



Universiteit
Leiden
The Netherlands

From symbolic constraint automata to Promela

Feng, H.; Bonsangue, M.M.; Lion, B.

Citation


Feng, H., Bonsangue, M. M., & Lion, B. (2022). From symbolic constraint automata to Promela. *Journal Of Logical And Algebraic Methods In Programming*, 128.
doi:10.1016/j.jlamp.2022.100794

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3485524>

Note: To cite this publication please use the final published version (if applicable).



From symbolic constraint automata to Promela

Hui Feng^a, Marcello Bonsangue^{a,*}, Benjamin Lion^b^a Leiden University, Leiden, the Netherlands^b CWI, Amsterdam, the Netherlands

ARTICLE INFO

Article history:

Received 31 March 2022

Received in revised form 11 July 2022

Accepted 13 July 2022

Available online 26 July 2022

ABSTRACT

In this paper, we study a subclass of constraint automata with local variables. The fragment denotes an executable subset of constraint automata for which synchronization and data constraints are expressed in an imperative guarded command style, instead of a denotational style as in the coordination language Reo. To demonstrate the executability property, we provide a translation scheme from symbolic constraint automata to Promela, the language of the model checker Spin. As a proof of concept, we model in Reo a software defined network circuit, and use the Spin model checker to verify that our model satisfies some temporal properties.

© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Over the past decades, the increasing complexity of software required massive investments from both academia and industry to improve the quality. Although capable of solving complex problems, current systems are a source of difficulties for developers and designers, as they both need a formal ground for specifying such systems and verifying the intricate structures ruling the interaction among inner components. The formal description of component-based systems has been one of the major interests of Luis Barbosa, who first realized how to combine parametric behavior into a calculus of state-based software components using monads, just like in functional programming languages [1].

In this paper, we continue Luis Barbosa work on component based systems using Reo as a language to specify interaction among components explicitly [2]. In Reo, a bit of interaction among components is modeled as a connector that constrains the data flowing at their ports. Formally, a connector is a relation over the set of ports in its interface that specifies which stream of data is allowed. Generally, connectors are built compositionally, using few synchronous and asynchronous primitive connectors that compose into more complex ones [3]. A state-based specification of a connector as an automaton with constraints enables reasoning, verifying, and eventually simulating the behavior of a coordinated system. An instance of such state-based specification first appeared in [4] using constraint automata. An extension of constraint automata, adding memories [5], simplifies the specification of stateful connectors in the presence of data. However, the declarative nature of the constraints labeling transitions of constraint automata renders its translation to an imperative language (e.g., for simulation) difficult.

In our work, we consider an imperative subset of constraint automata with memory, that we call symbolic constraint automata. Intuitively, transitions in symbolic constraint automata are structured as guarded commands, where the command (i.e., an update of local memories and output port values) is executed only if the guard (i.e., a predicate on local memory values and input port values) is true. Contrary to declarative model of [5], we do not allow guards to act on output values or

* Corresponding author.

E-mail addresses: h.feng@liacs.leidenuniv.nl (H. Feng), m.m.bonsangue@liacs.leidenuniv.nl (M. Bonsangue), b.lion@cwi.nl (B. Lion).



Fig. 1. Reo circuit of Sync.

modify input values, reflecting the imperative nature of our work. The model is tailored towards an automatic translation to Promela, our main result of this paper. Our starting point is the work [6], where the authors define a Promela encoding of Reo ports, and give few examples of translations of simple Reo connectors, but leave open a systematic translation from Reo to Promela. In fact, our translation in itself is non-trivial as it requires the identification of a subset of constraint automata with memories, and the use of some control variables to circumvent syntactic restrictions in Promela. It is outside the scope of this paper to define a composition operator at Promela level compositionally reflecting Reo synchronization operator. We only show an example of composition of two connectors in Promela pointing in that direction.

Our translation of connector specifications to Promela enables the use of the Spin [7] model checker to verify temporal properties of components. Spin is a powerful model checker that can simulate and verify complex concurrent systems. For instance, it has been used for verifying spacecraft control systems [8] and avionics systems [9] and also Internet communication protocols. In fact, as a proof of concept, we apply our framework on a Reo model of software defined networks. In particular, we consider a network with a switch, programmed by a controller, that acts as a firewall by forwarding only permitted messages to one or more hosts. The switch is modeled by a symbolic constraint automaton, while the host and the controllers are modeled directly as Promela processes. We use our translation to obtain a specification of our system in Promela, that we verify against some liveness properties.

We proceed as follows. In Section 2 we introduce symbolic constraint automata, their semantics, and their composition. We give few examples of standard Reo connectors and of their composition. Section 3 describes the translation from symbolic constraint automata to Promela. In Section 4 we present our case study based on a software defined network architecture. We conclude the paper with few related works and a summary.

2. Symbolic constraint automata

In this section, we give a brief overview of some Reo connectors before presenting their encoding as a symbolic constraint automata. Also, we show how the composition method of Reo applies to symbolic constraint automata.

Reo is a coordination language used in component-based systems where components communicate through (stateful) connectors. A connector imposes data and synchronization constraints on the data flowing at its ports by allowing data to move from input ports to output ports only when all constraints are satisfied. Two connectors may share the same port name, but with different input/output polarity, in which case the data flowing at that port must be equal in both connectors. By sharing ports, connectors compose to form more complex circuits.

Reo also comes with a graphical syntax, that depicts connectors as graphs and compose connectors by merging nodes. For instance, Fig. 1 denotes a simple Reo connector modeling synchronous data flow. It connects a single input port A to a single output port B . The data received at port A is sent synchronously and without any modification to port B , thus blocking both components connected to A and B until the data flow takes place. More generally, a Reo connector can have more input ports, and output ports, and the constraints may be rather complex and conditional on the data received at (some of the) input ports. Connectors can be stateful, which means that a datum can be stored in an internal state hidden from other components.

The specification of a Reo circuit is compositional, and given by forming the product of each connector's specification. A primitive connector, e.g., a Reo channel, is usually specified using constraint automata [4]. In that case, each transition is labeled by a synchronization and relational data constraints. Note that arbitrary primitive can be defined in Reo, which may involve arbitrary data constraints. Extensions of constraint automata with memories have been considered in [5]. Still, data constraints remain declarative, as the expressiveness of the constraint language allows for arbitrarily complex label on transitions. In order to define an executable class of Reo circuits, we consider a subset of constraint automata with transitions labeled by symbolic guarded actions. We show that, under some simple consistency conditions, symbolic constraint automata can be implemented and systematically translated into the Promela language. In fact, most of original Reo basic connectors can be implemented, with as only exception filter channels with predicates that constraint output ports or transformer channels that update input ports.

The basic building blocks of a symbolic constraint automata include a finite set $\mathbb{D}$ of data ranged over by d , and a set $\mathbb{V}$ of variables, ranged over by x, y and z . We use variables to denote Reo ports shared between a connector and its environment, with different read and write permissions. An input port is a variable that the environment writes to (i.e., *put* a value into it), and from which a Reo connector destructively reads (i.e., *take* a value from it). Symmetrically, an output port is a variable that a connector writes to (i.e., *put*) and from which the environment destructively reads (i.e., *take*). Local variables are internal to a channel and can be used to store data. An assignment of variables to values is a function $\sigma: \mathbb{V} \rightarrow \mathbb{D}$. We range over input, output, and local variables by i, o , and v , respectively.

To abstract from concrete actions, we use function symbols (ranged over by f) and predicate symbols (ranged over by P). As usual, each function symbol f comes equipped with an arity and coarity, i.e. the number of arguments it expects and it returns, respectively. Similarly, predicate symbols come with an arity. We assume the natural interpretation of such function and predicate symbols, as executable function on $\mathbb{D}^n \rightarrow \mathbb{D}^m$ and as decidable subset of $\mathbb{D}^n$, respectively. To simplify the

notation, we use the same symbol to denote both the predicate and function symbols and their interpretations. Here, n is the arity and m the coarity. Syntactic substitution in a term t of every occurrence of variable x for a term t_x is denoted as usual by $t[t_x/x]$.

Terms of symbolic constraint automata are defined by the grammar:

$$t ::= d \mid x \mid f(\bar{t})$$

Terms denote tuples of data values. Here the (local or port) variable x denotes the data value it stores, while $\bar{t}$ is a shorthand notation for a finite sequence of terms $t_1, \dots, t_n$, and $f(\bar{t})$ represents a tuple of values resulting from the computation f when executed with input values $\bar{t}$. The size of these tuples depends on the arity and coarity of f .

A *guarded action* α consists of a predicate and an assignment,

$$P(\bar{x}) \rightarrow \bar{y} := t.$$

We call $P(\bar{x})$ the guard of the guarded action α , and $\bar{y} := t$ the action that is executed when the guard is true. In general, we refer to the guard of a guarded action α by $g(\alpha)$, and the assignment at the right-hand side of α by $a(\alpha)$. We implicitly assume that the size of $\bar{y}$ corresponds to the coarity of t . To avoid problems with simultaneous assignments, and without loss of generality, different variables may occur only on the left-hand side of the assignment in an action. This way, for $z \in \bar{y}$, we can denote by $a(\alpha)_z$ the z -projection of the tuple of values resulting from evaluating the term $a(\alpha)$.

Given finite subsets I, O , and V of $\mathbb{V}$ denoting some input, output, and local variables, respectively, let $Act(I, O, V)$ be the set of actions α such that all variables occurring in $g(\alpha)$ are in $I \cup V$, all variables on the left-hand side of the assignment in $a(\alpha)$ are in $O \cup V$, and all variables on the right-hand side of the assignment in $a(\alpha)$ are either in V or occurring in $g(\alpha)$. The idea is that a guard $g(\alpha)$ constrains what value input may take, based on the current value of its local variables and pending values supplied by the environment. If the guard holds then the right-hand side of $a(\alpha)$ can be satisfied using values from the local variables and values given by the environment on its inputs. The result is assigned to the variables on the left-hand side of $a(\alpha)$ changing the local store and communicating the result of the computation to the environment via the output variables. Since output ports are only used to communicate a value to the environment, we assume no occurrence of them on the guard $g(\alpha)$ and in the term on the right-hand side of the assignment $a(\alpha)$. Dually, since input ports receive values only from the environment, we assume no occurrence of them on the left-hand side of the assignment $a(\alpha)$. We denote by $I(\alpha)$ the set of all input ports occurring in α . Similarly, we denote by $O(\alpha)$ and $V(\alpha)$ the sets of output ports and local variables, respectively occurring on the left-hand side of the assignment in α . All input ports occurring in $a(\alpha)$ must appear in the input of the guard $g(\alpha)$. The text $x = x$ denotes a guard that is true, and an assignment $x := x$ denotes the skip action. We often do not write them in a transition unless the context requires it.

For example, the guarded action $(i \leq v \rightarrow (o, v) := [i, i])$ is an action in $Act(I, O, V)$, with $I = \{i\}$, $O = \{o\}$ and $V = \{v\}$. But the guarded action $(i \leq v \rightarrow (o, i) := [i, o])$ is not because i and o appear both on the left-hand side and on the right-hand side of the action, respectively. Here $[-, -]$ is the pairing function, with arity 2 and coarity 2, mapping two inputs into their corresponding pair of values.

Definition 1. A symbolic constraint automaton is a tuple $(Q, q_0, I, O, V, \longrightarrow)$, where

- Q is a finite set of states including the initial state q_0 ,
- $I \subseteq \mathbb{V}$ is a finite set of input ports, $O \subseteq \mathbb{V}$ is a finite set of output ports, $V \subseteq \mathbb{V}$ is a finite set of local variables such that they are mutually disjoint, i.e., $I \cap O = I \cap V = O \cap V = \emptyset$,
- $\longrightarrow$ is a transition relation between states and labeled by actions in $Act(I, O, V)$.

A transition $q \xrightarrow{\alpha} q'$ denotes the possibility of executing the action α from the state q and moving to the state q' . In order for the actual execution of α to take place, the guard of the action α must hold upon evaluation in the current state. To simplify the notation, we will not write guards of actions that are always true. Note that, differently from [5], we do not need to specify pre- and post-values of a local variable, as our actions are imperative and thus the order is implicitly given by the assignment operator.

In Fig. 2, we show three symbolic constraint automata. The one on the left has no internal state, and the data received at the input port i is synchronously passed to the output o . This connector is called in Reo a *synchronous channel* and is depicted as in Fig. 1. The one in the middle has three variables: one input variable i , one output variable o , and an internal variable v . The automaton has two states: state 0 to indicate that the internal variable can be rewritten (i.e., the buffer is empty), and state 1 to indicate that it cannot (i.e., the buffer is full). The connector assigns to v the value taken from i if it is in the empty state 0, and puts to the port o the value from v if it is in the full state 1. The two states symbolic constraint automaton is simpler than the equivalent single state automaton (see Fig. 3), as it avoids guards on the internal variable and the use of an extra special value $\perp$ to indicate that a variable is ‘empty’. This connector is called a *Fifo1 channel* in Reo. Finally, the rightmost automaton has a non-trivial guard labeling each of its two transitions. If the predicate P holds when a value is available at an input port i , then the connector behaves like a synchronous connector and passes the input value to the output port o . Otherwise, $\neg P$ holds on the value of i and the value is taken from i and lost, meaning that the component waiting for a synchronization on port i is released.




Fig. 2. Three examples of symbolic constraint automata.




Fig. 3. An equivalent symbolic constraint automaton for the Fifo 1 connector.

An execution of a symbolic constraint automaton $(Q, q_0, I, O, V, \longrightarrow)$ is given in terms of an infinite sequence $(\sigma_i)_{i \in \mathbb{N}}$ of assignments of values to variables for which we can find an infinite sequence of states $(q_i)_{i \in \mathbb{N}}$ starting from the initial state q_0 and such that for all $n \geq 0$ there is a transition $q_n \xrightarrow{\alpha} q_{n+1}$ satisfying the following three conditions:

1. the interpretation of the guard $g(\alpha)$ holds in the assignment σ_n ,
2. for all $x \in O(\alpha) \cup V(\alpha)$ the value $\sigma_{n+1}(x)$ is the x -projection of the evaluation of the variables on the left-hand side of the assignment in $a(\alpha)$ in the state σ_n , that is $\sigma_{n+1}(x) = \sigma_n(a(\alpha)_x)$;
3. for all variables not involved in the guarded action α , the value does not change, that is, $\sigma_{n+1}(y) = \sigma_n(y)$ for all $y \in (I \cup O \cup V) \setminus (I(\alpha) \cup O(\alpha) \cup V(\alpha))$.

Consecutive assignments σ_n and σ_{n+1} in a sequence represent the change of the internal and observable state of the system. The above conditions guarantee that the guard must hold on the values assigned to input variables in the current state before an action is taken, and that, after execution of an action, the output variables and the local variables involved in the action change according to the action taken. The last condition guarantees that only variables that occur free in an action get modified (for example by the environment) when executing that action. Note that after the execution of an action, input variables in $I(\alpha)$ change to a new value assigned by the environment. Dually, the value assigned to an output variable before the execution of an action is presumably taken by the environment before the new output value, assigned by the action, overwrites it. In other words, a transition in symbolic constraint automata, likewise constraint automata, assigns values to its output and local variables while is constrained on input and local variables. Also, the environment is allowed to change the values assigned to other variables not declared in the automaton, i.e., any variable in $\mathbb{V} \setminus (I \cup O \cup V)$. This represents the effect of an independent action executed by the environment in parallel with the automaton.

Local variables are used to store externally unobservable information. Only communication via input and output ports should be observable. Therefore, we define the semantics of a symbolic constraint automaton as the set of all possible executions $(\sigma_i)_{i \in \mathbb{N}}$ defined as above, but projected only on their input and output variables. As usual, two automata are then equivalent if they have the same semantics, i.e. they generate the same set of finite traces of assignments of input and output variables.

Consider the symbolic constraint automaton in Fig. 2.(b). The following is an example of a sequence of assignments recognized by that automaton:

$$[i = 1, o = 0, v = 0][i = 2, o = 0, v = 1][i = 2, o = 1, v = 1][i = 3, o = 1, v = 2] \dots$$

Initially, the automaton is in state 0, for example, with value 1 on the input port i . The values of the two other variables do not matter at this point, and can be seen as previous values that remained stored but not accessible. By taking the transition to the state 1, the connector assigns the value of i to the internal variable v . The output port o is blocked and cannot be changed while executing this transition, while the input port is free and here is assumed to get the value 2 from the environment. When taking the next transition the content of the variable v is put in the port o , the input is blocked and the cycle can start again. For each such a sequence we can find an equivalent sequence for the automaton in Fig. 3, by using the extra value $\perp$ to check if the buffer is empty

$$[i = 1, o = 0, v = \perp][i = 2, o = 0, v = 1][i = 2, o = 1, v = \perp][i = 3, o = 1, v = 2] \dots$$

Conversely, for every sequence representing the behavior of the automaton in Fig. 3 we can find an equivalent sequence of assignments for the automaton in Fig. 2.(b) by copying the previous value of v instead of $\perp$, and assigning an arbitrary initial value for v . In other words, the two automata are equivalent. Note that the initial state of the automaton in Fig. 3 forces the connector to start with an empty buffer. Without this initial transition, the two automata would not be equivalent as one could start to output on port o the value stored in v .

The central operation on symbolic constraint automata is synchronization via their shared ports that are input ports for one automaton and output ports for another. Shared ports become internal local variables in the automaton resulting from the composition. No other synchronization by shared variables is allowed, as local variables are only visible within the scope of a connector. Our definition is similar in spirit to that of [5], but, in addition to that work, our symbolic constraint automata could be automatically translated to Promela. The explicit input and output variables, and the guarded command structure on the label impose some prerequisites on the product to avoid inconsistencies. In fact, we define composition only for pairs of symbolic constraint automata A_1 and A_2 such that (1) no local variables are in common, and (2) for every pair of actions α_1 and α_2 of the two automata, they synchronize only on some input ports used by one action and some output ports used by the other, but not on both input and output ports at the same time. More formally, we assume that, for all actions α_1 labeling a transition in A_1 and α_2 labeling a transition in A_2 , the following holds

$$I(\alpha_1) \cap O(\alpha_2) \neq \emptyset \Rightarrow O(\alpha_1) \cap I(\alpha_2) = \emptyset \quad \text{or, equivalently,} \\ O(\alpha_1) \cap I(\alpha_2) \neq \emptyset \Rightarrow I(\alpha_1) \cap O(\alpha_2) = \emptyset$$

We call two automata with these two properties *consistent*. The intuition behind synchronizing two guarded actions α_1 and α_2 is that their data value should agree on their shared ports so that it can flow from the output of one to the input of the other actions. The above condition together with the fact that the two automata do not share local variables - and thus $V(\alpha_1) \cap V(\alpha_2) = \emptyset$ - in fact impose a causality in the execution of their actions as input is needed to update the internal state and to be passed to output ports. Next we define formally the synchronization of two guarded actions α_1 and α_2 . Assume $I(\alpha_1) \cap O(\alpha_2) = \bar{u} \neq \emptyset$, and let

$$\alpha_1 = P_1(\bar{x}_1 \cup \bar{u}) \rightarrow \bar{y}_1 := t_1(\bar{z}_1 \cup \bar{u}) \quad \text{and} \quad \alpha_2 = P_2(\bar{x}_2) \rightarrow \bar{y}_2 \cup \bar{u} := t_2(\bar{z}_2),$$

where $\bar{x}_1 \cup \bar{u}$ is the sequence of input and local variables occurring in P_1 , $\bar{z}_1 \cup \bar{u}$ the sequence of input and local variables occurring in t_1 , and $\bar{y}_2 \cup \bar{u}$ the sequence of output and local variables occurring at the left-hand side of the assignment $a(\alpha_2)$. Note that variables in $\bar{u}$ cannot occur in P_2 nor in t_2 as they are output variables for α_2 . Similarly, they cannot occur in $\bar{y}_1$, as they are input variables for α_1 . By definition of guarded action, they can occur in P_1 . Under these circumstances, we can define the synchronization $\alpha_1 \otimes \alpha_2$ as the following guarded action

$$\alpha_1 \otimes \alpha_2 = P_1[t_{2_{\bar{u}}}/\bar{u}] \wedge P_2 \rightarrow \bar{y}_1, \bar{y}_2 \cup \bar{u} := t_1[t_{2_{\bar{u}}}/\bar{u}], t_2(\bar{z}_2).$$

Since the guard P_2 does not depend on output variables, we can evaluate it. If it holds we can then assign the values returned by t_2 to the output and local variables at the left-hand side of $a(\alpha_2)$. The values assigned to the shared output values $\bar{u}$ are then used in P_1 as constant replacing the input variables $\bar{u}$. If this predicate holds, then the same substitution is applied to t_1 so that we can compute the assignment. Note that $I(\alpha_1 \otimes \alpha_2) = (I(\alpha_1) \setminus O(\alpha_2)) \cup I(\alpha_2)$, $O(\alpha_1 \otimes \alpha_2) = O(\alpha_1) \cup (O(\alpha_2) \setminus I(\alpha_1))$, and $V(\alpha_1 \otimes \alpha_2) = V(\alpha_1) \cup V(\alpha_2) \cup (I(\alpha_1) \cap O(\alpha_2))$. The definition of $\alpha_1 \otimes \alpha_2$ for the symmetric case when $O(\alpha_1) \cap I(\alpha_2) \neq \emptyset$ is similar.

Definition 2. Without loss of generality, let Q_1 and Q_2 be two disjoint sets of states, and V_1 and V_2 be two disjoint sets of local variables. The composition of two consistent symbolic constraint automata $A_1 = (Q_1, q_1, I_1, O_1, V_1, \longrightarrow_1)$ and $A_2 = (Q_2, q_2, I_2, O_2, V_2, \longrightarrow_2)$ is defined as the automaton $A_1 \bowtie A_2 = (Q, q, I, O, V, \longrightarrow)$ where:

- $Q = Q_1 \times Q_2$,
- $q = \langle q_1, q_2 \rangle$,
- $I = (I_1 \setminus O_2) \cup (I_2 \setminus O_1)$,
- $O = (O_1 \setminus I_2) \cup (O_2 \setminus I_1)$,
- $V = V_1 \cup V_2 \cup (I_1 \cap O_2) \cup (I_2 \cap O_1)$, and
- $\longrightarrow$ is defined by the following rules:

$$\frac{q_1 \xrightarrow{\alpha_1}_1 q'_1 \text{ and } q_2 \xrightarrow{\alpha_2}_2 q'_2 \text{ and } IO(\alpha_1) \cap IO(A_2) = IO(\alpha_2) \cap IO(A_1)}{\langle q_1, q_2 \rangle \xrightarrow{\alpha_1 \otimes \alpha_2} \langle q'_1, q'_2 \rangle} \\ \frac{q_1 \xrightarrow{\alpha_1}_1 q'_1 \text{ and } IO(\alpha_1) \cap IO(A_2) = \emptyset}{\langle q_1, q_2 \rangle \xrightarrow{\alpha_1} \langle q'_1, q_2 \rangle} \quad \frac{q_2 \xrightarrow{\alpha_2}_2 q'_2 \text{ and } IO(\alpha_2) \cap IO(A_1) = \emptyset}{\langle q_1, q_2 \rangle \xrightarrow{\alpha_2} \langle q_1, q'_2 \rangle}$$

where $IO(A_i) = I_i \cup O_i$ and $IO(\alpha_i) = I(\alpha_i) \cup O(\alpha_i)$, for $i = 1, 2$.

Similar to the join operation on constraint automata the above synchronization operation on symbolic constraint automata synchronizes actions on shared ports and allow independent parallel behavior for actions with no shared ports. The composition of consistent symbolic constraint automata is symmetric and, when defined, associative. More examples of symbolic constraint automata can be found in the Appendix.

3. From Reo to Promela

Promela is a formal language [10] widely used to specify concurrent systems and is supported by Spin, an LTL model checker. A Promela program is composed of a set of processes. Processes run concurrently and interact through shared channels. Both synchronous and asynchronous communication between processes are supported. For synchronous communication, a channel works in a rendezvous mode with no buffer (zero capacity). For asynchronous communication, channels work as a Fifo buffer with a non-zero user specified capacity.

Spin is a model checker with Promela as specification language. In a nutshell, each Promela process is transformed by Spin into a finite state automaton. Processes are then synchronized into a single system automaton. Similarly, linear temporal properties expressed in the usual LTL syntax, are transformed into finite state automata. The automata representing both the Promela program and the LTL properties are exploited by Spin to verify and assert satisfaction of the properties [10].

Building on the work presented in [6], we translate some Reo connectors into Promela programs. We use symbolic constraint automata as a specification of Reo connectors, and use the resulting Promela process as a protocol to coordinate messages exchanged through the ports of other processes.

3.1. Implementing Reo ports in Promela

In Promela, a Reo port is expressed as a structure, as shown in Listing 1. It contains two Promela channels of capacity one: a data and a trig channel. A port in Reo is directional. We call putter the component that puts an element on the port, and getter the component that gets an element from the port. Operations on a port are blocking unless both a put and a get are performed at the same time. In which case the port *fires*, and the data is forwarded from the putter to the getter.

The data channel in the Promela implementation of a Reo port is used to forward the data message from the putter to the getter, while the trig channel is used to synchronize the putter and the getter. The reason of using two channels of size one instead of a rendezvous channel of size zero is that it is impossible in Promela to query whether a process is currently waiting on a rendezvous channel. We will later see that querying the state of a port is necessary for the protocol to coordinate the boundary processes.

```

1 typedef port {
2   chan data = [1] of {Data};
3   chan trig = [1] of {int}; }

```

Listing 1: Definition of a Reo port in Promela.

As described in Listing 2, two actions can be performed on a port: put and take. The function call `put(q, a)` atomically fills the data channel of `q` with the datum `a`, and blocks on the trig channel, waiting to synchronize with the component on the output side of `q`. The integer variable `x` is used to get a value from the trig channel, and hence to synchronize to it; the actual value communicated does not matter.

The function call `take(q, a)` atomically notifies, by outputting on the trig channel, that there is a component willing to take data, and blocks on the data channel, until a datum can be read and stored into the variable `a`. The integer value of `-1` written into the synchronization channel is arbitrary, as trig is used only for signaling.

```

1 inline put(q, a) {
2   int x;
3   q.data!a;
4   q.trig?x }
5
6 inline take(q, a) {
7   q.trig!-1; q.data?a }

```

Listing 2: put and take functions.

We describe two temporal properties in Listing 3 that reflect the synchronous behavior of a port. We say that a port *fires* whenever a data is exchanged between the putter and the getter. If a port does not fire, it is *silent*. In the case of an implementation of a port with two buffers, the firing property occurs whenever a port has both an input and an output request, i.e. both channels are full. We then know, due to the definition of the put and take operations, that the putter and the getter will be released from blocking on the port, and the getter will get the value from the data channel. We define some macros in Listing 3 to encode firing and silent property of port `p` as an LTL property.

Ports are not typed as input or output, as that depend on the component/connector that use them. We have seen in the previous section that within a connector, a port used in a guard must be an input port, whereas a port used on the left-hand side of the assignment of an action is an output port.


```

1 #define p_fires (
2   !(len(p.data) == 0) && !(len(p.trig) == 0) &&
3   X((len(p.data) == 0) || (len(p.trig) == 0)))
4
5 #define p_silent    (! p_fires)

```

Listing 3: Macros for firing of ports.

3.2. Implementing Reo connectors in Promela

Next we describe how to implement Reo connectors expressed as symbolic constraint automata in Promela. Let $A = (Q, q_0, I, O, V, \longrightarrow)$ be a symbolic constraint automaton.

A symbolic constraint automaton is encoded, in Promela, as a proctype. A Promela proctype has a name, a signature, and a body. The Spin model checker executes each proctype of its main concurrently, while taking into account blocking operations on channels. As we will see, input/output variables of a symbolic constraint automaton lead to shared port channels in Promela between the protocol and the boundary processes.

Input ports I and output ports O are passed as parameters to the Promela process resulting from the translation. As expected, local variables in V are declared locally to the Promela process, i.e., within the proctype body. For each port $P \in I \cup O$, a new local variable $_p$ is declared so to store the value taken, if P is an input, or passed, if P is an output. We use `mytype` as a generic type for input, output, and memory variables.

Each state in Q is encoded as a special value for the state variable. The state variable therefore models the control flow between the states of the automaton. For simplicity, and without loss of generality, here we assume that $q_0 = 0$ is the initial state and $Q = \{0, 1, \dots, n\}$.

We use the Promela non-deterministic `do – od` construct to model concurrent applications of transitions of a symbolic constraint automaton. The guard (respectively, the command) of the statements in the `do – od` results from the translation of the guard (respectively, the command) in the action labeling the corresponding transition. As a result, each guard in the `do – od` loop contains a clause that controls that variable state has the value corresponding to the pre-state, and updates the value to the post-state in the command. The Promela language allows for non-destructive reads of channel's value: the operation `A.data? <_a>` assigns to the variable `_a` the value stored in channel `A.data` without actually removing the values from that channel. However, in Promela, this operation cannot be executed in a guard within a `do – od` statement. We circumvent this problem by using a control variable f_i for each transition $t_i \in \longrightarrow$ of the automaton. The variable f_i can take two values. By default, f_i is 1. If the synchronization constraint of the guard of the i -th transition is satisfied but the data constraint is not, then the value of f_i is set to 0. Every transition, if taken successfully, resets all the f_i to 1.

More formally, for a symbolic constraint automaton $(Q, q, I, O, V, \longrightarrow)$ we show in Listing 4 its translation to Promela:

```

1 proctype SCA(port  $\overline{P}$ ) {
2    $\forall p \in P, \text{ mytype } \_p;$ 
3    $\forall v \in V, \text{ mytype } v;$ 
4    $\forall f \in \{f_1, \dots, f_k\}. \text{ bool } f = 1;$ 
5   int state=0;
6   do
7     :: transition 1;
8     :
9     :: transition k;
10  od }

```

Listing 4: Promela code generated for a symbolic constraint automaton.

Here $Q = \{0, \dots, n\}$, $q = 0$, $P = I \cup O$, $\overline{v} = V$, and the remaining overlined variables are just consecutive sequences of them. For every (input or output) port $p \in \overline{P}$ there is a variable $_p$ associated with it on which we store the value communicated via the port. The control variable `state` is used to store the current state of the automaton (thus it ranges between 0 and n), and variables $\{f_1, \dots, f_k\}$ are used when evaluating predicates on values received at input ports. Here k is the number of transitions. Let $q \xrightarrow{\alpha} q'$ be the j -th transition of the symbolic constraint automaton, with $1 \leq j \leq k$, and remember that the input ports in α are $I(\alpha) = \{i_1, \dots, i_m\}$ the output ports are $O(\alpha) = \{o_1, \dots, o_l\}$, and the local variables occurring in α are $\{v_1, \dots, v_h\}$. Then transition j in the listing above is given by

```

1 state==q && fj==1 &&  $\bigwedge_{i \in I(\alpha)} \text{Full}(i.data) \text{ \&\& } \bigwedge_{o \in O(\alpha)} \text{Full}(o.trig)$ 
2 -> Atomic {  $i_1.data?<\_i1>; \dots; i_m.data?<\_im>;$ 
3   if
4     ::  $P(i_1, \dots, i_m, v_1, \dots, v_h) == \text{True} \rightarrow \text{take}(i_1, \_i1); \dots; \text{take}(i_m, \_im);$ 
5       A( $\_i1, \dots, \_im, v_1, \dots, v_h, \_o1, \dots, \_ol$ );
6       put( $o_1, \_o1$ );  $\dots; \text{put}(o_l, \_ol);$ 
7       state= $q'$ ;  $f_1=1; \dots; f_k=1;$ 
8     :: else -> fj=0;

```


9 *fi* }

where P is a function encoding the guard of α and A is a function encoding the assignment action of α . When the control variables are not used in a transition (for example because there is no predicate on input variables), then it is possible to simplify the generated Promela code by removing such control variables. Similarly, when there is only one state, the variable state can be removed as its value is constant. In the Appendix, we show more examples of standard Reo connectors encoded as symbolic constraint automata and translated to Promela. It is important to remark that the control variables f_i do not introduce fairness or priority among the transitions, they only control the flow so that Promela will not choose the same transition again with a guard that has already been evaluated to false.

We list several examples of translation of symbolic constraint automata to Promela. Listing 5 gives the Promela code resulting from the translation of the symbolic constraint automaton in Fig. 2.(a) that models the *Sync* connector.

```

1 proctype Sync(port A; port B){
2   mytype _a; mytype _b; //internal values in Sync
3   int state = 0; //initial state
4   do
5     :: state==0 && Full(A.data) && Full(B.trig)
6       -> Atomic{take(A, _a); _b =_a; put(B,_b); state = 0};
7   od }
```

Listing 5: Promela code generated for a Sync connector.

Here the parameter portA is the input port of the Sync channel and portB is the output port. There are no local variables except those associated with the ports, and the control variable state that we will discuss later. The condition `full(A.data)&&full(B.trig)` is satisfied if and only if there is an ongoing `put(A, a)` operation on the input port A and an ongoing `take(B, x)` operation on the output port B. In this case, the Sync process atomically takes the data from port A, stores it in a local variable `_a`, executes the action of the associated transition of the symbolic constraint automata, i.e. `_a = _b`, and puts the value in port B. Of course, the single state automaton of the Sync connector could have been modeled by a much simpler Promela process without such an extra variable, but, as for the other connectors, we keep it for generality.

Listing 6 shows the Promela code for a two states symbolic constraint automaton of the *Fifo1* Reo connector given in Fig. 2.(b).

```

1 proctype Fifo1(port A; port B){
2   mytype _a; mytype _b; int state = 0;
3   mytype v; //v is the buffer of the Fifo 1 connector
4   // _a,_b,v are the input, output and local variables, respectively
5   do
6     :: Full(A.data) && state==0
7       -> Atomic{take(A,_a); v=_a; state=1};
8     :: Full(B.trig) && state==1
9       -> Atomic{_b=v; put(B,_b); state=0};
10  od }
```

Listing 6: Promela code generated for Fifo1 connector.

Besides an input and an output port, this time we also have a local variable `v` for the buffer of the *Fifo1* connector. The control variable state is initially set to 0, corresponding to the initial state of the automaton.

This time, the `do – od` loop contains two transitions, one for each transition of the symbolic constraint automaton. One statement corresponds to the transition that moves the control from `state = 0` to `state = 1` if there is a pending data on the input port A. The value is then taken and assigned to the local variable `v`. The other statement corresponds to the transition that moves the control from `state = 1` to `state = 0` when there is a pending request at the output port. In which case, the stored value is forwarded to the output port B.

Next, we show how to translate the Filter connector of Fig. 2.(c) that, unlike the other two examples, contains two transitions labeled with a non-trivial predicate on the input variable. The Promela code of the Filter connector is presented in Listing 7.

Initially the control variables `f1` and `f2` associated to the two transitions are set to true, meaning that any transition can be potentially selected. As before, the satisfiability of each guard depends on the presence of some data at the input port A and the presence of some signal at the output port B. The predicate P is evaluated only after one of the two statements of the `do – od` loop is chosen. If true, the action of the transition is taken and the two control variables `f1` and `f2` are reset to true. Otherwise the value of the associated control variable `fi` is set to false and the control goes back to the loop statement. In this way the i -th transition associated with `fn` will not be selected anymore, even if all other predicates in the guard of the statement are true (for $i = 1, 2$), which removes some undesirable livelocks.

```

1 proctype Filter(port A; port B){
2   mytype _a; mytype _b; int state=0; bool f1=1; bool f2=1;
3   do
4     :: state==0 && f1==1 && Full(A.data) && Full(B.trig)
5     -> Atomic{ A.data ? <_a>;
6       if
7         :: P(_a)==True -> take(A,_a); _b=_a; put(B,_b); state=0; f1=1; f2=1;
8         :: else -> f1=0;
9       fi
10    :: Full(A.data) && state==0 && f2==1
11    -> Atomic{ A.data ? <_a>;
12      if
13        :: P(_a)==False -> take(A,_a); state=0; f1=1; f2=1;
14        :: else -> f2=0;
15      fi
16    od }

```

Listing 7: Promela code generated for Filter connector.

We leave for future work a detailed description of the encoding in Promela of a composition operator mimicking that of symbolic constraint automata. We just give an example below corresponding to the composition of the Promela code generated for the Fifo1 and the Filter connectors in Listings 6 and 7, respectively.

```

1 proctype FilterFifo1(port A; port C){
2   mytype _a; mytype _c;
3   mytype v; //v is the buffer of the Fifo 1 connector
4   mytype b; //b is a shared port that becomes a local variable
5   int state=0; //there are 1 x 2 states in total
6   bool f1=1; bool f2=1; bool f3=1; //there are 3 transitions in total
7   do
8     :: state==0 && f1==1 && Full(A.data)
9     -> Atomic{ A.data ? <_a>;
10      if
11        :: P(_a)==True -> take(A,_a); _b=_a; v=_b; state=1; f1=1; f2=1; f3=1;
12        :: else -> f1=0;
13      fi
14     :: state==0 && f2==1 && Full(A.data)
15     -> Atomic{ A.data ? <_a>;
16      if
17        :: P(_a)==False -> take(A,_a); state=0; f1=1; f2=1; f3 =1;
18        :: else -> f2=0;
19      fi
20     :: state==1 && f3==1 && Full(C.trig)
21     -> Atomic{ _c=v; put(C,_c); state=0; f1=1; f2=1; f3 =1;};
22   od }

```

Listing 8: Promela code generated by composing Filter(portA;portB) with Fifo1(portB;portC).

4. A case study: verifying a software defined network

In this section, we apply our translation of symbolic constraint automata to Promela on a software defined network model, and use the Spin model checker to verify several temporal properties. We use the model of software defined networks introduced in [11], where all components of an SDN are represented as Reo connectors, and thus as symbolic constraint automata. The model is stateful and reflects the SDN separation between switches, controllers, and network. For each part, we briefly describe the Promela code obtained from the Reo components, and how we model the basic data flow operations of an SDN: PktIn, PktOut, FlowMod.

4.1. Preliminaries

An SDN consists of three planes: the application plane, the data plane, and the control plane. The data plane contains switches and the network infrastructure for distributing and receiving packets. The control plane consists of a set of operations that can be used to program the network switches. In short, the control planes are formed by a set of reactive controllers that communicate with the switches to program the behavior of the network. Controllers receive packets from the switches via the PktIn command, and send messages to them via the PktOut and the FlowMod commands. The former consists of a packet together with the set of ports of each switch the packet needs to be sent, the latter is used to update the flow table of the switch. A flow table is a list of pairs matching-action, where matching is a predicate evaluated on the header of incoming packets, and action is a set of ports where the packets need to be forwarded to after the




Fig. 4. Reo circuit of one switch.




Fig. 5. A simple example of an SDN architecture.

matching. When the set of ports is empty, the packet is dropped. Unlike traditional switches, a flow tables in an SDN can be dynamically updated via `FlowMod` operations.

In [11] we presented a generic SDN model built with a set of stateful and typed Reo connectors. Here we briefly summarize the model of a switch and refer the reader for more technical details to the original work. The Reo circuit of a SDN switch given in Fig. 4 is parametric on the input ports $P_i, i \in \{0, \dots, n\}$ and output ports $Q_j, j \in \{0, \dots, m\}$. It is composed out of several transformers and filter channels that add relevant information and route the packets received at input ports to the right output port. The core of the switch is the flow table, here represented by the τ in the variable channel. The table can be updated by `FlowMod` operations or simply consulted for routing packets. Switches are connected with SDN controllers via synchronous channel between P_0 and the output port of the controller, and a bounded queue channel between Q_0 and the input port of the controller. `FlowMod` and `PktOut` messages are received by the switch at port P_0 and normal packets are received from all other input ports. `PktIn` messages need not be evaluated by the flowtable and can immediately be redirected to the relevant output ports, whereas `FlowMod` messages follow another route and will update the flow table. Normal packets are evaluated according to the rules in the flow table τ , and accordingly they will be either deleted or sent to one or more output ports. `PktIn` messages are generated after the evaluation of packets at the flow table and redirected to the controller via the output port Q_0 , for example to request a table update.

4.2. A Promela SDN model via symbolic constraint automata

Next we briefly describe the generated Promela code for the Reo model of a switch `Switch(P0, P1, P2, Q0, Q1, Q2)` with two input ports $P1$ and $P2$, two output ports $Q1$ and $Q2$ as shown in Fig. 5. As explained above, those ports form the interface of the switch with the rest of the network. Additionally, the switch interface is extended with an input port $P0$ and an output port $Q0$ that serve to exchange flow messages with the SDN controller that we can model directly in Promela code. Flow messages from the controller `Controller(P, Q)` to the switch are exchanged synchronously, via the synchronous connector `Sync(P, P0)`. On the other direction, flow messages are exchanged asynchronously via a queue connector `Queue(Q0, Q)` from the port $Q0$ of the switch to the port Q of the controller.

Informally, the symbolic constraint automaton of a switch consists of a single state and few transitions labeled by the following type of guarded actions:

- α_0 is executed when a `FlowMod` message is received from the input port $P0$. The action here consists in updating the flow table according to the information sent by the controller.

- α_1 is enabled when a PktOut message is received by the switch from port P0. The resulting action forwards a packet to a subset (possibly empty) of output ports of the switch. This subset is contained in the PktOut message from the controller.
- α_2 is enabled when one of the input port P1 or P2 of the switch receives a normal packet that is then forwarded to a subset (possibly empty) of output ports according to the current information in the flow table of the switch.

The Promela code of the switch obtained via the translation from a symbolic constraint automaton is presented (in a simplified manner for reason of space) below. The full code and the results of its verification can be found in [12].

```

1 proctype Switch(port P0, P1, P2, Q0, Q1, Q2){
2   Message m;    // either FlowMod or PktOut
3   Packet p;
4   Flowtable ft_old; ft_new
5   bool f1=1; f2=1; ...;fn=1; //control variables
6   int state=0;   //automaton current state
7   do
8     //packet from P1 to Q1: this is a packet-forwarding  $\alpha_2$  type of action
9     :: state==0 && f1==1 && Full(P1.data) && Full(Q1.trig)
10    -> atomic{ take(P1,p); Match(ft_old,p,[ Q1 ]); put(Q1,p); f1=1;...fn=1}
11    //packet from P1 to both Q0 (PktIn message to controller) and Q2:
12    // this is a  $\alpha_2$  type of action
13    :: state==0 && f2==1 && Full(P1.data) && Full(Q0.trig) && Full(Q1.trig)
14    -> atomic{ take(P1,p); Match(ft_old,p,[ Q0,Q1 ]); put(Q0,p);put(Q1,p);
15              f1=1;...fn=1 }
16    //message from P0 to Q1 and Q2: this is an PktOut  $\alpha_1$  type of action
17    :: state==0 && f3==1 && Full(P0.data) && Full(Q1.trig)&& Full(Q2.trig)
18    -> atomic{ P0.data?<m>;
19              if
20                :: m==<p,[Q1,Q2]> -> take(P0,m);put(Q1,p);put(Q2,p); f1=1;...fn=1;
21                :: else -> f3=0;
22              fi }
23    //message from P0 to update flow table: this is a FlowMod  $\alpha_0$  type of action
24    :: state==0 && f4==1 && Full(P0.data)
25    -> atomic{ P0.data?<m>;
26              if
27                :: m==<p,ft_new> -> take(P0,p);update(ft_old,ft_new); f1=1;...fn=1;
28                :: else -> f4=0;
29              fi }
30    //Similar actions follows here....
31    ...
32  od}

```

Each transition synchronizes some of the actors in a network (hosts, switches, controllers). Packet forwarding is done on the basis of the result of a function Match() of the Reo transformer channel between ports *D* and *E* in Fig. 4. Another non-trivial function used here is update(), belonging to the transformer between ports *F* and *E* in Fig. 4 and used to model a FlowMod operation.

In our case study as described in Fig. 5, the network consists of a switch programmed by a controller, where hosts *A* and *B* produce packets, and hosts *C* and *D* consume them. We abstract from the specific behavior of the hosts and model them simply as producers and consumers of messages, respectively.

```

1 proctype HostA(port A){
2   packet p1;
3   atomic{p1.header = 11; p1.ipt = P1; put(A,p1)}
4 }

```

Host *A* produces a single packet that is sent to port *P1*. Here we assume a packet contains a header (with information such as the tcp/ip source or destination), and the port of the switch it is supposed to be received directly from the host.

```

1 proctype HostB(port B){
2   packet p2;
3   p2.ipt = P2;
4   do
5     :: atomic{p2.header = 11; put(B,p2)};
6     :: atomic{p2.header = 22; put(B,p2)}
7   od
8 }

```

The above Promela code for host *B* is similar to that of *A*, except that host *B* repeatedly sends packets with header 11 or 22 to port *P2*. Hosts *C* and *D* are consumers that repeatedly execute the take action from their ports *C* and *D* respectively. Once a packet is received, they update their own local counter storing the number of packets with header 11 that they receive. Below we show the Promela pseudocode of host *C*, that of *D* is similar.

```

1 proctype HostC(port C){
2   packet q1;
3   int counter=0;
4   do
5     :: atomic{take(C,q1); if q1.header==11 -> counter++;}
6   od
7 }

```

Finally we give the Promela code of the controller. The controller takes a packet from port Q if available. If the packet originally passed through port $P1$ then the controller adds a rule in the flow table to forward similar messages (i.e., coming from the same address as in the header) to port $Q1$. Packets with header 22 need to be forwarded to both ports $Q1$ and $Q2$ (thus to hosts C and D). Finally, the following firewall is installed: packets with header 11 that have passed through port $P2$ must be dropped. Note that the controller will insert rules into the flow table of the switch to execute the above commands, and will apply the action itself only the first time a packet is not matching any rule in the flow table, i.e., the first time the packet is forwarded to the controller.

```

1 proctype Controller(port Q, P){
2   Flowtable ft;
3   Packet p;
4   Message m;
5   do
6     :: Full(Q.data) && (Q.data.ipt==P1) ->
7       take(Q, p);
8       ft.cond = p.header; ft.action = [Q1];
9       //create a rule: if match header then forward to Q1
10      m=<p,ft>; put(P, m); //update flow table
11      m=<p,[Q1]>; put(P, m); // Forward p to port Q1;
12     :: Full(Q.data) && (Q.data.header==11) && (Q.data.ipt==P2) ->
13       //insert a firewall rule: no message from P2
14       take(Q,p);
15       ft.condition = p.header; f.action = []; // drop package if comes from P2
16       m=<p,ft>; put(P,m); //update flow table
17     :: Full(Q.data) && (Q.data.header==22) ->
18       take(Q, p);
19       ft.cond = p.header; ft.action = [Q1,Q2];
20       //create rule: if match header then forward to Q1 and Q2
21       m=<p,ft>; put(P, m); //update flow table
22       m=<p,[Q1,Q2]>; put(P, m); // Forward p to port Q1 and Q2;
23   od
24 }

```

Listing 9: An example of controller.

Port Q of the controller is linked to the output port of the queue connector $Queue(Q0, Q)$ which may store at most 10 packets. This number is reasonable as we do not expect many $PktIn$ messages to be forwarded to the Controller.

4.3. Verification and simulation

Fig. 6 describes the scenario for which host A sends a packet with header 11 *after* host B sent a packet with header 22. Here the packet of A will arrive to C by first passing through the controller, and the packet from host B will arrive to both hosts C and D . The LTL properties associated with this scenario are

```

prop1 { [] ((p1.header==11 && p1.ipt==P1) -> <> (q1.header==11)) }
/* satisfied */
prop2 { [] ((p2.header==22) -> (<> (q1.header==22) && <> (q2.header==22))) }
/* satisfied */

```

Intuitively they say: The message of host A will receive B (no loop holes) and always, if B sends a 22-message, then both eventually C receives a 22-message (but not necessarily the same one), and eventually the same for D . Assuming that only B sends 22-messages then this means that every message with header 22 received from hosts C or D is originated from host B .

Fig. 7 describes the scenario when the packet from host A is received at the switch. As the flow table of the switch is empty, the switch forwards the packet to the controller. However, the packet with the same header arriving from B is dropped. Of course, if the packet B arrives at the switch after the updates of the flow table due to the first message from host A , then the packet of B will match the flow table and be redirected to host C , violating the firewall rule. We can verify this formally in Spin via the following LTL property:

```

prop3 { <> (msg1.header==11) && (<> [] (HostC_counter==1)) } /* unsatisfied */

```




Fig. 6. Packets from A and B arriving both to C.




Fig. 7. Are packets from B with header=11 always dropped?

Table 1

Verification results.

Property	Errors found	Time usage	Depth reached	States stored	States matched	Transitions	State-vector
prop1	0	10.5 s	479	214292	660445	1058424	2088 byte
prop2	0	8.64 s	479	152057	373096	767167	2088 byte
prop3	2520	14.3 s	500	183077	924560	1346688	2184 byte

This property together with the code of the three hosts A, B and C, states that there is a state in the system where a packet with header 11 is received (either from host A or B) and eventually C receive a message with header 11 and no other such a message afterwards.

We use the Spin model checker to verify the three LTL properties, with the following parameters: Extra Compile-Time Directives is set to 20700; the number of hash-functions in Bitstate mode to 5, and the Physical Memory Available to 102400 Mbytes. We used depth-first search with partial order reduction and Bitstate/Supertrace in order to achieve better performance [7]. The results of the verification are shown in Table 1.

We simulate the results of prop3 with the first founded error written in the ".trail" file, the result of this simulation is shown in Fig. 8, where we see the packet forwarded by host B (here internally called prod2) to the switch (here called Protocol1). The flow table in the switch is updated successfully (see in action 25!11, 2) but soon the packet is dropped (see in action 25?11, 2). Before C (cons1) receive any other packets, the one sent from host A matches the new flow table (header = 11) and thus will be dropped (in action 1?11, 0) instead of being forwarded to C.

5. Related work

The first automata based model for Reo connectors appeared in [4] where constraint automata have been introduced. The authors define a product on constraint automata that implements the Reo composition operator on timed-data streams semantics [3]. Since then, several other operational models followed, such as Büchi automata of records [13], guarded




Fig. 8. Simulation result.

automata [14], and Reo automata [15]. Those models extend constraint automata by allowing some context-dependent reasonings. See [16] for an overview of the main operational and denotational models of Reo. Constraint automata with memory and their composition have been thoroughly studied in [5]. Our work is based on a similar model and semantics, but while we aim for finding a subset of it that can be easily implemented, the work in [5] concentrates on the efficient computation of the main composition operator.

Also model checking of Reo connectors has been a very active area of research. Vereofy is a dedicated model checker developed explicitly to verify linear time and branching time temporal properties of Reo connectors expressed as constraint automata [17]. Vereofy, however does not allow for explicit data to be handled in the automata and properties. For this reason, in [18] the authors encode Reo connectors as communicating processes, and use μ CRL2 to check some behavioral properties. The model checker μ CRL2 [19], is based on Algebra of communicating processes [20] with properties expressed as formulas in the modal μ -calculus with strong and branching bisimulation as equivalences as well as strong and weak trace equivalence. UPPAAL is a model checker [21], for linear time temporal property of networks of timed automata [22]. UPPAAL can perform reachability analysis, as well as simulation and error reports. See [23] for a recent use of UPPAAL to verify behavioral properties of real time Reo connectors. Our work differs from previous works on model checking Reo Connector in the following points. First of all, Spin is a data sensitive model checker. Contrary to Vereofy, we can verify temporal properties on connectors that involve data values and local memory. Second, our use of the Spin model checker is designed specifically to verify LTL properties of Reo connectors, contrary to the strong bisimulation equivalence used in proving equalities in the mCRL2 encoding of Reo channels in [18]. Thus is more in line with Vereofy and the semantic basis of Reo that is trace based, without branching properties. In addition, the encoding of Reo described in [18] does not encompass memory. Finally, our translation into Promela differs from the above works in the use of atomic statements in order to enforce synchrony. Of course extension of our work to real time systems and UPPAAL are imaginable, and can be pursued in the near future. All in all our work extends the literature by providing another tool chain to compile connectors to Promela and use the Spin model checker.

In [24], Ali Hussein et al. check network consistency properties in the model of SDN topology by using UPPAAL, the approach of them can detect an inconsistency or verify a flow in real-time. Elvira Albert et al. build an SDN model with barriers (which is Queue in our approach) in [25] by encoding it to language ABS based on top of the SYCO tool, then verify properties about safety policies and network loops. [26] uses Flow-LTL to specify the data flow in the extension of Petri nets with transits to check Petri Nets with transits by a hardware model checker ABC, and verifies concurrent updates and packets coherence in the network.

6. Conclusion

In this paper, we presented a full automatic translation from symbolic constraint automata to Promela. Symbolic constraint automata are a characterization of constraint automata with memory that can be used to compactly model almost all connectors of the coordination language Reo. In particular we restrict ourselves to an executable subset of Reo, assuming predicate and actions to be decidable. On the one hand, symbolic constraint automata cannot characterize relational constraint involving, for example, output ports in predicates. On the other hand, our symbolic constraint automata and their Promela translation easily allow for a generalization to lossy connectors by testing the non-presence of a trigger in an output ports before executing a lossy transition, as shown, for example, in the following Promela code for a lossy synchronous connector:

```

1 proctype LossySync(port A; port B){
2   mytype _a; mytype _b; //internal values in Sync
3   int state = 0; //initial state
4   do
5     :: state = 0 && Full(A.data) && Full(B.trig) ->
6       Atomic{take(A, _a); _b = _a; put(B, _b); state = 0};
7     :: state = 0 && Full(A.data) && Empty(B.trig) ->
8       Atomic{take(A, _a); state = 0};

```


9 `od }`

The translation to Promela is interesting in itself as we had to circumvent the problem of Promela not allowing for checking complex predicates on input ports in a guard of a statement. This is however a crucial feature for Reo connectors, that we have taken care of via delayed input and using control variable for recording transition already taken (and thus guaranteeing not only liveness, but also fairness in the transition selection process).

As a proof of concepts, we run our translation for the verification of an SDN model in Spin. The intuitive and modular Promela code is internally translated to a large transition system with more than 200.000 states and 100.000 transitions. It would be interesting to look for abstraction mechanism at the level of symbolic constraint automata to help reducing the state explosion. An obvious candidate is the combination of partial order reduction techniques at the symbolic level of the automaton itself.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Promela code of some Reo connectors

In this Appendix we give the Promela code generated by our translation for some basic Reo connectors. We present both an optimized code with no unnecessary local variables and the uniform code generated systematically.

• Synchronous Drain

Reo circuit: 

Symbolic constraint automaton: 


Optimized Promela code:


```
1 proctype Syncdrain(port A; port B){
2   mytype _a; mytype _b;
3   do
4     :: Full(A.data) && Full(B.data) -> Atomic{take(A, _a); take(B, _b)};
5   od }
```

Original Promela code:

```
1 proctype Syncdrain(port A; port B){
2   mytype _a; mytype _b; mytype x; int state=0; bool f1=1;
3   do
4     :: state==0 && Full(A.data) && Full(B.data) && f1==1 ->
5       Atomic{ True(A, B, sem);
6         if
7           :: sem==1 -> take(A, _a); take(B, _b); skip; state=0; f1=1;
8           :: else -> f1=0;
9         fi
10      od }
11   inline True(P,Q,x){
12     x = 1
13   }
```

• Non-deterministic Merger

Reo circuit: 

Symbolic constraint automaton: 

Optimized Promela code:

```
1 proctype Merger(port A; port B; port C){
2   mytype _a; mytype _b; mytype _c;
3   do
4     :: Full(A.data) && Full(C.trig) -> Atomic{take(A, _a); _c = _a; put(C, _c)};
```

```

5 :: Full(B.data) && Full(C.trig) -> Atomic{take(B, _b); _c = _b; put(C, _c)};
6 od }

```



Original Promela code:

```

1 proctype Merger(port A; port B; port C){
2   mytype _a; mytype _b; mytype _c; int state=0; bool f1=1; bool f2=1; int sem;
3   do
4     :: state==0 && Full(A.data) && Full(C.trig) && f1=1 -> Atomic{ True(A, sem);
5       if
6         :: sem==1 -> take(A, _a); _c = _a; put(C, _c); state=0; f1=1; f2=1;
7         :: else -> f1=0;
8       fi };
9     :: state==0 && Full(B.data) && Full(C.trig) && f2=1 -> Atomic{ True(B, sem);
10      if
11        :: sem==2 -> take(B, _b); _c = _b; put(C, _c); state=0; f1=1; f2=1;
12        :: else -> f2=0;
13      fi };
14    od }
15 inline True(P, x){
16   x = 1
17 }

```

• Replicator



Optimized Promela code:

```

1 proctype Replicator(port A; port B; port C){
2   mytype _a; mytype _b; mytype _c;
3   do
4     :: Full(A.data) && Full(B.trig) && Full(C.trig)
5       -> Atomic{take(A, _a); _b = _a; _c = _a; put(B, _b); put(C, _c)};
6   od }

```



Original Promela code:

```

1 proctype Replicator(port A; port B; port C){
2   mytype _a; mytype _b; mytype _c; int state=0; bool f1=1;
3   do
4     :: state==0 && Full(A.data) && Full(B.trig) && Full(C.trig) && f1==1
5       -> Atomic{ True(A, sem);
6         if
7           :: sem==1 -> take(A, _a); _b = _a; _c = _a; put(B, _b); put(C, _c);
8           state=0; f1=1;
9           :: else -> f1=0;
10          fi };
11    od }
12 inline True(P, x){
13   x = 1
14 }

```

• Transformer



Optimized Promela code:

```

1 proctype Transformer(port A; port B){
2   mytype _a; mytype _b;
3   do
4     :: Full(A.data) && Full(B.trig) ->

```

```

5   Atomic{take(A, _a); _b = f(_a); put(B, _b)};
6   od
7   inline f(x){
8   % user defined
9   }

```


Original Promela code:

```

1   proctype Transformer(port A; port B){
2   mytype _a; mytype _b; bool f1=1; int state=0;
3   do
4   :: state==0 && Full(A.data) && Full(B.trig) && f1==1 -> Atomic{True(A, sem);
5   if
6   :: sem==1 -> take(A, _a); _b=f(_a); put(B, _b); state=0; f1=1;
7   :: else -> f1=0;
8   fi};
9   od }
10  inline True(P, x){
11  x = 1
12  }
13  inline f(a){
14  % user defined
15  }

```

• PairMerger



$$C := \langle A, B \rangle$$

Symbolic constraint automaton:



Optimized Promela code:

```

1   proctype PairMerger(port A; port B; port C){
2   mytype _a; mytype _b; mytype _c; int state=0;
3   do
4   :: Full(A.data) && Full(B.data) && Full(C.trig) ->
5   Atomic{take(A, _a); take(B, _b); _c=<_a,_b>; put(C, _c)};
6   od }

```


Original Promela code:

```

1   proctype PairMerger(port A; port B; port C){
2   mytype _a; mytype _b; mytype _c; bool f1=1; int state=0;
3   do
4   :: state==0 && Full(A.data) && Full(B.data) && Full(C.trig) && f1==1 ->
5   Atomic{ True(A, B, sem);
6   if
7   :: sem==1 -> take(A, _a); take(B, _b); _c=<_a,_b>; put(C, _c);
8   state=0; f1=1;
9   :: else -> f1=0;
10  fi};
11  od }
12  inline True(P, Q, x){
13  x = 1
14  }


```

• Variable



$$B, \tau := \tau, A$$

Symbolic constraint automaton



Optimized Promela code:

```

1   proctype Variable(port A; port B){
2   mytype _a; mytype _b; mytype tau;
3   do
4   :: Full(A.data) && Full(B.trig) -> Atomic{ take(A, _a); _b = tau; tau = _a;

```

```

5   put(B, _b)
6   :: Full(A.data) -> Atomic{ take(A, _a);  $\tau$  = _a
7   :: Full(B.trig) ->   Atomic{ _b =  $\tau$ ; put(B, _b) }
8   od}

```

Original Promela code:

```

1  proctype Variable(port A; port B){
2  mytype _a; mytype _b; mytype  $\tau$ ; bool f1=1; bool f2=1; bool f3=1;
3  int sem; int state=0;
4  do
5  :: state==0 && Full(A.data) && Full(B.trig) && f1==1 ->
6     Atomic{True2(A,  $\tau$ , sem);
7     if
8     :: sem==1 -> take(A, _a); _b =  $\tau$ ;  $\tau$  = _a; put(B, _b);
9     state=0; f1=1; f2=1; f3=1;
10    :: else -> f1=0;
11    fi
12  :: state==0 && Full(A.data) && f2==1 -> Atomic{ True2(A,  $\tau$ , sem);
13  if
14  :: sem==1 -> take(A, _a);  $\tau$  = _a; state=0; f1=1; f2=1; f3=1;
15  :: else -> f2=0;
16  fi
17  :: state==0 && Full(B.trig) && f3==1 -> Atomic{ True1( $\tau$ , sem);
18  if
19  :: sem==1 -> _b =  $\tau$ ; put(B, _b); state=0; f1=1; f2=1; f3=1;
20  :: else -> f3=0;
21  fi
22  od}
23
24 inline True1(x, y){
25   y = 1
26 }
27 inline True2(P, x, y){
28   y = 1
29 }

```

References

- [1] Luís Soares Barbosa, Towards a calculus of state-based software components, *J. Univers. Comput. Sci.* 9 (8) (2003) 891–909.
- [2] Farhad Arbab, Puff, The Magic Protocol, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 169–206.
- [3] Farhad Arbab, Jan J.M.M. Rutten, A coinductive calculus of component connectors, in: Martin Wirsing, Dirk Pattinson, Rolf Hennicker (Eds.), Recent Trends in Algebraic Development Techniques, 16th International Workshop, WADT 2002, Frauenchiemsee, Germany, September 24–27, 2002, Revised Selected Papers, in: Lecture Notes in Computer Science, vol. 2755, Springer, 2002, pp. 34–55.
- [4] Christel Baier, Marjan Sirjani, Farhad Arbab, Jan Rutten, Modeling component connectors in Reo by constraint automata, *Sci. Comput. Program.* 61 (2) (2006) 75–113.
- [5] Sung-Shik T.Q. Jongmans, Tobias Kappé, Farhad Arbab, Constraint automata with memory cells and their composition, *Sci. Comput. Program.* 146 (2017) 50–86.
- [6] Benjamin Lion, Samir Chouali, Farhad Arbab, Compiling protocols to Promela and verifying their LTL properties, in: Proceedings of MODELS 2018 Workshops, 2018, pp. 31–39.
- [7] Gerard J. Holzmann, An analysis of bitstate hashing, *Form. Methods Syst. Des.* 13 (3) (1998) 289–307.
- [8] Klaus Havelund, Mike Lowry, John Penix, Formal analysis of a space-craft controller using spin, *IEEE Trans. Softw. Eng.* 27 (8) (2001) 749–765.
- [9] Gerald Lüttgen, Victor Carreno, Analyzing mode confusion via model checking, in: International SPIN Workshop on Model Checking of Software, Springer, 1999, pp. 120–135.
- [10] Gerard Holzmann, The Spin Model Checker: Primer and Reference Manual, first edition, Addison-Wesley Professional, 2003.
- [11] Hui Feng, Farhad Arbab, Marcello M. Bonsangue, A Reo model of software defined networks, in: Yamine Ait Ameur, Shengchao Qin (Eds.), Formal Methods and Software Engineering - 21st International Conference on Formal Engineering Methods, ICFEM 2019, Shenzhen, China, November 5–9, 2019, Proceedings, in: Lecture Notes in Computer Science, vol. 11852, Springer, 2019, pp. 69–85.
- [12] Link to the code and results, <https://github.com/githubifeng/Reo2Promela.git>.
- [13] Mohammad Izadi, Neil Birkbeck, Balu Adsumilli, Mutual noise estimation algorithm for video denoising, in: 2019 IEEE International Conference on Image Processing, ICIP 2019, Taipei, Taiwan, September 22–25, 2019, IEEE, 2019, pp. 2424–2428.
- [14] Marcello M. Bonsangue, Dave Clarke, Alexandra Silva, A model of context-dependent component connectors, *Sci. Comput. Program.* 77 (6) (2012) 685–706.
- [15] Alexandra Silva, A specification language for Reo connectors, in: Farhad Arbab, Marjan Sirjani (Eds.), Fundamentals of Software Engineering - 4th IPM International Conference, FSEN 2011, Tehran, Iran, April 20–22, 2011, Revised Selected Papers, in: Lecture Notes in Computer Science, vol. 7141, Springer, 2011, pp. 368–376.
- [16] Sung-Shik T.Q. Jongmans, Farhad Arbab, Overview of thirty semantic formalisms for Reo, *Sci. Ann. Comput. Sci.* 22 (1) (2012) 201–251.
- [17] Christel Baier, Tobias Blechmann, Joachim Klein, Sascha Klüppelholz, Wolfgang Leister, Design and verification of systems with exogenous coordination using Vereofy, in: Tiziana Margaria, Bernhard Steffen (Eds.), Leveraging Applications of Formal Methods, Verification, and Validation - 4th International Symposium on Leveraging Applications, ISoLA 2010, Heraklion, Crete, Greece, October 18–21, 2010, Proceedings, Part II, in: Lecture Notes in Computer Science, vol. 6416, Springer, 2010, pp. 97–111.
- [18] Natalia Kokash, Christian Krause, Erik P. de Vink, Reo + mcrl2: a framework for model-checking dataflow in service compositions, *Form. Asp. Comput.* 24 (2) (2012) 187–216.

- [19] Olav Bunte, Jan Friso Groote, Jeroen J.A. Keiren, Maurice Laveaux, Thomas Neele, Erik P. de Vink, Wieger Wesselink, Anton Wijs, Tim A.C. Willemse, The mcrl2 toolset for analysing concurrent systems - improvements in expressivity and usability, in: Tomás Vojnar, Lijun Zhang (Eds.), Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings, Part II, in: Lecture Notes in Computer Science, vol. 11428, Springer, 2019, pp. 21–39.
- [20] Jos C.M. Baeten, A brief history of process algebra, Theor. Comput. Sci. 335 (2–3) (2005) 131–146.
- [21] Johan Bengtsson, Kim Guldstrand Larsen, Fredrik Larsson, Paul Pettersson, Wang Yi, UPPAAL - a tool suite for automatic verification of real-time systems, in: Rajeev Alur, Thomas A. Henzinger, Eduardo D. Sontag (Eds.), Hybrid Systems III: Verification and Control, Proceedings of the DIMACS/SYCON Workshop on Verification and Control of Hybrid Systems, October 22–25, 1995, Rutgers University, New Brunswick, NJ, USA, in: Lecture Notes in Computer Science, vol. 1066, Springer, 1995, pp. 232–243.
- [22] Rajeev Alur, David L. Dill, The theory of timed automata, in: J.W. de Bakker, Cornelis Huizing, Willem P. de Roever, Grzegorz Rozenberg (Eds.), Real-Time: Theory in Practice, REX Workshop, Mook, the Netherlands, June 3–7, 1991, Proceedings, in: Lecture Notes in Computer Science, vol. 600, Springer, 1991, pp. 45–73.
- [23] Guillermina Cledou, José Proença, Bernhard H.C. Spath, Eric Verhulst, Hubs for VirtuosoNext: online verification of real-time coordinators, Sci. Comput. Program. 203 (2021) 102566.
- [24] Ali Hussein, Imad H. Elhajj, Ali Chehab, Ayman Kayssi, Sdn verification plane for consistency establishment, in: 2016 IEEE Symposium on Computers and Communication (ISCC), IEEE, 2016, pp. 519–524.
- [25] Elvira Albert, Miguel Gomez-Zamalloa, Miguel Isabel, Albert Rubio, Matteo Sammartino, Alexandra Silva, Actor-based model checking for software-defined networks, J. Log. Algebraic Methods Program. 118 (2021) 100617.
- [26] Bernd Finkbeiner, Manuel Giesekeing, Jesko Hecking-Harbusch, Ernst-Rüdiger Olderog, Model checking data flows in concurrent network updates (full version), arXiv:1907.11061 [abs], 2019.