

Foundations of Boolean Stream Runtime Verification* (Extended Version)

Laura Bozzelli¹ and César Sánchez^{2,3}

¹ Technical University of Madrid (UPM), Madrid, Spain.

² IMDEA Software Institute, Madrid, Spain

³ Institute for Information Security, CSIC, Spain.

Abstract. Stream runtime verification (SRV), pioneered by the tool LOLA, is a declarative approach 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 (1) the evaluation algorithms, and (2) 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.

1 Introduction

Runtime verification (RV) has emerged in the last decades as an applied formal technique for software reliability. In RV, a specification, expressing correctness requirements, is automatically translated into a *monitor*. Such a monitor is then used 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 *online* monitoring, while the latter one is called *offline* monitoring. Online monitoring is used to detect and possibly handle (e.g., by the execution of additional repair code) violations of the specification when the system is in operation. On the other hand, offline monitoring is used in post-mortem analysis and it is convenient for testing large systems before deployment. Unlike static verification (such as model-checking) which formally checks that all the (infinite)

* This work was funded in part by Spanish MINECO Project “TIN2012-39391-C04-01 STRONGSOFT” and by Spanish MINECO Project “TIN2012-38137-C02 VIVAC”.

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 [18, 15] for modern surveys on runtime verification.

Stream runtime verification. The first specification formalisms proposed for runtime verification were based on specification languages for static verification, typically LTL [19] or past LTL adapted for finite paths [16, 9, 5]. Other formalisms for expressing monitors include regular expressions [24], rule based specifications as proposed in the logic Eagle [1], or rewriting [23]. Stream runtime verification (SRV), first proposed in the tool LOLA [8], 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 [22].

Stream runtime verification offers two advantages to 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 (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 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 [10]. 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 [4, 2, 3]. 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 efficient 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 [7] – like Esterel [6], Lustre [13] or Signal [11] – specifically designed for observing traces of systems, 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, 12].

When used for synthesizing monitors, SRV specifications need to be *well-defined*: for every input there is a unique corresponding output stream. However, as with many synchronous languages, the declarative style of SRV allows

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, and in [8], a syntactical constraint called *well-formedness* is introduced in order to ensure the absence of circular dependencies, and guarantee well-definedness.

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 (BSRV). 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 (like semantic equivalence, implication and redundancy) can be reduced to the decision problems above.
- BSRV specifications can be naturally interpreted as language recognizers, where one selects 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 [8].

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. Section 5 is devoted to the decision problems for BSRV specifications. Finally, Section 6 concludes. Due to lack of space, some proofs are omitted and are included in the appendix.

2 Stream Runtime Verification (SRV)

In this Section, we recall the SRV framework [8]. We focus on SRV specifications over stream variables of the same type (with emphasis on the Boolean type).

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 and \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 and for all $1 \leq i \leq |w|$, $w(i)$ 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, $\tau[\ell|c]$ refers to the value of τ offset ℓ positions from the current position, and the constant c is the *default* value of type T assigned to positions from which the offset is 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 $1 \leq i \leq |\sigma|$:

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

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* and *last* for the Boolean stream expressions $0[-1|1]$ and $0[+1|1]$, respectively. Note that for a Boolean stream, *first* is 1 precisely at the first position, and *last* is 1 precisely at the last position.

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

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

For every *Boolean* stream valuation σ over Z such that $\sigma(Z) \in (01)^+$, the valuation of E with respect to σ is the uniform Boolean stream $1^{|\sigma|}$.

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 φ of type T over X and Y is a set of 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. A stream valuation of φ is a

stream valuation of type T over $X \cup Y$, while an *input* (resp., *output*) of φ is a stream valuation of type T over X (resp., 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 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 φ :

- *Under-definedness*: for some input σ_X , there are at least two distinct valuation models of φ associated with σ_X .
- *Over-definedness*: for some input σ_X , there is no valuation model of φ associated with 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 φ may be both under-defined and over-defined, 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 φ is *well-defined with respect to Z* if for each input σ_X , there is exactly one stream valuation σ_Z over Z having 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 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 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. 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 $\llbracket E \rrbracket$ the

offset ℓ if E is a stream expression of the form $\tau[\ell|c]$; otherwise, $\|E\|$ is 1. 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 .

Example 2. Consider the following BSRV 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.

3 BSRV as Language Recognizers

BSRV can be interpreted as a simple declarative formalism to specify languages of non-empty finite words. Formally, we associate to a BSRV 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, i.e.,

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

Example 3. 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 iff the valuation of the stream expression E w.r.t. $\sigma_X \cup \sigma_Y$ is in 1^+ iff $\sigma_X(x)(i) = 1$ for all odd positions i . Hence, $\mathcal{L}(\varphi)$ is the set of Boolean streams which assume the value 1 at the odd positions.

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. 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 .

Fix a BSRV 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* $f(\varphi)$:

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

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 such that $i - b(\varphi) \leq j \leq i + f(\varphi)$. 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. 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 . For 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* , intuitively represents the value of some stream valuation σ at some time step i , while $a_{-b(\varphi)}, \dots, a_{-1}$ (resp., $a_1, \dots, a_{f(\varphi)}$) represent the values of σ at the previous $b(\varphi)$ (resp., next $f(\varphi)$) time steps, if any (the symbol \perp is used to denote the absence of a previous or next 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 iff $\tau \in a$. 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:

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

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 . By construction, we easily obtain that $\sigma(w)$ is a valuation model of φ and, more precisely, the following lemma holds.

Lemma 1. *The mapping assigning 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 valuation 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. 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 satisfies $a \cap X = \iota$.

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 $\varphi_{\mathcal{A}}$ over the set of input variables X as follows. First, for each input symbol ι , we use 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:

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

Thus, we associate to each state $q \in Q$, an output variable \mathbf{q} , whose associated equation is the trivial one given by $\mathbf{q} = \mathbf{q}$. The equation for the output variable control is given by

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

where the boolean stream expression E_{ev} describes 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, it easily follows that 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^+ *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 issues. It turns out that the singly exponential blow-up in Theorem 1 cannot be avoided. To prove this we first show a linear time translation from standard 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 (we omit the rules for the boolean connectives and the atomic propositions, which are standard):

$$\begin{aligned}
(w, i) \models \bigcirc \psi & \Leftrightarrow i + 1 \leq |w| \text{ and } (w, i + 1) \models \psi \\
(w, i) \models \ominus \psi & \Leftrightarrow i > 1 \text{ and } (w, i - 1) \models \psi \\
(w, i) \models \psi_1 \mathcal{U} \psi_2 & \Leftrightarrow \exists i \leq j \leq |w|, (w, j) \models \psi_2 \text{ and } \forall i \leq h < j, (w, h) \models \psi_1 \\
(w, i) \models \psi_1 \mathcal{S} \psi_2 & \Leftrightarrow \exists 1 \leq j \leq i, (w, j) \models \psi_2 \text{ and } \forall j < h \leq i, (w, h) \models \psi_1
\end{aligned}$$

The language $\mathcal{L}(\psi)$ of a 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 . Then, the set of output variables Y of φ is defined as follows.

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

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

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

We need to show that $\mathcal{L}(\varphi) = \mathcal{L}(\psi)$. We prove the two set inclusions separately:

- 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 $(\sigma_X, 1) \models \psi$. By construction and structural induction, 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, $(\sigma_X, 1) \models \psi$, and we are done.
- For the converse inclusion $\mathcal{L}(\psi) \subseteq \mathcal{L}(\varphi)$, let $\sigma_X \in \mathcal{L}(\psi)$, hence, $(\sigma_X, 1) \models \psi$. 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, $\sigma_X \in \mathcal{L}(\varphi)$.

This concludes the proof. \square

It is well-known that there is a singly exponential succinctness gap between LTL with past and NFA [17]. Consequently, we obtain the following result.

Theorem 3. *BSRV are singly exponentially more succinct than NFA, that is, there is a finite set X of input variables and a family $(\varphi_n)_{n \geq 1}$ of BSRV such that for all $n \geq 1$, φ_n has input variables in X and 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} .

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

Fig. 1. Constructions for intersection, union, and reversal.

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 done in linear time.*

Proof. We illustrate the constructions for the considered language operations.

Details on their correctness of these constructions can be found in Appendix A.1.

Intersection, Union, and Reversal. The constructions are illustrated in Fig. 1. For the intersection, assuming w.l.o.g. that the BSRV φ and φ' have no output variable in common, the BSRV recognizing $\mathcal{L}(\varphi) \cap \mathcal{L}(\varphi')$ is simply the joint set of the equations of φ and φ' . For the union, we use two new output variables *check* and *main*. Intuitively, *check* is a uniform output variable used to guess whether the input has to be considered an input for φ or for φ' . The equation for *check* ensures that the streams for *check* range over all the uniform Boolean streams. Depending on the uniform value of *check* (if it is in 0^+ or 1^+), the equation for the output variable *main* ensures that the input is recognized iff either the equations of φ are fulfilled or the equations of φ' are fulfilled. For the reversal, the BSRV recognizing $\mathcal{L}(\varphi)^R$ is obtained from φ by replacing each subexpression $\tau[k|d]$ (resp., $\tau[-k|d]$) with $k > 0$ with the subexpression $\tau[-k|d]$ (resp., $\tau[k|d]$).

Positive Kleene closure. The construction is given in Fig. 2.

The BSRV recognizing $[\mathcal{L}(\varphi)]^+$ uses two new output variables: *wbegin* and *wend*. Intuitively, *wbegin* and *wend* are used for guessing a decomposition in the given input σ_X of the form $\sigma_X = \sigma_{X,1} \dots \sigma_{X,N}$ 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* (resp., *wend*) is used to mark the first (resp., the last) positions of the components $\sigma_{X,i}$. Moreover, 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

Positive Kleene closure for $\varphi = \{y_1 = E_1, \dots, y_k = E_k\}$

$$\varphi^+ = \{y_1 = E_1^+, \dots, y_k = E_k^+, \text{wbegin} = E_{\text{wbegin}}, \text{wend} = E_{\text{wend}}\}$$

$$E_{\text{wbegin}} = \text{if } (\text{first} \rightarrow \text{wbegin}) \wedge (\text{wbegin} \rightarrow \text{wend}[-1|1]) \text{ then wbegin else } \neg\text{wbegin}$$

$$E_{\text{wend}} = \text{if } (\text{last} \rightarrow \text{wend}) \wedge (\text{wend} \rightarrow \text{wbegin}[+1|1]) \text{ then wend else } \neg\text{wend}$$

and E_i^+ is obtained from E_i by replacing each stream subexpression $\tau[k|d]$ with $E_{\tau,k,d}$:

$$E_{\tau,k,d} = \begin{cases} \text{if } \bigvee_{j=k}^{j=1} \text{wbegin}[j|1] \text{ then } d \text{ else } \tau[k|d] & \text{if } k > 0 \\ \text{if } \bigvee_{j=1}^{j=-k} \text{wend}[-j|1] \text{ then } d \text{ else } \tau[k|d] & \text{if } k < 0 \end{cases}$$

Fig. 2. Construction for positive Kleene closure

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

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

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

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

\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:

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

Fig. 3. Construction for concatenation

guessed decomposition of the input σ_X , whether $k + j$ is still a position inside $\sigma_{X,i}$.

Concatenation. The construction is given in Fig. 3. We assume w.l.o.g. that the BSRV φ and φ' have no output variables in common. The BSRV recognizing $\mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$ uses a new output variable: **wmark**. 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 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

```

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  $p$ 

```

Fig. 4. Offline monitoring algorithm for well-defined BSRV

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 . \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. Additionally, we partially solve a question left open in [8] for the case of 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 $1 \leq i \leq k$. Moreover, 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, the set of accepting runs of \mathcal{A}_φ over an input σ_X is the set of expanded valuation models of φ encoding the valuation models of φ associated with the input σ_X . Hence, the following holds.

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

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 a 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 . Then, 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 traversing of the input trace is performed, and the algorithm outputs a stream valuation of φ . Since φ is well-defined, by using

Proposition 2, we easily deduce that the uniqueness conditions in the second phase of the algorithm are satisfied. Moreover, 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. (for a detailed proof of Theorem 5, see Appendix B.1).

Theorem 5. *One can construct an offline monitoring algorithm for well-defined BSRV running 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 [8], a syntactical condition for general SRV, called *well-formedness*, is introduced, which can be checked in polynomial time and implies well-definedness. 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 [8], 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* [8], such that the monitoring algorithm processes each position of the input trace exactly $ad(\varphi)+1$ times. An important question left open in [8] is whether for a well-formed SRV φ , it is possible to construct a φ -equivalent SRV 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 the level 1. In particular, we establish the following result (the full proof can be found in Appendix B.2).

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

5 Decision Problems

We investigate complexity issues for some relevant decision problems on BSRV. 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 BSRV:*

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. *Checking semantic equivalence is EXPSPACE-complete.*
3. *When interpreted as language recognizers, language emptiness is PSPACE-complete, while language universality, language inclusion, and language equivalence are EXPSPACE-complete.*

Here, we illustrate the upper bounds of Theorem 7(1). Detailed proofs of all results can be found in Appendix 5. We need a preliminary result (Proposition 3). 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.*

Upper Bounds of Theorem 7(1). Let φ be a BSRV over X and Y , and \mathcal{A}_φ be the NFA of Theorem 1 accepting $\mathcal{L}(\varphi)$ and whose size is *singly exponential* in the size of φ .

Under-definedness: by Theorem 1 and Lemma 1, φ is under-defined iff \mathcal{A}_φ is not unambiguous. Thus, since \mathcal{A}_φ can be constructed on the fly and PSPACE = NPSPACE, by Proposition 3 (with Υ being the identity map), it follows that the under-definedness problem is in PSPACE.

Over-definedness: 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 a well-known PSPACE-complete problem [20], membership in EXPSPACE for checking over-definedness follows.

Well-definedness: it is well-known that checking universality of unambiguous NFA can be done in polynomial time [25]. By Proposition 2, φ is well-defined iff \mathcal{A}_φ is universal and unambiguous. Thus, since checking that \mathcal{A}_φ is unambiguous can be done in PSPACE (in the size of φ), membership in EXPTIME for checking well-definedness follows.

Generalized Well-definedness: let $Z \subseteq Y$. 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. 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 $\Upsilon_Z(\pi) = \Upsilon_Z(\pi')$ iff the restrictions of σ and σ' to Z coincide. By Theorem 1, we obtain that φ is well-defined with respect to Z iff \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.

6 Conclusion

In this paper, we have studied some 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: it 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 [26]—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. H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Rule-based runtime verification. In *Proc. of VMCAI'04*, LNCS 2937, pages 44–57. Springer, 2004.
2. D. Basin, M. Harvan, F. Klaedtke, and E. Zălinescu. MONPOLY: Monitoring usage-control policies. In *Proc. of RV'12*, LNCS 7687. Springer, 2012.
3. D. Basin, F. Klaedtke, and S. Müller. Policy monitoring in first-order temporal logic. In *Proc. of CAV'10*, LNCS 6174, pages 1–18. Springer, 2010.
4. A. Bauer, R. Gore, and A. Tiu. A first-order policy language for history-based transaction monitoring. In *Proc. of ICTAC'09*, LNCS 5684, pages 96–111. Springer, 2009.
5. A. Bauer, M. Leucker, and C. Schallhart. Runtime verification for LTL and TLTL. *ACM Transactions on Software Engineering and Methodology*, 20(4):14, 2011.
6. G. Berry. *Proof, language, and interaction: essays in honour of Robin Milner*, chapter The foundations of Esterel, pages 425–454. MIT Press, 2000.
7. P. Caspi and M. Pouzet. Synchronous Kahn Networks. In *Proc. of ICFP'96*, pages 226–238. ACM Press, 1996.
8. B. D'Angelo, S. Sankaranarayanan, C. Sánchez, W. Robinson, B. Finkbeiner, H. B. Sipma, S. Mehrotra, and Z. Manna. LOLA: Runtime monitoring of synchronous systems. In *Proc. of TIME'05*, pages 166–174. IEEE CS Press, 2005.
9. C. Eisner, D. Fisman, J. Havlicek, Y. Lustig, A. McIsaac, and D. V. Campenhout. Reasoning with temporal logic on truncated paths. In *Proc. of CAV'03*, volume 2725 of LNCS 2725, pages 27–39. Springer, 2003.
10. B. Finkbeiner, S. Sankaranarayanan, and H. B. Sipma. Collecting statistics over runtime executions. *ENTCS*, 70(4):36–54, 2002.

11. T. Gautier, P. Le Guernic, and L. Besnard. SIGNAL: A declarative language for synchronous programming of real-time systems. In *Proc. of FPCA'87*, LNCS 274, pages 257–277. Springer, 1987.
12. A. E. Goodloe and L. Pike. Monitoring distributed real-time systems: A survey and future directions. Technical report, NASA Langley Research Center, 2010.
13. N. Halbwachs, P. Caspi, D. Pilaud, and J. Plaice. Lustre: a declarative language for programming synchronous systems. In *Proc. of POPL'87*, pages 178–188. ACM Press, 1987.
14. D. Harel. *The spirit of Computing*. Addison-Wesley, 2nd edition, 1992.
15. K. Havelund and A. Goldberg. Verify your runs. In *Proc. of VSTTE'05*, LNCS 4171, pages 374–383. Springer, 2005.
16. K. Havelund and G. Roşu. Synthesizing monitors for safety properties. In *Proc. of TACAS'02*, LNCS 2280, pages 342–356. Springer, 2002.
17. F. Laroussinie, N. Markey, and P. Schnoebelen. Temporal logic with forgettable past. In *Proc. of LICS'02*, pages 383–392. IEEE CS Press, 2002.
18. M. Leucker and C. Schallhart. A brief account of runtime verification. *The Journal of Logic and Algebraic Programming*, 78(5):293–303, 2009.
19. Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer, New York, 1995.
20. A. R. Meyer and L. J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *Proc. of FOCS'72*, pages 125–129. IEEE CS Press, 1972.
21. L. Pike, A. Goodloe, R. Morisset, and S. Niller. Copilot: A hard real-time runtime monitor. In *Proc. of RV'10*, LNCS 6418. Springer, 2010.
22. A. Pnueli and A. Zaks. PSL model checking and run-time verification via testers. In *Proc. of FM'06*, LNCS 4085, pages 573–586. Springer, 2006.
23. G. Roşu and K. Havelund. Rewriting-based techniques for runtime verification. *Autom. Softw. Eng.*, 12(2):151–197, 2005.
24. K. Sen and G. Roşu. Generating optimal monitors for extended regular expressions. *ENTCS*, 89(2):226–245, 2003.
25. R. E. Stearns and H. B. Hunt. On the equivalence and containment problems for unambiguous regular expressions, regular grammars and finite automata. *SIAM J. Comput.*, 14(3):598–611, 1985.
26. M. Veanes, P. Hooimeijer, B. Livshits, D. Molnar, and N. Bjrner. Symbolic finite state transducers: algorithms and applications. In *Proc. of POPL'12*, pages 137–150. ACM, 2012.

Appendix

A Proofs from Section 3

A.1 Full proof of Theorem 4

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 done in linear time.*

Proof. Intersection. Let φ be a BSRV over X and Y and φ' be a BSRV over X and Y' . Without loss of generality, we assume that $Y \cap Y' = \emptyset$. Then, the BSRV φ'' 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. Let $\varphi = \{y_1 = E_1, \dots, y_k = E_k\}$ be a BSRV over X and $Y = \{y_1, \dots, y_k\}$, and $\varphi' = \{y'_1 = E'_1, \dots, y'_h = E'_h\}$ be a BSRV over X and $Y' = \{y'_1, \dots, y'_h\}$. We construct in linear time a BSRV φ'' with set of input variables X and recognizing $\mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$ as follows. The set of output variables of φ'' is given by $Y \cup Y' \cup \{\text{check}, \text{main}\}$, where **check** and **main** are two new output variables. Intuitively, **check** is a uniform output variable used to guess whether the input σ_X has to be considered an input for φ or φ' . The equation for **check** is given by

$$\text{check} = \text{if } \neg\text{last} \rightarrow (\text{check} \leftrightarrow \text{check}[+1|1]) \text{ then check else } \neg\text{check}$$

Note that the boolean streams for **check** satisfying the above equation range over all uniform boolean streams. For each $y \in Y \cup Y'$, the equation of φ'' associated to y is simply $y = y$. Finally, the equation for the new output variable **main** is defined as follows

$$\text{main} = \text{if } \left((\text{check} \rightarrow \bigwedge_{i=1}^{i=k} y_i \leftrightarrow E_i) \wedge (\neg\text{check} \rightarrow \bigwedge_{i=1}^{i=h} y'_i \leftrightarrow E'_i) \right) \text{ then main} \\ \text{else } \neg\text{main}$$

Now, we show that the construction is correct, i.e. $\mathcal{L}(\varphi'') = \mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$. For the inclusion $\mathcal{L}(\varphi'') \subseteq \mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$, let $\sigma_X \in \mathcal{L}(\varphi'')$. Hence, there is a valuation model σ'' of φ'' associated with the input σ_X . By construction, $\sigma''(\text{check}) \in 0^+ \cup 1^+$. Assume that $\sigma''(\text{check}) \in 1^+$ (the other case being similar). Then, by definition of the equation for **main**, it holds that for each $1 \leq i \leq k$, $\llbracket y_i, \sigma'' \rrbracket = \llbracket E_i, \sigma'' \rrbracket$. 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 $\mathcal{L}(\varphi) \cup \mathcal{L}(\varphi') \subseteq \mathcal{L}(\varphi'')$, 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 φ'' associated with the input σ_X defined as follows: for all variables z of φ'' , $\sigma''(z) = \sigma'(z)$ if $z \in X \cup Y'$, and $\sigma''(z) \in 0^+$

otherwise. By construction, it easily follows that σ'' is a valuation model of φ'' . Hence, membership of σ_X in $\mathcal{L}(\varphi'')$ directly follows.

Language Reversal. First, we need additional notation. 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 leading to position i . Moreover, for a stream expression E , we denote by E^R the stream expression obtained from E by replacing each subexpression $\tau[k|d]$ (resp. $\tau[-k|d]$) with $k > 0$ with the subexpression $\tau[-k|d]$ (resp., $\tau[k|d]$). By structural induction, one can easily show that for a stream valuation σ over the variables of E and for all $1 \leq i \leq |\sigma|$, it holds that $\llbracket E, \sigma \rrbracket(i) = \llbracket E^R, \sigma^R \rrbracket(R(i))$. Now, let φ be a BSRV over X and Y . The BSRV φ^R accepting the reversal of $\mathcal{L}(\varphi)$ is obtained from φ by converting each equation $y = E$ of φ in the equation $y = E^R$. Let σ be a stream valuation over $X \cup Y$. Since $\llbracket E, \sigma \rrbracket(i) = \llbracket E^R, \sigma^R \rrbracket(R(i))$ for all stream expressions E over $X \cup Y$ and position i of σ , by construction, it follows that σ is a valuation model of φ if and only if σ^R is a valuation model of φ^R . Hence, the equality $\mathcal{L}(\varphi^R) = [\mathcal{L}(\varphi)]^R$ directly follows.

Positive Kleene closure. Let φ be a BSRV over X and Y . We construct in linear time a BSRV φ'' with set of input variables X and recognizing $[\mathcal{L}(\varphi)]^+$ as follows. The set Y'' of output variables of φ'' is obtained from Y by adding two new output variables, `wbegin` and `wend`. Intuitively, these new output variables are used for guessing a decomposition in the 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` (resp., `wend`) is used to mark the first (resp., the last) positions of the components $\sigma_{X,i}$. Formally, the equations for the output variables `wbegin` and `wend` are defined as follows.

`wbegin` = *if* (`first` \rightarrow `wbegin`) \wedge (`wbegin` \rightarrow `wend` $[-1|1]$) *then* `wbegin` *else* \neg `wbegin`
`wend` = *if* (`last` \rightarrow `wend`) \wedge (`wend` \rightarrow `wbegin` $[+1|1]$) *then* `wend` *else* \neg `wend`

Note that for a stream valuation σ satisfying the above two equations, σ can be decomposed in the form $\sigma = \sigma_1 \cdot \dots \cdot \sigma_N$ for some $N \geq 1$ such that $\sigma_i(\text{wbegin}) \in 10^*$ and $\sigma_i(\text{wend}) \in 0^*1$ for all $1 \leq i \leq N$. Finally, for each output variable $y \in Y$, the equation for y of the specification φ'' is obtained from the equation $y = E$ of φ by replacing each stream subexpression $\tau[k|d]$ occurring in E , with the stream expression $E_{\tau,k,d}$ defined as follows:

$$E_{\tau,k,d} = \begin{cases} \text{if } \bigvee_{j=k}^{j=1} \text{wbegin}[j|1] \text{ then } d \text{ else } \tau[k|d] & \text{if } k > 0 \\ \text{if } \bigvee_{j=-k}^{j=1} \text{wend}[-j|1] \text{ then } d \text{ else } \tau[k|d] & \text{if } k < 0 \end{cases}$$

Now, we show that the construction is correct, i.e. $\mathcal{L}(\varphi'') = [\mathcal{L}(\varphi)]^+$. 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 construction, 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 $\sigma_\ell(\mathbf{wbegin})$ is in 10^* and the boolean stream $\sigma_\ell(\mathbf{wend})$ is in 0^*1 . 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. Fix $1 \leq \ell \leq N$. 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 $y = \mathbf{E}$ of φ and for each position $1 \leq i \leq |\sigma_\ell|$, the following holds, where $y = \mathbf{E}''$ is the equation of φ'' associated to the variable y : $\llbracket \mathbf{E}, (\sigma_\ell)_{X \cup Y} \rrbracket(i) = \llbracket \mathbf{E}'', \sigma \rrbracket(p(i))$. Evidently, by construction, we just need to prove that for each subexpression $\tau[k|d]$ of φ , $\llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i) = \llbracket \mathbf{E}_{\tau,k,d}, \sigma \rrbracket(p(i))$. 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 $\mathbf{E}_{\tau,k,d}$, it follows that $\llbracket \mathbf{E}_{\tau,k,d}, \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i)$, 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 $\mathbf{E}_{\tau,k,d}$, it follows that $\llbracket \mathbf{E}_{\tau,k,d}, \sigma \rrbracket(p(i)) = d$, 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 $\mathbf{E}_{\tau,k,d}$, it follows that $\llbracket \mathbf{E}_{\tau,k,d}, \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], \sigma \rrbracket(p(i)) = \llbracket \tau[k|d], (\sigma_\ell)_{X \cup Y} \rrbracket(i)$, 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 $\mathbf{E}_{\tau,k,d}$, it follows that $\llbracket \mathbf{E}_{\tau,k,d}, \sigma \rrbracket(p(i)) = d$, 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 $\sigma = \sigma_1 \cdot \dots \cdot \sigma_N$ for some $N \geq 1$, and σ_ℓ is a valuation model of φ for all $1 \leq \ell \leq N$. Let $\sigma'' = \sigma''_1 \cdot \dots \cdot \sigma''_N$ be the extension of σ over $X \cup Y''$, 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 $y = \mathbf{E}''$ of φ'' associated with an output variable $y \in Y$. By construction and the hypothesis, in order to show that $\llbracket y, \sigma'' \rrbracket = \llbracket \mathbf{E}'', \sigma'' \rrbracket$, 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 , $\llbracket \tau[k|d], \sigma_\ell \rrbracket(i) = \llbracket \mathbf{E}_{\tau,k,d}, \sigma'' \rrbracket(p(i))$, 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 $\mathcal{L}(\varphi'') \subseteq [\mathcal{L}(\varphi)]^+$, and we omit the details here.

Concatenation. Let φ be a BSRV over X and Y and φ' be a BSRV over X and Y' . Without loss of generality, we assume that $Y \cap Y' = \emptyset$. We construct in linear time a BSRV φ'' with set of input variables X and recognizing $\mathcal{L}(\varphi) \cup \mathcal{L}(\varphi')$ as

follows. The set Y'' of output variables of φ'' is given by $Y'' = Y \cup Y' \cup \{\text{wmark}\}$, where wmark is a new output variable. Intuitively, wmark is used for guessing a decomposition in the input σ_X'' of the form $\sigma_X'' = \sigma_X \cdot \sigma_X'$ in such a way that $\sigma_X \in \mathcal{L}(\varphi)$ and $\sigma_X' \in \mathcal{L}(\varphi')$. In particular, the output variable wmark assume the value 1 along all and only the positions of σ_X . Formally, the equation for the output variable wmark is defined as follows.

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

Note that a boolean stream for wmark is always in 1^+0^+ . The equations of φ'' for the remaining output variables are defined as follows. For each stream expression $\tau[k, d]$, let $E_{\tau, k, d}$ be the stream expression defined as follows:

$$E_{\tau, k, d} = \begin{cases} \text{if } \bigvee_{j=k} -\text{wmark}[j|0] \text{ then } d \text{ else } \tau[k|d] & \text{if } k > 0 \\ \text{if } \bigvee_{j=1}^{-k} \text{wmark}[-j|1] \text{ then } d \text{ else } \tau[k|d] & \text{if } k < 0 \end{cases}$$

First, we consider the output variables $y \in Y$. Let $y = E$ be the equation of φ associated with y . Then, the equation for y of the specification φ'' is given by $y = \text{if } \text{wmark} \text{ then } E'' \text{ else } y$, where E'' is obtained from E by replacing each stream subexpression $\tau[k|d]$ occurring in E such that $k > 0$ with the stream expression $E_{\tau, k, d}$. Now, we consider the output variables $y \in Y'$. Let $y = E$ be the equation of φ' associated with y . Then, the equation for y of the specification φ'' is given by $y = \text{if } \neg \text{wmark} \text{ then } E'' \text{ else } y$, where E'' is obtained from E by replacing each stream subexpression $\tau[k|d]$ occurring in E such that $k < 0$ with the stream expression $E_{\tau, k, d}$.

Now, we show that the construction is correct, i.e. $\mathcal{L}(\varphi'') = \mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$. For the inclusion $\mathcal{L}(\varphi'') \subseteq \mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi')$, let $\sigma_X \in \mathcal{L}(\varphi'')$. Hence, there is a valuation model σ'' of φ'' associated with the input σ_X . By construction, 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 $y = \text{if } \text{wmark} \text{ then } E'' \text{ else } y$ is the equation of φ'' associated to the variable y : $\llbracket E, (\sigma)_{X \cup Y} \rrbracket(i) = \llbracket E'', \sigma'' \rrbracket(i)$. Evidently, by construction, we just need to prove that for each subexpression $\tau[k|d]$ of φ such that $k > 0$, $\llbracket \tau[k|d], (\sigma)_{X \cup Y} \rrbracket(i) = \llbracket E_{\tau, k, d}, \sigma'' \rrbracket(i)$. This can be done similarly to the proof for the positive Kleene closure, and we omit the details here.

For the converse inclusion $\mathcal{L}(\varphi) \cdot \mathcal{L}(\varphi') \subseteq \mathcal{L}(\varphi'')$, 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, one can easily show that σ'' is a valuation model of φ'' . Hence, membership of σ''_X in $\mathcal{L}(\varphi'')$ follows. \square

B Proofs from Section 4

B.1 Proof of Theorem 5

Theorem 5. *One can construct an offline monitoring algorithm for well-defined BSRV running 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.*

Proof. The offline monitoring algorithm for well-defined BSRV is given in Figure 4, where we assume that the input trace σ_X is available on a 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) resulting 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 . Then, 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 traversing 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 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 $1 \leq i \leq |\sigma_X|$,

(*) $\Lambda(i) \neq \emptyset$ and for each state q of \mathcal{A}_φ , $q \in \Lambda(i)$ if and only if there is a run of \mathcal{A}_φ over $\sigma_X(1), \dots, \sigma_X(i)$ leading to state q .

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 . Then, by construction of the algorithm, there are a state $p \in \Lambda(i+1)$ and 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 (*) above, $\Lambda(i) \neq \emptyset$. Hence, by hypothesis, there are two distinct states $q, q' \in \Lambda(i)$ such that $p \in \delta(q, \sigma_X(i+1))$ and $p \in \delta(q', \sigma_X(i+1))$. By Condition (*) 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 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 , Theorem 5 follows. \square

B.2 Proof of Theorem 6

First, we recall the notion of well-formedness [8]. Given a general SRV 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: $y \xrightarrow{k} z$ is an edge of the graph iff for 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 φ is well-formed if its dependency graph G_φ has no cycle with weight zero. As shown in [8], well-formedness can be checked in polynomial time and *strictly* implies well-definedness. Moreover, for a well-formed SRV φ , a strongly connected component (*SCC*) of G_φ can be classified as positive or negative, where an *SCC* is positive (resp., negative) if each cycle in the *SCC* has weight strictly positive (resp., strictly negative). The alternation depth $ad(\varphi)$ of a well-formed SRV is then defined as the maximum over the number of alternations of positive and negative vertices along a path of G_φ , where a vertex is positive (resp., negative) if it belongs to a positive (resp., negative) *SCC*.

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

Proof. Let φ be a well-defined 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. We denote by \mathcal{D}_φ the DFA accepting $\mathcal{L}(\varphi)$ resulting from \mathcal{A}_φ by 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 φ' over X and $Y' \supseteq Y$ satisfying the theorem, where the set of output variables Y' is defined as follows

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

Thus, we associate to each non-initial \mathcal{A}_φ -state q , an output variable 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 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: (i) $\text{Acc}_q := 1$ if $q \in F$, and $\text{Acc}_q := 0$ otherwise, and (ii) $\text{Init}(P, \iota) = 1$ if $P = \delta_{\mathcal{D}}(\{q_0\}, \iota)$, and $\text{Init}(P, \iota) = 0$ otherwise. 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 , q , and y are defined as follows.

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

By construction, it easily follows that the BSRV φ' is well-formed (hence, well-defined too) and the alternation depth of φ' is exactly 1. It remains 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 (*): for each $1 \leq i \leq |\sigma_X|$, there is exactly one state $p \in Q \setminus \{q_0\}$ such that $\mathbf{p} \in \sigma'(i)$ (i.e., such that the boolean value of \mathbf{p} at position i with respect to σ' is 1); moreover, $p = q_i$. 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 $P \in \sigma'(i)$; moreover, $P = P_i$. 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 (*) easily follows, which concludes. \square

C Proofs from Section 5

C.1 Proof of Proposition 3

The main result of Section 5 is Theorem 7.

Theorem 7. *For BSRV:*

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. *Checking semantic equivalence is EXPSPACE-complete.*
3. *When interpreted as language recognizers, language emptiness is PSPACE-complete, while language universality, language inclusion, and language equivalence are EXPSPACE-complete.*

The upper-bounds of Theorem 7(1) were shown in Section 5. The lower-bounds are proved in Theorems 8 and 9 below. Theorem 10 corresponds to Theorem 7(3), and Theorem 11 corresponds to Theorem 7(2).

First, we include the full proof of the auxiliary result, Proposition 3.

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

Proof. The nondeterministic algorithm solving the considered problem, given the input $(\mathcal{A}, \mathcal{Y})$, proceeds as follows: 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 (resp. q') is the last state of π (resp., π'), and a flag f which is 1 iff the projections $\mathcal{Y}(\pi)$ and $\mathcal{Y}(\pi')$ 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 (hence, π and π' are two accepting runs over the same input and $\mathcal{Y}(\pi)$ and $\mathcal{Y}(\pi')$ are distinct), then 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 iff either $f = 1$ or $\mathcal{Y}(p)$ and $\mathcal{Y}(p')$ are distinct, and the whole procedure is repeated. \square

Theorem 8. *The over-definedness and generalized well-definedness problems for BSRV are both $EXPSPACE$ -hard.*

Proof. 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 and φ' be the BSRV obtained from φ by adding the equation $z = 0$, where z is a fresh output variable. Evidently, φ is not over-defined iff φ' is well-defined with respect to $\{z\}$. Hence, it suffices to prove $EXPSPACE$ -hardness for the over-definedness problem. The result is obtained by a polynomial-time reduction from a domino-tiling problem for grids with rows of singly exponential length [14]. 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:

- 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 $j < 2^m-1$, $[f(i, j)]_{right} = [f(i, j+1)]_{left}$;
- 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 $n > 0$. It is well-known that checking the existence of a tiling for \mathcal{I} is $EXPSPACE$ -complete [14]. In the

following, we construct in polynomial-time a BSRV φ such that there exists a tiling for \mathcal{I} iff φ is over-defined. The set X of input variables of φ is given by

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

Thus, we associate to each domino-type $d \in \Delta$ an input variable d . Moreover, the additional input variables $b_1^+, \dots, b_m^+, b_1^-, \dots, b_m^-$ are used to encode the value of a m -bits counter numbering the cells of one row of the grid (b_i^+ is 1 iff the i th bit is 1, and b_i^- is 1 iff the i th bit is 0). Thus, a cell is encoded as a sequence 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.¹ Moreover, a tiling is encoded as a sequence of rows, starting from the first row, and 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 otherwise.

Now, we illustrate the construction of φ ensuring that the unique inputs for which there is no output stream valuation are those encoding tilings of \mathcal{I} . We use a *uniform* output variable PTU in order to check that the input encodes a *pseudo tiling*, i.e. a tiling of \mathcal{I} , where the requirement that two adjacent cells in a column have the same color on the shared edge is relaxed. For this, we use an additional output variable PT which assumes the value 1 everywhere iff the input streams encode a pseudo tiling. The equation for variable PT is as follows.

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

¹ We assume that the first bit is the least significant one.

The equation for the *uniform* output variable PTU is as follows. We also use an extra variable test_1 in order to avoid situations where PT is everywhere 1 and the uniform value of PTU is 0.

$$\begin{aligned} \text{PTU} &= \text{if } (\neg \text{first} \rightarrow (\text{PTU}[-1] \leftrightarrow \text{PTU})) \wedge (\neg \text{PT} \rightarrow \neg \text{PTU}) \text{ then PTU else } \neg \text{PTU} \\ \text{test}_1 &= \text{if } \{((\text{first} \wedge \text{PT} \wedge \neg \text{PTU}) \rightarrow \neg \text{test}_1) \wedge (\text{last} \rightarrow \text{test}_1) \wedge \\ &\quad ((\neg \text{first} \wedge \neg \text{test}_1[-1] \wedge \text{PT} \wedge \neg \text{PTU}) \rightarrow \neg \text{test}_1)\} \text{ then test}_1 \text{ else } \neg \text{test}_1 \end{aligned}$$

Thus, the uniform value of PTU is 1 iff the input encodes a *pseudo* tiling. 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} iff there are two adjacent cells (i, j) and $(i + 1, j)$ in some column which have *different* color on the shared edge. In order to check this last condition, we use $O(m + |\Delta|)$ additional output variables. In particular, for each $h = 1, 2$, we use the output variable Bl_h for marking a cell of the pseudo tiling: if 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 ”).

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

Moreover, we use the output variables m_1 , m_2 , and test_2 in order to ensure that the cells marked by Bl_1 and Bl_2 belongs to two adjacent rows (with the Bl_2 's row following the Bl_1 's row). The stream for m_1 is in 0^*1^+ and the suffix in 1^+ starts with the first cell numbered 0 following the cell marked by Bl_1 (Bl_1 included), while the stream for m_2 is in 1^+0^+ and the prefix in 1^+ ends with the first bit of the last cell numbered 0 which precedes the cell marked by Bl_2 (Bl_2 excluded). Finally, the variable test_2 is used to check that the first 1-value bit of m_1 corresponds to the last 1-value bit of m_2 .

$$\begin{aligned} \text{m}_1 &= \text{if } \text{PTU} \rightarrow \{(\text{last} \rightarrow \text{m}_1) \wedge (\neg \text{m}_1 \rightarrow \neg \text{m}_1[-1|0]) \wedge \\ &\quad (\text{m}_1 \rightarrow \text{m}_1[+1]) \wedge ((\text{Bl}_1 \wedge \bigwedge_{i=1}^m \text{b}_i^-[i-1]) \rightarrow \text{m}_1) \wedge \\ &\quad ((\text{m}_1 \wedge \neg \text{m}_1[-1|0]) \rightarrow (\text{Bl}_1 \wedge \bigwedge_{i=1}^m \text{b}_i^-[i-1]))\} \text{ then m}_1 \text{ else } \neg \text{m}_1 \\ \text{m}_2 &= \text{if } \text{PTU} \rightarrow \{(\text{last} \rightarrow \neg \text{m}_2) \wedge (\text{first} \rightarrow \text{m}_2) \wedge (\text{m}_2 \rightarrow \text{m}_2[-1]) \wedge \\ &\quad (\neg \text{m}_2 \rightarrow \neg \text{m}_2[+1|0]) \wedge ((\neg \text{Bl}_2 \wedge \bigwedge_{i=1}^m \text{b}_i^-[i-1]) \rightarrow \text{m}_2) \wedge \\ &\quad ((\text{m}_2 \wedge \neg \text{m}_2[+1]) \rightarrow (\neg \text{Bl}_2 \wedge \bigwedge_{i=1}^m \text{b}_i^-[i-1]))\} \text{ then m}_2 \text{ else } \neg \text{m}_2 \\ \text{test}_2 &= \text{if } \text{PTU} \rightarrow \{(\text{m}_2 \wedge \neg \text{m}_2[+1]) \leftrightarrow (\text{m}_1 \wedge \neg \text{m}_1[-1|0])\} \\ &\quad \text{then test}_2 \text{ else } \neg \text{test}_2 \end{aligned}$$

Finally, we use the output variables $\text{ob}_1, \dots, \text{ob}_m$, test_3 , and od for each $d \in \Delta$, for ensuring that the cells marked by Bl_1 and Bl_2 have the same number (i.e., belong to the same column) and have different color on the shared edge. The

equation for ob_i ensures that there is a *uniform* stream for ob_i *iff* the i th bits of the cells marked by Bl_1 and Bl_2 have the same value. Moreover, for each $d \in \Delta$, the equation for od ensures that there is a *uniform* stream for od *iff* either the domino-type of the cell marked by Bl_1 is *not* d , or the cells marked by Bl_1 and Bl_2 have different color on the shared edge. Thus, the equation for variable test_3 requires that the output streams for $\text{ob}_1, \dots, \text{ob}_m, \text{od}$ ($d \in \Delta$) are uniform.

$$\begin{aligned} \text{ob}_i &= \text{if PTU} \rightarrow \bigwedge_{h=1}^2 \{(\text{Bl}_h \wedge \neg \text{Bl}_h[-1|0]) \rightarrow \\ &\quad ((\text{b}_i^+[i-1] \rightarrow \text{ob}_i[i-1]) \wedge (\text{b}_i^-[i-1] \rightarrow \neg \text{ob}_i[i-1]))\} \text{ then } \text{ob}_i \text{ else } \neg \text{ob}_i \\ \text{od} &= \text{if PTU} \rightarrow \{((\text{d}[m] \wedge \text{Bl}_1 \wedge \neg \text{Bl}_1[-1|0]) \rightarrow \text{od}[m]) \wedge \\ &\quad ((\bigvee_{d' \in \Delta: (d')_{\text{down}}=(d)_{\text{up}}} \text{d}'[m] \wedge \text{Bl}_2 \wedge \neg \text{Bl}_2[-1]) \rightarrow \neg \text{od}[m])\} \text{ then } \text{od} \text{ else } \neg \text{od} \\ \text{test}_3 &= \text{if PTU} \rightarrow \neg \text{first} \rightarrow (\bigwedge_{i=1}^m (\text{ob}_i[-1] \leftrightarrow \text{ob}_i) \wedge \bigwedge_{d \in \Delta} (\text{od}[-1] \leftrightarrow \text{od})) \\ &\quad \text{then } \text{test}_3 \text{ else } \neg \text{test}_3 \end{aligned}$$

Hence, the unique inputs for which the constructed BSRV φ has no output stream valuation are those encoding tilings of \mathcal{I} . Thus, φ is over-defined *iff* 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 Theorem 8. \square

Theorem 9. *The under-definedness and well-definedness problems for BSRV are both PSPACE-hard.*

Proof. By a polynomial-time reduction from a domino-tiling problem for grids with rows of polynomial length [14]. An instance $\mathcal{I} = \langle C, \Delta, m, d_{\text{init}}, d_{\text{final}} \rangle$ of this problem is defined as in the proof of Theorem 8. However, here, 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). It is well-known that checking the existence of a tiling for \mathcal{I} is PSPACE-complete [14]. We construct in polynomial time a BSRV φ such that the following holds:

- there exists a tiling of \mathcal{I} *iff* φ is under-defined;
- φ is *not* under-defined *iff* 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

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

Thus, we associate to each domino-type $d \in \Delta$ an input variable d . Moreover, 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 *iff* the input streams encode a tiling. Formally, the equation for variable PT is as follows.

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

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 iff φ is well-defined. Note that the size of φ is quadratic in the size of \mathcal{I} . Hence, the result follows, which concludes the proof of Theorem 9. \square

Theorem 10. *For BSRV, viewed as language recognizers, language emptiness is PSPACE-complete, while language universality, language inclusion, and language equivalence are EXPSPACE-complete.*

Proof. Recall that for NFA, emptiness is NLOGSPACE-complete, while universality, inclusion, and equivalence are PSPACE-complete. Hence, the upper bounds of the theorem directly follows from Theorem 1. The matching lower bounds for universality, inclusion, and equivalence directly follows from Theorem 8 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. It remains to show that language emptiness for BSRV is PSPACE-hard. We modify the polynomial-time reduction

given in the proof of Theorem 9 as follows: the equation for the output variable PTU of the BSRV φ is updated as follows:

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

Hence, PTU is a uniform output variable 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 Theorem 9 ensures that for each input, the output stream for the output variable PT is uniquely determined, and PT assumes the value 1 everywhere *iff* 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, and we are done. \square

Theorem 11. *Checking semantic equivalence for BSRV is EXPSPACE-complete.*

Proof. Let φ be a BSRV over X and Y , φ' be a BSRV 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 iff $\mathcal{L}(\varphi) = \mathcal{L}(\varphi')$. Hence, the lower bound directly follows from Theorem 10. Now, let us consider the upper bound. By a straightforward adaptation of Theorem 1, one can construct in singly exponential time an NFA \mathcal{A}_φ (resp., $\mathcal{A}_{\varphi'}$) over $2^{X \cup Z}$ such that $\mathcal{L}(\mathcal{A}_\varphi)$ (resp., $\mathcal{L}(\mathcal{A}_{\varphi'})$) is the set of stream valuations over $X \cup Z$ which can be extended to valuation models of φ (resp., φ'). It follows that φ and φ' are equivalent with respect to Z iff $\mathcal{L}(\mathcal{A}_\varphi) = \mathcal{L}(\mathcal{A}_{\varphi'})$. Thus, since language equivalence for NFA is PSPACE-complete, membership in EXPSPACE for the considered problem follows. \square