

Synthesis from multi-paradigm specifications

Ioannis Filippidis

Richard M. Murray

Gerard J. Holzmann

March 9, 2015

Abstract

This work proposes a language for describing reactive synthesis problems that integrates imperative and declarative elements. The semantics is defined in terms of two-player turn-based infinite games with full information. Currently, synthesis tools accept linear temporal logic (LTL) as input, but this description is less structured and does not facilitate the expression of sequential constraints. This motivates the use of a structured programming language to specify synthesis problems. Transition systems and guarded commands serve as imperative constructs, expressed in a syntax based on that of the modeling language PROMELA. The syntax allows defining which player controls data and control flow, and separating a program into assumptions and guarantees. These notions are necessary for input to game solvers. The integration of imperative and declarative paradigms allows using the paradigm that is most appropriate for expressing each requirement. The declarative part is expressed in the LTL fragment of generalized reactivity(1), which admits efficient synthesis algorithms. The implementation translates PROMELA to input for the SLUGS synthesizer and is written in PYTHON.

Keywords: Reactive synthesis, Generalized Reactivity(1), Linear temporal logic, Infinite games, PROMELA

© 2015 by the authors.
All rights reserved.
Any trademarks remain property of their respective owners.

Contents

1	Introduction	5
1.1	Preliminaries	6
1.1.1	Linear Temporal Logic	6
1.1.2	Turn-based games	7
1.1.3	Games in logic	7
2	Language definition	8
2.1	Variables	8
2.1.1	Ownership	8
2.1.2	Declarative and imperative semantics	8
2.1.3	Ranged integer data type	9
2.1.4	Notation	9
2.2	Game graphs as programs	9
2.2.1	Syntax	10
2.2.2	Example	11
2.3	Expressions	11
2.3.1	Functions, predicates, and connectives	11
2.3.2	State predicates and actions	13
2.3.3	Deconstraining	13
2.3.4	Executability	14
2.4	Assignments	17
2.4.1	Truncation semantics	18
2.4.2	Equality operator in logic	18
2.5	Automata products	18
2.6	Atomic statements	19
2.6.1	Controllability of process execution	20
2.6.2	Expressing atomicity in temporal logic	21
2.6.3	Visibility to properties	22
2.6.4	Stutter invariance	22
2.6.5	Alternatives that accommodate stutter-sensitivity	24
2.7	Verification in open <code>Promela</code>	24
3	Translation to logic	26
3.1	Notation	26
3.2	Control and data flow	27
3.2.1	Initial conditions	28
3.3	Invariance of variables	29
3.4	Process scheduler	30
3.5	Exclusive execution	30
3.5.1	Stuttering and visibility	31
3.5.2	Unbounded loops	31

4	Implementation	32
4.1	Program graph construction	35
4.2	Bitwise encoding	37
4.2.1	Modwrap and saturating semantics	38
4.2.2	Bitfield representation of signed ranged integers	39
4.2.3	Arithmetic as bitvector logic	39
A	Appendices	42
A.1	Open Promela syntax	42
A.1.1	Lexemes	42
A.1.2	Grammar	42
A.1.3	The keyword <code>step</code>	43
A.2	Relevant work	43
A.2.1	RPromela and RSPIN	44
A.2.2	KALEIDOSCOPE	44

Chapter 1

Introduction

Over the past three decades system verification has aided design and become practical for industrial application. In the past decade, synthesis of systems from specifications has seen significant development [1, 2], partially owing to the discovery of temporal logic fragments that admit efficient synthesis algorithms [3, 4, 5, 6]. Applications range from protocol synthesis for hardware circuits [4], to correct-by-construction controllers for hybrid systems [7, 8, 9].

Model checking has polynomial complexity in the size of a given Kripke structure, which has motivated the development of a large number of tools and languages for modeling systems. Unlike verification using model checking, the tools for synthesis have been developed much more recently. One reason is that centralized synthesis from linear temporal logic (LTL) has doubly exponential complexity in the length of the specification formula [10], a result that did not encourage further development initially.

As a result, currently LTL is the language used for describing specifications as input to synthesis tools. There are many benefits in using a logic for synthesis tasks, its declarative nature being a major one, because it allows expressing separately individual requirements in a precise way. It also makes explicit the implicit conventions present in programming languages [11]. Another aspect of synthesis problems that makes declarative descriptions appropriate is that we want to describe as large a set of possible designs as possible, in order to avoid overconstraining the search space.

Paradigms However, not all specifications are best described declaratively. There exist synthesis problems whose description involves graph-like structures that are cumbersome for humans to write in logic. For example, robotics problems typically involve graph constraints that originate from possible physical configurations. Properties that specify sequential behavior also lead to graph-like structures, and require use of auxiliary variables that serve as memory. Expressing sequential composition in logic leads to long, unstructured formulas that deemphasize the specifier’s intent. The resulting specifications are difficult to maintain, and writing them is error-prone. In addition, the specifier may need to explicitly write clauses that constrain variables to remain unchanged, in order to maintain imperative state. This leads to longer formulas in which the intent behind individual clauses is less readable.

Borel hierarchy Another motivation stems from the temporal logic, or Borel, hierarchy [12, 13]. The more expressive a syntactic fragment of logic is, the more computationally expensive it is to construct a model, or prove that none exists.

Currently, algorithms of time complexity polynomial in the size of the state space are known for the fragment of generalized reactivity of rank one (one Streett pair), known as GR(1). In the automata hierarchy, the GR(1) fragment corresponds to deterministic Büchi automata [3, 13, 14, 15]. From a game theoretic perspective, deterministic automata describe the winning condition in such a way, that existential branching during synthesis commits to a future strategy, without the ability to recall it.

A large subset of properties that are of practical interest in industrial applications [16, 17, 4] can be expressed in the GR(1) fragment. There do exist properties that cannot be represented by deterministic Büchi automata, e.g., persistence $\diamond\Box p$ (eventually p holds uninterruptedly). Of these properties, those with Rabin rank equal to one are still amenable to polynomial time algorithms, by solving parity games with

a limited number of colors [5]. Higher Rabin ranks are not expected to admit solution by algorithms of polynomial complexity, unless $P = NP$ [5]. This motivates formulating the required properties in GR(1).

Translating properties to deterministic automata can be done automatically, but may lead to more states than manually written properties [14]. So the ability to write deterministic automata directly in a structured and readable language avoids the need for conversions, either automatic, or manual.

Synthesis tools have been developed for the GR(1) fragment, so temporal logic can be viewed as a lower level language above the representation used inside a solver, e.g., binary decision diagrams [18]. In this context, one can think of temporal logic as an analog of assembly language. It follows that by translating a problem description to logic, it can then be given as input to a suitable synthesis tool.

Modularization Another reason why specifications are not always purely declarative is that in many cases we want to synthesize a system using *existing* components. In other words, we already have a *partial* model of reality, and we *declare* to our synthesis tool what properties the controller under design should satisfy with respect to this model. So the partial model is best described imperatively, whereas the goal declaratively, using temporal logic.

Other Educationally, the transition for engineering students from a general purpose programming language, like PYTHON or C, directly to temporal logic constitutes a significant leap. Using a multiparadigm language can make this transition smoother.

The preceding arguments on the need for a more structured specification language that spans paradigms do not allude as to whether `Promela` or some other syntactic variant is more appropriate for this purpose. The selection of syntax will be discussed throughout, and the results can be adapted to any suitable syntax.

Contents This work proposes a language that can describe synthesis problems for open systems that react to an adversarial environment. The syntax is derived from that of `Promela`, whereas the semantics interprets it as a two-player turn-based game of infinite duration. Both synchronous and asynchronously scheduled centralized systems with full information can be synthesized.

In Section 1.1, we review temporal logic and relevant notions about two-player games. The semantics of the language is defined in Chapter 2. The presence of two players requires declaring who controls each variable (Section 2.1), as well as the data flow, and control flow in transition systems (Section 2.2). In addition, the specification needs to be partitioned into assumptions about the environment, and guarantees that the system must satisfy (Section 1.1, Section 2.2). The integration of declarative and imperative semantics is obtained by defining imperative variables (Section 2.1), deconstraining, and executability of actions (Section 2.3.4). In order to be synthesized, the program is translated to temporal logic, as described in Chapter 3. Relevant work is collected in Appendix A.2.

Acknowledgments This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA. The first author was partially supported by a graduate research fellowship from the Jet Propulsion Laboratory, over the summer of 2014. The authors would like to thank Scott Livingston for providing helpful comments on an early draft of this document.

1.1 Preliminaries

1.1.1 Linear Temporal Logic

Linear temporal logic is an extension of Boolean logic used to reason about temporal modalities over sequences. The temporal operators “next” \bigcirc and “until” \mathcal{U} suffice to define the other operators [19, 20]. Let AP be a set of variable symbols that can take values over $\mathbb{B} \triangleq \{\perp, \top\}$. A model of an LTL formula is a sequence of variable valuations called a *word* $w : \mathbb{N} \rightarrow \mathbb{B}^{AP}$. The formula $\bigcirc p$ is modeled by a sequence where $w[1]$ maps variable p to true. The formula $p\mathcal{U}q$ is modeled by a sequence where there exists a $j \in \mathbb{N}$ such that $w[j]$ maps q to true and for all $i \in \mathbb{N}_{<j}$, $w[i]$ maps p to true. The operator $\diamond p \triangleq \text{true}\mathcal{U}p$ requires that p be “eventually” true, and the operator $\square p \triangleq \neg\diamond\neg p$ requires that p be true over the whole word.

1.1.2 Turn-based games

In many applications, we are interested in designing a system that does not have full control over its reality. Some problem variables represent the behavior of another entity, usually referred to as the “environment”. The system reads these *input* variables and reacts by writing to *output* variables that it controls, continuing indefinitely. Such a system is called *open* [21, 22], to distinguish it from closed systems that have no inputs, and so full control.

The synthesis of an open system can be formulated as a two-player adversarial game of infinite duration [23, 24, 25]. The two players in the game are usually called the protagonist (system) and antagonist (environment). We control the protagonist, but not the antagonist. If the players move in turns, then the game is called *alternating*. Each pair of consecutive moves by the two players is called a *turn* of the game. In each turn, player 0 moves first, without knowing how player 1 will choose to move in that turn of the game. Then player 1 moves, knowing how player 0 moved in that turn.

Depending on which player we control, there are two types of game. If the protagonist is player 1, then the game is called *Mealy*, otherwise *Moore* [26, 27]. Due to the difference in knowledge about the opponent’s next move between the two flavors of game, more specifications are realizable in a Mealy game, than in a Moore game. There exist solvers for both Moore and Mealy games. Here we will consider Mealy games only, but the results can be adapted to Moore games as well.

1.1.3 Games in logic

Temporal logic can be used to describe both the possible moves in a game (the *arena* or *game graph*), as well as the winning condition (e.g., fairness). Let \mathcal{X} and \mathcal{Y} be two sets of propositional variables, controlled by the environment and system, respectively. Let \mathcal{X}' and \mathcal{Y}' denote primed variables, where x' represents the next value $\bigcirc x$ of variable x . We abuse notation by using primed variables inside temporal formulae.

Synthesis from LTL specifications is in 2EXPTIME [28, 10], motivating the search for fragments that admit more efficient synthesis algorithms. Generalized reactivity of index one, abbreviated as GR(1), is a fragment of LTL that admits synthesis algorithms of time complexity polynomial in the size of the state space [4]. We use GR(1) in the following [29, 30, 31, 32, 33], but the results can be adapted to larger fragments of LTL, provided that another synthesizer be used [34, 35, 36, 37, 38].

The possible moves in a Mealy game can be specified by initial and transition conditions that constrain the environment and system. Initial conditions are described by propositional formulae over $\mathcal{X} \cup \mathcal{Y}$. Transition conditions are described by safety formulae of the form $\Box \varphi_i$ where, for the environment $i = e$ and φ is a formula over $\mathcal{X} \cup \mathcal{X}' \cup \mathcal{Y}$, and for the system $i = s$ and φ is a formula over $\mathcal{X} \cup \mathcal{X}' \cup \mathcal{Y} \cup \mathcal{Y}'$. Note that the system plays second in each turn, so it can see \mathcal{X}' , whereas the environment cannot see \mathcal{Y}' , because it represents future values. The winning condition in a GR(1) game is described using progress formulae of the form $\Box \Diamond \psi_i, i \in \{e, s\}$, where ψ_i is a propositional formulae over $\mathcal{X} \cup \mathcal{Y}$.

The overall specification of a GR(1) game is of the form

$$\underbrace{\left(\bigwedge_{i=0}^{n_0} \theta_{e,i} \wedge \bigwedge_{i=0}^{n_1} \Box \varphi_{e,i} \wedge \bigwedge_{i=0}^{n_2} \Box \Diamond \psi_{e,i} \right)}_{\text{assumption}} \rightarrow \underbrace{\left(\bigwedge_{i=0}^{m_0} \theta_{s,i} \wedge \bigwedge_{i=0}^{m_1} \Box \varphi_{s,i} \wedge \bigwedge_{i=0}^{m_2} \Box \Diamond \psi_{s,i} \right)}_{\text{assertion}} \quad (1.1)$$

Note that requirements that constraint the environment are called *assumptions* and guarantees that the system must satisfy are called *assertions*. Assumptions limit the set of admissible environments, because, in practice, it is impossible to satisfy the design requirements in arbitrarily adversarial environments [21]. The implication above is interpreted by prioritizing between safety and liveness, to prevent the system from violating its safety assertion in case this would allow it to prevent the environment from satisfying its liveness assumption.

The algorithm for GR(1) synthesis proposed in [3] has complexity $O((nm|\Sigma|)^3)$ and the one later proposed in [4] has complexity $O(nm|\Sigma|^2)$, where n is the size of the assumption formula, m the size of the assertion formula, and Σ the state space (all possible valuations of variables), with size $|\Sigma|$ is exponential in the number of variables.

Chapter 2

Language definition

The language we are about to define is syntactically an extension of PROMELA [39], but its semantics is defined by a translation to turn-based infinite games with full information. PROMELA is a guarded command language that can represent transition systems, non-deterministic execution, and guard conditions for determining whether statements are executable [39, 40]. Its syntax can be found in the language reference manual [39, 41]. Here we will introduce syntactic elements only as needed for the presentation. Briefly, we mention that a program comprises of transition systems and automata, whose control flow can be described with sequential composition, selection and iteration statements, `goto`, as well as blocks that group statements for atomic execution.

2.1 Variables

2.1.1 Ownership

In a modeling language for closed systems, one player controls all the variables. In contrast, in a game there are two players. Let \mathcal{X} denote the variables controlled by the environment, and \mathcal{Y} those controlled by the system. We call *owner* of a variable the player that controls it. We use the keywords `env` and `sys` to signify the owner of a variable. Variables can be of Boolean, bit, byte, (bounded) integer, or bitfield type. These types are discussed more in Sections 2.1.3 and 4.2.1.

So there are two flavors of variables, depending on who decides their values:

1. controlled variables (system variables, also known as outputs, or output ports)
2. uncontrolled variables (environment variables, also known as inputs, or input ports)

In the global context (i.e., outside `proctype` declarations), the default owner is the system. Let V_p denote the set of program variables declared as controlled by player $p \in \{e, s\}$, for environment and system, respectively. The set V_e is a subset of \mathcal{X} , because it does not include auxiliary variables that may be introduced by the compiler as described in Chapter 3. From a first-order logic viewpoint, ownership defines quantification [22]. System variables are existentially quantified, whereas environment variables universally.

2.1.2 Declarative and imperative semantics

In imperative languages, variables remain unchanged, unless explicitly assigned new values. In declarative languages, variables are free to change, unless explicitly constrained [42]. In verification, both declarative languages like TLA [43] and SMV [44] have been used, as well as imperative languages like PROMELA and DVE [45].

In a synthesis problem, there are variables whose behavior be succincter to describe declaratively, whereas others imperatively. For this reason, we combine the two paradigms, by introducing a new keyword `free` to distinguish between imperative and declarative variables. Variables whose declaration includes the keyword `free` are by default allowed to be assigned any value in their domain, unless explicitly constrained otherwise.

Variables without the keyword `free` have imperative semantics, so their value remains unchanged, unless otherwise explicitly stated. Let V_f denote free, and V_c the imperative program variables, respectively.

2.1.3 Ranged integer data type

Symbolic methods for synthesis use binary decision diagrams (BDDs) [18, 20]. Reordering variables can be expensive, so reducing the number of bits is a primary objective, because each extra propositional variable of the problem introduces two BDD variables (a primed and an unprimed one). This effect has a milder impact on enumerated methods, because it increases the memory required to store each state, but leaves unaffected the number of reachable states. But for BDDs, the number of BDD nodes is sensitive to the number, and ordering, of variables.

In addition, the complexity of GR(1) synthesis is polynomial in the number $|\Sigma|$ of variable valuations (Section 1.1.3), which grows exponentially with each additional bit. We can reduce the number of bits by using bitfields whose width is tailored to the problem at hand, together with safety constraints. For example, a variable x that can take values in the set $\{0, 1, 2\}$ is represented in logic by two propositional variables $x_0, x_1 \in \mathbb{B}$ and the constraint

$$\square(0 \leq x \leq 2) \iff \square((x_0 \neq 1) \vee (x_1 \neq 1)),$$

which prevents the valuation $x_0 = 1, x_1 = 1$, i.e., $x = 3$.

For convenience, the *ranged* integer type `int(MIN, MAX)` is introduced to define a variable

$$x \in \{\text{MIN}, \text{MIN} + 1, \dots, \text{MAX}\}.$$

This syntax is a slight extension of the C-like syntax of pure `Promela`. Ranged integer variables have saturating semantics [46]. Safety constraints are automatically imposed on the bitfield representing a ranged integer. Other numerical data types (bits, bytes, bitfields) have mod wrap semantics (modular arithmetic).

2.1.4 Notation

Summarizing the notation defined earlier:

- \mathcal{X} are all the variables controlled by the system,
- \mathcal{Y} are all the variables controlled by the environment,
- V is the set of all variables declared in a program source listing,
- $V_s \subseteq V$ are the program variables controlled by the system,
- $V_e \triangleq V \setminus V_s$ are the program variables controlled by the environment,
- V_f are the declarative program variables, and
- $V_c \triangleq V \setminus V_f$ are the imperative program variables.

For a set of variables V , V' denotes the set of primed variables u' , for each variable $u \in V$. The set $\text{Var}(\varphi)$ denotes the variables that appear in formula φ . Combinations of subscripts denote intersection, for example $V_{c,s} = V_c \cap V_s$ is the set of imperative system variables.

2.2 Game graphs as programs

In many synthesis problems, the specification includes graph-like constraints. These may originate from physical configurations in robotics problems, deterministic automata to express a formula in GR(1), or describe abstractions of existing components that are to be controlled. Program graphs can be used to describe graph-like constraints. A *program graph* is a rooted directed multi-graph $P_r \triangleq (V_r, E_r, u)$ whose edges $E_r \subset V_r \times V_r \times \mathbb{N}$ are labeled with program statements, and nodes V_r abstract states of the system [47, 20]. Execution starts from the graph's root $u \in V_r$.

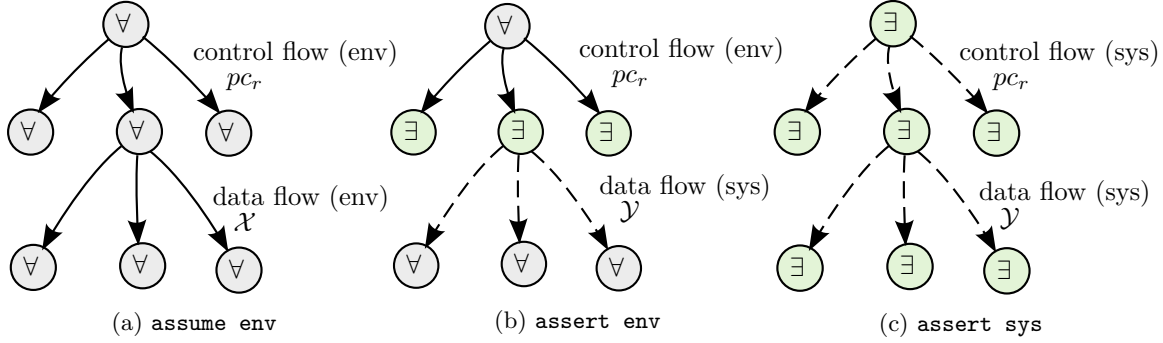


Figure 2.1: The keywords **assume**, **assert** declare the player whose data flow is constrained by the process program graph. The keywords **env**, **sys** declare the player that controls the program counter.

The notions of *control flow* and *data flow* are relevant to defining the behavior of a program graph. Control flow is the traversal of edges, whereas data flow is the behavior of program variables along this traversal. A *program counter* variable pc_r that stores the current node in V_r suffices to define control flow (together with a key variable to distinguish edges with identical source and target nodes). The environment (system) controls the flow of variables in \mathcal{X} (\mathcal{Y}).

A natural question to ask is who controls the program counter. Another question is whose data flow is constrained by the program graph P_r . If the same player controls the program counter and is constrained by P_r , then the program graph describes an automaton for the behavior of that player only. If one player controls the program counter, and the opponent reacts by choosing a compliant data flow, then the program graph describes a game.

As an example, suppose that the environment controls the program counter, and the system the data flow. At each node, the environment can pick any successor node, and the system must react by selecting a data flow compatible with the program statement that labels the edge that the environment picked. So *paths* in this program graph are *universally* quantified, whereas data flow is *existentially* quantified. This case is schematically depicted in Fig. 2.1b, and can be used to represent a verification problem Section 2.7.

The notion of path quantification corresponds to universal and existential nodes in alternating tree automata [48, 49, 50, 51, 52], although the program graphs presented here differ in how edges are labeled. If paths are universally quantified, then control flow nondeterminism is known as *demonic*, p.85 [53], [54], otherwise as *angelic* [55, 56, 57] The execution of such a program graph induces a *game graph* [58, 25].

2.2.1 Syntax

Program graphs are declared with the `proctype` keyword of `Promela` followed by statements enclosed in braces. We use the keyword **assume** (**assert**) to declare that a program graph constrains the environment’s (system’s) data flow. These keywords are common in theorem proving and program verification languages [59]. We use the keyword **env** (**sys**) to declare that the environment (system) controls the program counter pc_r . If the keywords are omitted, the control and data flow context are controlled by the system.

In Mealy games, the combination **assume sys** requires the “previous” operator of past LTL. Although reducible to GR(1) [4], past LTL is not supported by the `slugs` synthesizer used as backend by our implementation. So **assume sys** is currently not supported. We will call program graphs *processes*, noting that these processes have full information about each other, so they correspond to centralized synthesis, not distributed. The program counter *owner* is the player that controls variable pc_r . The *process player* is the player constrained by the program graph.

The example of Fig. 2.1b corresponds to

```

assert env proctype foo(){
    bit x;
    env bit y;
    ...
}

```

The default owner of variables declared in a `proctype` context is the player that controls the data flow. So in the previous code example, the variable `x` is controlled by the system, because of the keyword `assert`. The owner of local environment variables in an assertion needs to be explicitly declared, as shown for variable `y`.

2.2.2 Example

For example, the specification in Listing 2.1 defines a game between two players: the Bunny, and Taz, that move simultaneously. The Bunny must reach the carrot, without moving through a cell that Taz is in (`assert !t1`). Taz can only move between $xt \in \{1, 2\}$, has to keep visiting the lower row, and cannot go to cell (1,2) (`assume !t1`). Taz can move diagonally, but the Bunny only vertically or horizontally. Both players have an option to stay still (`skip`). Note that xt is a declarative variable.

Each process defines a program graph that constrains the respective player. In this example, each program counter is controlled by the corresponding constrained player. The repetition constructs (`do`) define mutually exclusive alternatives that the player has an option to select from in each turn.

Note how nondeterminism in process `taz` is demonic (universally quantified), whereas in `bunny` angelic (existentially quantified), i.e., the design freedom given to the synthesis tool. Each player has full information about all variables in the game, both local, as well as global. The solution is a strategy represented as a Mealy transducer [26] that the Bunny can use to win the game.

Note that the specification corresponds to a synchronous time interpretation, i.e., that the two players move simultaneously. The subformula $(x == xt) \ \&\& \ (y == yt)$ ensures that the Bunny and Taz are not in the same cell at the same turn of the game. It does *not* specify that the Bunny should avoid Taz moving to its cell in the *next* turn of the game.

Suppose that we desire to represent a game that in reality is played in alternating moves (like chess). Then the specification must be modified to ensure that Taz cannot reach the Bunny in its next move, i.e., before the Bunny plays its next turn. This requires using past LTL to refer to the last position of the Bunny, i.e., $\Box \neg((x = x_t) \wedge (y = y_t)) \wedge \Box \neg((\ominus x = x_t) \wedge (\ominus y = y_t))$. Past LTL is currently not supported, but is considered for implementation.

2.3 Expressions

In declarative languages like TLA, SMV, and the input syntax of GR(1) synthesis tools, the temporal operator “next” can be used to refer to values that variables will take in the next time step. Similarly, open `Promela` expressions can contain variable values at the next time step, extending the expressions of pure `Promela`. In temporal logic, the value of a variable x at the next time step is denoted by¹ $\bigcirc x$. In open `Promela` syntax, the value $\bigcirc x$ is denoted by priming the variable as x' . The prime is used as syntax for the “next” operator in TLA and for input to several synthesis tools. For readability, in the following we use primed variables also inside equations, so x' denotes $\bigcirc x$.

2.3.1 Functions, predicates, and connectives

Expressions are built from [61]:

- nullary function variables (integer-valued variables)
- unary function constants: -
- binary function constants: +, -, *, /, %, &, |, ^
- nullary predicate variables (Boolean-valued variables)
- binary predicate constants: <, <=, ==, !=, >=, > (<, ≤, =, ≠, ≥, >)
- nullary formula connectives `false`, `true` (\perp , \top)

¹There are two “next” operators, a function and a connective. By abusing notation, the symbol \bigcirc denotes the unary temporal connective “next” $\bigcirc : \text{wff} \rightarrow \text{wff}$ if x is a propositional variable, or the function “next” $\bigcirc : \mathcal{D} \rightarrow \mathcal{D}$ if x is a function variable. These are operators of unquantified first-order temporal logic [60].

Listing 2.1: The code that describes the example shown in Fig. 2.2.

```

1 free env int(1, 2) xt;
2 env int(0, 2) yt;
3
4 assume active env proctype taz() {
5     do
6         :: yt = yt - 1
7         :: yt = yt + 1
8         :: skip
9     od
10 }
11
12 assume ltl {
13     [] ((xt != 1) || (yt != 2)) &&
14     [] <> (yt == 0) }
15
16 sys int(0, 3) x;
17 sys int(0, 2) y;
18
19 assert active sys proctype bunny() {
20     do
21         :: x = x - 1
22         :: x = x + 1
23         :: y = y - 1
24         :: y = y + 1
25         :: skip
26     od
27 }
28
29 assert ltl {
30     [] ! ((x == xt) && (y == yt)) &&
31     [] <> ((x == 3) && (y == 1)) }

```

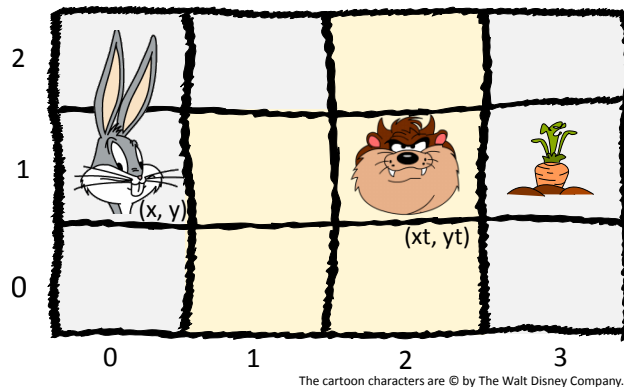


Figure 2.2: Synchronous game between two adversaries.

Table 2.1: Variables.

arity	Function	Predicate	Connective
0	bit, byte, int, short	bool	

Table 2.2: Constants.

arity	Function	Predicate	Connective
0	#define	#define	false, true
1	- (uminus)		!, ', [], <>
2	+, -, *, /, %, &, , ^	<, <=, ==, >=, >, !=	&&, , -, <->, U, W, R
3	ite(bool, num, num)		ite(bool, bool, bool),

- unary formula connectives ! (\neg)
- unary temporal formula connective ', [], <> ($\bigcirc, \square, \diamond$)
- binary formula connectives &&, ||, -, <->, U, W, R ($\wedge, \vee, \rightarrow, \leftrightarrow$)

The prime is used as postfix syntax for the temporal operator “next”. For example, x' is the value of variable x at the next time step. There are also temporal logic formulae (with “always”, “eventually”, and “until”) that can appear in `ltl` blocks, but here we focus on expressions that can appear as statements in a process.

2.3.2 State predicates and actions

Following TLA, we call *state predicate* an expression that does not contain primed variables [11]. Note that a state predicate does not contain temporal operators, so it has no side effects. In pure `Promela`, all expressions are state predicates.

An *action* is an expression that contains primed variables [11]. If controlled variables appear primed in an action, then the action can be regarded as generalization of an assignment, as described next.

2.3.3 Deconstraining

By default, imperative variables are constrained to remain unchanged. If any assumption (assertion) process executes a statement that contains a primed environment (system) variable, then that variable is *not* constrained to remain unchanged in that time step. Such a variable is called *deconstrained*.

For example, suppose that the following statements appear in an assertion:

```
sys bit x = 0;
x == 0;
x' == 1 - y
```

When the statement `x == 0` is executed, the variable x is constrained by $x' = x$. But when the statement `x' == 1 - y` is executed, the synthesizer is allowed to pick the next value of x as needed, in order to satisfy the equality $x' = 1 - y$.

Note that if a statement in an assumption (assertion) process contains primed imperative system (environment) variables, these are *not* deconstrained. The reason is that assumptions (assertions) are relevant only to the environment’s (system’s) data flow.

Moreover, primed environment variables *can* appear in system processes, because those are *past* values that the environment just selected. In a Mealy game, primed system variables cannot appear in an assumption process, because they are *future* values that the system has not yet selected (if the game is interpreted using an alternating time refinement).

For example, suppose that the environment controls variable x and the system controls y , and the following code appears in an assertion:

```
env int(0, 10) x;
sys int(0, 10) y;
(x' == 5) && (y' == 3)
```

The above expression is interpreted as the formula

$$(x' = 5) \wedge (y' = 3). \quad (2.1)$$

Note that the default constraint $y' = y$ has not been applied, otherwise the action would be satisfiable only if variable y already equals 3. Also, no constraint resulted for variable x , because an assertion does not constrain the environment's data flow.

In implementation, the imperative variables are constrained separately, by collecting all the edges that deconstrain them, as discussed in Chapter 3.

In some cases, we may desire to deconstrain an imperative variable, without adding any new constraints on their next value, leaving more freedom for the synthesizer to select the next value. We use the question mark as syntax for that purpose, as $x' == ?$, to remove the default imperative constraint for that time step.

Variables cannot appear multiply primed when using GR(1) synthesis, because nested next operators cannot appear in safety clauses of a GR(1) formula. If either a synthesis tool for full LTL is used [34, 36, 37, 38], bounded synthesis [62, 35], or the LTL formula that the program is translated to is converted, if possible, to GR(1) [14], then multiply primed variables can be used, but the complexity is less favorable, if not impractical.

Primed variables in a Mealy game

In this section, we describe the meaning of primed variables for a Mealy game. The game is alternating. If interpreted with a time refinement to alternating time, then a two-player alternating (turn-based) game is played between an output/input Moore machine and an input/output Mealy machine.

The system (Mealy transducer) looks at the current environment input that reacts to primed environment variables from \mathcal{X}' , and to the *previous* environment input (to which it has reacted in the previous turn of the game) by unprimed environment variables from \mathcal{X} . The environment refers to the current system output to which it reacts by using unprimed system variables from \mathcal{Y} . In order to refer to the last system output, the environment must use the unary temporal operator “previous” of past LTL.

The past fragment of LTL can be used in GR(1), as proved in [4]. However, we do not elaborate on that here. When supported by the implementation, syntax for a “previous” operator \ominus may be introduced in the future.

Nested “next”

If a GR(1) synthesis algorithm is used, then nested next operators cannot appear inside expressions (so no compound priming). Although such a formula can be expressed in the GR(1) fragment [14], this conversion is not implemented. An alternative is to use a synthesis algorithm for full LTL, as for example that implemented in Lily [34]. In that case, any temporal operators can, in theory, appear anywhere in the program's code, though this may have limited practical application.

2.3.4 Executability

A condition called *guard* is associated to each statement [40]. The process can execute a statement only if the guard evaluates to true. If a process currently has no executable statement, then it *blocks*. If no process has an executable statement, then the system will *deadlock*, and lose the game.

The guard of a state predicate is itself, as in PROMELA. Unprimed variables correspond to the last valuation, so the value of a state predicate does not depend on how the system or environment choose to play in the current turn of the game.

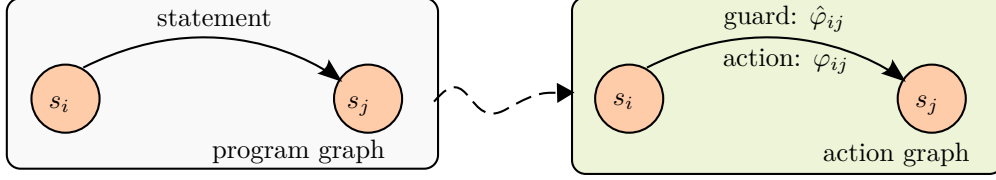


Figure 2.3: Transformation from a graph whose edges are labeled by program statements, to a graph whose edges are labeled by action formulae φ_{ij} (representing each statement) and guards $\hat{\varphi}_{ij}$, derived from the actions by using the syntactic approximation described in this section.

In contrast to a state predicate, an action includes primed variables, so it can constrain the values assigned next to variables. As noted earlier, actions can be considered as generalized assignments. Unlike assignments in pure PROMELA, actions are not always executable. An action is executable depending on an approximation of its satisfiability. The approximation is defined syntactically, as follows.

Syntactic approximation

Action realizability Consider an action in an assertion program graph. It can contain both unprimed variables from $\mathcal{X} \cup \mathcal{Y}$ and primed variables from $(\mathcal{X} \cup \mathcal{Y})'$. When the action is about to execute, the values of $\mathcal{X} \cup \mathcal{Y} \cup \mathcal{X}'$ are known. Only primed system variables from (\mathcal{Y}') remain free to be selected by the system (existentially quantified).

If the action *is* executable, then the process scheduler is allowed to select the process that will execute. The process must execute its next statement, otherwise the system deadlocks and loses the game. If, given a valuation in $\mathcal{X} \cup \mathcal{Y} \cup \mathcal{X}'$, the executable action is *not* satisfiable, then the system will deadlock. So unsatisfiability of an action can cause an undesired deadlock.

Let φ be an action and (x_A, y_A, x'_A) be a partial valuation at node A of the game graph associated to the Mealy game. If

$$\exists x \exists y \exists x' \exists y' \varphi \quad (2.2)$$

is true (φ is satisfiable), but

$$\exists y'_A \varphi[x_A/x, y_A/y, x'_A/x'] \quad (2.3)$$

is false (not satisfiable), then the past valuation of variables causes unsatisfiability. Therefore, it is not the system's "fault" that it cannot execute the action. We define executability of actions in an attempt to approximately distinguish these cases from those that φ is satisfiable.

In the following, we are interested in satisfiability of the formula

$$\varphi|_A \triangleq \varphi[x_A/x, y_A/y, x'_A/x']$$

(substitute x_A for x etc.), i.e., given a partial valuation that fixes values for x_A, y_A, y'_A . Note that y' are the only free variables in $\varphi|_A$.

We over-approximate the set of satisfying assignments of $\varphi|_A$, by defining an approximation $\hat{\varphi}$ that does not depend on \mathcal{Y}' . The approximation $\hat{\varphi}$ is defined later as a formula over the variables $\mathcal{X} \cup \mathcal{Y} \cup \mathcal{X}'$. Let

$$\hat{\varphi}|_A \triangleq \hat{\varphi}[x_A/x, y_A/y, x'_A/x'].$$

The approximation formula $\hat{\varphi}$ does not contain any free variables. If formula $\varphi|_A$ is satisfiable, then its approximation $\hat{\varphi}|_A$ is satisfiable. If the approximation $\hat{\varphi}$ is unsatisfiable, then $\varphi|_A$ is unsatisfiable. In summary

$$\left((\exists y' \varphi|_A) \rightarrow \hat{\varphi}|_A \right) \wedge \left((\neg \hat{\varphi}|_A) \rightarrow \neg \exists y' \varphi|_A \right) \quad (2.4)$$

Definition 1 Given the valuation (x_A, y_A, x'_A) , the action φ is executable if the approximation $\hat{\varphi}$ evaluates to \top .

There are three possible cases:

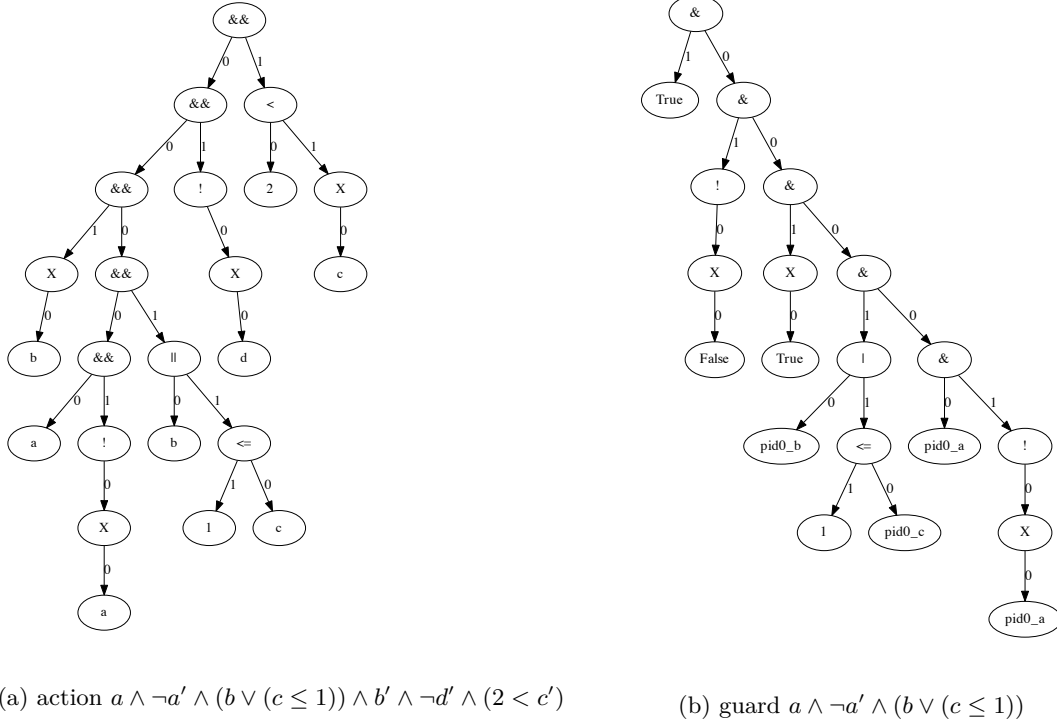


Figure 2.4: The syntactic approximation that defines guards for actions, i.e., for expressions that contain primed variables. The implementation counts negations while traversing the abstract syntax tree of the action.

1. The action $\varphi|_A$ is satisfiable, implying that the approximation $\hat{\varphi}|_A = \top$, so φ is executable. In this case, there exists some valuation over \mathcal{Y}' that the system can select as a reaction.
2. The action $\varphi|_A$ is unsatisfiable. In this case, there exists no reaction over \mathcal{Y}' that the system can select to satisfy φ given the state (x_A, y_A, x'_A) . There are two possible cases, depending on whether the approximation is also unsatisfiable:
 - (a) The approximation $\hat{\varphi}|_A = \perp$, so action φ is not executable. Since the action ϕ is not executable, the system cannot be forced to execute and fail to satisfy the action φ that is unsatisfiable at state (x_A, y_A, x'_A) .
 - (b) The approximation $\hat{\varphi}|_A = \top$, so action φ is executable. In this case, the process scheduler (controlled by the environment) can request that the system execute action φ . The system will fail to find a satisfying assignment over \mathcal{Y}' and thus deadlock, and lose the game.

The last case results due to the over-approximation of satisfiability. It can be the case only when the system can rename the syntactic occurrences of variables in \mathcal{Y}' by introducing fresh variables that it would then assign freely. In other words, a syntactic conflict that involves primed system variables does not discharge the system of its obligation to execute. So in such a case, the system can be forced to deadlock, exposing the logical inconsistency. Conversely, if the action remains unexecutable because of the past valuation, then the system can still be forced to block, in case it has no other process to execute.

Over-approximation The over-approximation $\hat{\varphi}$ is defined as follows. Consider the action φ in negation normal form (NNF). System variables can be Boolean or integer valued. We first define syntactically the approximation for Boolean variables, then for integer variables.

Each syntactic occurrence of a primed Boolean system variable $r' \in \mathcal{Y}'$ is replaced with a Boolean connective from $\{\perp, \top\}$. A variable in the scope of an even number of negations is replaced by \top , and a variable in the scope of an odd number of negations by \perp . For example, let a be a propositional environment variable, b, d propositional system variables and c an integer system variable ($\mathcal{X} \triangleq \{a\}$ and $\mathcal{Y} \triangleq \{b, c\}$). The formula

$$\varphi \triangleq a \wedge \neg a' \wedge (b \vee (c \leq 1)) \wedge b' \wedge \neg d' \wedge (2 < c')$$

is approximated by the formula

$$\begin{aligned} \hat{\varphi} &\triangleq a \wedge \neg a' \wedge (b \vee (c \leq 1)) \wedge \top \wedge \neg \perp \wedge \top \\ &= a \wedge \neg a' \wedge (b \vee (c \leq 1)) \wedge \top \wedge \top \wedge \top \\ &= a \wedge \neg a' \wedge (b \vee (c \leq 1)). \end{aligned}$$

The syntax trees of these two expressions are shown in Fig. 2.4. Flat variable names appear in the guard tree, because this is part of the logic formula that is provided to the synthesizer (after modular arithmetic is encoded in bitfield logic, as described in Section 4.2). Syntactic simplifications have not been applied to the trees shown.

Observe that there are four “next” operators (denoted by \mathbf{X}). Of these, three appear in arithmetic context ($\mathbf{X} \ a, \mathbf{X} \ b, \mathbf{X} \ d$) and one in propositional context ($2 < (\mathbf{X} \ c)$). The controlled variable identifiers b and d are substituted with \top and \perp , respectively, because b appears in positive context and d in negative context. The reference to the next value of variable c appears in the predicate “ $<$ ” that appears in positive context. So the expression $(2 < c')$ is replaced by \top .

The resulting guard formula is equisatisfiable with the formula obtained from φ , after replacing each syntactic occurrence of a primed system variable with a fresh variable. The models of $\hat{\varphi}$ form a superset of the models of φ .

So if $\hat{\varphi}$ remains unsatisfiable, then the system is not “responsible” for this lack of satisfiability, in the sense that even if it was allowed to rewrite the syntactic occurrences of primed variables that the system controls so as to decouple different occurrences of the same symbol, it cannot make the formula satisfiable. In this case, the approximation $\hat{\varphi}|_A$ evaluates to \perp , hence the action is not executable.

As noted previously, it may be the case that $\hat{\varphi}|_A = \top$, but $\varphi|_A$ is unsatisfiable. For example, the action $\varphi \triangleq y' \wedge \neg y'$ is unsatisfiable, but its approximation is $\hat{\varphi} \triangleq \top \wedge \neg \perp = \top$, so φ is always executable.

The arithmetic predicates $<, <=, ==, >=, >, !=$ occur either in a positive or negative context in NNF. Any arithmetic predicate that contains primed integer system variables is replaced by \top or \perp , depending on whether it appears in a positive or negative context, respectively. In implementation, bringing the formula to NNF is not necessary. Instead, it suffices to count the occurrences of negation from the abstract syntax tree root to the terminal of interest.

Alternatives

An alternative would be to define executability of constraints based on exact satisfiability. However, computing guards for such a notion of executability requires solving, for each statement, a quantified Boolean satisfiability problem. This is undesired, due to its computational cost. Although expressions written by humans are short for current SAT solvers, the presence of modular arithmetic can still pose difficulty to a SAT solver.

An `atomic` block can also be used to combine a Boolean expression that acts as guard, with an action. This block forms a compound statement that becomes executable, if, and only if, the guard evaluates to true. Atomic statements are discussed in Section 2.6.

2.4 Assignments

Each assignment statement is translated to a formula that has the same effect. In an assumption (assertion), only environment (system) variables can be assigned. The assigned variable is not constrained to remain unchanged, if it happens to be an imperative variable. All other non-free variables (that do not appear

inside the left operand) are subject to their default constraints. For example, the expression $\mathbf{x} = (\text{expr})$ is interpreted (roughly – see Chapter 3 for details) as the formula

$$(x' = (\text{expr})) \wedge \bigwedge_{y \in V_c \setminus \{x\}} (y' = y)$$

So the assignment operator overrides the default constraints that apply to imperative variables. The statement $\mathbf{x} = \mathbf{x}'$ is equivalent to the formula $x' = x'$, which is a tautology, so true irrespective of the value that variable x takes next.

As remarked in Section 2.3.3, primed system variables cannot appear in assumptions, so the following assumption statements

```
env bit y;
sys bit x;
x = y';
```

are not well-defined. In particular, y' is a value that the system will select, after the environment selects x' . Therefore, it can always observe what value the environment assigned to y' , and choose $x' = \neg y'$ to force the environment to lose the game.

Note that in an assumption (assertion), an imperative local variable of the system (environment) cannot be assigned to, because it is not visible to any assertion (assumption) scope of statements. Therefore, its default constraints imply that it will remain invariant. So such a local variable is of little use. Local opponent variables are useful only if declarative.

2.4.1 Truncation semantics

An assignment states that the next value of the assignee is equal to the value of the expression on the right hand side of the assignment symbol. If the assigned variable has mod-wrap semantics, then the value of the expression is truncated to the bitfield width of the assigned variable. As a result, an assignment to a variable with mod-wrap semantics cannot block.

If the assigned variable has saturation semantics, then the value of the expression is *not* truncated. So, an assignment to a variable with saturating semantics can block. Assignments are always executable, so if an assignment to a ranged integer blocks, then the corresponding player loses the game. The unconditional executability of assignments follows also from the executability of actions.

2.4.2 Equality operator in logic

There is no assignment in logic. So only the operator “=” appears in logic formulae. In the input syntax of a solver like `gr1c`, the code fragment $\mathbf{x} = \mathbf{x} + 1$ is interpreted as the formula $(x = x + 1)$, which is false, because the current value of variable x is different than the current value of x plus one. In contrast, in open `Promela` the token “=” signifies assignment, so the code fragment $\mathbf{x} = \mathbf{x} + 1$ is interpreted as the formula $(x' = x + 1)$, which is satisfiable.

2.5 Automata products

We introduce the keywords `async` and `sync` to define *asynchronous* and *synchronous*, respectively, products of program graphs. A product is defined by one of the keywords `async`, `sync`, followed by `proctype` declarations enclosed in braces `{...}`. The product is taken over the processes enclosed in braces. In general, each such block can contain `proctype` declarations, as well as other products of processes. We will refer to each encapsulated block or `proctype` contained inside a block as an *element* of the container block. So the syntax is of the form `async{ proctypes or sync products }`, and `sync{ proctypes }`.

A synchronous product blocks, if any program graph in it blocks. Otherwise, the scheduler can select the synchronous product for execution, so all program graphs must execute simultaneously. All program graphs in a synchronous product must constrain the same player.

```

1 sys proctype foo() {
2     ...
3 }
4
5 env proctype hoo() {
6     ...
7 }
8
9 sync{
10     sys proctype boo() {
11         ...
12     }
13 }
14     sys proctype woo() {
15         ...
16     }
17 }
18
19 sync{
20     env proctype qoo() {
21         ...
22     }
23
24     env proctype moo() {
25         ...
26     }
27 }

```

Figure 2.5: Products of program graphs.

An asynchronous product is the default top context, as in PROMELA. The scheduler can select any unblocked program graph (or `sync` product) inside an asynchronous product as the one that should execute next. The environment controls the scheduler, so if it picks a process whose guard evaluates to true, but its action is not satisfiable (Section 2.3), then the system fails to react and loses the game. If all program graphs of a player are blocked, then that player deadlocks and loses the game. In order to describe the process products discussed above, the `Promela` extended Backus-Naur form (EBNF) is augmented as described in Appendix A.1. An example of how asynchronous products can be nested is shown in Fig. 2.5.

Currently, asynchronous products inside synchronous products are not implemented. In the future, this can be obtained by adding the production rule `sync_unit ::= async`.

Note that nesting asynchronous (synchronous) products is equivalent to a single (flat) asynchronous (synchronous) product. Also, note that the asynchronous products between program graphs are in the context of full information, so the system is *not* asynchronous in the sense of [63, 64].

The asynchronous products defined above have universal quantification, because the environment schedules processes to execute. A possible extension is to define controlled asynchronous products, where the system selects the process that will execute. This is existential quantification of an asynchronous product. A prefix by `sys` can be used as syntax, i.e., `sys async{...}`. Currently, controlled asynchronous products are not implemented.

2.6 Atomic statements

The presence of two players and the target language (temporal logic) allow for a number of possible alternative interpretations for atomicity. We discuss them in this section.

Pure `Promela` includes two block statements that allow a process to request exclusive execution (of itself) by the process scheduler:

1. `atomic{...}`: the statements inside the block² that follows the keyword `atomic` are executed uninterruptedly, *unless* any one of them blocks. If a statement blocks, then atomicity is lost, allowing the scheduler to select the next process from among all the executable processes. Non-determinism is admissible inside an `atomic` statement³.

²In `Promela`, a *block* is defined as a set of statements contained in a pair of braces `{...}`.

³In verification, the system state is tracked during execution of an `atomic` block. The reason is to be “prepared” for losing atomicity, because, at that instance, the Kripke model’s state will need to be stored in the state space. Note that in `SPIN`, a `never` claim is *not* evaluated during the execution of an `atomic` statement. By modeling claims with a guarantee on a process controlled by the environment that observes the system execution, we can say that `never` claims in `SPIN` have the same semantics as global preemption, which is discussed later in this section.

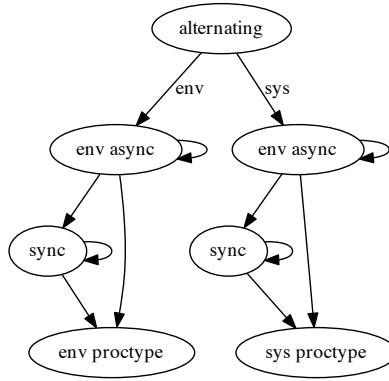


Figure 2.6: Nesting of synchronous and asynchronous process and player products. The prefix “env” signifies that the asynchronous product is universally quantified.

2. `d_step{...}`: the statements inside the block that follows the `d_step` keyword are executed strictly uninterruptedly. No statement is allowed to block. In other words, a `d_step` cannot lose atomicity. If it blocks, then an error is raised. Non-determinism is not interpreted as such, but resolved arbitrarily once, and that choice remains fixed over the entire state space (as obtained by unfolding by the semantics engine).

2.6.1 Controllability of process execution

First, we observe that in pure *Promela* an atomic statement prevents *all* other (executable) processes from executing. In open *Promela* though, there are two players that in turn execute steps of the processes that they control. So an atomic statement can preempt either:

- all other processes of the player that controls the process that requests atomic execution, or
- all processes, of both players.

The first alternative models time consumed by the system to execute atomic execution, by allowing the environment to keep playing against it in turns. This models the situation where the system can request atomic execution among its own processes, but not from the environment. In other words, the environment acts in a strictly adversarial way, and we assume that the atomic block cannot be ensured to execute in alternation with the environment.

The second alternative models system actions that can either

- be ensured to execute in negligible time⁴, within one turn of the system player, or
- a not totally adversarial environment, that cooperates slightly, by respecting preemption by system processes.

A unifying viewpoint is to consider these two cases of preemption power as another question of controllability:

- Preemption of system processes is controlled by the system in the first variant, but preemption of environment processes is uncontrolled by the system.
- Preemption of both environment and system processes is controlled by the system in the second variant.

⁴For example, due to time-scale separation between the system and its environment, similarly to the synchronous hypothesis [65].

The syntax `atomic(sys)...` can be used to define that only system processes are to be preempted. An `atomic...` block results in preemption of all processes, reducing the likelihood of confusion with pure `Promela`. Currently, preemption of both system and environment processes is implemented, but the syntax `atomic(sys)...` is a direct extension.

2.6.2 Expressing atomicity in temporal logic

An `atomic` statement can be represented in temporal logic by introducing auxiliary variables. Formally, atomicity is represented by Eqs. (3.5), (3.20), (3.21), (3.26) and (3.27) in Chapter 3, with some details discussed in Section 3.5. Here, we give an informal overview of how atomicity is represented in logic.

Consider a set of processes that execute asynchronously, and suppose that the scheduler selects a single process to execute next. The process executes its next statement, and can request to execute its next statement without interleaving with any other process, i.e., *exclusively*. We will call this process a *requestor*. Such a request is made when, in the program graph, the target node of the executed statement is in *atomic* context (Eq. (3.5)).

If the next statement is executable, then the scheduler grants the request for exclusive execution by selecting the requestor as the process that will execute next (Eq. (3.20)). Otherwise (i.e., when the next statement blocks), the scheduler selects another process to execute, one that is not blocked. The second case is called *loss of atomicity*.

In syntax, an `atomic` block defines nodes of the program graph that are in atomic context. If none of these statements blocks, then the block is executed uninterruptedly, without interleaving with other other processes (provided no statement is a `goto` with target outside the `atomic` block). Note that an `atomic` block cannot appear in a process that is part of a synchronous product.

Note that a request for exclusive execution is made when the target node is in atomic context. This can be the case when either the process enters atomic context, or when the process is already inside atomic context. A process can enter atomic context by either executing a statement that is syntactically the first in an `atomic` block, or by traversing a `goto` that jumps from outside, to inside an `atomic` block, or by resuming execution after losing atomicity.

It is interesting to note that an `atomic` context does *not* necessarily comprise of a *single* syntactic `atomic` context, because `goto` statements can jump to another syntactic `atomic` context, from within the current one, without interruption of atomic execution, as long as the target node is in the interior of an `atomic` block.

The requestor sets an auxiliary integer variable ex_s to notify the scheduler of its identity. Upon entering atomic context, the requestor sets $ex_s = j$, where j is its identity (pid). If the requestor wants to preempt processes of both players, then it sets the propositional variable pm_s , in addition to $ex_s = j$. The scheduler reacts by selecting j as the next pid to execute, or not, in case j is blocked.

A process sets ex_s to its last value if it does not request exclusive execution. The last value is selected to not correspond to an existing process (so if there are n processes, then ex_s can take $n + 1$ values). The last value is used in order to avoid introducing additional auxiliary Boolean variables, thus economizing on bits for representing the problem. Extending the domain of ex_s by one increases the number of bits required to represent it *only* if the number n is already a power of two. Otherwise, no new BDD variables are introduced. For example, if $n = 6$, then the domains:

- $\{0, 1, \dots, n - 1\} = \{0, 1, \dots, 5\}$ and
- $\{0, 1, \dots, n - 1, n\} = \{0, 1, \dots, 6\}$

both require 3 bits to represent them.

Notice that the environment controls the scheduler, so it can serve requests only in the next turn of the game. So if the system's process j requests atomic execution at turn k , then the scheduler responds in turn $k + 1$, by pausing the environment and selecting process j to execute in turn $k + 1$. The environment executes as normal in turn k , when the request was issued. This conforms to the order of play, i.e., the environment in turn k executed *before* the system issued its request. So indeed, this does not interleave with the statements in the atomic block.

Defining a translation of `atomic` blocks from `Promela` to temporal logic extends previous work that did not handle iteration constructs, loss of atomicity, nor non-determinism inside atomic blocks [66, 67, 68].

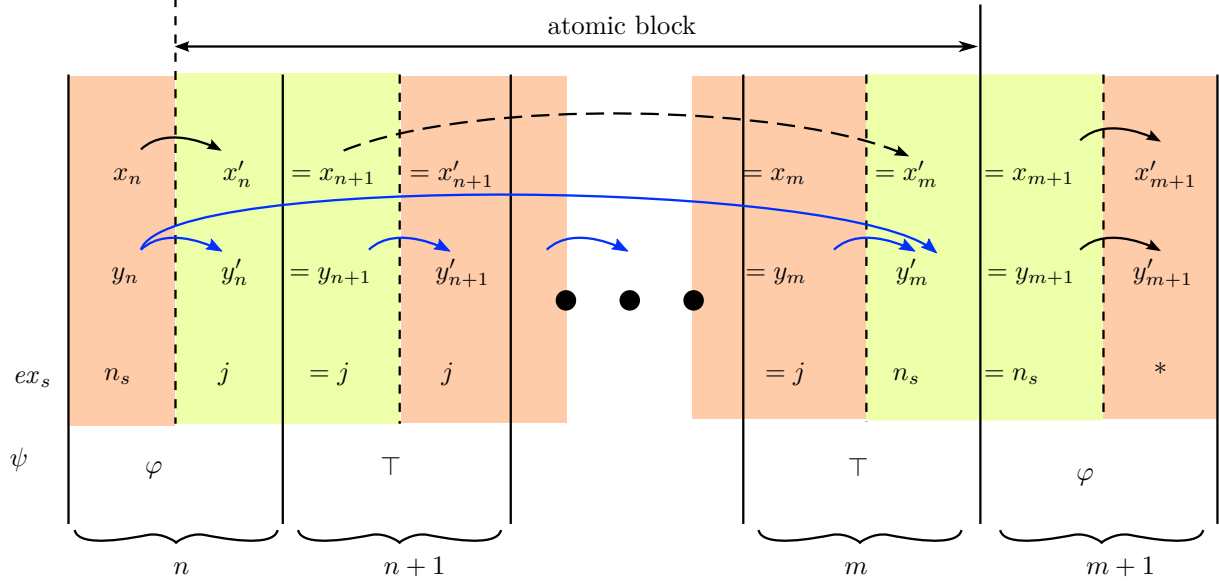


Figure 2.7: Timing diagram of atomic execution. The variables x are controlled by other processes and y by the atomically executing system process. The variables x “freeze” ($x'_{n+1} = x_{n+1}$) after atomic execution starts at time n . In addition, the formula $\psi \triangleq \varphi \vee (ps'_s = ex_s < n_s)$ deactivates the property φ , so that it won't observe any intermediate states produced during atomic execution.

2.6.3 Visibility to properties

If atomic statements execute over multiple time steps in temporal logic, then the intermediate states become visible in the game. A primary purpose of atomic transitions is to model changes of state that are ensured to happen indivisibly in implementation. Therefore, the intermediate states are a modeling artifact, so not relevant to LTL properties or other processes. For this reason, specification properties are not intended to take into account the states produced *during* atomic execution. In SPIN, these states are ignored in an atomic context, unless atomicity is lost (the last state produced becomes visible to LTL properties and other processes).

Consider an LTL safety property φ defined by an `ltl` block. For the reasons just described, the property needs to be “deactivated” if two conditions hold: some process just requested atomic execution, *and* the scheduler grants the request, i.e., the requestor did not blocked and lost atomicity. A property that does not contain “next” operators can be deactivated by disjoining it with these two conditions, as in Eqs. (3.26) and (3.27) that are introduced in Section 3.5. Fig. 2.7 shows how other processes “freeze”, and the property ψ is deactivated during atomic execution. Deactivation is necessary for safety properties, but not relevant to recurrence properties, because they are pure liveness properties.

Note that a variation is possible, by allowing a property to observe the intermediate states during atomic execution. Such a property represents a specification for the *internal* operation of that process, but not with respect to other processes. The implementation exposes an option to make atomic execution visible to safety properties.

2.6.4 Stutter invariance

The temporal logic representation of atomicity outlined in Section 2.6.2 remains valid under conditions that are discussed in this section. There are three cases to consider:

1. safety GR(1) properties that the user has defined as `ltl` blocks,
2. asynchronously executing other processes,

3. synchronously⁵ executing other processes that are preempted, e.g., Büchi automata.

As discussed next, in order to ensure correctness of the representation of atomicity in logic, some of the above cannot contain the “next” operator. We need to consider two categories of GR(1) safety formulae: those that contain the “next” operator, and those that do not. Note that formulae that contain only primed variables are almost⁶ equivalent to their unprimed counterpart.

As shown in Fig. 2.7, there are two different time scales:

- asynchronous time, which includes the time steps: $\dots n - 1, n, m + 1, \dots$ (macro steps)
- game time, which includes all time steps: $\dots n - 1, n, n + 1, \dots, m, m + 1, \dots$ (micro steps).

Let us consider how these different time scales appear to the properties and processes listed earlier.

Asynchronous processes Suppose that system process j executes in turn n , and requests exclusive execution. Any other system process $j + 1$ that executes asynchronously with process j has executed its last statement in turn $n - 1$. It is currently paused. If the request is granted, then the earliest it may resume is at turn $m + 1$. Therefore, expressions in process $j + 1$ can refer to primed global variables without being affected by how atomic execution is represented in logic. In other words, asynchronous processes are insensitive to external time refinement, as that introduced by simulating atomic transitions in game time.

Safety GR(1) properties In contrast, a safety assertion must refer only to unprimed global variables, i.e., it cannot contain the operator “next”. Referring to primed global variables can lead to incorrect results. For example, in Fig. 2.7 a property φ that refers to x' and y' will apply to valuation (x_n, x'_n, y_n, y'_n) . However, this valuation does not correspond to two consecutive time steps of asynchronous time, because the value y'_n corresponds to an intermediate state during atomic execution. In asynchronous time, the successor state of (x_n, y_n) is (x_{m+1}, y_{m+1}) , i.e., after the atomic transition has been completed (or exited atomic context due to loss of atomicity). So the correct value is $y'_m = y_{m+1}$, upon completion of the atomic transition of process j (note that $x_n = x'_m = x_{m+1}$). Therefore, the correct valuation to apply property φ to is (x_n, x'_n, y_n, y'_m) . Alternatives that are compatible with properties that include “next” are described in Section 2.6.5.

Not all properties that contain “next” are stutter-sensitive [69]. In other words, other formulae must be in a syntactic subset of stutter-invariant properties, namely those that do not include the operator “next” for non-local variables.

Preempted synchronous processes Similar observations apply to a preempted process $j + 2$ that executes synchronously with process j . A deterministic Büchi automaton may be such a synchronous process (so the “program graph” counterpart of a GR(1) safety property). Processes that execute synchronously with process j cannot include primed global variables in their statements. This ensures that intermediate states of atomic execution are not visible to guards of statements in those other processes. It also avoids interference of actions in these processes with the atomically executing process (e.g., if y'_n in Fig. 2.7 appeared in an action of a synchronously executing process that executes its last statement at time n). Only local variables can appear primed in expressions, or assigned to. Local variables are valid, because they are visible only to process $j + 2$, so they do not introduce coupling between different processes. For example, the program counter is a local variable in this sense (with respect to other processes, excluding the scheduler that is aware of all data flow).

Statements in atomic context If a statement that is executed atomically refers to environment variables, then these are equal to x'_n , as shown in Fig. 2.7. Therefore, such references are to be understood as the last environment reaction. In other words, only primed environment variables in atomically executing statements correspond to the primed environment variables in asynchronous time. Unprimed environment variables correspond to unprimed environment variables in asynchronous time only for the statement that enters atomic atomic context.

⁵Nesting of asynchronous products inside synchronous products and inclusion of atomic statements are not implemented yet, but nevertheless discussed in the analysis.

⁶The primed version excludes the initial time.

Remarks In SPIN, the operator “next” is discouraged, because it reduces the applicability of partial-order reduction. Another reason to avoid stutter-sensitivity in component specifications is to ensure that the specifications remain invariant under time refinement [11].

Some limitations to the use of atomic blocks have been described in this section. Note that these are not fundamental limitations of open **Promela** itself, but rather of the translation to temporal logic formulae. Instead, an enumerative synthesizer or model checker can interpret open **Promela** directly, without the need of an intermediate translation to formulae. The limitations about **atomic** statements that are described above do not apply to an enumerative approach. Also, they do not apply to a less efficient symbolic execution approach of Section 2.6.5 that would employ copies of variables as solver memory, in order to simulate correctly more use cases for atomic block usage.

2.6.5 Alternatives that accommodate stutter-sensitivity

There are two other options for representing atomic blocks in logic:

- Use auxiliary variables (like ex_s) to request from the process scheduler to preempt other processes.
- Use symbolic execution to derive a single action formula for the entire atomic block in temporal logic.

Duplicate variables All program variables can be duplicated and additional scheduling introduced to “forward” the values x_n, y_n to time $m + 1$ and use them as unprimed values for evaluating a property that contains both primed and unprimed variables, deactivating and activating the property accordingly. This approach increases the size of the state space exponentially in the number of variables. For this reason, it is not discussed further here.

Symbolic execution In the presence of stutter-sensitive properties, symbolic execution has to be used. Two limitations of symbolic execution are that:

- Non-determinism inside an **atomic** block can cause the formulae resulting from symbolic execution to grow large.
- An iteration construct inside an **atomic** block cannot be represented with symbolic execution if the number of iterations is not a priori bounded.

For example, the code fragment `atomic{ stmt0; stmt1; stmt2 }` would be translated to

$$\text{ite}(\hat{\varphi}_0, \text{ite}(\hat{\varphi}_1, \text{ite}(\hat{\varphi}_2, s[\varphi_0][\varphi_1][\varphi_2], s[\varphi_0][\varphi_1]), s[\varphi_0]), \text{no change}),$$

where $s[\varphi_0][\varphi_1]$ is postfix notation for the application of action φ_0 on state s , then action φ_1 on the result [11].

Instead of symbolic execution, the operator **&&** can be used to conjoin state predicates with expressions containing primed variables, creating arbitrary guarded commands. This avoids the need to write **atomic** blocks for expressing guarded reactions that involve assignments (a common use case). See Section 2.3.4 for the definition of executability for actions.

For atomic blocks that do not contain iteration statements, a possible compiler optimization is to check if any statement after the entry can block. If none of the interior statements can block, then symbolic execution can be applied without the need for **ite** operators, because atomicity cannot be lost in this special case.

2.7 Verification in open Promela

Verification can be approached by either checking that a desired (“positive”) property holds for all possible computations, or that there does not exist a satisfying computation for an undesired (“negative”) property.

Validity If checking whether a property holds, we desire to check all possible control flow paths. So paths are universally quantified. In verification, expressions in a model do not include primed variables (primed variables result only indirectly from assignments). Therefore, each path corresponds to a unique data flow.

The data flow needs to be existentially quantified. The reason is that when no feasible data flow exists, then a deadlock has been found. In this case, existential quantification over an empty set of successors yields the desired result (false), whereas universal does not. In other words, if the environment controlled the data flow, then deadlock would result in trivial realizability, instead of unrealizability (which is the correct result). The combination `assert env proctype` can express this use case.

When using a GR(1) synthesizer, the desired property must be in GR(1), so a recurrence property.

Satisfiability If checking whether a property does not hold, then a counterexample can be found by quantifying control and data flow existentially (closed system synthesis). In this case, unrealizability implies that the (undesired) property does not hold. When using a GR(1) synthesizer, the undesired property must be in GR(1), thus of the form $\psi \triangleq \varphi_i \wedge \square\varphi_s \wedge \square\Diamond\varphi_r$, since there no environment variables or assumptions are defined. In other words, realizability implies that there exists a system that satisfies ψ . This is a counterexample that violates the desired property $\neg\psi = \neg\varphi_i \vee \Diamond\neg\varphi_s \vee \Diamond\square\neg\varphi_r$.

Otherwise, a either a generalized Rabin(1) or full LTL synthesizer have to be used, in order to check persistence as negative properties, thus recurrence as positive properties.

Chapter 3

Translation to logic

3.1 Notation

Each program graph is identified by a unique integer process id $r \in \mathbb{N}$. Each set of entities inside the topmost asynchronous product is identified by a group id $m \in \text{gids}(p) \subset \mathbb{N}$, where $p \in \{e, s\}$ denotes the player. Let $n_p \triangleq \max\{\text{gids}(p)\} + 1$. The set $\text{pids}(m) \subset \mathbb{N}$ contains the identities of the program graphs that belong to group m .

Each group id corresponds to either a program graph, or a synchronous product. Therefore, each program graph is associated to both a pid and a gid. All program graphs in a synchronous product have the same gid. The scheduler can request simultaneous execution of all the program graphs in a synchronous product by selecting the corresponding gid.

A *program counter* variable $pc_r \in N_r \subset \mathbb{N}$ stores the current node of the program multi-digraph (N_r, E_r) . If the program graph is an assumption (assertion), then the variable pc_r is controlled by the environment (system). The set N_r of nodes is a contiguous subset of the natural numbers.

For each program graph r , the function $\text{key}(r)$ returns an auxiliary variable controlled by the player that controls pc_r . The key distinguishes multi-edges with same source and target nodes. The domain of variable $\text{key}(r)$ is sufficiently large to accommodate for the maximal multi-edge multiplicity of a multi-edge in any set of edges E_r . This number is typically small, and bounded by the maximal branching factor among program graphs.

The same key variable is reused for program graphs that execute asynchronously. Inside a synchronous product, there exist multiple program graphs that execute at the same time, so multiple key variables are needed in that case. For example, a synchronous product that contains three program graphs requires three key variables.

In addition, in order to avoid conflicts, there are three disjoint collections of variables that serve as keys: environment key variables for **assume env**, system key variables for **assert sys**, and environment key variables for **assert env**.

The environment variables $ps_p \in \mathbb{N}_{\leq n_p}$ represent the scheduler's choice of next program graph or synchronous product that will execute. The abbreviation "ps" stands for *process scheduler*. The system variable $ex_s \in \mathbb{N}_{\leq n_p}$ is used by a system process to request that it execute exclusively, and the system variable $pm_s \in \mathbb{B}$ to preempt the environment.

The indices i and j denote the source and target nodes, respectively, of an edge $(i, j, k) \in E_r$. For each edge $(i, j, k) \in E_r$, the function $\text{guard}(r, i, j, k)$ returns the guard formula defined in Section 2.3. Note that the value of $\text{key}(r)$ corresponds to k . Let the function

$$\text{inv}(x) \triangleq \begin{cases} x' = x, & \text{if } x \text{ is not an array} \\ \bigwedge_{l=0}^{\text{len}(x)-1} (x'_l = x_l), & \text{if } x \text{ is an array.} \end{cases} \quad (3.1)$$

where $\text{len}(x)$ is the length of the array variable x . Let $\text{ite}(x, y, z)$ denote the ternary conditional operator "if x , then y , else z ".

Each reference $\mathbf{a}[i]$ to an element in an array \mathbf{a} is encoded using the (appropriate) ternary conditional operator

$$\begin{aligned} & \text{ite}\left(i = \text{len}(a) - 1, a_{\text{len}(a)-1}, \right. \\ & \quad \text{ite}\left(i = \text{len}(a) - 2, a_{\text{len}(a)-2}, \right. \\ & \quad \quad \text{ite}\left(\dots, \right. \\ & \quad \quad \quad \left. \left. \text{ite}(i = 1, a_1, a_0)\right)\right) \end{aligned} \quad (3.2)$$

Array indexing errors are prevented by adding an additional safety assumption or assertion of the form

$$\square(0 \leq i \leq \text{len}(a) - 1) \quad (3.3)$$

A primed controlled array reference $\mathbf{a}[i]'$ or $\mathbf{a}[i']'$ is handled using the guard definition of Section 2.3.4. An uncontrolled array reference with a primed controlled index requires the following conversion to a disjunction as guard expression (the below is for an occurrence in the scope of an even number of negations)

$$\begin{aligned} \text{guard}(\mathbf{a}[i'] = w) &= \text{guard}((\exists j \in \mathbb{N}_{\text{len}(a)}. a_j = w) \wedge (i' = j)) \\ &= (\exists j \in \mathbb{N}_{\text{len}(a)}. a_j = w) \wedge \underbrace{\text{guard}(i' = j)}_{\top} = \bigvee_{j=0}^{j=\text{len}(a)-1} (a_j = w). \end{aligned} \quad (3.4)$$

Unless noted otherwise, the formulae that follow are safety assumptions or assertions. In the following, we use “=”, with the understanding that “ \leftrightarrow ” must be used instead, if the variables involved are propositional.

3.2 Control and data flow

Each process is parsed and converted to a program graph, according to the definitions of control flow constructs in PROMELA [39]. To each edge corresponds an action formula φ . The statements in each program graph are translated to logic formulae and guards (Section 2.3). Each graph, now labeled with formulae, is flattened to an equivalent formula of transitions (Eq. (3.6)), similarly to symbolic model checking [70]. This results in a data flow constraint (Eq. (3.7)).

$$\text{exclusive}(p, m, i, j, k) \triangleq \begin{cases} (ex'_s = m) \wedge pm'_s, & \text{if } p = s \text{ and } j \text{ in atomic context} \\ (ex'_s = n_s) \wedge \neg pm'_s, & \text{if } p = s \text{ and } j \text{ not in atomic context} \\ \top, & \text{if } p = e \end{cases} \quad (3.5)$$

$$\text{trans}(p, m, r) \triangleq \bigwedge_{i \in N_r} \left((pc_r = i) \rightarrow \bigvee_{(i,j,k) \in E_r} \left(\varphi_{r,i,j,k} \wedge (pc'_r = j) \wedge (\text{key}(r)' = k) \wedge \text{exclusive}(p, m, i, j, k) \right) \right) \quad (3.6)$$

We consider the case of static process instantiation. Dynamic process instantiation is currently not supported.

The following formula constrains the data flow of player p to comply to program graph r .

$$\text{dataflow}(p) \triangleq \bigwedge_{m \in \text{gids}(p)} \bigwedge_{r \in \text{pids}(m)} \left(\text{ite}(ps'_p = m, \text{trans}(p, m, r), \text{inv}(pc_r)) \right) \quad (3.7)$$

So the scheduler selects group m to execute by setting ps'_p to m . If program graph r is in group m , then it must execute a transition in $\text{trans}(p, m, r)$. Those program graphs that are not selected to execute must remain idle, i.e., $\text{inv}(pc_r)$. Note that $\text{inv}(pc_r)$ does not suffice to distinguish self-loops from an idle program graph. In addition, the variables in the scope of an idle program graph must remain unchanged, as ensured by Eqs. (3.12) and (3.15).

The data flow constraint of Eq. (3.7) suffices for graphs with same path and data flow quantification (`assume env` and `assert sys`), because it ensures that the player that controls pc_r and $\text{key}(r)$, and is constrained by action $\varphi_{r,i,j,k}$, will select appropriate values for them.

However, if the program graph is an assertion and the environment controls pc_r , then a separate control flow assumption is needed, Eq. (3.9). This assumption prevents the environment from selecting edges that have false guards. This case corresponds to program flow controlled by the environment (universal path quantification) and data flow controlled by the system (existential data quantification). Let `assert_env_pids` denote the assertions with pc_r and $\text{key}(r)$ controlled by the environment.

$$\begin{aligned} \text{guards}(r) \triangleq & \bigwedge_{i \in N_r} \left((pc_r = i) \rightarrow \left((\text{key}(r)' < \max_{(i,j,k) \in E_r} \{\text{multiplicity}(i,j,E_r)\}) \wedge \right. \right. \\ & \left. \left. \bigvee_{(i,j,k) \in E_r} \left((pc'_r = j) \wedge (\text{key}(r)' = k) \wedge \text{guard}(r,i,j,k) \right) \right) \right) \end{aligned} \quad (3.8)$$

In Eq. (3.8), the target node is determined by pc'_m , because this is the *latest* input from the environment to the system. The constraint $\text{key}(r)' < \max_{(i,j,k) \in E_r} \{\text{multiplicity}(i,j,E_r)\}$ prevents the environment from choosing a key value that exceeds the multiplicity of any outgoing multi-edge at node i . This ensures that the selected key value will correspond to an *existing* edge. The domain of $\text{key}(r)$ can be larger, in order to accommodate for the multiplicity of outgoing edges at other nodes.

$$\text{control_flow}_s \triangleq \bigwedge_{m \in \text{gids}(s)} \bigwedge_{r \in \text{pids}(m) \cap \text{assert_env_pids}} \left((ps'_s = m) \rightarrow \text{guards}(r) \right) \quad (3.9)$$

This control flow assumption also ensures that the environment selects an existing edge in the program graph, by selecting suitable values for the variables pc_r and key . For example, consider a system process with $\text{pid} = m$ that is currently at $pc_m = 0$ and has two outgoing edges leading to $pc_m = 1$ and $pc_m = 2$. Suppose that the edge $(0,1)$ is currently executable, but the other edge $(1,2)$ is not. If the environment selects $ps_s = m$ to execute, but sets $pc'_m = 2$, then the system is blocked. This environment behavior violates the guards of system transitions.

3.2.1 Initial conditions

Each variable has a default initial value, depending on its semantics. If a value is assigned to the variable upon declaration, then that value overrides the default initial value. For example, `free bit x = 1`. By default,

- a free variable is not constrained initially,
- a Boolean (numerical) imperative variable is initially `false` (zero),
- an imperative ranged integer is initially equal to the minimal value in its range.

Each element of an array variable is initialized separately.

Let `assert_sys_pids`, `assume_env_pids`, `asser_env_pids` denote the process identities declared with the corresponding path and data flow quantification. The system should initially satisfy

$$(ex_s = n_s) \wedge \neg pm_s \wedge (\forall r \in \text{assert_sys_pids}. (pc_r = P_r.\text{root}) \wedge (\text{key}(r) = 0)). \quad (3.10)$$

The environment is assumed to initially satisfy

$$\begin{aligned} & (\forall r \in \text{assume_env_pids}. (pc_r = P_r.\text{root}) \wedge (\text{key}(r) = 0)) \wedge \\ & (\forall r \in \text{assert_env_pids}. (pc_r = P_r.\text{root}) \wedge (\text{key}(r) = 0)). \end{aligned} \quad (3.11)$$

3.3 Invariance of variables

Eq. (3.12) ensures that local declarative variables remain invariant if their scope doesn't execute.

$$\text{local_free}(p) \triangleq \bigwedge_{m \in \text{gids}(p)} \bigwedge_{r \in \text{pids}(m)} ((ps'_p \neq m) \rightarrow \bigwedge_{x \in V_{f,p,r}} \text{inv}(x)) \quad (3.12)$$

where $V_{f,p,r}$ is the set of free variables that are controlled by player $p \in e, s$ and declared in the scope of program graph r .

An imperative variable x is constrained to remain unchanged *only* if none of the currently executed statements deconstrains x . Statements that deconstrain an imperative variable are

- an assignment to the variable, or
- an action that contains the primed value of the variable
- deconstrains the variable with $x = ?$.

Note that primed imperative variables that appear in an expression that comprises the right hand side of an assignment are not deconstrained (an assignment is understood as unidirectional, not as a balance equation).

It is possible that multiple statements are executed at the same time, by program graphs that are in the same synchronous product. So the constraints should be formulated at the scope that the variable x is defined, to account for all the program graphs that can refer to, and so also change, x . This is ensured by Eqs. (3.14) and (3.15). For brevity, in the following we use the auxiliary definition

$$\text{edge}(p, m, r, i, j, k) \triangleq (ps_p = m) \wedge (pc_r = i) \wedge (pc'_r = j) \wedge (\text{key}(r)' = k) \quad (3.13)$$

Eq. (3.14) ensures that primed references to elements in imperative arrays deconstrain only the referenced array element.

$$\text{array_inv}(p, m, r) \triangleq \bigwedge_{a \in \text{arrays}} \bigwedge_{(i,j,k) \in E_r} \left(\text{edge}(p, m, r, i, j, k) \rightarrow \bigwedge_{l=0}^{\text{len}(a)-1} \left(\left(\bigwedge_{e \in \text{idx}(i,l,k)} (e \neq l) \right) \rightarrow \text{inv}(a_l) \right) \right), \quad (3.14)$$

where $\text{idx}(i, l, k)$ is the set of expressions that appear as indices of references to elements of array a in action $\varphi_{r,i,l,k}$. This ensures that all array elements a_l other than those that appear primed in $\varphi_{r,i,l,k}$ remain invariant ($e \neq l$). The index expression e can be a variable, so its value is not fixed, thus conjunction over all array elements ($l \in \mathbb{N}_{<\text{len}(a)-1}$) is necessary.

Eq. (3.15) ensures that imperative variables remain invariant, unless a statement that deconstrains them (Section 2.3) is executed, by traversing an edge in $E_f(x)$.

$$\text{imperative_inv}(p, m, r) \triangleq \text{array_inv}(p, m, r) \wedge \bigwedge_{x \in V_{c,p,r}} \left(\text{inv}(x) \vee \bigvee_{(g,r,i,j,k) \in E_f(x)} \text{edge}(p, m, r, i, j, k) \right) \quad (3.15)$$

where $E_f(x)$ is the set of edges that deconstrain variable x , defined as

$$\begin{aligned} \text{is_free}(p, r, i, j, k) &\triangleq (x \in \text{Var}(\varphi_{r,i,j,k}) \cap V'_{p,c}) \\ E_f(x, p) &\triangleq \bigcup_{r \in \text{pids}(p)} \bigcup_{(i,j,k) \in E_r} \text{is_free}(p, r, i, j, k). \end{aligned} \quad (3.16)$$

Finally, a safety constraint is added for each ranged integer variable

$$\text{ranged_int}(p) \triangleq \bigwedge_{x \in \text{ranged_ints}(V_p)} \left(\min(\text{dom}(x)) \leq x \leq \max(\text{dom}(x)) \right), \quad (3.17)$$

where $\text{dom}(x)$ denotes the domain of variable x .

3.4 Process scheduler

Assume that the program comprises of an asynchronous product of synchronous products of sets of processes. Each set of processes contains one or more **proctype**. In pure **Promela** all sets are singletons, so the program is an asynchronous product of processes. Note that all processes inside a synchronous product constrain the same player (at least in the current formulation). Also, note that the grammar defined earlier does not allow asynchronous products to appear inside synchronous products. Let us assume for simplicity that the system does not control any asynchronous product. Nesting of asynchronous products is equivalent to one flat asynchronous product, and likewise for nesting of synchronous products. Therefore, we can always “flatten” the given program to an asynchronous product of synchronous products. Note that each **proctype** inside a synchronous product is limited to at most 1 active instance.

The scheduler cannot select for execution a blocked program graph (Eqs. (3.18) and (3.19)). This implies that a synchronous product blocks if at least one program graph in the product blocks.

$$\text{blocked}(r) \triangleq \bigvee_{i \in N_m} \left((pc_r = i) \wedge \bigwedge_{(j,k) \in \text{oute}(i, E_r)} \neg \text{guard}(r, i, j, k) \right) \quad (3.18)$$

$$\text{selectable} \triangleq \bigwedge_{p \in \{e, s\}} \bigwedge_{m \in \text{gids}(p)} \bigwedge_{r \in \text{pids}(m)} (\text{blocked}(r) \rightarrow (ps'_p \neq m)) \quad (3.19)$$

A system process can request to execute *atomically*, meaning that it will be granted uninterrupted execution, until it exits atomic context, because it either blocked (Eq. (3.20)) or reached non-atomic context (Eq. (3.5)).

$$\text{grant}_s \triangleq \bigwedge_{m \in \text{gids}(s)} \bigwedge_{r \in \text{pids}(m)} (((ex_s = m) \wedge \neg \text{blocked}(r)) \rightarrow (ps'_s = m)) \quad (3.20)$$

If no environment process is executing, then the environment has deadlocked and loses the game, unless the reason is that some system process currently paused the environment processes, in order to execute atomically (Eq. (3.21)).

$$\text{pause_env_if_req} \triangleq (ps'_e = n_e) \leftrightarrow (pm_s \wedge (ps'_s = ex_s < n_s)) \quad (3.21)$$

If all processes of player p have blocked, then the scheduler sets variable ps_p to its maximal value n_p (Eqs. (3.22) to (3.24)).

$$\text{sys_deadlock} \triangleq (ps'_s = n_s) \leftrightarrow \bigwedge_{m \in \text{gids}(s)} \bigvee_{r \in \text{pids}(m)} \text{blocked}(r) \quad (3.22)$$

$$\text{env_deadlock} \triangleq \left(\bigwedge_{m \in \text{gids}(e)} \bigvee_{r \in \text{pids}(m)} \text{blocked}(r) \right) \rightarrow (ps'_e = n_e) \quad (3.23)$$

$$\text{never_deadlock}_s \triangleq ps'_s \neq n_s \quad (3.24)$$

Labeled program statements induce labeled nodes in the program graph. If the label string contains the word “progress”, then a recurrence constraint is added to the liveness assumption or assertion, as in Eq. (3.25).

$$\text{progress}(p) \triangleq \bigwedge_{m \in \text{gids}(p)} \bigwedge_{r \in \text{pids}(m)} \bigwedge_{i \in N_{r, \text{progress}}} \Box \Diamond ((ps_p = m) \wedge (pc_r = i)) \quad (3.25)$$

A summary of the translation is given in Table 3.1, including the separation into assumption and assertion formulae.

3.5 Exclusive execution

If **atomic** blocks are included in system processes, then the following limitations apply, for the above translation to remain valid. All **ltl** properties and guards must not contain primed variables, thus be stutter invariant. Also, **atomic** blocks cannot appear inside synchronous products.

Table 3.1: Summary of translation to GR(1), separated into assumptions and assertions.

Equation	Assumption	Assertion	Purpose
Eq. (3.3)	if a is sys variable	if a is env variable	array index bound
Eq. (3.7)	for $p = e$	for $p = s$	data flow
Eq. (3.9)	yes		control flow of <code>assert env</code>
Eq. (3.10)		yes	sys init
Eq. (3.11)	yes		env init
Eq. (3.12)	for $p = e$	for $p = s$	local free var inv
Eq. (3.15)	for $p = e$	for $p = s$	imperative var default inv
Eq. (3.17)	for $p = e$	for $p = s$	ranged integer safety
Eq. (3.19)	yes		non-blocked processes
Eq. (3.20)	yes		grant if not blocked, else loses atomicity
Eq. (3.21)	yes		pause env only if preempted by sys
Eq. (3.24)		yes	sys paused means deadlocked
Eq. (3.22)	yes		pause sys iff all sys processes blocked
Eq. (3.23)	yes		pause env if all env processes blocked
Eq. (3.25)	for $p = e$	for $p = s$	progress labels
Eq. (3.26)	yes		mask env LTL
Eq. (3.27)		yes	mask sys LTL

3.5.1 Stuttering and visibility

Safety LTL formulae in `1t1` blocks are deactivated during atomic execution (Eqs. (3.26) and (3.27)), as for PROMELA. In other words, a safety requirement does not apply to atomic execution (the implementation exposes an option to make atomic execution visible to safety properties). Note that they are re-activated as soon as atomicity is lost.

$$\text{mask_env_1t1} \triangleq \bigwedge_{\psi \in \text{safety}(1\mathbf{t1},e)} \psi \vee (pm_s \wedge (ps'_s = ex_s < n_s)) \quad (3.26)$$

$$\text{mask_sys_1t1} \triangleq \bigwedge_{\psi \in \text{safety}(1\mathbf{t1},s)} \psi \vee (ps'_s = ex_s < n_s) \quad (3.27)$$

If $i = e$ ($i = s$), then the formula `safety(1t1, e)` denotes the safety assumptions (assertions) that are declared inside `1t1` blocks.

3.5.2 Unbounded loops

If unbounded loops appear inside atomic context, then there can be no liveness assumptions. The reason is that the system can “hide” in atomic execution forever, preventing the environment from satisfying its liveness assumptions, thus winning trivially. Requiring that the system infinitely often exit from its atomic execution is insufficient, because the system is discharged of the recurrence assertion if it can force the environment to violate the recurrence assumption. In order to avoid this, the environment recurrence goals must be disjoined with a persistence property to obtain $\Box\Diamond(\text{env goal}) \vee \Diamond\Box(\text{sys in atomic execution})$, a strong fairness formula. Persistence ($\Diamond\Box$) is outside GR(1).

Chapter 4

Implementation

The implementation is written in PYTHON and available [71, 72] under a BSD license. It comprises of a translator from open PROMELA to linear temporal logic with both (bounded) integer and Boolean variables, and an encoder of bounded integer arithmetic into propositional temporal logic. The result is used as input to the GR(1) synthesis tool `slugs` [33].

The front-end uses a lexer and parser written with PLY (Python lex-yacc) [73] that takes a `Promela` program source and returns an abstract syntax tree. The abstract syntax tree is then translated to a program graph. The program graph is a directed graph that has nodes that represent program states and edges labeled with program statements.

The program graph is flattened to logic in two passes by the compiler architecture as shown in Fig. 4.1. The first pass converts each statement that labels an edge to a guard formula and an action formula, producing a semi-symbolic representation. The second pass flattens the graph into a transition relation over primed and unprimed variables, introducing program counters and keys to distinguish multiedges. This flattening is applied to the program graph of each process to yield the formulae defined in Section 3.2.

These formulae are used to define the process scheduler, as described in Sections 3.4 and 3.5. Note that (orthogonally) the previous formulae are separated into an assumption clause and an assertion (guarantee) clause, as described in Table 3.1. The environment is obliged to satisfy the assumption and the system the assertion.

Up to this stage, the formulae are over both Boolean and (bounded) integer variables with modular arithmetic. In a second stage, the formulae are parsed and modular arithmetic is encoded into propositional logic, discussed in Section 4.2. The result is a propositional temporal logic formula that can be given as input to the `slugs` synthesizer. The synthesizer takes the formula and checks realizability using binary decision diagrams (BDDs) [18]. If the problem described by the program is realizable, then the synthesizer produces a winning strategy in the form of a Mealy machine. The Mealy machine can be represented either as a BDD, or as an enumerated graph (the latter is currently used). Note that the algorithm solves a two-player game of perfect information, in other words, every process can observe what other processes are doing internally and each player controls all its variables in a centralized way (even though these variables may be local to different system processes).

An example is shown in Figs. 4.5a and 4.5b for the process

```
assert active sys proctype foo(){
  do
    :: (x > 0); x = x - 1
    :: (x < 3); x = x + 1
    :: (y > 0); y = y - 1
    :: (y < 2); y = y + 1
    :: skip
  od
}
```

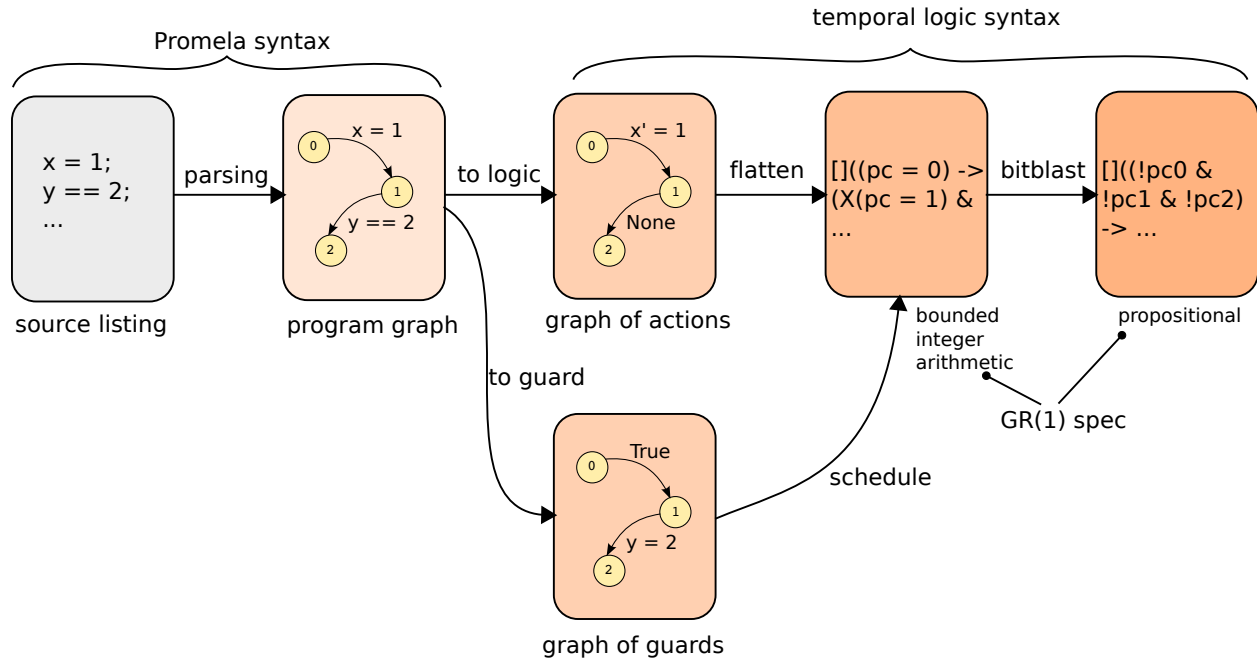



Figure 4.1: The compiler architecture that translates PROMELA source code to a propositional linear temporal logic specification that belongs to the fragment of generalized reactivity of rank 1.

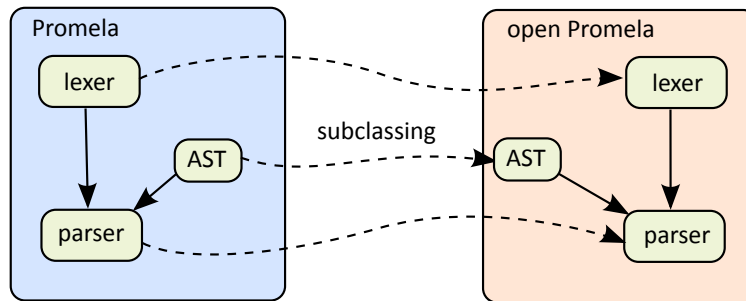


Figure 4.2: Parser front-end: the open Promela lexer and parser are subclasses of the pure Promela lexer and parser.

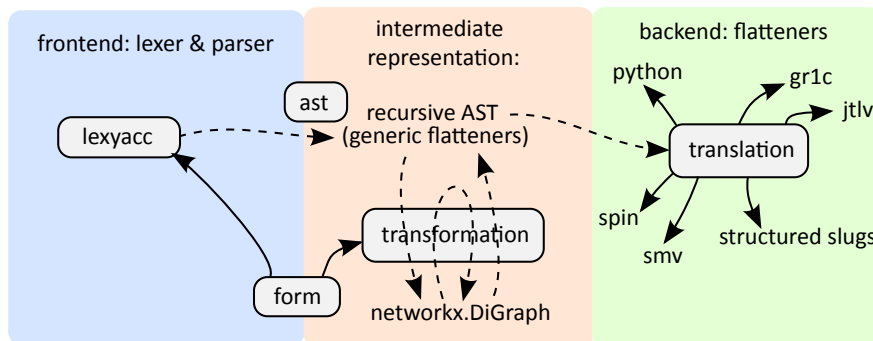


Figure 4.3: Architecture of subpackage tulip.spec.

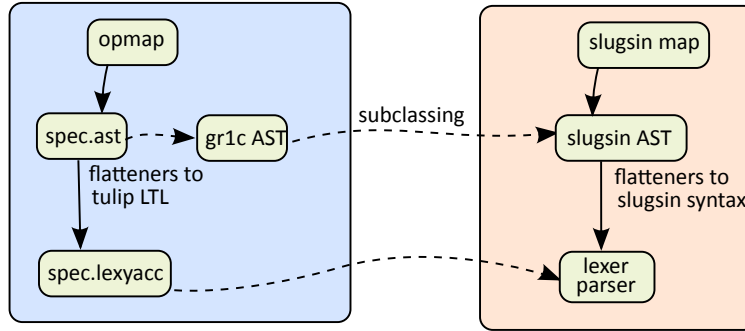
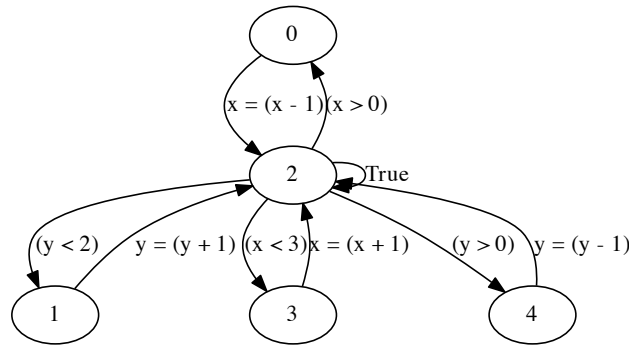
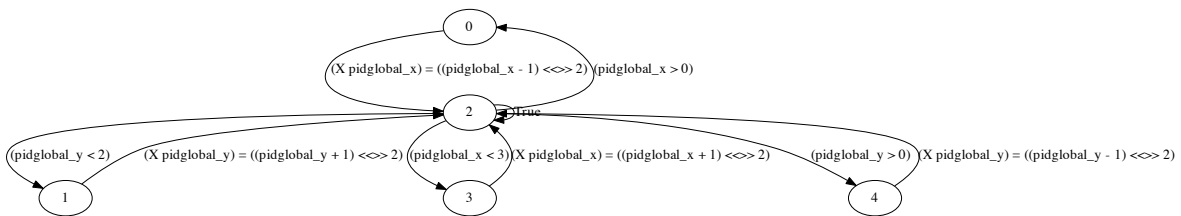


Figure 4.4: The logic AST in (sub)package `spec` is produced by a factory function. Passing a different operator map to this function results in AST nodes with flatteners to a different target syntax. In addition, node classes can be subclassed after being produced by the factory. The main reason for subclassing an AST node class is to override its `flatten` method.



(a) Program graph for an example process.



(b) Graph annotated with action formulae. The symbol “<<>>” denotes truncation of the left operand to the bitwidth defined by the right operand.

Figure 4.5: Program graph and graph of logic formulae.

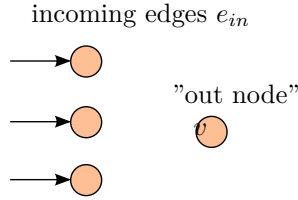


Figure 4.6: During recursive conversion to a program graph, each statement returns a set of incoming edges, together with an out-node.

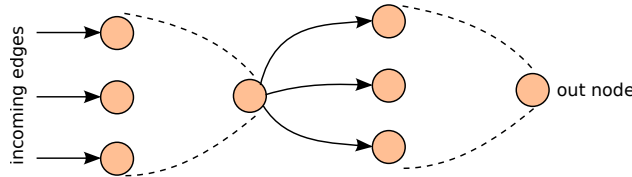


Figure 4.7: A sequence statement is the connection of its elements.

4.1 Program graph construction

The program graph is constructed by recursively converting the syntax tree to nodes and edges. There are two categories of AST nodes: those that define nodes and edges in the program graph, and those that represent expressions. Expressions comprise of operators (with arity 1 or 2) and terminals (variable references, integers, Boolean nullary connectives). Most other AST nodes represent part of the program graph.

We can distinguish roughly four types of AST nodes that define control flow:

1. edge: assignment, expression statement¹, assertion, **else** statement, **printf**, receive, send
2. simple paths: sequence
3. branching and merging: options (**do**, **if**)
4. jumps: **break**, **goto**, labels

The lists above are not exhaustive of all **Promela** constructs. Note that jumps can be regarded also as a special form of branching: two edges with common source node, and one edge guarded by **false**.

In order to convert an AST to a program graph, we need to define what nodes and edges result from each of the above four cases. During the recursion, each of the above returns a pair (e_{in}, v) of incoming edges $e_{in} \subseteq E_r$ and out-node v .

A statement that corresponds to an edge is converted to a fresh target node v , together with its incoming edge $(*, v, stmt)$ that has a yet unspecified source node and is labeled with the statement, Fig. 4.6.

A sequence $stmt_0, stmt_1, \dots, stmt_n$ assembles the edges that it contains, by connecting each out-node v with the incoming edges of its successor. The result is again a pair $(e_{in,0}, v_n)$, where $e_{in,0}$ are the incoming edges of the first element $stmt_0$ and v_n the out-node of the last element $stmt_n$, Fig. 4.7. Sequences are used to define a context (the body of an **atomic** or **d_step** block) and to mark option guards as such².

Selection and repetition statements have similarity in their conversion. A fresh target node v is created first. For each option sequence, its incoming edges $e_{in,j}$ are collected in e_{in} and the out-node v_j is connected to the target node v , by a new **goto** edge (v_j, v) . For a repetition construct, backward edges are also added, by connecting the target node v as source for the incoming edges $e_{in,j}$.

¹In pure **Promela**, expression statements are known as *conditions*, implying that they are guards. However, in open **Promela** an expression statement can contain primed variables, so it is an action. Actions are not pure guards, so they are called *expression statements*. See Section 2.3.4 for details of how guards are defined for actions.

²If a sequence is an option in a selection or repetition construct, then its first element is an *option guard*. Note that the first element may comprise of options, and in that case repetition and selection constructs yield slightly different nodes and edges.

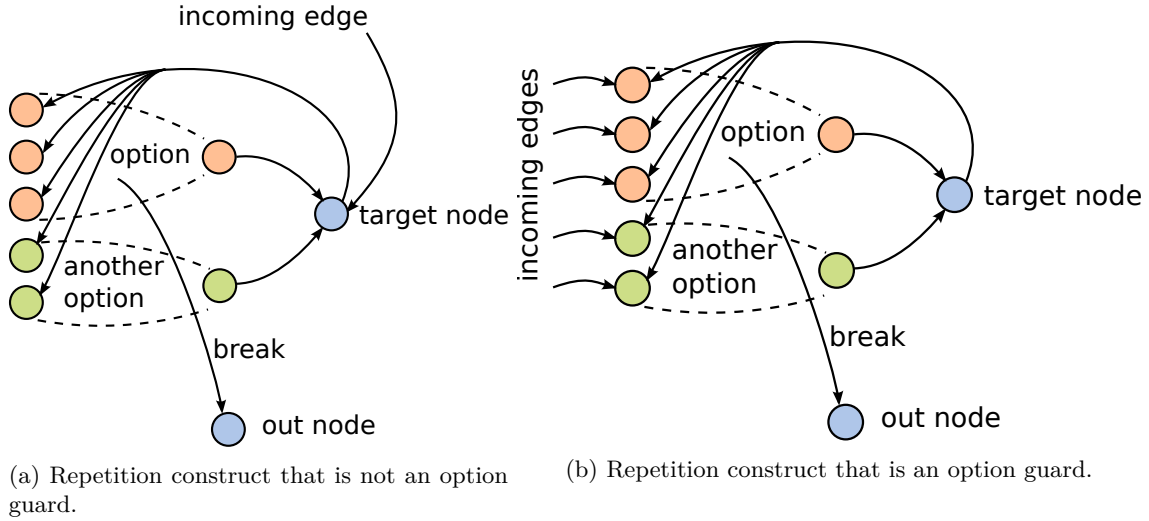


Figure 4.8: Possible cases for a repetition construct.

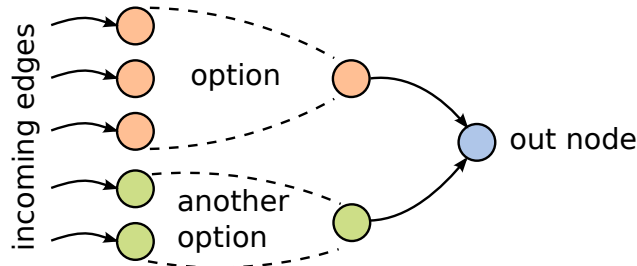


Figure 4.9: Selection construct.

The result is a pair of incoming edges with an out-node. For an `if` statement, the out-node is the target node v and the incoming edges are those collected in e_{in} for the options, Fig. 4.9. For a `do` statement, the out-node is a fresh node w different than the target node v , Fig. 4.8. Any `break` statements inside a `do` statement are equivalent to `goto` statements that lead to w (of the innermost `do` statement, in case of nested repetitions), Fig. 4.8. If a `do` statement is *not* itself an option guard, then a `goto` to the target node v is returned as incoming edge (instead of the collected e_{in}), Fig. 4.8a. Otherwise, the edge set e_{in} is returned, as for an `if` statement, Fig. 4.8b.

It remains to describe how jumps contribute to the program graph, Fig. 4.10. As already remarked, a `break` statement is equivalent to a `goto` to the out-node of the (immediate) enclosing repetition structure. A `goto S0` statement results in a pair (e_{in}, v) of fresh out-node v and a single incoming edge labeled with “goto” and the label-node u_{S0} as target, corresponding to the label `S0`, Fig. 4.10a.

Note that in the special case that a `break` or `goto` statement appears as option guard, then a condition statement `true` is used to annotate the incoming edge, instead of “goto”. This ensures that the edge will not be contracted later.

A labeled statement, as for example `S0: x = 1`, is treated very similarly to a `goto` statement. Let (e_{in}, v) be the incoming edges and out-node of the statement that is labeled, here $x = 1$. The label-node u_{S0} is created and connected as source node to the edges e_{in} . A single edge labeled with “goto” and with target node u_{S0} is returned as incoming edge for the labeled statement. The node v is returned as out-node for the labeled statement. The result is shown in Fig. 4.10b.

Some of the conversions described so far produce edges labeled with “goto”. In order to obtain the final program graph, these edges are contracted. A necessary condition for contraction is that the source node have a unique outgoing edge, so that after the contraction, no new outgoing edges be added to the target node (unless the target node also has a unique incoming edge – the edge that will be contracted). This

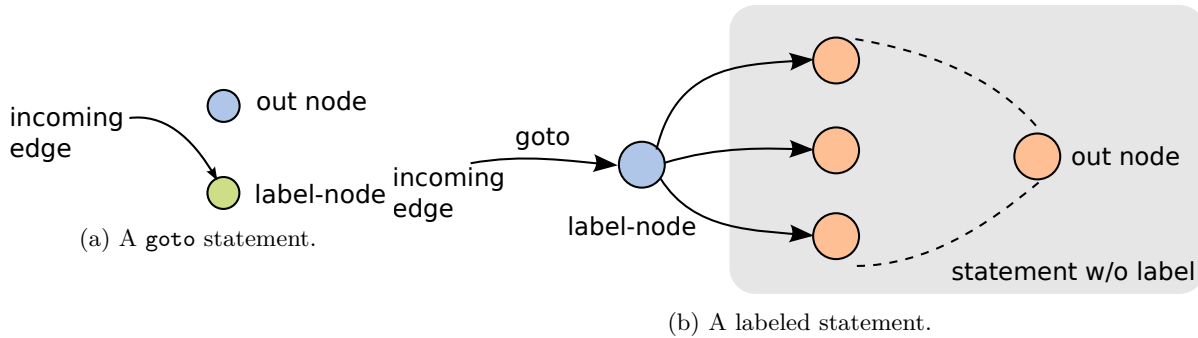


Figure 4.10: Jump statements.

condition is satisfied by the edges annotated with `goto`, because the only case that violates this condition is a `goto` statement that appears immediately after branching. Such a `goto` statement must be an option guard, and as noted earlier, it is replaced by a `true` condition statement.

Besides control flow, context can be defined by enclosing blocks of statements in braces preceded by the keywords `atomic` or `d_step`. Parsing represents a block of statements as a sequence, so the context is stored in the corresponding AST node. All program graph nodes produced by statements inside the sequence are marked with the enclosing context. Note that `goto` statements add edges to special label-nodes (u_{S0} above), but do not create the label-nodes. The context of a label-node is determined by the context of the corresponding labeled statement.

This completes the conversions of statements to program graph edges and nodes. In addition to control flow, variable definitions are collected and used to populate the table of symbols.

4.2 Bitwise encoding

The `Promela` program is translated to logic formulae that contain (bounded) integer arithmetic. They are encoded in logic formulae over bitfields by a translation commonly known as *bitblasting*. Both signed and unsigned integers are supported and represented in two's complement. The operations currently supported are:

- addition
- subtraction
- truncation to given constant width
- (dynamic) shift of a bitfield by a bitfield
- (static) shift of a bitfield by a given constant

The implementation reuses the abstract syntax tree (AST) for first-order temporal logic available in the `tulip.spec` package. For each AST node class, a subclass is defined that overrides the `flatten` method. These flattening methods translate temporal logic (temporal because it involves primed variables) with signed arithmetic of bounded integers to the input syntax of the synthesizer `slugs`.

The truncation operator is necessary for representing the effect of assignment to variables with mod-wrap semantics, as in pure `Promela`. The distinction between assignment and equality results from the language definition. In `Promela`, each expression is evaluated by:

- first, represent all values as signed integers in \mathbb{C} ,
- second, apply all operations in the expression to the precision allowed by the machine's arithmetic-logic unit (ALU) and the conventions defined by the `C` standard.

After the above steps for evaluating an expression have been completed, then there are two cases:

- the expression is a condition statement, so use its truth value (defined as equality to zero)
- a variable is assigned the value defined by this expression.

In the first case, it follows from the evaluation rules that full machine precision applies. So truncation occurs when the bitfield width required for operations reaches either the value 32 or 64, depending on the hardware and operating system.

In GR(1) synthesis applications we are typically interested in much smaller integer domains than 32 bits. Note that this does not refer to the state-space – rather to individual integer variables involved. For this reason, truncation will rarely occur during evaluation of an expression. In any case, the truncation operator can be applied as normal when this occurs.

The second case, involving assignment, is more interesting. It also constitutes the essential difference between the assignment:

$$x = y + 1$$

and the equality condition obtained by replacing the assigned variable with the corresponding primed variable, and the assignment statement by the equality operator:

$$x' == y + 1$$

In pure **Promela**, assignment *truncates* the right-hand side to the width of the datatype representing the variable x on the left-hand side. Let $\text{width}(x)$ denote the width of the bitfield representing variable x , so a constant defined by the program text. Let $\text{trunc}(a, b)$ denote the truncation operator, where a is the bitfield to be truncated and b a constant defining the result’s width. Then the assignment statement is equivalent to the logic formula

$$x' = \text{trunc}(y + 1, \text{width}(x))$$

It becomes evident that the assignment $x = y + 1$ differs from the action $x' == y + 1$, because, in general, the formula $\text{trunc}(y + 1, \text{width}(x))$ is different than the formula $y + 1$. From a synthesis perspective, the assignment is always realizable, because the result is truncated according to the width of x . In contrast, if the value of $y + 1$ does not fit the width of the bitfield representing x , then the formula $x' = y + 1$ is not realizable. The reason is that the system fails to match the value of $y + 1$ by x , because x cannot take sufficiently large values.

4.2.1 Modwrap and saturating semantics

Assignment in C and PROMELA truncates the right hand side expression, before setting the value of the left hand side variable to be equal to the right hand side expression. This is called mod-wrap semantics. For variables whose domain corresponds to a bitfield, modwrap semantics are computationally viable.

We defined a ranged integer data type. Except for ranged integers with minimal value 0 and maximal value that is a power of two, the rest do not correspond to all the values that the underlying bitfield can take. Some values are not possible. The question arises, how assignment to ranged integers should be interpreted.

We choose to not apply any truncation to a value assigned to a variable with saturating semantics. The motivation is computational efficiency of the induced BDD operations. In the following we discuss in detail the reasons that lead to this decision.

Suppose that we decided to truncate, using as base the interval width $\max(\text{dom}(x))$ (suppose for now that $\min(\text{dom}(x)) = 0$). Unlike the case for base two, the “modulo” operation for arbitrary base requires BDDs of exponential size [18]. Therefore, this natural choice for truncation semantics is inefficient.

The remaining option is to define the semantics using modulo operations of base two. One option is to truncate to the width of the underlying bitfield, and if the result is outside the range, then raise a safety violation. This semantics is confusing, because it is an indirect definition with little relation to the variable’s actual domain. The other option is to first truncate modulo the binary representation, then subtract the range (the result is ensured to belong to the variable’s domain, so assignment is always possible in this case). The problem with this approach is its very confusing semantics.

In summary, ranged integers serve as a shorthand for a longer bitfield, together with a safety formula, and deactivation of truncation during assignments.

4.2.2 Bitfield representation of signed integers

In this section, we describe how a ranged integer declared as $\text{int}(a, b)$ x is represented as a bitvector. Let $m \triangleq \max\{|a|, |b|\}$, where $a, b \in \mathbb{Z}$, and $\xi \triangleq \lceil \log_2 m \rceil$. Then the integer is represented by a bitfield of width $\xi + 1$ (that takes values over $\mathbb{N} \cap [-2^\xi, 2^\xi - 1]$, together with the safety constraint $\square(a \leq x \leq b)$).

4.2.3 Arithmetic as bitvector logic

All arithmetic operations can be computed using as main element an adder-subtractor, Fig. 4.13, that comprises of half-adders, Fig. 4.12. An adder-subtractor can do both addition and subtraction. A half-adder computes the functions

$$\begin{aligned} \text{sum}(a, b, c_{in}) &\triangleq a \oplus b \oplus c_{in} \\ \text{carry}(a, b, c_{in}) &\triangleq (a \wedge b) \vee ((a \oplus b) \wedge c_{in}), \end{aligned} \quad (4.1)$$

where \oplus denotes exclusive disjunction (XOR), i.e., $x \oplus y \triangleq (x \wedge \neg y) \vee (\neg x \wedge y)$. An adder-subtractor can perform both addition and subtraction (by setting the input signal c_{in}). The result is

$$\text{addsub}(a, b, c_{in}) \triangleq a + (1 - 2c_{in})b. \quad (4.2)$$

An adder-subtractor can be used to compute equality and strict inequality as follows

$$\begin{aligned} a = b &\iff \bigwedge_{i=0}^{l-1} \neg(a_i \oplus b_i) \\ a < b &\iff \text{addsub}(a, b, 1).c_{out} \oplus \neg(a_{l-1} \oplus b_{l-1}). \end{aligned} \quad (4.3)$$

Negated equality and non-strict inequality are obtained by negating the previous

$$\begin{aligned} a \neq b &\iff \neg(a = b) \\ a \geq b &\iff \neg(a < b). \end{aligned} \quad (4.4)$$

By symmetry, these suffice to compute $a > b$ and $a \leq b$. This is a brief overview (including errata) of the encoding described in [74], where more can be found, e.g., about sign extension.

The circuit of Fig. 4.13 becomes a subtractor if the carry input bit c_{in} is 1. This value of c_{in} triggers the computation of two's complement for the second input b . Computing two's complement is equivalent to negation in Boolean arithmetic. Therefore, the complement of the representation in two's complement of b is the representation in two's complement of $-b$. This turns addition of a and b to addition of a and $-b$, thus subtraction.

Note that 2^l requires $l + 1$ bits to be represented. It holds that

$$2^l = 0b \underbrace{10 \dots 0}_{l+1} = (0b \underbrace{011 \dots 1}_{l+1}) + 1 = (0b \underbrace{11 \dots 1}_l) + 1, \quad (4.5)$$

where the prefix $0b$ signifies base two notation. A positive number $0 < x < 2^{l-1}$ in two's complement representation is

$$0b \underbrace{x_{l-1} x_{l-2} \dots x_0}_l = 0b \underbrace{0 x_{l-2} \dots x_0}_l, \quad (4.6)$$

because $0 < x < 2^{l-1} \implies x_{l-1} = 0$. A negative number x with $-2^{l-1} \leq x < 0$ in two's complement representation is (Fig. 4.11)

$$\begin{aligned} 2^l - |x| &= (0b \underbrace{11 \dots 1}_l) - (0b \underbrace{|x|_{l-1} \dots |x|_0}_l) + 1 \\ &= (0b(1 - |x|_{l-1}) \dots (1 - |x|_0)) + 1 \\ &= (0b(\neg |x|_{l-1}) \dots (\neg |x|_0)) + 1. \end{aligned} \quad (4.7)$$

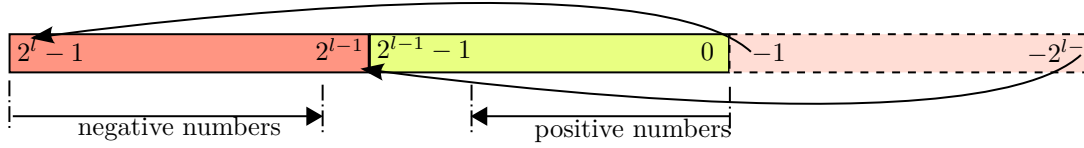


Figure 4.11: Two's complement on the line of integers.

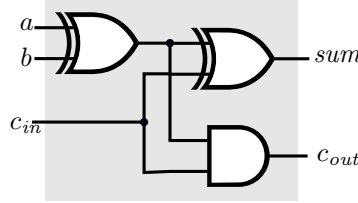


Figure 4.12: Full adder.

By the bounds on x , it is

$$\begin{aligned}
 -2^{l-1} \leq x < 0 &\iff 0 \leq 2^{l-1} + x < 2^{l-1} \iff 2^{l-1} \leq 2^{l-1} + 2^{l-1} - (-x) < 2^{l-1} + 2^{l-1} \iff \\
 2^{l-1} \leq 2^l - |x| < 2^l &\implies (2^l - |x|)_{l-1} = 1
 \end{aligned} \tag{4.8}$$

Up to this point we have noted the conversion from binary magnitude and separate bit sign representation to two's complement representation.

Next, it is shown that two's complement negates a number x . In other words, the two's complement of the representation y of x in two's complement yields the two's complement z of the negated $-x$. In addition, it is shown that it corresponds to the elementwise negation and carry-in $c_{in} = 1$ that turns a ripple-carry adder into a subtractor.

$$\begin{aligned}
 2^l - (0by_{l-1} \dots y_0) &= (0b\underbrace{1 \dots 1}_l) + 1 - (0by_{l-1} \dots y_0) \\
 &= (0b(1 - y_{l-1}) \dots (1 - y_0)) + 1 \\
 &= (0b(\neg y_{l-1}) \dots (\neg y_0)) + c_{in}.
 \end{aligned} \tag{4.9}$$

This is exactly the operation performed when $c_{in} = 1$. If x is positive, then its two's complement y has $y_{l-1} = 0$ and $|x| = 0by_{l-2} \dots y_0$. As already shown, the above (z) is the two's complement of $-|x|$. If x is negative, then its two's complement y has $y_{l-1} = 1 \implies \neg y_{l-1} = 0$, and $0b(\neg y_{l-2} \dots (\neg y_0)) + 1 = 2^{l-1} - (2^l - |x| - 2^{l-1}) = 2^{l-1} + 2^{l-1} - 2^l - |x| = -|x|$.

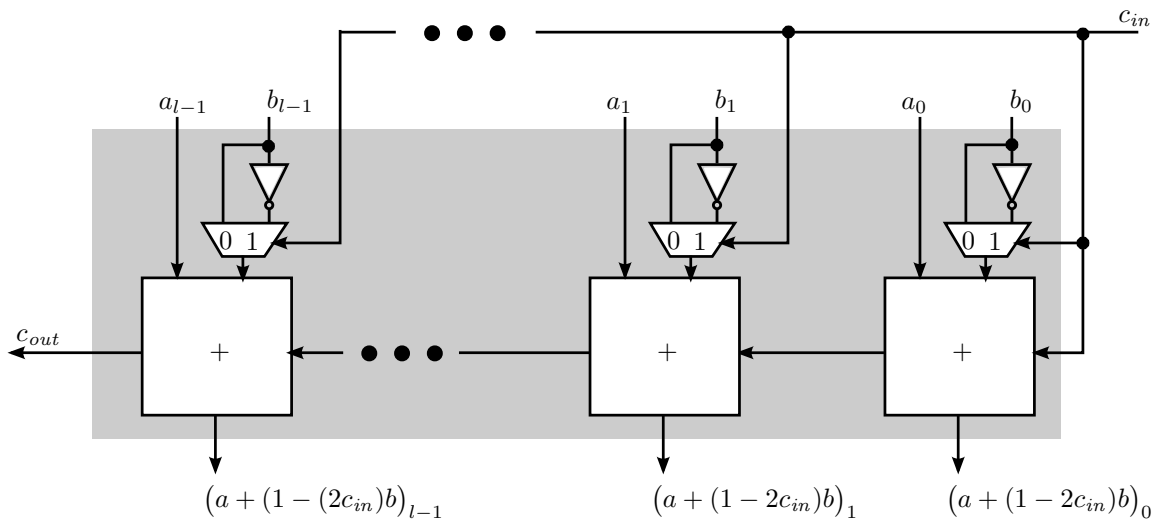


Figure 4.13: Ripple-carry adder-subtractor.

Appendix A

Appendices

A.1 Open Promela syntax

A.1.1 Lexemes

The following new keywords have been introduced:

- `env`, `sys` for declaring variables and control flow controlled by the environment and system, respectively,
- `free` for declaring symbolic variables,
- `assume`, `assert` for declaring data flow constrains on the environment and system, respectively,
- `sync`, `async` for declaring synchronous and asynchronous products,
- PRIME “`'`” as postfix “next” operator for variables (`○`).

A.1.2 Grammar

The following new parser production rules have been introduced (in union with existing production rules for the same nonterminals). Upper case denotes new keywords, or tokens like PRIME. Nonterminals are to be found in the language reference [39] and in the `promela.yacc` module of the parser in PYTHON.

$$\langle ltl \rangle ::= [\text{ASSUME} \mid \text{ASSERT}] \text{LTL} \{ ' \langle expr \rangle ' \}$$
$$\langle module \rangle ::= \langle async \rangle \mid \langle sync \rangle$$
$$\langle async \rangle ::= [\text{ENV} \mid \text{SYS}] \text{ASYNC} \{ ' \langle async-body \rangle ' \}$$
$$\langle async-body \rangle ::= [\langle async-body \rangle] \langle async-unit \rangle$$
$$\langle async-unit \rangle ::= \langle sync \rangle \mid \langle proc \rangle \mid \langle one-decl \rangle \mid \langle semi \rangle$$
$$\langle sync \rangle ::= [\text{ENV} \mid \text{SYS}] \text{SYNC} \{ ' \langle sync-body \rangle ' \}$$
$$\langle sync-body \rangle ::= [\langle sync-body \rangle] \langle sync-unit \rangle$$
$$\langle sync-unit \rangle ::= \langle proc \rangle \mid \langle once-decl \rangle \mid \langle semi \rangle$$
$$\langle proctype \rangle ::= [\text{ENV} \mid \text{SYS}] \text{PROCTYPE}$$
$$\langle prefix-proctype \rangle ::= [\text{ASSUME} \mid \text{ASSERT}] \langle prefix-proctype \rangle$$

Table A.1: The use case of a keyword `step` as a non-deterministic uninterruptible block.

		blocking	
		admissible	raises error
det/ic	No	<code>atomic</code>	<code>step</code>
	Yes	<code>atomic</code>	<code>d_step</code>

$\langle one\text{-}decl \rangle ::= [\text{FREE}] [\text{ENV} \mid \text{SYS}] [\langle \text{typename} \rangle \mid \langle \text{name} \rangle] \langle \text{var-list} \rangle$

$\langle \text{typename} \rangle ::= \text{'int' ' (' } \langle \text{number} \rangle \text{' , ' } \langle \text{number} \rangle \text{')'}$

$\langle \text{expr} \rangle ::= \langle \text{expr} \rangle \text{ PRIME}$

A.1.3 The keyword `step`

The two statements `atomic` and `d_step` serve three of the possible use cases, as depicted in Table A.1. The `d_step` without the prefix `d_` is a natural candidate for denoting strictly uninterrupted execution that can exhibit non-determinism, i.e., blocks of the form `verb{...}`. Such a block allows conditioning the selection of a block of statements on the a priori executability of all statements in the block. An `atomic` block cannot be used for this purpose, because it can be interrupted. A `d_step` may be undesirable, to leave design freedom to the synthesizer. The `step` keyword is currently not implemented, but may be introduced in the future.

A.2 Relevant work

The translation from PROMELA to declarative formalisms has been considered in [67, 68, 66] and decision diagrams in [75]. These translations aim at verification, do not have LTL as target language, and either have limited support for atomicity [66], no details [67], or do not handle program graphs correctly [68].

TLA [11, 43] subsumes the language proposed here, since it includes quantification, but is intended as a theorem proving activity, is declarative, and is aimed at verification. Nonetheless, one can view the proposed translation as from (open) PROMELA to TLA. SMV is a declarative language with “next” operator, as well as synchronous and asynchronous processes, but no notion of assumptions/assertions, controllability, or imperative variables [70, 44]. JTLV [30] takes as input an SMV-like language for synthesis specifications, but with no imperative constructs. ASPECTLTL is a further declarative extension for aspect-oriented programming [76].

RPROMELA is an extension of PROMELA that adds synchronous-reactive constructs (not in the sense of reactive synthesis) that include synchronous products and channels called ports [77, 78]. Its semantics are defined in terms of *stable states*, where the synchronous product blocks, waiting for message reception from its global ports. RPROMELA does not address modeling of the environment, nor declarative elements.

Synchronous-reactive languages can be considered as potential candidates for expressing synthesis specifications. These include imperative textual languages like ESTEREL, ESTEREL-C, JAVA-ESTEREL, QUARTZ, REACTIVE-C, imperative graphical languages like STATECHARTS, ARGOS, and SYNCCHARTS, and declarative textual ones like LUSTRE and LUCID SYNCHRONE, and declarative graphical ones like SIGNAL [79, 65]. These languages are by definition *deterministic*, intended for direct design of transducers [79, 65, 80]. In synthesis, non-determinism is an essential feature of the specification. Extending ESTEREL with non-determinism has been considered, but in the context of transducers and verification [81], and using guarded command languages like PROMELA as example.

The approach proposed here has common elements with *constraint imperative programming* (CIP), introduced with the experimental language KALEIDOSCOPE [82, 83, 84, 85], one of the first attempts to integrate the imperative and declarative constraint programming paradigms. An observation from [82], which applies also here, is that specifiers need to express two types of relations: long-lived (best described declaratively),

and sequencing relations (more naturally expressed in an imperative style). However, CIP does *not* ensure correct reactivity, because the constraints are solved online.

In the following sections, we discuss some of the above in more detail.

A.2.1 RPromela and RSPIN

An earlier extension of **Promela** that adds synchronous-reactive constructs has been proposed under the name **RPromela**, which stands for “reactive **Promela**” [77, 78]. The term “reactive” here refers to the synchronous-reactive semantics of languages like **ESTEREL** and *not* to reactivity as defined by Manna and Pnueli [60]. The seven new keywords it adds are

`automaton, in, proctype, inport, outport, link, external`

RSPIN is reported as a preprocessor to **SPIN** that translates **RPromela** models to **Promela** models. An example code fragment is

```
1 rproctype foo(  
2   inport a, b, c;  
3   outprt x, y, z  
4 ) (bit r) {  
5     automaton subproc(  
6       external inport a;  
7       inport x={bool};  
8       outport g={byte}  
9     ) () {...}  
10    ...  
11 }
```

A reactive process type, denoted by the keyword `rproctype`, encapsulates a number of **automata**. The automata inside a single `rproctype` are composed synchronously. The semantics of synchronous composition are based on the concept of “stable states”. A state is “stable” if it is a receive statement from channels external to the `rproctype`, and so can block. When a reactive process reacts, it receives from its **inports** (representing a transducer), computes until it blocks at the next stable state, and so produces its outputs by writing to its **outports**. The automata inside a single reactive process communicate exclusively by means of channels local to that `rproctype`. Not all automata need to execute during a synchronous reaction, i.e., at least one should be unblocked, but each of the remaining ones can be blocked or not. The synchronous product is taken statically by **RSPIN**, producing **Promela** code as output. The semantics of synchronous product are different than here (implicit “stable states”, linking via channels), as well as its syntax. **RPromela** is designed to use channels heavily, which quickly leads to state space explosion in explicit model checking. The semantics are based on transitions (explicit) and not “stable states” as in [78, 77] (implicit). Stable states associate implicit semantics to the synchronous products that can cause confusion more easily. Overall, **RPromela** is an interesting approach, that suggests a syntax for defining synchronous products in **Promela** and a useful resource for ideas that discusses the issues involved.

RPromela is not suitable for synthesis specifications, for the limitation of synchronous-reactive languages to be deterministic, with the intent of representing transducer implementations. Besides, there are several points that are treated differently here, because of differences in semantics, e.g., the synchronous product, additional semantics for open systems that are not considered in **RPromela**, (notions of imperative/declarative, primed variables and their owners/quantification, assertions/assumptions, and `pro`), as well as practical considerations for syntax, e.g., the keyword `sync` instead of `rproctype` that has small lexical distance from `proctype`.

A.2.2 Kaleidoscope

Kaleidoscope is an object-oriented constraint imperative programming language that was introduced by Bjorn Freeman-Benson in 1990 [82]. It was a research language and although the authors discuss an implementation, currently an implementation cannot be found online. Nonetheless, it constitutes one of the first

attempts to integrate the imperative and the declarative constraint programming paradigm. It introduces the term *Constraint Imperative Programming* (CIP). An observation from [82] that applies to some extent also here is that programmers (and specifiers alike) need to specify two types of relations:

1. Long-lived relations (more suitably described with a declarative paradigm)
2. Sequencing relations between program states (more naturally expressed in an imperative paradigm).

The integration of two paradigms serves this purpose in a single syntax.

Assignment in *Kaleidoscope* has some similarity to variable semantics as defined here. In particular, the `very_weak_stay` constraint is very similar to the default invariance of imperative variables here.

An innovative feature of *Kaleidoscope* is the distinction between statements and the advancement of time. The statement *separator* is “;”, whereas the variable valuation *streams* are advanced by “#”. All constraints between two consecutive hash marks hold simultaneously, much like the semantics of synchronous-reactive languages. Note the difference with synchronous-reactiveness though, in that the advancement of time is explicit, instead of implicit. Instead of explicit time advancement, synchronous constrains in open PROMELA can be expressed by conjoining them into a single expression.

An interesting observation is that constraints in CIP languages are expressed using keywords like `always` (“durable relations”), `once` (“transitory relations”), and `while ... assert`, which effectively express temporal logic specifications.

However, the approach in *Kaleidoscope* is based on solving constraints at runtime [83, 84, 86, 85] using a constraint solver and a constrain hypergraph. This solution approach does not ensure satisfaction of temporal logic specifications, as is the case here. In other words, it does *not* produce designs that are correct by construction.

In addition, there is no notion of reactivity in *Kaleidoscope* – no adversarial environment. Reactiveness is a significant motivation in this work, and distinguishes it from previous work in CIP languages.

The constraint imperative languages are full programming languages, with a lot of emphasis placed on type systems, constraints as objects, and the design of compilers and run-time constraint solvers interfaced to the virtual machine executing the program. They have been developed with the design of user interfaces in mind, as for example the graphical interface of an operating system.

This leads to another difference: *Kaleidoscope* is fundamentally object-oriented, whereas modeling Kripke structures in *Promela* is procedural (both are imperative) paradigms. Constraints in *Kaleidoscope* are intended to be applied to objects, so that object relations can be maintained. Note that this viewpoint is a direct consequence of aiming at the specification of behavior for geometric elements comprising a graphical user interface.

Another concept of *Kaleidoscope* and other CIP languages is that of a *constraint hierarchy*, that distinguishes between *required* constraints, and *optional* constraints that are ranked based on a subjective prioritization. Such constraint hierarchies cannot be expressed in a qualitative synthesis context and require quantitative games [87].

A recently proposed descendant of *Kaleidoscope* is *Babelsberg* [88], a new language that improves on those ideas (supported by the work on constraint satisfaction solvers, as for example *Cassowary*, that is used as the layout engine in *OS X Lion*).

Bibliography

- [1] O. Kupferman, “Recent challenges and ideas in temporal synthesis,” in *Proceedings of the 38th International Conference on Current Trends in Theory and Practice of Computer Science*, ser. SOFSEM’12. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 88–98. [Online]. Available: dx.org/10.1007/978-3-642-27660-6_8
- [2] I. Walukiewicz, “A landscape with games in the background,” *Logic in Computer Science, Symposium on*, vol. 0, pp. 356–366, 2004. [Online]. Available: doi.org/10.1109/LICS.2004.1319630
- [3] N. Piterman, A. Pnueli, and Y. Sa’ar, “Synthesis of reactive (1) designs,” in *Verification, Model Checking, and Abstract Interpretation (VMCAI)*. Charleston, SC, USA: Springer, January 2006, pp. 364–380. [Online]. Available: doi.org/10.1007/11609773_24
- [4] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa’ar, “Synthesis of Reactive(1) designs,” *Journal of Computer and System Sciences*, vol. 78, no. 3, pp. 911 – 938, 2012, in Commemoration of Amir Pnueli. [Online]. Available: doi.org/10.1016/j.jcss.2011.08.007
- [5] R. Ehlers, “Generalized Rabin(1) synthesis with applications to robust system synthesis,” in *NASA Formal Methods*, ser. LNCS, M. Bobaru, K. Havelund, G. J. Holzmann, and R. Joshi, Eds. Springer, 2011, vol. 6617, pp. 101–115. [Online]. Available: doi.org/10.1007/978-3-642-20398-5_9
- [6] R. Alur and S. La Torre, “Deterministic generators and games for ltl fragments,” *ACM Trans. Comput. Logic*, vol. 5, no. 1, pp. 1–25, Jan. 2004. [Online]. Available: doi.org/10.1145/963927.963928
- [7] H. Kress-Gazit, G. Fainekos, and G. Pappas, “Temporal-logic-based reactive mission and motion planning,” *Robotics, IEEE Transactions on*, vol. 25, no. 6, pp. 1370–1381, Dec 2009. [Online]. Available: doi.org/10.1109/TRO.2009.2030225
- [8] M. Kloetzer and C. Belta, “A fully automated framework for control of linear systems from temporal logic specifications,” *Automatic Control, IEEE Transactions on*, vol. 53, no. 1, pp. 287–297, Feb 2008. [Online]. Available: doi.org/10.1109/TAC.2007.914952
- [9] T. Wongpiromsarn, U. Topcu, and R. M. Murray, “Synthesis of control protocols for autonomous systems,” *Unmanned Systems*, vol. 1, no. 01, pp. 21–39, 2013. [Online]. Available: dx.org/10.1142/S2301385013500027
- [10] R. Rosner, “Modular synthesis of reactive systems,” Ph.D. dissertation, Weizmann Institute of Science, Rehovot, Israel, 1992. [Online]. Available: http://www.researchgate.net/publication/238759536_Modular_synthesis_of_reactive_systems/file/50463527f8b648c3ba.pdf
- [11] L. Lamport, “The temporal logic of actions,” *ACM Trans. Program. Lang. Syst.*, vol. 16, no. 3, pp. 872–923, May 1994. [Online]. Available: doi.org/10.1145/177492.177726
- [12] Z. Manna and A. Pnueli, “A hierarchy of temporal properties (invited paper, 1989),” in *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing*, ser. PODC ’90. New York, NY, USA: ACM, 1990, pp. 377–410. [Online]. Available: doi.org/10.1145/93385.93442
- [13] K. Schneider, *Verification of reactive systems: formal methods and algorithms*. Springer, 2004.

- [14] A. Morgenstern and K. Schneider, “A LTL fragment for GR(1)-synthesis,” *EPTCS*, vol. 50, pp. 33–45, Feb. 2011. [Online]. Available: doi.org/10.4204/EPTCS.50.3
- [15] S. Sohail, F. Somenzi, and K. Ravi, “A hybrid algorithm for ltl games,” in *Verification, Model Checking, and Abstract Interpretation*. Springer, 2008, pp. 309–323. [Online]. Available: doi.org/10.1007/978-3-540-78163-9_26
- [16] M. Dwyer, G. Avrunin, and J. Corbett, “Patterns in property specifications for finite-state verification,” in *Software Engineering, 1999. Proceedings of the 1999 International Conference on*, May 1999, pp. 411–420. [Online]. Available: doi.org/10.1145/302405.302672
- [17] Z. Manna and A. Pnueli, “Tools and rules for the practicing verifier,” Stanford University, Stanford, CA, USA, Tech. Rep., 1990. [Online]. Available: <ftp://reports.stanford.edu/pub/cstr/reports/cs/tr/90/1321/CS-TR-90-1321.pdf>
- [18] R. E. Bryant, “Graph-based algorithms for boolean function manipulation,” *IEEE Trans. Comput.*, vol. 35, no. 8, pp. 677–691, Aug. 1986. [Online]. Available: doi.org/10.1109/TC.1986.1676819
- [19] A. Pnueli, “The temporal logic of programs,” in *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, ser. SFCS ’77, 31 1977-nov. 2 1977, pp. 46–57. [Online]. Available: doi.org/10.1109/SFCS.1977.32
- [20] C. Baier and J.-P. Katoen, *Principles of Model Checking*. The MIT Press, May 2008.
- [21] M. Abadi and L. Lamport, “Open systems in tla,” in *Proceedings of the Thirteenth Annual ACM Symposium on Principles of Distributed Computing*, ser. PODC ’94. New York, NY, USA: ACM, 1994, pp. 81–90. [Online]. Available: doi.org/10.1145/197917.197960
- [22] A. Pnueli and R. Rosner, “On the synthesis of a reactive module,” in *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL ’89. New York, NY, USA: ACM, 1989, pp. 179–190. [Online]. Available: doi.org/10.1145/75277.75293
- [23] A. Church, “Application of recursive arithmetic to the problem of circuit synthesis,” in *Summaries of talks presented at the summer institute for symbolic logic*, 2nd ed. Cornell University: Communications Research Division, Institute for Defense Analyses, Princeton, N.J., (the summer institute was held in 1957, Church’s paper is a -significantly- extended version of the talk) 1960.
- [24] W. Thomas, “On the synthesis of strategies in infinite games,” in *STACS 95*. Springer, 1995, pp. 1–13. [Online]. Available: doi.org/10.1007/3-540-59042-0_57
- [25] —, “Solution of church’s problem: A tutorial,” *New Perspectives on Games and interaction*, vol. 5, 2008. [Online]. Available: <http://www.automata.rwth-aachen.de/download/papers/thomas/tho08c.pdf>
- [26] G. H. Mealy, “A method for synthesizing sequential circuits,” *Bell System Technical Journal*, vol. 34, no. 5, pp. 1045–1079, 1955. [Online]. Available: doi.org/10.1002/j.1538-7305.1955.tb03788.x
- [27] E. F. Moore, “Gedanken-experiments on sequential machines,” in *Automata Studies*, ser. Annals of mathematics studies, C. E. Ashby, W. R.; Shannon, Ed. Princeton University Press, 1956, vol. 34, pp. 129–153.
- [28] A. Pnueli and R. Rosner, “On the synthesis of an asynchronous reactive module,” in *Proceedings of the 16th International Colloquium on Automata, Languages and Programming*, ser. ICALP ’89. London, UK, UK: Springer-Verlag, 1989, pp. 652–671. [Online]. Available: doi.org/10.1007/BFb0035790
- [29] B. Jobstmann, S. Galler, M. Weiglhofer, and R. Bloem, “Anzu: A tool for property synthesis,” in *Computer Aided Verification*. Springer, 2007, pp. 258–262. [Online]. Available: doi.org/10.1007/978-3-540-73368-3_29

- [30] A. Pnueli, Y. Sa’ar, and L. D. Zuck, “JTLV: A framework for developing verification algorithms,” in *Computer Aided Verification*. Springer, 2010, pp. 171–174. [Online]. Available: doi.org/10.1007/978-3-642-14295-6_18
- [31] R. Bloem, A. Cimatti, K. Greimel, G. Hofferek, R. Könighofer, M. Roveri, V. Schuppan, and R. Seeber, “RATSY—A new requirements analysis tool with synthesis,” in *Computer Aided Verification*. Springer, 2010, pp. 425–429. [Online]. Available: doi.org/10.1007/978-3-642-14295-6_37
- [32] S. Livingston, R. Murray, and J. Burdick, “Backtracking temporal logic synthesis for uncertain environments,” in *ICRA*, 2012, pp. 5163–5170. [Online]. Available: dx.org/10.1109/ICRA.2012.6225208
- [33] R. Ehlers and V. Raman, “Low-effort specification debugging and analysis,” *EPTCS*, vol. 157, pp. 117–133, 07 2014. [Online]. Available: doi.org/10.4204/EPTCS.157.12
- [34] B. Jobstmann and R. Bloem, “Optimizations for LTL synthesis,” in *Formal Methods in Computer Aided Design, 2006. FMCAD’06*. IEEE, 2006, pp. 117–124. [Online]. Available: doi.org/10.1109/FMCAD.2006.22
- [35] R. Ehlers, “Unbeast: Symbolic bounded synthesis,” in *Proceedings of the 17th International Conference on Tools and Algorithms for the Construction and Analysis of Systems: Part of the Joint European Conferences on Theory and Practice of Software*, ser. TACAS’11/ETAPS’11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 272–275. [Online]. Available: doi.org/10.1007/978-3-642-19835-9_25
- [36] —, “Experimental aspects of synthesis,” *EPTCS*, vol. 50, 2011. [Online]. Available: doi.org/10.4204/EPTCS.50.1
- [37] A. Bohy, V. Bruyère, E. Filiot, N. Jin, and J.-F. Raskin, “ACACIA+, a tool for LTL synthesis,” in *Computer Aided Verification*. Springer, 2012, pp. 652–657. [Online]. Available: dx.org/10.1007/978-3-642-31424-7_45
- [38] I. Filippidis and contributors. (2013) List of verification and synthesis tools. [Online]. Available: https://github.com/johnyf/tool_lists/blob/master/verification_synthesis.md
- [39] G. J. Holzmann, *The SPIN Model Checker: Primer and Reference Manual*, 1st ed. Addison-Wesley Professional, 2003.
- [40] E. W. Dijkstra, “Guarded commands, nondeterminacy and formal derivation of programs,” *Commun. ACM*, vol. 18, no. 8, pp. 453–457, Aug. 1975. [Online]. Available: doi.org/10.1145/360933.360975
- [41] G. J. Holzmann. PROMELA language reference (<http://spinroot.com/spin/Man/promela.html>). [Online]. Available: <http://spinroot.com/spin/Man/promela.html>
- [42] P. Van-Roy and S. Haridi, *Concepts, techniques, and models of computer programming*. MIT press, 2004.
- [43] L. Lamport, *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2002.
- [44] R. Cavada, A. Cimatti, C. A. Jochim, G. Keighren, E. Olivetti, M. Pistore, M. Roveri, and A. Tchaltsev, “NuSMV 2.5 User Manual,” Fondazione Bruno Kessler, 18 Via Sommarive, 38055 Povo (Trento), Italy, Tech. Rep., 2010. [Online]. Available: nsmv.fbk.eu/NuSMV/userman/v25/nsmv.pdf
- [45] J. Barnat, L. Brim, V. Havel, J. Havlíček, J. Kriho, M. Lenčo, P. Ročkai, V. Štill, and J. Weiser, “DIVINE 3.0 – an explicit-state model checker for multithreaded C & C++ programs,” in *Computer Aided Verification (CAV 2013)*, ser. LNCS, vol. 8044. Springer, 2013, pp. 863–868. [Online]. Available: dx.org/10.1007/978-3-642-39799-8_60
- [46] J. Gennari, S. Hedrick, F. Long, J. Pincar, and R. Seacord, “Ranged integers for the C programming language,” Software Engineering Institute, Carnegie Mellon University, Technical Note CMU/SEI-2007-TN-027, Retrieved Oct 07, 2014 2007. [Online]. Available: <http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=8265>

- [47] R. M. Keller, “Formal verification of parallel programs,” *Commun. ACM*, vol. 19, no. 7, pp. 371–384, Jul. 1976. [Online]. Available: doi.org/10.1145/360248.360251
- [48] A. K. Chandra, D. C. Kozen, and L. J. Stockmeyer, “Alternation,” *J. ACM*, vol. 28, no. 1, pp. 114–133, Jan. 1981. [Online]. Available: doi.org/10.1145/322234.322243
- [49] D. E. Muller, A. Saoudi, and P. E. Schupp, “Alternating automata, the weak monadic theory of the tree, and its complexity,” in *Automata, Languages and Programming*. Springer, 1986, pp. 275–283. [Online]. Available: doi.org/10.1016/0304-3975(92)90076-R
- [50] M. Y. Vardi, “Alternating automata and program verification,” in *Computer Science Today*. Springer, 1995, pp. 471–485. [Online]. Available: doi.org/10.1007/BFb0015261
- [51] —, “An automata-theoretic approach to linear temporal logic,” in *Logics for concurrency: structure versus automata*. Springer, 1996, vol. 1043, pp. 238–266. [Online]. Available: doi.org/10.1007/3-540-60915-6_6
- [52] O. Kupferman and M. Vardi, “Safraless decision procedures,” in *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, Oct 2005, pp. 531–540. [Online]. Available: doi.org/10.1109/SFCS.2005.66
- [53] C. A. R. Hoare, *Communicating sequential processes*. Prentice-hall Englewood Cliffs, 1985–2004, vol. 178. [Online]. Available: http://www.usingcsp.com/cspbook.pdf
- [54] H. Søndergaard and P. Sestoft, “Non-determinism in functional languages,” *The Computer Journal*, vol. 35, no. 5, pp. 514–523, 1992. [Online]. Available: http://comjnl.oxfordjournals.org/content/35/5/514.abstract
- [55] J. McCarthy, “A basis for a mathematical theory of computation,” in *Computer Programming and Formal Systems*, P. Braffort and D. Hirschberg, Eds. North-Holland, 1963, pp. 33–70. [Online]. Available: doi.org/10.1145/1460690.1460715
- [56] R. W. Floyd, “Nondeterministic algorithms,” *J. ACM*, vol. 14, no. 4, pp. 636–644, Oct. 1967. [Online]. Available: doi.org/10.1145/321420.321422
- [57] M. Broy, “A theory for nondeterminism, parallelism, communication, and concurrency,” *Theoretical Computer Science*, vol. 45, no. 0, pp. 1 – 61, 1986. [Online]. Available: doi.org/10.1016/0304-3975(86)90040-X
- [58] R. McNaughton, “Infinite games played on finite graphs,” *Annals of Pure and Applied Logic*, vol. 65, no. 2, pp. 149 – 184, 1993. [Online]. Available: doi.org/10.1016/0168-0072(93)90036-D
- [59] K. R. M. Leino, “Dafny: An automatic program verifier for functional correctness,” in *Logic for Programming, Artificial Intelligence, and Reasoning*, ser. Lecture Notes in Computer Science, E. M. Clarke and A. Voronkov, Eds., vol. 6355. Springer, 2010, pp. 348–370. [Online]. Available: doi.org/10.1007/978-3-642-17511-4_20
- [60] Z. Manna and A. Pnueli, *The Temporal Logic of Reactive and Concurrent Systems*. New York, NY, USA: Springer-Verlag New York, Inc., 1992.
- [61] Z. Manna, *Mathematical theory of computation*. McGraw-Hill, 1974.
- [62] B. Finkbeiner and S. Schewe, “Bounded synthesis,” *International Journal on Software Tools for Technology Transfer*, vol. 15, no. 5-6, pp. 519–539, 2013. [Online]. Available: dx.org/10.1007/s10009-012-0228-z
- [63] U. Klein, N. Piterman, and A. Pnueli, “Effective synthesis of asynchronous systems from gr (1) specifications,” in *Verification, Model Checking, and Abstract Interpretation*. Springer, 2012, pp. 283–298. [Online]. Available: dx.org/10.1007/978-3-642-27940-9_19

- [64] A. Pnueli and U. Klein, “Synthesis of programs from temporal property specifications,” in *Formal Methods and Models for Co-Design, 2009. MEMOCODE '09. 7th IEEE/ACM International Conference on*, July 2009, pp. 1–7. [Online]. Available: doi.org/10.1109/MEMCOD.2009.5185372
- [65] N. Halbwachs, *Synchronous Programming of Reactive Systems*. Springer, 1992.
- [66] F. Ciesinski, C. Baier, M. Größer, and D. Parker, “Generating compact mtbdd-representations from probmela specifications,” in *Model Checking Software*, ser. Lecture Notes in Computer Science, K. Havelund, R. Majumdar, and J. Palsberg, Eds. Springer, 2008, vol. 5156, pp. 60–76. [Online]. Available: doi.org/10.1007/978-3-540-85114-1_7
- [67] M. Baldamus and J. Schröder-Babo, “P2B: A translation utility for linking promela and symbolic model checking (tool paper),” in *Proceedings of the 8th International SPIN Workshop on Model Checking of Software*, ser. SPIN '01. New York, NY, USA: Springer-Verlag New York, Inc., 2001, pp. 183–191. [Online]. Available: doi.org/10.1007/3-540-45139-0_11
- [68] Y. Jiang and Z. Qiu, “S2N: Model transformation from spin to nusmv,” in *Proceedings of the 19th International Conference on Model Checking Software*, ser. SPIN'12. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 255–260. [Online]. Available: doi.org/10.1007/978-3-642-31759-0_20
- [69] D. Peled and T. Wilke, “Stutter-invariant temporal properties are expressible without the next-time operator,” *Inf. Process. Lett.*, vol. 63, no. 5, pp. 243–246, Sep. 1997. [Online]. Available: doi.org/10.1016/S0020-0190(97)00133-6
- [70] K. L. McMillan, “Symbolic model checking: An approach to the state explosion problem,” Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, USA, 1992, uMI Order No. GAX92-24209. [Online]. Available: http://www.kenmcmil.com/pubs/thesis.pdf
- [71] “open PROMELA implementation.” [Online]. Available: http://github.com/johnyf/openpromela
- [72] “PROMELA parser in PYTHON.” [Online]. Available: http://github.com/johnyf/promela
- [73] D. M. Beazley. Ply (python lex-yacc) v3.4. [Online]. Available: http://www.dabeaz.com/ply/ply.html
- [74] D. Kroening and O. Strichman, *Decision procedures: an algorithmic point of view*. Springer Science & Business Media, 2008.
- [75] V. Beaudenon, E. Encrenaz, and S. Taktak, “Data decision diagrams for promela systems analysis,” *International journal on software tools for technology transfer*, vol. 12, no. 5, pp. 337–352, 2010. [Online]. Available: doi.org/10.1007/s10009-010-0135-0
- [76] S. Maoz and Y. Sa’ar, “AspectLTL: An aspect language for LTL specifications,” in *Proceedings of the Tenth International Conference on Aspect-oriented Software Development*, ser. AOSD '11. New York, NY, USA: ACM, 2011, pp. 19–30. [Online]. Available: doi.org/10.1145/1960275.1960280
- [77] E. Najm and F. Olsen, “Reactive EFSMs — Reactive Promela/RSPIN,” in *Tools and Algorithms for the Construction and Analysis of Systems*, ser. Lecture Notes in Computer Science, T. Margaria and B. Steffen, Eds. Springer, 1996, vol. 1055, pp. 349–368. [Online]. Available: doi.org/10.1007/3-540-61042-1_54
- [78] —, “Protocol verification with Reactive PROMELA/RSPIN,” in *Second SPIN Workshop*, 1996. [Online]. Available: spinroot.com/spin/Workshops/ws96/Ol.pdf
- [79] A. Gamatié, *Designing embedded systems with the Signal programming language: synchronous, reactive specification*. Springer, 2009. [Online]. Available: doi.org/10.1007/978-1-4419-0941-1
- [80] M. Jourdan, F. Lagnier, R. Maraninchi, and P. Raymond, “A multiparadigm language for reactive systems,” in *Computer Languages, 1994., Proceedings of the 1994 International Conference on*. IEEE, 1994, pp. 211–218. [Online]. Available: doi.org/10.1109/ICCL.1994.288379

- [81] M. Baldamus and K. Schneider, “Extending estereL by asynchronous concurrency,” in *Fachtagung zum Entwurf Integrierter Schaltungen*, 1999. [Online]. Available: <http://es.cs.uni-kl.de/publications/datarsg/BaSc99.pdf>
- [82] B. N. Freeman-Benson, “Kaleidoscope: Mixing objects, constraints, and imperative programming,” *SIGPLAN Not.*, vol. 25, no. 10, pp. 77–88, Sep. 1990. [Online]. Available: doi.org/10.1145/97946.97957
- [83] B. N. Freeman-Benson and A. Borning, “Integrating constraints with an object-oriented language,” in *Proceedings of the European Conference on Object-Oriented Programming*, ser. ECOOP ’92. London, UK, UK: Springer-Verlag, 1992, pp. 268–286. [Online]. Available: <http://dl.acm.org/citation.cfm?id=646150.679346>
- [84] B. Freeman-Benson and A. Borning, “The design and implementation of kaleidoscope’90—a constraint imperative programming language,” in *Computer Languages, 1992., Proceedings of the 1992 International Conference on*, Apr 1992, pp. 174–180. [Online]. Available: doi.org/10.1109/ICCL.1992.185480
- [85] G. Lopez, B. Freeman-Benson, and A. Borning, “Implementing constraint imperative programming languages: The KALEIDOSCOPE’93 virtual machine,” in *Proceedings of the Ninth Annual Conference on Object-oriented Programming Systems, Language, and Applications*, ser. OOPSLA ’94. New York, NY, USA: ACM, 1994, pp. 259–271. [Online]. Available: doi.org/10.1145/191080.191118
- [86] B. N. Freeman-Benson, “Constraint imperative programming,” Ph.D. dissertation, University of Washington, Seattle, WA, USA, 1992, uMI Order No. GAX92-03256.
- [87] L. de Alfaro and R. Majumdar, “Quantitative solution of omega-regular games,” *Journal of Computer and System Sciences*, vol. 68, no. 2, pp. 374 – 397, 2004, special Issue on {STOC} 2001. [Online]. Available: doi.org/10.1016/j.jcss.2003.07.009
- [88] T. Felgentreff, A. Borning, and R. Hirschfeld, “Babelsberg: Specifying and solving constraints on object behavior,” *Journal of Object Technology*, vol. 13, no. 4, pp. 1–38, September 2013. [Online]. Available: doi.org/10.5381/jot.2014.13.4.a1