

A Constraint-based Method to Compute Semantics of Channel-based Coordination Models

Behnaz Changizi, Natallia Kokash, and Farhad Arbab

Leiden Institute of Advanced Computer Science (LIACS)
P.O. Box 9512, 2300 RA Leiden, The Netherlands
CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
`b.changizi@umail.leidenuniv.nl`
`{natallia.kokash, farhad.arbab}@cwi.nl`
<http://www.liacs.nl/>

Abstract. Reo is an exogenous channel-based coordination language that acts as glue code to tie together software components and services. The building blocks of Reo models are connectors that impose constraints on the data-flow in component or service-based architectures in terms of data synchronization, buffering, mutual exclusion, etc. Several semantic models have been introduced to formalize the behavior of Reo. These models differ in terms of expressiveness, computation complexity and purposes that they serve. In this paper, we present a method and a tool for building formal automata-based semantics of Reo that unifies various aspects of existing semantics. We express the behavior of a Reo network as a mixed system of Boolean and numerical constraints constructed compositionally by conjuncting the assertions for its constituent parts. The solutions of this system are found with the help of off-the-shelf constraint solvers and are used to construct the constraint automaton with state memory that gives the sound and complete semantics of Reo with respect to existing models. Our approach is more efficient and scalable compared to the existing methods for generating formal semantics of Reo connectors.

1 Introduction

Service-oriented architecture [9] (SOA) is an indispensable solution for many of today's problems. The SOA implementation depends on a mesh of functionality units, called services. Services are loosely coupled and do not invoke or communicate with each other directly. Instead, they employ a pre-defined protocol, which specifies the way they can exchange messages amongst themselves. As a result, the correctness of a SOA implementation relies not only on the correctness of its involved services but also on the properness of its communication protocol.

Exogenous coordination languages provide dedicated frameworks to study the communication protocols as separate concerns. Such languages define the “glue code” that ties together the services to enable the message passing among the involved services. Reo [1] is one such coordination language. It realizes the coordination patterns in terms of its complex *connectors* also called *networks* that are built out of simple primitives called *channels*.

Each channel defines a form of coordination in terms of synchronizing, buffering, retaining data, etc., along with constraining its input and output data items. Reo allows hierarchical modeling where arbitrarily complex interaction connectors can be formed out of simpler networks. In our previous work [5] [12], we have presented the suitability of Reo to model behavioral patterns describable by business process models. We have also developed tools for automatic transformation of these models into Reo [12]. This enables the use of Reo analysis methods and tools on the coordination protocols that originally were not expressed in Reo.

To perform formal analysis on Reo networks, formal semantics of these models are necessary. Several operational semantics have been proposed for Reo [16] with various styles of I/O streams [1], automata, coloring [13] and constraints [19]. The most basic automata-based semantics of Reo is Constraint Automata (CA) [8]. An advantage of CA and its extensions is their support for data-constraints that are part of the coordination primitives in Reo. This is in contrast with the coloring semantics that abstracts data-flow and expresses the behavior of a connector only in terms of existence or lack of data-flow.

Constraint Automata with State Memory (CASM) [7] is an extension of CA that due to its state abstraction and data-awareness is suitable as a more compact semantic model for Reo in model checking. In this paper, we present a constraint-based technique and a tool to generate CASMs from Reo networks in a compositional manner. A shortcoming of this earlier work stems from its lack of support for data-dependent behavior. We overcome this shortcoming in the work we present in this paper.

The rest of this paper is organized as follows. In Section 2, we explain the basics of Reo. In Section 3, we introduce constraint automata with state memory. In Section 4, we formalize the mentioned semantic model of Reo in systems of constraints along with techniques to solve them. In Section 5, we show how we abstract from the internal ports in a Reo network in a minimized representation. In Section 6, we briefly discuss some properties of our presented constraint en-

coding of Reo networks. Finally, in Section 7, we conclude the paper and outline our future work.

2 Reo

Reo [1] is an exogenous channel-based coordination language, which can act as “glue code” to tie together software components and services. The building blocks of Reo models, *connectors*, impose constraints on the data-flow in terms of synchronization, buffering, mutual exclusion, etc. Every connector contains some *primitives*. The set of Reo primitives is open-ended, meaning that a user can define new primitives and extend the expressiveness of Reo.

The simplest Reo connector is a *channel* that has two *ends*, also called *ports*. Channel ends are either of type *source* that reads data into the channel or *sink* that writes the channel’s data out. Table 1 shows the most commonly used channels in Reo.

Table 1: Graphical representation of the more frequently used Reo channels

	A <i>Sync</i> channel accepts data from its source end iff it can dispense it simultaneously through its sink end.
	A <i>LossySync</i> reads a data item from its source end and writes it simultaneously to its sink end. If the sink end is not ready to accept the data item, the channel loses it.
	A <i>SyncDrain</i> reads data to discard through its two source ends iff both ends are ready to interact simultaneously.
	An <i>AsyncDrain</i> accepts and discards a data item from either of its source ends that offers one. If both ends offer data items simultaneously, then the channel chooses one non-deterministically.
	A <i>Filter</i> accepts a data item that does not matched its predefined filter pattern P from its source and loses it. For a data item that matches its filter pattern P , a filter channel behaves as a <i>Sync</i> channel; it accepts the data item iff it can dispense it simultaneously through its sink end.
	A <i>Transformer</i> acquires data at its source end, applies a predefined transform function f on it and simultaneously writes the result to its sink end, iff the data item is in the domain of the function f . Otherwise, the channel loses the data item.
	If a $FIFO_1$ is empty, it accepts incoming data from its source end and buffers it. Being full, the channel is ready to dispense data through its sink end and becomes empty. Because this channel has a buffer capacity of one data item, it is either full or empty at any given time, and thus the ends of this channel cannot interact simultaneously.

Table 2: Graphical representation of Reo nodes and two frequently used components

	A <i>Replicator</i> replicates the incoming data of its source to its sink ends simultaneously.
	A <i>Merger</i> non-deterministically chooses one of its source ends that is ready to communicate, take its incoming data item and writes it to its sink end.
	A <i>Router</i> accepts data from its source end and simultaneously writes it on one of its non-deterministically chosen sink ends that is ready to accept the data.
	A <i>Cross-product</i> reads one incoming data item from each of its incoming source ends, forms a tuple in which the data elements are set in the counter-clock-wise order with respect to its sink node, and simultaneously writes the resulting tuple on its sink end.

Reo nodes connect channels to each other to form Reo connectors, also called *circuits* or *networks*. Depending on whether all channel ends that coincide on a node are source ends, sink ends or a combination of both, nodes become *source*, *sink* or *mixed* nodes. A source node behaves as a synchronous *Replicator* that replicates the incoming data of its source to its sink ends simultaneously, while a sink node acts as a non-deterministic *Merger* that combines the flows of its source ends to its sink end. A mixed node is an atomic combination of a replicator and a non-deterministic merger. Each read and write action needs all of its involved source and sink ends to be able to interact synchronously. Otherwise, the action cannot take place. In addition, Reo allows hierarchical modeling and abstraction from inner structures by means of components. A connector can be converted to a component, exhibiting (part of) its inner logic as an observable behavioral interface. Table 2 shows the graphical representation of the Reo nodes and two of frequently used components, *Router* and *Cross-product*.

2.1 Formal Semantics of Reo

Coordination patterns that a Reo network imposes on data flow define the network's *behavior*. Formal analysis of a Reo network requires formal modeling of its behavior. The behavior of a Reo network consists of various dimensions that involve:

State: the states of elements forming a Reo network at any point of time define the state, also called *configuration*, of the network. The network state affects and is affected by the data flow at each step.

Synchronization/Exclusion: the term *synchronization* refers to atomic concurrent data flow over ports and nodes. Depending on its constituents and their arrangement inside the network, in each configuration, a Reo network

allows, requires, or forbids a group of ports to synchronize. We consider *mutual exclusion* as a special case of synchronization.

Data flow: the value of data items that the ports of a Reo network exchange can affect the network behavior, particularly if the network contains elements, such as filter or transformer channels that allow or forbid exchange of special data values.

Context dependency: the choices some Reo elements can take change non-monotonically as the context changes [10]. Context dependent semantics of Reo define context in terms of presence or absence of pending I/O requests on primitive ends. As an example of such behavior consider the original description of a *LossySync* channel, which writes on its sink end the data item it reads from its source end. The channel can lose a data item only if its sink end is not write-enabled [1].

Timing: the behavior of a Reo network can be subject to time constraints. For example, we can formalize a deadline for availability of some data using a *FIFO₁* channel, which associates an expiration time with data it buffers. If the channel’s sink end does not dispense the buffered data item before its expiration, the channel loses it [2].

Priority: presence of a prioritized element in a Reo network can influence some non-deterministic synchronization choices in the network by favoring data flow through ports related to that prioritized element.

Several formal semantic models have been proposed to specify the behavior of Reo. An extensive overview and comparison of these models is presented in [16].

Connector coloring (CC) [13] is a formal semantics for Reo that describes the behavior of a connector by assigning different colors to its ports to designate presence or absence of data flow. CC accounts for synchronization and context dependency. This model captures context dependency by propagating negative information about the absence of data flow inside the network.

The majority of Reo semantic models are based on automata. These semantic models allow the semantics of a large network to be constructed from the given automata for its constituents using the *product* of automata. The earliest automata-based Reo semantics is Constraint Automata (CA) [8] that accommodate synchronization, states, and data flow. Transitions in CA contain the set of synchronized ports and constraints over data that the ports exchange. An extension of CA, Constraint Automata with State Memory (CASM) [7] elaborates on states by introducing state memory cells and extends data constraints to accommodate them.

Extending CA with clock assignments and timing constraints, Timed Constraint Automata (TCA) [2] express the time-aware aspect of Reo networks. A SAT-based approach for bounded model checking of TCA is presented in [17]. Another extension of CA, Constraint Automata with Priority (CAP) [3] supports the propagation of prioritized requests.

The most recent semantics of Reo [19] that is based on constraint solving, deals with synchronization, data, and context dependency. It uses different constraints for each of these notions, which are added conjunctively to capture their

composition. This approach is the basis of the *DREAMS* [19] execution engine for Reo. Although in theory, this semantics accounts for data, the proposed implementation ignores data. This leads to incorrect results when dealing with data-sensitive elements such as *Filter* channel.

A translation of CA and CC to a process algebraic specification language called mCRL2 [15] has been done [18]. The mCRL2 toolset compiles the specifications into Labeled Transition Systems (LTS)s, which can be verified using the mCRL2 toolset. This approach represents certain aspects of existing semantic models for Reo in an automata-based form. For instance, it models the context dependency in the behavior of Reo networks using LTSs with labels corresponding to colors in CC.

Despite the abundance of the semantics that capture different aspects of Reo, there is an expressiveness gap between them. For instance, a network with priority dependent and time sensitive behavior cannot be expressed by a single formal semantics. In this paper, we present a unified symbolic constraint-based framework, where these different semantics coexist. This framework encodes the behavior of Reo networks in terms of constraints whose solutions form an existing automata-based semantic model of Reo. Unlike [18], our framework treats data symbolically, so that the state space does not blow up. Our framework extends *DREAMS* with time and priority. It also provides tool support using existing constraint solving programs to incorporate data and time constraints.

For a given configuration of a Reo network, *DREAMS* uses constraint solving to compute only a single solution for synchronous data flow in this configuration, which leads to the next configuration. In contrast, our approach considers all solutions and configurations which in turn enable us to map the solutions to an automata-based semantics that is well suited for model checking.

Generating the CA corresponding to a Reo network from the CA of its constituents using the CA *product* operator [8], is computationally expensive and memory inefficient. The complexity of composing m CA each having n transactions is $\mathcal{O}(n^m)$. However, we can convert this problem to the *NP*-complete problem of SAT-solving for which many efficient solvers exist.

Contribution: Instead of generating the CA corresponding to a Reo network directly by composing the CAs of its constituents, we encode each constituent CA into constraints whose conjunction forms the system of constraints that captures the full behavior of the network. The solutions of this system of constraints describe the behavior of the network in terms of variables representing different aspects of the behavior of Reo, such as synchronization, data flow, etc. From these solutions, we generate the CA corresponding to the network. We show that our solution is more efficient and more scalable compared to the previous attempts in generating the formal semantics of a Reo network.

We have developed a tool to automate our approach that is integrated in the Extensible Coordination Tools (ECT) [4]. ECT is a framework to design, develop, model check, test, and execute component-based software modeled by the Reo coordination language. The tools in the framework are integrated as Eclipse plug-ins and operate based on the operational semantics of Reo, most

notably, connector coloring and constraint automata. Our tool is the only tool that supports propagation of priorities in Reo networks. Furthermore, it unifies various behavioral aspects of Reo networks in a single encoding, which enables analysis of networks whose behavior spans over several semantic models.

3 Constraint Automata with State Memory

Constraint Automata with State Memory (CASM) [7] extends CA with variables that represent local memory cells of the states of the automata. Due to its elaboration on state information, we choose CASM as the main semantic model of Reo that our framework generates.

Definition 1 (Constraint Automaton with State Memory (CASM)). *A constraint automaton with state memory is a tuple $A = (Q, \mathcal{N}, \rightarrow, q_0, \mathcal{M})$ where*

- Q is a finite set of states.
- \mathcal{N} is a finite set of names.
- \rightarrow is a finite subset of $Q \times 2^{\mathcal{N}} \times DC(\mathcal{N}, \mathcal{M}, Data) \times Q$ is the transition relation of A , where $DC(\mathcal{N}, \mathcal{M}, Data)$ is the set of data constraints, defined below.
- $q_0 \in Q$ is an initial state.
- \mathcal{M} is a set of memory cell names, where $\mathcal{N} \cap \mathcal{M} = \emptyset$.

Every $n \in \mathcal{N}$ represents a node in a Reo connector. The set \mathcal{N} is partitioned into three mutually disjoint sets of source nodes \mathcal{N}^{src} , mixed nodes \mathcal{N}^{mix} , and sink nodes \mathcal{N}^{snk} . Because we make the replication and merge inherent in Reo nodes explicit as *replicator* and *merger* primitives (in Table 2), at most two primitive ends coincide on every node $n \in \mathcal{N}$. Thus, it follows that a source or a sink node contains only a single (source or sink) primitive end, and a mixed node contains exactly one source and one sink primitive ends.

We write $q \xrightarrow{N,g} p$ instead of $(q, N, g, p) \in \rightarrow$. For every transition $q \xrightarrow{N,g} p$, we require that $g \in DC(\mathcal{N}, \mathcal{M}, Data)$, where $DC(\mathcal{N}, \mathcal{M}, Data)$ is the language defined by the following grammar:

$$g ::= true \mid \neg g \mid g \wedge g \mid u = u \mid u < u ,$$

$$u ::= d(n) \mid m' \mid m \mid v$$

In this grammar, $=$ is the symmetric equality relation, $<$ is a full order relation, $n \in \mathcal{N} \subseteq \mathcal{N}$ denotes a port name, $d(n)$ represents the data item exchanged through port n , $m \in \mathcal{M}$ correspond to a memory cell in the current state, which is the source state of the transition, m' stands for the memory cell $m \in \mathcal{M}$ in the next state, which is the target state of the transition, and $v \in Data$. As usual, *false* stands for $\neg true$, $x > y$ stands for $y < x$, and other logical operators, such as \vee and \Rightarrow (the implication symbol) can be built from the given operators.

Transitions with data constraints that can be reduced to *false* using the Boolean laws are impossible and we omit them. A data constraint g that is

always *true* can be left out. We use \mathcal{M}_g to represent the set of all $m \in \mathcal{M}$ that syntactically appear as m in a data constraint g ; and \mathcal{M}'_g to refer to the set of all $m \in \mathcal{M}$ that syntactically appear as m' in g . The valuation function $\mathcal{V}_q : \mathcal{M} \rightarrow 2^{Data}$ designates the set of values $\mathcal{V}_q(m)$ of a memory cell $m \in \mathcal{M}$ in a state $q \in Q$, where $\mathcal{V}_{q_0}(m) = \emptyset$ for all $m \in \mathcal{M}$.

A transition $q \xrightarrow{N,g} p$ in a given constraint automaton with state memory is possible only if there exists a substitution for every syntactic element $d(n)$, m , and m' that appears in g to satisfy g . A substitution simultaneously replaces in g :

- every occurrence of $d(n)$ with the data value exchanged through the node $n \in \mathcal{N}$;
- every occurrence of m' of every $m \in \mathcal{M}$ with a value $v \in Data$;
- every occurrence $m \in \mathcal{M}$ with:
 - the special symbol ' \circ ' if $\mathcal{V}_q(m) = \emptyset$
 - a value $v \in \mathcal{V}_q(m)$, otherwise.

The guard g is satisfied if proper replacement values can be found to make g *true*. Making this transition, the automaton defines the valuation function \mathcal{V}_p for the target state p , as follows: for every $m \in \mathcal{M}'_g$, $\mathcal{V}_p(m)$ is the set of all $v \in Data$ whose replacements for m' satisfy g . For every other $m \in \mathcal{M}$, $\mathcal{V}_p(m) = \emptyset$.

A relational operator evaluates to *true* only if the values of its operands are in its respective relation. Thus, any operator with one or more \circ as an operand always evaluates to *false*. We call a CASM, *normalized* iff a) it does not have two states with the same set of state memory variables, and b) every two transitions differ at least in their start states, their target states, or their sets of synchronizing ports. For any arbitrary CASM that is not normalized, we can normalize it by a) introducing auxiliary variables, to make the set of state memory variables unique for each state, and b) by merging the transitions that have the same start and target states and synchronize the same ports. In the sequel, we consider only *normalized* CASMs.

Following are the definitions for *product* and *hiding* operations on CASM. Both definitions are adapted from [8].

Definition 2 (Product-automaton). For the CASMs $\mathcal{A}_1 = (Q_1, \mathcal{N}_1, \rightarrow_1, q_{0,1}, \mathcal{M}_1)$ and $\mathcal{A}_2 = (Q_2, \mathcal{N}_2, \rightarrow_2, q_{0,2}, \mathcal{M}_2)$, their product is defined as:

$$\mathcal{A}_1 \bowtie \mathcal{A}_2 = (Q_1 \times Q_2, \mathcal{N}_1 \cup \mathcal{N}_2, \rightarrow, q_{0,1} \times q_{0,2}, \mathcal{M}_1 \cup \mathcal{M}_2)$$

where the following rules define the transition relation \rightarrow :

$$\frac{q_1 \xrightarrow{N_1, g_1} p_1, q_2 \xrightarrow{N_2, g_2} p_2, N_1 \cap \mathcal{N}_2 = N_2 \cap \mathcal{N}_1}{\langle q_1, q_2 \rangle \xrightarrow{N_1 \cup N_2, g_1 \wedge g_2} \langle p_1, p_2 \rangle}$$

$$\frac{q_1 \xrightarrow{N_1, g_1} p_1, N_1 \cap \mathcal{N}_2 = \emptyset}{\langle q_1, q_2 \rangle \xrightarrow{N, g} \langle p_1, q_2 \rangle} \quad \frac{q_2 \xrightarrow{N_2, g_2} p_2, \mathcal{N}_1 \cap N_2 = \emptyset}{\langle q_1, q_2 \rangle \xrightarrow{N, g} \langle q_1, p_2 \rangle}$$

We can abstract from the data flow on certain Reo ports using the *hiding* operator defined as follows:

Definition 3 (Hiding). Let $\mathcal{A} = (Q, \mathcal{N}, \rightarrow, q_0, \mathcal{M})$ be a constraint automaton and $C \in \mathcal{N}$. The constraint automaton that results from hiding the node C in automaton \mathcal{A} is $\exists C[\mathcal{A}] = (Q, \mathcal{N} \setminus \{C\}, \rightarrow_C, q_0, \mathcal{M})$ and the transition relation \rightarrow_C is defined as follows:

$$p \xrightarrow{N, g} r, N' = N \setminus \{C\}, g' = \exists C[g], \text{ where } \exists C[g] = \bigvee_{d \in \text{Data}} g[\text{data}(C)/d].$$

$$p \xrightarrow{N', g'}_C r$$

4 Reo Constraint Satisfaction Problem

We extend the constraint-based framework in [19] to incorporate all behavioral aspects of a Reo network given by its various existing semantic models.

Let $\mathcal{N} = \mathcal{N}^{src} \cup \mathcal{N}^{mix} \cup \mathcal{N}^{snk}$ be the global set of nodes, \mathcal{M} the global set of memory state variables, and \mathcal{D} the global set of numerical data values. The set of primitive ends \mathcal{P} consists of all primitives p derived from \mathcal{N} by marking its elements with subscripts c and k , according to the following grammar:

$$p ::= r_c \mid s_k$$

where $r \in \mathcal{N}^{src} \cup \mathcal{N}^{mix}$ and $s \in \mathcal{N}^{snk} \cup \mathcal{N}^{mix}$. Observe that n_c and n_k belong to two different primitives that connect on the common node n .

Let $p \in \mathcal{P}$, $n \in \mathcal{N}$ and $m \in \mathcal{M}$ be a primitive end, a node, and a memory state variable, respectively. A free variable v that occurs in the constraints encoding the behavior of a Reo network has one of the following forms:

- \tilde{n} ranges over $\{\top, \perp\}$ to show presence or absence of flow on the node n .
- \hat{n} ranges over \mathcal{D} to represent the data value passing through the node n .
- \hat{m}, \hat{m}' range over $\{\top, \perp\}$ to denote whether or not the state memory variable m is defined in, respectively, the source and the target states of the transition to which the encoded guard belongs.
- \check{m}, \check{m}' range over \mathcal{D} to represent the values of the state memory variable m in, respectively, the source and the target states of the transition to which the encoded guard belongs.
- \vec{p} ranges over $\{\top, \perp\}$ to state that the reason for lack of data flow through the primitive end p originates from, respectively, the primitive to which p belongs or the context.

However, not all of the introduced variables are required for encoding the behavior of every Reo network. In presence of context dependent primitives like *LossySync* or in priority-sensitive networks, constraints include variables of the form \vec{p} . For the stateful elements such as *FIFO*₁, variables like $\hat{m}, \hat{m}', \check{m}$, and \check{m}' appear in the constraints.

Observe that the interpretation of some of the mentioned variables depends on the value of other variables. Referring to the variable \vec{p} makes sense only if

$\tilde{n} = \perp$, where the primitive end p belongs to the node n ; and \hat{n} , \check{m} and \check{m}' make sense only if $\tilde{n} = \top$, $\hat{m} = \top$ and $\hat{m}' = \top$, respectively.

The grammar for a constraint Ψ encoding the behavior of a Reo network is as follows:

$$\begin{aligned}
t &::= \hat{n} \mid \check{m} \mid \check{m}' \mid d \mid t \otimes d && \text{(terms)} \\
a &::= \tilde{n} \mid \vec{p} \mid \hat{n} \mid \hat{m}' \mid t = t \mid t < t && \text{(atoms)} \\
\Psi &::= a \mid \neg\Psi \mid \Psi \wedge \Psi && \text{(formulae)}
\end{aligned}$$

where $d \in \mathcal{D}$ is a constant, $\otimes \in \{+, -, *, /, \%, \wedge\}$,
and p is either of the form n_c or n_k .

Definition 4 (Reo Constraint Satisfaction Problem). *A Reo Constraint Satisfaction Problem (RCSP) is a tuple $\langle \mathcal{P}, \mathcal{M}, M_0, X, C \rangle$, where:*

- \mathcal{P} is a finite set of primitive ends.
- \mathcal{M} is a finite set of state memory variables.
- $M_0 \subseteq \mathcal{M}$ is a set of state memory variables that define the initial configuration of a Reo network.
- $X = \{\tilde{n}, \vec{p}, \hat{n}, \hat{m}', \hat{n}, \check{m}, \check{m}'\} : n \in \mathcal{N}, p \in \mathcal{P}, m \in \mathcal{M}$, where the values that the variables of the forms \hat{n}, \check{m} , and \check{m}' can assume are subsets of \mathcal{D} , and the other variables are Boolean, with values in $\{\top, \perp\}$.
- $C = \{C_1, C_2, \dots, C_m\}$ is a finite set of constraints, where each C_i is a constraint given by the grammar Ψ involving a subset of variables $V_i \subseteq X$.

Definition 5 (Composition). *We obtain the constraints corresponding to a Reo network compositionally by conjuncting the RCSPs of its constituents. We use the symbol \odot to denote the composition of RCSPs.*

However, connecting two Reo networks must not introduce incorrect data flow possibilities. This is done by enforcing a restriction on the possible solutions through the following axiom:

Axiom 1 (Mixed node axiom) *When two Reo networks connect on a common node x , where x_c is a source end in one network and x_k is a sink end in the other, the following constraint must hold:*

$$\neg x \leftrightarrow (\vec{x}_c \vee \vec{x}_k)$$

The *mixed node axiom*, which applies to all mixed nodes in a network, states that a node x cannot produce the reason for no flow all by itself.

4.1 Encoding Reo Elements in RCSPs

Table 3 summarizes the constraint encodings associated with commonly used Reo elements. If a Reo network does not contain any context dependent channel, the variables encoding the context dependency can be ignored in its RCSP. Table 4 shows the encoding of Reo elements from Table 3 where the context variables

are removed. Note that in the table, c and k denote the source and the sink ends of a primitive, respectively, and that dom refers to the data domain of the given function or predicate. In the case of elements with more than one source or sink ends, we use indices.

Table 3: Context dependent encoding of Reo primitives (extending [19])

Channel	Constraints
	$\psi_{Sync}(c, k) : \tilde{c} \leftrightarrow \tilde{k} \wedge \tilde{c} \rightarrow (\hat{c} = \hat{k}) \wedge \neg(\vec{c}_c \wedge \vec{k}_k)$
	$\psi_{SyncDrain}(c_1, c_2) : \tilde{c}_1 \leftrightarrow \tilde{c}_2 \wedge \neg(\vec{c}_{1c} \wedge \vec{c}_{2c})$
	$\psi_{AsyncDrain}(c_1, c_2) : \tilde{c}_1 \rightarrow (\neg\tilde{c}_2 \wedge \vec{c}_{2c}) \wedge \tilde{c}_2 \rightarrow (\neg\tilde{c}_1 \wedge \vec{c}_{1c})$
	$\psi_{LossySync}(c, k) : \tilde{k} \rightarrow \tilde{c} \wedge \tilde{k} \rightarrow (\hat{c} = \hat{k}) \wedge \neg\vec{c}_c \wedge \neg\tilde{c} \rightarrow \vec{k}_k$
	$\psi_{Merger}(c_{0..i}, k) : \tilde{c}_i \leftrightarrow \tilde{k} \wedge \tilde{c}_i \rightarrow (\hat{c}_i = \hat{k}) \wedge \neg\tilde{k} \rightarrow ((\neg\vec{k}_k \wedge \vec{c}_{ic}) \vee (\vec{k}_k \wedge \neg\vec{c}_{ic} \wedge \bigwedge_{j,j \neq i} \vec{c}_{jk}))$
	$\psi_{Replicator}(c, k_{0..i}) : \tilde{c} \leftrightarrow \bigwedge_i \tilde{k}_i \wedge (\tilde{c} \rightarrow \bigwedge_i (\hat{k}_i = \hat{c})) \wedge \neg\tilde{c} \rightarrow ((\neg\vec{c}_c \wedge \vec{k}_{ik}) \vee (\neg\vec{k}_{ik} \wedge \bigwedge_{j,j \neq i} \vec{k}_{jk} \wedge \vec{c}_c))$
	$\psi_{Router}(c, k_{0..i}) : \tilde{c} \leftrightarrow (\bigvee_i \tilde{k}_i) \wedge \bigwedge_{j,j \neq i} \neg(\tilde{k}_i \wedge \tilde{k}_j) \wedge \tilde{k}_i \rightarrow (\hat{k}_i = \hat{c}) \wedge \tilde{c} \leftrightarrow (\neg\vec{c}_c \vee \neg(\bigvee_i \vec{k}_{ik}))$
	$\psi_{FIFO_1}(c, k, m) : \tilde{c} \rightarrow (\neg\hat{m} \wedge \hat{m}' \wedge (\hat{m}' = \hat{c})) \wedge \tilde{k} \rightarrow (\hat{m} \wedge \neg\hat{m}' \wedge (\hat{m} = \hat{k})) \wedge (\neg\tilde{c} \wedge \neg\tilde{k}) \rightarrow (\hat{m} \leftrightarrow \hat{m}' \wedge \hat{m} \rightarrow (\hat{m} = \hat{m}')) \wedge \neg\hat{m} \rightarrow \vec{k}_k \wedge \hat{m} \rightarrow \vec{c}_c$
	$\psi_{Filter}(c, k, \mathcal{P}) : \tilde{k} \rightarrow (\tilde{c} \wedge \hat{c} \in dom(\mathcal{P}) \wedge \mathcal{P}(\hat{c})) \wedge \tilde{k} \rightarrow (\hat{c} = \hat{k}) \wedge (\neg\tilde{c} \rightarrow (\neg\vec{c}_c \leftrightarrow \vec{k}_k)) \wedge (\tilde{c} \wedge \neg\tilde{k} \rightarrow \vec{k}_k)$
	$\psi_{Transformer}(c, k, \mathcal{F}) : \tilde{k} \rightarrow (\tilde{c} \wedge \hat{c} \in dom(\mathcal{F})) \wedge \tilde{k} \rightarrow (\hat{k} = \mathcal{F}(\hat{c})) \wedge \neg(\vec{c}_c \wedge \vec{k}_k)$

Example 1. Figure 1 shows a Reo network that consists of a *transformer* channels with the function $3 * \hat{a}$, whose data domains is the set of numbers *Number* and a *filter* channel with the condition $\hat{b} \% 2 = 0$ and the domain *Number*.

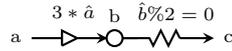


Figure 1: A data-aware Reo network

Since none of the Reo primitives in Figure 1 are context dependent, we use the constraints corresponding to this network as defined in Table 4.

Equation 1 states that flow occurs on the source end of the *transformer* channel iff it occurs on its sink end. In addition, flow can exist only if the data item that enters the source end of the channel is a number. In this case, the data item written on the sink end is treble the value of the source data item.

Table 4: Encoding Reo primitives (extending [19])

Channel	Constraints
	$\psi_{Sync}(c, k) : \tilde{c} \leftrightarrow \tilde{k} \wedge \tilde{c} \rightarrow (\hat{c} = \hat{k})$
	$\psi_{SyncDrain}(c_1, c_2) : \tilde{c}_1 \leftrightarrow \tilde{c}_2$
	$\psi_{AsyncDrain}(c_1, c_2) : \neg(\tilde{c}_1 \wedge \tilde{c}_2)$
	$\psi_{LossySync} : \tilde{k} \rightarrow \tilde{c} \wedge \tilde{k} \rightarrow (\hat{c} = \hat{k})$
	$\psi_{Merger}(c_{0..i}, k) : \tilde{k} \leftrightarrow (\bigvee_i \tilde{c}_i) \wedge_{j,j \neq i} \neg(\tilde{c}_i \wedge \tilde{c}_j) \wedge \tilde{c}_i \rightarrow (\hat{c}_i = \hat{k})$
	$\psi_{Replicator}(c, k_{0..i}) : \tilde{c} \leftrightarrow (\bigwedge_i \tilde{k}_i) \wedge \tilde{c} \rightarrow (\bigwedge_i (\hat{k}_i = \hat{c}))$
	$\psi_{Router}(c, k_{0..i}) : \tilde{c} \leftrightarrow (\bigvee_i \tilde{k}_i) \wedge_{j,j \neq i} \neg(\tilde{k}_i \wedge \tilde{k}_j) \wedge \tilde{k}_i \rightarrow (\hat{k}_i = \hat{c})$
	$\psi_{FIFO_1}(c, k, m) : \tilde{c} \rightarrow (\neg \hat{m} \wedge \hat{m}' \wedge (\hat{m}' = \tilde{c})) \wedge \tilde{k} \rightarrow (\hat{m} \wedge \neg \hat{m}' \wedge (\hat{m} = \hat{k})) \wedge (\neg \tilde{c} \wedge \neg \tilde{k}) \rightarrow (\hat{m} \leftrightarrow \hat{m}' \wedge \hat{m} \rightarrow (\hat{m} = \hat{m}'))$
	$\psi_{Filter}(c, k, \mathcal{P}) = \tilde{k} \rightarrow (\tilde{c} \wedge \hat{c} \in \text{dom}(\mathcal{P}) \wedge \mathcal{P}(\hat{c}) \wedge (\hat{c} = \hat{k}))$
	$\psi_{Transformer}(c, k, \mathcal{F}) = \tilde{k} \rightarrow (\tilde{c} \wedge \hat{c} \in \text{dom}(\mathcal{F})) \wedge \tilde{k} \rightarrow (\hat{k} = \mathcal{F}(\hat{c}))$

$$\psi_{Transformer}(a, b, 3 * \hat{a}) = \tilde{a} \leftrightarrow \tilde{b} \wedge \tilde{a} \rightarrow (\hat{a} \in \text{Number} \wedge \hat{b} = 3 * \hat{a}) \quad (1)$$

$$\psi_{Filter}(b, c, \hat{b}\%2 = 0) = \tilde{c} \rightarrow (\tilde{b} \wedge \hat{b} \in \text{Number} \wedge (\hat{b}\%2 = 0)) \quad (2)$$

Equation 2 expresses that flow on the source end of the *filter* channel leads to flow on its sink end, iff the data item belongs to the channel's accepting pattern (which is $\hat{b}\%2 = 0$). In this case, the value of data items passing through the ends are equivalent. No flow through the sink end c is either due to no flow on b or that the incoming data item does not satisfy the accepting pattern.

4.2 Solving RCSPs

In this section, we formalize the solutions of RCSPs and show how to obtain them.

Definition 6 (Solution). *A solution \mathcal{S} for a constraint C is function $\mathcal{S} : X \rightarrow 2^{\mathcal{D}}$ such that for all distinct $x_i \in X, 1 \leq i \leq n = |X|$, we have $z_i \in \mathcal{S}(x_i)$ implies $C[x_1, x_2, \dots, x_n \setminus z_1, z_2, \dots, z_n]$ is true.*

Since Reo Constraint Satisfaction Problems (RCSPs) have predicates with free variables of types Boolean and Data, a SAT-solver or a numeric constraint solver cannot solve them alone. Satisfiability Modulo Theories (SMT) solvers find solutions for propositional satisfiability problems where propositions are either Boolean or constraints in a specific theory. However, SMT-solvers are not applicable in our case either, because unlike SAT-solvers they find only an instance of a solution as opposed to the complete set of answers. Another drawback of most SAT- and SMT-solvers is that they work only on quantifier-free formulae.

To generate the CASM corresponding to a given Reo network, we need all solutions and thus resort to a hybrid approach that uses both SAT-solvers and Computer Algebra Systems (CASs). First, we form a pure Boolean constraint system by substituting data-dependent constraints with new Boolean variables. We find all solutions for the new constraints using a SAT-solver. Then, by substituting each such solution into the original constraints, we obtain a data dependent constraint satisfaction problem that a CAS can solve symbolically. From these solutions, we extract a CASM corresponding to the Reo network encoded by the original set of constraints. Our approach avoids state explosion by treating data constraints symbolically. In the following we elaborate on our approach.

In an RCSP $\langle \mathcal{P}, \mathcal{M}, M_0, X, C \rangle$, let X_B and X_D be the sets of free Boolean and free data variables of C , respectively, where $X = X_B \cup X_D$, and let A_D be the set of atomic predicates of C containing data variables. The following is our procedure for solving C .

1. We obtain C_B from C by replacing every occurrence of $a \in A_D$ with a unique new Boolean variable $n \notin X$. For example, for $C = (\tilde{f} \rightarrow \tilde{e}) \wedge (\tilde{f} \rightarrow (isNumber(\tilde{e}) \rightarrow \tilde{e}\%2 = 0))$ in Example 1, we obtain C_B as $(\tilde{f} \rightarrow \tilde{e}) \wedge (\tilde{f} \rightarrow (n_1 \rightarrow n_2))$ where n_1 and n_2 replace $isNumber(\tilde{e})$ and $\tilde{e}\%2 = 0$, respectively.
2. An off-the-shelf SAT-solver can find the set of solution S_B for C_B . We define the finite set of constraints $C[S_B] = \{C[x_1, x_2, \dots, x_n \setminus z_1, z_2, \dots, z_n] \mid \text{for all distinct } x_i \in X_B, 1 \leq i \leq n = |X_B|, z_i \in S(x_i), S \in S_B\}$.
3. Every $C_D \in C[S_B]$ is a numerical constraint satisfaction problem, which we (symbolically) solve using a Computer Algebra System. Every solution to each C_D along with the SAT solution $S \in S_B$ that produced $C_D \in C[S_B]$ in the previous step, constitute a solution to the RCSP.

Using the presented technique, we obtain the solutions for the RCSP corresponding to Examples 1 as follows:

1. $\langle \{\tilde{a} = \perp, \tilde{b} = \perp, \tilde{c} = \perp\}, \top \rangle$,
2. $\langle \{\tilde{a} = \top, \tilde{b} = \perp, \tilde{c} = \perp\}, \tilde{a} \notin Number \rangle$,
3. $\langle \{\tilde{a} = \top, \tilde{b} = \top, \tilde{c} = \perp\}, \tilde{a} \in Number \wedge \tilde{b} = 3 * \tilde{a} \wedge \tilde{b}\%2 \neq 0 \rangle$,
4. $\langle \{\tilde{a} = \top, \tilde{b} = \top, \tilde{c} = \top\}, \tilde{a} \in Number \wedge \tilde{b} = 3 * \tilde{a} \wedge \tilde{b}\%2 = 0 \wedge \tilde{b} = \tilde{c} \rangle$.

Example 2. Figure 2 depicts a Reo network that consists of a *LossySync* channel and a *FIFO*₁ channel connecting on the node b .

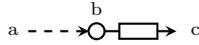


Fig. 2: A context sensitive Reo network

Since the Reo network in Figure 2 contains a *LossySync* that is a context dependent channel, we use the context-aware RCSP encoding from Table 3:

$$\psi_{LossySync}(a, b) = \tilde{b} \rightarrow (\tilde{a} \wedge (\hat{a} = \hat{b})) \wedge \neg \vec{a}_c \wedge \neg \tilde{a} \rightarrow \vec{b}_k \quad (3)$$

$$\begin{aligned} \psi_{FIFO_1}(b, c, m) = \tilde{b} \rightarrow (-\dot{m} \wedge \dot{m}' \wedge (\dot{m}' = \hat{b})) \wedge \tilde{c} \rightarrow (\dot{m} \wedge \neg \dot{m}' \wedge (\dot{m}' = \\ \hat{c})) \wedge (\neg \tilde{b} \wedge \neg \tilde{c}) \rightarrow ((\dot{m} \leftrightarrow \dot{m}') \wedge \dot{m} \rightarrow (\hat{m} = \hat{m}')) \wedge \neg \dot{m} \rightarrow \overrightarrow{c'_c} \wedge \dot{m} \rightarrow \overrightarrow{b'_k} \end{aligned} \quad (4)$$

Equation 3 states that flow on the sink end of the *LossySync* is due to flow on its source end. If there is flow on the sink end of the *LossySync*, the data items exchanged at the source and the sink ends are the same. However, it is possible that the source end has flow, but the sink end does not. In this case, the reason for no flow comes from the environment with which the sink end communicates. The third possible behavior of the channel is that there is no flow on the source end due to the environment, in which case the channel provides a reason for no flow on its sink end.

Equation 4 states the behavior of the *FIFO*₁ channel as follows: Flow on the source end of the channel states that the value of the variable representing the state memory (of the current state) is undefined. Flow on the source end defines the state memory variable for the next state to contain the value of the incoming data item. On the other hand, flow on the sink end means that the value of the state memory variable is defined. The data item leaving the sink end is equivalent to the buffer's data item. In addition, the value of the state memory variable becomes undefined in the next state. If there is no flow on the ends, the variables related to the states stay the same. Being empty, the *FIFO*₁ channels provides a reason for no flow on its sink end, while being full does so on the source end of the channel.

Following are the solutions for the RCSP of Example 2, where for brevity, we have omitted the values of the variables representing the context, such as $\overrightarrow{b'_c}$:

1. $\langle \{\tilde{a} = \perp, \tilde{b} = \perp, \tilde{c} = \perp, \dot{m} = \perp, \dot{m}' = \perp\}, \top \rangle$,
2. $\langle \{\tilde{a} = \top, \tilde{b} = \top, \tilde{c} = \perp, \dot{m} = \perp, \dot{m}' = \top\}, \hat{a} = \hat{b} \wedge \hat{m}' = \hat{b} \rangle$,
3. $\langle \{\tilde{a} = \top, \tilde{b} = \perp, \tilde{c} = \perp, \dot{m} = \top, \dot{m}' = \top\}, \hat{m} = \hat{m}' \rangle$,
4. $\langle \{\tilde{a} = \perp, \tilde{b} = \perp, \tilde{c} = \perp, \dot{m} = \top, \dot{m}' = \top\}, \hat{m} = \hat{m}' \rangle$,
5. $\langle \{\tilde{a} = \top, \tilde{b} = \perp, \tilde{c} = \perp, \dot{m} = \top, \dot{m}' = \perp\}, \hat{m} = \hat{c} \rangle$,
6. $\langle \{\tilde{a} = \perp, \tilde{b} = \perp, \tilde{c} = \top, \dot{m} = \top, \dot{m}' = \perp\}, \hat{m} = \hat{c} \rangle$.

4.3 CASM Construction

In order to construct the CASM from the set of solutions S for an RCSP $\langle \mathcal{P}, \mathcal{M}, M_0, X, C \rangle$, we first map each solution $\langle s, s_d \rangle \in S$ into a transition $t : q \xrightarrow{N, g} p$ as follows:

- $q = \langle \{m \mid m \in \mathcal{M}, s(\dot{m}) = \top \} \rangle$,
- $p = \langle \{m \mid m \in \mathcal{M}, s(\dot{m}') = \top \} \rangle$,
- $\mathcal{N} = \{n \mid n_c \in \mathcal{P} \vee n_k \in \mathcal{P}\}$,
- $N = \{n \mid n \in \mathcal{N}, s(\tilde{n}) = \top \}$,
- The data constraint g is (a syntactic variant of) the s_d .

We obtain the CASM $A = (Q, \mathcal{N}, \rightarrow, q_0, \mathcal{M})$ from the set T of all transitions generated above, where:

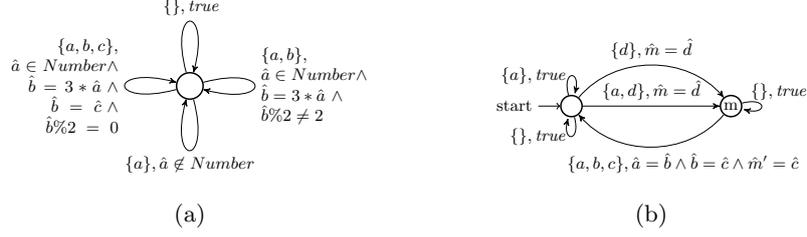


Fig. 3: CASMs generated for a) Example 1 and b) Example 2

- $Q = \{q \mid q \xrightarrow{N,g} p \in T \vee p \xrightarrow{N,g} q \in T\}$,
- $q_0 = \langle \{m \mid m \in M_0, s(\hat{m}) = \top \} \rangle$,
- \rightarrow is the set of all transitions, T ,
- \mathcal{M} is the same \mathcal{M} as in the RCSP.

Figures 3a and 3b show the CASMs derived from the RCSPs in Examples 1 and 2.

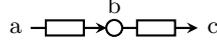
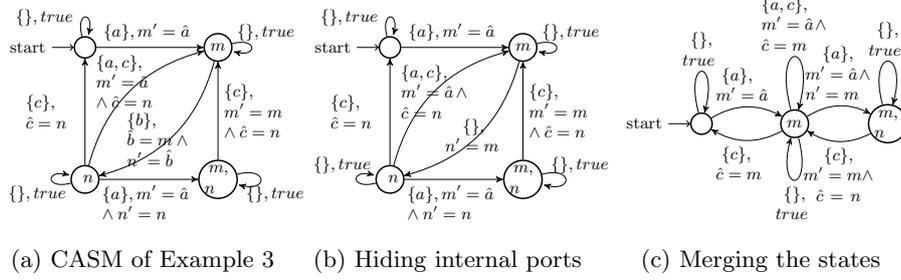
5 Hiding

We use *hiding* to abstract from internal transitions. The author in [19] proposes applying the existential quantifier to the constraints encoding of the behavior of a network to abstract from internal ports and their corresponding data variables. Similarly, we use existential quantifiers such as $\exists \hat{e}, \hat{e}, \vec{e} : C$, where C is the RCSP of a Reo network and e is an internal node to hide.

Although several algorithms exist for the problem of quantifier elimination in Boolean algebra and first order logic [11], [6] and [14], we are not aware of any working tool that does quantifier elimination on Boolean algebraic formulae. Therefore, our tool implements the *hiding* operator as defined for CASM.

Hiding the internal nodes on some transitions can make the set of their synchronized nodes empty. Here, we refer to such a transition as a *empty* transition, if the free variables of its guard are merely state memory variables. Under some circumstances, we can merge the source and the target states of empty transitions. Let q and p be two states in a CASM such that $q \xrightarrow{\emptyset, g} p$. The following is the conditions under which the state p can merge into the state q :

1. The states q and p have the same number of state memory variables.
2. The guard g consists of the conjunction of the predicates in the forms of $m = n'$. This way, g defines a correspondence relation between the state memory variables of the state q and those of the state p .
3. For each transition like $q \xrightarrow{N, g'} r$ where $r \notin \{p, q\}$, there is a transition $p \xrightarrow{N, g''} r$ such that $g' \leftrightarrow g''$, where g'' is obtained by replacing all the variables of the form n' with m' in g' , only if g contains $m = n'$.

Fig. 4: Two $FIFO_1$ s forming $FIFO_2$ Fig. 5: Hiding the empty transition and merging its source and target states for the CASM of $FIFO_2$ in Figure 4

4. For each transition like $r \xrightarrow{N, g'} p$ where $r \notin \{p, q\}$, there is a transition $r \xrightarrow{N, g''} q$ such that $g'' \leftrightarrow g'_g$, where g'_g is made by substituting all the variables of the form n with m in g' , only if g includes $m = n'$.

In the case that above conditions hold, the state p merges into the state q as follows:

1. We eliminate the transition $q \xrightarrow{\emptyset, g} p$.
2. We remove the state p after substituting n , n' , and p with m , m' , and q in all transitions. Observe that such substitutions convert the non-eliminated transitions between the states q and p into loops over the state q .

Example 3. Figure 4 shows a $FIFO_2$ that consists of two $FIFO_1$ s. The CASM corresponding to the $FIFO_2$ is in Figure 5a. Figure 5b depicts the CASM resulting from hiding the mixed node b . Figure 5c presents the result of eliminating the empty transitions.

6 Discussions

The existing semantic models for Reo are not all equally expressive, as some aspects of the (intended) semantics of Reo cannot be expressed in some of these models. A classical example of this sort is a *LossySync* channel connected to a $FIFO_1$ channel. Since CA are not context dependent, the CA model of this Reo connector considers it a valid behavior when the *LossySync* channel loses an incoming data item while the $FIFO_1$ is empty. However, according to the CC semantics a context dependent *LossySync* channel does not lose data in this situation.

On the other hand, because CC does not consider data constraints, there are cases where some flows that the CC semantic model considers legitimate, are not allowed in the CA model. An example of this type is a *Transformer* channel, which produces certain data items connected to a *Filter* channel not accepting those values. An advantage of our method is that it considers flows that are valid with respect to both semantics.

The RCSP encoding of Reo networks that we have presented in this paper has the properties of soundness and completeness with respect to the CA and the CC models. Moreover, it is compositional. In future work, we will prove these properties as the space restrictions do not allow us to address these questions in the current paper.

7 Conclusions and Future Work

In this paper, we have presented a constraint-based framework, which encodes the semantic models of Reo networks as constraint satisfaction problems whose predicates are either Boolean propositions or numerical constraints. We presented a hybrid approach to find the solutions for these problems. An advantage of our approach is that it treats data constraints symbolically to mitigate the state explosion problem. From this solution, we construct the semantic model corresponding to a Reo network in the form of constraint automata with state memory. Our framework supports *product* and *hiding* operations on constraint automata. We have implemented and integrated our approach in the ECT. As part of our ongoing work, we are using this framework to encode other aspects of the semantics of Reo, namely, priorities and timed behavior. In this way, our work will be the most expressive framework that exists to analyze Reo networks. Furthermore, we will prove soundness and completeness of the RCSP encoding of Reo networks along with its compositionality.

References

1. Farhad Arbab. Reo: a Channel-Based Coordination Model for Component Composition. *Mathematical Structures in Computer Science*, 14:329–366, 2004.
2. Farhad Arbab, Christel Baier, Frank De Boer, and Jan Rutten. Models and Temporal Logics for Timed Component Connectors. In *2nd International Conference on Software Engineering and Formal Methods (SEFM 2004)*, pages 198–207. IEEE Computer Society, 2004.
3. Farhad Arbab, Christel Baier, and Marjan Sirjani. Priority in Reo and Constraint Automata. Technical report, Centrum voor Wiskunde en Informatica. In preparation.
4. Farhad Arbab, Christian Koehler, Ziyang Maraikekar, Young-Joo Moon, and Jose Proenca. Modeling, Testing and Executing Reo Connectors with the Eclipse Coordination Tools. In *5th International Workshop on Formal Aspects of Component Software (FACS 2008)*, volume 8. ENTCS, 2008.
5. Farhad Arbab, Natallia Kokash, and Sun Meng. Towards Using Reo for Compliance-Aware Business Process Modeling. In *ISoLA*, pages 108–123, 2008.

6. Ayari, A. and Basin, D. QUBOS: Deciding Quantified Boolean Logic using Propositional Satisfiability Solvers. In *Formal Methods in Computer-Aided Design*, pages 187–201. Springer, 2002.
7. Bahman Pourvatan and Marjan Sirjani and Hossein Hojjat and Farhad Arbab. Symbolic Execution of Reo Circuits using Constraint Automata. *Sci. Comput. Program.*, 77(7-8):848–869, 2012.
8. Christel Baier, Marjan Sirjani, Farhad Arbab, and Jan J. M. M. Rutten. Modeling Component Connectors in Reo by Constraint Automata. *Science of Computer Programming*, 61(2):75–113, 2006.
9. Michael Bell. Introduction to Service-Oriented Modeling. *Service-Oriented Modeling: Service Analysis, Design, and Architecture*, page 3, 2008.
10. Marcello Bonsangue, Dave Clarke, and Alexandra Silva. Automata for Context-Dependent Connectors. In *Proceedings of the 11th International Conference on Coordination Models and Languages, COORDINATION '09*, pages 184–203, Berlin, Heidelberg, 2009. Springer-Verlag.
11. Lucas Bordeaux and Lintao Zhang. A Solver for Quantified Boolean and Linear Constraints. In *Proceedings of the 2007 ACM symposium on Applied computing, SAC '07*, pages 321–325, New York, NY, USA, 2007. ACM.
12. Behnaz Changizi, Natallia Kokash, and Farhad Arbab. A Unified Toolset for Business Process Model Formalization. In *7th International Workshop on Formal Engineering approaches to Software Components and Architectures (FESCA 2010)*, pages 147–156. ENTCS, 2010.
13. Dave Clarke, David Costa, and Farhad Arbab. Connector Colouring I: Synchronisation and Context Dependency. *Science of Computer Programming*, 66(3):205–225, 2007.
14. J. H. Davenport. Computer algebra applied to itself. *J. Symb. Comput.*, 6:127–132, August 1988.
15. Jan Friso Groote, Aad Mathijssen, Michel Reniers, Yaroslav Usenko, and Muck Van Weerdenburg. The Formal Specification Language mCRL2. In *Methods for Modelling Software Systems (MMOSS 2006)*, volume 06351 of *Dagstuhl Seminar Proceedings*. IBFI, 2006.
16. Sung-Shik T.Q. Jongmans and Farhad Arbab. Overview of Thirty Semantic Formalisms for Reo. *Scientific Annals of Computer Science*, 22:201–251, 2012.
17. Stephanie Kemper. SAT-based Verification for Timed Component Connectors. *Electr. Notes Theor. Comput. Sci.*, 255:103–118, 2009.
18. Natallia Kokash, Christian Krause, and Erik de Vink. Reo + mCRL2: A Framework for Model-checking Dataflow in Service Compositions. *Formal Aspects of Computing*, 2011.
19. José Proença. *Synchronous Coordination of Distributed Components*. PhD thesis, Institut for Programming research and Algorithms, 2011.