

Generalised Regular MSC Languages

Benedikt Bollig¹, Martin Leucker², and Thomas Noll¹

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

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

Abstract. We establish the concept of regularity for languages consisting of Message Sequence Charts (MSCs). To this aim, we formalise their behaviour by string languages and give a natural definition of regularity in terms of an appropriate Nerode right congruence. Moreover, we present a class of accepting automata and establish several decidability and closure properties of MSC languages. We also provide a logical characterisation by a monadic second-order logic interpreted over MSCs. In contrast to existing work on regular MSC languages, our approach is neither restricted to a certain class of MSCs nor tailored to a fixed communication medium (such as a FIFO channel). It explicitly allows MSCs with message overtaking and is thus applicable to a broad range of channel types like mixtures of stacks and FIFOs.

1 Introduction

Components of distributed systems usually communicate with each other via message passing: A sender process sends a message over a channel, from which it is taken by the receiver process. A prominent formalism to model this kind of systems is that of *Message Sequence Charts* (MSCs) [8,9]. They are standardised, can be denoted both textually and graphically, and are often employed in industry. Furthermore, they are quite similar to the notion of sequence charts of the Unified Modelling Language (UML) [2].

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 labelled arrow from one line to another corresponds to the communication event of sending the labelling value from the first process to the second. As the vertical lines are interpreted as time axes, there is the general rule that arrows must not go “upwards” because this would describe a situation that a message is received before it has been sent. Figure 1(a) gives an example of an MSC. Collections of MSCs are used to capture the scenarios that a designer might want the system to follow or to avoid.

When one considers the dynamic behaviour of an MSC, i.e., the sequences of actions that may be observed when the system is executed, one distinguishes between the so-called visual-order semantics and the causal-order semantics. The

* Most of the work was completed during the author’s employment at Lehrstuhl für Informatik II, Aachen University of Technology, Germany.




Fig. 1. An MSC and its formalisation

visual order assumes that the events are ordered as shown in the MSC. That is, the events on a single process line are linearly ordered, and sending events precede their corresponding receiving events. For example, Process Q in Figure 1(a) has to read the a symbol before it can read b . In the *causal order-based semantics*, a concrete communication medium between the processes is taken into account, e.g., a first-in-first-out (FIFO) channel. Furthermore, receiving events on the same process line are not ordered unless they are “causality dependent”. For instance, reading event b may occur before reading a : As Process P might have sent a after R has sent b and assuming a single FIFO incoming channel for Q , Q will potentially receive b before a . Note that, under the same communication assumption, reading c must occur before reading d . To simplify our presentation, we adopt the visual-order point of view in the following. But we would like to stress that—with minor modifications—our very general approach also works wrt. the causal order.

Given the system specification in the form of a collection of MSCs, one is interested in doing formal analysis to discover errors at the early stages of system design. Of course, the first question arising is which kinds of collections of MSCs are amenable to formal methods. In a pioneering work by Henriksen et al. [7], a definition of *regularity* of MSC languages is proposed. A characterisation in terms of message-passing automata and in terms of monadic second-order logic is also given. The paper explains in a convincing way the benefits of these alternative descriptions, arguing that this is the “right” notion of regularity for MSCs. For example, a characterisation in terms of finite devices (automata) gives evidence for a collection of MSCs to be *realisable*.

However, this approach has a serious limitation. So-called “MSCs with message overtaking” cannot be considered. But these are explicitly defined in the official standard [8] and must be taken into account. The limitation stems from the fact that, for establishing a link between MSCs and classical language theory, the graphical representation of an MSC has somehow to be mapped to the domain of strings. The straightforward approach, enumerating the possible linearisations of the events that occur in an MSC, only works for simple types of MSCs where the correspondence between a sending event and its receiving coun-

terpart can be derived from the order in which they occur in the string. Note that also [1] has to restrict the admissible class of MSCs in order to be able to relate MSCs and string languages.

Our solution to this problem is to associate with every communication event in the string representation of an MSC a natural number that explicitly establishes this correspondence. As it will become clear in the next section, this allows us to drop any restriction on the set of MSCs under consideration. The price to pay is that, for *arbitrary* collections of MSCs, we have to work with strings, automata, etc. over *infinite alphabets*. For practical applications though, the “simple” collections of MSCs are of interest. Therefore, within the domain of (MSC word) languages, we will spot the *regular* ones. These are defined in terms of a Nerode right congruence, which allows a straightforward generalisation to languages over infinite alphabets.

To support formal analysis, we introduce a new kind of automaton (MFA) accepting linearisations of MSCs. More precisely, our notion of MFAs guarantees that every accepted word is indeed a linearisation of an MSC. Moreover, we establish several closure properties and decidability results. In particular, we show that language inclusion is decidable, a crucial property for *model-checking* applications. Our concept of automata is similar to the one introduced by Kaminski and Francez [10]. Note, however, that in their setting the problem of language inclusion is undecidable [14]. Furthermore, our framework is well suited for extensions. In [6], compositional message sequence graphs (CMSGs) are introduced to describe larger classes of MSCs. Our automata model MFA is well prepared to accept languages of CMSGs, which can be characterised by MSC languages with *regular representative linearisations*, a concept defined and studied by Madhusudan and Meenakshi [12]. However, due to lack of space, this topic will be discussed elsewhere.

Subsequently, we follow the line of [7] and develop an alternative automata-theoretic characterisation based on message-passing automata as well as a description in terms of monadic second-order logic. Although the results are similar, the proofs are of a different nature because it is generally impossible to lift proofs directly from the setting of languages over finite alphabets to the infinite case.

Our main contribution is to develop a theory of regular collections of MSCs in terms of Nerode right congruences, finite automata, message-passing automata, and models of MSO formulas for the full class of MSCs. Thus, we provide the formal basis for subsequent verification questions. Our approach has already turned out to be useful in the setting of LTL model checking for MSCs [3].

Due to space constraints, we omit most proofs for the presented results. These can be found in the full version of the paper [4].

2 Message Sequence Charts and Their Linearisations

In this section, we present our formal model for MSCs and establish a string representation, which describes their behaviour in a linear way.

2.1 Message Sequence Charts

For $N \geq 2$, let $\mathcal{P}_N := \{1, \dots, N\}$ be a set of *processes* and A a finite *message alphabet*. Let further $\Sigma_S := \{S_p^q(\lambda) \mid p, q \in \mathcal{P}_N, p \neq q, \lambda \in A\}$ and $\Sigma_R := \{R_p^q(\lambda) \mid p, q \in \mathcal{P}_N, p \neq q, \lambda \in A\}$ denote the sets of *send* and *receive actions*, respectively, and $\Sigma := \Sigma_S \cup \Sigma_R$ their union. An action $S_p^q(\lambda)$ stands for sending a message λ from Process p to Process q , and $R_p^q(\lambda)$ represents the corresponding receive action, which is then executed by Process q . In this sense, $Corr := \{(S_p^q(\lambda), R_p^q(\lambda)) \mid p, q \in \mathcal{P}_N, p \neq q, \lambda \in A\}$ relates those actions that belong together. From now on, all premises and definitions are made wrt. a fixed set \mathcal{P}_N of processes and a fixed message alphabet A .

An MSC is a tuple of the form $M = (\{E_p\}_{p \in \mathcal{P}_N}, \{\preceq_p\}_{p \in \mathcal{P}_N}, f, L)$ where $\{E_p\}_{p \in \mathcal{P}_N}$ is a family of pairwise disjoint finite sets of *events*, each of which is totally ordered by a relation $\preceq_p \subseteq E_p \times E_p$. (For simplicity, we consider \preceq_p as a relation over $E := \bigcup_{p \in \mathcal{P}_N} E_p$, the set of all events.) Let $P : E \cup \Sigma \rightarrow \mathcal{P}_N$ yield the process an event or an action belongs to, i.e., $P(e) = p$ for any $e \in E_p$, $P(S_p^q(\lambda)) = p$, and $P(R_p^q(\lambda)) = q$. M is required to induce a partition $E = S \cup R$ of the events into send (S) and receive events (R) such that $f : S \rightarrow R$ is a bijective mapping satisfying the following:

- The *visual order* $\preceq \subseteq E \times E$ of M , i.e., the reflexive and transitive closure of $\bigcup_{p \in \mathcal{P}_N} \preceq_p \cup \{(e, f(e)) \mid e \in S\}$, is a partial order; in particular, it is antisymmetric.
- $L : E \rightarrow \Sigma$ provides information about the messages being interchanged by communicating events whereby, for all $e \in S$, there is some $\lambda \in A$ such that

$$L(e) = S_{P(e)}^{P(f(e))}(\lambda) \text{ and } L(f(e)) = R_{P(f(e))}^{P(e)}(\lambda).$$

Figure 1(b) presents a formal version of the MSC shown in Figure 1(a).

A partial execution (configuration) of an MSC can be described by a downwards closed subset of events, containing those events that occurred so far. Formally, given an MSC $M = (\{E_p\}_{p \in \mathcal{P}_N}, \{\preceq_p\}_{p \in \mathcal{P}_N}, f, L)$, a *configuration* of M is a subset E' of E satisfying $E' = \downarrow E' := \{e \in E \mid \exists e' \in E' : e \preceq e'\}$. Let $Conf(M)$ denote the set of configurations of M . The execution of M can be described by a transition relation $\longrightarrow_M \subseteq Conf(M) \times \Sigma \times Conf(M)$ where $c \xrightarrow{\sigma}_M c'$ iff there exists $e \in E - c$ such that $L(e) = \sigma$ and $c' = c \cup \{e\}$.

2.2 MSC Words

A suitable notion of regularity for a class of objects should have similarities with existing notions for regular sets of objects. We will therefore reduce regularity of collections of MSCs to regularity of word languages. Thus, we have to identify an MSC with a set of words, which will be called *linearisations* or *MSC words*. A linearisation represents a possible execution sequence of the events occurring in an MSC. To justify this view, it is necessary to guarantee that—up to isomorphism—from a set of linearisations a corresponding MSC can be unambiguously inferred and vice versa. We are then able to define an equivalence on MSC words whose equivalence classes on their own determine exactly one MSC and, as a whole, stand for the set of all MSCs.

So one of the main problems is how to define an MSC word. For example, $w = S_1^2(a)S_1^2(a)R_1^2(a)R_1^2(a) \in \Sigma^*$ might define the MSC M_1 given in Figure 2. But as w is also a correct linearisation of the MSC




Fig. 2. MSCs generated by α_1 and α_2

aside, we could likewise imagine that w represents M_2 , relating the first and the fourth position of w . We therefore cannot unambiguously correlate a word in Σ^* with an MSC. Faced with causal-order semantics, the problem of relating events will be even more involved. In particular, if we make use of nondeterministic channels (which might allow MSCs to behave both in a FIFO manner and as a stack, for example), we need some information about which positions belong together. For this purpose, each position of a word $w \in \Sigma^*$ is equipped with a natural number indicating the matching positions (namely those showing the same number). The words $\alpha_1, \alpha_2 \in (\Sigma \times \mathbb{N})^*$ from Figure 3 are such MSC words. Notice that α_1 will determine the MSC M_1 , whereas M_2 will emerge from α_2 . To avoid these difficulties, [7] and [1] do not allow an MSC like M_2 . However, M_2 is a perfect “MSC with message overtaking”, which is explicitly allowed in the MSC standard [8,9].

Developing our theory step by step, we first call a word $\alpha \in (\Sigma \times \mathbb{N})^*$

- *proper* iff for all $(\sigma, \tau) \in Corr$, $\pi \in \mathbb{N}$, and prefixes α' of α , $|\alpha'|_{(\tau, \pi)} \leq |\alpha'|_{(\sigma, \pi)} \leq |\alpha'|_{(\tau, \pi)} + 1$, and
- *complete* iff it is proper and for all $(\sigma, \tau) \in Corr$ and $\pi \in \mathbb{N}$, $|\alpha|_{(\sigma, \pi)} = |\alpha|_{(\tau, \pi)}$.

Thus, in a proper word every receiving event (we sometimes refer to positions of MSC words as events) must be preceded by a sending counterpart, and, for each number π and each send action, at most one “open” sending event is admitted.

Definition 1 (MSC Word). A word $\sigma_1^{\pi_1} \dots \sigma_\ell^{\pi_\ell} \in (\Sigma \times \mathbb{N})^*$ is called an MSC word iff it is complete. Let MW denote the set of all MSC words and PW the set of proper words.

To see some examples, look at the words $\alpha_1, \dots, \alpha_4 \in (\Sigma \times \mathbb{N})^*$ given in Figure 3. As mentioned before, α_1 and α_2 are MSC words, whereas α_3 is certainly proper but not complete and α_4 is not even proper. We will refer to α_1 and α_2 as exemplary MSC words throughout the rest of the paper.

Given a proper word $\alpha = \sigma_1^{\pi_1} \dots \sigma_\ell^{\pi_\ell} \in PW$, we determine which positions are *matching*. For $i, j \in \{1, \dots, \ell\}$, we write $i \searrow_\alpha j$ iff $i < j$, $(\sigma_i, \sigma_j) \in Corr$, and $j = \min\{k \mid k > i \text{ and } \pi_k = \pi_i \text{ and } (\sigma_i, \sigma_k) \in Corr\}$.

Referring to the previous example, $1 \searrow_{\alpha_1} 3$ and $2 \searrow_{\alpha_1} 4$ as well as $1 \searrow_{\alpha_2} 4$ and $2 \searrow_{\alpha_2} 3$.

$$\begin{array}{ll}
\alpha_1 = \begin{array}{cccc} S_1^2(a) & S_1^2(a) & R_1^2(a) & R_1^2(a) \\ 1 & 3 & 1 & 3 \end{array} & \alpha_2 = \begin{array}{cccc} S_1^2(a) & S_1^2(a) & R_1^2(a) & R_1^2(a) \\ 1 & 2 & 2 & 1 \end{array} \\
\alpha_3 = \begin{array}{cccc} S_1^2(a) & S_2^1(b) & R_2^1(b) & S_2^1(b) & R_1^2(a) \\ 2 & 1 & 1 & 1 & 2 \end{array} & \alpha_4 = \begin{array}{cccc} S_1^2(a) & S_2^1(b) & R_2^1(b) & S_2^1(b) & S_1^2(a) \\ 2 & 1 & 1 & 1 & 2 \end{array}
\end{array}$$

Fig. 3. Exemplary words

2.3 From MSC Words to MSCs

Let us show that MSC words indeed represent MSCs. Falling back on the matching relation, a word $\alpha = \frac{\sigma_1}{\pi_1} \dots \frac{\sigma_\ell}{\pi_\ell} \in \text{MW}$ generates an MSC $M(\alpha) := (\{E_p\}_{p \in \mathcal{P}_N}, \{\preceq_p\}_{p \in \mathcal{P}_N}, f, L)$ where

- $E_p = \{n \in \{1, \dots, \ell\} \mid P(\sigma_n) = p\}$,
- $S = \{n \in \{1, \dots, \ell\} \mid \sigma_n \in \Sigma_S\}$,
- $R = \{n \in \{1, \dots, \ell\} \mid \sigma_n \in \Sigma_R\}$,
- $n \preceq_p m$ iff $n, m \in E_p$ and $n \leq m$,
- $f(n) = m$ iff $n \searrow_\alpha m$, and
- $L(n) = \sigma_n$.

For example, α_1 generates the MSC M_1 illustrated in Figure 2, whereas α_2 generates M_2 .

Moreover, there is no problem in extending the above definition to proper words, which then determine prefixes of MSCs.

Note that two different proper words can stand—up to isomorphism—for one and the same MSC or configuration of an MSC, respectively: Since the naturals are only used for identifying matching positions, we have some freedom in choosing the actual value. Furthermore, we are free to choose the linearisation of independent events. Therefore, we define two equivalence relations $\approx \subseteq \text{PW} \times \text{PW}$ and $\sim \subseteq \text{MW} \times \text{MW}$. The first identifies words with equivalent projections onto the second component; the latter, as introduced further below, allows to permute the positions of an MSC word.

Thus, for $\alpha = \frac{\sigma_1}{\pi_1} \dots \frac{\sigma_\ell}{\pi_\ell} \in \text{PW}$ and $\beta = \frac{\tau_1}{\rho_1} \dots \frac{\tau_m}{\rho_m} \in \text{PW}$, let $\alpha \approx \beta$ iff $\sigma_1 \dots \sigma_\ell = \tau_1 \dots \tau_m$ and for all $i, j \in \{1, \dots, \ell\}$, $i \searrow_\alpha j$ iff $i \searrow_\beta j$.

For instance, let α_1^n emerge from α_1 by replacing 3 in the natural-number component with some $n \in \mathbb{N}$. Then, $\alpha_1^n \in \text{MW}$ iff $n \neq 1$, and $\alpha_1^n \in \text{MW}$ implies $\alpha_1 \approx \alpha_1^n$. But notice that $\alpha_1 \not\approx \alpha_2$ because the second condition in the definition of \approx is violated.

For a proper word $\alpha = \frac{\sigma_1}{\pi_1} \dots \frac{\sigma_\ell}{\pi_\ell} \in \text{PW}$, let $\text{open}(\alpha) \subseteq \Sigma_S \times \mathbb{N}$ denote the set of those send events that are not followed by a matching receive event, i.e., $\text{open}(\alpha) := \{(\sigma_i, \pi_i) \mid \sigma_i \in \Sigma_S \text{ and there is no } j > i \text{ such that } i \searrow_\alpha j\}$. We call the elements of $\text{open}(\alpha)$ *open events*. A word $\alpha \in \text{PW}$ is called in *normal form* iff for all prefixes $\frac{\sigma_1}{\pi_1} \dots \frac{\sigma_k}{\pi_k}$ of α , $\sigma_k \in \Sigma_S$ implies $\pi_k = \min\{\pi \in \mathbb{N} \mid (\sigma_k, \pi) \notin \text{open}(\frac{\sigma_1}{\pi_1} \dots \frac{\sigma_{k-1}}{\pi_{k-1}})\}$. Thus, for every sending event, the lowest available number is chosen. Note that every equivalence class in PW/\approx contains exactly one word

in normal form. For $\alpha \in \text{PW}$, let furthermore $nf(\alpha) = \beta$ iff $\alpha \approx \beta$ and β is in normal form. For instance, $nf(\alpha_1) = \alpha_1^2$, whereas α_2 is already in normal form so that $nf(\alpha_2) = \alpha_2$. nf is applied to sets of words in the expected manner.

In the following, we will not distinguish \approx -equivalent words.

Definition 2 (MSC Word Language). *A set $\mathcal{L} \subseteq \text{MW}$ is called an MSC word language iff $\mathcal{L} = \mathcal{L}^\approx$ where \mathcal{L}^\approx denotes the \approx -closure of \mathcal{L} .*

Note that, for any MSC word language \mathcal{L} , it holds $\mathcal{L} = nf(\mathcal{L})^\approx$.

Characterising regular languages within the scope of MSCs, a certain restriction of words and MSCs will prove to be important. Given a natural number B , $\alpha \in \text{MW}$ is called *B-bounded* iff for all prefixes α' of α and actions $\sigma \in \Sigma_S$, $|open(\alpha') \cap \{(\sigma, \pi) \mid \pi \in \mathbb{N}\}| \leq B$. This means that, for every send action, the number of open events is bounded by B . Examples for 2-bounded MSC words are α_1 and α_2 . Note that we could likewise call α *B-bounded* iff for all prefixes α' of α , $|open(\alpha')| \leq B$, i.e., the total number of open send events is bounded by B , or also iff for all prefixes α' of α and $p \in \mathcal{P}_N$, $|open(\alpha') \cap \{(\sigma, \pi) \mid \pi \in \mathbb{N}, P(\sigma) = p\}| \leq B$, which means that the number of open events per process is bounded by B .

The definitions differ in the concrete bound, and the appropriate definition taken may vary depending on an underlying channel type. However, all presented results hold for every of these definitions.

2.4 Linearisations of MSCs

To finally relate MSCs to the rich theories of languages and automata over words, the concept of linearisations of an MSC is essential. We call an MSC word $\alpha = \frac{\sigma_1}{\pi_1} \dots \frac{\sigma_\ell}{\pi_\ell} \in (\Sigma \times \mathbb{N})^*$ a *linearisation* of an MSC $M = (\{E_i\}_{i \in \mathcal{P}_N}, \{\preceq_i\}_{i \in \mathcal{P}_N}, f, L)$ with a set of events $E = \{e_1, \dots, e_\ell\}$ iff there are $c_1, \dots, c_\ell \in \text{Conf}(M)$ with $\emptyset \xrightarrow{M} c_1 \xrightarrow{M} c_2 \xrightarrow{M} \dots \xrightarrow{M} c_\ell$ and there is a bijective mapping $\chi : E \rightarrow \{1, \dots, \ell\}$ such that for all $e \in E$, $L(e) = \sigma_{\chi(e)}$, and for all $e \in S$, $e' \in R$, $f(e) = e'$ implies $\chi(e) \searrow_\alpha \chi(e')$. $Lin(M)$ denotes the set of linearisations of M . For a set \mathcal{M} of MSCs, we canonically define $Lin(\mathcal{M}) := \bigcup \{Lin(M) \mid M \in \mathcal{M}\}$. For instance, the exemplary word α_1 is a linearisation of the MSC M_1 shown in Figure 2, and α_2 is a linearisation of M_2 . When, above, we spoke of isomorphism of two MSCs, we actually meant “inducing the same set of linearisations” instead.

An MSC is called *B-bounded* iff all of its linearisations are *B-bounded*. A collection of MSCs (a collection of MSC words, respectively) is *B-bounded* iff all members are *B-bounded*. Furthermore, we speak of *boundedness* in general iff we deal with *B-boundedness* for an arbitrary B .

We now turn towards $\sim \subseteq \text{MW} \times \text{MW}$, the second natural equivalence relation to study on linearisations of MSCs because it takes permutations of positions into account. For example, in Figure 1, it makes no real difference whether $S_3^2(b)$ occurs before $R_1^2(a)$ or after it. Given Σ , we define the *dependence relation* $D(\Sigma) \subseteq (\Sigma \times \mathbb{N})^2$ and write $(\sigma, \pi)D(\Sigma)(\sigma', \pi')$ iff $P(\sigma) = P(\sigma')$ or $[(\sigma, \sigma') \in \text{Corr}$ and $\pi = \pi']$ or $[(\sigma', \sigma) \in \text{Corr}$ and $\pi = \pi']$. It turns out that the pair $(\Sigma \times \{1, \dots, B\}, D(\Sigma) \cap (\Sigma \times \{1, \dots, B\})^2)$ is a Mazurkiewicz trace alphabet

[5] for every natural B —a fact which was already used in [11] providing a *direct* link between Mazurkiewicz traces and MSCs.

We then define the relation \sim to be the least equivalence relation satisfying the following: If $\alpha = \beta_1(\sigma, \pi)(\sigma', \pi')\beta_2$ and $\alpha' = \beta_1(\sigma', \pi')(\sigma, \pi)\beta_2$ for suitable β_1, β_2 and not $(\sigma, \pi)D(\Sigma)(\sigma', \pi')$, then $\alpha \sim \alpha'$.

This section concludes with the following important properties of sets of linearisations that are induced by MSCs. In particular, they establish the expected connections between linearisations and the equivalence relations \approx and \sim .

Theorem 1. *For an MSC M and $\alpha \in \text{Lin}(M)$, $\text{Lin}(M) = \text{Lin}(M(\alpha))$.*

Theorem 2. *For $\alpha \in \text{MW}$, $\text{Lin}(M(\alpha)) = [\alpha]_{(\approx \cup \sim)^*}$.*

Theorems 1 and 2 can be shown by employing standard techniques taken, for example, from the Mazurkiewicz trace theory. The proofs are left to the reader.

3 Regular MSC Word Languages and Their Automata

We already mentioned that the regularity of collections of MSCs will be defined in terms of regular MSC word languages. But as MSC words are defined over the infinite alphabet $\Sigma \times \mathbb{N}$, we have to modify the usual notion of regularity. In [10], a definition of regular word languages over infinite alphabets is proposed by providing an extended automata model that employs a finite transition relation but generates a behaviour catering for the fact that we deal with an infinite alphabet. However, important questions for these automata are undecidable. Thus, we follow a different approach. We first constitute an algebraic characterisation of regularity by means of a slightly adapted version of the Nerode right congruence, which allows a straightforward extension to infinite alphabets. Then, we establish its equivalence to an automata model that has similarities with the one described in [10] but is better suited for MSCs and provides desired properties.

3.1 Regular MSC Word Languages

Given an MSC word language \mathcal{L} , recall the definition of the Nerode right congruence $\equiv_{\mathcal{L}} \subseteq \text{PW} \times \text{PW}$: $\alpha \equiv_{\mathcal{L}} \beta$ iff $\forall \gamma \in (\Sigma \times \mathbb{N})^*. (\alpha\gamma \in \mathcal{L} \text{ iff } \beta\gamma \in \mathcal{L})$. As we want to identify \approx -equivalent words, we define $\approx_{\mathcal{L}} \subseteq \text{PW} \times \text{PW}$ as an extension of the Nerode right congruence by $\alpha \approx_{\mathcal{L}} \beta$ iff $\text{nf}(\alpha) \equiv_{\mathcal{L}} \text{nf}(\beta)$.

Definition 3 (Regular MSC Word Language). *An MSC word language \mathcal{L} is called regular iff $\approx_{\mathcal{L}}$ has finite index.*

The next characterisation of regular MSC word languages prepares for proving their correspondence with a certain class of finite automata, which we introduce further below.

Theorem 3. *Let \mathcal{L} be an MSC word language. \mathcal{L} is regular iff $\text{nf}(\mathcal{L})$ is a regular word language over $\Sigma \times Q$ for a finite subset Q of \mathbb{N} .*

Corollary 1. *Regular MSC word languages are bounded.*

The next theorem will be useful when, in Section 4, we consider \sim -closed MSC word languages.

Theorem 4. *Let \mathcal{L} be a \sim -closed regular MSC word language. Then $\text{nf}(\mathcal{L})^\sim$ is a regular word language over a finite alphabet.*

The proof of Theorem 3 is technically involved while the one for Theorem 4 proceeds by establishing a link to Mazurkiewicz trace theory and borrows deeper results to show the claim. Cf. the long version of the paper for details.

3.2 MSC Finite-Memory Automata

We now present an automata model characterising the class of regular MSC word languages. Our definition is inspired by [10] but modified to suit the requirements for MSCs and to allow stronger decidability results. Our model can be described as a finite automaton that makes use of a finite window whose positions occur in the labellings of the transitions—as well as elements of Σ —and indicate where to store a symbol of the infinite alphabet $\Sigma \times \mathbb{N}$ (concerning send actions) and where to take it from (concerning receive actions), respectively. Normal forms of regular MSC word languages could also be accepted by “standard” finite automata and we can use this fact to establish certain closure properties. However, not every finite automaton accepts normal forms of MSC words so that we do not get a precise automata-theoretic characterisation of regular MSC word languages which is the basis for a powerful algorithmic support of the theory on MSCs.

Definition 4 (MSC Finite-Memory Automaton). *An MSC finite-memory automaton (MFA) is a quintuple of the form $\mathcal{A} = (S, r, \Delta, q_0, F)$ where*

- S is a nonempty finite set of states,
- $r \geq 1$ is a natural number called window length,
- $\Delta \subseteq S \times (\Sigma \times \{1, \dots, r\}) \times S$ is the transition relation,
- $q_0 \in S$ is the initial state, and
- $F \subseteq S$ is the set of final states.

Figure 4 shows two MFAs, each with a window of length two. Let thereby S stand for $S_1^2(a)$ and R for $R_1^2(a)$.

Given an MFA \mathcal{A} as above, a configuration of \mathcal{A} lists the current state and the current window entries, which are either numbered send events or empty (denoted by #). Thus, let $\text{Conf}_{\mathcal{A}} := S \times ((\Sigma_S \times \mathbb{N}) \cup \{\#\})^r$ denote the (infinite) set of configurations of \mathcal{A} . We define a transition relation $\Longrightarrow_{\mathcal{A}} \subseteq \text{Conf}_{\mathcal{A}} \times (\Sigma \times \mathbb{N}) \times \text{Conf}_{\mathcal{A}}$ as follows:




Fig. 4. Two MFAs

- For $\sigma \in \Sigma_S$, $(s, \mathbf{w}) \xrightarrow{(\sigma, \pi)}_{\mathcal{A}} (t, \mathbf{v})$ iff (σ, π) does not occur in \mathbf{w} and there is a transition $(s, (\sigma, k), t) \in \Delta$ such that $\mathbf{w}[k] = \#$, $\mathbf{v}[k] = (\sigma, \pi)$, and for each $l \neq k$, $\mathbf{w}[l] = \mathbf{v}[l]$.
- For $\sigma \in \Sigma_R$, $(s, \mathbf{w}) \xrightarrow{(\sigma, \pi)}_{\mathcal{A}} (t, \mathbf{v})$ iff there is a transition $(s, (\sigma, k), t) \in \Delta$ such that $\mathbf{w}[k] = (\tau, \pi)$ where $(\tau, \sigma) \in \text{Corr}$, $\mathbf{v}[k] = \#$, and for each $l \neq k$, $\mathbf{w}[l] = \mathbf{v}[l]$.

Thus, the meaning of a transition $(s, (\text{S}_p^q(\lambda), k), t)$ is the following: If \mathcal{A} is in state s , it is able to read an input symbol $(\text{S}_p^q(\lambda), \pi)$, $\pi \in \mathbb{N}$, iff the k th position of its window is currently free and, furthermore, $(\text{S}_p^q(\lambda), \pi)$ does not occur elsewhere in the window, i.e., there is no further open $(\text{S}_p^q(\lambda), \pi)$ -labelled send event. Taking the transition, the automaton stores $(\text{S}_p^q(\lambda), \pi)$ in the k th position and enters state t . If, in contrast, the automaton reads an input symbol $(\text{R}_p^q(\lambda), \pi)$, there has to be a transition $(s, (\text{R}_p^q(\lambda), k), t)$ such that the k th position of the window currently shows the corresponding send symbol $(\text{S}_p^q(\lambda), \pi)$. Replacing this symbol with $\#$, the automaton enters state t .

A *run* of \mathcal{A} on a word $\frac{\sigma_1}{\pi_1} \dots \frac{\sigma_\ell}{\pi_\ell} \in (\Sigma \times \mathbb{N})^*$ is a corresponding sequence $(s_0, \mathbf{w}_0)(s_1, \mathbf{w}_1) \dots (s_\ell, \mathbf{w}_\ell)$ of configurations such that $s_0 = q_0$, $\mathbf{w}_0 = \#^r$, and for each $i \in \{1, \dots, \ell\}$, $(s_{i-1}, \mathbf{w}_{i-1}) \xrightarrow{(\sigma_i, \pi_i)}_{\mathcal{A}} (s_i, \mathbf{w}_i)$. The run is *accepting* iff $s_\ell \in F$ and $\mathbf{w}_\ell = \#^r$. $\mathcal{L}(\mathcal{A}) := \{\alpha \mid \text{there is an accepting run of } \mathcal{A} \text{ on } \alpha\}$ forms the language defined by \mathcal{A} . We conclude that matching events in an accepted word use one and the same position of the window for their “agreement”.

Due to the conditions we laid down for making transitions and accepting words, an MFA will accept MSC words only. A receive symbol has to be preceded by a corresponding send symbol, which, on its part, has to wait for the corresponding receive symbol before repeating the identical send symbol. Thus, we make sure that an accepted word is proper. Furthermore, as a run is accepting as soon as it ends in a final configuration featuring an empty window, completeness of accepted words is ensured. Moreover, the recognised language is \approx -closed because matching symbols can be read with—up to the MSC-word condition—arbitrary natural numbers. Notice that a regular MSC word language is not necessarily \sim -closed, a key feature allowing [12] to model CMSGs in terms of MSC word languages. We sum up these considerations as follows:

Proposition 1. *Given an MFA \mathcal{A} , $\mathcal{L}(\mathcal{A})$ is an MSC word language.*

For example, let us consider the MFAs \mathcal{A}_1 and \mathcal{A}_2 illustrated by Figure 4 and behaving in a FIFO manner and as a stack, respectively. For the sake of clarity, let S stand for $\text{S}_1^2(a)$ and R for $\text{R}_1^2(a)$. Note that our MFAs permit only Process 1 to send and only Process 2 to receive a message a .

Recall our exemplary words α_1 and α_2 . In fact, $\alpha_1 \in \mathcal{L}(\mathcal{A}_1)$ and $\alpha_2 \in \mathcal{L}(\mathcal{A}_2)$, but $\alpha_1 \notin \mathcal{L}(\mathcal{A}_2)$ and $\alpha_2 \notin \mathcal{L}(\mathcal{A}_1)$. An accepting run of \mathcal{A}_1 on α_1 first writes $(\text{S}_1^2(a), 1)$ into the first position of the window and then $(\text{S}_1^2(a), 3)$ into the second, whereupon the window is cleared in the same order, reading first $(\text{R}_1^2(a), 1)$ and then $(\text{R}_1^2(a), 3)$.

Our notion of MFAs covers exactly the class of regular MSC word languages.

Theorem 5. *An MSC word language \mathcal{L} is regular iff there is an MFA \mathcal{A} such that $\mathcal{L} = \mathcal{L}(\mathcal{A})$.*

Proof. Exploiting Theorem 3, we specify respective automata.

(\implies) Let $\mathcal{A} = (S, \longrightarrow, q_0, F)$ be a finite automaton with $\longrightarrow \subseteq S \times (\Sigma \times Q) \times S$ for a finite set $Q \subseteq \mathbb{N}$ such that $\mathcal{L}(\mathcal{A}) = nf(\mathcal{L})$. The MFA $\mathcal{A}' = (S', r, \Delta, q'_0, F')$ satisfying $\mathcal{L}(\mathcal{A}') = \mathcal{L}(\mathcal{A})^\approx = \mathcal{L}$ is given by $S' = S$, $r = |Q|$, $(s, (\sigma, k), t) \in \Delta$ iff $s \xrightarrow{(\sigma, k)} t$, $q'_0 = q_0$, and $F' = F$.

(\impliedby) Given an MFA $\mathcal{A} = (S, r, \Delta, q_0, F)$, let first $Q = \{1, \dots, r\}$ and let $\mathcal{A}' = (S', \longrightarrow, q'_0, F')$ be the corresponding finite automaton satisfying $\mathcal{L}(\mathcal{A}') = nf(\mathcal{L}(\mathcal{A}))$, defined as follows: $S' = S \times ((\Sigma_S \times Q) \cup \{\#\})^r$, $\longrightarrow \subseteq S' \times (\Sigma \times Q) \times S'$ where $(s, \mathbf{w}) \xrightarrow{(\sigma, \pi)} (t, \mathbf{v})$ iff both $(s, \mathbf{w}) \xrightarrow{(\sigma, \pi)}_{\mathcal{A}} (t, \mathbf{v})$ and $\sigma \in \Sigma_S$ implies $\pi = \min\{\pi' \in \mathbb{N} \mid (\sigma, \pi')$ does not occur in $\mathbf{w}\}$, $q'_0 = (q_0, \#^r)$, and $F' = F \times \{\#^r\}$. \square

Given an MSC word language in terms of an MFA, the first natural question is whether it defines the trivial language.

Theorem 6. *It is decidable whether a regular MSC word language given by an MFA is empty.*

We obtain this result applying the construction shown in the proof of Theorem 5.

Theorem 7. *The class of regular MSC word languages is closed under union, intersection, concatenation, and Kleene star.*

Theorem 7 follows from Theorem 3 and the fact that the concatenation of two MSC words in normal form is again in normal form.

To support the algorithmic handling of MFAs, one could alternatively provide automata-theoretic constructions that also establish the above closure properties. These are explained in the full version of the paper.

The class of regular MSC word languages is not closed under complement because the complement of a regular MSC word language is always unbounded. Thus, the standard way to show decidability of language inclusion does not work. However, in contrast to the general case of regular languages over infinite alphabets where this problem is undecidable (see [14]), we can show decidability in our setting, again applying the construction used for proving Theorem 5. This is of great importance for the development of model-checking algorithms.

Theorem 8. *Given MFAs \mathcal{A}_1 and \mathcal{A}_2 , it is decidable whether $\mathcal{L}(\mathcal{A}_1) \subseteq \mathcal{L}(\mathcal{A}_2)$.*

4 Regular MSC Languages and Their Automata

4.1 Regular MSC Languages

We now extend our theory of regular MSC word languages to collections of MSCs. Regularity of such a collection is reduced to regularity of the set of corresponding linearisations.

Definition 5 (Regular MSC Language). *A collection \mathcal{M} of MSCs is called a regular MSC language iff $\text{Lin}(\mathcal{M})$ is a regular MSC word language.*

According to this definition, the set of linearisations of a regular MSC language is necessarily \sim -closed by Theorem 2. Hence, regular MSC languages cannot be characterised by MFAs because these accept also non- \sim -closed languages. We therefore develop a generalisation of message-passing automata [7] that accepts exactly regular MSC word languages corresponding to regular MSC languages. Note that a regular MSC language is bounded.

One might ask at this stage for the reason considering regular MSC word languages as well as regular MSC languages because the latter seem to be the first choice studying linearisations of MSCs. This is true when we abstract from a communication medium between the processes of an MSC. Consider for example the MSC presented in Figure 1(b). In the visual-order approach, there is no difference whether $S_3^2(b)$ occurs before $R_1^2(a)$ or vice versa. However, turning towards more complex semantics of MSCs, this might not be true any longer. Suppose the two processes communicate via a one element buffer. Then the only linear execution we will see is that $R_1^2(a)$ occurs before $S_3^2(b)$. Thus, the set of linearisations of an MSC is no longer necessarily \sim -closed. It is indeed possible to model communication mediums by means of certain MFAs, which enrich a specification in form of MSCs [3].

4.2 Generalised Message-Passing Automata

The following automata model employs different automata components, each of which executes the actions of one single process. They communicate with each other over a window roughly as featured by an MFA. The length of this window is still bounded by a natural number r . The crucial point is that the window entries are no longer single send events (each paired with a natural number) but sequences of send events (each paired with a natural number and an additional message). To preserve \sim -closedness of the recognised languages, the components rather have to restrict themselves, whereas the window is a communication medium only. For example, we could imagine an automata component that has to keep a send action waiting until it executes a certain receive action, which, in turn, has to be preceded by a corresponding send action executed by another component. In fact, our view generalises the model proposed in [7], which has its origins in [15].

Definition 6 (Generalised Message-Passing Automaton). *A generalised message-passing automaton (GMPA) is a family $\mathcal{A} = (\{\mathcal{A}_p\}_{p \in \mathcal{P}_N}, r, \bar{q}_{in}, F, \text{Mess})$ of local automata together with a natural number $r \geq 1$, a global initial state \bar{q}_{in} , a set of global final states F , and a nonempty finite set of messages Mess . A local automaton is of the form $\mathcal{A}_p = (S_p, \Delta_p)$ where S_p is a nonempty finite set of local states and $\Delta_p \subseteq S_p \times (\Sigma_p \times \{1, \dots, r\} \times \text{Mess}) \times S_p$ is a set of local transitions (Σ_p contains the actions belonging to Process p). \bar{q}_{in} is an element and F a subset of $S_{\mathcal{A}} := \times_{p \in \mathcal{P}_N} S_p$, the set of global states of \mathcal{A} .*

For a GMPA \mathcal{A} , the (infinite) set of its configurations is defined by $\text{Conf}_{\mathcal{A}} := S_{\mathcal{A}} \times \{\chi \mid \chi : \Sigma_S \times \{1, \dots, r\} \rightarrow (\mathbb{N} \times \text{Mess})^*\}$. Let $\bar{s}[p]$ be the p th component of

a global state $\bar{s} \in S_{\mathcal{A}}$. Furthermore, for $W : \Sigma_{\mathcal{S}} \times \{1, \dots, r\} \rightarrow (\mathbb{N} \times \text{Mess})^*$, let $W[(\sigma, k) / w]$ denote the function that coincides with W with the exception that, for (σ, k) , it yields w . We define $\Longrightarrow_{\mathcal{A}} \subseteq \text{Conf}_{\mathcal{A}} \times (\Sigma \times \mathbb{N}) \times \text{Conf}_{\mathcal{A}}$ as follows:

- For $\sigma \in \Sigma_{\mathcal{S}}$ with $P(\sigma) = p$, $(\bar{s}, W) \xrightarrow{(\sigma, \pi)}_{\mathcal{A}} (\bar{t}, V)$ iff for all $k' \in \{1, \dots, r\}$ and $m' \in \text{Mess}$, (π, m') does not occur in $W(\sigma, k')$, and there is a transition $(\bar{s}[p], (\sigma, k, m), \bar{t}[p]) \in \Delta_p$ such that $V = W[(\sigma, k) / W(\sigma, k) \cdot (\pi, m)]$ and, for all $l \in \mathcal{P}_N - \{p\}$, $\bar{s}[l] = \bar{t}[l]$.
- For $\sigma \in \Sigma_{\mathcal{R}}$ with $P(\sigma) = p$ and $(\tau, \sigma) \in \text{Corr}$, $(\bar{s}, W) \xrightarrow{(\sigma, \pi)}_{\mathcal{A}} (\bar{t}, V)$ iff there are a transition $(\bar{s}[p], (\sigma, k, m), \bar{t}[p]) \in \Delta_p$ and a word $w \in (\mathbb{N} \times \text{Mess})^*$ such that $W(\tau, k) = (\pi, m) \cdot w$, $V = W[(\tau, k) / w]$, and, for all $l \in \mathcal{P}_N - \{p\}$, $\bar{s}[l] = \bar{t}[l]$.

A *run* of \mathcal{A} on a word $\sigma_1^{\pi_1} \dots \sigma_{\ell}^{\pi_{\ell}} \in (\Sigma \times \mathbb{N})^*$ is defined in analogy to the MFA case. That is, we are dealing with a sequence $(\bar{s}_0, W_0)(\bar{s}_1, W_1) \dots (\bar{s}_{\ell}, W_{\ell})$ of configurations such that $\bar{s}_0 = \bar{q}_{in}$, $W_0(\sigma, k) = \varepsilon$ for all $(\sigma, k) \in \Sigma_{\mathcal{S}} \times \{1, \dots, r\}$, and $(\bar{s}_{i-1}, W_{i-1}) \xrightarrow{(\sigma_i, \pi_i)}_{\mathcal{A}} (\bar{s}_i, W_i)$ for each $i \in \{1, \dots, \ell\}$. The run is *accepting* iff $\bar{s}_{\ell} \in F$ and $W_{\ell}(\sigma, k) = \varepsilon$ for all $(\sigma, k) \in \Sigma_{\mathcal{S}} \times \{1, \dots, r\}$. Finally, $\mathcal{L}(\mathcal{A}) := \{\alpha \mid \text{there is an accepting run of } \mathcal{A} \text{ on } \alpha\}$ denotes the language defined by \mathcal{A} .

Let $\text{Reach}(\mathcal{A})$ denote the set of configurations reachable within a run of \mathcal{A} . For $B \in \mathbb{N}$, we call \mathcal{A} *B-bounded* iff for all $(\bar{s}, W) \in \text{Reach}(\mathcal{A})$ and $\sigma \in \Sigma_{\mathcal{S}}$, $\sum_{k \in \{1, \dots, r\}} |W(\sigma, k)| \leq B$. We call it *bounded* iff it is *B-bounded* for some B .

Let us formulate the fundamental result of this section.

Theorem 9. *Let $\mathcal{L} \subseteq \text{MW}$ be an MSC word language. The following statements are equivalent:*

1. *There is a regular MSC language \mathcal{M} with $\text{Lin}(\mathcal{M}) = \mathcal{L}$.*
2. *\mathcal{L} is a \sim -closed regular MSC word language.*
3. *There is a bounded GMPA \mathcal{A} such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}$.*

Proof. The equivalence of 1. and 2. immediately follows from the definitions. Given a bounded GMPA, it is an easy task to define an equivalent MFA which shows that 3. implies 2. The other direction, however, is more involved and requires some results on regular Mazurkiewicz trace languages and related automata due to Zielonka [15]. We give a sketch of the proof and refer to the full version for the details.

Given a \sim -closed regular MSC word language \mathcal{L} , we build a bounded GMPA \mathcal{A} with $\mathcal{L}(\mathcal{A}) = \mathcal{L}$. The outline of this construction is as follows: We first observe that, for a certain B , $\text{nf}(\mathcal{L})^{\sim}$ can be considered to be a regular Mazurkiewicz trace language over $\Sigma \times \{1, \dots, B\}$ with an appropriate dependence alphabet. Then we can find an asynchronous automaton recognising $\text{nf}(\mathcal{L})^{\sim}$. The underlying distributed alphabet will comprise, apart from alphabets for each process, some additional components, which guarantee that the induced dependence relation complies with $D(\Sigma)$ (see also the proof of Theorem 4). These additional components have to be factored into the process components and the transition relation, making the transformation of the asynchronous automaton into a GMPA complicated. Concretely, the transitions synchronously taken by several

local automata have to be simulated by message passing. For example, consider Process P_1 sending a message to Process P_2 by executing (σ, k) . Actually, an equally labelled transition would have to be taken on the part of an additional component, in which (σ, k) is involved. But as in the GMPA such a component is not at the disposal of P_1 , P_1 guesses a corresponding move and writes it, along with the original message, into the message pool. The receiving process can take this message if the guessed move corresponds to the actual state of the additional component, which P_2 carries along. Our construction is similar to the one in [7] and uses the time-stamping protocol for non-FIFO computations described in [13] to ensure boundedness of the constructed GMPA \square

Thus, bounded GMPAs characterise exactly the \sim -closed regular MSC word languages and therewith exactly the regular MSC languages. For example, the 2-bounded GMPA \mathcal{A} given in Figure 5 recognises the $(\approx \cup \sim)^*$ -closure of

$$\left\{ \left(\begin{array}{cc} S_1^2(a) & S_1^2(a) \\ 1 & 2 \end{array} \ R_1^2(a) \ R_1^2(a) \ S_2^1(b) \ R_2^1(b) \right)^n \mid n \geq 0 \right\}.$$

5 A Logical Characterisation

We formulate a monadic second-order logic that characterises exactly the class of regular MSC languages. Given supplies $\text{Var} = \{x, y, \dots\}$ of individual variables and $\text{VAR} = \{X, Y, \dots\}$ of set variables, the syntax of $\text{MSO}(\mathcal{P}_N, \mathcal{A})$ is defined by

$$\varphi ::= L_\sigma(x) \mid x \in X \mid x \preceq y \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \exists x\varphi \mid \exists X\varphi \in \text{MSO}(\mathcal{P}_N, \mathcal{A})$$

where $\sigma \in \Sigma$, $x, y \in \text{Var}$, and $X \in \text{VAR}$. Given an MSC, individual variables are interpreted as events and set variables as sets of events. $L_\sigma(x)$ is satisfied if the event of x is labelled with σ , \preceq is interpreted as the partial order of the MSC, and the remaining constructs are defined as usual. We only consider formulas without free variables. For $\varphi \in \text{MSO}(\mathcal{P}_N, \mathcal{A})$ and $B \in \mathbb{N}$, let $\mathcal{M}_\varphi^B := \{M \mid M \text{ is } B\text{-bounded, } M \models \varphi\}$. We conclude with the fundamental result of this section.

Theorem 10. *Given a set \mathcal{M} of MSCs, \mathcal{M} is a regular MSC language iff there exist a formula $\varphi \in \text{MSO}(\mathcal{P}_N, \mathcal{A})$ and $B \in \mathbb{N}$ such that $\text{Lin}(\mathcal{M}) = \text{Lin}(\mathcal{M}_\varphi^B)$.*

The proof follows the outline of [7] although the concrete steps are different.




Fig. 5. A 2-bounded GMPA

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 *Lecture Notes in Computer Science*, pages 114–129. Springer, 1999.
2. João 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. Benedikt Bollig and Martin Leucker. Modelling, Specifying, and Verifying Message Passing Systems. In Claudio Bettini and Angelo Montanari, editors, *Proceedings of the Symposium on Temporal Representation and Reasoning (TIME'01)*, pages 240–248. IEEE Computer Society Press, June 2001.
4. Benedikt Bollig, Martin Leucker, and Thomas Noll. Generalised Regular MSC Languages. Technical Report AIB-03-2002, RWTH Aachen, January 2002.
5. Volker Diekert and Yves Métivier. Partial commutation and traces. In G. Rozenberg and A. Salomaa, editors, *Handbook on Formal Languages*, volume III. Springer, Berlin-Heidelberg-New York, 1997.
6. Elsa Gunter, Anca Muscholl, and Doron Peled. Compositional message sequence charts. In Tiziana Margaria and Wang Yi, editors, *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'01)*, volume 2031 of *Lecture Notes in Computer Science*, pages 496–511. Springer, April 2001.
7. J. G. Henriksen, M. Mukund, K. Narayan Kumar, and P. S. Thiagarajan. Regular collections of message sequence charts. In *Proceedings of 25th International Symposium on Mathematical Foundations of Computer Science (MFCS'2000)*, volume 1893 of *Lecture Notes in Computer Science*, pages 405–414. Springer, 2000.
8. ITU-TS. ITU-TS Recommendation Z.120anb: Formal Semantics of Message Sequence Charts. Technical report, ITU-TS, Geneva, 1998.
9. ITU-TS. ITU-TS Recommendation Z.120: Message Sequence Chart 1999 (MSC99). Technical report, ITU-TS, Geneva, 1999.
10. Michael Kaminski and Nissim Francez. Finite-memory automata. *Theoretical Computer Science*, 134(2):329–363, November 1994.
11. Dietrich 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*, Lecture Notes in Computer Science. Springer, 2002.
12. P. Madhusudan and B. Meenakshi. Beyond message sequence graphs. In *Proceedings of the 21st Conference on Foundations of Software Technology and Theoretical Computer Science*, Lecture Notes in Computer Science. Springer, 2001.
13. M. Mukund, K. Narayan Kumar, and M. Sohoni. Keeping track of the latest gossip in message-passing systems. Technical Report TCS-95-3, School of Mathematics, SPIC Science Foundation, Madras, India, 1995.
14. Frank Neven, Thomas Schwentick, and Victor Vianu. Towards regular languages over infinite alphabets. In *Proceedings of 26th International Symposium on Mathematical Foundations of Computer Science (MFCS'01)*, Lecture Notes in Computer Science. Springer, 2001.
15. Wiesław Zielonka. Notes on finite asynchronous automata. *R.A.I.R.O. — Informatique Théorique et Applications*, 21:99–135, 1987.