

# Communicating automata

*Dietrich Kuske<sup>1</sup> and Anca Muscholl<sup>2</sup>*

<sup>1</sup>TU Ilmenau, Institut für Theoretische Informatik

<sup>2</sup> LaBRI, Université Bordeaux and CNRS

2010 Mathematics Subject Classification: Primary: 68Q68 Secondary: 68Q45, 68Q60

Key words: communicating automata, fifo channels, model checking, realizability

## 1 Introduction

Communicating automata are a key computational model for concurrent systems. In this simple model, a finite number of processes cooperate via asynchronous message passing using unbounded channels. Such automata are very popular for modeling and reasoning about communication protocols: they are used to define the semantics of standardized protocol specification languages such as the ITU specification and description language SDL (Specification and Description Language). Automatic verification of properties of communicating automata is yet challenging, since these machines are Turing powerful [18], in particular they subsume Post tag systems [65]. As a consequence, every non-trivial property, including state reachability, is undecidable. As a consequence, approximated or semi-algorithmic verification techniques are essential. The aim of this survey is to describe automated approaches developed for verifying properties of communicating automata.

The reachability problem is the simplest instance of the more general *model-checking* problem, that consists in verifying that all behaviors of an automaton satisfy a given property, usually described in some logical formalism such as e.g. temporal logics [54]. A more ambitious goal is *synthesis*: the aim here is to compute a correct automaton from the beginning. Given a specification, the synthesis problem asks for an automaton with the same set of behaviors as the specification.

In this survey we report on two analysis problems for communicating automata: the *model-checking* problem and a simpler variant of synthesis, called the *realizability* problem. For realizability, the automaton that is to be computed does not interact with the environment. This makes realizability easier than synthesis, where one also assumes an unpredictable environment. It is worth to note that realizability is already a challenging problem for communicating automata, since it amounts to compute the distribution of the specification on a given process architecture, that limits the way information is conveyed.

The synthesis problem is even more challenging, since it requires to solve some kind of distributed game.

Semi-algorithmic verification falls mostly into two categories, *under-* and *over-approximating* methods. Under-approximating methods restrict the behavior of a system, whereas over-approximating ones consider some relaxation thereof. For instance, ignoring the order of messages in the channels is an over-approximation that turns communicating automata into Petri nets. Bounding the capacity of channels is an under-approximation that turns communicating automata into finite automata. Acceleration methods using some finitary representation of possibly infinite sets of configurations (called symbolic representations), are another example of under-approximating methods. In the case of communicating automata, such symbolic representations are based on finite automata or some tractable, extended automata models [10, 11, 16].

An interesting instance of over-approximation is the setting where communication channels can lose messages. Lossy channel systems are a particular instance of well-structured transition systems [1, 32], and this implies that their reachability problem is decidable [2], albeit of non-primitive recursive complexity [66]. On the other hand, more general properties like liveness for lossy channel systems, are undecidable [1].

The above mentioned approaches to the reachability problem emphasize the languages of messages and present them symbolically. Alternatively, the *partial-order* approach emphasizes the executions of communicating automata. The partial-order approach is event-based and captures naturally the concurrent aspect of process communication. Its main advantage is that the realizability problem can be stated very easily, since the specifications and the behaviors of communicating automata are both captured by the same kind of objects, namely message sequence charts.

The partial-order approach offers solutions for the model-checking and the realizability problem, assuming some universal or existential bounds on channels [43, 41, 38]. Whereas communicating automata with universal channel bounds are finite-state systems, the existentially bounded version belongs to the class of infinite-state systems. An existential channel bound means that the events of any execution of the automaton can be scheduled in such a way that they can be executed with bounded channels. A simple example illustrating this idea is a pair of processes, a producer and a consumer, where the producer keeps sending messages to the consumer. Since there is no control on the relative speed of the two processes, there is no bound on the number of transiting messages. But for verifying many properties, like e.g. control-state reachability, it suffices to consider only schedules where the messages are consumed without delay, thus corresponding to a channel of size one. For communicating automata with universal and existential channel bounds the realizability problem has a solution, and this can be stated as a Kleene-Büchi theorem for languages of message sequence charts.

**Overview.** Section 2 introduces the automaton model and the questions we consider in this survey. In Part I we focus on the reachability problem and summarize results on symbolic representations (Section 3) and lossy channel systems (Section 4). In Part II we introduce various specification frameworks, ranging between logics and message sequence charts. Then we discuss solutions for the model-checking problem. In Part III we present solutions for the realizability problem w.r.t. the specifications introduced in Part II.

Several recent research directions are not included in this chapter. One particular question concerns the analysis of extensions of communicating automata by additional storage capabilities, including pushdown storage [8, 49, 44, 25, 26, 23, 24].

## 2 Communicating Automata and Verification

Communicating automata follow the simple paradigm of a network of automata cooperating asynchronously over point-to-point, fifo communication channels. They arise naturally as models for peer-to-peer interaction, as occurring e.g. in distributed protocols using asynchronous message passing.

We consider systems described by means of a fixed *communication network*, consisting of a finite set of concurrent *processes*  $\mathcal{P}$ , together with a set of *channels*  $Ch \subseteq \{(p, q) \in \mathcal{P}^2 \mid p \neq q\}$ , that stand for point-to-point links [18]. Following the classical definitions, we exclude multiple channels between a pair of processes, as well as self-linking channels. However, this restriction has no big impact on the kind of results we will present. In our model, processes act either by point-to-point communication or by local actions. Possible actions come from a (finite) *communication alphabet*  $\Sigma$ , that is parametrized by the network  $(\mathcal{P}, Ch)$ , a set  $Msg$  of message contents, and a set  $Act$  of local actions (for convenience we omit  $\mathcal{P}, Ch, Msg, Act$  from the notation  $\Sigma$ ). A send action denoted as  $p!q(m)$  means that process  $p$  sends a message with content  $m \in Msg$  to process  $q$  on channel  $(p, q) \in Ch$ . A receive action denoted as  $p?q(m)$  means that  $p$  receives from  $q$  a message with content  $m \in Msg$  on channel  $(q, p) \in Ch$ . A local action denoted as  $c_p$  means that process  $p$  performs the action  $c \in Act$ . The set of  $p$ -actions (equivalently, the  $p$ -local alphabet) equals  $\Sigma_p = \{p!q(m), p?q(m), c_p \mid (p, q) \in Ch, (q, p) \in Ch, m \in Msg, c \in Act\}$ , and we let  $\Sigma = \bigcup_{p \in \mathcal{P}} \Sigma_p$  denote the set of all actions. As usual,  $\Sigma^*$  and  $\Sigma^\omega$  denote the set of finite and infinite sequences over  $\Sigma$ , resp., and  $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$ .

**Definition 2.1.** A *communicating automaton* (CA for short) is a tuple  $\mathcal{A} = \langle (\mathcal{A}_p)_{p \in \mathcal{P}}, \Sigma, F \rangle$  where

- each  $\mathcal{A}_p = (S_p, \rightarrow_p, s_p^0)$  is a labeled transition system with state space  $S_p$ , transition relation  $\rightarrow_p \subseteq S_p \times \Sigma_p \times S_p$ , and initial state  $s_p^0 \in S_p$ ;
- $F \subseteq \prod_{p \in \mathcal{P}} S_p$  is a set of *global* final states.

We denote the product  $S := \prod_{p \in \mathcal{P}} S_p$  as set of *global states*.

The *size* of a CA  $\mathcal{A} = \langle (\mathcal{A}_p)_{p \in \mathcal{P}}, \Sigma, F \rangle$  with message contents  $Msg$  is defined as  $\sum_{p \in \mathcal{P}} |S_p| + |Msg| + |Act|$ .

The behavior of a CA is defined as the behavior of an infinite labeled transition system, by considering the possible (local) transitions on the set of configurations of the CA. A *configuration* of the CA  $\mathcal{A}$  consists of a global state, together with a word from  $Msg^*$  for each channel  $(p, q) \in Ch$ . We write  $C = \langle s, w \rangle$  for a configuration with global state  $s \in S$  and channel contents  $w \in (Msg^*)^{Ch}$ , and denote by  $s_p$  and  $w_{p,q}$ , the  $p$ -component of  $s$  and the  $(p, q)$ -component of  $w$ , respectively. The set of all configurations of  $\mathcal{A}$  is denoted  $\mathcal{C}_\mathcal{A}$  (or simply  $\mathcal{C}$  when there is no risk of confusion). The *initial configuration* of

$\mathcal{A}$  is  $C_0 = \langle s^0, \vec{\varepsilon} \rangle$  with  $\vec{\varepsilon}_{p,q} = \varepsilon$  (the empty word) for all  $(p, q) \in Ch$ . A configuration  $C = \langle s, w \rangle$  is *final* if  $s \in F$  is a final state of  $\mathcal{A}$  (note that the channels need not be empty in a final configuration).

For two configurations  $C = \langle s, w \rangle, C' = \langle s', w' \rangle$  and an action  $a \in \Sigma_p$ , we write  $C \xrightarrow{a} C'$  if the following hold:

- $s_p \xrightarrow{a} s'_p$  is a transition of  $\mathcal{A}_p$  (i.e.,  $(s_p, a, s'_p) \in \rightarrow_p$ ), and  $s'_q = s_q$  for all  $q \neq p$ ,
- if  $a = p!q(m)$  is a send action, then  $w'_{p,q} = w_{p,q}m$  (message  $m$  is inserted into the channel from  $p$  to  $q$ ) and  $w'_{r,s} = w_{r,s}$  for all  $(r, s) \neq (p, q)$  (all other channels are unchanged).
- if  $a = p?q(m)$  is a receive action, then  $w_{q,p} = mw'_{q,p}$  (message  $m$  is deleted from the channel from  $q$  to  $p$ ) and  $w'_{r,s} = w_{r,s}$  for all  $(r, s) \neq (q, p)$  (all other channels are unchanged).
- if  $a = c_p$  is a local action, then  $w = w'$ .

We say that  $C'$  is a *successor* of  $C$  (and write  $C \longrightarrow C'$ ) if there exists some  $a \in \Sigma$  with  $C \xrightarrow{a} C'$ . As usual, we write  $\xrightarrow{*}$  for the reflexive-transitive closure of  $\longrightarrow$ . For a set  $X \subseteq \mathcal{C}$  of configurations, we write  $post(X)$  for the set of successors of configurations from  $X$  and  $post^*(X)$  for the set of all *reachable* configurations from  $X$ :

$$post(X) := \{C' \in \mathcal{C} \mid \exists C \in X : C \longrightarrow C'\},$$

$$post^*(X) := \{C' \in \mathcal{C} \mid \exists C \in X : C \xrightarrow{*} C'\}.$$

The *reachability set* of a CA  $\mathcal{A}$ , denoted  $Reach(\mathcal{A})$ , is the set


$$Reach(\mathcal{A}) = post^* (\{C_0\})$$

(recall that  $C_0$  is the initial configuration of  $\mathcal{A}$ ).

A CA is *deadlock-free*, if from every reachable configuration it can reach a final configuration.

**Example 2.1.** The CA in the figure below describes the communication between two (finite-state) processes **C** and **S**, connected through one channel in each direction (process **C** on the left, and **S** on the right). The set of message contents is  $Msg = \{0, 1, \$\}$ . From the initial configuration  $\langle (c_0, s_0), (\varepsilon, \varepsilon) \rangle$  (say,  $(\mathbf{C}, \mathbf{S})$  is the first channel) the configurations  $\langle (c_1, s_0), (010, \varepsilon) \rangle$  and  $\langle (c_0, s_0), (101, \$) \rangle$  are reachable, but not  $\langle (c_0, s_0), (0101, \$) \rangle$ .

For instance,  $\langle (c_0, s_0), (\varepsilon, \varepsilon) \rangle \xrightarrow{C!S(0)} \langle (c_1, s_0), (0, \varepsilon) \rangle \xrightarrow{C!S(1)} \langle (c_0, s_0), (01, \varepsilon) \rangle$ .



A CA  $\mathcal{A}$  is called *deterministic*, if for every local state  $s \in S_p$  the following holds:

- $s \xrightarrow{a}_p s_1$  and  $s \xrightarrow{a}_p s_2$  implies  $s_1 = s_2$ , for every  $a \in \Sigma_p$ ,
- $s \xrightarrow{p!q(m_1)}_p s_1$  and  $s \xrightarrow{p!q(m_2)}_p s_2$  (for some  $s_1, s_2$ ) implies  $m_1 = m_2$ .

The notion of determinism used above originates from [43]. The second condition in this definition is motivated by viewing message contents as control information, that has to be chosen deterministically by the sending process.

**Definition 2.2.** A *run* of a CA  $\mathcal{A}$  is a (finite or infinite) sequence of transitions:  $\rho = C_1 \xrightarrow{a_1} C_2 \xrightarrow{a_2} C_3 \cdots$ , with  $C_i \in \mathcal{C}_{\mathcal{A}}$  configurations and  $a_i \in \Sigma$  actions. The *labeling* of the run  $\rho$ , denoted  $\ell(\rho)$ , is the sequence of actions  $a_1 a_2 \cdots \in \Sigma^\infty$ .

A finite run  $\rho = C_1 \xrightarrow{a_1} C_2 \xrightarrow{a_2} \cdots \xrightarrow{a_{n-1}} C_n$  is *accepting* if  $C_1 = \langle s^0, \vec{\varepsilon} \rangle$  is the initial configuration and  $C_n \in F \times (Msg^*)^{Ch}$  is a final configuration.

The *language* of a CA  $\mathcal{A}$ , denoted  $\mathcal{L}(\mathcal{A}) \subseteq \Sigma^*$ , is the set

$$\mathcal{L}(\mathcal{A}) = \{w \in \Sigma^* \mid w = \ell(\rho) \text{ for some accepting run } \rho\}.$$

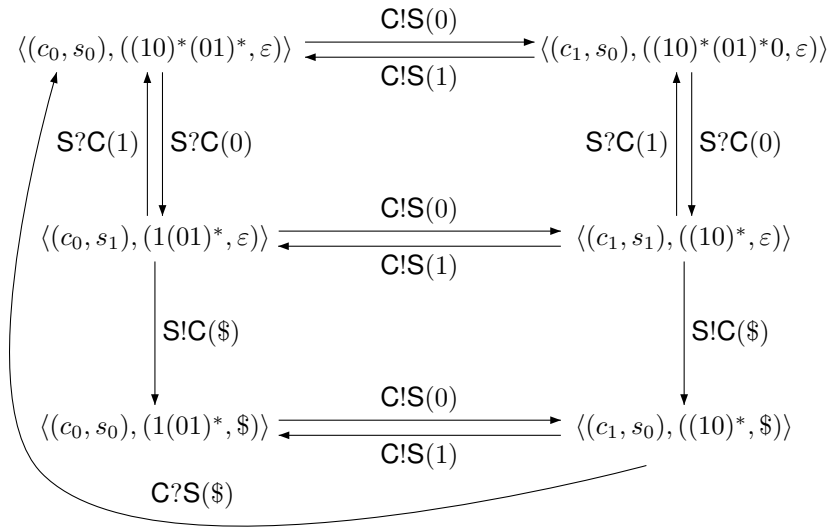
**Example 2.2.** An abstract view of the transition system associated to the CA  $\mathcal{A}$  from Example 2.1 is given in Figure 1. There, we grouped certain configurations and connected them by action-labeled edges such that the initial configuration belongs to the group at the top-left corner and any  $a$ -successor of any configuration of some group belongs to the  $a$ -successors of that group. One can easily verify that the groups of configurations form the least solution to this constraint system. Hence the reachability set  $Reach(\mathcal{A})$  is the union of all the configurations mentioned in that figure.

Assuming that all states of  $\mathcal{A}$  are final, the language  $\mathcal{L}(\mathcal{A})$  of  $\mathcal{A}$  consists of words over  $\{C!S(0), S?C(0), C!S(1), S?C(1)S!C(\$), C?S(\$)\}$  of the following form: the projection on  $\Sigma_C$  belongs to the prefix closure of  $((C!S(0) C!S(1))^* C!S(0) C?S(\$))^*$ , the projection on  $\Sigma_S$  belongs to the prefix closure of  $((S?C(0) S?C(1))^* S?C(0) S!C(\$))^*$ , and the words are well-formed w.r.t. the send/receives (cf. definition of “valid word” in Section 5). In particular, the process  $S$  sends at most one more message than process  $C$  receives.

Notice that we did not impose in the definition of a CA  $\mathcal{A} = \langle (A_p)_{p \in \mathcal{P}}, \Sigma, F \rangle$  any restriction on the local automaton  $A_p$ . In general, we might be interested in various kinds of (possibly infinite-state) automata, such as pushdown automata [49, 44]. However, a basic kind of CA is obtained by requiring that every  $A_p$  is a finite-state automaton, and then we denote  $\mathcal{A}$  as *communicating finite-state machine* (CFM for short). Most of the research done in the past decades on CAs focused on CFMs, and we will concentrate on them in the next sections.

### Basic verification questions

The simplest property considered in automated verification is *safety*, which amounts to ask a reachability question. The *reachability problem* for CFMs in its most general form consists in computing the set  $post^*(Init)$  for some set of configurations  $Init$ . A more specific reachability question consists in asking for two configurations  $C_1, C_2 \in \mathcal{C}$  of



**Figure 1.** An abstract view of the transition system.

some CFM  $\mathcal{A}$ , whether  $C_1 \xrightarrow{*} C_2$ . A further instantiation of the reachability problem is control-state reachability:

*Control-state reachability problem:* Given a CFM  $\mathcal{A}$ , a configuration  $C$  and a global state  $s \in S$ , is some configuration with global state  $s$  reachable from  $C$ ?

All reachability problems above are undecidable for CFMs [18], and in Part I and Section 7.1 of this survey we present three methods to tackle this problem.

Obviously, there are many more verification questions that one can ask about a given CFM, some of them being specific to the model. We give some examples below:

- *Termination:* Given a CFM  $\mathcal{A}$  and a configuration  $C$ , are all runs of  $\mathcal{A}$  from  $C$  finite?
- *Structural termination:* Given a CFM  $\mathcal{A}$ , are all runs of  $\mathcal{A}$  from *any* configuration finite?
- *Absence of deadlocks:* Given a CFM, is it deadlock-free?
- *Boundedness:* Given a CFM  $\mathcal{A}$  and a configuration  $C$ , are only finitely many configurations reachable from  $C$ ?
- *Repeated reachability (Büchi acceptance):* Given a CFM  $\mathcal{A}$ , a configuration  $C$  and a global state  $s$ , is there some run from  $C$  that meets  $s$  infinitely often?
- *Model-checking:* Given a CFM  $\mathcal{A}$  and a property  $P$ , do all words from  $\mathcal{L}(\mathcal{A})$  satisfy  $P$ ?

The last of these questions is very general, as the term “property  $P$ ” can be instantiated in very different ways. Here, we will concentrate on the case where  $P$  is given by some formula (see Part II for details). This survey does not discuss the branching-time model-checking problem – that talks about properties of the collection of all runs – but only

the linear-time case, see also Chapter 38 for a detailed discussion about linear- versus branching-time specifications.

## Part I

# Reachability

This part presents two methods to tackle the undecidability of the reachability problem for CFM: *Symbolic representations* of the reachability set help to accelerate the naive enumeration of all reachable configurations, and the *lossy channel* method allows to compute an over-approximation of the reachability set. A third approach based on *channel bounds* will be used to compute under-approximations of the reachability set in Section 7.1.

## 3 Symbolic Representations

The set of all reachable configurations of some CFM can be easily enumerated by a breadth-first search which requires to handle finite sets of configurations. Eventually, every reachable configuration is computed. In this section, we discuss a technique that speeds up this process. There are two main ideas involved:

- (1) Consider possibly infinite sets of configurations, that are given symbolically through a finite description (see also Chapter 32 for more details related to symbolic representations of sets of numbers).
- (2) Compute the effect of “meta-transitions”, that correspond to sequences of actions of the CFM.

In the following, let  $\mathcal{A} = \langle (\mathcal{A}_p)_{p \in \mathcal{P}}, \Sigma, F \rangle$  be some CFM with global state space  $S$ . Let us enumerate the channels in some fixed way, say as  $ch_1$  to  $ch_k$ , such that a configuration  $\langle s, w \rangle$  with  $w = (w_{ch_i})_{i=1}^k$  can be identified with the word  $s w_{ch_1} \# w_{ch_2} \# \dots w_{ch_k} \#$  (assuming that  $\# \notin \text{Msg} \cup S$ ).

One of the first attempts to use finite automata for symbolic representations can be found in [63]. There, the CFM  $\mathcal{A}$  is said to have *recognizable channels* if the set  $\{sw \in S(\text{Msg}^* \#)^k \mid \langle s, w \rangle \in \text{Reach}(\mathcal{A})\}$  is a regular language. For CFMs with recognizable channels, reachability is decidable [63]. One can check this by using two semi-algorithms: one is the plain enumeration of all reachable configurations, and the other one enumerates all regular languages  $L \subseteq S(\text{Msg}^* \#)^k$ , looking for one that contains the initial configuration  $C_0$  (i.e., with  $s^0 \#^k \in L$ ), excludes the target configuration, and is closed under the one-step transition relation. Notice however that recognizable channels are very restrictive: a CFM containing one process that sends the same (arbitrary) sequence of messages to two other processes, does not have recognizable channels.

*Queue-content Decision Diagrams* [10] (QDD for short) describe (possibly infinite) sets of configurations of a CFM by finite automata. The main idea here is to accelerate the computation of the reachability set – which however might be non-representable by a QDD. To this aim, one analyzes the effect of iterating a loop in the transition system

associated with the CFM.

More precisely, a QDD is a finite automaton accepting words from  $(Msg^*\#)^k$ . A set of configurations  $X$  is *QDD-representable*, if all configurations from  $X$  have the same global state  $s \in S$  and there is a QDD accepting the set  $\{w \in (Msg^*\#)^k \mid \langle s, w \rangle \in X\}$  (i.e., if this set is regular). Now let  $\sigma \in \Sigma^*$ ,  $s \in S$  a global state,  $X$  a set of configurations with global state  $s$ . Then  $post_\sigma^*(X)$  is the set of channel contents  $w'$  such that the configuration  $\langle s, w' \rangle$  is reachable from some configuration  $\langle s, w \rangle \in X$  by a path whose label belongs to  $\sigma^*$ .

A word  $\sigma$  is (effectively) *QDD-preserving*, if  $post_\sigma^*(X)$  is (effectively) QDD-representable whenever  $X$  is a QDD-representable set of configurations. It is *non-counting w.r.t. the channel*  $(p, q)$  if one of the following conditions holds:

- $\sigma$  contains no letter  $p!q$ ,
- the message alphabet of channel  $(p, q)$  is unary and  $\sigma$  contains the same number of occurrences of  $p!q$  as of occurrences of  $q?p$ .

**Theorem 3.1** ([11]). *The following assertions are equivalent for any word  $\sigma \in \Sigma^*$ :*

- (1)  $\sigma$  is QDD-preserving,
- (2)  $\sigma$  is effectively QDD-preserving,
- (3)  $\sigma$  is non-counting w.r.t. all but at most one channel.

**Example 3.1.** We consider again the CFM from Example 2.1 (page 1032). The word  $\sigma = C!S(0)C!S(1)$  is non-counting w.r.t. channel  $(S, C)$  (but not w.r.t. the channel  $(C, S)$ ), so the theorem above can be applied with  $X = \{C_0\}$ , yielding  $post_\sigma^*(X) = \{((01)^n, \varepsilon) \mid n \geq 0\}$ . The sequence  $\sigma' = C!S(0)S?C(0)S!C(\$)C?S(\$)C!S(0)C!S(1)$  is non-counting w.r.t.  $(S, C)$  and  $post_{\sigma'}^*(X) = \{(01, \varepsilon), (1001, \varepsilon)\}$ .

A semi-algorithm for model-checking LTL properties (with atomic propositions referring to global states), based on the QDD-representation, was proposed in [11].

Symbolic representations going beyond regular ones were introduced in [16]: *Constrained Queue-content Decision Diagrams* (CQDD for short). A CQDD consists of a restricted finite automaton  $\mathcal{B}$ , together with a Presburger constraint  $\pi$ . An accepting run in a CQDD is an accepting run of the automaton  $\mathcal{B}$  whose Parikh image satisfies  $\pi$ . A word  $\sigma \in \Sigma^*$  is (effectively) CQDD-preserving, if  $post_\sigma^*(X)$  is (effectively) CQDD-representable whenever  $X$  is a CQDD-representable set of configurations. It should be noted that CQDDs and QDDs are incomparable: CQDDs add power via the Presburger constraints, but lose on the level of the regular set since only a restricted form of automata is allowed. The following theorem shows that CQDDs can be used for a semi-algorithm computing the reachability set from any given initial configuration (since any finite set is CQDD-representable).

**Theorem 3.2** ([16]). *Every word  $\sigma \in \Sigma^*$  is effectively CQDD-preserving.*

As mentioned before, the above results allow to speed up the enumeration of all reachable configurations of a CFM. Suppose that we want to know whether some configuration with global state  $s$  is reachable from the initial configuration  $C_0$ . The trivial semi-algorithm mentioned at the beginning of this section performs a breath-first search in



the transition system, thereby enumerating all sequences of actions that can be performed from the initial configuration. The improved semi-algorithm handles at once all sequences of the form  $\rho_1 \sigma^* \rho_2$  with  $\rho_1, \sigma, \rho_2 \in \Sigma^*$  and computes the set  $post_{\rho_2}(post_{\sigma}^*(post_{\rho_1}(C_0)))$ .

## 4 Faulty Channel Systems

In this section, we show a technique that allows to compute a superset of the reachability set of a CFM. The idea is to add transitions to the labeled transition system derived from the CFM. This can be done in many more or less natural ways. The first possibility is to ignore the order of messages in the channels. The CFM becomes thus a Petri net, whose reachability problem is decidable [55, 45, 50] (albeit of high complexity, since no primitive recursive algorithm is known so far). However, the main problem with this approximation technique is that it is too coarse for a realistic modeling.

The additional transitions considered here can either lose messages from or introduce messages into channels. Such a model is interesting in its own right since it allows us to model imperfect channels. *Lossy machines* are CFMs where channels can lose an arbitrary number of messages, at any time. For CFMs with *insertion errors*, new messages can be inserted in channels, at any time. Although these two models have different flavor, the techniques used to manipulate them are quite similar. However, we will see a minor qualitative difference between the results obtained for these two models (cf. Remark 4.6 below).

Lossy CFMs (or *lossy channel systems*) represent a special instance of a more general class of infinite-state systems, namely *well-structured transition systems* (WSTS for short). WSTS were considered independently by Finkel and Schnoebelen [31, 32], and by Abdulla and Jonsson [1, 2]. The basic idea behind a WSTS  $\langle \mathcal{S}, \rightarrow \rangle$  is to endow the set of configurations  $\mathcal{S}$  with a *well quasi-order* (wqo for short) in order to manipulate certain infinite subsets of  $\mathcal{S}$  symbolically. A wqo  $\preceq$  on  $\mathcal{S}$  is a quasi-order without infinite anti-chains and without infinite decreasing chains. A transition system  $\langle \mathcal{S}, \rightarrow, \preceq \rangle$  is a WSTS if it satisfies the following *monotonicity* property: for every  $s \rightarrow s'$  and every  $s_1 \in \mathcal{S}$  with  $s \preceq s_1$ , it is required that some  $s'_1 \in \mathcal{S}$  exists with  $s_1 \rightarrow s'_1$  and  $s' \preceq s'_1$ . Monotonicity appears with different flavors in [32]. One variant consists in requiring  $s_1 \rightarrow^+ s'_1$  in the definition above. This is then denoted as *transitive monotonicity*.

Two properties are crucial for WSTS. The first one is that every subset  $X \subseteq \mathcal{S}$  has a *finite* set of minimal elements, denoted  $\min(X)$ . The second property is that the predecessor relation preserves *upward-closed*<sup>1</sup> sets, i.e.,  $pre(X) := \{s \mid s \rightarrow s' \text{ for some } s' \in X\}$  is upward-closed whenever  $X$  is upward-closed. As a consequence, reachability of upward-closed sets  $X$  of states can be decided by a backward algorithm, that computes the set  $pre^*(X) = \bigcup_{n \geq 0} pre^n(X)$  as least fixpoint of the operation  $Y \mapsto pre(Y) \cup X$ . Intersecting the result with the set of initial configurations solves the reachability problem for WSTS.

**Theorem 4.1** ([2]). *Let  $\langle \mathcal{S}, \rightarrow, \preceq \rangle$  be a WSTS such that  $\prec$  is decidable and  $\min(pre(X))$*

<sup>1</sup> $X$  is upward-closed if  $X = \{s' \mid s \preceq s' \text{ for some } s \in X\}$ .

is computable from  $\min(X)$  for every upward-closed set  $X \subseteq \mathcal{S}$ . Then, for  $s, s' \in \mathcal{S}$ , it is decidable whether there exists  $s'' \in \mathcal{S}$  with  $s \rightarrow^* s''$  and  $s' \preceq s''$ .

Termination for WSTS can be decided by a forward algorithm, computing the *finite reachability tree*  $FRT(s)$  from a state  $s \in \mathcal{S}$ : this is the prefix of the reachability tree built from state  $s$ , obtained by defining a node  $t'$  as leaf if there is some node  $t \preceq t'$  on the path from  $s$  to  $t'$ . Note that the tree  $FRT(s)$  is finite if  $\langle \mathcal{S}, \rightarrow \rangle$  has finite branching.

**Theorem 4.2** ([32]). *Termination is decidable for finitely branching WSTS  $\langle \mathcal{S}, \rightarrow, \preceq \rangle$  with transitive monotonicity such that  $\prec$  is decidable and  $\{s' \in \mathcal{S} \mid s \rightarrow s'\}$  is computable from  $s \in \mathcal{S}$ .*

For lossy CFMs, the choice of a wqo is very natural. One starts with the subword ordering: let  $x \preceq y$  if  $x = x_1 \cdots x_n$  and  $y = y_0 x_1 y_1 \cdots y_{n-1} x_n y_n$  for some  $x_i, y_i \in \text{Msg}^*$ . This wqo extends to configurations of the CFM: for two configurations  $C = \langle s, w \rangle$ ,  $C' = \langle s', w' \rangle$ , let  $C \preceq C'$  if  $s = s'$  and  $w_{p,q} \preceq w'_{p,q}$  for all channels  $(p, q)$ . This proves the following result:

**Corollary 4.3.** *Control-state reachability and termination for lossy CFMs are decidable problems.*

However, the complexity of both problems is non-primitive recursive [66]. A precise complexity characterization was obtained in [21] in terms of a hierarchy of recursive functions.

On the negative side, more complex properties such as repeated reachability of a given global state, are undecidable for lossy CFMs [2]. Extensions of this undecidability result were obtained in [56], by considering lossy counter machines. These are usual counter machines (with zero tests) where the counters can be decremented spontaneously. The main result in the latter paper is that it is undecidable whether a lossy counter machine has an infinite run from some initial configuration:

**Theorem 4.4** ([56]). *Structural termination and boundedness for lossy counter machines are undecidable problems.*

Through a simulation of lossy counter machines by lossy CFMs, the above result extends to the latter model (and yields undecidability for two other problems):

**Corollary 4.5** ([2, 56]). *The following questions about lossy CFMs are undecidable: structural termination, boundedness, and repeated reachability.*

A different picture arises when the source of faults are insertion errors, as considered e.g. in [19, 17]. As defined there, insertion CFMs (called insertion channel machines) are CFMs where channels can acquire additional contents spontaneously. In addition, in [17], such machines are endowed with tests for channel emptiness and for non-occurrence of a given content.

**Remark 4.6.** Notice that for an insertion CFM, its reachability set (defined as in [19]) is recognizable, since it is upward-closed w.r.t. the wqo from the proof of Cor. 4.3. The same holds for lossy CFMs, since the reachability set is here downward-closed and therefore the complement of an upward-closed set. However, an essential difference is that a finite automaton accepting the reachability set of an insertion CFM is effectively computable, whereas this is not the case for lossy CFMs. For the first assertion one can use a similar argument as in the fixpoint computation of the backward reachability algorithm mentioned above [19]. For lossy CFMs the reachability set cannot be effectively computable, otherwise the boundedness problem for lossy CFMs would be decidable, contradicting Theorem 4.4 (just compute a finite automaton for the reachability set and check it for finiteness).

Together with the wqo from the proof of Cor. 4.3, insertion CFMs are WSTS. From Thm. 4.1, it follows that their control-state reachability problem is decidable as well. As for lossy CFMs, this problem is not primitive recursive. In contrast, insertion CFMs behave better w.r.t. the termination problem:

**Theorem 4.7** ([17]). *The (structural) termination problem for insertion CFMs has non-elementary, yet primitive recursive complexity.*

**Further reading.** The paper [19] also considers channels with duplication errors. Systems mixing lossy channels and error-free ones are considered in [20], providing a (polynomial) characterization of those architectures where reachability is decidable. Adding probabilities to lossy CFMs has been considered in several papers. For instance, non-deterministic CFMs with probabilistic lossiness were shown in [9] to have a decidable reachability problem, and so do some instances of the repeated reachability problem (e.g. the “almost surely” instance). On the other hand, repeated reachability with positive probability remains undecidable, by a reduction from the boundedness problem for lossy CFMs.

## Part II

# Specifications and Model-Checking

In this part, we first introduce formalisms to describe properties of runs of CFMs from a language-theoretic viewpoint. We start with sequential specifications – and show that model-checking CFMs against such specifications is undecidable. We then introduce message sequence charts, a graphical way of presenting the causal dependencies in a run of a CFM, and logics like monadic second order logic and propositional dynamic logic, that express properties of message sequence charts, i.e., causal properties of runs. In Section 7, we will investigate the model-checking problem for such logics. The results in this part support the “rule of thumb” that the only way to avoid undecidability of the model-checking problem is to use specifications that are compatible with the causal structure.

## 5 Sequential Specifications

We start with some notation first. A word over the communication alphabet  $\Sigma$  is called valid, if it can potentially be executed from some configuration with empty channels. To make this notion strict, let  $\pi_{p,q}^!, \pi_{p,q}^? : \Sigma^* \rightarrow Msg^*$  be the homomorphisms defined by  $\pi_{p,q}^!(p!q(m)) = \pi_{p,q}^?(q?p(m)) = m$  and  $\pi_{p,q}^!(a) = \pi_{p,q}^?(b) = \varepsilon$  for  $a$  not of type  $p!q$  and  $b$  not of type  $q?p$ . Then a word  $u \in \Sigma^*$  is *valid* if  $\pi_{p,q}^!(v)$  is a prefix of  $\pi_{p,q}^?(v)$  for any prefix  $v$  of  $u$  and any channel  $(p, q)$  (i.e., the sequence of messages sent from  $p$  to  $q$  is consistent with the sequence received, and a receive never precedes its matching send).

A simple formalism to describe a property of words over  $\Sigma$  are finite automata. Another one is linear time temporal logic LTL [54], that can refer to matching send and receive events. The syntax of this extended LTL is:

$$\sigma ::= true \mid \sigma \mid X\varphi \mid X_{msg}\varphi \mid \varphi U \varphi \mid \varphi \vee \varphi \mid \neg\varphi,$$

where  $\sigma \in \Sigma$ .

Let  $u = a_1 a_2 \cdots a_n$  be a valid word with  $a_i \in \Sigma$  and let  $1 \leq i \leq n$ . Then the satisfaction relation is defined inductively:

$$\begin{aligned} u, i \models \sigma &\Leftrightarrow a_i = \sigma \\ u, i \models X\varphi &\Leftrightarrow u, i+1 \models \varphi \text{ and } i < n \\ u, i \models X_{msg}\varphi &\Leftrightarrow a_i = p!q(m) \text{ and there exists } i < j \leq n \text{ with} \\ &\quad u, j \models \varphi, a_j = q?p(m), \text{ and } |\pi_{p,q}^!(a_1 \cdots a_i)| = |\pi_{p,q}^?(a_1 \cdots a_j)| \\ u, i \models \varphi_1 U \varphi_2 &\Leftrightarrow \text{there exists } i \leq k \leq n \text{ with } u, k \models \varphi_2 \\ &\quad \text{and } u, j \models \varphi_1 \text{ for all } i \leq j < k \end{aligned}$$

( $u, i \models \varphi_1 \vee \varphi_2$  and  $u, i \models \neg\varphi$  are defined in the obvious way.) The modalities  $X$  (next) and  $U$  (until) are standard,  $X_{msg}$  refers to the matching receive of the current send event. We write  $u \models \varphi$  for  $u, 1 \models \varphi$ .

**Example 5.1.** Unsurprisingly, the following example shows that the additional  $X_{msg}$  operator can express non-regular properties, even restricted to valid words over  $\Sigma$ . In order to keep it simple, we assume that  $\mathcal{P} = \{p, q, r\}$ ,  $Ch = \{(p, q), (p, r), (r, q)\}$  and that  $Msg$  is a singleton (therefore omitted).

Consider the LTL formula  $\varphi = true U (X c_p \wedge X_{msg} X c_q)$  expressing that there is a matching pair of send and receive events on channel  $(p, q)$ , that are immediately followed by the local actions  $c_p$  and  $c_q$ , resp. Let  $L$  be a regular language that contains all valid words satisfying  $\varphi$ . Notice that words of the form

$$u = (p!q)^{x_1} c_p (p!q)^{x_2} (p!r r?p r!q q?r) (q?p)^{y_1} c_q (q?p)^{y_2}$$

with  $x_i, y_i \geq 0$ , are valid if and only if  $x_1 + x_2 \geq y_1 + y_2$ . A valid word  $u$  as above belongs to  $L$  if and only if  $x_1 = y_1 > 0$ . Provided  $y_1$  is sufficiently large, a pumping argument allows us to find  $z_1 < y_1$  such that also

$$v = (p!q)^{x_1} c_p (p!q)^{x_2} (p!r r?p r!q q?r) (q?p)^{z_1} c_q (q?p)^{y_2}$$

belongs to  $L$ . But the valid word  $v$  does not satisfy  $\varphi$ . Hence, any regular language that contains all valid words satisfying  $\varphi$  also contains some valid word not satisfying  $\varphi$ .

The most serious problem about sequential specifications is that the model-checking problem for CFMs is undecidable.

**Proposition 5.1.** *Model-checking CFMs against LTL properties is undecidable.*

An obvious reduction from Post’s correspondence problem yields a fixed CFM with an undecidable model-checking problem: Let  $\mathcal{P} = \{p, q, r, s\}$ ,  $Ch = \{(p, q), (r, s)\}$ , and  $Msg = \{a, b, 1, 2, 3, 4, 5, \$\}$ . The CFM  $\mathcal{A}$  we want to model-check is the “universal” CFM on this architecture, and its language is the set of all valid words. Now let  $\mathcal{I} = (u_i, v_i)_{1 \leq i \leq 5}$  be five pairs of words over the alphabet  $A = \{a, b\}$ , i.e., an instance of Post’s correspondence problem. Then the language

$$\left[ \bigcup_{1 \leq i \leq 5} p!q(i)r!s(u_i) \right]^+ p!q(\$)r!s(\$) \left[ \bigcup_{1 \leq i \leq 5} q?p(i)s?r(v_i) \right]^+ q?p(\$)s?r(\$)$$

can be expressed by an LTL formula  $\varphi$  (here,  $r!s(a_1 a_2 \dots a_n)$  means the sequence of sends  $r!s(a_1) \dots r!s(a_n)$ , and  $s?r(a_1 \dots a_n)$  is to be understood similarly).

Suppose  $u$  belongs to the language of  $\varphi$  and is valid. Then  $\pi_{p,q}^!(u) = \pi_{p,q}^?(u) = i_1 i_2 \dots i_n \$$  implies that  $i_1, \dots, i_n$  is a solution for  $\mathcal{I}$ . Conversely, any solution to  $\mathcal{I}$  gives a valid word satisfying  $\varphi$ . Hence, all words from the language of  $\mathcal{A}$  satisfy  $\neg\varphi$  if and only if  $\mathcal{I}$  has no solution – but this is undecidable [62].

## 6 Partial Order Specifications

A CFM is meant to model a network of independent processes. In particular, distinct processes can act at the same time. For instance, the processes  $p$  and  $p'$  can, at the very same moment, send messages  $m$  and  $m'$  to processes  $q$  and  $q'$ , resp.

Recall that from a CFM, we defined a labeled infinite transition system such that executions of the CFM are considered as paths in this transition system. This forces us to linearly order the two actions  $p!q(m)$  and  $p'!q'(m')$  from above, giving rise to two different paths in the transition system.

Although simple, the interleaving semantics is not satisfactory for at least two reasons: one is the undecidability of the model-checking problem, the other is that causal properties, e.g. races, are hard to formulate in the interleaving semantics. In this section, we model computations of CFMs by particular partial orders, called message sequence charts. We will propose two logical formalisms that can describe causal properties of such objects.

### 6.1 Message sequence charts

Message sequence charts are  $\Sigma$ -labeled posets  $\langle E, \leq, \lambda \rangle$ , with  $E$  a set of events,  $\leq$  a partial order on  $E$ , and  $\lambda : E \rightarrow \Sigma$  a labeling function. We write  $P(e)$  for the process on which an event  $e$  is located. That is, we let  $P(e) = p$  if  $\lambda(e) \in \Sigma_p$ , call  $e$  a  $p$ -event, and let  $E_p = P^{-1}(p)$  be the set of all  $p$ -events. An event labeled by some  $p!q(m)$  ( $q?p(m)$ , resp.) is called an event of *type*  $p!q$  ( $q?p$ , resp.).

We need to define two relations  $\leq_P$  and  $<_{msg}$  on events:

- $e \leq_P f$  if  $P(e) = P(f)$  and  $e \leq f$ .
- $e <_{msg} f$  if for some  $(p, q) \in Ch$ , the two following conditions hold:
  - (1)  $e$  is an event of type  $p!q$  and  $f$  of type  $q?p$ .
  - (2) The number of events  $e' \leq e$  of type  $p!q$  equals the number of events  $f' \leq f$  of type  $q?p$ .

The meaning of the two relations  $\leq_P$  and  $<_{msg}$  is as follows. The relation  $\leq_P$  describes the order of the events executed by each process, whereas  $<_{msg}$  describes matching pairs of send and receive events (under the assumption that channels are FIFO). By  $<_P$  we denote the immediate process successor relation:  $e <_P f$  if  $e <_P f$  and  $e \leq_P g <_P f$  implies  $e = g$ .

**Definition 6.1.** A *message sequence chart* (MSC for short) is a finite  $\Sigma$ -labeled poset  $M = \langle E, \leq, \lambda \rangle$  – up to isomorphism – satisfying the conditions below:

- $\leq = (\leq_P \cup <_{msg})^*$ .
- The set of  $p$ -events  $E_p \subseteq E$  is linearly ordered for any process  $p \in \mathcal{P}$ .
- For any event  $f$  of type  $q?p$ , there exists an event  $e$  and a message content  $m \in Msg$  such that  $e <_{msg} f$ ,  $\lambda(e) = p!q(m)$ , and  $\lambda(f) = q?p(m)$ .


The relation  $<_{msg}$  matches sends and receives. By the last requirement, matching send- and receive-events handle the same message content. It can be understood as saying “any message received has been sent”. In the literature, one often finds in addition the complementary requirement “any message sent will be received”, the above defined concept is then known as *prefix MSCs*. We prefer to use this more general definition and still speak of MSCs.

**Example 6.1.** An example MSC is shown in the left half of Figure 2; this MSC is an execution of the CFM of Example 2.1, page 1032. All events of process C are drawn as circles on the *process line* C (and similarly for S). For simplicity, not all event labels are shown in the picture. Vertical edges denote the relation  $<_P$  (oriented downwards) and diagonal arrows denote the relation  $<_{msg}$ ; the label of such an arrow denotes the message contents transmitted (hence, e.g., the action performed by the source event of the first diagonal edge is C!S(0)). The partial order is therefore the reflexive and transitive closure of all the arrows. The last two events of type C!S are unmatched since the messages sent are not received.

A *linearization* of a  $\Sigma$ -labeled partial order  $\langle E, \leq, \lambda \rangle$  is the sequence of labels of a linear extension of this order. Thus, a linearization is a word over the alphabet  $\Sigma$  and the set  $Lin(M)$  of linearizations of an MSC  $M$  is a subset of  $\Sigma^*$ . For a set (or language) of partial orders  $X$ , we write  $Lin(X) = \bigcup_{M \in X} Lin(M)$ .

For an MSC, the relation between the partial order and its linearizations is very tight:

**Lemma 6.1.** *Any linearization of an MSC is a valid word and any valid word  $u$  is the linearization of some unique MSC that we denote  $msc(u)$ .*



**Figure 2.** A message sequence chart (left) and a symbolic representation of the set of executions of the CFM in Example 2.1 (right).

Let  $\mathcal{A} = \langle (\mathcal{A}_p)_{p \in \mathcal{P}}, \Sigma, F \rangle$  be a CFM and let  $\mathcal{M}(\mathcal{A})$  denote the set of MSCs that admit at least one linearization in  $\mathcal{L}(\mathcal{A})$ , i.e.,  $\mathcal{M}(\mathcal{A}) = \{ msc(u) \mid u \in \mathcal{L}(\mathcal{A}) \}$ . Since a CFM cannot distinguish between different linearizations of an MSC this set equals  $\{ M \mid Lin(M) \subseteq \mathcal{L}(\mathcal{A}) \}$ .

We next explain a more direct characterization of the set  $\mathcal{M}(\mathcal{A})$  of MSCs accepted by the CFM  $\mathcal{A} = \langle (\mathcal{A}_p)_{p \in \mathcal{P}}, \Sigma, F \rangle$ . Let  $M = \langle E, \leq, \lambda \rangle$  be an MSC. A *run of  $\mathcal{A}$  on  $M$*  is a mapping  $\rho : E \rightarrow \bigcup_{p \in \mathcal{P}} S_p$  labeling events of  $M$  by local states, such that for any  $p \in \mathcal{P}$  and any  $e, f \in E_p$  we have

- if  $e \leq_P f$ , then  $\rho(e) \xrightarrow{\lambda(f)}_p \rho(f)$  and
- if  $f = \min(E_p, \leq_P)$ , then  $s_p^0 \xrightarrow{\lambda(f)}_p \rho(f)$ .

For  $p \in \mathcal{P}$ , let  $s_p = \rho(\max(E_p, \leq_P))$  if  $E_p \neq \emptyset$  and the initial state  $s_p^0$  otherwise. Then the run  $\rho$  is *accepting* if the global state  $(s_p)_{p \in \mathcal{P}}$  belongs to the set of accepting states  $F$ . Now one can show that  $M \in \mathcal{M}(\mathcal{A})$  if and only if  $\mathcal{A}$  admits some accepting run  $\rho$  on  $M$ .

**Example 6.2.** The right part of Figure 2 is another abstract view of the executions of the CFM in Example 2.1, this time in a more succinct way using MSCs. The message labeled  $(01)^*0$  stands for any sequence of message arrows from C to S with contents  $0, 1, \dots, 0$ . The event  $C!S((10)^*)$  stands for a sequence of sends  $C!S$  with contents  $1, 0, \dots, 1, 0$ , similarly for  $C!S((01)^*(0 + \varepsilon))$ . The upper half of this picture can be iterated. Recall that process C and S alternate between states  $c_0, c_1$ , and  $s_0, s_1$  respectively. With this picture

it is rather easy to check that the set of reachable configurations with state  $(c_0, s_0)$  has channel contents in one of the following sets (cf. also Figure 1):

- $((01)^*, \varepsilon)$ : **C** and **S** both in upper half, either **S** before sending \$, or both at the end of upper half,
- $((10)^*1, \$)$ : **C** and **S** both in lower half, after **S** sends \$ (and before **C** receives it),
- $((10)^+(01)^*, \varepsilon)$ : **C** and **S** both in lower half, after **C** receives \$.

In the following, it is more convenient to consider the set  $\mathcal{M}(\mathcal{A})$  as the semantics of the CFM  $\mathcal{A}$ . Properties of the semantics of  $\mathcal{A}$  can therefore be expressed as properties of MSCs. The following two sections therefore introduce logical formalisms that can describe properties of MSCs.

## 6.2 Monadic second order logic

We fix supplies  $\text{Var}$  of *individual* and  $\text{VAR}$  of *set variables*. We will use the convention that lower case variables belong to  $\text{Var}$  and upper case variables to  $\text{VAR}$ . The set  $\text{MSO}$  of *monadic second order* (or  $\text{MSO}$ ) formulas is given by the following grammar:

$$\begin{aligned} \varphi ::= & \sigma(x) \mid x \prec_P y \mid x \prec_{msg} y \mid x \leq y \mid x = y \mid x \in X \mid \\ & \exists x \varphi \mid \exists X \varphi \mid \neg \varphi \mid \varphi \vee \varphi, \end{aligned}$$

where  $\sigma \in \Sigma$  is an action,  $x, y$  are individual variables from  $\text{Var}$ , and  $X \in \text{VAR}$  is a set variable. A formula is *first order* if it does not contain any subformula  $\exists X \psi$  for  $X \in \text{VAR}$ , it is *existential* if it is of the form  $\exists X_1 \exists X_2 \dots \exists X_n \varphi$  where  $\varphi$  is first order. We write  $\text{FO}$  and  $\text{EMSO}$  for the sets of first order and existential  $\text{MSO}$ -formulas.

Let  $M = \langle E, \leq, \lambda \rangle$  be an MSC,  $f : \text{Var} \rightarrow E$  and  $g : \text{VAR} \rightarrow 2^E$  be interpretations of variables, and  $\varphi$  a formula from  $\text{MSO}$ . We then define  $M, f, g \models \varphi$  (read “the MSC  $M$  satisfies the formula  $\varphi$  under the interpretations of variables  $f$  and  $g$ ”) by induction:

$$\begin{aligned} M, f, g \models \sigma(x) & \Leftrightarrow \lambda(f(x)) = \sigma, \\ M, f, g \models x \prec_P y & \Leftrightarrow f(x) \prec_P f(y), \\ M, f, g \models x \in X & \Leftrightarrow f(x) \in g(X), \\ M, f, g \models \exists x \varphi & \Leftrightarrow \text{there exists } v \in E \text{ with } M, f[x/v], g \models \varphi, \end{aligned}$$

and similarly for the formulas  $x \prec_{msg} y$ ,  $x \leq y$ ,  $x = y$ , and  $\exists X \varphi$ . Here,  $f[x/v]$  denotes the function  $\text{Var} \rightarrow E$  that sends  $x$  to  $v$  and agrees with  $f$  for all other variables. The semantics of  $\neg \varphi$  and  $\varphi_1 \vee \varphi_2$  are standard. For a subset  $R$  of  $\{\prec_P, \prec_{msg}, \leq\}$ , let  $\text{MSO}(R)$  be the restriction of the set of  $\text{MSO}$ -formulas to those that mention at most the binary relations from  $R$ .

As a first example, consider the following formula

$$\text{Succ} = x < y \wedge \bigvee_{p \in \mathcal{P}} (\Sigma_p(x) \wedge \Sigma_p(y) \wedge \forall z (\Sigma_p(z) \rightarrow z \leq x \vee y \leq z))$$

from  $\text{MSO}(\leq)$  with two free individual variables  $x$  and  $y$ , where  $\Sigma_p(x)$  is an abbreviation of  $\bigvee_{\sigma \in \Sigma_p} \sigma(x)$ . Then we have  $M, f, g \models \text{Succ}$  if and only if  $f(x) \prec_P f(y)$ . Hence any formula from  $\text{MSO}(\{\prec_P\} \cup R)$  can be translated into an equivalent formula from



$\text{MSO}(\{\leq\} \cup R)$ , independently from what the set  $R$  is. Note that the same holds if we only consider first order or existential formulas. But such an argument cannot work for  $<_{msg}$ , more precisely, there is a formula from  $\text{MSO}(<_P, <_{msg})$  without equivalent counterpart in  $\text{MSO}(\leq)$ . Such an example is the LTL formula from Example 5.1, that we rephrase in  $\text{MSO}(<_P, <_{msg})$  as

$$\varphi = \exists x, x', y, y' : x <_P x' \wedge y <_P y' \wedge x <_{msg} y \wedge c_p(x') \wedge c_q(y').$$

Suppose that  $\psi$  is an equivalent formula from  $\text{MSO}(\leq)$ , and let  $u$  be a valid word as in Example 5.1. Then  $\text{msc}(u)$  is linearly ordered by  $\leq$ , so we can interpret  $\psi$  directly on the word  $u$ . Let  $L$  be the *regular* set of all words over  $\Sigma$  that satisfy  $\psi$ . Then  $L$  contains in particular all valid words  $u$  such that  $\text{msc}(u)$  is linearly ordered and satisfies  $\varphi$ . Arguments as in Example 5.1 show that  $L$  also contains some valid word  $v$  with  $\text{msc}(v)$  linear and  $\text{msc}(v) \not\models \varphi$ . But  $v \in L$  and  $\text{msc}(v)$  linear imply  $\text{msc}(v) \models \psi$ , so  $\psi$  and  $\varphi$  cannot be equivalent.

Next consider the following formula  $\text{Leq}(x, y)$  from  $\text{MSO}(<_P, <_{msg})$ :

$$\exists X : (y \in X \wedge \forall z' (z' \in X \leftrightarrow (z' = x \vee \exists z \in X : z <_{msg} z' \vee z <_P z'))).$$

Then the formula starting with  $\forall z'$  holds in  $M$  (with respect to  $f$  and  $g$ ) if and only if  $g(X)$  is the set of nodes  $e \in E$  where  $f(x) \leq e$ . Hence  $\text{Leq}(x, y)$  holds if and only if  $f(x) \leq f(y)$ . As a consequence, any formula  $\varphi$  from  $\text{MSO}$  can be translated into an equivalent formula  $\bar{\varphi}$  from  $\text{MSO}(<_P, <_{msg})$ . Since the above formula uses an additional set variable,  $\bar{\varphi}$  is in general not first order or existential, even if  $\varphi$  is first order or existential, respectively.

### 6.3 Propositional dynamic logic

In this section, we introduce a temporal logic that can describe causal properties. This logic is adapted from classical Propositional Dynamic Logic [33]. Like the logic TLC [7], which is interpreted over Mazurkiewicz traces (cf. Section 9.1), our variant of PDL is interpreted directly on MSCs, and extends the logic  $\text{TLC}^-$  considered in [64].

The logic PDL has three types of formulas: local formulas express properties of single events in an MSC, path expressions express properties of sequences of events, and global formulas express properties of MSCs as a whole. As example, the local formula  $p!q(m)$  is true if and only if the current event sends the message  $m$  from  $p$  to  $q$ . The path expression  $(msg; proc)^*$  holds true of a sequence of events  $e_0, e_1, e_2, e_3, \dots, e_{2n}$  if and only if  $e_{2i} <_{msg} e_{2i+1} <_P e_{2i+2}$  for all  $0 \leq i < n$ . That is,  $e_0 < e_{2n}$  via a path that alternates between message arcs and process successor arcs. Then, the local formula  $\langle (msg; proc)^* \rangle p!q(m)$  holds if, from the current event, there is a path as described, leading to an event that sends message content  $m$  from  $p$  to  $q$ . Finally, the global formula  $E(q!p(n) \wedge \langle (msg; proc)^* \rangle p!q(m))$  holds if and only if the MSC contains some event labeled  $q!p(n)$  where the previous formula holds.

*Path expressions*  $\pi$  and *local formulas*  $\alpha$  are inductively (with  $\sigma \in \Sigma$  below):

$$\begin{aligned} \pi &::= proc \mid proc^{-1} \mid msg \mid msg^{-1} \mid \{\alpha\} \mid \pi; \pi \mid \pi + \pi \mid \pi^*, \\ \alpha &::= true \mid \sigma \mid \alpha \vee \alpha \mid \neg \alpha \mid \langle \pi \rangle \alpha, \end{aligned}$$

Local formulas express properties of single events in MSCs. To define the semantics of local formulas, let  $M = \langle E, \leq, \lambda \rangle$  be an MSC and  $e \in E$  an event from  $M$  with  $P(e) = p$ . Then  $M, e \models \sigma$  if and only if  $\lambda(e) = \sigma$ ;  $M, e \models \alpha_1 \vee \alpha_2$  and  $M, e \models \neg\alpha$  are defined canonically.

The semantics of modalities  $\langle \pi \rangle \alpha$  indexed by a path expression  $\pi$  is to reach via a program (path)  $\pi$  an event where the local formula  $\alpha$  is satisfied:

$$\begin{aligned} M, e \models \langle proc \rangle \alpha &\Leftrightarrow \text{there exists } e' \in E \text{ with } e \leq_P e' \text{ and } M, e' \models \alpha, \\ M, e \models \langle proc^{-1} \rangle \alpha &\Leftrightarrow \text{there exists } e' \in E \text{ with } e' \leq_P e \text{ and } M, e' \models \alpha, \\ M, e \models \langle msg \rangle \alpha &\Leftrightarrow \text{there exists } e' \in E \text{ with } e <_{msg} e' \text{ and } M, e' \models \alpha, \\ M, e \models \langle msg^{-1} \rangle \alpha &\Leftrightarrow \text{there exists } e' \in E \text{ with } e' <_{msg} e \text{ and } M, e' \models \alpha, \\ M, e \models \langle \{ \alpha \} \rangle \beta &\Leftrightarrow M, e \models \alpha \text{ and } M, e \models \beta, \\ M, e \models \langle \pi_1; \pi_2 \rangle \alpha &\Leftrightarrow M, e \models \langle \pi_1 \rangle \langle \pi_2 \rangle \alpha, \\ M, e \models \langle \pi_1 + \pi_2 \rangle \alpha &\Leftrightarrow M, e \models \langle \pi_1 \rangle \alpha \vee \langle \pi_2 \rangle \alpha, \\ M, e \models \langle \pi^* \rangle \alpha &\Leftrightarrow \text{there exists } n \geq 0 \text{ with } M, e \models (\langle \pi \rangle)^n \alpha. \end{aligned}$$

*Global formulas* are Boolean combinations of properties of the form “there exists an event satisfying the local formula  $\alpha$ ”. Their syntax is given by:

$$\varphi ::= E\alpha \mid A\alpha \mid \varphi \vee \varphi \mid \varphi \wedge \varphi,$$

where  $\alpha$  ranges over the set of local formulas. The semantics is defined by:

$$\begin{aligned} M \models E\alpha &\Leftrightarrow \text{there exists an event } e \text{ with } M, e \models \alpha, \\ M \models A\alpha &\Leftrightarrow M, e \models \alpha \text{ for all events } e. \end{aligned}$$

The definitions of  $M \models \varphi_1 \vee \varphi_2$  and  $M \models \varphi_1 \wedge \varphi_2$  are obvious.

Semantically, a local formula of the form  $\langle \{ \alpha \}; (proc + msg)^* \rangle \beta$  corresponds to the until construct  $\alpha \text{U} \beta$  in the temporal logic  $\text{TLC}^-$  [64], i.e.,  $\text{TLC}^-$  is a semantic fragment of PDL. In  $\text{TLC}^-$ , however, one cannot express properties such as “there is an odd number of messages from  $p$  to  $q$ ”, which is expressible in PDL by the global formula below.

**Example 6.3.** We use the usual abbreviations for Boolean and rational expressions, as well as  $p!q$  for  $\bigvee_{m \in \text{Msg}} p!q(m)$ . The local formula  $p!q \wedge \neg \langle (proc^{-1})^+ \rangle p!q$  is satisfied by the first event on process  $p$  of type  $p!q$ . Paths on process  $p$  with an even number of events of type  $p!q$  (ending with such an event, and not counting the first event) are described by the path formula  $[((proc; \{-p!q\})^*; proc; \{p!q\})^2]^*$ . The global formula  $A((p!q \wedge \neg \langle (proc^{-1})^+ \rangle p!q) \rightarrow [((proc; \{-p!q\})^*; proc; \{p!q\})^2]^* \neg \langle proc^+ \rangle p!q)$  says that the number of messages from  $p$  to  $q$  is odd.

## 7 Model-Checking

Since even reachability is an undecidable problem for CFMs, model-checking properties of such automata requires some restrictions. Given that one main reason for undecidability are the channels, we focus in this section on two variants of channel bounds. Notice

that channel bounds are well justified in practice, since any concrete implementation of a communication protocol will have to use channels of some given size. The existential version of channel bounds is optimistic and considers only those executions that can be rescheduled with bounded channels. The universal version is pessimistic and considers only those executions that, independently of the scheduling, use bounded channels. Although these two variants of bounds look rather similar, it is important to stress that certain communication networks (e.g., acyclic ones) have existentially, but no universally bounded channels (cf. Example 7.3).

## 7.1 Channel bounds

Let  $B \in \mathbb{N}$  be a fixed positive integer. A valid word  $u$  is *B-bounded* if, for every prefix  $v$  of  $u$  the difference between the number of *matched* events of type  $p!q$  and those of type  $q?p$  in  $v$  is bounded by  $B$ :

$$\min\{|\pi_{p,q}^1(v)|, n\} - |\pi_{p,q}^2(v)| \leq B, \quad \text{where } n = |\pi_{p,q}^2(u)|.$$

We write  $\Sigma^*|_B$  for the set of valid  $B$ -bounded words. It is not hard to see that this set is regular.

We explain now what “rescheduling a run” means. Let  $\sim$  be the least equivalence relation on the set of *valid* words such that  $u \sim v$  whenever  $u = xaby$  and  $v = xbay$  with  $x, y \in \Sigma^*$ ,  $a \in \Sigma_p$  and  $b \in \Sigma_q$  for two distinct processes  $p$  and  $q$ . We say that  $u$  can be rescheduled into  $v$  if  $u \sim v$ . For a valid word  $u$ , let  $[u]_{\sim}$  denote its equivalence class with respect to  $\sim$ , i.e., the set of valid words that can be rescheduled to  $u$ .

A valid word  $u$  is *universally B-bounded* if  $[u]_{\sim} \subseteq \Sigma^*|_B$  and it is *existentially B-bounded* if  $[u]_{\sim} \cap \Sigma^*|_B \neq \emptyset$ .

**Example 7.1.** The word  $u = C!S(0) C!S(1) C!S(0) S?C(0) S?C(1) C!S(1)$  is valid and 2-bounded (note that  $|\pi_{C,S}^2(u)| = 2$ ), but not 1-bounded because of the prefix  $C!S(0) C!S(1)$  that consists of (matched) sends, only. But we have

$$\begin{aligned} u &\sim C!S(0) C!S(1) S?C(0) C!S(0) S?C(1) C!S(1) \\ &\sim C!S(0) S?C(0) C!S(1) C!S(0) S?C(1) C!S(1) \\ &\sim C!S(0) S?C(0) C!S(1) S?C(1) C!S(0) C!S(1) \end{aligned}$$

and the final word is 1-bounded. Consequently,  $u$  is existentially 1-bounded. Since  $u$  is not 1-bounded, it cannot be universally 1-bounded, but the reader is invited to check that  $u$  is universally 2-bounded.

## 7.2 Bounded CFMs

As a preparation for model-checking, we first consider the control-state reachability problem relative to channel bounds:

- *Universally bounded control-state reachability problem:* Given a CFM  $\mathcal{A}$ , a bound  $B \in \mathbb{N}$ , and a global state  $s \in S$ , is there some universally  $B$ -bounded valid word that leads from the initial configuration to some configuration with global state  $s$ ?

- The *existentially bounded control-state reachability problem* is similar, one requires the valid word to be existentially  $B$ -bounded.

To solve the second of these problems, we modify the transition system associated with the CFM  $\mathcal{A}$  as follows: configurations consist of a global state, together with a word from  $\text{Msg}^{\leq B}$  or the special channel content  $\perp$  indicating that the channel is inactive, for each channel  $(p, q) \in \text{Ch}$ . The definition of transitions labeled  $p!q(m)$  from page 1032 is altered to

- if  $a = p!q(m)$  is a send action, then either  $w'_{p,q} = w_{p,q}m$  (message  $m$  is inserted into the channel from  $p$  to  $q$ ) or  $w'_{p,q} = \perp$  (the channel from  $p$  to  $q$  is inactive). As before,  $w'_{r,s} = w_{r,s}$  for all  $(r, s) \neq (p, q)$  (all other channels are unchanged).

In effect, send actions to channel  $(p, q)$  behave as before or, if the channel is inactive, are ignored. On the other hand, no read actions can be executed on inactive channels. The result is a finite automaton and in this automaton, some configuration with global state  $s$  can be reached if and only if the existentially bounded control-state reachability problem has an affirmative answer. Since the size of the finite automaton is exponential in  $B$  and  $|\mathcal{P}|$ , this is a PSPACE-algorithm (with  $B$  given in unary).

The universally bounded control-state reachability problem is slightly more involved. Here, one first observes that the set of universally  $B$ -bounded valid words is regular. Hence, it suffices to take the finite automaton from the previous paragraph, intersect its language with the language of universally  $B$ -bounded valid words, and check this intersection for emptiness. Thus, we have:

**Proposition 7.1.** *The universally and the existentially bounded control-state reachability problem are PSPACE-complete (with  $B$  given in unary).*

PSPACE-hardness is easily shown by reducing from the halting problem of linearly space bounded Turing machines: to simulate such a machine, one associates a process with each tape cell. At each instant, exactly one process is active, corresponding to the position of the head. This process can take over activity to one of its neighbors, thereby simulating the moves of the head. This CFM can be defined in such a way that its language consists of universally 1-bounded words, only.

A CFM  $\mathcal{A}$  is *universally  $B$ -bounded* if all words from its language  $\mathcal{L}(\mathcal{A})$  are universally  $B$ -bounded:  $\mathcal{L}(\mathcal{A}) \subseteq \Sigma^*|_B$ . Similarly, a CFM  $\mathcal{A}$  is *existentially  $B$ -bounded* if all words from its language  $\mathcal{L}(\mathcal{A})$  are existentially  $B$ -bounded. We call  $\mathcal{A}$  *universally* or *existentially bounded* if it is universally or existentially  $B$ -bounded for some  $B \in \mathbb{N}$ .

**Example 7.2.** The CFM in Example 2.1 (page 1032) is not universally bounded, however existentially 1-bounded. For instance, the word  $(\text{C!S}(0) \text{C!S}(1))^k (\text{S?C}(0) \text{S?C}(1))^k$  from its language is not  $2k - 1$ -bounded (and hence not universally  $2k - 1$ -bounded), but can be rescheduled into the 1-bounded word  $(\text{C!S}(0) \text{S?C}(0) \text{C!S}(1) \text{S?C}(1))^k$ .

**Example 7.3.** Assuming that the undirected graph underlying the communication network  $(\mathcal{P}, \text{Ch})$  is a forest, every CFM over such a network is existentially 1-bounded [44], but not universally bounded, in general. Actually, universal channel bounds can only be guaranteed if the network has suitable cycles. This is related to the notion of *loop-connected* MSC-graphs [61].

**Remark 7.2.** The above definition of channel bounds for CFMs is based on accepting runs. Alternatively, we can also define universal/existential bounds over all possible runs. This makes only a difference at the level of deciding such bounds, but not in the expressivity results presented later.

The next result shows that  $B$ -boundedness of CFMs is an undecidable problem, but becomes decidable assuming deadlock-freeness.

**Proposition 7.3** ([39]). (1) *Given a positive integer  $B$  and a deadlock-free CFM  $\mathcal{A}$ , it is decidable whether  $\mathcal{A}$  is universally (existentially, resp.)  $B$ -bounded. This problem is PSPACE-complete if the bound  $B$  is given in unary.*

(2) *The following problems are undecidable:*

- *Is a deterministic CFM universally (existentially, resp.)  $B$ -bounded (for a given bound  $B$ )?*
- *Is a deterministic and deadlock-free CFM universally (existentially, resp.) bounded?*

### 7.3 Model-checking with channel bounds

With universally bounded channels a decision procedure for the model-checking problem of CFMs w.r.t. e.g. regular specifications is straightforward, since such a machine is equivalent to a (exponentially larger) finite automaton. We can make this statement more precise, by observing that many natural decision problems about universally  $B$ -bounded CFMs are PSPACE-complete. This observation could be declared as another “rule of thumb”, that holds as well for models like 1-safe Petri nets [30]. As already mentioned, PSPACE-hardness follows from the simulation of linearly bounded automata by a universally 1-bounded CFM. For the upper bound, let us consider model-checking a universally  $B$ -bounded CFM  $\mathcal{A}$  against an LTL property, see Section 5. The idea is to consider words over the extended alphabet  $\Sigma \times \{0, \dots, B-1, \$\}$  whose projection to  $\Sigma^*$  belongs to the language  $\mathcal{L}(\mathcal{A})$  of  $\mathcal{A}$  and whose second component indexes actions as follows for each channel  $(p, q)$ :

- the occurrences of actions of type  $q?p$  are numbered modulo  $B$ ,
- the occurrences of *matched* actions of type  $p!q$  are numbered modulo  $B$ ,
- the occurrences of *unmatched* actions of type  $p!q$  are indexed by \$.

Then, in the LTL formula  $\varphi$ , replace every occurrence of  $a \in \Sigma$  by  $\bigvee_{i \in \{0, \dots, B-1, \$\}} (a, i)$  and of  $X_{msg}\psi$  by

$$\bigvee_{\substack{0 \leq i < B \\ (p,q) \in Ch}} (p!q, i) \wedge (\neg(q?p, i)) \text{U}((q?p, i) \wedge \psi).$$

Then a  $B$ -bounded word satisfies  $\varphi$  if and only if its “indexed version” satisfies this new formula. Hence we reduced model-checking of universally bounded CFMs against LTL-formulas to classical model-checking of LTL (without the operator  $X_{msg}$ ).

**Proposition 7.4.** *Model-checking universally  $B$ -bounded CFMs against LTL and CTL properties is PSPACE-complete (with  $B$  in unary encoding).*

The upper bound for CTL is shown by simulating universally  $B$ -bounded CFMs by 1-safe Petri nets and applying [30]. The simulation uses a place for each local state of  $\mathcal{A}$ , plus  $|Ch| \cdot |Msg| \cdot B$  places, one for each tuple from  $Ch \times Msg \times \{0, \dots, B-1\}$ . The set of transitions of the Petri net equals the set of local transitions of the CFM  $\mathcal{A}$ , and each transition of  $\mathcal{A}$  is simulated as expected.

We turn now to the model-checking problem for existentially  $B$ -bounded CFMs. As already shown in Section 5, model-checking CFMs against sequential specifications such as LTL, is undecidable. From the proof of Proposition 5.1, we can see that undecidability already applies to existentially 1-bounded CFMs. We can also observe that no *regular* set of words describes the same set of MSCs as the LTL property from Proposition 5.1. This is a key observation that will make model-checking existentially bounded CFMs possible using representatives:

**Definition 7.1.** Let  $K \subseteq \Sigma^*$ . A language  $R \subseteq \Sigma^*|_B$  is a *set of  $B$ -representatives for  $K$*  if

$$\{msc(w) \mid w \in K\} = \{msc(v) \mid v \in R\}.$$

The next proposition describes a class of regular properties for which model-checking existentially  $B$ -bounded CFMs is decidable:

**Proposition 7.5.** *Given  $B, B' \in \mathbb{N}$ , an existentially  $B$ -bounded CFM  $\mathcal{A}$  and a regular set  $R$  of  $B'$ -representatives for a property  $K \subseteq \Sigma^*$ , it is decidable whether  $\mathcal{L}(\mathcal{A}) \cap K \neq \emptyset$ . The complexity is PSPACE if  $R$  is described by an automaton, and  $B, B'$  are in unary encoding.*

For the proof of the proposition above we note first that an existentially  $B$ -bounded CFM is also existentially  $B'$ -bounded for  $B' \geq B$ . So we can assume that  $B' \leq B$ . Let  $\mathcal{B}$  be a finite automaton obtained from  $\mathcal{A}$  and  $B'$ , such that  $\mathcal{L}(\mathcal{B}) = \mathcal{L}(\mathcal{A}) \cap \Sigma^*|_{B'}$ . Notice that for every  $u, v \in \Sigma^*|_{B'}$  with  $msc(u) = msc(v)$ , we have  $u \in \mathcal{L}(\mathcal{B})$  if and only if  $v \in \mathcal{L}(\mathcal{B})$ . Therefore,  $\mathcal{L}(\mathcal{A}) \cap K \neq \emptyset$  is equivalent to  $\mathcal{L}(\mathcal{B}) \cap R \neq \emptyset$ . Since  $\mathcal{B}$ 's size is exponential in  $\mathcal{A}$  and  $B'$ , the claim follows.

A prime example of properties satisfying the assumption of Proposition 7.5 are partial order properties, as shown by the two results below:

**Proposition 7.6** ([53]). *Given  $B \in \mathbb{N}$  and a sentence  $\varphi$  from MSO, the language*

$$\{w \in \Sigma^*|_B \mid msc(w) \models \varphi\}$$

*is (effectively) regular.*

Model-checking a CFM  $\mathcal{A}$  against a formula  $\varphi$  of some *partial-order* logics is the question whether all  $M \in \mathcal{M}(\mathcal{A})$  satisfy  $\varphi$ .

**Corollary 7.7** ([38]). *Let  $B \in \mathbb{N}$  be given. Model-checking existentially  $B$ -bounded CFMs against MSO is decidable.*

The decidable problem above is known to be of non-elementary complexity. The situation changes to the better if we consider the temporal logic PDL.

**Proposition 7.8** ([14, 58]). *Given  $B \in \mathbb{N}$  and a PDL sentence  $\varphi$ , one can construct a finite automaton of size exponential in  $\varphi$ ,  $B$ , and the number of channels, whose language is a set of  $B$ -representatives of  $\{u \in \Sigma^* \mid \text{msc}(u) \models \varphi\}$ .*

The original proof from [14] (for a restricted version of the logic) goes via realizations, see Section 8.3. An alternative proof [58] translates PDL-formulas into automata that walk along paths in an MSC, using then a translation of such automata into alternating, 2-way automata on words (only this second step uses the bound  $B$ ). Alternatively, Proposition 7.8 can be shown by translating a PDL sentence  $\varphi$  over MSCs into an ETL sentence  $\tilde{\varphi}$  over  $\Sigma^*_B$ . The logic ETL is an extension of LTL by regular operators described by finite automata [70]. Besides atomic statements  $a \in \Sigma \cup \{p!q, q?p \mid (p, q) \in Ch\}$  and the Boolean connectives, it uses operators of the kind  $\mathcal{B}(\varphi_1, \dots, \varphi_n)$ , where  $\mathcal{B}$  is an automaton over the alphabet  $\{a_1, \dots, a_n\}$  and  $\varphi_i$  are ETL formulas. Formulas of this logic are evaluated over positions  $i$  in sequences  $w \in \Sigma^*$ . Then  $w, i \models \mathcal{B}(\varphi_1, \dots, \varphi_n)$  if there is a word  $a_{i_1}a_{i_2} \dots a_{i_m} \in \mathcal{L}(\mathcal{B})$  such that  $w, i + j \models \varphi_{i_j}$  for all  $1 \leq j \leq m$ . For the simulation we also need the 2-way version of ETL from [47]. The latter paper shows how to translate ETL formulas into Büchi automata with an exponential blow-up, even with operators using 2-way, alternating Büchi automata.

We sketch the translation from PDL to ETL. Given a global PDL formula  $\alpha$ , we define  $\tilde{\alpha}$  in such a way that for every MSC  $M$  and every  $B$ -bounded linearization  $w$  of  $M$ , we have  $M, e \models \alpha$  if and only if  $w, i \models \tilde{\alpha}$ , where  $i$  is the position of  $w$  corresponding to event  $e$  in  $M$ . As already observed for LTL with  $X_{\text{msg}}$ , one can express the  $\text{msg}$  relation by usual LTL over  $B$ -bounded words with the modulo counters, i.e., over  $B$ -bounded words on the alphabet  $\Sigma \times \{0, \dots, B-1, \$\}$ . For instance, the PDL formula  $\langle \text{msg} \rangle \alpha$  translates to  $\bigvee_{(p,q) \in Ch, 0 \leq i < B} (p!q, i) \wedge \neg(q?p, i) \text{U}(q?p, i) \tilde{\alpha}$ . It is easy to encode this LTL formula into an ETL formula with a non-deterministic finite automaton of size linear in  $B$ , that accepts at a position satisfying  $\tilde{\alpha}$ . More generally, we can proceed as previously described for every formula  $\langle \pi \rangle \alpha$ , by using a non-deterministic automaton of polynomial size in  $\pi$ ,  $B$  and  $\mathcal{P}$ , over local formulas  $\alpha'$  occurring as atoms in  $\pi$ . Global formulas can be handled by the LTL operators F and G. Note that  $|\tilde{\alpha}| \in O(B \cdot |Ch| \cdot |\alpha|)$ .

**Corollary 7.9** ([14, 58]). *For given  $B \in \mathbb{N}$ , model-checking existentially  $B$ -bounded CFMs against PDL is decidable in PSPACE.*

**Further reading.** The paper [58] considers non-terminating CFMs and infinite MSCs. It then makes sense to also consider formulas of the form  $\langle \pi^\omega \rangle \top$  that hold at some event if it is the starting point of some infinite path that can be split into infinitely many sub-paths each conforming to the path expression  $\pi$ . A comprehensive study of complexity of the model checking problem of CFMs against different temporal logics can be found in [59], including MSO-definable temporal logics, PDL, and more concrete logics.

Timed CFMs were considered in [22, 46, 3]. Both [22, 46] study local timed automata with communication channels, whereas [3] studies a timed extension of communicating automata based on event clocks. The paper [22] defined the semantics of the automata in terms of timed MSCs and proposed a solution to an MSC-matching problem using the tool UPPAAL. In [3], a timed extension of Büchi's equivalence between automata and MSO was obtained, and it was shown that the logic is decidable over existentially

bounded channels.

Systems of pushdown automata communicating via FIFO-channels have been considered in [49, 44]. In that setting, bounding the channels does not suffice to make control-state reachability decidable. Instead, [49] shows decidability for acyclic networks and [44] characterizes networks with a decidable existentially channel-bounded control-state reachability problem. More recently, [25, 26, 23, 24] considers a more general model than CFMs that includes communication and pushdown storage, and proposes the notion of split-width, that guarantees the decidability of the model-checking problem.

## Part III

# Realizability

This part discusses the synthesis problem, that is, how to convert a specification into an “equivalent” CFM. As already mentioned, we assume that our systems are closed (i.e., without environment), and the problem we look at is called *realizability*. We present two solutions for realizability, one for unbounded channels (Section 8), and another one for universally- and existentially-bounded channels (Section 9.3). Both solutions use additional message contents, and we will see in Section 8.1 that this is indeed required. We formalize this variant of “equivalence” between specifications and realizations below.

Let  $\Sigma$  be some communication alphabet on the network  $(\mathcal{P}, Ch)$  with message contents from  $Msg$ . Furthermore, let  $\Sigma'$  be the analogous communication alphabet with message contents from  $Msg'$  and let  $\pi : Msg' \rightarrow Msg$  be a function. This function is extended naturally to a homomorphism  $\Sigma'^* \rightarrow \Sigma^*$  by setting  $\pi(p!q(m')) = p!q(\pi(m'))$  and similarly for receive actions.

**Definition 7.2.** A specification  $K \subseteq \Sigma^*$  is *realizable* if there exists some CFM  $\mathcal{A}$  with communication alphabet  $\Sigma'$  and a projection  $\pi$  from its message contents  $Msg'$  to  $Msg$ , such that  $\pi(\mathcal{L}(\mathcal{A})) = K$ .

*Realizability problem:* Given a specification  $K \subseteq \Sigma^*$ , is  $K$  realizable? If so, produce a CFM that realizes  $K$ .

A set of MSCs  $X$  is called *realizable*, if  $Lin(X)$  is so.

## 8 Realizability with Unrestricted Channels

### 8.1 Local acceptance and sequential specifications

Let  $\Sigma$  be some communication alphabet and  $K \subseteq \Sigma^*$  a language (the specification). The question in this section is whether  $K$  is the language of some CFM  $\mathcal{A}$ . This is a more restrictive question than the one in Definition 7.2, since the message alphabet is already defined by the specification. However, we will see that this question is undecidable. The



reason for undecidability is – once again – the incompatibility between the specification, which is sequential-like, and the concurrent target realization.

A necessary condition for  $K \subseteq \Sigma^*$  being the language of a CFM is that  $K$  must be closed under rescheduling:  $K \subseteq \Sigma^*$  is called *closed*, if for all valid words  $u \sim v$ , we have  $u \in K$  if and only if  $v \in K$  (see page 1047 for the definition of the rescheduling relation  $\sim$ ).

Note that the language of any CFM is the union of finitely many languages of CFMs with only one global accepting state. A slightly more general class is that of CFMs with local acceptance: A CFM  $\mathcal{A} = \langle (\mathcal{A}_p)_{p \in \mathcal{P}}, \Sigma, F \rangle$  has *local acceptance* if there are sets  $F_p \subseteq S_p$  for  $p \in \mathcal{P}$  such that  $F = \prod_{p \in \mathcal{P}} F_p$ .

**Proposition 8.1.** *A regular specification  $K \subseteq \Sigma^*$  is the language of some CFM with local acceptance if and only if it satisfies the conditions below:*

- (1)  $K$  is closed,
- (2)  $K$  consists of valid words, only, and
- (3) every valid word  $v \in \Sigma^*$  such that  $\pi_p(v) \in \pi_p(K)$ , for every  $p \in \mathcal{P}$ , belongs to  $K$  where  $\pi_p : \Sigma^* \rightarrow \Sigma_p^*$  is the natural projection homomorphism.

In the situation of condition (3), [5, 6] use the terminology “the valid word  $v$  is weakly implied by  $K$ ”. For the proof of Proposition 8.1, let us assume the existence of some CFM  $\mathcal{A}$  with local acceptance, such that  $\mathcal{L}(\mathcal{A}) = K$ . Even without local acceptance, we immediately get (1) and (2). Now consider valid words  $v \in \Sigma^*$  and  $v_p \in K$  such that  $\pi_p(v) = \pi_p(v_p)$  for every process  $p \in \mathcal{P}$ . It is not difficult to construct an accepting run of  $\mathcal{A}$  on  $v$ , from some accepting runs of  $\mathcal{A}$  on the words  $v_p$ : combining the  $p$ -projection of a run on  $v_p$ , over  $p \in \mathcal{P}$ , to a run on  $v$  is possible since the communication alphabet is the same, and acceptance is guaranteed by the local final states. Conversely, suppose  $K$  is regular and satisfies the conditions above. Then we can construct a CFM whose process  $p$  executes the words from  $\pi_p(K)$  (for all  $p \in \mathcal{P}$ ). This CFM has local acceptance and accepts  $K$ .

If  $\mathcal{A}$  is a finite automaton over  $\Sigma$ , then for deciding whether its language consists of valid words only, it suffices to check two conditions: the first requires that there is no  $uv \in \mathcal{L}(\mathcal{A})$  with  $|u|_{p!q} < |u|_{q?p}$ . The second one says that there is no  $u p!q(m) v q?p(m') w \in \mathcal{L}(\mathcal{A})$  with  $m \neq m'$  and  $|u|_{p!q} = |uv|_{q?p}$ . Both conditions can be decided e.g. using a 1-counter automaton. If this test succeeds, we can decide easily whether the language of  $\mathcal{A}$  is closed. Hence, two of the three conditions from Proposition 8.1 are decidable, but the last one is not:

**Theorem 8.2 ([6]).** *It is undecidable whether a regular (even LTL) specification  $K \subseteq \Sigma^*$  is the language of some CFM with local acceptance.*

In [6], the theorem is only stated for regular specifications, but one can check that the reduction also works for LTL specifications (where the additional modality  $X_{msg}$  is not used). The proof relies on condition (3) from Proposition 8.1 and therefore on the local acceptance. It seems open whether the class of regular languages of CFMs is decidable. We should also stress that an important argument in the undecidability proof above is that the specification is complete, i.e., that we ask for a CFM with *language*  $K$  as opposed to

the realizability problem where additional message contents are allowed. In Section 9.3 we will see that the realizability problem for regular languages is decidable.

## 8.2 MSO specifications

The theorem below is an extension of the well-known Büchi correspondence between regular and EMSO-definable sets of words (or trees, Mazurkiewicz traces, etc), to MSCs without any channel restriction. This characterization should be compared with those provided in Section 9.3 for universally and existentially bounded CFMs. The main difference is that for automata with universal/existential channel bounds, one can use more expressive logics.

**Theorem 8.3** ([15]). *For  $K \subseteq \Sigma^*$ , the following are equivalent:*

- (1)  *$K$  is realizable.*
- (2)  *$\text{msc}(K)$  is the language of some EMSO( $\leq_P, <_{\text{msg}}$ ) formula,  $K$  is closed and consists of valid words, only.*

The proof of the implication (1)  $\rightarrow$  (2) is very similar to the “usual” proof, using the partial order definition of CFM runs given at the end of Section 6.1. The existence of the mapping  $\rho: E \rightarrow \bigcup_{p \in \mathcal{P}} S_p$  is expressed by existential set quantifiers. Additional existential set quantifiers are used to express the message contents sent and received by the CFM. The message predicate  $<_{\text{msg}}$  is used for guaranteeing the consistency of these message contents. The local consistency of  $\rho$  is expressed using the successor predicate  $\leq_P$ .

The “usual” proof of the converse implication proceeds by induction on the structure of the formula. This requires the class of realizable specifications to be closed under complementation. But [15] shows that this is not the case! Therefore, the proof from [15] proceeds very differently (it is inspired by W. Thomas’ graph acceptors [68]): first, it uses Hanf’s theorem [42] from model theory to transform the first order kernel of the given formula into a specific normal form. Intuitively, a formula in this normal form is a positive Boolean combination of statements of the form “there are at most/at least  $k$  nodes of the MSC whose neighborhood of radius  $r$  looks as follows ...”. Then it provides an explicit CFM that computes, for every node of the MSC, its neighborhood of radius  $r$ . This handles the first order kernel of the formula, the outermost existential quantification is dealt with by a projection, as in Büchi’s proof.

The use of Hanf’s theorem also explains why this proof only works if we exclude the binary relation  $\leq$  from formulas: with only  $<_{\text{msg}}$  and  $\leq_P$ , any node has at most three neighbors and therefore the size of neighborhoods of radius  $r$  is at most  $3^r$ . Without such a finite bound, Hanf’s theorem is not applicable.

The proof of the implication (2)  $\rightarrow$  (1) is not only effective, but it even constructs a CFM whose size is elementary (although multiply exponential) in that of the formula [12, 13]. The reason is, again, the restriction of formulas to those mentioning only  $<_{\text{msg}}$  and  $\leq_P$ , but not  $\leq$ .

It seems not to be known whether Theorem 8.3 also holds for EMSO, i.e., whether we could allow the order relation  $\leq$  in formulas (this would necessarily result in a non-elementary CFM [60]).

### 8.3 PDL specifications

In general, the interest in temporal logics is their rather good complexity compared with MSO. The next theorem complies with this observation, in that it provides an exponential-time solution for the realizability problem for restricted PDL specifications.

In the following, we restrict PDL in such a way that no path expression contains both, forward modalities *proc* or *msg* and backward modalities *proc*<sup>-1</sup> or *msg*<sup>-1</sup>. More formally *forward* and *backward path expressions*  $\pi_+$  and  $\pi_-$  and *local formulas*  $\alpha$  are defined by simultaneous induction:

$$\begin{aligned}\pi_+ &::= \text{proc} \mid \text{msg} \mid \{\alpha\} \mid \pi_+; \pi_+ \mid \pi_+ + \pi_+ \mid \pi_+^*, \\ \pi_- &::= \text{proc}^{-1} \mid \text{msg}^{-1} \mid \{\alpha\} \mid \pi_-; \pi_- \mid \pi_- + \pi_- \mid \pi_-^*, \\ \alpha &::= \text{true} \mid \sigma \mid \alpha \vee \alpha \mid \neg\alpha \mid \langle \pi_+ \rangle \alpha \mid \langle \pi_- \rangle \alpha,\end{aligned}$$

where  $\sigma$  ranges over the alphabet  $\Sigma$ . Global formulas of rPDL are then defined analogously to those of PDL.

**Theorem 8.4** ([14]). *rPDL specifications are realizable in exponential time. More precisely, from a global rPDL formula  $\varphi$ , one can compute in exponential time a CFM  $\mathcal{A}$  realizing  $\varphi$ .*

The proof follows the inductive approach from [34, 35]: Let  $\alpha_{n+1}$  be a local formula of PDL whose top-most local subformulas are  $\alpha_1, \dots, \alpha_n$ . Then one builds a CFM that accepts MSCs whose events carry  $n + 1$  additional bits (one for each top-most subformula and the last one for the formula  $\alpha_{n+1}$ ). Let  $M = (E, \leq, \lambda)$  be an MSC and, for  $1 \leq i \leq n + 1$ , let  $X_i \subseteq E$  be the set of events where bit  $i$  is set to 1. Let furthermore  $\alpha'_{n+1}$  be obtained from  $\alpha_{n+1}$  by replacing every top-most occurrence of  $\alpha_i$  by the informal statement “bit  $i$  is set to 1” (for  $1 \leq i \leq n$ ). Then the CFM accepts  $(M, X_1, \dots, X_n, X_{n+1})$  if and only if  $X_{n+1} = \{e \in E \mid M, e \models \alpha'_{n+1}\}$ . For example, if  $\alpha_2 = \neg\alpha_1$ , then the CFM accepts MSCs  $M$  with two additional bits per node that have to be distinct at every node. The crucial point is that this CFM does not depend on the formulas  $\alpha_1, \dots, \alpha_n$ , but only on the outermost operator of  $\alpha_{n+1}$ . The construction is rather obvious for the operators of the form  $\sigma \in \Sigma$  (without subformulas),  $\vee$  and  $\langle \{.\} \rangle$ . (with two subformulas), and  $\neg$ ,  $\langle \text{proc} \rangle$ . and the like (with one subformula). But also the path formulas  $\alpha_2 = \langle \pi_+ \rangle \alpha_1$  and  $\langle \pi_- \rangle^{-1} \alpha_1$  cause no deep problem (we discuss the forward-path formula, only). At each and every event, one starts a copy of a finite automaton accepting the word language that is described by  $\pi_+$ . Depending on whether the event was claimed to satisfy or not to satisfy  $\alpha_2$ , this automaton is simulated along one or all outgoing edges and is required to accept or not to accept at an event satisfying  $\alpha_1$ .<sup>2</sup>

Combining Theorems 8.3 and 8.4, we infer that any rPDL formula can be translated into an equivalent formula from EMSO( $\leq_P, <_{msg}$ ). We do not know whether this is also possible for full PDL, where path expressions can mix forward and backward modalities. The converse translation, from EMSO( $\leq_P, <_{msg}$ ) to rPDL, is not always possible: Since the class of realizable MSC-languages is not closed under complementation [15], Theorem 8.3 implies that EMSO( $\leq_P, <_{msg}$ ) is (semantically) not closed under complementation. But the set of PDL formulas is closed under negation (using the usual de

<sup>2</sup>The main complication in [14] stems from the more general setting of infinite runs considered there.

Morgan rules). Hence, PDL is a proper fragment of  $\text{EMSO}(\prec_P, \prec_{msg})$ , but no non trivial description of this fragment is known. We do not even know whether the set of formulas from  $\text{EMSO}(\prec_P, msg)$  that can be translated into PDL is decidable.

## 9 Channel Bounds, Mazurkiewicz Traces and Realizability

This section presents a detailed study of CFMs with universal and existential channel bounds, extending the notions introduced in Section 7.1. The ultimate goal is to show a Büchi-like equivalence between automata and MSO in the setting of CFMs with channel bounds. In other words, one obtains a solution for the realizability problem for regular specifications (universal bounds), but also some non-regular ones (existential bounds). We start with a brief survey on the well-studied partial order model of Mazurkiewicz traces, then present the link between traces and channel bounds, and finally conclude with the realizability theorems.

### 9.1 Mazurkiewicz traces

Trace monoids [57], known for a long time in combinatorics as partially commutative monoids, were introduced in computer science in the late seventies by A. Mazurkiewicz for describing the behavior of Petri nets. Their essential feature is to express the semantics of a concurrent system by a (static) relation of independence between actions. Formally, a *trace alphabet* is a pair  $(A, I)$  consisting of an alphabet  $A$  and a symmetric and irreflexive relation  $I \subseteq A^2$ . The relation  $I$  will be referred to as the *independence relation*; its complement  $D = A^2 \setminus I$  is the *dependence relation*.

A *Mazurkiewicz trace* is a finite  $A$ -labeled partial order  $\langle E, \leq, \lambda \rangle$  – up to isomorphism – with the labeling  $\lambda : E \rightarrow A$  satisfying both conditions below, for any events  $e, f \in E$ :

- if  $e$  is an immediate predecessor of  $f$  (denoted as  $e < f$ ), then  $(\lambda(e), \lambda(f)) \in D$ ,
- if  $e$  and  $f$  are incomparable, then  $(\lambda(e), \lambda(f)) \in I$ .

Automata and logics over Mazurkiewicz traces have yielded many nice results, generalizing results for word languages. One prominent example is the Kleene-Büchi theorem stating the equivalence between automata, monadic second order logic and regular expressions, which was generalized to languages of finite and infinite traces (see [28] for a collection of surveys on trace theory). A second example is the equivalence between first order logic and linear temporal logics over traces [67, 27]. However, the deepest result of trace theory is undoubtedly the construction of deterministic *Zielonka automata* (also known as asynchronous automata) from finite automata [71]. It is this particular result which will be used for constructing CFMs with bounded channels in Section 9.3. For sake of completeness we introduce these automata and Zielonka's theorem in the rest of this section.

A *Zielonka automaton* over a trace alphabet  $(A, I)$  is a product of local automata synchronizing over common actions. The trace alphabet is described by a distribution

of the actions on processes, given by a function  $dom : A \rightarrow (2^{\mathcal{P}} \setminus \{\emptyset\})$  associating each action  $a \in A$  with a (non-empty) set of processes  $dom(a)$ . Moreover,  $(a, b) \in I$  if and only if the respective domains are disjoint, i.e.,  $dom(a) \cap dom(b) = \emptyset$ . A Zielonka automaton  $\mathcal{A} = \langle (S_p)_{p \in \mathcal{P}}, (\delta_a)_{a \in A}, s^0, F \rangle$  consists of a finite set of states  $S_p$  for each process  $p$ , a transition relation  $\delta_a \subseteq (\prod_{p \in dom(a)} S_p)^2$  for each letter, an initial state  $s^0 \in \prod_{p \in \mathcal{P}} S_p$  and a set of final states  $F \subseteq \prod_{p \in \mathcal{P}} S_p$ . The language of such an automaton is defined as the language of the finite automaton  $\mathcal{B}$  with set of states  $S = \prod_{p \in \mathcal{P}} S_p$  and transition relation  $\delta \subseteq S \times A \times S$  given by  $(s, a, s') \in \delta$  if  $((s_p)_{p \in dom(a)}, (s'_p)_{p \in dom(a)}) \in \delta_a$  and  $s'_q = s_q$  for all  $q \notin dom(a)$ . The automaton is deterministic if  $\delta_a$  is a function from  $\prod_{p \in dom(a)} S_p$  to  $\prod_{p \in dom(a)} S_p$ , for each  $a$ . The language  $L = \mathcal{L}(\mathcal{A})$  of such an automaton is  $I$ -closed:  $uabv \in L$  implies  $ubav \in L$ , for every  $u, v \in A^*$ ,  $(a, b) \in I$ .

**Theorem 9.1** ([71]). *A language  $L \subseteq A^*$  is  $I$ -closed and regular if and only if it is accepted by a deterministic Zielonka automaton.*

In terms of complexity, the original construction of Zielonka has been improved stepwise over the years: currently the best upper bound on the overall number of local states of the Zielonka automaton is  $O(2^{|\mathcal{P}|^4} \cdot |\mathcal{A}|^{|\mathcal{P}|^2})$ , with  $\mathcal{A}$  the minimal deterministic automaton of  $L$  [37]. Surprisingly, an exponential lower bound in the number of processes is known only for a subclass of Zielonka automata [37]. When  $L$  is given by a non-deterministic automaton  $\mathcal{A}$ , then [40] gives an asymptotically optimal construction which is simply exponential in both  $|\mathcal{A}|$  and  $|\mathcal{P}|$ .

## 9.2 Bounds and traces

Recall the definition of a  $B$ -bounded valid word from Section 7.1, as well as the fact that linearizations of MSCs are necessarily valid. It makes therefore sense to say that an MSC  $M$  is existentially (universally, resp.)  $B$ -bounded if and only if one (all, resp.) linearizations of  $M$  are  $B$ -bounded.

This matches the definition of existentially or universally  $B$ -bounded valid words: an MSC  $M$  is existentially or universally  $B$ -bounded if and only if  $Lin(M)$  consists of existentially or universally  $B$ -bounded words, only. A set of MSCs is existentially or universally  $B$ -bounded if and only if each of its members is so.

We describe now encodings of universally (existentially, resp.)  $B$ -bounded MSCs into traces, following [48] and [38]. The trace alphabet  $(A, I)$  is given by  $A = \Sigma \times \{0, \dots, B-1, \$\}$  and the following dependence relation  $D \subseteq A \times A$ : let  $(a, i)D(b, j)$  if either  $a, b \in \Sigma_p$  for some  $p \in \mathcal{P}$ , or  $i = j$  and  $\{a, b\} = \{p!q(m), q?p(m)\}$  for some  $(p, q) \in Ch, m \in Msg$ . Clearly,  $I = A^2 \setminus D$  is symmetric and irreflexive, hence  $(A, I)$  is a trace alphabet.


Since the universally bounded case is easier, we start with this case as a preparation. The encoding  $tr(M)$  of a *universally*  $B$ -bounded MSC  $M = \langle E, \leq, \lambda \rangle$  is obtained similarly to the indexing on page 1049: If  $e \in E$ , then  $\lambda_B(e) = (\lambda(e), n)$  such that

- $e$  is of type  $q?p$  and  $n = |\{e' < e \mid e' \text{ is of type } q?p\}| \bmod B$ ,
- $e$  is of type  $p!q$ , there is some  $f \in E$  with  $(e, f) \in msg$ , and  $n = |\{e' < e \mid e' \text{ is of type } p!q\}| \bmod B$ , or

- $e$  is of type  $p!q$ , there is no  $f \in E$  with  $(e, f) \in msg$ , and  $n = \$$ .

This yields a trace: if  $e$  is an immediate predecessor of  $f$ , then they belong to the same process or they are matching send and receive (in which case  $\lambda_B(e) = (p!q(m), n)$  and  $\lambda_B(f) = (q?p(m), n)$  are dependent). If  $\lambda_B(e)$  and  $\lambda_B(f)$  are dependent, then they belong to the same process or they equal  $(p!q(m), n)$  and  $(q?p(m), n)$  or vice versa. But then, universal  $B$ -boundedness of  $M$  implies that  $e$  and  $f$  are comparable.

For an *existentially  $B$ -bounded* MSC  $M = \langle E, \leq, \lambda \rangle$  we need to introduce an additional binary relation  $rev_B$  on events: let  $(r, s') \in rev_B$  if, for some  $(p, q) \in Ch$ : event  $r$  has type  $q?p$  with matching receive  $s$ , event  $s' > s$  has type  $p!q$  and  $|\{s'' \mid s < s'' \leq s' \text{ and } s'' \text{ has type } p!q\}| = B$ . That is, the relation  $rev_B$  matches a receive event  $r$  with the  $B$ -th send of the same type after the matching send of  $r$ . Let  $\preceq_B$  be the reflexive-transitive closure of  $\leq \cup rev_B$  and let  $tr(M) = \langle E, \preceq_B, \lambda_B \rangle$ , with  $\lambda_B$  as above. It is easy to check that we have  $rev_B \subseteq \leq$  if and only if  $M$  is universally  $B$ -bounded, so the two definitions of  $tr(M)$  coincide in this case. Figure 3 shows an example for a trace  $tr(M)$  associated with an existentially 2-bounded MSC. The dashed edges are  $rev_2$ -edges.



**Figure 3.** Trace  $tr(M)$  associated with an existentially 2-bounded MSC. Diagonal edges are messages, dashed edges are  $rev_2$ -edges. On each vertical line, events are ordered from top to bottom.

The relation  $\preceq_B$  is, in general, not a partial order since it could contain a cycle.

**Lemma 9.2** ([52, 38]). *Let  $M$  be an MSC.*

- $M$  is existentially  $B$ -bounded if and only if  $\preceq_B$  is acyclic, if and only if  $tr(M)$  is a trace over the alphabet  $(A, I)$ .
- If  $M$  is existentially  $B$ -bounded, then  $Lin(M) \cap \Sigma^*|_B = Lin(proj_\Sigma(tr(M)))$ , where  $proj_\Sigma$  is the projection of the labeling on the first component of  $A$ .
- If  $M$  is universally  $B$ -bounded, then  $M = proj_\Sigma(tr(M))$ .

We summarize the connection between languages of CFMs with channel bounds and trace languages in the proposition below. We say that  $K \subseteq \Sigma^*$  is *B-closed*<sup>3</sup>, if for every valid words  $u, v \in \Sigma^*|_B$  with  $u \sim v$ , it holds that  $u \in K$  if and only if  $v \in K$ . Closed regular languages of valid words are those of universally bounded CFMs (cf. Theorem 9.5), whereas closures of *B-closed* regular languages of valid words are those of existentially *B-bounded* CFMs (cf. Theorem 9.7).

It can be decided in polynomial space whether  $L$  is *B-closed*, for  $L$  given by an automaton and  $B$  in unary encoding.

**Proposition 9.3** ([48, 38]). *Let  $B \in \mathbb{N}$  and let  $K \subseteq \Sigma^*|_B$  be a regular language of valid  $B$ -bounded words that is  $B$ -closed. Set  $L = \text{Lin}(\text{tr}(\text{msc}(K))) \subseteq A^*$ .*

*Then  $L$  is a regular and  $I$ -closed language over  $(A, I)$ . Moreover,  $K = \text{proj}_\Sigma(L)$ .*

**Remark 9.4.** Let  $\mathcal{A}$  be a universally *B-bounded* CFM. Then  $\mathcal{L}(\mathcal{A})$  is a closed and regular set of valid *B-bounded* words. Hence, by the proposition above, there exists a regular and *I-closed* language  $L$  with  $\mathcal{L}(\mathcal{A}) = \text{proj}_\Sigma(L)$ .

Next let  $\mathcal{A}$  be an existentially *B-bounded* CFM. Then  $\mathcal{L}(\mathcal{A})$  is a closed set of valid words, but  $\mathcal{L}(\mathcal{A})$  is in general neither regular nor a set of *B-bounded* words. But the set  $\mathcal{L}(\mathcal{A}) \cap \Sigma^*|_B$  of *B-bounded* words from  $\mathcal{L}(\mathcal{A})$  is *B-closed* and regular. Hence, by the proposition above, there is a regular *I-closed* language  $L$  with  $K = \text{proj}_\Sigma(L)$ . Since  $\mathcal{M}(\mathcal{A})$  is existentially *B-bounded*, we get  $\mathcal{M}(\mathcal{A}) = \text{msc}(K) = \text{msc}(\text{proj}_\Sigma(L))$ .

To see that  $\mathcal{L}(\mathcal{A}) = \text{proj}_\Sigma(L)$  cannot be achieved for  $\mathcal{A}$  existentially 2-bounded, consider a CFM  $\mathcal{A}$  with two processes,  $p, q$ . Each of these processes has a unique state, and  $p$  ( $q$ , resp.) loops on its state by sending the message  $m$  to  $q$  (receiving from  $p$ , resp.). Then  $\mathcal{L}(\mathcal{A}) \cap (p!q(m)^* q?p(m)^*) = \{p!q(m)^k q?p(m)^\ell \mid k \leq \ell\}$ , hence  $\mathcal{L}(\mathcal{A})$  is not regular, consequently, it cannot be the projection of any regular language  $L$ .

### 9.3 Realizability with bounded channels

As already mentioned, universal or existential channel bounds give a robust framework for MSC languages, in which we can state the Büchi-like equivalence between automata and monadic second order logic, cf. Theorems 9.5 and 9.7. In particular, the results show that realizability for specifications based on regular sets or on full monadic second order logic (with any of the relations  $\leq, \leq_P, <_{\text{msg}}$ ) is decidable. For regular specifications, realizability can be solved rather efficiently in exponential time, see Corollary 9.6 and 9.8.

A word  $w \in \Sigma^*$  is *complete* if  $|w|_{p!q(m)} = |w|_{q?p(m)}$  for all channels  $(p, q)$  and all message contents  $m$ . A language  $K \subseteq \Sigma^*$  is *message deterministic* if, whenever  $u p!q(m)$  and  $u p!q(n)$  are prefixes of some elements of  $K$ , then  $m = n$ . Note that the language of every deterministic CFM is message deterministic. It follows that only message deterministic languages can be realized by deterministic CFMs. A language is *deterministically* realizable, if it can be realized by a deterministic CFM.

**Theorem 9.5** ([43]). *Let  $B \in \mathbb{N}$  and let  $K \subseteq \Sigma^*|_B$  be closed. Then the following assertions are equivalent:*

<sup>3</sup>The reader may compare this notion of *B-closed* with the definition of closed languages (see page 1053).

- (1)  $K$  is regular.
- (2)  $\text{msc}(K)$  is the language of some formula of the logic  $\text{MSO}(\prec_P, \prec_{\text{msg}})$  (or, equivalently, of  $\text{MSO}(\leq, \prec_{\text{msg}})$ , or of  $\text{MSO}(\leq)$ ).
- (3)  $\text{msc}(K)$  is the language of some formula of the logic  $\text{EMSO}(\prec_P, \prec_{\text{msg}})$  (or, equivalently, of  $\text{EMSO}(\leq, \prec_{\text{msg}})$ , or of  $\text{EMSO}(\leq)$ ).
- (4)  $K$  is realizable.

Now suppose, in addition, that all words from  $K$  are complete. Then the following are equivalent:

- (1')  $K$  is regular and message deterministic.
- (5)  $K$  is deterministically realizable.

Note that the realizing CFMs in (4) and (5) are necessarily universally  $B$ -bounded since any word from  $K$  is universally  $B$ -bounded. Note also that the implication (4)  $\rightarrow$  (3) for  $\text{EMSO}(\prec_P, \prec_{\text{msg}})$  is a special instance of Theorem 8.3, which also implies (3) for  $\text{EMSO}(\leq, \prec_{\text{msg}})$ . For the logic  $\text{EMSO}(\leq)$ , the implication (4)  $\rightarrow$  (3) was shown in [48] using the fact that  $\text{msc}(K)$  consists of universally  $B$ -bounded MSCs, only. The implication (2)  $\rightarrow$  (1) follows by translating the formula that talks about MSCs to an MSO formula talking about Mazurkiewicz traces, the equivalence between MSO and regular  $I$ -closed languages [69, 29], and the closure of regular languages under projections.

The proofs of the implications (1)  $\rightarrow$  (4) and (1')  $\rightarrow$  (5) uses the encoding described in Section 9.2. It was first presented in [48], that extends the Büchi characterization above from finite MSCs to infinite ones. The main idea is to use Proposition 9.3 in order to obtain a regular and  $I$ -closed language  $L \subseteq A^*$  with  $K = \text{proj}_\Sigma(L)$ . Then, Theorem 9.1 gives a deterministic Zielonka automaton  $\mathcal{B}$  for  $L$ , which can be simulated by some CFM. Suppose  $K$  contains some incomplete word. Then the simulating CFM has to guess which send events are matched and which are unmatched (since the former get as additional index a number  $0 \leq n < B$  and the latter the symbol  $\$$ ). It turns out that this is the only source of non-determinism.

**Corollary 9.6.** *Let  $B \in \mathbb{N}$  be fixed. Then any closed, regular specification  $K \subseteq \Sigma^*_B$  is realizable in exponential time. If  $K$  is given by a DFA  $\mathcal{B}$  and contains only complete words, then one can construct a deterministic, universally  $B$ -bounded CFM  $\mathcal{A}$  realizing  $K$ , of size polynomial in the size of  $\mathcal{B}$  and exponential in  $|Ch|$  and  $B$ .*

The complexity stated in the corollary uses the construction of Zielonka automata given in [37], see also Section 9.1. Notice that we apply the construction to the “modulo  $B$  relabeling” of  $K$  (i.e., the word language  $\text{Lin}(\text{tr}(\text{msc}(K)))$ ), so we assume that we start with a DFA of size polynomial in  $B$  and exponential in  $Ch$ . Then we obtain a deterministic Zielonka automaton of size that is still polynomial in  $B$ , and (simply) exponential in  $Ch$  and  $|\overline{\mathcal{P}}| = |Ch| \cdot B$ , hence in both  $B$  and  $Ch$ . The resulting automaton includes the time-stamping needed in the CFM simulation.

The generalization of Theorem 9.5 to the existentially bounded case is less obvious. The reason is that although Proposition 9.3(2) offers a similar encoding by traces as in the universally bounded case, a CFM cannot easily simulate the corresponding Zielonka automaton. The reader might try to define a CFM that realizes the set of existentially  $B$ -bounded MSCs (that is, the set  $\text{msc}(\Sigma^*_B)$ ). Whereas this task is not very difficult



to accomplish (even deterministically) in the universally bounded case, it becomes much more involved in the existentially bounded case, see [38].

**Theorem 9.7** ([38]). *Let  $B \in \mathbb{N}$  and let  $K \subseteq \Sigma^*|_B$  be  $B$ -closed. Then the following assertions are equivalent:*

- (1)  $K$  is regular.
- (2)  $\text{msc}(K)$  is the language of some formula of one of the logics  $\text{MSO}(\leq_P, <_{\text{msg}})$ ,  $\text{MSO}(\leq, <_{\text{msg}})$ , or  $\text{MSO}(\leq)$ , respectively.
- (3)  $\text{msc}(K)$  is the language of some formula of one of the logics  $\text{EMSO}(\leq_P, <_{\text{msg}})$ ,  $\text{EMSO}(\leq, <_{\text{msg}})$ , or  $\text{EMSO}(\leq)$ , respectively.
- (4)  $\text{msc}(K)$  is realizable.

Note that the realizing CFM in (4) is necessarily existentially  $B$ -bounded since any word from  $K$  is existentially  $B$ -bounded. The main difference between this and Theorem 9.5 is that  $K$  is only  $B$ -closed (and not necessarily closed). Given this more general language  $K$ , we cannot construct a deterministic CFM as in Theorem 9.5, as the following example shows:

**Example 9.1.** Let  $\mathcal{P} = \{p, q, r\}$  and consider the following existentially 1-bounded set  $X$  of MSCs. Let  $M \in X$  if the following properties hold:

- Process  $p$  alternates between sends to  $q$  and to  $r$  (with fixed content  $m$ ).
- Between every pair of consecutive receives from  $p$ , process  $q$  ( $r$ , respectively) executes 0 or 1 local actions  $a_q$  ( $a_r$ , respectively).
- Let  $u_q$  ( $u_r$ , respectively) be the 0-1 sequence of local actions  $a_q$  ( $a_r$ , respectively) executed by  $q$  ( $r$ , respectively) between two consecutive receives from  $p$ . Then  $u_q = u_r$ .

It is fairly easy to see that  $X$  can be realized by a non-deterministic CFM where process  $p$  chooses between sending (to both  $q$  and  $r$ ) a bit  $b \in \{0, 1\}$ , that tells  $q$  and  $r$  to execute  $b$  local actions before the next receive. But any deterministic CFM fails to realize  $X$ . Suppose by way of contradiction that a deterministic CFM  $\mathcal{A}$  realizes  $X$ . Notice first that there is a unique infinite run  $\rho_0$  on the process  $p$ , such that every finite run  $\rho$  of  $\mathcal{A}$  is a prefix of  $\rho_0$  (when projected to process  $p$ ). If  $\rho$  is large enough, then there exist two MSCs  $M_1, M_2 \in X$  with the same projection on  $p$ , but different projections on  $q$ , that reach the same global state in  $\mathcal{A}$  (with empty channels). For these two MSCs, the runs on  $p$  are identical and we can combine them to an accepting run on the MSC  $M$  that coincides with  $p, q$  on  $M_1$  and with  $p, r$  on  $M_2$ . This yields  $M \in X$ , a contradiction.

The proofs of Theorems 9.5 and 9.7 are very similar except for the implication (1)  $\rightarrow$  (4), i.e., the transformation of a regular language  $K$  into a CFM. It is shown again using a Zielonka automaton accepting the recognizable trace language corresponding to  $K$ . However, one faces two additional problems here. First, recall that the conversion of an existentially bounded MSC into a Mazurkiewicz trace requires  $\text{rev}_B$ -edges. Therefore, the simulation of a Zielonka automaton by the CFM is non-deterministic, since the information conveyed by the  $\text{rev}_B$ -edges in the runs of the Zielonka automaton has to be guessed in the CFM (recall that these edges are “virtual” in the MSC). Second, one needs to construct a CFM accepting the set of all existentially  $B$ -bounded MSCs. The test that

the relation  $\leq \cup rev_B$  is acyclic is the most intricate part of the proof, and makes use of non-determinism, too. The main idea is to look for particular cycles of bounded length, and to check the non-existence of such cycles by guessing a suitable labeling of events and counting particular events. The details can be found in [38].

A further difference between the proofs of Theorems 9.5 and 9.7 concerns the translation of a formula from  $EMSO(\leq, <_{msg})$  into an equivalent one from  $EMSO(\leq)$ . The problem occurs since one needs to express the set of existentially  $B$ -bounded MSCs by an  $EMSO(\leq)$ -formula. For this, one can use the CFM realizing precisely  $msc(\Sigma^*|_B)$ , see [38].

Although the constructions in [38] are quite involved, the asymptotic complexity is the same as in the universally bounded case:

**Corollary 9.8.** *Let  $B \in \mathbb{N}$  be fixed. Then for any  $B$ -closed, regular specification  $K \subseteq \Sigma^*|_B$ , the set  $msc(K)$  is realizable in exponential time. If  $K$  is given by a DFA  $\mathcal{B}$  then one can construct a non-deterministic, existentially  $B$ -bounded CFM  $\mathcal{A}$  realizing  $msc(K)$  of size polynomial in the size of  $\mathcal{B}$  and exponential in  $|\mathcal{P}|$  and  $B$ .*

**Further reading.** The negative result of Theorem 8.2 is counterbalanced by the following positive result (see [4] and [51] for the matching lower bound): it is decidable in exponential space whether a regular specification is the language of a deadlock-free CFM with local acceptance.

Throughout this section we considered only CFMs with finite behavior. We briefly mention some related results on CFMs with infinite behavior. Theorem 9.5 has been extended to universally bounded sets of infinite MSCs and CFMs with Büchi and Muller acceptance in [48]. Subsequently, [12] has generalized the logical characterization of arbitrary CFMs from [15] by showing that for infinite behaviors, Muller and Büchi acceptance are equivalent, and also equivalent to  $EMSO(\leq_P, <_{msg}) + \exists^\infty$  that extends  $EMSO(\leq_P, <_{msg})$  by the quantifier  $\exists^\infty$  expressing that there exist infinitely many nodes with a given property. The question whether EMSO is strictly more expressive than arbitrary CFMs is open, both in the finitary and the infinitary case.

Theorems 9.5 and 9.7 state Büchi equivalences between automata and logics over universally and existentially bounded MSCs. We did not mention a third characterization, in terms of CMSC-graphs, that corresponds roughly to regular expressions over MSCs. In that setting one imposes restrictions on loops of such graphs (known as local synchronization and global cooperation, resp.) to capture universally and existentially bounded “regular” sets of MSCs, resp., see [61, 43, 38].

A different notion of realizability was considered in [36], where the specification talks only about local actions, so that a CFM realization amounts to synthesize the necessary communication over a fixed architecture. Several decidable notions of deadlock-free realizability are proposed in the latter paper.

## References

- [1] P. A. Abdulla, K. Cerans, B. Jonsson, and Y.-K. Tsay. General decidability theorems for infinite state systems. In *LICS'96*, pages 313–323. IEEE Computer Society, 1996. [1030](#), [1037](#)
- [2] P. A. Abdulla and B. Jonsson. Verifying programs with unreliable channels. *Inform. and Comput.*, 127(2):91–101, 1996. [1030](#), [1037](#), [1038](#)
- [3] S. Akshay, B. Bollig, and P. Gastin. Automata and logics for timed message sequence charts. In *FSTTCS'07*, number 4855 in LNCS, pages 290–302. Springer, 2007. [1051](#)
- [4] R. Alur, K. Etessami, and M. Yannakakis. Realizability and verification of MSC graphs. In *ICALP'01*, number 2076 in LNCS, pages 797–808. Springer, 2001. [1062](#)
- [5] R. Alur, K. Etessami, and M. Yannakakis. Inference of message sequence charts. *IEEE Trans. Software Eng.*, 29(7):623–633, 2003. [1053](#)
- [6] R. Alur, K. Etessami, and M. Yannakakis. Realizability and verification of MSC graphs. *Theor. Comp. Science*, 331(1):97–114, 2005. [1053](#)
- [7] R. Alur, D. Peled, and W. Penczek. Model-checking of causality properties. In *LICS'95*, pages 90–100. IEEE Computer Society Press, 1995. [1045](#)
- [8] M. F. Atig, A. Bouajjani, and T. Touili. On the reachability analysis of acyclic networks of pushdown systems. In *CONCUR'08*, number 5201 in LNCS, pages 356–371. Springer, 2008. [1031](#)
- [9] C. Baier, N. Bertrand, and Ph. Schnoebelen. Verifying nondeterministic probabilistic channel systems against omega-regular linear-time properties. *ACM Transactions on Computational Logic*, 9(1), 2007. [1039](#)
- [10] B. Boigelot and P. Godefroid. Symbolic verification of communication protocols with infinite state spaces using QDDs. In *CAV'96*, number 1102 in LNCS, pages 1–12. Springer, 1996. [1030](#), [1035](#)
- [11] B. Boigelot, P. Godefroid, B. Willems, and P. Wolper. The power of QDDs. In *SAS'97*, number 1302 in LNCS. Springer, 1997. [1030](#), [1036](#)
- [12] B. Bollig and D. Kuske. Muller message-passing automata and logics. *Inform. and Comput.*, 206:1084–1094, 2008. [1054](#), [1062](#)
- [13] B. Bollig and D. Kuske. An optimal construction of Hanf sentences. *Journal of Applied Logic*, 10:179–186, 2012. [1054](#)
- [14] B. Bollig, D. Kuske, and I. Meinecke. Propositional Dynamic Logic for message-passing systems. *Logical Methods in Computer Science*, 6(3:16):1–31, 2010. [1051](#), [1055](#)
- [15] B. Bollig and M. Leucker. Message-passing automata are expressively equivalent to EMSO logic. *Theor. Comp. Science*, 358(2-3):150–172, 2006. [1054](#), [1055](#), [1062](#)
- [16] A. Bouajjani and P. Habermehl. Symbolic reachability analysis of FIFO-channel systems with nonregular sets of configurations. *Theor. Comp. Science*, 221(1-2):211–250, 1999. [1030](#), [1036](#)
- [17] P. Bouyer, N. Markey, J. Ouaknine, Ph. Schnoebelen, and J. Worrell. On termination for faulty channel machines. In *STACS'08*, volume 1 of *LIPICs*, pages 121–132. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008. [1038](#), [1039](#)
- [18] D. Brand and P. Zafiropulo. On communicating finite-state machines. *J. ACM*, 30(2):323–342, 1983. [1029](#), [1031](#), [1034](#)
- [19] G. Cécé, A. Finkel, and S. Iyer. Unreliable channels are easier to verify than perfect channels. *Inform. and Comput.*, 124:20–31, 1996. [1038](#), [1039](#)

- [20] P. Chambart and Ph. Schnoebelen. Mixing lossy and perfect Fifo channels. In *CONCUR'08*, number 5201 in LNCS, pages 340–355. Springer, 2008. [1039](#)
- [21] P. Chambart and Ph. Schnoebelen. The ordinal recursive complexity of lossy channel systems. In *LICS'08*, pages 205–216. IEEE Computer Society Press, 2008. [1038](#)
- [22] P. Chandrasekaran and M. Mukund. Matching scenarios with timing constraints. In *FORMATS'06*, number 4202 in LNCS, pages 98–112. Springer, 2006. [1051](#)
- [23] L. Clemente, F. Herbreteau, and G. Sutre. Decidable topologies for communicating automata with FIFO and bag channels. In *CONCUR'14*, number 8704 in LNCS, pages 281–296. Springer, 2014. [1031](#), [1052](#)
- [24] A. Cyriac. *Verification of Communicating Recursive Programs via Split-width*. PhD thesis, LSV, ENS Cachan, 2014. [1031](#), [1052](#)
- [25] A. Cyriac, P. Gastin, and K. N. Kumar. MSO decidability of multi-pushdown systems via split-width. In *CONCUR'12*, number 7454 in LNCS, pages 547–561. Springer, 2012. [1031](#), [1052](#)
- [26] A. Cyriac, P. Gastin, and K. N. Kumar. Controllers for the verification of communicating multi-pushdown systems. In *CONCUR'14*, number 8704 in LNCS, pages 297–311. Springer, 2014. [1031](#), [1052](#)
- [27] V. Diekert and P. Gastin. Pure future local temporal logics are expressively complete for Mazurkiewicz traces. *Inform. and Comput.*, 204:1597–1619, 2006. [1056](#)
- [28] V. Diekert and G. Rozenberg, editors. *The Book of Traces*. World Scientific, Singapore, 1995. [1056](#)
- [29] W. Ebinger and A. Muscholl. Logical definability on infinite traces. *Theor. Comp. Science*, 154:67–84, 1996. [1060](#)
- [30] J. Esparza. Decidability and complexity of Petri net problems - an introduction. In W. Reisig and G. Rozenberg, editors, *Advanced Course on Petri Nets 1996*, number 1491 in LNCS, pages 374–428. Springer, 1998. [1049](#), [1050](#)
- [31] A. Finkel. A generalization of the procedure of Karp and Miller to well structured transition systems. In *ICALP'87*, number 267 in LNCS, pages 499–508. Springer, 1987. [1037](#)
- [32] A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theor. Comp. Science*, 256(1-2):63–92, 2001. [1030](#), [1037](#), [1038](#)
- [33] M. Fischer and R. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18:194–211, 1979. [1045](#)
- [34] P. Gastin and D. Kuske. Satisfiability and model checking for MSO-definable temporal logics are in PSPACE. In *CONCUR'03*, number 2761 in LNCS, pages 222–236. Springer, 2003. [1055](#)
- [35] P. Gastin and D. Kuske. Uniform satisfiability problem for local temporal logics over Mazurkiewicz traces. *Inform. and Comput.*, 208:797–816, 2010. [1055](#)
- [36] B. Genest. On implementation of global concurrent systems with local asynchronous controllers. In *CONCUR'05*, number 3653 in LNCS, pages 443–457. Springer, 2005. [1062](#)
- [37] B. Genest, H. Gimbert, A. Muscholl, and I. Walukiewicz. Optimal Zielonka-type construction of deterministic asynchronous automata. In *ICALP'10*, number 6199 in LNCS, pages 52–63. Springer, 2010. [1057](#), [1060](#)
- [38] B. Genest, D. Kuske, and A. Muscholl. A Kleene theorem and model checking algorithms for existentially bounded communicating automata. *Inform. and Comput.*, 204(6):920–956, 2006. [1030](#), [1050](#), [1057](#), [1058](#), [1059](#), [1061](#), [1062](#)

- [39] B. Genest, D. Kuske, and A. Muscholl. On communicating automata with bounded channels. *Fundam. Inform.*, 80(1-3):147–167, 2007. [1049](#)
- [40] B. Genest and A. Muscholl. Constructing exponential-size deterministic Zielonka automata. In *ICALP'06*, number 4052 in LNCS, pages 565–576, 2006. [1057](#)
- [41] B. Genest, A. Muscholl, H. Seidl, and M. Zeitoun. Infinite-state high-level MSCs: Model-checking and realizability. *J. Comput. Syst. Sci.*, 72(4):617–647, 2006. [1030](#)
- [42] W. Hanf. The theory of models. In J. W. Addison, L. Henkin, and A. Tarski, editors, *Model-Theoretic Methods in the Study of Elementary Model-Theoretic Methods in the Study of Elementary Model-Theoretic Methods in the Study of Elementary Logic*. North-Holland, Amsterdam, 1965. [1054](#)
- [43] J. G. Henriksen, M. Mukund, K. N. Kumar, M. Sohoni, and P. Thiagarajan. A theory of regular MSC languages. *Inform. and Comput.*, 202(1):1–38, 2005. [1030](#), [1033](#), [1059](#), [1062](#)
- [44] A. Heußner, J. Leroux, A. Muscholl, and G. Sutre. Reachability analysis of communicating pushdown systems. *Logical Methods in Computer Science*, 8(3), 2012. [1031](#), [1033](#), [1048](#), [1052](#)
- [45] S. Kosaraju. Decidability of reachability in vector addition systems. In *STOC'82*, pages 267–281. ACM, 1982. [1037](#)
- [46] P. Krcal and W. Yi. Communicating timed automata: The more synchronous, the more difficult to verify. In *CAV'06*, number 4144 in LNCS, pages 249–262. Springer, 2006. [1051](#)
- [47] O. Kupferman, N. Piterman, and M. Y. Vardi. Extended temporal logic revisited. In *CONCUR'01*, number 2154 in LNCS, pages 519–535. Springer, 2001. [1051](#)
- [48] D. Kuske. Regular sets of infinite message sequence charts. *Inform. and Comput.*, 187:80–109, 2003. [1057](#), [1059](#), [1060](#), [1062](#)
- [49] S. La Torre, P. Madhusudan, and G. Parlato. Context-bounded analysis of concurrent queue systems. In *TACAS'08*, number 4963 in LNCS, pages 299–314. Springer, 2008. [1031](#), [1033](#), [1052](#)
- [50] J. Leroux and S. Schmitz. Demystifying reachability in vector addition systems. In *LICS'15*, pages 56–67. IEEE Computer Society Press, 2015. [1037](#)
- [51] M. Lohrey. Realizability of high-level message sequence charts: closing the gaps. *Theor. Comp. Science*, 309(1-3):529–554, 2003. [1062](#)
- [52] M. Lohrey and A. Muscholl. Bounded MSC communication. *Inf. and Comput.*, 189(2):160–181, 2004. [1058](#)
- [53] P. Madhusudan and B. Meenakshi. Beyond message sequence graphs. In *FSTTCS'01*, number 2245 in LNCS, pages 256–267. Springer, 2001. [1050](#)
- [54] Z. Manna and A. Pnueli. Verification of the concurrent programs: the temporal framework. In R. Boyer and J. Moore, editors, *The Correctness Problem in Computer Science*, pages 215–273. Academic Press, 1981. [1029](#), [1040](#)
- [55] E. W. Mayr. An algorithm for the general Petri net reachability problem. *SIAM J. of Comput.*, 13:441–460, 1984. [1037](#)
- [56] R. Mayr. Undecidable problems in unreliable computations. *Theor. Comp. Science*, 297(1-3):337–354, 2003. [1038](#)
- [57] A. Mazurkiewicz. Concurrent program schemes and their interpretations. DAIMI Rep. PB 78, Aarhus University, Aarhus, 1977. [1056](#)

- [58] R. Mennicke. Propositional dynamic logic with converse and repeat for message-passing systems. *Logical Methods in Computer Science*, 9(2:12):1–35, 2013. [1051](#)
- [59] R. Mennicke. *Model Checking Concurrent Systems using Temporal Logics*. PhD thesis, TU Ilmenau, 2015. [1051](#)
- [60] A. R. Meyer. Weak monadic second order theory of one successor is not elementary recursive. In *Proc. of Logic Colloquium*, volume 453 of *Lecture Notes in Mathematics*, pages 132–154. Springer, 1975. [1054](#)
- [61] A. Muscholl and D. Peled. Message sequence graphs and decision problems on Mazurkiewicz traces. In *MFCS'99*, number 1672 in LNCS, pages 81–91. Springer, 1999. [1048](#), [1062](#)
- [62] T. Neary. Undecidability of binary tag systems and the Post correspondence problem for five pairs of words. In *STACS'15*, volume 30 of *LIPICs*, pages 649–661. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. [1041](#)
- [63] J. K. Pachl. Reachability problems for communicating finite state machines. Research Report CS-82-12, Dep. of Comp. Sci., Univ. of Waterloo, 1982. see also <http://arxiv.org/abs/cs/0306121>. [1035](#)
- [64] D. Peled. Specification and verification of message sequence charts. In *FORTE'00*, number 183 in IFIP Conference Proceedings, pages 139–154. Kluwer, 2000. [1045](#), [1046](#)
- [65] E. Post. Formal reductions of the combinatorial decision problem. *Amer. J. of Math.*, 65(2):197–215, 1943. [1029](#)
- [66] Ph. Schnoebelen. Verifying lossy channel systems has nonprimitive recursive complexity. *Inform. Proc. Lett.*, 83(5):251–261, 2002. [1030](#), [1038](#)
- [67] P. S. Thiagarajan and I. Walukiewicz. An expressively complete linear time temporal logic for Mazurkiewicz traces. In *LICS'97*, pages 183–194. IEEE Computer Society Press, 1997. [1056](#)
- [68] W. Thomas. On logics, tilings, and automata. In *ICALP'91*, number 510 in LNCS, pages 441–454. Springer, 1991. [1054](#)
- [69] W. Thomas and W. Zielonka. Logical definability of trace languages. Unpublished manuscript, 1990. [1060](#)
- [70] M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Inf. and Comput.*, 115(1):1–37, 1994. [1051](#)
- [71] W. Zielonka. Notes on finite asynchronous automata. *RAIRO - Informatique Théorique et Applications*, 21:99–135, 1987. [1056](#), [1057](#)