



Centrum voor Wiskunde en Informatica

**REPORT**RAPPORT

*SEN*

Software Engineering



*Software ENgineering*

Models and Temporal Logics for Timed Component  
Connectors

F. Arbab, C. Baier, F.S. de Boer, J.J.M.M. Rutten

**REPORT SEN-R0411 JULY 2004**

CWI is the National Research Institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organization for Scientific Research (NWO).

CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

**Software Engineering (SEN)**

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2004, Stichting Centrum voor Wiskunde en Informatica

P.O. Box 94079, 1090 GB Amsterdam (NL)

Kruislaan 413, 1098 SJ Amsterdam (NL)

Telephone +31 20 592 9333

Telefax +31 20 592 4199

ISSN 1386-369X

# Models and Temporal Logics for Timed Component Connectors

## ABSTRACT

Component-based software engineering advocates construction of software systems through composition of coordinated autonomous components. Significant benefits of this approach include software reuse, simpler and faster construction, enhanced reliability, and dramatic reductions in the complexity of construction of provably correct critical systems, many of which involve real-time concerns. Effective, flexible component composition by itself still poses a challenge today and yet the special nature of real-time constraints makes component-based construction of real-time systems even more demanding. The coordination language Reo supports compositional system construction through connectors that exogenously coordinate the interactions among the constituent components which unawaresly comprise a complex system, into a coherent collaboration. The simple, yet surprisingly rich, calculus of channel composition that underlies Reo offers a flexible framework for compositional construction of coordinating component connectors with real-time properties. In this paper, we present an operational semantics for the channel-based component connectors of Reo in terms of Timed Constraint Automata and introduce a temporal-logic for specification and verification of their real-time properties.

*1998 ACM Computing Classification System:* C.2.4, D.1.3, D.2.4, D.2.6, D.2.11, D.2.13, D.3.2, D.3.3, F.1.2, F.3.1, F.3.2, F.3.3

*Keywords and Phrases:* Coordination, Real-time, Composition, Reo, Constraint Automata, Timed Automata, Linear Temporal Logic, Timed Data Streams

# Models and Temporal Logics for Timed Component Connectors

Farhad Arbab<sup>1,3</sup>, Christel Baier<sup>2</sup>, Frank de Boer<sup>1,3</sup>, Jan Rutten<sup>1,4</sup>

<sup>1</sup>Centrum voor Wiskunde en Informatica, Department of Software Engineering, Amsterdam, The Netherlands

<sup>2</sup>Universität Bonn, Institut für Informatik I, Germany

<sup>3</sup> Universiteit Leiden, The Netherlands

<sup>4</sup> Vrije Universiteit Amsterdam, The Netherlands

## Abstract

Component-based software engineering advocates construction of software systems through composition of coordinated autonomous components. Significant benefits of this approach include software reuse, simpler and faster construction, enhanced reliability, and dramatic reductions in the complexity of construction of provably correct critical systems, many of which involve real-time concerns. Effective, flexible component composition by itself still poses a challenge today and yet the special nature of real-time constraints makes component-based construction of real-time systems even more demanding.

The coordination language Reo supports compositional system construction through connectors that exogenously coordinate the interactions among the constituent components which unawaresly comprise a complex system, into a coherent collaboration. The simple, yet surprisingly rich, calculus of channel composition that underlies Reo offers a flexible framework for compositional construction of coordinating component connectors with real-time properties. In this paper, we present an operational semantics for the channel-based component connectors of Reo in terms of Timed Constraint Automata and introduce a temporal-logic for specification and verification of their real-time properties.

Keywords and phrases: Coordination, Real-time, Composition, Reo, Constraint Automata, Timed Automata, Linear Temporal Logic, Timed Data Streams

1998 ACM Computing Classification: C.2.4, D.1.3, D.2.4, D.2.6, D.2.11, D.2.13, D.3.2, D.3.3, F.1.2, F.3.1, F.3.2, F.3.3

## 1 Introduction

The task of designing a complex concurrent system with several components requires a *coordination model* that formalizes their mutual interactions. The internals of black-box components cannot be modified to implement such coordinated interactions. Coordination, therefore, becomes the responsibility of the “glue-code” that inter-connects the constituent components of a composite system, and of its underlying run-time middle-ware. Reo [5] offers a powerful glue language for implementation of coordinating component connectors based on a calculus of mobile channels.

In this paper, we consider the real-time aspects of Reo when the behavior specification of channels and component interfaces can involve *timing constraints*. Because connectors, not components, are the primary concern in Reo, our primary interest here is with channels whose behavior involves temporal constraints; and with their composition. For instance, a deadline  $t$  for the availability of some data can be formalized as the behavior of a FIFO channel that associates an *expiration date*,  $t$ , with every data item that enters its buffer: the channel loses a data item in its buffer  $t$  units of time after it enters through its source (unless, of course, it is dispensed through its sink in the meanwhile). Another example is a timer channel that gets activated by a data item through its source, after which it returns a timeout signal through its sink, after a specified delay of exactly  $t$  units of time.

As the operational model for Reo connector circuits, we use *timed constraint automata* (TCA) which extend their untimed version [7] with the concepts borrowed from classical timed automata with location invariants [1, 18]. TCA have two kinds of transitions: (1) internal changes of the locations caused by some

time constraints and (2) transitions that represent the synchronized execution of I/O-operations at some of the ports. Using ideas similar to [7], the construction of a timed constraint automaton from a given timed Reo circuit can be performed in a *compositional* manner, using composition operators on TCA that model Reo’s operators *join* and *hide* to build complex connectors out of instances of basic channel types.

One conceptual difference between TCA and classical timed automata is the treatment of immediate actions or urgent synchronous channels, as they are used e.g., in the tools [20, 17, 32]. The assumption that synchronous I/O-operations must be executed as soon as they become enabled makes no sense in our framework. For instance, assume that we have a FIFO channel carrying data from node *A* to node *B* and a synchronous channel from *B* to another node *C*. As soon as *A* places a value in the FIFO buffer it becomes available for consumption through node *B*, and thus, the synchronous communication between *B* and *C* become enabled. On the other hand, the input and output of the same data item must not occur simultaneously through a FIFO channel, by its definition. Thus, we need a delay for the synchronization between *B* and *C*. Moreover, Reo allows to explicitly specify deadlines of “shortly delayed” activities or other time constraints (e.g., lower bounds for the delay) using an appropriate combination of timed channels.

The semantics of the TCA and timed Reo circuits relies on *timed data streams* as in [8, 7], comprising a formalization of the possible data-flow at each node over time. To specify a desired coordination mechanism, we use a variant of linear temporal logic (LTL) [25, 22] with real-time constraints, which we call *timed scheduled-data-stream logic* (TSDSL) and has a semantics based on timed data streams. TSDSL essentially relies on a combination of the time-abstract temporal modalities in LTL and timed regular expressions [9]. We show through a series of examples how TSDSL can serve as a specification formalism for (timed) Reo circuits, sketch the ideas of a model checking algorithm, and explain the relation of TSDSL with refinement relations.

**Related models.** There are several other related real-time models that also focus on aspects of coordination. Timed interface automata (TIA) [11] or real-time variants of I/O-automata, e.g., [23, 14, 19], are related to TCA in the same way as their untimed versions. I/O-automata rely on the assumption of input-enabledness which is not required (and would not make sense) in constraint automata.

The purpose of TIA is orthogonal to our approach involving timed Reo connectors and TCA. (There are some conceptual differences, e.g., TIA use action labels rather than port names, but these are not important as the formal definition of TIA and TCA can be adapted to eliminate these differences.) The major goal of TIA is to provide a formalism to specify and to check the compatibility of real-time components by means of their interfaces. Our focus is on compositional reasoning about (design and analysis of) channel-based coordination mechanisms, based on their data-flow. Thus, our framework allows to design and analyze a coordination context in which certain components are used and to construct their interfaces, while the approach of interface automata allows to check a-posteriori whether a design makes the components work together in the desired way.

Although compositionality in timed Reo and TCA is in the spirit of real-time process algebras, e.g., [26, 31, 21], Reo’s philosophy of composing connectors out of a variety of basic channel types via *join* and *hide* and supporting any kind of synchronous or asynchronous communication differs from classical process algebra approaches which provide operators for modeling choice, parallelism, and recursion (all of which are implicit in Reo).

**Organization of the paper.** Timed constraint automata are introduced in Section 2. In Section 3 we explain the main features of Reo circuits and how timed constraint automata can serve as their operational model. Timed scheduled-data-stream logic (TSDSL) is introduced in Section 4. Section 5 concludes the paper.

The submitted version has an appendix which contains some technical details and pictures for larger TCA. To meet the length restriction, the appendix will be removed in the final version.

## 2 Timed constraint automata

Edges in timed constraint automata are labeled with tuples  $(N, dc, cc, C)$  where  $N$  is a set of ports/nodes that synchronously perform certain I/O-operations,  $dc$  is a data constraint that specifies the concrete values that are transferred through those I/O-operations,  $cc$  is a clock constraint, and  $C$  is a set of clocks that are reset to 0. If  $N = \emptyset$  then the edge represents an internal move (in which case  $dc = \text{true}$ ). Before presenting the formal definition, we give a simple example. Fig. 1 shows on its left a Reo circuit with a 1-bounded FIFO-channel with expiration connecting nodes  $A$  and  $B$  and a synchronous channel connecting nodes  $B$  and  $C$ . A FIFO channel “with expiration” is a lossy channel that loses any data item that remains in its buffer longer than its “expiration date” which in this case is 3 time units after it enters the buffer of the channel. Thus, in this example, there is an implicit deadline for the data transfer operation at node  $B$ . The picture on the right shows the TCA corresponding to this Reo circuit. In the




Figure 1: Reo circuit and timed constraint automaton

TCA on the right-hand-side in Fig. 1, location  $s$  stands for the initial configuration where the buffer is empty, while location  $\bar{s}(d)$  represents the configuration where the buffer is filled with data element  $d$ . If nodes  $B$  and  $C$  are ready for I/O-operations within 3 time units, in location  $\bar{s}(d)$  then we assume that  $B$  takes an element  $d$  from the buffer and immediately forwards it to  $C$ . This corresponds to the transition labeled with the set  $\{B, C\}$  and the data constraint  $d_B = d_C = d$ . Although there is no explicit lower time bound for the delay of the  $\{B, C\}$ -transition, our semantics forces some time elapse in location  $\bar{s}(d)$  before the  $\{B, C\}$ -transition can fire, even if  $B$  and  $C$  are waiting for an input value. This is different in ordinary timed automata, but is needed here because a FIFO channel (by its definition) does not allow for the synchronous transfer of data from its source to its sink end. If  $B$  cannot transfer the element out of the FIFO buffer (because no I/O operation is available on  $C$  to synchronize with  $B$ ), the message is lost 3 time units after entering  $\bar{s}(d)$ . This is modeled by the invariance condition  $x \leq 3$  at location  $\bar{s}(d)$  which forces the automaton to leave  $\bar{s}(d)$  if the current value of  $x$  is 3.

**Notation 2.1 (Data assignments, data constraints)** In the sequel, we assume finite and non-empty sets  $Data$  consisting of data items that can be transferred through channels, and  $\mathcal{N}$  consisting of node names. A data assignment denotes a function  $\delta : N \rightarrow Data$  where  $\emptyset \neq N \subseteq \mathcal{N}$ . We use notations like  $\delta = [A \mapsto \delta_A : A \in N]$  to describe the data-assignment that assigns the value  $\delta_A \in Data$  to every node  $A \in N$ . Data constraints can be viewed as a symbolic representation of *sets* of data assignments. Formally, data constraints (denoted  $dc$ ) are propositional formulas built from the atoms “ $d_A \in P$ ” and “ $d_A = d_B$ ” where  $A, B \in \mathcal{N}$  and  $P \subseteq Data$  (plus the standard boolean connectors  $\wedge, \vee, \neg$ , etc.). For  $N \subseteq \mathcal{N}$ ,  $DA(N)$  denotes the set of all data assignments for the node-set  $N$  and  $DC(N)$  the set of data constraints that at most refer to the terms  $d_A$  for  $A \in N$ . We write  $DA$  for  $\bigcup_{\emptyset \neq N \subseteq \mathcal{N}} DA(N)$  and  $DC$  for  $DC(\mathcal{N})$ .  $\square$

**Notation 2.2 (Clock assignments, clock constraints)** Let  $\mathcal{C}$  be a finite set of clocks. A clock assignment means a function  $\nu : \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ . If  $t \in \mathbb{R}_{\geq 0}$  then  $\nu + t$  denotes the clock assignment that assigns the value  $\nu(x) + t$  to every clock  $x \in \mathcal{C}$ . If  $C \subseteq \mathcal{C}$  then  $\nu[C := 0]$  stands for the clock assignment that returns the value 0 for every clock  $x \in C$  and the value  $\nu(x)$  for every clock  $x \in \mathcal{C} \setminus C$ . A clock constraint (denoted  $cc$ ) for  $\mathcal{C}$  is a conjunction of atoms of the form “ $x \bowtie n$ ” where  $x \in \mathcal{C}$ ,  $\bowtie \in \{<, \leq, >, \geq, =\}$  and  $n \in \mathbb{N}$ .  $CA(\mathcal{C})$  (or  $CA$ ) denotes the set of all clock assignments and  $CC(\mathcal{C})$  (or  $CC$ ) the set of all clock constraints.  $\square$

The symbol  $\models$  stands for the obvious satisfaction relation for data (or clock) constraints which results from interpreting data (clock) constraints over data (clock) assignments. Satisfiability, validity, logical equivalence  $\equiv$  and logical implication  $\leq$  of data (clock) constraints are defined as usual. For data constraints, we often use simplified notations such as “ $d_A = d$ ” rather than “ $d_A \in \{d\}$ ”.

**Definition 2.3 (Timed constraint automata)** A TCA is a tuple  $T = (S, \mathcal{C}, \mathcal{N}, \mathcal{E}, S_0, ic)$  where  $S$  is a finite set of control states (also called locations),  $\mathcal{C}$  a finite set of clocks,  $\mathcal{N}$  a finite set of node names, and  $S_0 \subseteq S$  a set of initial locations.  $ic : S \rightarrow CC$  is a function that assigns to any location  $s$  an invariance condition  $ic(s)$ . The edge relation  $\mathcal{E}$  is a subset of  $S \times 2^{\mathcal{N}} \times DC \times CC \times 2^{\mathcal{C}} \times S$  such that  $dc \in DC(\mathcal{N})$  for any edge  $e = (s, N, dc, cc, C, \bar{s}) \in \mathcal{E}$ . Moreover, we assume that all data and clock guards on the edges and the invariance conditions are satisfiable. (For edges with the empty node-set, we require a data constraint  $dc$  with  $dc \equiv \text{true}$ .)  $\square$

The automaton in Fig. 1 is a simplified picture for a TCA where  $d$  is used as a data parameter. The presented TCA has the location space  $S = \{s\} \cup \{\bar{s}(d) : d \in \text{Data}\}$ . The assignment “ $d := d_A$ ” in the parametric version stands for the data constraint  $d_A = d$  in the TCA. An interface specification of a *timed sequencer* that coordinates the data-flow of two components via synchronous channels is shown in Fig. 2. We assume the deadline  $t = 3$  for the write-operations, that is, the sequencer in location  $s$  waits up to  $t$  time units to synchronize with component 1. If it fails then the sequencer moves via the edge labeled with the empty set to location  $\bar{s}$  and tries to synchronize with component 2, and so on.





Figure 2: Timed sequencer

**Example 2.4 (Alternating Bit Protocol)** We consider a variant of the ABP where two components (the sender and the receiver) are connected via lossy synchronous channels. We follow here essentially the description in [24] but do not assume unreliable channels that may lose data in an unpredictable way. Instead, we assume lossy synchronous channels (as in Reo, see Section 3) where a data item written to the source end of such a channel is lost if the sink end of the channel cannot perform a matching I/O-operation to consume it.



Via its input port  $I$ , the sender is fed with some input which it delivers to the receiver via the channel connecting ports  $A$  and  $C$ . The receiver acknowledges the receipt of the message via the channel between  $D$  and  $B$  and outputs the message through its port  $O$ . The sender attaches a bit to the messages and expects the corresponding control bit as acknowledgment. If the expected control bit  $b$  arrives through port  $B$  then the sender switches its mode and sends the next message together with the bit  $\neg b$ . If a certain deadline ( $t_S$  in our example) expires then the sender resends the message with the same control bit  $b$  with a delay of at most  $\rho_S$ . The same upper bound  $\rho_S$  is assumed for the time interval between the receipt of a message  $d$  on input port  $I$  and the sending a message from output port  $A$ . Acknowledgments that contain a non-expected control bit are ignored as they belong to the previous message.

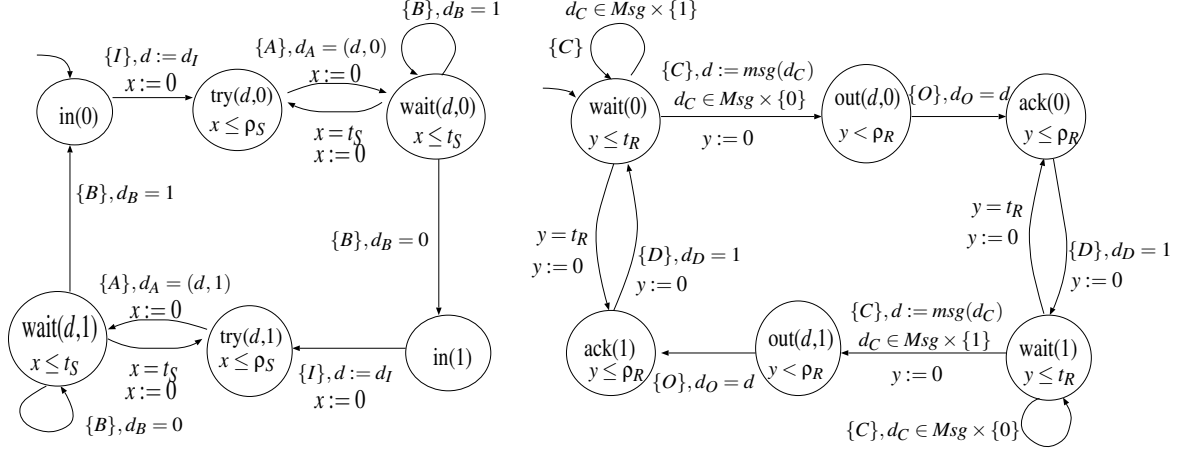


Figure 3: TCA for the sender and the receiver of the ABP

The behavior of the receiver is complementary to that of the sender. In mode  $b$ , the receiver waits for the arrival of an input  $(d, b)$  through its port  $C$  and acknowledges its receipt with the bit  $b$ , while messages of the form  $(d, -b)$  are ignored. The receiver resends the acknowledgment if the next message with the expected control bit  $b$  does not arrive within  $t_R$  time units. (In particular, the receiver resends the control bit of the last message infinitely many times if data-flow at port  $I$  eventually terminates.) Moreover, we assume the upper time bound  $\rho_R$  for the success of the write-operation on output port  $O$  as well as the receiver's acknowledgment by sending the control bit. Fig. 3 shows the interface specifications for the sender and the receiver by (data parametrized) TCA. We assume here the data domain  $Data = \{0, 1\} \cup Msg \cup Msg \times \{0, 1\}$  and write  $msg$  to denote the projection of the pairs  $(d, b)$  to the message-component (i.e.,  $msg(d, b) = d$ ). Fig. 9 shows the “combined” TCA  $\mathcal{T}_{ABP}$  for the ABP.<sup>1</sup>  $\square$

**Definition 2.5 (State-transition graph of a TCA)** *Given a TCA  $\mathcal{T}$  as above,  $\mathcal{T}$  induces a state-transition graph  $\mathcal{A}_{\mathcal{T}} = (Q, \longrightarrow, Q_0)$  as follows. The states are pairs  $q = \langle s, \nu \rangle$  consisting of a location  $s$  and a clock assignment  $\nu$ . Thus, the state space is  $Q = S \times CC$ . The set of initial states is  $Q_0 = \{\langle s_0, \mathbf{0} \rangle : s_0 \in S_0, \mathbf{0} \models ic(s_0)\}$  where  $\mathbf{0}$  stands for the clock assignment that returns the value 0 for all clocks. The transition relation  $\longrightarrow \subseteq Q \times 2^N \times DA \times \mathbb{R}_{\geq 0} \times Q$  is defined by the following rules:*

$$\frac{\begin{array}{l} (s, N, dc, cc, C, \bar{s}) \in \mathcal{E}, \\ t > 0 \text{ s.t. } \nu + \bar{t} \models ic(s) \text{ for all } 0 < \bar{t} \leq t \\ (\nu + t)[C := 0] \models ic(\bar{s}) \text{ and } \nu + t \models cc \\ \delta \in DA(N) \text{ s.t. } \delta \models dc \end{array}}{\langle s, \nu \rangle \xrightarrow{N, \delta, t} \langle \bar{s}, (\nu + t)[C := 0] \rangle}$$

If  $N = \emptyset$ , we use in addition the same rule with  $t = 0$ :

$$\frac{(s, \emptyset, \text{true}, cc, C, \bar{s}) \in \mathcal{E}, \nu[C := 0] \models ic(\bar{s}), \nu \models cc}{\langle s, \nu \rangle \xrightarrow{\emptyset, \emptyset, t} \langle \bar{s}, \nu[C := 0] \rangle}$$

A state  $q = \langle s, \nu \rangle$  is called *terminal* iff it has no outgoing transitions, but allows the possibility for unbounded passage of time, i.e.,  $\nu + t \models ic(s)$  for all  $t > 0$ . A *time-lock* refers to a state  $q = \langle s, \nu \rangle$  that has no outgoing transitions and there exists a  $t > 0$  with  $\nu + t \not\models ic(s)$ .  $\mathcal{T}$  is called *time-lock free* iff  $\mathcal{A}_{\mathcal{T}}$  does not contain a reachable time-lock.  $\square$

Edges with non-empty node-sets can fire only after some positive delay. This reflects the general idea of constraint automata where all observable activities that occur at the same time instant (i.e., atomically) are collapsed into a single transition.

<sup>1</sup>Essentially, this TCA is obtained by the join operator (cf. Def. 3.4), while taking care of the special semantics of lossy synchronous channels, which forces its sink and the source ends to synchronize if both can perform I/O-operations.



**Notation 2.6 (Runs, time divergence)** Let  $\mathcal{T}$  be a TCA as before and  $q = \langle s, \nu \rangle$  a state in  $\mathcal{A}_{\mathcal{T}}$ . A  $q$ -run (or briefly run) in  $\mathcal{T}$  denotes any (finite or infinite) sequence of successive transitions in  $\mathcal{A}_{\mathcal{T}}$  starting in state  $q$ . Formally, a  $q$ -run has the form

$$\mathbf{q} = q_0 \xrightarrow{N_0, \delta_0, t_0} q_1 \xrightarrow{N_1, \delta_1, t_1} \dots$$

where  $q_0 = q$ .  $\mathbf{q}$  is called initial if  $q_0 \in Q_0$ .  $\mathbf{q}$  is called time divergent if  $\mathbf{q}$  is infinite and  $t_0 + t_1 + \dots = \omega$ . Maximality of a run means that it is either time divergent or finite and ends in a terminal state.  $\square$

Intuitively,  $N_i$  is the set of nodes in state  $q_i$  that are scheduled to synchronously perform the next I/O-operations, while  $\delta_i$  represents the concrete values that are exchanged through those operations at the nodes  $A \in N_i$ . The value  $t_i$  stands for the delay.

**Notation 2.7 (TSD stream)** A timed scheduled data stream for a node-set  $\mathcal{N}$  denotes any (finite or infinite) sequence  $\Theta = (N_0, \delta_0, t_0), (N_1, \delta_1, t_1), \dots \in (2^{\mathcal{N}} \times DA \times \mathbb{R}_{\geq 0})^\infty$  such that  $\delta_i \in DA(N_i)$ ,  $0 < t_0 < t_1 < \dots$  and  $\lim_{i \rightarrow \infty} t_i = \omega$  if  $\mathbf{q}$  is infinite. The empty TSD stream is denoted by the symbol  $\varepsilon$ . The length  $|\Theta| \in \mathbb{N} \cup \{\omega\}$  is defined as the number of triples  $(N, \delta, t)$  in  $\Theta$ . The execution time  $\tau(\Theta)$  is  $\omega$  if  $\Theta$  is infinite,  $t_k$  if  $|\Theta| = k + 1$ , and 0 if  $\Theta = \varepsilon$ . We write  $TSDS(\mathcal{N})$  or simply  $TSDS$  to denote the set of all TSDS for node-set  $\mathcal{N}$ .  $\square$

**Notation 2.8 (TSDS-language of a TCA)** If  $\mathbf{q}$  is a run in a TCA  $\mathcal{T}$  as above then the induced TSD stream  $\Theta(\mathbf{q}) = (N_{i_0}, \delta_{i_0}, \bar{t}_{i_0}), (N_{i_1}, \delta_{i_1}, \bar{t}_{i_1}), \dots$  is obtained from  $\mathbf{q}$  by (1) removing all transitions in  $\mathbf{q}$  with the empty node set, (2) building the projection on the transition labels, and (3) replacing the sojourn times  $t_i$  by the absolute time points  $\bar{t}_i = t_0 + \dots + t_i$ . The generated language of a state  $q$  in  $\mathcal{A}_{\mathcal{T}}$  is  $\mathcal{L}(\mathcal{T}, q) = \{\Theta(\mathbf{q}) : \mathbf{q} \text{ is a maximal } q\text{-run}\}$ . The language  $\mathcal{L}(\mathcal{T})$  consists of all TSD streams  $\Theta(\mathbf{q})$  where  $\mathbf{q}$  is a maximal and initial run.  $\square$


For instance, the language of the timed sequencer in Fig. 2 consists of all TSD streams  $\Theta = ((N_i, \delta_i, \bar{t}_i))_i$  where  $N_i \in \{\{A\}, \{B\}\}$  and  $\bar{t}_{i+1} - \bar{t}_i > 3$  if  $N_{i+1} = N_i$ .

### 3 Timed Reo circuits

Reo [5] is a channel-based exogenous coordination model wherein complex coordinators, called *connectors*, are built from instances of basic channel types using certain composition operators. In this paper, as in [8, 7], we do not consider the dynamic behavior of components in creating and composing connectors. We concentrate here on connectors that have graphical representations as *Reo circuits* which express the mechanisms that coordinate the data-flow through the channels connecting the input/output ports of some components.

Reo's notion of *channel* is far more general than its common interpretation and encompasses any primitive communication medium with exactly two ends. Channel ends are classified into *source* ends through which data enter and *sink* ends through which data leave their respective channels. A write operation can be performed on the source end of a channel, providing data to enter into the channel, while a take operation can be performed on the sink end of a channel to obtain data out of the channel. We explain the workings of Reo with a few examples of its basic channel types and formalize their behavior by TCA.

**FIFO channels.** The simplest form of an asynchronous channel is a FIFO channel with one buffer cell, which we denote as *FIFO1*. A *FIFO1* channel is graphically represented by a small box in the middle of an arrow. The buffer is assumed to be initially empty if no data item is shown in the box in its graphical representation (as in the example below). The graphical representation of a *FIFO1* channel whose buffer initially contains a data element  $d$  is the same except that it also shows a  $d$  inside the box representing its buffer.



On the left in this figure, we have a normal *FIFO1* channel which keeps a data item in its buffer until it is taken out through its sink. On the right we show a *lossy* variant, called *expiring FIFO1*, where a data item is lost if it is not taken out of the buffer through the sink end of the channel within  $t$  time units after it enters through its source end.






Figure 4: TCA for a normal and an expiring FIFO1 channels


**Synchronous channels.** A synchronous channel, depicted as a solid arrow, has one source- and one sink-end. Write and take operations must occur simultaneously on the two ends of this channel, which is formalized by a TCA with a single location:



A *P-producer* is a synchronous channel that, like a normal synchronous channel, allows write and take operations to succeed atomically on its source and sink ends, respectively, except that the value dispensed through this channel's sink end is always a data element  $d \in P$ , regardless of the value it consumes through its source end.




A *lossy synchronous channel* (depicted as a dashed arrow) is similar to a normal synchronous channel, except that it always accepts all data items through its source end. If it is possible for it to simultaneously dispense the data item through its sink (e.g., there is a take operation pending on its sink) the channel transfers the data item; otherwise the data item is lost.




The above figure shows a TCA that captures the general “possible” behavior of a lossy synchronous channel. To model the context-sensitive behavior of a lossy channel where the  $\{A\}$ -transition is impossible if  $B$  is


ready to synchronize, the concept of priorities can be used as we explain in the forthcoming paper [6]. More exotic channels permitted in Reo include the *synchronous drain* that has two source ends. Because a drain has no sink end, no data value can ever be obtained from this channel. Thus, all data accepted by this channel are lost. A synchronous drain accepts a data item through one of its ends iff a data item is also available for it to simultaneously accept through its other end as well.



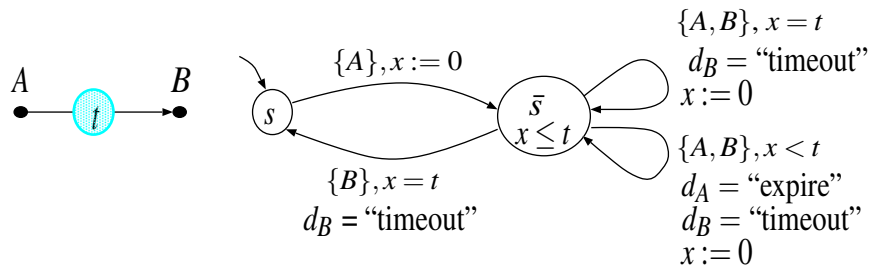
**Timer.** The source end of a  $t$ -timer channel accepts any input value  $d \in Data$  and returns on its sink end a timeout signal after a delay of  $t$  time units.



A  $t$ -timer with the *off-option* allows the timer to be stopped before the expiration of its delay when a special “off” value is consumed through its source end. Similarly, the *reset-option* allows the timer to be reset to 0 after it has been activated when a special “reset” value is consumed through its source end. The following figure shows a  $t$ -timer with both the reset- and the off-options.



A *timer with early expiration* makes the timer produce its timeout signal through its sink and reset itself when it consumes a special “expire” value through its source.



In some cases, it is useful to have a timer that is initially activated. In the graphical representation of this timer, we simply put the word “on” under its circle-symbol. In its TCA, we declare  $\bar{s}$  as the initial location (rather than  $s$ ).




Figure 5: Example construction of a Reo circuit

**Reo circuits.** Complex connectors have graphical representations, called *Reo circuits*, which can be generated by applying certain composition operators to channels. We may think of a Reo-circuit as a finite graph where the *nodes* are labeled with pairwise disjoint, non-empty sets of channel ends and where the edges represent the established channels. The major operations to create Reo connector circuits are join and hiding.

To construct a Reo circuit, we start with several instances of basic channels and organize them in a graph where initially each channel end constitutes a separate node, and each pair of nodes are connected by an edge representing their respective channel. We then apply a series of join operations that take as input two nodes  $A$  and  $B$  and combine them into a new node  $C$ . In this way, several channel ends may coincide on one node. If all channel ends coincident on a node  $C$  are source ends,  $C$  is called a *source node* and it acts as a *replicator*: writing a data item to a source node succeeds when all of its coincident channel ends are capable of accepting the data item simultaneously, in which case the data item is atomically copied into every one of the source ends coincident on  $C$ . If all channel ends coincident on  $C$  are sink ends,  $C$  is called a *sink node* and it behaves as a *merger*: an attempt to take a data item from a sink node succeeds when at least one of its coincident channel ends has a suitable value to offer, in which case the suitable value available through one of these channel ends is non-deterministically selected for the take operation. If  $C$  contains both source and sink channel ends then  $C$  is called a *mixed node* and it behaves as a self-contained pumping station, combining the replicator and merger behavior of source and sink nodes. No take or write operation can be performed on a mixed node; a mixed node autonomously selects suitable values available through its coincident sink ends (merger behavior) and copies them to its coincident source ends (replicator behavior).


The hiding operator allows to create “components” by putting a thick box around a circuit, insulating all of its mixed nodes inside the box and allowing access to its sink and source nodes, placed on the border of the box, only. The idea is that the mixed nodes are internal to the component and no other component can modify or connect to them. Formally, we make hidden (mixed) nodes invisible and abstract their names away.

Fig. 5 demonstrates how to build a Reo circuit via join and hiding. Mixed node  $I$  serves as an initializer which activates the timer. Either  $A$  and  $B$  synchronize before the timer expires or the timeout signal occurs at  $T$  (after exactly  $t$  time units). In either case, the buffer is refilled and the whole procedure restarts.


When modeling Reo circuits by (timed) constraint automata the locations stand for the configurations of the circuits (e.g., contents of the FIFO channels) while the transitions stand for the possible data-flow at one time instance and its effect on the configuration. Intuitively, if we regard a circuit itself as a component, the source nodes of the circuit act as the input ports, and its sink nodes as the output ports of the component. The data-flow through mixed nodes is totally specified by the circuit.

There is a subtle difference between the roles of the sink and source nodes on the one hand and the mixed nodes on the other hand. If an edge contains at least one sink or source node  $A$  then the transition must be regarded as conditional: it can be taken under the condition that the environment that controls the data-flow at node  $A$  (the component that uses  $A$  as an in- or output port) performs the corresponding I/O-operation. On the other hand, any transition with a node-set consisting of mixed nodes only can be taken without any involvement by the environment.

**Example 3.1** The following figure shows on its left how an expiring FIFO1 channel can be constructed out of a normal FIFO1 channel and a timer set to expire after  $t$  time units. On the right we have a circuit that ensures the lower bound “ $>t$ ” for a take operation on  $B$ ; it yields a FIFO1 channel that guarantees every data item will remain in its buffer at least  $t$  time units.



We may also control the frequency of data transfer in synchronous channels with time-constrained channels. In the following figure, on the left, data-flow from  $A$  to  $B$  is possible only once every  $\geq t$  time units.



The  $t$ -timer with early expiration in the circuit on the right ensures that as long as data items are available at  $A$ , they will be consumed at least once every  $t$  time units. Whenever a take operation is performed on  $C$ , the data item available at  $A$  is transferred through  $B$  to  $C$  via the synchronous and the lossy synchronous channels that connect these nodes. The transfer at  $A$  simultaneously produces an “expire” signal (through the  $P$ -producer connected to  $A$ , where  $P$  is the singleton data set  $\{\text{expire}\}$ ) which prematurely fires the timer channel, enabling the synchronous drain to allow the data transfer at  $B$ . If no take operation occurs at  $C$ , the timer produces its timeout-signal after  $t$  time units, enabling the transfer of a data item from  $A$  to  $B$ , because the lossy synchronous channel at  $B$  always accepts (and in this case loses this data item). (Because the two ends of the timer always have to synchronize in this circuit, the assumption that the timer is initially on is essential, since otherwise it can never be started.)  $\square$





Figure 6: Reo circuit for a timed sequencer

**Example 3.2 (Timed sequencer)** The timed sequencer in Fig. 2 can be realized by the Reo circuit shown in Fig. 6 (and hiding all nodes except for  $A$  and  $B$ ). Here, we use a  $t$ -timer with early expiration

which is assumed to be initially switched on.  $A$  can transfer a value only if  $D$  simultaneously also takes a value from the upper buffer. The expiring FIFO1 channel allows this to happen only at some point in time  $t_0 < t$ . If this happens, an expire-signal is sent (via the  $P$ -producer from  $D$  to  $G$  where  $P$  is the singleton data set  $\{\text{expire}\}$ ) which forces the timeout-signal to become available at  $H$ . Because the buffer of the left FIFO1 channel is full and it is connected at  $E$  through a synchronous drain and a lossy synchronous channel via  $J$  to  $H$ , the availability of the timeout-signal at  $H$  triggers the synchronous transfer of the contents of the left FIFO1 channel into the right FIFO1. The replication behavior of  $H$  also attempts to simultaneously write a copy of the timeout-signal into the top lossy synchronous channel connected to  $H$ . However, because at this point in time (i.e.,  $t_0$ ), there is no data available at  $C$ , the synchronous drain connected to  $C$  prevents  $I$  from participating in the transfer of this copy of the timeout-signal from  $H$ ; therefore, the lossy synchronous channel connecting  $H$  to  $I$  loses this data. At this point, the same behavior symmetrically repeats with  $B$ .

If  $A$  has no value to transfer within the first  $t$  time units then  $D$  does not transfer the data element out the buffer but the timeout signal becomes available at  $H$  at time  $t$ . Simultaneously, the message in the buffer of the upper expiring FIFO1 channel is lost. At this point in time (i.e.,  $t$ ), there is no data available at  $C$ , and the synchronous drain connected to  $C$  prevents  $I$  from participating in the transfer of a copy of the timeout-signal from  $H$ ; the lossy synchronous channel connecting  $H$  to  $I$  loses this data.

On the other hand, node  $E$  can take the data element out of the buffer of the left FIFO1 channel. Also  $G$  is ready to start the timer again. Thus,  $H$  synchronizes with the nodes  $J$ ,  $E$  and  $G$  which yields a configuration symmetric to the initial one with  $B$  instead of  $A$ .



The above figure shows the TCA (before hiding) where we skip the data constraints.<sup>2</sup> □





Figure 7: When does  $B$  perform a take-operation?

**Remark 3.3 (Time-constraints for the I/O-operations)** In the Reo circuit in Fig. 7, node  $B$  is a mixed node which is “always” ready to consume a message from the buffer of the expiring FIFO1 channel

<sup>2</sup>In addition to the node-names used in the circuit, we use the names  $G_E$ ,  $G_C$ ,  $G_D$  and  $G_F$  to make clear which take-operation is performed on node  $G$ . Such auxiliary names will also be used in the compositional approach to model the merge semantics.

because the synchronous drain on its right is “always” ready to dispose of any value. The TCA for this circuit has a TSD stream of the form  $(\{A\}, [A \mapsto d], 0), (\{A\}, [A \mapsto d], 4), (\{A\}, [A \mapsto d], 8), \dots$  where  $A$  continuously transfers data items into the buffer of the expiring FIFO1 channel, which in turn loses them all because the data transfer at  $B$  takes longer than the specified expiration bound of 3 time units (e.g., because the synchronous drain is too slow). In fact, the above circuit makes no assumptions about the possible delay of  $B$ ’s data transfer operation. Its TCA involves an enabled transition with a node-set consisting of a mixed node with an unbounded delay.

One possibility to avoid such scenarios is to assign *deadlines* to edges  $e = (s, N, dc, cc, C, \bar{s})$  where  $N$  consists of mixed nodes. For instance, assigning a deadline of 2 to the  $\{B\}$ -edge in the above example ensures that all values transferred by  $A$  are eventually taken out of the buffer by  $B$ . However, the timing behavior of the nodes (deadlines or lower time bounds for I/O-operations) can also be made explicit at the syntax level of Reo circuits, using an appropriate combination of Reo’s timed channels. For instance, the deadline of 2 in the above example can be guaranteed by a 2-timer with the off-option as follows:



□

We now define the join operator on TCA which captures the replicator semantics of source (or mixed) nodes. It can serve as the semantic operator for the join of two nodes where at least one of them is a source node. We assume that we are given the TCA  $\mathcal{T}_1$  and  $\mathcal{T}_2$  for two fragments  $R_1$  and  $R_2$  of a Reo circuit and that we want to perform the join operations for the nodes  $B_i$  (in  $\mathcal{T}_1$ ) and  $\tilde{B}_i$  (in  $\mathcal{T}_2$ ),  $i = 1, \dots, n$ , where at least one of the nodes  $B_i$  or  $\tilde{B}_i$  is a source node (i.e., has no coincident sink channel end). We first rename  $\tilde{B}_i$  into  $B_i$  and then apply the following join operator to  $\mathcal{T}_1$  and  $\mathcal{T}_2$ .

**Definition 3.4 (Join for TCA)** *Given two TCA  $\mathcal{T}_i = (S_i, C_i, N_i, \mathcal{E}_i, S_{0,i}, ic_i)$ ,  $i = 1, 2$ , with disjoint clock sets, the product  $\mathcal{T}_1 \bowtie \mathcal{T}_2$  is defined as an TCA with the location space  $S = S_1 \times S_2$ , the set  $S_0 = S_{0,1} \times S_{0,2}$  of initial locations, the node-set  $\mathcal{N} = \mathcal{N}_1 \cup \mathcal{N}_2$ , and the clock set  $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2$ . The location invariance is given by  $ic(\langle s_1, s_2 \rangle) = ic_1(s_1) \wedge ic_2(s_2)$ . The edge relation  $\mathcal{E}$  is obtained through the following rules. The first rule concerns the “synchronization case” where two edges with common nodes are combined as well as the case where two edges with non-empty “local” node-sets are taken simultaneously:*

$$\frac{\begin{array}{l} (s_1, N_1, dc_1, cc_1, C_1, \bar{s}_1) \in \mathcal{E}_1, \\ (s_2, N_2, dc_2, cc_2, C_2, \bar{s}_2) \in \mathcal{E}_2, \\ N_1 \cap N_2 = N_2 \cap N_1, N_1 \neq \emptyset, N_2 \neq \emptyset, dc_1 \wedge dc_2 \neq \text{false} \end{array}}{\langle \langle s_1, s_2 \rangle, N_1 \cup N_2, dc_1 \wedge dc_2, cc_1 \wedge cc_2, C_1 \cup C_2, \langle \bar{s}_1, \bar{s}_2 \rangle \rangle \in \mathcal{E}}$$


The second rule applies to edges all of whose involved nodes are local to only one of the automata:

$$\frac{(s_1, N_1, dc_1, cc_2, C_1, \bar{s}_1) \in \mathcal{E}_1, N_1 \cap N_2 = \emptyset}{\langle \langle s_1, s_2 \rangle, N_1, dc_1, cc_1, C_1, \langle \bar{s}_1, s_2 \rangle \rangle \in \mathcal{E}}$$

and its symmetric rule. In particular, the latter rule applies to transitions with empty node-sets. □

A correctness result for the join operator is presented in the appendix (Section A).

To mimic the merge semantics of sink (or mixed) nodes we use the same technique as in [8, 7]. To join two nodes  $A$  and  $B$  where each of them contains at least one sink end we (1) choose a new node-name, say  $C$ , and (2) return  $\mathcal{T}_{Merger}(A, B, C) \bowtie \mathcal{T}_A \bowtie \mathcal{T}_B$  where  $\mathcal{T}_A$  and  $\mathcal{T}_B$  are the TCA that model the sub-circuits containing  $A$  and  $B$  respectively, and  $\mathcal{T}_{Merger}(A, B, C)$  has the following form:



Hiding a node-set  $M$  in a TCA removes all  $M$ -nodes from its edges. However, given an edge with a node-set consisting of  $M$ -nodes only, we must ensure that this edge can be taken only after some positive delay. We model this by using an additional clock.


**Definition 3.5 (Hiding for TCA)** Given a TCA  $\mathcal{T} = (S, \mathcal{C}, \mathcal{N}, \mathcal{E}, S_0, ic)$ , a new clock  $y \notin \mathcal{C}$ , and  $M \subseteq \mathcal{N}$ , we define  $\exists M[\mathcal{T}] = (S, \mathcal{C} \cup \{y\}, \mathcal{N} \setminus M, \mathcal{E}', S_0, ic)$  where  $\mathcal{E}'$  is obtained by the rule:

$$\frac{(s, N, dc, cc, C, \bar{s}) \in \mathcal{E}, (N = \emptyset \vee N \setminus M \neq \emptyset)}{(s, N \setminus M, \bigvee_{\delta \in DA(M)} dc[A/\delta_A : A \in M], cc, C \cup \{y\}, \bar{s}) \in \mathcal{E}'}$$

$$\frac{(s, N, dc, cc, C, \bar{s}) \in \mathcal{E}, \emptyset \neq N \subseteq M}{(s, \emptyset, \text{true}, cc \wedge (y > 0), C \cup \{y\}, \bar{s}) \in \mathcal{E}'}$$

Here,  $dc[A/\delta_A : A \in M]$  is derived from  $dc$  by the syntactic replacement of the term  $d_A$  with the value  $\delta_A \in \text{Data}$  for all  $A \in M$ . (More precisely, we replace “ $d_A \in P$ ” with **true** or **false**, depending on whether or not  $\delta_A$  belongs to  $P$ .)  $\square$

**Example 3.6** The TCA for the circuit in Fig. 5 can be obtained by joining the TCA for all of its involved channels together with  $T_{\text{Merger}}(F_1, F_2, F)$ .



The above figure shows the resulting TCA before and after hiding. (For simplicity, we skip the data constraints and irrelevant resettings of  $y$ .)  $\square$

Of course, using arbitrary combinations of timed channels can lead to TCA with time-locks. However, using (modifications of) standard region- or zone-graph algorithms [1, 18] we may check the time-lock freedom of a given Reo circuit.

## 4 Timed Scheduled-Data-Stream Logic

To specify the behavior of timed Reo circuits, one can use a TCA  $\mathcal{T}$  and require that the TSD-language generated by a given Reo circuit is contained in  $\mathcal{L}(\mathcal{T})$ . In this sense,  $\mathcal{T}$  specifies the “legal” behavior of the circuit. However, it is often easier to use a logical formalism to express the desired properties rather than using an automata model. In this section, we introduce Time Scheduled-Data-Stream Logic (TSDSL) which is a real-time variant of LTL and allows to reason about the observable data-flow of a Reo circuit by means of the TSD streams generated by its underlying TCA. Instead of the modality  $\bigcirc$  (next step),



TSDSL uses formulas of the type  $\langle \alpha \rangle \varphi$  which consist of a so-called timed scheduled-data expression  $\alpha$  and a formula  $\varphi$ . This type of formulas is inspired by propositional dynamic logic [12] and extended temporal logic [29]. The timed scheduled-data expressions are variants of timed regular expressions [9] built from atoms of the form  $\langle N, dc \rangle$ . The TSD expressions specify *sets of finite TSD streams*. The intuitive meaning of  $\langle \alpha \rangle \varphi$  is that every initial run has a finite prefix generating a word of the language of  $\alpha$  such that  $\varphi$  holds for its corresponding suffix.

**Syntax of TSDSL.** In the sequel, we assume a fixed finite and non-empty set  $\mathcal{N}$  of nodes. The abstract syntax of TSDSL-formulas is given by the following grammar:

$$\varphi ::= \text{true} \mid \varphi_1 \wedge \varphi_2 \mid \neg \varphi \mid \langle \alpha \rangle \varphi \mid \varphi_1 \cup \varphi_2$$

where  $\alpha$  is a timed scheduled-data expression (TSD expression) built by the grammar:

$$\alpha = \langle N, dc \rangle \mid \alpha_1 \vee \alpha_2 \mid \alpha_1 \wedge \alpha_2 \mid \alpha_1 ; \alpha_2 \mid \alpha^* \mid \alpha^I$$

Here,  $N$  is a non-empty node-set,  $dc$  a satisfiable data constraint for  $N$ , and  $I \subseteq \mathbb{R}_{\geq 0} \cup \{\omega\}$  a (possibly unbounded) time interval with its upper-bound in  $\mathbb{N} \cup \{\omega\}$ . The meanings of  $\alpha_1 \vee \alpha_2$  (union, choice),  $\alpha_1 \wedge \alpha_2$  (intersection)<sup>3</sup>,  $\alpha_1 ; \alpha_2$  (concatenation, sequential composition), and  $\alpha^*$  (Kleene closure, finitely many repetitions) are obvious.  $\alpha^I$  has the same meaning as  $\alpha$ , except for the additional requirement that the total execution time falls in the time interval  $I$ .

Intuitively,  $\langle \alpha \rangle \varphi$  holds for a TCA iff all its TSD streams have a finite prefix that generates an  $\alpha$ -stream and  $\varphi$  holds for its remaining suffix. The dual operator for  $\langle \alpha \rangle \varphi$  is  $\llbracket \alpha \rrbracket \varphi = \neg \langle \alpha \rangle \neg \varphi$  which holds for a TCA iff for each of its TSD streams  $\Theta$  and all prefixes of  $\Theta$  that generate an  $\alpha$ -word, the formula  $\varphi$  holds for the corresponding suffix of  $\Theta$ . Other boolean connectives, like disjunction  $\vee$  or implication  $\rightarrow$ , are derived in the usual way.

**Remark 4.1** We can also allow for  $\omega$ -regular TSD expressions that result from adding an  $\omega$ -operator. Although this increases expressiveness, we skip this option here. In contrast to the real-time extensions of LTL, as, e.g., in [16, 3, 2], TSDSL does not use time-constrained temporal modalities such as  $\mathbf{U}^{\leq t}$ . These can be added to TSDSL, but in the examples (see below) it turned out that the time-constraints in the TSD expressions are sufficient to formulate the relevant properties of Reo circuits.  $\square$

**Simplified notation.** We often skip the semicolon for the concatenation operator (i.e.,  $\alpha\beta$  stands short for  $\alpha;\beta$ ). We simply write  $\langle N \rangle$  for  $\langle N, \text{true} \rangle$  and often omit brackets: e.g.,  $\langle A, dc \rangle$  is short-hand for  $\langle \{A\}, dc \rangle$  and  $\langle N \rangle$  for  $\langle \langle N \rangle \rangle$ . We write  $\langle \dots A \dots \rangle$  to denote the disjunction of the expressions  $\langle N \rangle$  where  $N$  ranges over all subsets of  $\mathcal{N}$  that contain the node  $A$ .  $\langle \neg A \rangle$  stands for the disjunction of all expressions  $\langle N \rangle$  where  $N$  ranges over all non-empty node-sets that do not contain  $A$ .  $\langle \cdot \rangle$  denotes the disjunction of all atoms  $\langle N \rangle$  where  $N$  is an arbitrary non-empty node-set.  $\langle \cdot \rangle \varphi$  stands for  $\langle \langle \cdot \rangle \rangle \varphi$ . We also often skip  $\text{true}$  and write  $\langle \alpha \rangle$  for  $\langle \alpha \rangle \text{true}$ : e.g., the TCA for the normal FIFO1 channel (Fig. 4) satisfies the formula

$$\llbracket (\langle A \rangle \langle B \rangle)^* \rrbracket \langle A \rangle \wedge \llbracket (\langle A \rangle \langle B \rangle)^* \langle A \rangle \rrbracket \langle B \rangle$$

which states that the data-flows at nodes  $A$  and  $B$  alternate, starting with  $A$ .

---

<sup>3</sup>Standard regular expressions do not contain an intersection operator (although regular languages are closed under intersection). However, as pointed out in [9], in timed settings, the class of timed languages induced by timed regular expressions without an explicit intersection operator is not closed under intersection.

**Derived operators.** The standard *next step* operator is derived as  $\circ\varphi = \langle \cdot \rangle \varphi$ . In particular,  $\circ\text{true}$  asserts the occurrence of some observable data-flow, while  $\neg\circ\text{true}$  states that data-flow has stopped. The modalities *eventually* and *always* can be derived as usual by definitions  $\diamond\varphi = \text{true}\mathbf{U}\varphi$  and  $\square\varphi = \neg\diamond\neg\varphi$ . For instance, the following TSDSL formula specifies the behavior of a normal FIFO1 channel (cf. Fig. 4):

$$\square\left(\bigwedge_{d \in \text{Data}} \llbracket \langle A, d_A = d \rangle \rrbracket \langle \langle B, d_B = d \rangle \rangle \right) \wedge \square(\langle B \rangle \rightarrow \circ\langle A \rangle)$$

The expiring FIFO1 channel in Fig. 4 satisfies the TSDSL formula

$$\square\left(\bigwedge_{d \in \text{Data}} \llbracket \langle A, d_A = d \rangle \rrbracket (\langle \langle B, d_B = d \rangle \rangle^{<t} \vee \neg\langle \cdot \rangle^{<t})\right)$$

which expresses the fact that within  $t$  time units after  $A$ 's write-operation either  $B$  takes the element from the buffer or there is no observable data-flow. For the timed sequencer (Fig. 2 and Example 3.2) the following formula holds

$$\square\llbracket A \rrbracket (\langle \langle B \rangle \rangle^{\leq t} \vee \neg\langle \cdot \rangle^{\leq t})$$

stating that whenever data-flow is observed at  $A$ , within the next  $t$  time units there is either data-flow at  $B$  or no observable data-flow at all.

The weak variant  $\tilde{\mathbf{U}}$  of until is obtained as  $\varphi_1 \tilde{\mathbf{U}} \varphi_2 = (\varphi_1 \mathbf{U} \varphi_2) \vee (\square\varphi_1)$ . For instance, the  $t$ -timer with reset-option (but without the off-option) fulfills the formula

$$\square\llbracket A \rrbracket (\langle \langle A, d_A = \text{reset} \rangle \rangle^{<t} \tilde{\mathbf{U}} \langle \langle B, d_B = \text{timeout} \rangle \rangle).$$

To provide the formal definition of the semantics of a TSD expressions and TSDSL-formulas we need some additional notation for working with TSD streams.

**Notation 4.2 (Time cuts, concatenation, Kleene closure)** Let  $\Theta = (N_0, \delta_0, t_0), (N_1, \delta_1, t_1), \dots$  be a TSD stream as in Notation 2.7. For a point in time  $t \in \mathbb{R}_{\geq 0}$ , we define  $\Theta \uparrow t$  as the suffix of  $\Theta$  that ignores every data-flow that occurs before  $t$  and formalizes the observable behavior in the time interval  $[t, \infty[$ . That is,  $\Theta \uparrow t = \varepsilon$  if  $|\Theta| = k + 1 < \omega$  and  $t_k < t$ . Otherwise,  $\Theta \uparrow t = (N_k, \delta_k, t_k), \dots$  where  $k$  is the smallest index such that  $t_k \geq t$ .

$\Theta \downarrow t$  is the TSD stream that describes the data-flow in the time interval  $[0, t[$ . That is,  $\Theta \downarrow t = \varepsilon$  if  $\Theta = \varepsilon$  or  $t_0 \geq t$ . Otherwise,  $\Theta \downarrow t = (N_0, \delta_0, t_0), \dots, (N_k, \delta_k, t_k)$  where  $k$  is the largest index such that  $t_k < t$ .

The concatenation of finite TSD streams is defined as follows. We define  $\Theta; \varepsilon = \varepsilon; \Theta = \Theta$ . If  $\Theta_1 = (N_0, \delta_0, t_0), \dots, (N_n, \delta_n, t_n)$  and  $\Theta_2 = (M_0, \sigma_0, \rho_0), \dots, (M_m, \sigma_m, \rho_m)$  then  $\Theta_1; \Theta_2$  is  $(N_0, \delta_0, t_0), \dots, (N_n, \delta_n, t_n), (M_0, \sigma_0, t_n + \rho_0), \dots, (M_m, \sigma_m, t_n + \rho_m)$ . If  $L$  and  $\tilde{L}$  are TSDS-languages with the same node-set  $\mathcal{N}$  then  $L; \tilde{L} = \{\Theta; \tilde{\Theta} : \Theta \in L, \tilde{\Theta} \in \tilde{L}\}$  and  $L^* = \bigcup_{n \geq 0} L^n$  where  $L^0 = \{\varepsilon\}$ ,  $L^{n+1} = L^n; L$ .  $\square$

**Semantics of TSD expressions and TSDSL-formulas.** We define  $\mathcal{L}(\alpha) \subseteq \text{TSDS}$  by structural induction.  $\mathcal{L}(\langle N, dc \rangle)$  is the set of all TSD streams of length 1 that have the form  $(N, \delta, t)$  where  $\delta \models dc$ . We define  $\mathcal{L}(\alpha_1 \vee \alpha_2) = \mathcal{L}(\alpha_1) \cup \mathcal{L}(\alpha_2)$ ,  $\mathcal{L}(\alpha_1 \wedge \alpha_2) = \mathcal{L}(\alpha_1) \cap \mathcal{L}(\alpha_2)$ ,  $\mathcal{L}(\alpha_1; \alpha_2) = \mathcal{L}(\alpha_1); \mathcal{L}(\alpha_2)$  and  $\mathcal{L}(\alpha^*) = \mathcal{L}(\alpha)^*$ . The semantics of time-constrained expressions is formalized by  $\mathcal{L}(\alpha^I) = \{\Theta \in \mathcal{L}(\alpha) : \tau(\Theta) \in I\}$ .<sup>4</sup> The satisfaction relation  $\models$  for TSDSL-formulas and TSD streams is defined by structural induction as shown in Fig. 8. For the derived  $\llbracket \cdot \rrbracket$ -operator, we obtain  $\Theta \models \llbracket \alpha \rrbracket \varphi$  iff for all  $t \geq 0$  we have:  $\Theta \downarrow t \in \mathcal{L}(\alpha)$  implies  $\Theta \uparrow t \models \varphi$ . We define  $\mathcal{L}(\varphi) = \{\Theta \in \text{TSDS}(\mathcal{N}) : \Theta \models \varphi\}$  and define logical equivalence  $\equiv$  of TSDSL-formulas as  $\varphi_1 \equiv \varphi_2$  iff  $\mathcal{L}(\varphi_1) = \mathcal{L}(\varphi_2)$ . If  $\mathcal{T}$  is a TCA and  $q$  a state in  $\mathcal{A}_{\mathcal{T}}$  then  $q \models \varphi$  iff  $\mathcal{L}(\mathcal{T}, q) \subseteq \mathcal{L}(\varphi)$ . Moreover, we define  $\mathcal{T} \models \varphi$  iff  $\mathcal{L}(\mathcal{T}) \subseteq \mathcal{L}(\varphi)$ .

<sup>4</sup>Recall that  $\tau(\Theta)$  denotes the execution time of  $\Theta$  (see Notation 2.7).

$\Theta \models \text{true}$	
$\Theta \models \varphi_1 \wedge \varphi_2$	iff $\Theta \models \varphi_1$ and $\Theta \models \varphi_2$
$\Theta \models \neg\varphi$	iff $\Theta \not\models \varphi$
$\Theta \models \varphi_1 \mathbf{U} \varphi_2$	iff $\exists t \in \mathbb{R}_{\geq 0}$ s.t. $\Theta \uparrow t \models \varphi_2$ and $\Theta \uparrow \rho \models \varphi_1$ for all $\rho$ with $0 \leq \rho < t$
$\Theta \models \langle \alpha \rangle \varphi$	iff $\exists t \in \mathbb{R}_{\geq 0}$ s.t. $\Theta \downarrow t \in \mathcal{L}(\alpha) \wedge \Theta \uparrow t \models \varphi$

Figure 8: Satisfaction relation for TSDSL-formulas

**Example 4.3 (Alternating bit protocol)** The properties of the ABP (see Example 2.4 and Fig. 9) can be specified by the formula

$$\varphi_{ABP}(t) = \bigwedge_{d \in \text{Data}} \square[\langle I, d_I = d \rangle] \langle \langle \neg I \rangle^* \langle O, d_O = d \rangle \rangle^{\leq t}$$

for some time bound  $t$ .  $\varphi_{ABP}(t)$  states that whenever the sender receives a message  $d$  at port  $I$ , within its next  $t$  time units the receiver will output  $d$  at port  $O$  during which time the sender does not accept a new input message through port  $I$ .<sup>5</sup>

For an arbitrary choice of the time-parameters  $t_S, t_R, \rho_S$  and  $\rho_R$  we cannot expect that  $\mathcal{T}_{ABP} \models \varphi_{ABP}(t)$ . For instance, if  $\rho_R = 5$  and  $t_R = t_S = 2$  then the following behavior is possible. The first input at  $I$  arrives at time instant 2.5. The receiver may move to location  $ack(1)$  earlier, say at time instant 1.5. With the take-operation at input port  $I$ , the automaton moves from state  $\langle in(0), ack(1), x = 2.5, y = 2 \rangle$  to  $\langle try(d, 0), ack(1), x = 0, y = 2 \rangle$ . At time 3.5, the sender tries to send  $(d, 0)$  through port  $A$  and moves to location  $wait(d, 0)$ . Now, clock  $x$  has the value 0, clock  $y$  the value 2. After 1 time unit the receiver sends the control bit  $b = 1$  which the sender ignores. Thus, we enter the global state  $q = \langle wait(d, 0), wait(0), x = 1, y = 0 \rangle$ . The sender is forced to move to location  $try(d, 0)$ . One time unit later, clock  $y$  has the value 2 and forces the receiver to leave location  $wait(0)$ . We enter now the global state  $\langle try(d, 0), ack(1), x = 1, y = 0 \rangle$ . After waiting for 1 time unit, the sender resends the pair  $(d, 0)$  which leads to the global state  $\langle wait(d, 0), ack(1), x = 0, y = 1 \rangle$ . One time unit later, the receiver resends the control bit 1 which the sender ignores again. We now reenter state  $q$  and may continue in the same way, without ever producing an output at port  $O$ . Hence, for this choice of the time-parameter we obtain  $\mathcal{T}_{ABP} \not\models \square(\langle I \rangle \rightarrow \diamond \langle O \rangle)$ . In particular, there is no  $t$  such that  $\varphi_{ABP}(t)$  holds for  $\mathcal{T}_{ABP}$ .

Assuming  $\rho_R < \rho_S < t_R$  and  $\rho_R < t_S$  then no message sent via the lossy channel connecting  $A$  and  $C$  will be lost. In fact, it can only happen that the receiver acknowledges more than once the receipt of the last message (because no upper time bound is assumed for the arrival of messages at input port  $I$ ). The reachable fragment of the TCA is shown in Fig. 10. We obtain  $\mathcal{T}_{ABP} \models \varphi_{ABP}(\rho_S + \rho_R)$ , stating that the delay for the output at  $O$  is bounded above by the maximal sojourn time of the sender in location  $wait(d, b)$  plus the maximal delay  $\rho_R$  for the receiver to send the acknowledgment after it receives a message through port  $C$ . (This is the best bound we can expect.) The fact that messages along the  $A$ - $C$  channel are never lost can be formalized by the TSDSL formula  $\neg \diamond \langle A \rangle$  which states that it is not possible to observe a data-flow at node  $A$  only (not together with  $C$ ).


When  $\rho_R < t_S < t_R$  and  $\rho_S < t_R - t_S$ , messages sent from  $A$  to  $C$  may get lost. However, when  $A$  resends the message the receiver accepts the message through port  $C$ . In this case, we have  $\mathcal{T}_{ABP} \models \varphi_{ABP}(\rho_S + \rho_R + t_S)$ , stating that the delay for the output at port  $O$  is at most the maximal delay for the sender and receiver to send their messages along the lossy channels connecting them plus the deadline  $t_S$  which the sender uses for resending message-bit pairs. The reachable part of the TCA under these assumptions is shown in Fig. 11. The property that a message sent along the  $A$ - $C$  channel can be lost only once can be formalized by the TSDSL formula  $\neg \diamond \langle \langle A \rangle \langle \neg I \rangle^* \langle A \rangle \rangle$ .  $\square$

<sup>5</sup>As input on  $I$  can occur simultaneously with the receiver resending its acknowledgment of the previous message via port  $D$ , the atom  $\langle I, d_I = d \rangle$  can be replaced with the expression  $\langle I, d_I = d \rangle \vee \langle \{I, D\}, d_I = d \rangle$ .

**The TSDSL Model Checking problem** addresses the question of whether  $\mathcal{T} \models \varphi$  holds for a given TCA  $\mathcal{T}$  and TSDSL formula  $\varphi$ . We briefly sketch the main ideas of a TSDSL model checking algorithm that relies on variants of standard automata-based algorithms for LTL and (timed) regular expressions.


First, we switch from  $\varphi$  to  $\neg\varphi$  which we regard as a formula of (untimed) LTL with action labels. Here,  $\langle\alpha\rangle$  is treated as a next step operator with the label  $\alpha$ . Then, we may apply standard techniques, e.g., [30, 15, 27, 13], modified for the action-labeled case, to construct a nondeterministic Büchi automaton  $\mathcal{B}$  for  $\neg\varphi$ , whose transitions are labeled with the expressions  $\alpha$  that occur in sub-formulas  $\langle\alpha\rangle\psi$  of  $\varphi$ . We now turn  $\mathcal{B}$  into a TCA  $\mathcal{T}_{\mathcal{B}}$  with Büchi acceptance condition. (See Appendix A.)

For this, we first construct a TCA  $\mathcal{T}_{\alpha}$  for every TSD expression  $\alpha$  that occurs in  $\mathcal{B}$  as a transition-label.  $\mathcal{T}_{\alpha}$  has a unique initial location, called  $start(\alpha)$ , and a location  $stop(\alpha)$  such that  $\mathcal{L}(\alpha)$  is the set of all TSD streams  $\Theta$  that are induced by a finite run in  $\mathcal{T}_{\alpha}$  starting in  $start(\alpha)$  and ending in  $stop(\alpha)$ . The construction of the TCA  $\mathcal{T}_{\alpha}$  is by structural induction, essentially as described in [9]. For instance, for  $\alpha = \gamma^I$  we introduce one new clock  $x$  that is not used in  $\mathcal{T}_{\gamma}$  and perform the following construction for  $\mathcal{T}_{\alpha}$ :



The invariance condition “ $x \in I$ ” ensures that location  $stop(\alpha)$  can be entered only in runs where the execution time lies within the time interval  $I$ . (Here, the edges from  $stop(\gamma)$  to  $stop(\alpha)$  are labeled with the empty node-set and data and clock constraint true.)

The TCA  $\mathcal{T}_{\mathcal{B}}$  is now obtained as follows. The locations in  $\mathcal{T}_{\mathcal{B}}$  consist of the states in the Büchi automaton  $\mathcal{B}$  and the locations in the TCA  $\mathcal{T}_{\alpha}$ .<sup>6</sup> We then replace every transition  $q \xrightarrow{\alpha} p$  in  $\mathcal{B}$  with the following fragment of  $\mathcal{T}_{\mathcal{B}}$ :



We then have  $\mathcal{L}(\mathcal{T}_{\mathcal{B}}) = \mathcal{L}(\neg\varphi)$  where Büchi acceptance is assumed for  $\mathcal{T}_{\mathcal{B}}$ . Thus, by Corollary A.7,  $\mathcal{T} \models \varphi$  iff  $\mathcal{L}(\mathcal{T} \bowtie \mathcal{T}_{\mathcal{B}}) = \mathcal{L}(\mathcal{T}) \cap \mathcal{L}(\varphi) = \emptyset$ . Hence, we may apply (modifications of) the standard region graph algorithms to check for emptiness of timed automata [1].

**TSDSL versus refinement relations.** Let  $\mathcal{T}_1$  and  $\mathcal{T}_2$  be two TCA with the same node-set  $\mathcal{N}$ . Clearly, if  $\mathcal{L}(\mathcal{T}_1) \subseteq \mathcal{L}(\mathcal{T}_2)$  then, for any TSDSL-formula  $\varphi$ ,  $\mathcal{T}_2 \models \varphi$  implies  $\mathcal{T}_1 \models \varphi$ . Thus, if  $\mathcal{L}(\mathcal{T}_1) = \mathcal{L}(\mathcal{T}_2)$  then  $\mathcal{T}_1$  and  $\mathcal{T}_2$  satisfy exactly the same TSDSL-formulas. A sufficient decidable criterion for checking (TSDLS- or) language-equivalence of two TCA is to switch to a coarser equivalence corresponding to timed bisimulation for ordinary timed automata [10]. In our setting, a timed bisimulation for a TCA  $\mathcal{T}$  is the coarsest equivalence  $\sim$  on the state space  $Q$  of the induced state-transition graph  $\mathcal{A}_{\mathcal{T}}$  such that for all  $q_1, q_2 \in Q$  with  $q_1 \sim q_2$  and all  $N \subseteq \mathcal{N}$ ,  $\delta \in DA(N)$ ,  $t \in \mathbb{R}_{\geq 0}$ :

$$\forall q_1 \xrightarrow{N, \delta, t} p_1 \exists p_2 \in Q \text{ s.t. } q_1 \xrightarrow{N, \delta, t} p_2 \text{ and } p_1 \sim p_2.$$

The simulation relation is defined as the coarsest binary relation  $\preceq$  on the state space  $Q$  of  $\mathcal{A}_{\mathcal{T}}$  such that for all  $q_1, q_2 \in Q$  with  $q_1 \preceq q_2$  and all  $N \subseteq \mathcal{N}$ ,  $\delta \in DA(N)$ ,  $t \in \mathbb{R}_{\geq 0}$ :

<sup>6</sup>We assume that the state spaces and clock sets are disjoint and that for any TSD expression  $\alpha$  that occurs more than once in  $\mathcal{B}$  a copy of  $\mathcal{T}_{\alpha}$  is used.

$$\forall q_1 \xrightarrow{N,\delta,t} p_1 \exists p_2 \in Q \text{ s.t. } q_1 \xrightarrow{N,\delta,t} p_2 \text{ and } p_1 \preceq p_2.$$

The relation  $\preceq$  is finer than language-inclusion, and thus, preserves all TSDSL formulas in the sense that if  $q_1 \preceq q_2$  and  $q_2 \models \varphi$  then  $q_1 \models \varphi$ . The question of whether one state of a TCA simulates another one can be answered with the help of the region graph construction as in [28].

## 5 Conclusion

In this paper, we introduced a formal model to reason about timing constraints for Reo component connectors. We presented composition operators for join and hiding that can serve as a basis for the automated construction of an automata-model from a given (timed) Reo circuit and as a starting point for its formal verification. In particular, (slightly modified versions of) well-known algorithms for checking time-lock freedom in ordinary timed automata can serve for checking the realizability of the coordination mechanisms of a Reo circuit with timing constraints. Moreover, we suggested a linear-time temporal logic for reasoning about the real-time behavior of component connectors by means of their timed scheduled-data streams and explained how the standard region- or zone-graphs model checking algorithms for timed automata can be adapted for our setting.

Our future work includes an implementation of the presented model checking algorithms and case studies. Moreover, we intend to study an alternating-time logic in the style of [4] that allows to reason about the possibility for certain components to cooperate such that a given (real-time) property holds.

## References

- [1] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- [2] R. Alur, T. Feder, and T. Henzinger. The benefits of relaxing punctuality. *Journal of the ACM*, 43(1):116–146, 1996.
- [3] R. Alur and T. A. Henzinger. A really temporal logic. *Journal of the ACM*, 41:181–204, 1994.
- [4] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
- [5] F. Arbab. Reo: A channel-based coordination model for component composition. *Mathematical Structures in Computer Science*, 14(3):1–38, 2004.
- [6] F. Arbab, C. Baier, F. de Boer, J.J.M.M. Rutten, and M. Sirjani. Modeling context-sensitive behaviors of component connectors with priorities. Forthcoming paper, 2004.
- [7] F. Arbab, C. Baier, J.J.M.M. Rutten, and M. Sirjani. Modeling component connectors in reo by constraint automata. In *FOCLASA'03*, Electronic Notes in Theoretical Computer Science, 2003. To appear. For the full version see <http://web.informatik.uni-bonn.de/I/baier/publikationen.html>.
- [8] F. Arbab and J.J.M.M. Rutten. A coinductive calculus of component connectors. In D. Paterson, M. Wirsing and R. Hennicker, editors, *Recent Trends in Algebraic Development Techniques, Proceedings of 16th International Workshop on Algebraic Development Techniques (WADT 2002)*, volume 2755 of *Lecture Notes in Computer Science*, pages 35–56. Springer-Verlag, 2003. <http://www.cwi.nl/ftp/CWIreports/SEN/SEN-R0216.pdf>.
- [9] E. Asarin, P. Caspi, and O. Maler. Timed regular expressions. *Journal of the ACM*, 49(2):172–206, 2002.
- [10] K. Cerans. Decidability of bisimulation equivalences for parallel timer processes. In *Proc. CAV*, volume 663 of *LNCS*, pages 302–315, 1993.

- [11] L. de Alfaro, T. A. Henzinger, and M. Stoelinga. Timed interfaces. In *Proc. EMSOFT*, volume 2491 of *LNCS*, pages 108–122, 2002.
- [12] M. J. Fischer and R.J. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Science*, 8:194–211, 1979.
- [13] P. Gastin and D. Oddoux. Fast LTL to Büchi automata translation. In *Proc. 13th International Conference on Computer Aided Verification (CAV)*, volume 2102 of *Lecture Notes in Computer Science*, pages 53–65, 2001.
- [14] R. Gawlick, R. Segala, J. Soegaard-Andersen, and N. Lynch. Liveness in timed and untimed systems. *Information and Computation*, 141(2):119–171, 1998.
- [15] R. Gerth, D. Peled, M. Vardi, and P. Wolper. Simple on-the-fly automatic verification of linear temporal logic. In *Protocol Specification Testing and Verification*, pages 3–18. Chapman & Hall, 1995.
- [16] E. Harel, O. Lichtenstein, and A. Pnueli. Explicit clock temporal logic. In *Proc. LICS*, pages 402–413. IEEE Computer Society Press, 1990.
- [17] T.A. Henzinger, P.-H. Ho, and H. Wong-Toi. Hytech: A model checker for hybrid systems. *Software Tools for Technology Transfer*, 1:110–122, 1997.
- [18] T.A. Henzinger, X. Nicollin, J. Sifakis, and S. Yovine. Symbolic Model Checking for Real-Time Systems. *Information and Computation*, 111(2):193–244, 1994.
- [19] D.K. Kaynar, N.A. Lynch, R. Segala, and F.W. Vaandrager. A framework for modelling timed systems with restricted hybrid automata. In *Proceedings 24th IEEE International Real-Time Systems Symposium (RTSS'03)*, pages 166–177. IEEE Computer Society, 2003.
- [20] K. Larsen, P. Pettersson, and W. Yi. UPPAAL in a nutshell. *International Journal on Software Tools for Technology Transfer*, 1(1-2):134–152, 1997.
- [21] L. Leonard and G. Leduc. An enhanced version of timed lotos and its application to a case study. In *Proc. Formal Description Techniques VI*, pages 483–498. North-Holland, Amsterdam, 1994.
- [22] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer, New York, 1992.
- [23] M. Merritt, F. Modugno, and M. R. Tuttle. Time-constrained automata (extended abstract). In *Proc. CONCUR*, volume 527 of *LNCS*, pages 408–423, 1991.
- [24] R. Milner. *Communication and Concurrency*. Prentice Hall International Series in Computer Science. Prentice Hall, 1989.
- [25] A. Pnueli. The temporal logic of programs. In *Proc. FOCS*, pages 46–57. IEEE Computer Society Press, 1977.
- [26] G. M. Reed and A. W. Roscoe. A timed model for communication sequential processes. *Theoretical Computer Science*, 58:249–261, 1988.
- [27] F. Somenzi and R. Bloem. Efficient Büchi automata from LTL formulae. In *Proc. 12th International Conference on Computer Aided Verification (CAV)*, volume 1855 of *Lecture Notes in Computer Science*, 2000.
- [28] S. Tasiran, R. Alur, R. Kurshan, and R. Brayton. Verifying abstractions of timed systems. In *Proc. CONCUR*, volume 1119 of *LNCS*, pages 546–562, 1996.
- [29] P. Wolper. Specification and synthesis of communicating processes using an extended temporal logic. In *Proc. POPL*, pages 20–33, 1982.
- [30] P. Wolper, M. Vardi, and A. Sistla. Reasoning about infinite computation paths. In *Proc. FOCS*, pages 185–194. IEEE Computer Society Press, 1983.

- [31] W. Yi. CCS + time = an interleaving model for real time systems. In *Proc. ICALP*, volume 510 of *LNCS*, pages 217–228. Springer-Verlag, 1991.
- [32] S. Yovine. Kronos: A verification tool for real-time systems. *Software Tools for Technology Transfer*, pages 123–133, 1997.

## A Join, hiding and Büchi acceptance

**Notation A.1 (Join and hiding for TSD streams)** Let  $\Theta$  be a TSD stream over  $\mathcal{N}$  and  $B \in \mathcal{N}$ . The projection  $\Theta|_B \in (Data \times \mathbb{R}_{\geq 0})^\infty$  of  $\Theta$  on  $B$  denotes the sequence of pairs  $(d, t) \in Data \times \mathbb{R}_{\geq 0}$  that is obtained from  $\Theta$  by (1) removing all triples  $(N, \delta, t)$  where  $B \notin N$ ; and (2) replacing any remaining triples  $(N, \delta, t)$  with the pair  $(\delta_B, t)$ .

If  $M \subseteq \mathcal{N}$  then  $hide(\Theta, M)$  denotes the unique TSD stream  $\bar{\Theta} \in TSDS(M)$  such that  $\bar{\Theta}|_B = \Theta|_B$  for all  $B \in M$ . Given two TSD streams  $\Theta_1 \in TSDS(\mathcal{N}_1)$  and  $\Theta_2 \in TSDS(\mathcal{N}_2)$ , their join is undefined if there is a node  $B \in \mathcal{N}_1 \cap \mathcal{N}_2$  such that  $\Theta_1|_B \neq \Theta_2|_B$ . Otherwise we define their join  $\Theta_1 \bowtie \Theta_2 \in TSDS(\mathcal{N}_1 \cup \mathcal{N}_2)$  as the unique TSD stream such that  $(\Theta_1 \bowtie \Theta_2)|_A = \Theta_i|_A$  if  $A \in \mathcal{N}_i$ .  $\square$

**Notation A.2 (Join and hiding for TSDS-languages)** Given two TSDS-languages  $L_1 \subseteq TSDS(\mathcal{N}_1)$  and  $L_2 \subseteq TSDS(\mathcal{N}_2)$ , their join  $L_1 \bowtie L_2 \subseteq TSDS(\mathcal{N}_1 \cup \mathcal{N}_2)$  consists of all TSD streams  $\Theta$  that can be obtained by joining the TSD streams  $\Theta_1 \in L_1$  and  $\Theta_2 \in L_2$ . If  $M \subseteq \mathcal{N}$  and  $L \subseteq TSDS(\mathcal{N})$  then  $\exists M[L] = \{hide(\Theta, M) : \Theta \in L\}$ .  $\square$

The following lemma can be proved using similar arguments as in the untimed case (cf. [7]):

**Lemma A.3** *Let  $\mathcal{T}$ ,  $\mathcal{T}_1$  and  $\mathcal{T}_2$  be TCA. Then,*

- (a)  $\mathcal{L}(\mathcal{T}_1 \bowtie \mathcal{T}_2) = \mathcal{L}(\mathcal{T}_1) \bowtie \mathcal{L}(\mathcal{T}_2)$ .
- (b)  $\mathcal{L}(\exists M[\mathcal{T}]) = \exists M[\mathcal{L}(\mathcal{T})]$

The join of TSDS-languages with the same node-set agrees with their intersection. Thus, we obtain:

**Corollary A.4** *If  $\mathcal{T}_1$  and  $\mathcal{T}_2$  are TCA with the same node-set then  $\mathcal{L}(\mathcal{T}_1 \bowtie \mathcal{T}_2) = \mathcal{L}(\mathcal{T}_1) \cap \mathcal{L}(\mathcal{T}_2)$ .*

**Definition A.5 (TCA with Büchi acceptance)** *A Büchi TCA denotes a pair  $\mathcal{F} = (\mathcal{T}, S_{acc})$  consisting of a TCA  $\mathcal{T} = (S, Q, \mathcal{N}, \mathcal{E}, S_0, ic)$  and a set  $S_{acc} \subseteq S$  of accepting locations. A  $q$ -run in  $\mathcal{T}$  is called accepting iff it is either finite and ends in an accepting location, or visits infinitely often an accepting location. The language  $\mathcal{L}(\mathcal{F})$  denotes the set of TSD streams that can be generated by an accepting maximal run.  $\square$*

Note that for any TCA  $\mathcal{T}$  we have  $\mathcal{L}(\mathcal{T}) = \mathcal{L}(\mathcal{F}_{\mathcal{T}})$  where  $\mathcal{F}_{\mathcal{T}}$  is the Büchi TCA that results from  $\mathcal{T}$  by declaring all locations to be accepting.

The join of two Büchi TCAs  $\mathcal{F}_1$  and  $\mathcal{F}_2$  with disjoint clock-sets is defined as the standard join operator (Def. 3.4) where the accepting locations  $\langle s_1, s_2 \rangle$  in  $\mathcal{F}_1 \bowtie \mathcal{F}_2$  are those such that location  $s_1$  is accepting in  $\mathcal{F}_1$  and location  $s_2$  is accepting in  $\mathcal{F}_2$ . The hiding operator  $\exists M[\mathcal{F}]$  for Büchi TCA relies on the hiding operator for TCA (Def. 3.5) and does not change the accepting locations. We then have:

**Lemma A.6**  $\mathcal{L}(\mathcal{F}_1 \bowtie \mathcal{F}_2) = \mathcal{L}(\mathcal{F}_1) \bowtie \mathcal{L}(\mathcal{F}_2)$  and  $\mathcal{L}(\exists M[\mathcal{F}]) = \exists M[\mathcal{L}(\mathcal{F})]$

For Büchi TCA with the same node-set we obtain:

**Corollary A.7** *If  $\mathcal{F}_1$  and  $\mathcal{F}_2$  are two Büchi TCA with the same node-set then  $\mathcal{L}(\mathcal{F}_1 \bowtie \mathcal{F}_2) = \mathcal{L}(\mathcal{F}_1) \cap \mathcal{L}(\mathcal{F}_2)$ .*



## B TCA for the alternating bit protocol

Fig. 9 shows the full TCA  $\mathcal{T}_{ABP}$  that is obtained by joining the automata for the sender and the receiver. Figs. 10 and 11 show the relevant fragments for the assumptions  $(\rho_R < \min\{\rho_S, t_S\}) \wedge (\rho_S < t_R)$  and  $(\rho_R < t_S < t_R) \wedge (\rho_S < t_R - t_S)$ , respectively.

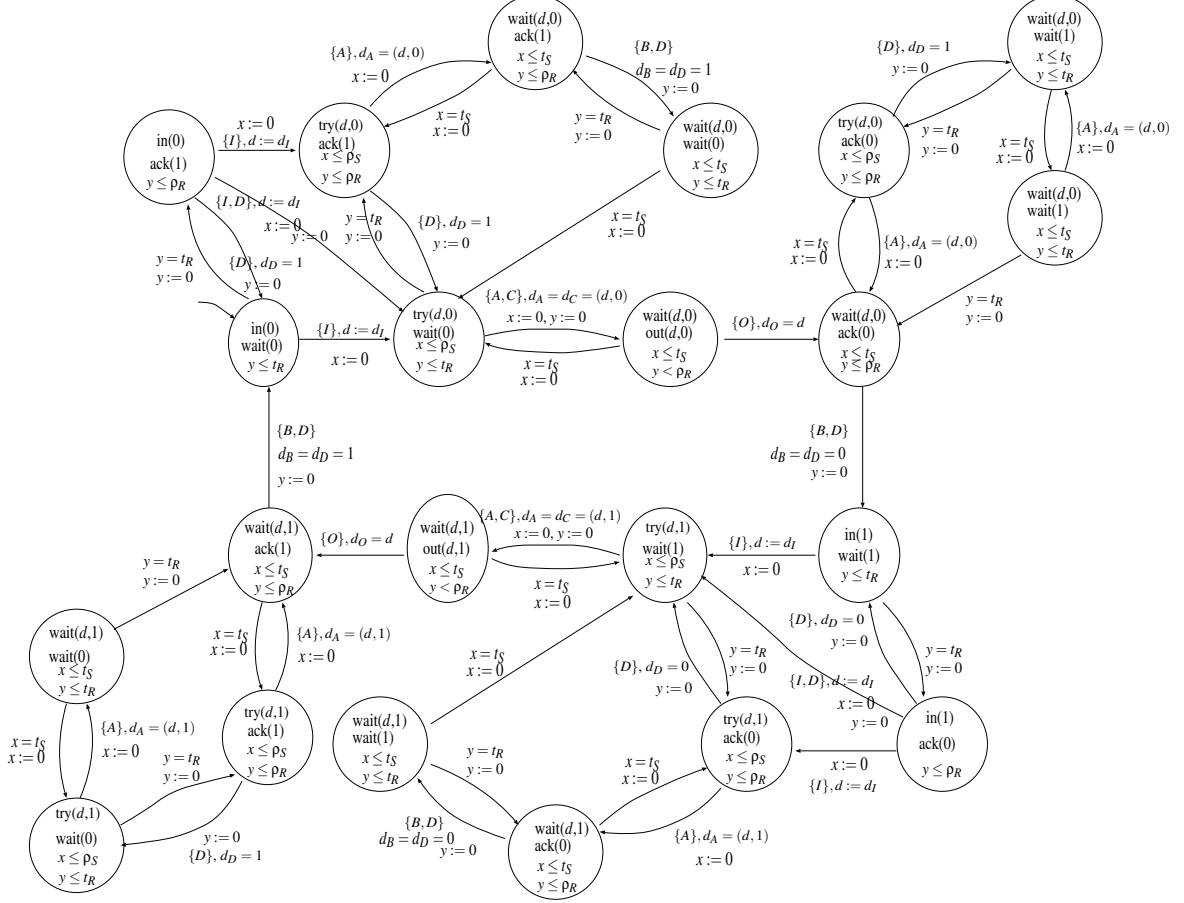


Figure 9: TCA  $\mathcal{T}_{ABP}$  for the ABP

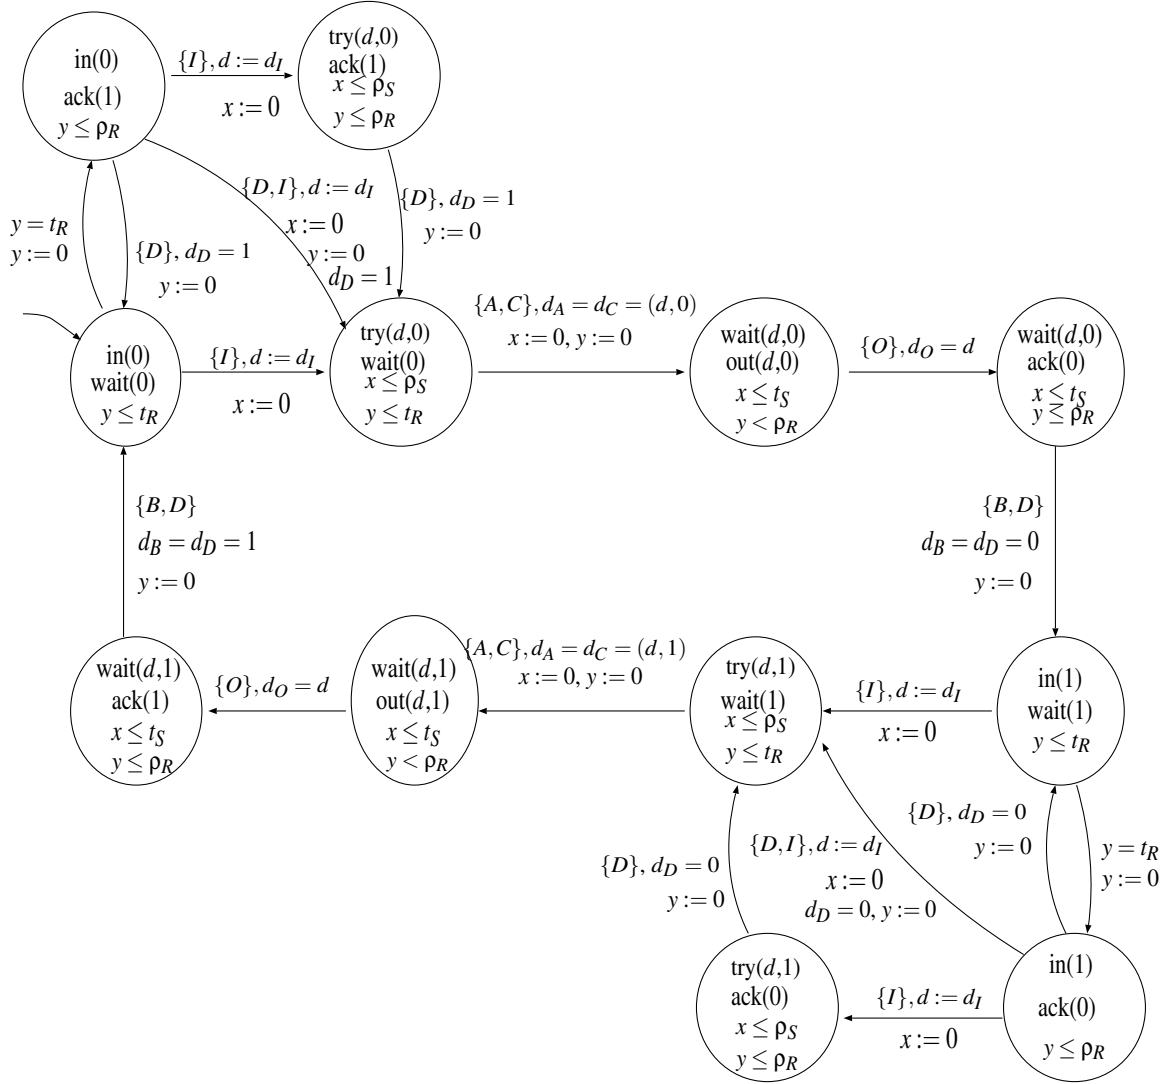


Figure 10: TCA for the ABP for  $\rho_R < \min\{\rho_S, t_S\}$  and  $\rho_S < t_R$

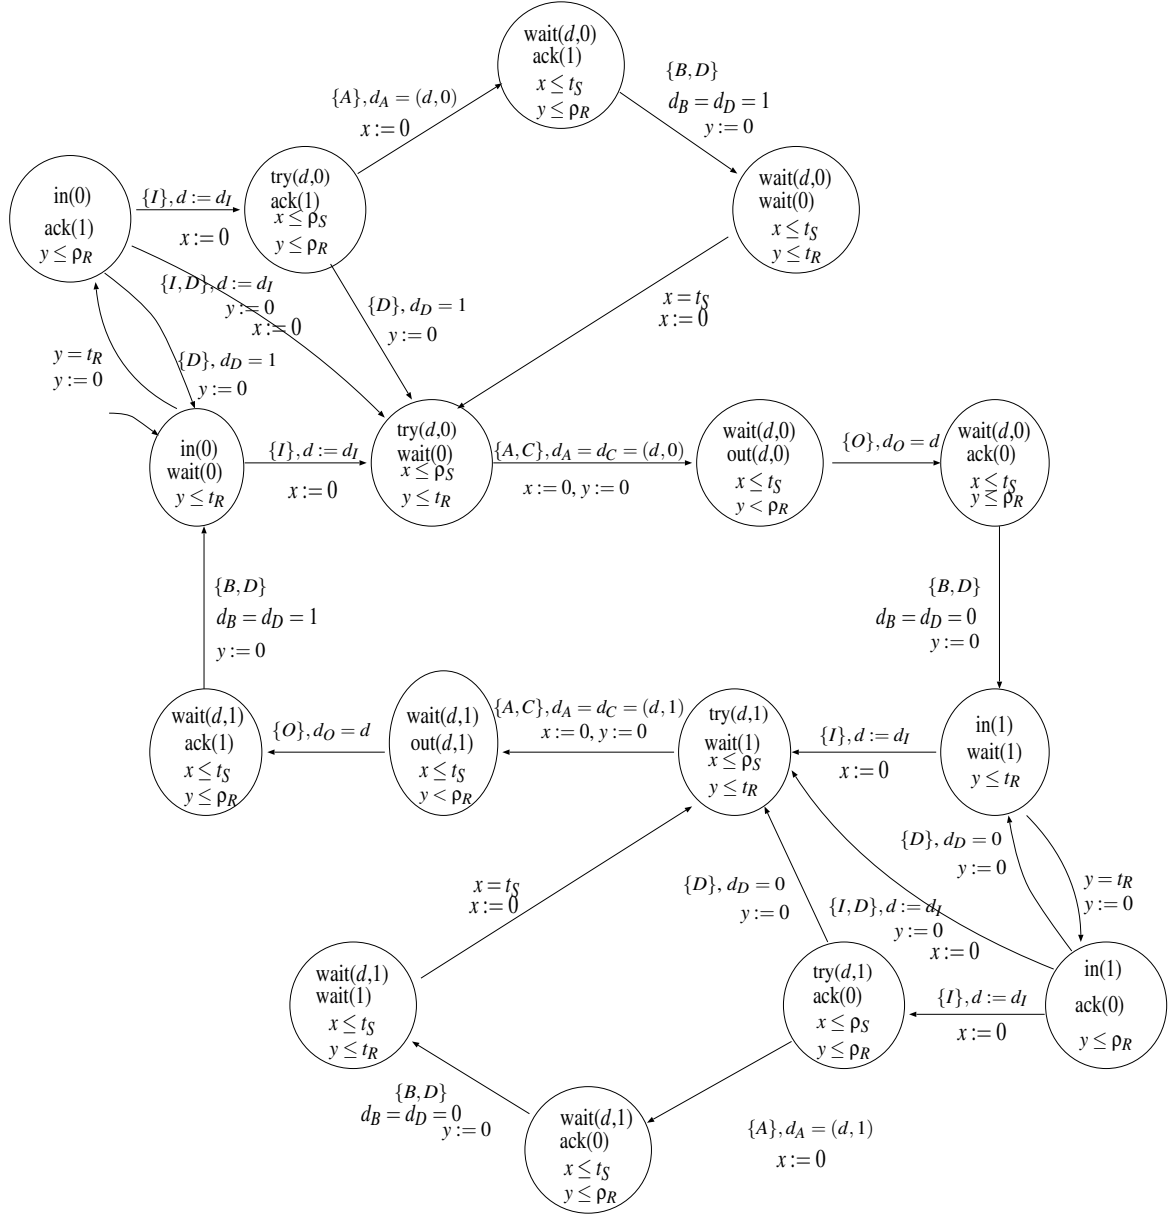


Figure 11: TCA for the ABP for  $\rho_R < t_S < t_R$  and  $\rho_S < t_R - t_S$