

Foundations of Boolean Stream Runtime Verification[☆]

Laura Bozzelli^a, César Sánchez^b

^a*Technical University of Madrid (UPM), Madrid, Spain*

^b*IMDEA Software Institute, Madrid, Spain*

Abstract

Stream runtime verification (SRV), pioneered by the tool LOLA, is a declarative formalism to specify synchronous monitors. In SRV, monitors are described by specifying dependencies between output streams of values and input streams of values. The declarative nature of SRV enables a separation between the evaluation algorithms, and the monitor storage and its individual updates. This separation allows SRV to be lifted from conventional failure monitors into richer domains to collect statistics of traces. Moreover, SRV allows to easily identify specifications that can be efficiently monitored online, and to generate efficient schedules for offline monitors.

In spite of these attractive features, many important theoretical problems about SRV are still open. In this paper, we address complexity, expressiveness, succinctness, and closure issues for the subclass of Boolean SRV (BSRV) specifications. Additionally, we show that for this subclass, offline monitoring can be performed with only two passes (one forward and one backward) over the input trace in spite of the alternation of past and future references in the BSRV specification.

Keywords: Complexity and expressiveness, Succinctness, Efficiency in closure operations, Offline monitoring

1. Introduction

Runtime verification (RV) has emerged in the last decades as an applied formal technique for software reliability. In RV, a specification expresses correctness requirements and is automatically translated into a *monitor*. Such a monitor is then used

[☆]This paper is an extended version of the results published in the Proceedings of the 5th International Conference on Runtime Verification (RV'14), LNCS 8734, pages 64–79, Springer, 2014.

Email addresses: laura.bozzelli@fi.upm.es (Laura Bozzelli), cesar.sanchez@imdea.org (César Sánchez)

to check either the current execution of a running system, or a finite set of recorded executions with respect to the given specification. The former scenario is called *on-line* monitoring, while the latter one is called *offline* monitoring. Online monitoring is used to detect and possibly handle violations of the specification when the system is in operation, for example, by the execution of additional repair code. On the other hand, offline monitoring is used in post-mortem analysis and it is convenient for testing large systems before deployment, or to inspect system logs. Unlike static verification techniques like model-checking, which formally checks that all the infinite executions or traces of a system satisfy the specification, RV only considers a single finite trace. Thus, this methodology sacrifices completeness guarantees to obtain an immediately applicable and formal extension of testing. See [1, 2] for modern surveys on runtime verification.

Stream runtime verification and related work. The first specification formalisms proposed for runtime verification were based on specification languages for static verification, typically LTL [3] or past LTL adapted for finite paths [4, 5, 6].

Other formalisms for expressing monitors include regular expressions [7], rule based specifications as proposed in the system Eagle [8], or rewriting [9]. Stream runtime verification (SRV), first proposed in the tool LOLA [10], is an alternative to define monitors for synchronous systems.

In SRV, specifications declare explicitly the dependencies between *input* streams of values (representing the observable behavior of the system) and *output* streams of values (describing error reports and diagnosis information). These dependencies can relate the current value of an output stream with the values of the same or other streams in the present moment, in past instants (like in past temporal formulas) or in future instants. A similar approach to describe temporal relations as streams was later introduced as temporal testers [11]. More modernly, the semantics of some temporal logics for continuous signals, like STL (see e.g. [12]) are defined in terms of the relation between the signals defined for an STL formula and the signals assumed for the subexpressions, in a similar manner as for SRV.

Stream runtime verification offers two advantages for the description of monitors. First, SRV separates the algorithmic aspects of the runtime evaluation—by explicitly declaring the data dependencies—from the specific individual operations performed at each step in these evaluation algorithms—which depend on the type of data being observed, manipulated and stored. In this manner, well-known evaluation algorithms for monitoring Boolean observations—for example those adapted from temporal logics—can be generalized to richer data domains, producing monitors that collect statistics about traces. Similarly to the Boolean case, the first approaches for collecting statistics from running traces were based on extensions of LTL or automata [13]. SRV can be viewed as a generalization of these approaches to streams. Other modern approaches to

the runtime verification for statistic collection extend first-order LTL [14, 15, 16]. Moreover, the declarative nature of SRV allows to identify specifications that are amenable for efficient online monitoring, essentially those specifications whose values can be resolved by past and present observations. Additionally, the analysis of dependencies also allows to generate offline monitors by scheduling passes over the dumped traces, where the number of passes (back and forth) depends on the number of alternations between past and future references in the specification.

SRV can be seen as a variation of synchronous languages [17]—like Esterel [18], Lustre [19] or Signal [20]—but specifically designed for observing traces of systems, by removing the causality assumption. In synchronous languages, stream values can only depend on past or present values, while in SRV a dependency on future values is additionally allowed to describe future temporal observations. In recent years, SRV has also been extended to real-time systems [21, 22] in the system Copilot, developed by Galois and NASA.

When used for synthesizing monitors, SRV specifications need to be *well-defined*: for every input there must be a unique corresponding output stream. However, as with many synchronous languages, the declarative style of SRV permits to write syntactically correct specifications that are not well-defined: for some observations, either there is no possible output (*over-definedness*) or there is more than one output (*under-definedness*). This anomaly is caused by circular dependencies. In [10], a syntactical constraint called *well-formedness* was introduced in order to ensure the absence of circular dependencies, and guarantee well-definedness. Natural specifications, written by engineers or translated from specifications in temporal logic and similar formalisms, are usually well-formed and hence well-defined. However, many practical questions that specification engineers ask—like whether a specification is consistent, universal or satisfiable, or whether two specifications are equivalent—can be reduced to the decision problems that we study in this paper, including checking well-definedness.

Symbolic transducers [23, 24] have been recently introduced as an extension to finite state automata and transducers that annotate transitions with logical formulae on the input values, to model sets of concrete transitions. Like SRV, symbolic transducers allow to model complex data in the input and to produce complex data as output. The main difference is that symbolic transducers do not allow to relate or compare inputs produced at different instants, as the only information that the transducer is allowed to store is its finite state. On the other hand, SRV allows to manipulate and store intermediate streams and relate stream values at different instants. The price to pay is, of course, undecidability when manipulating complex data. Extending symbolic transducers with the ability to relate inputs at different positions leads to undecidable decision problems [24]. Similarly, decision problems for general SRV are also undecidable when one allows rich enough data to be manipulated and stored in the streams. In this paper, we limit the data to Boolean values, but (unlike symbolic

transducers) we allow to relate streams at different instants, and study the complexity of the corresponding problems.

Another related line of work introduced recently is regular string transformations [25] and the language DReX, which allows to define in a controlled manner a subset of string transformations. SRV (on characters as input data) can also be used to define string transformations, but the synchronous model of computation of SRV restricts the output to have the same length as the input. However, general SRV is not restricted to characters and allows to define richer (synchronous) specifications. Moreover, the evaluation algorithms of DReX require space depending on the size of the input string, while for SRV it is known how to schedule the offline evaluation of a given specification using an amount of memory that is independent on the size of the input, and is only constant on the size of the spec. More generally, extending stream runtime verification to define string transformations, and in particular non-synchronous string transformations, is an unexplored area of research.

Our contribution. In spite of its applicability, several foundational theoretical problems of SRV have not been studied so far. In this paper, we address complexity, expressiveness, succinctness, and closure properties for Boolean SRV. Our results can be summarized as follows.

- We establish the complexity of checking whether a specification is under-defined, over-defined or well-defined. Apart from the theoretical significance of these results, many important practical properties of specifications can be reduced to the decision problems above. For example, our results provide algorithms to check whether two specifications are equivalent, or whether a part of a specification is redundant because it is subsumed by another part of the specification.
- BSRV specifications can be naturally interpreted as language recognizers, where one defines a language by selecting the inputs for which the specification admits some output. We prove that in this setting, BSRV captures precisely the class of regular languages. We also show efficient closure constructions for many language operations. Additionally, BSRV specifications can be exponentially more succinct than nondeterministic finite-state automata (NFA).
- Finally, based on the construction of the NFA associated with a *well-defined* BSRV specification, we show how to schedule an offline algorithm with only two passes, one forward and one backward. This gives a partial answer (for the Boolean case) to the open problem of reducing the number of passes in offline monitoring for *well-formed* SRV specifications [10].

The rest of the paper is structured as follows. Section 2 revisits SRV. In Section 3 we establish expressiveness, succinctness, and closure results for BSRV specifications

when interpreted as language recognizers. In Section 4, we describe the two-pass offline monitoring algorithm for well-defined BSRV specifications. Section 5 is devoted to the decision problems for BSRV specifications. Finally, Section 6 concludes.

2. Stream Runtime Verification (SRV)

In this Section, we recall the SRV framework [10]. We focus on SRV specifications over stream variables of the same type (with emphasis on the Boolean type).

Informally, an SRV specification describes a relation between input and output streams, where a stream is a finite sequence of values from a given type T (the i^{th} value in the sequence represents the value of the stream at time step i). In the specification, each stream is referred by a variable (denoting the value of the stream at the current time step), and the relation between input and output streams is described by making use of stream expressions, whose building blocks are:

- variables and constants;
- function application;
- offset expressions for referring to the value of a stream at a future/past time with a specified offset from the current time.

Now, we proceed with the formal definition of the SRV specification language.

For all natural numbers i and j with $i \leq j$, we denote by $[i, j]$ the set of natural numbers h such that $i \leq h \leq j$.

A type T is a tuple $T = \langle D, F \rangle$ consisting of a countable value domain D and a finite collection F of interpreted function symbols f , where f denotes a computable function from D^k to D and $k \geq 0$ is the specific arity of f . Note that 0-ary function symbols (constants) are associated with individual values. In particular, we consider the *Boolean type*, where $D = \{0, 1\}$ and F consists of the Boolean operators \wedge , \vee and \neg . A *stream of type T* is a non-empty *finite* word w over the domain D of T . Given such a stream w , $|w|$ is the length of w . For all positions $1 \leq i \leq |w|$, $w(i)$ is the i^{th} letter of w (the value of the stream at time step i). The stream w is *uniform* if there is $d \in D$ such that w is in d^* .

For a finite set Z of (stream) variables, a *stream valuation of type T over Z* is a mapping σ assigning to each variable $z \in Z$, a stream $\sigma(z)$ of type T such that the streams associated with the different variables in Z have the same length N for some $N \geq 1$. We also say that N is the length of σ , which is denoted by $|\sigma|$.

Remark 1. Note that for the Boolean type, a stream valuation σ over Z can be identified with the non-empty word over 2^Z of length $|\sigma|$ whose i^{th} symbol, written $\sigma(i)$, is the set of variables $z \in Z$ such that $\sigma(z)(i) = 1$.

Stream Expressions. Given a finite set Z of variables, the set of *stream expressions* E of type T over Z is inductively defined by the following syntax:

$$E := \tau \mid \tau[\ell|c] \mid f(E_1, \dots, E_k)$$

where τ is either a constant of type T or a variable in Z , ℓ is a non-null integer, c is a constant of type T , and $f \in F$ is a function of type T and arity $k > 0$. Informally, $\tau[\ell|c]$ is an offset expression which refers to the value of τ offset ℓ positions from the current position, and the constant c is the *default* value of type T , which is assigned to positions from which the offset falls after the end or before the beginning of the stream. Stream expressions E of type T over Z are interpreted over stream valuations σ of type T over Z . The *valuation* of E with respect to σ , written $\llbracket E, \sigma \rrbracket$, is the stream of type T and length $|\sigma|$ inductively defined as follows for all $1 \leq i \leq |\sigma|$:

- **Constants:** $\llbracket c, \sigma \rrbracket(i) = c$
- **Variables:** $\llbracket z, \sigma \rrbracket(i) = \sigma(z)(i)$ for all $z \in Z$
- **Offsets:** $\llbracket \tau[\ell|c], \sigma \rrbracket(i) = \begin{cases} \llbracket \tau, \sigma \rrbracket(i + \ell) & \text{if } 1 \leq i + \ell \leq |\sigma| \\ c & \text{otherwise} \end{cases}$
- **Expressions:** $\llbracket f(E_1, \dots, E_k), \sigma \rrbracket(i) = f(\llbracket E_1, \sigma \rrbracket(i), \dots, \llbracket E_k, \sigma \rrbracket(i))$

For the *Boolean type*, we use some shortcuts: $E_1 \rightarrow E_2$ stands for $\neg E_1 \vee E_2$, $E_1 \leftrightarrow E_2$ stands for $(E_1 \rightarrow E_2) \wedge (E_2 \rightarrow E_1)$, and *if* E *then* E_1 *else* E_2 stands for $(E \wedge E_1) \vee (\neg E \wedge E_2)$. Additionally, we use *first* for the Boolean stream expressions $0[-1|1]$ and we use *last* for $0[+1|1]$. Note that for a Boolean stream, *first* is 1 precisely at the first position (and 0 elsewhere), and *last* is 1 precisely at the last position (and 0 elsewhere).

Example 1. Consider the following Boolean stream expression E over $Z = \{x\}$:

$$E := \text{if } x \text{ then } x \text{ else } x[1|0]$$

Consider a Boolean stream valuation σ over $\{x\}$ such that $\sigma(x) \in (01)^+$. The valuation of E with respect to σ is the uniform Boolean stream $1^{|\sigma|}$. \square

Stream Runtime Verification specification language (SRV). Given a finite set X of *input* variables and a set $Y = \{y_1, \dots, y_n\}$ of *output* variables with $X \cap Y = \emptyset$, an SRV specification φ of type T over X and Y is a set of *defining* equations

$$\varphi : \{y_1 := E_1, \dots, y_n := E_n\}$$

where E_1, \dots, E_n are stream expressions of type T over $X \cup Y$. Note that there is exactly one equation for each output variable, and that each defining equation can use both

input and output variables. The intended meaning of a defining equation is to provide the value of the corresponding output stream at every given position. Before formally defining the semantics of SRV, we give some examples as motivation.

Example 2. Consider the specification “*every request must be eventually followed by a grant before the trace ends.*” This specification can be expressed by the following SRV specification with Boolean input streams `request` and `grant`, output streams `reqgrant` and `evgrant`, and the following defining equations:

$$\begin{aligned} \text{reqgrant} &:= \text{if request then evgrant else } 0 \\ \text{evgrant} &:= \text{grant} \vee \text{evgrant}[+1|0] \end{aligned}$$

Essentially, the stream `reqgrant` holds at a given position whenever either there is not a request, or there is a grant in a future position, as stated by the stream `evgrant`. This specification corresponds to the LTL expression $\Box(\text{request} \rightarrow \Diamond \text{grant})$. For a given input trace, the stream `reqgrant` is true exactly at those positions where the LTL formula holds. \square

Example 3. The specification in the previous example contains a single positive offset. It is possible, as shown in [10], to write an equivalent specification that only uses negative offsets:

$$\begin{aligned} \text{waitgrant} &:= \neg \text{grant} \wedge (\text{request} \vee \text{waitgrant}[-1|0]) \\ \text{ok} &:= \text{last} \rightarrow \neg \text{waitgrant} \end{aligned}$$

The output stream `waitgrant` captures whether no grant has been produced since the last request. By evaluating `waitgrant` at the end of the trace, one can obtain whether there is a pending request. \square

Example 4. Let us revisit the specification from the previous examples. One criticism is that a single grant suffices to match several previous requests. A refined version of the intended specification states that “*every request has a response, and every response happens to satisfy a request*”. One way to express this specification is by using integers as a type for intermediate streams `countreq` and `countgrant`, with defining equations:

$$\begin{aligned} \text{countreq} &:= \text{countreq}[-1|0] + (\text{if request then } 1 \text{ else } 0) \\ \text{countgrant} &:= \text{countgrant}[-1|0] + (\text{if grant then } 1 \text{ else } 0) \\ \text{ok} &:= (\text{countgrant} \leq \text{countreq}) \wedge (\text{last} \rightarrow (\text{countreq} - \text{countgrant}) = 0) \end{aligned}$$

Essentially, the stream `countreq` with type integer counts the number of requests that have been seen in the past, and similarly `countgrant` counts the number of grants. The Boolean stream `ok` states that at any point, the number of grants cannot be higher than the number of requests and at the end of the trace requests and grants must match. \square

Example 5. The previous specification can be refined to store the pending requests, that are removed when responses with the same identifier are received. Let `reqid` and `grantid` be input streams of identifiers for actual requests and grants. We define an output stream `pending` whose type is the range of finite sets of identifiers, and a Boolean output stream `ok` with the following defining equations:

$$\begin{aligned} \text{pending} &:= \text{pending}[-1|\emptyset] \cup (\text{if request then } \{\text{reqid}\} \text{ else } \emptyset) \\ &\quad \setminus (\text{if grant then } \{\text{grantid}\} \text{ else } \emptyset) \\ \text{ok} &:= \text{last} \rightarrow (\text{pending} = \emptyset) \end{aligned}$$

Essentially, the stream `pending` stores at each position those requests that have not been granted yet. Hence, one requires that, at the end of the trace, the set of pending requests is empty. In this manner, this specification computes the set of pending requests. Additionally, one can ask that grants only arrive on pending requests using the following output stream `goodgrant`:

$$\begin{aligned} \text{goodgrant} &:= (\neg \text{grant}) \vee \left(\begin{array}{c} \text{grantid} = \text{reqid} \\ \wedge \\ \text{grantid} \notin \text{pending} \end{array} \right) \vee \left(\begin{array}{c} \text{grantid} \neq \text{reqid} \\ \wedge \\ \text{grantid} \in \text{pending} \end{array} \right) \\ \text{ok}_2 &:= \text{ok} \wedge \text{goodgrant} \end{aligned}$$

The stream `goodgrant` guarantees that every grant satisfies a unique request, by checking that the grant satisfies either an immediate request or a past request (but not both).
□

Example 6. (From [26]) We consider a simple latch, as described in [27] with a single Boolean input and a single Boolean output. Whenever the input is true the output is reversed with respect to the previous state. This can be accomplished with the following specification with input `x`, output `y`:

$$y := \text{if } x \text{ then } \neg y[-1|0] \text{ else } y[-1|0]$$

□

Example 7. (Resettable counter, From [26]) Consider a resettable counter with two Boolean inputs, `inc` and `reset`. The input `inc` increments the counter and the input `reset` resets the counter. The counter is modeled by a stream `cnt` of type integer that is initially set to zero. The defining equation for `cnt` is:

$$\text{cnt} := \text{if reset then } 0 \text{ else } (\text{cnt}[-1|0] + \text{if inc then } 1 \text{ else } 0)$$

□

Semantics of SRV. A stream valuation of a specification $\varphi : \{y_1 = E_1, \dots, y_n = E_n\}$ is a stream valuation of type T over $X \cup Y$, while an *input* of φ is a stream valuation of type T over X and an *output* of φ is a stream valuation of type T over Y . Given an input σ_X of φ and an output σ_Y of φ such that σ_X and σ_Y have the same length, $\sigma_X \cup \sigma_Y$ denotes the stream valuation of φ defined in the obvious way. The SRV specification φ describes a relation, written $\llbracket \varphi \rrbracket$, between inputs σ_X of φ and outputs σ_Y of φ , defined as follows: $(\sigma_X, \sigma_Y) \in \llbracket \varphi \rrbracket$ iff $|\sigma_X| = |\sigma_Y|$ and for each equation $y_j = E_j$ of φ ,

$$\llbracket y_j, \sigma \rrbracket = \llbracket E_j, \sigma \rrbracket \quad \text{where } \sigma = \sigma_X \cup \sigma_Y$$

If $(\sigma_X, \sigma_Y) \in \llbracket \varphi \rrbracket$, we say that the stream valuation $\sigma_X \cup \sigma_Y$ is a *valuation model* of φ (associated with the input σ_X). Note that in general, for a given input σ_X , there may be zero, one, or multiple valuation models associated with σ_X . This leads to the following notions for an SRV specification φ :

- *Under-definedness*: for some input σ_X , there are at least two distinct valuation models of φ associated with σ_X . In this case we also say that φ is under-defined for σ_X .
- *Over-definedness*: for some input σ_X , there is no valuation model of φ associated with σ_X . In this case we also say that φ is over-defined for σ_X .
- *Well-definedness*: for each input σ_X , there is exactly one valuation model of φ associated with σ_X .

Note that an SRV specification φ may be both under-defined and over-defined (for different inputs), and φ is well-defined iff it is neither under-defined nor over-defined. For runtime verification, SRV serves as a query language on program behaviors (input streams) from which one computes a unique answer (the output streams). In this context, a specification is useful only if it is well-defined. However, in practice, it is convenient to distinguish *intermediate* output variables from *observable* output variables separating output streams that are of interest to the user from those that are used only to facilitate the computation of other streams. This leads to a more general notion of well-definedness. Given a subset $Z \subseteq Y$ of output variables, an SRV specification φ is *well-defined with respect to Z* if for each input σ_X , there is exactly one stream valuation σ_Z over Z of the same length as σ_X such that $\sigma_X \cup \sigma_Z$ can be extended to some valuation model of φ (*uniqueness of the output streams over Z*).

Analogously, we consider a notion of semantic equivalence between SRV specifications of the same type and having the same input variables, which is parameterized by a set of output variables. Formally, given an SRV φ of type T over X and Y , an SRV specification φ' of type T over X and Y' , and $Z \subseteq Y \cap Y'$, we say that φ and φ' are equivalent with respect to Z if for each valuation model σ of φ , there is a valuation

model σ' of φ' such that σ and σ' coincide on $X \cup Z$, and vice-versa. Moreover, if $Y' \supseteq Y$, then we say that φ' is φ -equivalent if φ and φ' are equivalent with respect to Y .

Remark 2. In the rest of the paper, we focus on Boolean SRV (BSRV for short). Thus, in the following, we omit the reference to the type T in the various definitions. For the complexity analysis, we assume that the offsets ℓ in the subexpressions $\tau[\ell|c]$ of a BSRV are encoded in unary. For a Boolean stream expression E , we denote by $\|E\|$ the absolute value of offset ℓ if E is a stream expression of the form $\tau[\ell|c]$, and we let $\|E\|$ be defined as 1 otherwise. The size $|\varphi|$ of a BSRV φ is defined as $|\varphi| := \sum_{E \in SE(\varphi)} \|E\|$, where $SE(\varphi)$ is the set of stream subexpressions of φ . Essentially the size of an expression φ is the size of its encoding when offsets are written in unary.

Example 8. Consider the following BSRV specifications over $X = \{x\}$ and $Y = \{y\}$:

$$\varphi_1 : \{y := x \wedge y\} \quad \varphi_2 : \{y := x \wedge \neg y\} \quad \varphi_3 : \{y := \text{if } x \text{ then } x[2|0] \text{ else } x[-2|0]\}$$

The specification φ_1 is under-defined since $(1^N, 0^N)$ and $(1^N, 1^N)$ are two valuation models for each $N \geq 1$. On the other hand, the specification φ_2 is over-defined since for each $N \geq 1$, there is no valuation model associated with the input 1^N . Finally, the specification φ_3 is well-defined. \square

Example 9. Consider a scenario where we have two input variables `start` and `end` and an observable output variable `y`. We intend that for every input σ_X , where $X = \{\text{start}, \text{end}\}$, the output Boolean stream σ_Y for `y` gets the value 1 exactly at those positions i such that i belongs to a *session* of the input, that is an interval of positions $I = [h, k]$ satisfying the following:

- within I , the Boolean stream for `start` gets the value 1 exactly at position h . That is, $\sigma_X(\text{start})(h) = 1$ and $\sigma_X(\text{start})(\ell) = 0$ for all $h < \ell \leq k$;
- within I , the Boolean stream for `end` gets the value 1 exactly at position k . That is, $\sigma_X(\text{end})(k) = 1$ and $\sigma_X(\text{end})(\ell) = 0$ for all $h \leq \ell < k$.

The above requirement is accomplished by a BSRV specification φ which uses two additional intermediate output variables `havestarted` and `willend`, whose associated equations are:

$$\begin{aligned} \text{havestarted} & := (\neg \text{start} \wedge \neg \text{end}) \wedge (\text{havestarted}[-1|0] \vee (\text{start}[-1|0] \wedge \neg \text{end}[-1|1])) \\ \text{willend} & := (\neg \text{start} \wedge \neg \text{end}) \wedge (\text{willend}[+1|0] \vee (\text{end}[+1|0] \wedge \neg \text{start}[+1|1])) \end{aligned}$$

The Boolean stream for the output variable `havestarted` assumes the value 1 exactly at the *inner* positions i (i.e., positions which are neither *start* nor *end* positions) such that

the greatest *start* position j which precedes i exists, and j is also a *non-end* position. Analogously, the Boolean stream for the output variable *willend* assumes the value 1 exactly at the *inner* positions i such that the smallest *end* position j which follows i exists, and j is also a *non-start* position. Finally, the equation for the output variable y (capturing the positions inside the sessions of the input) is defined as follows:

$$y := (\text{start} \wedge \text{end}) \vee (\text{start} \wedge \text{willend}[+1|0]) \vee (\text{end} \wedge \text{havestarted}[-1|0]) \vee (\text{willend} \wedge \text{havestarted})$$

□

3. BSRV as Language Recognizers

BSRV can be interpreted as a simple declarative formalism to specify languages of non-empty finite words. We associate to a BSRV specification φ over X and Y , the language $\mathcal{L}(\varphi)$ of non-empty finite words over 2^X (or, equivalently, input stream valuations) for which the specification φ admits a valuation model. Formally,

$$\mathcal{L}(\varphi) = \{\sigma_X \mid (\sigma_X, \sigma_Y) \in \llbracket \varphi \rrbracket \text{ for some } \sigma_Y\}$$

Example 10. Let $X = \{x\}$, $Y = \{y\}$, and $\varphi : \{y := \text{if } E \text{ then } y \text{ else } \neg y\}$, where

$$E := (\text{first} \rightarrow (x \wedge y)) \wedge (y \rightarrow \neg y[+1|0]) \wedge (\neg y \rightarrow (x[+1|1] \wedge y[+1|1]))$$

A pair (σ_X, σ_Y) is a valuation model of φ precisely whenever the valuation of the stream expression E w.r.t. $\sigma_X \cup \sigma_Y$ is in 1^+ . This happens when $\sigma_X(x)(i) = 1$ for all odd positions i . Hence, $\mathcal{L}(\varphi)$ is the set of Boolean streams which assume the value 1 at the odd positions. Note that this is the finite version of the language that Wolper used to exhibit the limitation in expressive power of LTL [28]. □

In the following, we show that BSRV, as language recognizers, are effectively equivalent to nondeterministic finite automata (NFA) on finite words. While the translation from NFA to BSRV can be done in polynomial time, the converse translation involves an *unavoidable* singly exponential blowup. Moreover, BSRV turn out to be effectively and *efficiently* closed under many language operations.

In order to present our results, we shortly recall the class of NFA on finite words. An NFA \mathcal{A} over a finite input alphabet I is a tuple $\mathcal{A} = \langle Q, q_0, \delta, F \rangle$, where Q is a finite set of states, $q_0 \in Q$ is the initial state, $\delta : Q \times I \rightarrow 2^Q$ is the transition function, and $F \subseteq Q$ is a set of accepting states. The NFA \mathcal{A} is *deterministic* if for all $(q, \iota) \in Q \times I$, $\delta(q, \iota)$ is either empty or a singleton. Given an input word $w \in I^*$, a run π of \mathcal{A} over w is a sequence of states $\pi = q_1, \dots, q_{|w|+1}$ such that q_1 is the initial state and for all $1 \leq i \leq |w|$, $q_{i+1} \in \delta(q_i, w(i))$. The run π is accepting if it leads to an accepting state

(i.e, $q_{|w|+1} \in F$). The language $\mathcal{L}(\mathcal{A})$ accepted by \mathcal{A} is the set of non-empty finite words w over I such that there is an accepting run of \mathcal{A} over w . \mathcal{A} is *universal* if $\mathcal{L}(\mathcal{A}) = I^+$. A language over non-empty finite words is *regular* if it is accepted by some NFA. An NFA is *unambiguous* if for each input word w , there is at most one accepting run on w .

Let us fix a BSRV specification φ on X and Y . In order to build an NFA accepting $\mathcal{L}(\varphi)$, we define an encoding of the valuation models of φ . For this, we associate to φ two parameters, the *back reference distance* $b(\varphi)$ and the *forward reference distance* $f(\varphi)$, which are defined as follows:

$$\begin{aligned} b(\varphi) &= \max(0, \{\ell \mid \ell > 0 \text{ and } \varphi \text{ contains a subexpression of the form } z[-\ell, c]\}) \\ f(\varphi) &= \max(0, \{\ell \mid \ell > 0 \text{ and } \varphi \text{ contains a subexpression of the form } z[\ell, c]\}) \end{aligned}$$

The meaning of $b(\varphi)$ and $f(\varphi)$ is to capture for a stream valuation σ of φ and an expression E of φ , the value of E w.r.t. σ at a time step i is completely specified by the values of σ at time steps j within $i - b(\varphi) \leq j \leq i + f(\varphi)$. We define the following alphabets:

$$A = 2^{X \cup Y} \quad A_{\perp} = A \cup \{\perp\} \quad P_{\varphi} = (A_{\perp})^{b(\varphi)} \times A \times (A_{\perp})^{f(\varphi)}$$

where \perp is a special symbol such that $\perp \notin A$. Note that a stream valuation of φ corresponds to a non-empty finite word over the alphabet A , and the cardinality of P_{φ} is singly exponential in the size of φ . Given an element $p = (a_{-b(\varphi)}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)})$ of P_{φ} , the component a_0 , called the *main value of p* , which intuitively represents the value of some stream valuation σ at some time step. The prefix $a_{-b(\varphi)}, \dots, a_{-1}$ represent the values of σ at the previous $b(\varphi)$ time steps. The suffix $a_1, \dots, a_{f(\varphi)}$ denotes the values of σ in the following $f(\varphi)$ steps. The symbol \perp is used to denote the absence of a previous or future time step. Let τ be either a Boolean constant or a variable in $X \cup Y$, and $a \in A$. Then, the Boolean *value* of τ in a is τ if τ is a constant, otherwise the value is 1 precisely when $\tau \in a$ (that is, when the variable τ is present in a). For a Boolean stream expression E over $X \cup Y$ and an element $p = (a_{-b(\varphi)}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)})$ of P_{φ} , the *value* $\llbracket E, p \rrbracket$ of E with respect to p is the computable Boolean value inductively defined as follows:

- **Constants:** $\llbracket c, p \rrbracket = c$
- **Variables:** $\llbracket z, p \rrbracket =$ the value of z in a_0
- **Offsets:** $\llbracket \tau[\ell|c], p \rrbracket = \begin{cases} \text{the value of } \tau \text{ in } a_{\ell} & \text{if } -b(\varphi) \leq \ell \leq f(\varphi) \text{ and} \\ & a_{\ell} \neq \perp \\ c & \text{otherwise} \end{cases}$
- **Expressions:** $\llbracket f(E_1, \dots, E_k), p \rrbracket = f(\llbracket E_1, p \rrbracket, \dots, \llbracket E_k, p \rrbracket)$

We denote by Q_φ the subset of P_φ consisting of the elements p of P_φ such that for each equation $y = E$ of φ , the value of y with respect to p coincides with the value of E with respect to p . Let $\#$ be an additional special symbol (which will be used as initial state of the NFA associated with φ). An *expanded valuation model* of φ is a word of the form $\# \cdot w$ such that w is a non-empty finite word w over the alphabet Q_φ satisfying the following:

- $w(1)$ is of the form $(\perp, \dots, \perp, a_0, a_1, \dots, a_{f(\varphi)})$;
- $w(|w|)$ is of the form $(a_{-b(\varphi)}, \dots, a_{-1}, a_0, \perp, \dots, \perp)$;
- if $1 \leq i < |w|$ and $w(i) = (a_{-b(\varphi)}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)})$, then there is $d \in A_\perp$ such that $w(i+1)$ is of the form $(a_{-b(\varphi)+1}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)}, d)$.

For an expanded valuation model $\# \cdot w$ of φ , the *associated stream valuation* $\sigma(w)$ is the stream valuation of φ of length $|w|$ whose i -th element is the main value of the i -th element of w .

Example 11. Consider the BSRV φ over $X = \{x\}$ and $Y = \{y\}$ of Example 10. We have that $b(\varphi) = 0$ and $f(\varphi) = 1$. The following word

$$\#, \left(\begin{array}{c|c} \boxed{x:1} & x:0 \\ \hline \boxed{y:1} & y:0 \end{array} \right), \left(\begin{array}{c|c} \boxed{x:0} & x:1 \\ \hline \boxed{y:0} & y:1 \end{array} \right), \left(\begin{array}{c|c} \boxed{x:1} & x:1 \\ \hline \boxed{y:1} & y:0 \end{array} \right), \left(\begin{array}{c|c} \boxed{x:1} & \perp \\ \hline \boxed{y:0} & \perp \end{array} \right)$$

(where at each position the boxed element represents the central value) is an expanded valuation model of φ , whose associated stream valuation is given by

$$\left(\begin{array}{c} x:1 \\ y:1 \end{array} \right), \left(\begin{array}{c} x:0 \\ y:0 \end{array} \right), \left(\begin{array}{c} x:1 \\ y:1 \end{array} \right), \left(\begin{array}{c} x:1 \\ y:0 \end{array} \right)$$

□

By construction, we easily obtain that for an expanded valuation model $\# \cdot w$ of φ , $\sigma(w)$ is a valuation model of φ . More precisely, the following lemma holds.

Lemma 1. The map that assigns to each expanded valuation model $\# \cdot w$ of φ the associated stream valuation $\sigma(w)$ is a bijection between the set of expanded valuation models of φ and the set of valuation models of φ .

By the above characterization of the set of valuations models of a BSRV φ , we easily obtain the following result.

Theorem 1 (From BSRV to NFA). Given a BSRV φ over X and Y , one can construct in singly exponential time an NFA \mathcal{A}_φ over the alphabet 2^X accepting $\mathcal{L}(\varphi)$ whose set of states is $Q_\varphi \cup \{\#\}$. Moreover, for each input σ_X , the set of accepting runs of \mathcal{A}_φ over σ_X is the set of expanded valuation models of φ encoding the valuation models of φ associated with the input σ_X .

PROOF. Recall that Q_φ is the subset P_φ consisting of consistent valuations (all output valuations receive the same value as their defining equations), which constrain the set of values of future and past elements of the tuples.

The NFA \mathcal{A}_φ is defined as $\mathcal{A}_\varphi = \langle Q_\varphi \cup \{\#\}, \#, \delta_\varphi, F_\varphi \rangle$, where F_φ is the set of elements of Q_φ of the form $(a_{-b(\varphi)}, \dots, a_{-1}, a_0, \perp, \dots, \perp)$, and $\delta(p, \iota)$ is defined as follows for all states p and input symbol $\iota \in 2^X$:

- if $p = \#$, then $\delta_\varphi(p, \iota)$ is the set of states of the form $(\perp, \dots, \perp, a_0, a_1, \dots, a_{f(\varphi)})$ such that $a_0 \cap X = \iota$;
- if $p = (a_{-b(\varphi)}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)}) \in Q_\varphi$, then $\delta_\varphi(p, \iota)$ is the set of states of the form $(a_{-b(\varphi)+1}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)}, d)$ for some $d \in A_\perp$ whose main value a_1 satisfies $a_1 \cap X = \iota$.

By construction, for each input σ_X , the set of accepting runs of \mathcal{A}_φ over σ_X coincides with the set of expanded valuation models $\# \cdot w$ of φ such that the stream valuation $\sigma(w)$ is associated with the input σ_X . Thus, by Lemma 1, the result follows. \square

For the converse translation from NFA to BSRV, we show the following.

Theorem 2 (From NFA to BSRV). Given an NFA \mathcal{A} over the input alphabet 2^X , one can construct in polynomial time a BSRV $\varphi_{\mathcal{A}}$ with set of input variables X such that $\mathcal{L}(\varphi_{\mathcal{A}}) = \mathcal{L}(\mathcal{A})$.

PROOF. Let $\mathcal{A} = \langle Q, q_0, \delta, F \rangle$. We construct a BSRV specification $\varphi_{\mathcal{A}}$ over the set of input variables X as follows. First, for each input symbol ι , we introduce a Boolean expression E_ι over X , encoding the input symbol ι , defined as $E_\iota := (\bigwedge_{x \in \iota} x) \wedge (\bigwedge_{x \in X \setminus \iota} \neg x)$. The set Y of output variables of $\varphi_{\mathcal{A}}$ is defined as follows:

$$Y = \bigcup_{q \in Q} \{q\} \cup \{\text{control}\}$$

We associate to each state $q \in Q$, an output variable q , whose defining equation is the trivial one given by $q = q$. The equation for the output variable control is given by

$$\text{control} := \text{if } E_{ev} \text{ then control else } \neg \text{control}$$

where the Boolean stream expression E_{ev} captures precisely the accepting runs of the NFA \mathcal{A} and is defined as follows:

$$\begin{aligned}
E_{\text{ev}} = & \underbrace{\bigvee_{q \in Q} (q \wedge \bigwedge_{p \in Q \setminus \{q\}} \neg p)}_{\text{at each step, } \mathcal{A} \text{ is exactly in one state}} \wedge \underbrace{(\text{first} \longrightarrow q_0)}_{\text{a run of } \mathcal{A} \text{ starts at the initial state}} \wedge \\
& \underbrace{\bigwedge_{q \in Q} \bigwedge_{\iota \in I} ((q \wedge E_\iota) \longrightarrow \bigvee_{p \in \delta(q, \iota)} p[+1|1])}_{\text{the evolution of } \mathcal{A} \text{ is } \delta\text{-consistent}} \wedge \underbrace{(\text{last} \longrightarrow \bigvee_{(q, \iota) \in \{(q, \iota) \mid \delta(q, \iota) \cap F \neq \emptyset\}} (q \wedge E_\iota))}_{\text{the run of } \mathcal{A} \text{ is accepting}}
\end{aligned}$$

By construction, given an input stream valuation σ_X , there is a valuation model of $\varphi_{\mathcal{A}}$ associated with the input σ_X if and only if there is a stream valuation σ associated with the input σ_X such that the valuation of E_{ev} with respect to σ is a uniform stream in 1^+ . In turn, the valuation of E_{ev} with respect to σ is a uniform stream in 1^+ if and only if there is an accepting run of \mathcal{A} over the input σ_X . Hence, the result follows. \square

As a corollary of Theorems 1 and 2, we obtain the following result.

Corollary 1. BSRV, when interpreted as language recognizers, capture the class of regular languages over non-empty finite words.

Succinctness. It turns out that the singly exponential blow-up in Theorem 1 cannot be avoided. To prove this succinctness result we first show a linear time translation from linear temporal logic LTL with past over *finite words*—which captures a subclass of regular languages—into BSRV. Recall that formulas ψ of LTL with past over a finite set AP of atomic propositions are defined as follows:

$$\psi ::= p \mid \neg \psi \mid \psi \vee \psi \mid \bigcirc \psi \mid \ominus \psi \mid \psi \mathcal{U} \psi \mid \psi \mathcal{S} \psi$$

where $p \in AP$ and \bigcirc , \ominus , \mathcal{U} , and \mathcal{S} are the ‘next’, ‘previous’, ‘until’, and ‘since’ temporal modalities. For a finite word w over 2^{AP} and a position $1 \leq i \leq |w|$, the satisfaction relation $(w, i) \models \psi$ is defined as follows:

$$\begin{aligned}
(w, i) \models p & \Leftrightarrow p \in w[i] \\
(w, i) \models \neg \psi & \Leftrightarrow (w, i) \not\models \psi \\
(w, i) \models \psi_1 \vee \psi_2 & \Leftrightarrow (w, i) \models \psi_1 \text{ or } (w, i) \models \psi_2 \\
(w, i) \models \bigcirc \psi & \Leftrightarrow i + 1 \leq |w| \text{ and } (w, i + 1) \models \psi \\
(w, i) \models \ominus \psi & \Leftrightarrow i > 1 \text{ and } (w, i - 1) \models \psi \\
(w, i) \models \psi_1 \mathcal{U} \psi_2 & \Leftrightarrow \text{for some } j \text{ with } i \leq j \leq |w|, (w, j) \models \psi_2 \text{ and} \\
& \text{for all } h \text{ with } i \leq h < j, (w, h) \models \psi_1 \\
(w, i) \models \psi_1 \mathcal{S} \psi_2 & \Leftrightarrow \text{for some } j \text{ with } 1 \leq j \leq i, (w, j) \models \psi_2 \text{ and} \\
& \text{for all } h \text{ with } j < h \leq i, (w, h) \models \psi_1
\end{aligned}$$

The language $\mathcal{L}(\psi)$ of an LTL formula ψ is the set of non-empty finite words w over 2^{AP} such that $(w, 1) \models \psi$.

Proposition 1. LTL with past can be translated in linear time into BSRV.

PROOF. Let ψ be a formula of LTL with past over a finite set AP of atomic propositions. We construct in linear time a BSRV specification φ over the set of input variables $X = AP$ such that $\mathcal{L}(\varphi) = \mathcal{L}(\psi)$. Let $SF(\psi)$ be the set of subformulas of ψ . We introduce the following set of output variables Y in φ :

$$Y = \bigcup_{\theta \in SF(\psi)} \{y_\theta\} \cup \{\text{init}\}$$

Essentially, we associate to each subformula θ of ψ , an output variable y_θ . The intended meaning of these variables is that for an input σ_X , corresponding to a non-empty finite word over 2^{AP} , and a valuation model σ associated with σ_X , and for each time step i , the value of variable y_θ is 1 precisely when θ holds at position i along σ_X . The equations for the output variables are defined as follows, where $p \in AP = X$:

$$\begin{aligned} \text{init} &:= \text{first} \rightarrow (y_\psi \vee \neg \text{init}) \\ y_p &:= p \\ y_{\neg\theta} &:= \neg y_\theta \\ y_{\theta_1 \vee \theta_2} &:= y_{\theta_1} \vee y_{\theta_2} \\ y_{\circ\theta} &:= y_\theta[+1|0] \\ y_{\ominus\theta} &:= y_\theta[-1|0] \\ y_{\theta_1 \mathcal{U} \theta_2} &:= y_{\theta_2} \vee (\neg \text{last} \wedge y_{\theta_1} \wedge y_{\theta_1 \mathcal{U} \theta_2}[+1|1]) \\ y_{\theta_1 \mathcal{S} \theta_2} &:= y_{\theta_2} \vee (\neg \text{first} \wedge y_{\theta_1} \wedge y_{\theta_1 \mathcal{S} \theta_2}[-1|1]) \end{aligned}$$

We show now that the construction is correct by showing $\mathcal{L}(\varphi) = \mathcal{L}(\psi)$. For the inclusion $\mathcal{L}(\varphi) \subseteq \mathcal{L}(\psi)$, let $\sigma_X \in \mathcal{L}(\varphi)$. Hence, there is a valuation model σ of φ associated with the input σ_X . We need to show that $(\sigma_X, 1) \models \psi$. One can easily show by construction and structural induction that for all $\theta \in SF(\psi)$ and positions i along σ_X , $(\sigma_X, i) \models \theta$ if and only if $\sigma(y_\theta)(i) = 1$. Moreover, the equation for the output variable init ensures that $\sigma(y_\psi)(1) = 1$. Hence, $(\sigma_X, 1) \models \psi$. Consequently, $\mathcal{L}(\varphi) \subseteq \mathcal{L}(\psi)$.

For the converse inclusion $\mathcal{L}(\psi) \subseteq \mathcal{L}(\varphi)$, let $\sigma_X \in \mathcal{L}(\psi)$ and consequently $(\sigma_X, 1) \models \psi$. We define a stream valuation σ associated with the input σ_X as follows: $\sigma(\text{init}) = 1^{|\sigma|}$ and for all $\theta \in SF(\psi)$ and positions i along σ_X , $\sigma(y_\theta)(i) = 1$ if $(\sigma_X, i) \models \theta$, and $\sigma(y_\theta)(i) = 0$ otherwise. Since $(\sigma_X, 1) \models \psi$, by construction, it easily follows that σ is a valuation model of φ associated with the input σ_X . Hence, $\sigma_X \in \mathcal{L}(\varphi)$ and consequently $\mathcal{L}(\psi) \subseteq \mathcal{L}(\varphi)$. \square

It is well-known [29] that there is a singly exponential succinctness gap between LTL with past and NFA. Consequently, by Proposition 1, we obtain the following result.

Theorem 3. BSRV specification are singly exponentially more succinct than NFA. In particular, there is a finite set X of input variables and a family $(\varphi_n)_{n \geq 1}$ of BSRV specifications with input variables X such that for all $n \geq 1$, φ_n has size polynomial in n , and every NFA accepting $\mathcal{L}(\varphi_n)$ has at least $2^{\Omega(n)}$ states.

Effective closure under language operations. An interesting feature of the class of BSRV is that, when interpreted as language recognizers, BSRV are effectively and *efficiently* closed under many language operations. For two languages \mathcal{L} and \mathcal{L}' of finite words, \mathcal{L}^R denotes the reversal of \mathcal{L} , $\mathcal{L} \cdot \mathcal{L}'$ denotes the concatenation of \mathcal{L} and \mathcal{L}' , and \mathcal{L}^+ denotes the positive Kleene closure of \mathcal{L} .

For a BSRV φ , we say that an output variable y of φ is *uniform* if for each valuation model of φ , the stream for y is uniform.

Theorem 4. BSRV are effectively closed under the following language operations: intersection, union, reversal, positive Kleene closure, and concatenation. Additionally, the constructions for these operations can be performed in *linear time*.

PROOF. Let $\varphi : \{y_1 := E_1, \dots, y_k := E_k\}$ and $\varphi' : \{y'_1 := E'_1, \dots, y'_h := E'_h\}$ be two BSRV specifications over the same set X of input variables. Let $Y = \{y_1, \dots, y_k\}$ be the set of output variables of φ and $Y' = \{y'_1, \dots, y'_h\}$ the set of output variables of φ' . We assume without loss of generality that the BSRV specifications φ and φ' have no output variable in common (i.e. $Y \cap Y' = \emptyset$). We describe each construction individually.

Intersection. The construction for intersection is illustrated in Fig. 1. The BSRV specification recognizing $\mathcal{L}(\varphi) \cap \mathcal{L}(\varphi')$ is simply the joint set of the equations of φ and φ' . Correctness of the construction immediately follows.

Union. The construction of the BSRV $\varphi \cup \varphi'$ recognizing $\mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$ is given in Fig. 1. Note that we use two new additional output variables: **switch** and **main**. These variables use a common gadget, built by assigning to an output variable y an equation of the form:

$$y := \text{if } \alpha \text{ then } y \text{ else } \neg y$$

This gadget forces the sub-expression α to be true at all positions in every valuation model. In the case of the construction for *union*, **switch** is forced by the sub-expression to be a uniform output variable, which is then used to guess whether the input has to be considered an input for φ or for φ' . Depending on the uniform value of **switch** (if it is in 0^+ or 1^+), the equation for the output variable **main** ensures that the input is

recognized precisely when either the equations of φ are fulfilled or the equations of φ' are fulfilled.

Now, we show that the construction is correct by showing that $\mathcal{L}(\varphi \cup \varphi') = \mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$. For the inclusion $\mathcal{L}(\varphi \cup \varphi') \subseteq \mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$, let $\sigma_X \in \mathcal{L}(\varphi \cup \varphi')$. Hence, there is a valuation model σ'' of $\varphi \cup \varphi'$ associated with the input σ_X . By construction, $\sigma''(\text{switch}) \in 0^+ \cup 1^+$. Assume that $\sigma''(\text{switch}) \in 1^+$ (the other case is analogous). By definition of the equation for `main`, for every $1 \leq i \leq k$, $\llbracket y_i, \sigma'' \rrbracket = \llbracket E_i, \sigma'' \rrbracket$. Thus, the restriction of σ'' to $X \cup Y$ is a valuation model of φ . Hence, $\sigma_X \in \mathcal{L}(\varphi)$ and the result follows. For the converse inclusion $\mathcal{L}(\varphi) \cup \mathcal{L}(\varphi') \subseteq \mathcal{L}(\varphi \cup \varphi')$, let $\sigma_X \in \mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$. Assume that $\sigma_X \in \mathcal{L}(\varphi')$ (the other case being similar). Hence, there is a valuation model σ' of φ' associated with the input σ_X . Let σ'' be the stream valuation of $\varphi \cup \varphi'$ associated with the input σ_X defined as follows: for all variables z of $\varphi \cup \varphi'$, $\sigma''(z) = \sigma'(z)$ if $z \in X \cup Y'$, and $\sigma''(z) \in 0^+$ otherwise. By construction, it follows that σ'' is a valuation model of $\varphi \cup \varphi'$. Hence $\sigma_X \in \mathcal{L}(\varphi \cup \varphi')$.

Reversal. The construction for *reversal* is very simple and appears in Fig. 1. The BSRV φ^R recognizing $\mathcal{L}(\varphi)^R$ is obtained from φ by replacing each equation $y = E$ with the equation $y = E^R$, where E^R denotes the stream expression obtained from E by replacing each sub-expression $\tau[k|d]$ with $k > 0$ for $\tau[-k|d]$, and replacing each sub-expression $\tau[-k|d]$ with $k > 0$ for $\tau[k|d]$.

We prove now that the construction is correct. For a finite word w and a position $1 \leq i \leq |w|$, we denote by w^R the reverse of w , and by $R(i)$ the position of w given by $|w| - i + 1$. Note that the suffix of w^R from position $R(i)$ is the reverse of the prefix of w

$\varphi : \{y_1 := E_1, \dots, y_k := E_k\} \quad \varphi' : \{y'_1 := E'_1, \dots, y'_h := E'_h\}$
<p>Intersection: $\varphi \cap \varphi' : \{y_1 := E_1, \dots, y_k := E_k, y'_1 := E'_1, \dots, y'_h := E'_h\}$ where $\{y_1, \dots, y_k\} \cap \{y'_1, \dots, y'_h\} = \emptyset$.</p>
<p>Union: $\varphi \cup \varphi' : \{y_1 := y_1, \dots, y'_h := y'_h, \text{switch} := E_{\text{switch}}, \text{main} := E_{\text{main}}\}$</p>
$E_{\text{switch}} = \text{if } \neg \text{last} \rightarrow (\text{switch} \leftrightarrow \text{switch}[+1 1]) \text{ then switch else } \neg \text{switch}$
$E_{\text{main}} = \text{if } ((\text{switch} \rightarrow \bigwedge_{i=1}^{i=k} y_i \leftrightarrow E_i) \wedge (\neg \text{switch} \rightarrow \bigwedge_{i=1}^{i=h} y'_i \leftrightarrow E'_i)) \text{ then main else } \neg \text{main}$
<p>Reversal: $\varphi^R : \{y_1 := E_1^R, \dots, y_k := E_k^R\}$ E_i^R is obtained from E_i by converting each offset k in its opposite $-k$.</p>

Figure 1: Constructions for intersection, union, and reversal.

leading to position i . It follows by structural induction that for all stream expressions E , stream valuations σ over the variables of E , and $1 \leq i \leq |\sigma|$:

$$\llbracket E, \sigma \rrbracket(i) = \llbracket E^R, \sigma^R \rrbracket(R(i))$$

Then for all stream valuations σ over $X \cup Y$, σ is a valuation model of φ if and only if σ^R is a valuation model of φ^R . Hence, $\mathcal{L}(\varphi^R) = [\mathcal{L}(\varphi)]^R$.

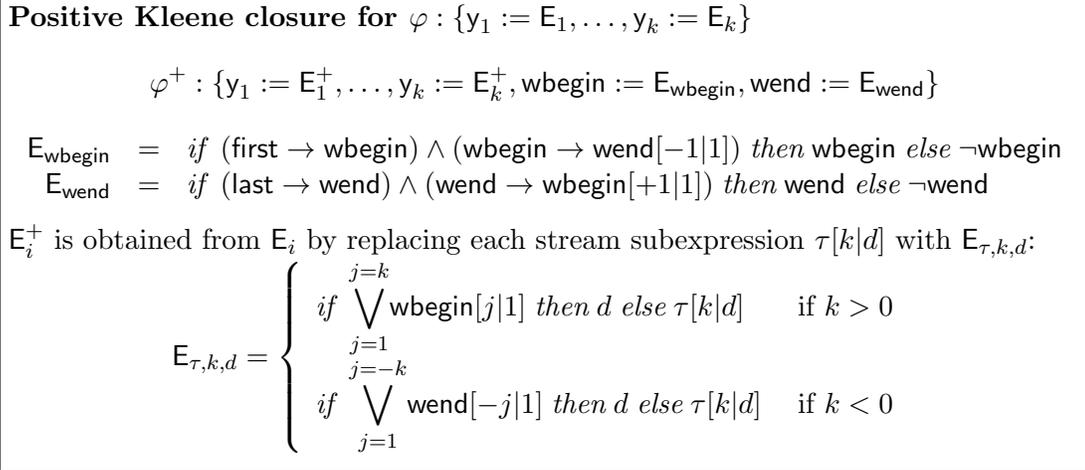


Figure 2: Construction for positive Kleene closure

Positive Kleene closure. The construction is given in Fig. 2. The BSRV specification φ^+ that recognizes $[\mathcal{L}(\varphi)]^+$ uses two new additional output variables: `wbegin` and `wend`, again defined using the gadget described above. Intuitively, `wbegin` and `wend` are used for guessing a decomposition in the given input σ_X of the form $\sigma_X = \sigma_{X,1} \cdot \dots \cdot \sigma_{X,N}$ for some $N \geq 1$ in such a way that each component $\sigma_{X,i}$ is in $\mathcal{L}(\varphi)$. In particular, the output variable `wbegin` is used to mark the first positions of the components $\sigma_{X,i}$, and `wend` is used to mark the last position. The equations for the output variables of φ are modified to allow checking for an offset k of φ and a position j inside a component $\sigma_{X,i}$ in the guessed decomposition of the input σ_X , whether $k + j$ is still a position inside $\sigma_{X,i}$.

Now, we show that the construction is correct by proving that $\mathcal{L}(\varphi^+) = [\mathcal{L}(\varphi)]^+$. For the inclusion $\mathcal{L}(\varphi^+) \subseteq [\mathcal{L}(\varphi)]^+$, let $\sigma_X \in \mathcal{L}(\varphi^+)$. Hence, there is a valuation model σ of φ^+ associated with the input σ_X . By the equations for the output variables `wbegin` and `wend`, there is $N \geq 1$ such that σ can be written in the form $\sigma = \sigma_1 \cdot \dots \cdot \sigma_N$ and for all $1 \leq \ell \leq N$, the Boolean stream $\sigma_\ell(\text{wbegin})$ is in 10^* and the Boolean stream $\sigma_\ell(\text{wend})$ is in 0^*1 . Fix $1 \leq \ell \leq N$. We show that the restriction $(\sigma_\ell)_{X \cup Y}$ of σ_ℓ to

the set of variables $X \cup Y$ is a valuation model of φ . Hence, membership of σ_X in $[\mathcal{L}(\varphi)]^+$ follows. For all positions $1 \leq i \leq |\sigma_\ell|$ along the stream valuation σ_ℓ , we denote by $p(i)$ the corresponding position along σ . In order to prove that σ_ℓ is a valuation model of φ , by hypothesis, it suffices to show that for each equation $y = E$ of φ and for each position $1 \leq i \leq |\sigma_\ell|$, the following holds, where $y = E^+$ is the equation of φ^+ associated to the output variable y :

$$\llbracket E, (\sigma_\ell)_{X \cup Y} \rrbracket(i) = \llbracket E^+, \sigma \rrbracket(p(i))$$

We just need to prove that for each subexpression $\tau[k|d]$ of φ , $\llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i) = \llbracket E_{\tau,k,d}, \sigma \rrbracket(p(i))$, where the stream expression $E_{\tau,k,d}$ is as in Fig. 2. There are two cases:

- $k > 0$: first, assume that $i+k \leq |\sigma_\ell|$. Since $\sigma_\ell(\mathbf{wbegin})$ is in 10^* , we obtain that for all $1 \leq j \leq k$, $\sigma(\mathbf{wbegin})(p(i)+j) = 0$. Hence, by definition of $E_{\tau,k,d}$, it follows that $\llbracket E_{\tau,k,d}, \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i)$, and the result follows in this case. Now, assume that $i+k > |\sigma_\ell|$. Hence, $\llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i) = d$. Then, either $p(i) + k > |\sigma|$, or there is $1 \leq j \leq k$ such that $\sigma(\mathbf{wbegin})(p(i) + j) = 1$. By definition of $E_{\tau,k,d}$, it follows that $\llbracket E_{\tau,k,d}, \sigma \rrbracket(p(i)) = d$, and the result follows in this case as well.
- $k < 0$: first, assume that $i-k \geq 1$. Since $\sigma_\ell(\mathbf{wend})$ is in 0^*1 , we obtain that for all $1 \leq j \leq -k$, $\sigma(\mathbf{wend})(p(i) - j) = 0$. Hence, by definition of $E_{\tau,k,d}$, it follows that $\llbracket E_{\tau,k,d}, \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i)$, and the result follows in this case. Now, assume that $i-k < 1$. Hence, $\llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i) = d$. Then, either $p(i) - k < 1$, or there is $1 \leq j \leq -k$ such that $\sigma(\mathbf{wend})(p(i) - j) = 1$. By definition of $E_{\tau,k,d}$, it follows that $\llbracket E_{\tau,k,d}, \sigma \rrbracket(p(i)) = d$, and the result follows in this case as well.

For the converse inclusion $[\mathcal{L}(\varphi)]^+ \subseteq \mathcal{L}(\varphi^+)$, let $\sigma_X \in [\mathcal{L}(\varphi)]^+$. Hence, there is a stream valuation σ of φ associated with the input σ_X such that σ is of the form $\sigma = \sigma_1 \cdots \sigma_N$ for some $N \geq 1$, and σ_ℓ is a valuation model of φ for all $1 \leq \ell \leq N$. Let $\sigma'' = \sigma''_1 \cdots \sigma''_N$ be the extension of σ over $X \cup Y \cup \{\mathbf{wbegin}, \mathbf{wend}\}$, where the Boolean streams $\sigma''(\mathbf{wbegin})$ and $\sigma''(\mathbf{wend})$ are defined as follows: for all $1 \leq \ell \leq N$, $\sigma''_\ell(\mathbf{wbegin})$ is in 10^* and the Boolean stream $\sigma''_\ell(\mathbf{wend})$ is in 0^*1 . We show that σ'' is a valuation model of φ^+ , hence, membership of σ_X in $\mathcal{L}(\varphi^+)$ follows. By construction, the equations for the output variables \mathbf{wbegin} and \mathbf{wend} are satisfied with respect to the stream valuation σ'' . Now, let us consider an equation $y = E$ of φ^+ associated with an output variable $y \in Y$. By construction, in order to show that $\llbracket y, \sigma'' \rrbracket = \llbracket E, \sigma'' \rrbracket$, it suffices to prove that for all $1 \leq \ell \leq N$, $1 \leq i \leq |\sigma_\ell|$, and subexpression $\tau[k|d]$ of φ , $\llbracket \tau[k|d], \sigma_\ell \rrbracket(i) = \llbracket E_{\tau,k,d}, \sigma'' \rrbracket(p(i))$, where for $1 \leq i \leq |\sigma_\ell|$, $p(i)$ denotes the corresponding position along σ . This can be shown as for the proof of the inclusion $\mathcal{L}(\varphi^+) \subseteq [\mathcal{L}(\varphi)]^+$.

$$\varphi : \{y_1 := E_1, \dots, y_k := E_k\} \quad \varphi' = \{y'_1 := E'_1, \dots, y'_h := E'_h\}$$

Concatenation: $\{y_1, \dots, y_k\} \cap \{y'_1, \dots, y'_h\} = \emptyset$

$$\begin{aligned} \varphi \cdot \varphi' : \{y_1 &:= \text{if wmark then } \tilde{E}_1 \text{ else } y_1, \\ &\dots, \\ y_k &:= \text{if wmark then } \tilde{E}_k \text{ else } y_k, \\ y'_1 &:= \text{if } \neg\text{wmark then } \tilde{E}'_1 \text{ else } y'_1, \\ &\dots, \\ y'_h &:= \text{if } \neg\text{wmark then } \tilde{E}'_h \text{ else } y'_h, \\ \text{wmark} &:= E_{\text{wmark}}\} \end{aligned}$$

$$E_{\text{wmark}} = \text{if } (\text{first} \rightarrow \text{wmark}) \wedge (\text{last} \rightarrow \neg\text{wmark}) \wedge (\text{wmark} \rightarrow \text{wmark}[-1|1]) \wedge (\neg\text{wmark} \rightarrow \neg\text{wmark}[+1|0]) \text{ then wmark else } \neg\text{wmark}$$

\tilde{E}_i is obtained from E_i by replacing each stream subexpression $\tau[k|d]$ s.t. $k > 0$ with:

$$\text{if } \bigvee_{j=1}^{j=k} \neg\text{wmark}[j|0] \text{ then } d \text{ else } \tau[k|d]$$

\tilde{E}'_i is obtained from E'_i by replacing each stream subexpression $\tau[k|d]$ s.t. $k < 0$ with:

$$\text{if } \bigvee_{j=1}^{j=-k} \text{wmark}[-j|1] \text{ then } d \text{ else } \tau[k|d]$$

Figure 3: Construction for concatenation

Concatenation. The construction is given in Fig. 3. The BSRV specification $\varphi \cdot \varphi'$ recognizing $\mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$ uses a new additional output variable, **wmark**, once again based on the same gadget. This variable is used for guessing a decomposition in the given input of the form $\sigma_X \cdot \sigma'_X$ in such a way that $\sigma_X \in \mathcal{L}(\varphi)$ and $\sigma'_X \in \mathcal{L}(\varphi')$. In particular, the stream for the output variable **wmark** assumes the value 1 along all and only the positions of σ_X (the equation for **wmark** ensures that a Boolean stream for **wmark** is always in 1^+0^+). Moreover, the equations for the output variables of φ are modified in order to allow to check for a positive offset $k > 0$ of φ and a position j inside σ_X in the guessed decomposition $\sigma_X \cdot \sigma'_X$ of the input, whether $k + j$ is still a position inside σ_X . Analogously, the equations for the output variables of φ' are modified to allow checking for a negative offset $k < 0$ of φ' and a position j inside σ'_X in the guessed decomposition $\sigma_X \cdot \sigma'_X$ of the input, whether $k + j$ is still a position inside σ'_X .

Now, we show that the construction is correct: $\mathcal{L}(\varphi \cdot \varphi') = \mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$. For the

inclusion $\mathcal{L}(\varphi \cdot \varphi') \subseteq \mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$, let $\sigma_X \in \mathcal{L}(\varphi \cdot \varphi')$. Hence, there is a valuation model σ'' of $\varphi \cdot \varphi'$ associated with the input σ_X . By the equation for the output variable `wmark`, σ'' can be written in the form $\sigma'' = \sigma \cdot \sigma'$ such that $\sigma(\text{wmark})$ is in 1^+ and $\sigma'(\text{wmark})$ is in 0^+ . We show that the restriction $(\sigma)_{X \cup Y}$ of σ to the set of variables $X \cup Y$ is a valuation model of φ , and the restriction $(\sigma')_{X \cup Y'}$ of σ' to the set of variables $X \cup Y'$ is a valuation model of φ' . Hence, membership of σ_X in $\mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$ follows. We consider the stream valuation $(\sigma)_{X \cup Y}$ (the proof for $(\sigma')_{X \cup Y'}$ is similar). In order to prove that $(\sigma)_{X \cup Y}$ is a valuation model of φ , by hypothesis and construction, it suffices to show that for each equation $y = E$ of φ and for each position $1 \leq i \leq |\sigma|$, the following holds, where $y = \text{if wmark then } E'' \text{ else } y$ is the equation of $\varphi \cdot \varphi'$ associated to the variable y :

$$\llbracket E, (\sigma)_{X \cup Y} \rrbracket(i) = \llbracket E'', \sigma'' \rrbracket(i)$$

By construction, we just need to prove that for each subexpression $\tau[k|d]$ of φ such that $k > 0$, $\llbracket \tau[k|d], (\sigma)_{X \cup Y} \rrbracket(i) = \llbracket E_{\tau,k,d}, \sigma'' \rrbracket(i)$, where the stream expression $E_{\tau,k,d}$ is as in Fig. 3. This is analogous to the proof for the positive Kleene closure.

For the converse inclusion $\mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi') \subseteq \mathcal{L}(\varphi \cdot \varphi')$, let $\sigma''_X \in \mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$. Hence, $\sigma''_X = \sigma_X \cdot \sigma'_X$ and there are a valuation model σ_m of φ associated with the input σ_X and a valuation model σ'_m of φ' associated with the input σ'_X . Let σ'' be the stream valuation over $X \cup Y \cup Y' \cup \{\text{wmark}\}$ associated with the input σ''_X defined as follows: $\sigma'' = \sigma \cdot \sigma'$ with $|\sigma| = |\sigma_X|$, where (i) $\sigma(z) = \sigma_m(z)$ if $z \in X \cup Y$, and $\sigma(z) = 1$ otherwise, and (ii) $\sigma'(z) = \sigma'_m(z)$ if $z \in X \cup Y'$, and $\sigma'(z) = 0$ otherwise. By construction σ'' is a valuation model of $\varphi \cdot \varphi'$. Hence, membership of σ''_X in $\mathcal{L}(\varphi \cdot \varphi')$ follows.

This concludes the proof of Theorem 4. \square

4. Offline Monitoring for Well-defined BSRV

In this section, we propose an offline monitoring algorithm for well-defined BSRV based on Theorem 1. The algorithm runs in time linear in the length of the input trace (input streams) and singly exponential in the size of the specification.

Our algorithm also solves the following question which is left open in [10]. In offline monitoring, one can afford to traverse the input stream several times, in the backward and in the forward directions. However, efficient algorithms must perform each traversal using only a bounded amount of memory that does not depend on the length of the trace. The algorithm in [10] for offline monitoring starts by analyzing the specification, calculating the offset dependencies between output streams by looking at their defining equations. Then, the algorithm performs a forward traversal to resolve (in a memory-less manner) the values of an output stream variable based on the values of other output variables that have been already resolved and which appear with past offsets. Similarly, backward traversals are performed to resolve expressions with future

```

Monitoring( $\varphi, \sigma_X$ )    /**  $\varphi$  is a well-defined BSRV and  $\mathcal{A}_\varphi = \langle Q, q_0, \delta, F \rangle$  **/
 $\Lambda \leftarrow \{q_0\}$ 
for  $i = 1$  upto  $|\sigma_X|$  do
  update  $\Lambda \leftarrow \{q \in Q \mid q \in \delta(p, \sigma_X(i)) \text{ for some } p \in \Lambda\}$ 
  store  $\Lambda$  at position  $i$  on the tape
for  $i = |\sigma_X|$  downto 1 do
  let  $\Lambda$  be the set of states stored at position  $i$  on the tape
  if  $i = |\sigma_X|$  then  $p \leftarrow$  the unique accepting state in  $\Lambda$ 
  else let  $q$  be the unique state in  $\Lambda$  such that  $p \in \delta(q, \sigma_X(i+1))$ ;
    update  $p \leftarrow q$ 
  output at position  $i$  the main value of  $p$ 

```

Figure 4: Offline monitoring algorithm for well-defined BSRV

offsets. In summary, the algorithm in [10] performs a number of passes proportional to the number of backward and forward references in the defining equations. Intuitively speaking, this number of passes corresponds to the number of alternations between future and past operators in a temporal logic specification. The open question is whether a specification can be modified into an equivalent specification that only requires a constant number of forward and backward passes, each of which uses an amount of memory that does not depend on the length of the trace. We show here that only two passes (one forward and one backward) are required for BSRV.

Let φ be a BSRV over X and Y , and $\mathcal{A}_\varphi = \langle Q, q_0, \delta, F \rangle$ be the NFA over 2^X accepting $\mathcal{L}(\varphi)$ of Theorem 1. Recall that $Q \setminus \{q_0\}$ is contained in $(A_\perp)^{b(\varphi)} \times A \times (A_\perp)^{f(\varphi)}$, where $A = 2^{X \cup Y}$ and $A_\perp := A \cup \{\perp\}$, and an expanded valuation model of φ is of the form $\pi = q_0, q_1, \dots, q_k$, where $q_i \in Q \setminus \{q_0\}$ for all $1 \leq i \leq k$. The valuation model of φ encoded by π is the sequence of the main values of the states q_i visited by π . By Theorem 1, for every input σ_X , the set of accepting runs of \mathcal{A}_φ over σ_X is the set of expanded valuation models of φ encoding the valuation models of φ associated with the input σ_X . Hence, we obtain the following.

Proposition 2. A BSRV φ is well-defined if and only if the NFA \mathcal{A}_φ is universal and unambiguous.

PROOF. The specification φ is well-defined iff for every input σ_X , φ admits one and only one (expanded) valuation model iff for every input σ_X , there is one and only one accepting run of \mathcal{A}_φ over σ_X iff \mathcal{A}_φ is universal and unambiguous. \square

The offline monitoring algorithm for well-defined BSRV is given in Fig. 4, where we assume that the input trace σ_X is available on an input tape. The algorithm operates

in two phases. In the first phase, a forward traversing of the input trace is performed, and the algorithm simulates the unique run over the input σ_X of the deterministic finite state automaton (DFA) that would result from \mathcal{A}_φ by the classical powerset construction. Let $\{q_0\}, \Lambda(1), \dots, \Lambda(|\sigma_X|)$ be the run of this DFA over σ_X . At each step i , the state $\Lambda(i)$ of the run resulting from reading the input symbol $\sigma_X(i)$ is stored in the i th position of the tape. In the second phase, a backward traversal of the input trace is performed, and the algorithm outputs a stream valuation of φ .

We claim that the uniqueness conditions in the second phase of the algorithm are satisfied, and the output is the unique valuation model of the *well-defined* BSRV specification φ associated with the input σ_X . By Proposition 2, \mathcal{A}_φ is universal. Thus, by the classical construction of the DFA associated with \mathcal{A}_φ , it holds that for all $1 \leq i \leq |\sigma_X|$,

$$\begin{aligned} \Lambda(i) \neq \emptyset \text{ and for each state } q \text{ of } \mathcal{A}_\varphi, \\ q \in \Lambda(i) \text{ if and only if} \\ \text{there is a run of } \mathcal{A}_\varphi \text{ over } \sigma_X(1), \dots, \sigma_X(i) \text{ leading to state } q. \end{aligned} \tag{1}$$

We assume that the uniqueness conditions are not satisfied, and derive a contradiction. Let i be the greatest position along σ_X such that the uniqueness condition at step i is not satisfied. Assume that $i < |\sigma_X|$ (the other case being simpler). For each $q \in Q$, let \mathcal{A}_φ^q be the NFA obtained from \mathcal{A}_φ by replacing the initial state q_0 with q . By Condition (1) above, $\Lambda(i) \neq \emptyset$. Hence, by hypothesis, there are a state $p \in \Lambda(i+1)$ and two distinct states $q, q' \in \Lambda(i)$ such that $p \in \delta(q, \sigma_X(i+1))$ and $p \in \delta(q', \sigma_X(i+1))$. Moreover, by construction of the algorithm, there is a run of \mathcal{A}_φ^p over $\sigma_X(i+2), \dots, \sigma_X(|\sigma_X|)$ leading to an accepting state q_{acc} . By Condition (1) above, we deduce that there are two distinct accepting runs of \mathcal{A}_φ over σ_X . This is a contradiction because by Proposition 2, \mathcal{A}_φ is unambiguous. Hence, the uniqueness conditions in the second phase of the algorithm are satisfied. Moreover, by construction, it follows that the sequence of states computed by the algorithm in the second phase is the unique accepting run π of \mathcal{A}_φ over σ_X . Therefore, the algorithm outputs the valuation model of φ encoded by π , which is the unique valuation model of φ associated with the input σ_X . Thus, since the size of the NFA \mathcal{A}_φ is singly exponential in the size of φ , we obtain the following result.

Theorem 5. One can construct an offline monitoring algorithm for a well-defined BSRV specification, that runs in time linear in the length of the input trace and singly exponential in the size of the specification. Additionally, the algorithm processes a position of the input trace exactly twice.

In [10], a syntactical condition for general SRV, called *well-formedness*, is introduced, which can be checked in polynomial time and guarantees that the semantic condition given by well-definedness is met. Well-formedness ensures the absence of circular definitions by requiring that a dependency graph of the output variables have not

zero-weight cycles. As illustrated in [10], for the restricted class of well-formed SRV, it is possible to construct an offline monitoring algorithm which runs in time linear in the length of the input trace and the size of the specification. Moreover, one can associate to a well-formed SRV φ a parameter $ad(\varphi)$, called *alternation depth* [10], such that the monitoring algorithm processes each position of the input trace exactly $ad(\varphi)+1$ times. An important question left open in [10] is whether for a well-formed SRV specification φ , it is possible to construct a φ -equivalent SRV specification whose alternation depth is minimal. Here, we settle partially this question for the class of BSRV. By using the same ideas for constructing the algorithm of Fig. 4, we show that for the class of BSRV, the semantic notion of well-definedness coincides with the syntactical notion of well-formedness (modulo BSRV-equivalence), and the hierarchy of well-formed BSRV induced by the alternation depth collapses to 1.

We give now the formal details. First, we revisit the notion of well-formedness [10]. Given a general SRV specification φ over X and Y , the *dependency graph* G_φ of φ is the finite weighted directed graph whose set of vertices is Y and whose set of weighted edges is defined as follows. There is an edge $y \xrightarrow{k} z$ in the graph whenever the equation $y = E$ of φ associated with y , either $k = 0$ and z occurs in E , or $k \neq 0$ and $z[k|d]$ occurs in E for some d . The weight of a finite path of G_φ is the sum of the weights of its edges. An SRV specification φ is *well-formed* if its dependency graph G_φ has no cycle with weight zero. As shown in [10], for a well-formed SRV specification φ , a strongly connected component (*SCC*) of G_φ can be classified as positive or negative, where an *SCC* is positive if every cycle in the *SCC* has weight strictly positive, or negative if every cycle is strictly negative. Clearly, every *SCC* is positive or negative because otherwise one can build a zero weight path by traversing a negative and a positive cycle a sufficient number of times to cancel each other. The alternation depth $ad(\varphi)$ of a well-formed SRV is then defined as the maximum over the number of alternations between positive and negative vertices along a path of G_φ , where a vertex is positive if it belongs to a positive *SCC* and negative if it belongs to a negative *SCC*.

We establish the following result for the class of BSRV.

Theorem 6. Given a well-defined BSRV specification φ , one can build in doubly exponential time a φ -equivalent BSRV specification which is well-formed and has alternation depth 1.

PROOF. Let φ be a well-defined BSRV specification over X and Y , and $\mathcal{A}_\varphi = \langle Q, q_0, \delta, F \rangle$ be the NFA over 2^X accepting $\mathcal{L}(\varphi)$ built in the proof of Theorem 1. We denote by \mathcal{D}_φ the DFA accepting $\mathcal{L}(\varphi)$ resulting by determinizing \mathcal{A}_φ using the classical power set construction. Recall that $\mathcal{D}_\varphi = \langle 2^Q, \{q_0\}, \delta_{\mathcal{D}}, F_{\mathcal{D}} \rangle$, where $F_{\mathcal{D}} = \{P \in 2^Q \mid P \cap F \neq \emptyset\}$ and for all $P \in 2^Q$ and $\iota \in 2^X$, $\delta_{\mathcal{D}}(P, \iota) = \{q \in Q \mid q \in \delta(p, \iota) \text{ for some } p \in P\}$. For each output variable $y \in Y$, we denote by $Q(y)$ the set of \mathcal{A}_φ -states $q =$

$(a_{-b(\varphi)}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)}) \in Q \setminus \{q_0\}$ such that y is in the main value a_0 of q . In the following, we construct a BSRV specification φ' over X and $Y' \supseteq Y$ satisfying the statement of the theorem, where the set of output variables Y' is defined as follows

$$Y' = Y \cup \bigcup_{q \in Q \setminus \{q_0\}} \{\mathbf{q}\} \cup \bigcup_{P \in 2^Q} \{P\}$$

Thus, we associate to each non-initial \mathcal{A}_φ -state q , an output variable \mathbf{q} , and to each \mathcal{D}_φ -state $P \in 2^Q$, an output variable P .

For each $\iota \in 2^X$, let E_ι be the Boolean expression over X encoding precisely the input symbol ι given by $E_\iota := (\bigwedge_{x \in \iota} x) \wedge (\bigwedge_{x \in X \setminus \iota} \neg x)$. Additionally, for all $q \in Q$ and $P \in 2^Q$, let Acc_q and $\text{Init}(P, \iota)$ be the Boolean constants defined as follows:

$$\text{Acc}_q = \begin{cases} 1 & \text{if } q \in F \\ 0 & \text{otherwise} \end{cases} \quad \text{Init}(P, \iota) = \begin{cases} 1 & \text{if } P = \delta_{\mathcal{D}}(\{q_0\}, \iota) \\ 0 & \text{otherwise} \end{cases}$$

Then, for all $P \in 2^Q$, $q \in Q \setminus \{q_0\}$, and $y \in Y$, the equations of the BSRV φ' for the output variables P , \mathbf{q} , and y are defined as follows.

$$\begin{aligned} P &:= \text{if first then } \bigwedge_{\iota \in 2^X} (E_\iota \rightarrow \text{Init}(P, \iota)) \text{ else } \bigvee_{(P', \iota) \in \{(P', \iota) \mid P = \delta_{\mathcal{D}}(P', \iota)\}} E_\iota \wedge P'[-1|0] \\ \mathbf{q} &:= \text{if } \bigvee_{P \in \{P \in 2^Q \mid q \in P\}} P \text{ then } E_q \text{ else } 0 \\ &\quad \text{where } E_q := \text{if last then } \text{Acc}_q \text{ else } \bigvee_{\iota \in 2^X} \bigvee_{q' \in \delta(q, \iota)} (E_\iota[+1|0] \wedge \mathbf{q}'[+1|0]) \\ y &:= \bigvee_{q \in Q(y)} \mathbf{q} \end{aligned}$$

By construction, the specification BSRV φ' is well-formed—and consequently well-defined—and the alternation depth of φ' is exactly 1. We still need to show that φ' is φ -equivalent. Let σ_X be an input stream valuation, and σ and σ' be the unique valuation models of φ and φ' , respectively, associated with the input σ_X . We need to prove that the restrictions of σ and σ' to Y coincide. Since φ is well-defined, by Proposition 2, \mathcal{A}_φ is universal and unambiguous. Let $\pi = q_0, q_1, \dots, q_{|\sigma_X|}$ be the unique accepting run of \mathcal{A}_φ over σ_X (which encodes σ). Then, by the equations of φ' associated with the output variables $y \in Y$, it suffices to prove the following condition:

$$\text{for each } 1 \leq i \leq |\sigma_X|, \text{ there is exactly one state } p \in Q \setminus \{q_0\} \text{ such that} \quad (2) \\ p \in \sigma'(i) \text{ and } p = q_i$$

Let $\pi_{\mathcal{D}} = \{q_0\}, P_1, \dots, P_{|\sigma_X|}$ be the run of \mathcal{D}_φ over σ_X . First, we observe that the equations for the output variables P ensure that for each $1 \leq i \leq |\sigma_X|$, there is exactly

one \mathcal{D}_φ -state $P \in 2^Q$ such that $P \in \sigma'(i)$, and $P = P_i$. By using this observation and the fact that \mathcal{A}_φ is universal and unambiguous, and proceeding as in the proof of correctness of the algorithm of Figure 4, Condition (2) easily follows, which concludes the proof. \square

5. Decision Problems

In this section, we show complexity results for some relevant decision problems related to BSRV specifications. In particular, we establish that while checking well-definedness is in EXPTIME, checking for a given BSRV φ and a given subset Z of output variables, whether φ is well-defined with respect to Z (*generalized well-definedness problem*) is instead EXPSPACE-complete. Our results can be summarized as follows.

Theorem 7. For the class of BSRV, the following hold:

1. The under-definedness problem is PSPACE-complete, the well-definedness problem is in EXPTIME and at least PSPACE-hard, while the over-definedness problem and the *generalized* well-definedness problem are both EXPSPACE-complete.
2. When BSRV are interpreted as language recognizers, language emptiness is PSPACE-complete, while language universality, language inclusion, and language equivalence are EXPSPACE-complete.
3. Checking semantic equivalence is EXPSPACE-complete.

In the following Subsections 5.1 and 5.2, we establish the upper bounds and the lower bounds of Theorem 7.

5.1. Upper bounds of Theorem 7

We first need a preliminary result (Proposition 3 below). For an NFA $\mathcal{A} = \langle Q, q_0, \delta, F \rangle$, a *state projection* of \mathcal{A} is a mapping $\Upsilon : Q \rightarrow P$ for some finite set P such that for all $q \in Q$, $\Upsilon(q)$ is computable in logarithmic space (in the size of Q). The mapping Υ can be extended to sequences of states in the obvious way. We say that the NFA \mathcal{A} is *unambiguous with respect to Υ* if for all $w \in \mathcal{L}(\mathcal{A})$ and accepting runs π and π' of \mathcal{A} over w , their projections $\Upsilon(\pi)$ and $\Upsilon(\pi')$ coincide.

Proposition 3. Given an NFA \mathcal{A} and a state projection Υ of \mathcal{A} , checking whether \mathcal{A} is not unambiguous with respect to Υ can be done in NLOGSPACE.

PROOF. The following non-deterministic algorithm solves the problem, given the input (\mathcal{A}, Υ) : at each step, the algorithm guesses two runs π and π' of \mathcal{A} over the same input. The algorithm keeps in memory only the pair of states (q, q') , where q is the last state of π and q' is the last state of π' , and a flag f which is 1 whenever the projections

$\Upsilon(\pi)$ and $\Upsilon(\pi')$ of the two runs π and π' guessed so far are distinct. Initially, q and q' coincide with the initial state, and $f = 0$. If $f = 1$, and q and q' are both accepting (and then π and π' are two accepting runs over the same input and $\Upsilon(\pi)$ and $\Upsilon(\pi')$ are distinct), the algorithm terminates with success. Otherwise, the algorithm guesses two transitions of \mathcal{A} from q and q' reading the same input symbol, leading to states p and p' , respectively, re-writes the memory by replacing the pair (q, q') with the new pair (p, p') , and the flag f with the new flag f' , where f' is 1 precisely whenever either $f = 1$ or $\Upsilon(p)$ and $\Upsilon(p')$ are distinct, and the whole procedure is repeated. \square

Now, we provide the upper bounds of Theorem 7. Fix a BSRV φ over X and Y , and let \mathcal{A}_φ be the NFA of Theorem 1 accepting $\mathcal{L}(\varphi)$ and whose size is *singly exponential* in the size of φ .

Membership in PSPACE for under-definedness of BSRV. By Theorem 1 and Lemma 1, for every input σ_X , there is a bijection between the set of accepting runs of \mathcal{A}_φ over σ_X and the set of valuation models of φ associated with σ_X . Hence, φ is under-defined if and only if \mathcal{A}_φ is *not* unambiguous. Since \mathcal{A}_φ can be constructed on the fly and $\text{PSPACE} = \text{NPSPACE}$, by Proposition 3 (with Υ as the identity map), it follows that the under-definedness problem is in PSPACE.

Membership in EXPTIME for well-definedness of BSRV. Checking universality of unambiguous NFA can be done in polynomial time [30]. By Proposition 2, φ is well-defined if and only if \mathcal{A}_φ is universal and unambiguous. By Proposition 3, checking that \mathcal{A}_φ is unambiguous can be done in PSPACE (in the size of φ). Thus, since the size of \mathcal{A}_φ is singly exponential in the size of φ , we obtain that checking well-definedness for BSRV is in EXPTIME.

Membership in EXPSPACE for over-definedness and generalized well-definedness of BSRV. Since \mathcal{A}_φ accepts $\mathcal{L}(\varphi)$, φ is over-defined iff \mathcal{A}_φ is not universal. Thus, since checking universality for NFA is PSPACE-complete [31], membership in EXPSPACE for checking over-definedness follows. Now, let us consider the generalized well-definedness problem. Let $Z \subseteq Y$ be a subset of the output variables. Recall that the set of non-initial states of \mathcal{A}_φ is contained in $(A_\perp)^{b(\varphi)} \times A \times (A_\perp)^{f(\varphi)}$, where $A = 2^{X \cup Y}$ and $A_\perp := A \cup \{\perp\}$. Let Υ_Z be the state projection of \mathcal{A}_φ assigning to the initial state q_0 of \mathcal{A}_φ q_0 itself, and assigning to each non-initial state $(a_{-b(\varphi)}, \dots, a_{-1}, a_0, a_1, \dots, a_{f(\varphi)})$ of \mathcal{A}_φ the tuple $(d_{-b(\varphi)}, \dots, d_{-1}, d_0, d_1, \dots, d_{f(\varphi)})$, where for all $b(\varphi) \leq i \leq f(\varphi)$, $d_i = a_i$ if $a_i = \perp$, and $d_i = a_i \cap Z$ otherwise. By Theorem 1, given an input σ_X , the accepting runs of \mathcal{A}_φ over σ_X are the expanded valuation models of φ encoding the valuation models of φ associated with the input σ_X . Now, let σ and σ' be two valuation models of φ associated with an input σ_X , and π and π' be the *expanded* valuation models encoding σ and σ' , respectively. By construction, it follows that $\Upsilon_Z(\pi) = \Upsilon_Z(\pi')$ if

and only if the restrictions of σ and σ' to Z coincide. By Theorem 1, we obtain that φ is well-defined with respect to Z if and only if \mathcal{A}_φ is unambiguous with respect to Υ_Z and \mathcal{A}_φ is universal. Thus, since checking universality for NFA is PSPACE-complete, by Proposition 3, membership in EXPSPACE for checking generalized well-definedness follows.

Membership in PSPACE for language emptiness. Since checking language emptiness for NFA is NLOGSPACE-complete, by Theorem 1, the result follows.

Membership in EXPSPACE for language universality, language inclusion, and language equivalence of BSRV. Recall that for NFA, universality, inclusion, and equivalence are PSPACE-complete [31]. Hence, by Theorem 1, the result follows.

Membership in EXPSPACE for semantic equivalence of BSRV. Let φ be a BSRV over X and Y , φ' be a BSRV over X and Y' , and $Z \subseteq Y \cap Y'$. By a straightforward adaptation of Theorem 1, one can construct in singly exponential time two NFA $\mathcal{A}_{\varphi,Z}$ and $\mathcal{A}_{\varphi',Z}$ over $2^{X \cup Z}$ such that $\mathcal{L}(\mathcal{A}_{\varphi,Z})$ is the set of stream valuations over $X \cup Z$ which can be extended to valuation models of φ , and $\mathcal{L}(\mathcal{A}_{\varphi',Z})$ is the set of stream valuations over $X \cup Z$ which can be extended to valuation models of φ' . It follows that φ and φ' are equivalent with respect to Z if and only if $\mathcal{L}(\mathcal{A}_{\varphi,Z}) = \mathcal{L}(\mathcal{A}_{\varphi',Z})$. Thus, since language equivalence for NFA is PSPACE-complete, membership in EXPSPACE for the considered problem follows.

5.2. Lower bounds of Theorem 7

First, we consider the EXPSPACE-hardness results of Theorem 7.

EXPSPACE-hardness for over-definedness of BSRV. The result is obtained by a polynomial-time reduction from a domino-tiling problem for grids with rows of singly exponential length [32]. An instance \mathcal{I} of this problem is a tuple $\mathcal{I} = \langle C, \Delta, m, d_{init}, d_{final} \rangle$, where C is a finite set of colors, $\Delta \subseteq C^4$ is a set of tuples $\langle c_{down}, c_{left}, c_{up}, c_{right} \rangle$ of four colors, called *domino-types*, $m > 0$ is a natural number (written in unary), and $d_{init}, d_{final} \in \Delta$ are two domino-types. For $n > 0$, a *tiling of \mathcal{I} for the $n \times 2^m$ -grid* is a mapping $f : [0, n - 1] \times [0, 2^m - 1] \rightarrow \Delta$ satisfying the following:

- *row requirement:* two adjacent cells in a row have the same color on the shared edge: for all $(i, j) \in [0, n - 1] \times [0, 2^m - 1]$ with $j < 2^m - 1$, $[f(i, j)]_{right} = [f(i, j + 1)]_{left}$;
- *column requirement:* two adjacent cells in a column have the same color on the shared edge: for all $(i, j) \in [0, n - 1] \times [0, 2^m - 1]$ with $i < n - 1$, $[f(i, j)]_{up} = [f(i + 1, j)]_{down}$;

- $f(0, 0) = d_{init}$ and $f(n - 1, 2^m - 1) = d_{final}$.

A tiling of \mathcal{I} is a tiling of \mathcal{I} for the $n \times 2^m$ -grid for some $n > 0$.

Checking the existence of a tiling for \mathcal{I} is EXPSPACE-complete [32]. We also use the notion of *pseudo-tiling* for \mathcal{I} , which is similar to the notion of tiling but the column requirement is relaxed. Note that the column requirement is the crucial feature which makes the considered domino tiling problem EXPSPACE-hard.¹ In the following, we construct in polynomial-time a BSRV specification φ such that there exists a tiling for \mathcal{I} if and only if φ is over-defined. The set X of input variables of φ is given by

$$X = \{d \mid d \in \Delta\} \cup \{b_1^+, \dots, b_m^+, b_1^-, \dots, b_m^-\}$$

We associate to each domino-type $d \in \Delta$ an input variable d . The additional input variables $b_1^+, \dots, b_m^+, b_1^-, \dots, b_m^-$ are used to encode the value of an m -bit counter numbering the cells of one row of the grid (b_i^+ is 1 whenever the i -th bit is 1, and b_i^- is 1 whenever the i -th bit is 0). Thus, a cell is encoded as a finite word over 2^X of length $m + 1$, the first m positions giving the binary encoding of the column number and the last position giving the associated domino-type.² More precisely, a cell with content $d \in \Delta$ and column number $j \in [0, 2^m - 1]$ is encoded by the word $\{d\}\{b_1\} \dots \{b_m\}$, where $b_k \in \{b_k^+, b_k^-\}$ for all $k \in [1, m]$, and $b_k = b_k^+$ iff the k^{th} bit in the binary encoding of the column number j is 1. A tiling is then encoded as a sequence of rows, starting from the first row, where a row lists the encodings of cells from left to right.

In the following, for a stream variable y and an integer k , we use $y[k]$ for the stream expression $y[k]1$ if $k \neq 0$, and for the stream expression y if $k = 0$.

We illustrate now the construction of φ that ensures that the unique inputs for which there is no output stream valuation are those encoding tilings of \mathcal{I} .

The preliminary step in the construction of φ is to enforce a designated output variable PTU to be uniform with the uniform value 1 characterizing the inputs encoding pseudo-tilings. For this, we need two additional output variables PT and test₁.

The equation for the output variable PT ensures that PT assumes the value 1

¹One can easily show that checking the existence of a *pseudo-tiling* is just PSPACE-complete.

²We assume that the first bit is the least significant one.

everywhere *iff* the input streams encode a pseudo-tiling:

$$\begin{aligned}
\text{PT} := & \underbrace{\bigvee_{x \in X} (x \wedge \bigwedge_{x' \in X \setminus \{x\}} \neg x')}_{\text{exactly one input variable has value 1}} \wedge \\
& \underbrace{\bigwedge_{i=1}^{m-1} ((b_i^+ \vee b_i^-) \rightarrow (\neg \text{last} \wedge (b_{i+1}^+[+1] \vee b_{i+1}^- [+1]))) \wedge ((b_m^+ \vee b_m^-) \rightarrow (\neg \text{last} \wedge \bigvee_{d \in \Delta} d[+1]))}_{\text{the input is a list of numbered cells}} \wedge \\
& \underbrace{(\text{first} \rightarrow (\bigwedge_{i=1}^m b_i^- [i-1] \wedge d_{init}[+m]))}_{\text{the input starts with a } d_{init}\text{-cell numbered 0}} \wedge \underbrace{(\text{last} \rightarrow (\bigwedge_{i=1}^m b_i^+ [i-m-1] \wedge d_{final}))}_{\text{the input ends with a } d_{final}\text{-cell numbered } 2^m - 1} \wedge \\
& \left\{ (\neg \text{last} \wedge \bigvee_{d \in \Delta} d) \longrightarrow \left((b_1^+ [+1] \vee b_1^- [+1]) \wedge \right. \right. \\
& \left. \left. (b_1^+ [-m] \leftrightarrow b_1^- [+1]) \wedge \bigwedge_{i=1}^{m-1} (b_{i+1}^+ [-m+i] \leftrightarrow b_{i+1}^- [i+1]) \leftrightarrow (b_i^+ [-m+i-1] \wedge b_i^- [i]) \right) \right\} \wedge \\
& \underbrace{\left(\bigwedge_{d \in \Delta} d \rightarrow (\text{last} \vee \bigwedge_{i=1}^{i=m} b_i^- [i] \vee \bigvee_{d' \in \Delta: (d')_{\text{left}}=(d)_{\text{right}}} d'[m+1]) \right)}_{\text{two adjacent cells in a row have the same color on the shared edge}}
\end{aligned}$$

The equation for the output variable PTU is given by

$$\text{PTU} := \text{if } (\neg \text{first} \rightarrow (\text{PTU}[-1] \leftrightarrow \text{PTU})) \wedge (\neg \text{PT} \rightarrow \neg \text{PTU}) \text{ then PTU else } \neg \text{PTU}$$

Hence, PTU is a uniform output variable, and the uniform value of PTU is 0 whenever PT assumes the value 0 at some position. We exploit the additional variable test_1 in order to avoid situations where PT is everywhere 1 and the uniform value of PTU is 0. The equation for test_1 is as follows:

$$\text{test}_1 := \text{if } \{ ((\text{first} \wedge \text{PT} \wedge \neg \text{PTU}) \rightarrow \neg \text{test}_1) \wedge (\text{last} \rightarrow \text{test}_1) \wedge ((\neg \text{first} \wedge \neg \text{test}_1[-1] \wedge \text{PT} \wedge \neg \text{PTU}) \rightarrow \neg \text{test}_1) \} \text{ then } \text{test}_1 \text{ else } \neg \text{test}_1$$

Thus, the uniform value of PTU is 1 iff the input encodes a pseudo-tiling (note that there is an output valuation of PT, PTU, and test_1 for each input).

Now, we describe the crucial step of the construction of φ . Assume that the input encodes a pseudo-tiling, i.e. the uniform value of PTU is 1. Then, the input is *not* a tiling of \mathcal{I} if and only if the followign condition holds:

$$\text{there are two adjacent cells } (i, j) \text{ and } (i+1, j) \text{ in some column which have } \textit{different} \text{ color on the shared edge.} \quad (3)$$

In order to check Condition (3), we use $O(m + |\Delta|)$ additional output variables: the equations for these variables ensure that for each input σ_X , there is a valuation model over σ_X if and only if whenever σ_X encodes a pseudo-tiling, the above Condition (3) is satisfied. Hence, the unique inputs for which there is no stream valuation are those encoding tilings of \mathcal{I} .

The fulfillment of Condition (3) for an input which is a pseudo-tiling is ensured in three steps.

- *First step:* we use two output variables, Bl_1 and Bl_2 , for ‘marking’ two cells c_1 and c_2 of the pseudo-tiling.
- *Second step:* we use three additional output variables, namely, m_1 , m_2 , and test_2 in order to check that the two cells c_1 and c_2 marked by Bl_1 and Bl_2 , respectively, belong to two adjacent rows (with the c_2 ’s row following the c_1 ’s row).
- *Third step:* finally, we use the output variables $\text{ob}_1, \dots, \text{ob}_m$, test_3 , and od for each $d \in \Delta$ to guarantee that the cells marked by Bl_1 and Bl_2 have the same column number but different color on the shared edge.

Now, we proceed with the formal definition of the equations for the output variables of Steps 1–3.

Equations for Step 1. For each $h = 1, 2$, the equation for the output variable Bl_h requires that whenever the uniform value of PTU is 1, then the stream for Bl_h is in 0^*1^+ and the suffix in 1^+ starts with a cell (“the cell marked by Bl_h ”).

$$\text{Bl}_h := \text{if PTU} \rightarrow \{(\text{last} \rightarrow \text{Bl}_h) \wedge ((\text{Bl}_h \wedge \neg \text{Bl}_h[-1|0]) \rightarrow (\text{b}_1^+ \vee \text{b}_1^-)) \wedge (\neg \text{Bl}_h \rightarrow \neg \text{Bl}_h[-1|0]) \wedge (\text{Bl}_h \rightarrow \text{Bl}_h[+1])\} \text{ then } \text{Bl}_h \text{ else } \neg \text{Bl}_h$$

Equations for Step 2. The output variable m_1 is used to mark the first cell of the row following the one containing the cell marked by Bl_1 . Formally, the equation for m_1 ensures that whenever the input is a pseudo-tiling, then the stream for m_1 is in 0^+1^+ and the suffix in 1^+ starts with a cell (“the cell marked by m_1 ”) which is the first cell numbered 0 following the cell marked by Bl_1 .

$$\begin{aligned} \text{m}_1 := \text{if PTU} \rightarrow & \{(\text{last} \rightarrow \text{m}_1) \wedge (\text{first} \rightarrow \neg \text{m}_1) \wedge (\neg \text{m}_1 \rightarrow \neg \text{m}_1[-1|0]) \wedge \\ & (\text{m}_1 \rightarrow \text{m}_1[+1]) \wedge ((\text{Bl}_1 \wedge \bigwedge_{i=1}^m \text{b}_i^-[i]) \rightarrow \text{m}_1[1]) \wedge \\ & ((\neg \text{m}_1 \wedge \text{m}_1[1]) \rightarrow (\text{Bl}_1 \wedge \bigwedge_{i=1}^m \text{b}_i^-[i]))\} \text{ then } \text{m}_1 \text{ else } \neg \text{m}_1 \end{aligned}$$

The output variable m_2 is used to mark the first cell of the row containing the cell marked by Bl_2 . Formally, the equation for m_2 ensures that whenever the input is a pseudo-tiling, then the stream for m_2 is in 0^+1^+ and the prefix in 0^+ ends with the

first bit of a cell (“the cell marked by m_2 ”) which is the last cell numbered 0 preceding the cell marked by Bl_2 (Bl_2 included).

$$\begin{aligned} m_2 := & \text{if PTU} \rightarrow \{(\text{last} \rightarrow m_2) \wedge (\text{first} \rightarrow \neg m_2) \wedge (\neg m_2 \rightarrow \neg m_2[-1|0]) \wedge \\ & (m_2 \rightarrow m_2[+1]) \wedge ((\neg Bl_2 \wedge \bigwedge_{i=1}^m b_i^-[i]) \rightarrow \neg m_2[+1|0]) \wedge \\ & ((\neg m_2[+1|1] \wedge m_2[+2|0]) \rightarrow (\neg Bl_2 \wedge \bigwedge_{i=1}^m b_i^-[i]))\} \text{ then } m_2 \text{ else } \neg m_2 \end{aligned}$$

Thus, the cells marked by Bl_1 and Bl_2 belong to two adjacent rows (with the Bl_2 's row following the Bl_1 's row) if and only if m_1 and m_2 mark the same cell if and only if the first 1-value position of m_1 corresponds to the last 0-value position of m_2 . The equation for the output variable $test_2$ ensures this last condition.

$$\begin{aligned} test_2 := & \text{if PTU} \rightarrow \{(\neg m_2 \wedge m_2[+1]) \leftrightarrow (m_1 \wedge \neg m_1[-1|0])\} \\ & \text{then } test_2 \text{ else } \neg test_2 \end{aligned}$$

Equations for Step 3. Finally, we define the equations for the output variables ob_1, \dots, ob_m , $test_3$, and od ($d \in \Delta$) ensuring that the cells marked by Bl_1 and Bl_2 have the same column number but different color on the shared edge. The equation for ob_i guarantees that there is a *uniform* stream for ob_i if and only if the i -th bits of the cells marked by Bl_1 and Bl_2 have the same value.

$$\begin{aligned} ob_i := & \text{if PTU} \rightarrow \bigwedge_{h=1}^2 \{ (Bl_h \wedge \neg Bl_h[-1|0]) \rightarrow \\ & ((b_i^+[i-1] \rightarrow ob_i[i-1]) \wedge (b_i^-[i-1] \rightarrow \neg ob_i[i-1])) \} \text{ then } ob_i \text{ else } \neg ob_i \end{aligned}$$

For each $d \in \Delta$, the equation for od ensures that there is a uniform stream for od if and only if whenever the domino-type of the cell marked by Bl_1 is d , then the cells marked by Bl_1 and Bl_2 have different color on the shared edge.

$$\begin{aligned} od := & \text{if PTU} \rightarrow \{((d[m] \wedge Bl_1 \wedge \neg Bl_1[-1|0]) \rightarrow od[m]) \wedge \\ & ((\bigvee_{d' \in \Delta: (d')_{down}=(d)_{up}} d'[m] \wedge Bl_2 \wedge \neg Bl_2[-1]) \rightarrow \neg od[m])\} \text{ then } od \text{ else } \neg od \end{aligned}$$

Thus, the equation for variable $test_3$ enforces the output streams for ob_1, \dots, ob_m, od ($d \in \Delta$) to be uniform.

$$\begin{aligned} test_3 := & \text{if PTU} \rightarrow \neg \text{first} \rightarrow (\bigwedge_{i=1}^m (ob_i[-1] \leftrightarrow ob_i) \wedge \bigwedge_{d \in \Delta} (od[-1] \leftrightarrow od)) \\ & \text{then } test_3 \text{ else } \neg test_3 \end{aligned}$$

Hence, the unique inputs for which the constructed BSRV φ has no output stream valuation are those encoding tilings of \mathcal{I} . In other words, φ is over-defined precisely when there is a tiling of \mathcal{I} . Note that the size of φ is quadratic in the size of \mathcal{I} . This concludes the proof of EXPSPACE-hardness of checking over-definedness for BSRV.

EXPSPACE-hardness for generalized well-definedness of BSRV. First, we observe that the complement of the over-definedness problem can be reduced in linear time to the generalized well-definedness problem. Indeed, let φ be a BSRV specification and φ' be the BSRV specification obtained from φ by adding the equation $z = 0$, where z is a fresh output variable. Clearly, φ is not over-defined if and only if φ' is well-defined with respect to $\{z\}$. Hence, by the obtained EXPSPACE-completeness result for over-definedness of BSRV, EXPSPACE-hardness of checking generalized well-definedness of BSRV follows.

EXPSPACE-hardness for language universality, language inclusion, and language equivalence of BSRV. The results directly follow from EXPSPACE-completeness of checking over-definedness for BSRV and the facts that language universality for BSRV is the complement of the over-definedness problem, and language inclusion and language equivalence can be reduced in linear time to language universality.

EXPSPACE-hardness for semantic equivalence of BSRV. Let φ be a BSRV specification over X and Y , φ' be a BSRV specification over X and Y' , and $Z \subseteq Y \cap Y'$. Note that if $Z = \emptyset$, then φ and φ' are equivalent with respect to Z if and only if $\mathcal{L}(\varphi) = \mathcal{L}(\varphi')$. Hence, by EXPSPACE-completeness of checking language equivalence for BSRV, the result follows.

The only missing parts to complete the proof of Theorem 7 are the PSPACE-hardness of checking under-definedness, well-definedness, and language emptiness for BSRV.

PSPACE-hardness for under-definedness and well-definedness of BSRV. We proceed by a polynomial-time reduction from a domino-tiling problem for grids with rows of polynomial length [32]. An instance $\mathcal{I} = \langle C, \Delta, m, d_{init}, d_{final} \rangle$ of this problem is defined as in the proof of EXPSPACE-hardness for over-definedness of BSRV. However, in this case a tiling of \mathcal{I} is a tiling of \mathcal{I} for the $n \times m$ -grid for some $n > 0$ (i.e., the length of any row is m). The existence of a tiling for \mathcal{I} is an PSPACE-complete problem [32]. We construct in polynomial time a BSRV φ such that the following holds:

- there exists a tiling of \mathcal{I} if and only if φ is under-defined;
- φ is *not* under-defined if and only if φ is well-defined.

Hence, the result follows. Now, we illustrate the construction of φ . The set X of input variables of the specification φ is given by

$$X = \{d \mid d \in \Delta\} \cup \{\text{mk}\}$$

We associate to each domino-type $d \in \Delta$ an input variable d . The additional input variable mk is used as a separator between two adjacent rows. Thus, a tiling is encoded as a sequence of rows separated by the special marker, starting from the first row. Additionally, the first row is preceded by the special marker, and the last row is followed by the special marker. The specification φ has two output variables: PT and PTU . The output variable PT is used to check that the input encodes a tiling of \mathcal{I} : in particular, PT assumes the value 1 everywhere if and only if the input streams encode a tiling. Formally, the equation for variable PT is as follows.

$$\begin{aligned} \text{PT} := & \underbrace{\bigvee_{x \in X} (x \wedge \bigwedge_{x' \in X \setminus \{x\}} \neg x')}_{\text{exactly one input variable has value 1}} \wedge \\ & \underbrace{(\text{first} \rightarrow (\text{mk} \wedge d_{\text{init}}[+1|0]))}_{\text{the input starts with a } d_{\text{init}}\text{-cell}} \wedge \underbrace{(\text{last} \rightarrow (\text{mk} \wedge d_{\text{final}}[-1|0]))}_{\text{the input ends with a } d_{\text{final}}\text{-cell}} \wedge \\ & \underbrace{\left((\text{mk} \wedge \neg \text{last}) \rightarrow (\text{mk}[m+1|0] \wedge \bigwedge_{i=1}^m \neg \text{mk}[i|1]) \right)}_{\text{adjacent rows are separated by the marker}} \wedge \\ & \underbrace{\bigwedge_{d \in \Delta} \left(d \rightarrow (\text{mk}[+1|0] \vee \bigvee_{d' \in \Delta: (d')_{\text{left}}=(d)_{\text{right}}} d'[+1|0]) \right)}_{\text{two adjacent cells in a row have the same color on the shared edge}} \wedge \\ & \underbrace{\bigwedge_{d \in \Delta} \left(d \rightarrow \bigvee_{d' \in \Delta: (d')_{\text{down}}=(d)_{\text{up}}} d'[m+1|1] \right)}_{\text{two adjacent cells in a column have the same color on the shared edge}} \end{aligned}$$

Finally, the equation for the *uniform* output variable PTU is as follows.

$$\text{PTU} := \text{if } (\neg \text{first} \rightarrow (\text{PTU}[-1|1] \leftrightarrow \text{PTU})) \wedge (\neg \text{PT} \rightarrow \neg \text{PTU}) \text{ then PTU else } \neg \text{PTU}$$

Note that if the input does *not* encode a tiling (i.e., for some position, PT assumes the value 0), then the uniform value of PTU is 0. Otherwise, the uniform value of PTU may be 0 or 1. Since for each input, the stream valuation for the other output variable PT is uniquely determined, it follows that there is a tiling for \mathcal{I} iff φ is under-defined. Moreover, since for each input, there is some output stream valuation, it follows that φ is *not* under-defined if and only if φ is well-defined. Note that the size of φ is quadratic in the size of \mathcal{I} . Hence, the result follows.

PSPACE-hardness for language emptiness of BSRV. We modify the polynomial-time reduction given in the proof of PSPACE-hardness of checking under-definedness for BSRV as follows: the equation for the output variable PTU of the BSRV φ is updated as follows:

$$\text{PTU} := \textit{if } (1 \leftrightarrow \text{PT}) \wedge (\neg \text{PT} \rightarrow \neg \text{PTU}) \textit{ then PTU else } \neg \text{PTU}$$

Hence, PTU is a uniform output variable whose uniform value is always 1. Moreover, the output stream for PTU is defined iff the output stream for PT is in 1^+ . Additionally, the construction in the proof of PSPACE-hardness of checking under-definedness for BSRV ensures that for each input, the output stream for the output variable PT is uniquely determined, and PT assumes the value 1 everywhere if and only if the input streams encode a tiling. Hence, the updated construction is a polynomial-time reduction from a PSPACE-complete problem to language emptiness for BSRV.

6. Conclusion

In this paper, we have studied some fundamental theoretical problems for the class of Boolean SRV. We have also presented an offline monitoring algorithm for well-defined BSRV that only requires two passes over the dumped trace. An open question is the exact complexity of checking well-definedness for BSRV. We only show here that this problem lies somewhere between PSPACE and EXPTIME. Future work includes the theoretical investigation and the development of monitoring algorithms for SRV over richer data types, such as counters and stacks. In particular, the emerging field of symbolic automata and transducers [23]—that extend the classical notions from discrete alphabets to theories handled by solvers—seems very promising to study in the context of SRV, which in turn can extend automata from states and transitions to stream dependencies. The combination of these two extensions has the potential to provide a rich but *tractable* foundation for the runtime verification of values from rich types. Additionally, we are studying the extension to the monitoring of visibly pushdown systems, where SRV is extended to deal with traces containing calls and returns.

Finally, we plan to study the monitorability of well-definedness of specifications. If one cannot determine well-definedness statically, a plausible alternative would be to use a monitor that *assumes* well-definedness in tandem with a monitor that *detects* non-well-definedness (and hence, the incorrectness of the first monitor).

References

- [1] M. Leucker, C. Schallhart, A brief account of runtime verification, *The Journal of Logic and Algebraic Programming* 78 (5) (2009) 293–303.

- [2] K. Havelund, A. Goldberg, Verify your runs, in: Proc. of VSTTE'05, LNCS 4171, Springer, 2005, pp. 374–383.
- [3] Z. Manna, A. Pnueli, Temporal Verification of Reactive Systems: Safety, Springer, New York, 1995.
- [4] K. Havelund, G. Roşu, Synthesizing monitors for safety properties, in: Proc. of TACAS'02, LNCS 2280, Springer, 2002, pp. 342–356.
- [5] C. Eisner, D. Fisman, J. Havlicek, Y. Lustig, A. McIsaac, D. V. Campenhout, Reasoning with temporal logic on truncated paths, in: Proc. of CAV'03, Vol. 2725 of LNCS 2725, Springer, 2003, pp. 27–39.
- [6] A. Bauer, M. Leucker, C. Schallhart, Runtime verification for LTL and TLTL, ACM Transactions on Software Engineering and Methodology 20 (4) (2011) 14.
- [7] K. Sen, G. Roşu, Generating optimal monitors for extended regular expressions, ENTCS 89 (2) (2003) 226–245.
- [8] H. Barringer, A. Goldberg, K. Havelund, K. Sen, Rule-based runtime verification, in: Proc. of VMCAI'04, LNCS 2937, Springer, 2004, pp. 44–57.
- [9] G. Roşu, K. Havelund, Rewriting-based techniques for runtime verification., Automated Software Engineering 12 (2) (2005) 151–197.
- [10] B. D'Angelo, S. Sankaranarayanan, C. Sánchez, W. Robinson, B. Finkbeiner, H. B. Sipma, S. Mehrotra, Z. Manna, LOLA: Runtime monitoring of synchronous systems, in: Proc. of TIME'05, IEEE CS Press, 2005, pp. 166–174.
- [11] A. Pnueli, A. Zaks, PSL model checking and run-time verification via testers, in: Proc. of FM'06, LNCS 4085, Springer, 2006, pp. 573–586.
- [12] A. Donzé, O. Maler, E. Bartocci, D. Nickovic, R. Grosu, S. A. Smolka, On temporal logic and signal processing, in: Proc. of ATVA'12, Vol. 7561 of LNCS, Springer, 2012, pp. 92–106.
- [13] B. Finkbeiner, S. Sankaranarayanan, H. B. Sipma, Collecting statistics over run-time executions, ENTCS 70 (4) (2002) 36–54.
- [14] A. Bauer, R. Gore, A. Tiu, A first-order policy language for history-based transaction monitoring, in: Proc. of ICTAC'09, LNCS 5684, Springer, 2009, pp. 96–111.
- [15] D. Basin, M. Harvan, F. Klaedtke, E. Zălinescu, MONPOLY: Monitoring usage-control policies, in: Proc. of RV'12, LNCS 7687, Springer, 2012.

- [16] D. Basin, F. Klaedtke, S. Müller, Policy monitoring in first-order temporal logic, in: Proc. of CAV'10, LNCS 6174, Springer, 2010, pp. 1–18.
- [17] P. Caspi, M. Pouzet, Synchronous Kahn Networks, in: Proc. of ICFP'96, ACM Press, 1996, pp. 226–238.
- [18] G. Berry, Proof, language, and interaction: essays in honour of Robin Milner, MIT Press, 2000, Ch. The foundations of Esterel, pp. 425–454.
- [19] N. Halbwachs, P. Caspi, D. Pilaud, J. Plaice, Lustre: a declarative language for programming synchronous systems, in: Proc. of POPL'87, ACM Press, 1987, pp. 178–188.
- [20] T. Gautier, P. Le Guernic, L. Besnard, SIGNAL: A declarative language for synchronous programming of real-time systems, in: Proc. of FPCA'87, LNCS 274, Springer, 1987, pp. 257–277.
- [21] L. Pike, A. Goodloe, R. Morisset, S. Niller, Copilot: A hard real-time runtime monitor, in: Proc. of RV'10, LNCS 6418, Springer, 2010.
- [22] A. E. Goodloe, L. Pike, Monitoring distributed real-time systems: A survey and future directions, Tech. rep., NASA Langley Research Center (2010).
- [23] M. Veanes, P. Hooimeijer, B. Livshits, D. Molnar, N. Bjørner, Symbolic finite state transducers: algorithms and applications., in: Proc. of POPL'12, ACM, 2012, pp. 137–150.
- [24] L. D'Antoni, M. Veanes, Extended symbolic finite automata and transducers, Formal Methods in System Design 47 (1) (2015) 93–119.
- [25] R. Alur, L. D'Antoni, M. Raghothaman, DReX: A declarative language for efficiently evaluating regular string transformations, in: Proc. of POPL'15, ACM, 2015, pp. 125–137.
- [26] L. Pike, N. Wegmann, S. Niller, A. Goodloe, Copilot: Monitoring embedded systems, Tech. rep., NASA, NASA/CR2012-217329 (2012).
- [27] Digital design and computer organization, CRC Press, 2004.
- [28] P. Wolper, Temporal logic can be more expressive, Information and Control 56 (1983) 72–99.
- [29] F. Laroussinie, N. Markey, P. Schnoebelen, Temporal logic with forgettable past, in: Proc. of LICS'02, IEEE CS Press, 2002, pp. 383–392.

- [30] R. E. Stearns, H. B. Hunt, On the equivalence and containment problems for unambiguous regular expressions, regular grammars and finite automata, *SIAM J. Comput.* 14 (3) (1985) 598–611.
- [31] A. R. Meyer, L. J. Stockmeyer, The equivalence problem for regular expressions with squaring requires exponential space, in: *Proc. of FOCS'72*, IEEE CS Press, 1972, pp. 125–129.
- [32] D. Harel, *The spirit of Computing*, 2nd Edition, Addison-Wesley, 1992.