

Chapter 1

The Theory of Message Sequence Charts

K. Narayan Kumar

*Chennai Mathematical Institute
H1 SIPCOT IT Park, Siruseri, India.*

Message Sequence Charts or MSCs are a visual formalism used in the specification of systems in many domains including telecommunications, object oriented design and forms a part of the UML language. Consequently, the formal study of MSCs has received considerable attention over the last decade. We survey some of the key results in this area with particular emphasis on the notion of regularity and its relationship to automata, logics and model-checking.

1.1. Message sequence charts

Message Sequence Charts (MSCs) [37] are an appealing visual formalism used in a number of software engineering notational frameworks such as SDL [58] and UML [16]. An MSC is a representation of a single behaviour or run of a system consisting of a collection of processes communicating with each other asynchronously via buffered channels. In this paper we shall restrict ourselves to systems with first-in-first-out (FIFO) channels. Fig. 1.1 shows an MSC involving three processes p, q and r and three messages. An MSC is to be read from top to bottom, the vertical lines denote processes and the arrows represent messages. Formally, MSCs are defined as labelled partial orders.

Let $\mathcal{P} = \{p, q, r, \dots\}$ be a finite set of processes that communicate with each other through messages via reliable FIFO channels using a finite set of message types \mathcal{M} . For $p \in \mathcal{P}$, let $\Sigma_p = \{p!q(m), p?q(m) \mid p \neq q \in \mathcal{P}, m \in \mathcal{M}\}$ be the set of communication actions in which p participates. The action $p!q(m)$ is read as p sends the message m to q and the action $p?q(m)$ is read as p receives the message m from q . We set $\Sigma = \bigcup_{p \in \mathcal{P}} \Sigma_p$. We also denote the set of channels by $Ch = \{(p, q) \in \mathcal{P}^2 \mid p \neq q\}$. Whenever the set of processes \mathcal{P} is clear from the context, we write Σ instead of $\Sigma_{\mathcal{P}}$, etc. Observe that our notation restricts us to a single channel from a process p to any other process q , however this is merely for notational convenience and the results do not change if we permit multiple channels between pairs of processes.

Labelled posets A $\Sigma_{\mathcal{P}}$ -labelled poset is a structure $M = (E, \leq, \lambda)$ where (E, \leq)

is a partially ordered set and $\lambda : E \rightarrow \Sigma_{\mathcal{P}}$ is a labelling function. For $e \in E$, let $\downarrow e = \{e' \mid e' \leq e\}$. For $p \in \mathcal{P}$ and $a \in \Sigma_{\mathcal{P}}$, we set $E_p = \{e \mid \lambda(e) \in \Sigma_p\}$ and $E_a = \{e \mid \lambda(e) = a\}$, respectively. For $(p, q) \in Ch$, we define the relation $<_{pq}$:

$$e <_{pq} e' \stackrel{\text{def}}{=} \exists m \in \mathcal{M} \text{ such that } \lambda(e) = p!q(m), \lambda(e') = q?p(m) \text{ and} \\ |\downarrow e \cap E_{p!q(m)}| = |\downarrow e' \cap E_{q?p(m)}|$$

The relation $e <_{pq} e'$ says that channels are FIFO with respect to *each message*—if $e <_{pq} e'$, the message m read by q at e' is the one sent by p at e .

Finally, for each $p \in \mathcal{P}$, we define the relation $\leq_{pp} = (E_p \times E_p) \cap \leq$, with $<_{pp}$ standing for the largest irreflexive subset of \leq_{pp} . We write $Chn(e) = p!q$ if $\lambda(e) = p!q(m)$ and $Chn(e) = p?q$ if $\lambda(e) = p?q(m)$.

Definition 1.1. An MSC over \mathcal{P} is a *finite* $\Sigma_{\mathcal{P}}$ -labelled poset $M = (E, \leq, \lambda)$ where:

- (1) Each relation \leq_{pp} is a linear (total) order.
- (2) If $p \neq q$ then for each $m \in \mathcal{M}$, $|E_{p!q(m)}| = |E_{q?p(m)}|$.
- (3) If $e <_{pq} e'$, then $|\downarrow e \cap (\bigcup_{m \in \mathcal{M}} E_{p!q(m)})| = |\downarrow e' \cap (\bigcup_{m \in \mathcal{M}} E_{q?p(m)})|$.
- (4) The partial order \leq is the reflexive, transitive closure of $\bigcup_{p, q \in \mathcal{P}} <_{pq}$.

The second condition ensures that every message sent along a channel is received. The third condition says that every channel is FIFO across all messages.

In diagrams, the events of an MSC are presented in *visual order*. The events of each process are arranged in a vertical line and messages are displayed as horizontal or downward-sloping directed edges. Fig. 1.1 shows an example with three processes $\{p, q, r\}$ and six events $\{p_1, p_2, q_1, q_2, r_1, r_2\}$ corresponding to three messages— m_1 from p to q , m_2 from q to r and m_3 from p to r .




Fig. 1.1. An MSC

For an MSC $M = (E, \leq, \lambda)$, we let $\text{Lin}(M) = \{\lambda(\pi) \mid \pi \text{ is a linearization of } (E, \leq)\}$. For instance, $p!q(m_1) q?p(m_1) q!r(m_2) p!r(m_3) r?q(m_2) r?p(m_3)$ is one linearization of the MSC in Fig. 1.1. We write $\text{Lin}(M)$ for the set of linearizations of an MSC M .

Note that under the FIFO assumption an MSC can be reconstructed from any one linearization — the relation $<_{pp}$ is determined by the order of the p events in the linearization while $<_{pq}$ is determined by matching the i^{th} $p!q$ event with the i^{th} $q?p$ event.

Definition 1.2. Let M be an MSC and $B \in \mathbb{N}$. We say that $w \in \text{Lin}(M)$ is *B-bounded* if for every prefix v of w and for every channel $(p, q) \in Ch$, $\sum_{m \in \mathcal{M}} |v \upharpoonright \{p!q(m)\}| - \sum_{m \in \mathcal{M}} |v \upharpoonright \{q?p(m)\}| \leq B$, where $v \upharpoonright \Gamma$ denotes the projection of v on $\Gamma \subseteq \Sigma_{\mathcal{P}}$.




Fig. 1.2. A 3-bounded MSC

This means that along the sequential execution of M described by w , no channel ever contains more than B -messages. Consider the MSC in Fig. 1.2. The linearization $p!q \ q?p \ p!r \ p!q \ q?p \ q!r \ r?q \ p!q \ q?p \ r?p$ is 1-bounded while the linearization $p!q \ p!r \ p!q \ p!q \ q?p \ q?p \ p!r \ q?p \ r?q \ r?p$ is 3-bounded.

Definition 1.3. We say that M is *universally B -bounded* if every $w \in \text{Lin}(M)$ is B -bounded and that M is *existentially B -bounded* if there is a $w \in \text{Lin}(M)$ which is B -bounded. (We sometimes write B -bounded to mean universally B -bounded.)

The MSC in Fig. 1.2 is universally 3-bounded. The 1-bounded linearization listed earlier shows that this MSC is also existentially 1-bounded. The following proposition (paraphrased) from [34] characterizes universal B -boundedness for MSCs.

Proposition 1.1. *An MSC M is not B -bounded if and only if there are processes p, q and $p!q$ labelled events $e_1 <_{pp} e_2, \dots <_{pp} e_{B+1}$ such that the $q?p$ labelled event e' with $e_1 <_{pq} e'$ is not below e_{B+1} .*

Optimal linearizations: Given an MSC M the smallest B for which it is B -bounded can be computed in linear time as shown in [27]. In particular, an optimal linearization w.r.t. boundedness can be computed by the following greedy strategy — at each step, extend the linearization, if possible, by a receive event; otherwise extend it by picking from the set of candidate send events the one that minimizes the maximum number of undelivered messages in any channel. Applying this greedy strategy to the MSC in Fig. 1.2 we get $p!q \ q?p \ p!r \ p!q \ q?p \ q!r \ r?q \ r?p \ p!q \ q?p$ which is a 1-bounded linearization.

1.1.1. Concatenation of MSCs

The concatenation of M_1 and M_2 written as $M_1 \circ M_2$ denotes the behaviour where each process participates in its events in M_1 and then follows it by participating in its events in M_2 . It is not necessary for all processes to complete the events in M_1 before any process enters M_2 .

Formally, the (asynchronous) concatenation of MSCs is defined as follows.




Fig. 1.3. MSC Concatenation

Definition 1.4. Let $M_1 = (E^1, \leq^1, \lambda_1)$ and $M_2 = (E^2, \leq^2, \lambda_2)$ be a pair of MSCs such that E^1 and E^2 are disjoint. The (asynchronous) concatenation of M_1 and M_2 yields the MSC $M_1 \circ M_2 = (E, \leq, \lambda)$ where $E = E^1 \cup E^2$, $\lambda(e) = \lambda_i(e)$ if $e \in E^i$, $i \in \{1, 2\}$, and $\leq_{pp} = \leq_{pp}^1 \cup \leq_{pp}^2 \cup (E_p^1 \times E_p^2)$ and for $(p, q) \in Ch$, $\leq_{pq} = \leq_{pq}^1 \cup \leq_{pq}^2$.

MSCs over a given alphabet form a monoid with the empty MSC as the identity (i.e. concatenation is an associative operation). Fig. 1.3 describes the concatenation of two MSCs. In the linearization $p!r \ p!q \ q?p \ q!r \ r?q \ p!q \ q?p \ r?p$ of $A_1 \circ A_2$ observe that the last event of A_1 occurs after both the events of A_2 .

We can repeatedly decompose MSCs into the concatenation of smaller (nontrivial) MSCs till we are left with MSCs that cannot be decomposed.

Definition 1.5. An MSC M is said to be an *atom* if it cannot be expressed as the concatenation of two nontrivial MSCs.

The MSCs A_1 and A_2 in Fig. 1.3 are atoms. The MSC M in Fig. 1.2 is the concatenation $A_2 \circ A_1 \circ A_2$. In this case, the decomposition of M into atoms is unique. In general this need not be the case. To understand the exact nature of the decomposition of MSCs into atoms we need some terminology from the theory of traces.

1.1.2. MSCs and Traces

A *dependence alphabet* is a pair (Σ, D) where the alphabet Σ is a finite set of actions and the *dependence relation* $D \subseteq \Sigma \times \Sigma$ is reflexive and symmetric. The *independence relation* I is the complement of D . For $A \subseteq \Sigma$, the set of letters independent of A is denoted by $I(A) = \{b \in \Sigma \mid (a, b) \in I \text{ for all } a \in A\}$ and the set of letters depending on (some action in) A is denoted by $D(A) = \Sigma \setminus I(A)$.

A *Mazurkiewicz trace* is a labelled partial order $t = (V, \leq, \lambda)$ where V is a set of vertices labelled by $\lambda : V \rightarrow \Sigma$ and \leq is a partial order over V satisfying the following conditions: For all $x \in V$, the downward set $\downarrow x = \{y \in V \mid y \leq x\}$ is finite, $(\lambda(x), \lambda(y)) \in D$ implies $x \leq y$ or $y \leq x$, and $x < y$ implies $(\lambda(x), \lambda(y)) \in D$, where $< = \leq \setminus <^2$ is the immediate successor relation in t .

Let \equiv_I be the equivalence relation on Σ^* given by the reflexive transitive closure of the relation $uabv \sim ubav$ whenever $(a, b) \in I$. The set of linearizations of a trace t is an equivalence class of \equiv_I . Conversely, it is possible to reconstruct the trace from any given linearization. For a more detailed introduction to Trace theory, the reader is referred to ??.

We are now in a position to characterize the different decompositions of an MSC into atoms. Let A be any finite collection of atoms over the set of processes \mathcal{P} . One can equip A with natural dependence alphabet structure (A, D) where $(a, b) \in D$ if and only if there is a process p that is active in both a and b .

Proposition 1.2. *Let M be an MSC, A be a finite set of atoms that contains the atoms that appear in M and (A, D) be the corresponding dependence alphabet. Then, if $M = a_1 \circ a_2 \circ \dots \circ a_k$ and $M = a'_1 \circ a'_2 \circ \dots \circ a'_l$ then $a_1 a_2 \dots a_k \equiv_I a'_1 a'_2 \dots a'_l$.*

Thus, the decomposition of an MSC into atoms is unique up to commutations of independent atoms. In particular, the set of atoms that are necessary to decompose an MSC M , denoted by $\mathcal{A}(M)$, is unambiguously defined. We shall write $\mathcal{A}_t(M)$ to denote the trace (or equivalently the linearizations of the trace) over atoms associated with M .

There is a second dependence alphabet associated with MSCs which yields an alternative characterization of B -bounded MSCs. Let $\Sigma_B = \Sigma_{\mathcal{P}} \times \{0, 1, \dots, B-1\}$ for some natural number B . Let D be the dependence relation given by $(x, i) D (y, j)$ if either x and y occur in the same process or if $Chn(x) = p!q$, $Chn(y) = q?p$ and $i = j$. In a B -bounded MSC, the i^{th} receive on a channel must necessarily occur before the $(i+B)^{th}$ send on the same channel justifying the demand for an ordering between them in the dependency alphabet.

Given an MSC $M = (E, \leq, \lambda)$ we transform it into a Σ_B -labelled partial order $tr(M)$. Let $tr(M) = (E, \leq, \lambda')$ where $\lambda'(e) = (\lambda(e), i)$ where $i = |\{e' \mid e' < e, Chn(e') = Chn(e)\}|$. That is, we tag each event by its channel count modulo B . Does this yield a trace w.r.t the alphabet (Σ_B, D) ?

There are two kinds of dependencies in the definition of the relation D . A dependency of the first kind, within a process, is also enforced in the MSC. An event labelled $p!q$ is guaranteed to be below the corresponding receive event labelled $q?p$ in the MSC. However, no ordering is necessary between other occurrences of events labelled by these two letters. Thus, the second kind of dependency demanded by D is in general stronger than the ordering in the MSC. Fig. 1.4 describes the trace $tr(M)$ over the




Fig. 1.4.

alphabet $\Sigma_{\mathcal{P}} \times \{0, 1\}$ corresponding to the MSC in Fig. 1.2. Observe that there is no ordering between the first $q?p$ and the third $p!q$ in the MSC whilst they are ordered in the trace.

However, for B -bounded MSCS we have the following result.

Proposition 1.3. [38; 39] *Let M be an MSC. Then $tr(M)$ is a trace over the alphabet Σ_B if and only if M is a universally B -bounded MSC.*

Thus, universally B -bounded MSCs are traces over the aforementioned dependence alphabet. This allows us to exploit the well-developed theory of traces in the study of B -bounded MSCs.

Finally, let us examine existentially B -bounded MSCs and their relationship to traces. By definition, an MSC is not existentially B -bounded iff it cannot be linearized into a B -bounded word, and what rules out such a linearization is the following scenario — a $p!q$ event whose associated receive is above the next B $p!q$ -labelled events under \leq . This can be formalized as follows.

Definition 1.6. Let (E, \leq, λ) be a given MSC and B a natural number. The relation $<_{rev}$ on E , defined below, relates the i^{th} receive on a channel with the $(i + B)^{th}$ send on the same channel.

$$e <_{rev} e' \stackrel{\text{def}}{=} \exists f. f <_{pq} e \ \& \ \lambda(f) = \lambda(e') \ \& \ |\{e'' | f < e'' \leq e', \lambda(e'') = \lambda(f)\}| = B$$

Existentially B -bounded MSCs are precisely those that do not violate $<_{rev}$.

Proposition 1.4. [43] *Let $M = (E, \leq, \lambda)$ be an MSC and $<_{rev}$ be as defined above. M is existentially B -bounded if and only if $\leq \cup <_{rev}$ is acyclic. Existential B -boundedness of an MSC M can be decided in linear-time.*

In the MSC in Fig. 1.5, with $B = 2$, we find that the $<_{rev}$ edge marked with the dotted line induces a cycle. For $B = 3$, $<_{rev}$ is consistent with the ordering of this MSC implying that this MSC is existentially 3-bounded.

Let $M = (E, \leq, \lambda)$ and $\leq_B = (\leq \cup <_{rev})^*$. If M is existentially B -bounded then (E, \leq_B, λ) is a labelled partial order. Let $tr'(M)$ be the labelled partial order (E, \leq_B, λ') where λ' is the labelling function described earlier.

Proposition 1.5. [43] *If an MSC M is existentially B -bounded then the labelled partial order $tr'(M)$ is a trace over the alphabet (Σ_B, D) . (If M is not existentially B -bounded then $tr'(M)$ is not even a partial order, leave alone a trace.)*

It worth noting that for a universally B -bounded MSC M , $<_{rev} \subseteq \leq$ and therefore $tr'(M) = tr(M)$. Thus, Prop. 1.5 can be thought of as a generalization of Prop. 1.3.




Fig. 1.5.

1.1.3. MSC Languages and regularity

An *MSC language* is a (finite or infinite) collection of MSCs over a given set of processes \mathcal{P} and messages \mathcal{M} . Given the correspondence between MSCs and their linearizations, we may also regard a language of MSCs as a collection of words over $\Sigma_{\mathcal{P}}$ given by the linearizations of the MSCs in the language. In what follows we shall use these two notions interchangeably.

Definition 1.7. [36] A language of MSCs is said to be *regular* if the word language of its linearizations is a regular language.

In any prefix of a linearization of an MSC the number of $p!q$ events is at least as many as the number of $q?p$ events for any pair of processes p and q . We say that a word over $\Sigma_{\mathcal{P}}$ is *proper* if it satisfies this property. In any linearization of an MSC, there are as many $p!q$ events as there are $q?p$ events for every pair of processes p and q . We use *complete* to denote this property. Thus linearizations of MSCs are proper and complete words while their prefixes are proper words.

Let $(Q, \Sigma, \delta, s, F)$ be a deterministic finite automaton accepting a regular MSC language L . We further assume that every state is reachable and that a final state is reachable from every state. Suppose, u and v are two proper words that lead to the same state q from s . Let w be any word that leads from q to some final state. Thus, uw and vw are both complete words. Thus $|\#_{p!q}(v)| - |\#_{q?p}(v)| = |\#_{p!q}(v)| - |\#_{q?p}(v)|$ for each pair of processes p and q (where $\#_a(w)$ denotes the number of a 's in the word w). This leads to the following result from [36] which assures us that regular languages are bounded.

Lemma 1.1. *Every regular MSC language L is B -bounded for some B . In particular, B can be chosen to be smaller than the size of the minimal automaton accepting the linearizations of L .*

The converse of this lemma is obviously false — for instance, consider the language $\{(p!q \ q?p)^i \ (r!s \ s?r)^i \mid i \geq 0\}$. We end this section with the definition of finitely generated MSC languages.

Definition 1.8. A language \mathcal{L} of MSCs is said to be finitely generated if $\bigcup_{M \in \mathcal{L}} \mathcal{A}(M)$ is a finite set.

The MSC language given by the complete words $p!q \ (q!p \ p?q)^* \ q?p$ is not finitely generated. As a matter of fact, every word in this language is an atom.

1.2. Message Sequence Graphs

The ITU standard Z.120 describing MSCs also proposes a mechanism to describe collections of MSCs. This mechanism called HMSC (or High-level Message Sequence Charts) or Message Sequence Graphs (MSGs) allows branching, concatenation and iteration.

Definition 1.9. A *Message Sequence Graph* is a structure $\mathcal{G} = (Q, \rightarrow, Q_{in}, F, \Phi)$, where Q is a finite and nonempty set of states, $\rightarrow \subseteq Q \times Q$, $Q_{in} \subseteq Q$ is a set of initial states, $F \subseteq Q$ is a set of final states and Φ labels each state with an MSC.

A *path* π through an MSG \mathcal{G} is a sequence $q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_n$ such that $(q_{i-1}, q_i) \in \rightarrow$ for $i \in \{1, 2, \dots, n\}$. The MSC generated by π is $M(\pi) = M_0 \circ M_1 \circ M_2 \circ \dots \circ M_n$, where $M_i = \Phi(q_i)$. A path $\pi = q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_n$ is a *run* if $q_0 \in Q_{in}$ and $q_n \in F$. The language of MSCs accepted by \mathcal{G} is $L(\mathcal{G}) = \{M(\pi) \mid \pi \text{ is a run through } \mathcal{G}\}$. We say that an MSC language \mathcal{L} is *MSG-definable* if there exists an MSG \mathcal{G} such that $\mathcal{L} = L(\mathcal{G})$.

An example of an MSG is depicted in Fig. 1.6. The initial state is marked \Rightarrow and the final state has a double line. The MSC M corresponding to the path $q_0 \rightarrow q_1 \rightarrow q_0 \rightarrow q_2 \rightarrow q_0$ is also given in the figure.




Fig. 1.6. A message sequence graph

It can be verified that the language generated by the MSG in Fig. 1.6 is a regular language. However, this need not always be the case. There are two reasons why MSGs can generate non-regular languages.

The first reason, as illustrated in Fig. 1.7, is its combination of concurrency and iteration. The language \mathcal{L} generated by this MSG is $\{(M_1 \circ M_2)^n \mid n \geq 0\}$. The events in M_1 and M_2 are completely independent (concurrent) of each other. Thus, we may choose to linearize an MSC of the form $(M_1 \circ M_2)^n$ by first listing all the events involving p and q and then listing the events involving r and s . As a consequence, \mathcal{L} projected to $\{p!q(m), r!s(m)\}^*$ consists of $\sigma \in \{p!q(m), r!s(m)\}^*$ such that $|\sigma|_{p!q(m)} = |\sigma|_{r!s(m)}$, which is not a regular string language. Hence \mathcal{L} is not a regular MSC language.

The second reason, as illustrated by the producer-consumer example in Fig. 1.8, is that the buffers can be unbounded. The linearized language of this MSG is $\{w \mid \#_{p!q}w = \#_{q?p}w \ \& \ \forall v \leq w. \#_{p!q}v \geq \#_{q?p}v\}$.

However, the language of an MSG is always finitely generated since every MSC in the language can be decomposed using the MSCs that label the nodes of the MSG.




Fig. 1.7.

Proposition 1.6. [35] Let \mathcal{G} be a MSG. Let L_N be the set of MSCs that label the nodes of \mathcal{G} and let \mathcal{L} be the language generated by \mathcal{G} . Then, \mathcal{L} is a finitely generated MSC language and in particular, $\mathcal{A}(\mathcal{L}) = \mathcal{A}(L_N)$.

This also means that MSGs are not sufficient to describe every regular MSC language. For example, the language $p!r \ p!q \ q?p \ (q!r \ r?q \ r!q \ q?r)^* \ r?p$ is a regular MSC language that is not finitely generated and hence not MSG definable.




Fig. 1.8.

1.2.1. Communication Graph

The key to understanding the non-regularity of MSGs lies in studying their *communication graphs*.

Definition 1.10. For an MSC $M = (E, \leq, \lambda)$, CG_M , the *communication graph* of M , is the directed graph (\mathcal{P}, \mapsto) where:

- \mathcal{P} is the set of processes of the system.
- $(p, q) \in \mapsto$ iff there exists an $e \in E$ with $\lambda(e) = p!q(m)$.




Fig. 1.9.

The communication graph of the MSC M in Fig. 1.2 is in Fig. 1.9. This graph is not strongly connected. This means that M^* , the iteration of M , is not a bounded language. The reasoning goes as follows: After participating in the events in the first copy of M , the process p can go ahead and participate

in its events in the second copy and then the third copy and so on, before q or r participate in any event at all, forcing the channel from p to q to be unbounded. Suppose we modify the MSC M by adding a message from q to p . The iteration of M would still be unbounded — now the processes p and q can participate in all their events in copy one, and then copy two and so on before process r completes any event, forcing the channel from q to r to be unbounded. However the addition of an event from r to p would force the iteration of M to be bounded. But, this also makes the communication graph strongly connected.

Definition 1.11. A (communication) graph is *locally strongly connected* if the graph is the disjoint union of a collection of strongly connected components. An

MSG \mathcal{G} is locally strongly connected if the communication graph of the MSC generated by every cycle (simple loop) in \mathcal{G} is locally strongly connected.

The communication graph of every word in the language generated by the MSG in Fig. 1.7 is the same and is described in the Fig. 1.10. It is locally strongly connected and thus the MSG is locally strongly connected.

Suppose that the communication graph of an MSC M is locally strongly connected and let $X \subseteq \mathcal{P}$ be one of the strongly connected components. Let $p, q \in X$. Since there is a path from q to p in the communication



Fig. 1.10.

graph of length at most $|X|$, in any MSC of the form $M^{|X|}$ (i.e. $M \circ M \circ \dots \circ M$, $|X|$ times), there is a p event which is above the q events in the first M . Thus, p can at most be in the $(|X| + 1)^{st}$ copy before q completes its events in the first copy, ensuring that the channel from p to q is bounded. This argument does not rely on the fact that all the “copies” are identical but only uses the fact that all the copies have the same communication graph. Extending this argument gives

Lemma 1.2. [48; 43] *Let \mathcal{G} be an MSG and M be an MSC.*

- (1) *If the communication graph of M is locally strongly connected then M^* is a B -bounded MSC language for $B \geq |M| \times |\mathcal{P}|$.*
- (2) *If \mathcal{G} is locally strongly connected then $L(\mathcal{G})$ is a B -bounded language for any $B \geq |\mathcal{G}| \cdot |\mathcal{P}| \cdot \text{Max}$ where Max is the maximum number of send events in a MSC labelling any one node of \mathcal{G} .*
- (3) *If every node in \mathcal{G} is reachable from an initial node and in turn can reach an accepting node and $L(\mathcal{G})$ is a bounded MSC language then \mathcal{G} is locally strongly connected.*

Boundedness, a necessary condition for regularity, by itself does not guarantee regularity — the MSG in Fig 1.7 has a locally strongly connected communication graph yet generates a non-regular language (In particular iterating $M_1 \circ M_2$ generates a non-regular language). A further structural restriction on MSGs is needed to rule out the other reason for non-regularity in MSG definable languages — iterations of concurrent behaviours.

Intuitively, the reason for the non-regularity of $(M_1 \circ M_2)^*$ is that $(M_1 \circ M_2)^N = M_1^N \circ M_2^N$, due to the independence of the events in M_1 and M_2 , thus implicitly maintaining a counter.

Suppose the communication graph of an MSC M consists of a single nontrivial (i.e. of size at least 2) strongly connected component and a collection of trivial components (corresponding to each process that does not participate in any event in M). If p and q are two processes that participate in M , then in any segment of M^N of the form $M^{|\mathcal{P}|+1}$ there is a p event that depends on a previous q event within the segment and vice versa. Thus all independence is within small segments

(of size $M^{|\mathcal{P}|+1}$ or less) and the buffers are also bounded. This is sufficient ([6]) to ensure that M^* is a regular MSC language. This leads to the following definition.

Definition 1.12. An MSC M is said to be *locally synchronized* if its communication graph contains only one nontrivial strongly connected component and a collection of isolated vertices.

An MSG is said to be *locally synchronized* if the communication graph of the MSC generated by any loop in the MSG is locally synchronized. (Locally synchronized MSGs have also been called *bounded* MSGs or *com-connected* MSGs in literature.)

The following result shows that local synchronization is sufficient to guarantee regularity.

Lemma 1.3. [5; 54] *The language of any locally synchronized MSG is regular.*

Note that the definition of locally synchronized MSGs places a demand on all loops in the MSG and not just the cycles. Consider the MSG in Fig 1.11, adapted from [54]. The states q_0 and q_2 are labelled by the empty MSC. The states q_1 , q_3 and q_4 are labelled by the MSCs M_1 , M_3 and M_4 (described in the figure) respectively. Observe that every cycle in this MSG generates a locally synchronized MSC, however, the language accepted by this MSG is not regular. Every time the processes p and q switch from exchanging the message m to exchanging the message n (or vice versa) the processes r and s exchange a pair of messages. It is easy to derive the non-regularity from this observation.




Fig. 1.11.

Interestingly, if we prohibit the labelling of nodes by the empty MSC, then the definition of a locally synchronized MSG can be weakened.

Proposition 1.7. *Let \mathcal{G} be an MSG in which every node is labelled by a nontrivial MSC. If every cycle in \mathcal{G} describes a locally synchronized MSC then \mathcal{G} is locally synchronized.*

This can be seen as follows: Note that every loop that is not a cycle must properly contain a cycle. Let $p_0p_1 \dots p_k = q_0q_1 \dots q_m = p_{k+m} \dots p_0$ be a loop that contains the cycle $q_0q_1 \dots q_m$. The loop $p_0p_1 \dots p_kp_{k+m+1} \dots p_0$ is smaller and by the induction hypothesis generates a locally synchronized MSC. So does the cycle $q_0q_1 \dots q_m$. Any process that participates in the MSC labelling q_0 (and there is at

least one such process) is in the single nontrivial SCC of the communication graph of the smaller loop as well as the single nontrivial SCC of the communication graph of the cycle. Thus, the communication graph of the union of the loop and the cycle is also locally synchronized. Thus,

Lemma 1.4. *Let \mathcal{G} be an MSG in which every node is labelled by a nontrivial MSC. If every cycle in \mathcal{G} generates a locally synchronized MSC then the language of \mathcal{G} is regular.*

Earlier, we remarked that not all regular MSC languages can be described using MSGs. The following result characterizes the collection of MSG definable regular MSC languages.

Lemma 1.5. *[35; 48] A regular MSC language L is definable using MSGs if and only if it is finitely generated. Any such language can also be described using a locally synchronized MSG.*

The proof in [48] (which incidentally defines MSGs as regular expressions constructed using MSCs) exploits the translation from MSCs to traces over the underlying set of atoms and pulls back the corresponding result for traces ([57]).

We have seen that locally strongly connected MSGs are guaranteed to be bounded but place no restrictions on independent iterations. Interestingly, we can also exclude the complications of independent iterations without forcing boundedness (or regularity).

Definition 1.13. [28; 49] An MSC is said to be *globally cooperative* if the symmetric closure of its communication graph is the union of a single strongly connected component and a collected of isolated vertices.

An MSG \mathcal{G} is said to be *globally cooperative* if the MSC generated by any loop in \mathcal{G} is globally cooperative. (In [49], globally cooperative MSGs are called *c-HMSCs*.)

The motivation for this definition comes from trace theory. A word w over a dependence alphabet (Σ, D) is said to be *connected* if the graph $(\Sigma_w, D \downarrow \Sigma_w)$, on the letters that appear in w and the dependency relation restricted to these letters, consists of a single connected component. Fix a dependence alphabet (Σ, D) . A finite automaton over Σ in which every loop generates a connected word is said to be *loop connected*. Let the language L be the language of a loop connected finite automaton. Then, the language $\{tr(w) \mid w \in L\}$ of traces represented by words in L is a regular trace language. Equivalently the *trace-closure* of L , $\{w' \mid \exists w \in L. w' \sim w\}$, is a regular language ([18]).

We can easily transform an MSG into an equivalent one where every state is labelled by an atom — simply replace each node labelled by the MSC $A_1 \circ A_2 \dots \circ A_k$, by a sequence of k nodes each labelled by an atom. (This transformation takes a globally cooperative MSG to a globally cooperative MSG.) Thus, MSGs can be

thought of as finite automata over the alphabet of atoms where the states are labelled by letters instead of transitions.

Lemma 1.6. *Let \mathcal{G} be an MSG labelled by atoms and let $\mathcal{A}(\mathcal{G})$ be the set of atoms labelling the nodes of \mathcal{G} . Then, \mathcal{G} is globally cooperative if and only if it is loop connected as a finite automaton over the dependency alphabet $(\mathcal{A}(\mathcal{G}), D)$ where aDb whenever there is a process that is active in both a and b .*

This leads to the following regularity theorem for a globally cooperative MSGs. This proves to be a very useful tool in resolving a number of decision problems for globally cooperative MSGs.

Theorem 1.1. *[29] Let \mathcal{G} be a globally cooperative MSG. Then, the language $\{\mathcal{A}_t(M) \mid M \in L(\mathcal{G})\}$ is a regular language and a finite automaton accepting this language with size at most $2^{|\mathcal{G}| \cdot |\mathcal{P}|}$ can be constructed.*

The following result from [29] shows that the if you take away independent iteration what is needed to ensure regularity is the boundedness of channels.

Proposition 1.8. *An MSG \mathcal{G} is locally synchronized if and only if it is globally cooperative and accepts a bounded language.*

1.2.2. Decision problems

The language of an MSG is nonempty if and only if there is a path in the MSG from a start state to a final state. Thus, emptiness is decidable.

By Lemma 1.2, to decide whether a MSG accepts a universally bounded language it suffices to check if it is locally strongly connected. Every MSG accepts an existentially bounded language: in particular, if all the MSCs labelling the nodes of an MSG \mathcal{G} are existentially B -bounded, then the language of \mathcal{G} is also existentially B -bounded. The paper by Lohrey and Muscholl [43] establishes a comprehensive collection of the decidability results for a variety boundedness problems for MSGs and MSCs. Most importantly they demonstrate lower bounds for a variety of problems.

Theorem 1.2. *Let \mathcal{G} be an MSG.*

- (1) \mathcal{G} is always existentially bounded.
- (2) Given \mathcal{G} and B we can check whether \mathcal{G} is existentially B -bounded in linear time ([43]).
- (3) Checking whether \mathcal{G} is universally bounded is decidable ([48]). This problem is co-NP complete ([43]).
- (4) Given \mathcal{G} and B checking whether \mathcal{G} is universally B -bounded is co-NP complete. This problem is co-NP complete even if B is fixed to the constant 1 ([43]).

There are many factors that contribute to the size of an MSC or an MSG — the number of processes in \mathcal{P} , the size of the message alphabet \mathcal{M} and the number of events. Theorem 1.2 holds as it is even if \mathcal{M} is a singleton set ([43]). More often than not, the number of events is likely to be several orders of magnitude larger than the number of processes. Fortunately, if we fix the size of \mathcal{P} , all the above problems become efficiently solvable.

Theorem 1.3. [43] *Fix a set \mathcal{P} of processes. The problem of checking whether a MSG over \mathcal{P} is universally bounded (universally B -bounded for a given B) is in NL.*

The following lower-bound for checking structural properties of MSGs is from in [29]. Note that this result uses the assumption that the number of processes is part of the input.

Proposition 1.9. *Deciding whether a given MSG is locally synchronized (globally cooperative) is co-NP complete.*

Local synchronization provides a sufficient condition for regularity. However, there is no hope of obtaining an exact characterization.

Proposition 1.10. [34; 48] *Checking whether an MSG accepts a regular language is an undecidable problem.*

The reason for this undecidability is the following: It is easy to translate a dependency alphabet (Σ, D) into a collection of atoms \mathcal{A} over a set of processes \mathcal{P} by assigning an atom $\mathcal{A}(a)$ for each letter $a \in \Sigma$ in such a way that aDb if and only if $\mathcal{A}(a)$ and $\mathcal{A}(b)$ share an active process. For instance, the MSCs M_1 , M_3 and M_4 (from Fig. 1.11) can be used to represent the dependence alphabet $(\{a, b, c\}, (b, c))$. Using this representation we can transform finite automata over Σ into MSGs over (\mathcal{A}, D) . The language of this MSG is regular if and only if the trace-closure of the original language is regular. However, checking the regularity of trace-closure is in general an undecidable problem (see for eg. [56; 59]).

Model-checking: A specification describes a collection of behaviours. These could be a set of allowed behaviours that the system should conform to, or a set of disallowed behaviours that the system must avoid. This results in two versions of the model checking problem, the positive and negative model checking problems. In the positive model checking problem the task is to verify that the set of behaviours of the system, L_{sy} , is a subset of the set of behaviours, L_{sp} , described by the specification. In the negative model checking problem, the task is to verify that $L_{sy} \cap L_{sp} = \emptyset$.

Suppose the specification as well as the system are described by MSGs.

Theorem 1.4. [55; 5] *For MSGs the following problems are undecidable:*

- (1) Given \mathcal{G}_1 and \mathcal{G}_2 , is $L(\mathcal{G}_1) \subseteq L(\mathcal{G}_2)$?
- (2) Given \mathcal{G}_1 and \mathcal{G}_2 , is $L(\mathcal{G}_1) \cap L(\mathcal{G}_2) = \emptyset$?

Once again, translations from the corresponding problems in trace theory suffices to prove the undecidability. As a matter of fact if we rule out independent iterations these problems become decidable.

Theorem 1.5. [55] *For globally cooperative MSGs the following problems are decidable:*

- (1) Given \mathcal{G}_1 and \mathcal{G}_2 , is $L(\mathcal{G}_1) \subseteq L(\mathcal{G}_2)$?
- (2) Given \mathcal{G}_1 and \mathcal{G}_2 , is $L(\mathcal{G}_1) \cap L(\mathcal{G}_2) = \emptyset$?

The proofs exploit the regular representation via atoms provided by Theorem 1.1 — $L(\mathcal{G}_1) \subseteq L(\mathcal{G}_2)$ if and only if $\mathcal{A}_t(L(\mathcal{G}_1)) \subseteq \mathcal{A}_t(L(\mathcal{G}_2))$ and $L(\mathcal{G}_1) \cap L(\mathcal{G}_2) = \emptyset$ if and only if $\mathcal{A}_t(L(\mathcal{G}_1)) \cap \mathcal{A}_t(L(\mathcal{G}_2)) = \emptyset$. Theorem 1.1 makes globally cooperative MSGs perhaps the most general of the classes of MSGs amenable to algorithmic analysis.

At this point we turn our attention to a regularity property that holds for all MSG definable languages.

Definition 1.14. A set X of linearizations is a *set of representatives* for an MSC language L if $\{M \mid \text{Lin}(M) \cap X \neq \emptyset\} = L$.

Languages that have a regular set of representatives are needless to say interesting. If L is a regular MSC language then $\text{Lin}(L)$ is a regular set of representatives for L . However, as we shall see, the class of languages with regular set of representatives is much larger.

Let \mathcal{G} be an MSG. For each state $q \in \mathcal{G}$, fix a linearization w_q of the MSC labelling q . For any path $\pi = q_1 \rightarrow q_2 \rightarrow \dots \rightarrow q_k$, let $w(\pi) = w_{q_1} w_{q_2} \dots w_{q_k}$. Then, $\{w(\pi) \mid \pi \text{ is a run through } \mathcal{G}\}$ is a regular set as well as a set of representatives for $L(\mathcal{G})$.

Proposition 1.11. [44] *Every MSG definable language has a regular set of representatives.*

As such regular representations for L_{sy} and L_{sp} by themselves do not render the model-checking problems effective (as is clear from Theorem 1.4). Even if $L_1 \subseteq L_2$, it is easy to find representative sets X_1 and X_2 respectively in such a way that $X_1 \cap X_2 = \emptyset$. However,

Theorem 1.6. [25] *Suppose L_{sy} is given by a regular set of representatives X_{sy} and L_{sp} is a regular MSC language. Then, the positive and negative model checking problems are decidable. Thus, MSGs can be model-checked w.r.t. regular MSG specifications.*

In proof note that the positive model-checking problem boils down to checking if $X_{sy} \subseteq \text{Lin}(L_{sp})$ which is merely the containment of regular languages, and the negative model-checking involves deciding if $X_{sy} \cap \text{Lin}(L_{sp}) = \emptyset$.

An argument identical to the one used to prove Lemma 1.1 ensures that all the words in any regular representation of an MSC language are B -bounded for some B . Thus,

Proposition 1.12. *If L has a regular set of representatives then L is existentially B -bounded for some B .*

Let $\text{Lin}^B(L)$ be the set of linearizations of L that are B -bounded. Theorem 1.6 can be strengthened as follows: from a regular set of representatives X_{sy} for the system we can derive a bound B such that every word in X_{sy} is B -bounded. Now, $X_{sy} \subseteq \text{Lin}(L_{sp})$ if and only if $X_{sy} \subseteq \text{Lin}^B(L_{sp})$ and $X_{sy} \cap \text{Lin}(L_{sp}) = X_{sy} \cap \text{Lin}^B(L_{sp})$. Thus, it suffices that $\text{Lin}^B(L_{sp})$ be an (effectively constructible) regular set.

Theorem 1.7. [25] *Let L_{sy} be given by a regular set of representatives X_{sy} and let L_{sp} be such that $\text{Lin}^B(L_{sp})$ is effectively regular for some B such that every word in X_{sy} is B -bounded. Then, the positive and negative model-checking problems are decidable.*

Later in this section we shall see that B -bounded linearizations of any globally cooperative MSG language is a regular language. Moreover, as we shall see in Section 1.4 there is another natural class of systems for which B -bounded linearizations are regular for any B . These results, drawn from [25; 26] and [29] show that the model-checking problem for MSGs is decidable for a fairly generous class of specifications.


1.2.3. Compositional MSGs

There have been several attempts at extending the definition of MSGs to increase their expressive power. For instance, the *netcharts* model ([52; 9]) attempts at combining the features of MSGs and petri-nets to obtain a model that can generate all regular MSC languages. In this section we consider a natural weakening of the definition of MSCs and MSGs that results in a richer specification language and yet retains most of the useful properties of MSGs.

A compositional MSC is essentially a segment of an MSC, and thus may contain receive events without matching send events and send events without matching receive events. Any association between sends and receives included must satisfy the FIFO assumption. Formally,

Definition 1.15. [30] A CMSC over \mathcal{P} is a *finite* $\Sigma_{\mathcal{P}}$ -labelled poset $M = (E, \leq, \lambda, \text{msg})$ (with the notation defined in section 1.1) where

- (1) Each relation \leq_{pp} is a linear (total) order.

Fig. 1.12. A CMSC M and two elements of $M \circ M$

(2) msg is a partial injective mapping from S to R where

- $S = \{e \in E \mid \lambda(e) = p!q(m) \text{ for some } p, q, m\}$
- $R = \{e \in E \mid \lambda(e) = p?q(m) \text{ for some } p, q, m\}$

satisfying

- (a) if $msg(s) = r$ then $s = p!q(m)$ and $r = q?p(m)$ for some p, q, m . We write $s <_{pq} msg(s)$ in that case.
- (b) if $s_1 \leq_{pp} s_2$, $Ch(s_1) = Ch(s_2) = p!q$ and $msg(s_1)$ and $msg(s_2)$ are defined then $msg(s_1) \leq_{qq} msg(s_2)$.
- (3) $\leq = (\cup_{p \in \mathcal{P}} \leq_{pp} \cup \cup_{(p,q) \in Ch} <_{pq})^*$.

In Fig. 1.12 we have a CMSC with one unmatched send and one unmatched receive events.

Every MSC is a CMSC. It is easy to check that a CMSC is an MSC if and only if msg is total and onto. Following [26], we define the concatenation of two CMSCs M_1 and M_2 as a set of CMSCs: for each process, the events in M_1 precede its events in M_2 , further send events in M_1 may be matched with receive events in M_2 , as long as it does not violate the FIFO condition. The result is a set as we are not obliged to match up unmatched sends in M_1 with unmatched receives in M_2 and there may be more than one way to match up unmatched sends in M_1 with unmatched receives in M_2 .

Definition 1.16. Let $M_i = (E_i, \leq_i, \lambda_i, msg_i)$, $i = 1, 2$ be CMSCs with $E_1 \cap E_2 = \emptyset$. The concatenation $M_1 \circ M_2$ is the collection of CMSCs of the form $M = (E_1 \cup E_2, \leq, \lambda, msg)$ where

- (1) $M \downarrow_{E_i} = M_i$ for $i = 1, 2$, where $M \downarrow_F$ is the restriction of \leq , λ and msg to the events in F .
- (2) For each $e \in E_2$, if $e \leq e'$ then $e' \in E_2$ (i.e.) send events of E_2 cannot be matched with receive events in E_1 .

Figure 1.12 illustrates CMSC concatenation by listing two CMSCs that belong to $M \circ M$. The FIFO assumption ensures that if $M_1 \circ M_2$ contains an MSC then it is unique. The operation \circ can be extended to sets of CMSCs, $S \circ T = \{M \circ M' \mid M \in S \ \& \ M' \in T\}$. On sets of MSCs, the operation \circ is associative, i.e., $S \circ (T \circ U) = (S \circ T) \circ U$. We generalize MSGs to CMSGs in the obvious manner.

Definition 1.17. [30] A *Compositional Message Sequence Graph* is a structure $\mathcal{G} = (Q, \rightarrow, Q_{in}, F, \Phi)$, where Q is a finite and nonempty set of states, $\rightarrow \subseteq Q \times Q$, $Q_{in} \subseteq Q$ is a set of initial states, $F \subseteq Q$ is a set of final states and Φ labels each state with a CMSC.

A path $\pi = q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_n$ is an *accepting run* if $q_0 \in Q_{in}$ and $q_n \in F$. The language of MSCs accepted by \mathcal{G} , $L(\mathcal{G})$, is the set of MSCs in the set

$$\{\Phi(q_0) \circ \Phi(q_1) \circ \dots \circ \Phi(q_k) \mid q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_k \text{ is an accepting run of } \mathcal{G}\}$$

We say that an CMSC language \mathcal{L} is *CMSG-definable* if there exists an CMSG \mathcal{G} such that $\mathcal{L} = L(\mathcal{G})$.

First of all note that some CMSGs may generate an empty MSC language even though there are paths from the initial state to final states. For instance, a CMSG with a single state that is initial and final, a self-loop, and labelled by the CMSC M from Fig. 1.12 is one such CMSG. Any CMSC generated by this CMSG has unmatched sends (and receives).

Observe that the CMSC $M_1 \circ M^i \circ M_2$, where M is the CMSC from Fig 1.12, M_1 is the CMSC with just a single $p!q$ event and M_2 is the CMSC with a single $q?p$ event, is an atom for all $i \geq 0$. Thus, CMSG recognizable languages need not be finitely generated. As a matter of fact, as shown in [30], it is quite easy to show that any regular MSC language is CMSG-definable — roughly speaking, we may replace each edge labelled $p!q$ (or $q?p$) in the finite automaton for this language by a node labelled by an MSC with a single $p!q$ (or $q?p$) event. This construction actually implies something stronger:

Proposition 1.13. *If L is an MSC language with a regular set of representatives then L is a CMSG-definable language.*

Interestingly, since every path in this CMSG corresponds to a valid linearization of an MSC in L , every path in this CMSG generates at least one MSC. Checking whether a CMSG generates any MSC or not is undecidable ([30]). Thus, CMSGs are somewhat unrestrained for a specification language.

Definition 1.18. [26] A CMSG \mathcal{G} is said to be *safe* if for any accepting path $q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_k$ the set $\Phi(q_0) \circ \Phi(q_1) \circ \dots \circ \Phi(q_k)$ contains at least one MSC.

Every MSG is a safe CMSG. The property of being safe is decidable and safety gives a sufficient condition to ensure analyzability of CMSGs. Fix a linearization w_q for each node q of a safe CMSG \mathcal{G} . It is easy to see that the language

$\{w_{q_0}w_{q_1}\dots w_{q_k} \mid q_0 \rightarrow q_1 \rightarrow \dots q_k \text{ is an accepting run}\}$ is regular and a set of representatives of the language of \mathcal{G} . Combining this with the observation following Proposition 1.13 gives

Proposition 1.14. [45] *An MSC language L has a regular set of representatives if and only if it is the language of a safe CMSG. Thus, every safe CMSG language is existentially B -bounded for some B .*

Safe CMSGs should be as analyzable as the class of MSGs; after all, they enjoy the only regularity property that we have been able to associate with MSGs! To verify this, let us generalize some of the structural restrictions on MSGs to CMSGs.

Definition 1.19. [26] The *communication graph* of a CMSC is the directed graph whose vertices are the elements of \mathcal{P} and there is an edge from p to q whenever there is a $p!q$ event and a $q?p$ event (the two need not necessarily be matched) in the CMSC.

The communication graph of the CMSC M in Fig. 1.12 is the complete directed graph on two vertices. The notions of *locally synchronized* CMSCs and *globally cooperative* CMSCs is defined as before. Finally, a CMSG \mathcal{G} is *locally synchronized* if it is safe and every CMSC generated by every loop in \mathcal{G} is locally synchronized. A CMSG \mathcal{G} is *globally cooperative* if it is safe and every CMSC generated by every loop in \mathcal{G} is globally cooperative.

Theorem 1.8. [26] *Let \mathcal{G} be a globally cooperative CMSG and let B be an integer such that $B \geq |\mathcal{G}|$. Then, the set of B -bounded linearizations of $L(\mathcal{G})$ is a regular language (recognized by a finite automaton whose size is $O(|G|^{Poly(|G|, B, |\mathcal{P}|)})$).*

The proof relies on the relationship to Mazurkiewicz traces. Combining this with Theorem 1.7 we get

Theorem 1.9. [26] *The positive and negative model checking problems are decidable when L_{sy} is the language of a safe CMSG and L_{sp} is the language of a globally cooperative CMSG.*

The essence of this story, which took some time to develop and culminate in the papers [29; 26], is the following: of the two suspected reasons for the non-analyzability of MSG (CMSG) based specifications, the culprit is independent (concurrent) iterations and eliminating that via a structural restriction (globally cooperative MSGs) delivers a generous decidability result for model-checking.

1.3. Monadic second order logic over MSCs

Monadic second order logic (or MSO) is the logical counter part to automata. Büchi and Elgot ([17; 22]) showed that MSO over finite words has the same expressive power as finite automata, Büchi then extended this to MSO over infinite words

and automata over infinite words that bear his name. These connections are not isolated and a host of similar connections have been established between MSO and automata. The most relevant of these results to our context is the one relating regular trace languages and MSO over traces due to Wolfgang Thomas ([61]) and extended to infinite traces by Ebinger and Muscholl ([21]).

Definition 1.20. Fix a set \mathcal{P} of processes and messages. The formulas of the monadic second order logic over MSCs (MSO) are as follows:

$\varphi ::= a(x) \mid x \in X \mid x \leq y \mid x \leq_{pp} y \mid x \leq_p y \mid x \leq_{pq} y \mid \neg\varphi \mid \varphi \wedge \varphi \mid \exists X.\varphi \mid \exists x.\varphi$
where $a \in \Sigma_{\mathcal{P}}$ and $(p, q) \in Ch$.

We will also be interested in the fragment of existential monadic second-order formulas (EMSO) which are of the form $\exists X_1 \exists X_2 \dots \exists X_k. \varphi$ where φ is a first order formula. We will also be interested in restricted versions of these logics obtained by permitting only a subset of the 4 relational symbols in the syntax, and this will be made explicit by listing the allowed subset: for eg. $\text{MSO}(\leq_p, \leq_{pq})$ to stand for the fragment that does not use \leq and \leq_{pp} .

An MSO formula is interpreted over an MSC. The first order variables range over the events in the MSC, second order variables over sets of events and the relational operators have the obvious interpretation: \leq is the ordering on the events of the MSC, \leq_{pp} is the ordering on events in process p , \leq_p is the immediate successor relation within a process p , and \leq_{pq} is the message induced ordering between a send from p and the corresponding receive in q . Observe that the first two relations are ordering relations while the latter two are not. Finally $a(x)$ is true with $x = e$ if $\lambda(e) = a$. The interpretation of the logical operators and quantifiers is as usual. It is quite easy to see that \leq_{pp} and \leq_p can be defined using \leq , but it turns out that \leq_{pq} cannot be so defined. Sentences in MSO define languages of MSCs, $L(\varphi) = \{M \mid M \models \varphi\}$ and in this case we say that L is MSO-definable.

Here is a sentence that characterises universally 2-bounded MSCs.

$$\bigwedge_{(p,q) \in Ch} \forall x. \forall y. \forall z. (p!q(x) \wedge p!q(y) \wedge p!q(z) \wedge (x < y) \wedge (y < z)) \implies \exists w. (x <_{pq} w) \wedge (w < z))$$

It asserts that in any sequence of 3 sends, the receive corresponding to the first send must be in the past of the third send (see Prop. 1.1). This can be generalized to describing universally B -bounded MSCs for any fixed B .

It is well-known that the transitive closure of a relation is definable in any monadic second order logic: xR^*y if and only if the smallest set containing x and closed under R also contains y . However, such a definition makes essential use of universal quantification over sets and consequently, such a translation is not always possible in the existential fragment of monadic second order logic. Since \leq and

\leq_{pp} can be defined as transitive closure of $\leq_p \cup <_{pq}$ and \leq_p respectively, MSO over MSCs is equivalent to $\text{MSO}(\leq_p, <_{pq})$.

In [36; 34] the following characterization theorem is presented.

Theorem 1.10. *A B -bounded language L is regular if and only if $L = L(\varphi) \cap \{M \mid M \text{ is universally } B\text{-bounded}\}$ for some $\text{MSO}(\leq)$ formula φ .*

The result holds even if the logic is restricted to be $\text{EMSO}(\leq)$. The proof ([36; 34]) is based on the ideas used in similar results for MSO over traces in [61; 60]. In one direction, given a $\text{MSO}(\leq)$ formula φ , we construct a formula ψ in MSO over words such that $w \models \psi$ if and only if (i) w is a B -bounded complete and proper word and (ii) the MSC M_w generated by w satisfies φ and (iii) M_w is universally B -bounded. Establishing (i) and (ii) shows that the set of B -bounded linearizations of $L(\varphi)$ is a regular language (which implies Theorem 1.11 below).

Part (i) is easy as the set of B -bounded proper and complete words is a regular language and one may appeal to the Büchi-Elgot theorem. For part (ii), the proof proceeds by defining, in MSO over words, a binary relation \preceq on the positions of the word in such a way that for any B -bounded w and positions i and j in w , $i \preceq j$ if and only if the corresponding events (say e_i and e_j) are ordered under \leq in M_w . Clearly, $e_i \leq_{pp} e_j$ in M_w if and only if $i < j$ and i and j are p events in w . We still have to show that $<_{pq}$ is definable. This makes essential use of the B -boundedness of w . The formula asserts that there is a $k \in \{0, 1, \dots, B-1\}$ such that i is labelled by $p!q$, the number of positions to its left labelled by $p!q$ is $k(\text{modulo } B)$, j is to the right of i , it is labelled by $q?p$ and the number of positions to its left labelled by $q?p$ is $k(\text{modulo } B)$ and there is no position between i and j labelled by $q?p$ for which the number of positions to its left labelled $q?p$ is $k(\text{modulo } B)$. Since transitive closure is definable in MSO the result follows. (The transitive closure can be avoided by a slightly more elaborate argument using the fact that one has to hop across processes at most $|\mathcal{P}|$ number of times.) Finally, part (iii) follows from the fact that with \preceq and \leq , universally B -boundedness can be defined as explained in the example above.

The other direction is somewhat more involved. Given a regular MSC language and B , use Büchi-Elgot theorem to pick a formula ψ in MSO over words describing the linearizations of this language. Then show that in $\text{MSO}(\leq)$ one can define a relation \preceq over the MSC that fixes a canonical linearization of the MSC (using techniques from [60]), and then interpret ψ over this linearization.

This result generalizes to MSO over infinite MSCs and regular languages of infinite MSCs as shown by D. Kuske ([39]). As a matter of fact, [39] provides a complete theory of regular languages of infinite MSCs. For the relationship between existentially B -bounded languages and MSO we have the following theorem:

Theorem 1.11. *[45] For any MSO formula φ and any B , $\text{Lin}^B(L(\varphi))$ is a regular language.*

An immediate consequence is that any existentially B -bounded language described by an MSO formula has a regular set of representatives. As a corollary to the previous two theorems we have

Corollary 1.1. *The problem of checking whether an $MSO(\leq)$ formula (or a MSO formula) is satisfiable over universally B -bounded MSCs is decidable. Similarly, checking satisfiability over existentially B -bounded MSCs is decidable.*

The natural question then is to ask “Can we model check safe CMSGs (or MSGs) w.r.t. to MSO?”. The answer is affirmative. In proof note that, any safe CMSG has a regular set of representatives, Theorem 1.11 implies that $\text{Lin}^B(L(\varphi))$ is regular for any B and any φ in MSO and thus Theorem 1.7 is applicable.

Theorem 1.12. [44; 45] *The problem of deciding whether every MSC generated by a safe CMSG (or MSG) satisfies a MSO formula is decidable.*

B.Bollig and M. Leucker [15] study the expressiveness of MSO and EMSO over MSCs and using techniques from [46] show that

Theorem 1.13.

- (1) *The monadic quantifier alternation hierarchy of the logic MSO (over MSCs) is infinite. Thus, $MSO(<_p, <_{pq})$ is strictly more expressive than $EMSO(<_p, <_{pq})$.*
- (2) *The logics $MSO(\leq)$ and $EMSO(<_p, <_{pq})$ are incomparable.*

We shall return to the expressive power of MSO over MSCs a little later after we introduce an implementation model for MSCs.

1.4. Message Passing Automata

Safe CMSGs, MSGs and MSO are elegant and expressive languages to describe collections of MSCs. However, they are far removed from an execution model where each process is situated at a different location and there are limitations on what each process actually knows of the global state. The natural execution model for MSCs is that of *message passing automata* (also referred to as *communicating finite-state machines*).

A message passing automaton consists of a collection of finite state processes which communicate with each other by sending messages on FIFO channels. Each transition in a process involves either sending a message to some process or consuming a message from one of its input channels. It is possible to enrich these automata with local moves, however since it has no effect on the results in this section, we work without them. Fig. 1.13 illustrates a message passing automaton implementing a producer-consumer system and an MSC accepted by it. Formally, an MPA is defined as follows:




Fig. 1.13. An MPA

Definition 1.21. [13] Let $\Sigma_{\mathcal{P}}$ be the communication alphabet over the set of processes \mathcal{P} and message alphabet \mathcal{M} . A *message-passing automaton (MPA)* over $\Sigma_{\mathcal{P}}$ is a structure $\mathcal{A} = (\{\mathcal{A}_p\}_{p \in \mathcal{P}}, \Delta, s_{in}, F)$ where:

- Δ is a finite alphabet of *auxiliary messages*.
- Each component \mathcal{A}_p is of the form (S_p, \rightarrow_p) where S_p is a finite set of p -local states and $\rightarrow_p \subseteq S_p \times \Sigma_p \times \Delta \times S_p$ is the p -local transition relation.
- $s_{in} \in \prod_{p \in \mathcal{P}} S_p$ is the global initial state.
- $F \subseteq \prod_{p \in \mathcal{P}} S_p$ is the set of global final states.

Observe that our definition allows the *tagging* of each message with auxiliary contents drawn from the set Δ .

The local transition relation \rightarrow_p specifies how the process p sends and receives messages. The transition $(s, p!q(m), x, s')$ says that in state s , p can send the message m to q tagged with auxiliary information x and move to state s' . Similarly, the transition $(s, p?q(m), x, s')$ signifies that at state s , p can receive the message m from q tagged with information x and move to state s' .

A global state of \mathcal{A} is an element of $\prod_{p \in \mathcal{P}} S_p$. For a global state s , s_p denotes the p th component of s . A *configuration* is a pair (s, χ) where s is a global state and $\chi : Ch \rightarrow (\mathcal{M} \times \Delta)^*$ is the *channel state* describing the message queue in each channel c . The *initial configuration* of \mathcal{A} is $(s_{in}, \chi_\varepsilon)$ where $\chi_\varepsilon(c)$ is the empty string ε for every channel c . The set of *final configurations* of \mathcal{A} is $F \times \{\chi_\varepsilon\}$. Observe that in a final configuration all the channels must be empty.

A global move of the automaton involves one of the process depositing a message into a channel (sending a message) or consuming a message from the channel (receiving a message) according to its local transition relation. Suppose, (s, χ) is a configuration and $(s_p, p!q(m), x, s'_p) \in \rightarrow_p$. Then, $(s, \chi) \xrightarrow{p!q(m)} (s', \chi')$ where for $r \neq p$, $s_r = s'_r$, for each $r \in \mathcal{P}$, $\chi'((p, q)) = \chi((p, q)) \cdot (m, x)$, and for $c \neq (p, q)$, $\chi'(c) = \chi(c)$. Similarly, if (s, χ) is a configuration and $(s_p, p?q(m), x, s'_p) \in \rightarrow_p$, then there is a global move $(s, \chi) \xrightarrow{p?q(m)} (s', \chi')$ where for $r \neq p$, $s_r = s'_r$, for each $r \in \mathcal{P}$, $\chi'((q, p)) = (m, x) \cdot \chi'((q, p))$, and for $c \neq (q, p)$, $\chi'(c) = \chi(c)$.

A run of the automaton is a sequence of such global moves and we write $Conf_{\mathcal{A}}$




Fig. 1.14. An MPA accepting infinite number of atoms

for the set of reachable configurations of \mathcal{A} . A run is accepting if it ends in a final configuration. For instance

$$((p, q), \varepsilon) \xrightarrow{p!q(m)} ((p, q), m) \xrightarrow{p!q(m)} ((p, q), mm) \xrightarrow{q?p(m)} ((p, q), m) \xrightarrow{q?p(m)} ((p, q), \varepsilon)$$

is an accepting run the automaton in Fig. 1.13 on the word $p!q(m)p!q(m)q?p(m)q?p(m)$.

We define $L(\mathcal{A}) = \{\sigma \mid \mathcal{A} \text{ has an accepting run over } \sigma\}$. Since all channels are empty in the initial and final configurations and a message can be received only if it has already been sent, it is easy to check that any word accepted by an MPA is proper and complete. It is not difficult to see that if $L(\mathcal{A})$ contains one linearization of an MSC M then it contains all the linearizations of the MSC M . As a matter of fact, it is quite easy to define runs of MPAs directly on MSCs as a mapping from events on the MSC to global states of the automaton. Thus $L(\mathcal{A})$ is the set of linearizations of an MSC language. As usual we shall use $L(\mathcal{A})$ to denote the MSC language accepted by \mathcal{A} as well as its linearizations.

Each configuration of an MPA records the messages sent and as yet undelivered. For $B \in \mathbb{N}$, a configuration (s, χ) is B -bounded if $|\chi(c)| \leq B$ for every channel $c \in Ch$ and \mathcal{A} is a B -bounded automaton if every reachable configuration $(s, \chi) \in Conf_{\mathcal{A}}$ is B -bounded. Clearly a B -bounded automaton accepts a universally B -bounded language. The global state space of any B -bounded MPA is therefore finite and consequently, every B -bounded MPA accepts a regular MSC language. The converse is also true but we shall get to that a little later. The MPA in Fig. 1.13 accepts a unbounded language.

From any MPA \mathcal{A} and a natural number B we can generate a finite automaton accepting precisely those B -bounded words that are accepted by \mathcal{A} and thus

Proposition 1.15. *For any MPA \mathcal{A} and any natural number B , the language $\text{Lin}^B(L(\mathcal{A}))$ is a regular language.*

The language of an MPA need not be finitely generated. Fig. 1.14 describes an MPA and one of the MSCs it accepts. This MPA accepts a regular MSC language




Fig. 1.15. An MPA with an existentially unbounded language

and every MSC accepted by this automaton is an atom. Thus, there are MPA acceptable languages that cannot be described by MSGs.

Finally, MPAs can also accept languages that are not existentially bounded (for any B), thus MPAs are capable of accepting languages that cannot be described using safe CMSGs. The MPA in Fig. 1.15 accepts the language of MSCs generated by the words $(p!q)^n p!r r?p r!q q?r (q?p)^n$ and it is easy to verify that this is not an existentially bounded language.

Definition 1.22. An MPA is said have local accepting states if $F = \prod_{p \in \mathcal{P}} F_p$ for some $F_p \subseteq S_p$.

1.4.1. MPAs without auxiliary messages

We first examine the power of MPAs without auxiliary message alphabets (i.e.) MPAs whose auxiliary alphabet is singleton.

Definition 1.23. An MSC language L is said to be *weakly realizable* if it is the language of an MPA with a singleton auxiliary message alphabet and with local accepting states.

Consider the language $\{M_1, M_2\}$ of MSCs (from Fig. 1.16). This set is not weakly realizable — Suppose \mathcal{A} is an MPA accepting this language. From the accepting run on M_1 , we know that there are runs $p_0 \xrightarrow{p!q(m)} p_1 \xrightarrow{p!s(m)} p_2$ and $q_0 \xrightarrow{q?p(m)} q_1$ for p and q ending in accepting states. Similarly, from the accepting run on M_2 , we know that there are runs $r_0 \xrightarrow{r!s(m)} r_1$ and $s_0 \xrightarrow{s?r(m)} s_1 \xrightarrow{s?p(m)} s_2$ ending in accepting states. This means that the MSC M is also accepted, since p and q may behave exactly as they do in accepting M_1 and r and s may behave exactly as they do in accepting M_2 and all four processes end up in an accepting states on M .

The language of an MPA with local accepting states and over a singleton auxiliary message alphabet is merely the shuffle or free product of the local languages




Fig. 1.16. An implied scenario

of the processes and this is formalized as follows:

Given an MSC M (or any linearization w of M) and a process p , $M|_p$ is the word over Σ_p consisting of the projection of M (or equivalently w) to the events in process p . For eg. $M_1|_p = p!q(m)p!s(m)$ for the MSC M_1 in Fig. 1.16. For a language L , $L_p = \{M|_p \mid M \in L\}$. Finally, given word languages L_p over Σ_p for each $p \in \mathcal{P}$, let $\prod_{p \in \mathcal{P}} L_p = \{M \mid M|_p \in L_p\}$ (the usual free product).

Definition 1.24. Given a set L of MSCs its *implied closure* $\text{Imp}(L)$ is defined as follows

$$\text{Imp}(L) = \{M \mid \forall p \in \mathcal{P}. \exists M_p \in L. M|_p = M_p|_p\}$$

If $M \in \text{Imp}(L) \setminus L$ then we say that L has an *implied scenario* and that M is an implied scenario of L .

In Fig 1.16, M is an implied scenario of $\{M_1, M_2\}$. The following characterization (albeit non-effective) is easy to prove.

Proposition 1.16. [6] *If L is weakly realizable then $L = \text{Imp}(L)$. Conversely, suppose L is an MSC language and $L|_p$ is a regular language for each p , then L is weakly realizable only if $L = \text{Imp}(L)$.*

Implied scenarios are of practical interest. Often, a designer specifies a system as a collection of MSCs using say an MSG. The existence of an implied scenario indicates that an implementation by MPAs (w/o auxiliary tagging) would result in behaviours not foreseen by the designer (these might or might not be *bad*). So it would be useful to check if a given MSC language L has any implied scenarios at all, and construct a representation for $\text{Imp}(L)$. The implied closure of a B -bounded language may contain MSCs that are not B -bounded. As




Fig. 1.17.

a matter of fact, the implied closure of a B -bounded regular language need not be bounded at all.

In Fig. 1.17 observe that the two MSCs M_1 and M_2 have complete communication graphs. Therefore, the language $(M_1 + M_2)^*$ is a regular MSG-definable language. On the other hand, for each $k \in \mathbb{N}$, the MSC in which the p -projection matches $M_1^{2k} M_2^k$ and the q -projection matches $M_2^k M_1^{2k}$ has a global cut where the channel (p, q) has capacity $k + 1$. The figure shows the case $k = 2$. The dotted line marks the global cut where the channel (p, q) has maximum capacity. Thus, $\text{Imp}(L)$ need not be a regular MSC language when L is a regular MSC language. It gets worse.

Theorem 1.14. [7] *The problem of checking whether $\text{Imp}(L) = L$ is undecidable even for regular MSC languages presented as locally synchronized MSGs.*

Let us examine this result a little. Let B be the bound on the channels in L . From the definition of $\text{Imp}(L)$ it is not difficult to check that if L is weakly realizable then an MPA \mathcal{A} implementing L can be constructed as follows: For each process p pick a minimal finite automaton (w/o dead or unreachable states) accepting the language $L|_p$ as \mathcal{A}_p and set $F = \prod_{p \in \mathcal{P}} F_p$. Thus we have a candidate implementation. Yet, weak realizability is undecidable because it is not possible to check whether the language accepted by \mathcal{A} equals L . By restricting \mathcal{A} to runs where no channel has more than B messages, we have a finite automaton accepting the set of B -bounded words in $\text{Imp}(L)$. Thus, checking if there is a B -bounded implied scenario for L is decidable. The difficulty is in finding if there are implied scenarios violating the B -bound. In particular, the existence of a partial run reaching a configuration where some channel has more than B messages does not mean that there are implied scenarios. This is because such a partial run may not be extendable to an accepting run. (However, this means that this partial run has ended in a configuration from where no final configuration is reachable. We shall return to this point a little later.)

Theorem 1.14 has been strengthened [12] to show that this problem is undecidable even when L is a 1-bounded language and undecidability holds even with just 2 processes. These undecidability arguments make essential use of the fact that the channels are FIFO (and therefore requires the message alphabet to have at least two messages.) Restriction to the trivial message alphabet yields the first positive result.

Proposition 1.17. [49] *The problem of checking whether the language of a locally synchronized MSG over a singleton message alphabet is weakly realizable is decidable.*

Since weak realizability is impossible to analyze, it is natural to look for stronger notions of implementability. Alur et al. propose a notion called *safe realizability* which is amenable to algorithmic analysis.

Definition 1.25. A configuration χ of an MPA \mathcal{A} is a *deadlock* if there are no reachable final configurations. An MPA is said to be *deadlock-free* if it has no reachable deadlock configurations. A language is said to be *safely realizable* if it is the language of a deadlock-free MPA with local accepting states and over a singleton auxiliary message alphabet.

It is also possible to characterize safely realizable languages as a closure property akin to Prop. 1.16 (see [6; 42]).

Proposition 1.18. *L is safely realizable if and only if it is weakly realizable and satisfies the following closure property:*

$$\{w \mid \forall p. \exists u_p \in L. w \upharpoonright p \leq u_p\} \subseteq \{w \mid \exists u \in L. w \leq u\}$$

The closure condition demands that any partial MSC (or proper word) whose projections on every process is consistent with some accepting run of the process must be extendable to an MSC (or proper complete word) in L . Fortunately, safe realizability is analyzable.


Theorem 1.15. *The problem of checking whether a given globally cooperative MSG generates a safely realizable language or not is decidable in EXPSPACE and the problem is EXPSPACE-complete even for locally synchronized MSGs. However, the safe realizability problem for arbitrary MSGs is undecidable.*

We go back to our analysis of why the availability of a candidate implementation does not suffice to ensure decidability of weak implementation. The analysis there ends with a conclusion placed within parenthesis. This conclusion stated in our recently acquired terminology states that the candidate implementation \mathcal{A} has deadlocks or implied scenarios whenever it reaches a configuration violating the B -bound on some channel. Thus, if \mathcal{A} ever reaches a configuration violating the B -bound on some channel, it cannot be a safe implementation of L . Finally, it is not difficult to verify, using Prop. 1.18 that if at all L is safely realizable then \mathcal{A} is such a realization. That completes our sketch of the decidability argument.

The decidability for locally synchronized MSGs appears in [7], that for globally cooperative MSGs in [49] and the exact complexity result as well as the undecidability result is from [42].

1.4.2. MPAs with auxiliary messages

In this section we study the expressive power of MPAs with no restriction on the use of auxiliary message contents. We begin by showing that the collection $\{M_1, M_2\}$ from Fig. 1.16 can be implemented using auxiliary messages. An MPA with auxiliary messages drawn from the set $\{1, 2\}$ is described in figure 1.18 The process p signals the process s using the auxiliary information regarding the identity of the MSC (The auxiliary information in all the other messages can be ignored.)

Fig. 1.18. An MPA accepting $\{M_1, M_2\}$

This example illustrates the use of auxiliary information: it allows a process to convey information about its past (in this case p conveys the information “I sent a message to q ” to s by tagging the message with a 1 instead of a 2.) However, since the auxiliary message alphabet is a finite set, it only allows a bounded amount of information about the past to be conveyed in any message. It turns out that this ability to forward bounded amount of information is very powerful.

Theorem 1.16. [36; 34] *An MSC language L over a set of process \mathcal{P} and messages \mathcal{M} is regular if and only if there is an MPA \mathcal{A} , over the same alphabet and with an auxiliary message alphabet Δ , such that $L = L(\mathcal{A})$.*

The proof of this theorem is beyond the scope of this paper. However we provide a brief sketch of the difficulties and main ideas involved.

This theorem is an example of a distributed synthesis theorem — it states that given the global description of a regular MSC language, it is possible to construct a distributed implementation as an MPA. The key ingredients that go into the proof are drawn from the celebrated result of W. Zielonka [64] showing that every regular trace language is recognized by an *asynchronous automaton*.

Let us examine the main difficulty in proving such a result. Suppose the global description is a finite automaton G . We can equip each local process with a copy of G if necessary. Yet, after a sequence of events w , which process is to keep track of the current state of G ? Observe that every event takes place only in one process and each process directly observes only the events that it participates in, so it is not possible for any one process to maintain the state of G correctly. For the moment assume that we may tag the messages with unbounded amount of auxiliary information. Then, every process can send the entire history of all the events it has participated in as well as all the events about which it has learnt from others through messages it has received. So, whenever a message is received, the receiving process knows the entire set of events that are below this event in the MSC order. However, an MSC could have up to $|\mathcal{P}|$ maximal events and thus even

with this passing around of unbounded amount of information, no one process has information about the entire MSC.

Suppose the processes are $1, 2, \dots, K$. Process 1 has with it all the events that occur below the maximal 1 event in the MSC. Now, we would like process 2 to provide us with not all the events in its past, but only those that appear in its past but not in the past of process 1. For this, we need information about the events in each process $j \in \{2, 3, 4, \dots, K\}$, that are in the past of 2 but not in the past of 1 (called the 1 residue at 2). With this information we can piece together all the events in the past of 1 and 2 without any ambiguity. Then, we need to obtain from process 3 information about events in $\{3, 4, \dots, K\}$ that do not appear in the past of 1 and 2 (the $\{1, 2\}$ residue at 3) and so on. Then, the MSC can be reconstructed from the residues available at $1, 2, \dots, K$.

The reduction from unbounded to bounded information hinges on the fact that instead of keeping any partial MSC M (or a linearization w) of the history of a process, we might as well keep the transition function that this word or partial MSC defines on the state space of G .

The ability to program each process to maintain its residues w.r.t. to other processes requires a sophisticated time stamping algorithm from [51] which can perform comparisons such as “is my information about j more recent than i ’s information about j ?”. Most importantly this time-stamping algorithm requires each process to maintain only a bounded amount of information and tags each message with only a bounded amount of information. The proof in [50; 34] proceeds along these lines.

An alternative is to use the connection to regular trace languages, then use Zielonka’s construction to obtain an asynchronous automaton and then translate back such an automaton into an MPA. This is the structure of the proof in [38]. The above theorem can be strengthened.

Definition 1.26. [50] An MPA is said to be deterministic if

- If $(s, p!q, m_1, s'_1) \in \rightarrow_p$ and $(s, p!q, m_2, s'_2) \in \rightarrow_p$ then $m_1 = m_2$ and $s'_1 = s'_2$.
- If $(s, p?q, m, s'_1) \in \rightarrow_p$ and $(s, p?q, m, s'_2) \in \rightarrow_p$ then $s'_1 = s'_2$.

Determinacy requires that the nature of the message sent from p to q depends only on the local state of the sender p . Note, however, that from the same state, p may have the possibility of sending messages to more than one process. When receiving a message, the new state of the receiving process is fixed uniquely by its current local state and the content of the message. Once again, a process may be willing to receive messages from more than one process in a given state. This definition ensures that a deterministic automaton has at most one run on any $w \in \Sigma_{\mathcal{P}}^*$. Now, we are in a position to state the main characterization theorem for regular MSC languages.

Theorem 1.17. [34] *Let B be any integer and L be a language of universally B -bounded MSCs. Then the following are equivalent:*

- (1) *L is a regular MSC language*
- (2) *L is definable in $MSO(\leq)$ (or $EMSO(\leq)$)*
- (3) *L is the language of a deterministic B -bounded MPA.*
- (4) *L is the language of a B -bounded MPA.*

This theorem has been generalized to infinite MSCs by D.Kuske (see [39]). There are couple of other points worth noting: first of all the definition of MPAs uses a global set of final states and the physical realizability of such a global set of acceptance states is debatable. Secondly, the automata constructed in the proof of the above theorem may deadlock, and once again this makes its usability somewhat limited. In a recent series of papers, N. Baudru and R. Morin [8; 10], have shown that every regular MSC language can be implemented using (nondeterministic) MPAs that are deadlock free and whose acceptance set is local (i.e. $F = \prod_{p \in \mathcal{P}} F_p$ for some collection $(F_p)_{p \in \mathcal{P}}$).

1.4.3. Implementing existentially bounded languages

The characterization in Theorem 1.17 shows that every regular MSC language can be implemented using a deterministic MPA. The corresponding question for MSC languages with a regular set of representatives was solved by Genest, Kuske and Muscholl [25; 26]. An earlier paper [29] set the stage for such a result through a host of results on the virtues of existentially bounded languages including the analyzability of globally cooperative CMSGs (Theorem 1.7) and efficient implementability (as MPAs) for a subclass (locally cooperative MSGs) of globally cooperative CMSGs. The proof of this characterization uses the translation to traces from existentially B -bounded MSCs (described in Section 1.1) to Mazurkiewicz traces.

Theorem 1.18. [26] *Let B be an integer and let L be a language of existentially B -bounded MSCs. Then the following statements are equivalent.*

- (1) *$\text{Lin}^B(L)$ is a regular set of representatives for L .*
- (2) *L is generated by a globally cooperative CMSG.*
- (3) *L is MSO ($EMSO$) definable.*
- (4) *L is implementable using an MPA.*

We have already seen some of the relationships: Theorem. 1.8 shows that one can go from globally cooperative CMSGs to languages with $\text{Lin}^B(L)$ as a regular set of representatives. Theorem 1.11 allows us to move from MSO definable languages to languages with $\text{Lin}^B(L)$ as a regular set of representatives. Prop. 1.13 shows that any language with a regular set of representatives can be translated to a safe CMSG and Proposition 1.14 shows that every safe CMSG generates a language with a regular set of representatives. These relationships are strengthened to restrict the

class of CMSGs to globally cooperative CMSGs in [26] using Ochmanski's theorem ([57]). Finally by Prop. 1.15 the set $\text{Lin}^B(L)$ is regular for any MPA.

This leaves the difficult part: a decomposition theorem showing that every existentially B -bounded language has a distributed implementation as an MPA and this involves a fairly complex argument via Mazurkiewicz trace theory and is beyond the scope of this article. The automaton constructed is a nondeterministic automaton and this is unavoidable.

Proposition 1.19. [27] *There are existentially bounded languages recognized by nondeterministic MPAs that cannot be recognized by deterministic MPAs.*

Translating MPAs to EMSO is quite routine, but the converse is not. A remarkable result due to Bollig and Leucker shows the following:

Theorem 1.19. [15] *An MSC language L is $\text{EMSO}(<_p, <_{pq})$ definable if and only if it is the language of an MPA.*

In contrast to Theorems 1.17 and 1.18 this theorem applies to arbitrary MSCs. This theorem in combination with Theorem 1.13 shows that

Theorem 1.20. [15] *The class of languages accepted by MPAs is not closed under complementation.*

1.4.4. Decision Problems

Every channel is a queue and it is quite easy to simulate counter machines with MPAs. Consequently, general MPAs are too expressive for any sort of analysis:

Theorem 1.21. [27]

- (1) *The language emptiness problem for deterministic MPAs is undecidable.*
- (2) *Given a B checking whether a deterministic MPA accepts a universally B -bounded language is undecidable for every $B > 0$.*
- (3) *Given a deterministic, deadlock-free automaton in which every global state is accepting, checking whether it accepts a universally bounded language is undecidable.*
- (4) *Given a B checking whether a deterministic MPA accepts a existentially B -bounded language is undecidable for every $B > 0$.*
- (5) *Given a deterministic, deadlock-free automaton in which every global state is accepting, checking whether it accepts a existentially bounded language is undecidable.*

Items (3) and (5) become decidable for a fixed B , since for a deadlock-free automaton it suffices to check if a configuration with $B + 1$ messages in the channel is reachable.

1.5. Conclusion

In this article we have surveyed a selection of the results from the theory of MSCs. These are by no means exhaustive. A number of results that have been omitted here due to lack of space — to name a few, there is a host of decidability (and undecidability) results for the so called “pattern matching” problems on MSCs (see for instance [53; 24]), the related problem of whether a class of MSCs can be implemented by MPAs with additional messages (not merely the addition of auxiliary content to existing messages; see for instance [23]), branching time specification and analysis (see for instance [19; 32; 33] and finally there have been some recent results in extending the theory of MSCs with time (see for instance [1; 2; 3]).

Acknowledgments

This work was partially supported by *Timed-DISCOVERI*, a project under the Indo-French Networking Project.

References

1. S. Akshay, Benedikt Bollig, Paul Gastin: Automata and Logics for Timed Message Sequence Charts. *Proceedings of the 27th International Conference on Foundations of Software Technology and Theoretical and Theoretical Computer Science*, Lecture Notes in Computer Science **4855**, Springer-Verlag (2007) 290–302
2. S. Akshay, Madhavan Mukund, K. Narayan Kumar: Checking Coverage for Infinite Collections of Timed Scenarios. *Proceedings of the 18th International Conference on Concurrency Theory (CONCUR’07)*, Lecture Notes in Computer Science, Springer-Verlag **4703** (2007) 181–196
3. S. Akshay, Benedikt Bollig, Paul Gastin, Madhavan Mukund, K. Narayan Kumar: Distributed Timed Automata with Independently Evolving Clocks. *Proceedings of the 19th International Conference on Concurrency Theory (CONCUR’07)*, Lecture Notes in Computer Science, Springer-Verlag **5201** (2008) 82–97
4. Alur, R., Holzmann, G. J., Peled, D.: An analyzer for message sequence charts. *Software Concepts and Tools* **17**(2) (1996) 70–77
5. Alur, R., Yannakakis, M.: Model checking of message sequence charts. *Proceedings of the 10th International Conference on Concurrency Theory (CONCUR’99)*, Lecture Notes in Computer Science **1664**, Springer-Verlag (1999) 114–129
6. Alur, R., Etassami, K., Yannakakis, M.: Inference of message sequence graphs. *Proceedings of the 22nd International Conference on Software Engineering (ICSE 2000)*, Association for Computing Machinery (2000) 304–313.
7. Alur, R., Etassami, K., Yannakakis, M.: Realizability and Verification of MSC Graphs. *Proceedings Automata, Languages and Programming, 28th International Colloquium (ICALP 2001)*, Lecture Notes in Computer Science **2076**, Springer-Verlag (2001) 797–808.
8. N.Baudru, R. Morin: Safe Implementability of Regular Message Sequence Chart and Specifications. *Proceedings of the ACIS Fourth International Conference on Software*

- Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing*, (SNPD'03), (2003) 210–217.
9. N. Baudru, R. Morin: The Synthesis Problem of Netcharts. *Proceedings of the 27th International Conference on Applications and Theory of Petrinets and other Models of Concurrency*, (ICATPN'06), Lecture Notes in Computer Science **4024**, Springer-Verlag (2006) 84–104
 10. N. Baudru, R. Morin: Synthesis of Safe Message-Passing Systems. *Proceedings of the 27th International Conference on Foundations of Software Technology and Theoretical and Theoretical Computer Science*, (FSTTCS'07), Lecture Notes in Computer Science **4855**, Springer-Verlag (2007) 277–289
 11. Ben-Abdallah, H., Leue, S.: Syntactic detection of process divergence and non-local choice in message sequence charts. *Proceedings of the 3rd Workshop on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'97)*, Lecture Notes in Computer Science **1217**, Springer-Verlag (1997) 259–274
 12. Puneet Bhateja, Paul Gastin, Madhavan Mukund, K. Narayan Kumar: Local Testing of Message Sequence Charts Is Difficult *16th International Symposium on the Fundamentals of Computation Theory*, (FCT'07), Lecture Notes in Computer Science, Springer-Verlag **4639** (2007) 76–87
 13. D. Brand, P. Zafiropulo: On Communicating Finite-state Machines. *Journal of the ACM* **30** (2) (1983) 323–342
 14. Bollig, B., Leucker, M., Noll, T.: Generalised regular MSC languages. *Proceedings of the 5th International Conference on Foundations of Software Science and Computation Structures (FOSSACS'02)*, Lecture Notes in Computer Science **2303**, Springer-Verlag (2002) 52–66
 15. Bollig, B., Leucker, M.: Message-passing automata are expressively equivalent to EMSO Logic. *Theoretical Computer Science* **358** (2-3) (2006) 150–172
 16. Booch, G., Jacobson, I., Rumbaugh, J.: *Unified Modeling Language User Guide*. Addison-Wesley (1997)
 17. Büchi, J. R.: Weak second-order arithmetic and finite automata. *Z. Math. Logik Grundl. Math.* **6** (1960) 66–92
 18. Clerbout, M., Latteux, M.: Semi-commutations. *Information and Computation* **73**(1) (1987) 59–74
 19. Damm, W., Harel, D.: LSCs: Breathing life into message sequence charts. *Formal Methods in System Design* **19**(1) (2001) 45–80.
 20. Diekert, V., Rozenberg, G. (Eds.): *The Book of Traces*. World Scientific (1995)
 21. Ebinger, W., Muscholl, A.: Logical definability on infinite traces. *Theoretical Computer Science* **154**(1) (1996) 67–84
 22. Elgot, C.C.: Decision problems of finite automata design and related arithmetics. *Transactions of the American Mathematical Society*. **98** (1960) 21–52.
 23. Blaise Genest: On Implementation of Global Concurrent Systems with Local Asynchronous Controllers. *Proceedings of the 16th International Conference on Concurrency* (CONCUR'05), Lecture Notes in Computer Science, Springer-Verlag **3653** (2005) 443–457
 24. Blaise Genest, Anca Muscholl: Pattern Matching and Membership for Hierarchical Message and Sequence Charts. *Proceedings of the 5th Latin American Symposium on Theoretical Computer Science* (LATIN'02), Lecture Notes in Computer Science, Springer-Verlag **2286** (2002) 326–340
 25. Blaise Genest, Dietrich Kuske and Anca Muscholl: A Kleene Theorem for a Class of Communicating Automata with Effective Algorithms. *Developments in Language Theory*, Lecture Notes in Computer Science, Springer-Verlag **3340** (2004) 30–48

26. Blaise Genest, Dietrich Kuske and Anca Muscholl: A Kleene Theorem for a Class of Communicating Automata with Effective Algorithms. *Information and Computation* **204** (6) (2006) 920–956
27. Blaise Genest, Dietrich Kuske and Anca Muscholl: On Communicating Automata with Bounded Channels. *Fundamenta Informaticae* **80**(1-3) (2007) 147–167
28. Blaise Genest, Anca Muscholl, Helmut Seidl, Marc Zeitoun: Infinite-State High-Level MSCs: Model-Checking and Realizability. *Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP'02)*, Lecture Notes in Computer Science **2380**, Springer-Verlag (2002) 657–668
29. Blaise Genest, Anca Muscholl, Helmut Seidl, Marc Zeitoun: Infinite-State High-Level MSCs: Model-Checking and Realizability. *Journal of Computer and System Sciences* **72** (4) (2006) 617–647
30. Elsa L. Gunter, Anca Muscholl, Doron Peled: Compositional Message Sequence Charts. *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'01)*, Lecture Notes in Computer Science, Springer-Verlag **2031** (2001) 496–511.
31. Harel, D., Gery, E.: Executable object modeling with statecharts. *IEEE Computer*, July 1997 (1997) 31–42
32. D. Harel and R. Marelly: Specifying and Executing Behavioral Requirements: The Play In/Play-Out Approach. *Software and System Modeling (SoSyM)* (to appear)
33. Harel, D., Kugler, H., Marelly, R., Pnueli, A.: Smart Play-out of Behavioral Requirements. *Proceedings Formal Methods in Computer-Aided Design, 4th International Conference (FMCAD 2002)*, Lecture Notes in Computer Science **2517** Springer-Verlag (2002) 378–398.
34. J.G. Henriksen, M. Mukund, K. Narayan Kumar, M.A. Sohoni and P.S. Thiagarajan: A Theory of Regular MSC Languages. *Information and Computation* **202**(1) (2005) 1–38
35. J.G. Henriksen, M. Mukund, K. Narayan Kumar and P.S. Thiagarajan: On Message Sequence Graphs and Finitely Generated Regular MSC Languages. *Proceedings of the International Colloquium on Automata, Languages and Programming 2000 (ICALP'00)*, Lecture Notes in Computer Science **1854**, Springer-Verlag (2000) 675–686
36. J.G. Henriksen, M. Mukund, K. Narayan Kumar and P.S. Thiagarajan: Regular Collections of Message Sequence Charts. *Proceedings of the Conference on the Mathematical Foundations of Computer Science 2000 (MFCS'00)*, Lecture Notes in Computer Science **1893**, Springer-Verlag (2000) 405–414
37. ITU-TS Recommendation Z.120: *Message Sequence Chart (MSC)*. ITU-TS, Geneva (1997)
38. Kuske, D.: A further step towards a theory of regular MSC languages. *Proceedings of the Symposium on the Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science **2285**, Springer-Verlag (2002) 489–500
39. Kuske, D.: Regular sets of infinite message sequence charts. *Information and Computation* **187**(1) (2003) 80–109
40. Ladkin, P. B., Leue, S.: Interpreting message flow graphs. *Formal Aspects of Computing* **7**(5) (1995) 473–509
41. Levin, V., Peled, D.: Verification of message sequence charts via template matching. *Proceedings of the 7th International Conference on Theory and Practice of Software Development (TAPSOFT'97)*, Lecture Notes in Computer Science **1214**, Springer-Verlag (1997) 652–666
42. Lohrey, M.: Realizability of high-level message sequence charts: closing the gaps.

- Theoretical Computer Science* **309** (1-3) (2003) 529–554
43. M. Lohrey, A. Mushcoll: Bounded MSC communication. *Information and Computation* **189**(2) (2004) 160–181
 44. Madhusudan, P.: Reasoning about sequential and branching behaviours of message sequence graphs. *Proceedings of the 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*, Lecture Notes in Computer Science **2076**, Springer-Verlag (2001) 396–407
 45. Madhusudan, P., Meenakshi, B.: Beyond message sequence graphs. *Proceedings of the 21st Conference on the Foundations of Software Technology and Theoretical Computer Science (FSTTCS'01)*, Lecture Notes in Computer Science **2245**, Springer-Verlag (2001) 256–267
 46. Matz, O., Thomas, W.: The Monadic Quantifier Alternation Hierarchy over Graphs is Infinite. *Proceedings of the Symposium on Logic in Computer Science (LICS)*, IEEE (1997) 236–244
 47. Mauw, S., Reniers, M. A.: High-level message sequence charts, *Proceedings of the 8th SDL Forum, SDL'97: Time for Testing — SDL, MSC and Trends*, Elsevier (1997).
 48. R. Morin: On Regular Message Sequence Chart Languages and Relationships to Mazurkiewicz Trace Theory. *Proceedings of the 4th International Conference on Foundations of Software Science and Computation Structures*, FOSSACS'01, Lecture Notes in Computer Science **2030**, Springer-Verlag (2001) 332–346
 49. R. Morin: Recognizable Sets of Message Sequence Charts. *Proceedings of the 19th Annual Symposium on the Theoretical Aspects of Computer Science*, STACS'02, Lecture Notes in Computer Science **2285**, Springer-Verlag (2002) 523–534
 50. Mukund, M., Narayan Kumar, K., Sohoni, M.: Synthesizing distributed finite-state systems from MSCs. *Proceedings of the 11th International Conference on Concurrency Theory (CONCUR 2000)*, Lecture Notes in Computer Science, Springer-Verlag (2000).
 51. Mukund, M., Narayan Kumar, K., Sohoni, M.: Bounded time-stamping in message-passing systems. *Theoretical Computer Science* **290**(1) (2003) 221–239
 52. Mukund, M., Narayan Kumar, K., Thiagarajan, P.S.: Netcharts: Bridging the gap between HMSCs and executable specifications. *Proceedings of the 14th International Conference on Concurrency Theory (CONCUR'03)*, Lecture Notes in Computer Science, Springer-Verlag **2761** (2003) 293–307
 53. Muscholl, A.: Matching specifications for message sequence charts. *Proceedings of the 2nd International Conference on Foundations of Software Science and Computation Structures (FOSSACS'99)*, Lecture Notes in Computer Science **1578**, Springer-Verlag (1999) 273–287
 54. Muscholl, A., Peled, D.: Message sequence graphs and decision problems on Mazurkiewicz traces. *Proceedings of the 24th International Symposium on Mathematical Foundations of Computer Science (MFCS'99)*, Lecture Notes in Computer Science **1672**, Springer-Verlag (1999) 81–91
 55. Muscholl, A., Peled, D., Su, Z.: Deciding properties for message sequence charts. *Proceedings of the 1st International Conference on Foundations of Software Science and Computation Structures (FOSSACS'98)*, Lecture Notes in Computer Science **1378**, Springer-Verlag (1998) 226–242
 56. Mushcoll, A., Petersen, H.: Note on the Commutative Closure of Star-Free Languages. *Information Processing Letters* **57** (2) (1996) 71–74
 57. Edward Ochmanski: Regular behaviour of concurrent systems. *Bulletin of the EATCS* **27** (1985) 56–67
 58. Rudolph, E., Graubmann, P., Grabowski, J.: Tutorial on message sequence charts. In *Computer Networks and ISDN Systems — SDL and MSC* **28** (1996).

59. Sakarovitch, J.: The “last” decision problem for rational trace languages. *Proceedings of the 1st Latin American Symposium on Theoretical Informatics (LATIN’92)*, Lecture Notes in Computer Science **583**, Springer-Verlag (1992) 460–473
60. Thiagarajan, P. S., Walukiewicz, I.: An expressively complete linear time temporal logic for Mazurkiewicz traces. *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science (LICS’97)*, IEEE Computer Society Press (1997) 183–194
61. Thomas, W.: Automata on infinite objects. In van Leeuwen, J. (Ed.): *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*, Elsevier Science Publishers (1990) 133–191
62. Thomas, W.: Languages, automata, and logic. In Rozenberg, G., Salomaa, A. (Eds.): *Handbook of Formal Language Theory, Volume III*, Springer-Verlag (1997) 389–455
63. Vardi, M. Y., Wolper, P.: An automata-theoretic approach to automatic program verification. *Proceedings of the 1st Annual IEEE Symposium on Logic in Computer Science (LICS’86)*, IEEE Computer Society Press (1986) 332–345
64. Zielonka, W.: Notes on finite asynchronous automata. *R.A.I.R.O. Informatique Théorique et Applications* **21** (1987) 99–135

Index

- CMSCs, 16
 - communication graph, 19
 - concatenation, 17
- CMSGs, 18
 - globally cooperative, 19
 - locally synchronized, 19
- Compositional Message Sequence Graphs,
 - see* CMSGs
- Compositional MSCs, *see* CMSCs
- Mazurkiewicz Traces, 4
 - dependence alphabet, 4
 - independence relation, 4
- Message Sequence Charts, *see* MSCs
- Message Sequence Graphs, *see* MSGs
- Message-passing automata, *see* MPAs
- MPAs, 23
 - B -bounded, 24
 - deadlock, 28
- MSCs, 2
 - atom, 4
 - communication graph, 9
 - locally strongly connected, 9
 - concatenation, 3
 - existentially B -bounded, 3
 - implied scenario, 26
 - language, 7
 - finitely generated, 7
 - implied closure, 26
 - regular set of representatives, 15
 - safely realizable, 28
 - set of representatives, 15
 - weakly realizable, 25
 - linearization
 - B -bounded, 2
 - linearizations, 2
 - optimal, 3
 - Monadic second order logic, 19
 - MPAs
 - local acceptance, 25
 - MSO, 19
 - satisfiability, 22
 - netcharts, 16
 - regular MSC language, 7
 - universally B -bounded, 3
- MSGs, 8
 - bounded, 11
 - com-connected, 11
 - decidability, 13
 - globally cooperative, 12
 - locally synchronized, 11
 - model-checking, 14