January 2002

Extending Compositional Message Sequence Graphs

Benedikt Bollig

Aachen University of Technology (RWTH)

Martin Leucker

University of Pennsylvania

Philipp Lucas

Aachen University of Technology (RWTH)

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

Benedikt Bollig, Martin Leucker, and Philipp Lucas, "Extending Compositional Message Sequence Graphs", . January 2002.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-02-09.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/171
For more information, please contact repository@pobox.upenn.edu.

Extending Compositional Message Sequence Graphs

Abstract

We extend the formal developments for message sequence charts (MSCs) to support scenarios with lost and found messages. We define a notion of *extended compositional message sequence charts* (ECMSCs) which subsumes the notion of compositional message sequence charts in expressive power but additionally allows to define lost and found messages explicitly. As usual, ECMSCs might be combined by means of choice and repetition towards (extended) compositional message sequence graphs. We show that - despite extended expressive power - model checking of monadic second-order logic (MSO) for this framework remains to be decidable. The key technique to achieve our results is to use an extended notion for linearizations.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-02-09.

Extending Compositional Message Sequence Graphs

Benedikt Bollig¹, Martin Leucker^{2*}, and Philipp Lucas¹

¹ Lehrstuhl für Informatik II, Aachen University of Technology (RWTH), Germany
{bollig,phlucas}@i2.informatik.rwth-aachen.de

² Dept. of Computer and Information Science, University of Pennsylvania, USA
leucker@cis.upenn.edu

Abstract. We extend the formal developments for message sequence charts (MSCs) to support scenarios with lost and found messages. We define a notion of *extended compositional message sequence charts* (ECMSCs) which subsumes the notion of compositional message sequence charts in expressive power but additionally allows to define lost and found messages explicitly. As usual, ECMSCs might be combined by means of choice and repetition towards (extended) compositional message sequence graphs. We show that—despite extended expressive power—model checking of monadic second-order logic (MSO) for this framework remains to be decidable. The key technique to achieve our results is to use an extended notion for linearizations.

1 Introduction

In the development process of a large and complex communication protocol, one usually starts with describing scenarios showing aspects of its desired behavior. The scenarios may be combined to get a more complete view of the system. This partial specification guides the realization of the resulting piece of hardware or software, which, in turn, might be checked according to these scenarios.

Message Sequence Charts (MSCs), which are similar to UML’s sequence diagrams [2], are a prominent, standardized [10, 11], and graphical as well as textual formalism to further this approach. An MSC defines a set of processes and a set of communication actions between these processes. In the visual representation of an MSC, processes are drawn as vertical lines. A labeled arrow from one line to another corresponds to the communication event of sending the labeling value from the first process to the second. Figure 1(a) gives an example of an MSC.

MSCs can be combined towards so-called *high-level MSCs* or *MSC graphs* (MSGs)¹ by means of choice and repetition. An MSG is a directed graph or finite automaton whose edges are labeled with MSCs. Every path through this graph corresponds to the sequential composition of the basic MSCs. In this way, arbitrarily long and alternative scenarios may be obtained out of the given ones.

Since MSCs and MSGs are employed in the early design phase of a system, it is essential to provide decent formal methods for analyzing them. In recent years,

* This research was supported in part by NSF CCR-9988409, NSF CCR-0086147, NSF CISE-9703220, ARO DAAD19-01-1-0473, DARPA ITO MOBIES F33615-00-C-1707, and ONR N00014-97-1-0505.

¹ In the following, we will use the more suggestive notion of MSGs.

a lot of work was done to support this issue, however, usually for a restricted class of MSCs.

The first question which arises in this context is whether a *set* of MSCs is *realizable*. In the seminal work of [9] and [8], regular sets of MSCs were defined and studied for this reason. The main idea was to study linearizations of MSCs. A linearization of an MSC is a sequence of all its send and receive events whereby the partial order imposed by the graphical representation of the MSC is respected. It was shown that regular sets of MSCs are the ones which are definable by monadic second-order logic (MSO) over MSCs. This work was later extended in [5] to a larger class of MSCs namely one which also allows MSCs with message overtaking. Therefore, a conceptually new way to model MSCs and their linearizations was introduced.

Another approach is to analyze MSGs. Model checking of MSGs with respect to an automaton specification was studied in [1]. More specifically, the question whether the sequences of send and receive actions obtained by linearizing the MSCs established by finite paths through the MSG adhere to some MSO formula (over words) was tackled. This problem is also known as *linearization model checking*. It turned out that model checking becomes undecidable in general, unless some *boundedness* assurance is guaranteed (in the case of the more natural asynchronous composition of MSCs).

On the contrary, the model-checking problem of MSGs against MSO formulas which are interpreted directly on MSCs and not on its linearizations was shown to be decidable [13].

A weakness in the expressiveness of MSGs was pointed out in [7] and a suitable extension towards so-called high-level *compositional* MSCs (HCMSCs) was proposed. It was shown that basic properties of these HCMSCs are undecidable. [14] slightly restricts the class of HCMSCs towards compositional message sequence graphs (CMSGs), still maintaining the main features of HCMSCs. They show that monadic second-order logic over these CMSGs is decidable. However, not every CMSC-labeled graph is a CMSG. While it is decidable whether such a structure is a CMSG, it might bother the user of CMSGs to understand the restrictions imposed to obtain a CMSG.

In this paper, we extend the notion of CMSGs to support aspects of the message-sequence-chart standard [11] which have been excluded in the previous works. Our goal is to allow so-called *black* and *white holes*, which refer to unmatched send and receive events. This concept allows the modeling of lost and found messages, respectively. Thus, one is able to describe behaviors of a protocol that catches up after a message has been lost. We present scenarios for the Alternating-Bit Protocol giving evidence for the adequacy of our formalism. Furthermore, our formalism also supports message overtaking.

Note that in some definitions for compositional message sequence charts, one might end up with unmatched send or receive events. However, this can only happen to a limited extent. For example, for *left-closed* CMSCs defined in [7], it

is possible that a send event lacks its corresponding receive event. But in that case, no further matched send event of the same type will be able to follow.

In the following, we abbreviate our extended compositional message sequence graphs by ECMSGs and call our extended version of a compositional message sequence chart ECMSC. It turns out that all scenarios definable in terms of CMSGs can also be defined by ECMSGs. Due to the fact that ECMSGs support the definition of holes and message overtaking, we easily obtain the result that ECMSGs are strictly more expressive than CMSGs.

Despite the additional expressive power, we show that model checking of MSO formulas remains to be decidable in our extended setting. We combine ideas of [5] and proof techniques of [14] and [5] to obtain our results. As opposed to CMSGs defined in [14], every ECMSC-labeled graph is an ECMSG. Thus, no further test whether the setting given by a user is indeed an (E)CMSG has to be applied.

Finally, we show in which way our model can be related to so-called Mazurkiewicz-trace-closed languages. As pointed out by [14], this is the cornerstone for applying the rich theory of local temporal logics over traces such as TrPTL [19].

Organization of the paper In the next section, we define extended compositional MSCs and ECMSGs. We give a set of scenarios for the Alternating-Bit Protocol in Section 3. In Section 4, we study the correspondence of ECMSGs and regular word languages. Model checking procedures for ECMSGs are considered in Section 5.

Fig. 1. An MSC and its formalization

2 Message Sequence Charts and Graphs

Before we present our extensions, let us start with the customary formal definition of message sequence charts, in which an MSC is understood of as a partial order of send and receive events, which can be matched subject to a bijective mapping:

Let $\mathcal{P}$ be a finite, nonempty set of *process names* (or just *processes*) and *Mess* a finite *message alphabet*. Let further $\Sigma_S := \{S_p^q(a) \mid p, q \in \mathcal{P}, p \neq q, a \in \text{Mess}\}$ and $\Sigma_R := \{R_p^q(a) \mid p, q \in \mathcal{P}, p \neq q, a \in \text{Mess}\}$ denote the sets of *send* and *receive actions*, respectively, and $\Sigma := \Sigma_S \cup \Sigma_R$ the set of *actions*. An action

$S_p^q(a)$ stands for sending a message a from process p to process q , and $R_p^q(a)$ represents the corresponding receive action, which is then executed by process q . In this sense, $Corr := \{(S_p^q(a), R_p^q(a)) \mid p, q \in \mathcal{P}, p \neq q, a \in Mess\}$ relates those actions that are corresponding. From now on, all premises and definitions are made with respect to a fixed set $\mathcal{P}$ of processes and a fixed message alphabet $Mess$. Thus, Σ and the derived symbols are also fixed. Throughout the paper, $Dom(f)$ and $Range(f)$ will furthermore denote the domain and the range of a function f , respectively.

A *Message Sequence Chart (MSC)* is a tuple $m = (E, \preceq, t, o, \mu, \ell)$ where

- E is the finite set of *events*.
- $t : E \rightarrow \{S, R\}$ is the *type function* indicating whether we deal with a *send* or *receive event*, i.e., $t(e) = S$ or respectively $t(e) = R$.
- $o : E \rightarrow \mathcal{P}$, the *process function*, binds each event to a process.
- $\mu : t^{-1}(S) \rightarrow t^{-1}(R)$ is the bijective *matching function*.
- $\ell : E \rightarrow \Sigma$ is the *labeling function* such that, for each $e \in E$, $t(e) = S$ implies

$$\ell(e) = S_{o(e)}^{o(\mu(e))}(a) \text{ and } \ell(\mu(e)) = R_{o(e)}^{o(\mu(e))}(a)$$

for some $a \in Mess$. In other words, events matched by μ are labeled by ℓ with corresponding action labels.

- $\preceq \subseteq E \times E$ is a partial order that is total on each process line and orders receive events after their originating send event. Thus, we require $\preceq$ to satisfy the following:
 - $\preceq_p := \preceq \cap (o^{-1}(p) \times o^{-1}(p))$ is a total order for each $p \in \mathcal{P}$
 - $\preceq = (\bigcup_{p \in \mathcal{P}} \preceq_p \cup \{(e, \mu(e)) \mid t(e) = S\})^*$

To simplify our notation, we will also use o to denote a mapping associating a label from Σ with its corresponding process, i.e., $o(S_p^q(a)) = o(R_q^p(a)) = p$.

The formalization of an MSC, as illustrated in Figure 1(b), is a directed graph in which the nodes, according to ℓ , are labeled with elements from Σ and where the arrows depict μ . Furthermore, a vertical line corresponds to $\preceq_p$ for a suitable $p \in \mathcal{P}$.

An MSC is just a finite object showing a scenario for an execution of an underlying system. To derive a more sophisticated description of the system, one aims at combining several MSCs. Therefore, however, it turns out to be helpful, if one allows so-called “unmatched” send events, which then might be gathered in a subsequent MSC by a corresponding “unmatched” receive event, as pointed out in [7]. Furthermore, we want to be able to describe scenarios in which, on the one hand, there has been a send event but no corresponding receive event², or, on the other hand, a message is received that has not been sent³. In terms of MSC standard [10], we want to support *black* and *white holes*, respectively. Thus, we extend our definition accordingly:

² e.g., because the message got lost

³ e.g., because some message from some previous session remains in the input channel

Definition 1 (Extended Compositional Message Sequence Chart). An Extended Compositional Message Sequence Chart (ECMSC) is just an MSC $(E, \preceq, t, o, \mu, \ell, \lambda)$ as above except for t , which is henceforth a function $E \rightarrow \{S, R, U, \perp\}$ such that an event e with $t(e) \in \{S, R\}$ is labeled as stipulated above and $t(e) \in \{U, \perp\}$ implies $o(\ell(e)) = o(e)$. In addition, λ is a mapping $t^{-1}(U) \rightarrow \mathbb{N}$, called matching information, such that

1. for each $\sigma \in \Sigma$ and each natural number n , $|\{e \in E \mid t(e) = U, \ell(e) = \sigma, \text{ and } \lambda(e) = n\}| \leq 1$, and
2. for all $e \in E$ with $t(e) = U$ and $\ell(e) \in \Sigma_S$, there is no event e' with $t(e') = U$, $(\ell(e), \ell(e')) \in \text{Corr}$, $e' \not\preceq e$, and $\lambda(e) = \lambda(e')$.

Fig. 2. Two counter-examples and an example

Let us study the previous definition in more detail. As before, if the type of an event is S or R, we deal with a send or receive event corresponding to a single message. The symbol $\perp$ can be employed to explicitly define a hole (a *lost* message or event). The symbol U denotes an unmatched message event which can possibly find a suitable communication partner in a subsequent or previous ECMSC. To give information in which way unmatched events of type U should be combined for two given ECMSCs, we employ the functions ℓ and λ . Obviously, a white hole (an unmatched receive event), can only be matched with a black hole (an unmatched send event). We do not want to model misdirected messages, so a send message from p to q may only be received by q —to ensure this, the labeling function ℓ provides the required information.

To allow for a more complex modeling, the matching information λ is used to identify corresponding send and receive events. As we will see below, an unmatched send event with number k usually matches the first suitable unmatched receive event that has number k , which should be one event of a subsequent ECMSC. Therefore, we require λ to be in the way that, within a given ECMSC, there is no corresponding receive event which could be matching to a send event (cf. condition 2 of Definition 1). For instance, look at Figure 2, where, in each diagram, the use of only one message $a \in \text{Mess}$ is assumed and a possibly unmatched

message, as indicated, is sent to (received from, respectively) the neighbored process line. The diagram in the middle is no ECMSC, because condition 2) is not fulfilled. There is no reason not to match the send event with number 2 with the receive event numbered 2. However, the third diagram in this line, though it employs unmatched messages of type U with corresponding labelings and the same matching information, is an ECMSC. The unmatched receive event with number 2 can now only be matched with one of a previous ECMSC.

For the same reason, we do not allow to carry two unmatched events of type U with the same labeling and matching information (cf. condition 1 of Definition 1). Reusing the same number k for a send event, for example, means intuitively that the first message got lost and should therefore be labeled with $\perp$. Thus, the diagram on the left of Figure 2 is not an ECMSC.

Another approach to deal with lost and found messages formally, is to add one process for each channel, i.e. pair of processes. For example, for the ECMSC shown in Figure 3(a), one might add a process PQ and simulate every message from P to Q by a message from P to PQ followed by one from PQ to Q (cf. Figure 3(b)). Lost or found messages could then be represented by leaving out the second or the first message, respectively. So, one could specify and reason in the framework of CMSCs instead of ECMSCs. However, finding a similar message to the one lost could not be distinguished from passing on the original one (see lower part in Figure 3(a) and Figure 3(b)), which is undesirable especially when the messages of the scenarios are obtained by abstracting concrete message values to a finite number of messages. While the MSC on the left hand side has lost messages, the second does not. Furthermore, message overtaking cannot be modeled. Finally, explaining why a formula is not valid for the given scenario is more difficult, if the underlying scenario is changed and $n(n - 1)$ processes are added to n processes. Thus, we are convinced that our direct approach is preferable to a simulating one, last but not least, for complexity reasons.

Fig. 3. Simulating lost and found messages

Before we turn towards the formal definition of the concatenation of ECMSCs, let us introduce some abbreviations: Given an ECMSC m as above, let

$S_m := t^{-1}(S)$, $S_m^u := t^{-1}(U) \cap \ell^{-1}(\Sigma_S)$, and $U_m, \perp_m, R_m$ and R_m^u be defined analogously. Furthermore, we set $\lambda(m) := \text{Range}(\lambda)$ as well as, for a collection $M = \{m_1, \dots, m_k\}$ of ECMSCs, $\lambda(M) := \bigcup_{i \in \{1, \dots, k\}} \lambda(m_i)$.

Let us now consider the concatenation of ECMSCs. The main idea is that we wish to specify the behavior of a system that first follows ECMSC m_1 and then ECMSC m_2 by defining a single ECMSC $m_1 \cdot m_2$ that allows exactly the desired behaviors. In our setting, we use *asynchronous concatenation*, i.e., any process that has completed all events in m_1 may proceed to m_2 . To concatenate black and white holes, we employ a function μ_{m_1, m_2} , which formally captures the previously mentioned ideas.

For ECMSCs $m_i = (E_i, \preceq_i, t_i, o_i, \mu_i, \ell_i, \lambda_i)$, $i = 1, 2$, with disjoint sets of events (if this is not the case, the events have to be renamed accordingly), let $\mu_{m_1, m_2} : S_{m_1}^u \rightarrow R_{m_2}^u$ be a partial function given by $\mu_{m_1, m_2}(e) = e'$ iff both $(\ell_1(e), \ell_2(e')) \in \text{Corr}$ and $\lambda_1(e) = \lambda_2(e')$. Observe that, due to conditions 1) and 2) of Definition 1, μ_{m_1, m_2} is in fact well-defined. The *product* of m_1 and m_2 , denoted by $m_1 \cdot m_2$, is defined to be the ECMSC $m = (E, \preceq, t, o, \mu, \ell, \lambda)$ where

- $E = E_1 \cup E_2$
- $\preceq = (\preceq_1 \cup \preceq_2 \cup \{(e, e') \in E_1 \times E_2 \mid o_1(e) = o_2(e')\} \cup \mu_{m_1, m_2})^*$
- $S_m = S_{m_1} \cup S_{m_2} \cup \text{Dom}(\mu_{m_1, m_2})$
- $R_m = R_{m_1} \cup R_{m_2} \cup \text{Range}(\mu_{m_1, m_2})$
- $\perp_m = \perp_{m_1} \cup \perp_{m_2}$
- $\cup \{e \in S_{m_1}^u \setminus \text{Dom}(\mu_{m_1, m_2}) \mid \exists e' \in S_{m_2}^u : \ell_1(e) = \ell_2(e'), \lambda_1(e) = \lambda_2(e')\}$
- $\cup \{e \in R_{m_2}^u \setminus \text{Range}(\mu_{m_1, m_2}) \mid \exists e' \in R_{m_1}^u : \ell_2(e) = \ell_1(e'), \lambda_2(e) = \lambda_1(e')\}$
- $U_m = E \setminus (S_m \cup R_m \cup \perp_m)$
- $o = o_1 \cup o_2$,
- $\mu = \mu_1 \cup \mu_2 \cup \mu_{m_1, m_2}$
- $\ell = \ell_1 \cup \ell_2$
- $\lambda = (\lambda_1 \cup \lambda_2)|_{U_m}$

In other words, the events of m comprise the ones of m_1 and m_2 , and the partial order is obtained by the reflexive and transitive closure of the orders of m_1 and m_2 , ordering the events of m_2 of a process p after the events of m_1 of process p , and the order imposed by the newly matching send and receive events, denoted by (the graph of) the function μ_{m_1, m_2} . The send and receive events of m are augmented by the events related by μ_{m_1, m_2} . The lost events are the previous ones together with the unmatched send events of type U in m_1 for which a subsequent event with the same label and matching information (in m_2) exists, and, dually, the unmatched receive events of type U in m_2 which are preceded by events with the same label and number (in m_1). All remaining events remain unmatched ones of type U and may be used for further concatenation.

Consider the ECMSCs shown in Figure 4 where, as before, the use of only one message $a \in \text{Mess}$ is assumed and a possibly unmatched message is sent to the neighbored process line. We have $m_1 \cdot m_2 = m_4$, resulting in an ECMSC with message overtaking (i.e., with equally labeled arrows which are not arranged in a

FIFO manner).⁴ Observe that the unmatched send and receive events numbered 2 have been matched and that the first send event numbered 1 turned into a lost one as it is followed by an equal send event of m_2 that is likewise numbered 1. Furthermore, $m_4 \cdot m_3 = m_5$, matching the events numbered 1.

Fig. 4. Some ECMSCs

$$\begin{aligned}
 w_1 &= S_P^Q(a)_1 \ S_P^Q(a)_2 \ S_P^Q(a)_3 \ R_P^Q(a)_3 & w_2 &= S_P^Q(a)_1 \ R_P^Q(a)_2 \\
 w'_1 &= S_P^Q(a)_1 \ S_P^Q(a)_2 \ S_P^Q(a)_1 \ R_P^Q(a)_1 & w_4 &= S_P^Q(a)_\perp \ S_P^Q(a)_4 \ S_P^Q(a)_3 \ S_P^Q(a)_1 \ R_P^Q(a)_3 \ R_P^Q(a)_4
 \end{aligned}$$

Fig. 5. Exemplary linearizations

We easily see that $\cdot$ is not associative. For example, while $(m_1 \cdot m_2) \cdot m_3 = m_5$, $m_1 \cdot (m_2 \cdot m_3)$ yields m_6 . In the following, we let $\cdot$ associate to the left, i.e., $m_1 \cdot m_2 \cdot m_3$ denotes $(m_1 \cdot m_2) \cdot m_3$.

Another equivalence will turn out to be important, which does not distinguish between ECMSCs that only differ in the type and matching information of unmatched events: Given ECMSCs $m_i = (E_i, \preceq_i, t_i, o_i, \mu_i, \ell_i, \lambda_i)$, $i = 1, 2$, we write $m_1 \equiv m_2$ iff $(E_1, \preceq_1, t_1|_{S_{m_1} \cup R_{m_1}}, o_1, \mu_1, \ell_1)$ and $(E_2, \preceq_2, t_2|_{S_{m_2} \cup R_{m_2}}, o_2, \mu_2, \ell_2)$ are isomorphic. For example, referring to Figure 4, $m_5 \equiv m_6$. Furthermore, m_5 is equivalent to the ECMSC which we obtain from m_5 by replacing 2 with 3. The

⁴ Given ECMSCs m and m' , we write $m = m'$ iff m and m' are isomorphic.

equivalences $=$ and $\equiv$ over ECMSCs are extended to sets as customary. In particular, $M_1 \equiv M_2$ iff $\{[m]_{\equiv} \mid m \in M_1\} = \{[m]_{\equiv} \mid m \in M_2\}$.

As we have defined the concatenation of ECMSCs, we are now ready to define ECMSGs:

Definition 2 (Extended Compositional Message Sequence Graph). *An Extended Compositional Message Sequence Graph (ECMSG) is a finite automaton $\mathcal{H} = (States, \Pi, \delta, s_{in}, F, \chi)$, i.e., $States$ is a nonempty finite set of states, Π is a finite alphabet, $\delta \subseteq States \times \Pi \times States$ is the set of transitions, $s_{in} \in States$ is the initial state, and $F \subseteq States$ is the set of final states. In addition, χ maps each letter h from Π to an ECMSC $\chi(h)$.*

The word language $L(\mathcal{H}) \subseteq \Pi^*$ is defined in the expected manner. Moreover, $\mathcal{H}$ defines an ECMSC language $M(\mathcal{H}) := \{\chi(h_1 \dots h_n) \mid h_1 \dots h_n \in L(\mathcal{H}), h_i \in \Pi\}$ where, for $n \geq 1$, $\chi(h_1 \dots h_n) := \chi(h_1) \cdot \dots \cdot \chi(h_n)$, and $\chi(\varepsilon)$ is set to be the empty ECMSC. Let furthermore $\lambda(\mathcal{H}) := \lambda(Range(\chi))$. We give a large example for an ECMSG in the next section. To get an idea of the previous definition for the moment, look at $\mathcal{H}'$ in Figure 6.

Note that, since FIFO matching is assumed in CMSCs, a CMSC lacks a matching information. As a CMSC in the style of [14] only traces a bounded number of unmatched events within a path—every accepting path must guarantee that the behavior is *complete*, i.e., without unmatched events—we can, involving some combinatorial considerations, find an equivalent ECMSG for it. This justifies our notion of *extended CMSCs*.

More precisely, given a CMSC $\mathcal{H}$, we are looking for an ECMSG $\mathcal{H}'$ such that $M(\mathcal{H}') = M(\mathcal{H})$. Basically, the states of $\mathcal{H}'$ are copies of the ones of $\mathcal{H}$ but enriched by some additional information. Namely, for each send action, they contain a sequence of natural numbers that depicts the matching information of currently unmatched messages, respectively, which is unique for each state. Similarly, the ECMSCs used in $\mathcal{H}'$ are copies of the CMSCs from $\mathcal{H}$ whereby each unmatched message in a CMSC (without any matching information) becomes an unmatched event in $\mathcal{H}'$ of type U with a matching information according to the additional parameters of the current state.

For example, look at Figure 6, which refers to the prominent producer-and-consumer protocol that cannot be described by means of a simple MSG [7]. Again, we assume the use of one message only. In $\mathcal{H}'$, whenever a path goes through state s'_1 , the first unmatched send event in the corresponding ECMSC so far will be labeled 1 (recall that we simulate a FIFO behavior). Thus, in each outgoing ECMSC, the first unmatched receive event is likewise labeled 1, respectively. Furthermore, as the matching information of state s'_1 is depicted by the sequence (1, 2), the first unmatched send event in $m_2^{3,1}$ contains the matching information 3. Apart from the above restrictions, the transitions of $\mathcal{H}'$ could be understood as an unwinding of $\mathcal{H}$.

As we did not define CMSGs formally and as the detailed algorithm of transforming CMSGs into ECMSGs is rather technical and would not provide further insights, we confine ourselves to giving the above example.

Fig. 6. From CMSGs to ECMSGs

3 A System Specification

Let us formalize scenarios for a variant of the well-known *Alternating-Bit Protocol* (ABP) [16] to exemplify how to use our framework.

The ABP was designed to guarantee a reliable data transmission through insecure channels, which may lose or duplicate a message. The general idea is that a sender process *Send* sends a message together with a control bit b to a receiver process *Receive*, which, in turn, sends the control bit back. As soon as *Send* receives the control bit b , it sends the next chunk of data, now together with $1 - b$. If, however, the control bit has not been returned to the sender within a given time frame, the sender will resend the message using b , because it thinks that the data has not been delivered. The control bit, on which we concentrate in our further analysis, might be used to detect errors in the communication.

Let us now consider several scenarios in a more structured way: First, a sender *Send* sends a bit b and sets a timer, which we describe in the ECMSC

11

S^b shown in Figure 8.⁶ The sending might be followed by a timeout or by an acknowledgement followed by a timeout, as shown respectively in the ECMSCs TO_1 and $RAck^b$. Thus, we may combine the scenarios as shown in Figure 8 and obtain the ECMSC $\mathcal{H}_1$. The receiver behaves dually. We identify the scenarios R^b , $SAck^b$, and TO_2 , which we combine towards $\mathcal{H}_2$.

We specify the ABP simply as the parallel product $\mathcal{H}_1 \parallel \mathcal{H}_2$ of the ECMSCs $\mathcal{H}_1$ and $\mathcal{H}_2$ from Figure 8, where we define the parallel product of two ECMSCs $(States_1, \Pi, \delta_1, s_{in}^1, F_1, \chi)$ and $(States_2, \Pi, \delta_2, s_{in}^2, F_2, \chi)$, as one might expect, to be $(States_1 \times States_2, \Pi, \delta, (s_{in}^1, s_{in}^2), F_1 \times F_2, \chi)$. The set δ is hereby given as $\{((s_1, s_2), h, (s'_1, s'_2)) \mid (s_1, h, s'_1) \in \delta_1, s_2 \in States_2\} \cup \{((s_1, s_2), h, (s_1, s'_2)) \mid s_1 \in States_1, (s_2, h, s'_2) \in \delta_2\}$.

A (partial) run of the product automaton is shown in Figure 7(b), whose sequence of ECMSCs yields the ECMSC aside (Figure 7(a)). Thus, the ECMSC of Figure 7(a) is a prefix of an ECMSC from $M(\mathcal{H}_1 \parallel \mathcal{H}_2)$.

Anticipating the developments to come, let us mention that we now might analyze our scenarios with respect to MSO formulas. For example, we might ask whether we can derive scenarios in which we have unmatched receive events (white holes) $(\exists y \bigvee_{p,q,a} (L_{R_p^q(a)}(y) \wedge \neg \exists x (x \rightarrow y)))$. If we are faced with a setting that our channel will never produce white holes, we can simply analyze a formula φ by model checking our protocol against the formula $(\forall y \exists x ((\bigvee_{p,q,a} L_{R_p^q(a)}(y)) \Rightarrow x \rightarrow y)) \Rightarrow \varphi$, i.e., ruling out those ECMSCs that do not employ an unmatched receive event.

4 ECMSCs and Regular Languages

In this section, we provide notions to handle ECMSCs by means of *linearizations*. Our goal is to achieve a characterization as in [14], that is, the language of an ECMSC can be *represented* by a regular word language. This characterization will be the basis for the model checking procedure. The main idea of our approach is to use linearizations which are indexed by natural numbers.

A word $w = \frac{\sigma_1}{u_1} \dots \frac{\sigma_n}{u_n} \in (\Sigma \times (\mathbb{N} \cup \{\perp\}))^*$ is called *ECMSC word*. Its positions can be considered to be events of an ECMSC, and together with its letters, we can obtain the matching information (cf. Figures 4 and 5). More specifically, w defines an ECMSC $m_w := (E, \preceq, t, o, \mu, \ell, \lambda)$ given as follows:

- $E = \{1, \dots, n\}$
- $\preceq$ is the reflexive and transitive closure of $\sqsubset$ where $i \sqsubset j$ iff $i < j$ and
 - $o(\sigma_i) = o(\sigma_j)$ or
 - $(\sigma_i, \sigma_j) \in Corr$ and $u_i = u_j \neq \perp$
- $S_{m_w} = \{i \mid \sigma_i \in \Sigma_S \text{ and } \exists j > i : (\sigma_i, \sigma_j) \in Corr \text{ and } u_i = u_j \neq \perp \text{ and } (\sigma_k, u_k) \neq (\sigma_i, u_i) \text{ for } i < k < j\}$

⁶ Note that we do not show the second timer process in scenarios for the sender and the first timer process in scenarios involving the receiver for lack of space. Furthermore, note that S^b represents two scenarios: S^0 and S^1 .

Fig. 8. The Alternating-Bit Protocol

$$\begin{aligned}
R_{m_w} &= \{i \mid \sigma_i \in \Sigma_{\mathcal{R}} \text{ and} \\
&\quad \exists j < i : (\sigma_j, \sigma_i) \in Corr \text{ and} \\
&\quad u_i = u_j \neq \perp \text{ and } (\sigma_k, u_k) \neq (\sigma_i, u_i) \text{ for } j < k < i\} \\
U_{m_w} &= \{i \mid \sigma_i \in \Sigma_{\mathcal{S}}, u_i \neq \perp, \text{ and } \nexists j > i : (\sigma_i, u_i) = (\sigma_j, u_j) \text{ or} \\
&\quad [(\sigma_i, \sigma_j) \in Corr \text{ and } u_i = u_j]\} \\
&\quad \cup \{i \mid \sigma_i \in \Sigma_{\mathcal{R}}, u_i \neq \perp, \text{ and } \nexists j < i : (\sigma_i, u_i) = (\sigma_j, u_j) \text{ or} \\
&\quad [(\sigma_j, \sigma_i) \in Corr \text{ and } u_i = u_j]\} \\
\perp_{m_w} &= E \setminus (S_{m_w} \cup R_{m_w} \cup U_{m_w}) \\
- o(i) &= o(\sigma_i)
\end{aligned}$$

- $\mu(i) = \min\{j \mid i < j \text{ and } (\sigma_i, \sigma_j) \in \text{Corr} \text{ and } u_i = u_j\}$
- $\ell(i) = \sigma_i$
- $\lambda(i) = u_i$ (for i with $t(i) = \text{U}$)

The set of ECMSC words (recall that Σ is fixed) is denoted by $\mathcal{W}$. We call $w = \sigma_1 \dots \sigma_n \in \mathcal{W}$ a *linearization* of an ECMSC m iff $m_w = m$. Let $\text{Lin}(m)$ denote the set of linearizations of m . Furthermore, we will make use of $\bar{\lambda}(w) := \{u_i \mid i \in S_{m_w}\}$ describing the set of numbers used in w for matched events (positions).

To give some examples, look at Figures 4 and 5. While w_1 is a linearization of m_1 , w'_1 is not.⁷ In fact, any ECMSC is sufficiently described by a single linearization. The word w_2 is one of two possible linearizations of m_2 , and w_4 is a linearization of m_4 . Observe that $m_{w_1 w_2} = m_1 \cdot m_2 = m_4$. Generally, we can state:

Proposition 1. *Given a nonempty finite collection $M = \{m_1, \dots, m_k\}$ of ECMSCs and linearizations $w_i \in \text{Lin}(m_i)$ with $\bar{\lambda}(w_i) \cap \lambda(M) = \emptyset$, $i \in \{1, \dots, k\}$, we have $m_{w_1 \dots w_k} \equiv m_1 \cdot \dots \cdot m_k$.*

An important result of [14] also holds in our setting, saying—informally—that languages of ECMSGs can be represented by regular word languages. We call a language $L \subseteq \mathcal{W}$ *regular* iff there is a $B \in \mathbb{N}$ such that $L \subseteq (\Sigma \times (\{1, \dots, B\} \cup \{\perp\}))^*$ and L is regular in the usual sense.

Theorem 1. *Let M be a set of ECMSGs.*

1. *If M is the language of an ECMSG, then there is a regular language $L \subseteq \mathcal{W}$ with $\{m_w \mid w \in L\} \equiv M$.*
2. *If there is a regular language $L \subseteq \mathcal{W}$ with $\{m_w \mid w \in L\} = M$, then M is $\equiv$ -equivalent to the language of an ECMSG.*

Proof. 1. We first extend the usual notion of finite automata towards using words instead of letters in the transition function, i.e., in the following, we deal with a structure $\mathcal{A} = (\text{States}, \Pi, \delta, s_{\text{in}}, F)$ such that δ is a finite subset of $\text{States} \times \Pi^* \times \text{States}$. This is just for notational convenience and does not alter the expressive power of finite automata.

For an ECMSG $\mathcal{H} = (\text{States}, \Pi, \delta, s_{\text{in}}, F, \chi)$, we can build an (extended) finite automaton $\mathcal{A}_{\mathcal{H}} = (\text{States}', \Sigma \times (\{1, \dots, B\} \cup \{\perp\}), \delta', s'_{\text{in}}, F')$ (with B a natural) such that $\{m_w \mid w \in L(\mathcal{A}_{\mathcal{H}})\} = M(\mathcal{H})$ as follows: Initially, for each ECMSC $m \in \text{Range}(\chi)$, we pick a linearization $w_m \in \text{Lin}(m)$ with $\bar{\lambda}(w_m) \cap \lambda(\mathcal{H}) = \emptyset$. (As a general rule, for each ECMSC, one would determine a *canonical* linearization, for example, in the sense of [5].) $\mathcal{A}_{\mathcal{H}}$ is then given as follows: B is the maximal index used in a linearization. Furthermore, $\text{States}' = \text{States}$, $s'_{\text{in}} = s_{\text{in}}$, $F' = F$, and $(q, w, q') \in \delta'$ iff there is $h \in \Pi$ such that both $(q, h, q') \in \delta$ and $w = w_{\chi(h)}$. The language equivalence (with respect to $\equiv$) follows from Proposition 1.

⁷ replace 1 by $\perp$ in m_1

2. We first assign to a symbol $(\sigma, u) \in \Sigma \times (\mathbb{N} \cup \{\perp\})$ the ECMSC $m_{(\sigma, u)} := (E, \preceq, t, o, \mu, \ell, \lambda)$ defined by $E = \{e\}$, $\preceq = \{(e, e)\}$, $t(e) = \mathbb{U}$ if $u \in \mathbb{N}$, $t(e) = \perp$ otherwise, $o(e) = o(\sigma)$, $\mu = \emptyset$, $\ell(e) = \sigma$, and, if $t(e) = \mathbb{U}$, $\lambda(e) = u$. I.e., $m_{(\sigma, u)}$ consists of exactly one unmatched event of type $\mathbb{U}$ or type $\perp$. (Note that, in fact, this definition is according to the ECMSC we obtain if (σ, u) is considered as a word.) Given a (usual) finite automaton $\mathcal{A} = (States, \Sigma \times (\{1, \dots, B\} \cup \{\perp\}), \delta, s_{in}, F)$ whose language is L , the ECMSG $\mathcal{H}_L$ satisfying $M(\mathcal{H}_L) \equiv \{m_w \mid w \in L\}$ is given by $(States, \Sigma \times (\{1, \dots, B\} \cup \{\perp\}), \delta, s_{in}, F, \chi)$ where $\chi((\sigma, u)) = m_{(\sigma, u)}$. $\square$

5 Model Checking for ECMSGs

We have now set out the scene to consider model checking for ECMSGs.

5.1 MSO Model Checking

Following [14], we give a decision procedure for a monadic second-order logic interpreted over ECMSCs, using the concept of a regular representative linearization.

Given a supply $\text{Var} = \{x, y, \dots\}$ of individual variables, which are going to be interpreted over events of an ECMSC, and a supply $\text{VAR} = \{X, Y, \dots\}$ of set variables, which are interpreted over sets of events, the syntax of $\text{MSO}(\mathcal{P}, \text{Mess})$ is defined by the following grammar:

$$\varphi ::= L_\sigma(x) \mid x \rightarrow y \mid x \in X \mid x \preceq y \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \exists x\varphi \mid \exists X\varphi$$

where $\sigma \in \Sigma$, $x, y \in \text{Var}$, and $X \in \text{VAR}$. Moreover, we allow the usual abbreviations. The intuitive meaning of $x \rightarrow y$ is that x is send event and y its corresponding receive event. $x \preceq y$ holds iff y follows x in the order of the ECMSC at hand. The remaining constructs are as usual. We formalize the intuitive study: Let $m = (E, \preceq, t, o, \mu, \ell, \lambda)$ be an ECMSC. Given an interpretation function $\mathcal{I}$, which assigns to an individual variable x an event $\mathcal{I}(x) \in E$ and to a set variable X a set of events $\mathcal{I}(X) \subseteq E$, the satisfaction relation $m \models_{\mathcal{I}} \varphi$ for a formula $\varphi \in \text{MSO}(\mathcal{P}, \text{Mess})$ is inductively defined as follows:

- $m \models_{\mathcal{I}} L_\sigma(x)$ iff $\ell(\mathcal{I}(x)) = \sigma$
- $m \models_{\mathcal{I}} x \rightarrow y$ iff $\mathcal{I}(x) \in S_m$, $\mathcal{I}(y) \in R_m$, and $\mu(\mathcal{I}(x)) = \mathcal{I}(y)$
- $m \models_{\mathcal{I}} x \in X$ iff $\mathcal{I}(x) \in \mathcal{I}(X)$
- $m \models_{\mathcal{I}} x \preceq y$ iff $\mathcal{I}(x) \preceq \mathcal{I}(y)$
- $m \models_{\mathcal{I}} \neg\varphi$ iff $m \not\models_{\mathcal{I}} \varphi$
- $m \models_{\mathcal{I}} \varphi \vee \psi$ iff $m \models_{\mathcal{I}} \varphi$ or $m \models_{\mathcal{I}} \psi$
- $m \models_{\mathcal{I}} \exists x\varphi$ iff $\exists e \in E : m \models_{\mathcal{I}[x/e]} \varphi$
- $m \models_{\mathcal{I}} \exists X\varphi$ iff $\exists E' \subseteq E : m \models_{\mathcal{I}[X/E']} \varphi$

We only consider formulas without free variables in the following and accordingly write $m \models \varphi$ instead of $m \models_{\mathcal{I}} \varphi$. For $\varphi \in \text{MSO}(\mathcal{P}, \text{Mess})$, let $M_\varphi := \{m \mid m \models \varphi\}$. Note that M_φ is closed with respect to $\equiv$, i.e., for any ECMSCs m and m' with $m \equiv m'$, it holds $m \models \varphi$ iff $m' \models \varphi$.

Theorem 2. *Given $\varphi \in \text{MSO}(\mathcal{P}, \text{Mess})$ and an ECMSG $\mathcal{H}$, we can decide whether $M(\mathcal{H}) \subseteq M_\varphi$.*

Proof. Let $\mathcal{A}_\mathcal{H}$ be the automaton of $\mathcal{H}$ from the first part of the proof of Theorem 1 with $L(\mathcal{A}_\mathcal{H}) \subseteq (\Sigma \times (\{1, \dots, B\} \cup \{\perp\}))^* =: \mathcal{W}_B$. From φ , we inductively build an MSO formula $\|\varphi\|$ interpreted over elements of $\mathcal{W}_B$ such that $\{w \in \mathcal{W}_B \mid w \models \|\varphi\|\} = \{w \in \mathcal{W}_B \mid m_w \models \varphi\} =: L_\varphi^B$. The fact that $M(\mathcal{H}) \subseteq M_\varphi$ iff $L(\mathcal{A}_\mathcal{H}) \subseteq L_\varphi^B$ then leads to a decision procedure in the obvious manner.

Let us determine $\|\varphi\|$ inductively as follows: For $L_\sigma(x)$, we have to guess the right number, thus, we define $\|L_\sigma(x)\| := \bigvee_{u \in \{1, \dots, B, \perp\}} L_{(\sigma, u)}(x)$. Element relation, negation, disjunction and existential quantification carry through, i.e., $\|x \in X\| := x \in X$, $\|\neg\psi\| := \neg\|\psi\|$, $\|\psi_1 \vee \psi_2\| := \|\psi_1\| \vee \|\psi_2\|$, $\|\exists x\psi\| := \exists x\|\psi\|$, and $\|\exists X\psi\| := \exists X\|\psi\|$. x and y are corresponding send and receive events if y follows x , their labelings are correlated, and there is no correlated element inbetween. So we define $\|x \rightarrow y\| := \searrow(x, y)$ where the matching predicate is given by

$$\searrow(x, y) := \bigvee_{\substack{(\sigma, \tau) \in \text{Corr} \\ u \in \{1, \dots, B\}}} [x < y \wedge L_{(\sigma, u)}(x) \wedge L_{(\tau, u)}(y) \wedge \nexists z(x < z < y \wedge (L_{(\sigma, u)}(z) \vee L_{(\tau, u)}(z)))]$$

Finally, y follows x causally if there is a corresponding sequence of matching send and receive events or events on the same process line. Thus,

$$\|x \preceq y\| := \exists X [x \in X \wedge y \in X \wedge \forall z(z \in X \wedge z \neq y \Rightarrow \exists z' (z' \in X \wedge z < z' \wedge \text{Proc}(z) = \text{Proc}(z') \vee \searrow(z, z')))]$$

and $\text{Proc}(x)$ is defined as expected. $\square$

5.2 Temporal Logics for ECMSCs

While MSO is a powerful and useful specification logic on its own, the result of the previous subsection provides simple decidability results for further (temporal) logics by encoding. For example, *TLC* [18] can easily be encoded into $\text{MSO}(\mathcal{P}, \text{Mess})$ and model checking can be shown to be decidable in this way. However, this approach does not seem to be reasonable for practical issues, due to the non-elementary complexity of model checking MSO formulas. In the domain of Mazurkiewicz traces [6], several temporal logics have been studied and sophisticated model-checking procedures were developed.

These logics usually define *trace-closed* languages. Therefore, as pointed out by Madhusudan and Meenakshi [14], it is desirable to find the existence of temporal logics interpreted over ECMSC words that are trace-closed in some sense, i.e., whose formulas are satisfied by either all linearizations of an ECMSC or none of them. First attempts in this regard were done in [15] and [4].

We will now establish a strong connection between ECMSCs and the theory of Mazurkiewicz traces, which subsequently yields trace-closed logics as well as corresponding decision procedures in a natural manner. With respect to Σ , the *dependence relation* $D(\Sigma) \subseteq (\Sigma \times (\mathbb{N} \cup \{\perp\}))^2$ is given by $(\sigma, u)D(\Sigma)(\sigma', u')$ iff

- $o(\sigma) = o(\sigma')$ or
- $(\sigma, \sigma') \in \text{Corr}$ and $u = u' \neq \perp$ or
- $(\sigma', \sigma) \in \text{Corr}$ and $u = u' \neq \perp$.

For each natural B , the pair $\tilde{\Sigma} := (\Sigma \times \{1, \dots, B, \perp\}, D(\Sigma) \cap (\Sigma \times \{1, \dots, B, \perp\})^2)$ turns out to be a *Mazurkiewicz trace alphabet* [6]. Kuske and Morin investigate relations like this in detail [12, 17].

A relation $\sim_{\tilde{\Sigma}} \subseteq \mathcal{W}_B \times \mathcal{W}_B$ provides information about which ECMSC words are seen to be equivalent with respect to $D(\Sigma)$. So let $\sim_{\tilde{\Sigma}}$ be the least equivalence relation satisfying the following: If $w = w_1(\sigma, u)(\sigma', u')w_2$ and $w' = w_1(\sigma', u')(\sigma, u)w_2$ for suitable w_1, w_2 and not $(\sigma, u)D(\Sigma)(\sigma', u')$, then $w \sim_{\tilde{\Sigma}} w'$. Thus, within an equivalence class of $\sim_{\tilde{\Sigma}}$, we are allowed to permute two neighbored positions in an ECMSC word that are labeled with independent actions (actions which are not dependent).

The above connection between ECMSCs and Mazurkiewicz traces is useful due to the following fact: For each ECMSC m , $\text{Lin}(m) \cap \mathcal{W}_B$ is the finite union of equivalence classes of $\sim_{\tilde{\Sigma}} \cap (\mathcal{W}_B \times \mathcal{W}_B)$. It is then effortlessly possible to suit temporal logics for Mazurkiewicz traces like TrPTL [19] or LTrL [20] to ECMSCs and to employ corresponding decision procedures, for example in the style of [3].

6 Conclusion

In this paper, we presented a formal definition of extended compositional message sequence charts (ECMSCs) with *black* and *white holes*, which supports methods of formal analysis. Furthermore, we allow message overtaking as defined in the MSC standard. The ECMSCs can be combined by means of choice and repetition towards (extended) compositional message sequence graphs (ECMSGs). Giving scenarios and ECMSGs for the Alternating-Bit Protocol, we validated that our formalism is useful for practical applications.

Despite its extended expressive power, we have shown that model checking monadic second-order formulas is decidable. Similar as in [14], our decision procedure uses the idea to *represent* the language of an ECMSG by a regular language. To obtain the result that languages of ECMSGs indeed can be represented by regular word languages, we introduce *indexed* linearizations.

These indexed linearizations give furthermore a link between MSCs and Mazurkiewicz traces in a natural way, as requested by [14]. This link is the cornerstone for easily carrying over the sophisticated work on temporal logics for Mazurkiewicz traces to the domain of MSCs.

References

1. R. Alur and M. Yannakakis. Model checking of message sequence charts. In *Proceedings of the 10th International Conference on Concurrency Theory*, volume 1664 of *LNCS*, pages 114–129. Springer, 1999.

2. J. Araújo. Formalizing sequence diagrams. In *Proceedings of the OOPSLA'98 Workshop on Formalizing UML. Why? How?*, volume 33, 10 of *ACM SIGPLAN Notices*, New York, 1998. ACM Press.
3. B. Bollig and M. Leucker. Deciding LTL over Mazurkiewicz traces. In *Proceedings of the Symposium on Temporal Representation and Reasoning (TIME'01)*, pages 189–197. IEEE Computer Society Press, 2001.
4. B. Bollig and M. Leucker. Modelling, Specifying, and Verifying Message Passing Systems. In *Proceedings of the Symposium on Temporal Representation and Reasoning (TIME'01)*, pages 240–248. IEEE Computer Society Press, 2001.
5. B. Bollig, M. Leucker, and T. Noll. Generalised regular MSC languages. In M. Nielsen, editor, *Proceedings of the 5th International Conference on Foundations of Software Science and Computation Structures (FOSSACS'02)*, volume 2303 of *LNCS*, pages 52–66. Springer, 2002.
6. V. Diekert and Y. Métivier. Partial commutation and traces. In G. Rozenberg and A. Salomaa, editors, *Handbook on Formal Languages*, volume III. Springer, 1997.
7. E. Gunter, A. Muscholl, and D. Peled. Compositional message sequence charts. In *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'01)*, volume 2031 of *LNCS*, pages 496–511. Springer, 2001.
8. J. G. Henriksen, M. Mukund, K. N. Kumar, and P. S. Thiagarajan. On message sequence graphs and finitely generated regular msc languages. In *Proceedings of 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*, volume 1853 of *LNCS*, pages 675–686. Springer, 2000.
9. J. G. Henriksen, M. Mukund, K. N. Kumar, and P. S. Thiagarajan. Regular collections of message sequence charts. In *Proceedings of 25th International Symposium on Mathematical Foundations of Computer Science (MFCS'00)*, volume 1893 of *LNCS*, pages 405–414. Springer, 2000.
10. ITU-TS. ITU-TS Recommendation Z.120anb: Formal Semantics of Message Sequence Charts. Technical report, ITU-TS, Geneva, 1998.
11. ITU-TS. ITU-TS Recommendation Z.120: Message Sequence Chart 1999 (MSC99). Technical report, ITU-TS, Geneva, 1999.
12. D. Kuske. Another step towards a theory of regular MSC languages. In *Proceedings of the 19th International Symposium on Theoretical Aspects of Computer Science (STACS'02), 2002*, volume 2285 of *LNCS*. Springer, 2002.
13. P. Madhusudan. Reasoning about sequential and branching behaviours of message sequence graphs. In *Proceedings of 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*, volume 2076 of *LNCS*, pages 396–407. Springer, 2001.
14. P. Madhusudan and B. Meenakshi. Beyond message sequence graphs. In *Proceedings of the 21st Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 2245 of *LNCS*, pages 256–267. Springer, 2001.
15. B. Meenakshi and R. Ramanujam. Reasoning about message passing in finite state environments. In *Proceedings of 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*, volume 1853 of *LNCS*. Springer, 2000.
16. R. Milner. *Communication and Concurrency*. International Series in Computer Science. Prentice Hall, 1989.
17. R. Morin. Recognizable sets of message sequence charts. In *Proceedings of the 19th International Symposium on Theoretical Aspects of Computer Science (STACS'02), 2002*, volume 2285 of *LNCS*. Springer, 2002.
18. D. Peled. Specification and verification of message sequence charts. In *Proc. IFIP FORTE/PSTV*, 2000.
19. P. S. Thiagarajan. A trace consistent subset of PTL. In I. Lee and S. A. Smolka, editors, *Proceedings of the 6th International Conference on Concurrency Theory (CONCUR'95)*, volume 962 of *LNCS*, pages 438–452, Philadelphia, Pennsylvania, 21–24 Aug. 1995. Springer.
20. P. S. Thiagarajan and I. Walukiewicz. An expressively complete linear time temporal logic for Mazurkiewicz traces. In *Proceedings of 12th IEEE Symposium on Logic in Computer Science*, pages 183–194, Warsaw, Poland, IEEE Computer Society Press.