 shows model checking complexities for the logics described in the previous subsections. We can make the following observations:

- There is a clear correspondence between expressiveness of a logic and its model checking complexity (as expected).
- The complexity of model checking of all logics discussed in this section is P with respect to the size of the model \mathcal{I} .
- The complexity of model checking LTL and CTL* is exactly the same. Nevertheless, there are not many practical tools for CTL* verification. This is probably due to the fact that (i) LTL model checking is faster than CTL* model checking in practice despite the theoretical complexity results, and (ii) branching-time logics are less natural than linear-time ones [84].
- While both CTL and ATL formulas can be checked in time $O(|\varphi| \times |\mathcal{I}|)$, the theoretical complexity of CTL is lower because it is related to graph reachability, whereas ATL model checking is related to AND-OR graph reachability [18].

It is important to emphasise that the size of the model $|\mathcal{I}|$ refers to the number of states (not variables). In general, adding one Boolean variable to a system doubles the size of the model. This is sometimes referred to as the *state explosion problem*. The size of a model with m Boolean variables is $O(2^m)$. More generally, the size of a model with n variables where the i -th variable has v_i possible values is $O(\prod_{i=1}^n v_i)$. It is often the case that they are internally represented using $m = \sum_{i=1}^n \lceil \log_2 v_i \rceil$ Boolean variables (e.g. when using BDDs, see Subsections 2.3.1 and 2.3.2). Since models are usually described implicitly in the form of a program (based on variables), CTL complexity is sometimes stated as PSPACE-COMPLETE with respect to program size [64].

2.2.8 Summary


In this section, we have described several well-established temporal logics as well as the relatively new SL formalism, which is central to our project. We gave the syntax and semantics of all logics. We also briefly outlined how they can be model checked and compared their relative expressiveness on a fixed set of properties. The following diagram summarises the relationships between the logics ($L_1 \rightarrow L_2$ means that L_2 is strictly more expressive than L_1):



We also discussed epistemic modalities expressing agents' knowledge and provided the model checking complexities of all logics.

Verification technique	Number of states
Explicit approaches	10^6
Binary decision diagrams	10^{25}
Bounded model checking	10^{100}
Symmetry reduction	10^{150}
Predicate abstraction	10^{1000}

Table 2.2: Approximate state spaces that can be handled by various techniques [62].

Figure 2.1: Truth table, a binary decision tree (BDT), and a binary decision diagram (BDD) for the propositional formula $\neg(x \wedge y)$.

2.3 Verification Methods

Once we have converted a system S and a property P to a model \mathcal{M}_S and a formula φ_P respectively (as described in the previous sections), we can verify whether the system indeed satisfies the property by checking [62]:

$$\mathcal{M}_S, s_0 \models \varphi_P$$

where s_0 is the initial state of the system. As we have already mentioned, it is often more convenient to calculate the set of states $\|\varphi_P\|_{\mathcal{M}_S}$ at which φ_P is true first:

$$\|\varphi_P\|_{\mathcal{M}_S} \triangleq \{s \mid \mathcal{M}_S, s \models \varphi_P\}$$

Verifying that the system satisfies the property is then equivalent to checking $s_0 \in \|\varphi_P\|_{\mathcal{M}_S}$.

As an example demonstrating this approach, the recursive model-checking algorithm SAT_{CTL} for CTL was presented in the previous section (see Definition 2.12). We can use it directly to verify an arbitrary property of an arbitrary system (so-called explicit model checking). However, models are rarely given explicitly, but instead, in the form of a program. Moreover, we would like to avoid direct calculation of the set operations, which could be very slow. Therefore, other more efficient techniques have been proposed and implemented. Table 2.2 shows the maximum state spaces that can be handled by various techniques.

2.3.1 Binary Decision Diagrams

Binary decision diagrams [49] are a technique for representing propositional formulas (similarly to truth tables). Figure 2.1 shows a truth table (Figure 2.1a) and a binary decision tree (Figure 2.1b) for the propositional formula $\neg(x \wedge y)$.

Definition 2.33 (Binary Decision Trees). A *binary decision tree* (BDT) is a complete binary tree whose non-terminal nodes are labelled with Boolean variables (x_1, x_2, \dots, x_n) and whose terminal nodes are labelled with Boolean values 0 and 1. Each non-terminal node of a BDT has two outgoing edges – one solid line and one dotted line (representing that the corresponding Boolean variable is 1 and 0 respectively). A BDT has one layer of non-terminal nodes for every Boolean variable.

A binary decision tree f represents a Boolean function $f(x_1, x_2, \dots, x_n)$. Given an assignment to all variables x_1, x_2, \dots, x_n occurring in f , we determine the value of the function as follows. We start at the root of f . If the current node is non-terminal with label x_i , we follow the solid line when $x_i = 1$ and the dotted line when $x_i = 0$. If the current node is terminal, the value of the function is the label of the terminal node (0 or 1).

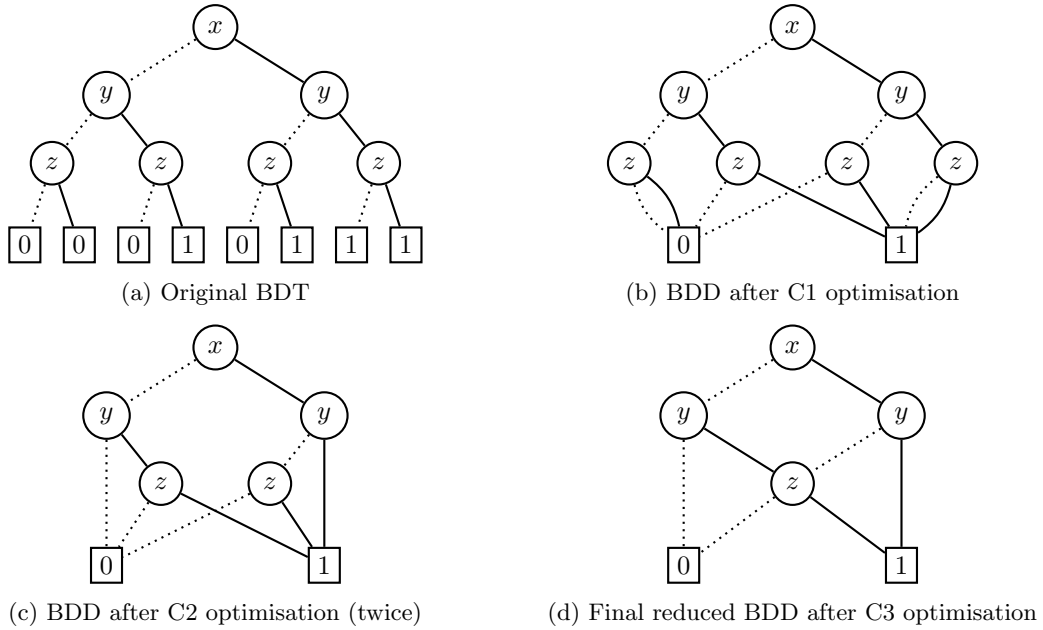


Figure 2.2: All reduction steps of the BDT for the formula $(x \wedge y) \vee (y \wedge z) \vee (z \wedge x)$ using optimisations C1–C3.

A BDT for a formula $\varphi(x_1, x_2, \dots, x_n)$ with n variables has $2^n - 1$ non-terminal nodes and 2^n terminal nodes. Binary decision trees are thus no more efficient than truth tables (the truth table for $\varphi(x_1, x_2, \dots, x_n)$ has 2^n entries). Fortunately, BDTs can often be compressed. Consider the BDT in Figure 2.1b. The left y node is redundant since both its edges point to terminal nodes labelled with 1 (although distinct). The dotted edge from the node labelled with x can thus point to a terminal node labelled with 1 directly instead. Intuitively, once we know that x is false, the formula $\neg(x \wedge y)$ is true regardless of the value of y . The resulting graph-based structure (Figure 2.1c) is a binary decision diagram [49, p. 364]. Binary decision diagrams (BDDs) are a generalisation of binary decision trees.

There are three different ways for reducing a BDD [49, p. 363]:

- C1. Removal of duplicate terminals.** We merge all terminal 0-nodes into one 0-node. Similarly, we merge all terminal 1-nodes into one 1-node.
- C2. Removal of redundant tests.** If both outgoing edges of a non-terminal node n point to the same non-terminal node m , we can remove node n and send all its incoming edges to node m .
- C3. Removal of duplicate non-terminals.** If two distinct non-terminal nodes m, n are roots of two structurally identical sub-BDDs, we can merge m and n .

The optimisations C1–C3 are demonstrated in Figure 2.2.

Definition 2.34 (Binary Decision Diagram). A *binary decision diagram* (BDD) is a finite directed acyclic graph with a unique initial node whose non-terminal nodes are labelled with Boolean variables (x_1, x_2, \dots, x_n) and whose terminal nodes are labelled with Boolean values 0 and 1. Each non-terminal node of a BDD has two outgoing edges – one solid line and one dotted line (representing that the corresponding Boolean variable is 1 and 0 respectively).

A BDD is *reduced* iff none of the optimisations C1–C3 can be applied to it.

BDDs are often much more space efficient than BDTs [49, p. 366]. However, they still have several drawbacks, (i) there may be multiple occurrences of the same variable along a path (which is inefficient), and (ii) there is no simple way of comparing two BDDs for equivalence. Figure 2.3 shows 6 different reduced BDDs representing the same propositional formula (demonstrating the second problem). We can solve this problem by imposing an ordering on the variables [49, p. 367].




Figure 2.3: All possible ROBDD orderings for the propositional formula $x \vee (y \wedge z)$.

Definition 2.35 (Ordered Binary Decision Diagram). Let $[x_1, x_2, \dots, x_n]$ be an ordered list of variables without duplications and f be a BDD all of whose variables occur in the list. The BDD f has the *ordering* $[x_1, x_2, \dots, x_n]$ iff for every occurrence of x_i followed by x_j along any path in f , we have $i < j$.

A BDD is an *ordered binary decision diagram* (OBDD) iff it has some ordering. An OBDD is a *reduced ordered binary decision diagram* (ROBDD) iff it is reduced.

Figure 2.3 shows ROBDD with different orderings for the same propositional formula. In general, the ordering of the chosen variable ordering has a significant impact on the size and shape of a ROBDD [49]. Note that the BDT in Figure 2.1b is ordered although it is not reduced. Although two structurally different reduced BDDs can be semantically equivalent, once we fix the ordering of the variables, the ROBDD for a given Boolean expression is unique. All OBDDs which are semantically equivalent thus have a *canonical form* – the reduced form, which can be obtained by applying optimisations C1–C3 [62].

To sum up, ROBDDs have the following advantages (compared to truth tables) [49]:

1. They provide a *compact representation* of complex Boolean formulas. Nevertheless, the size of the ROBDD representing certain Boolean functions (e.g. integer multiplication [49, p. 381]) is still exponential.
2. They have a *canonical form* allowing very simple equivalence, satisfiability, and validity checking.
3. They support efficient calculation of various *Boolean operations* (e.g. conjunction). Please refer to [49, p. 372] for more details about the corresponding algorithms.

These properties make ROBDDs suitable for use in verification as we describe in the next section.

2.3.2 Symbolic Model Checking

We will now describe how BDDs described in the previous subsection can be used for symbolic model checking. We will proceed as follows:

1. **Translate states to conjunctions of literals.** Consider m Boolean variables x_1, \dots, x_m . They can encode 2^m different values¹⁰. Given a set S of states, we need n Boolean variables such that $2^n \leq |S|$ [62]. The smallest n satisfying this condition is in general $n = \lceil \log_2 |S| \rceil$. Each state is then uniquely represented by a conjunct of the variables (e.g. $f_{s_1} = (\neg x_1 \wedge \neg x_2 \wedge \dots \wedge \neg x_n)$). Table 2.3a shows how this can be done.
2. **Translate sets of states to Boolean formulas.** We can encode a set of states $X = \{s_1, s_2, \dots, s_m\}$ as a Boolean formula $f_X = \bigvee_{i=1}^m f_i = f_1 \vee f_2 \vee \dots \vee f_m$. Since each disjunct f_i is a conjunction of literals, the resulting formula f_X is in DNF¹¹. Table 2.3b shows how this can be done.
3. **Translate set operations to Boolean operations.** Let $A, B \subseteq S$ be arbitrary sets of states represented by Boolean formulas f_A, f_B . We can express set operations as follows: $f_{A \cup B} = f_A \vee f_B$ (union), $f_{A \cap B} = f_A \wedge f_B$ (intersection), $f_{A \setminus B} = f_A \wedge \neg f_B$ (difference), and $f_{A'} = f_S \wedge \neg f_A$ (complement).

¹⁰We can think of these values as binary numbers from $x_1 = 0, \dots, x_m = 0$ up to $x_1 = 1, \dots, x_m = 1$.

¹¹Disjunctive normal form.

State	Equivalent formula
s_1	$\neg x_1 \wedge \neg x_2 \wedge \neg x_3$
s_2	$\neg x_1 \wedge \neg x_2 \wedge x_3$
s_3	$\neg x_1 \wedge x_2 \wedge \neg x_3$
s_4	$\neg x_1 \wedge x_2 \wedge x_3$
s_5	$x_1 \wedge \neg x_2 \wedge \neg x_3$
s_6	$x_1 \wedge \neg x_2 \wedge x_3$

Set	Equivalent formula (simplified)
\emptyset	0
$\{s_1\}$	$\neg x_1 \wedge \neg x_2 \wedge \neg x_3$
$\{s_1, s_2\}$	$\neg x_1 \wedge \neg x_2$
$\{s_2, s_3\}$	$\neg x_1 \wedge ((\neg x_2 \wedge x_3) \vee (x_2 \wedge \neg x_3))$
$\{s_4, s_5, s_6\}$	$(\neg x_1 \wedge x_2 \wedge x_3) \vee (x_1 \wedge \neg x_2)$
S	$\neg x_1 \vee \neg x_2$

(a) States.

(b) Sets of states.

Table 2.3: Possible encoding of states using Boolean formulas ($S = \{s_1, s_2, s_3, s_4, s_5, s_6\}$). Formulas for sets of states are obtained by disjunction, e.g. $f_{\{s_1, s_2\}} = f_{s_1} \vee f_{s_2} = (\neg x_1 \wedge \neg x_2 \wedge \neg x_3) \vee (\neg x_1 \wedge \neg x_2 \wedge x_3) = \neg x_1 \wedge \neg x_2$. Note that since certain conjuncts do not represent any states (e.g. $x_1 \wedge x_2 \wedge x_3$), we do not care if they are in a set and thus we can for example encode the set of all states simply as $f_S = 1$. Also note that this is not the only possible encoding of states.

4. Translate relations to Boolean formulas Let $R = \{(s_{1a}, s_{1b}), (s_{2a}, s_{2b}), \dots, (s_{ma}, s_{mb})\} \subseteq S \times S$ be a binary relation. We can encode it as a Boolean expression $f_R = \bigvee_{i=1}^m (f_{s_{ia}} \wedge f'_{s_{ib}}) = (f_{s_{1a}} \wedge f'_{s_{1b}}) \vee (f_{s_{2a}} \wedge f'_{s_{2b}}) \vee \dots \vee (f_{s_{ma}} \wedge f'_{s_{mb}})$, where f'_{s_i} is obtained from f_{s_i} by replacing variables x_1, x_2, \dots, x_n with variables x'_1, x'_2, \dots, x'_n . Figure 2.4 shows how a transition relation can be encoded. Note that we need $2n$ variables to represent the binary relation. Relations of higher arities can be encoded in a similar way.

This allows us to express a model \mathcal{M}_S using BDDs for sets of states, transition relations, protocols, etc. The function $\text{SAT}^{\mathcal{M}_S}(\varphi_P)$ which recursively calculates the set of states of \mathcal{M}_S satisfying a formula φ_P can be easily modified to perform the corresponding set operations on BDDs (instead of explicitly on sets of states).

To give a complete picture, we explain how the predecessor functions $\text{pre}_{\exists}, \text{pre}_{\forall} : 2^S \rightarrow 2^S$ are translated to BDD operations. Intuitively, given a set of states $X \subseteq S$, $\text{pre}_{\exists}(X)$ is the set of states that can reach some state in X in one step and $\text{pre}_{\forall}(X)$ is the set of states such that all states reachable in one step from them are in X . Formally, the general definitions of the functions are¹²:

$$\begin{aligned} \text{pre}_{\exists}(X) &\triangleq \{s \in S \mid \exists s' \in S. R(s, s') \wedge s' \in X\} \\ \text{pre}_{\forall}(X) &\triangleq \{s \in S \mid \forall s' \in S. R(s, s') \rightarrow s' \in X\} \end{aligned}$$

We can observe that:

$$\begin{aligned} \text{pre}_{\forall}(X) &\triangleq \{s \in S \mid \forall s' \in S. R(s, s') \rightarrow s' \in X\} \\ &= S \setminus \{s \in S \mid \neg(\forall s' \in S. R(s, s') \rightarrow s' \in X)\} \\ &= S \setminus \{s \in S \mid \exists s' \in S. \neg(R(s, s') \rightarrow s' \in X)\} \\ &= S \setminus \{s \in S \mid \exists s' \in S. R(s, s') \wedge s' \notin X\} \\ &= S \setminus \{s \in S \mid \exists s' \in S. R(s, s') \wedge s \in S \setminus X\} \\ &= S \setminus \text{pre}_{\exists}(S \setminus X) \end{aligned}$$

Hence, we only need to explain how pre_{\exists} can be computed using BDDs. Let $X \subseteq S$ be an arbitrary set of states, $R \subseteq S \times S$ a transition relation, and f_X, f_R their Boolean representation (e.g. BDDs). We calculate $f_{\text{pre}_{\exists}(X)}$ as follows¹³ [62]:

1. Obtain f'_X by replacing all variables x_i with their primed versions x'_i . Intuitively, we convert X from the set of *current states* to the set of *next states*.
2. Combine the set of next states with the transition relation $f_R \wedge f'_X$. The result now already contains the set of new current states $\text{pre}_{\exists}(X)$. However, it also still contains the next states X .

¹²We defined the predecessor functions on interpreted systems in Definition 2.11.

¹³Please refer to [49] for details on how these operations can be efficiently performed on BDDs.

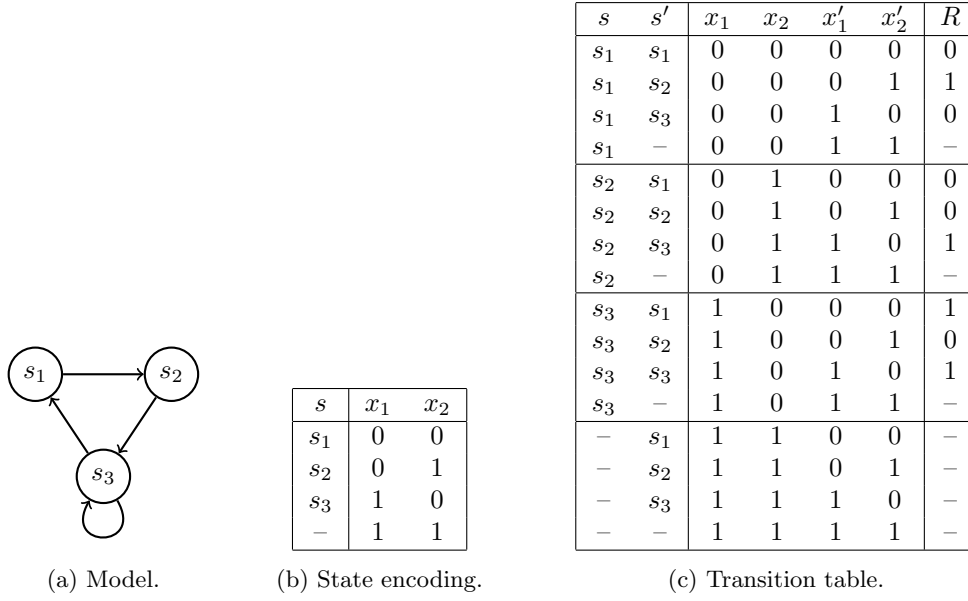


Figure 2.4: Sample encoding of the transition relation $R = \{(s_1, s_2), (s_2, s_3), (s_3, s_1), (s_3, s_3)\}$ of a model with states $S = \{s_1, s_2, s_3\}$ [62]. The Boolean formula representing the transition is $f_R = (\neg x_1 \wedge \neg x_2 \wedge \neg x'_1 \wedge x'_2) \vee (\neg x_1 \wedge x_2 \wedge x'_1 \wedge \neg x'_2) \vee (x_1 \wedge \neg x_2 \wedge \neg x'_1 \wedge \neg x'_2) \vee (x_1 \wedge \neg x_2 \wedge x'_1 \wedge \neg x'_2)$. The value “-” (so-called “don’t care”) represents combinations encoding invalid pairs of states.

3. Calculate $f_{\text{pre}\exists}(X) = \exists x'_1 \exists x'_2 \dots \exists x'_n (f_R \wedge f'_X)$. We are effectively removing the dependence of the current states on the next states.

2.3.3 Summary

In this section, we discussed how model checking can be performed symbolically using efficient operations on BDDs. Compared to explicit approaches, which can handle up to 10^6 states, BDDs support state spaces with up to 10^{25} states (see Table 2.2). We will use BDDs in later parts of this report to describe efficient symbolic implementations of our novel model checking algorithms for fragments of SL (see Subsections 4.2.4 and 5.2.4).

2.4 Automata and Games

Model checking and strategy synthesis for a specification in a particular language (e.g. LTL) can often be reduced to solving an infinite two-player game [77]. In fact, we will use this approach to provide a novel algorithm for the verification of SL[1G] in Chapter 5. This section presents the relevant concepts.

2.4.1 ω -Automata

Automata on infinite words have played a central role in specifying and verifying reactive systems (systems which do not terminate, e.g. operating systems) since their introduction in the 1960’s [43]. Informally, finite ω -automata are finite state machines which accept *infinite words* (i.e. infinite sequences of symbols) according to some acceptance condition. Firstly, we formally define the concept of finite and infinite words [43].

Definition 2.36 (Words). Let Σ be a finite alphabet with symbols $a, b, c, \dots \in \Sigma$. Then,

- Σ^* is the set of all *finite words* over Σ . We use small letters $u, v, w, \dots \in \Sigma^*$ to denote finite words and write $u = u(0)u(1) \dots u(n)$ with $u(i) \in \Sigma$ and length $|u| = n$.
- Σ^ω is the set of all *infinite words* over Σ . We use small greek letters $\alpha, \beta, \gamma, \dots \in \Sigma^\omega$ to denote infinite words and write $\alpha = \alpha(0)\alpha(1) \dots$ with $\alpha(i) \in \Sigma$ and length $|\alpha| = \omega$.

Given an ω -word $\alpha \in \Sigma^\omega$, the set of symbols occurring in α is:

$$\text{Occ}(\alpha) = \{a \in \Sigma \mid \exists i. \alpha(i) = a\}$$

and the set of symbols occurring infinitely often in α is:

$$\text{Inf}(\alpha) = \{a \in \Sigma \mid \forall i \exists j > i. \alpha(j) = a\}$$

An automaton (both finite and infinite) can be decomposed into a semi-automaton and an acceptance condition [77]. We will introduce the two concepts separately. Intuitively, a semi-automaton is a finite state machine whose edges are labelled with the symbols of a finite alphabet.

Definition 2.37 (Semi-automaton). Let S be a finite set of states and Σ an alphabet. Then $\mathfrak{A} = (S, \Sigma, I, R)$ is a *semi-automaton* where $I \subseteq S$ is a set of initial states and $R \subseteq S \times \Sigma \times S$ a transition relation.

A semi-automaton \mathfrak{A} is *deterministic* iff (i) $I = \{s_I\}$ (there is a unique initial state) and (ii) for all $s \in S, a \in \Sigma$ we have $|\{s' \in S \mid R(s, a, s')\}| = 1$ (there is exactly one next state for every state and input). Otherwise, \mathfrak{A} is *non-deterministic*.

The transitions of a semi-automaton can be (equivalently) defined as a transition function $\delta : S \times \Sigma \rightarrow 2^S$. In the case of a *deterministic* semi-automaton, the transition function is simply $\delta : S \times \Sigma \rightarrow S$. Therefore, we will sometimes write a deterministic semi-automaton as $\mathfrak{A} = (S, \Sigma, s_I, \delta)$.

The acceptance component of an automaton is defined with respect to the set of runs of the underlying semi-automaton on an infinite word.

Definition 2.38 (Run). Let $\mathfrak{A} = (S, \Sigma, I, R)$ be a semi-automaton and $\alpha \in \Sigma^\omega$ an infinite word. A *run* of \mathfrak{A} on α is an infinite word $\beta \in S^\omega$ such that (i) $\beta(0) \in I$ and (ii) for all $i \geq 0$ we have $R(\beta(i), \alpha(i), \beta(i+1))$.

There are various acceptance conditions with different expressive power and translation complexity¹⁴. Here we present only the most common ones [43].

Definition 2.39 (Acceptance). Let $\mathfrak{A} = (S, \Sigma, I, R)$ be a semi-automaton. An infinite word $\alpha \in \Sigma^\omega$ is accepted (according to the relevant acceptance condition) iff there exists a run $\beta \in S^\omega$ of \mathfrak{A} on α such that:

- **Büchi acceptance condition** $F \subseteq S$ (a set of states):

$$\text{Inf}(\beta) \cap F \neq \emptyset$$

- **Generalised Büchi acceptance condition** $F_1, \dots, F_n \subseteq S$ (sets of states):

$$\forall i \in \{1, \dots, n\}. \text{Inf}(\beta) \cap F_i \neq \emptyset$$

- **Muller acceptance condition** $\mathcal{F} \subseteq 2^S$ (a set of state sets):

$$\text{Inf}(\beta) \in \mathcal{F}$$

- **Rabin acceptance condition** $\Omega = \{(E_1, F_1), \dots, (E_k, F_k)\}$ where $E_i, F_i \subseteq S$ (a family of pairs of state sets):

$$\exists i \in \{1, \dots, k\}. (\text{Inf}(\beta) \cap E_i = \emptyset) \wedge (\text{Inf}(\beta) \cap F_i \neq \emptyset)$$

- **Streett acceptance condition** $\Omega = \{(E_1, F_1), \dots, (E_k, F_k)\}$ where $E_i, F_i \subseteq S$ (a family of pairs of state sets):

$$\forall i \in \{1, \dots, k\}. (\text{Inf}(\beta) \cap E_i \neq \emptyset) \vee (\text{Inf}(\beta) \cap F_i = \emptyset)$$

¹⁴All automata in Definition 2.39 except for deterministic (generalised) Büchi automata have the same expressive power. However, translations between the classes often incur an exponential blow-up. Please refer to [43, Section 1.7] and [83] for more details.

- **Parity acceptance condition** $c : S \rightarrow \{0, \dots, k\}$ with $k \in \mathbb{N}$ (a colouring function):

$$\min \{c(s) \mid s \in \text{Inf}(\beta)\} \text{ is even}$$

- **Generalised parity acceptance condition** $c_1, \dots, c_n : V \rightarrow \{0, \dots, k_i\}$ with $k_i \in \mathbb{N}$ (a family of colouring functions):

$$\forall i \in \{1, \dots, n\}. \min \{c_i(s) \mid s \in \text{Inf}(\beta)\} \text{ is even}$$

Note that the acceptance condition F of Büchi automata is often referred to as a *fairness condition*. Finally, we are ready to define an ω -automaton as a combination of a semi-automaton and an acceptance condition:

Definition 2.40 (ω -Automata). Let (S, Σ, I, R) be a semi-automaton and \mathcal{F} an acceptance condition. Then $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ is an (existential¹⁵) ω -automaton. $\text{Lang}(\mathfrak{A}) \triangleq \{\alpha \in \Sigma^\omega \mid \mathfrak{A} \text{ accepts } \alpha\}$ is the language of \mathfrak{A} .

Clearly, Büchi automata can be transformed trivially to equivalent generalised Büchi automata by setting $k = 1$ and $F_1 = F$. We also show how a generalised Büchi automaton can be converted to an equivalent Büchi automaton. Let $\mathfrak{A} = (S, \Sigma, I, R, \langle F_1, \dots, F_k \rangle)$ be a generalised Büchi automaton [77, 83]. We construct a deterministic Büchi automaton $\mathfrak{D} = (\{q_1, \dots, q_k, q_a\}, S, q_1, \delta, \{q_a\})$ with the transition function δ defined as:

$$\delta(q_a, s) \triangleq q_1$$

$$\delta(q_i, s) \triangleq \begin{cases} q_a & \text{if } s \in F_i \wedge i = k \\ q_{i+1} & \text{if } s \in F_i \wedge i < k \\ q_i & \text{if } s \notin F_i \end{cases}$$

The automaton \mathfrak{D} has $k + 1$ states and is shown in Figure 2.5. Whenever it is in state q_i , it “waits” until some state in F_i is encountered. Hence, if q_a is visited infinitely often, some state in each of the sets F_1, \dots, F_k must be visited infinitely often. Conversely, if some state in each of the sets F_1, \dots, F_k is visited infinitely often, the state q_a will be visited infinitely often.

The Büchi automaton \mathfrak{B} equivalent to \mathfrak{A} is obtained using the *automaton product* of \mathfrak{A} (without the generalised Büchi acceptance condition) and \mathfrak{D} :

$$\mathfrak{B} = \mathfrak{A} \times \mathfrak{D} = \langle S \times \{q_1, \dots, q_n, q_a\}, \Sigma, I \times \{q_1\}, R \times \delta, S \times \{q_a\} \rangle$$

Note that since \mathfrak{D} is deterministic, the resulting Büchi automaton \mathfrak{B} is deterministic if and only if \mathfrak{A} is deterministic. Some other properties (e.g. unambiguity [77], see Definition 2.44) are also preserved.

2.4.2 Symbolically Represented ω -Automata

ω -automata can soon become too large to store explicitly. Therefore, it is often useful to represent them symbolically using BDDs (see Subsection 2.3.1). In this subsection, we explain how to do that.

Let $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ be an ω -automaton. We proceed as follows:

1. We represent the sets S and Σ using finite disjoint sets of variables V_S and V_Σ respectively. Furthermore, we introduce another set of variables $V'_S = \{v' \mid v \in V_S\}$ so that we could represent the binary transition relation. Subsection 2.3.2 explains how to do this in more detail.
2. We encode the initial set of states $I \subseteq S$ as a propositional formula over V_S .
3. We encode the transition relation $R \subseteq S \times \Sigma \times S$ as a propositional formula over $V_S \cup V_\Sigma \cup V'_S$.

¹⁵All ω -automata in this thesis are *existential* because we require the acceptance condition to be satisfied by *some* run (see Definition 2.39). Some authors consider \forall -automata instead, which accept an infinite word if *all* runs of the automaton on the word are accepting [67].




Figure 2.5: The deterministic Büchi automaton $\mathfrak{D} = (\{q_1, \dots, q_k, q_a\}, S, \delta, q_1, \{q_a\})$ for converting a generalised Büchi automaton $\mathfrak{A} = (S, \Sigma, R, I, (F_1, \dots, F_k))$ to a Büchi automaton. Note that accepting states have double borders and edges are labelled with sets of states (rather than individual states).

4. We encode the acceptance condition \mathcal{F} as a propositional formula (or a set thereof depending on the class of the automaton) over V_S .

The following definition summarises the symbolic representation of ω -automata:

Definition 2.41 (Automaton Formula). Let V_S and V_Σ be two finite disjoint sets of Boolean variables, \mathcal{I} a propositional formula over V_S , \mathcal{R} a propositional formula over $V_S \cup V_\Sigma \cup \{v' \mid v \in V_S\}$, and \mathcal{F} a propositional formula (or a set thereof) over V_S . Then $\mathcal{A}_\exists(V_S, V_\Sigma, \mathcal{I}, \mathcal{R}, \mathcal{F})$ is an (existential) automaton formula.

We often need to combine automata by means of automaton product. This can be done easily when they are represented symbolically. Let $\mathfrak{A}_1 = \mathcal{A}_\exists(V_{S_1}, V_{\Sigma_1}, \mathcal{I}_1, \mathcal{R}_1, \mathcal{F}_1)$ and $\mathfrak{A}_2 = \mathcal{A}_\exists(V_{S_2}, V_{\Sigma_2}, \mathcal{I}_2, \mathcal{R}_2, \mathcal{F}_2)$ be two automaton formulas. The automaton formula representing their product is¹⁶:

$$\mathfrak{A}_1 \times \mathfrak{A}_2 = \mathcal{A}_\exists(V_{S_1} \cup V_{S_2}, V_{\Sigma_1} \cup V_{\Sigma_2}, \mathcal{I}_1 \wedge \mathcal{I}_2, \mathcal{R}_1 \wedge \mathcal{R}_2, \mathcal{F}_1 \cup \mathcal{F}_2)$$

2.4.3 Translating LTL Formulas to ω -automata

There is clearly a resemblance between LTL formulas and ω -automata. Both identify infinite paths which are accepted according to some criterion. Therefore, it should not come as a surprise that arbitrary LTL formulas can be translated to finite ω -automata which are equivalent in the sense that they accept precisely those infinite paths which satisfy the corresponding formulas [86]. In fact, as we have already mentioned in Subsection 2.2.1, many modern model checkers including NuSMV (discussed in Subsection 2.5.4) first translate an LTL formula to a generalised Büchi automaton (using the so-called tableau construction [34]) and then perform CTL model checking with fairness constraints.

In this subsection we describe the *standard translation* [83] of LTL formulas to non-deterministic generalised Büchi automata. As explained in Subsection 2.4.1, it is straightforward to translate a generalised Büchi automaton to a Büchi automaton. Firstly, we define what it means for an LTL formula and an ω -automaton to be equivalent.

Definition 2.42 (LTL- ω -Automaton Equivalence). Fix a finite non-empty set of atomic propositions AP . Let $\varphi \in LTL$ be an LTL formula and $\mathfrak{A} = (S, 2^{AP}, I, R, \mathcal{F})$ an ω -automaton. We say that φ and \mathfrak{A} are *equivalent* if for every interpreted system \mathcal{I} and infinite path $\pi \in Pth$ in \mathcal{I} , we have:

$$\mathcal{I}, \pi \models_{LTL} \varphi \quad \text{iff} \quad \mathfrak{A} \text{ accepts } h^{-1}(\pi)$$

where $h^{-1}(\pi) = h^{-1}(\pi(0))h^{-1}(\pi(1))\dots$ and $h^{-1}(g) = \{p \in AP \mid g \in h(p)\}$ for $g \in G$.

¹⁶We assume that \mathfrak{A}_1 and \mathfrak{A}_2 are generalised Büchi automata. Otherwise, the combined acceptance condition might be different from $\mathcal{F}_1 \cup \mathcal{F}_2$.

We are now ready to describe how LTL formulas are translated to equivalent ω -automata. Figure 2.6 shows a recursive bottom-up algorithm (referred to as *standard translation*) which takes as input an LTL formula ϕ and returns a non-deterministic generalised Büchi automaton formula $\mathcal{A}_{\exists}(Q_{\phi}, AP, \mathcal{I}_{\phi}, \mathcal{R}_{\phi}, \mathcal{F}_{\phi})$ equivalent to ϕ .

Note that the standard translation can violate the constraints on the formulas \mathcal{I} and \mathcal{F} in Definition 2.41, which require that they range over the state variables V_S only. It turns out that the constraint on \mathcal{I} can be temporarily relaxed during the construction as it can be satisfied again at any time by simply adding an extra state variable q :

$$\mathcal{A}_{\exists}(V_S \cup \{q\}, V_S, q, \mathcal{R} \wedge (q \leftrightarrow \mathcal{I}), \mathcal{F})$$

This is necessary only (i) when a temporal operator is encountered¹⁷ and (ii) at the very end so that the final automaton formula satisfies the constraints. Thanks to temporarily relaxing the constraints, we get the result summarised in the following proposition¹⁸.

Proposition 2.2. Let $\varphi \in LTL$ be an arbitrary LTL formula with t temporal operators. We can construct a symbolic representation of a non-deterministic generalised Büchi automaton $\mathcal{A}_{\exists}(Q, AP, \mathcal{I}_{\varphi}, \mathcal{R}_{\varphi}, \mathcal{F}_{\varphi})$ equivalent to φ in time $O(|\varphi|)$ with $|Q| \leq 2t + 1$ state variables. The automaton will have at most 2^{2t+1} states.

In order to give some intuition into how the standard translation works, we will now discuss two examples. Firstly, we consider the LTL formula Xp . The automaton formula equivalent to p is $\mathfrak{A}_p = \mathcal{A}_{\exists}(\emptyset, AP, p, \top, \emptyset)$. Since we will apply the temporal operator X , we have to ensure that the initial state formula contains only state variables. Therefore, we add an extra state variable q_1 and obtain the automaton formula (also equivalent to p) $\mathfrak{A}'_p = \mathcal{A}_{\exists}(\{q_1\}, AP, q_1, q_1 \leftrightarrow p, \emptyset)$. Finally, we apply the rule for X and obtain the automaton formula $\mathfrak{A}_{Xp} = \mathcal{A}_{\exists}(\{q_1, q_2\}, AP, q_2, (q_1 \leftrightarrow p) \wedge (q_2 \leftrightarrow q'_1), \emptyset)$. The automaton is shown in Figure 2.7a. Intuitively, the variables q_1 and q_2 stand for the LTL formulas p and Xp respectively. This is expressed in the transition relation by the fact that q_1 is true whenever p is true and q_2 is true whenever q_1 will be true in the next state.

Secondly, we consider the LTL formula aUb . The automaton formulas equivalent to the atoms a and b are $\mathfrak{A}_a = \mathcal{A}_{\exists}(\emptyset, AP, a, \top, \emptyset)$ and $\mathfrak{A}_b = \mathcal{A}_{\exists}(\emptyset, AP, b, \top, \emptyset)$ respectively. The initial state formula of the automaton equivalent to the right subformula of aUb must contain only state variables (so that the associated fairness formula would contain only state variables). Therefore, we introduce an extra state variable q_1 and obtain the automaton formula $\mathfrak{A}'_b = \mathcal{A}_{\exists}(\{q_1\}, AP, q_1, q_1 \leftrightarrow b, \emptyset)$. Before applying the operator U , we calculate the automaton product $\mathfrak{A}_a \times \mathfrak{A}'_b = \mathcal{A}_{\exists}(\{q_1\}, AP, a \wedge q_1, q_1 \leftrightarrow b, \emptyset)$. We now introduce another extra variable q_2 which should be true whenever aUb is true. Intuitively, q_2 should be true when b is true (equivalent to q_1) or a is true and q_2 is true in the next state. Hence, we get the following recursive relation¹⁹:

$$q_2 \leftrightarrow q_1 \vee a \wedge q'_2$$

However, this specification is not sufficient because it is also satisfied by a path where a holds forever from the current state onwards. In other words, the specification admits all paths satisfying the LTL formula $aWb \equiv (aUb) \vee Ga$. In order to enforce that b (q_1) is eventually true (in case q_2 is true in the current state), we must add the fairness constraint:

$$q_2 \rightarrow q_1$$

We show that given an arbitrary infinite path along which the fairness condition is satisfied infinitely often, q_2 is true in the first state of the path if and only if aUb holds on the path:

\Rightarrow : Assume that q_2 is true in the first state and the fairness constraint is satisfied infinitely often. Then, either²⁰ q_2 is infinitely often false (in which case b must be true after a finite number of steps

¹⁷Note that when calculating the automaton formula for $\varphi U \psi$ (in Figure 2.6), \mathcal{I}_{φ} need not contain only state variables because it is neither used in a fairness constraint (like \mathcal{I}_{ψ} for U and \mathcal{I}_{φ} for F, G) nor are its variables switched between the current and the next state (like \mathcal{I}_{φ} for X). Hence, every temporal operator will incur at most two new state variables in total.

¹⁸The upper bound for the number of state variables (and consequently states) is slightly tighter than in Proposition 1 of [77].

¹⁹Note that this relation is very similar to the fixpoint rule for $\text{SAT}_{CTL}^T(E[\varphi_1 U \varphi_2])$ in Definition 2.12.

²⁰We use the equivalence $q_2 \rightarrow q_1 \equiv \neg q_2 \vee q_1$ and the fact that a disjunction fairness constraint $GF(\varphi \vee \psi)$ is equivalent to the disjunction of fairness constraints $(GF\varphi) \vee (GF\psi)$.


```

1 function GENBÜCHI( $\Phi$ )
2   switch  $\Phi$  do
3     case  $p \in AP$ :
4       return  $\mathcal{A}_{\exists}(\emptyset, AP, p, \top, \emptyset)$ 
5     case  $\neg\varphi$ :
6        $\mathcal{A}_{\exists}(Q_{\Phi}, AP, \mathcal{I}_{\varphi}, \mathcal{R}_{\Phi}, \mathcal{F}_{\Phi}) \equiv \text{GENBÜCHI}(\varphi)$ 
7       return  $\mathcal{A}_{\exists}(Q_{\Phi}, AP, \neg\mathcal{I}_{\varphi}, \mathcal{R}_{\Phi}, \mathcal{F}_{\Phi})$ 
8     case  $\varphi \wedge \psi$ :
9        $\mathcal{A}_{\exists}(Q_{\Phi}, AP, \mathcal{I}_{\varphi} \wedge \mathcal{I}_{\psi}, \mathcal{R}_{\Phi}, \mathcal{F}_{\Phi}) \equiv \text{GENBÜCHI}(\varphi) \times \text{GENBÜCHI}(\psi)$ 
10      return  $\mathcal{A}_{\exists}(Q_{\Phi}, AP, \mathcal{I}_{\varphi} \wedge \mathcal{I}_{\psi}, \mathcal{R}_{\Phi}, \mathcal{F}_{\Phi})$ 
11     case  $\varphi \vee \psi$ :
12       $\mathcal{A}_{\exists}(Q_{\Phi}, AP, \mathcal{I}_{\varphi} \vee \mathcal{I}_{\psi}, \mathcal{R}_{\Phi}, \mathcal{F}_{\Phi}) \equiv \text{GENBÜCHI}(\varphi) \times \text{GENBÜCHI}(\psi)$ 
13      return  $\mathcal{A}_{\exists}(Q_{\Phi}, AP, \mathcal{I}_{\varphi} \vee \mathcal{I}_{\psi}, \mathcal{R}_{\Phi}, \mathcal{F}_{\Phi})$ 
14     case  $X\varphi$ :
15       $\mathcal{A}_{\exists}(Q_{\varphi}, AP, \mathcal{I}_{\varphi}, \mathcal{R}_{\varphi}, \mathcal{F}_{\varphi}) \equiv \text{GENBÜCHI}(\varphi)$ 
16       $q := \text{NEWVAR}$ 
17      return  $\mathcal{A}_{\exists}(Q_{\varphi} \cup \{q\}, AP, q, \mathcal{R}_{\varphi} \wedge (q \leftrightarrow X\mathcal{I}_{\varphi}), \mathcal{F}_{\varphi})$ 
18     case  $F\varphi$ :
19       $\mathcal{A}_{\exists}(Q_{\varphi}, AP, \mathcal{I}_{\varphi}, \mathcal{R}_{\varphi}, \mathcal{F}_{\varphi}) \equiv \text{GENBÜCHI}(\varphi)$ 
20       $q := \text{NEWVAR}$ 
21      return  $\mathcal{A}_{\exists}(Q_{\varphi} \cup \{q\}, AP, q, \mathcal{R}_{\varphi} \wedge (q \leftrightarrow \mathcal{I}_{\varphi} \vee Xq), \mathcal{F}_{\varphi} \cup \{q \rightarrow \mathcal{I}_{\varphi}\})$ 
22     case  $G\varphi$ :
23       $\mathcal{A}_{\exists}(Q_{\varphi}, AP, \mathcal{I}_{\varphi}, \mathcal{R}_{\varphi}, \mathcal{F}_{\varphi}) \equiv \text{GENBÜCHI}(\varphi)$ 
24       $q := \text{NEWVAR}$ 
25      return  $\mathcal{A}_{\exists}(Q_{\varphi} \cup \{q\}, AP, q, \mathcal{R}_{\varphi} \wedge (q \leftrightarrow \mathcal{I}_{\varphi} \wedge Xq), \mathcal{F}_{\varphi} \cup \{\mathcal{I}_{\varphi} \rightarrow q\})$ 
26     case  $\varphi U \psi$ :
27       $\mathcal{A}_{\exists}(Q_{\Phi}, AP, \mathcal{I}_{\varphi} \wedge \mathcal{I}_{\psi}, \mathcal{R}_{\Phi}, \mathcal{F}_{\Phi}) \equiv \text{GENBÜCHI}(\varphi) \times \text{GENBÜCHI}(\psi)$ 
28       $q := \text{NEWVAR}$ 
29      return  $\mathcal{A}_{\exists}(Q_{\Phi} \cup \{q\}, AP, q, \mathcal{R}_{\Phi} \wedge (q \leftrightarrow \mathcal{I}_{\psi} \vee \mathcal{I}_{\varphi} \wedge Xq), \mathcal{F}_{\Phi} \cup \{q \rightarrow \mathcal{I}_{\psi}\})$ 
30   end switch
31 end function

```

Figure 2.6: Recursive algorithm for translating LTL formulas to non-deterministic generalised Büchi automata (standard translation) [83]. The function NEWVAR returns a new unused Boolean variable. For an arbitrary Boolean formula φ , $X\varphi$ refers to the formula obtained from φ by swapping all Boolean variables with their primed versions (i.e. switching from the current state variables to the next state variables).

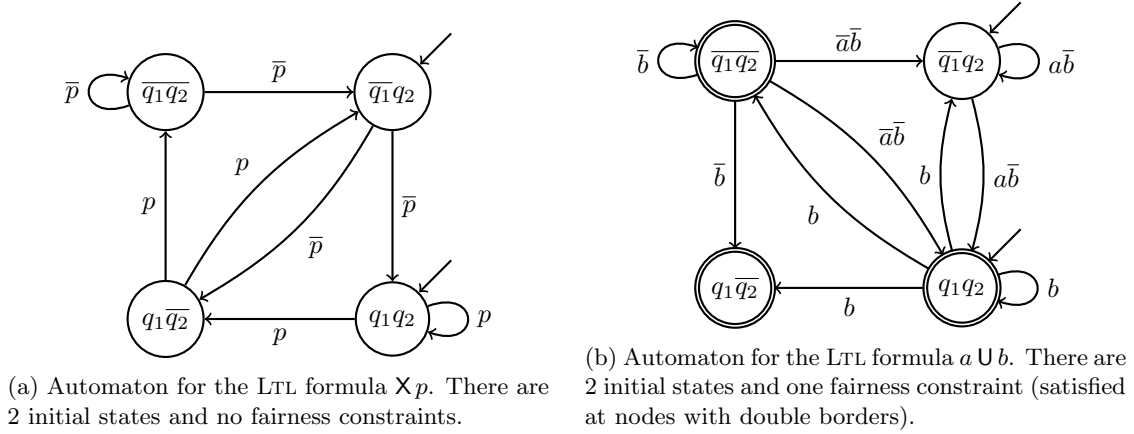


Figure 2.7: Automata generated using standard translation. Negated variables are overlined for readability (e.g. $\overline{q_1 q_2}$ stands for $\neg q_1 \wedge \neg q_2$).

in order for q_2 to be true in the first state) or b is infinitely often true (so again, b will be true after a finite number of steps). In both cases, $a U b$ holds (thanks to the fact that b will be true for the first time in a finite number of steps, until which point a must be true since q_2 is true in the first state).

\Leftarrow : Assume that $a U b$ holds on the path. Then b is true at some point in the (finite) future. Moreover, a is true in every state until that point. Hence, by the recursive relation, q_2 is true in the first state.

The resulting generalised Büchi automaton formula equivalent to the LTL formula $a U b$ (obtained by adding the recursive relation and fairness constraint to the automaton product $\mathfrak{A}_a \times \mathfrak{A}'_b$) is:

$$\mathfrak{A}_{a U b} = \mathcal{A}_{\exists}(\{q_1, q_2\}, AP, q_2, (q_1 \leftrightarrow b) \wedge (q_2 \leftrightarrow q_1 \vee a \wedge q'_2), \{q_2 \rightarrow q_1\})$$

Note that the exactly same automaton formula is obtained by applying the standard translation. The corresponding automaton is shown in figure 2.7b.

2.4.4 Determinisation

The Büchi²¹ automata obtained by the standard translation, discussed in Subsection 2.4.3, are non-deterministic. In order to be able to synthesise strategies which enforce the corresponding LTL formulas, we need to further transform the non-deterministic automata into equivalent deterministic automata.

Translation of non-deterministic automata to deterministic ω -automata is typically performed using Safra's algorithm [82]. Although the procedure is optimal for Büchi automata, it is very difficult to implement because it generates deterministic Rabin automata over trees, which cannot be efficiently represented symbolically using BDDs [77].

In [77, Section 5], a new determinisation procedure for *non-confluent* Büchi automata, which we will use in our novel model checking algorithm for SL[1G] in Chapter 5, was proposed. Intuitively, an ω -automaton is non-confluent if every finite prefix of a run is uniquely determined by the last visited state [77].

Definition 2.43 (Non-confluent ω -Automata). Let $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ be an ω -automaton. We say that \mathfrak{A} is *non-confluent* if for every infinite word α we have: if β_1 and β_2 are two runs of \mathfrak{A} on α that intersect at a position t_0 (i.e. $\beta_1(t_0) = \beta_2(t_0)$), then $\beta_1(t) = \beta_2(t)$ for all $t \leq t_0$.

We show how to convert the automata generated by the standard translation to equivalent non-confluent automata before describing the procedure itself. Firstly, we note that the standard translation generates *unambiguous* automata [25]. An automaton is unambiguous if every accepted word has a unique accepting run (although there may be other non-accepting runs).

²¹Standard translation constructs non-deterministic *generalised* Büchi automata. However, as explained in Subsection 2.4.1, it is straightforward to transform a generalised Büchi automaton into an equivalent Büchi automaton.

Definition 2.44 (Unambiguous ω -Automata). Let $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ be an ω -automaton and set $\mathfrak{A}^* = (S, \Sigma, S, R, \mathcal{F})$. We say that \mathfrak{A} is *unambiguous* iff for every infinite word $\alpha \in \Sigma^\omega$, there is at most one accepting run $\beta \in S^\omega$ of \mathfrak{A}^* on α ²².

Proposition 2.3. The generalised Büchi automata generated by the standard translation in Figure 2.6 are unambiguous.

It turns out that there is a connection between unambiguous and non-confluent automata: An unambiguous automaton without *useless states* is non-confluent [77]. Intuitively, a state is *not* useless if (i) it is reachable from one of the initial states and (ii) there exists some accepting run starting from it.

Definition 2.45 (Useless States). Let $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ be an ω -automaton and $q \in S$ a state. We say that q is *useless* if (i) there is no initial state $q_I \in I$ such that q is reachable from q_I in \mathfrak{A} or (ii) there is no infinite word $\alpha \in \Sigma^\omega$ accepted by the ω -automaton $\mathfrak{A}' = (S, \Sigma, \{q\}, R, \mathcal{F})$.

Proposition 2.4. Let $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ be an *unambiguous* automaton with no useless states. Then \mathfrak{A} is non-confluent.

Therefore, if we remove useless states from an unambiguous automaton, we obtain a non-confluent automaton. Given an automaton $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$, the set of reachable states which are not dead ends (i.e. states which are *not* useless) can be calculated as an intersection of two fixpoints:

$$Q = \underbrace{\text{lfp}_X [I \cup \{s \in S \mid \exists x \in X \exists a \in \Sigma. R(x, a, s)\}]}_{\text{calculate reachable states}} \cap \underbrace{\text{gfp}_X \{s \in S \mid \exists x \in X \exists a \in \Sigma. R(s, a, x)\}}_{\text{remove dead ends}}$$

This is very simple to implement symbolically using operations on BDDs (see Subsection 2.3.2).

To sum up, we can convert an LTL formula to an equivalent non-confluent non-deterministic Büchi automaton as follows (all steps can be performed symbolically):

1. We convert the LTL formula to a generalised Büchi automaton using the standard translation (Subsection 2.4.3). The automaton is unambiguous (Proposition 2.3).
2. We convert the generalised Büchi automaton to a Büchi automaton using automaton product as explained in Subsection 2.4.1. The transformation preserves unambiguity [77].
3. We remove useless states from the Büchi automaton using the fixpoint computation above. The resulting automaton is non-confluent (Proposition 2.4).

Determinisation Procedure for Non-confluent Büchi Automata

A non-confluent Büchi automaton can be determinised using a new procedure presented in [77, 78]²³, which is a specialisation of Safra's algorithm for non-confluent automata. Let $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ be a non-confluent Büchi automaton with n reachable states. The procedure constructs a *deterministic parity automaton* \mathfrak{B} equivalent to \mathfrak{A} . The states of \mathfrak{B} are $(n + 1)$ -tuples of pairs (S_i, m_i) , where $S_i \subseteq S$ is a set of states and $m_i \in \{0, 1\}$ a Boolean flag for the visit of accepting states. The initial state of \mathfrak{B} is $((I, 0), (\emptyset, 0), \dots, (\emptyset, 0))$. To compute the successor of a state $((S_0, m_0), \dots, (S_n, m_n))$ under input $a \in \Sigma$, we proceed as follows (we treat the tuple as a list):

1. **Delete empty sets.** Let e be the lowest index such that $S_e = \emptyset$. We delete the pair (S_e, m_e) and shift all pairs (S_i, m_i) with $i > e$ to the left. In addition, we delete all pairs from the right end with $S_i = \emptyset$ except for (S_0, m_0) .
2. **Calculate successors.** We replace each subset S_i by its existential successors $\text{succ}_{\exists}^{R,a}(S_i)$:

$$\text{succ}_{\exists}^{R,a}(S_i) = \{s' \in S \mid \exists s \in S_i. R(s, a, s')\}$$

Note that S_0 is the set of states that \mathfrak{A} can reach after having read a finite input word $w \in \Sigma^*$ starting in one of the initial states.

²²i.e. there is at most one accepting run $\beta \in S^\omega$ of \mathfrak{A} on α starting in an arbitrary state ($\beta(0) \in S$).

²³The first version of the new procedure was introduced in [78]. The updated version, presented in [77], is slightly different (as explained in the footnote on page 82 in [77]). Our implementation of the updated version did not work properly on some examples, whereas the first one did. Therefore, we describe here the first version. We gained some useful insights about the procedure from [89], where it is tailored for the verification of parametric linear temporal logic (PLTL).

3. **Append.** We append $(S_0 \cap \mathcal{F}, 0)$ to the end of the list to remember that the paths leading to states $S_0 \cap \mathcal{F}$ have already visited \mathcal{F} .
4. **Mark and clean up.** Let ℓ be the length of the list. For each $i \in \{0, \dots, \ell - 1\}$, if $S_i \setminus \mathcal{F} \subseteq \bigcup_{j=i+1}^{\ell-1} S_j$ and $S_i \neq \emptyset$, then we set $m_i = 1$ and remove S_i from all S_j with $j > i$. Otherwise, we set $m_i = 0$.

As \mathfrak{A} is non-confluent, we know that a finite run is uniquely characterised by its last state. Hence, if a state occurs in two sets S_i and S_j with $j > i$, then we know that both sets are following the same run. Thus, whenever a state set S_i contains only states which are accepting (i.e. in \mathcal{F}) or in any of the state sets S_j with $j > i$, we know that all runs in S_i have visited an accepting state recently. Hence, we also mark S_i as accepting and remove the states S_i from all S_j with $j > i$.

5. **Pad.** Pad the list with $n + 1 - \ell$ pairs $(\emptyset, 0)$ at the end. The length of the resulting list is $n + 1$ and hence is a valid state of \mathfrak{P} .

After the clean up, every non-empty state set S_i must contain some state $q \in S_i$ such that $q \notin S_j$ for all $j > i$. Hence, the padded list with $n + 1$ pairs must have at least one empty state set. Thus, at least one pair will be removed in the step 1 and the list will again have length at most $n + 1$ after step 3.

We now explain how to determine the colour of an arbitrary state $((S_0, m_0), \dots, (S_n, m_n))$ of \mathfrak{P} . Let e be the lowest index such that $S_e = \emptyset$ and m the lowest index i such that $m_i = 1$ (or ∞ if there is no such index). The colouring function $c : S \rightarrow \{0, \dots, 2n - 1\}$ is defined as²⁴:

$$c(q) = \begin{cases} 1 & \text{if } e = 0 \\ 2m & \text{if } m < e \\ 2e - 1 & \text{if } 0 < e < m \end{cases}$$

The idea of the construction is as follows: If an entry S_i never becomes empty after a certain position on a run π of \mathfrak{P} on $\alpha \in \Sigma^\omega$ and is marked infinitely often, then we know that a run ξ of \mathfrak{A} on α contained in π is introduced infinitely often in step 3. Hence, ξ is accepting²⁵ and π contains an accepting run of \mathfrak{A} . Please refer to Section 3 of [78] for a more detailed explanation. The procedure described above is formalised in the following definition.

Definition 2.46 (Non-confluent Automaton Determinisation). Let $\mathfrak{A} = (S, \Sigma, I, R, \mathcal{F})$ be a non-confluent Büchi automaton with $|S| = n$ states. A deterministic parity automaton $\mathfrak{P} = (T, \Sigma, t_I, \delta, c)$ equivalent to \mathfrak{A} is constructed as follows:

- The states of \mathfrak{P} are $(n + 1)$ -tuples of pairs containing a subset of S and a Boolean flag: $T = \{((S_0, m_0), \dots, (S_n, m_n)) \mid S_i \subseteq S \wedge m_i \in \{0, 1\}\}$.
- The initial state is $t_I = ((I, 0), (\emptyset, 0), \dots, (\emptyset, 0))$.
- The transition function δ is defined as follows: Given a state $t = ((S_0, m_0), \dots, (S_n, m_n))$, the next state $t' = ((S'_0, m'_0), \dots, (S'_n, m'_n))$ under input $a \in \Sigma$ satisfies:

$$m'_i = \begin{cases} \left(\text{suc}_{\exists}^{R,a}(S_i) \setminus \mathcal{F} \right) \subseteq \left(\bigcup_{j=i+1}^n \text{suc}_{\exists}^{R,a}(S_j) \right) & \text{if } i < e \wedge \text{suc}_{\exists}^{R,a}(S_i) \neq \emptyset \\ \left(\text{suc}_{\exists}^{R,a}(S_{i+1}) \setminus \mathcal{F} \right) \subseteq \left(\bigcup_{j=i+2}^n \text{suc}_{\exists}^{R,a}(S_j) \right) & \text{if } i \geq e \wedge \text{suc}_{\exists}^{R,a}(S_{i+1}) \neq \emptyset \\ 0 & \text{else} \end{cases}$$

$$S'_0 = \text{suc}_{\exists}^{R,a}(S_0)$$

$$S'_i = \left(\begin{cases} \text{suc}_{\exists}^{R,a}(S_i) & \text{if } i < e \\ \text{suc}_{\exists}^{R,a}(S_{i+1}) & \text{if } e \leq i < E \\ \text{suc}_{\exists}^{R,a}(S_0) \cap \mathcal{F} & \text{if } e < i = E \vee e = i = E + 1 \\ \emptyset & \text{else} \end{cases} \right) \setminus \left(\bigcup_{\substack{j=0 \\ m'_j=1}}^{i-1} S'_j \right) \quad \text{for } i > 0$$

²⁴Note that $e \leq n$ (at least one state set S_i must be empty) and $m \neq e$ for the reachable states. By considering all three cases in the definition of the colouring function, we see that the maximum possible colour for any reachable state is $2n - 1$ because it cannot be the case that both $m = n$ and $m < e$.

²⁵Since the introduction of ξ in step 3 occurs only when an accepting state is encountered.

where e is the lowest index such that $S_e = \emptyset$ and E is the largest index such that $S_E \neq \emptyset$.

- The colour of a state $t = ((S_0, m_0), \dots, (S_n, m_n))$ is:

$$c(t) = \begin{cases} 1 & \text{if } e = 0 \\ 2m & \text{if } m < e \\ 2e - 1 & \text{if } 0 < e < m \end{cases}$$

where e is the lowest index such that $S_e = \emptyset$ and m is the lowest index i such that $m_i = 1$ (or ∞ if there is no such index).

The complexity of the procedure is provided in the following proposition [77]²⁶:

Proposition 2.5. Given a non-confluent Büchi automaton with n states, the construction in Definition 2.46 yields a deterministic parity automaton with at most $2^{(n+1)^2}$ states and $2n$ colours.

Examples of the Determinisation Procedure

In order to give some intuition into how the new determinisation procedure works, we will now discuss two examples. Firstly, we consider the non-deterministic unambiguous generalised Büchi automaton \mathfrak{A}_{Xp} equivalent to the LTL formula Xp in Figure 2.7a. Since it has no fairness constraints, we can treat it as a non-deterministic unambiguous Büchi automaton where all states are accepting²⁷. We start by calculating the set of reachable states which are not dead ends. There are four such states: $s_1 = \overline{q_1}q_2$, $s_2 = \overline{q_1}q_2$, $s_3 = q_1\overline{q_2}$, and $s_4 = q_1q_2$. Hence, the states of the equivalent deterministic parity automaton \mathfrak{P}_{Xp} will be 5-tuples.

The initial state of the parity automaton is $t_I = ((\{s_2, s_4\}, 0), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0))$. The colour of the initial state is $c(t_I) = 1$ since $e = 1$ and $m = \infty$. In order to determine $\delta(t_I, p)$, we perform steps 1–5 of the determinisation procedure and obtain the next state $t_1 = ((\{s_3, s_4\}, 1), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0))$. Also, we find that $\delta(t_I, \overline{p}) = t_1$, i.e. the second state of \mathfrak{P}_{Xp} is the same regardless of the input. This is in line with the meaning of Xp (only the second input “matters”).

The final parity automaton \mathfrak{P}_{Xp} is shown in Figure 2.8a. It has four reachable states and uses two colours. It is indeed equivalent to the LTL formula Xp : If the second input is $\neg p$, the run will stay forever in the state $((\emptyset, 0), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0))$ with colour 1. Hence, the minimal colour occurring infinitely often will be odd so the run will not be accepted (see Definition 2.39). Conversely, if the second input is p , the run will stay forever in the state $((\{s_1, s_2, s_3, s_4\}, 1), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0), (\emptyset, 0))$ with colour 0. Hence, the minimal colour occurring infinitely often will be even so the run will be accepted. Note that the last four state sets S_1 – S_4 are always empty in all reachable states of \mathfrak{P}_{Xp} . Thus, we could have used single pairs (S_0, m_0) as states of \mathfrak{P}_{Xp} (instead of 5-tuples of pairs).

As a second example, we consider the non-deterministic unambiguous generalised Büchi automaton \mathfrak{A}_{aUb} equivalent to the LTL formula aUb in Figure 2.7b. Since it has exactly one fairness constraint, we can treat it as a non-deterministic unambiguous Büchi automaton straightaway²⁸. There are only three reachable states of \mathfrak{A}_{aUb} which are not dead ends: $s_1 = \overline{q_1}q_2$, $s_2 = \overline{q_1}q_2$, and $s_3 = q_1q_2$. Hence, the states of the equivalent deterministic parity automaton \mathfrak{P}_{aUb} will be 4-tuples. The parity automaton \mathfrak{P}_{aUb} is shown in Figure 2.8b. It has eight reachable states and uses three colours. Clearly, a much smaller deterministic parity automaton, also equivalent to aUb , can be constructed by hand²⁹. Note that in this case, all four pairs $(S_0, m_0), \dots, (S_3, m_3)$ are required in the tuple.

²⁶Note that since we use the original version of the procedure introduced in [78], the tuples consist of $n+1$ pairs (S_i, m_i) (rather than n as in the updated version [77]). There are 2 possible values for m_i and 2^n possible values for each S_i . Hence, there are $(2 \times 2^n)^{n+1} = 2^{(n+1)^2}$ possible states of the automaton. We also provide a slightly tighter bound on the number of colours because we have shown that the maximum possible colour of any reachable state is $2n - 1$.

²⁷If we apply the general procedure for converting generalised Büchi automata to Büchi automata in Subsection 2.4.1, the auxiliary deterministic Büchi automaton will be $\mathfrak{D} = (\{q_a\}, S, q_a, \delta, \{q_a\})$ with $\delta(q_a, s) = q_a$ for all $s \in S$. Observe that \mathfrak{D} has only one state, which is accepting, and accepts all infinite paths $\pi \in S^\omega$. Hence, the automaton product will have identical states and transitions as the original generalised Büchi automaton. Moreover, all its states will be accepting.

²⁸The general procedure for converting generalised Büchi automata to Büchi automata in Subsection 2.4.1 would actually double the number of states. This is due to the fact that the auxiliary deterministic Büchi automaton \mathfrak{D} would have two states.

²⁹However, constructing automata manually is only feasible for simple LTL formulas. For a more systematic approach, please refer to Chapter 3 of [77] for an introduction to minimisation of ω -automata.


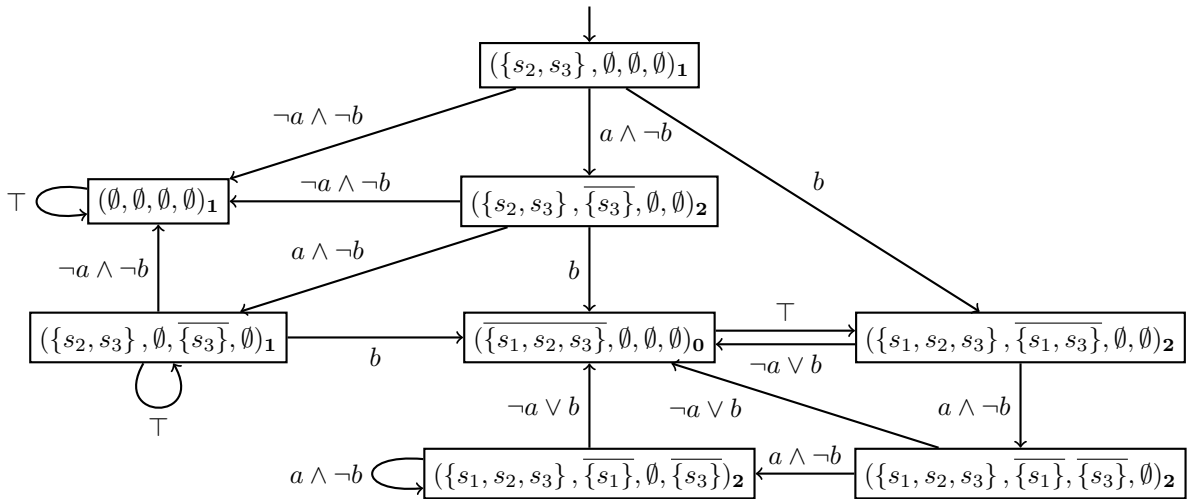
(a) Deterministic parity automaton for the LTL formula Xp .(b) Deterministic parity automaton for the LTL formula $a U b$.

Figure 2.8: Deterministic parity automata constructed using the new determinisation procedure. Marked sets are overlined for readability (e.g. $\overline{\{s_3, s_4\}}$ stands for $(\{s_3, s_4\}, 1)$). The colours of the states are written as subscripts (e.g. the colour of the state $(\emptyset, \emptyset, \emptyset, \emptyset, \emptyset)_1$ is 1).

Symbolic Implementation of the Determinisation Procedure

Although the determinisation procedure is quite complicated, it is possible to implement it symbolically [77, 78]. Let $\mathfrak{A} = \mathcal{A}_{\exists}(V, AP, \mathcal{I}, \mathcal{R}, \mathcal{F})$ be a non-deterministic non-confluent Büchi automaton with m state variables $v_0, \dots, v_{m-1} \in V$ and n reachable states³⁰ $s_0, s_1, \dots, s_{n-1} \in S \subseteq 2^V$. Hence, the states of the equivalent parity automaton $\mathfrak{P} = \mathcal{A}_{\exists}(Q, AP, \mathcal{I}_P, \mathcal{R}_P, \langle C_i \rangle_{i \in \{0, \dots, 2n+1\}})$ are $(n+1)$ -tuples $((S_0, m_i), \dots, (S_n, m_n))$. For each subset S_i , we allocate n variables q_i^0, \dots, q_i^{n-1} such that q_i^k represents that $s_k \in S_i$. Furthermore, we introduce $n+1$ variables m_0, \dots, m_n to represent the markings. The set of state variables of \mathfrak{P} is thus:

$$Q \triangleq \{q_i^k \mid 0 \leq k < n \wedge 0 \leq i \leq n\} \cup \{m_i \mid 0 \leq i \leq n\}$$

In order to be able to express the transition function, we also need the next versions of all variables in Q (i.e. $q_i^{k'}$ and m_i'). The initial state of \mathfrak{P} is encoded as:

$$\mathcal{I}_P \triangleq \underbrace{\left[\bigwedge_{s_k \in \mathcal{I}} q_0^k \right] \wedge \left[\bigwedge_{s_k \notin \mathcal{I}} \neg q_0^k \right]}_{S_0 = \mathcal{I}} \wedge \underbrace{\left[\bigwedge_{i=1}^n \bigwedge_{k=0}^{n-1} \neg q_i^k \right]}_{S_i = \emptyset \text{ for all } i > 0} \wedge \underbrace{\left[\bigwedge_{i=0}^n \neg m_i \right]}_{m_i = 0 \text{ for all } i}$$

Before encoding the transition relation, we need to define several auxiliary formulas. Firstly, we need to express the successor function $\text{succ}_{\exists}^{\mathcal{R}, a}(S_i)$. For this purpose we introduce a family of formulas φ_i^k over $Q \cup AP$, such that φ_i^k is true iff $s_k \in \text{succ}_{\exists}^{\mathcal{R}, a}(S_i)$:

$$\varphi_i^k \triangleq \bigvee_{j=0}^{n-1} q_i^j \wedge \underbrace{\exists v_0 \dots v_{m-1} v'_0 \dots v'_{m-1} s_j \wedge \mathcal{R} \wedge s'_k}_{\text{propositional formula over } AP \text{ for transition from } s_j \text{ to } s_k}$$

Secondly, we need to express the indices e and E . We introduce a family of formulas $\Gamma_{i \geq e}$ and $\Gamma_{i=E}$, such that Γ_{ε} is true iff the expression ε is true:

$$\Gamma_{i \geq e} \triangleq \underbrace{\bigvee_{j=0}^i \bigwedge_{k=0}^{n-1} \neg q_j^k}_{S_j = \emptyset \text{ for some } j \leq i} \quad \Gamma_{i=E} \triangleq \underbrace{\bigvee_{k=0}^{n-1} q_i^k}_{S_i \neq \emptyset} \wedge \underbrace{\left[\bigwedge_{j=i+1}^n \bigwedge_{k=0}^{n-1} \neg q_j^k \right]}_{S_j = \emptyset \text{ for all } j > i}$$

We are now ready to encode the transition relation. Recall that a state set S_i is marked (i.e. $m_i = 1$) iff it is non-empty and every non-accepting state $s_k \in S_i \setminus \mathcal{F}$ appears in some state set S_j with $j > i$. Thus, the equation for the next value of each of the variables m_i is:

$$\Xi_i \triangleq m_i' \leftrightarrow \underbrace{\left[\bigvee_{k=0}^{n-1} q_i^{k'} \right]}_{S_i' \neq \emptyset} \wedge \underbrace{\left[\bigwedge_{s_k \in S \setminus \mathcal{F}} \left(q_i^{k'} \rightarrow \bigvee_{j=i+1}^n \varphi_j^k \right) \right]}_{(S_i' \setminus \mathcal{F}) \subseteq (\bigcup_{j=i+1}^n S_j')}$$

There are three possible cases for the next state set S_i' : (i) successors of the local state set $\text{succ}_{\exists}^{\mathcal{R}, a}(S_i)$ (when $i < e$), (ii) successors of the right state set $\text{succ}_{\exists}^{\mathcal{R}, a}(S_{i+1})$ (when $e \leq i < E$), or (iii) accepting successors of the first state set $\text{succ}_{\exists}^{\mathcal{R}, a}(S_0) \cap \mathcal{F}$ (when $e < i = E$ or $e = i = E + 1$). If none of the conditions is satisfied, $S_i' = \emptyset$. The first case always applies to the first set S_0' (regardless of the values of e and E):

$$\Delta_0^k \triangleq q_0^{k'} \leftrightarrow \varphi_0^k$$

³⁰Note that although the set of reachable states can be calculated symbolically, we need to explicitly enumerate them. Therefore, strictly speaking, the procedure is *semi-symbolic*.

For $i > 0$, when a state $s_k \notin \mathcal{F}$ is not accepting, the first two cases apply to S'_i :

$$\Phi_i^k \triangleq q_i^{k'} \leftrightarrow \underbrace{(\neg\Gamma_{i \geq e} \wedge \varphi_i^k)}_{\text{first case}} \vee \underbrace{\Gamma_{i \geq e} \wedge \varphi_{i+1}^k}_{\text{second case}} \wedge \neg \underbrace{\bigvee_{j=0}^{i-1} (q_j^{k'} \wedge m'_j)}_{\text{remove marked state sets on the left}}$$

For $i > 0$, when a state $s_k \in \mathcal{F}$ is accepting, all three cases apply to S'_i :

$$\Psi_i^k \triangleq q_i^{k'} \leftrightarrow (\neg\Gamma_{i \geq e} \wedge \varphi_i^k \vee \Gamma_{i \geq e} \wedge \varphi_{i+1}^k \vee \underbrace{[\Gamma_{i \geq e} \wedge \Gamma_{i=E} \vee \neg\Gamma_{i-1 \geq e} \wedge \Gamma_{i-1=E}] \wedge \varphi_0^k}_{\text{third case}}) \wedge \neg \bigvee_{j=0}^{i-1} (q_j^{k'} \wedge m'_j)$$

Finally, the complete transition relation is:

$$\mathcal{R}_P \triangleq \underbrace{\left[\bigwedge_{k=0}^{n-1} \Delta_0^k \wedge \Xi_0 \right]}_{(S'_0, m'_0)} \wedge \bigwedge_{i=1}^n \underbrace{\left[\left(\bigwedge_{s_k \notin \mathcal{F}} \Phi_i^k \right) \wedge \left(\bigwedge_{s_k \in \mathcal{F}} \Psi_i^k \right) \wedge \Xi_i \right]}_{(S'_i, m'_i) \text{ for all } i > 0}$$

It remains to show how to encode the colouring function. We represent it by a family of formulas $\mathcal{C}_0, \dots, \mathcal{C}_{2n+1}$ such that a state $t \in 2^Q$ has colour c iff the expression $\exists q_0 \dots q_{|Q|-1}. \mathcal{C}_c \wedge t$ is true. Recall, that the colouring function refers to the lowest indices e and m such that $S_i = \emptyset$ and $m_i = 1$, respectively:

$$\Gamma_{i=e} \triangleq \Gamma_{i \geq e} \wedge \neg\Gamma_{i-1 \geq e} \quad \Gamma_{i \geq m} \triangleq \bigvee_{j=0}^i m_j \quad \Gamma_{i=m} \triangleq \Gamma_{i \geq m} \wedge \neg\Gamma_{i-1 \geq m}$$

The formulas \mathcal{C}_{2k} and \mathcal{C}_{2k-1} are defined as follows (note that \mathcal{C}_1 has an extra term for the special case $e = 0$):

$$\begin{aligned} \mathcal{C}_1 &\triangleq \Gamma_{1=e} \wedge \neg\Gamma_{0 \geq m} \vee \Gamma_{0=e} \\ \mathcal{C}_{2k-1} &\triangleq \Gamma_{k=e} \wedge \neg\Gamma_{k-1 \geq m} \quad \text{with } 2 \leq k \leq n+1 \\ \mathcal{C}_{2k} &\triangleq \Gamma_{k=m} \wedge \neg\Gamma_{k \geq e} \quad \text{with } 0 \leq k \leq n \end{aligned}$$

This completes the symbolic implementation of the determinisation procedure.

When considering the deterministic parity automaton \mathfrak{A}_{X_p} (shown in Figure 2.8a) equivalent to the LTL formula X_p , we found that most often not all pairs (S_i, m_i) are necessary. As suggested in [77], we can construct the parity automaton for a fixed bound $b \leq n+1$ (i.e. use b -tuples instead of $(n+1)$ -tuples) and check if the bound is sufficient: We calculate the set of reachable states and check if any of them requires more than b pairs under some input. More formally, let X be a formula representing the set of reachable states when using b -tuples. The bound is *sufficient* iff the following logical expression is equivalent to \perp (essentially, we are testing the third case in Ψ_i^b):

$$X \wedge \mathcal{R}_P \wedge \neg\Gamma_{b-1 \geq e} \wedge \Gamma_{b-1=E} \wedge \bigvee_{s_k \in \mathcal{F}} \left[\varphi_0^k \wedge \neg \bigvee_{j=0}^{b-1} (q_j^{k'} \wedge m'_j) \right]$$

If the bound is not sufficient, we increase it and try again. It can be shown that the range of the colouring function on reachable states for a sufficient bound b is $\{0, \dots, \min(2b-1, 2n-1)\}$ ³¹.

³¹We have already explained why the colour cannot be greater than $2n-1$. If the bound b is sufficient, then the state set S_b would be empty in all reachable states (that is exactly why we did not increase the bound any further). Hence, we have $e \leq b$ in all reachable states. By considering all three cases in the definition of the colouring function, we see that the maximum colour cannot exceed $2b-1$.

2.4.5 Games

As we have mentioned already, many problems in the area of software verification can be reduced to solving infinite two-player games. In fact, our novel SL[1G] model checking algorithm presented in Chapter 5 relies on such a reduction. Therefore, we give a brief overview of infinite two-player games on directed graphs in this subsection. Please refer to [68], which this subsection is largely based on, for a more detailed introduction.

A game consists of an arena and a winning set³². Intuitively, an *arena* represents the “board” on which a game is played. A formal definition follows.

Definition 2.47 (Arenas). Let V_0 and V_1 be sets of *0-vertices* and *1-vertices*, which are disjoint (i.e. $V_0 \cap V_1 = \emptyset$). Furthermore, let $E \subseteq V \times V$ be an *edge relation*, where $V = V_0 \cup V_1$ is the set of all vertices. Then an *arena* is a triple $\mathcal{A} = (V_0, V_1, E)$.

The set of *successors* of a vertex $v \in V$ is defined as $vE \triangleq \{v' \in V \mid (v, v') \in E\}$. A vertex $\bar{v} \in V$ is a *dead end* iff $\bar{v}E = \emptyset$ (i.e. it has no successors).

The games we are considering have only two players, player 0 and player 1. We often talk about player σ , where $\sigma \in \{0, 1\}$. The other player is referred to as player $\bar{\sigma}$ (i.e. $\bar{\sigma} = 1 - \sigma$).

The game is played as follows: A token is initially placed on some vertex $v \in V$. If it is a 0-vertex ($v \in V_0$), player 0 moves the token to one of its successors vE . Symmetrically, if it is a 1-vertex ($v \in V_1$), player 1 moves the token to one of its successors vE . This pattern is repeated forever or until the current vertex v is a dead end. Such finite and infinite sequences of vertices are called *plays*³³.

Definition 2.48 (Game Plays). Let $\mathcal{A} = (V_0, V_1, E)$ be an arena. Then a *play* in \mathcal{A} is either:

- a non-empty finite path $\pi = v_0v_1 \dots v_n \in V^+$ such that $v_i \in v_{i-1}E$ for all $0 < i \leq n$ and $v_nE = \emptyset$ (*finite play*) or
- an infinite path $\pi = v_0v_1 \dots \in V^\omega$ such that $v_i \in v_{i-1}E$ for all $i > 0$ (*infinite play*).

Having defined an arena and a play, we are ready to define a *game* and the *winner of a play*. Intuitively, player 0 wins a game if he forces player 1 into a dead end or if the infinite play satisfies some winning condition.

Definition 2.49 (Games). Let $\mathcal{A} = (V_0, V_1, E)$ be an arena and $Win \subseteq V^\omega$ a *winning set*. Then the pair $\mathfrak{G} = (\mathcal{A}, Win)$ is a *game*.

Definition 2.50 (Winner of a Play). Let $\mathfrak{G} = (\mathcal{A}, Win)$ be a game and π a play in \mathcal{A} . Then player 0 is the *winner* of π in \mathfrak{G} iff:

- π is a finite play $\pi = v_0v_1 \dots v_n$ where v_n is a 1-vertex ($v_n \in V_1$) and a dead end ($v_nE = \emptyset$) or
- π is an infinite play and $\pi \in Win$.

Otherwise, player 1 is the winner of π in \mathfrak{G} .

Similarly to acceptance conditions for ω -automata (see Definition 2.39), there is a multitude of winning conditions for games with different expressive power and translation complexity. In addition, the conditions differ in the amount of memory that a player needs in order to be able to win. Here we present only the basic winning conditions [68].

Definition 2.51 (Winning Conditions). Let $\mathcal{A} = (V_0, V_1, E)$ be an arena. The winning set Win of the game $\mathfrak{G} = (\mathcal{A}, Win)$ is the set of all infinite plays $\pi \in V^\omega$ in \mathcal{A} such that:

- **Büchi winning condition** $F \subseteq V$ (a set of vertices):

$$\text{Inf}(\pi) \cap F \neq \emptyset$$

³²Notice the analogy with ω -automata (Definition 2.40), which consist of a semi-automaton and an acceptance condition.

³³We also refer to *plays* in the context of SL semantics (see Definition 2.26). It should always be clear from the context which concept we are referring to.

- **Muller winning condition** $\mathcal{F} \subseteq 2^V$ (a set of vertex sets):

$$\text{Inf}(\pi) \in \mathcal{F}$$

- **Rabin winning condition** $\Omega = \{(E_1, F_1), \dots, (E_k, F_k)\}$ where $E_i, F_i \subseteq V$ (a family of pairs of vertex sets):

$$\exists i \in \{1, \dots, k\}. (\text{Inf}(\pi) \cap E_i = \emptyset) \wedge (\text{Inf}(\pi) \cap F_i \neq \emptyset)$$

- **Streett winning condition** $\Omega = \{(E_1, F_1), \dots, (E_k, F_k)\}$ where $E_i, F_i \subseteq V$ (a family of pairs of vertex sets):

$$\forall i \in \{1, \dots, k\}. (\text{Inf}(\pi) \cap E_i \neq \emptyset) \vee (\text{Inf}(\pi) \cap F_i = \emptyset)$$

- **Parity winning condition** $c : V \rightarrow \{0, \dots, k\}$ with $k \in \mathbb{N}$ (a colouring function):

$$\min \{c(v) \mid v \in \text{Inf}(\pi)\} \text{ is even}$$

- **Generalised parity winning condition** $c_1, \dots, c_n : V \rightarrow \{0, \dots, k_i\}$ with $k_i \in \mathbb{N}$ (a family of colouring functions):

$$\forall i \in \{1, \dots, n\}. \min \{c_i(v) \mid v \in \text{Inf}(\pi)\} \text{ is even}$$

Notice the correspondence between winning conditions for games and acceptance conditions for ω -automata (see Definition 2.39). We shall use this observation in Chapter 5 to combine an ω -automaton and an arena into a two-player game (see Definition 5.7).

Given a game \mathfrak{G} , we usually want to find out which states of the underlying arena are winning for each player (i.e. from which states a player can force a winning condition to be true). The set of all such states is referred to as the *winning region*. Moreover, we want to construct the corresponding *winning strategies*. Intuitively, a strategy³⁴ is a function which maps histories of the game (finite prefixes of plays) to moves (next vertices).

Definition 2.52 (Game Strategy). Let $\mathcal{A} = (V_0, V_1, E)$ be an arena and $\sigma \in \{0, 1\}$ a player index. Then a *strategy* for player σ is a partial function $f_\sigma : V^*V_\sigma \rightarrow V$, such that for all $\pi \in V^*$ and $v \in V_\sigma$, if $\pi v \in \text{dom}(f_\sigma)$, then $(v, f_\sigma(\pi v)) \in E$.

A prefix of a play $\pi = v_0v_1 \dots v_n$ is *conform* with f_σ iff for every $0 < i \leq n$, if $v_0v_1 \dots v_{i-1} \in \text{dom}(f_\sigma)$, then $v_i = f_\sigma(v_0v_1 \dots v_{i-1})$. A play is conform with f_σ iff each of its prefixes is conform with f_σ .

Definition 2.53 (Winning Strategies). Let $\mathfrak{G} = (\mathcal{A}, \text{Win})$ be a game, $U \subseteq V$ a set of vertices, and f_σ a strategy for player σ . f_σ is a *winning strategy* for player σ on U iff all plays conform with f_σ starting in U are winning for player σ .

Definition 2.54 (Winning Regions). Let $\mathfrak{G} = (\mathcal{A}, \text{Win})$ be a game. The winning region $W_\sigma(\mathfrak{G}) \subseteq V$ (or W_σ for short) is the set of all vertices $v \in V$, such that player σ has a winning strategy f_σ on $\{v\}$.

Clearly, given a game $\mathfrak{G} = (\mathcal{A}, \text{Win})$ the winning regions $W_0(\mathfrak{G})$ and $W_1(\mathfrak{G})$ are always disjoint. If it is also the case that $W_0(\mathfrak{G})$ and $W_1(\mathfrak{G})$ form a partition of V (i.e. $W_0(\mathfrak{G}) \cup W_1(\mathfrak{G}) = V$), then the game is *determined*. Intuitively, a game is determined iff each vertex of the underlying arena is winning for one of the players. It turns out that all the winning conditions in Definition 2.51 have this property [68].

Definition 2.55 (Determinacy). Let $\mathfrak{G} = (\mathcal{A}, \text{Win})$ be a game with winning regions $W_0(\mathfrak{G})$ and $W_1(\mathfrak{G})$. We say that \mathfrak{G} is *determined* iff $\{W_0(\mathfrak{G}), W_1(\mathfrak{G})\}$ is a partition of V .

Proposition 2.6. All winning conditions in Definition 2.51 are determined.

The definition of a strategy is very general. In particular, it refers to the whole history of the game, which is finite but arbitrarily long. Thus, in theory, a player might need infinite memory to implement such a strategy. Clearly, this is not practical. Therefore, we consider two special types of strategies: *finite memory strategies* (also called *forgetful*) and *memoryless strategies* (also called *positional*).

³⁴Again, we also refer to *strategies* in the context of ATL and SL (see Definitions 2.16 and 2.21). The main difference is that ATL/SL strategies map sequences of states to actions, whereas game strategies map sequences of states to next states. It should be clear from the context which of the concepts we are referring to.

winning condition	player 0	player 1
reachability	yes	yes
Büchi	yes	yes
Müller	no	no
Rabin	yes	no
Streett	no	yes
parity	yes	yes
generalised parity	no	yes

Table 2.4: Memoryless determinacy of various infinite games [28, 68]. “yes” means that the player wins every game with the corresponding winning condition using a memoryless strategy.

Definition 2.56 (Finite Memory Strategies). Let $\mathfrak{G} = (\mathcal{A}, Win)$ be a game. A strategy f_σ is *finite memory* if there exists a finite set M , an element $m_I \in M$, and partial functions $\delta : V \times M \rightarrow M$ and $g : V_\sigma \times M \rightarrow V$ such that the following holds: For each prefix of a play $\pi = v_0v_1 \dots v_n$ and sequence $m_0m_1 \dots m_n$ where (i) $m_0 = m_I$ and (ii) $(v_i, m_i) \in \text{dom}(\delta)$ implies $m_{i+1} = \delta(v_i, m_i)$ for all $0 \leq i < n$, we have either (a) $\pi \notin \text{dom}(f_\sigma)$ and $(v_n, m_n) \notin \text{dom}(g)$, or (b) $\pi \in \text{dom}(f_\sigma)$, $(v_n, m_n) \in \text{dom}(g)$, and $f_\sigma(\pi) = g(v_n, m_n)$.

Intuitively, the set M represents the possible memory states of player σ and m_I is his initial memory state. The strategy is executed by player σ as follows:

Let the current vertex be v_i and the current memory state of player σ m_i .

- Move.** If $v_i \in V_\sigma$ and $(v_i, m_i) \in \text{dom}(g)$, then move to $v_{i+1} = g(v_i, m_i)$. Otherwise, $v_{i+1} \in v_iE$ is arbitrary or selected by player $\bar{\sigma}$.
- Update memory.** If $(v_i, m_i) \in \text{dom}(\delta)$, update the memory state of player σ to $m_{i+1} = \delta(v_i, m_i)$. Otherwise, $m_{i+1} \in M$ is arbitrary. Player σ can now forget m_i . Note that the update is carried out even when the move is decided by player $\bar{\sigma}$.

If we set M to be the singleton set $\{m_I\}$, then the strategy becomes memoryless. Intuitively, when using a memoryless strategy, the next vertex depends only on the current vertex. Hence, we can view it as a partial function $V_\sigma \rightarrow V$.

Definition 2.57 (Memoryless Strategy). Let $\mathfrak{G} = (\mathcal{A}, Win)$ be a game. A strategy f_σ is *memoryless* iff for every prefixes of plays $\pi_1v, \pi_2v \in V^*V_\sigma$, we have either (i) $\pi_1v \notin \text{dom}(f_\sigma)$ and $\pi_2v \notin \text{dom}(f_\sigma)$, or (ii) $\pi_1v \in \text{dom}(f_\sigma)$, $\pi_2v \in \text{dom}(f_\sigma)$, and $f_\sigma(\pi_1v) = f_\sigma(\pi_2v)$.

Finite memory is often sufficient for winning strategies in infinite games. In fact, all games considered in this thesis have this property, referred to as *finite memory determinacy* [68]. Furthermore, some games require no memory at all (so-called *memoryless determinacy*). These include reachability and parity games [68], both of which will be discussed shortly. Table 2.4 shows which players always have a memoryless winning strategy for various winning conditions.

Proposition 2.7. In every game with a winning condition from Definition 2.51, both players win using finite memory strategies.

Proposition 2.8. In every reachability and parity game, both players win using memoryless strategies.

We will now discuss solving reachability games, parity games, and generalised parity games in more detail, as they are relevant for the new model checking algorithm for SL[1G] we propose in Chapter 5.

Reachability Games

We start by considering *reachability games* [68]. We will use them to demonstrate some concepts which will be useful for solving more complex games later. By *solving* we mean finding the winning regions and strategies for both players.

In reachability games, the aim (of player 0) is simply to reach some vertex of a designated set of vertices (or force player 1 into a dead end).

Definition 2.58 (Reachability Game). Let $\mathcal{A} = (V_0, V_1, E)$ be an arena and $X \subseteq V$ a set of vertices. Then a *reachability game* $R(\mathcal{A}, X)$ is a game in which a play π is winning iff (i) some vertex from X occurs in π or (ii) a dead end for player 1 occurs in π .

Note that this differs from all games studied so far because a finite path ending in a dead end for player 0 can still be winning for player 0 as long as some vertex from X is visited. The game is solved by calculating the *attractor* of X for player 0. Intuitively, an attractor of a set of vertices X for player σ is the set of all vertices from which player σ can force a visit of X (regardless of what player $\bar{\sigma}$ does).

Definition 2.59 (σ -Attractor). Let $\mathcal{A} = (V_0, V_1, E)$ be an arena and $X \subseteq V$ a set of vertices. The σ -*attractor* of X in \mathcal{A} is the set of vertices $\text{attr}_\sigma(\mathcal{A}, X) \triangleq \bigcup_{i \geq 0} X^i$ where:

$$\begin{aligned} X^0 &\triangleq X \\ X^{i+1} &\triangleq \{v \in V_\sigma \mid vE \cap X^i \neq \emptyset\} \cup \{v \in V_{\bar{\sigma}} \mid vE \subseteq X^i\} \end{aligned}$$

A (memoryless) *attractor strategy* for player σ is a partial function $f_\sigma : V_\sigma \rightarrow V$ such that $\text{dom}(f_\sigma) = V_\sigma \cap (\text{attr}_\sigma(X) \setminus X)$ and, for each vertex $v \in X^{i+1} \cap V_\sigma$, we have $f_\sigma(v) \in X^i$.

Intuitively, the attractor strategy ensures that the play is moving “towards” X . On the other hand, player σ is “trapped” inside $V \setminus \text{attr}_\sigma(\mathcal{A}, X)$, i.e. player $\bar{\sigma}$ can ensure that the game stays outside $\text{attr}_\sigma(\mathcal{A}, X)$ forever. Such a set of vertices is referred to as a *trap*.

Definition 2.60 (σ -Trap). Let $\mathcal{A} = (V_0, V_1, E)$ be an arena. A set of vertices $Y \subseteq V$ is a σ -*trap* iff (i) $vE \subseteq Y$ for every $v \in Y \cap V_\sigma$ and (ii) $vE \cap Y \neq \emptyset$ for every $v \in Y \cap V_{\bar{\sigma}}$. A (memoryless) *trapping strategy* for player $\bar{\sigma}$ is a partial function $f_{\bar{\sigma}} : V_{\bar{\sigma}} \rightarrow V$ such that $\text{dom}(f_{\bar{\sigma}}) = Y \cap V_{\bar{\sigma}}$ and, for each vertex $v \in Y \cap V_{\bar{\sigma}}$, we have $f_{\bar{\sigma}}(v) \in V_{\bar{\sigma}}$.

Notice that all σ -attractors (including $\text{attr}_\sigma(\mathcal{A}, \emptyset)$) contain all dead ends for player $\bar{\sigma}$ (since $vE = \emptyset \subseteq X^i$). Hence, the winning regions of a reachability game $R(\mathcal{A}, X)$ are:

$$(W_0, W_1) = (\text{attr}_0(\mathcal{A}, X), V \setminus \text{attr}_0(\mathcal{A}, X))$$

The winning strategies for the players are as follows: Player 0 uses an attractor strategy in his winning region $W_0 = \text{attr}_0(\mathcal{A}, X)$, i.e. he ensures that X will be reached eventually. Conversely, player 1 uses a trapping strategy in this winning region $W_1 = V \setminus \text{attr}_0(\mathcal{A}, X)$, through which he ensures that the attractor of X will never be visited. The following proposition follows from the fact that the attractor of any set in a finite arena with n vertices and m vertices can be calculated in time $O(m + n)$ [68]:

Proposition 2.9. Let $\mathcal{A} = (V_0, V_1, E)$ be an arena with $n = |V| = |V_0| + |V_1|$ vertices and $m = |E|$ edges. The reachability game $R(\mathcal{A}, X)$ for an arbitrary set of vertices $X \subseteq V$ can be calculated in time $O(m + n)$.

Parity Games

The parity winning condition is the “most fundamental” winning condition [46]. It has the following desirable properties:

- Every other winning condition for two-player infinite games can be reduced to it [68].
- It enjoys memoryless determinacy (Proposition 2.8).
- It is easily dualisable³⁵. The dual of a parity winning condition c is again a parity winning condition c^D with $c^D(v) \triangleq c(v) + 1$.

Recall that player 0 wins a play π in a parity game $\mathfrak{G} = (\mathcal{A}, c)$ iff either (i) the minimum³⁶ colour occurring infinitely often along π is even, or (ii) π leads to a dead end for player 1. The solution of a parity game is obtained recursively by solving subgames first and then combining the subresults. A subgraph and a subgame are defined as follows.

³⁵Let $\mathfrak{G} = ((V_0, V_1, E), \text{Win})$ be a game. The *dual* of \mathfrak{G} is a game $\mathfrak{G}^D \triangleq ((V_1, V_0, E), V^\omega \setminus \text{Win})$, i.e. the roles of the players are swapped and the winning condition is complemented.

³⁶Equivalently, the parity winning condition is sometimes defined with respect to the *maximum* colour occurring infinitely often.

Definition 2.61 (Subgraph and Subgame). Let $\mathfrak{G} = (\mathcal{A}, c)$ be a parity game with arena $\mathcal{A} = (V_0, V_1, E)$ and $U \subseteq V$ a set of vertices. The *subgraph* of \mathfrak{G} induced by U is defined as $\mathfrak{G}[U] \triangleq (\mathcal{A}|_U, c|_U)$, where $\mathcal{A}|_U \triangleq (V_0 \cup U, V_1 \cup U, E \cup (U \times U))$, $\text{dom}(c|_U) \triangleq U$, and $c|_U(u) \triangleq c(u)$ for all $u \in U$.

A subgraph $\mathfrak{G}[U]$ is a *subgame* of \mathfrak{G} iff every dead end in $\mathfrak{G}[U]$ is also a dead end in \mathfrak{G} .

The algorithm for solving parity games shown in Figure 2.9 is originally due to McNaughton and Zielonka [70, 88]. It returns the winning regions of the game and can also be used to construct the corresponding winning strategies. It proceeds as follows:

1. **Case 1.** If the game has no vertices, the algorithm simply returns empty regions (line 3).
2. The algorithm finds the lowest colour among all vertices p , the associated player σ ($\sigma = 0$ if p is even, $\sigma = 1$ if p is odd), and the set of all such vertices X . Then it calculates their attractor A .
3. The algorithm recursively finds the winning regions (W_0, W_1) of the subgame of \mathfrak{G} with all vertices in A removed (line 9).
4. **Case 2.** If no vertices of the subgame are winning for player $\bar{\sigma}$ (line 10), then we are done. Player σ will either keep winning in the subgame from some point onwards, or player $\bar{\sigma}$ can infinitely often move out of the subgame into the attractor A , in which case player σ can ensure that nodes with priority p are visited infinitely often. Hence, all vertices of the game are winning for player σ . The winning strategy is the combination of the winning strategy for the subgame and the attractor strategy for A .

For strategy synthesis, it is important to ensure that the moves made by player σ in vertices X stay within \mathfrak{G} . Remember that \mathfrak{G} might be a subgame of a larger game that we are trying to solve.

5. Since $V \setminus A$ is a σ -trap in the game (otherwise, player σ could reach A from some vertex $v \in V \setminus A$ and since A is an attractor we would have $v \in A$, which is a contradiction), we have $W_{\bar{\sigma}}(\mathfrak{G}[V \setminus A]) \subseteq W_{\bar{\sigma}}(\mathfrak{G})$. Thus, the algorithm calculates the attractor B of $W_{\bar{\sigma}}$ (line 13). Clearly, B is a subset of the winning region for player $\bar{\sigma}$ in \mathfrak{G} .
6. The algorithm recursively finds the winning regions (W_0, W_1) of the subgame of \mathfrak{G} with all vertices in B removed (line 14).
7. **Case 3.** We have $W_{\sigma}(\mathfrak{G}) = W_{\sigma}(\mathfrak{G}[V \setminus B])$ and $W_{\bar{\sigma}}(\mathfrak{G}) = W_{\sigma}(\mathfrak{G}[V \setminus B]) \cup B$ since $V \setminus B$ is a trap for player $\bar{\sigma}$ and player σ has no incentive to move to the attractor B (because it is winning for player $\bar{\sigma}$). The winning strategy for player σ is the winning strategy for the subgame $\mathfrak{G}[V \setminus B]$. The winning strategy for player $\bar{\sigma}$ is a combination of his winning strategy for the subgame $\mathfrak{G}[V \setminus B]$, the attractor strategy for B , and his winning strategy for the subgame $\mathfrak{G}[V \setminus A]$.

Given a finite parity game, the computational complexity of the algorithm is linear in the number of edges, polynomial in the number of vertices, and exponential in the number of colours [51]:

Proposition 2.10. Let $\mathfrak{G} = (\mathcal{A}, c)$ be a parity game with arena $\mathcal{A} = (V_0, V_1, E)$. Further more, let $n = |V|$ be the number of vertices, $m = |E|$ the number of edges, and $d = |\text{img}(c)|$ the number of colours. The worst-case time complexity of the algorithm in Figure 2.9 is $O(m \times (n/d)^d)$.

Although a number of algorithms for solving parity games with better worst-case complexity have been proposed [51, 79], Zielonka's algorithm beats them in practice [44]. As an interesting side note, the parity game solving problem belongs to the $\text{UP} \cap \text{co-UP}$ complexity class³⁷ [50] and it is unknown whether a polynomial algorithm for solving it exists [54].

Generalised Parity Games

Generalised parity games were introduced in [28]. They refer to games with conjunctions or disjunctions of parity conditions. Here we consider the case that the goal of player 0 is a *conjunction* of parity conditions, i.e. he wants to ensure that for each colouring function, the colour occurring infinitely often is even. We can immediately see that the goal of player 1 is then a *disjunction* of parity conditions, i.e.

³⁷Note that $\text{P} \subseteq \text{UP} \cap \text{co-UP} \subseteq \text{NP} \cap \text{co-NP}$.

```

1 function SOLVEPARITY( $\mathfrak{G}$ )
2   if  $V = \emptyset$  then
3     return  $(\emptyset, \emptyset)$  ▷ Case 1.
4   end if
5    $p := \min \{c(v) \mid v \in V\}$ 
6    $\sigma := p \bmod 2$ 
7    $X := c^{-1}(p)$ 
8    $A := \text{attr}_\sigma(\mathcal{A}, X)$ 
9    $(W_0, W_1) := \text{SOLVEPARITY}(\mathfrak{G}[V \setminus A])$ 
10  if  $W_\sigma = \emptyset$  then ▷ Case 2.
11     $W_\sigma = V$ 
12  else ▷ Case 3.
13     $B := \text{attr}_{\bar{\sigma}}(\mathcal{A}, W_\sigma)$ 
14     $(W_0, W_1) := \text{SOLVEPARITY}(\mathfrak{G}[V \setminus B])$ 
15     $W_\sigma := W_\sigma \cup B$ 
16  end if
17  return  $(W_0, W_1)$ 
18 end function

```

Figure 2.9: Algorithm for solving parity games [79]. It returns the winning regions (W_0, W_1) of a parity game $\mathfrak{G} = (\mathcal{A}, c)$ with arena $\mathcal{A} = (V_0, V_1, E)$ and colouring function $c : V \rightarrow \{0, \dots, k\}$ where $k \in \mathbb{N}$.

he wins if for at least one of the colouring functions, the colour occurring infinitely often is odd. Hence, the dual of a conjunctive parity game is a disjunctive parity game (and vice versa). This *asymmetry* between the two players is also reflected in the fact that player 1 (disjunction) always wins memoryless while player 0 (conjunction) requires finite memory in general.

The classical algorithm for solving generalised parity games is shown in Figure 2.10. Again, it returns the winning regions of the game. The winning strategies are as follows [28]:

- **Case 1.** The game has no vertices so there are no strategies.
- **Case 2.** Player 0 plays according to his winning strategy for the subgame $\mathfrak{G}[V \setminus A]$ when in W'_0 . The winning strategy for player 1 depends on the current vertex v :
 - $v \in W'_1$: winning strategy for player 1 for the subgame $\mathfrak{G}[V \setminus A]$;
 - $v \in c_i^{-1}(p_i + 1) \cap G$: any move as long as the game stays within G ;
 - $v \in \text{attr}_1(\mathcal{A}|_G, c_i^{-1}(p_i + 1) \cap G) \setminus (c_i^{-1}(p_i + 1) \cap G)$: (memoryless) attractor strategy;
 - $v \in W_1$: winning strategy for player 1 for the subgame $\mathfrak{G}[H]$;
 - $v \in A \setminus G$: (memoryless) attractor strategy.
- **Case 3.** All vertices are winning for player 0. In order to satisfy all conditions infinitely often, he needs to use the conjunct number i as memory (i.e. $M = \{1, \dots, n\}$). When his current memory state is i and the current vertex is v , player 0 uses the following strategy:
 - $v \in c_i^{-1}(p_i)$: choose any successor in G and update memory to $i + 1$;
 - $v \in \text{attr}_0(\mathcal{A}, c_i^{-1}(p_i)) \setminus c_i^{-1}(p_i)$: attractor strategy (keep memory set to i);
 - otherwise, the following two cases apply. They refer to the value of W_0 and its attractor calculated on lines 14 and 15 in *some* iteration (not just the last one).
 - * $v \in W_0$: winning strategy for player 0 for the subgame $\mathfrak{G}[H]$;
 - * $v \in \text{attr}_0(\mathcal{A}|_G, W_0) \setminus W_0$: (memoryless) attractor strategy.

Given a finite generalised parity game, the computational complexity of the algorithm is linear in the number of edges, polynomial in the number of vertices, and exponential in the number of colours of each colouring function [28]:

```

1 function SOLVEGENPARITY( $\mathfrak{G}$ )
2   if  $V = \emptyset$  then
3     return  $(\emptyset, \emptyset)$  ▷ Case 1.
4   end if
5   for  $i := 1, \dots, k$  do
6      $p_i := 2 \lfloor (\min \{c_i(v) \mid v \in V\}) / 2 \rfloor$  ▷ Ensure that  $p_i$  is even.
7   end for
8   for  $i := 1, \dots, k$  do
9     if  $p_i < k_i$  then
10       $G := V \setminus \text{attr}_0(\mathcal{A}, c_i^{-1}(p_i))$ 
11       $H := G \setminus \text{attr}_1(\mathcal{A}|_G, c_i^{-1}(p_i + 1) \cap G)$ 
12      repeat
13         $(W_0, W_1) := \text{SOLVEGENPARITY}(\mathfrak{G}[H])$ 
14         $G := G \setminus \text{attr}_0(\mathcal{A}|_G, W_0)$ 
15         $H := G \setminus \text{attr}_1(\mathcal{A}|_G, c_i^{-1}(p_i + 1) \cap G)$ 
16      until  $W_1 = \emptyset \vee W_1 = H$ 
17      if  $W_1 = H$  then
18         $A := \text{attr}_1(\mathcal{A}, G)$ 
19         $(W'_0, W'_1) := \text{SOLVEGENPARITY}(\mathfrak{G}[V \setminus A])$ 
20        return  $(W'_0, W'_1 \cup A)$  ▷ Case 2.
21      end if
22    end if
23  end for
24  return  $(V, \emptyset)$  ▷ Case 3.
25 end function

```

Figure 2.10: Algorithm for solving generalised parity games [28]. It returns the winning regions (W_0, W_1) of a generalised parity game $\mathfrak{G} = (\mathcal{A}, \langle c_1, \dots, c_n \rangle)$ with arena $\mathcal{A} = (V_0, V_1, E)$ and colouring functions $c_i : V \rightarrow \{0, \dots, k_i\}$ where $k_i \in \mathbb{N}$.

Proposition 2.11. Let $\mathfrak{G} = (\mathcal{A}, \langle c_1, \dots, c_k \rangle)$ be a generalised parity game with arena $\mathcal{A} = (V_0, V_1, E)$. Furthermore, let $n = |V|$ be the number of vertices, $m = |E|$ the number of edges, and $d_i = |\text{img}(c_i)|$ the number of colours of colouring function c_i for all $i \in [1..k]$. The worst-case time complexity of the algorithm in Figure 2.10 is:

$$O(m \times n^d) \times \binom{\lceil d/2 \rceil}{\lceil d_1/2 \rceil, \dots, \lceil d_k/2 \rceil} = O(m \times n^d) \times \frac{\lceil d/2 \rceil!}{\prod_{i=1}^k \lceil d_i/2 \rceil!}$$

where $d \triangleq \sum_{i=1}^k d_i$.

2.4.6 Summary

In this section, we gave a brief overview of concepts related to ω -automata and two-player games, which we will use to describe our model checking algorithm for SL[1G] in Chapter 5. We also explained how an arbitrary LTL formula can be converted to an equivalent deterministic parity automaton symbolically. Finally, we discussed how the winning regions and strategies of reachability games, parity games, and generalised parity games can be calculated.

2.5 Existing Tools

This section gives an overview of existing model checkers. We provide their key features, file formats, and supported specification languages. We also indicate whether they use BDD-based or SAT-based techniques (or both). The tools are presented in alphabetical order.

2.5.1 MCK

MCK is a model-checking tool for multi-agent systems developed at the School of Computer Science and Engineering at the University of New South Wales [7,14]. It uses mainly BDDs but also supports bounded model checking (BMC) and explicit state model checking (ESMC). The BDD packages supported by MCK are David Long's Binary Decision BDD package, BuDDy, and CUDD [1,60,85]. The basic logics for specifying properties supported by MCK are:

- LTL, CTL, CTL*;
- μ -calculus;
- Epistemic modalities including common knowledge;
- Fairness constraints.

In addition to the (traditional) observational semantics, MCK supports clock, asynchronous perfect recall, and synchronous perfect recall semantics. It also provides a GUI and counterexample generation. MCK is implemented in Haskell.

2.5.2 MCMAS

MCMAS is a *BDD-based* model-checker for the verification of multi-agent systems (MAS) developed at Imperial College London released under GNU General Public License (GPL) [5,8,63]. It uses the CUDD BDD package [85]. MAS descriptions are given in the form of programs in ISPL (Interpreted Systems Programming Language). MCMAS supports the following logics for specifying properties:

- CTL with fairness constraints;
- ATL with fairness constraints;
- Epistemic modalities;
- Deontic modalities expressing correctness of agents' behaviour³⁸ [66].

³⁸The syntax is extended with the deontic modality $O_i \varphi$. Informally, $O_i \varphi$ is true iff φ is true in all states in which agent $i \in \text{Agt}$ is functioning correctly.

MCMAS also supports generating counterexample traces and provides a GUI in the form of an Eclipse plug-in, whose functionalities include ISPL program editing with dynamic syntax checking, interactive execution mode, and counterexample/witness display [63]. MCMAS is implemented in C++.

As part of this project, we implemented two extensions for MCMAS which add support for SLK and SL[1G] model checking (see Sections 6.2 and 6.3). Therefore, MCMAS functionality, usage, and architecture is described in much more detail in Section 6.1. Nevertheless, we decided to include it in this section as well so that it could be easily compared with other existing tools.

2.5.3 Mocha

Mocha is a BDD-based model checker developed jointly at the University of California at Berkeley, the University of Pennsylvania, and the State University of New York at Stony Brook [17]. It differs from conventional model checkers in that its main purpose is to facilitate the development of new verification techniques. Systems are modelled using reactive modules, which represent their synchronous, asynchronous, and real-time components. Mocha supports the following formalisms for specifying properties [16]:

- ATL;
- Invariants (propositional formulas which are intended to be true in all reachable states³⁹);
- Abstract modules, which should be implemented by the system module (refinement).

Mocha provides a GUI for interactive simulation. It is currently available in two versions, cMocha and jMocha, implemented in C and Java respectively. The latter currently does not support ATL model checking.

2.5.4 NuSMV

NuSMV is an open-source tool for model checking published under GNU Lesser General Public License 2.1 (LGPL) which supports both BDD-based and SAT-based model checking [6, 10, 30]. NuSMV is a reimplement and extension of SMV [69]. It uses the CUDD BDD package and supports the Z-Chaff and MiniSat SAT-solvers [3, 36, 85].

NuSMV processes files in the SMV format, which it first translates to a finite state machine (FSM) and then performs either BDD-based model checking, or bounded model checking (BMC) [32]. Figure 2.11 shows an SMV file and the corresponding finite state machine. The tool supports the following logics for specifying properties [31]:

- LTL using either the tableau method [34], in which it is automatically converted to a CTL formula and verified on a tableau with fairness constraints (BDD-based), or using BMC (SAT-based);
- CTL with fairness constraints (BDD-based);
- RTCTL (Real Time CTL [40]) which augments CTL with real-time constraints⁴⁰ (BDD-based);
- A subset of PSL (Property Specification Language [15]).

NuSMV also supports generating counterexample traces and provides an interactive shell [31]. It is implemented in C.

2.5.5 PRISM

PRISM is a probabilistic model checker released under the GNU General Public License [5, 11, 57]. It uses various techniques including BDDs, MTBDDs⁴¹, discrete-event simulation, quantitative abstraction refinement, and symmetry reduction. PRISM supports a wide range of models:

³⁹Note that checking an invariant φ is equivalent to checking the CTL formula $\text{AG } \varphi$.

⁴⁰For example, the RTCTL formula $\text{AG}(\alpha \rightarrow \text{AF}^{0..4} \beta)$ expresses that whenever α is true, β is true in 0 to 4 steps [31].


⁴¹Multi-terminal binary decision diagrams (MTBDDs) are a generalisation of BDDs in which terminal nodes are labelled with arbitrary real values (not just 0 and 1) [56, Section 4].

```

MODULE main
  VAR
    b0 : boolean;
    b1 : boolean;
    n : {0, 1, 2, 3};
  ASSIGN
    init(b0) := FALSE;
    init(b1) := FALSE;
    init(n) := 0;
    next(b0) := !b0;
    next(b1) := b1 xor (b0 & !next(b0));
    next(n) := (n + 1) mod 4;
  LTLSPEC
    G (b0 xor X b0);
  CTLSPEC
    AG (b0 <-> (n mod 2) = 1);
  CTLSPEC
    AG (b1 <-> (n / 2) = 1);

```

(a) SMV file.



(b) Finite state machine.

Figure 2.11: Sample SMV file (with one LTL property and two CTL properties) and the corresponding finite state machine.

- Discrete deterministic – Discrete-time Markov chains (DTMCs);
- Discrete non-deterministic – Markov decision processes (MDPs) and probabilistic automata (PAs);
- Continuous deterministic – Continuous-time Markov chains (CTMCs);
- Continuous non-deterministic – Probabilistic timed automata (PTAs) and Priced probabilistic timed automata (PPTAs).

and the extensions of these models with costs and rewards [57]. PRISM’s property specification language⁴² subsumes the following logics:

- PCTL (Probabilistic computation tree logic [22]) and PCTL* which augment CTL and CTL* with probability bounds⁴³;
- CSL (Continuous stochastic logic [20]) for continuous systems inspired by CTL⁴⁴;
- LTL;
- Most of CTL.

PRISM also provides optimal adversary/strategy generation for nondeterministic models and a GUI. It is implemented in a mixture of Java and C++.

2.5.6 VerICS

VerICS is a *SAT-based* model-checking tool for multi-agent systems developed at the Institute of Computer Science, Polish Academy of Sciences [52]. It uses bounded model checking (BMC) and unbounded

⁴²For example the property “ $P < 0.1 [F <= 100 \text{ num_errors} > 5]$ ” expresses that “the probability that more than 5 errors occur within the first 100 time units is less than 0.1” [11].

⁴³The syntax is extended with the probabilistic operator P . Intuitively, $P_{\geq p} \varphi$ ($P_{\leq p} \varphi$) means that φ holds with probability at least (at most) p .

⁴⁴Informally, CSL extends propositional logic with formulas of the form $P_{>p} \phi$ where $\phi = (\varphi_1 U_{[a_1, b_1]} \varphi_2 U_{[a_2, b_2]} \dots \varphi_n)$. ϕ is true on a path π iff there exist real numbers t_1, \dots, t_{n-1} such that for each integer $i \in [1..n-1]$ we have $a_i \leq t_i \leq b_i$ and $\forall t' \in [t_{i-1}, t_i]. \pi(t') \models_{\text{CSL}} \varphi_i$ where $t_0 = 0$.

model checking (UMC) to verify temporal, epistemic, and deontic properties on timed automata and timed Petri nets. BMC uses the MiniSat or RSat SAT-solver and UMC uses a modified version of the Z-Chaff SAT-solver [3, 12, 36]. The logics supported by VerICS include:

- CTL_pK (CTL with past and knowledge operators) using UMC;
- ECTLKD (existential fragment of CTL with knowledge and deontic operators) using BMC;
- TECTLK (existential fragment of timed CTL with knowledge operators) using BMC.

VerICS provides a GUI for modelling timed automata and timed Petri nets. It is implemented in Java.

2.6 Summary

In this chapter, we provided the background theory that the rest of this report is based on. We explained that formal verification techniques typically comprise three components, a framework for modelling systems, a specification language, and a verification method. *Model checking*, which is a model-based verification method, refers to the process of determining whether a property P encoded as a formula ϕ_P holds in a system S represented by a model \mathcal{M}_S , i.e. $\mathcal{M}_S \models \phi_P$. We then discussed each of the three components of verification using model checking.

We first compared the most common frameworks for modelling systems and explained why interpreted systems are best suited for modelling multi-agent systems. As a second step, we described several logics of increasing expressiveness for specifying temporal properties of a system and demonstrated the power of SL. We also gave a brief overview of epistemic modalities, which express agents' knowledge. We concluded our tour of specification languages by pointing out the close relationship between expressiveness and model checking complexity.

We then discussed verification methods. We introduced BDDs, which are commonly used for symbolic model checking because of their ability to handle large state spaces efficiently, and explained how they can represent sets and relations. A large portion of the chapter was devoted to the theory of ω -automata and two-player infinite games, which are central to the model checking algorithms of many temporal logics including LTL.

Finally, we gave an overview of state of the art model checkers and compared their functionality and architecture. Before moving on to the next chapter, we would like to point out that none of them supports either ATL^* or SL.

Chapter 3

Fragment Selection

Our aim is to build a model checker for SL. As explained in Subsection 2.2.5, SL is very expressive. Unlike ATL*, which subsumes all other logics presented in Section 2.2, SL can express complex game-theoretic concepts like Nash equilibria. However, as we have already pointed out, such power comes at a cost. In fact, there are several problems with the original SL defined on perfect recall semantics:

1. **Complexity.** The model checking of SL is NONELEMENTARY with respect to the size of the formula. More precisely, model checking an SL sentence φ with alternation number k is k -EXPSpace-HARD and $(k+1)$ -EXPTIME [72]. This is far worse than the 2EXPTIME-COMPLETE model checking complexity of ATL* [18]. This suggests that, given the computational power of today's computers, any model checking algorithm would be able to verify only very simple SL formulas, rendering SL expressiveness in practice very limited. Nevertheless, it is worth pointing out that the complexity is still P-COMPLETE with respect to the size of the model. Hence, it should be possible to check such simple SL formulas against very large models.
2. **Non-behavioural strategies.** SL admits *non-behavioural strategies* where an agent's action might depend on another agent's action in another counterfactual game [75]. Given that an agent clearly cannot predict another agent's behaviour, such strategies are *not synthesisable* in practice. This is a major drawback since we want to know not only *whether* a formula is true or false, but also *why* it is the case, i.e. what strategies should the agents employ in order to either enforce it, or violate it.
3. **Undecidability under incomplete information.** SL is undecidable under perfect recall and incomplete information¹. This means that we cannot augment SL with epistemic modalities, which require incomplete information, unless agents have imperfect recall. Moreover, without incomplete information, the synthesised strategies might not be *uniform*, i.e. an agent might have to perform different actions in two indistinguishable local states. Such strategies are of limited use in models where agents do not have complete knowledge of the system (e.g. interpreted systems discussed in Subsection 2.1.3).

To sum up, model checking the original SL (with complete information) would probably have very low performance and limited application. Therefore, we decided to consider *fragments* of SL, which do not suffer from these limitations. Given the undecidability of SL under perfect recall and incomplete information, we had basically two options²:

1. SL with *imperfect recall* (and incomplete information) and
2. SL with *complete information* (and perfect recall).

We considered both options and identified two fragments of SL, namely *Epistemic Strategy Logic* (SLK) and *One-Goal Strategy Logic* (SL[1G]), for which we designed model checking algorithms presented in Chapters 4 and 5.

¹A proof of undecidability of ATL under incomplete information and perfect recall was presented in [35]. Since SL (strictly) subsumes ATL, our claim follows.

²SL with *complete* information and imperfect recall is a subset of SL with *incomplete* information and imperfect recall. Therefore, we do not discuss this option separately.

3.1 Imperfect Recall

The first way to tackle the issues outlined at the beginning of this chapter is to use imperfect recall semantics with incomplete information, under which agents have no memory of the past and do not have complete knowledge of the global state. Hence, their actions depend solely on their current local states, i.e. they use memoryless strategies $f_i : L_{iE} \rightarrow Act_i$. This set up, also referred to as *memoryless semantics*, addresses most of the main problems associated with full SL:

1. The model checking complexity with respect to the size of the formula is PSPACE (see Theorem 4.1). This is a *massive improvement* compared to SL with perfect recall. However, the model checking complexity with respect to the size of the model is now also PSPACE.
2. SL with imperfect recall still admits non-behavioural strategies. Moreover, it appears that under incomplete information, behavioural semantics cannot be obtained even by restricting the syntax of SL as we will do in Section 3.2.
3. Incomplete information is decidable under imperfect recall. This should not come as a surprise given that the number of local states, actions, and hence memoryless strategies is finite. Intuitively, we could explicitly enumerate all possible strategies each time a quantifier is encountered.

Although SL with imperfect recall admits non-behavioural strategies, we believe that it has great theoretical value as it supports reasoning about game-theoretic concepts under incomplete information. Furthermore, we can increase its expressiveness by adding epistemic modalities (see Subsection 2.2.6), which turn out to have no impact on either complexity or decidability. We refer to this new variant of SL with imperfect recall, incomplete information, and epistemic modalities as *Epistemic Strategy Logic*, or Strategy Logic with Knowledge (SLK). We formally define it and provide a model checking algorithm for it in Chapter 4.

3.2 Complete Information

Another approach for addressing the undecidability problem is to use complete information semantics, under which agents have complete knowledge of the whole model. Their actions thus depend on the whole history of global states, i.e. they use strategies³ $f_i : Trk \rightarrow Act_i$. This set up was used by the original SL, introduced in [76], and it still suffers from the first two problems, namely NONELEMENTARYSPACE-HARD model checking complexity and non-behavioural strategies.

To address this issue, several *syntactic fragments* of SL have been proposed [72, 75]. The fragments use LTL syntax (see Definition 2.7) augmented with the quantification rule $\wp\psi$ where \wp is a quantification prefix⁴ for all free variables in ψ , ψ is an SL formula consisting of agent-closed *goals* of the form $b\varphi$, b is a binding prefix, and φ is an LTL formula. The syntactic fragments differ in the combinations of goals they permit. For example, the following is a Boolean-Goal Strategy Logic formula:

$$\underbrace{\llbracket x \rrbracket \langle\langle y \rangle\rangle \llbracket z \rrbracket}_{\text{quant. prefix } \wp} \underbrace{\left[(a, x)(b, y)(c, z) \mathbf{X} p \vee (a, y)(b, z)(c, x) (\mathbf{F} \mathbf{G} q \rightarrow \mathbf{F} \mathbf{G} r) \right]}_{\text{goal } b_1 \varphi_1} \underbrace{\left[(a, y)(b, z)(c, x) (\mathbf{F} \mathbf{G} q \rightarrow \mathbf{F} \mathbf{G} r) \right]}_{\text{goal } b_2 \varphi_2}$$

binding prefix b_1
LTL formula φ_2

Informally, the fragments differ in which operators can occur “between” a quantification prefix \wp and goals $b\varphi$. The following syntactic fragments have been proposed in [72, 75] (k is alternation number of the formula):

- **Nested-Goal Strategy Logic** (SL_[NG]).

Syntax: $\psi ::= \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \mathbf{X}\psi \mid \psi \mathbf{U}\psi \mid b\psi \mid \varphi$ (any LTL operator or bind. prefix)
 Complexity: NONELEMENTARY (k -EXPSpace-HARD lower b., $(k + 1)$ -EXPTIME upper b.)
 Semantics: non-behavioural
 Satisfiability: undecidable

³Recall that $Trk \subseteq G^+$ is the set of all possible *tracks* (non-empty finite sequences of global states) in the underlying interpreted system (see Definition 2.6).

⁴We define the quantification and binding prefixes formally when we introduce SL[1G] syntax (see Definition 5.1).

- **Boolean-Goal Strategy Logic** (SL[BG]).
 - Syntax: $\psi ::= \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid b\varphi$ (any Boolean operator)
 - Complexity: open question (2EXPTIME-HARD lower b., $(k + 1)$ -EXPTIME upper b.)
 - Semantics: non-behavioural
 - Satisfiability: undecidable
- **Conjunctive-Goal Strategy Logic** (SL[CG]).
 - Syntax: $\psi ::= \psi \wedge \psi \mid b\varphi$ (conjunction only)
 - Complexity: 2EXPTIME-COMPLETE
 - Semantics: behavioural
 - Satisfiability: open problem
- **Disjunctive-Goal Strategy Logic** (SL[DG]).
 - Syntax: $\psi ::= \psi \vee \psi \mid b\varphi$ (disjunction only)
 - Complexity: 2EXPTIME-COMPLETE
 - Semantics: behavioural
 - Satisfiability: open problem
- **One-Goal Strategy Logic** (SL[1G]).
 - Syntax: $\psi ::= b\varphi$ (quantification coupled with binding)
 - Complexity: 2EXPTIME-COMPLETE
 - Semantics: behavioural
 - Satisfiability: decidable

Note that SL[DG] and SL[CG] are duals of each other. Therefore, we will refer to them jointly⁵ as SL[DG/CG]. The complete expressiveness chain of all syntactic fragments, the original SL, and ATL* together with their model checking complexities is [72, 75]:

$$\underbrace{\text{ATL}^* < \text{SL}[1\text{G}] < \text{SL}[\text{DG}/\text{CG}]}_{2\text{EXPTIME-COMPLETE}} < \underbrace{\text{SL}[\text{BG}]}_{?} \leq \underbrace{\text{SL}[\text{NG}] \leq \text{SL}}_{\text{NONELEMENTARY-SPACE-HARD}}$$

Clearly, there are two candidates for our purposes: SL[1G] and SL[DG/CG]. Both of them have a much lower model checking complexity than SL and admit only behavioural strategies. In addition, they have the same model checking complexity as ATL* while being strictly more expressive. Although there is a difference between SL[1G] and SL[DG/CG], namely that the satisfiability of SL[1G] is known to be decidable but still remains an open problem for SL[DG/CG], it is not relevant for our purposes and thus not a strong criterion for fragment selection.

Given the amount of time for the project and the theoretical complexity of model checking and strategy synthesis under perfect recall, we decided to focus on SL[1G]. We provide a novel model checking algorithm for SL[1G] with perfect recall and complete information in Chapter 5. It is perfectly possible that only minor modifications of the algorithm would be necessary in order to support SL[DG/CG] as well.

3.3 Toy Model

Some of the model checking concepts presented in Chapters 4 and 5 are very abstract and might be difficult to grasp on first reading. In order to make them easier to understand for the reader, we will demonstrate them on a very simple toy model, a game of *Rock-Paper-Scissors*. It is a traditional two-player game, where both players simultaneously show a gesture (using their hands) representing one of the three objects. This pattern is repeated, until each of the players shows a different gesture. Once their gestures differ, the winner is determined according to the following rules: rock beats scissors, scissors beat paper, and paper beats rock.

The simplicity of this game will allow us to provide detailed examples for almost all concepts and procedures introduced in this thesis. However, it is very important to stress that *this toy model is in*

⁵This does *not* mean that both disjunctions and conjunctions can be used within one prefix. Such a fragment would be equivalent to SL[BG] since $\neg b\varphi \equiv b\neg\varphi$ (i.e. we could propagate all negations using De Morgan's laws into the goals).

no way representative of the expressive power of SL or its fragments. Unfortunately, if we used a more realistic example, some of the constructions could not be fully presented in this thesis as they would have hundreds or even thousands of states. In fact, as we will see in Chapter 5, even this simple model (with only 3 states) generates a parity game with more than 50 states when model checking the SL[1G] formula Xp (see Figure 5.2). More compelling examples, including specifications expressing Nash equilibria and agents' knowledge, will be presented in Section 6.4.

3.3.1 Formal Definition

Let us now define the model more formally as an interpreted system $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ (see Definition 2.5). To model the two-player game, we use three agents $Agt = \{1, 2, E\}$, where E represents the environment. The internal states L_i and actions Act_i of the agents are as follows:

$$\begin{array}{ll} L_1 \triangleq \{s\} & Act_1 \triangleq \{r, p, s, i\} \\ L_2 \triangleq \{s\} & Act_2 \triangleq \{r, p, s, i\} \\ L_E \triangleq \{s_g, s_1, s_2\} & Act_E \triangleq \{i\} \end{array}$$

Both agents have only one internal state s , i.e. all state information is kept in the environment (and hence agents have complete information). Conversely, the environment can only perform the idle action i . The environment's states s_g, s_1 , and s_2 stand for "game", "player 1 victory", and "player 2 victory" respectively. There are 3 global states⁶ $G \triangleq L_1 \times L_2 \times L_E = \{(s, s, s_g), (s, s, s_1), (s, s, s_2)\}$, which we will denote g_g, g_1 , and g_2 for conciseness. The set of initial states is $I \triangleq \{g_g\}$. The protocols of the agents are as follows:

$$\begin{array}{lll} P_1(s, s_g) \triangleq \{r, p, s\} & P_1(s, s_1) \triangleq \{i\} & P_1(s, s_2) \triangleq \{i\} \\ P_2(s, s_g) \triangleq \{r, p, s\} & P_2(s, s_2) \triangleq \{i\} & P_2(s, s_2) \triangleq \{i\} \\ P_E(s_g) \triangleq \{i\} & P_E(s_1) \triangleq \{i\} & P_E(s_2) \triangleq \{i\} \end{array}$$

Furthermore, there are 16 joint actions, only 10 of which are possible:

$$Act \triangleq Act_1 \times Act_2 \times Act_E = \{(r, r, i), (r, p, i), (r, s, i), \cancel{(r, i, i)}, (p, r, i), (p, p, i), (p, s, i), \cancel{(p, i, i)}, (s, r, i), (s, p, i), (s, s, i), \cancel{(s, i, i)}, \cancel{(i, r, i)}, \cancel{(i, p, i)}, \cancel{(i, s, i)}, (i, i, i)\}$$

To avoid cluttered notation, we will denote them rr, rp, \dots, ii (using only the actions of the two original players). The transition functions of agents 1 and 2 are defined as $t_i(s, a) \triangleq s$ for all joint actions $a \in Act$. The environment evolution function is defined as follows:

$$\begin{array}{llll} t_E(s_g, rr) \triangleq s_g & t_E(s_g, pr) \triangleq s_1 & t_E(s_g, sr) \triangleq s_2 & t_E(s_1, ii) \triangleq s_1 \\ t_E(s_g, rp) \triangleq s_2 & t_E(s_g, pp) \triangleq s_g & t_E(s_g, sp) \triangleq s_1 & t_E(s_2, ii) \triangleq s_2 \\ t_E(s_g, rs) \triangleq s_1 & t_E(s_g, ps) \triangleq s_2 & t_E(s_g, ss) \triangleq s_g & \end{array}$$

Finally, we introduce two propositional variables p_1 and p_2 with $h(p_1) \triangleq \{g_1\}$ and $h(p_2) \triangleq \{g_2\}$. The complete interpreted system is shown in Figure 3.1.

3.3.2 Symbolic Implementation

Here we describe how the interpreted system for the toy problem formalised in Subsection 3.3.1 can be encoded symbolically (see Subsection 2.3.2 for a brief overview of symbolic model checking). The approach we describe here is used by the MCMAS model checker (see Subsection 2.5.2).

We start by describing how to encode internal states and actions. A set of internal states L_i can be encoded using $\log_2 \lceil |L_i| \rceil$ variables. Since both L_1 and L_2 contain only one element, we do not need any

⁶In Definition 2.5, we use G to denote the set of *reachable* states, which is a subset of the set of global states, i.e. $G \subseteq L_1 \times \dots \times L_n \times L_E$. However, it turns out that the set of reachable states is equal to the set of global states in this model.




Figure 3.1: Representation of the toy model as an interpreted system.

variables to encode them. L_E can be represented using $\log_2 [|L_E|] = \log_2 [3] = 2$ Boolean variables. Thus, the global states G can be encoded using a Boolean vector $\bar{v} = [v_0, v_1]$:

$$g_g(\bar{v}) = \neg v_0 \wedge \neg v_1 \quad g_1(\bar{v}) = v_0 \wedge \neg v_1 \quad g_2(\bar{v}) = \neg v_0 \wedge v_1$$

In addition, we need 2 Boolean variables v'_0 and v'_1 to represent the next state in the transition relation. Since g_g is the only initial state, we have $I(\bar{v}) = \neg v_0 \wedge \neg v_1$. Actions are encoded in a very similar manner. Hence, we need no variables to represent the action of the environment and each of the two original players requires 2 Boolean variables to represent their 4 possible actions. A joint action can thus be encoded using a Boolean vector $\bar{w} = [w_0, w_1, w_2, w_3]$. For example, the encoding of the joint action sp is:

$$\text{sp}(\bar{w}) = \underbrace{\neg w_0 \wedge w_1}_s \wedge \underbrace{w_2 \wedge \neg w_3}_p$$

In general, a protocol P_i and an evolution function t_i are encoded as follows [65]:

$$P_i(\bar{v}, \bar{w}) = \bigvee_{l_{iE} \in L_{iE}} \left[l_{iE}(\bar{v}_{iE}) \wedge \bigvee_{a_i \in Act_i} a_i(\bar{w}_i) \right]$$

$$t_i(\bar{v}, \bar{w}, \bar{v}') = \bigvee_{l_{iE} \in L_{iE}} \bigvee_{a \in Act} [l_{iE}(\bar{v}_i) \wedge a(\bar{w}) \wedge t_i(l_{iE}, a)(\bar{v}')]]$$

For example, the encoding of the protocol P_1 is:

$$P_1(\bar{v}, \bar{w}) = \underbrace{\neg v_0 \wedge \neg v_1 \wedge (\neg w_0 \wedge \neg w_1 \vee w_0 \wedge \neg w_1 \vee \neg w_0 \wedge w_1)}_{P_1(s, s_g) = \{r, p, s\}} \vee \underbrace{v_0 \wedge \neg v_1 \wedge w_0 \wedge w_1}_{P_1(s, s_1) = \{i\}} \vee \underbrace{v_0 \wedge \neg v_1 \wedge w_0 \wedge w_1}_{P_1(s, s_2) = \{i\}}$$

If we define the global protocol and evolution function as $P(\bar{v}, \bar{w}) = \bigwedge_{i \in Agt} P_i(\bar{v}, \bar{w})$ and $t(\bar{v}, \bar{w}, \bar{v}') = \bigwedge_{i \in Agt} t_i(\bar{v}, \bar{w}, \bar{v}')$ respectively, we can derive an expression for the binary transition relation:

$$R(\bar{v}, \bar{v}') = \exists \bar{w}. P(\bar{v}, \bar{w}) \wedge t(\bar{v}, \bar{w}, \bar{v}')$$

Finally, we can calculate the fixpoint for reachable states⁷ $G = \text{lfp}_Q [I \cup \text{succ}_{\exists}(Q)]$ symbolically as:

$$G(\bar{v}) = \text{lfp}_{\Theta} [I(\bar{v}) \vee \underbrace{\exists \bar{v}'. (\Theta' \wedge R(\bar{v}', \bar{v}))}_{\text{succ}_{\exists}(Q)}]$$

where Θ' is a Boolean formula equal to Θ with variables in \bar{v} and \bar{v}' swapped.

This completes the symbolic implementation of the toy model. ISPL code⁸ for the model is provided in Appendix A.

⁷ succ_{\exists} is the opposite of pre_{\exists} (see Definition 2.11). It calculates the set of *successors* of a set of global states $X \subseteq G$ in an interpreted system \mathcal{I} : $\text{succ}_{\exists}(X) \triangleq \{g \in G \mid \exists g' \in X \exists a \in Act. t(g', a) = g\}$.

⁸As explained in Subsection 2.5.2, ISPL is a language for modelling multi-agent systems supported by MCMAS. The ISPL syntax is described in Subsection 6.1.2.

3.4 Summary

In this short chapter, we discussed the problems associated with model checking full SL, namely NONELEMENTARYSPACE-HARD model checking complexity, non-behavioural strategies, and undecidability under incomplete information. In order to tackle these issues, we identified two *fragments* of SL, Epistemic Strategy Logic (SLK) and One-Goal Strategy Logic (SL[1G]), which we will further investigate in Chapters 4 and 5 respectively.

Furthermore, we provided a very simple *toy model* of the game of Rock-Paper-Scissors, which we will use throughout the rest of this report to demonstrate new concepts. We encoded it as an interpreted system and explained how it can be implemented symbolically.

Chapter 4

Epistemic Strategy Logic

In this chapter, we introduce *Epistemic Strategy Logic*, or Strategy Logic with Knowledge (SLK), a modification of SL (see Subsection 2.2.5) which is defined on imperfect recall semantics with incomplete information (i.e. agents have no memory of the past and do not have a complete knowledge of the global state of the system). We show that the corresponding model checking problem belongs to the PSPACE complexity class. We also provide an efficient model checking algorithm for it and prove its correctness. Both the logic and the algorithm constitute original material developed as part of this individual project. This chapter is split into two parts: Section 4.1 introduces the logic and Section 4.2 describes the model checking algorithm.

4.1 Logic

SLK is a logic which combines the ability of SL to express game-theoretic concepts with the epistemic framework for describing agents' knowledge. It is a unique blend of three well-established formalisms: (i) LTL operators X, F, G, and U, (ii) SL strategy quantifiers $\langle\langle x \rangle\rangle$, $\llbracket x \rrbracket$ and agent binding (a, x) , and (iii) epistemic modalities K_i , E_A , D_A , and C_A . This structure allows us to reason about the temporal, behavioural, and epistemic aspects of a model separately. As explained in Section 3.1, the logic is defined on imperfect recall semantics with incomplete information in order for the model checking problem to be decidable.

This section provides formal definitions of SLK syntax, semantics, and other related concepts. The definitions will provide us with a solid foundation for the development of the model checking algorithm presented in the next section. We will also discuss several examples and demonstrate some limitations of SLK.

4.1.1 Syntax

The syntax of SLK extends the syntax of SL (Definition 2.19) with epistemic modalities representing individual knowledge K_i , group knowledge E_A , distributed knowledge D_A , and common knowledge C_A , where $i \in \text{Agt}$ is an agent and $A \subseteq \text{Agt}$ a set of agents (see Subsection 2.2.6 for a brief introduction to epistemic modalities).

Definition 4.1 (SLK Syntax). SLK *formulas* are built inductively using the following grammar, where $p \in AP$ is an atomic proposition, $x \in Var$ a variable, $i \in \text{Agt}$ an agent, and $A \subseteq \text{Agt}$ a set of agents:

$$\begin{aligned} \varphi ::= & p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid X\varphi \mid F\varphi \mid G\varphi \mid \varphi U \varphi \mid \\ & \langle\langle x \rangle\rangle\varphi \mid \llbracket x \rrbracket\varphi \mid (i, x)\varphi \mid K_i\varphi \mid E_A\varphi \mid D_A\varphi \mid C_A\varphi \end{aligned}$$

where epistemic operators are applied to sentences only, i.e. $\text{free}(\varphi) = \emptyset$ in $K_i\varphi$, $E_i\varphi$, $D_A\varphi$, and $C_A\varphi$. SLK denotes the infinite set of formulas generated by the above rules.

We use the set of free agents/variables as in Definition 2.20 with an extra rule: $\text{free}(K_i\varphi) \triangleq \text{free}(E_A\varphi) \triangleq \text{free}(D_A\varphi) \triangleq \text{free}(C_A\varphi) \triangleq \emptyset$, for all $i \in \text{Agt}$ and $A \subseteq \text{Agt}$. The reason for this is that φ

must be a sentence according to the definition above. For example, given the set of agents $Agt = \{E, a, b\}$ and an agent $i \in Agt$, $\llbracket e \rrbracket \langle x \rangle \llbracket y \rrbracket (E, e)(a, x)(b, y) K_i X p$ and $\llbracket e \rrbracket \langle x \rangle \llbracket y \rrbracket K_i (E, e)(a, x)(b, y) X p$ are not well-formed SLK formulas since $\text{free}(X p) = \{E, a, b\}$ and $\text{free}((E, e)(a, x)(b, y) X p) = \{e, x, y\}$. On the other hand, $\llbracket e \rrbracket \langle x \rangle \llbracket y \rrbracket (E, e)(a, x)(b, y) X K_i p$ and $K_i \llbracket e \rrbracket \langle x \rangle \llbracket y \rrbracket (E, e)(a, x)(b, y) X p$ are well-formed because both p and $\llbracket e \rrbracket \langle x \rangle \llbracket y \rrbracket (E, e)(a, x)(b, y) X p$ are sentences, i.e. $\text{free}(p) = \text{free}(\llbracket e \rrbracket \langle x \rangle \llbracket y \rrbracket (E, e)(a, x)(b, y) X p) = \emptyset$.

Observe that the syntax of SLK strictly subsumes all logics in the ATL* hierarchy with imperfect recall and incomplete information including CTLK and ATLK. Moreover, it allows us to express properties which could not be expressed by any of the previous logics (including the original SL). For example, the SLK formula:

$$E_{\{a,b\}} \llbracket e \rrbracket \llbracket x \rrbracket (E, e)(a, x)(b, x) [F G p \wedge F G \neg p]$$

expresses that “the agents a and b both know that if they use the same strategy, p will be infinitely often true and infinitely often false”. This formula demonstrates that the large expressive power of SLK comes from the combination of temporal, strategic, and epistemic modalities.

4.1.2 Basic Concepts

Before defining the semantics of SLK, we have to introduce some basic concepts similar to the ones introduced in Subsection 2.2.5. Essentially, we have to cater for the fact that agents now have incomplete information and imperfect recall.

Let us consider agents’ strategies. Since agents now have imperfect recall (no memory), their next actions depend on their current state only. Furthermore, in order for a strategy to be executable by an agent with incomplete information, it must assign the same action to all global states that the agent cannot distinguish, i.e. when $g_1, g_2 \in G$ are two global states such that $g_1 \sim_i g_2$, it must be the case that the strategy of agent $i \in Agt$ satisfies $f_i(g_1) = f_i(g_2)$. This property is sometimes referred to as Γ -uniformity [65]. Since SLK strategies can be shared among multiple agents, it must apply to all of them. Hence, an SLK strategy must (i) comply with the protocols of all agents that share it and (ii) map any two global states that cannot be distinguished by any of the agents to the same action. A formal definition follows.

Definition 4.2 (Uniform Shared Memoryless Strategies). A *uniform shared memoryless strategy* of a non-empty set of agents $A \subseteq Agt$ in an interpreted system \mathcal{I} is a function $f : G \rightarrow Act_A$ such that the following two conditions hold:

1. **Shared.** For all global states $g \in G$ and agents $i \in A$, $f(g) \in P_i(l_{iE}(g))$, i.e. agents comply with their protocols.
2. **Uniform.** For all global states $g_1, g_2 \in G$ and agents $i \in A$, if $g_1 \sim_i g_2$, then we have $f(g_1) = f(g_2)$, i.e. the strategy is a mapping on local states $f : L_{iE} \rightarrow Act_i$.

The set of all uniform shared memoryless strategies of the set of agents A satisfying the above properties is denoted $UStr_A$. The set of all possible shared strategies is defined as $UStr \triangleq \bigcup_{A \subseteq Agt}^{A \neq \emptyset} UStr_A$.

We augment the definition of sharing (Definition 2.22) with an extra case for the epistemic modalities: $\text{sharing}(K_i \varphi) \triangleq \text{sharing}(E_A \varphi) \triangleq \text{sharing}(D_A \varphi) \triangleq \text{sharing}(C_A \varphi) \triangleq \text{sharing}(\varphi)$ for all $i \in Agt$ and $A \subseteq Agt$. The definitions of *assignment* and *play* in SLK are analogous to the ones for SL (Definitions 2.23 and 2.26). The only difference is that they refer to uniform shared *memoryless* strategies and thus do not require translation (Definitions 2.24 and 2.25). We will omit them for conciseness.

4.1.3 Semantics

The semantics of SLK is defined with respect to interpreted systems (see Definition 2.5). For an interpreted system \mathcal{I} , a global state $g \in G$, and an assignment χ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$, we write $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi$ to indicate that the formula $\varphi \in \text{SLK}$ holds at g in \mathcal{I} under χ . A formal definition follows.

Definition 4.3 (SLK Semantics). Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an interpreted system and AP a set of propositional formulas. For all SLK formulas $\varphi \in \text{SLK}$, global states $g \in G$, and assignments $\chi \in \text{Asg}$ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$, the modelling relation $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi$ is inductively defined as follows:

1. $\mathcal{I}, \chi, g \models_{\text{SLK}} \top$.
2. $\mathcal{I}, \chi, g \models_{\text{SLK}} p$ iff $g \in h(p)$, with $p \in AP$.
3. For all formulas $\varphi, \varphi_1, \varphi_2 \in \text{SLK}$, it holds that:
 - (a) $\mathcal{I}, \chi, g \models_{\text{SLK}} \neg\varphi$ iff $\mathcal{I}, \chi, g \not\models_{\text{SLK}} \varphi$;
 - (b) $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_1 \wedge \varphi_2$ iff $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_1$ and $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_2$.
 - (c) $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_1 \vee \varphi_2$ iff $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_1$ or $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_2$.
4. For a variable $x \in \text{Var}$ and a formula $\varphi \in \text{SLK}$, it holds that:
 - (a) $\mathcal{I}, \chi, g \models_{\text{SLK}} \langle\langle x \rangle\rangle\varphi$ iff there exists a strategy $f \in \text{UStr}_{\text{sharing}(\varphi, x)}$ such that $\mathcal{I}, \chi[x \mapsto f], g \models_{\text{SLK}} \varphi$;
 - (b) $\mathcal{I}, \chi, g \models_{\text{SLK}} \llbracket x \rrbracket\varphi$ iff for all strategies $f \in \text{UStr}_{\text{sharing}(\varphi, x)}$ it holds that $\mathcal{I}, \chi[x \mapsto f], g \models_{\text{SLK}} \varphi$.
5. For an agent $i \in \text{Agt}$, a variable $x \in \text{Var}$, and a formula $\varphi \in \text{SLK}$, it holds that $\mathcal{I}, \chi, g \models_{\text{SLK}} (i, x)\varphi$ iff $\mathcal{I}, \chi[i \mapsto \chi(x)], g \models_{\text{SLK}} \varphi$.
6. Let $\pi = \text{play}(g, \chi)$. For all formulas $\varphi, \varphi_1, \varphi_2 \in \text{SLK}$, it holds that:
 - (a) $\mathcal{I}, \chi, g \models_{\text{SLK}} \mathbf{X}\varphi$ iff $\mathcal{I}, \chi, \pi(1) \models_{\text{SLK}} \varphi$;
 - (b) $\mathcal{I}, \chi, g \models_{\text{SLK}} \mathbf{F}\varphi$ iff there exists $i \geq 0$ such that $\mathcal{I}, \chi, \pi(i) \models_{\text{SLK}} \varphi$;
 - (c) $\mathcal{I}, \chi, g \models_{\text{SLK}} \mathbf{G}\varphi$ iff for all $i \geq 0$ we have $\mathcal{I}, \chi, \pi(i) \models_{\text{SLK}} \varphi$;
 - (d) $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_1 \mathbf{U} \varphi_2$ iff there is an index $i \in \mathbb{N}$ such that $\mathcal{I}, \chi, \pi(i) \models_{\text{SLK}} \varphi_2$ and, for all indices $j \in \mathbb{N}$ with $0 \leq j < i$, it holds that $\mathcal{I}, \chi, \pi(j) \models_{\text{SLK}} \varphi_1$.
7. For an agent $i \in \text{Agt}$, a set of agents $A \subseteq \text{Agt}$, and a formula $\varphi \in \text{SLK}$, it holds that:
 - (a) $\mathcal{I}, \chi, g \models_{\text{SLK}} \mathbf{K}_i \varphi$ iff for all $g' \in G$ such that $g \sim_i g'$, it holds that $\mathcal{I}, \emptyset, g' \models_{\text{SLK}} \varphi$;
 - (b) $\mathcal{I}, \chi, g \models_{\text{SLK}} \mathbf{E}_A \varphi$ iff for all $g' \in G$ such that $g \sim_A^E g'$, it holds that $\mathcal{I}, \emptyset, g' \models_{\text{SLK}} \varphi$;
 - (c) $\mathcal{I}, \chi, g \models_{\text{SLK}} \mathbf{D}_A \varphi$ iff for all $g' \in G$ such that $g \sim_A^D g'$, it holds that $\mathcal{I}, \emptyset, g' \models_{\text{SLK}} \varphi$;
 - (d) $\mathcal{I}, \chi, g \models_{\text{SLK}} \mathbf{C}_A \varphi$ iff for all $g' \in G$ such that $g \sim_A^C g'$, it holds that $\mathcal{I}, \emptyset, g' \models_{\text{SLK}} \varphi$.

Although it is very similar, the SLK semantics differs from SL semantics (see Definition 2.29) in two aspects: (i) it is defined on uniform shared memoryless strategies and (ii) it supports epistemic operators. Both modifications are consequences of using incomplete information. Notice that we assume that agents are not aware of the strategy assignment. In fact, they are not even aware of their own strategy for the purposes of epistemic accessibility. This assumption is expressed by the empty assignments in the epistemic cases.

In order to complete the description of the SLK semantics, we define satisfiability and a model of an SLK sentence. Intuitively, a formula is satisfiable if it is true in some interpreted system.

Definition 4.4 (SLK Satisfiability). We say that an interpreted system \mathcal{I} is a *model* for an SLK sentence φ , in symbols $\mathcal{I} \models_{\text{SLK}} \varphi$, iff, for all initial global states $g_i \in I$, it holds that $\mathcal{I}, \emptyset, g_i \models_{\text{SLK}} \varphi$. More generally, we say that an interpreted system \mathcal{I} is a *model* for an SLK sentence φ on a global state $g \in G$, in symbols $\mathcal{I}, g \models_{\text{SLK}} \varphi$, iff $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi$. An SLK sentence φ is *satisfiable* iff there is a model for it.

Finally, we define implication and equivalence between SLK formulas.

Definition 4.5 (SLK Implication and Equivalence). Given two SLK 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 $\varphi_1 \Rightarrow \varphi_2$, iff, for all interpreted systems \mathcal{I} , global states $g \in G$ and assignments $\chi \in \text{Asg}$ with $\text{free}(\varphi) \subseteq \text{dom}(\chi)$, it holds that if $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_1$ then $\mathcal{I}, \chi, g \models_{\text{SLK}} \varphi_2$. We say that φ_1 is *equivalent* to φ_2 , in symbols $\varphi_1 \equiv \varphi_2$, iff both $\varphi_1 \Rightarrow \varphi_2$ and $\varphi_2 \Rightarrow \varphi_1$ hold.

We will now consider two sample SLK formulas and explain why they do or do not hold in the toy model in Figure 3.1:

- $\mathcal{I}, \emptyset, g_g \models_{\text{SLK}} \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle (\text{E}, e)(1, x)(2, y) \times p_2$. This formula expresses that “whatever strategies the environment and player 1 use, there is a strategy for player 2, such that he will win in the next round”. Intuitively, this is true. To see why it is formally true, let d be an arbitrary strategy for the environment and f an arbitrary strategy for player 1. Define a strategy h for player 2 as follows:

$$h(g_g) \triangleq \begin{cases} r & \text{if } f(g_g) = s \\ p & \text{if } f(g_g) = r \\ s & \text{if } f(g_g) = p \end{cases}$$

We now need to show that $\mathcal{I}, \chi, g_g \models_{\text{SLK}} \times p_2$ where $\chi = \{(e, d), (x, f), (y, h), (\text{E}, d), (1, f), (2, h)\}$. The resulting play is $\pi = g_g g_2^\omega$ (player 2 always beats player 1 in the first round by the construction of h). The formula $\times p_2$ is true at g_g iff $\mathcal{I}, \chi, \pi(1) \models_{\text{SLK}} p_2$. Since we have $\pi(1) = g_2 \in h(p_2)$ and d and f were arbitrary, we are done.

- $\mathcal{I}, \emptyset, g_g \not\models_{\text{SLK}} \langle\langle y \rangle\rangle \llbracket e \rrbracket \llbracket x \rrbracket (\text{E}, e)(1, x)(2, y) \times p_2$. This formula expresses that “there is a strategy for player 2, such that whatever strategies the environment and player 1 use, he will win in the next round”. Again, our intuition tells us that this formula should be false. To see why this is indeed false, let h an arbitrary strategy for player 2. We pick some strategy d for the environment and define a strategy f for player 1 as follows:

$$f(g_g) \triangleq \begin{cases} r & \text{if } h(g_g) = s \\ p & \text{if } h(g_g) = r \\ s & \text{if } h(g_g) = p \end{cases}$$

We now need to show that $\mathcal{I}, \chi, g_g \not\models_{\text{SLK}} \times p_2$ where $\chi = \{(e, d), (x, f), (y, h), (\text{E}, d), (1, f), (2, h)\}$. The resulting play is $\pi = g_g g_1^\omega$ (player 1 always beats player 2 in the first round by the construction of f). The formula $\times p_2$ is true at g_g iff $\mathcal{I}, \chi, \pi(1) \models_{\text{SLK}} p_2$. Since we have $\pi(1) = g_1 \notin h(p_2)$ and h was arbitrary, we are done.

4.1.4 Comparison with Strategy Logic

Let us now reiterate the main differences between the two logics. SLK differs from SL in the following ways (please refer to Subsection 2.2.5 for more information about the original SL):

1. **Imperfect recall.** Intuitively, SLK agents have no memory of the past and make their decisions based on the current states only. Formally, SLK strategies are mappings from *local* states to actions ($f : L_{iE} \rightarrow Act_i$). Compare this with ATL and SL strategies (see Definitions 2.16 and 2.21), which map non-empty finite sequences of *global* states (tracks) to actions ($f : Trk \rightarrow Act_i$).
2. **Incomplete information.** Agents do not have complete knowledge of the whole system in SLK, i.e. they do not “see” other agents’ variables. This puts the uniformity constraint on strategies (see Definition 4.2) so that they could be executed by agents (see Subsection 4.1.2). As explained in Chapter 3, SL is undecidable under perfect recall semantics with incomplete information.
3. **Epistemic modalities.** SLK supports epistemic modalities expressing agents’ knowledge. The four epistemic modalities supported by SLK are individual knowledge $K_i \varphi$, group knowledge $E_A \varphi$, distributed knowledge $D_A \varphi$, and common knowledge $C_A \varphi$. See Subsection 2.2.6 for more details about epistemic modalities.
4. **Underlying framework.** SLK is based on interpreted systems whereas SL was *originally* defined on concurrent game structures. Note that we redefined SL on interpreted systems in Subsection 2.2.5 so that it would be easier to compare the two logics. As explained in Subsection 2.1.3, interpreted systems are more natural for expressing incomplete information. Both frameworks are defined in Section 2.1.

While epistemic modalities dramatically increase the expressiveness of SLK, imperfect recall puts heavy constraints on agents’ behaviour. The impact of no memory is discussed in the next subsection.

4.1.5 Limitations

As we explained in Section 3.1, SLK solves the undecidability of SL under incomplete information using imperfect recall semantics. Effectively, we are removing agents' memory and forcing them to base their decisions purely on their current local states. While memoryless strategies are easier to reason about and more compact, they are also less powerful and non-behavioural. We will now explain these limitations. For a start, we show that some simple properties may not be achievable using memoryless strategies. Consider the toy model in Figure 3.1 and the following SLK specification:

$$\langle\langle e \rangle\rangle \langle\langle x \rangle\rangle \langle\langle y \rangle\rangle (E, e)(1, x)(2, y)[X (\neg p_1 \wedge \neg p_2) \wedge XX (p_1 \vee p_2)]$$

The rough meaning of the formula is “*There exist strategies for the two players such that they draw first and then one of them wins*”. Clearly, there exist memoryful strategies which satisfy the specification above (starting from the initial state g_g): The agents simply perform the same action in the first round (e.g. rock-rock) and different actions in the second round (e.g. rock-scissors). Formally, the memoryful strategies f_1, f_2 are:

$$\begin{array}{ll} f_1(g_g) \triangleq r & f_2(g_g) \triangleq r \\ f_1(g_g g_g) \triangleq r & f_2(g_g g_g) \triangleq s \end{array}$$

However, there are *no memoryless strategies* that would fulfil the specification. The reason is that since both agents perform an action $a \in \{r, p, s\}$ in the first round, they have to do it again in the next round because their local state has not changed. Hence, they will keep drawing forever.

We can see that memoryless strategies have less power than memoryful strategies. This can be useful in certain scenarios. Consider the situation where we are synthesising the behaviour of a simple hardware device (e.g. a thermostat) with a fixed amount of memory. In this case, we consider each possible memory settings of the device to be one local state and allow arbitrary transitions between them. Informally, we let the agent decide what it wants to remember. It is appropriate to treat the agent as memoryless since its memory is already encoded in its local state space. We then try to synthesise a memoryless strategy for it. If we succeed, then the agent's fixed-size memory is sufficient to achieve the goal. Hence, SLK can be used to answer questions like “*Does agent a have enough memory to enforce φ ?*”.

On the other hand, there are many scenarios when it is clearly inappropriate to assume that the agent has no memory. Consider the following security protocol specification:

$$\langle\langle e \rangle\rangle [x](E, e)(\text{intruder}, x)G \neg K_{\text{intruder}} \text{password}$$

It asserts that there is some strategy e for the environment such that whatever strategy x the intruder uses, they will never know the password. The major flaw with this specification is that it assumes that the intruder has no memory, i.e. it *underestimates* their power. Hence, we can be certain about a “yes” answer only if we know exactly what they are capable of, which is rarely the case. Conversely, if the answer is “no”, it does not necessarily mean that there is no memoryful strategy for the environment that would ensure security. This simple example demonstrates that SLK is not appropriate for the verification of security protocols¹.

As we pointed out in Section 3.1, imperfect recall does not address the problem of non-behavioural strategies, in which an agent's action depends on another agent's action in another counterfactual play. Consider again the toy model in Figure 3.1 and the following SLK specification:

$$\langle\langle e \rangle\rangle (E, e)[y][z]\langle\langle x \rangle\rangle [((1, x)(2, y)X p_1) \leftrightarrow ((1, y)(2, z)X p_2)]$$

This formula is more complicated than the previous ones. Its meaning roughly is as follows: “*For all strategies y, z there is a strategy x , such that x beats y iff z beats y* ”. Given this, it should be clear that this formula is true because we can always set $x = z$. Now the question is, how do we determine the strategy x ? The action of player 1 in $(1, x)(2, y)X p_1$ depends on the action of player 2 in $(1, y)(2, z)X p_2$. In other words, the strategy of player 1 depends on the strategy of player 2 in another counterfactual play, i.e. it is *non-behavioural*. While the formula above might seem quite convoluted, it is precisely this ability to express complicated interdependencies of agents' behaviour why SL was introduced in the first place. Please refer to [75] if you want to know more about behavioural and non-behavioural fragments of SL.

¹Nevertheless, we will do exactly that in Subsection 6.4.1.

4.2 Model Checking

In this section, we focus on model checking SLK. Recall that the model checking problem for SLK is as follows: Given a system S represented by an interpreted system \mathcal{I}_S and a property P expressed as an SLK formula $\varphi_P \in \text{SLK}$, we want to determine whether the formula is true in the model:

$$\mathcal{I}_S \stackrel{?}{\models}_{\text{SLK}} \varphi_P$$

We provide a model checking algorithm for SLK which calculates the set $\|\varphi_P\|_{\mathcal{I}_S}$ of global states of \mathcal{I}_S in which the formula φ_P holds. The problem above is then decided by checking if all initial states I of \mathcal{I}_S satisfy φ_P :

$$I \stackrel{?}{\subseteq} \|\varphi_P\|_{\mathcal{I}_S}$$

We first discuss the complexity of the problem. Then we present an algorithm which admits an efficient symbolic implementation.

4.2.1 Complexity

Before developing an algorithm for model checking SLK, we consider the complexity of the SLK model checking (decision) problem:

Given an interpreted system \mathcal{I} , a global state $g \in G$, and an SLK formula $\varphi \in \text{SLK}$, determine whether $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi$.

We claim that the problem can be solved in a polynomial amount of space (PSPACE) with respect to both the size of the interpreted system² $|\mathcal{I}|$ and the size of the formula $|\varphi|$. The idea is to determine the problem recursively with a function that accepts a global state $g' \in G$ and an assignment $\chi' \in \text{Asg}$ as input. Since both g' and χ' require only a polynomial amount of space and there will be at most $|\varphi|$ nested calls to the function, we obtain the desired complexity.

Theorem 4.1. The SLK model checking problem is PSPACE with respect to both the size of the model $|\mathcal{I}|$ and the size of the formula $|\varphi|$.

Proof. An arbitrary state $g' \in G$ can be encoded using $\lceil \log_2 |G| \rceil$ Boolean variables (how this can be done is explained in Subsection 2.3.2). An arbitrary uniform shared memoryless strategy $f_A : G \rightarrow \text{Act}_A$ for a set of agents $A \subseteq \text{Agt}$ can be stored using $\lceil \log_2 \left[\left| \prod_{i \in A} \text{Act}_i \right|^G \right] \rceil \leq |G| \times \lceil \log_2 |\max_{i \in A} \text{Act}_i| \rceil$ space. Note that we assumed in Definition 2.22 that every variable is quantified at most once within a formula. Therefore, we do not need to remember the full assignment on $\text{Var} \cup \text{Agt}$ as the entries for agents are merely copies of the entries for the variables they are bound to. Hence, an arbitrary assignment $\chi' \in \text{Asg}$ can be stored using at most

$$\underbrace{|\text{vars}(\varphi)| \times (|G| \times \lceil \log_2 |\max_{i \in A} \text{Act}_i| \rceil + 1)}_{\text{partial mapping from variables in } \varphi \text{ to strategies}} + \underbrace{|\text{Agt}| \times \lceil \log_2 |\text{vars}(\varphi)| \rceil}_{\text{“pointers” for agents into the variable assignment}} = O(|\mathcal{I}|^2 \times |\varphi|)$$

space. The semantics of SLK (Definition 4.3) can be easily transformed into a function $\text{CHECK}(\varphi', g', \chi')$. The important cases are (we omit the other cases for conciseness):

- $\text{CHECK}(\langle\langle x \rangle\rangle \varphi'', g', \chi')$. We store the current strategy for the variable x , $s := \chi'(x)$. Then we perform a **for** loop over all possible strategies $\chi'(x) := f_{\text{sharing}(\varphi, x)}$ and call $\text{CHECK}(\varphi'', g', \chi')$. After the loop we restore the current strategy for the variable x , $\chi'(x) := s$. If any of the calls within the loop succeeded, we return “yes”. Otherwise, we return “no”.

The temporary storage s for one strategy and some counter variables for the loop fit within the $O(|\mathcal{I}|^2 \times |\varphi|)$ space bound.

- $\text{CHECK}(\text{X} \varphi'', g', \chi')$. Given g' and χ'' , it is easy to calculate the successor g'' (by calculating $\pi(1)$ of $\pi = \text{play}(\chi'', g')$ as in Definition 2.26). We then call $\text{CHECK}(\varphi'', g'', \chi')$.

$O(|\mathcal{I}|^2 \times |\varphi|)$ space will be sufficient for calculating the successor.

²To be precise, we define $|\mathcal{I}| = |G| + |\text{Agt}| + |\text{Act}|$.

- $\text{CHECK}(\varphi'' \cup \varphi''', g', \chi')$. We introduce a counter i and perform the following (pseudocode):


```

for  $i := 0$  to  $|G|$  do
  if  $\text{CHECK}(\varphi''', g', \chi')$  then
    return “yes”
  else if  $\neg \text{CHECK}(\varphi'', g', \chi')$  then
    return “no”
  end if
   $g' := \text{SUCCESSOR}(g', \chi')$ 
end for
return “no”
      
```

We traverse the path starting from g' as long as φ'' holds. If we reach a state where φ''' holds, the formula $\varphi'' \cup \varphi'''$ holds so we return “yes”. If we do not reach such a state within $|G|$ steps, we must be in a cycle (since there are $|G|$ global states) where φ''' never holds. Therefore, we return “no”.

The extra counter (and perhaps some auxiliary variables) will surely fit within $O(|\mathcal{I}|^2 \times |\varphi|)$ space.

The depth of the call stack will be at most $|\varphi|$. Hence, if we reuse space, $\text{CHECK}(\varphi, g, \emptyset)$ will use at most $O(|\mathcal{I}|^2 \times |\varphi|^2)$ space. Our claim follows. \square

The model checking complexity of ATL^* with imperfect recall is PSPACE-COMplete [23]. Therefore, it is tempting to immediately conclude that the problem for SLK is also PSPACE-COMplete . However, there is an important difference between ATL^* and SLK semantics: ATL^* assigns memoryless strategies only to the existentially quantified agents (see Definition 2.18), whereas SLK assigns memoryless strategies to all agents. Informally, an ATL^* expression $\langle\langle A \rangle\rangle \psi$ means that “There exist memoryless strategies for agents A such that no matter what the other agents do, ψ will be true”. The SLK formula $[\langle\langle x_i \rangle\rangle(i, x_i)]_{i \in A} [[\langle\langle y_i \rangle\rangle(i, y_i)]_{i \in \text{Agt} \setminus A} \varphi$ has a slightly different meaning: “There exist memoryless strategies for agents A , such that for all memoryless strategies of the other agents, ψ will be true”. Intuitively, SLK restricts the universally quantified agents more than ATL^* does³.

Nevertheless, we conjecture that the SLK model checking problem is indeed PSPACE-COMplete . We believe that this could be shown by reducing the problem of evaluating a Quantified Boolean Formula⁴, which is PSPACE-HARD [45], to the SLK model checking problem.

4.2.2 Algorithm

The model checking algorithm SAT_{SLK} for SLK , which calculates the set of global states in which a given formula is true, is a modification of the existing model checking algorithm for ATLK ⁵ used by MCMAS [65, 81]. It differs in two ways:

1. It has an extra parameter, which represents the *binding* of agents to variables. When model checking a formula, we start with an empty binding and augment it whenever an agent binding operator $(i, x)\varphi$ is encountered.
2. Unlike the original algorithm, which merely returns the set of states in which a given formula holds, the modified algorithm returns a set of *extended states*. Intuitively, an extended state is a pair of (i) a global state and (ii) a variable assignment (mapping variables to strategies) subject to which the formula holds in that state.

We will now define the aforementioned concepts of a binding and variable assignment more formally. For simplicity, we will fix Var to always be the set of variables quantified in the SLK formula we are considering (e.g. if the formula to be checked is $\varphi = \langle\langle x \rangle\rangle [[y]](a, x)(b, y)Xp$, then we set $\text{Var} = \text{vars}(\varphi) = \{x, y\}$). This will allow us to define variable assignments as total functions, which will make the theory and proofs much simpler.

³Note that this subtle difference in semantics has no effect on the relationship between $\text{SL}[1G]$ and ATL^* with perfect recall ($\text{SL}[1G]$ strictly subsumes ATL^* [72]). Imposing memoryful strategies on agents does not constrain their behaviour in any way.

⁴A *Quantified Boolean Formula* has the form $Q_0 x_0 \cdots Q_{n-1} x_{n-1}. E(x_0, \dots, x_{n-1})$ where E is a propositional formula over propositional variables x_0, \dots, x_{n-1} and for all $0 \leq i < n$, Q_i is an existential quantifier \exists or a universal quantifier \forall .

⁵A similar model checking algorithm for CTL is shown in Definition 2.12.

Definition 4.6 (Bindings). Let \mathcal{I} be an interpreted system. Then a *binding* is a partial function $b : \text{Agt} \rightarrow \text{Var}$ which maps agents in its domain to variables. $\text{Bnd} \triangleq \text{Agt} \rightarrow \text{Var}$ denotes the set of all bindings.

Definition 4.7 (Variable Assignments). Let \mathcal{I} be an interpreted system. Then a *variable assignment* is a function $v : \text{Var} \rightarrow \text{UStr}$ which maps variables in its domain to uniform shared memoryless strategies. $\text{VAsg} \triangleq \text{Var} \rightarrow \text{UStr}$ denotes the set of all variable assignments.

Note that variable assignments (Definition 4.7) are also assignments (Definition 2.23 modified for SLK strategies as explained in Subsection 4.1.2), i.e. $\text{VAsg} \subseteq \text{Asg}$.

As we have already explained, an extended state is a pair of a (i) global state and (ii) a variable assignment subject to which a formula is true in that state. For example, $(g_1, \{(x, f), (y, g)\})$ is an extended state which means roughly: “(the formula is true in) global state g_1 when agents bound to variables x and y act according to the strategies f and g respectively.” A formal definition follows.

Definition 4.8 (Extended States). Let \mathcal{I} be an interpreted system, $g \in G$ a global state and $v \in \text{VAsg}$ a variable assignment. Then an *extended state* is a pair $\langle g, v \rangle \in G \times \text{VAsg}$. $\text{Ext} \triangleq G \times \text{VAsg}$ denotes the set of all extended states.

The meaning of the variable assignment in an extended state is probably slightly unclear right now. Intuitively, an extended state $\langle g, v \rangle \in \text{Ext}$ *guarantees* a formula $\varphi \in \text{SLK}$ iff all assignments which agree with v make the formula true in the state g . Since this explanation is still quite vague, we define the notion of *guarantee* more formally.

Definition 4.9 (Guarantee). Let \mathcal{I} be an interpreted system, $\langle g, v \rangle \in \text{Ext}$ an extended state, $b \in \text{Bnd}$ a binding, and $\varphi \in \text{SLK}$ an SLK formula with $\text{free}(\varphi) \cap \text{Agt} \subseteq \text{dom}(b)$. We say that $\langle g, v \rangle$ *guarantees* φ in \mathcal{I} under b iff for the assignment $\chi_v \in \text{Asg}$ defined as:

$$\chi_v \triangleq v \cup \{(a, v(b(a))) \mid a \in \text{dom}(b)\}$$

we have $\mathcal{I}, \chi_v, g \models_{\text{SLK}} \varphi$. We write this as $\mathcal{I}, \langle g, v \rangle, b \vdash \varphi$.

Now that we have covered the basic structures, we can define the concepts of negation and predecessors of extended states. These will be necessary for the model checking algorithm presented at the end of this subsection. Let us start with negation. Assume that we have calculated the set of extended states E which guarantee the formula φ under a binding b . We want to find the set E' of extended states which guarantee the formula $\neg\varphi$. Intuitively, E' should contain all extended states that somehow disagree with E . E' is calculated as follows:

$$E' = \text{Ext} \setminus E$$

We shall now prove that our claim is correct. We assume that E is the set of all extended states that guarantee φ under b and show that $\text{Ext} \setminus E$ is the set of all extended states that guarantee $\neg\varphi$ under b .

Lemma 4.1. Let \mathcal{I} be an interpreted system, $b \in \text{Bnd}$ a binding, and $\varphi \in \text{SLK}$ an SLK formula with $\text{free}(\varphi) \cap \text{Agt} \subseteq \text{dom}(b)$. Let $E \subseteq \text{Ext}$ be the set of all extended states which guarantee φ in \mathcal{I} under b . Then $\text{Ext} \setminus E$ is the set of all extended states which guarantee $\neg\varphi$ in \mathcal{I} under b .

Proof. Let $E' \subseteq \text{Ext}$ be the set of all extended states which guarantee $\neg\varphi$ in \mathcal{I} under b . We show that $E' = \text{Ext} \setminus E$.

\Rightarrow : Take an arbitrary extended state $\langle g, v \rangle \in E'$. Since $\langle g, v \rangle$ guarantees $\neg\varphi$, we have $\mathcal{I}, \chi_v, g \models_{\text{SLK}} \neg\varphi$. By SLK semantics (Definition 4.3), $\mathcal{I}, \chi_v, g \not\models_{\text{SLK}} \varphi$. Thus $\langle g, v \rangle \notin E$, so we have $\langle g, v \rangle \in \text{Ext} \setminus E$.

\Leftarrow : Take an arbitrary extended state $\langle g, v \rangle \in \text{Ext} \setminus E$. Since $\langle g, v \rangle$ does not guarantee φ (otherwise, we would have $\langle g, v \rangle \in E$), we have $\mathcal{I}, \chi_v, g \not\models_{\text{SLK}} \varphi$. By SLK semantics (Definition 4.3), $\mathcal{I}, \chi_v, g \models_{\text{SLK}} \neg\varphi$. Thus, $\langle g, v \rangle \in \text{Ext}$ guarantees $\neg\varphi$.

□

Having covered negation, it remains to explain how to calculate the set of previous extended states, i.e. given a set of extended states $E \subseteq Ext$ which guarantee φ under a binding b , determine the set of extended states $E' \subseteq Ext$ which guarantee $\mathbf{X}\varphi$ under the same binding.

We first define the transition relation on global states implied by a binding and a variable assignment. Intuitively, there is a transition from state $g_1 \in G$ to a state $g_2 \in G$ (given the binding and variable assignment) if there exists a joint action between them such that each agent acts according to the strategy assigned to the variable it is bound to.

Definition 4.10 (Implied Transition Relation). Let \mathcal{I} be an interpreted system, $g_1, g_2 \in G$ two global states, $b \in Bnd$ a binding, and $v \in VAsg$ a variable assignment. Then the *transition relation* $\rightarrow_v^b \subseteq G \times G$ implied by b and v is defined by $g_1 \rightarrow_v^b g_2$, iff $\text{dom}(b) = Agt$ and there exists a joint action $a \in Act$ such that $t(g_1, a) = g_2$, and, for all agents $i \in Agt$, it holds that $a_i(a) = v(b(i))(g_1)$.

We are now ready to explain how the set of previous extended states is calculated. Intuitively, given an extended state $\langle g, v \rangle \in Ext$, the previous extended states are pairs $\langle g', v \rangle$ where g' is the successor of g when all agents act according to their strategies in v .

Definition 4.11 (Previous Extended States). Let \mathcal{I} be an interpreted system, $E \subseteq Ext$ a set of extended states, and $b \in Bnd$ a binding such that $\text{dom}(b) = Agt$. Then the function $\text{pre} : 2^{Ext} \times Bnd \rightarrow 2^{Ext}$, which returns the set of *previous extended states*, is defined as $\text{pre}(E, b) \triangleq \{\langle g, v \rangle \in Ext \mid \exists g' \in G. \langle g', v \rangle \in E \wedge g \rightarrow_v^b g'\}$.

Again, we will show that the function pre is correct. We assume that E is the set of all extended states which guarantee φ under b and show that $\text{pre}(E)$ is the set of all extended states which guarantee $\mathbf{X}\varphi$ under b .

Lemma 4.2. Let \mathcal{I} be an interpreted system, $b \in Bnd$ a binding with $\text{dom}(b) = Agt$, and $\varphi \in SLK$ an SLK formula. Let $E \subseteq Ext$ be the set of all extended states which guarantee φ in \mathcal{I} under b . Then $\text{pre}(E) \subseteq Ext$ is the set of all extended states which guarantee $\mathbf{X}\varphi$ in \mathcal{I} under b .

Proof. Let $E' \subseteq Ext$ be the set of all extended states which guarantee $\mathbf{X}\varphi$ in \mathcal{I} under b . We show that $\text{pre}(E) = E'$.

\Rightarrow : Take an arbitrary extended state $\langle g, v \rangle \in \text{pre}(E, b)$. By construction, there is a global state $g' \in G$ such that $\langle g', v \rangle \in E$ and $g \rightarrow_v^b g'$. Since $\langle g', v \rangle \in E$ guarantees φ , we have $\mathcal{I}, \chi_v, g' \models_{SLK} \varphi$. The implied transition relation implies $g' = t(g, a)$ where $a_i(a) = v(b(i))(g)$ for all agents $i \in Agt$. This can be also rewritten as $g' = t(g, \langle \chi_v(a)(g) : a \in Agt \rangle)$ because $\chi_v(i) = v(b(i))$ for all agents $i \in \text{dom}(b) = Agt$.

We want to show that $\langle g, v \rangle$ guarantees $\mathbf{X}\varphi$, i.e. $\mathcal{I}, \chi_v, g \models_{SLK} \mathbf{X}\varphi$. This is the case iff $\mathcal{I}, \chi_v, \pi(1) \models_{SLK} \varphi$ where $\pi = \text{play}(\chi_v, g)$ (Definition 4.3). From Definition 2.26, we get $\pi(1) = t(g, \langle \chi_v(a)(g) : a \in Agt \rangle)$, so $\pi(1) = g'$. Since we have already shown $\mathcal{I}, \chi_v, g' \models_{SLK} \varphi$, we have $\mathcal{I}, \chi_v, g \models_{SLK} \mathbf{X}\varphi$ as required.

\Leftarrow : Take an arbitrary extended state $\langle g, v \rangle \in E'$. Thus, we have $\mathcal{I}, \chi_v, g \models_{SLK} \mathbf{X}\varphi$. By Definition 4.3, this means $\mathcal{I}, \chi_v, \pi(1) \models_{SLK} \varphi$ where $\pi = \text{play}(\chi_v, g)$. From Definition 2.26, we get $\pi(1) = t(g, \langle \chi_v(a)(g) : a \in Agt \rangle)$. This can be equivalently written as $\pi(1) = t(g, a)$ where $a_i(a) = \chi_v(a)(g) = v(b(i))(g)$ for all agents $i \in Agt$.

As $\text{dom}(b) = Agt$ by assumption, we have $g \rightarrow_v^b \pi(1)$. Moreover, since $\mathcal{I}, \chi_v, \pi(1) \models_{SLK} \varphi$, we have $\langle \pi(1), v \rangle \in E$ (because it guarantees φ). Therefore, we have $\langle g, v \rangle \in \text{pre}(E, b)$ as required. \square

Finally, we have all the ingredients to define the model checking algorithm $\text{SAT}_{SLK}^{\mathcal{I}}(\cdot, \cdot)$ for SLK.

Definition 4.12 (SLK Model Checking Algorithm). Let \mathcal{I} be an interpreted system, $\varphi \in SLK$ an SLK formula and $b \in Bnd$ a binding, such that $\text{free}(\varphi) \cap Agt \subseteq \text{dom}(b)$. Then the model checking function $\text{SAT}_{SLK}^{\mathcal{I}} : SLK \times Bnd \rightarrow 2^{Ext}$ is inductively defined as follows:

1. $\text{SAT}_{SLK}^{\mathcal{I}}(\top, b) \triangleq Ext$.
2. $\text{SAT}_{SLK}^{\mathcal{I}}(p, b) \triangleq \{\langle g, v \rangle \mid g \in h(p)\}$, with $p \in AP$.

3. For all formulas $\varphi, \varphi_1, \varphi_2 \in SLK$, it is defined as:
 - (a) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\neg\varphi, b) \triangleq \text{Ext} \setminus \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi, b)$;
 - (b) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_1 \wedge \varphi_2, b) \triangleq \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_1, b) \cap \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_2, b)$;
 - (c) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_1 \vee \varphi_2, b) \triangleq \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_1, b) \cup \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_2, b)$.
4. For an agent $i \in \text{Agt}$, a variable $x \in \text{Var}$, and a formula $\varphi \in SLK$, $\text{SAT}_{\text{SLK}}^{\mathcal{I}}((i, x)\varphi, b) \triangleq \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi, b[i \mapsto x])$.
5. For a variable $x \in \text{Var}$ and an SLK formula $\varphi \in SLK$, it is defined as:
 - (a) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\langle\langle x \rangle\rangle\varphi, b) \triangleq \{ \langle g, v \rangle \mid \exists f \in \text{UStr}_{\text{sharing}(\varphi, x)}. \langle g, v[x \mapsto f] \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi, b) \}$;
 - (b) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\llbracket x \rrbracket\varphi, b) \triangleq \{ \langle g, v \rangle \mid \forall f \in \text{UStr}_{\text{sharing}(\varphi, x)}. \langle g, v[x \mapsto f] \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi, b) \}$.
6. For all formulas $\varphi, \varphi_1, \varphi_2 \in SLK$, it is defined as:
 - (a) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{X}\varphi, b) \triangleq \text{pre}(\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi, b), b)$;
 - (b) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{F}\varphi, b) \triangleq \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{T}\text{U}\varphi, b)$;
 - (c) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{G}\varphi, b) \triangleq \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\neg\text{F}\neg\varphi, b)$;
 - (d) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_1 \text{U}\varphi_2, b) \triangleq \text{lf}_X[\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_2, b) \cup (\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi_1, b) \cap \text{pre}(X, b))]$.
7. For an agent $i \in \text{Agt}$, a set of agents $A \subseteq \text{Agt}$, and a formula $\varphi \in SLK$, it is defined as:
 - (a) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{K}_i\varphi, b) \triangleq \text{Ext} \setminus \{ \langle g, v \rangle \in \text{Ext} \mid \exists \langle g', v' \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\neg\varphi, \emptyset). g' \sim_i g \}$;
 - (b) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{E}_A\varphi, b) \triangleq \text{Ext} \setminus \{ \langle g, v \rangle \in \text{Ext} \mid \exists \langle g', v' \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\neg\varphi, \emptyset). g' \sim_A^E g \}$;
 - (c) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{D}_A\varphi, b) \triangleq \text{Ext} \setminus \{ \langle g, v \rangle \in \text{Ext} \mid \exists \langle g', v' \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\neg\varphi, \emptyset). g' \sim_A^D g \}$;
 - (d) $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\text{C}_A\varphi, b) \triangleq \text{Ext} \setminus \{ \langle g, v \rangle \in \text{Ext} \mid \exists \langle g', v' \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\neg\varphi, \emptyset). g' \sim_A^C g \}$.

The correctness of the algorithm is asserted in the following theorem.

Theorem 4.2. Let \mathcal{I} be an interpreted system and $\varphi \in SLK$ an SLK sentence. Then the set of all states at which φ holds is: $\{g \in G \mid \mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi\} = \{g \in G \mid \exists v \in \text{VAsg}. \langle g, v \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi, \emptyset)\}$.

Proof (Sketch). We prove by induction that for an arbitrary interpreted system \mathcal{I} , SLK formula $\varphi \in SLK$, and binding $b \in \text{Bnd}$ such that $\text{free}(\varphi) \subseteq \text{dom}(b)$, $\text{SAT}_{\text{SLK}}^{\mathcal{I}}(\varphi, b)$ is the set of all extended states that guarantee φ in \mathcal{I} under b . The two important cases were proved in Lemmas 4.1 (negation) and 4.2 (temporal step).

Since the topmost formula φ is a sentence, it will either hold, or not hold in each state (regardless of the variable assignment). Therefore, we existentially quantify over variable assignments. \square

Observe that the SLK model checking algorithm is decidable because the sets UStr , Var , VAsg , Agt , Bnd , G , and, consequently, Ext are finite. An efficient symbolic implementation of the algorithm using BDDs is presented in Subsection 4.2.4.

4.2.3 Strategy Synthesis

As we have explained in Subsection 4.1.5, one of the main limitations of SLK are *non-behavioural strategies*, which depend on counterfactual scenarios. Consequently, SLK strategies are very difficult⁶ to synthesise. To see why this is the case, consider the SLK sentence $\varphi \triangleq \llbracket x \rrbracket \llbracket y \rrbracket \langle\langle z \rangle\rangle \psi$. Assume that φ holds at a particular state $g \in G$ in an interpreted system \mathcal{I} , i.e. $\mathcal{I}, g \models_{\text{SLK}} \varphi$. We would now like to synthesise a uniform shared strategy $f_z : G \rightarrow \text{Act}_z$ for the variable z depending on the uniform shared strategies $f_x : G \rightarrow \text{Act}_x$ and $f_y : G \rightarrow \text{Act}_y$ for the variables x and y respectively where $\text{Act}_v \triangleq \text{Act}_{\text{sharing}(\psi, v)}$

⁶Unlike perfect recall SL strategies, which cannot be synthesised in general due to being non-behavioural, imperfect recall SLK strategies can always be synthesised. This should not be surprising given that the set of uniform strategies is *finite* for a given interpreted system and SLK formula. Hence, we can explicitly enumerate all possible combinations of strategies.

for $v \in \text{Var}$. If SLK strategies were behavioural, there *would* exist a mapping m_1 (also referred to as *elementary dependence map* [75]) from the next actions of f_x and f_y in g to the next action of f_z in g :

$$m_1 : G \rightarrow \underbrace{(Act_x \times Act_y \rightarrow Act_z)}_{(g, f_x(g), f_y(g)) \mapsto f_z(g)}$$

There would be at most $|Act_x| \times |Act_y|$ possible inputs to m_1 to determine $f_z(g)$ as it depends only on $f_x(g)$ and $f_y(g)$. Unfortunately, such a mapping does not exist in general because SLK strategies are non-behavioural. Instead, a more general mapping m_2 (also referred to as *dependence map* [75]) from strategies f_x and f_y to the strategy f_z must be considered:

$$m_2 : \underbrace{(G \rightarrow Act_x) \times (G \rightarrow Act_y)}_{(f_x, f_y, g) \mapsto f_z(g)} \rightarrow (G \rightarrow Act_z)$$

Informally, determining $f_z(g)$ requires the same amount of information as constructing the whole strategy f_z . In order to synthesise the action $f_z(g)$ or the strategy f_z , we possibly need to know the complete strategies f_x and f_y . More importantly, the maximum number of *entries* in the mappings m_1 and m_2 are:

$$m_1 : |G| \times |Act_x| \times |Act_y| \qquad m_2 : |G| \times |Act_x|^{|G|} \times |Act_y|^{|G|}$$

Assume that the interpreted system \mathcal{I} is quite small and has $|G| = 10$ global states and $|Act_x| = |Act_y| = |Act_z| = 10$ actions. While the mapping m_1 for *behavioural* strategies would require at most 1000 entries of the form $(g, f_x(g), f_y(g)) \mapsto f_z(g)$, the mapping m_2 for *non-behavioural* strategies might have up to 10^{21} entries of the form $(f_x, f_y, g) \mapsto f_z(g)$. Furthermore, if we encode each output of f_z using only $\lceil \log_2 |Act_z| \rceil = 3$ bits and store the whole mapping in a large array, m_1 will use at most 375 *bytes* while m_2 might need up to 375 *exabytes*. To put this number into perspective, the total amount of stored information in the whole world was 295 exabytes (optimally compressed) in 2007 [47]. This example demonstrates the high impracticality of general SLK strategy synthesis on any non-trivial models.

While SLK strategy synthesis is infeasible in general, it can be performed efficiently on certain types of formulas. We will now explain the concepts of *witness* and *counterexample strategies* and describe how these can be synthesised using the model checking algorithm discussed in Subsection 4.2.2. Let \mathcal{I} be an interpreted system, $g \in G$ a global state, and $\varphi_w \triangleq \langle\langle x \rangle\rangle \psi_w$ and $\varphi_c \triangleq \llbracket y \rrbracket \psi_c$ two SLK sentences. Furthermore, assume that φ_w holds at g while φ_c does not, i.e. $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi_w$ and $\mathcal{I}, \emptyset, g \not\models_{\text{SLK}} \varphi_c$. By SLK semantics (see Definition 4.3), there is a memoryless uniform shared strategy f_w for x which makes ψ_w true at g . Conversely, there must be a memoryless shared strategy f_c for y which makes ψ_c false at g . f_w and f_c are referred to as a *witness* and a *counterexample strategy* respectively. Intuitively, the strategy f_w is a “witness” to φ_w being true at g while f_c is a “counterexample” for φ_c at g . A slightly more general form of the two concepts is provided in the following definition.

Definition 4.13 (Witness and Counterexample Strategies). Let \mathcal{I} be an interpreted system, $g \in G$ a global state, and $\varphi_w \triangleq \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{m-1} \rangle\rangle \psi_w$ and $\varphi_c \triangleq \llbracket y_0 \rrbracket \dots \llbracket y_{n-1} \rrbracket \psi_c$ two SLK sentences. Then:

- Memoryless uniform shared strategies $f_{w0}, \dots, f_{w(m-1)}$ are *witness strategies* for φ_w at g iff (i) $f_{wi} \in \text{UStr}_{\text{sharing}(\psi_w, x_i)}$ for all $0 \leq i < m$ and (ii) $\mathcal{I}, \chi_w, g \models_{\text{SLK}} \psi_w$ where $\chi_w \triangleq \{(x_i, f_{wi}) \mid 0 \leq i < m\}$.
- Memoryless uniform shared strategies $f_{c0}, \dots, f_{c(n-1)}$ are *counterexample strategies* for φ_c at g iff (i) $f_{ci} \in \text{UStr}_{\text{sharing}(\psi_c, y_i)}$ for all $0 \leq i < n$ and (ii) $\mathcal{I}, \chi_c, g \not\models_{\text{SLK}} \psi_c$ where $\chi_c \triangleq \{(y_i, f_{ci}) \mid 0 \leq i < n\}$.

Consider the toy model in Figure 3.1 and the SLK formulas $\varphi_1 \triangleq \langle\langle e \rangle\rangle \langle\langle x \rangle\rangle \langle\langle y \rangle\rangle (E, e)(1, x)(2, y) \times p_1$ and $\varphi_2 \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \llbracket y \rrbracket (E, e)(1, x)(2, y) \times p_2$. The formula φ_1 means that “*there exist strategies for all the agents, such that player 1 wins in the next round*”. The formula φ_2 expresses that “*for all strategies of the agents, player 2 wins in the next round*”. By SLK semantics, we have $\mathcal{I}, \emptyset, g_g \models_{\text{SLK}} \varphi_1$ and $\mathcal{I}, \emptyset, g_g \not\models_{\text{SLK}} \varphi_2$. Incidentally, the following strategies are both *witness strategies* for φ_1 and *counterexample strategies*

for φ_2 at g_g :

$$\begin{array}{lll} f_e(g_g) \triangleq i & f_x(g_g) \triangleq p & f_y(g_g) \triangleq r \\ f_e(g_1) \triangleq i & f_x(g_1) \triangleq i & f_y(g_1) \triangleq i \\ f_e(g_2) \triangleq i & f_x(g_2) \triangleq i & f_y(g_2) \triangleq i \end{array}$$

The two concepts are *duals* of each other in the sense that if f is a witness strategy for $\langle\langle x \rangle\rangle\psi$ at g , then it is a counterexample strategy for $\llbracket x \rrbracket\neg\psi$ at g (and vice versa).

Lemma 4.3. Let \mathcal{I} be an interpreted system, $g \in G$ a global state, ψ an agent-closed SLK formula with $\text{free}(\psi) = \{x_0, \dots, x_{n-1}\}$, and f_0, \dots, f_{n-1} memoryless uniform shared strategies such that $f_i \in UStr_{\text{sharing}(x_i, \psi)}$ for $0 \leq i < n$. Then f_0, \dots, f_{n-1} are witness strategies of $\langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi$ at g iff f_0, \dots, f_{n-1} are counterexample strategies of $\llbracket x_0 \rrbracket \dots \llbracket x_{n-1} \rrbracket\neg\psi$ at g .

Proof. We prove both directions of the equivalence separately:

\Rightarrow : Assume that f_0, \dots, f_{n-1} are witness strategies of $\langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi$ at g . By Definition 4.13, we have $f_i \in UStr_{\text{sharing}(\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, \chi, g \models \psi$ where $\chi \triangleq \{(x_i, f_i) \mid 0 \leq i < n\}$. By Definition 2.22, we have $\text{sharing}(\neg\varphi, x) = \text{sharing}(\varphi, x)$ for all SLK formulas $\varphi \in SLK$ and variables $x \in Var$ so $f_i \in UStr_{\text{sharing}(\neg\psi, x_i)}$ for all $0 \leq i < n$. Since $\mathcal{I}, \chi, g \models_{\text{SLK}} \psi$, $\mathcal{I}, \chi, g \not\models_{\text{SLK}} \neg\psi$ holds by SLK semantics (Definition 4.3). Hence, by Definition 4.13, f_0, \dots, f_{n-1} are counterexample strategies for $\llbracket x_0 \rrbracket \dots \llbracket x_{n-1} \rrbracket\neg\psi$ at g .

\Leftarrow : Assume that f_0, \dots, f_{n-1} are counterexample strategies of $\llbracket x_0 \rrbracket \dots \llbracket x_{n-1} \rrbracket\neg\psi$ at g . By Definition 4.13, we have $f_i \in UStr_{\text{sharing}(\neg\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, \chi, g \not\models_{\text{SLK}} \neg\psi$ where $\chi \triangleq \{(x_i, f_i) \mid 0 \leq i < n\}$. Again, we have $f_i \in UStr_{\text{sharing}(\psi, x_i)}$ for all $0 \leq i < n$ by Definition 2.22. Since $\mathcal{I}, \chi, g \not\models_{\text{SLK}} \neg\psi$, $\mathcal{I}, \chi, g \models_{\text{SLK}} \psi$ holds by SLK semantics (Definition 4.3). Hence, by Definition 4.13, f_0, \dots, f_{n-1} are witness strategies of $\langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi$ at g . □

This duality will allow us to focus on witness strategies and their synthesis only: Suppose an SLK sentence $\varphi \triangleq \llbracket x_0 \rrbracket \dots \llbracket x_{n-1} \rrbracket\psi$ does not hold at a state $g \in G$ in an interpreted system \mathcal{I} and we want to construct counterexample strategies for φ at g . Since $\mathcal{I}, \emptyset, g \not\models_{\text{SLK}} \varphi$, we must have $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \neg\varphi$ by SLK semantics (see Definition 4.3). Using the equivalence $\neg \llbracket x_0 \rrbracket \dots \llbracket x_{n-1} \rrbracket\psi \equiv \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\neg\psi$, we get $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\neg\psi$. Finally, we construct witness strategies for $\neg\varphi$ at g , which are also counterexample strategies for φ at g by Lemma 4.3.

The procedure for constructing a counterexample strategy relies on witness strategy synthesis. Before describing how witness strategies can be retrieved, we need to show that they always exist when an SLK formula holds.

Lemma 4.4. Let \mathcal{I} be an interpreted system, ψ an agent-closed SLK formula such that $\text{free}(\psi) = \{x_0, \dots, x_{n-1}\}$, and $\varphi \triangleq \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi$ an SLK sentence. Then the following holds: $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi$ iff there exist witness strategies f_0, \dots, f_{n-1} for φ at g .

Proof. We prove both directions of the equivalence separately:

\Rightarrow : Assume that $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi$. By SLK semantics (Definition 4.3), there exist strategies f_0, \dots, f_{n-1} such that $f_i \in UStr_{\text{sharing}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, \chi, g \models_{\text{SLK}} \psi$ where $\chi = \{(x_i, f_i) \mid 0 \leq i < n\}$. By Definition 2.22, we have $\text{sharing}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi, x_i) = \text{sharing}(\psi, x_i)$ for all $0 \leq i < n$. Hence f_0, \dots, f_{n-1} satisfy both conditions for being witness strategies of φ at g (see Definition 4.13).

\Leftarrow : Assume that there exist witness strategies f_0, \dots, f_{n-1} for φ at g . By Definition 4.13, we have $f_i \in UStr_{\text{sharing}(\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, \chi, g \models_{\text{SLK}} \psi$ where $\chi \triangleq \{(x_i, f_i) \mid 0 \leq i < n\}$. By Definition 2.22, we have $\text{sharing}(\psi, x_i) = \text{sharing}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi, x_i)$. Hence, we get $f_i \in UStr_{\text{sharing}(\langle\langle x_{i+1} \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle\psi, x_i)}$ for all $0 \leq i < n$. Therefore, $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi$ by SLK semantics (Definition 4.3).

□

It remains to explain how witness strategies can be synthesised using the SLK model checking algorithm we introduced in Subsection 4.2.2. Consider an SLK sentence $\varphi \triangleq \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ that holds at a global state $g \in G$ in some interpreted system \mathcal{I} . Since $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi$, by Lemma 4.4, there must be some witness strategies f_0, \dots, f_{n-1} for φ at g , which we want to synthesise. The corresponding assignment $\chi \in \text{Asg}$ on the variables x_0, \dots, x_{n-1} satisfies $\mathcal{I}, \chi, g \models_{\text{SLK}} \psi$. Let $E \triangleq \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi, \emptyset)$ be the set of extended states which guarantee ψ in \mathcal{I} under the empty binding. E contains all possible extended states $\langle g', v' \rangle$ such that $\mathcal{I}, g', v' \models_{\text{SLK}} \psi$. Hence, it must be the case that $\langle g, v \rangle \in E$ where v is some variable assignment which extends χ .

Therefore, in order to synthesise witness strategies for φ at g , it suffices to find an extended state $\langle g, v \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi, \emptyset)$ such that the strategies $v(x_0), \dots, v(x_{n-1})$ are uniform with respect to the agents which share them. The witness strategies for φ at g are then $v(x_0), \dots, v(x_{n-1})$.

Lemma 4.5. Let \mathcal{I} be an interpreted system, ψ an agent-closed SLK formula such that $\text{free}(\psi) = \{x_0, \dots, x_{n-1}\}$, and $\varphi \triangleq \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ an SLK sentence. Then the following properties hold:

1. For all variable assignments $\langle g, v \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi, \emptyset)$ such that $v(x_i) \in \text{UStr}_{\text{sharing}(\psi, x_i)}$ for $0 \leq i < n$, $v(x_0), \dots, v(x_{n-1})$ are witness strategies for φ at g .
2. If there exist witness strategies for φ at g , then there exists a variable assignment $v \in \text{VAsg}$ such that $v(x_i) = f_i$ for $0 \leq i < n$ and $\langle g, v \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi, \emptyset)$.

Proof. We prove both properties separately:

1. Take an arbitrary extended state $\langle g, v \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi, \emptyset)$ such that $v(x_i) \in \text{UStr}_{\text{sharing}(\psi, x_i)}$ for $0 \leq i < n$. As we have explained in the proof of Theorem 4.2, $\langle g, v \rangle$ guarantees ψ in \mathcal{I} under \emptyset . By Definition 4.9, $\mathcal{I}, g, v \models_{\text{SLK}} \psi$. By SLK semantics (Definition 4.3), we have $\mathcal{I}, \chi, g \models_{\text{SLK}} \psi$ where $\chi \triangleq \{(x_i, v(x_i)) \mid 0 \leq i < n\}$ since $\text{free}(\psi) \subseteq \{x_0, \dots, x_{n-1}\} \subseteq \text{dom}(v)$. Hence, $v(x_0), \dots, v(x_{n-1})$ satisfy both conditions for being witness strategies for φ at g (see Definition 4.13).
2. Assume that f_0, \dots, f_{n-1} are witness strategies for φ at g . By Definition 4.13, $f_i \in \text{UStr}_{\text{sharing}(\psi, x_i)}$ for all $0 \leq i < n$ and $\mathcal{I}, \chi, g \models_{\text{SLK}} \psi$ where $\chi \triangleq \{(x_i, f_i) \mid 0 \leq i < n\}$. Since $\text{dom}(\chi) = \text{free}(\psi)$, by SLK semantics (Definition 4.3), we have $\mathcal{I}, g, v \models_{\text{SLK}} \psi$ for all $v \in \text{VAsg}$ such that $\chi \subseteq v$. There must exist at least one such variable assignment v because $\text{dom}(\chi) \subseteq \text{Var}$ (simply set $v(x) \triangleq \chi(x)$ for $x \in \text{dom}(\chi)$ and assign arbitrary strategies to variables $y \in \text{Var} \setminus \text{dom}(\chi)$). By Definition 4.9, $\langle g, v \rangle \in \text{Ext}$ guarantees ψ in \mathcal{I} under \emptyset . Therefore, $\langle g, v \rangle \in \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi, \emptyset)$ (see proof of Theorem 4.2).

□

Informally, the first property in Lemma 4.5 expresses *soundness* of the approach, i.e. that it will return only witness strategies for φ at g . Conversely, the second property asserts *completeness* of the approach, i.e. that it will return witness strategies for φ at g , if they exist.

4.2.4 Symbolic Implementation

In this subsection, we discuss how the algorithm presented in Subsection 4.2.2 can be implemented symbolically using BDDs. As explained in Subsection 2.3.2, BDDs are a very efficient representation for manipulation of Boolean formulas. BDDs are used by many existing model checkers including MCMAS (see Subsection 2.5.2), which uses them for model checking CTL and ATL [65]. We present here a modification of the symbolic implementation for SLK.

We start by representing the parameters of the interpreted system by means of Boolean formulas [81]. Given an interpreted system:

$$\mathcal{I} = \left\langle (L_i, \text{Act}_i, P_i, t_i)_{i \in \text{Agt}}, I, h \right\rangle$$

we can represent global states and joint actions as follows [65, 81]:

- For every agent $i \in \text{Agt}$, we can encode its set of internal states L_i with $\text{nv}(i) = \lceil \log_2 |L_i| \rceil$ Boolean variables. Thus, a *global state* $g \in G$ can be encoded as a Boolean vector $\bar{v} = (v_0, \dots, v_{N-1})$, where $N = \sum_{i \in \text{Agt}} \text{nv}(i)$.

- For every agent $i \in \text{Agt}$, we can encode its set of actions Act_i with $\text{na}(i) = \lceil \log_2 |\text{Act}_i| \rceil$ Boolean variables. Thus, a *joint action* $a \in \text{Act}$ can be encoded as a Boolean vector $\bar{w} = (w_0, \dots, w_{M-1})$, where $M = \sum_{i \in \text{Agt}} \text{na}(i)$.

An example of how this encoding can be done is provided in Subsection 3.3.2.

Encoding of Extended States

In order to implement our new SLK model checking algorithm, we need to represent sets of *extended states* (see Definition 4.8), which consist of a global state and a variable assignment. We thus represent the variable assignment explicitly using Boolean variables as well.

Since the number of strategy variables (and hence the domain of a variable assignment) depends on the SLK formula we are checking, the total number of Boolean variables also depends on the formula. Let $\varphi \in \text{SLK}$ be an SLK sentence and $\text{vars}(\varphi) \subseteq \text{Var}$ the set of variables quantified in φ , e.g. $\text{vars}(\langle\langle e \rangle\rangle \llbracket x \rrbracket \langle\langle y \rangle\rangle (\text{E}, e)(1, x)(2, y) \text{G} [\neg p_1 \wedge \neg p_2 \wedge \langle\langle z \rangle\rangle (2, z) \text{X} p_2]) = \{e, x, y, z\}$. For each variable $x \in \text{vars}(\varphi)$, we represent the strategy associated with x using a number of Boolean variables.

Let us now consider an arbitrary variable $x \in \text{vars}(\varphi)$. In the SLK model checking algorithm (Definition 4.12), the variable quantifies over strategies $f \in \text{UStr}_{\text{sharing}(\varphi, x)}$. The domain and the range of a strategy $f \in \text{UStr}_A$ are G and Act_A respectively (see Definition 4.2). Hence, we could easily represent it using $|G| \times \lceil \log_2 |\text{Act}_A| \rceil$ Boolean variables, i.e. we would encode the action associated with each global state. For example, the following strategy for player 1 ($A = \{1\}$) in the toy model (see Figure 3.1):

$$f(g_g) \triangleq r \qquad f(g_1) \triangleq i \qquad f(g_2) \triangleq i$$

could be encoded using $3 \times \lceil \log_2 |r, p, s, i| \rceil = 6$ Boolean variables:

$$\underbrace{\neg b_0 \wedge \neg b_1}_{f(g_g)=r} \wedge \underbrace{b_2 \wedge b_3}_{f(g_1)=i} \wedge \underbrace{b_4 \wedge b_5}_{f(g_2)=i}$$

We can see that this representation is not the most compact one. Namely, we do not need 2 Boolean variables to encode the action in state g_1 because $P_1(s, s_1) = \{i\}$. In fact, we do not need any Boolean variables to encode $f(g_1) = f(g_2) = i$ because i is the only available action in both g_1 and g_2 so agent 1 always has to perform it in these states. Hence, a more compact representation can be obtained by using only $\sum_{g \in G} \lceil \log_2 |\bigcap_{i \in A} P_i(l_{iE}(g))| \rceil$ Boolean variables. The strategy f for the toy model would now be encoded using $\lceil \log_2 |\{r, p, s\}| \rceil + 2 \lceil \log_2 |\{i\}| \rceil = 2$ Boolean variables as $\neg b_0 \wedge \neg b_1$, which is optimal since there are exactly three possible strategies available to agent 1. More generally, this encoding is almost optimal⁷ for shared memoryless strategies.

However, our strategies are also *uniform*. If have a strategy $f \in \text{UStr}_{\{i\}}$ for an agent $i \in \text{Agt}$, then for all states $g_1, g_2 \in G$, $g_1 \sim_i g_2$ implies $f(g_1) = f(g_2)$. Hence, there is no point in storing both actions $f(g_1)$ and $f(g_2)$. Consider some interpreted system with 2 agents a, b and 3 reachable global states g_1, g_2, g_3 where $g_1 \sim_a g_2$ and $g_2 \sim_b g_3$. Furthermore, assume the following protocols:

$$\begin{aligned} P_a(l_{aE}(g_1)) &= P_a(l_{aE}(g_2)) \triangleq \{a_1, a_2, a_3\} & P_b(l_{bE}(g_1)) &\triangleq \{a_1, a_2\} \\ P_a(l_{aE}(g_3)) &\triangleq \{a_2, a_3\} & P_b(l_{bE}(g_2)) &= P_b(l_{bE}(g_3)) \triangleq \{a_1, a_2, a_3\} \end{aligned}$$

How many Boolean variables do we need to encode a strategy $f \in \text{UStr}_{\{a,b\}}$ for both agents? Let us forget for a while that f is uniform. The more compact representation we found would require $\sum_{g \in G} \lceil \log_2 |P_a(l_{aE}(g)) \cap P_b(l_{bE}(g))| \rceil = \lceil \log_2 |\{a_1, a_2, a_3\}| \rceil + \lceil \log_2 |\{a_1, a_2\}| \rceil + \lceil \log_2 |\{a_2, a_3\}| \rceil = 4$ Boolean variables. What are the possible strategies f ? We require $f(g) \in P_i(l_{iE}(g))$ for all agents $i \in \{a, b\}$ and global states $g \in G$: $f(g_1) \in \{a_1, a_2\}$, $f(g_2) \in \{a_1, a_2, a_3\}$, and $f(g_3) \in \{a_2, a_3\}$. So there are 12 possible shared memoryless strategies. But f must be *uniform*. Since $g_1 \sim_a g_2$, and $g_2 \sim_b g_3$, we must have $f(g_1) = f(g_2) = f(g_3)$. Therefore, there is only one possible uniform strategy defined as $f(g) \triangleq a_2$ for all $g \in G$. This answers our first question: We do not need any Boolean variables.

⁷It is optimal as long as we store each action individually. Consider the situation when an agent can perform 3 actions in state s_1 and 5 actions in s_2 . Our encoding will require $\lceil \log_2 3 \rceil + \lceil \log_2 5 \rceil = 5$ Boolean variables to represent a strategy. Since there are $3 \times 5 = 15$ possible strategies, it is possible to encode a strategy using only $\lceil \log_2 15 \rceil = 4$ Boolean variables.

The example above showed us that for a given set of agents $A \subseteq \text{Agt}$, the epistemic accessibility relations \sim_i with $i \in A$ induce regions of the global state space, to which the uniform strategies $f \in \text{UStr}_A$ must assign the same action. It turns out that these regions are equivalence classes with respect to the common epistemic accessibility relation \sim_A^C . We shall now prove this statement.

Lemma 4.6. Let \mathcal{I} be an interpreted system and $A \subseteq \text{Agt}$ a set of agents. Then a memoryless strategy $f : G \rightarrow \text{Act}_A$ is *uniform* iff for each set of global states in the quotient set $S \in G/\sim_A^C$, we have $f(s_1) = f(s_2)$ for all $s_1, s_2 \in S$.

Proof. We shall prove both directions of the equivalence separately.

\Rightarrow : Assume that f is uniform and take an arbitrary set $S \in G/\sim_A^C$. Furthermore, take arbitrary global states $g_1, g_2 \in S$. By the definition of quotient set, we have $g_1 \sim_A^C g_2$. There are two cases:

- $g_1 = g_2$. Trivially, $f(g_1) = f(g_2)$.
- By the definition of common epistemic accessibility relation (Definition 2.31), there is a chain of global states $g'_1, g'_2, \dots, g'_n \in G$ and agents $i_1, i_2, \dots, i_{n+1} \in A$ with $n \geq 0$ such that $g_1 \sim_{i_1} g'_1 \sim_{i_2} g'_2 \dots g'_n \sim_{i_{n+1}} g_2$. By uniformity of f (Definition 4.2), we get $f(g_1) = f(g'_1) = \dots = f(g'_n) = f(g_2)$.

\Leftarrow : Assume that for each $S \in G/\sim_A^C$, we have $f(s_1) = f(s_2)$ for all $s_1, s_2 \in S$. Take an arbitrary agent $i \in A$ and global states $s_1, s_2 \in G$ such that $s_1 \sim_i s_2$. To prove uniformity of f , we need to show that $f(s_1) = f(s_2)$. Since $s_1 \sim_i s_2$, we also have $s_1 \sim_A^C s_2$. Hence, s_2 belongs to the equivalence class $[s_1]_{\sim_A^C}$. By definition of an equivalence class, it must be the case that $s_1 \in [s_1]_{\sim_A^C}$ and $[s_1]_{\sim_A^C} \in G/\sim_A^C$. Since $s_1, s_2 \in [s_1]_{\sim_A^C}$ and $[s_1]_{\sim_A^C} \in G/\sim_A^C$, we get $f(s_1) = f(s_2)$ by the initial assumption as required. □

This allows us to present an even more compact representation of a strategy $f \in \text{UStr}_A$ with $A \subseteq \text{Agt}$ in an interpreted system \mathcal{I} . We only need to store one action for each *shared local state* $S \in G/\sim_A^C$. Thus, we can represent f using $\sum_{S \in G/\sim_A^C} \left[\log_2 \left| \bigcap_{g \in S} \bigcap_{i \in A} P_i(l_{iE}(g)) \right| \right]$ Boolean variables. Finally, a variable assignment $v \in \text{VAsg}$ for a formula $\varphi \in \text{SLK}$ can be represented using a Boolean vector $\bar{u} = (u_0, \dots, u_{K-1})$ such that:

$$K = \sum_{x \in \text{vars}(\varphi)} \sum_{S \in G_x} \left[\log_2 \left| \bigcap_{g \in S} \bigcap_{i \in \text{sharing}(\varphi, x)} P_i(l_{iE}(g)) \right| \right]$$

where $G_x \triangleq G/\sim_{\text{sharing}(\varphi, x)}^C$ is the set of shared local states for variable x . Note that despite both optimisations, the worst case still remains $K = |\text{vars}(\varphi)| \times |G| \times \lceil \log_2(\max_{i \in \text{Agt}} |\text{Act}_i|) \rceil$, i.e. we need polynomially many Boolean variables with respect to both the size of the model $|\mathcal{I}|$ and the number of strategy variables in the formula $|\text{vars}(\varphi)|$.

An extended state $\langle g, v \rangle \in \text{Ext}$ can be represented by a concatenation of the Boolean vectors \bar{v}_g and \bar{u}_v , which can in turn be identified with Boolean formulas, represented by conjunctions of literals (as in Subsection 2.3.2). A set of extended states $E \subseteq \text{Ext}$ can be expressed as the disjunction of the Boolean formulas encoding each extended state in E .

Encoding of the Algorithm

Given a binding $b \in \text{Bnd}$ such that $\text{dom}(b) = \text{Agt}$, we define a formula $S^b(e, a)$, where $e \in \text{Ext}$ and $a \in \text{Act}$, representing the *strategy restrictions* of the implied transition relation (see Definition 4.10). The formula asserts that all agents act according to their strategies:

$$S^b(\langle g, v \rangle, a) \triangleq \forall i \in \text{Agt}. v(b(i))(g) = a_i(a)$$

Let \bar{v} , \bar{v}' , \bar{w} , and \bar{u} be the Boolean vectors for representing current global states, next global states, joint actions, and variable assignments respectively. We can encode strategy restrictions as a Boolean formula $S^b(\bar{v}, \bar{w}, \bar{u})$:

$$S^b(\bar{v}, \bar{w}, \bar{u}) = \bigwedge_{i \in \text{Agt}} \bigvee_{l \in L_{i\text{E}}} \left[l(\bar{v}_{i\text{E}}) \wedge \bigvee_{a \in P_i(l)} a(\bar{w}_{i\text{E}}) \wedge a(\bar{u}_{v(b(i)), l}) \right]$$

where:

- $l(\bar{v}_{i\text{E}})$ is the Boolean formula representing that the local state of agent $i \in \text{Agt}$ is $l \in L_{i\text{E}}$;
- $a(\bar{w}_{i\text{E}})$ is the Boolean formula representing that the action of agent $i \in \text{Agt}$ is $a \in \text{Act}_i$;
- $a(\bar{u}_{v(x), l})$ is the Boolean formula representing the fact that the action in local state⁸ l assigned by the strategy mapped to variable x is a .

We can represent the global protocol and evolution function as Boolean formulas $P(\bar{v}, \bar{w})$ and $t(\bar{v}, \bar{w}, \bar{v}')$ by taking the conjunctions of Boolean formulas representing the individual agent's protocols P_i and evolution functions t_i ($i \in \text{Agt}$), like we did in Subsection 3.3.2. The Boolean formula $R_t^b(\bar{v}, \bar{w}, \bar{v}')$ for the *implied transition relation* $\rightarrow_v^b \subseteq G \times G$ (see Definition 4.10) is then constructed from the conjunction of the Boolean formulas representing the global protocol, global evolution function, and strategy restrictions:

$$R_t^b(\bar{v}, \bar{v}', \bar{u}) = \exists \bar{w}. P(\bar{v}, \bar{w}) \wedge t(\bar{v}, \bar{w}, \bar{v}') \wedge S^b(\bar{v}, \bar{w}, \bar{u})$$

Note that we quantify over actions, encoded as \bar{w} , as in Subsection 3.3.2, but we keep the variable assignment in the extra parameter⁹ \bar{u} . Quantification over the variable assignment is performed when a strategy quantifier ($\langle\langle x \rangle\rangle$, $\llbracket x \rrbracket$) is encountered. Also note that the strategy restrictions S^b depend on the binding b and thus have to be recomputed when the binding is updated, i.e. when the agent binding operator (i, x) with $i \in \text{Agt}$ and $x \in \text{Var}$ is encountered. This is not the case for the other Boolean formulas ($P(\bar{v}, \bar{w})$ and $t(\bar{v}, \bar{w}, \bar{v}')$), which are constant for a given interpreted system.

The *individual epistemic accessibility relation* $R_i^K(\bar{v}, \bar{v}')$ for each agent $i \in \text{Agt}$ and the other epistemic accessibility relations are encoded in a similar way [65]. The computation of the set of *reachable states* G is explained at the end of Subsection 3.3.2.

Finally, the algorithm $\text{SAT}_{\text{SLK}} : \text{SLK} \times \text{Bnd} \rightarrow 2^{\text{Ext}}$ can be translated into operations on BDDs representing encoded sets of extended states.

Strategy Synthesis

Synthesising *witness* and *counterexample strategies* (see Definition 4.13) using the symbolic implementation is very simple because our representation of variable assignments ensures that all strategies are uniform (see Definition 4.2).

Assume that we want to synthesise witness strategies for an SLK sentence $\varphi \triangleq \langle\langle x_0 \rangle\rangle \dots \langle\langle x_{n-1} \rangle\rangle \psi$ at a global state $g \in G$. Since $\mathcal{I}, \emptyset, g \models_{\text{SLK}} \varphi$, there exist witness strategies for φ at g (see Lemma 4.4). We calculate $E = \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi, \emptyset)$ and pick an arbitrary extended state $\langle g, v \rangle \in E$. Note that all strategies in v are uniform due to our representation. By Lemma 4.5, there exists at least one such extended state $\langle g, v \rangle$ (because there exist witness strategies for φ at g) and for all such extended states, $v(x_0), \dots, v(x_{n-1})$ are witness strategies for φ at g . Witness strategy synthesis for φ at g can be implemented symbolically as follows:

1. Calculate $E := \text{SAT}_{\text{SLK}}^{\mathcal{I}}(\psi)$. This is performed using the symbolic implementation presented earlier in this subsection.
2. Filter out extended states with a different global state $E' := \{(g', v') \in E \mid g' = g\}$. The symbolic representation of this operation is $E'(\bar{v}, \bar{u}) := E(\bar{v}, \bar{u}) \wedge g(\bar{v})$.

⁸Strictly speaking, the strategy maps *global* states to actions. However, as explained earlier, there exists a set $S \subseteq G / \sim_{\text{sharing}(\varphi, x)}^c$ such that, for all global states $g \in G$, if $l_{i\text{E}}(g) = l$ then $g \in S$. Thus, we can also interpret the strategy as a mapping from *local* states to actions.

⁹Observe that we do not need to allocate BDD variables \bar{v}' for the next value of the variable assignment because it does not change in a temporal transition (see Definition 4.11).

3. Pick an arbitrary extended state $\langle g, v \rangle \in Ext$. This is equivalent to selecting one conjunct (also referred to as *minterm*) $C(\bar{v}, \bar{u})$ from $E'(\bar{v}, \bar{u})$. BDD packages usually provide a built-in function for this operation¹⁰.
4. The conjunct $C(\bar{v}, \bar{u})$ encodes the extended state $\langle g, v \rangle$, where v contains the witness strategies. If we want to find the next action $f_i(g')$ of a strategy $f_i = v(x_i)$ for $0 \leq i < n$ at a global state $g' \in G$, we iterate over all possible actions $a \in \bigcap_{j \in \text{sharing}(\psi, x_i)} P_j(l_{jE}(g'))$ and pick one such that $C(\bar{v}, \bar{u}) \wedge a(\bar{u}_{v(x_i), g'})$ is not equivalent to falsity, where $a(\bar{u}_{v(x_i), g'})$ is the Boolean formula representing the fact that the next action of the strategy mapped to variable x_i at the global state g' is a .

Complexity

Having explained how the SLK model checking algorithm can be implemented symbolically, it is natural to ask what its complexity is. As we have just shown, the symbolic encoding requires polynomially many Boolean variables. Therefore, the algorithm runs in *exponential time* with respect to both the size of the model $|Z|$ and the number of quantified variables $|\text{vars}(\varphi)|$.

Theorem 4.3. Let $\mathcal{I} = \langle (L_i, Act_i, P_i, t_i)_{i \in Agt}, I, h \rangle$ be an arbitrary interpreted and $\varphi \in SLK$ be an SLK sentence. The worst case time complexity of the symbolic implementation is:

$$2^{O(|G| \times |\text{vars}(\varphi)| \times \log_2 |Act|)}$$

where G is the set of reachable global states of \mathcal{I} .

Proof (Sketch). The symbolic implementation uses:

- $O(\log_2 |G|)$ Boolean variables to represent the *current global state*;
- $O(\log_2 |G|)$ Boolean variables to represent the *next global state*;
- $O(\log_2 |Act|)$ Boolean variables to represent the *joint actions*;
- $O(|G| \times |\text{vars}(\varphi)| \times \log_2 (\max_{i \in Agt} |Act_i|))$ Boolean variables to represent the *variable assignment*.

The total number of Boolean variables needed is thus $O(|G| \times |\text{vars}(\varphi)| \times \log_2 |Act|)$. Since BDD operations take polynomial time with respect to the size of the relevant BDDs in the worst case [49], they take at most exponential¹¹ time in the number of Boolean variables. Hence, our claim follows. \square

The proof shows that the complexity of the algorithm is dominated by the encoding of the variable assignments. One of the main future directions in SLK verification is thus finding more compact representations of extended states.

Observe that the symbolic algorithm can use more than a polynomial amount of space because it performs operations on sets of extended states. While the procedure described in Subsection 4.2.1 uses only a polynomial amount of space (as it operates on individual states and assignments), we believe it is unlikely that such an explicit approach would outperform the symbolic algorithm.

4.3 Summary

In this chapter, we introduced a novel fragment of SL called Epistemic Strategy Logic (SLK), which is defined on imperfect recall semantics with incomplete knowledge (i.e. agents have no memory of the past and do not have complete knowledge of the global state). SLK combines LTL temporal operators, SL strategy quantifiers, and epistemic modalities. The result is a formalism which is even more expressive than SL because it can express combined concepts such as knowledge about Nash equilibria. We showed that, despite its expressive power, the SLK model checking problem belongs to the PSPACE complexity

¹⁰If no such function is available, we can skip step 3 ($C(\bar{v}, \bar{u}) := E'(\bar{v}, \bar{u})$) and refine the conjunct upon each lookup in step 4 ($C(\bar{v}, \bar{u}) := C(\bar{v}, \bar{u}) \wedge a(\bar{u}_{v(x_i), g'})$).

¹¹Given v Boolean variables, a BDD depending only on these variables can have at most 2^v nodes.

class. Furthermore, we provided an exponential-time model checking algorithm for SLK which admits an efficient symbolic implementation and supports witness/counterexample strategy synthesis. Hence, the algorithm can be used to automatically synthesise agents' behaviour which ensures an arbitrary SLK specification.

The first part of the chapter focused on the basic SLK *theory*. We gave the syntax and semantics of the logic, which is based on *uniform* strategies where the next actions assigned to two global states indistinguishable by an agent must be the same. This property ensures that the strategies are executable by the agents despite incomplete information. We then compared the logic to the original SL and discussed its limitations due to memoryless and non-behavioural strategies.

The second part of the chapter discussed SLK *model checking*. We started by defining the model checking problem formally and proved that it belongs to the PSPACE complexity class with respect to both the size of the model and the formula. We then provided a practical model checking algorithm for SLK based on the concept of extended states and showed that it is correct. Furthermore, we showed how it can be used to synthesise witness and counterexample strategies for SLK sentences of the form $\langle\langle x_0 \rangle\rangle \cdots \langle\langle x_{m-1} \rangle\rangle \psi$ and $\llbracket y_0 \rrbracket \cdots \llbracket y_{n-1} \rrbracket \psi$ respectively. Finally, we described how the algorithm can be implemented symbolically and showed that it has worst-case exponential time complexity.

In Section 6.2, we will describe how we developed an extension of MCMAS which implements the SLK model checking algorithm introduced in this chapter using BDDs.

Chapter 5

One-Goal Strategy Logic

One-Goal Strategy Logic (SL[1G]) is another syntactic fragment of SL (see Subsection 2.2.5), which was introduced in [72]. Unlike SLK, SL[1G] is defined on perfect recall semantics with complete information (i.e. agents have perfect memory of the past and complete knowledge of the global state). We provide the first practical model checking algorithm for SL[1G]. We prove its correctness and show that it has optimal worst-case time complexity. We explain how it can be extended to support general strategy synthesis and provide an efficient symbolic implementation. This chapter is split into two parts: Section 5.1 constitutes a very short revision of SL[1G] and Section 5.2 describes the new model checking algorithm we have developed.

5.1 Logic

As we have already mentioned in Section 3.2, SL[1G] is one of the syntactic fragments of full SL introduced in [72]. Since it is the least expressive fragment, it does not have the full power of SL and *cannot* express game-theoretic concepts like Nash equilibria. On the other hand, the “rewards” for this reduction in expressiveness are some very desirable properties [73]:

1. **Behavioural strategies.** Informally, SL[1G] strategies depend only on the history of the game and the next actions of other agents in the current state. More importantly they do *not* depend on other agents’ actions in the future or in other counterfactual games. While this might not seem very important, it is a fundamental property on which our model checking algorithm relies.
2. **2EXPTIME-COMPLETE model checking complexity.** While this is still a very high complexity, it is a massive improvement over the NONELEMENTARYSPACE-HARD complexity of full SL. More importantly, SL[1G] has exactly the *same complexity* as ATL* while being *strictly more expressive* [72]. Therefore, our new algorithm presented in Section 5.2 is also optimal¹ for model checking ATL*.

Note that we are referring to the complexity with respect to the size of the *formula*. The complexity of ATL* and all perfect recall fragments of SL with respect to the size of the *model* is P. Recall that SLK, which has imperfect recall, is in PSPACE with respect to *both* the size of the model and the formula (see Theorem 4.1).

3. **Decidable satisfiability.** While this property is not important for our purposes, it is still a desirable result. Note that SL[1G] is the only fragment of SL for which it is known that satisfiability is decidable (see Section 3.2 for an overview of the fragments and their properties).

Recall that SL[1G] uses LTL syntax augmented with an extra rule $\wp b\varphi$, where \wp is a quantification prefix (e.g. $\llbracket e \rrbracket \langle\langle x \rangle\rangle \llbracket y \rrbracket$), b is a binding prefix (e.g. $(E, e)(a, x)(b, y)$), and φ is an LTL formula. We first define the concept of a quantification and binding prefix formally [72].

¹However, this is not the case for certain fragments of ATL* like LTL, which is PSPACE-COMPLETE with respect to the size of the formula (see Table 2.1). While our algorithm can be used for model checking LTL, it is far from optimal.

Definition 5.1 (SL[1G] Prefixes). A *quantification prefix* over a set of variables $V \subseteq \text{Var}$ is a finite word $\wp \in \{\langle\langle x \rangle\rangle, \llbracket x \rrbracket \mid x \in V\}^{|V|}$ of length $|V|$ such that each variable $x \in V$ occurs in \wp exactly once. $QPre_V$ denotes the set of all quantification prefixes over V . $QPre \triangleq \bigcup_{V \subseteq \text{Var}} QPre_V$ is the set of all quantification prefixes.

A *binding prefix* over a set of variables $V \subseteq \text{Var}$ is a finite word $b \in \{(a, x) \mid a \in \text{Agt} \wedge x \in V\}^{|\text{Agt}|}$ of length $|\text{Agt}|$ such that each agent $i \in \text{Agt}$ occurs in b exactly once. $BPre_V$ denotes the set of all binding prefixes over V . $BPre \triangleq \bigcup_{V \subseteq \text{Var}} BPre_V$ is the set of all binding prefixes.

For example, $\wp_1 = \langle\langle z \rangle\rangle \llbracket x \rrbracket \langle\langle y \rangle\rangle$ is a valid quantification prefix over the set of variables $V = \{x, y, z\}$, i.e. $\wp_1 \in QPre_V$. However, $\wp_2 = \langle\langle x \rangle\rangle \langle\langle y \rangle\rangle$ and $\wp_3 = \langle\langle x \rangle\rangle \llbracket y \rrbracket \llbracket z \rrbracket \langle\langle x \rangle\rangle$ are not valid quantification prefixes because \wp_2 does not quantify $z \in V$ and \wp_3 contains $x \in V$ twice (also their lengths are not $|V| = 3$). Similarly, $b_1 = (b, x)(a, x)$ is a valid binding prefix over V when the set of agents is $\text{Agt} = \{a, b\}$, whereas $(a, x)(a, y)$ and $(a, x)(b, u)$ are not. Note that $BPre = BPre_{\text{Var}}$ by construction.

We are now ready to provide a formal definition of SL[1G] syntax [72].

Definition 5.2 (SL[1G] Syntax). SL[1G] *formulas* are built inductively from the set of atomic propositions AP , strategy variables Var , agents Agt , quantification prefixes $QPre$, and binding prefixes $BPre$, by using the following grammar, with $p \in AP$, $x \in \text{Var}$, $a \in \text{Agt}$, $b \in BPre$, and $\wp \in QPre$:

$$\wp ::= p \mid \top \mid \neg \wp \mid \wp \wedge \wp \mid \wp \vee \wp \mid \mathbf{X} \wp \mid \mathbf{F} \wp \mid \mathbf{G} \wp \mid \wp \mathbf{U} \wp \mid \wp b \wp$$

where $\wp \in QPre_{\text{free}(b\wp)}$. $SL[1G]$ denotes the infinite set of formulas generated by the above rules.

The conditions on \wp and b above ensure that $\wp b \wp$ is an SL[1G] sentence, i.e. $\text{free}(\wp b \wp) = \emptyset$. We must simply make sure that each agent is bound to a variable in b and that the variable is quantified in \wp . Assume that the set of agents is $\text{Agt} = \{a, b\}$. Then $\llbracket x \rrbracket \langle\langle y \rangle\rangle (a, x)(b, y) \mathbf{F} \mathbf{G} p$ is a well-formed SL[1G] formula (in fact, it is an SL[1G] sentence). On the other hand, $\llbracket x \rrbracket (a, x) q$ is not well-formed because it does not bind agent b to any strategy variable. A prefix is sometimes more readable when it interleaves quantifiers and bindings (e.g. $\llbracket x \rrbracket (a, x) \langle\langle y \rangle\rangle (b, y) \mathbf{X} r$). While this violates the definition of SL[1G] syntax, we will allow it as long as each binding occurs after the corresponding quantifier (e.g. $(a, x) \langle\langle x \rangle\rangle (b, x) s$ is not a well-formed SL[1G] formula because (a, x) appears before $\langle\langle x \rangle\rangle$).

SL[1G] is defined with respect to memoryful strategies. As we have explained in Chapter 3, incomplete information is undecidable under perfect recall semantics. Therefore, we have (again) two options: (i) consider SL[1G] with imperfect recall, or (ii) consider SL[1G] with complete information. The first option turns out not to be very interesting because we obtain a logic which is strictly less expressive than SLK while most probably still having the same complexity², namely PSPACE with respect to both the size of the model and the formula. Therefore, we consider SL[1G] with *perfect recall* and *complete information*. Consequently, SL[1G] uses *shared (memoryful) strategies* (Definition 2.21) instead of uniform shared memoryless strategies (Definition 4.2). Finally, since SL[1G] is a syntactic fragment of SL, the semantics remains the same (Definition 2.29).

5.2 Model Checking

In this section, we present a novel model checking algorithm for SL[1G], which we obtain by reducing the model checking problem to solving an infinite two-player parity game. The structure of this section is as follows: Subsection 5.2.1 introduces the algorithm based on the reduction, Subsection 5.2.2 discusses the complexity of the algorithm, Subsection 5.2.3 explains how the algorithm can be modified to support general strategy synthesis, Subsection 5.2.4 describes an efficient symbolic implementation of the algorithm, and Subsection 5.2.5 presents an optimisation technique for the algorithm called separate determinisation.

²There is a chance that a more efficient algorithm might be found for SL[1G] with imperfect recall. However, recall that although ATL with imperfect recall is not completely subsumed by SL[1G] with imperfect recall (see Subsection 4.2.1), it is less expressive. Since the complexity of model checking ATL with imperfect recall is P^{NP} -COMPLETE [23], there is most likely very little to be gained.

5.2.1 Algorithm

We introduce here the new algorithm for model checking arbitrary $SL[1G]$ formulas. Consider an interpreted system \mathcal{I} and an $SL[1G]$ sentence $\varphi \in SL[1G]$. To make the algorithm easier to understand, we assume that φ is a *principal sentence* of the form $\wp b\psi$ with $\wp \in QPre_{\text{free}(b\psi)}$ and $b \in BPre$. If this is not the case, we can simply add one quantifier and agent binding for each agent (it will not matter because φ is a sentence). Since our discussion will be based largely on $SL[1G]$ principal sentences, we shall define the concept formally.

Definition 5.3 ($SL[1G]$ Principal Sentence). Let \mathcal{I} be an interpreted system. An $SL[1G]$ formula $\varphi \in SL[1G]$ is a *principal sentence* iff (i) it is a sentence and (ii) it is of the form $\wp b\psi$ where $\wp \in QPre_{\text{free}(b\psi)}$ is a quantification prefix, $b \in BPre$ a binding prefix, and $\psi \in SL[1G]$ an $SL[1G]$ formula. Furthermore, φ is a *basic principal sentence* iff $\psi \in LTL$ is an LTL formula.

An $SL[1G]$ formula $\varphi' \in SL[1G]$ is a *principal subsentence* of φ iff φ' is a principal sentence and a subformula of φ . Furthermore, φ' is a *strict principal subsentence* of φ iff φ' is a principal subsentence of φ and $\varphi' \neq \varphi$. Finally, φ' is a *direct principal subsentence* of φ iff (i) φ' is a strict principal subsentence of φ and (ii) there exists no strict principal subsentence φ'' of φ such that φ' is a strict principal subsentence of φ'' .

Although we introduced a considerable number of new terms in the definition above, they should be easy to understand. To make sure that the precise meanings of all the concepts are clear, we provide a single comprehensive example³:

$$\varphi = \llbracket x \rrbracket(1, x)(2, x)G \underbrace{[-p_1 \wedge \neg p_2 \wedge \llbracket y \rrbracket \langle \langle z \rangle \rangle(1, y)(2, z)X \underbrace{(p_2 \wedge \neg \langle \langle u \rangle \rangle \langle \langle v \rangle \rangle(1, u)(2, v)F \neg p_2)}_{\varphi'}]}_{\varphi'}$$

There are three principal sentences, namely φ , φ' , and φ'' . φ'' is also the only *basic* principal sentence (both φ and φ' have non-LTL subformulas). The principal subsentences of φ are φ , φ' , and φ'' . However, only φ' and φ'' are *strict* principal subsentences of φ . Finally, both φ and φ' each have exactly one *direct* principal subsentence, namely φ' and φ'' (although in general a formula might have several direct principal subsentences). Note that φ'' cannot have any strict or direct principal subsentences because it is a basic principal sentence.

Recursive Approach

Our aim is now to find the set of all states $\|\varphi\|_{\mathcal{I}} \subseteq G$ at which an $SL[1G]$ principal sentence $\varphi = \wp b\psi$ holds, i.e. $\|\varphi\|_{\mathcal{I}} \triangleq \{g \in G \mid \mathcal{I}, \emptyset, g \models_{SL} \varphi\}$. We proceed in a recursive manner over the structure of φ . According to $SL[1G]$ syntax (Definition 5.2), ψ is a formula which combines atoms and *direct principal subsentences* of the form $\varphi' = \wp' b' \psi'$ using only Boolean and temporal connectives. For example:

$$\varphi = \underbrace{\langle \langle x \rangle \rangle}_{\wp} \underbrace{\llbracket y \rrbracket(a, x)(b, y)}_b \underbrace{[p \rightarrow X \underbrace{\langle \langle u \rangle \rangle \langle \langle v \rangle \rangle(a, u)(b, v)}_{\psi} F G q]}_{\varphi'}$$

Since φ' is a principal subsentence, we have $\text{free}(\varphi') = \emptyset$. Therefore, for all $g \in G$ and $\chi \in \text{Asg}$, we have $\mathcal{I}, \chi, g \models_{SL} \varphi'$ iff $\mathcal{I}, \emptyset, g \models_{SL} \varphi'$. Using our definition of $\|\varphi'\|_{\mathcal{I}}$, we get $\mathcal{I}, \chi, g \models_{SL} \varphi'$ iff $g \in \|\varphi'\|_{\mathcal{I}}$ for all $g \in G$ and $\chi \in \text{Asg}$. In other words, φ' holds in all states $\|\varphi'\|_{\mathcal{I}}$ (and no other states) regardless of the assignment χ . Therefore, we can do the following for each direct principal subsentence $\varphi' = \wp' b' \psi'$ of φ :

1. Calculate $\|\varphi'\|_{\mathcal{I}}$ (recursively);
2. Replace φ' in φ with a new atom $p_{\varphi'} \in AP$;
3. Update the assignment $h := h \cup \{p_{\varphi'}, \|\varphi'\|_{\mathcal{I}}\}$.

³The example is based on the toy model in Figure 3.1. We omit strategies for the environment for conciseness (it does not matter whether we use existential or universal quantifiers anyway). The meaning of the formula is roughly: "As long as both players perform the same actions, neither of them wins and for each player 1 strategy there is a player 2 strategy which ensures that player 2 wins in the next round and nothing can take his victory away after that".

This procedure preserves the truth value of φ . The following lemma asserts this formally.

Lemma 5.1. Let \mathcal{I} be an interpreted system and $\varphi = \wp b \psi$ an arbitrary SL[1G] principal sentence. Furthermore, let $\varphi^* = \wp b \psi^*$ be the principal sentence with all direct principal subsentences φ' of φ replaced with atoms $p_{\varphi'}$ such that $h^*(p_{\varphi'}) \triangleq \|\varphi'\|_{\mathcal{I}}$ and \mathcal{I}^* the interpreted system with the updated assignment h^* . Then for all global states $g \in G$ and assignments $\chi \in \text{Asg}$, we have $\mathcal{I}, \chi, g \models_{\text{SL}} \varphi$ iff $\mathcal{I}^*, \chi, g \models_{\text{SL}} \varphi^*$.

Proof (Sketch). The proof is performed by induction on the structure of φ . We show here only the case when two corresponding subformulas of φ and φ^* are different, i.e. when a principal subsentence φ' is replaced with a new atom $p_{\varphi'}$. Let φ' be an arbitrary principal subsentence of ψ . Then we have for all states $g \in G$ and assignments $\chi \in \text{Asg}$:

$$\begin{aligned} \mathcal{I}, \chi, g \models_{\text{SL}} \varphi' &\text{ iff } \mathcal{I}, \emptyset, g \models_{\text{SL}} \varphi' \\ &\text{ iff } g \in \|\varphi'\|_{\mathcal{I}} \\ &\text{ iff } g \in h^*(p_{\varphi'}) \\ &\text{ iff } \mathcal{I}^*, \emptyset, g \models_{\text{SL}} p_{\varphi'} \\ &\text{ iff } \mathcal{I}^*, \chi, g \models_{\text{SL}} p_{\varphi'} \end{aligned}$$

The first and the last equivalence follow from the fact that φ' and $p_{\varphi'}$ are SL[1G] sentences. The second one is the very definition of $\|\varphi'\|_{\mathcal{I}}$. The third one is true by construction of h^* . The fourth one follows from SL semantics (Definition 2.29). \square

Once we have replaced all principal subsentences of ψ , we are left with an SL[1G] *basic principal sentence* $\varphi^* = \wp b \psi^*$ where $\psi^* \in \text{LTL}$ is an LTL formula. Hence, we have reduced the problem of model checking an arbitrary SL[1G] *principal sentence* to recursive model checking of SL[1G] *basic principal sentences*. We shall now see how to solve this simpler problem.

Algorithm Overview

We outline here how the problem of model checking SL[1G] basic principal sentences $\varphi = \wp b \psi$ can be reduced to the problem of solving a two-player parity game. Let \mathcal{I} be an interpreted system. To calculate the set of global states at which ψ holds in \mathcal{I} , we proceed as follows:

1. **Formula automaton.** We construct a deterministic parity automaton $\mathfrak{B}_{\mathcal{I}}^{\psi}$ equivalent to the LTL formula ψ .
2. **Arena construction.** We construct a two-player arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ representing the global state space G and the interdependency of strategies in the prefix $\wp b$.
3. **Game combination.** We combine the arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ and parity automaton $\mathfrak{B}_{\mathcal{I}}^{\psi}$ into a two-player parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$. Solving the parity game yields its winning regions, which can be used to calculate $\|\varphi\|_{\mathcal{I}}$.

We shall now expand on each of the steps above. Unlike the SLK model checking algorithm we proposed in Subsection 4.2.2, our new SL[1G] model checking algorithm relies on the theory of ω -automata and infinite games. Please refer to Section 2.4 for a brief introduction into these areas.

Formula Automaton

The first step of our new model checking algorithm for SL[1G] is the construction of a deterministic parity automaton equivalent to the underlying LTL formula ψ of an SL[1G] basic principal formula $\varphi = \wp b \psi$. This is a very standard procedure, which is usually performed in three steps:

1. We construct a *non-deterministic generalised Büchi automaton* \mathfrak{A}_{ψ} equivalent to ψ . We do this using the standard translation presented in Figure 2.6. An explanation of the procedure with examples is provided in Subsection 2.4.3.

2. We convert \mathfrak{A}_ψ to a *non-deterministic Büchi automaton* \mathfrak{B}_ψ . As explained in Subsection 2.4.1, this can be done easily using an automaton product with the deterministic automaton shown in Figure 2.5.
3. We transform \mathfrak{B}_ψ into an equivalent *deterministic parity automaton* \mathfrak{P}_ψ using the determinisation procedure outlined in Subsection 2.4.4.

The automata \mathfrak{A}_ψ , \mathfrak{B}_ψ , and \mathfrak{P}_ψ above use sets of atomic propositions $A \in 2^{AP}$ as their alphabet. Given an interpreted system \mathcal{I} , we can convert them into automata $\mathfrak{A}_\mathcal{I}^\psi$, $\mathfrak{B}_\mathcal{I}^\psi$, and $\mathfrak{P}_\mathcal{I}^\psi$ on global states G by modifying their transition relations accordingly:

$$R_\mathcal{I}(s, g, s') \text{ iff } R(s, \{p \in AP \mid g \in h(p)\}, s') \text{ for all } (s, g, s') \in S \times G \times S'$$

It is easy to see that $\mathfrak{A}_\mathcal{I}^\psi$, $\mathfrak{B}_\mathcal{I}^\psi$, and $\mathfrak{P}_\mathcal{I}^\psi$ accept all infinite paths $\pi \in Pth$ in \mathcal{I} which satisfy ψ , i.e. $\mathcal{I}, \pi \models_{\text{LTL}} \psi$. Note that the resulting automata depend only on the model \mathcal{I} and the underlying LTL formula ψ , i.e. they are independent of the prefixes \wp and \flat . Nevertheless, we will sometimes refer to them as $\mathfrak{A}_\mathcal{I}^\wp$, $\mathfrak{B}_\mathcal{I}^\wp$, and $\mathfrak{P}_\mathcal{I}^\wp$ because it simplifies the notation.

To give the reader a more concrete idea of SL[1G] model checking and strategy synthesis, we will now introduce an example, which we will use several times to demonstrate key concepts. Consider the toy model in Figure 3.1 and the SL[1G] basic principal sentence $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle (E, e)(1, x)(2, y) \mathbf{G} [\neg p_1 \wedge \neg p_2]$, which means roughly: “*Whichever action player 1 performs, there exists an action for player 2 such that neither player will ever win*”. The quantification and binding prefixes are $\wp = \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle$ and $\flat = (E, e)(1, x)(2, y)$ respectively. The underlying LTL formula is $\psi = \mathbf{G} [\neg p_1 \wedge \neg p_2]$. The deterministic parity automaton $\mathfrak{P}_\mathcal{I}^\gamma$ equivalent to ψ in \mathcal{I} is shown in Figure 5.1a. We obtained it using the procedure outlined above. Note that it is not minimal because the states t_I and t_1 could be merged (with colour 0).

We would like to point out that the recursive step, i.e. replacing direct principal subsentences \wp' of \wp with atoms $p_{\wp'}$, can be incorporated into the procedure for constructing the non-deterministic generalised Büchi automaton $\mathfrak{A}_\mathcal{I}^\psi$. Assume that we are currently model checking an SL[1G] principal sentence $\varphi = \wp \flat \psi$ where ψ is not necessarily an LTL formula. We can construct the non-deterministic generalised Büchi automaton $\mathfrak{A}_\mathcal{I}^\psi$ straightaway (despite ψ possibly not being an LTL formula) by adding an extra case to the standard translation in Figure 2.6:

```

function GENBÜCHI( $\Phi$ )
  switch  $\Phi$  do
    ...
    case  $\wp' \flat' \psi'$ :
       $h := h \cup \{(p_\Phi, \llbracket \Phi \rrbracket_\mathcal{I})\}$            ▷ Solve recursively and update assignment.
       $\Phi := p_\Phi$                                    ▷ Replace subsentence with the new atom.
      return  $\mathcal{A}_\exists(\emptyset, AP, p_\Phi, \top, \emptyset)$ 
    ...
  end switch
end function


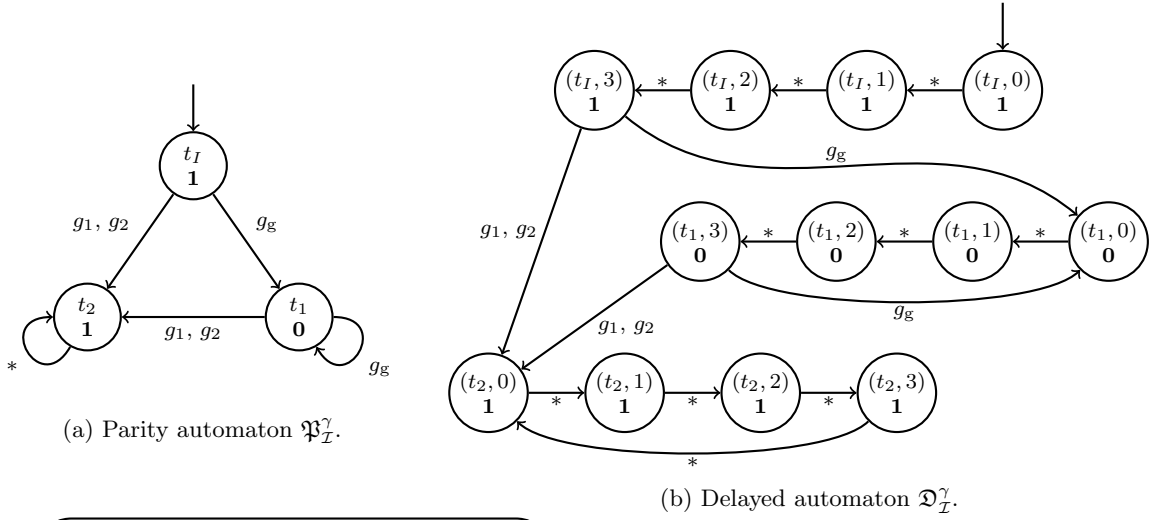
```

Hence, we actually do not need to transform φ to an SL[1G] basic principal sentence separately because the modified procedure will take care of all direct principal subformulas $\wp' \flat' \psi'$ automatically.

Arena Construction

The main idea of the whole transformation is to convert the model checking problem to a two-player game (see Subsection 2.4.5 for a brief introduction to games). Intuitively, model checking an SL[1G] basic principal sentence $\varphi = \wp \flat \psi$, i.e. deciding whether $\mathcal{I}, \emptyset, g \models_{\text{SL}} \varphi$ for some global state $g \in G$, can be imagined as a simple finite game between two players. The players take turns depending on the order of quantifiers in \wp and select the (whole) strategies for each variable in the quantifier:

- The *existential player*, selecting strategies for existentially quantified variables (e.g. $\langle\langle x \rangle\rangle$), is trying to satisfy the formula ψ ;
- The *universal player*, selecting strategies for universally quantified variables (e.g. $\llbracket y \rrbracket$), is trying to falsify the formula ψ .



(c) Formula arena \mathcal{A}_T^γ . The existential and universal player states are represented by squares and circles respectively.

Figure 5.1: Parity automaton \mathfrak{P}_T^γ , delayed automaton \mathfrak{D}_T^γ , and formula arena \mathcal{A}_T^γ of the toy model for the SL[1G] basic principal sentence $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle \langle y \rangle \rangle (E, e)(1, x)(2, y)G [\neg p_1 \wedge \neg p_2]$. A star represents an arbitrary global state.

For example, if $\varphi = \langle\langle x \rangle\rangle \llbracket y \rrbracket \langle\langle z \rangle\rangle$, then the game is as follows: the existential player selects a strategy for x , the universal player selects a strategy for y , and finally the existential player selects a strategy for z . Once all strategies have been selected, we construct the corresponding complete assignment $\chi \in \mathit{CAsg}$ (from the players' choices and the binding prefix \flat), and check the LTL formula ψ . Given that the way in which the strategies were selected by the players is in accordance with SL semantics (Definition 2.29), we have $\mathcal{I}, \emptyset, g \models_{\text{SL}} \varphi$ iff $\mathcal{I}, \chi, g \models_{\text{SL}} \psi$.

The finite game we outlined above is very simple and short. In fact, it will have exactly $|\varphi|$ rounds, one for each strategy variable. However, there are possibly infinitely many strategies and hence infinitely many moves that one of the players can make in each round. We see that the finite game idea does not make our problem any easier. Ideally, we would want the players to make simple choices (e.g. select a single action) as the game progresses. Instead of selecting the whole strategies up front, they would somehow build them gradually throughout the game. After all, this is much closer to the concept of a strategy in “games” that people play for fun. The question is now: Is such a reduction possible for SL[1G] model checking?

The answer is “yes”, we can reduce SL[1G] model checking of φ to solving an infinite two-player game where the players pick an action for each strategy in a given state. We will explain this reduction in more detail shortly. We first have to justify why we can do this. It turns out that this is the very reason why *behavioural semantics* of SL[1G] is so important. It was shown in [75] that for an arbitrary SL (including SL[1G]) sentence $\varphi\psi_{\text{SL}} \in \text{SL}$ where $\varphi \in \mathit{QPre}_{\text{free}(\psi_{\text{SL}})}$ and ψ_{SL} is an agent-closed SL formula, we have for each global state $g \in G$, $\mathcal{I}, \emptyset, g \models_{\text{SL}} \varphi\psi_{\text{SL}}$ iff there exists a dependence map θ_{SL} for φ over strategies such that $\mathcal{I}, g, \theta_{\text{SL}}(\chi) \models_{\text{SL}} \psi_{\text{SL}}$ for all $\chi \in \mathit{Asg}$ with $\llbracket \varphi \rrbracket = \text{dom}(\chi)$. A dependence map θ for the quantification prefix φ over strategies is a function from strategies for universally quantified variables in φ (denoted $\llbracket \varphi \rrbracket$) to strategies for all quantified variables in φ (denoted φ):

$$\left(\llbracket \varphi \rrbracket \rightarrow \underbrace{\left(\text{Trk} \rightarrow \bigcup_{i \in \text{Agt}} \text{Act}_i \right)}_{\text{strategies for universally quantified variables}} \right) \rightarrow \left(\varphi \rightarrow \underbrace{\left(\text{Trk} \rightarrow \bigcup_{i \in \text{Agt}} \text{Act}_i \right)}_{\text{strategies for all variables}} \right)$$

which satisfies the independence of variables. For example, if $\varphi = \llbracket [x] \langle\langle y \rangle\rangle \llbracket [z] \rrbracket$ then y is independent of z (because it appears before it) and x, z are independent of all variables (because they are universally quantified). This must be ensured by the dependence map (e.g. $\theta_{\text{SL}}(\chi)(z) = \chi(z)$). The problem with θ_{SL} is that the action of a strategy in one game might depend on the action of another strategy in another game (non-behavioural semantics). Fortunately, a stronger result holds for SL[1G] principal sentences $\varphi\flat\psi_{\text{SL}[1G]}$ [74]: There exists an *elementary* dependence map $\theta_{\text{SL}[1G]}$ for φ over strategies (iff the principal sentence holds in the state). A dependence map is elementary iff it can be transformed into an equivalent function (called *adjoint*) of the form:

$$\text{Trk} \rightarrow \left(\left(\llbracket \varphi \rrbracket \rightarrow \underbrace{\bigcup_{i \in \text{Agt}} \text{Act}_i}_{\substack{\text{next actions} \\ \text{for universally} \\ \text{quantified variables}}} \right) \rightarrow \left(\varphi \rightarrow \underbrace{\bigcup_{i \in \text{Agt}} \text{Act}_i}_{\substack{\text{next actions} \\ \text{for all variables}}} \right) \right)$$

Intuitively, this means that given the history of a play Trk , the next actions of strategies depend only on the next actions of the universally quantified strategies. This is exactly what we want! Before we return to the reduction, it should be pointed out that this (existence of an elementary dependence map) is the very definition of behavioural semantics⁴. Our explanation of dependence maps was quite informal. Please refer to [74, 75] for a more formal explanation and proofs.

We can therefore reduce the model checking of an SL[1G] basic principal sentence $\varphi\flat\psi$ to solving the following infinite game: Let $g \in G$ be the current global state. For each $0 \leq k < |\varphi|$ (in increasing order), the existential or universal player selects an action $a_k \in \text{Act}_{\text{sharing}(\flat\psi, \varphi(k))}$ for the strategy assigned to

⁴Formally, behavioural semantics are defined as follows [75]: Let \mathcal{I} be an interpreted system and $\varphi = \flat\psi$ an SL sentence with $\varphi \in \mathit{QPre}_{\text{free}(\psi)}$. Then for each global state $g \in G$, we have $\mathcal{I}, \emptyset, g \models_{\text{SLB}} \varphi$ iff there exists an elementary dependence map θ for the quantification prefix φ over strategies such that $\mathcal{I}, g, \theta(\chi) \models_{\text{SL}} \psi$ for all $\chi \in \mathit{Asg}$ with $\llbracket \varphi \rrbracket = \text{dom}(\chi)$. Notice the two different modelling relations “ \models_{SLB} ” (behavioural semantics) and “ \models_{SL} ” (SL semantics).

$\wp(k)$. Once all actions have been selected, a temporal transition according to the resulting joint action $a \in Act$ is performed and the new current state is $g' = t(g, a)$. This pattern of selecting actions and performing a temporal transition is repeated forever. Let $\pi = gg' \dots \in Pth$ be the infinite path obtained by this procedure. We define that the existential player wins the game iff the LTL formula ψ holds along π , i.e. $\mathcal{I}, \pi \models_{LTL} \psi$. The behavioural semantics of $SL[1G]$ then imply:

$$\mathcal{I}, \emptyset, g \models_{SL} \wp b \psi \quad \text{iff} \quad g \text{ is in the winning region of the existential player}$$

Essentially, we have just reduced $SL[1G]$ model checking to solving a two-player infinite game with an LTL winning condition.

There is one rather cosmetic issue with our formulation of the game, namely the fact that multiple decisions are made by both players within the same state (selecting actions for all strategies) before a move occurs. We shall make it more standard by inflating the state space and constructing an *arena* whose states are pairs (g, d) where $g \in G$ is the current state and d is a tuple of actions selected so far. Initially, d is empty and is gradually appended by the players as they choose actions for strategies. Once all actions have been selected, d represents a joint action, a temporal transition occurs, and the new state is $(t(g, d), \llbracket \rrbracket)$. A formal definition of the formula arena follows.

Definition 5.4 (Formula Arena). Let \mathcal{I} be an interpreted system and $\varphi \in SL[1G]$ a $SL[1G]$ principal sentence of the form $\wp b \psi$. We construct a *formula arena* $\mathcal{A}_{\mathcal{I}}^{\wp b}$ = (V_0, V_1, E) of \mathcal{I} for φ where:

- The states V of $\mathcal{A}_{\mathcal{I}}^{\wp b}$ are pairs $(g, d) \in G \times Dec_{\mathcal{I}}^{\wp b}$ such that for all $0 \leq i < |d|$ we have $d(i) \in \bigcap_{i \in \text{sharing}(b\top, \wp(k))} P_i(l_{iE}(g))$
- The existential player states $V_0 \subseteq V$ are states $(g, d) \in V$ such that $|d| < |\wp|$ and $\wp(|d|)$ is an existential strategy quantifier.
- The universal player states $V_1 \subseteq V$ are states $(g, d) \in V$ such that $|d| = |\wp|$ or $\wp(|d|)$ is a universal strategy quantifier.
- The edge relation is defined as:

$$E = \{((g, d), (g, d \cdot a)) \in V \times V \mid |d| < |\wp|\} \cup \{((g, d), (t(g, d^{Act}), \llbracket \rrbracket)) \in V \times V \mid |d| = |\wp|\}$$

where $Dec_{\mathcal{I}}^{\wp b} \triangleq \bigcup_{\ell=0}^{|\wp|} \prod_{k=0}^{\ell-1} \bigcup_{i \in \text{sharing}(b\top, \wp(k))} Act_i$ is the set of decisions and $d^{Act} \in Act$ is a joint action such that for all $0 \leq k < |\wp|$ and $i \in \text{sharing}(b\top, \wp(k))$ we have $a_i(d^{Act}) = d(k)$.

Note that we could allow decisions of length at most $|\wp| - 1$ (instead of $|\wp|$) and perform temporal transition immediately when the action for the last strategy is selected. The reason why we chose to include the extra step is that it nicely models non-determinism in case we use an evolution *relation* (instead of an evolution function). Intuitively, this is the same as adding an extra universally quantified variable at the end of \wp representing the final evolution choice. Again, we shall sometimes refer to a formula arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ as $\mathcal{A}_{\mathcal{I}}^{\wp b \psi}$ although it is independent of the underlying LTL formula ψ . Finally, notice that the definition uses the identity $\text{sharing}(b\top, \wp(k)) = \text{sharing}(b\psi, \wp(k))$ which can be shown for all $SL[1G]$ principal sentences $\wp b \psi$ and indices $0 \leq k < |\wp|$.

For example, the formula arena $\mathcal{A}_{\mathcal{I}}^{\gamma}$ of the toy model in Figure 3.1 and the $SL[1G]$ principal sentence $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle (E, e)(1, x)(2, y)G [\neg p_1 \wedge \neg p_2]$ introduced earlier is shown in Figure 5.1c. Let $f \in SStr_{\{E\}}$, $g \in SStr_{\{1\}}$, and $h \in SStr_{\{2\}}$ be the underlying strategies of the agents. Consider a game that starts in the global state g_g . The corresponding state of the arena is $(g_g, \llbracket \rrbracket)$, which is a universal player state corresponding to the quantifier $\llbracket e \rrbracket$. Only one action is available to the environment so the universal player must move to the state $(g_g, [i])$, i.e. set $f(g_g) = i$. Now it is again the universal player's turn. He has three options for $\llbracket x \rrbracket$, namely r, p, and s. Suppose he picks r. Thus, we move to the state $(g_g, [i, r])$ and have $g(g_g) = r$. The existential player now has the same three options for $\langle\langle y \rangle\rangle$. Suppose he picks s. We end up in the state $(g_g, [i, r, s])$ and have $h(g_g) = s$. Since all actions have been selected, we convert the decision $[i, r, s]$ to the joint action $a = rs$. As we are using an evolution function, there is exactly one next state that the universal player must move to, namely $(t(g_g, a), \llbracket \rrbracket) = (g_1, \llbracket \rrbracket)$. The whole process starts again. The universal player has to choose the action i for $\llbracket e \rrbracket$ so we move to the state $(g_1, [i])$ and

set $f(g_g g_1) = i$. Then the universal player has to choose an action for $\llbracket x \rrbracket$ and so on. We hope that this detailed example gives the reader a good understanding of the formula arena concept.

It remains to explain what the *winning condition* of the new game is and how it relates to SL[1G] model checking. Intuitively, the existential player wins an infinite play $(g_0, d_0^0) \dots (g_0, d_0^{\ell_0}) (g_1, d_1^0) \dots \in V^\omega$ in the formula arena $\mathcal{A}_{\mathcal{I}}^{\varphi^b}$ for an SL[1G] basic principal sentence $\varphi^b \psi$ iff the infinite path $g_0 g_1 \dots \in Pth$ satisfies the LTL formula ψ . The following definition expresses this formally.

Definition 5.5 (Pseudo-LTL Game). Let \mathcal{I} be an interpreted system and $\varphi^b \psi$ an SL[1G] basic principal sentence. The *pseudo-LTL game* $\mathfrak{L}_{\mathcal{I}}^{\varphi^b \psi}$ of \mathcal{I} for $\varphi^b \psi$ is a game $(\mathcal{A}_{\mathcal{I}}^{\varphi^b}, Win)$ where $\mathcal{A}_{\mathcal{I}}^{\varphi^b}$ is the formula arena of \mathcal{I} for $\varphi^b \psi$ and $Win \subseteq V^\omega$ is the winning set which contains all possible paths $\pi \in V^\omega$ such that $\mathcal{I}, \pi_{\mathcal{I}} \models_{LTL} \psi$. $\pi_{\mathcal{I}} \in Pth$ refers to the underlying path in \mathcal{I} , i.e. $\forall i \geq 0 \exists d \in Dec_{\mathcal{I}}^{\varphi^b} . (\pi_{\mathcal{I}}(i), d) = \pi((|\varphi| + 1)i)$.

Informally, $\pi_{\mathcal{I}}$ is equal to π “modulo” $(|\varphi| + 1)$ without the decisions. Consider again the formula arena $\mathcal{A}_{\mathcal{I}}^{\gamma}$ in Figure 5.1c of the toy model in Figure 3.1 and the SL[1G] principal sentence $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle (E, e)(1, x)(2, y)G [\neg p_1 \wedge \neg p_2]$. Let $\pi \in V^\omega$ be an infinite path in $\mathcal{A}_{\mathcal{I}}^{\gamma}$ which stays forever within the box labelled g_g , e.g. $\pi = (g_g, \square) (g_g, [i]) (g_g, [i, p]) (g_g, [i, p, p]) (g_g, \square) (g_g, [i]) (g_g, [i, s]) (g_g, [i, s, s]) \dots$. The underlying path in \mathcal{I} is $\pi_{\mathcal{I}} = g_g g_g \dots = g_g^\omega$. Since $\mathcal{I}, \pi_{\mathcal{I}} \models_{LTL} G [\neg p_1 \wedge \neg p_2]$, π is winning in $\mathfrak{L}_{\mathcal{I}}^{\gamma}$ by the definition above.

The important relationship between pseudo-LTL games and SL[1G] model checking is stated in the following lemma.

Lemma 5.2. Let \mathcal{I} be an interpreted system and $\varphi = \varphi^b \psi$ an SL[1G] basic principal sentence. Then for all global states $g \in G$ we have: $\mathcal{I}, \emptyset, g \models_{SL} \varphi$ iff (g, \square) is a winning state for the existential player in the pseudo-LTL game $\mathfrak{L}_{\mathcal{I}}^{\varphi^b \psi}$.

Proof (Sketch). The problem of model checking SL[1G] can be reduced to solving a so-called *dependence-vs-valuation game* [72]. In this game, the players alternate as follows: the existential player chooses a dependence map $\theta : (\llbracket \varphi \rrbracket \rightarrow \bigcup_{i \in Agt} Act_i) \rightarrow (\varphi \rightarrow \bigcup_{i \in Agt} Act_i)$ for φ over actions in the current state $g \in G$. Then the universal player chooses a valuation $v : \llbracket \varphi \rrbracket \rightarrow \bigcup_{i \in Agt} Act_i$. The combination $\theta(v) : \varphi \rightarrow \bigcup_{i \in Agt} Act_i$ assigns actions to all variables and determines the next state $g' \in G$. The existential player wins the game iff the infinite path $\pi = gg' \dots \in Pth$ satisfies the LTL formula ψ .

Instead of choosing the whole dependence map and valuation at once, the players in $\mathfrak{L}_{\mathcal{I}}^{\varphi^b \psi}$ assign actions to strategies one by one for each quantifier in φ . The order of players’ moves in $\mathfrak{L}_{\mathcal{I}}^{\varphi^b \psi}$ ensures that the independence constraints of the dependence map are satisfied (while not imposing any extra restrictions). Given this and the fact that the winning conditions are the same, the two games are equivalent. Hence, our claim follows. \square

Game Combination

We now have all the ingredients to construct the parity game. Let \mathcal{I} be an interpreted system and $\varphi = \varphi^b \psi$ an SL[1G] *basic* principal sentence. We have done the following so far:

1. We constructed a deterministic parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi} = (T, G, t_I, \delta, c)$ which accepts exactly those paths which satisfy the LTL formula ψ .
2. We constructed a formula arena $\mathcal{A}_{\mathcal{I}}^{\varphi^b} = (V_0, V_1, E)$ of \mathcal{I} for φ . The states of the arena are possible pairs of global states and decisions $V \subseteq G \times Dec_{\mathcal{I}}^{\varphi^b}$. The winning set of the corresponding pseudo-LTL game $\mathfrak{L}_{\mathcal{I}}^{\varphi^b \psi}$ contains all infinite paths $\pi \in V^\omega$ such that the underlying path in the interpreted system $\pi_{\mathcal{I}} \in Pth$ satisfies the LTL formula ψ .

Clearly, we want to somehow combine the automaton and the arena because $\mathfrak{P}_{\mathcal{I}}^{\psi}$ represents the winning condition of $\mathfrak{L}_{\mathcal{I}}^{\varphi^b \psi}$. However, we cannot simply take their product because, informally, they work at different, albeit constant, “speeds”. While $\mathfrak{P}_{\mathcal{I}}^{\psi}$ performs a temporal transition at every step, it takes exactly $|\varphi| + 1$ turns before a different underlying global state is reached by $\mathcal{A}_{\mathcal{I}}^{\varphi^b}$. To cater for this asynchrony, we can make the parity automaton “wait” for $|\varphi| + 1$ steps before each actual transition.

We do this by extending the state of $\mathfrak{P}_{\mathcal{I}}^{\psi}$ with a simple counter. A path $t_0 t_1 t_2 \dots \in T^{\omega}$ in $\mathfrak{P}_{\mathcal{I}}^{\psi}$ will then become the following path in the *delayed automaton* $\mathfrak{D}_{\mathcal{I}}^{\varphi\psi}$:

$$\underbrace{(t_0, 0) (t_0, 1) \dots (t_0, |\varphi| - 1) (t_0, |\varphi|)}_{|\varphi|+1 \text{ steps before } \delta \text{ transition}} (t_1, 0) (t_1, 1) \dots \in (T \times \{0, \dots, |\varphi|\})^{\omega}$$

This idea of adding a counter to the state of the parity automaton is captured in the following definition.

Definition 5.6 (Delayed Automaton). Let \mathcal{I} be an interpreted system, $\mathfrak{P}_{\mathcal{I}}^{\psi} = (T, G, t_I, \delta, c)$ a deterministic parity automaton for an $SL[1G]$ formula $\psi \in SL[1G]$ in \mathcal{I} , and $\varphi \in QPre$ a quantification prefix. We define a deterministic parity automaton $\mathfrak{D}_{\mathcal{I}}^{\varphi\psi} \triangleq (T \times \{0, \dots, |\varphi|\}, G, (t_I, 0), \delta_{\mathfrak{D}}, c_{\mathfrak{D}})$, called the *delayed automaton* for ψ over φ in \mathcal{I} , where the transition and colouring function are defined for all $t \in T$, $g \in G$, and $0 \leq k \leq |\varphi|$ as:

$$\delta_{\mathfrak{D}}((t, k), g) \triangleq \begin{cases} (t, k+1) & \text{if } k < |\varphi| \\ (\delta(t, g), 0) & \text{if } k = |\varphi| \end{cases} \quad c_{\mathfrak{D}}((t, k)) \triangleq c(t)$$

Similarly to the other automata, we shall sometimes refer to a delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\varphi\psi}$ as $\mathfrak{D}_{\mathcal{I}}^{\varphi}$ although it is independent of the binding prefix \flat . Informally, the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\varphi\psi}$ is now working at the “same speed” as the arena $\mathcal{A}_{\mathcal{I}}^{\varphi\flat}$. For example, Figure 5.1b shows the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\gamma}$ for the $SL[1G]$ principal sentence $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle (E, e)(1, x)(2, y)G [\neg p_1 \wedge \neg p_2]$ in the toy model in Figure 3.1. It was obtained by adding a counter with range 0–3 to the parity automaton $\mathfrak{P}_{\mathcal{I}}^{\gamma}$ in Figure 5.1a. Note that both the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\gamma}$ and the formula arena $\mathcal{A}_{\mathcal{I}}^{\gamma}$ in Figure 5.1c perform the same number of intermediate steps before a true transition occurs.

Let $(g_0, \square) \in G \times Dec_{\mathcal{I}}^{\varphi\flat}$ be some state of $\mathcal{A}_{\mathcal{I}}^{\varphi\flat}$. If we run both $\mathcal{A}_{\mathcal{I}}^{\varphi\flat}$ and $\mathfrak{D}_{\mathcal{I}}^{\varphi\psi}$ in parallel starting from (g_0, \square) and $(t_I, 0)$ respectively, we obtain the following infinite paths:

$$\begin{array}{ccc} (g_0, \square) \rightarrow (g_0, [a_0]) \rightarrow \dots \rightarrow (g_0, [a_0, \dots, a_{|\varphi|-1}]) & \xrightarrow{\text{temporal transition}} & (t(g_0, a), \square) \rightarrow \dots \\ (t_I, 0) \rightarrow (t_I, 1) \rightarrow \dots \rightarrow (t_I, |\varphi|) & \xrightarrow{\text{temporal transition}} & (\delta(t_I, g_0), 0) \rightarrow \dots \end{array}$$

(Dashed arrows indicate that the state $(g_0, [a_0, \dots, a_{|\varphi|-1}])$ is used for both transitions.)

where $a \triangleq [a_0, \dots, a_{|\varphi|-1}]^{Act}$ is the joint action implied by the accumulated decision. We can observe that when a temporal transition occurs, the underlying global state of the arena g_0 is used for two purposes simultaneously (dashed arrows): (i) it determines the next state of the arena together with the joint action a and (ii) it serves as an input to the parity automaton. Notice that the length of the decision tuple is always equal to the value of the counter. We are now ready to define the *combined parity game* $\mathfrak{G}_{\mathcal{I}}^{\varphi\psi}$.

Definition 5.7 (Combined Parity Game). Let \mathcal{I} be an interpreted system and $\varphi\flat\psi \in SL[1G]$ an $SL[1G]$ basic principal formula. Furthermore, let $\mathcal{A}_{\mathcal{I}}^{\varphi\flat} = (V_0, V_1, E)$ be the formula arena of \mathcal{I} for $\varphi\flat\psi$ with states $V \subseteq G \times Dec_{\mathcal{I}}^{\varphi\flat}$ and $\mathfrak{D}_{\mathcal{I}}^{\varphi\psi} = (T \times \{0, \dots, |\varphi|\}, G, (t_I, 0), \delta_{\mathfrak{D}}, c_{\mathfrak{D}})$ the delayed automaton for ψ over φ in \mathcal{I} . We construct a two-player *combined parity game* $\mathfrak{G}_{\mathcal{I}}^{\varphi\psi} \triangleq ((V_0 \times T, V_1 \times T, E_{\mathfrak{G}}), c_{\mathfrak{G}})$ of \mathcal{I} for $\varphi\flat\psi$ where for all states $(g, d, t) \in V \times T$, the colouring function is defined as $c_{\mathfrak{G}}((g, d, t)) \triangleq c_{\mathfrak{D}}((t, |d|)) = c(t)$. The transition relation is defined as follows:

$$E_{\mathfrak{G}} \triangleq \{((g, d, t), (g', d', t')) \in (V \times T) \times (V \times T) \mid E((g, d), (g', d')) \wedge \delta_{\mathfrak{D}}((t, |d|), g)\}$$

For all $g \in G$, we define $\text{start}(g) \triangleq (g, \square, t_I)$.

Figure 5.2 shows the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\gamma}$ of the toy model in Figure 3.1 for the $SL[1G]$ principal sentence $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle (E, e)(1, x)(2, y)G [\neg p_1 \wedge \neg p_2]$ introduced earlier. We obtained $\mathfrak{G}_{\mathcal{I}}^{\gamma}$ by combining the formula arena $\mathcal{A}_{\mathcal{I}}^{\gamma}$ in Figure 5.1c and the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\gamma}$ in Figure 5.1b. It highlights the winning regions and strategies of the existential player⁵, which we will discuss shortly.

Finally, we have everything we need to explain how $SL[1G]$ model checking can be reduced to solving two-player parity games. This reduction (together with its implementation) is most probably the *biggest achievement* of the whole project.

⁵Note that the existential and universal player are players 0 and 1 of the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\varphi\psi}$ respectively.

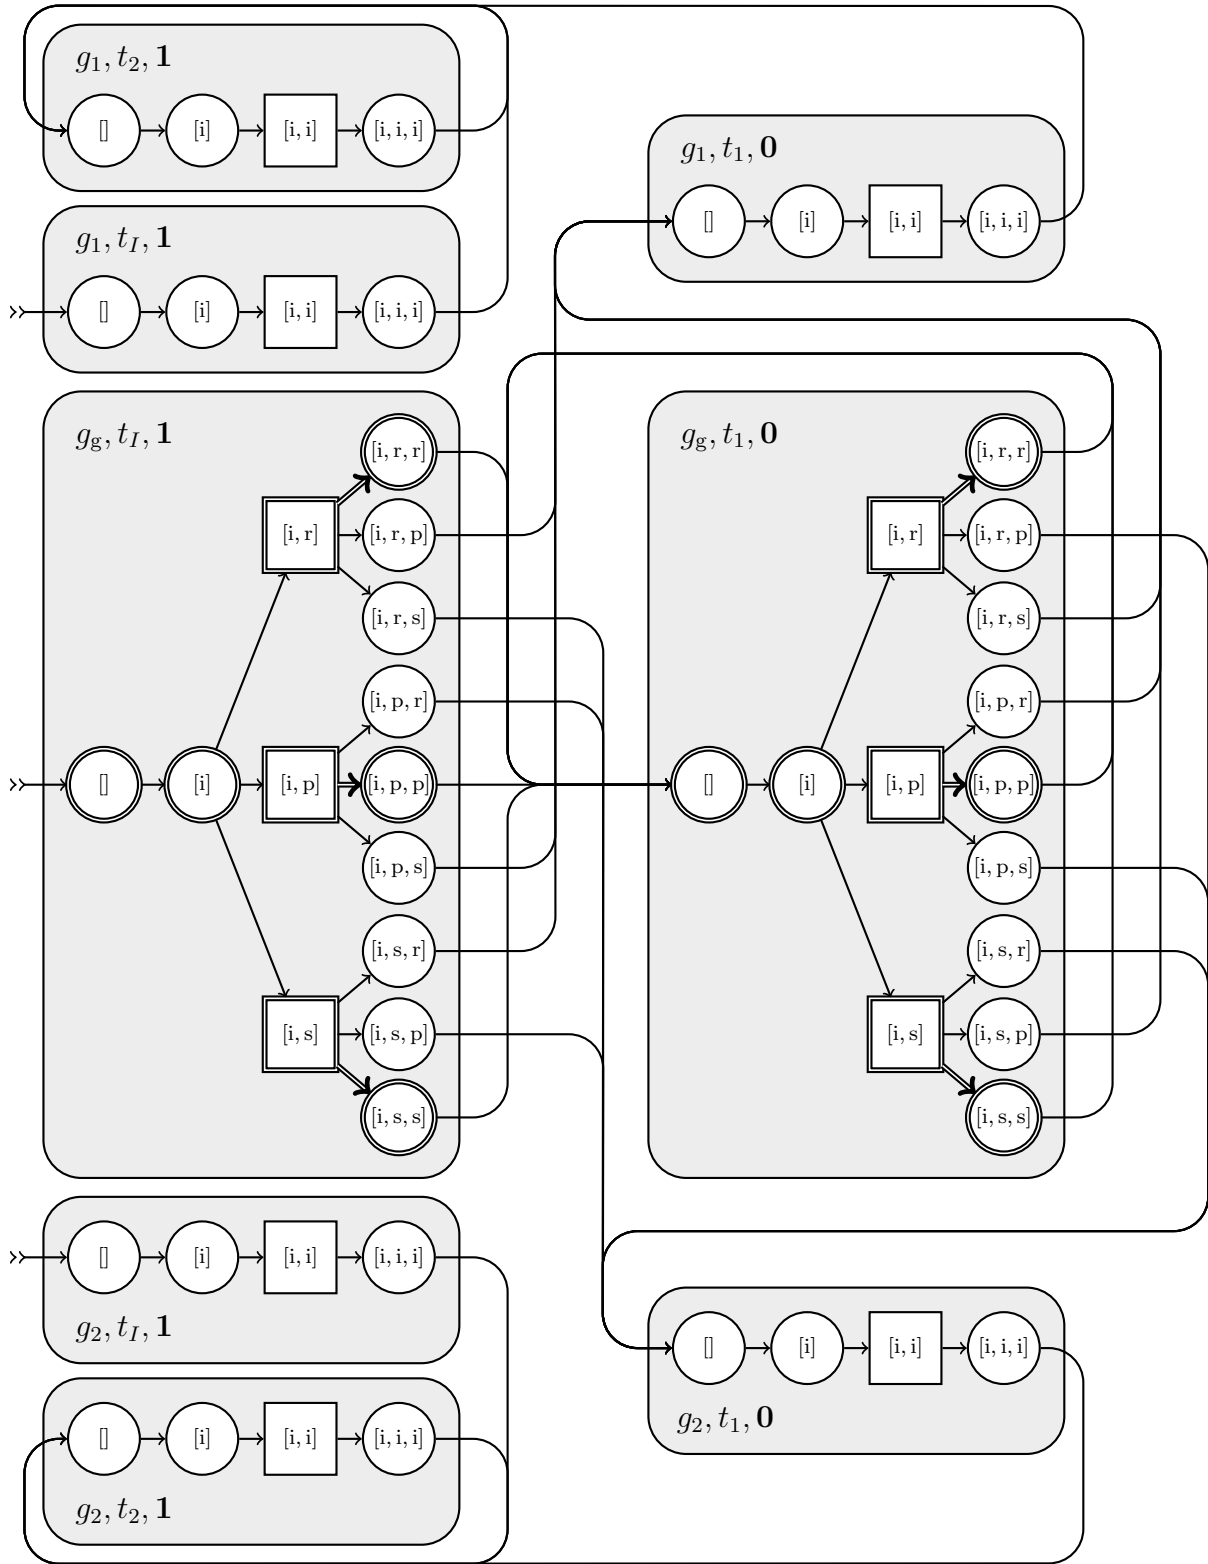


Figure 5.2: Combined parity game \mathfrak{G}_T^γ of the toy model for the SL[1G] basic principal sentence $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle y \rangle (E, e)(1, x)(2, y)G [-p_1 \wedge \neg p_2]$. Existential and universal player states are represented by squares and circles respectively. Winning states and strategies for the existential player have double edges. The universal player can play an arbitrary strategy in his winning region in order to win. The bold numbers (**0** and **1**) refer to the colours assigned by the colouring function. The three states on the left with incoming arrows are $\text{start}(g_1)$, $\text{start}(g_g)$, and $\text{start}(g_2)$ respectively.

Theorem 5.1. Let \mathcal{I} be an interpreted system, $\wp b\psi$ an SL[1G] basic principal sentence, and $g \in G$ an arbitrary global state. Then the following holds:

$$\mathcal{I}, \emptyset, g \models_{\text{SL}} \wp b\psi \quad \text{iff} \quad \text{start}(g) \in W_0(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi})$$

where $W_0(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi})$ is the winning region of the existential player in the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$.

Proof. The claim follows directly from our construction of the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$. However, given the importance of this result, we feel that a more direct proof should be provided. Essentially, we use the following two facts:

1. The delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp b\psi}$ accepts all infinite paths $\tau \in Pth$ such that the infinite path $\pi = \tau(|\wp|)\tau(2|\wp|+1)\dots\tau((k+1)|\wp|+k)\dots \in Pth$ satisfies the LTL formula ψ , i.e. $\mathcal{I} \models_{\text{LTL}} \psi$. This comes directly from the correctness of the existing techniques (the standard translation, transformation from generalised Büchi automata to Büchi automata, and the determinisation procedure) presented in Section 2.4.
2. The winning condition for a path π in the pseudo-LTL game $\mathfrak{L}_{\mathcal{I}}^{\wp b\psi}$, based on the arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$, is that the underlying path $\pi_{\mathcal{I}}$ in \mathcal{I} satisfies the LTL formula ψ . Lemma 5.2 asserts that (g, \square) is winning for the existential player in $\mathfrak{L}_{\mathcal{I}}^{\wp b\psi}$ iff $\mathcal{I}, \emptyset, g \models_{\text{SL}} \wp b\psi$.

We prove both directions of the equivalence separately:

- \Rightarrow : Assume that $\mathcal{I}, \emptyset, g \models_{\text{SL}} \wp b\psi$. Then by Lemma 5.2, (g, \square) is winning for the existential player in $\mathfrak{L}_{\mathcal{I}}^{\wp b\psi}$. This means that the existential player can enforce that the infinite path $\pi \in V^\omega$ starting in (g, \square) in the arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ is such that the underlying path $\pi_{\mathcal{I}} \in Pth$ in \mathcal{I} satisfies the LTL formula ψ . By correctness of the existing procedures, $\pi_{\mathcal{I}}$ is accepted by $\mathfrak{P}_{\mathcal{I}}^{\psi}$. This means that the infinite path constructed from the first elements of the tuples in π (i.e. a path in which each state from $\pi_{\mathcal{I}}$ occurs $|\wp|+1$ times in a row) is accepted by the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp b\psi}$ (by construction). Hence, the existential player can enforce an infinite path in $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$ starting in $\text{start}(g)$, which is winning for him (since the winner is decided by the parity condition of $\mathfrak{D}_{\mathcal{I}}^{\wp b\psi}$). Therefore, $\text{start}(g) \in W_0(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi})$.
- \Leftarrow : Assume that $\text{start}(g) \in W_0(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi})$. By construction of $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$, this means that the existential player can enforce an infinite path $\pi \in V^\omega$ starting in (g, \square) in the arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$, which is accepted by the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp b\psi}$. Given the shape of π (the first element of the tuple is always repeated $|\wp|+1$ times in a row) and the acceptance condition of $\mathfrak{D}_{\mathcal{I}}^{\wp b\psi}$ explained above, the underlying path $\pi_{\mathcal{I}} \in Pth$ in \mathcal{I} satisfies the LTL formula ψ . By definition of the winning condition (namely that $\mathcal{I}, \pi_{\mathcal{I}} \models_{\text{LTL}} \psi$), the existential player can enforce the path π in $\mathfrak{L}_{\mathcal{I}}^{\wp b\psi}$ (as it is based on $\mathcal{A}_{\mathcal{I}}^{\wp b}$), which is winning for him. Hence $\pi(0) = (g, \square)$ is winning for the existential player in $\mathfrak{L}_{\mathcal{I}}^{\wp b\psi}$. Finally, by Lemma 5.2, we have $\mathcal{I}, \emptyset, g \models_{\text{SL}} \wp b\psi$ as required. □

Therefore, calculating the winning regions $W_0(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi})$ and $W_1(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi})$ of the existential and universal player in the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$ yields the solution to the model checking problem. This can be done using existing algorithms for solving parity games such as the one in Figure 2.9.

Corollary 5.1. Let \mathcal{I} be an interpreted system and $\wp b\psi$ an SL[1G] basic principal sentence. The set of all global states $\|\wp b\psi\|_{\mathcal{I}} \subseteq G$ of \mathcal{I} in which $\wp b\psi$ holds is $\|\wp b\psi\|_{\mathcal{I}} = \left\{ g \in G \mid \text{start}(g) \in W_0(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}) \right\}$.

Consider again the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\gamma}$ in Figure 5.2 of the toy model in Figure 3.1 for the SL[1G] principal formula $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle \langle y \rangle \rangle (E, e)(1, x)(2, y)G [\neg p_1 \wedge \neg p_2]$. We have $\text{start}(g_g) \in W_0(\mathfrak{G}_{\mathcal{I}}^{\gamma})$ but $\text{start}(g_1), \text{start}(g_2) \notin W_0(\mathfrak{G}_{\mathcal{I}}^{\gamma})$. Therefore, by Theorem 5.1, $\mathcal{I}, \emptyset, g_g \models_{\text{SL}} \gamma$, $\mathcal{I}, \emptyset, g_1 \not\models_{\text{SL}} \gamma$, and $\mathcal{I}, \emptyset, g_2 \not\models_{\text{SL}} \gamma$. Hence, we have $\|\gamma\|_{\mathcal{I}} = \{g_g\}$. This result agrees with the intuitive meaning of the formula presented earlier.

We have now covered all steps of our new model checking algorithm for SL[1G]. The complete algorithm is shown in Figure 5.3. It expects an SL[1G] *principal sentence* (not necessarily basic). As


```

1 function CHECKSL[1G]( $\wp b\psi$ )
2    $\mathfrak{A}_{\mathcal{I}}^{\psi} := \text{GENBÜCHI}(\psi)$  ▷ Standard translation (Figure 2.6).
3    $\mathfrak{B}_{\mathcal{I}}^{\psi} := \text{BÜCHI}(\mathfrak{A}_{\mathcal{I}}^{\psi})$  ▷ Convert to a Büchi automaton (Subsection 2.4.1).
4    $\mathfrak{P}_{\mathcal{I}}^{\psi} := \text{PARITY}(\mathfrak{B}_{\mathcal{I}}^{\psi})$  ▷ Determinise (Subsection 2.4.4).
5    $\mathfrak{D}_{\mathcal{I}}^{\wp b\psi} := \text{DELAYED}(\mathfrak{P}_{\mathcal{I}}^{\psi}, \wp)$  ▷ Convert to a delayed automaton (Definition 5.6).
6    $\mathcal{A}_{\mathcal{I}}^{\wp b\psi} := \text{ARENA}(\mathcal{I}, \wp, \flat)$  ▷ Arena construction (Definition 5.4).
7    $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi} := \text{COMBINE}(\mathcal{A}_{\mathcal{I}}^{\wp b\psi}, \mathfrak{D}_{\mathcal{I}}^{\wp b\psi})$  ▷ Combine into a parity game (Definition 5.7).
8    $(W_0, W_1) := \text{SOLVEPARITY}(\mathfrak{G}_{\mathcal{I}}^{\wp b\psi})$  ▷ Solve parity game (Figure 2.9).
9   return  $\{g \in G \mid \text{start}(g) \in W_0\}$  ▷ Return  $\|\wp b\psi\|_{\mathcal{I}}$  (Corollary 5.1).
10 end function

```

Figure 5.3: The model checking algorithm for SL[1G] with references to the relevant parts of this thesis. The argument $\wp b\psi$ is an arbitrary SL[1G] principal sentence. The procedure GENBÜCHI on line 2 calls CHECKSL[1G] recursively on direct principal subformulas of $\wp b\psi$ and replaces them with new atoms as explained on page 89. Therefore, $\wp b\psi$ can be treated as a *basic* principal sentence from line 3 onwards.

explained earlier, an arbitrary SL[1G] sentence ψ can be converted to an equivalent principal sentence as follows:

$$[\langle\langle x_i \rangle\rangle]_{i \in \text{Agt}} [(i, x_i)]_{i \in \text{Agt}} \psi$$

For example, the SL[1G] sentence $p_1 \vee p_2$ for the toy model in Figure 3.1 can be converted into an equivalent principal sentence $\langle\langle x_E \rangle\rangle \langle\langle x_1 \rangle\rangle \langle\langle x_1 \rangle\rangle (E, x_E)(1, x_1)(2, x_2)(p_1 \vee p_2)$, which can then be passed as input to the algorithm in Figure 5.3.

The algorithm is further discussed in the rest of this chapter: Its complexity is discussed in Subsection 5.2.2. Subsection 5.2.3 explains how it can be used for the purposes of strategy synthesis. An efficient symbolic implementation of the algorithm is provided in Subsection 5.2.4. Finally, an optimisation for improving its performance called separate determinisation is presented in Subsection 5.2.5.

5.2.2 Complexity

In this subsection, we will discuss the complexity of the model checking algorithm presented in the previous subsection and compare it with the theoretical complexity of the SL[1G] model checking problem. Recall that the decision problem is as follows:

Given an interpreted system \mathcal{I} , a global state $g \in G$, and an SL[1G] formula $\varphi \in \text{SL}[1G]$, determine whether $\mathcal{I}, \emptyset, g \models_{\text{SL}[1G]} \varphi$.

It has been shown that the theoretical complexity of this problem is 2EXPTIME-COMPLETE with respect to the size of the formula $|\varphi|$ and P-COMPLETE with respect to the size of the model $|\mathcal{I}|$. We argue that our new algorithm satisfies both bounds.

Theorem 5.2. Let \mathcal{I} be an interpreted system and $\varphi \in \text{SL}[1G]$ an SL[1G] sentence. Our model checking algorithm calculates the set of all global states $\|\varphi\|_{\mathcal{I}} \subseteq G$ satisfying φ in time $|\mathcal{I}|^{2^{O(|\varphi|)}}$.

Proof. We start by considering an arbitrary SL[1G] *basic principal sentence* $\wp b\psi$ and show that our claim holds. The model checking algorithm (see Figure 5.3) proceeds as follows:

1. It constructs a non-deterministic generalised Büchi automaton $\mathfrak{A}_{\mathcal{I}}^{\psi}$ with $O(2^{|\psi|})$ states (see Proposition 2.2).
2. It converts $\mathfrak{A}_{\mathcal{I}}^{\psi}$ to a non-deterministic Büchi automaton $\mathfrak{B}_{\mathcal{I}}^{\psi}$. $\mathfrak{A}_{\mathcal{I}}^{\psi}$ can have up to $O(|\psi|)$ fairness constraints, so the resulting automaton will have $O(|\psi| \times 2^{|\psi|}) = 2^{O(|\psi|)}$ states.
3. It transforms $\mathfrak{B}_{\mathcal{I}}^{\psi}$ to a deterministic parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi}$ with $2^{(2^{O(|\psi|)} + 1)^2} = 2^{2^{O(|\psi|)}}$ states and $2 \times 2^{O(|\psi|)} = 2^{O(|\psi|)}$ colours (see Proposition 2.5).
4. It transforms $\mathfrak{P}_{\mathcal{I}}^{\psi}$ to the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp b\psi}$ with $|\wp| \times 2^{2^{O(|\psi|)}}$ states.

5. It constructs the arena $\mathcal{A}_{\mathcal{I}}^{b\varphi}$ with $O(|\mathcal{I}| \times |\varphi|)$ states.
6. It combines $\mathcal{A}_{\mathcal{I}}^{b\varphi}$ and $\mathcal{D}_{\mathcal{I}}^{\varphi\psi}$ into the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\varphi\psi}$ with $|\mathcal{I}| \times |\varphi| \times 2^{2^{O(|\psi|)}}$ states and $2^{O(|\psi|)}$ colours.
7. Finally, it solves the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\varphi\psi}$. The algorithm for solving parity games in Figure 2.9 has time complexity $O(m \times (n/d)^d)$ (Proposition 2.10), where m is the number of edges, n the number of vertices, and d the number of colours. Hence, solving $\mathfrak{G}_{\mathcal{I}}^{\varphi\psi}$ will take time:

$$\underbrace{\left[|\mathcal{I}| \times |\varphi| \times 2^{2^{O(|\psi|)}} \right]^2}_{\text{edges}} \times \underbrace{\left[\frac{|\mathcal{I}| \times |\varphi| \times 2^{2^{O(|\psi|)}}}{2^{O(|\psi|)}} \right]^{2^{O(|\psi|)}}}_{\text{vertices/colours}} = |\mathcal{I}|^{2^{O(|\varphi\psi|)}}$$

All the construction steps (1–6) of the algorithm fit into the $|\mathcal{I}|^{2^{O(|\varphi\psi|)}}$ time bound. Therefore, the time complexity of our model checking algorithm on SL[1G] *basic principal sentences* is $|\mathcal{I}|^{2^{O(|\varphi\psi|)}}$.

Let us now come back to φ , which is an arbitrary SL[1G] sentence (not necessarily principal). We will model check φ in a recursive bottom-up manner as explained in Subsection 5.2.1. Hence, we will model-check at most $|\varphi|$ SL[1G] basic principal sentences of size at most $|\varphi|$. Furthermore, if φ is not a principal sentence, then it must be a Boolean combination⁶ of some principal sentences, the results of which we can combine using set operations (instead of adding quantifiers and bindings so that $\varphi' b' \varphi$ would be a principal sentence). Therefore, φ can be checked in time:

$$\underbrace{|\varphi| \times \underbrace{|\mathcal{I}|^{2^{O(|\varphi|)}}}_{\text{basic principal sentence}}}_{\text{arbitrary sentence}} = |\mathcal{I}|^{(\log_{|\mathcal{I}|} |\varphi|) + 2^{O(|\varphi|)}} \leq |\mathcal{I}|^{(\log_{|\mathcal{I}|} |\varphi|) \times 2^{O(|\varphi|)}} = |\mathcal{I}|^{2^{(\log_2 \log_{|\mathcal{I}|} |\varphi|) + O(|\varphi|)}} = |\mathcal{I}|^{2^{O(|\varphi|)}} \quad (5.1)$$

Hence, our claim follows. \square

This is a very positive result because our algorithm has the same time complexity as the theoretical algorithm for model checking ATL* presented in [23]. However, it should be noted that the proof above is rather “pessimistic” in the sense that it assumes that each principal subsentence has size $|\varphi|$. In fact, a lower worst-case complexity can be derived:

$$|\varphi| + \sum_{\varphi' \in \text{sub}(\varphi)} |\mathcal{I}|^{2^{O(\text{temp}(\varphi'))}}$$

where $\text{sub}(\varphi)$ is the set of all principal subsentences of φ (including φ itself) and $\text{temp}(\varphi')$ is the number of temporal operators in φ' (excluding strict principal subformulas). The reason why we can use $\text{temp}(\varphi')$ instead of $|\varphi'|$ is that a non-deterministic generalised Büchi automaton for an LTL formula with $\text{temp}(\varphi')$ temporal operators has at most $2^{2 \times \text{temp}(\varphi') + 1}$ states (see Proposition 2.2).

5.2.3 Strategy Synthesis

The model checking algorithm presented in Subsection 5.2.1 decides *whether* a given SL[1G] formula holds in a given model. In addition, we would also like to know *why* that is the case, i.e. what are the strategies that make the formula true. For example, if the formula $\llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle (E, e)(1, x)(2, y) \mathbf{G} [\neg p_1 \wedge \neg p_2]$ is true in a given state, we want to know how the strategy assigned to y depends on the strategies assigned to e and x . Conversely, if the formula is false in some state, we want to know the corresponding strategies for the universally quantified variables e and x .

Let \mathcal{I} be an interpreted system and $\varphi b \psi$ an SL[1G] principal sentence. Our algorithm reduces the SL[1G] model checking problem to solving a parity game $\mathfrak{G}_{\mathcal{I}}^{\varphi b \psi}$. We then find its winning regions

⁶This follows from SL[1G] syntax (Definition 5.2) and the fact that $\text{free}(\varphi) = \emptyset$ (if φ combined its direct principal subsentences using temporal operators, we would have $\text{free}(\varphi) = \text{Agt}$).

(W_0, W_1) using the existing algorithm for solving parity games presented in Figure 2.9. As explained in Subsection 2.4.5, the algorithm can also be used to construct the winning strategies for both players. By Proposition 2.8, the winning strategies for $\mathfrak{G}_T^{\wp b \psi}$ are *memoryless*, i.e. they are partial functions $w_\sigma : V_\sigma \times T \rightarrow V \times T$ such that $\text{dom}(w_\sigma) = W_\sigma \cap (V_\sigma \times T)$ for players $\sigma \in \{0, 1\}$. Since the domains of w_0 and w_1 are disjoint, we can regard them as one partial function $w : V \times T \rightarrow V \times T$ with $\text{dom}(w) = W_0 \cap (V_0 \times T) \cup W_1 \cap (V_1 \times T)$ defined as $w \triangleq w_0 \cup w_1$. Intuitively, w assigns a winning move $w(s) \in V \times T$ to every state $s \in V \times T$ that is in the winning region of the player that it belongs to, i.e. either $W_0 \cap (V_0 \times T)$, or $W_1 \cap (V_1 \times T)$.

Clearly, the winning strategies for the game somehow encode the $\text{SL}[1G]$ shared strategies, which we want to synthesise. For example, the winning strategy for the existential player in the combined parity game \mathfrak{G}_T^γ in Figure 5.2 requires that he matches the universal player's move (notice the emphasised transitions of the form $[i, a] \Rightarrow [i, a, a]$ with $a \in \{r, p, s\}$). Nevertheless, there are important differences between the two types of strategies:

1. The *game strategies* are memoryless, partial, and map game states to game states ($w_\sigma : V_\sigma \times T \rightarrow V \times T$ with $\sigma \in \{0, 1\}$). There are always exactly two of them, one for the existential player w_0 and one for the universal player w_1 .
2. The *agent strategies* are memoryful, total, and map non-empty finite sequences of global states (tracks) to actions ($f_k : \text{Trk} \rightarrow \text{Act}_{\text{sharing}(b\psi, \wp(k))}$ with $0 \leq k < |\wp|$). There is one for each quantifier in the quantifier prefix \wp .

Let us start with the memory. Could the agent strategies also always be memoryless? No, unfortunately not. To see why, have a look at the example in Subsection 4.1.5, where a very simple SLK formula (which is also an $\text{SL}[1G]$ formula) cannot be satisfied using memoryless strategies. Intuitively, the difference is that the game already “contains” the deterministic parity automaton \mathfrak{P}_T^ψ , which defines the winning condition. The players can thus use its underlying state to select their moves “for free” because it is a part of the arena. This means that if the agents simulate the progress of \mathfrak{P}_T^ψ in their memory, they will have almost the same information as they have in $\mathfrak{G}_T^{\wp b \psi}$. However, they might also need to remember some of the other agents' actions, which are part of the game state space as well.

The issue of w_σ being a partial function is a rather cosmetic one. If it is undefined in some state $s \in V_\sigma \times T$, i.e. $s \notin \text{dom}(w_\sigma)$, it just means that player σ cannot enforce a victory from s . In that case, we can assign an arbitrary move $s' \in V \times T$ such that $E_{\mathfrak{G}}(s, s')$ to player σ in state s . We shall denote this a *total game strategy* \hat{w}_σ .

Definition 5.8 (Total Strategy). Let \mathfrak{G} be a memoryless determined game with an underlying arena $\mathcal{A} = (V_0, V_1, E)$ with no dead ends, i.e. E is serial. Furthermore, let the winning strategies of \mathfrak{G} be (w_0, w_1) . Then a *total strategy* for player $\sigma \in \{0, 1\}$ in \mathfrak{G} is a total function $\hat{w}_\sigma : V_\sigma \rightarrow V$ such that for all $v \in V_\sigma$, we have $E(v, \hat{w}_\sigma(v))$ and if $v \in \text{dom}(w_\sigma)$, then $\hat{w}_\sigma(v) = w_\sigma(v)$. The *total solution* of \mathfrak{G} is a total function $\hat{w} : V \rightarrow V$ such that $\hat{w} \triangleq \hat{w}_0 \cup \hat{w}_1$.

The third difference is that the output of the game strategies are game states, whereas the agent strategies return actions. Let us have a closer look at what the input an output of \hat{w} look like when selecting an action for the quantifier $\wp(k)$ with $0 \leq k < |\wp|$:

$$\hat{w}((g, [a_0, \dots, a_{k-1}], t)) = (g, [a_0, \dots, a_{k-1}, a_k], t)$$

It merely appends the action $a_k \in \text{Act}_{\text{sharing}(b\psi, \wp(k))}$, which the strategy $\wp(k)$ will do in the current state $g \in G$, to the tuple of actions $[a_0, \dots, a_{k-1}]$ of the strategy variables it depends on $(\wp(0), \dots, \wp(k-1))$. Intuitively, the equality above implies that the next action which the strategy assigned to the variable $\wp(k)$ selects will depend only on the global state $g \in G$, the underlying parity automaton state $t \in T$, and the next actions selected by the preceding strategies in the quantifier prefix. Recall that the underlying global and parity automaton states g and t in the equation above change only when a temporal transition occurs (i.e. every $|\wp| + 1$ steps in the game). Moreover, the next parity automaton state $t' = \delta(t, g)$ does not depend on the accumulated decisions $[a_0, \dots, a_{|\wp|-1}]$ of the players. Hence, if we regard the states of the parity automaton T as a memory, we get the following *memory update* and *next action* functions

for variable $\wp(k)$ with $0 \leq k < |\wp|$:

$$\delta_k : T \times G \rightarrow T \quad \alpha_k : T \times G \times \underbrace{\left[\prod_{l=0}^{k-1} Act_{\text{sharing}(b\psi, \wp(l))} \right]}_{\text{next actions of strategies } \wp(0), \dots, \wp(k-1)} \rightarrow \underbrace{Act_{\text{sharing}(b\psi, \wp(k))}}_{\text{next action of strategy } \wp(k)}$$

This is very similar to a finite memory strategy (Definition 2.56). However, α_k depends on the next actions of other strategies in the current state. Therefore, strictly speaking, the strategy might require an infinite amount of memory. To see why this is the case, consider again the SL[1G] formula for the toy model in Figure 3.1 $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle \langle y \rangle \rangle G \neg (p_1 \vee p_2)$. It clearly holds in the initial state g_g : Player 2 just always plays the same action as player 1. Hence, for all strategies x for player 1, there exists a corresponding strategy y for player 2, namely $y = x$, which ensures this property. However, consider the case when player 1 plays rock when the current round number is a square and scissors otherwise:

0	1		4		9	
r	r	s	s	r	s	s
			s	s	s	s
			r		s	r
						s
						...

Clearly, player 2 can satisfy the property by applying the exact same strategy. However, such a strategy is not finite memory because it must somehow keep track of the number of rounds until the next square, which can become arbitrarily large. Nevertheless, the strategy can be regarded as having finite memory in some sense if we include the next actions of other strategies in the input to the next action function (exactly like we did for α_k).

To sum up, the next action $f_k(\pi)$ of the synthesised strategy for variable $\wp(k)$ depends on: (i) the track through the interpreted system so far $\pi \in Trk$ and (ii) the next actions $f_0(\pi), \dots, f_{k-1}(\pi)$ of the strategies for variables $\wp(0), \dots, \wp(k-1)$. The following definition expresses this idea formally.

Definition 5.9 (SL[1G] Strategy Synthesis). Let \mathcal{I} be an interpreted system, $\wp b\psi$ an SL[1G] principal sentence, $0 \leq k < |\wp|$ an integer, and f_0, \dots, f_{k-1} shared strategies for variables $\wp(0), \dots, \wp(k-1)$. Furthermore, let $\mathfrak{P}_{\mathcal{I}}^{\psi} = (T, G, t_I, \delta, c)$ be the deterministic parity automaton for ψ and (w_0, w_1) the winning strategies of the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$. Then we define the strategy $f_k : Trk \rightarrow Act_{\text{sharing}(b\psi, \wp(k))}$ for variable $\wp(k)$ for all tracks $\pi \in Trk$ implicitly as:

$$\hat{w}(\langle \text{last}(\pi), [f_0(\pi), \dots, f_{k-1}(\pi)], \delta(t_I, \pi_{\leq |\pi|-2}) \rangle) = \langle \text{last}(\pi), [f_0(\pi), \dots, f_{k-1}(\pi), \underline{f_k(\pi)}], \delta(t_I, \pi_{\leq |\pi|-2}) \rangle$$

where $\delta(t_I, \pi_{\leq |\pi|-2}) \triangleq \delta(\dots \delta(\delta(t_I, \pi(0)), \pi(1)) \dots, \pi(|\pi|-2))$.

While the definition above might look very complex, it is simply reusing the winning strategies of the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$. Once more, consider the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\gamma}$ in Figure 5.2 of the toy model in Figure 3.1 for the formula $\gamma \triangleq \llbracket e \rrbracket \llbracket x \rrbracket \langle \langle y \rangle \rangle (E, e)(1, x)(2, y)G [\neg p_1 \wedge \neg p_2]$. Assume that the history of the game so far is $\pi = g_g g_g g_g$. Hence, the current global state is $\text{last}(\pi) = \pi(2) = g_g$. As the environment has only one available action in g_g , we know that $f_0(\pi) = i$. Let us further assume that player 1 has played scissors, i.e. $f_1(\pi) = s$. The decision accumulated so far is thus $[f_0(\pi), f_1(\pi)] = [i, s]$ and the state of the parity automaton is $\delta(t_I, \pi) = \delta(\delta(t_I, g_g), g_g) = t_1$. The current implied state of the game $\mathfrak{G}_{\mathcal{I}}^{\gamma}$ in Figure 5.2 is therefore $(g_g, [i, s], t_1)$. We can see that this state is winning for the existential player and we have $w_0((g_g, [i, s], t_1)) = (g_g, [i, s, \underline{s}], t_1)$. Hence, $f_2(\pi) = s$, i.e. player 2 should also play scissors. This agrees with our expectation.

We shall now state the relationship between SL[1G] model checking and strategy synthesis. Informally, if an SL[1G] principal formula $\wp b\psi$ holds in a global state $g \in G$ and we are given arbitrary strategies for the universally quantified variables in \wp , the strategies for the existentially quantified variables synthesised⁷ according to Definition 5.9 ensure that the formula ψ holds in g .

Theorem 5.3. Let \mathcal{I} be an interpreted system and $\wp b\psi$ an SL[1G] principal formula. Furthermore, for each index $0 \leq k < |\wp|$ (in increasing order), let f_k be (i) an arbitrary shared strategy $f_k \in$

⁷Note that although we assume that we are given strategies for all universally quantified variables up front (before we synthesise the existential ones), Definition 5.9 ensures that the synthesised strategy for a variable $\wp(i)$ does not depend on the strategy for a variable $\wp(j)$ with $i < j$.

$SStr_{\text{sharing}(b\psi, \wp(k))}$ if $\wp(k)$ is a *universal* quantifier or (ii) a strategy synthesised according to Definition 5.9 if $\wp(k)$ is an *existential* quantifier. Then for all global states $g \in G$:

$$\text{if } \mathcal{I}, \emptyset, g \models_{\text{SL}} \wp b\psi \quad \text{then } \mathcal{I}, \chi, g \models_{\text{SL}} \psi$$

where $\chi \in \text{CAsg}$ is a complete assignment with $\text{dom}(\chi) = \text{Agt}$ such that for all indices $0 \leq k < |\wp|$ and agents $i \in \text{sharing}(b\psi, \wp(k))$ we have $\chi(i) \triangleq f_k$.

Proof. Let $\mathfrak{P}_{\mathcal{I}}^{\psi} = (T, G, t_I, \delta, c)$ be the deterministic parity automaton equivalent to the LTL formula ψ and $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi} = ((V_0 \times T, V_1 \times T, E_{\mathfrak{G}}), c_{\mathfrak{G}})$ the combined parity game in \mathcal{I} for $\wp b\psi$ with winning regions (W_0, W_1) and strategies (w_0, w_1) . Note that we can treat ψ as an LTL formula because all direct principal subsentences are removed even before determinisation (as explained in Figure 5.3). We shall now prove the implication.

Take an arbitrary state $g \in G$ and assume that $\mathcal{I}, \emptyset, g \models_{\text{SL}} \wp b\psi$. Let $\pi = \text{play}(\chi, g)$ be the infinite play of the game. We want to show that the LTL formula ψ holds along π . By the definition of a play (Definition 2.26) and the structure of the assignment χ , we have $\pi(0) = g$ and for all $i \geq 0$, $\pi(i+1) = t\left(\pi(i), [f_0(\pi_{\leq i}), \dots, f_{|\wp|-1}(\pi_{\leq i})]^{Act}\right)$. The LTL formula ψ holds along π iff $\mathfrak{P}_{\mathcal{I}}^{\psi}$ accepts π . By construction of the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$ (Definition 5.7) and the fact that $[f_0(\pi_{\leq i}), \dots, f_{|\wp|-1}(\pi_{\leq i})]$ is a valid decision, π is accepted by $\mathfrak{P}_{\mathcal{I}}^{\psi}$ iff the following path is winning for the existential player in the game:

$$\begin{aligned} \tau \triangleq & (\pi(0), [], t_I) (\pi(0), [f_0(\pi_{\leq 0})], t_I) \dots \\ & (\pi(0), [f_0(\pi_{\leq 0}), \dots, f_{|\wp|-1}(\pi_{\leq 0})], t_I) (\pi(1), [], \delta(t_I, \pi(0))) \dots \end{aligned}$$

We will now argue that $\tau \in (V \times T)^{\omega}$ must be winning for the existential player in $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$.

Firstly, we show that τ always stays within the winning region W_0 of the existential player. Assume (for contradiction) that this is not the case, i.e. that there is some least position $k \geq 0$ such that $\tau(k) \notin W_0$. By Theorem 5.1, $\tau(0) = (g, [], t_I) \in W_0$ so $k > 0$. Consider the preceding state $\tau(k-1)$. If $\tau(k-1)$ belongs to the universal player, it would be in his winning region W_1 (as he can clearly move to another winning state $\tau(k) \in W_1 = V \setminus W_0$). This is in contradiction with the assumption that k is the first index such that $\tau(k) \notin W_0$. So $\tau(k-1)$ must be a state which belongs to the existential player. This state must be of the form $\tau(k-1) = (\pi(j), [f_0(\pi_{\leq j}), \dots, f_{l-1}(\pi_{\leq j})], \delta(t_I, \pi_{\leq j-1}))$ for some $j \geq 0$ and $0 \leq l < |\wp|$ such that $\wp(l)$ is an existential quantifier. Since $\tau(k-1) \in V_0 \cap W_0$, we must have $\tau(k-1) \in \text{dom}(w_0)$, i.e. the winning game strategy is defined for $\tau(k-1)$. Moreover, by Definition 5.9, we have $w_0(\tau(k-1)) = (\pi(j), [f_0(\pi_{\leq j}), \dots, f_l(\pi_{\leq j})], \delta(t_I, \pi_{\leq j-1})) = \tau(k)$. But then $\tau(k)$ must be a winning state for the existential player. Therefore, by contradiction, τ stays always in the winning region of the existential player W_0 .

Secondly, we show that for all $k \geq 0$, if $\tau(k) \in V_0$, then $\tau(k+1) = w_0(\tau(k))$, i.e. the path τ follows the winning strategy. Take an arbitrary index $k \geq 0$ such that $\tau(k) \in V_0$. This state must be of the form $\tau(k) = (\pi(j), [f_0(\pi_{\leq j}), \dots, f_{l-1}(\pi_{\leq j})], \delta(t_I, \pi_{\leq j-1}))$ for some $j \geq 0$ and $0 \leq l < |\wp|$ such that $\wp(l)$ is an existential quantifier. By Definition 5.9, we have $w_0(\tau(k)) = (\pi(j), [f_0(\pi_{\leq j}), \dots, f_l(\pi_{\leq j})], \delta(t_I, \pi_{\leq j-1})) = \tau(k+1)$ as required.

We have just shown that τ always stays within the winning region of the existential player W_0 and follows the winning strategy w_0 . Therefore, it is a winning path for the existential player in $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$. This implies that $\mathfrak{P}_{\mathcal{I}}^{\psi}$ accepts π . Hence, the LTL formula ψ is true along π . Our claim follows. \square

This approach can also be used to synthesise universally quantified strategies when a SL[1G] principal formula $\wp b\psi$ does *not hold* in a state $g \in G$ of an interpreted system \mathcal{I} , i.e. when $\mathcal{I}, \emptyset, g \not\models_{\text{SL}} \wp b\psi$. Since we have defined interpreted systems on transition *functions* (see Definition 2.5), by SL semantics (see Definition 2.29), this is equivalent to $\mathcal{I}, \emptyset, g \models_{\text{SL}} \overline{\wp} b\neg\psi$, where $\overline{\wp}$ is equal to \wp with swapped quantifiers (e.g. for $\wp = \llbracket e \rrbracket \llbracket x \rrbracket \langle\langle y \rangle\rangle$ we have $\overline{\wp} = \langle\langle e \rangle\rangle \langle\langle x \rangle\racket \llbracket y \rrbracket$). We can thus synthesise universally quantified strategies in $\wp b\psi$ by synthesising existentially quantified strategies in $\overline{\wp} b\neg\psi$ (using Theorem 5.3). Observe that this is not true in general when using transition *relations* as we might have both $\mathcal{I}, \emptyset, g \not\models_{\text{SL}} \wp b\psi$ and $\mathcal{I}, \emptyset, g \not\models_{\text{SL}} \overline{\wp} b\neg\psi$. To see how this can happen, consider the case when the LTL formula is $\psi = Xp$ and the transition relation is completely non-deterministic. Then no existential (universal) strategies can ensure p will be true (false) in the next state. One way to tackle this issue is to quantify the non-determinism

explicitly, i.e. set $\wp' = \wp \llbracket n \rrbracket$ and $\flat' = \flat(N, n)$. Essentially, we transform the transition relation into an agent with a universally quantified strategy and a transition function. Then $\mathcal{I}, \emptyset, g \not\models_{\text{SL}} \wp' \flat' \psi$ implies $\mathcal{I}, \emptyset, g \models_{\text{SL}} \overline{\wp'} \flat' \neg \psi$ again so we can perform strategy synthesis. Informally, in doing so, we are assuming that the nondeterminism is “playing on the universal quantifiers’ side”.

Since the SL[1G] strategy synthesis merely reuses the result generated by the model checking algorithm, namely the winning strategies (w_0, w_1) of the existential and universal player in the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp \flat \psi}$, it has the same complexity.

Note that SL[1G] strategy synthesis performs essentially the same task as the elementary dependence map $\theta_{\text{SL}[1G]}$ for a quantification prefix over strategies mentioned in Subsection 5.2.1. Similarly to SLK, it is possible to define SL[1G] witness and counterexample strategies (see Definition 4.13). We do not discuss this in detail because the SL[1G] strategy synthesis algorithm presented in this subsection is more general and can handle both concepts without any modifications.

5.2.4 Symbolic Implementation

In this subsection we discuss how the model checking algorithm (and therefore also strategy synthesis) presented in Subsection 5.2.1 can be implemented symbolically using BDDs (see Subsection 2.3.1). This subsection requires a good understanding of symbolically represented interpreted systems. Most importantly, it assumes that the reader knows how global states and joint actions can be represented using Boolean vectors. Please refer to Subsections 2.3.2, 3.3.2, and 4.2.4 for more details about this topic.

Delayed Automaton

We first explain how the delayed automaton (see Definition 5.6) is constructed symbolically as it is a relatively straightforward process. Let \mathcal{I} be an interpreted system and $\wp \flat \psi \in \text{SL}[1G]$ an SL[1G] principal sentence. Furthermore, assume that the Boolean vectors for representing current global states, next global states, and joint actions are $\bar{v} = (v_0, \dots, v_{N-1})$, $\bar{v}' = (v'_0, \dots, v'_{N-1})$, and $\bar{w} = (w_0, \dots, w_{M-1})$. The delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp \flat \psi}$ is constructed as follows (lines 2–5 in the algorithm in Figure 5.3):

1. **Construction of a non-deterministic generalised Büchi automaton $\mathfrak{A}_{\mathcal{I}}^{\psi}$.** The standard translation, which we use for this purpose, is already defined symbolically in Figure 2.6. It returns an automaton formula $\mathfrak{A}_{\mathcal{I}}^{\psi} = \mathcal{A}_{\exists}(\bar{q}, \bar{v}, \mathcal{I}_G, \mathcal{R}_G, \mathcal{F}_G)$ with a Boolean vector of new variables $\bar{q} = (q_0, \dots, q_{K-1})$ and their next versions $\bar{q}' = (q'_0, \dots, q'_{K-1})$. As explained on page 5.2.1, recursive solution of direct principal subformulas of $\wp \flat \psi$ is performed at this stage.

Note that we define already the non-generalised Büchi automaton on global states. Therefore, the formula representing an atom $p \in AP$ is $\bigvee_{g \in h(p)} g(\bar{v})$. In the case of the new atoms $p_{\wp'}$ representing direct principal subsentences, this becomes $\bigvee_{g \in \|\wp'\|_{\mathcal{I}}} g(\bar{v})$.

2. **Conversion to a non-deterministic Büchi automaton $\mathfrak{B}_{\mathcal{I}}^{\psi}$.** This is a very straightforward procedure. We introduce $L = \lceil \log_2(|\mathcal{F}_G| + 1) \rceil$ Boolean variables $\bar{r} = (r_0, \dots, r_{L-1})$ (together with their next versions \bar{r}') to represent the counter. The result is an automaton formula $\mathfrak{B}_{\mathcal{I}}^{\psi} = \mathcal{A}_{\exists}(\bar{q}, \bar{v}, \mathcal{I}_B, \mathcal{R}_B, \mathcal{F}_B)$ such that:

$$\begin{aligned} \mathcal{I}_B &= \mathcal{I}_G \wedge 0(\bar{r}) \\ \mathcal{R}_B &= \mathcal{R}_G \wedge \left(\left[\bigvee_{i=0}^{|\mathcal{F}_G|-1} i(\bar{r}) \wedge (\mathcal{F}_G(i) \wedge (i+1)(\bar{r}') \vee \neg \mathcal{F}_G(i) \wedge i(\bar{r}')) \right] \vee [|\mathcal{F}_G|(\bar{r}) \wedge 0(\bar{r}')] \right) \\ \mathcal{F}_B &= |\mathcal{F}_G|(\bar{r}) \end{aligned}$$

where $i(\bar{r})$ and $i(\bar{r}')$ are the binary representation of $0 \leq i \leq |\mathcal{F}_G|$ using the Boolean vectors \bar{r} and \bar{r}' respectively.

3. **Determinisation to a parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi}$.** A symbolic implementation of the determinisation procedure is provided at the end of Subsection 2.4.4. The result is an automaton formula $\mathfrak{P}_{\mathcal{I}}^{\psi} = \mathcal{A}_{\exists}(\bar{t}, \bar{v}, \mathcal{I}_P, \mathcal{R}_P, \mathcal{C})$ with a Boolean vector of new variables $\bar{t} = (t_0, \dots, t_{P-1})$ (and their next versions).

Note that the Boolean vectors \bar{q} and \bar{r} are not used from the next step onwards.

4. **Transformation to a delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp\psi}$.** This is again a very straightforward procedure. We introduce $D = \lceil \log_2(|\wp| + 1) \rceil$ Boolean variables $\bar{d} = (d_0, \dots, d_{D-1})$ (together with their next versions \bar{d}') to represent the counter. The result is an automaton formula $\mathfrak{D}_{\mathcal{I}}^{\wp\psi} = \mathcal{A}_{\exists}(\bar{t}\bar{d}, \bar{v}, \mathcal{I}_D, \mathcal{R}_D, \mathcal{C})$ where:

$$\begin{aligned} \mathcal{I}_D &= \mathcal{I}_P \wedge 0(\bar{d}) \\ \mathcal{R}_D &= \mathcal{R}_P \wedge |\wp|(\bar{d}) \wedge 0(\bar{d}') \vee \left[\bigvee_{i=0}^{|\wp|-1} i(\bar{d}) \wedge (i+1)(\bar{d}') \right] \wedge \left[\bigwedge_{j=0}^{P-1} t_j \leftrightarrow t'_j \right] \end{aligned}$$

We shall now discuss how the formula arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ can be constructed symbolically and then combined with the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp\psi}$ into the parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$.

Formula Arena

As we have stated in Subsection 5.2.1, the formula arena $\mathcal{A}_{\mathcal{I}}^{\wp b} = (V_0, V_1, E)$ (see Definition 5.4) is one of the key concepts of our new model checking algorithm for $\text{SL}[1G]$. Providing a thorough explanation of its symbolic representation is therefore of utmost importance.

We shall first explain how the state space of the arena can be represented symbolically. Recall that the states of $\mathcal{A}_{\mathcal{I}}^{\wp b}$ are pairs $(g, d) \in V$ where $g \in G$ is a global state and $d \in \text{Dec}_{\mathcal{I}}^{\wp b}$ a decision. Since the symbolic representation of global states was already discussed in Subsection 4.2.4, we will focus on encoding the decisions. According to Definition 5.4, the decision d is a tuple of actions such that:

1. The length of d is $0 \leq |d| \leq |\wp|$. As we have pointed out in Subsection 5.2.1, when combining the formula arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ and the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp\psi}$ into the parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b\psi}$, the length of the decision is always equal to the value of the counter. Therefore, we can use the Boolean variables \bar{d} for the counter of the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp\psi}$ to encode the length of the decision d .
2. For each index $0 \leq k < |d|$, we have $d(k) \in \text{Act}_{\text{sharing}(b\top, \wp(k))}$, i.e. $d(k)$ must be some action that all agents sharing the variable $\wp(k)$ can perform. We can encode this by introducing $E = \sum_{k=0}^{|\wp|-1} \lceil \log_2 |\text{Act}_{\text{sharing}(b\top, \wp(k))}| \rceil$ Boolean variables $\bar{e} = (e_0, \dots, e_{E-1})$. We will use \bar{e}^k to denote the Boolean variables representing the action assigned to $\wp(k)$ for $0 \leq k \leq |\wp|$. While it is not strictly necessary, we will set the value of \bar{e}^k to zero for $|d| \leq k < |\wp|$ in order to reduce the set of reachable states.

Therefore, an arbitrary state of the arena $(g, d) \in V$ is represented symbolically as follows:

$$(g, d)(\bar{v}, \bar{d}, \bar{e}) = g(\bar{v}) \wedge |d|(\bar{d}) \wedge \left[\bigwedge_{k=0}^{|d|-1} d(k)(\bar{e}^k) \right] \wedge \left[\bigwedge_{k=|d|}^{|\wp|-1} 0(\bar{e}^k) \right]$$

Consider the arena $\mathcal{A}_{\mathcal{I}}^{\gamma}$ in Figure 5.1c of the toy model in Figure 3.1 for the $\text{SL}[1G]$ basic principal sentence $\gamma = \llbracket e \rrbracket \llbracket x \rrbracket \langle y \rangle (\mathbb{E}, e)(1, x)(2, y)G \llbracket \neg p_1 \wedge \neg p_2 \rrbracket$. The global states of \mathcal{I} are represented using the Boolean vector $\bar{v} = (v_0, v_1)$ (see Subsection 3.3.2). Since $|\wp| = 3$, we need $D = \lceil \log_2(3 + 1) \rceil = 2$ Boolean variables $\bar{d} = (d_0, d_1)$ to represent the length of the decision. Finally, we need $E = \lceil \log_2 \llbracket \{i\} \rrbracket \rrbracket + 2 \lceil \log_2 \llbracket \{r, p, s, i\} \rrbracket \rrbracket = 4$ Boolean variables $\bar{e} = (e_0, e_1, e_2, e_3)$ to represent the decision. Note that $e^0 = ()$, $e^1 = (e_0, e_1)$, and $e^2 = (e_2, e_3)$. For example, the state $(g, d) = (g_g, [i, p])$ in $\mathcal{A}_{\mathcal{I}}^{\gamma}$ is encoded as:

$$(g, d)(\bar{v}, \bar{d}, \bar{e}) = \underbrace{\neg v_0 \wedge \neg v_1}_{g=g_g} \wedge \underbrace{\neg d_0 \wedge d_1}_{|d|=2} \wedge \underbrace{e_0 \wedge \neg e_1}_{d(1)=p} \wedge \underbrace{\neg e_2 \wedge \neg e_3}_{d(2)=\perp}$$

We are now ready to discuss the symbolic representation of the edge relation $E \subseteq V \times V$ of the arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$. There are two cases we have to consider:

1. An action a is *appended* to the decision d in an arena state (g, d) . Although this is relatively straightforward, we must make sure that all agents can actually perform the action⁸, i.e. $a \in \bigcap_{i \in \text{sharing}(b\top, \wp(|d|))} P_i(l_{iE}(g))$. We can do this by reusing the symbolic implementation of the individual protocols $P_i(\bar{v}, \bar{w})$ (see Subsection 3.3.2). For each $0 \leq k < |\wp|$ this case is encoded as:

$$E_1^k(\bar{v}, \bar{d}, \bar{e}, \bar{v}', \bar{d}', \bar{e}') = k(\bar{d}) \wedge (k+1)(\bar{d}') \wedge \\ 0(\bar{e}^k) \wedge \left[\bigwedge_{l=k+1}^{|\wp|} 0(\bar{e}^l) \wedge 0(\bar{e}^l) \right] \wedge \left[\bigwedge_{j=0}^{k-1} \bigwedge_{m=0}^{|\bar{e}^j|-1} e_m^j \leftrightarrow e_m^j \right] \wedge \\ \left[\bigvee_{a \in \text{Act}_{S_k}} \bigwedge_{i \in S_k} \exists \bar{w}_i. P_i(\bar{v}_{iE}, \bar{w}_i) \wedge a(\bar{w}_i) \wedge a(\bar{e}^k) \right]$$

where $S_k = \text{sharing}(b\top, \wp(k))$. The first line increments the length of the decision tuple. The second line asserts that the decision does not change arbitrarily. The third line ensures that only possible actions are appended.

2. The decision is complete and thus a *temporal transition* occurs. While this case is slightly more complicated, we can simplify it by reusing the global evolution function $t(\bar{v}, \bar{w}, \bar{v}')$ (see Subsection 3.3.2). Intuitively, we force all agents to perform the action that the decision assigns to them. This case is encoded as:

$$E_2(\bar{v}, \bar{d}, \bar{e}, \bar{v}', \bar{d}', \bar{e}') = |\wp|(\bar{d}) \wedge 0(\bar{d}') \wedge \left[\bigwedge_{j=0}^{|\wp|-1} 0(\bar{e}^j) \right] \wedge \\ \left[\bigvee_{a \in \text{Act}} \exists \bar{w}. a(\bar{w}) \wedge t(\bar{v}, \bar{w}, \bar{v}') \wedge \bigwedge_{k=0}^{|\wp|-1} \bigwedge_{i \in S_k} a_i(a)(\bar{e}^k) \right]$$

where $S_k = \text{sharing}(b\top, \wp(k))$ for $0 \leq k < |\wp|$. Again, the first line ensures that the decision (and its successor) have the correct form. The second line finds a joint action $a \in \text{Act}$ which agrees with the decision and performs the temporal transition.

Finally, the *edge relation* E of the arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ is implemented as a disjunction of the two cases above:

$$E(\bar{v}, \bar{d}, \bar{e}, \bar{v}', \bar{d}', \bar{e}') = \left[\bigvee_{k=0}^{|\wp|-1} E_1^k(\bar{v}, \bar{d}, \bar{e}, \bar{v}', \bar{d}', \bar{e}') \right] \vee E_2(\bar{v}, \bar{d}, \bar{e}, \bar{v}', \bar{d}', \bar{e}')$$

We are only interested in paths in $\mathcal{A}_{\mathcal{I}}^{\wp b}$ starting in states (g, \square) where $g \in G$ is a reachable global state. A symbolic representation of all *vertices* V of the arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ can thus be obtained as the least fixed point $V = \text{lfp}_Q [\{(g, \square) \mid g \in G\} \cup \text{suc}_{\exists}(Q)]$ (similarly to the calculation of the set of reachable states at the end of Subsection 3.3.2):

$$V(\bar{v}, \bar{d}, \bar{e}) = \text{lfp}_{\Theta} \left[\left(\bigvee_{g \in G} g(\bar{v}) \wedge 0(\bar{d}) \wedge 0(\bar{e}) \right) \vee \exists \bar{v}' \bar{d}' \bar{e}'. (\Theta' \wedge E(\bar{v}', \bar{d}', \bar{e}', \bar{v}, \bar{d}, \bar{e})) \right]$$

where Θ' is a Boolean formula equal to Θ with variables in $\bar{v}, \bar{d}, \bar{e}$ and $\bar{v}', \bar{d}', \bar{e}'$ swapped. The vertices

⁸If we do not ensure this, the game might end up in a dead end (g, d_{\perp}) when a temporal transition should occur as the accumulated decision d_{\perp} might not represent a possible joint action in g . By our construction of $\mathcal{A}_{\mathcal{I}}^{\wp b}$, (g, d_{\perp}) would be a dead end for the universal player. Hence, it would always be a winning state for the existential player. Informally, this would “encourage” the existential player to choose actions which cannot be performed in the current state.

of the *existential and universal player* in $\mathcal{A}_{\mathcal{I}}^{\wp b}$ can in turn be calculated symbolically as follows:

$$V_0(\bar{v}, \bar{d}, \bar{e}) = V(\bar{v}, \bar{d}, \bar{e}) \wedge \bigvee_{k=0}^{|\wp|-1} k(\bar{d}) \wedge \text{existential}(k)$$

$$V_1(\bar{v}, \bar{d}, \bar{e}) = V(\bar{v}, \bar{d}, \bar{e}) \wedge \bigvee_{k=0}^{|\wp|-1} k(\bar{d}) \wedge \neg \text{existential}(k)$$

where $\text{existential}(k)$ is true iff $\wp(k)$ is an existentially quantified variable for $0 \leq k < |\wp|$.

Combined Game

The parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$ is a combination of the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp \psi}$ and the formula arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ (see Definition 5.7). This is very simple because our symbolic representations of both structures uses the same counter \bar{d} . The *edge relation* $E_{\mathfrak{G}}$ of the game is implemented symbolically as:

$$E_{\mathfrak{G}}(\bar{v}, \bar{d}, \bar{e}, \bar{t}, \bar{v}', \bar{d}', \bar{e}', \bar{t}') = E(\bar{v}, \bar{d}, \bar{e}, \bar{v}', \bar{d}', \bar{d}') \wedge \mathcal{R}_D(\bar{t}, \bar{d}, \bar{v}, \bar{t}', \bar{d}')$$

Again, the sets of *all vertices* $V_{\mathfrak{G}}$, *existential player vertices* $V_{\mathfrak{G}0}$, and *universal player vertices* $V_{\mathfrak{G}1}$ can be calculated symbolically as follows:

$$V_{\mathfrak{G}}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = \text{lfp}_{\Theta} \left[\left(\bigvee_{g \in G} g(\bar{v}) \wedge 0(\bar{e}) \wedge \mathcal{I}_D(\bar{d}, \bar{t}) \right) \vee \overline{\exists v' d' e' t'}. (\Theta' \wedge E(\bar{v}', \bar{d}', \bar{e}', \bar{t}', \bar{v}, \bar{d}, \bar{e}, \bar{t})) \right]$$

$$V_{\mathfrak{G}0}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = V_{\mathfrak{G}}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) \wedge \bigvee_{k=0}^{|\wp|-1} k(\bar{d}) \wedge \text{existential}(k)$$

$$V_{\mathfrak{G}1}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = V_{\mathfrak{G}}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) \wedge \bigvee_{k=0}^{|\wp|-1} k(\bar{d}) \wedge \neg \text{existential}(k)$$

The *colouring function*⁹ $\mathcal{C}_{\mathfrak{G}}$ is represented as a family of $|\mathcal{C}|$ Boolean formulas:

$$\mathcal{C}_{\mathfrak{G}i}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = V_{\mathfrak{G}}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) \wedge \mathcal{C}_i(\bar{t})$$

where \mathcal{C} is the family of Boolean formulas representing the colouring function of the parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi}$ and $0 \leq i < |\mathcal{C}|$ a colour.

Once we have a symbolic representation of the parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$, we solve it using the existing algorithm in Figure 2.9. The symbolic implementation of the algorithm is quite straightforward because it uses only basic set operations (Subsection 2.3.2 explains how set operations can be represented symbolically) and attractors (see Definition 2.59). We shall therefore only discuss how *attractors* and *attractor strategies* can be calculated symbolically. Recall that an attractor $\text{attr}_{\sigma}(\mathcal{A}, X) \subseteq V$ of a set of nodes $X \subseteq V$ for a player $\sigma \in \{0, 1\}$ in an arena $\mathcal{A} = (V_0, V_1, E)$ is defined as:

$$\text{attr}_{\sigma}(\mathcal{A}, X) \triangleq \text{lfp}_Q [X \cup \{v \in V_{\sigma} \mid vE \cap Q \neq \emptyset\} \cup \{v \in V_{\bar{\sigma}} \mid vE \subseteq Q\}]$$

Intuitively, $\text{attr}_{\sigma}(\mathcal{A}, X)$ is the set of all vertices of \mathcal{A} from which player σ can force a visit to X . If the vertices of \mathcal{A} are represented using a Boolean vector¹⁰ \bar{u} (and its next version \bar{u}'), the attractor can be calculated symbolically as:

$$\text{attr}_{\sigma}(\mathcal{A}, X)(\bar{u}) = \text{lfp}_{\Theta} [X(\bar{u}) \vee (V_{\sigma}(\bar{u}) \wedge \exists \bar{u}'. \Theta' \wedge E(\bar{u}, \bar{u}')) \vee (V_{\bar{\sigma}}(\bar{u}) \wedge \neg \exists \bar{u}'. \neg \Theta' \wedge E(\bar{u}, \bar{u}'))] \quad (5.2)$$

We are also interested in the winning strategies of $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$ for the purposes of strategy synthesis (see Subsection 5.2.3). In parity games, these are constructed from attractor strategies. Informally, a (memoryless) attractor strategy for $\text{attr}_{\sigma}(\mathcal{A}, X)$ assigns to states $\text{attr}_{\sigma}(\mathcal{A}, X) \cap V_{\sigma}$ moves “towards” the target

⁹More formally, $\mathcal{C}_{\mathfrak{G}i}$ is a symbolic representation of the set of all vertices of $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$ with colour $0 \leq i < |\mathcal{C}_{\mathfrak{G}}|$, i.e. $\mathcal{C}_{\mathfrak{G}i} = \{v \in V_{\mathfrak{G}} \mid c_{\mathfrak{G}}(v) = i\}$.

¹⁰In the case of the combined game $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$, we have $\bar{u} = \bar{v}\bar{d}\bar{e}\bar{t}$.

set X . We can represent a strategy as a set of pairs (v, v') such that $v \in \text{attr}_\sigma(\mathcal{A}, X) \cap V_\sigma$ and $v' \in vE$. Moreover, we can represent both the attractor and the strategy together as a disjunction of the attractor strategy (which contains all vertices $\text{attr}_\sigma(\mathcal{A}, X) \cap V_\sigma$) and the vertices of the other player in the attractor $\text{attr}_\sigma(\mathcal{A}, X) \cap V_{\bar{\sigma}}$. This can be calculated as:

$$\text{attr}_\sigma(\mathcal{A}, X)(\bar{u}, \bar{u}') = \text{lfp}_\Theta [X(\bar{u}, \bar{u}') \vee (V_\sigma(\bar{u}) \wedge \neg(\exists \bar{u}'. \Theta) \wedge (\exists \bar{u}. \Theta') \wedge E(\bar{u}, \bar{u}')) \vee (V_{\bar{\sigma}}(\bar{u}) \wedge \neg \exists \bar{u}'. \neg(\exists \bar{u}. \Theta') \wedge E(\bar{u}, \bar{u}'))]$$

Note that X can be a combination of vertices $v \in V$ and strategy mappings $(v, v') \in V \times V$ as well. This fixpoint is very similar to the one in Equation 5.2. The main difference is that Θ is now a Boolean formula over both \bar{v} and \bar{v}' . Hence, we have to existentially quantify it *before* combining it with the edge relation (unlike in Equation 5.2), i.e. $(\exists \bar{u}. \Theta') \wedge E(\bar{u}, \bar{u}')$ instead of $\exists \bar{u}'. \Theta' \wedge E(\bar{u}, \bar{u}')$. More importantly, there is one extra term, namely $\neg(\exists \bar{u}'. \Theta)$, in the existential player case. It ensures that we add each state to the attractor only once. This is in line with the definition of the attractor strategy (see Definition 2.59). If we did not add this restriction, we could end up in the following situation: Let $\mathcal{A} = (V_0, V_1, E)$ be an arena with 3 vertices $V = \{a, b, c\}$ such that all of them belong to the existential player, i.e. $V_0 = V$ and $V_1 = \emptyset$. Furthermore, assume that the edge relation is $E = \{(b, a), (c, a), (b, c), (c, b)\}$. Consider the calculation of $\text{attr}_0(\mathcal{A}, \{a\})$. Firstly, the pairs (b, a) and (c, a) are added to the initial set $\{a\}$. If we do not require that each vertex is added at most once, (b, c) and (c, b) will be added to the attractor in the second iteration. But then we would have $(b, c), (c, b) \in \text{attr}_0(\mathcal{A}, \{a\})$ and $b, c \notin \{a\}$, i.e. a strategy which cycles between b and c forever and never reaches the target set $\{a\}$.

Model Checking and Strategy Synthesis

The symbolic implementation of the algorithm for solving parity games (see Figure 2.9) will return Boolean formulas representing both the winning regions and strategies $(W_{\mathfrak{G}_0}, W_{\mathfrak{G}_1})$ of the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$. The formula $\|\wp b \psi\|_{\mathcal{I}}(\bar{v})$ encoding the set of all global states of \mathcal{I} at which the SL[1G] principal formula $\wp b \psi$ holds is finally calculated as (see Theorem 5.1):

$$\|\wp b \psi\|_{\mathcal{I}}(\bar{v}) = \exists \bar{d} \bar{e} \bar{t} \bar{v}' \bar{d}' \bar{e}' \bar{t}'. W_{\mathfrak{G}_0}(\bar{v}, \bar{d}, \bar{e}, \bar{t}, \bar{v}', \bar{d}', \bar{e}', \bar{t}') \wedge \underbrace{\mathcal{I}_D(\bar{d}, \bar{t})}_{\text{start}} \wedge 0(\bar{e})$$

The game $\mathfrak{G}_{\mathcal{I}}^{\wp b \psi}$ together with the winning regions and strategies $(W_{\mathfrak{G}_0}, W_{\mathfrak{G}_1})$ encodes the synthesised strategies for all variables in the quantification prefix \wp . Assume that we want to determine the next action $f_k(\pi)$ that the strategy f_k for variable $\wp(k)$ with $0 \leq k < |\wp|$ should assign to the track $\pi \in \text{Trk}$ in \mathcal{I} . Given the next actions $f_0(\pi), \dots, f_{k-1}(\pi)$ of strategies for variables $\wp(0), \dots, \wp(k-1)$, we proceed as follows (see Definition 5.9):

1. We start in the state $s_0^0(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = \text{start}(\pi(0))(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = \pi(0)(\bar{v}) \wedge \mathcal{I}_D(\bar{d}, \bar{t}) \wedge 0(\bar{e})$. We basically set the initial state t_I of the underlying parity automaton $\mathfrak{B}_{\mathcal{I}}^{\psi}$ of the game ($t_0 = t_I$).
2. For $i = 1$ to $|\pi| - 1$: We set $s_i^0(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = \exists \bar{v}' \bar{d}' \bar{e}' \bar{t}'. E_{\mathfrak{G}}(\bar{v}', \bar{d}', \bar{e}', \bar{t}', \bar{v}, \bar{d}, \bar{e}, \bar{t}) \wedge \pi(i-1)(\bar{v}) \wedge 0(\bar{d}) \wedge [\exists \bar{d}' \bar{e}'. s_{i-1}^0(\bar{v}', \bar{d}', \bar{e}', \bar{t}')]]$. Essentially, we are traversing the automaton $\mathfrak{B}_{\mathcal{I}}^{\psi}$ ($t_i = \delta(t_{i-1}, \pi(i-1))$).
3. For $j = 0$ to $k-1$: We set $s_{|\pi|-1}^{j+1}(\bar{v}, \bar{d}, \bar{e}, \bar{t}) = \exists \bar{v}' \bar{d}' \bar{e}' \bar{t}'. E_{\mathfrak{G}}(\bar{v}', \bar{d}', \bar{e}', \bar{t}', \bar{v}, \bar{d}, \bar{e}, \bar{t}) \wedge s_{|\pi|-1}^j(\bar{v}', \bar{d}', \bar{e}', \bar{t}') \wedge f_j(\pi)(\bar{e}^j)$. We accumulate the decision $[f_0(\pi), \dots, f_{k-1}(\pi)]$.
4. $s_{|\pi|-1}^k$ is the implied current state of the combined parity game $\mathfrak{G}_{\mathcal{I}}^{\wp b}$ (given π and $f_0(\pi), \dots, f_{k-1}(\pi)$). Hence, we can use the winning regions and strategies $(W_{\mathfrak{G}_0}, W_{\mathfrak{G}_1})$ to determine the next action $f_k(\pi)$: We select an action $a \in \bigcap_{i \in \text{sharing}(b \top, \wp(k))} P_i(l_{iE}(\text{last}(\pi)))$ such that the expression $\exists \bar{v} \bar{d} \bar{e} \bar{t} \bar{v}' \bar{d}' \bar{e}' \bar{t}'. E_{\mathfrak{G}}(\bar{v}', \bar{d}', \bar{e}', \bar{t}', \bar{v}, \bar{d}, \bar{e}, \bar{t}) \wedge (W_{\mathfrak{G}_0}(\bar{v}', \bar{d}', \bar{e}', \bar{t}', \bar{v}, \bar{d}, \bar{e}, \bar{t}) \vee W_{\mathfrak{G}_1}(\bar{v}', \bar{d}', \bar{e}', \bar{t}', \bar{v}, \bar{d}, \bar{e}, \bar{t})) \wedge a(\bar{e}^k)$ is not false. If there is no such action, we select any possible action a .

This concludes our symbolic implementation of SL[1G] model checking and strategy synthesis.

Complexity

The time complexity of our symbolic implementation is the same as that of the model checking algorithm discussed in Subsection 5.2.2, namely polynomial in the size of the model $|\mathcal{I}|$ and doubly exponential in the size of the formula $|\varphi|$.

Theorem 5.4. Let \mathcal{I} be a an interpreted system and φ an SL[1G] sentence. The worst-case time complexity of the symbolic implementation is:

$$|\mathcal{I}|^{2^{O(|\varphi|)}}$$

Proof (Sketch). We only need to prove the statement for an arbitrary *basic principal sentence* $\varphi = \wp \wp \psi$ (see the proof of Theorem 5.2). The symbolic implementation of our algorithm uses:

- $O(\log_2 |G|)$ Boolean variables to represent the *global state*;
- $O(\sum_{i \in \text{Agt}} \log_2 |Act_i|)$ Boolean variables to represent *joint actions*;
- $O(|\psi|)$ Boolean variables to represent the states of the *non-deterministic (generalised) Büchi automaton*;
- $2^{O(|\psi|)}$ Boolean variables to represent the *deterministic parity automaton*;
- $O(\log_2 |\wp|)$ Boolean variables to represent the *delayed automaton and arena counter*;
- $O(\sum_{i \in \text{Agt}} \log_2 |Act_i|)$ Boolean variables to represent *decisions*.

The total number of Boolean variables needed is thus $O(\log_2 |G| + \sum_{i \in \text{Agt}} \log_2 |Act_i| + |\wp|) + 2^{O(|\psi|)} = O(\log_2 |\mathcal{I}|) + 2^{O(|\varphi|)}$. Since BDD operations take polynomial time with respect to the size of the relevant BDDs in the worst case [49], the worst-case time complexity of each BDD operation is $O(|\mathcal{I}|^k)2^{O(|\varphi|)}$ for some fixed constant $k > 0$. As we have stated in the proof of Theorem 5.2, the model checking algorithm has at most $|\mathcal{I}|^{2^{O(|\varphi|)}}$ steps. By multiplying these two expressions, we get the desired worst-case time complexity $|\mathcal{I}|^{2^{O(|\varphi|)}}$. \square

5.2.5 Separate Determinisation

The SL[1G] model checking algorithm complexity proof (see Theorem 5.2) and the experimental results (see Section 6.4) indicate that one major bottleneck of the algorithm is parity game solving. In this section, we propose an optimisation technique, which addresses this problem by reducing the size of the game. Our experimental results demonstrate that it significantly reduces model checking time as well as memory usage (see Table 6.4b).

We motivate the optimisation by means of a simple example. Consider the following SL[1G] basic principal sentence (for some interpreted system \mathcal{I} with agents $\text{Agt} \triangleq \{a, b\}$):

$$\beta \triangleq \llbracket x \rrbracket \langle \langle y \rangle \rangle (a, x)(b, y)[(p \rightarrow \text{GF } p) \wedge (q \rightarrow \text{GF } q) \wedge (r \rightarrow \text{GF } r)]$$

The underlying LTL formula of β is $\psi_\beta = (p \rightarrow \text{GF } p) \wedge (q \rightarrow \text{GF } q) \wedge (r \rightarrow \text{GF } r)$. The non-deterministic Büchi automaton $\mathfrak{B}_{\mathcal{I}}^{\psi_\beta}$ for ψ_β obtained on line 3 of our SL[1G] model checking algorithm in Figure 5.3 has 508 reachable states which are not dead ends. Therefore, 509 Boolean variables will be required to encode a single pair (S_i, m_i) in a state of the deterministic parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi_\beta}$ on line 4 of our SL[1G] model checking algorithm (the determinisation procedure we use is described in Subsection 2.4.4). As we shall see, at least three such pairs will be required to encode a state of $\mathfrak{P}_{\mathcal{I}}^{\psi_\beta}$. Hence, at least 1527 Boolean variables will be required to encode a state of the deterministic parity automaton. To put this into perspective, $\mathfrak{P}_{\mathcal{I}}^{\psi_\beta}$ will have at least $2^{1527} \approx 4.7 \times 10^{459}$ states. Comparing this with Table 2.2, we see that our symbolic implementation using BDDs will *not* be able to check β . This is a very negative result given that β is a relatively concise specification. It renders our SL[1G] model checking algorithm presented in Subsection 5.2.1 completely infeasible for verification and synthesis of large multi-agent systems with multiple constraints.

We propose an optimisation technique, which we refer to as *separate determinisation*, which allows us to tackle the problem above. Specifications of systems are usually given as conjunctions of LTL formulas [77]. Hence, we are interested in the verifications of SL[1G] principal sentences of the form:

$$\wp b [\psi_0 \wedge \psi_1 \wedge \dots \wedge \psi_{n-1}]$$

We proceed as before and recursively replace all direct principal subformulas of the formula above with atoms (see Subsection 5.2.1). Hence we can assume $\psi_0, \dots, \psi_{n-1}$ are LTL formulas without any loss of generality. Our aim is to construct a deterministic automaton $\mathfrak{S}_{\mathcal{I}}^{\psi}$ which accepts precisely those infinite paths $\pi \in Pth$ along which the LTL formula $\psi \triangleq \psi_0 \wedge \dots \wedge \psi_{n-1}$ holds. Take an arbitrary infinite path $\pi \in Pth$. We have the following chain of equivalences:

$$\begin{aligned} \mathcal{I}, \pi \models_{\text{LTL}} \psi & \text{ iff } \forall i \in \{0, \dots, n-1\}. \mathcal{I}, \pi \models_{\text{LTL}} \psi_i \\ & \text{ iff } \forall i \in \{0, \dots, n-1\}. \pi \in \text{Lang}(\mathfrak{P}_{\mathcal{I}}^{\psi_i}) \\ & \text{ iff } \pi \in \bigcap_{i=0}^{n-1} \text{Lang}(\mathfrak{P}_{\mathcal{I}}^{\psi_i}) \end{aligned}$$

Therefore, ψ holds along π iff it is accepted by the deterministic parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi_i} = (T_i, G, t_{I_i}, \delta_i, c_i)$ for each conjunct ψ_i with $0 \leq i < n$. Consider the following *generalised parity automaton* (see Definition 2.39) obtained as a product of $\mathfrak{P}_{\mathcal{I}}^{\psi_0}, \dots, \mathfrak{P}_{\mathcal{I}}^{\psi_{n-1}}$:

$$\mathfrak{S}_{\mathcal{I}}^{\psi} \triangleq \prod_{i=0}^{n-1} \mathfrak{P}_{\mathcal{I}}^{\psi_i} = \left(\prod_{i=0}^{n-1} T_i, G, (t_{I_0}, \dots, t_{I_{(n-1)}}), \delta, \{c_1, \dots, c_{n-1}\} \right)$$

where $\delta((t_0, \dots, t_{n-1}), g) \triangleq (\delta_0(t_0, g), \dots, \delta_{n-1}(t_{n-1}, g))$ for all $(t_0, \dots, t_{n-1}) \in \prod_{i=0}^{n-1} T_i$ and $g \in G$. We claim that $\mathfrak{S}_{\mathcal{I}}^{\psi}$ accepts exactly those infinite paths in \mathcal{I} which satisfy ψ .

Lemma 5.3. Let \mathcal{I} be an interpreted system, $\psi_0, \dots, \psi_{n-1}$ arbitrary LTL formulas such that $\psi \triangleq \bigwedge_{i=0}^{n-1} \psi_i$, and $\pi \in Pth$ a path in \mathcal{I} . Furthermore, let $\mathfrak{P}_{\mathcal{I}}^{\psi_i}$ be a deterministic parity automaton equivalent to ψ_i for each $0 \leq i < n$. Then the following holds for the deterministic generalised parity automaton $\mathfrak{S}_{\mathcal{I}}^{\psi} \triangleq \prod_{i=0}^{n-1} \mathfrak{P}_{\mathcal{I}}^{\psi_i}$: $\pi \in \text{Lang}(\mathfrak{S}_{\mathcal{I}}^{\psi})$ iff $\mathcal{I}, \pi \models_{\text{LTL}} \psi$.

Proof. We can prove this statement directly using equivalences: By construction of $\mathfrak{S}_{\mathcal{I}}^{\psi}$, $\pi \in \text{Lang}(\mathfrak{S}_{\mathcal{I}}^{\psi})$ iff $\pi \in \bigcap_{i=0}^{n-1} \text{Lang}(\mathfrak{P}_{\mathcal{I}}^{\psi_i})$. By definition of set intersection, $\pi \in \bigcap_{i=0}^{n-1} \text{Lang}(\mathfrak{P}_{\mathcal{I}}^{\psi_i})$ iff $\pi \in \text{Lang}(\mathfrak{P}_{\mathcal{I}}^{\psi_i})$ for all $0 \leq i < n$. By our initial assumption, for all $0 \leq i < n$, we have $\pi \in \text{Lang}(\mathfrak{P}_{\mathcal{I}}^{\psi_i})$ iff $\mathcal{I}, \pi \models_{\text{LTL}} \psi_i$. By LTL semantics (Definition 2.8), $\mathcal{I}, \pi \models_{\text{LTL}} \psi$ iff $\mathcal{I}, \pi \models_{\text{LTL}} \psi_i$ for all $0 \leq i < n$. Our claim follows. \square

Consider again the LTL formula $\psi_{\beta} = (p \rightarrow \text{GF} p) \wedge (q \rightarrow \text{GF} q) \wedge (r \rightarrow \text{GF} r)$. Each of the non-deterministic Büchi automata $\mathfrak{B}_{\mathcal{I}}^{\psi_0}$, $\mathfrak{B}_{\mathcal{I}}^{\psi_1}$, and $\mathfrak{B}_{\mathcal{I}}^{\psi_2}$ for the subformulas $\psi_0 = p \rightarrow \text{GF} p$, $\psi_1 = q \rightarrow \text{GF} q$, and $\psi_2 = r \rightarrow \text{GF} r$ has 13 reachable states which are not dead ends. The resulting deterministic parity automata $\mathfrak{P}_{\mathcal{I}}^{\psi_0}$, $\mathfrak{P}_{\mathcal{I}}^{\psi_1}$, and $\mathfrak{P}_{\mathcal{I}}^{\psi_2}$ each have 42 variables¹¹. Therefore, the deterministic generalised parity automaton $\mathfrak{S}_{\mathcal{I}}^{\psi_{\beta}}$ equivalent to ψ_{β} has 126 Boolean variables and $2^{126} \approx 8.5 \times 10^{37}$ states. While this is still a very large number, it is much more manageable than the size of the deterministic parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi_{\beta}}$, which would have at least $2^{1527} \approx 4.7 \times 10^{459}$ states.

Once we have constructed the deterministic generalised parity automaton $\mathfrak{S}_{\mathcal{I}}^{\psi}$, the optimised SL[1G] model checking algorithm proceeds in a very similar manner to the original one in Figure 5.3: We construct the formula arena $\mathcal{A}_{\mathcal{I}}^{\wp b}$ and the delayed automaton $\mathfrak{D}_{\mathcal{I}}^{\wp b}$. Note that $\mathcal{A}_{\mathcal{I}}^{\wp b}$ is unchanged but $\mathfrak{D}_{\mathcal{I}}^{\wp b}$ is now a generalised parity automaton. Finally, we construct the combined *generalised parity game* $\mathfrak{G}_{\mathcal{I}}^{\wp b}$ and solve it using the existing algorithm in Figure 2.10. The complete optimised algorithm is shown in Figure 5.4.

¹¹Their states require at least three pairs (S_i, m_i) (see Subsection 2.4.4). Hence, they each have $3 \times (13 + 1) = 42$ variables. This is also the reason why we know that the original parity automaton $\mathfrak{P}_{\mathcal{I}}^{\psi_{\beta}}$ equivalent to ψ_{β} would require at least three pairs (S_i, m_i) to encode a state.

```

1 function CHECKSL[1G]OPTIMISED( $\wp\psi$ )
2    $\psi_0 \wedge \dots \wedge \psi_{n-1} := \psi$  ▷ Split conjuncts.
3   for  $i := 0, \dots, n-1$  do
4      $\mathfrak{A}_{\mathcal{I}}^{\psi_i} := \text{GENBÜCHI}(\psi_i)$ 
5      $\mathfrak{B}_{\mathcal{I}}^{\psi_i} := \text{BÜCHI}(\mathfrak{A}_{\mathcal{I}}^{\psi_i})$ 
6      $\mathfrak{P}_{\mathcal{I}}^{\psi_i} := \text{PARITY}(\mathfrak{B}_{\mathcal{I}}^{\psi_i})$ 
7   end for
8    $\mathfrak{S}_{\mathcal{I}}^{\psi} := \text{GENPARITY}(\mathfrak{P}_{\mathcal{I}}^{\psi_0}, \dots, \mathfrak{P}_{\mathcal{I}}^{\psi_{n-1}})$  ▷ Calculate automaton product.
9    $\mathfrak{D}_{\mathcal{I}}^{\wp\psi} := \text{DELAYED}(\mathfrak{S}_{\mathcal{I}}^{\psi}, \wp)$  ▷ Convert to a delayed generalised parity automaton.
10   $\mathcal{A}_{\mathcal{I}}^{\wp\psi} := \text{ARENA}(\mathcal{I}, \wp, b)$ 
11   $\mathfrak{G}_{\mathcal{I}}^{\wp\psi} := \text{COMBINE}(\mathcal{A}_{\mathcal{I}}^{\wp\psi}, \mathfrak{D}_{\mathcal{I}}^{\wp\psi})$  ▷ Combine into a generalised parity game.
12   $(W_0, W_1) := \text{SOLVEGENPARITY}(\mathfrak{G}_{\mathcal{I}}^{\wp\psi})$  ▷ Solve generalised parity game (Figure 2.10).
13  return  $\{g \in G \mid \text{start}(g) \in W_0\}$ 
14 end function

```

Figure 5.4: The optimised model checking algorithm for SL[1G] which uses *separate determinisation*. The steps which differ from the original algorithm (see Figure 5.3) have comments.

This optimised algorithm does not improve the general $|\mathcal{I}|^{2^{O(|\varphi|)}}$ time complexity of SL[1G] model checking (see Subsection 5.2.2) because it has no effect when ψ has only one conjunct, in which case the generalised parity game is a parity game with one colouring function, i.e. $\mathfrak{S}_{\mathcal{I}}^{\psi} = \mathfrak{P}_{\mathcal{I}}^{\psi}$. Nevertheless, the following more precise worst-case time complexity can be derived¹²:

$$|\varphi| + \sum_{\varphi' \in \text{sub}(\varphi)} |\mathcal{I}|^{D(\varphi')} \binom{D(\varphi')}{D_0(\varphi'), \dots, D_{N(\varphi')-1}(\varphi')}$$

where $N(\varphi')$ is the number of conjuncts in $\varphi' = \wp b [\psi_0 \wedge \dots \wedge \psi_{N(\varphi')-1}]$, $D_i(\varphi') \triangleq 2^{O(\text{temp}(\psi_i))}$ for $0 \leq i < N(\varphi')$, $D \triangleq \sum_{i=0}^{N(\varphi')-1} D_i(\varphi')$, $\binom{n}{k_0, \dots, k_{n-1}} \triangleq \frac{n!}{\prod_{i=0}^{n-1} k_i!}$, and $\text{temp}(\varphi')$ is the number of temporal operators in φ' (without strict principal subformulas). Note that this expression is equal to Equation 5.1 when each principal subsentence of φ has exactly one conjunct.

The symbolic implementation of the model checking algorithm (see Subsection 5.2.4) is modified appropriately. There is one extra issue, which is only relevant to SL[1G] *strategy synthesis* (see Subsection 5.2.3). Recall that generalised parity games are not memoryless determined for the existential player in general (see Subsection 2.4.5), i.e. the winning strategy for the existential player requires finite memory, namely a counter for the current conjunct $0 \leq i < n$. We solve this problem by (i) augmenting the Boolean vectors representing the current and next game state with $\lceil \log_2 n \rceil$ Boolean variables for the counter and (ii) adding an auxiliary state after every temporal transition in which the existential player has the option to keep or increment the value of the counter:

$$(g, [a_0, \dots, a_{|\varphi|-1}], t, i) \xrightarrow{\text{temporal transition}} (g', -, t', i) \begin{cases} \xrightarrow{\text{keep } i} (g', [], t', i) \\ \xrightarrow{\text{increment } i} (g', [], t', (i+1) \bmod n) \end{cases}$$

5.3 Summary

In this chapter, we first gave a brief overview of SL[1G] and then introduced a practical model checking algorithm for it. This is a very important result because there is currently no such algorithm. In fact, we are not even aware of any practical model checking algorithms for ATL*, which is strictly subsumed

¹²We use the fact that the time complexity of solving a generalised parity game with n vertices, m edges, and maximum colours d_0, \dots, d_{k-1} is $O(mn^d) \binom{\lceil d/2 \rceil}{\lceil d_0/2 \rceil, \dots, \lceil d_{k-1}/2 \rceil}$ where $d = \sum_{i=0}^{k-1} d_i$ (see Proposition 2.11). Note that we ignore the coefficients because $D(\varphi')/2 = \sum_{i=0}^{k-1} D_i(\varphi')/2 = \sum_{i=0}^{k-1} 2^{O(\text{temp}(\psi_i))-1} = \sum_{i=0}^{k-1} 2^{O(\text{temp}(\psi_i))} = D(\varphi')$.

by SL[1G]. Therefore, we believe that we have put forward the first practical model checking algorithm for both SL[1G] and ATL*. Moreover, we showed that it has optimal worst-case time complexity, proved its correctness, and provided an efficient symbolic implementation. Since it supports general strategy synthesis, it can be used to automatically generate agents' behaviour from SL[1G] specifications.

The algorithm reduces the model checking of an SL[1G] basic principal sentence $\wp b \psi$ to the problem of solving a combined parity game, which is obtained as the product of a formula arena, which represents the interdependency of the quantified strategies in the prefix $\wp b$, and a delayed automaton, which accepts all paths through the arena which satisfy the underlying LTL formula ψ . We showed how the winning strategies of the combined parity game can be used for SL[1G] strategy synthesis. We also provided an optimisation technique called separate determinisation, which improves the performance of the algorithm on SL[1G] formulas of the form $\wp b [\psi_0 \wedge \dots \wedge \psi_{n-1}]$ by determinising each conjunct separately and reducing the problem to solving a generalised parity game instead.

In Section 6.3, we will describe how we developed an extension of MCMAS which implements the SL[1G] model checking algorithm introduced in this chapter using BDDs.

Chapter 6

Implementation

A very important part of this project was the implementation of the novel model checking algorithms for SLK and SL[1G] presented in Sections 4.2 and 5.2 respectively. We developed both algorithms as extensions for the MCMAS model checker described in Subsection 2.5.2. To our best knowledge, there are *no other tools* which would support SL or any of its fragments such as SLK or SL[1G]. In our opinion, this makes our contribution even more valuable, as it demonstrates the feasibility of SL as a practical specification language rather than a mere theoretical construct. Our belief is further supported by the experimental results described in Section 6.4.

Extending an existing tool relieved us of the burden to design and implement all low-level procedures from scratch. Instead, it allowed us to focus on the important aspects of the new model checking algorithms and spend more time optimising their performance. On the other hand, we had to follow the existing design pattern of the tool and accept its shortcomings. We shall start by describing the existing functionality, usage, and architecture of MCMAS. We will then describe the new functionality of the SLK and SL[1G] extensions and explain which parts of the tool we had to modify in order to support each of the fragments. Finally, we present the experimental results we obtained on several scalable scenarios.

6.1 Existing Tool

MCMAS is a Model Checker for Multi-Agent Systems (MAS) developed at Imperial College London released under GNU Public Licence (GPL) [5, 8, 63]. A brief overview of its existing functionality has already been provided in Subsection 2.5.2, where it was compared with other existing verification tools. Here we shall describe it in more detail to give the reader a better idea of the original tool as well as the extensions which we developed as part of this project.

6.1.1 Functionality

MCMAS has the following features:

1. **Underlying framework.** The multi-agent systems in MCMAS are modelled as *interpreted systems* (see Definition 2.5), which naturally capture temporal, strategic, deontic (correct behaviour), and epistemic (knowledge) properties of a system. An important consequence of using this formalism is that all agents evolve synchronously. The semantics of the specifications languages supported by MCMAS (described below) are also defined with respect to interpreted systems (see Section 2.2).

The reader should be familiar with the concept of interpreted systems by now as we have been using it heavily throughout Chapters 4 and 5. A very simple toy model and its formal representation as an interpreted system is provided in Section 3.3.

2. **Input format.** The description of interpreted systems is provided in the form of ISPL files. Such files use special syntax to describe the whole model including the agents, protocols, evolution functions, local variables, observable variables, initial states, and specifications to be checked. Please refer to [9] for the complete ISPL syntax. ISPL code for the toy model introduced in Section 3.3 is provided in Appendix A.

3. **Specification languages.** MCMAS supports CTL (see Subsection 2.2.2) and ATL (see Subsection 2.2.4) augmented with epistemic operators expressing agents' knowledge (see Subsection 2.2.6) and deontic modalities for correct behaviour (not relevant for our project). Moreover, MCMAS admits arbitrary nesting of CTL and ATL operators as well as the non-temporal modalities. In addition, basic fairness conditions¹ are supported. LTL, CTL*, ATL*, and SL (including all syntactic fragments) are currently *not supported* by MCMAS.
4. **Model checking.** The main function of MCMAS is model checking of specifications against a model of a system. The tool automatically verifies and reports whether each formula holds in all initial states of the model or not. Thanks to its underlying efficient symbolic implementation (see Subsection 2.3.2), MCMAS can easily handle state spaces with more than 10^{20} possible states [63].
5. **Witness and counterexample executions.** In addition to reporting whether a formula is true or false in the given model, MCMAS can provide the user with sample executions justifying the truth value of the formula. If the formula is true, a *witness execution* is generated. Conversely, if the formula is false, a *counterexample execution* is generated. This is a very useful feature as it allows the user to identify and then fix the parts of the model in which the specification is not satisfied.
6. **Interactive simulation.** The tool provides an interactive simulation mode, in which the user selects an initial state and follows a possible evolution of the model by choosing the actions of all agents at each time step. The user can always also backtrack to the previous state. This feature does not provide automatic verification. Instead, it allows the user to manually check that the model is correct.
7. **Graphical interface.** Although MCMAS is a purely command-line tool with textual input and output, an Eclipse plugin supporting most of its features has been developed. It can currently perform model checking, interactive simulation, and witness/counterexample analysis. Its most prominent feature is visualisation of witness/counterexample executions.

The tool runs on most architectures including Linux, Windows, and Mac OS.

The latest publicly available release of MCMAS at the time of writing is version 1.2.1. We describe here version 1.1.0, which was available in October 2013 when our project forked off. The main difference between the two versions is the support for uniform ATL semantics in version 1.2.1 (uniformity is explained in Subsection 4.1.2).

6.1.2 Usage

We provide here a short introduction to the ISPL file format and the MCMAS command-line tool. Note that this subsection is in no way intended to be a complete manual for MCMAS. Our aim is to give the reader a basic idea of how the tool can be used for purposes relevant to our project. Please refer to the MCMAS manual [9] for a complete description of the ISPL Syntax and the command-line tool.

ISPL Syntax

An ISPL file contains both the multi-agent system and the properties to be verified. These are internally represented as an interpreted system (see Definition 2.5) and logic formulas (see Section 2.2) respectively. The file consists of the following sections [9]:

- **Semantics** (*optional*) – Statement which defines the evolution function semantics (see below). Two options are available: `MultiAssignment` (default) and `SingleAssignment`.
- **Agent** – Every agent of the multi-agent system is defined separately. Defining the environment is optional. Each specification has the following fields for an agent $i \in \text{Agt}$:

¹A fairness condition is an arbitrary formula φ which rules out paths along which φ is not true infinitely often. This increases the expressiveness of CTL and ATL with the ability to assume that certain unwanted behaviour will not occur (e.g. that no agent will hold a lock forever) [9]. Note that fairness constraints can be expressed naturally in LTL as $[GF\varphi] \rightarrow \psi$ where ψ is the formula that we want to check.

- **Vars** (*optional for environment*) – Internal variables of the agent. Three types of variables are supported in ISPL: Boolean ($x: \text{boolean}$), enumeration ($y: \{a, b, c\}$), and bounded integer ($z: 1 \dots 4$). These variables represent the set of internal states L_i of the agent.
- **Obsvars** (*environment only, optional*) – Variables observable by all agents. The types of variables are the same as in **Vars**. The set of local states L_{jE} of each agent $j \in \text{Agt}$ is induced by the the observable variables (of the environment) and the agent’s internal variables.
- **Lobsvars** (*other agents only, optional*) – Set of local variables of the environment that the agent can observe. This makes ISPL syntax more expressive than interpreted systems because different agents can observe different internal variables of the environment.
- **RedStates** (*optional*) – Constraints on the correct behaviour of the agent for deontic modalities. This field is not relevant for our project.
- **Actions** – Set of actions Act_i available to the agent.
- **Protocol** – Protocol P_i of the agent, which maps local states L_{iE} to sets of actions 2^{Act_i} . It consists of lines of the form *condition*: {*actions*}, where *condition* is a condition on the local state and *actions* is a list of actions available to agent i if the condition holds. Note that the conditions on different lines need not be mutually exclusive. The keyword **Other** represents a condition which is true iff all other conditions fail.
- **Evolution** – Evolution function t_i of the agent. Similarly to the protocol, it consists of lines of the form *next_state* if *condition*, where *condition* is a condition on the local state of the agent and the joint action of all agents.

If multi-assignment semantics is used (default), *next_state* refers to all local variables that change and exactly one condition line applies. If single-assignment semantics is used instead, *next_state* can refer to only one variable and, for each local variable, exactly one condition line applies. Please refer to the MCMAS manual [9] for the exact meaning of the two options.

- **Evaluation** – Atomic propositions AP and their assignment h over global states. For every atomic proposition, the evaluation contains an entry of the form *atom* if *condition*, where *condition* is a condition on the current global state.
- **InitStates** – Boolean formula over all variables which defines the initial states I .
- **Groups** (*optional*) – List of groups of agents. These are used in some epistemic modalities (see Subsection 2.2.6) and ATL temporal operators (see Subsection 2.2.4).
- **Fairness** (*optional*) – List of fairness formulas. This field is not relevant to our project.
- **Formulae** – List of formulas to be verified. Table 6.1 shows how logic formulas are translated into ISPL syntax.

Command-Line Tool

The general usage of the MCMAS command-line tool is as follows:

```
$ ./mcmas [OPTIONS] FILE
```

where *[OPTIONS]* is an optional list of command-line options and *FILE* is the name of the ISPL file. For example, the file `examples/toy_model_existing.ispl` contains the toy model introduced in Section 3.3, a CTL specification $EF p$, and an ATL specification $\langle\langle\{1, 2\}\rangle\rangle G \neg (p_1 \vee p_2)$ (see Subsections 2.2.2 and 2.2.4 for a description of CTL and ATL respectively). These formulas are expressed in the ISPL file as `EF p` and `<g>G !(p1 or p2)` respectively, where **g** is a group containing agents `Player1` and `Player2`. The complete ISPL file is provided in Appendix A. The following command checks the specifications and generates witness/counterexample executions where possible:

```
$ ./mcmas -c 1 examples/toy_model_existing.ispl
```

The `-c` flag for displaying witness/counterexample executions has a numeric parameter 1–3 which specifies the format of the executions. This command generates the following output:

Logic	Formula	ISPL syntax
propositional	p	<i>atom</i>
	$\neg\varphi$	<i>!formula</i>
	$\varphi_1 \wedge \varphi_2$	<i>formula1 and formula2</i>
	$\varphi_1 \vee \varphi_2$	<i>formula1 or formula2</i>
	$\varphi_1 \rightarrow \varphi_2$	<i>formula1 -> formula2</i>
CTL (Subsection 2.2.2)	$AX\varphi$	<i>AX formula</i>
	$EX\varphi$	<i>EX formula</i>
	$AF\varphi$	<i>AF formula</i>
	$EF\varphi$	<i>EF formula</i>
	$AG\varphi$	<i>AG formula</i>
	$EG\varphi$	<i>EG formula</i>
	$A[\varphi_1 U \varphi_2]$	<i>A(formula1 U formula2)</i>
	$E[\varphi_1 U \varphi_2]$	<i>E(formula1 U formula2)</i>
ATL (Subsection 2.2.4)	$\langle\langle A \rangle\rangle X\varphi$	<i><group>X formula</i>
	$\langle\langle A \rangle\rangle F\varphi$	<i><group>F formula</i>
	$\langle\langle A \rangle\rangle G\varphi$	<i><group>G formula</i>
	$\langle\langle A \rangle\rangle[\varphi_1 U \varphi_2]$	<i><group>(formula1 U formula2)</i>
epistemic (Subsection 2.2.6)	$K_i\varphi$	<i>K(agent, formula)</i>
	$E_A\varphi$	<i>GK(group, formula)</i>
	$D_A\varphi$	<i>DK(group, formula)</i>
	$C_A\varphi$	<i>GCK(group, formula)</i>
deontic (not relevant)	$O_i\varphi$	<i>O(agent, formula)</i>

Table 6.1: Translation between logic formulas and ISPL syntax.

```

1 *****
2           MCMAS v1.1.0
3
4   This software comes with ABSOLUTELY NO WARRANTY, to the extent
5   permitted by applicable law.
6
7   Please check http://vas.doc.ic.ac.uk/tools/mcmas/ for the latest release.
8   Please send any feedback to <mcmas@imperial.ac.uk>
9 *****
10
11 Command line: ./mcmas -c 1 examples/toy_model_existing.ispl
12
13 examples/toy_model.ispl has been parsed successfully.
14 Global syntax checking...
15 Done
16 Encoding BDD parameters...
17 Building partial transition relation...
18 Building BDD for initial states...
19 Building reachable state space...
20 Checking formulae...
21 Verifying properties...
22   Formula number 1: (EF p1), is TRUE in the model
23   The following is a witness for the formula:
24     < 0 1 >
25   States description:
26   ----- State: 0 -----
27 Agent Environment
28   state = game
29 Agent Player1

```

```

30 Agent Player2
31 -----
32 ----- State: 1 -----
33 Agent Environment
34   state = p1win
35 Agent Player1
36 Agent Player2
37 -----
38   Formula number 2: (<g>G (! (p1 || p2))), is TRUE in the model
39   The following is a witness for the formula:
40   A witness exists but could not be generated.
41   done, 2 formulae successfully read and checked
42   execution time = 0.007
43   number of reachable states = 3
44   BDD memory in use = 8968368

```

Lines 22 and 38 indicate that the both the CTL formula $EF p_1$ and the ATL formula $\langle\langle\{1, 2\}\rangle\rangle G \neg (p_1 \vee p_2)$ are true in all initial states of the model. A witness execution with 2 states, namely g_g and g_1 (see Section 3.3), for the CTL formula is provided on lines 24–37. According to line 40, MCMAS could not generate a witness execution for the ATL formula.

Other flags supported by MCMAS include deadlock checking (`-k`), arithmetic overflow checking (`-a`), interactive simulation (`-s`), and verbosity level (`-v` with a numeric parameter 1–5). The complete list of flags supported by MCMAS is provided on the tool’s help screen:

```
$ ./mcmas -h
```

6.1.3 Architecture

MCMAS uses binary decision diagrams (see Subsection 2.3.1) to represent the models of systems and implements all algorithms using very efficient symbolic operations on this data structure (see Subsection 2.3.2). The tool is written in C++ (except for the Eclipse plugin, which is implemented in Java) and uses the CUDD BDD package [85]. Over the years, it has been developed collaboratively by members of the Verification of Autonomous Systems Group (VAS) in the Department of Computing at Imperial College London [13].

The aim of this subsection is to provide a high-level description of MCMAS internals and bridge the gap between the theoretical background presented in Chapter 2 and the existing MCMAS functionality described in Subsection 6.1.1. We will first give an overview of the model checking and witness/counterexample generation process and then describe the overall structure of the MCMAS source code. Finally, we will comment on the software design of MCMAS.

Model Checking Process

Assume that MCMAS is executed with the name of an ISPL file as a command-line argument (see Subsection 6.1.2 for a short introduction to MCMAS usage). The tool performs the following steps:

1. **Parser.** The ISPL file is scanned and parsed into an interpreted system \mathcal{I} representation of the model and an abstract syntax tree for each formula $\varphi_0, \dots, \varphi_{n-1}$. The scanner and parser were automatically generated using Flex [4] and Bison [2] respectively.
2. **Syntax check.** The program checks that the the model and the specifications are well-formed. This includes type checking and name resolution. The tool ensures that all expressions (e.g. an agent’s protocol) refer only to existing variables observable by the corresponding agent. If the model or a formula violates any of the constraints, an error message is printed and the whole execution is aborted.
3. **Symbolic encoding.** Each local variable (of an agent) with k possible values is allocated $2 \times \lceil \log_2 k \rceil$ BDD variables to represent its current and next value. Once all local variables of all agents have been allocated BDD variables, the BDD vectors \bar{v}, \bar{v}' for representing the current and next

global state are constructed as their concatenation. Similarly, the BDD vector \bar{u} for representing joint actions is composed of the BDD variables for representing individual actions. Consequently, protocols, evolution functions, initial states, and epistemic accessibility relations are calculated. A detailed example of how all of this can be done is provided in Subsection 3.3.2.

4. **Reachability analysis.** The set of states G reachable from the set of initial states I of the interpreted system \mathcal{I} is calculated. As explained at the end of Subsection 3.3.2, this is performed using a simple fixpoint calculation over the temporal transition relation.

Checks for deadlocks and/or arithmetic overflow in the reachable states are carried out at this stage if the user requests them via command-line flags.

Finally, if there is at least one fairness condition in the ISPL file, the set of reachable fair states, at which some path satisfying all fairness conditions starts, is calculated.

5. **Model checking.** For each formula φ_i with $0 \leq i < n$, MCMAS checks if all initial states I of the interpreted system \mathcal{I} satisfy φ_i , i.e. $I \subseteq \|\varphi_i\|_{\mathcal{I}}$. It does this by introducing an atom $init \in AP$ which holds in the initial states, i.e. $h(init) = I$, and then checking that the formula $init \rightarrow \varphi_i$ is true in all reachable states, i.e. $\|init \rightarrow \varphi_i\|_{\mathcal{I}} = G$.

The tool uses a recursive bottom-up model checking algorithm similar to the one for CTL in Definition 2.12.

6. **Witness/counterexample generation.** If a formula starting with an E path quantifier is true in the model (e.g. $\mathcal{I} \models EGp$), a witness execution is presented. Conversely, if a formula starting with an A path quantifier does not hold in the model (e.g. $\mathcal{I} \not\models AFp$), a counterexample execution is presented. Both witness and counterexample executions are provided for Boolean combinations of formulas. Please refer to [9] for more details about the combinations of CTL and ATL operators supported by witness/counterexample generation in MCMAS.

Note that witness/counterexample generation is performed only when explicitly requested by the user via a command-line flag.

Please refer to the MCMAS source code available at [8] for low-level details of the individual steps.

Source Code Structure

The structure of the MCMAS source code directory is as follows (not all files are listed):

- **main.cc** – Main file of the whole tool. It uses functions and classes defined in the other files to implement the model checking process we discussed earlier. It is responsible for the majority of output and defines global variables used by the other source files.
- **Makefile** – Project build file. This file describes how MCMAS is built from the source files. The whole tool can be compiled easily by executing the **make** command. Build rules for a wide range of platforms including Linux, Windows (using Cygwin), and Mac OS are provided. Both 32 and 64-bit architectures are supported.
- **cudd-2.5.0-exp/** – The CUDD package for manipulating BDDs.
- **doc/** – MCMAS documentation,
- **examples/** – Sample ISPL files.
- **include/** – Header files.
- **parser/** – Files relevant to ISPL file parsing (e.g. token and grammar files) and syntax checking.
- **utilities/** – All other source files.
 - **computereach.cc** – Functions for calculating the set of reachable states.
 - **modal_formula.cc** – Class representing modal formulas. This source file also implements the recursive model checking algorithm and provides the witness/counterexample generation functionality.

- `read_options.cc` – Command-line arguments handler.
- `simulation.cc` – Interactive simulation.
- `utilities.cc` – Model checking functions and custom BDD operations.

Software Design

Overall, MCMAS source code uses a mixture of the *object-oriented paradigm* and *procedural paradigm*:

- Object-oriented paradigm: classes, objects, inheritance, encapsulation
- Procedural paradigm: functions, global variables

As we have already pointed out in Section 1.2, the quality of the code is quite low from a software engineering point of view, despite being developed as open source. We have encountered the following issues during the development of the extensions:

1. **Complete lack of testing.** We believe that this is the *biggest problem* by far as there are no guarantees about the correctness of individual functions, classes, or the program as a whole. Furthermore, it makes any large-scale code refactoring virtually impossible because the programmer has no easy way of verifying that they did not break any existing functionality. It is almost paradoxical that a software verification tool is not verified in any structured way.
2. **Global variables.** The source code of MCMAS uses many global variables, which are generally believed to be a bad programming practice. Consequently, it is difficult to test individual functions and reason about their correctness because of side-effects.
3. **Inconsistent style.** Different parts of the codebase use different indent styles. In addition, some files contain lines with more than 100 characters. This makes the source code difficult to read and understand.
4. **Long functions.** Some of the functions are very long. For example, the body of the method for generating witness/counterexample executions has 2053 lines. It is very difficult to understand and modify such functions.

Ideally, the existing source code should be refactored and augmented with tests. Unfortunately, the current situation is a vicious cycle: Refactoring is difficult without having tests and adding tests is difficult without refactoring first. Therefore, we believe that the whole tool should be redesigned and rewritten from scratch in the future.

Despite these shortcomings, MCMAS is one of the leading tools in the area of symbolic model checking of multi-agent systems. More importantly, its existing source code provided us with solid foundations on which we could build the SLK and SL[1G] extensions. In the next two sections, we will describe how we extended MCMAS with support for SLK and SL[1G], which were discussed in Chapters 4 and 5 respectively.

6.2 Epistemic Strategy Logic Extension

Epistemic Strategy Logic (SLK) is a new fragment of SL, which we introduced in Chapter 4. As it is defined on imperfect recall semantics with incomplete information, it can combine various epistemic and game-theoretic concepts. Consequently, SLK can express complex specifications such as an agent's knowledge about Nash equilibria. In Section 4.2 we provided a model checking algorithm for SLK as well as an efficient symbolic implementation of it. We have developed an extension of MCMAS which puts the algorithm into practice and thereby adds support for SLK model checking to MCMAS. We will now discuss the new functionality of the extension, its usage, and the modified architecture.

SLK Formula	ISPL syntax
$\langle\langle x \rangle\rangle \varphi$	<code><<variable>> formula</code>
$[[x]] \varphi$	<code>[[variable]] formula</code>
$(i, x) \varphi$	<code>(agent, variable) formula</code>
$X \varphi$	<code>X formula</code>
$F \varphi$	<code>F formula</code>
$G \varphi$	<code>G formula</code>
$\varphi_1 \cup \varphi_2$	<code>formula1 U formula2</code>

Table 6.2: Translation between SLK formulas and ISPL syntax. Translation of the logics supported by the original version of MCMAS is shown in Table 6.1.

6.2.1 Functionality

The SLK extension of MCMAS, which we developed as part of this project, has the following new features (see Subsection 6.1.1 for the list of existing features):

1. **Specification languages** (*updated*). Our extension adds support for SLK specifications to MCMAS. Moreover, it can handle arbitrary nesting of CTL, ATL, and SLK operators.
2. **Model checking** (*updated*). We have incorporated the efficient symbolic implementation of our SLK model checking algorithm (see Subsection 4.2.4) in the extension.
In order to improve its performance, we have developed an *experimental parallel implementation* of the SLK model checking algorithm which can be enabled using a command-line flag (see Subsection 6.2.2).
3. **Witness and counterexample executions** (*updated*). We have added support for SLK operators to the witness/counterexample execution generator.
4. **Witness and counterexample strategies** (*new*). We have implemented witness and counterexample strategy synthesis for SLK formulas as explained in Subsection 4.2.3. If an SLK formula $\langle\langle x_0 \rangle\rangle \dots \langle\langle x_{m-1} \rangle\rangle \psi$ holds in the model, witness strategies for the variables x_0, \dots, x_{m-1} are synthesised (when requested by the user). Conversely, if an SLK formula $[[y_0]] \dots [[y_{n-1}]] \phi$ does not hold in the model, counterexample strategies for the variables y_0, \dots, y_{n-1} are synthesised.

To sum up, the extension adds support for SLK together with witness and counterexample strategy synthesis. This is a major enhancement of MCMAS functionality as it provides the possibility to automatically synthesise agents' behaviour which satisfies an SLK specification.

6.2.2 Usage

The SLK extension augments ISPL syntax with the new SLK operators (see Definition 4.1). The translation from SLK formulas to ISPL syntax is shown in Table 6.2. Apart from that, the usage of the SLK extension is the same as that of the original tool (see Subsection 6.1.2).

Consider the file `examples/toy_model_slk.ispl` containing the toy model introduced in Section 3.3 and the SLK specification $\langle\langle e \rangle\rangle (E, e) \langle\langle x \rangle\rangle (1, x) \langle\langle y \rangle\rangle (2, y) G \neg (p_1 \vee p_2)$, which means that “*there exist strategies for all agents such that neither player will ever win*”. This formula is expressed in ISPL syntax as `<<e>> (Environment, e) <<x>> (Player1, x) <<y>> (Player2, y) G !(p1 or p2)`. The complete ISPL file is provided in Appendix A. The following command checks the specification and generates a witness/counterexample execution with strategies (where possible):

```
$ ./mcmas -c 1 examples/toy_model_slk.ispl
```

The SLK extension generates the following output:

```
1 *****
2                               MCMAS-SL v1.1.0
3
```

```

4   This software comes with ABSOLUTELY NO WARRANTY, to the extent
5   permitted by applicable law.
6
7   Please check http://vas.doc.ic.ac.uk/tools/mcma/ for the latest release.
8   Please send any feedback to <mcma@imperial.ac.uk>
9   *****
10
11  Command line: ./mcma -c 1 examples/toy_model_slk.ispl
12
13  examples/toy_model.ispl has been parsed successfully.
14  Global syntax checking...
15  Done
16  Encoding BDD parameters...
17  Building partial transition relation...
18  Building BDD for initial states...
19  Building reachable state space...
20  Checking formulae...
21  Verifying imperfect recall properties...
22   Formula number 1: <<e>> (Environment, e) <<x>> (Player1, x) <<y>> (Player2, y)
23   G (! (p1 || p2)), is TRUE in the model
24   The following is a witness for the formula:
25   < 0 0 >
26   States description:
27   ----- State: 0 -----
28   Agent Environment
29   state = game
30   Agent Player1
31   Agent Player2
32   -----
33   Strategies:
34   ----- Strategy e [Environment] -----
35   Agent Environment
36   Environment.state=game (0): idle
37   Environment.state=p1win (1): idle
38   Environment.state=p2win (2): idle
39   -----
40   ----- Strategy x [Player1] -----
41   Agent Player1
42   Environment.state=game (0): rock
43   Environment.state=p1win (1): idle
44   Environment.state=p2win (2): idle
45   -----
46   ----- Strategy y [Player2] -----
47   Agent Player2
48   Environment.state=game (0): rock
49   Environment.state=p1win (1): idle
50   Environment.state=p2win (2): idle
51   -----
52  done, 1 imperfect recall formulae successfully read and checked
53  execution time = 0.006
54  number of reachable states = 3
55  BDD memory in use = 8989648

```

Line 23 indicates that the SLK formula is true. A witness execution, namely a cycle in the state g_g , is provided on lines 25–32. More importantly, the following *witness strategies* for the SLK specification are

provided on lines 34–51:

$$\begin{array}{lll}
 f_e(g_g) \triangleq i & f_x(g_g) \triangleq r & f_y(g_g) \triangleq r \\
 f_e(g_1) \triangleq i & f_x(g_1) \triangleq i & f_y(g_1) \triangleq i \\
 f_e(g_2) \triangleq i & f_x(g_2) \triangleq i & f_y(g_2) \triangleq i
 \end{array}$$

The extension adds a new command-line flag `-t`, which enables the experimental parallel implementation of the SLK model checking algorithm. It has a non-negative numeric parameter which defines the maximum depth of thread branching, e.g. `-t 3` allows at most $2^3 = 8$ threads. This functionality is disabled by default.

6.2.3 Architecture

In order to add SLK support to MCMAS, we had to modify its existing source code and add new files. We will first give a high-level overview of how we updated the steps of the original model checking process. Then we shall list the modified and new source files. Finally, we will discuss the experimental parallel implementation of the model checking algorithm.

Model Checking Process

The original MCMAS model checking process (see Subsection 6.1.3) has been modified in the SLK extension as follows:

1. **Parser.** The ISPL token and grammar files were augmented with SLK syntax (see Table 6.2).
2. **Syntax check.** The updated syntax checking procedure verifies that all SLK specifications are sentences (see Definition 4.1), i.e. that all agents are bound to a strategy when a temporal transition occurs.
3. **Symbolic encoding.** The symbolic encoding of the global states and joint actions is unchanged. Since the number of BDD variables needed to represent variable assignments (see Definition 4.7) depends on the SLK formula, they are allocated for each SLK specification separately in the recursive model checking algorithm (step 5). More precisely, BDD variables for representing the strategy assigned to a variable x are allocated when the corresponding quantifier $\langle\langle x \rangle\rangle$ or $\llbracket x \rrbracket$ is encountered.
4. **Reachability analysis.** This step is completely unchanged in the extension.
5. **Model checking.** We have augmented the existing recursive symbolic model checking algorithm of MCMAS with the symbolic implementation of our SLK model checking algorithm presented in Subsection 4.2.4.
6. **Witness/counterexample generation.** We added support for SLK to the existing procedure for generating witness/counterexample executions. More importantly, we implemented SLK witness/counterexample strategy synthesis, which we discussed in Subsection 4.2.3.

Please refer to the attached project source code for low-level details of the individual steps.

Source Code Structure

We modified and augmented the original MCMAS codebase (see Subsection 6.1.3) while developing the SLK extension. The following files were added or updated (not all files are listed):

- `main.cc` – We made only minor changes to the main file (e.g. added a function for printing witness strategies).
- `Makefile` – We added build rules for the new files.
- `examples/` – We added several sample ISPL files which demonstrate the new functionality of the extension.

- `include/` – We included prototype declarations for new classes in the existing header files. In addition, we added separate header files for benchmarking and the SLK model checking algorithm functions.
- `parser/` – We modified the ISPL token and grammar files.
- `utilities/` – All new features were implemented in this folder:
 - `benchmark.cc` (*new*) – Toolkit for analysing the performance of both extensions (SLK and SL[1G]). We developed it specifically for this project in order to find bottlenecks in the algorithms.
 - `modal_formula.cc` – We extended the syntax checking procedure, the recursive model checking algorithm, and the function for generating witnesses/counterexamples.
 - `read_options.cc` – We added the `-t` command-line flag, which enables experimental parallelisation of SLK model checking.
 - `sl_imperfect_recall.cc` (*new*) – This file contains the core SLK model checking procedures, which are invoked by the recursive model checking algorithm in `modal_formula.cc`. The core functionality includes calculating the shared state space of a quantifier, allocating BDD variables for strategies, and SLK temporal operators. This file also provides the experimental parallel implementation of the model checking algorithm.
 - `strategy.cc` (*new*) – Class representing SLK strategies.

Parallel Implementation

As we have explained in Subsection 4.2.4, our symbolic implementation of the SLK model checking algorithm represents a strategy $f \in UStr_A$ for a set of agents $A \subseteq Agt$ as a mapping from shared local states $S \in G/\sim_A^C$ to actions. In other words, for every shared local state $S \in G/\sim_A^C$, we store the corresponding action using a separate set of BDD variables. Therefore, we must explicitly enumerate the shared local state space. The algorithm which performs this computation for a given set of agents $A \subseteq Agt$ is shown in Figure 6.1.

Our experimental results indicate that the enumeration of local states (function LOCALSTATES in Figure 6.1) is a *major bottleneck* of the algorithm. In order to improve the performance of our algorithm, we created an experimental *parallel implementation* of the procedure which uses multiple threads. It spawns a separate thread for one of the two recursive function calls on line 17 of the LOCALSTATES function in Figure 6.1. Each time such a branching occurs, both threads are left with roughly one half of the original state space. In order not to create too many threads, branching only occurs in the t top-most recursive calls, where t is the maximum depth of thread branching (set via the `-t` command-line flag). More precisely, a new thread is spawned only when $n < t$, where n is the fourth argument of LOCALSTATES and represents the current recursion depth. The impact of this optimisation on the performance of the model checking algorithm is shown in Table 6.3.

While the results indicate that the optimisation does reduce model checking time, the approach appears not to be very scalable. The problem is that the CUDD BDD package has *no support for concurrency*. In order to avoid race conditions, each thread has to have a separate BDD manager². After all the local states are enumerated by the threads, the corresponding BDDs (one per local state) must be *transferred* to a single BDD manager. Since each manager might be using a completely different variable ordering (see Subsection 2.3.1), this transfer soon becomes the bottleneck when multiple threads are used. The only solution to this problem, apart from rewriting the CUDD library, is to use a different BDD package which supports concurrency. Given the fact that MCMAS already uses CUDD extensively and the limited time for our project, we decided not pursue this direction any further. Instead, we devoted all of our remaining time to the development of a model checking algorithm for SL[1G], which we then implemented as another extension for MCMAS (see the next section).

²A BDD manager stores a hash table for BDD nodes, which ensures the canonicity of the BDDs, and other auxiliary data structures. Every BDD is handled by a BDD manager. More importantly, the arguments of BDD operations must have the same manager.

```

1 function SHAREDSTATES( $A, \bar{v}, G$ )                                ▷ Enumerate shared local states of agents  $A$ .
2   for  $k := 0$  to  $|A| - 1$  do
3      $\hat{L}_{A(k)} := \text{LOCALSTATES}(A(k), \bar{v}, G(\bar{v}), 0)$ 
4   end for
5    $S := \hat{L}_{A(0)}$ 
6   for  $k := 1$  to  $|A| - 1$  do
7      $S := \text{MERGE}(S, \hat{L}_{A(k)})$ 
8   end for
9   return  $S$ 
10 end function

11 function LOCALSTATES( $i, \bar{v}, \lambda, n$ )                            ▷ Enumerate local states of agent  $i$ .
12   if  $\lambda = \perp$  then
13     return  $\emptyset$                                               ▷  $\lambda$  is not reachable.
14   else if  $n = |\overline{v_{iE}}|$  then
15     return  $\{\lambda\}$ 
16   else
17     return  $\text{LOCALSTATES}(i, \bar{v}, \lambda \wedge \neg \overline{v_{iE}}(n), n + 1) \cup \text{LOCALSTATES}(i, \bar{v}, \lambda \wedge \overline{v_{iE}}(n), n + 1)$ 
18   end if
19 end function

20 function MERGE( $S_1, S_2$ )                                          ▷ Merge state spaces  $S_1$  and  $S_2$ .
21   for  $j := 0$  to  $|S_2| - 1$  do
22      $S := \{s \in S_1 \mid s(\bar{v}) \wedge S_2(j)(\bar{v}) \neq \perp\}$       ▷ Find states  $S \subseteq S_1$  compatible with  $S_2(j)$ 
23      $S_1 := S_1 \setminus S$                                           ▷ Remove the compatible states  $S$ .
24      $S_1 := S_1 \cup \{S_2(j)(\bar{v}) \vee \bigvee_{s \in S} s(\bar{v})\}$         ▷ Merge  $S_2(j)$  and the compatible states  $S$ .
25   end for
26   return  $S_1$ 
27 end function

```

Figure 6.1: Functions for enumerating the set of *shared local states* $S = G / \sim_A^C$ for a set of agents $A \subseteq \text{Agt}$. \bar{v} is a Boolean vector for representing the current global state and G is the set of global states.

6.3 One-Goal Strategy Logic Extension

One-goal Strategy Logic (SL[1G]) is another fragment of SL, which we discussed in Chapter 5. Unlike SLK, it is defined with respect to perfect recall semantics and complete information, i.e. agents have both perfect memory of the past and complete knowledge of the system. In Section 5.2 we provided a model checking algorithm for SL[1G] as well as an efficient symbolic implementation of it. Furthermore, we have developed an extension of MCMAS which puts the algorithm into practice and thereby adds support for SL[1G] model checking and strategy synthesis to MCMAS. We will now discuss the new functionality of the extension, its usage, and the modified architecture.

6.3.1 Functionality

The SL[1G] extension of MCMAS, which we developed as part of this project, has the following new features (see Subsection 6.1.1 for the list of existing features):

1. **Specification languages** (*updated*). Our extension adds support for SL[1G] specifications to MCMAS. The new SL[1G] syntax is separate from the other logics (CTL, ATL, and SLK) because the fragment does not support epistemic modalities (due to complete information) and uses a very different model checking algorithm based on ω -automata (see Section 5.2).
2. **Model checking** (*updated*). We have incorporated the efficient symbolic implementation of our SL[1G] model checking algorithm (see Subsection 5.2.4) in the extension. Hence, it now supports the verification of SL[1G] specifications.

We implemented both the original and the *optimised version* of our algorithm (see Subsection 5.2.5).

3. **Strategy synthesis** (*new*). We implemented full SL[1G] strategy synthesis, which we discussed in Subsection 5.2.3 and explained how it can be implemented symbolically in Subsection 5.2.4. Unlike the SLK extension (see Subsection 6.2.1), the SL[1G] extension supports strategy synthesis for *arbitrary* SL[1G] formulas. However, unlike SLK witness and counterexample strategies, the synthesised strategies may be inter-dependent, i.e. the next action of one strategy in the current state might depend on the next actions of other strategies (see Subsection 5.2.3).
4. **Witness and counterexample strategies** (*new*). Although witness and counterexample strategies for formulas of the form $\langle\langle x_0 \rangle\rangle \dots \langle\langle x_{m-1} \rangle\rangle \psi$ and $\llbracket y_0 \rrbracket \dots \llbracket y_{n-1} \rrbracket \phi$ are technically subsumed by full strategy synthesis (previous point), the extension provides the option to present them differently. Rather than generating a full solution with inter-dependent strategies for all variables (see Figure 5.2), a separate strategy is constructed for each variable in the existential or universal quantification prefix. Both representations are demonstrated in Subsection 6.3.2.

To sum up, the extension adds support for SL[1G] model checking and strategy synthesis. Similarly to the SLK extension presented in Section 6.2, this is a major enhancement of MCMAS functionality as it provides the possibility to automatically synthesise agents' behaviour which satisfies an SL[1G] specification.

As we have pointed out in Section 3.2, SL[1G] strictly subsumes ATL*. Therefore, all logics in the ATL* hierarchy which were previously not supported by MCMAS, notably LTL, CTL*, and ATL* (see Subsections 2.2.1, 2.2.3, and 2.2.4), are now indirectly³ supported thanks to our extension. Moreover, the extension has optimal time complexity⁴ for ATL* model checking and strategy synthesis (see Table 2.1 and Theorem 5.4). This is a major achievement because there are *no existing tools* that would support ATL* specifications (as far as we know).

³The formulas must be rewritten into equivalent SL[1G] formulas in order to be checked by MCMAS, e.g. an LTL formula ψ has to be transformed to the equivalent SL[1G] formula $\llbracket x_0 \rrbracket \dots \llbracket x_{|Agt|-1} \rrbracket (Agt(0), x_0) \dots (Agt(|Agt|-1), x_{|Agt|-1}) \psi$. This is only a matter of adding extra ISPL syntax which would automatically perform the translation.

⁴This is not the case for LTL and CTL*, which both have PSPACE-COMplete model checking complexity with respect to the size of the formula (see Table 2.1).

6.3.2 Usage

The $SL[1G]$ extension augments ISPL syntax with the new $SL[1G]$ operators (see Definition 5.2). The translation from $SL[1G]$ formulas to ISPL syntax is the same as for SLK formulas (see Table 6.2). Unlike SLK formulas, $SL[1G]$ formulas cannot contain operators of the other modal logics supported by MCMAS. In order to avoid ambiguity (SLK and $SL[1G]$ use the same operators), every $SL[1G]$ formula must be preceded by a #PR (perfect recall) tag in an ISPL file:

```

1  ...
2  Formulae
3  -- SLK formula
4  <<e>> (Environment, e) <<x>> (Player, x) <<y>> (Player2, y) G !(p1 or p2)
5  -- SL[1G] formula
6  #PR <<e>> (Environment, e) <<x>> (Player, x) <<y>> (Player2, y) G !(p1 or p2)
7  end Formulae
8  ...

```

The `Formulae` section contains the formula $\langle\langle e \rangle\rangle(E, e)\langle\langle x \rangle\rangle(1, x)\langle\langle y \rangle\rangle(2, y)G \neg (p_1 \vee p_2)$ twice: Line 4 represents an SLK formula (imperfect recall with incomplete information) whereas line 6 represents an $SL[1G]$ formula (perfect recall with complete information). Apart from the #PR tag, the usage of the $SL[1G]$ extension is the same as that of the original tool and the SLK extension (see Subsections 6.1.2 and 6.2.2).

Witness/Counterexample Strategy Synthesis

We first demonstrate the ability of the $SL[1G]$ extension to synthesise witness/counterexample strategies and its limitations. Consider the file `examples/toy_model_sl1g.ispl` containing the toy model introduced in Section 3.3 and the $SL[1G]$ specification $\langle\langle e \rangle\rangle(E, e)\langle\langle x \rangle\rangle(1, x)\langle\langle y \rangle\rangle(2, y)G \neg (p_1 \vee p_2)$, which means that “*there exist strategies for all agents such that neither player will ever win*”. This formula is expressed in ISPL syntax as #PR <<e>> (Environment, e) <<x>> (Player1, x) <<y>> (Player2, y) G !(p1 or p2). The complete ISPL file is provided in Appendix A. The following command checks the specification and generates witness/counterexample strategies:

```
$ ./mcmas -c 1 examples/toy_model_sl1g.ispl
```

The $SL[1G]$ extension generates the following output (additional line breaks were inserted so that the output would fit on a page):

```

1  *****
2  MCMAS-SL v1.1.0
3
4  This software comes with ABSOLUTELY NO WARRANTY, to the extent
5  permitted by applicable law.
6
7  Please check http://vas.doc.ic.ac.uk/tools/mcmas/ for the latest release.
8  Please send any feedback to <mcmas@imperial.ac.uk>
9  *****
10
11 Command line: ./mcmas -c 1 examples/toy_model_sl1g.ispl
12
13 examples/toy_model_sl1g.ispl has been parsed successfully.
14 Global syntax checking...
15 Done
16 Global reachable states:
17   g0: Environment.state=game (initial)
18   g1: Environment.state=p1win
19   g2: Environment.state=p2win
20 Perfect recall formula 1: <<e>> (Environment, e) <<x>> (Player1, x) <<y>>
21   (Player2, y) G (! (p1 || p2)), is TRUE in the model

```

```

22 Witness 1/1 (<<e>> (Environment, e) <<x>> (Player1, x) <<y>> (Player2, y)
23       G (! (p1 || p2))):
24 Strategy <<e>>:
25   start:
26     g0->s0 (initial)
27     g1->s1
28     g2->s2
29   s0 (winning):
30     action: idle
31     g0->s3
32     g1->s4
33     g2->s5
34   s1:
35     action: <no winning action>
36     g1->s6
37   s2:
38     action: <no winning action>
39     g2->s7
40   s3 (winning):
41     action: idle
42     g0->s3
43     g1->s4
44     g2->s5
45   s4:
46     action: <no winning action>
47     g1->s6
48   s5:
49     action: <no winning action>
50     g2->s7
51   s6:
52     action: <no winning action>
53     g1->s6
54   s7:
55     action: <no winning action>
56     g2->s7
57 Strategy <<x>>:
58   start:
59     g0->s0 (initial)
60     g1->s1
61     g2->s2
62   s0 (winning):
63     action: paper
64     g0->s3
65     g1->s4
66     g2->s5
67   s1:
68     action: <no winning action>
69     g1->s6
70   s2:
71     action: <no winning action>
72     g2->s7
73   s3 (winning):
74     action: paper
75     g0->s3
76     g1->s4
77     g2->s5

```

```

78     s4:
79         action: <no winning action>
80         g1->s6
81     s5:
82         action: <no winning action>
83         g2->s7
84     s6:
85         action: <no winning action>
86         g1->s6
87     s7:
88         action: <no winning action>
89         g2->s7
90 Strategy <<y>>:
91     start:
92         g0->s0 (initial)
93         g1->s1
94         g2->s2
95     s0 (winning):
96         action: paper
97         g0->s3
98         g1->s4
99         g2->s5
100    s1:
101        action: <no winning action>
102        g1->s6
103    s2:
104        action: <no winning action>
105        g2->s7
106    s3 (winning):
107        action: paper
108        g0->s3
109        g1->s4
110        g2->s5
111    s4:
112        action: <no winning action>
113        g1->s6
114    s5:
115        action: <no winning action>
116        g2->s7
117    s6:
118        action: <no winning action>
119        g1->s6
120    s7:
121        action: <no winning action>
122        g2->s7
123 done, 1 perfect recall formulae successfully read and checked
124 execution time = 0.015
125 number of reachable states = 3
126 BDD memory in use = 9163152

```

Line 21 indicates that the $SL[1G]$ formula is true. Witness strategies for the variables e , x , and y are provided on lines 24–122. Unlike SLK witness strategies (see Subsection 6.2.2), which are memoryless, $SL[1G]$ witness strategies require *finite memory*. Thus, they are presented to the user in the form of automata where each automaton state corresponds to one memory state. Since the $SL[1G]$ formula is true in the initial state g_g , the underlying LTL formula $G \neg (p_1 \vee p_2)$ will be true along the path if all agents follow the corresponding witness strategies.

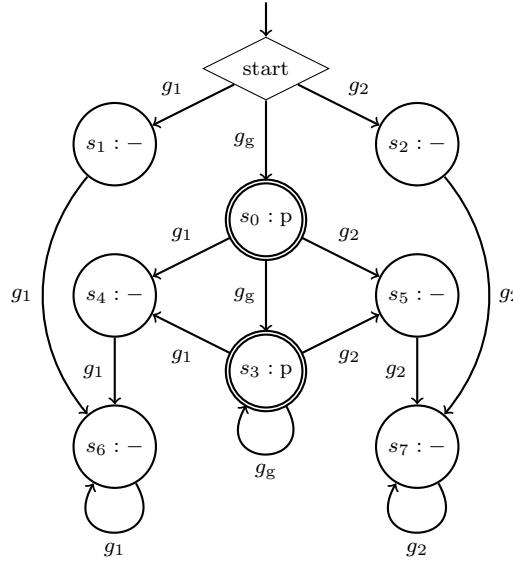


Figure 6.2: Finite memory witness strategy for variable x in the $\text{SL}[1\text{G}]$ formula $\langle\langle e \rangle\rangle(\text{E}, e)\langle\langle x \rangle\rangle(1, x)\langle\langle y \rangle\rangle(2, y)\text{G}\neg(p_1 \vee p_2)$ synthesised by the $\text{SL}[1\text{G}]$ extension for the toy model. Winning memory states have double borders and contain the actions that should be performed next, e.g. paper should be played in the memory state s_3 . Memory updates are represented by transitions labelled with the new global states.

We shall now explain how the witness strategies synthesised by the extension work. The automaton representation of the strategy for variable x is shown in Figure 6.2. The initial memory state of the automaton depends on the initial global state. This dependency is represented by the start node. For example, if the initial global state is g_g , the initial memory state is s_0 . If the $\text{SL}[1\text{G}]$ formula can be enforced from a memory state, it is referred to as a winning memory state and is assigned the next action to perform, e.g. the next action in s_3 is paper. The memory state is updated upon every temporal transition with the new global state. For example, if the current memory state is s_3 and a temporal transition to the global state g_1 occurs, the new memory state is s_4 . Note that neither the next action nor the memory update depends directly on the current global state.

Witness strategies can only be synthesised for existential quantification prefixes $\langle\langle x_0 \rangle\rangle \dots \langle\langle x_{m-1} \rangle\rangle$ in states where the formula holds. Conversely, counterexample strategies can only be synthesised for universal quantification prefixes $\llbracket y_0 \rrbracket \dots \llbracket y_{n-1} \rrbracket$ in states where the formula does not hold. Consider the much more interesting $\text{SL}[1\text{G}]$ sentence $\gamma \triangleq \llbracket e \rrbracket(\text{E}, e)\llbracket x \rrbracket(1, x)\langle\langle y \rangle\rangle(2, y)\text{G}\neg(p_1 \vee p_2)$ used throughout Chapter 5, which means roughly: “*Whichever action player 1 performs, there exists an action for player 2 such that neither player will ever win*”. Counterexample strategies can be synthesised for the universal quantification prefix $\llbracket e \rrbracket\llbracket x \rrbracket$ of γ in states g_1 and g_2 , where it does not hold. However, we are more interested in synthesising a strategy for player 2 (variable y). Unfortunately, this cannot be done using witness strategy synthesis because the strategy depends on the strategies for the environment and player 1 (variables e and x). General strategy synthesis, which we describe next, has to be used in this case.

General Strategy Synthesis

We now demonstrate the more general ability of the $\text{SL}[1\text{G}]$ extension to synthesise arbitrary strategies. Consider the file `examples/toy_model_sl1g2.ispl` containing the toy model introduced in Section 3.3 and the $\text{SL}[1\text{G}]$ specification $\gamma \triangleq \llbracket e \rrbracket(\text{E}, e)\llbracket x \rrbracket(1, x)\langle\langle y \rangle\rangle(2, y)\text{G}\neg(p_1 \vee p_2)$ discussed earlier. This formula is expressed in ISPL syntax as `#PR [[e]] (Environment, e) [[x]] (Player1, x) <<y>> (Player2, y) G !(p1 or p2)`. The complete ISPL file is provided in Appendix A. The following command checks the specification γ and synthesises strategies for all variables in it:

```
$ ./mcmas -solutions examples/toy_model_sl1g2.ispl
```

where `-solutions` is a new command-line flag introduced by the `SL[1G]` extension which enables general strategy synthesis. The following output is generated by the extension (additional line breaks were inserted so that the output would fit on a page):

```

1 *****
2           MCMAS-SL v1.1.0
3
4   This software comes with ABSOLUTELY NO WARRANTY, to the extent
5   permitted by applicable law.
6
7   Please check http://vas.doc.ic.ac.uk/tools/mcmas/ for the latest release.
8   Please send any feedback to <mcmas@imperial.ac.uk>
9 *****
10
11 Command line: ./mcmas -solutions examples/toy_model_sl1g2.ispl
12
13 examples/toy_model_sl1g2.ispl has been parsed successfully.
14 Global syntax checking...
15 Done
16 Encoding BDD parameters...
17 Building partial transition relation...
18 Building BDD for initial states...
19 Building reachable state space...
20 Checking formulae...
21 Verifying imperfect recall properties...
22 done, 0 imperfect recall formulae successfully read and checked
23 Verifying perfect recall properties...
24   Global reachable states:
25     g0: Environment.state=game (initial)
26     g1: Environment.state=p1win
27     g2: Environment.state=p2win
28   Perfect recall formula 1: [[e]] (Environment, e) [[x]] (Player1, x) <<y>>
29                           (Player2, y) G (! (p1 || p2)), is TRUE in the model
30   Solution 1/1 ([[e]] (Environment, e) [[x]] (Player1, x) <<y>> (Player2, y)
31                 G (! (p1 || p2))):
32   start:
33     g0->s0 (initial)
34     g1->s1
35     g2->s2
36   s0 (winning):
37     temporal transition:
38     g0->s3
39     g1->s4
40     g2->s5
41   strategies:
42     [[e]]=idle:
43     [[x]]=paper:
44       <<y>>=paper*: g0
45       <<y>>=rock: g1
46       <<y>>=scissors: g2
47     [[x]]=rock:
48       <<y>>=paper: g2
49       <<y>>=rock*: g0
50       <<y>>=scissors: g1
51     [[x]]=scissors:
52       <<y>>=paper: g1

```



```

53         <<y>>=rock: g2
54         <<y>>=scissors*: g0
55     s1:
56         temporal transition:
57         g1->s6
58         strategies:
59         [[e]]=idle*:
60         [[x]]=idle*:
61         <<y>>=idle: g1
62     s2:
63         temporal transition:
64         g2->s7
65         strategies:
66         [[e]]=idle*:
67         [[x]]=idle*:
68         <<y>>=idle: g2
69     s3 (winning):
70         temporal transition:
71         g0->s3
72         g1->s4
73         g2->s5
74         strategies:
75         [[e]]=idle:
76         [[x]]=paper:
77         <<y>>=paper*: g0
78         <<y>>=rock: g1
79         <<y>>=scissors: g2
80         [[x]]=rock:
81         <<y>>=paper: g2
82         <<y>>=rock*: g0
83         <<y>>=scissors: g1
84         [[x]]=scissors:
85         <<y>>=paper: g1
86         <<y>>=rock: g2
87         <<y>>=scissors*: g0
88     s4:
89         temporal transition:
90         g1->s6
91         strategies:
92         [[e]]=idle*:
93         [[x]]=idle*:
94         <<y>>=idle: g1
95     s5:
96         temporal transition:
97         g2->s7
98         strategies:
99         [[e]]=idle*:
100        [[x]]=idle*:
101        <<y>>=idle: g2
102     s6:
103        temporal transition:
104        g1->s6
105        strategies:
106        [[e]]=idle*:
107        [[x]]=idle*:
108        <<y>>=idle: g1

```

```

109     s7:
110         temporal transition:
111             g2->s7
112         strategies:
113             [[e]]=idle*:
114             [[x]]=idle*:
115             <<y>>=idle: g2
116     done, 1 perfect recall formulae successfully read and checked
117     execution time = 0.046
118     number of reachable states = 3
119     BDD memory in use = 9195888

```

Line 29 indicates that γ is true in the initial state g_g . The synthesised strategies are provided on lines 32–115 in the form of one large *automaton* which encodes both the memory and interdependencies between strategies. Consider the memory state s_0 , which is winning, i.e. γ can be enforced by the existentially quantified strategies. It has two fields:

1. **temporal transition.** This field defines how the memory state is updated upon temporal transition. The memory update is the same as for witness SL[1G] strategies presented earlier.

For example, if the current memory state is s_0 and the next global state is g_0 , the next memory state is s_3 .

2. **strategies** This field represents how the next action of each strategy depends on the actions of other strategies. If an existential strategy can enforce γ from its current memory state, its next action is marked with *. Conversely, if a universal strategy can falsify γ from its current memory state, its next action is also marked with *. The next global state for each combination of actions is also provided for convenience.

For example, if the current memory state (of the strategy for variable y) is s_0 and the next actions of strategies for variables e and x are idle and scissors respectively, the next action of the strategy for variable y is scissors. The next global state will then be g_0 , which in turn determines the next memory state s_3 (according to the **temporal transition** field).

Observe that the synthesised solution has a very similar structure to the combined parity game \mathfrak{G}_I^γ in Figure 5.2. This is a direct consequence of the SL[1G] strategy synthesis procedure discussed in Subsection 5.2.3.

Separate Determinisation

The separate determinisation optimisation technique for the SL[1G] model checking algorithm proposed in Subsection 5.2.5 is *enabled* by default. It can be disabled using the command-line flag `-sd 0`.

6.3.3 Architecture

In order to add SL[1G] support to MCMAS, we had to modify its existing source code and add new files. We will first give a high-level overview of how we updated the steps of the original model checking process and then list the modified and new source files.

Model Checking Process

The original MCMAS model checking process (see Subsection 6.1.3) has been modified in the SL[1G] extension as follows:

1. **Parser.** The ISPL token and grammar files were augmented with SL[1G] syntax (see Table 6.2).
2. **Syntax check.** Since SL[1G] formulas cannot be mixed with other logics and have a distinct syntax (thanks to the `#PR` tag), we implemented the syntax check as a separate procedure. It verifies that all SL[1G] specifications are sentences that satisfy the syntactic constraints (see Definition 5.2).

3. **Symbolic encoding.** The symbolic encoding of the global states and joint actions is unchanged. Since the number of BDD variables needed to represent the auxiliary model checking data structures (see Subsection 5.2.1) depends on the SL[1G] formula, they are allocated for each SL[1G] specification separately in the model checking algorithm (step 5).
4. **Reachability analysis.** This step is completely unchanged in the extension.
5. **Model checking.** We implemented the symbolic model checking algorithm for SL[1G] presented in Subsection 5.2.4. Similarly to the syntax checking procedure (step 2), the algorithm is separate from the one for the other logics (including SLK).
6. **Witness/counterexample generation.** The method for SL[1G] witness/counterexample strategy synthesis is again separate from the one for the other logics supported by MCMAS. It reuses the solution calculated by the procedure for general strategy synthesis (step 7).
7. **General strategy synthesis.** We implemented general strategy synthesis symbolically as explained in Subsection 5.2.4. Since this functionality is not provided by either the original tool or the SLK extension, it was developed as a completely new procedure.

Please refer to the attached project source code for low-level details of the individual steps.

Source Code Structure

We modified and augmented the original MCMAS codebase (see Subsection 6.1.3) while developing the SL[1G] extension. The following files were added or updated (not all files are listed):

- `main.cc` – We made only minor changes to the main file (e.g. added a function for printing reachable states).
- `Makefile` – We added build rules for the new files.
- `examples/` – We added several sample ISPL files which demonstrate the new functionality of the extension.
- `include/` – We included prototype declarations for new classes in the existing header files. In addition, we added separate header files for benchmarking and the SL[1G] model checking algorithm functions and classes.
- `parser/` – We modified the ISPL token and grammar files.
- `utilities/` – All new features were implemented in this folder:
 - `arena.cc`, `automaton.cc`, `game.cc`, `generalized_automaton.cc`, `intermediate_automaton.cc`, `nondeterministic_automaton.cc`, `parity_automaton.cc` (*new*) – Classes representing auxiliary data structures used by the SL[1G] model checking algorithm (see Subsection 5.2.1). A UML diagram for all these classes is shown in Figure 6.3.
 - `benchmark.cc` (*new*) – Toolkit for analysing the performance of both extensions (SLK and SL[1G]). We developed it specifically for this project in order to find bottlenecks in the algorithms.
 - `modal_formula.cc` – We implemented the (separate) syntax checking procedure for SL[1G] formulas.
 - `quantifier.cc` (*new*) – Class representing a strategy quantifier (e.g. $\langle\langle x \rangle\rangle$).
 - `read_options.cc` – We added the `-solutions` and `-sd` command-line flags, which enable general strategy synthesis and toggle separate determinisation respectively.
 - `sl_perfect_recall.cc` (*new*) – This is the main file of the SL[1G] extension. It contains the implementation of the recursive model checking algorithm with strategy synthesis (see Subsection 5.2.4).
 - `sl_solution.cc` (*new*) – Class representing the solution of an SL[1G] formula. It contains (*i*) the set of global states in which the formula holds and (*ii*) the synthesised strategies with interdependencies represented as one large automaton (see Subsection 6.3.2).

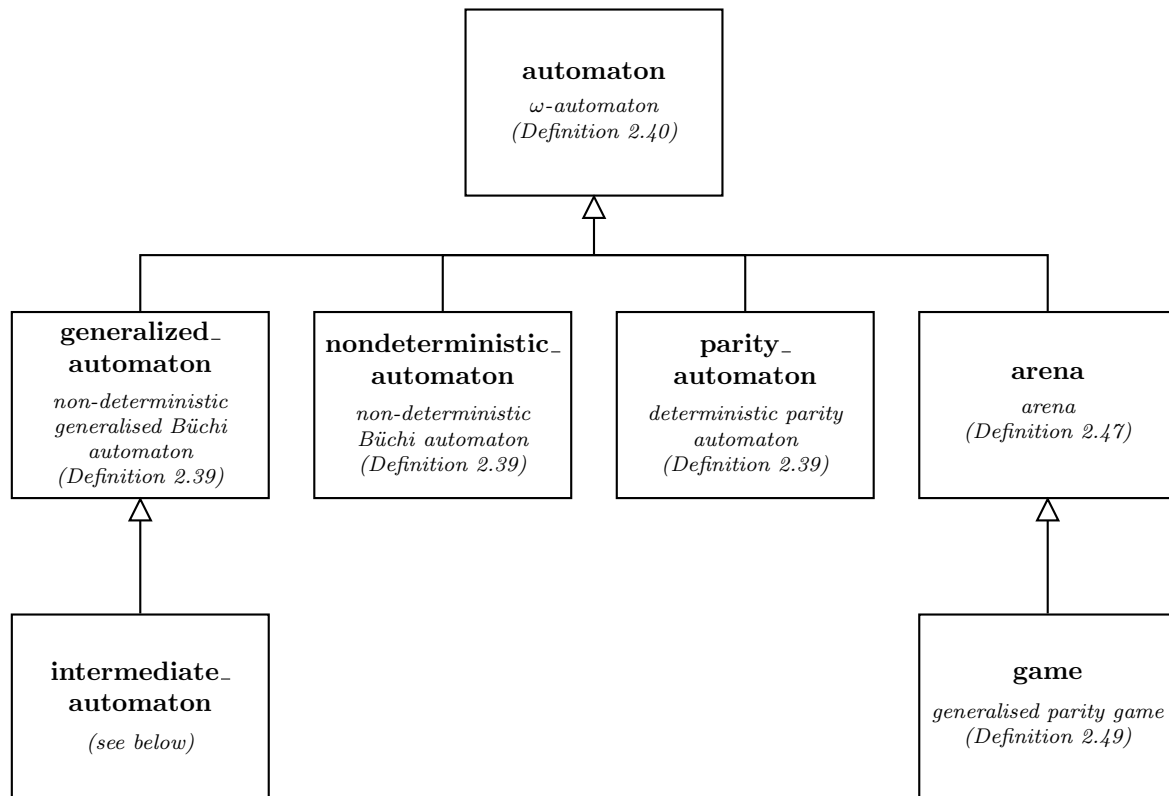


Figure 6.3: UML diagram for classes representing auxiliary data structures used by the SL[1G] model checking algorithm. The `intermediate_automaton` class represents a more general form of a non-deterministic generalised Büchi automaton, which we use for intermediate results in the standard translation (see Subsection 2.4.3).

6.4 Experimental Results

We present here the experimental results obtained using the SLK and SL[1G] extensions discussed in Sections 6.2 and 6.3 respectively on several scalable real-life scenarios. We measure the amount of time and memory used by both extensions and compare it with the performance of the original tool on CTL and ATL. We also demonstrate the differences in expressiveness and strategy synthesis between the two fragments.

The following scalable real-life scenarios are discussed in this section:

1. **Dining Cryptographers.** A simple security protocol which ensures anonymity of participants. We use this example to demonstrate the ability of SLK to express agents' knowledge and compare the performance of the tool on equivalent CTLK, ATLK, and SLK formulas. We also use it to evaluate the impact of the experimental parallel implementation of the SLK model checking algorithm, which we discussed at the end of Subsection 6.2.3.
2. **Cake cutting.** A problem where a cake of a certain size needs to be divided fairly among a group of agents. It demonstrates not only the ability of SLK to express Nash equilibria, but also the fact that our SLK extension can synthesise protocols which achieve them.
3. **Scheduler.** A process scheduler which ensures mutual exclusion of a shared resource while preventing starvation. We use this example to compare the performance of CTL, ATL, SLK, and SL[1G] model checking. We also evaluate the impact of the separate determinisation optimisation technique for the SL[1G] model checking algorithm proposed in Subsection 5.2.5.
4. **Nim.** A traditional game in which players take turns to remove objects from heaps until all heaps are empty. We use this example to compare the performance of ATL, SLK, and SL[1G] model checking on a fixed-size specification.

The experiments were run on an Intel® Core™ i7-3770 CPU 3.40GHz lab machine with 16GB RAM running Linux kernel version 3.8.0-35-generic. ISPL files and results for all scenarios can be found in the `examples/benchmarks` folder in the attached project archive. In addition, the source code of automatic ISPL code generators for some of the scenarios can be found in the `examples/generators` folder.

6.4.1 Dining Cryptographers

The dining cryptographers protocol is a hypothetical protocol introduced in [29] where the anonymity of participants is preserved due to lack of knowledge:

“Three cryptographers are sitting down to dinner at their favourite three-star restaurant. Their waiter informs them that arrangements have been made with the maître d’hôtel for the bill to be paid anonymously. One of the cryptographers might be paying for the dinner, or it might have been NSA (U.S. National Security Agency). The three cryptographers respect each other’s right to make an anonymous payment, but they wonder if NSA is paying. They resolve their uncertainty fairly by carrying out the following protocol:

Each cryptographer flips an unbiased coin behind his menu, between him and the cryptographer on his right, so that only the two of them can see the outcome. Each cryptographer then states aloud whether the two coins he can see—the one he flipped and the one his left-hand neighbour flipped—fell on the same side or on different sides. If one of the cryptographers is the payer, he states the opposite of what he sees. An odd number of differences uttered at the table indicates that a cryptographer is paying; an even number indicates that NSA is paying (assuming that the dinner was paid for only once). Yet if a cryptographer is paying, neither of the other two learns anything from the utterances about which cryptographer it is.”

The protocol described above works for an arbitrary number of cryptographers $n \geq 3$. We model it as an interpreted system (see Definition 2.5) with agents $\text{Agt} \triangleq \{E, c_1, \dots, c_n\}$. We model each cryptographer as an agent whose actions represent the possible announcements (“same” and “different”). The coins are internal variables of the Environment which we make selectively observable to the agents via their `Lobsvars` fields (see Subsection 6.1.2):

```

1 Agent DinCrypt1
2   Lobsvars = {coin1, coin2};
3   ...
4 end Agent

```

The required property of the protocol—if a cryptographer did not pay and an odd number of differences is uttered at the table, then (i) he knows that another cryptographer paid for the dinner but (ii) he does not know which one it was—is expressed in CTLK, ATLK, and SLK as:

$$\varphi_{\text{CTLK}} \triangleq \text{AG } \psi \quad \varphi_{\text{ATLK}} \triangleq \langle\langle \emptyset \rangle\rangle \text{G } \psi \quad \varphi_{\text{SLK}} \triangleq \llbracket x_e \rrbracket \llbracket x_1 \rrbracket \cdots \llbracket x_n \rrbracket (\text{E}, x_e)(c_1, x_1) \cdots (c_n, x_n) \text{G } \psi \quad (6.1)$$

where⁵:

$$\psi \triangleq (\text{odd} \wedge \neg \text{paid}_1) \rightarrow \underbrace{\left[\text{K}_{c_1} \bigvee_{i=2}^n \text{paid}_i \right]}_{\text{cryptographer } c_1 \text{ knows that another cryptographer paid}} \wedge \underbrace{\left[\bigwedge_{i=2}^n \neg \text{K}_{c_1} \text{paid}_i \right]}_{\text{cryptographer } c_1 \text{ does not know which cryptographer paid}}$$

The model checking results for the formulas above with $3 \leq n \leq 18$ cryptographers in Table 6.3 and Figure 6.4 indicate that reasonably large state spaces can be verified by the SLK extension. We can observe that it has similar running time and peak memory usage to CTLK and ATLK for $n < 10$ cryptographers. However, its performance drops considerably faster for $n \geq 10$ cryptographers because of the large number of shared local states that need to be enumerated (see the discussion at the end of Subsection 6.2.3). Consequently, the SLK model checking procedure runs out of physical memory (16GB) for $n > 18$ cryptographers. In contrast, CTLK and ATLK require no assignments and their performance⁶ is dominated by the calculation of the reachable state space.

Table 6.3 and Figure 6.4 also show the impact of the *experimental parallel implementation* of the SLK extension discussed at the end of Subsection 6.2.3. More precisely, we evaluated the performance of the SLK algorithm with 1, 2, and 4 threads (thread branching 0, 1, and 2 set using the `-t` command-line flag). We can observe that the 2-thread version is approximately 16% faster than the sequential one for $n \geq 10$ cryptographers⁷. Although the 4-thread version is slightly faster than the sequential one for $n \geq 11$ cryptographers as well, it does not outperform the 2-thread version for $n < 18$ cryptographers. This result confirms our prior belief that the optimisation approach is not very scalable (see Section 6.2.3). The execution of the 4-thread version on the problem with $n = 18$ cryptographers was prematurely terminated⁸, possibly because it used too much memory (3.8GB) and/or CPU (166%).

Note that the protocol specification cannot be expressed in SL[1G] because it does not support epistemic modalities expressing agents' knowledge due to complete information semantics (see Section 5.1).

6.4.2 Cake Cutting

The cake-cutting problem is a well-known mathematical puzzle [41]:

“Two persons own a cake which they want to split into two parts to be allotted between them. The cake may be made of different ingredients with different values to each person, i.e. each person has his own measure for evaluating any given part of the cake. Is there any procedure or protocol which will enable the two persons to cut the cake into two pieces such that each person will get at least $\frac{1}{2}$ of the cake by his own measure?”

We consider a generalisation of the problem [26, 27] where $n \in \mathbb{N}$ agents want to split a cake of size $d \in \mathbb{N}$. The problem is modelled as a multi-player game with n agents and an environment, which has to provide a protocol to slice the cake in a fair way. The game has *hybrid* turns: at even rounds the n

⁵By symmetry of the problem, we can consider the first cryptographer without loss of generality.


⁶Surprisingly, the ATLK model checking algorithm is always faster than the one for CTLK (despite ATLK being strictly more expressive than CTLK). However, this difference is negligible compared with the running time of the reachable state space calculation.

⁷Incidentally, this is the same threshold from which SLK performance becomes considerably lower than that of CTLK and ATLK.


⁸We ran the experiment (with 4 threads and $n = 18$ cryptographers) repeatedly on different lab machines and it was killed each time by the operating system.

crypts n	possible states	reachable states	reachability time (s)	CTLK		ATLK		SLK					
				time (s)	mem (MB)	time (s)	mem (MB)	1 thread		2 threads		4 threads	
								time (s)	mem (MB)	time (s)	mem (MB)	time (s)	mem (MB)
3	82944	128	0.00	0.00	2.69	0.00	2.69	0.00	2.79	0.01	4.91	0.02	9.23
4	1.99×10^6	320	0.02	0.00	2.71	0.00	2.70	0.00	2.91	0.02	4.97	0.04	9.24
5	4.78×10^7	768	0.04	0.00	2.75	0.00	2.75	0.01	3.07	0.03	5.10	0.05	9.67
6	1.15×10^9	1792	0.08	0.00	2.79	0.00	2.77	0.02	3.41	0.04	5.37	0.07	9.94
7	2.75×10^{10}	4096	0.21	0.00	3.03	0.00	2.97	0.05	4.47	0.08	5.88	0.12	10.71
8	6.60×10^{11}	9216	0.40	0.00	3.19	0.00	3.27	0.18	6.37	0.21	7.42	0.24	12.56
9	1.58×10^{13}	20480	0.43	0.00	3.38	0.00	3.42	0.38	9.59	0.38	9.82	0.46	15.96
10	3.80×10^{14}	45056	4.19	0.28	13.02	0.16	12.75	2.29	18.62	2.23	22.33	2.43	28.21
11	9.13×10^{15}	98304	1.69	0.04	5.49	0.01	4.75	5.01	23.55	4.37	27.51	4.77	39.59
12	2.19×10^{17}	212992	2.32	0.01	4.40	0.01	4.20	11.61	39.25	9.61	40.36	10.32	60.82
13	5.26×10^{18}	458752	2.05	0.10	6.50	0.03	5.62	32.63	81.95	26.30	79.75	28.39	118.27
14	1.26×10^{20}	983040	1.96	0.08	6.87	0.03	5.00	89.46	169.95	67.99	169.68	73.57	243.80
15	3.03×10^{21}	2.10×10^6	21.35	0.35	12.73	0.08	9.29	163.77	363.66	139.42	354.18	156.57	500.96
16	7.27×10^{22}	4.46×10^6	6.66	0.09	6.65	0.04	4.83	422.03	758.43	345.83	736.02	377.50	977.79
17	1.74×10^{24}	9.44×10^6	9.10	0.13	6.67	0.08	6.04	734.44	1360.59	643.03	1270.93	705.44	1359.73
18	4.19×10^{25}	1.99×10^7	65.94	0.50	12.66	0.14	12.67	2654.46	3550.66	2129.54	3680.89	<i>process killed</i>	

Table 6.3: Experimental model checking results for the dining cryptographers protocol.



(a) Running time.



(b) Peak memory usage.

Figure 6.4: Experimental model checking results for the dining cryptographers protocol.

agents concurrently choose how they want to split the cake, while at odd rounds the environment decides how to actually slice the cake and assigns the obtained pieces to a subset of the agents. Therefore, the problem of splitting a cake of size d among n agents is divided into simpler problems in which cakes of size $d' < d$ have to be split between $n' < n$ agents. The game terminates once each agent receives a slice.

We model the game as an interpreted system (see Definition 2.5) with agents $Ag_t \triangleq \{E, 1, \dots, n\}$ where [26]:

- **Agents.** The set of *internal states* of an agent $i \in [1..n]$ is $L_i \triangleq [1..n] \times [1..d] \times \{P, E\}$, where the first component represents the number of agents $n' \leq n$ among which the current cake has to be split, the second component denotes the current size of the cake $d' \leq d$, and the flags P and E indicate whether it is the agents' or the environment's turn.

The set of *actions* available to an agent $i \in [1..n]$ is $Act_i \triangleq \{\perp\} \cup [1..d]$ where \perp indicates no choice (the agent already got his piece of cake or it is the environment's turn) and a numeric value represents the size of the slice he wants. The agent's *protocol* is defined as $P_i(\langle(m, k, E), l_E\rangle) \triangleq P_i(\langle(1, k, P), l_E\rangle) \triangleq \{\perp\}$ and $P_i(\langle(m, k, P), l_E\rangle) \triangleq [1..k]$ for $m > 1$.

- **Environment.** The internal states of the environment represent the possible divisions of the cake among subsets of agents. To represent them, we need to introduce some auxiliary concepts. Firstly, we define the set of functions $F \triangleq \{f : [1..n] \rightarrow [1..n] \mid \exists h \in [1..d]. \text{img}(f) = [1..h]\}$ which assign agents to pieces of cake. Secondly, the sizes of the pieces are encoded by means of partial functions $S \triangleq [1..n] \rightarrow [1..d]$. The possible divisions of the cake are then pairs $(f, s) \in H \triangleq \left\{ (f, s) \in F \times S \mid \text{dom}(s) = \text{img}(f) \wedge \sum_{i \in \text{dom}(s)} s(i) \leq d \right\}$ where f maintains the information about the splitting whose size is contained in s .

In order to define the protocol of the environment, we need to define a partial order on H which allows us to identify all possible continuations from a given state. For all $(f, s), (f', s') \in H$, we write that $(f', s') \preceq (f, s)$ iff the following two conditions hold:

1. if $f'(i) = f'(j)$, then $f(i) = f(j)$, for all $i, j \in [1..n]$;
2. $\sum_{j \in [1..n], f(j)=f(i)} s'(f'(j)) \leq s(f(i))$, for all $i \in [1..n]$.

Intuitively, $(f', s') \preceq (f, s)$ if f' is a refinement of f , i.e. f' represents a finer partition than f , and s' is coherent with the sizes of the pieces s before the last cuts were made.

At this point, we can define the *actions* of the environment $Act_E \triangleq H$ and its *local states* $L_E \triangleq Act = H \times \prod_{i=1}^n Act_i$. Intuitively, the Cartesian product stores the previous actions of agents $i \in [1..n]$ (so that the environment could base its decision on their actions). Finally, the protocol of the environment is defined as⁹:

$$P_E(\langle a_E, a_1, \dots, a_n \rangle) \triangleq \begin{cases} \{a_E\} & \text{if } a_i = \perp \text{ for all } i \in [1..n] \\ \{(f, s) \in H \mid (f, s) \preceq a_E\} & \text{otherwise} \end{cases}$$

Intuitively, when it is the turn of the players, the environment has to preserve the previous splitting. Otherwise, it can use all possible refinements on its previous move a_E .

- **Evolution.** Since the states of the environment maintain both the information on the splitting of the cake (environment's previous action) and the choices made by the agents in the previous round, the *environment evolution function* is simply $t_E(a, a') \triangleq a'$.

There are two cases for the *agents' evolution functions*:

1. States of the form (m, k, P) (agents' turn) merely allow a subset of $m \leq n$ agents to make a proposal to the environment for the cut of their piece of cake of size $k \leq d$. Therefore, we have $t_i(\langle(m, k, P), a\rangle, a') \triangleq (m, k, E)$ for all $i \in [1..n]$.

⁹The protocol might seem reversed because the environment makes no move when $a_i = \perp$ for all $i \in [1..n]$. However, it is important to realise that a_1, \dots, a_n refer to the *previous* actions of the agents. Hence, if the agents did not make any move in the previous round, it is now their turn (unless the game is over, in which case the environment should not do anything either).

2. States of the form (m, k, E) (environments' turn) are used to adopt the decision of the environment $a_E(a) = (f, s)$. To do this, we set $t_i(\langle(m, k, E), a\rangle, a') \triangleq (m', k', P)$ where $m' \triangleq |\{j \in [1..n] \mid f(j) = f(i)\}| \leq m$ and $k' \triangleq s(f(i))$ for all $i \in [1..n]$.

• **Initial states.** The set of initial states $I \triangleq \{g\}$ contains a unique global state $g \in G$ defined as:

1. $l_i(g) = (n, d, P)$ for all $i \in [1..n]$;
2. $l_E(g) = ((f, s), \perp, \dots, \perp)$ where $f(i) = 1$ for all $i \in [1..n]$ and $s(1) = d$.

• **Assignment.** We fix a set of *atomic propositions* $AP \triangleq [1..n] \times [1..d]$, where each atomic proposition $\langle i, c \rangle \in AP$ indicates that agent i got piece of cake of size c . The *assignment* $h : AP \rightarrow 2^G$ is formally defined as $h(\langle i, c \rangle) \triangleq \{g \in G \mid l_i(g) = (1, c, P)\}$ for all $\langle i, c \rangle \in AP$.

The existence of a protocol for the environment is then given by the following SLK specification:

$$\varphi \triangleq \langle\langle x \rangle\rangle (\varphi_F \wedge \varphi_S) \quad (6.2)$$

where:

- $\varphi_F \triangleq \llbracket y_1 \rrbracket \cdots \llbracket y_n \rrbracket (\psi_{NE} \rightarrow \psi_E)$ ensures that the protocol x is *fair*, i.e. all possible Nash equilibria (y_1, \dots, y_n) of the agents guarantee equality of splitting;
- $\varphi_S \triangleq \langle\langle y_1 \rangle\rangle \cdots \langle\langle y_n \rangle\rangle \psi_{NE}$ ensures that the protocol has a *solution*, i.e. there is at least one Nash equilibrium;
- $\psi_{NE} \triangleq \bigwedge_{i=1}^n \left(\bigwedge_{v=1}^d (\langle\langle z \rangle\rangle b_i p_i(v)) \rightarrow \left(\bigvee_{c=v}^d b p_i(c) \right) \right)$ ensures that if agent i has a strategy z that allows him to obtain a piece of cake of size v once the strategies of the other agents are fixed, he is already able to obtain a slice of size $c \geq v$ by means of his original strategy y_i , i.e. the original strategies y_1, \dots, y_n form a *Nash equilibrium*;
- $\psi_E \triangleq b \bigwedge_{i=1}^n p_i(\lfloor d/n \rfloor)$ ensures that agent i is able to obtain a piece of size $\lfloor d/n \rfloor$;
- $b \triangleq (E, x)(1, y_1) \cdots (n, y_n)$, $b_i \triangleq (E, x)(1, y_1) \cdots (i, z) \cdots (n, y_n)$, and $p_i(c) \triangleq F \langle i, c \rangle$ are auxiliary abbreviations.

We were able to verify the formula φ defined above on a system with $n = 2$ agents and a cake of size $d = 2$. Moreover, the SLK extension can automatically synthesise a witness strategy h_x for the environment as well as the Nash equilibrium (h_{y_1}, h_{y_2}) of the agents¹⁰:

```

1  ----- Strategy x [Environment] -----
2  Agent Environment
3  E.a1=0, E.a2=0, E.f1=1, E.f2=1, E.s1=1, E.s2=0, E.turn=player (8): f_1_1_s_1
4  E.a1=0, E.a2=0, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=player (7): f_1_1_s_2
5  E.a1=0, E.a2=0, E.f1=1, E.f2=2, E.s1=1, E.s2=1, E.turn=env (5): f_1_2_s_1_1
6  E.a1=0, E.a2=0, E.f1=1, E.f2=2, E.s1=1, E.s2=1, E.turn=player (9): f_1_2_s_1_1
7  E.a1=0, E.a2=0, E.f1=2, E.f2=1, E.s1=1, E.s2=1, E.turn=env (6): f_1_2_s_1_1
8  E.a1=0, E.a2=0, E.f1=2, E.f2=1, E.s1=1, E.s2=1, E.turn=player (10): f_2_1_s_1_1
9  E.a1=1, E.a2=1, E.f1=1, E.f2=1, E.s1=1, E.s2=0, E.turn=env (4): f_1_1_s_1
10 E.a1=1, E.a2=1, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env (3): f_1_1_s_2
11 E.a1=1, E.a2=2, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env (2): f_1_1_s_1
12 E.a1=2, E.a2=1, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env (1): f_2_1_s_1_1
13 E.a1=2, E.a2=2, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env (0): f_1_2_s_1_1
14  -----
15  ----- Strategy y1 [P1] -----
16  Agent P1
17  E.a1=0, E.a2=0, E.f1=1, E.f2=1, E.s1=1, E.s2=0, E.turn=player, k=1, m=2 (8): a1
18  E.a1=0, E.a2=0, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=player, k=2, m=2 (7): a2
19  E.a1=0, E.a2=0, E.f1=1, E.f2=2, E.s1=1, E.s2=1, E.turn=env, k=1, m=1 (5): a0

```

¹⁰Environment is abbreviated to E. A witness execution was also generated by the extension but we omit it for conciseness.

```

20 E.a1=0, E.a2=0, E.f1=1, E.f2=2, E.s1=1, E.s2=1, E.turn=player, k=1, m=1 (9): a0
21 E.a1=0, E.a2=0, E.f1=2, E.f2=1, E.s1=1, E.s2=1, E.turn=env, k=1, m=1 (6): a0
22 E.a1=0, E.a2=0, E.f1=2, E.f2=1, E.s1=1, E.s2=1, E.turn=player, k=1, m=1 (10): a0
23 E.a1=1, E.a2=1, E.f1=1, E.f2=1, E.s1=1, E.s2=0, E.turn=env, k=1, m=2 (4): a0
24 E.a1=1, E.a2=1, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (3): a0
25 E.a1=1, E.a2=2, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (2): a0
26 E.a1=2, E.a2=1, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (1): a0
27 E.a1=2, E.a2=2, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (0): a0
28 -----
29 ----- Strategy y2 [P2] -----
30 Agent P2
31 E.a1=0, E.a2=0, E.f1=1, E.f2=1, E.s1=1, E.s2=0, E.turn=player, k=1, m=2 (8): a1
32 E.a1=0, E.a2=0, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=player, k=2, m=2 (7): a1
33 E.a1=0, E.a2=0, E.f1=1, E.f2=2, E.s1=1, E.s2=1, E.turn=env, k=1, m=1 (5): a0
34 E.a1=0, E.a2=0, E.f1=1, E.f2=2, E.s1=1, E.s2=1, E.turn=player, k=1, m=1 (9): a0
35 E.a1=0, E.a2=0, E.f1=2, E.f2=1, E.s1=1, E.s2=1, E.turn=env, k=1, m=1 (6): a0
36 E.a1=0, E.a2=0, E.f1=2, E.f2=1, E.s1=1, E.s2=1, E.turn=player, k=1, m=1 (10): a0
37 E.a1=1, E.a2=1, E.f1=1, E.f2=1, E.s1=1, E.s2=0, E.turn=env, k=1, m=2 (4): a0
38 E.a1=1, E.a2=1, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (3): a0
39 E.a1=1, E.a2=2, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (2): a0
40 E.a1=2, E.a2=1, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (1): a0
41 E.a1=2, E.a2=2, E.f1=1, E.f2=1, E.s1=2, E.s2=0, E.turn=env, k=2, m=2 (0): a0
42 -----

```

For example, line 13 represents the following constraint on the strategy h_x for the environment:

$$h_x(\underbrace{((\{(1,1), (2,1)\})}_{E.f1=1 \ E.f2=1}, \underbrace{\{(1,2)\}}_{E.s1=2}, \underbrace{2}_{E.a1=2}, \underbrace{2}_{E.a2=2})) = \underbrace{((\{(1,1), (2,2)\})}_{f'}, \underbrace{\{(1,1), (2,1)\}}_{s'})_{f_{1_2_s_{1_1}}}$$

where $2 \notin \text{dom}(s)$ is encoded by $E.s2=0$. Observe that $(f', s') \preceq (f, s)$ holds as required.

Unfortunately, we were not able to verify larger examples because the lab machine ran out of memory. For example, consider the scenario with $n = 2$ agents and $d = 3$ pieces of cake, which has 29 reachable states. The encoding of extended states, which contain strategies mapping each reachable state to an action, for the formula φ requires 105 Boolean variables, 54 of which represent the strategy h_x we wish to synthesise. The Boolean variables in turn generate a state space with 8.09×10^{20} possible *extended* states¹¹. We found that the intermediate BDDs have order of 10^9 nodes, each of which has 32 bytes. A single intermediate BDD thus requires approximately 32GB of memory and the whole algorithm would use *several hundreds of gigabytes*.

This negative result should not be surprising given the theoretical difficulty of the cake-cutting problem. Moreover, we are synthesising the entire protocol and not just the agents' optimal behaviour. In our opinion, this is still a great achievement because our tool can automatically *synthesise Nash equilibria* of arbitrary memoryless games (given enough memory and time).

Note that the specification in Equation 6.2 again cannot be expressed in SL[1G]. Unlike the dining cryptographers specification in Equation 6.1, which contains epistemic modalities unsupported by SL[1G], this one cannot be translated to an equivalent SL[1G] formula because of the restrictions on syntax, namely that every quantification prefix φ must be coupled with a goal $b\varphi$ where b binds all agents to variables in φ (see Definition 5.2).

¹¹We were able to handle even larger state spaces (up to 4.19×10^{25} possible *global* states) in the dining cryptographers scenario (see Table 6.3). There are three important differences between the problems which we believe explain this discrepancy: (i) Equation 6.1 contains only one temporal operator (unlike Equation 6.2); (ii) the global state changes only during the first temporal transition in the dining cryptographers model, whereas the cake-cutting game might take several rounds; and (iii) unlike the environment protocol in the cake cutting problem, which we aim to synthesise, the dining cryptographers protocol is deterministic (once the coins are flipped).

6.4.3 Scheduler

Most operating systems nowadays support the concept of concurrency where several processes (or threads) run at the same time¹². While parallelism can greatly improve the performance of many applications and simplify software design, it can also lead to race conditions where multiple processes use a shared resource simultaneously and its final state depends on the relative timing, which is non-deterministic and thus difficult to debug. Correctness of concurrent systems is crucial because they often run indefinitely and provide services that other systems depend on (e.g. web servers). Since humans generally tend to have difficulties reasoning about them and designing them, automatic verification and synthesis of such systems is of vital importance.

We consider a simple *preemptive scheduler system* [26] composed of $n \in \mathbb{N}$ processes and an arbiter (represented by the environment). The processes try to access a shared resource which is mutually exclusive. When more than one process requests the resource, the arbiter has to decide who will obtain it. We are interested in the following two properties:

1. **Mutual exclusion.** This property (also referred to as a *safety* property) asserts that at most one process owns the resource at any given point in time.
2. **Absence of starvation.** This property (also referred to as a *fairness* property) requires that every request is eventually satisfied, i.e. no process will wait for the resource forever.

Again, we model the scheduler system as an interpreted system (see Definition 2.5) with agents $\text{Agt} \triangleq \{\text{E}, 1, \dots, n\}$ where [26]:

- **Processes.** Every process $i \in [1..n]$ has three possible *internal states* $L_i \triangleq \{\text{in}, \text{wt}, \text{rs}\}$ which stand for “internal computation”, “waiting for resource”, and “owns resource” respectively. Furthermore, the *actions* of the process are $\text{Act}_i \triangleq \{\text{id}, \text{rq}, \text{rn}, \text{rl}\}$ meaning “idle”, “request resource”, “relinquish resource”, and “release resource” respectively. The third action (“relinquish resource”) is used by the process to stop waiting for the resource. Finally, the *protocol* of the process is defined as $P_i(\text{in}) \triangleq \{\text{id}, \text{rq}\}$, $P_i(\text{wt}) \triangleq \{\text{id}, \text{rn}\}$, and $P_i(\text{rs}) \triangleq \{\text{id}, \text{rs}\}$.
- **Arbiter.** A *state* of the environment either indicates that the resource is free, or represents the set of processes that require the resource but do not own it $L_E \triangleq \{\text{fr}\} \cup 2^{[1..n]}$. The *actions* of the environment are $\text{Act}_E \triangleq \{\text{id}\} \cup [1..n]$ where *id* indicates that the environment does not do any visible action (no arbitrage is required) and a numeric value $i \in [1..n]$ represents the fact that the arbiter has granted access to the shared resource to process i . The *protocol* of the environment is defined as $P_E(\text{fr}) \triangleq P_E(\emptyset) \triangleq \{\text{id}\}$ and $P_E(A) \triangleq A$ for all $A \subseteq [1..n]$ with $A \neq \emptyset$.
- **Evolution.** Before formally defining the evolution functions, we introduce some auxiliary notation. Let $a \in \text{Act}$ be a joint action. The sets of agents that request, relinquish, and release the resource are defined as $\text{Rq}(a) \triangleq \{i \in [1..n] \mid a_i(a) = \text{rq}\}$, $\text{Rn}(a) \triangleq \{i \in [1..n] \mid a_i(a) = \text{rn}\}$, and $\text{Rl}(a) \triangleq \{i \in [1..n] \mid a_i(a) = \text{rl}\}$ respectively. Note that $|\text{Rl}(a)| \leq 1$ for all possible joint actions $a \in \text{Act}$.

There are three cases to consider for the *arbiter's evolution function*:

1. Starting from the state *fr*, if no process makes a request, the state does not change. Otherwise, if there is exactly one process requesting the resource, the environment transits to \emptyset . In the case of concurrent requests, their set $\text{Rq}(a)$ is recorded instead:

$$t_E(\text{fr}, a) \triangleq \begin{cases} \text{fr} & \text{if } |\text{Rq}(a)| = 0 \\ \emptyset & \text{if } |\text{Rq}(a)| = 1 \\ \text{Rq}(a) & \text{otherwise} \end{cases}$$

2. Starting from the state \emptyset , if the process that owns the resource releases it and no other process requests it, the environment transits to *fr*. If the owner does not release the resource and no

¹²Note that we do not care whether the system uses real or apparent concurrency. We are interested in its high-level behaviour instead.

process makes a request or if the owner releases it and exactly one process makes a request, the state does not change. In all other cases, the set of concurrent requests $Rq(a)$ is recorded:

$$t_E(\emptyset, a) \triangleq \begin{cases} \text{fr} & \text{if } |Rl(a)| = 1 \text{ and } |Rq(a)| = 0 \\ \emptyset & \text{if } |Rl(a)| = |Rq(a)| \\ Rq(a) & \text{otherwise} \end{cases}$$

3. Finally, starting from a set $A \neq \emptyset$ of processes that are waiting for the resource, the environment transits to the new set $A \setminus (\{a_E(a)\} \cup Rn(a)) \cup Rq(a)$ maintaining the set of processes that do not relinquish the request and are not chosen by the arbiter, augmented with the new ones that want to access the shared resource:

$$t_E(A, a) \triangleq A \setminus (\{a_E(a)\} \cup Rn(a)) \cup Rq(a)$$

The *evolution function of a process* $i \in [1..n]$ is defined analogously:

$$\begin{aligned} t_i(\langle \text{in}, \text{fr} \rangle, a) &\triangleq \begin{cases} \text{in} & \text{if } a_i(a) = \text{id} \\ \text{rs} & \text{if } Rq(a) = \{i\} \\ \text{wt} & \text{otherwise} \end{cases} \\ t_i(\langle \text{in}, \emptyset \rangle, a) &\triangleq \begin{cases} \text{in} & \text{if } a_i(a) = \text{id} \\ \text{rs} & \text{if } |Rl(a)| = 1 \text{ and } Rq(a) = \{i\} \\ \text{wt} & \text{otherwise} \end{cases} \\ t_i(\langle \text{in}, A \rangle, a) &\triangleq \begin{cases} \text{in} & \text{if } a_i(a) = \text{id} \\ \text{wt} & \text{otherwise} \end{cases} \\ t_i(\langle \text{rs}, \emptyset \rangle, a) &\triangleq \begin{cases} \text{in} & \text{if } a_i(a) = \text{rl} \\ \text{rs} & \text{otherwise} \end{cases} \\ t_i(\langle \text{rs}, A \rangle, a) &\triangleq \begin{cases} \text{in} & \text{if } a_i(a) = \text{rl} \text{ or } a_E(a) \neq i \\ \text{rs} & \text{otherwise} \end{cases} \\ t_i(\langle \text{wt}, A \rangle, a) &\triangleq \begin{cases} \text{in} & \text{if } a_i(a) = \text{rn} \\ \text{rs} & \text{if } a_i(a) = \text{id} \text{ and } a_E(a) = i \\ \text{wt} & \text{otherwise} \end{cases} \end{aligned}$$

- **Initial states.** The set of initial states $I \triangleq \{g\}$ contains a unique global state $g \in G$ such that $l_E(g) = \text{fr}$ and $l_i(g) = \text{in}$ for all $i \in [1..n]$.
- **Assignment.** We fix a set of *atomic propositions* $AP \triangleq \{\text{rs}, \text{wt}\} \times [1..n]$, where $\langle \text{wt}, i \rangle$ and $\langle \text{rs}, i \rangle$ indicate that process $i \in [1..n]$ is waiting for and owns the resource respectively. The *assignment* $h : AP \rightarrow 2^G$ is formally defined as $h(\langle s, i \rangle) \triangleq \{g \in G \mid l_i(g) = s\}$ for all $i \in [1..n]$ and $s \in \{\text{wt}, \text{rs}\}$.

The desired properties of the system can be expressed using the following $\text{SL}[1G]$ ¹³ specifications:

1. **Mutual exclusion.** The following specification asserts that *all strategies* for the scheduler already ensure mutual exclusion, i.e. we want to *verify* the property:

$$\varphi_{\text{ME}} \triangleq \llbracket x \rrbracket \llbracket y_1 \rrbracket \cdots \llbracket y_n \rrbracket (E, x)(1, y_1) \cdots (n, y_n) \mathbf{G} \neg \bigvee_{i=1}^n \bigvee_{j=i+1}^n \langle \text{rs}, i \rangle \wedge \langle \text{rs}, j \rangle$$

This property can be equivalently expressed in CTL and ATL using the AG and $\langle\langle \emptyset \rangle\rangle \mathbf{G}$ operators respectively (see Equation 6.1).

¹³Recall that $\text{SL}[1G]$ syntax (Definition 5.2) is a subset of SLK syntax (Definition 4.1). Thus, $\text{SL}[1G]$ specifications are also SLK specifications.

The model checking results for the formula with $2 \leq n \leq 10$ processes are shown in Table 6.4a. We can see that the CTL and ATL formulas are checked almost instantaneously. Similarly to the dining cryptographers scenario (see Subsection 6.4.1), the SLK model checking performance on φ_{ME} drops quickly for $n \geq 8$ processes. However, the SL[1G] model checking performance on φ_{ME} drops even faster and the procedure already takes more than 2 hours and 1.4GB memory on the model with $n = 7$ processes. Moreover, it runs out of physical memory (16GB) for $n > 7$ processes.

The individual steps of the SL[1G] model checking algorithm (see Figure 5.3) which take more than one second, namely arena construction, parity game combination, and parity game solution, are also shown in Table 6.4a. We can see that the SL[1G] model checking performance is dominated by parity game solving. For example, calculating the parity game solution for φ_{ME} with $n = 7$ processes takes approximately 1 hour and 50 minutes, which is roughly 77% of the whole model checking time (2 hours and 23 minutes). This is in line with our expectation that parity game solving is one of the bottlenecks of our SL[1G] model checking algorithm.

2. **Absence of starvation.** The following specifications asserts the *existence* of a strategy for the scheduler which ensures absence of starvation, i.e. we want to *synthesise*¹⁴ the strategy for:

$$\varphi_{\text{AS}} \triangleq \langle\langle x \rangle\rangle [y_1] \dots [y_n] (\text{E}, x) (1, y_1) \cdots (n, y_n) \bigwedge_{i=1}^n \text{G} (\langle \text{wt}, i \rangle \rightarrow \text{F} \neg \langle \text{wt}, i \rangle)$$

This property cannot be expressed in CTL because it refers to the capability of an agent to enforce a property. Moreover, it cannot be expressed in ATL either because it contains nested temporal operators (G and F) which depend on the same strategies. However, it can be translated to an ATL* formula by replacing the whole quantification and binding prefix with the quantifier $\langle\langle \{E\} \rangle\rangle$.

We can use the equivalence $\bigwedge_{i=1}^n \text{G} \psi_i \equiv \text{G} \bigwedge_{i=1}^n \psi_i$ to reduce the number of temporal operators in φ_{AS} in order to improve SLK and SL[1G] model checking performance. Since the underlying LTL formula $\bigwedge_{i=1}^n \text{G} \psi_i$ is a conjunction, we can also apply the *separate determinisation* optimisation technique for SL[1G] discussed in Subsection 5.2.5. Therefore, we empirically evaluate the following three approaches: (i) the reduced formula with the SLK algorithm, (ii) the reduced formula with the unoptimised SL[1G] algorithm, and (iii) the original formula with the optimised SL[1G] algorithm.

The model checking results for the formulas with $2 \leq n \leq 3$ processes are shown in Table 6.4b. Unlike for φ_{ME} , SL[1G] in this case outperforms SLK by an order of magnitude for $n \geq 2$ processes. Moreover, the SLK extension runs out of physical memory for $n > 3$ processes. We believe that the performance reversal reflects the different structures of φ_{ME} (safety condition) and φ_{AS} (fairness condition). A further investigation needs to be done in order to identify classes of formulas which can be checked efficiently by one of the two algorithms. The optimised SL[1G] model checking algorithm is almost twice faster than the unoptimised one for $3 \leq n \leq 4$ processes and almost 5 times faster for $n = 5$ processes, i.e. the relative speedup increases with the number of processes. Moreover, the optimised algorithm can handle a system with $n = 6$ processes¹⁵, whereas the unoptimised one runs out of memory. This is in line with our expectations (see Subsection 5.2.5).

Again, the individual steps of the SL[1G] model checking algorithms (see Figures 5.3 and 5.4) which take more than one second, namely automaton determinisation, arena construction, parity game combination, and parity game solution, are shown in Table 6.4b. Similarly to φ_{ME} , SL[1G] model checking performance is dominated by parity game solving. Unlike the optimised determinisation procedure, which is almost instantaneous, the unoptimised one is even slower than arena construction for $n \geq 4$ processes. This is due to the different sizes of the non-deterministic Büchi automata (see Subsection 5.2.5): The unoptimised algorithm determinises 1 automaton with

¹⁴Some of the SL[1G] witness strategies for the formula have more than 10^5 states. Therefore, we do not actually request witness strategy synthesis (`-c 1` flag) because printing the strategies would become a major bottleneck of the algorithm which would skew the performance results. Omitting the flag does not have any impact on the performance of the algorithms themselves, which perform (symbolic) strategy synthesis internally anyway.

¹⁵Although it takes almost 24 hours.

$2^{2(n+1)+\lceil\log_2(n+2)\rceil}$ states¹⁶, whereas the optimised one determinises n automata with 2^6 states¹⁷. This—together with the fact that the worst-case time complexity of the determinisation procedure is exponential with respect to the size of the non-deterministic Büchi automaton—explains why the optimised algorithm outperforms the unoptimised one and why the relative speedup increases with the number of processes.

6.4.4 Nim

Nim is a classical mathematical game for two players [77]:

“Nim is a two-player turn-based game in which players alternately remove objects from heaps. On each turn, one of the two players chooses a heap and a non-zero number of objects that he removes from this heap. The player who removes the last object wins the game.”

A complete theory for the game was developed already in 1901 [77]. Its key concept is the so-called *Nim-sum*, which is calculated as an exclusive or (xor) of the heap sizes. For example, if there are three heaps with sizes 3, 4, and 5 respectively, the Nim-sum is $011 \oplus 100 \oplus 101 = 010$. A game state is winning (for any of the two players) iff its Nim-sum is *not* 0. Consequently, the winning strategy is to finish every move with a Nim-sum of 0. In the previous example, the state is winning for the current player and his winning strategy is to remove 2 objects from the first heap.

We consider a variant of the game where an *upper bound* is imposed on the number of objects that can be removed in a turn¹⁸. We model the game with $n \in \mathbb{N}$ heaps with $k_1, \dots, k_n \in \mathbb{N}$ elements and an upper bound $b > 0$ as an interpreted system (see Definition 2.5) with agents $\text{Agt} \triangleq \{\text{E}, 1, 2\}$:

- **Environment.** The set of *internal states* of the environment is defined as $L_E \triangleq \{t_1, t_2, w_1, w_2\} \times \prod_{j=1}^n [0..k_j]$ where t_1, t_2, w_1 , and w_2 stand for “first player’s turn”, “second player’s turn”, “first player won”, and “second player won” respectively. The environment has only one available *action* $\text{Act}_E \triangleq \{\text{id}\}$ meaning “idle”. The environment *protocol* is simply $P_E(l_E) \triangleq \{\text{id}\}$ for all $l_E \in L_E$.
- **Players.** Both players $i \in \{1, 2\}$ have only one available *internal state* $L_i \triangleq \{\text{in}\}$, i.e. all information is in the internal state of the environment and thus both players have complete information (since their local states are $L_{iE} \triangleq L_i \times L_E$, see Definition 2.5). The set of actions available to each player is $\text{Act}_i \triangleq \{\text{id}\} \cup \{\langle m, l \rangle \in \mathbb{N} \times \mathbb{N} \mid 1 \leq m \leq n \wedge 1 \leq l \leq \min(k_m, b)\}$ where id stands for “idle” and $\langle m, l \rangle$ indicates that the player removes l objects from heap m . The *protocol* of a player $i \in \{1, 2\}$ is defined as $P_i(\text{in}, \langle w_1, 0, \dots, 0 \rangle) \triangleq P_i(\text{in}, \langle w_2, 0, \dots, 0 \rangle) \triangleq P_i(\text{in}, \langle t_1, l_1, \dots, l_n \rangle) \triangleq \{\text{id}\}$ and $P_i(\text{in}, \langle t_i, l_1, \dots, l_n \rangle) \triangleq \{\langle m, l \rangle \in \text{Act}_i \setminus \{\text{id}\} \mid l \leq l_m\}$ for all $(l_1, \dots, l_n) \in \prod_{j=1}^n [0..k_j] \setminus \{(0, \dots, 0)\}$.
- **Evolution.** The *environment evolution* is defined as follows. Firstly, if the a player $i \in \{1, 2\}$ has already won, the state does not change:

$$t_E(\langle w_i, 0, \dots, 0 \rangle, a) \triangleq \langle w_i, 0, \dots, 0 \rangle$$

for all joint actions $a \in \text{Act}$. Secondly, if the next move of a player removes all objects from the last non-empty heap, the player wins:

$$\begin{aligned} t_E(\langle t_1, 0, \dots, l_m, \dots, 0 \rangle, (\text{id}, \langle m, l_m \rangle, \text{id})) &\triangleq \langle w_1, 0, \dots, 0 \rangle \\ t_E(\langle t_2, 0, \dots, l_m, \dots, 0 \rangle, (\text{id}, \text{id}, \langle m, l_m \rangle)) &\triangleq \langle w_2, 0, \dots, 0 \rangle \end{aligned}$$

¹⁶Recall that a non-deterministic Büchi automaton \mathfrak{B} equivalent to a non-deterministic generalised Büchi automaton \mathfrak{A} is obtained as a product $\mathfrak{B} = \mathfrak{A} \times \mathfrak{D}$ with an auxiliary deterministic Büchi automaton \mathfrak{D} with a counter (see the end of Subsection 2.4.1). In this particular case, $2(n+1)$ BDD variables represent the state of the underlying generalised non-deterministic Büchi automaton and $\lceil\log_2(n+2)\rceil$ variables represent the counter.

¹⁷4 BDD variables represent the state of the underlying generalised non-deterministic Büchi automaton and 2 variables represent the counter. Moreover, each of the n automata has only 6 reachable states which are not dead ends (recall that the number of BDD variables needed by the equivalent deterministic parity automaton depends on the number of reachable states which are not dead ends, see Subsection 2.4.4).

¹⁸This variant of Nim is referred to as *subtraction game* at <http://en.wikipedia.org/wiki/Nim>. According to the article, the winning regions of the game can be calculated using the Nim-sum of heap sizes modulo $b+1$ where $b > 0$ is the upper bound.

procs n	possible states	reach. states	reach. time (s)	CTL		ATL		SLK		SL[1G]				
				time (s)	mem (MB)	time (s)	mem (MB)	time (s)	mem (MB)	time (s)				mem (MB)
										total	arena	combine	solve	
2	72	9	0.00	0.00	2.64	0.00	2.64	0.00	2.67	0.01	0.00	0.00	0.01	2.99
3	432	21	0.00	0.00	2.72	0.00	2.72	0.00	2.81	0.14	0.02	0.03	0.09	5.23
4	2592	49	0.01	0.00	2.81	0.00	2.81	0.00	3.11	3.42	0.27	0.76	2.39	15.15
5	15552	113	0.02	0.00	3.02	0.00	3.03	0.02	4.20	65.81	2.82	11.58	51.40	33.37
6	93312	257	0.02	0.00	3.40	0.00	3.40	0.11	8.98	819.45	39.46	154.20	625.67	233.77
7	559872	577	0.08	0.00	4.47	0.00	4.48	0.86	23.18	8572.79	351.56	1597.84	6622.77	1500.69
8	3.36×10^6	1281	0.17	0.00	6.73	0.00	6.72	4.88	91.26	<i>out of memory</i>				
9	2.02×10^7	2817	0.20	0.00	11.62	0.00	11.62	36.71	561.69	<i>out of memory</i>				
10	1.21×10^8	6145	0.10	0.00	22.23	0.00	22.23	222.36	2769.64	<i>out of memory</i>				

(a) Mutual exclusion property φ_{ME} .

procs n	SLK		SL[1G]											
			unoptimised						optimised					
	time (s)	mem (MB)	time (s)					mem (MB)	time (s)					mem (MB)
			total	determ.	arena	combine	solve		total	determ.	arena	combine	solve	
2	0.00	2.67	0.09	0.00	0.00	0.02	0.06	4.44	0.12	0.00	0.00	0.02	0.10	4.52
3	116.78	124.90	10.11	0.01	0.02	1.06	9.02	15.14	6.39	0.00	0.01	0.65	5.72	14.54
4	<i>out of memory</i>		631.11	0.33	0.26	86.47	544.04	41.80	338.16	0.00	0.24	23.33	314.57	40.70
5	<i>out of memory</i>		29593.56	5.86	2.61	4323.81	25261.15	2792.22	6131.43	0.00	2.65	444.06	5684.69	306.41
6	<i>out of memory</i>		<i>out of memory</i>						85976.57	0.00	38.27	8012.11	77925.96	2688.93

(b) Absence of starvation property φ_{AS} .

Table 6.4: Experimental model checking results for the preemptive scheduler system.

where $1 \leq m \leq n$ and $1 \leq l_m \leq \min(k_m, b)$. Finally, if none of the previous cases apply, the players' turns alternate:

$$\begin{aligned} t_E(\langle t_1, l_1, \dots, l_m, \dots, l_n \rangle, (\text{id}, \langle m, l \rangle, \text{id})) &\triangleq \langle t_2, l_1, \dots, l_m - l, \dots, l_n \rangle \\ t_E(\langle t_2, l_1, \dots, l_m, \dots, l_n \rangle, (\text{id}, \text{id}, \langle m, l \rangle)) &\triangleq \langle t_1, l_1, \dots, l_m - l, \dots, l_n \rangle \end{aligned}$$

where $0 \leq l_j \leq k_j$ for all $j \in [1..n]$, $1 \leq m \leq n$, $1 \leq l \leq \min(l_m, b)$, and $(l_1, \dots, l_m - l, \dots, l_n) \neq (0, \dots, 0)$.

The *player evolution* is simply defined as $t_i(\text{in}, a) \triangleq \text{in}$ for both players $i \in \{1, 2\}$ and all joint actions $a \in \text{Act}$.

- **Initial states.** The set of initial states $I \triangleq \{g\}$ contains a unique global state $g \in G$ such that $l_E(g) = \langle t_1, k_1, \dots, k_n \rangle$ and $l_1(g) = l_2(g) = \text{in}$.
- **Assignment.** We fix a set of *atomic propositions* $AP \triangleq \{p_1, p_2\}$ where p_1 and p_2 mean that the first and the second player won respectively. The *assignment* is formally defined as $h(p_i) \triangleq \{g \in G \mid l_E(g) = \langle w_i, 0, \dots, 0 \rangle\}$ for $i \in \{1, 2\}$.

The existence of a winning strategy for the first player is expressed in ATL and SL[1G] as:

$$\varphi_{\text{ATL}} \triangleq \langle\langle \{1\} \rangle\rangle F p_1 \qquad \varphi_{\text{SL}[1G]} \triangleq \langle\langle s_1 \rangle\rangle \llbracket s_2 \rrbracket \llbracket e \rrbracket (1, s_1)(2, s_2)(E, e) F p_1$$

Again, $\varphi_{\text{SL}[1G]}$ is also an SLK formula. Note that the specification cannot be expressed in CTL because it refers to the capability of an agent to enforce a property.

The model checking results for the formulas on one heap ($n = 1$) of various sizes and an upper bound on the number of removed objects $b = 3$ are shown in Table 6.5. We can observe that ATL can easily handle heaps with 5×10^5 objects in less than 1 minute. Although the SL[1G] performance is worse (as expected given its model checking complexity), the algorithm can still model check heaps with 5×10^4 objects in a reasonable amount of time (less than one hour)¹⁹. This is mainly due to the fact that the specification has a fixed size. In contrast, SLK runs out of memory when there are more than 16 objects on the heap, i.e. its performance is several orders of magnitude worse than that of both ATL and SL[1G]. This might be surprising given that SLK outperformed SL[1G] on the mutual exclusion property in the scheduler scenario (see Table 6.4a). The reason is that in the Nim game, all agents have complete information. Consequently, SLK strategies map individual states to actions (rather than shared local states to actions as in the previous scenarios), i.e. different BDD variables must be allocated for every single reachable global state. Hence, incomplete information actually makes SLK model checking easier²⁰ because it reduces the number of BDD variables needed to encode strategies (see Subsection 4.2.4).

Again, the individual steps of the SL[1G] model checking algorithm (see Figure 5.3) which take more than one second, namely arena construction, parity game combination, and parity game solution, are shown in Table 6.5 as well. Unlike in the previous scenarios, arena construction takes a significant amount of time. In fact, it dominates SL[1G] performance when the heap has more than 2×10^4 objects. This is an unexpected result, which we suspect might be caused by the size of the intermediate BDDs and the operations of the underlying BDD package (e.g. variable reordering) triggered during the arena construction. Observe that, unlike SLK, SL[1G] has approximately the same memory usage as ATL.

6.4.5 Analysis

In general, both the SLK and the SL[1G] extension have significantly lower performance than the original tool on CTL and ATL. We expected this outcome given the high model checking complexities of the SL fragments. However, this does not mean that the extensions are not useful in practice. They can verify specifications which cannot be checked by the original tool. We have discussed two examples of such properties: (i) Nash equilibria (see Subsection 6.4.2) and (ii) fairness conditions (see Subsection 6.4.3). Moreover, both our extensions support witness strategy synthesis, which can be used to automatically

¹⁹Surprisingly, SL[1G] model checking of the heap with 2×10^4 objects takes less than half the time of model checking the heap with 10^5 objects. We believe this might be caused by the operations of the underlying BDD package (e.g. variable reordering) as we could not find any other explanation.

²⁰Paradoxically, incomplete information makes SL[1G] model checking undecidable.

heap size	possible states	reachable states	reachability time (s)	ATL		SLK		SL[1G]				mem (MB)
				time (s)	mem (MB)	time (s)	mem (MB)	time (s)				
								total	arena	combine	solve	
10	44	20	0.00	0.00	2.64	0.02	3.10	0.03	0.00	0.00	0.03	3.24
11	48	22	0.00	0.00	2.65	0.13	4.30	0.03	0.00	0.00	0.03	3.15
12	52	24	0.00	0.00	2.66	1.05	11.50	0.04	0.00	0.01	0.03	3.17
13	56	26	0.00	0.00	2.67	10.03	39.34	0.04	0.00	0.01	0.03	3.26
14	60	28	0.00	0.00	2.67	74.38	181.08	0.04	0.00	0.00	0.03	3.19
15	64	30	0.00	0.00	2.67	458.11	821.16	0.03	0.00	0.00	0.02	3.24
16	68	32	0.00	0.00	2.67	1423.07	1874.01	0.04	0.00	0.01	0.03	3.20
17	72	34	0.00	0.00	2.68	<i>out of memory</i>		0.04	0.00	0.01	0.03	3.26
18	76	36	0.00	0.00	2.68	<i>out of memory</i>		0.05	0.00	0.01	0.04	3.21
19	80	38	0.00	0.00	2.70	<i>out of memory</i>		0.04	0.00	0.00	0.04	3.20
20	84	40	0.00	0.00	2.71	<i>out of memory</i>		0.05	0.00	0.00	0.04	3.21
50	204	100	0.00	0.00	2.85	<i>out of memory</i>		0.09	0.00	0.00	0.08	3.36
100	404	200	0.00	0.00	3.06	<i>out of memory</i>		0.14	0.00	0.00	0.14	3.56
200	804	400	0.00	0.01	3.41	<i>out of memory</i>		0.31	0.00	0.00	0.30	3.95
500	2004	1000	0.01	0.01	4.55	<i>out of memory</i>		0.82	0.01	0.00	0.80	5.16
1000	4004	2000	0.01	0.03	6.13	<i>out of memory</i>		1.66	0.07	0.00	1.59	7.28
2000	8004	4000	0.03	0.06	9.23	<i>out of memory</i>		7.58	1.33	0.02	6.23	11.39
5000	20004	10000	0.08	0.16	19.27	<i>out of memory</i>		30.37	7.97	0.02	22.28	23.78
10000	40004	20000	0.20	1.39	35.31	<i>out of memory</i>		625.60	33.46	0.25	591.88	45.36
20000	80004	40000	0.44	0.89	67.93	<i>out of memory</i>		305.42	163.51	0.03	141.88	72.70
50000	200004	100000	1.12	2.62	160.76	<i>out of memory</i>		1894.03	1337.15	0.09	556.79	171.07
100000	400004	200000	2.26	3.56	312.25	<i>out of memory</i>		7606.22	7152.87	0.01	453.34	324.56
200000	800004	400000	39.62	59.35	617.69	<i>out of memory</i>		59235.59	32838.33	4.82	26392.43	649.08
500000	2000004	1.00×10^6	43.99	57.43	1538.71	<i>out of memory</i>		<i>runs for more than 48 hours</i>				

Table 6.5: Experimental model checking results for the Nim game with one heap ($n = 1$) and at most 3 objects removed every turn ($b = 3$).

generate agents' behaviour that ensures the properties. In other words, the aim of the extensions is not to replace the existing tool, but to add support for completely new features. Therefore, the results presented in this section should be regarded as an illustration of the new functionality (rather than a benchmark with negative performance results).

The experimental results presented in this section confirm that the main performance bottleneck of the SLK model checking algorithm is the BDD encoding of the extended states, which allocates separate BDD variables for each shared local state of a strategy (see Subsection 4.2.4). As we have seen in the Nim scenario (see Subsection 6.4.4), this problem is more pronounced when all agents have complete information, in which case shared local states collapse to reachable global states. Informally, the less information agents have the better the SLK model checking performance is. This is quite a rare situation because incomplete information usually leads to higher model checking complexity (ATL with imperfect recall [23]) or even undecidability (ATL with perfect recall [35]). Nevertheless, the SLK extension can handle reasonably large state spaces for certain specifications as we have seen in the dining cryptographers scenario (see Subsection 6.4.1). Its performance can be slightly improved by using multiple threads (see Table 6.3). However, as we discussed at the end of Subsection 6.2.3, the potential of this approach seems to be limited because the CUDD BDD package used by MCMAS has no built-in support for concurrency.

The performance of the SL[1G] extension depends strongly on the specification. This is a direct consequence of the high SL[1G] model checking complexity (2EXPTIME-COMPLETE with respect to the size of the formula). As we have seen in the Nim scenario (see Subsection 6.4.4), the SL[1G] extension performs very well on large state spaces when the formula has a fixed size (see Table 6.5). Furthermore, the optimisation technique discussed in Subsection 5.2.5 significantly improves the model checking performance of SL[1G] conjunctions (see Table 6.4b). The relative performance of the SLK and SL[1G] extension also depends strongly on the type of the formula (compare Tables 6.4a and 6.4b). We found that the main bottlenecks of the SL[1G] model checking algorithms (see Figures 5.3 and 5.4) are parity game solving and arena construction. While the former is in line with our expectations, the reason for the latter remains unclear (see the discussion at the end of Subsection 6.4.4).

6.5 Summary

In this chapter, we discussed the existing MCMAS model checker and our two extensions for it which implement the novel model checking algorithms for SLK and SL[1G] introduced in Chapters 4 and 5 respectively. The experimental results we obtained on several scalable scenarios demonstrate the feasibility of practical model checking and strategy synthesis for both SL fragments despite the high theoretical complexities.

The first part of the chapter gave an overview of the existing tool and our two extensions. We started by describing the functionality, usage, and underlying architecture of the original tool. We then explained how each extension modified it in order to add support for the new features. We also demonstrated the ability of the extensions to synthesise strategies for SLK and SL[1G] specifications. Furthermore, we described an experimental parallel implementation of the SLK extension and discussed its limitations caused by the fact that the CUDD BDD package used by MCMAS does not support concurrency.

The second part of the chapter focused on the experimental results of the extensions. We introduced several scalable scenarios which we used to compare the performance of the original tool on CTL and ATL with the performance of the extensions on SLK and SL[1G]. The scenarios also demonstrated the greater expressiveness of the SL fragments (compared to CTL and ATL) as well as their relative performance, which strongly depends on the formula to be verified. We found that although the extensions have lower performance than the original tool on CTL and ATL, they are able to handle reasonably large state spaces.

Chapter 7

Evaluation

In Chapters 4 and 5, we introduced novel model checking algorithms for Epistemic Strategy Logic (SLK) and One-Goal Strategy Logic (SL[1G]) respectively. Furthermore, in Chapter 6, we presented two extensions of the existing model checker MCMAS (see Subsection 2.5.2) which implement the new model checking algorithms. In this chapter, we evaluate both the algorithms and their implementation. We then summarise the differences between the two SL fragments.

7.1 Theory

The first part of this project involved designing novel model checking algorithms for fragments of SL which can be implemented symbolically. We have done the following:

1. **Epistemic Strategy Logic** (Chapter 4). We defined a new fragment of SL, Epistemic Strategy Logic (SLK), on imperfect recall semantics with incomplete information. SLK is a specification language for multi-agent systems which combines LTL temporal connectives, SL strategy operators, and epistemic modalities. On the one hand, this combination results in a very powerful formalism for reasoning about game-theoretic concepts as well as agents' knowledge. On the other hand, SLK specifications suffer from certain limitations due to memoryless strategies, as discussed in Subsection 4.1.5.

We have defined the SLK model checking problem and proved that it is in the PSPACE complexity class with respect to both the size of the model and the formula (see Theorem 4.1). Furthermore, we have provided a model checking algorithm for SLK (see Subsection 4.2.2) which admits an efficient symbolic implementation (see Subsection 4.2.4). Finally, we discussed the problem of general SLK strategy synthesis and provided a procedure for constructing witness and counterexample strategies (see Subsection 4.2.3).

2. **One-Goal Strategy Logic** (Chapter 5). We redefined the syntax and semantics of an existing fragment of SL, One-Goal Strategy Logic (SL[1G]), on interpreted systems (see Definition 2.5). Despite being the least expressive syntactic fragment of full SL, SL[1G] subsumes many well-established temporal logics including those in the ATL* hierarchy while still having the same model checking complexity as ATL*, namely 2EXPTIME-COMPLETE with respect to the size of the specification and P-COMPLETE with respect to the size of the model.

We provided a novel procedure which reduces SL[1G] model checking to the problem of solving a parity game (see Subsection 5.2.1). We proved its correctness and showed that it is optimal in the sense that it achieves the lower complexity bound (see Theorem 5.2). Moreover, we provided an efficient symbolic implementation of the procedure with the same time complexity (see Subsection 5.2.4). Finally, we have explained how it can be extended to support general strategy synthesis for arbitrary SL[1G] formulas (see Subsection 5.2.3).

In addition, we proposed an optimisation technique for the SL[1G] model checking algorithm called *separate determinisation* (see Subsection 5.2.5) and described its impact on the complexity. We also explained how the symbolic implementation can be modified to incorporate the optimisation.

We shall now discuss the strengths and weaknesses of our theoretical contributions.

7.1.1 Strengths

Our theoretical contributions possess the following strengths:

1. **First practical SL algorithm.** All algorithms for model checking SL (or its fragments) introduced so far are very theoretical. They are mostly intended to serve as proofs of model checking decidability rather than to be implemented and used in practice. In contrast, our novel SLK and SL[1G] model checking algorithms were designed with implementation in mind from the very beginning. Furthermore, they both admit efficient symbolic representations (see Subsections 4.2.4 and 5.2.4). In our opinion, this is the first step towards practical SL model checking.
2. **Optimal SL[1G] complexity.** Not only have we designed the first practical algorithm for model checking SL[1G], but the algorithm also has optimal SL[1G] model checking time complexity, i.e. its running time is polynomial in the size of the model and doubly-exponential in the size of the specification (see Subsection 5.2.2). Moreover, it is optimal for ATL* model checking as well because ATL* has the same model checking complexity as SL[1G] (despite being strictly less expressive). However, the algorithm does not have optimal complexity for all logics in the ATL* hierarchy (e.g. LTL and CTL* are PSPACE-COMplete with respect to the size of the formula).
3. **Symbolic implementation.** Our model checking and strategy synthesis algorithms for both SLK and SL[1G] can be implemented symbolically using BDDs (see Subsections 2.3.1 and 2.3.2). While the theoretical complexity is not improved by the data structure¹, BDDs generally outperform explicit approaches in practice by several orders of magnitude (see Table 2.2). This is especially important because of the high model checking complexity of both fragments.
4. **Modularity.** Our novel model checking algorithms for SLK and SL[1G] are highly modular in the sense that individual steps can be replaced by alternative procedures with minimal impact on the rest of the algorithm. For example, the method for translating a non-deterministic Büchi automaton to a deterministic parity automaton (see Subsection 2.4.4) could be replaced with a different determinisation procedure by changing only one line of the SL[1G] model checking algorithm. We hope that our algorithm (or some of its parts) might serve as a basis for model checking algorithms for other SL fragments like Disjunctive/Conjunctive-Goal Strategy Logic (SL[DG/CG], see Section 3.2).
5. **Formally correct.** All concepts used in our model checking and strategy synthesis algorithms are formally defined. Moreover, all steps of the algorithms either reuse existing techniques (e.g. standard translation in Figure 2.6) or apply novel procedures which we proved are correct (e.g. combined parity game equivalence in Theorem 5.1). Therefore, our solutions of the model checking and strategy synthesis problems for SLK and SL[1G] are both sound and complete.

The main strength of our novel model checking algorithms for SLK and SL[1G] is the fact that they can be implemented symbolically. Moreover, we proved that both algorithms are correct and that the worst-case time complexity of our SL[1G] model checking algorithm is optimal.

7.1.2 Weaknesses

The main weakness of our theoretical contributions is the *suboptimal complexity* of the SLK model checking algorithm. We have shown that the SLK model checking problem belongs to the PSPACE class in Subsection 4.2.1. However, as we have already pointed out, the symbolic implementation of the SLK model checking algorithm presented in Subsection 4.2.4 might use more than a polynomial amount of space. Therefore, our SLK algorithm does not have optimal space complexity. Instead, we argued that it has singly-exponential worst-case time complexity in both the size of the model and the formula. We strongly believe that the space suboptimality is outweighed by the empirical efficiency of binary decision diagrams (see point 3 in Subsection 7.1.1).

¹In fact, our symbolic implementation of the SLK model checking algorithm has suboptimal space complexity (see Subsection 7.1.2).

7.2 Implementation

The second part of our project involved implementing the novel model checking algorithms and testing their performance. We have done the following:

1. **Implementation** (Sections 6.2 and 6.3). We implemented our model checking algorithms for SLK and SL[1G] as extensions of the existing model checker MCMAS (see Subsection 2.5.2). In both cases, the existing ISPL syntax was extended to accommodate for the new operators and the algorithm was translated to a sequence of symbolic operations on binary decision diagrams (see Subsections 2.3.1 and 2.3.2). In addition, the enhanced tool adds support for SLK witness/counterexample strategy synthesis (see Subsection 6.2.2) and SL[1G] witness/counterexample and general strategy synthesis (see Subsection 6.3.2).

In addition, we provided an *experimental parallel implementation* of the SLK model checking algorithm (see Subsection 6.2.3). The separate determinisation optimisation technique (see Subsection 5.2.5) is also supported by the SL[1G] extension.

2. **Experimental results** (Section 6.4). We tested both extensions on several scalable real-life scenarios and compared their performance with the performance of MCMAS on equivalent CTL and ATL formulas (see Tables 6.3, 6.4a, and 6.5). We also demonstrated the expressiveness of the fragments together with witness strategy synthesis (see Subsection 6.4.2). Finally, we compared the relative performance of the two SL fragments and evaluated the impact of the proposed optimisation techniques (experimental parallel implementation for SLK and separate determinisation for SL[1G]).

We shall now discuss the strengths and weaknesses of our implementation.

7.2.1 Strengths

Our implementation possess the following strengths:

1. **First SL tool.** As we have already mentioned several times, we are not aware of any existing tool that would support SL or any of its fragments. Therefore, MCMAS (with our extensions) is the first tool that supports SLK and SL[1G]. In our opinion, this is the biggest strength of our project because it demonstrates that model checking and strategy synthesis for both fragments is feasible in practice.

In fact, we are not even aware of any existing tool that would support ATL* (see Subsection 2.2.4). Therefore, we believe that our product is the first tool to support² ATL*. Consequently, MCMAS can now handle all logics in the ATL* hierarchy which it did not support before including LTL and CTL* (see Subsections 2.2.1 and 2.2.3).

2. **Symbolic implementation.** Since our novel model checking algorithms for SLK and SL[1G] admit efficient symbolic representations (see item 3 in Subsection 7.1.1), the extensions implement them using operations on BDDs (see Subsections 2.3.1 and 2.3.2), which outperform explicit approaches in practice by several orders of magnitude (see Table 2.1).
3. **Modularity.** Similarly to the new model checking algorithms (see item 4 in Subsection 7.1.1), our implementation of both MCMAS extensions is highly modular. This is very good from a software engineering perspective because our code can be easily reused or modified in the future by other people.
4. **Ease of use.** Despite its wide range of functionality, MCMAS is a simple command-line tool which is easy to use. The user only has to specify the name of the ISPL file to check and possibly pass a command-line flag if extra functionality is required (e.g. counterexample/witness generation). Both the input (ISPL syntax) and output of the tool is mostly self-explanatory and therefore simple to understand.

²As we have already explained in Subsection 6.3.1, the ATL* formula needs to be rewritten to an equivalent SL[1G] formula first. However, this is a very straightforward translation, which could be automatised by extending the ISPL syntax with ATL* strategy quantifiers.

In terms of input, our extensions merely *(i)* add three new command-line flags (listed in MCMAS help) and *(ii)* naturally extend the syntax with SL strategy operators (see Table 6.2). We believe that the format of the output is also quite easy to understand: SLK strategies are presented directly as mappings from states to actions and SL[1G] strategies are represented as automata (see Subsections 6.2.2 and 6.3.2).

The main strength of our solution is its novelty. We developed the first tool which supports model checking and strategy synthesis for fragments of SL. Moreover, the underlying algorithms for both fragments are implemented symbolically using efficient operations on BDDs.

7.2.2 Weaknesses

Our implementation has the following weaknesses:

1. **Performance.** The experimental results presented in Section 6.4 indicate that the biggest weakness of both extensions is their relatively low performance (compared to the performance of the original tool on CTL and ATL). This should not be surprising given the high expressiveness and therefore high model checking complexity of both SLK and SL[1G]:

While CTL and ATL belong to the P complexity class (see Table 2.1), the SLK and SL[1G] model checking problems are PSPACE³ and 2EXPTIME-COMPLETE respectively with respect to the size of the *specification*. For this reason, the tool cannot handle large SLK and SL[1G] formulas in general. Nevertheless, the performance of the extensions can be highly optimised for certain types of formulas: We implemented an optimisation technique for SL[1G] model checking called *separate determinisation* that speeds up the model checking of conjunctions of LTL formulas, which occur often in practice (see Subsection 5.2.5). As we can see in Table 6.4b, the technique significantly reduces both model checking time and memory usage and thereby increases the size of the state space that the extension can handle. Other possible optimisation techniques are discussed in Section 8.2.

Although our SLK model checking algorithm has exponentially lower time complexity with respect to the size of the formula than the one for SL[1G], it appears to be less feasible in practice. The time complexity of the model checking algorithm with respect to the size of the *model* is much more important because specifications are usually reasonably concise even for large state spaces. This turns out to be the main problem of our SLK model checking algorithm, which cannot handle large state spaces because it runs in (worst-case) exponential time in the size of the model (unlike the algorithms for CTL, ATL, and SL[1G]). This problem manifested itself in the Nim game (see Subsection 6.4.4), where the SLK performance was several orders of magnitude worse than that of both ATL and SL[1G]. Unfortunately, it appears that this is an unavoidable consequence of imperfect recall semantics (rather than a shortcoming of our algorithm).

2. **Limited testing.** As we have discussed in Subsection 6.1.3, the source code of MCMAS has a very bad design from a software engineering perspective (e.g. it has no tests whatsoever). While we tried to apply object-oriented programming principles during the development of the extensions (e.g. we used class hierarchy to represent different types of ω -automata as shown in Figure 6.3), we had to reuse some existing patterns in order to avoid “reinventing the wheel”. In particular, our code relies on global structures storing the BDD parameters, which makes it extremely difficult to write unit tests for individual procedures.

We have partially mitigated this issue by developing a framework for *system testing* of MCMAS and creating a test suite with several pairs of sample input and output files instead. Both model checking and strategy synthesis (for the new SL fragments) were tested by the framework. However, we think that unit tests for all MCMAS functions (not only our extensions) should be written if the tool is ever redesigned in the future and the dependency on global variables is reduced.

Overall, the main weakness of the implementation is its low performance, which is not surprising given the novelty of the algorithms and the high complexity of SLK and SL[1G] model checking. This issue can be mitigated to a large extent by applying optimisation techniques for various types of specifications. The development of such techniques is one of the most promising future directions of this project (see Section 8.2).

³As we have explained in Subsection 4.2.1, we conjecture that the SLK model checking problem is also PSPACE-HARD.

7.3 Fragment Comparison

Throughout this project, we were developing model checking frameworks for two fragments of SL in parallel. Our choice of the fragments, namely SLK and SL[1G], was explained in Chapter 3. The relevant theory and our novel model checking algorithms for both logics were presented in Chapters 4 and 5 respectively. The implementations of the algorithms in the form of MCMAS extensions were discussed in Sections 6.2 and 6.3 respectively. Finally, the experimental results of both extensions were provided in Section 6.4:

	SLK	SL[1G]
Fragment selection	Chapter 3	
Theory	Section 4.1	Section 5.1
Model checking	Section 4.2	Section 5.2
Implementation	Section 6.2	Section 6.3
Experimental results	Section 6.4	

We shall now summarise the differences and similarities between the two fragments as well as the corresponding model checking algorithms, implementations, and experimental results:

1. **Syntax.** Both fragments use SL strategy operators ($\langle\langle x \rangle\rangle$, $\llbracket x \rrbracket$, (i, x) with $x \in Var$ and $i \in Agt$, see Definition 2.19) and LTL temporal operators (X, F, G, U , see 2.7). While there are no restrictions⁴ on the structure of SLK formulas (e.g. an agent can be bound to a different strategy while the strategies of the other agents remain the same, see Definition 4.1), SL[1G] quantification and binding prefixes must be coupled together (i.e. all agents are always bound to newly quantified strategies simultaneously, see Definition 5.2). Furthermore, SLK adds support for epistemic modalities expressing agents knowledge (K_i, E_A, D_A, C_A with $i \in Agt$ and $A \subseteq Agt$, see Subsection 2.2.6), which SL[1G] does not support because its semantics is defined with respect to complete information. Therefore, all SL[1G] formulas are also SLK formulas (but not the other way round), i.e. $SL[1G] \subset SLK$.
2. **Semantics.** We defined the semantics of both fragments on interpreted systems (see Definition 2.5). SLK uses *imperfect recall semantics with incomplete information* (i.e. agents have no memory of the past and do not have complete knowledge of the system, see Definition 4.3), whereas SL[1G] uses *perfect recall semantics with complete information* (i.e. agents remember the whole history and have complete knowledge of the system, see Definition 2.29). As we explained in Chapter 3, we considered these two semantics because perfect recall semantics with incomplete information—which we would have ideally used instead—leads to an undecidable model checking problem.
3. **Strategies.** A key concept of both fragments are strategies which map possible configurations of the system to actions. The types of strategies that both fragments use are closely connected to their semantics (item 2):

SLK uses *uniform shared memoryless strategies* (see Definition 4.2), which map shared local states (of a group of agents $A \subseteq Agt$) to actions, i.e. $f_A : G / \sim_A^C \rightarrow Act_A$. SLK strategies are *non-behavioural*, i.e. the next action of an agent in one state might depend on the actions of other agents in counterfactual evolutions of the system (see Subsection 4.1.5).

SL[1G] uses *shared memoryful strategies* (see Definition 2.21), which map non-empty finite sequences of global states (tracks) to actions, i.e. $f_A : Trk \rightarrow Act_A$. Unlike SLK strategies, SL[1G] strategies are generally *not uniform* because a single local state history of an agent might be mapped to two different actions. However, SL[1G] strategies are *behavioural* (see Subsection 5.2.1).
4. **Expressiveness.** Both fragments can express the ability of agents to enforce temporal properties of the system by means of *shared interdependent strategies* (e.g. “*For all strategies of agent a, there exists a shared strategy for agents b and c, such that for all strategies ...*”). Therefore, SLK and SL[1G] are more expressive⁵ than ATL^* with imperfect and perfect recall respectively, which only

⁴As long as the formulas are *sentences* (see Definition 2.20).

⁵Recall that there is a subtle difference between SLK and ATL^* semantics in the way they handle universally quantified agents (see Subsection 4.2.1). Therefore, strictly speaking, SLK does *not* subsume ATL^* with imperfect recall. However, SL[1G] still strictly subsumes ATL^* with perfect recall (despite the difference).

Logic		Complexity w.r.t. $ \varphi $	Complexity w.r.t. $ \mathcal{I} $
imperfect recall	ATL*	PSPACE-COMplete	PSPACE-COMplete
	SLK	PSPACE	PSPACE
perfect recall	ATL*	2EXPTIME-COMplete	P-COMplete
	SL[1G]	2EXPTIME-COMplete	P-COMplete

Table 7.1: Model checking problem complexity of ATL*, SLK (see Theorem 4.1), and SL[1G] with respect to the size of the model $|\mathcal{I}|$ and the formula $|\varphi|$ [23, 72]. We conjecture that SLK is PSPACE-COMplete with respect to both the size of the model and the formula (see Subsection 4.2.1).

refer to *individual independent* strategies (e.g. “*There exists a strategy for agent d , such that no matter what the other agents do, . . .*”). This difference is due to the fact that SL handles strategies explicitly via quantifiers, whereas ATL* strategies are only implicit.

Since SLK syntax is not restricted and supports epistemic modalities, SLK is more expressive⁶ than SL[1G]: As we have seen in the dining cryptographers and scheduler scenarios (see Subsections 6.4.1 and 6.4.3), SLK can express agents’ knowledge and game-theoretic concepts like Nash equilibria. We explained that neither of these properties can be expressed in SL[1G].

5. **Model checking *problem* complexity.** The theoretical model checking problem complexity of both SLK and SL[1G] as well as ATL* is shown in Table 7.1. We can see that the SL fragments have the same worst-case complexities as their ATL* counterparts despite being more expressive. Moreover, SL[1G] has exactly the same model checking complexity as ATL* with perfect recall, which it strictly subsumes. As we have explained in Subsection 4.2.1, we conjecture that SLK has the same model checking complexity as its ATL* counterpart as well.

Both SL fragments have high model checking complexities, which have a negative impact on the performance of our extensions (see Section 6.4).

6. **Model checking *algorithm* complexity.** The worst-case time complexity of our SL[1G] model checking algorithm is *optimal* in the sense that it achieves the lower bounds on the problem complexity, i.e. doubly-exponential in the size of the formula and polynomial in the size of the model (see Theorems 5.2 and 5.4):

$$|\mathcal{I}|^{2^{O(|\varphi|)}}$$

Slightly tighter bounds for the unoptimised and optimised SL[1G] algorithm are provided in Subsections 5.2.2 and 5.2.5 respectively. Moreover, we argued that the algorithm is also optimal for ATL* (see item 2 in Subsection 7.1.1).

The worst-case time complexity of our SLK algorithm is exponential in both the size of the model and the formula (see Theorem 4.3):

$$2^{O(|\mathcal{I}| \times |\varphi|)}$$

While we are uncertain whether a sub-exponential time complexity can be achieved (in which case the time complexity of our algorithm would be suboptimal), we know that the space complexity of our algorithm is *suboptimal* because the problem belongs to the PSPACE complexity class (see Theorem 4.1), but our algorithm may use more than a polynomial amount of space (see Subsection 4.2.4).

7. **Optimisation techniques.** We have designed optimisation techniques for both SL fragments in order to improve the performance of the algorithms and extensions:

We developed an *experimental parallel implementation* of our SLK extension which uses multiple threads to enumerate the local states of an agent (see Subsection 6.2.3). Although it does have an impact on SLK performance (two threads reduced model checking time by roughly 16% in the dining cryptographers scenario, see Table 6.3 and Figure 6.4), it is not very scalable because the CUDD BDD package used by MCMAS has no built-in support for concurrency.

⁶This does *not* mean that SLK subsumes SL[1G]. Recall that there is a fundamental difference between the semantics of the fragments: SLK is defined with respect to *memoryless* strategies, whereas SL[1G] is defined with respect to *memoryful* strategies.

We proposed an optimisation technique for the SL[1G] model checking algorithm called *separate determinisation* (see Subsection 5.2.5), which improves the performance of the algorithm on SL[1G] principal sentences containing conjunctions of LTL formulas, i.e. SL[1G] sentences of the form $\wp b \bigwedge_{i=1}^n \psi_i$. Our experimental results indicate that the optimisation has a major impact on model checking time and memory: It reduced the time to verify an SL[1G] formula with 5 conjuncts on a model with 1.56×10^4 states almost 5 times and increased the state space that the SL[1G] extension can handle (see Table 6.4b). Moreover, the positive effect of the optimisation appears to increase with the number of conjuncts.

8. **Strategy synthesis.** Both extensions support *witness/counterexample strategy synthesis* for formulas of the form $\langle\langle x_0 \rangle\rangle \cdots \langle\langle x_{m-1} \rangle\rangle \psi$ and $\llbracket y_0 \rrbracket \cdots \llbracket y_{n-1} \rrbracket \phi$ (see Subsections 6.2.1 and 6.3.1). In addition, the SL[1G] extension supports *general strategy synthesis* for arbitrary SL[1G] sentences, which is demonstrated in Subsection 6.3.2. This functionality is not provided by the SLK extension because SLK strategies are non-behavioural, which makes general strategy synthesis extremely difficult as explained in Subsection 4.2.3.
9. **Command-line flags.** Both extensions reuse the existing `-c FORMAT` flag for enabling witness/counterexample strategy synthesis (see Subsections 6.2.2 and 6.3.2). Note that the extensions ignore the `FORMAT` parameter, which takes values 1–3. In addition, they introduce extra command-line flags specific to their functionality:

The SLK extension adds the `-t DEPTH` flag which enables thread branching of the experimental parallel implementation (with a given depth).

The SL[1G] extension adds the `-solutions` and `-sd SWITCH` flags which enable general strategy synthesis and toggle separate determinisation respectively (the `SWITCH` parameter can take values 0 and 1).
10. **Performance.** The performance of both extensions is relatively low compared to that of MCMAS on CTL and ATL (see item 1 in Subsection 7.2.2). This is not surprising given the high model checking complexities of the two SL fragments. We found that their performance depends on different factors: SLK performance depends mainly on the size of the state space and the amount of information available to the agents, whereas SL[1G] performance depends on the size and shape of the formula. These correlations are in line with the theoretical complexities of our model checking algorithms.

We can see that all the differences listed above stem from the different semantics used by the two fragments (item 2).

7.4 Summary

In this chapter, we evaluated both the theoretical and implementational contributions of our project. We first summarised our contributions in both areas and then discussed their strengths and weaknesses. The main strength of our solution is its novelty both on the theoretical and practical side as we have designed and implemented the first practical model checking algorithms for fragments of SL. The main weakness, which is low performance, can be overcome to a large extent by using various optimisation techniques. The impact of separate determinisation on SL[1G] model checking performance shows that the development of such techniques is one of the most promising future directions of this project (see Section 8.2).

We also provided a comprehensive comparison of SLK and SL[1G], both of which played a central role in our project. We found that all of the differences between the two fragments, the corresponding algorithms, and their implementations stem from their different semantics (SLK has imperfect recall with incomplete information, whereas SL[1G] has perfect recall with complete information).

Chapter 8

Conclusions

We tackled the problem of verifying *Strategy Logic* (SL) specifications. To do this, we designed the first practical model checking algorithms for two fragments of SL and implemented them as extensions of the existing MCMAS model checker. Our experimental results demonstrate that SL specifications, which can express various game-theoretic concepts including Nash equilibria, can be used in practice to verify and synthesise agents' behaviour. We have thereby made the first step towards bridging the gap between game-theory and practical model checking.

8.1 Summary of Work

To conclude the report, we summarise the work we have done and comment on whether we have met our objectives. The main goal of our project was to develop a model checker for SL. In order to achieve this goal, we first had to identify fragments of SL that are suitable for model checking. We found that the main obstacles to our goal were the undecidability of perfect recall semantics under incomplete information and the high model checking complexity of full SL. We addressed these obstacles by identifying two fragments of SL with desirable properties.

We first put forward a novel fragment of SL called *Epistemic Strategy Logic* (SLK), which is defined on imperfect recall semantics with incomplete information and adds support for epistemic modalities expressing agents' knowledge. We discussed its increased expressiveness due to the combination of strategic, temporal, and epistemic operators as well as its limitations caused by memoryless strategies. We proved that the model checking problem for SLK is PSPACE in both the size of the model and the formula. Finally, we designed an exponential-time model checking algorithm for SLK which we showed is correct, can be used to synthesise witness strategies, and admits an efficient symbolic implementation.

We then considered an existing fragment of SL called *One-Goal Strategy Logic* (SL[1G]), which is defined on perfect recall semantics with complete information. Although it is strictly more expressive than all logics in the ATL* hierarchy, it has the same model checking complexity. We designed the first practical model checking algorithm for SL[1G] by reducing the problem to solving a two-player parity game. We proved that procedure is correct and has optimal worst-case time complexity. Furthermore, we explained how it can be used for general strategy synthesis and provided an efficient symbolic implementation. We also provided an optimisation technique for the class of SL[1G] formulas containing conjunctions of LTL formulas, which often occur in practice.

We implemented our novel model checking algorithms for both SLK and SL[1G] as extensions for the existing model checker MCMAS. We have thus developed the first model checker for fragments of SL and thereby achieved our main goal. The functionality of the extended tool includes witness strategy synthesis for both fragments and general strategy synthesis for SL[1G]. We also implemented the optimisation technique for SL[1G] as well as an experimental parallel implementation of the SLK model checking algorithm. Our experimental results demonstrate that the use of the fragments for verification and strategy synthesis is feasible in practice despite their high model checking complexities. Given the novelty of our solution and the promising results, we deem our project successful.

8.2 Future Work

Although we have achieved all our objectives outlined in Section 1.1 and implemented the first model checker for (a subset of) SL, the journey to practical SL model checking is nowhere near its end. The main obstacle, which needs to be overcome, remains the low performance of both fragments on large models. Possible future directions of this project include:

- **Automaton minimisation.** Both the standard translation (see Subsection 2.4.3) and the determinisation procedure (see Subsection 2.4.4), which we use in the SL[1G] model checking algorithm (see Figure 5.3), generate automata which are not necessarily minimal. One way to improve SL[1G] model checking performance is therefore to always minimise intermediate automata before performing the next step of the algorithm. A wide range of techniques for minimising both non-deterministic and deterministic automata which can be implemented symbolically have been presented in [77].
- **Determinisation of specific formula classes.** An alternative approach for reducing the size of intermediate automata generated by the SL[1G] model checking algorithm is to use different determinisation procedures (such as subset and breakpoint construction) for specific classes of LTL formulas (e.g. safety conditions). The feasibility of this approach has been demonstrated in [77].
- **Better parity game algorithm.** Our SL[1G] model checking procedure currently uses a very simple exponential-time algorithm for solving parity games. A wide variety of other more efficient algorithms, which we could use instead, are discussed in [79]. We expect that such a modification would have a significant impact on SL[1G] model checking performance since parity game solving is one of the main bottlenecks of the procedure (see Subsection 6.4.5). Likewise, more efficient algorithms for solving generalised parity games, which we use in our optimised SL[1G] model checking algorithm (see Subsection 5.2.5), have been proposed [28].
- **Alternative verification techniques.** In this project we focused purely on symbolic model checking using BDDs. However, there are many other verification techniques (see Table 2.2), some of which can handle much larger state spaces than BDDs. It is possible that a more efficient model checker for SLK, SL[1G], or a different SL fragment could be based on one of them. In our opinion, the most likely candidate is *bounded model checking*, which is used by some existing tools (e.g. NuSMV, see Subsection 2.5.4) for LTL model checking.
- **Alternative BDD packages.** As we have explained in Subsection 6.2.3, the main obstacle to improving SLK model checking performance through parallelism is the fact that the CUDD BDD package used by MCMAS has no built-in support for concurrency. There are two possible solutions to this problem: We could either modify the package, or replace it with a different one that supports parallelism (e.g. the BDDNOW package [71]).
- **Other SL fragments.** The next step in terms of augmenting the functionality of our extensions is to add support for other more expressive fragments of SL (see Section 3.2). The most likely candidate is the Disjunctive/Conjunctive-Goal Strategy Logic (SL[DG/CG]), which admits behavioural strategies and still has 2EXPTIME-COMplete model checking complexity despite being strictly more expressive than SL[1G]. Given the similar properties of the two fragments, we hope that support for SL[DG/CG] could be added by modifying our SL[1G] model checking algorithm. On the other hand, the high model checking complexity and non-behavioural semantics of the other SL fragments (SL[BG] and SL[NG]) might turn out to be a too big obstacle to practical model checking. Even if that is not the case, we expect the corresponding model checking algorithms to be very different from the one we proposed here for SL[1G].
- **Incomplete information.** While perfect recall semantics with incomplete information is undecidable in general (see Chapter 3), there exist special settings in which the problem becomes decidable. One example is the class of *two-player parity games with imperfect information*, in which the first player has incomplete knowledge of the game, whereas the second player has complete knowledge. An algorithm for solving such games together with a symbolic implementation was proposed in [21]. Recall that our model checking algorithm for SL[1G] constructs a parity game where the first player

represents the existentially quantified strategy variables (see Subsection 5.2.1). Hence, we could use alternative $SL[1G]$ semantics where the agents bound to the existentially quantified strategy variables in a formula have (the same) incomplete information and verify the formula by reducing the problem to solving a parity game with imperfect recall.

- **Improve support for subsumed logics.** Although $SL[1G]$ subsumes all logics in the ATL^* hierarchy, our $SL[1G]$ extension does not have direct support for them, i.e. an ATL^* formula must be translated (by the user) to an equivalent $SL[1G]$ formula before it can be verified. Extending ISPL syntax (see Subsection 6.1.2) with ATL^* formulas should be rather easy and it would greatly enhance the usability of the extension. Moreover, given that we have already implemented the procedures for translating LTL formulas to equivalent non-deterministic Büchi automata, it should not be too difficult to add support for efficient LTL and CTL^* model checking as well.

Bibliography

- [1] The BDD library. <http://www.cs.cmu.edu/~modelcheck/bdd.html>. Accessed 26/01/2014.
- [2] Bison – GNU parser generator. <http://www.gnu.org/software/bison/>. Accessed 02/06/2014.
- [3] Boolean satisfiability research group at Princeton. <http://www.princeton.edu/~chaff/software.html>. Accessed 26/01/2014.
- [4] flex: The fast lexical analyzer. <http://flex.sourceforge.net/>. Accessed 02/06/2014.
- [5] GNU general public license. <http://www.gnu.org/licenses/gpl.html>. Accessed 25/01/2014.
- [6] GNU lesser general public license. <http://www.gnu.org/licenses/lgpl.html>. Accessed 24/01/2014.
- [7] MCK. <http://cgi.cse.unsw.edu.au/~mck/pmck/>. Accessed 26/01/2014.
- [8] MCMAS home page. <http://vas.doc.ic.ac.uk/software/mcmas/>. Accessed 25/01/2014.
- [9] MCMAS v1.2.1: User manual. <http://vas.doc.ic.ac.uk/software/mcmas/documentation/>. Accessed 02/06/2014.
- [10] NuSMV home page. <http://nusmv.fbk.eu/>. Accessed 24/01/2014.
- [11] PRISM – probabilistic symbolic model checker. <http://www.prismmodelchecker.org/>. Accessed 26/01/2014.
- [12] RSat homepage. <http://reasoning.cs.ucla.edu/rsat/>. Accessed 26/01/2014.
- [13] VAS – verification of autonomous systems. <http://vas.doc.ic.ac.uk/>. Accessed 02/06/2014.
- [14] MCK 1.0.0 – user manual. <http://cgi.cse.unsw.edu.au/~mck/pmck/mcks/docDownload/manual>, Feb. 2012. Accessed 26/01/2014.
- [15] Accellera. Property specification language - reference manual - version 1.1. <http://www.eda.org/vfv/docs/PSL-v1.1.pdf>. Accessed 24/01/2014.
- [16] R. Alur, L. de Alfaro, T. A. Henzinger, S. C. Krishnan, F. Y. C. Mang, S. Qadeer, S. K. Rajamani, and S. Tasiran. Mocha user manual. <http://mtc.epfl.ch/software-tools/mocha/doc/>. Accessed 15/06/2014.
- [17] R. Alur, T. Henzinger, F. Mang, S. Qadeer, S. Rajamani, and S. Tasiran. Mocha: Modularity in model checking. In A. J. Hu and M. Y. Vardi, editors, *Computer Aided Verification*, volume 1427 of *Lecture Notes in Computer Science*, pages 521–525. Springer Berlin Heidelberg, 1998.
- [18] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, Sept. 2002.
- [19] S. Anthony. Amazon unveils 30-minute Prime Air quadcopter delivery service, but it’s completely impractical. *ExtremeTech*, Dec. 2013. Accessed 25/01/2014.

- [20] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton. Verifying continuous time Markov chains. In R. Alur and T. Henzinger, editors, *Proc. 8th International Conference on Computer Aided Verification (CAV'96)*, volume 1102 of *LNCS*, pages 269–276. Springer, 1996.
- [21] D. Berwanger, K. Chatterjee, M. De Wulf, L. Doyen, and T. A. Henzinger. Alpaga: A tool for solving parity games with imperfect information. *CoRR*, abs/0901.4728, 2009.
- [22] A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In P. Thiagarajan, editor, *Proc. 15th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'95)*, volume 1026 of *LNCS*, pages 499–513. Springer, 1995.
- [23] N. Bulling, J. Dix, and W. Jamroga. Model checking logics of strategic ability: Complexity. In M. Dastani, K. V. Hindriks, and J.-J. C. Meyer, editors, *Specification and Verification of Multi-agent Systems*, pages 125–159. Springer US, 2010.
- [24] R. Carroll. CES 2014: driverless cars are coming and they want to be your friends. *The Guardian*, Jan. 2014. Accessed 25/01/2014.
- [25] O. Carton and M. Michel. Unambiguous Büchi automata. *Theoretical Computer Science*, 297(1–3):37–81, 2003. Latin American Theoretical Informatics.
- [26] P. Čermák, A. Lomuscio, F. Mogavero, and A. Murano. Few examples of interpreted systems. Personal communication, 2014.
- [27] P. Čermák, A. Lomuscio, F. Mogavero, and A. Murano. MCMAS-SLK: A model checker for the verification of strategy logic specifications. *CoRR*, abs/1402.2948, 2014.
- [28] K. Chatterjee, T. Henzinger, and N. Piterman. Generalized parity games. Technical Report UCB/EECS-2006-144, University of California, Berkeley, Nov. 2006.
- [29] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
- [30] A. Cimatti, E. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella. NuSMV version 2: An opensource tool for symbolic model checking. In *Proc. International Conference on Computer-Aided Verification (CAV 2002)*, volume 2404 of *LNCS*, Copenhagen, Denmark, July 2002. Springer.
- [31] A. Cimatti, E. Clarke, F. Giunchiglia, and M. Roveri. NUSMV: A new symbolic model checker. *International Journal on Software Tools for Technology Transfer*, 2(4):410–425, 2000.
- [32] A. Cimatti, M. Pistore, M. Roveri, and R. Sebastiani. Improving the encoding of LTL model checking into SAT. In *Proc. workshop on Verification Model Checking and Abstract Interpretation VMCAI'02*, volume 2294 of *LNCS*, Venice, Italy, Jan. 2002. Springer.
- [33] E. M. Clarke and E. A. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In D. Kozen, editor, *Logics of Programs*, volume 131 of *Lecture Notes in Computer Science*, pages 52–71. Springer Berlin Heidelberg, 1982.
- [34] E. M. Clarke, O. Grumberg, and K. Hamaguchi. Another look at LTL model checking. *Formal Methods in System Design*, 10(1):47–71, 1997.
- [35] C. Dima and F. L. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
- [36] N. Eén and N. Sörensson. MiniSat page. <http://www.minisat.se/>. Accessed 26/01/2014.
- [37] E. A. Emerson and J. Y. Halpern. “sometimes” and “not never” revisited: On branching versus linear time temporal logic. *J. ACM*, 33(1):151–178, Jan. 1986.
- [38] E. A. Emerson, C. S. Jutla, and A. P. Sistla. On model-checking for fragments of μ -calculus. In C. Courcoubetis, editor, *Computer Aided Verification*, volume 697 of *Lecture Notes in Computer Science*, pages 385–396. Springer Berlin Heidelberg, 1993.

- [39] E. A. Emerson and C.-L. Lei. Modalities for model checking: branching time logic strikes back. *Science of Computer Programming*, 8(3):275–306, 1987.
- [40] E. A. Emerson, A. K. Mok, A. P. Sistla, and J. Srinivasan. Quantitative temporal reasoning. In E. M. Clarke and R. P. Kurshan, editors, *Computer-Aided Verification*, volume 531 of *Lecture Notes in Computer Science*, pages 136–145. Springer Berlin Heidelberg, 1991.
- [41] S. Even and A. Paz. A note on cake cutting. *Discrete Applied Mathematics*, 7(3):285–296, 1984.
- [42] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*, volume 1 of *MIT Press Books*. The MIT Press, 2003.
- [43] B. Farwer. ω -automata. In E. Grädel, W. Thomas, and T. Wilke, editors, *Automata Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*, pages 3–21. Springer Berlin Heidelberg, 2002.
- [44] O. Friedmann and M. Lange. Solving parity games in practice. In Z. Liu and A. Ravn, editors, *Automated Technology for Verification and Analysis*, volume 5799 of *Lecture Notes in Computer Science*, pages 182–196. Springer Berlin Heidelberg, 2009.
- [45] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.
- [46] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2002.
- [47] M. Hilbert and P. López. The worlds technological capacity to store, communicate, and compute information. *Science*, 332(6025):60–65, 2011.
- [48] I. Hodkinson. C499: Modal and temporal logic. University Lecture, 2013. Department of Computing, Imperial College London.
- [49] M. Huth and M. Ryan. *Logic in Computer Science: Modelling and Reasoning about Systems*. Cambridge University Press, New York, NY, USA, 2004.
- [50] M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, 1998.
- [51] M. Jurdziński. Small progress measures for solving parity games. In H. Reichel and S. Tison, editors, *STACS 2000*, volume 1770 of *Lecture Notes in Computer Science*, pages 290–301. Springer Berlin Heidelberg, 2000.
- [52] M. Kacprzak, W. Nabiałek, A. Niewiadomski, W. Penczek, A. Pólrola, M. Szreter, B. Woźna, and A. Zbrzezny. VerICS 2007 – a model checker for knowledge and real-time. *Fundamenta Informaticae*, 85(1–4):313–328, 2008.
- [53] R. Kaivola, R. Ghughal, N. Narasimhan, A. Telfer, J. Whittmore, S. Pandav, A. Slobodová, C. Taylor, V. Frolov, E. Reeber, and A. Naik. Replacing testing with formal verification in Intel® Core™ i7 processor execution engine validation. In A. Bouajjani and O. Maler, editors, *Computer Aided Verification*, volume 5643 of *Lecture Notes in Computer Science*, pages 414–429. Springer Berlin Heidelberg, 2009.
- [54] H. Klauck. Algorithms for parity games. In *Automata, Logics, and Infinite Games*, pages 107–129, 2001.
- [55] S. Kripke. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83–94, 1963.
- [56] M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic symbolic model checking with PRISM: A hybrid approach. *International Journal on Software Tools for Technology Transfer (STTT)*, 6(2):128–142, 2004.

- [57] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In G. Gopalakrishnan and S. Qadeer, editors, *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*, volume 6806 of *LNCS*, pages 585–591. Springer, 2011.
- [58] F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. In H. Seidl, editor, *Foundations of Software Science and Computational Structures*, volume 4423 of *Lecture Notes in Computer Science*, pages 243–257. Springer Berlin Heidelberg, 2007.
- [59] N. G. Leveson and C. S. Turner. An investigation of the Therac-25 accidents. *Computer*, 26(7):18–41, 1993.
- [60] J. Lind-Nielsen. BuDDy – a binary decision diagram package. <http://vlsicad.eecs.umich.edu/BK/Slots/cache/www.itu.dk/research/buddy/>. Accessed 26/01/2014.
- [61] A. Lomuscio. *Knowledge Sharing among Ideal Agents*. PhD thesis, School of Computer Science, University of Birmingham, Birmingham, UK, June 1999.
- [62] A. Lomuscio. C303: Systems verification. University Lecture, 2013. Department of Computing, Imperial College London.
- [63] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: A model checker for the verification of multi-agent systems. In *Computer Aided Verification*, volume 5643 of *Lecture Notes in Computer Science*, pages 682–688. Springer Berlin Heidelberg, 2009.
- [64] A. Lomuscio and F. Raimondi. The complexity of model checking concurrent programs against CTLK specifications. In M. Baldoni and U. Endriss, editors, *Declarative Agent Languages and Technologies IV*, volume 4327 of *Lecture Notes in Computer Science*, pages 29–42. Springer Berlin Heidelberg, 2006.
- [65] A. Lomuscio and F. Raimondi. Model checking knowledge, strategies, and games in multi-agent systems. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, AAMAS '06, pages 161–168, New York, NY, USA, 2006.
- [66] A. Lomuscio and M. Sergot. Deontic interpreted systems. *Studia Logica*, 75(1):63–92, 2003.
- [67] Z. Manna and A. Pnueli. Specification and verification of concurrent programs by \forall -automata. In B. Banieqbal, H. Barringer, and A. Pnueli, editors, *Temporal Logic in Specification*, volume 398 of *Lecture Notes in Computer Science*, pages 124–164. Springer Berlin Heidelberg, 1989.
- [68] R. Mazala. Infinite games. In E. Grädel, W. Thomas, and T. Wilke, editors, *Automata Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*, pages 23–38. Springer Berlin Heidelberg, 2002.
- [69] K. L. McMillan. *Symbolic Model Checking*. Kluwer Academic Publishers, Norwell, MA, USA, 1993.
- [70] R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65(2):149–184, 1993.
- [71] K. Milvang-Jensen, , and A. J. Hu. BDDNOW: A parallel BDD package. In *Proc. FMCAD*, pages 501–507. Springer, 1998.
- [72] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. Reasoning about strategies: On the model-checking problem. *CoRR*, abs/1112.6275, 2011.
- [73] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. A decidable fragment of strategy logic. *CoRR*, abs/1202.1309, 2012.
- [74] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. What makes ATL* decidable? a decidable fragment of strategy logic. In M. Koutny and I. Ulidowski, editors, *CONCUR 2012 Concurrency Theory*, volume 7454 of *Lecture Notes in Computer Science*, pages 193–208. Springer Berlin Heidelberg, 2012.

- [75] F. Mogavero, A. Murano, and L. Sauro. On the boundary of behavioral strategies. In *Logic in Computer Science (LICS), 2013 28th Annual IEEE/ACM Symposium on*, volume 0, pages 263–272, June 2013.
- [76] F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In K. Lodaya and M. Mahajan, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, volume 8 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 133–144, Dagstuhl, Germany, 2010. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [77] A. Morgenstern. *Symbolic controller synthesis for LTL specifications*. PhD thesis, 2010.
- [78] A. Morgenstern and K. Schneider. From LTL to symbolically represented deterministic automata. In F. Logozzo, D. A. Peled, and L. D. Zuck, editors, *Verification, Model Checking, and Abstract Interpretation*, volume 4905 of *Lecture Notes in Computer Science*, pages 279–293. Springer Berlin Heidelberg, 2008.
- [79] J. Obdržálek. *Algorithmic Analysis of Parity Games*. PhD thesis, University of Edinburgh, 2006. Submitted: January 31, 2006. Examined: May 29, 2006.
- [80] A. Pnueli. The temporal logic of programs. In *Foundations of Computer Science, 1977., 18th Annual Symposium on*, pages 46–57, Oct 1977.
- [81] F. Raimondi and A. Lomuscio. Automatic verification of multi-agent systems by model checking via OBDDs. *Journal of Applied Logic*, 2005.
- [82] S. Safra. On the complexity of ω -automata. In *Foundations of Computer Science, 1988., 29th Annual Symposium on*, pages 319–327, Oct 1988.
- [83] K. Schneider. *Verification of Reactive Systems*. Texts in Theoretical Computer Science. Springer Berlin Heidelberg, 2004.
- [84] P. Schnoebelen. The complexity of temporal logic model checking. *Advances in Modal Logic*, 4:393–436, 2002.
- [85] F. Somenzi. CUDD: CU decision diagram package. <http://vlsi.colorado.edu/~fabio/CUDD/>. Accessed 26/01/2014.
- [86] M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Inf. Comput.*, 115(1):1–37, 1994.
- [87] B. Woźna and A. Zbrzezny. Bounded model checking for the universal fragment of CTL, 2002.
- [88] W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1–2):135–183, 1998.
- [89] M. Zimmermann. Optimal Bounds in Parametric LTL Games. *ArXiv e-prints*, June 2011.

Appendix A

Toy Model ISPL Code

We provide here an ISPL description of the toy model introduced in Section 3.3.

```
1 Agent Environment
2   Obsvars:
3     state : {game, p1win, p2win};
4   end Obsvars
5   Actions = {idle};
6   Protocol:
7     Other : {idle};
8   end Protocol
9   Evolution:
10    state = p1win if state = game and
11      (Player1.Action = rock and Player2.Action = scissors or
12       Player1.Action = paper and Player2.Action = rock or
13       Player1.Action = scissors and Player2.Action = paper);
14    state = p2win if state = game and
15      (Player2.Action = rock and Player1.Action = scissors or
16       Player2.Action = paper and Player1.Action = rock or
17       Player2.Action = scissors and Player1.Action = paper);
18  end Evolution
19 end Agent
20
21 Agent Player1
22   Vars:
23   end Vars
24   Actions = {rock, paper, scissors, idle};
25   Protocol:
26     Environment.state = game: {rock, paper, scissors};
27     Other: {idle};
28   end Protocol
29   Evolution:
30   end Evolution
31 end Agent
32
33 Agent Player2
34   Vars:
35   end Vars
36   Actions = {rock, paper, scissors, idle};
37   Protocol:
38     Environment.state = game: {rock, paper, scissors};
39     Other: {idle};
```

```
40   end Protocol
41   Evolution:
42   end Evolution
43 end Agent
44
45 Evaluation
46   p1 if Environment.state = p1win;
47   p2 if Environment.state = p2win;
48 end Evaluation
49
50 InitStates
51   Environment.state = game;
52 end InitStates
53
54 Groups
55   g = {Player1, Player2};
56 end Groups
57
58 Formulae
59   -- Existing Tool (examples/toy_model_existing.ispl)
60   EF p1;
61   <g>G !(p1 or p2);
62
63   -- SLK Extension (examples/toy_model_slk.ispl)
64   <<e>> (Environment, e) <<x>> (Player1, x) <<y>> (Player2, y) G !(p1 or p2);
65
66   -- SL[1G] Extension (examples/toy_model_sl1g.ispl, examples/toy_model_sl1g2.ispl)
67   #PR <<e>> (Environment, e) <<x>> (Player1, x) <<y>> (Player2, y) G !(p1 or p2);
68   #PR [[e]] (Environment, e) [[x]] (Player1, x) <<y>> (Player2, y) G !(p1 or p2);
69 end Formulae
```