

Muller Message-Passing Automata and Logics

Benedikt Bollig¹ and Dietrich Kuske²

¹ LSV, ENS Cachan, CNRS, France
bollig@lsv.ens-cachan.fr

² Institut für Informatik, Universität Leipzig, Germany
kuske@informatik.uni-leipzig.de

Abstract. We study nonterminating message-passing automata whose behavior is described by infinite message sequence charts. As a first result, we show that Muller, Büchi, and termination-detecting Muller acceptance are equivalent for these devices. To describe the expressive power of these automata, we give a logical characterization. More precisely, we show that they have the same expressive power as the existential fragment of a monadic second-order logic featuring a first-order quantifier to express that there are infinitely many elements satisfying some property. Our result is based on a new extension of the classical Ehrenfeucht-Fraïssé game to cope with infinite structures and the new first-order quantifier.

1 Introduction

The study of the relation between logical formalisms and operational automata devices has been a fascinating area of computer science and has produced some splendid results. From a logicians point of view, this relation allows us to decide logical theories effectively, from a system developer's point of view, the logical formalism might be considered as a specification language formalizing essential properties of a system, whereas the automaton appears as a model of the system itself.

The probably most famous connection between automata theory and classical logic has been established by Büchi and Elgot, who showed that finite automata and monadic second-order (MSO) logic are expressively equivalent [5, 8]. The sequential nature of finite automata, however, limits their use in the modeling of distributed systems which called for more general automata models that employ some communication mechanism between their components. This communication can be ensured by synchronous actions (e.g., asynchronous automata whose behavior is described by Mazurkiewicz traces) or by the exchange of messages along channels (e.g., message passing automata whose behavior is described by message sequence charts). For terminating behaviors, the expressive power of these models has been related to that of some sort of MSO logic [3, 7, 10, 12, 17].

One single execution of a distributed system is often modeled as a directed acyclic graph (V, E) with a set of events V and a binary relation E that describes the causal dependency between events. Any MSO property of words, Mazurkiewicz traces, or (existentially) bounded message sequence charts can be equivalently expressed by the appropriate automata model (and vice versa). It should be noted that the transitive closure

of the causal dependency E can be described in MSO and forms the temporal precedence relation. It can therefore also be used in the above cases. Since message-passing automata can in general not be complemented, MSO is too powerful in the context of unbounded message sequence charts [3]; but the restriction of MSO to its existential fragment (EMSO) is equivalent to message-passing automata without any channel bounds [3].

When modeling reactive systems, one is rather interested in infinite behaviors. Indeed, Büchi showed that MSO logic over infinite words is still as expressive as finite automata that require at least one final state to be visited infinitely often. Such an acceptance condition comes in many flavors, and variations thereof give rise to Büchi, Muller, Rabin, and Streett automata, which, in the nondeterministic case, are all equivalent [18]. The same applies to the settings of asynchronous (cellular) automata over infinite Mazurkiewicz traces [9, 16]. Kuske proposed message-passing automata with a Muller acceptance condition to make it capable of accepting infinite MSCs. As it turns out, the resulting automata model is equivalent to MSO logic over MSCs, provided the channel capacity is bounded [14].

It is the aim of this paper to lift the boundedness condition in this result, i.e., to characterize nonterminating behaviors of Muller message-passing automata with unbounded channels. After introducing the necessary notions, Section 2.2 shows that *Muller-, Büchi-, and even termination-detecting Muller-MPAs all have the same expressive power.* The proofs of these equivalence results use direct automata constructions. Contrary to the setting of terminating behaviors, EMSO is weaker than Muller message-passing automata: the set of infinite MSCs that send infinitely many messages from the first to the second component cannot be described by any EMSO formula. To overcome this deficiency, we introduce the additional first-order quantifier $\exists^\infty x\varphi(x)$ requesting infinitely many events x to satisfy some property $\varphi(x)$.

Our main result states that the extension of existential monadic second-order logic by an infinity quantifier is expressively equivalent to message-passing automata with nonterminating behaviors. Our proof follows the route of [3] that dealt with finite message sequence charts and EMSO and could therefore build on powerful results on this logic and its first-order fragment. Since these results are not available for our extension of EMSO, we first have to develop an analogous theory. This is the theme of Section 3 where we develop some model theory of the extension of first-order logic by the infinity quantifier; namely, we present an appropriate variant of Ehrenfeucht-Fraïssé games and threshold equivalence which leads to a Hanf-type theorem. As a result, any first-order sentence with infinity quantifier can be translated into some conditions on the number of realizations of spheres. Building on work by Bollig and Leucker [3], Section 4 shows that these conditions can be checked by message-passing automata equipped with a (termination-detecting) Muller condition. It also characterizes the expressive power of existential monadic second-order logic without the infinity quantifier by message-passing automata and the termination-detecting Staiger-Wagner acceptance condition.

A full version of this paper, which actually works on the more general model of asynchronous cellular automata with types, is available [2].

2 Message-Passing Automata with Nonterminating Behavior

We consider communicating systems where several sequential agents exchange messages through channels, executing send and receive actions. A send action is of the form $i!j$ indicating that agent i sends a message to agent j . The complementary receive action is denoted $j?i$. Here, agent j can read a message provided it has been sent through the corresponding channel from i to j . So let us, throughout the paper, fix a finite set Ag of agents. For an agent i , we denote by Σ_i the set $\{i!j, i?j \mid j \in Ag \setminus \{i\}\}$ of actions that are available to i . The union $\bigcup_{i \in Ag} \Sigma_i$ of all the actions is denoted by Σ .

2.1 Message-Passing Automata and Their Behavior

Let us make precise our model of a reactive system with a message-passing mechanism, which goes back to Brand and Zafiropulo [4] and was later extended by Kuske [14] to deal with infinite scenarios (see also [2]).

Definition 2.1. A message-passing automaton over Ag (or, for short, MPA) is a structure $\mathcal{A} = ((\mathcal{A}_i)_{i \in Ag}, \mathcal{D}, \iota)$ where

- \mathcal{D} is a nonempty finite set of synchronization data,
- for each $i \in Ag$, \mathcal{A}_i is a pair (Q_i, Δ_i) where
 - Q_i is a finite set of local states and
 - $\Delta_i \subseteq Q_i \times \Sigma_i \times \mathcal{D} \times Q_i$ is the set of local transitions, and
- $\iota \in \prod_{i \in Ag} Q_i$ is the global initial state.

The operational behavior of an MPA proceeds as one might expect. An agent i can execute send and receive actions according to its specification in terms of \mathcal{A}_i . Executing $i!j$ has the effect of writing a message into the fifo channel (i, j) (from i to j). Actually, this message is supplemented by some synchronization data from \mathcal{D} to extend the expressive power of MPAs. The benefit of synchronization data will become clear when we define the behavior of MPAs formally. Accordingly, $j?i$, which is executed by agent j , receives the message from i that is located at the top of the channel (i, j) . Thus, any two local machines \mathcal{A}_i and \mathcal{A}_j with $i \neq j$ are connected by two channels, the first for sending messages from i to j and the second for the reverse direction.

To describe the behavior of an MPA formally, we use the standardized formalism of *message sequence charts* (MSCs, [13]). There, the sequential behavior of an agent i is described by a vertical time-line, which will be modeled as a sequence of edges in a graph whose nodes are labeled with actions from Σ_i and referred to as *events*. Moreover, a send node and the corresponding receive node are joined by a (horizontal) message arrow. The edge relation of an MSC gives rise to a partial order relation constraining the execution order of the nodes. Moreover, edges are labeled with elements from $C = Ag \cup \{\text{msg}\}$, which provide some control information to identify message arrows and process arrows.

Definition 2.2. A message sequence chart over Ag (MSC, for short) is a graph $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$ where

- V is the set of events,

- $\lambda : V \rightarrow \Sigma$ is the labeling function,
- E^* (with $E = \bigcup_{\ell \in C} E_\ell$) is a partial order on V (we write $u \leq v$ for $(u, v) \in E^*$),
- for any $v \in V$, $\{u \in V \mid u \leq v\}$ is a finite set,
- for any $i \in Ag$, E_i is the cover relation³ of some total order on $V_i = \lambda^{-1}(\Sigma_i)$,
- for any $(u, v) \in E_{\text{msg}}$, there exist $i, j \in Ag$ distinct such that $\lambda(u) = i!j$ and $\lambda(v) = j?i$,
- for any $u \in V$, there is $v \in V$ such that $(u, v) \in E_{\text{msg}}$ or $(v, u) \in E_{\text{msg}}$, and
- for any $(u, v), (u', v') \in E_{\text{msg}}$ with $\lambda(u) = \lambda(u')$, we have $u \leq u'$ iff $v \leq v'$.

The last condition in the definition above expresses that messages are received in the same order in which they have been sent. Hence it reflects that we deal with fifo-channels only.

Let $\mathcal{A} = ((\mathcal{A}_i)_{i \in Ag}, \mathcal{D}, \iota)$ with $\mathcal{A}_i = (Q_i, \Delta_i)$ be an MPA and, moreover, let $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$ be an MSC. For a mapping $\rho : V \rightarrow \bigcup_{i \in Ag} Q_i$ (which is a candidate for a run of \mathcal{A} on M), we define the mapping $\rho^- : V \rightarrow \bigcup_{i \in Ag} Q_i$ as follows. Let $i \in Ag$ and $v \in V_i$. If we can find $u \in V_i$ such that $(u, v) \in E_i$, then we set $\rho^-(v) = \rho(u)$. If there is no such u , we let $\rho^-(v) = \iota[i]$. A run of \mathcal{A} on M is a pair (ρ, μ) of mappings $\rho : V \rightarrow \bigcup_{i \in Ag} Q_i$ and $\mu : V \rightarrow \mathcal{D}$ such that

- for any $(u, v) \in E_{\text{msg}}$, $\mu(u) = \mu(v)$ and
- for any $v \in V_i$, $(\rho^-(v), \lambda(v), \mu(v), \rho(v)) \in \Delta_i$.

2.2 Muller, Büchi, and Staiger-Wagner Message-Passing Automata

We will now extend our automata model by some acceptance modes that originate from the work on automata on infinite words. All except the Staiger-Wagner acceptance depend on those states that appear infinitely often in a run. So let us first give the following definitions. Let $\mathcal{A} = (((Q_i, \Delta_i))_{i \in Ag}, \mathcal{D}, \iota)$ be an MPA (we set $Q = \bigcup_{i \in Ag} Q_i$) and let $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$ be an MSC. For a mapping $\rho : V \rightarrow Q$, we define functions $\text{Inf}_\rho : Ag \rightarrow 2^Q$ and $\text{Inf}_\rho^+, \text{Occ}_\rho^+ : Ag \rightarrow 2^Q \times \{\text{inf}, \overline{\infty}\}$ as follows (with $i \in Ag$):

$$\text{Inf}_\rho[i] = \begin{cases} \{q \mid \forall u \in V_i \exists v \in V_i : u \leq v \text{ and } q = \rho(v)\} & \text{if } V_i \neq \emptyset \\ \{\iota[i]\} & \text{otherwise} \end{cases}$$

$$\text{Inf}_\rho^+[i] = \begin{cases} (\text{Inf}_\rho[i], \overline{\infty}) & \text{if } V_i \text{ is finite} \\ (\text{Inf}_\rho[i], \text{inf}) & \text{otherwise} \end{cases} \quad \text{Occ}_\rho^+[i] = \begin{cases} (\rho^{-1}(V_i), \overline{\infty}) & \text{if } V_i \text{ is finite} \\ (\rho^{-1}(V_i), \text{inf}) & \text{otherwise} \end{cases}$$

If V_i is finite, then $\text{Inf}_\rho[i]$ describes the state assumed at the event that is maximal in V_i (which is the local state $\iota[i]$ if V_i is even empty). If V_i is infinite, then $\text{Inf}_\rho[i]$ is the set of states assumed infinitely often. If $\text{Inf}_\rho[i]$ is a singleton, we do not know whether V_i is finite or not – this additional information is present in $\text{Inf}_\rho^+[i]$. Similarly, the definition of Occ_ρ^+ will capture the (termination-detecting) Staiger-Wagner acceptance condition.

Definition 2.3. A Büchi MPA or Muller MPA is a structure $\mathcal{A} = ((\mathcal{A}_i)_{i \in Ag}, \mathcal{D}, \iota, \mathcal{F})$ with $\mathcal{A}_i = (Q_i, \Delta_i)$ such that $((\mathcal{A}_i)_{i \in Ag}, \mathcal{D}, \iota)$ is an MPA and $\mathcal{F} \subseteq \prod_{i \in Ag} 2^{Q_i}$.

Now let (ρ, μ) be some run of \mathcal{A} on the MSC $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$.

³ The cover relation of a total or partial order \preceq on V_i is its direct successor relation $\prec \setminus \prec^2$.

- (1) If \mathcal{A} is a Büchi MPA, then the run (ρ, μ) is accepting if there is $\bar{q} \in \mathcal{F}$ such that $\bar{q}[i] \cap \text{Inf}_\rho[i] \neq \emptyset$ for all $i \in \text{Ag}$.
- (2) If \mathcal{A} is a Muller MPA, then the run (ρ, μ) is accepting if $\text{Inf}_\rho \in \mathcal{F}$.

Definition 2.4. A termination-detecting Staiger-Wagner MPA or termination-detecting Muller MPA is a structure $\mathcal{A} = ((\mathcal{A}_i)_{i \in \text{Ag}}, \mathcal{D}, \iota, \mathcal{F})$ with $\mathcal{A}_i = (Q_i, \Delta_i)$ such that $((\mathcal{A}_i)_{i \in \text{Ag}}, \mathcal{D}, \iota)$ is an MPA and $\mathcal{F} \subseteq \prod_{i \in \text{Ag}} (2^{Q_i} \times \{\text{inf}, \overline{\text{inf}}\})$.

Let (ρ, μ) be some run of \mathcal{A} on the MSC $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$.

- (1) If \mathcal{A} is a termination-detecting Staiger-Wagner MPA, then (ρ, μ) is accepting if $\text{Occ}_\rho^+ \in \mathcal{F}$.
- (2) If \mathcal{A} is a termination-detecting Muller MPA, then (ρ, μ) is accepting if $\text{Inf}_\rho^+ \in \mathcal{F}$.

If \mathcal{A} is some of these MPA, then the language $L(\mathcal{A})$ accepted by \mathcal{A} is the set of those MSCs that admit an accepting run of \mathcal{A} .

The generalized model of termination-detecting Muller MPAs will turn out to be tremendously helpful when, in Section 4, we study the relationship between logic and MPAs. Let us first prove that termination-detecting Muller MPAs are not more expressive than Muller or Büchi MPAs (whereas termination-detecting Staiger-Wagner MPAs are strictly weaker).

2.3 Muller and Büchi MPAs vs. Termination-Detecting Muller MPAs

In a termination-detecting Muller MPA, the acceptance condition can distinguish between the infinite repetition of a local state and the appearance of this state as the final one. This distinction is not directly possible in a Muller MPA. To solve this problem, we first state that a Muller MPA can determine whether a particular agent performs finitely or infinitely many actions. This is done by adding a flag to each local process that alternates between 0 and 1 (see [2] for details).

Lemma 2.5. Let $k \in \text{Ag}$. There are Muller MPAs \mathcal{A} and \mathcal{B} such that, for any MSC $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$, $M \in L(\mathcal{A})$ iff V_k is infinite and $M \in L(\mathcal{B})$ iff V_k is finite.

Theorem 2.6. Let L be a set of MSCs. Then, the following are equivalent:

1. there exists a termination-detecting Muller MPA \mathcal{A} such that $L = L(\mathcal{A})$.
2. there exists a Muller MPA \mathcal{A}' such that $L = L(\mathcal{A}')$.

Proof. Suppose $\mathcal{A} = ((\mathcal{A}_i)_{i \in \text{Ag}}, \mathcal{D}, \iota, \mathcal{F})$ with $\mathcal{A}_i = (Q_i, \Delta_i)$ is a Muller MPA and let π_1 denote the projection of $2^{\bigcup_{i \in \text{Ag}} Q_i} \times \{\text{inf}, \overline{\text{inf}}\}$ onto the first component. Then, let \mathcal{F}' comprise all tuples $\bar{q} \in \prod_{i \in \text{Ag}} 2^{Q_i} \times \{\text{inf}, \overline{\text{inf}}\}$ with $\pi_1 \circ \bar{q} \in \mathcal{F}$. This defines a termination-detecting Muller MPA $\mathcal{A}' = ((\mathcal{A}_i)_{i \in \text{Ag}}, \mathcal{D}, \iota, \mathcal{F}')$ that certainly accepts the same language as \mathcal{A} does. For the other implication, let $\mathcal{A}' = ((\mathcal{A}_i)_{i \in \text{Ag}}, \mathcal{D}, \iota, \mathcal{F}')$ be some termination-detecting Muller MPA. For $\bar{q} \in \mathcal{F}'$, consider the Muller MPA $\mathcal{A}_{\bar{q}} = ((\mathcal{A}_i)_{i \in \text{Ag}}, \mathcal{D}, \iota, \{\pi_1 \circ \bar{q}\})$. Then, the language of the termination-detecting Muller MPA $\mathcal{A}' = ((\mathcal{A}_i)_{i \in \text{Ag}}, \mathcal{D}, \iota, \{\bar{q}\})$ is an intersection of $L(\mathcal{A}_{\bar{q}})$ with some sets of the form $\{M \mid V_k \text{ is infinite}\}$ and $\{M \mid V_k \text{ is finite}\}$. Since any of these sets can be accepted by a Muller MPA and since Muller MPAs are closed under union and intersection [2], the result follows immediately. \square

We now show that Büchi MPAs are as expressive as termination-detecting Muller MPAs, too. Surprisingly, this is not the case in the word setting when considering both finite and infinite words. In our distributed setting, however, the distinction between the infinite repetition of a local state and the appearance of this state as the final one is possible.

Lemma 2.7. *Let $k \in Ag$. There exist Büchi MPAs \mathcal{A} and \mathcal{B} such that, for any MSC $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$, we have $M \in L(\mathcal{A})$ iff V_k is infinite and $M \in L(\mathcal{B})$ iff V_k is finite.*

Proof. The construction of \mathcal{B} is straightforward. Instead of \mathcal{A} , we restrict to building, for $\sigma \in \Sigma_k$, a Büchi MPA \mathcal{A}_σ that accepts those MSCs in which σ is executed infinitely often. Let σ' be the communication action complementing σ , which is executed by some k' (e.g., if σ is of the form $k!k'$, then $\sigma' = k'?k$). The idea is that k and k' work together to detect that, in fact, σ and σ' occur infinitely often. Both agents toggle between states 0 and 1 when executing σ and σ' , respectively. However, in the acceptance condition, k requires 0 to be taken infinitely often, whereas k' claims to visit 1 infinitely often. Formally, we set $\mathcal{A}_\sigma = (((Q_i, \Delta_i))_{i \in Ag}, \mathcal{D}, \iota, \mathcal{F})$ with $\mathcal{D} \neq \emptyset$ arbitrary, $Q_i = \{0, 1\}$ for any $i \in Ag$, and Δ_i contains any tuple $(q, \tau, m, q') \in Q_i \times \Sigma_i \times \mathcal{D} \times Q_i$ such that $\tau \in \{\sigma, \sigma'\}$ iff $q' = 1 - q$. Moreover, $\iota = (0)_{i \in Ag}$, $\mathcal{F}[i] = \{0\}$ for any $i \neq k'$, and $\mathcal{F}[k'] = \{1\}$. \square

We conclude that Büchi and termination-detecting Muller MPAs are equally expressive.

Theorem 2.8. *Let L be a set of MSCs. Then the following are equivalent:*

1. *there exists a termination-detecting Muller MPA \mathcal{A} such that $L = L(\mathcal{A})$.*
2. *there exists a Büchi MPA \mathcal{A}' such that $L = L(\mathcal{A}')$.*

3 Structures, Logic, and the Ehrenfeucht-Fraïssé Game

Towards a logical characterization of our automata models, we first study more general structures than MSCs. Obviously, MSCs can be seen as relational structures whose signature contains binary relations E_ℓ for $\ell \in C$ and unary relations R_σ for $\sigma \in \Sigma$. In this section, we actually consider more general structures over finite and function-free signatures σ . In the following, \mathfrak{A} and \mathfrak{B} will denote σ -structures, whereas A refers to the universe of \mathfrak{A} and B to that of \mathfrak{B} .

3.1 Structures and Monadic Second-Order Logic

The *Gaifman graph* $G(\mathfrak{A})$ of a structure \mathfrak{A} is a graph (A, E) with universe A (i.e., the universe of the structure \mathfrak{A}). Two elements $a, b \in A$ are connected by an edge (i.e., $(a, b) \in E$) if they belong to some tuple in some relation, i.e., if there is a relation symbol $P \in \sigma$ and a tuple $(a_1, \dots, a_n) \in P$ such that $a, b \in \{a_1, a_2, \dots, a_n\}$. We will speak of the degree of a in \mathfrak{A} whenever we actually mean the degree of a in the Gaifman graph of \mathfrak{A} . If all elements of \mathfrak{A} have degree at most l , then we say that \mathfrak{A} has degree at most l . Now let $a, b \in A$. Then the *distance* $d_{\mathfrak{A}}(a, b)$ (or $d(a, b)$ if \mathfrak{A} is understood) denotes the minimal length of a path connecting a and b in the Gaifman graph $G(\mathfrak{A})$. For

$\bar{a} = (a_1, \dots, a_n) \in A^n$ and $b \in A$, we write $d(\bar{a}, b) = \min\{d(a_1, b), \dots, d(a_n, b)\}$. Let $r \in \mathbb{N}$ and \bar{c} denote the sequence of constants in the structure \mathfrak{A} . The r -sphere $r\text{-Sph}(\mathfrak{A})$ of \mathfrak{A} is the substructure of \mathfrak{A} generated by the universe $\{b \in A \mid d_{\mathfrak{A}}(\bar{c}, b) \leq r\}$ (note that, if \mathfrak{A} does not contain any constants, this set is empty and the sphere is the empty structure). For a tuple \bar{a} of elements in \mathfrak{A} , the r -sphere of \mathfrak{A} around \bar{a} is the r -sphere of the extension (\mathfrak{A}, \bar{a}) of \mathfrak{A} by constants \bar{a} .

We fix supplies $\text{Var} = \{x, y, x_1, x_2, \dots\}$ of *individual* and $\text{VAR} = \{X, Y, \dots\}$ of *set variables*. The set $\text{MSO}^\infty(\sigma)$ of *extended monadic second-order* (or MSO^∞) formulas over σ is given by the following grammar:

$$\varphi ::= P(x_1, \dots, x_n) \mid x_1 = x_2 \mid x_1 \in X \mid \neg\varphi \mid \varphi \vee \psi \mid \exists x\varphi \mid \exists X\varphi \mid \exists^\infty x\varphi$$

where $n \in \mathbb{N}$, $P \in \sigma$ is an n -ary predicate symbol, x_i is a variable from Var or a constant symbol, $x \in \text{Var}$, and $X \in \text{VAR}$.

Let \mathfrak{A} be a σ -structure, $\varphi(x_1, \dots, x_m, X_1, \dots, X_n) \in \text{MSO}^\infty$ be a formula, and $\bar{a} = (a_1, \dots, a_m) \in A^m$ and $\bar{A} = (A_1, \dots, A_n) \in (2^A)^n$ be tuples of elements and subsets of A . Then the *satisfaction* relation $\mathfrak{A} \models \varphi(\bar{a}, \bar{A})$ is defined as usual such that $\mathfrak{A} \models (\exists^\infty y\psi)(\bar{a}, \bar{A})$ iff $\mathfrak{A} \models \psi(a, \bar{a}, \bar{A})$ for infinitely many $a \in A$ (for $\psi(y, x_1, \dots, x_m, X_1, \dots, X_n) \in \text{MSO}^\infty$).

We define the following fragments of $\text{MSO}^\infty(\sigma)$:

1. the *first-order fragment* $\text{FO}^\infty(\sigma)$ comprises those MSO^∞ formulas that do not contain any set quantifier
2. the *existential fragment* $\text{EMSO}^\infty(\sigma)$ comprises the MSO^∞ formulas of the form $\exists X_1 \dots \exists X_n \varphi$ with $\varphi \in \text{FO}^\infty(\sigma)$
3. the *monadic second-order fragment* $\text{MSO}(\sigma)$ comprises those MSO^∞ formulas that do not contain the quantifier \exists^∞
4. *first-order logic* $\text{FO}(\sigma)$ equals $\text{MSO}(\sigma) \cap \text{FO}^\infty(\sigma)$
5. *existential monadic second-order logic* $\text{EMSO}(\sigma)$ equals $\text{MSO}(\sigma) \cap \text{EMSO}^\infty(\sigma)$

In the following, we may omit the reference to σ and write, for example, FO instead of $\text{FO}(\sigma)$.

The *quantifier-rank* $\text{qr}(\varphi)$ of a formula φ in FO^∞ is the nesting depth of quantifiers in φ . More precisely, $\text{qr}(\varphi) = 0$ if φ is atomic, $\text{qr}(\neg\varphi) = \text{qr}(\varphi)$, $\text{qr}(\varphi \vee \psi) = \max\{\text{qr}(\varphi), \text{qr}(\psi)\}$, and $\text{qr}(\exists x\varphi) = \text{qr}(\exists^\infty x\varphi) = \text{qr}(\varphi) + 1$. For $k \in \mathbb{N}$, we denote by $\text{FO}^\infty[k]$ the set of first-order formulas of quantifier rank at most k . Any formula $\varphi \in \text{FO}^\infty[k+1]$ is logically equivalent to a Boolean combination of formulas of the form $\exists^\infty x\psi$ and $\exists x\psi$ with $\psi \in \text{FO}^\infty[k]$. Hence, for a fixed tuple of free variables, there are up to logical equivalence only finitely many formulas in $\text{FO}^\infty[k]$.

Let $k, m \in \mathbb{N}$, \mathfrak{A} be a σ -structure, and \bar{a} an m -ary vector of elements of \mathfrak{A} . The *rank- k m -type* of \bar{a} in \mathfrak{A} comprises those $\text{FO}^\infty[k]$ formulas that hold true for \bar{a} : $\text{type}_k(\mathfrak{A}, \bar{a}) := \{\varphi \in \text{FO}^\infty[k] \mid \mathfrak{A} \models \varphi(\bar{a})\}$. Hence the number of different rank- k m -types is finite. Moreover, for any rank- k m -type T , there is a formula $\alpha_T(\bar{x}) \in \text{FO}^\infty[k]$ (with \bar{x} an m -tuple) such that, for every σ -structure \mathfrak{A} and $\bar{a} \in A^m$, $\mathfrak{A} \models \alpha_T(\bar{a})$ iff $\text{type}_k(\mathfrak{A}, \bar{a}) = T$.

3.2 Ehrenfeucht-Fraïssé Games

The classical Ehrenfeucht-Fraïssé game is used to characterize the expressive power of the logic $\text{FO}[k]$. Here, we propose an extension to similarly capture $\text{FO}^\infty[k]$.

The FO^∞ -game is played between two players, the *spoiler* and the *duplicator*. A *game position* is a triple $(\mathfrak{A}, \mathfrak{B}, k)$ where \mathfrak{A} and \mathfrak{B} are structures over the same function-free signature σ and $k \in \mathbb{N}$. This position is *winning (for duplicator)* if $k = 0$ and the binary relation $\{(c^{\mathfrak{A}}, c^{\mathfrak{B}}) \mid c \text{ constant symbol from } \sigma\}$ is a partial isomorphism from \mathfrak{A} to \mathfrak{B} . If $k > 0$, the game proceeds as follows:

- (1) Spoiler chooses to proceed with (2) or (2').
- (2) Spoiler chooses $a \in A$ or $b \in B$.
- (3) Duplicator chooses an element in the other structure (i.e., $b \in B$ or $a \in A$).
- (4) The game proceeds with $((\mathfrak{A}, a), (\mathfrak{B}, b), k - 1)$.
- (2') Spoiler chooses an infinite subset Z of A or of B .
- (3') Duplicator chooses an infinite subset of the set Z and an infinite subset of the other structure (i.e., after this step, we have infinite subsets A' and B' of A and B , resp.).
- (4') Spoiler chooses elements $a \in A'$ and $b \in B'$.
- (5') The game proceeds with $((\mathfrak{A}, a), (\mathfrak{B}, b), k - 1)$.

For σ -structures \mathfrak{A} and \mathfrak{B} and $k \in \mathbb{N}$, we write $\mathfrak{A} \equiv_k^\infty \mathfrak{B}$ if duplicator can force the FO^∞ -play started in $(\mathfrak{A}, \mathfrak{B}, k)$ into a winning position. The classical Ehrenfeucht-Fraïssé game is obtained by not allowing the spoiler to go to (2'). If duplicator can force this classical Ehrenfeucht-Fraïssé game from $(\mathfrak{A}, \mathfrak{B}, k)$ into a winning position, we write $\mathfrak{A} \equiv_k \mathfrak{B}$.

Theorem 3.1. *Let \mathfrak{A} and \mathfrak{B} be σ -structures and let $k \in \mathbb{N}$.*

- (1) *(Ehrenfeucht-Fraïssé) \mathfrak{A} and \mathfrak{B} agree on $\text{FO}[k]$ iff $\mathfrak{A} \equiv_k \mathfrak{B}$.*
- (2) *\mathfrak{A} and \mathfrak{B} agree on $\text{FO}^\infty[k]$ iff $\mathfrak{A} \equiv_k^\infty \mathfrak{B}$.*

Proof. We adapt the proof of the first statement from [15, 18] to also show the second. The equivalence is shown by induction on k , the case $k = 0$ is obvious.

Now let \mathfrak{A} and \mathfrak{B} be σ -structures that agree on $\text{FO}^\infty[k + 1]$. We consider the case that, in step (2'), spoiler chooses an infinite set $Z \subseteq A$. Since there are only finitely many rank- k 1-types, there exist a rank- k 1-type T and an infinite set $A' \subseteq Z$ with $T = \text{type}_k(\mathfrak{A}, a)$ for all $a \in A'$. Then, $\mathfrak{A} \models \exists^\infty x \alpha_T(x)$, which implies $\mathfrak{B} \models \exists^\infty x \alpha_T(x)$ since $\exists^\infty x \alpha_T$ has quantifier rank at most $k + 1$. Thus the set $B' = \{b \in B \mid T = \text{type}_k(\mathfrak{B}, b)\}$ is infinite. Whatever $(a, b) \in A' \times B'$ spoiler chooses in step (4'), we have by the induction hypothesis $(\mathfrak{A}, a) \equiv_k^\infty (\mathfrak{B}, b)$. If spoiler chooses to go to (2), similar arguments apply. Hence duplicator can force the game into a position $((\mathfrak{A}, a), (\mathfrak{B}, b), k)$ from where he has a winning strategy.

Now suppose $\mathfrak{A} \equiv_{k+1}^\infty \mathfrak{B}$. Since any formula of quantifier rank $k + 1$ is a Boolean combination of formulas of the form $\exists x \varphi$ or $\exists^\infty x \varphi$ with $\varphi \in \text{FO}^\infty[k]$, it suffices to consider these two cases. Suppose $\mathfrak{A} \models \exists^\infty x \varphi$. Spoiler can in step (2') choose an infinite set $Z = \{a \in A \mid \mathfrak{A} \models \varphi(a)\}$. Since duplicator can force the play into a winning position, there are infinite sets $A' \subseteq Z$ and $B' \subseteq B$ such that, for any $(a, b) \in A' \times B'$, we have $(\mathfrak{A}, a) \equiv_k^\infty (\mathfrak{B}, b)$. Hence, by the induction hypothesis, (\mathfrak{A}, a) and (\mathfrak{B}, b) agree on $\text{FO}^\infty[k]$. In particular, $\mathfrak{B} \models \varphi(b)$ for any $b \in B'$. Thus, $\mathfrak{B} \models \exists^\infty x \varphi$. The argument for $\exists x \varphi$ is similar. \square

3.3 Threshold Equivalence

In the context of structures of bounded degree, *threshold equivalence* provides a refinement of \equiv_k and, finally, a normal form of FO formulas that restricts to counting of spheres up to a certain threshold [15, 18].

For a structure \mathfrak{A} and an isomorphism type τ of a sphere, let $|\mathfrak{A}|_\tau$ denote the number of occurrences of τ in \mathfrak{A} .

Definition 3.2. Let $r, t \in \mathbb{N}$ and \mathfrak{A} and \mathfrak{B} be σ -structures. We write $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r,t} \mathfrak{B}$ if, for any isomorphism type τ of an r -sphere around a single element, $|\mathfrak{A}|_\tau = |\mathfrak{B}|_\tau$ or both $t < |\mathfrak{A}|_\tau$ and $t < |\mathfrak{B}|_\tau$.

Similarly, we write $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r,t} \mathfrak{B}$ if, for any isomorphism type τ of an r -sphere of a single element, $|\mathfrak{A}|_\tau = |\mathfrak{B}|_\tau$ or both $t < |\mathfrak{A}|_\tau < \infty$ and $t < |\mathfrak{B}|_\tau < \infty$.

In other words, $\stackrel{\infty}{\leftrightarrow}_{r,t}$ and $\stackrel{\infty}{\leftrightarrow}_{r,t}$ distinguish structures on the basis of the number of realizations of r -spheres up to some threshold t . But the former does not distinguish between “many” and “infinitely many” realizations of a sphere. The latter identifies all natural numbers $t + 1, t + 2, \dots$, but makes a difference between them and infinity.

Theorem 3.3. For any $k, l \geq 0$, there are $r, t \geq 0$ such that, for any σ -structures \mathfrak{A} and \mathfrak{B} of degree at most l , $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r,t} \mathfrak{B}$ implies $\mathfrak{A} \equiv_k \mathfrak{B}$.

Replacing $\stackrel{\infty}{\leftrightarrow}_{r,t}$ by $\stackrel{\infty}{\leftrightarrow}_{r,t}$ and \equiv_k by \equiv_k in this theorem, one obtains Hanf’s theorem [11]. It states that \mathfrak{A} and \mathfrak{B} agree on $\text{FO}[k]$ whenever $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r,t} \mathfrak{B}$. Similarly, \mathfrak{A} and \mathfrak{B} agree on $\text{FO}^\infty[k]$ whenever $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r,t} \mathfrak{B}$.

Proof. Let $r_0 = 0$ and, for $i \in \mathbb{N}$, set $r_{i+1} = 3r_i + 1$. Moreover, we set t_i to be $i \cdot \text{size}$ where size is the maximal size of an r_i -sphere whose elements have degree at most l .

One first shows that duplicator can force the FO^∞ -play from $(\mathfrak{A}, \mathfrak{B}, k + 1)$ with $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r_{k+1}, t_{k+1}} \mathfrak{B}$ into some position $((\mathfrak{A}, a), (\mathfrak{B}, b), k)$ with $(\mathfrak{A}, a) \stackrel{\infty}{\leftrightarrow}_{r_k, t_k} (\mathfrak{B}, b)$.

To this aim, first assume that spoiler chooses in (2’) an infinite set $Z \subseteq A$. Since there are only finitely many isomorphism types of spheres around single elements in \mathfrak{A} , there is an infinite set $A' \subseteq Z$ such that, for any $a, a' \in A'$, we have $r_{k+1}\text{-Sph}(\mathfrak{A}, a) \cong r_{k+1}\text{-Sph}(\mathfrak{A}, a') =: S$. Since $r_{k+1}\text{-Sph}(\mathfrak{A})$ is finite, we can even assume that $d(a, c) > 2r_{k+1} + 1$ for any constant c from \mathfrak{A} and any $a \in A'$. From $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r_{k+1}, t_{k+1}} \mathfrak{B}$, we obtain the existence of infinitely many $b \in B$ with $r_{k+1}\text{-Sph}(\mathfrak{B}, b) \cong S$. Let B' denote this infinite set. Then duplicator chooses these two sets A' and B' in step (3’).

To show $(\mathfrak{A}, a) \stackrel{\infty}{\leftrightarrow}_{r_k, t_k} (\mathfrak{B}, b)$ for any $(a, b) \in A' \times B'$, one uses arguments as in [15, Sect. 4.2].

Thus, duplicator can force any play from $(\mathfrak{A}, \mathfrak{B}, k)$ with $\mathfrak{A} \stackrel{\infty}{\leftrightarrow}_{r_k, t_k} \mathfrak{B}$ into a position $(\mathfrak{A}', \mathfrak{B}', 0)$ with $\mathfrak{A}' \stackrel{\infty}{\leftrightarrow}_{0,0} \mathfrak{B}'$. Let a_1, \dots, a_n be the constants from \mathfrak{A}' and similarly for \mathfrak{B}' . Then the 0-sphere around a_1 in \mathfrak{A}' has $n + 1$ constants where the first and last coincide. Since $\mathfrak{A}' \stackrel{\infty}{\leftrightarrow}_{0,0} \mathfrak{B}'$, it is also realized in \mathfrak{B}' which is only possible by b_1 . This implies that $0\text{-Sph}(\mathfrak{A}')$ and $0\text{-Sph}(\mathfrak{B}')$ (i.e., the restrictions of \mathfrak{A}' and \mathfrak{B}' to their constants) are isomorphic. Since quantifier-free formulas cannot distinguish structures with isomorphic restrictions to constants, we obtain $\mathfrak{A}' \equiv_0 \mathfrak{B}'$. Hence the position $(\mathfrak{A}', \mathfrak{B}', 0)$ is winning for duplicator. \square

4 Message-Passing Automata and Logics

MPAs can be used to compute the sphere around any node of an MSC. This feature, described formally in the following proposition, is the key connection between these automata and the logical characterization of first-order expressible properties.

Proposition 4.1 (cf. [3]). *Let r be a natural number. There are a termination-detecting Muller / termination-detecting Staiger-Wagner MPA $\mathcal{A}_r = (((Q_i, \Delta_i))_{i \in Ag}, \mathcal{D}, \iota, \mathcal{F})$ and a mapping η from $\bigcup_{i \in Ag} Q_i$ into the set of spheres of radius r around a single vertex such that $L(\mathcal{A}_r)$ is the set of all MSCs and, for any MSC $M = (V, \{E_\ell\}_{\ell \in C}, \lambda)$, for any run (ρ, μ) of \mathcal{A}_r on M , and for any $u \in V$, we have $\eta(\rho(u)) = r\text{-Sph}(M, u)$.*

A first attempt to construct such an MPA would be to make each event guess its sphere and then to show that these guesses can be verified by relating the guessed spheres of neighboring events (i.e., the previous and next event on the process line and the communication partner). However, this straightforward approach does not work; the least counterexample we know of uses radius $r = 4$ and an MSC of seven agents and 74 events [1, pages 137–145]. Instead, the MPA \mathcal{A}_r from [3] guesses its own as well as spheres of nodes nearby and its own position in these additional spheres. Adding some global information allows us to locally check whether all the guesses are correct.

We now relate the expressive power of all types of MPAs and the extended logic. Recall that MSCs can be considered as relational structures of bounded degree, whose signature contains binary relations E_ℓ for $\ell \in C$ and unary relations R_σ for $\sigma \in \Sigma$. As expected, we will write the formula $R_\sigma(x)$ as $\lambda(x) = \sigma$. Moreover, we write $\text{EMSO}^\infty(Ag, C)$ for $\text{EMSO}^\infty((E_\ell)_{\ell \in C}, (R_\sigma)_{\sigma \in \Sigma})$, $\text{FO}^\infty(Ag, C)$ etc. are to be understood similarly.

4.1 Termination-Detecting Muller MPAs and Logic

Lemma 4.2. *Let $r \in \mathbb{N}$, $t \in \mathbb{N} \cup \{\infty\}$, $i \in Ag$, and S be some r -sphere in some MSC around a single vertex from V_i . There exist termination-detecting Muller MPAs recognizing the sets of MSCs M with $|M|_S = t$ and $t < |M|_S < \infty$, respectively.*

Proof. Exemplarily, we consider the cases $|M|_S = \infty$ and $t < |M|_S < \infty$ for $t \in \mathbb{N}$. One starts from the termination-detecting Muller MPA $\mathcal{A}_r = (((Q_i, \Delta_i))_{i \in Ag}, \mathcal{D}, \iota, \mathcal{F})$ and η from Prop. 4.1.

To detect $|M|_S = \infty$, we just keep those accepting tuples $(F_j, \theta_j)_{j \in Ag}$ from \mathcal{F} that satisfy $\theta_i = \text{inf}$ and $F_i \cap \eta^{-1}(S) \neq \emptyset$.

To detect $t < |M|_S < \infty$ for $t \in \mathbb{N}$, we extend the states of \mathcal{A}_i by a counter that counts the number of realizations of S up to $t + 1$, i.e., the new local state space of agent i is $Q_i \times \{0, \dots, t + 1\}$ with initial state $(\iota[i], 0)$. To distinguish “at least $t + 1$ ” from “infinitely many” realizations of S , the acceptance condition is the set of pairs $(F_j, \theta_j)_{j \in Ag}$ such that $\theta_i = \overline{\infty}$ implies $F_i \subseteq Q_i \times \{t + 1\}$, $\theta_i = \infty$ implies $F_i \subseteq (Q_i \setminus \eta^{-1}(S)) \times \{t + 1\}$, and $(\pi_1(F_j), \theta_j)_{j \in Ag} \in \mathcal{F}$. \square

Theorem 4.3. *Let L be a set of MSCs. Then, the following are equivalent:*

1. *there exists a termination-detecting Muller MPA \mathcal{A} such that $L = L(\mathcal{A})$.*
2. *there exists an EMSO $^\infty(Ag, C)$ sentence φ such that $L = \{M \mid M \models \varphi\}$.*

Proof. By Theorem 2.8, it is sufficient to translate a Büchi MPA into an equivalent EMSO $^\infty(Ag, C)$ sentence. This construction follows similar instances of that problem, e.g., [6]. Second order variables X_q for $q \in \mathcal{D} \cup \bigcup_{i \in Ag} Q_i$ encode an assignment of messages and states to vertices. The first-order part then expresses that this assignment is a run. In addition, we have to take care of the acceptance condition. Any such condition $\bar{q} \in \prod_{i \in Ag} 2^{Q_i}$ is translated into the conjunction of the following formulas for any $i \in Ag$

$$\bigvee_{q \in \bar{q}[i]} \left(\begin{array}{l} \exists^\infty x (x \in X_q \wedge \lambda(x) \in \Sigma_i) \\ \vee \exists x (x \in X_q \wedge \lambda(x) \in \Sigma_i \wedge \neg \exists y ((x, y) \in E_i)) \end{array} \right)$$

(supplemented by $\dots \vee \forall x \neg \lambda(x) \in \Sigma_i$ if $\iota[i] \in \bar{q}[i]$). The kernel of this formula expresses that the state q is assumed infinitely often by process i or, alternatively, it is assumed by the last event of this process.

Consider the other implication. Since termination-detecting Muller MPAs are closed under projection, it suffices to consider the case $\varphi \in \text{FO}^\infty$. By Theorem 3.3, L is a finite union of $\stackrel{\infty}{\simeq}_{r,t}$ -equivalence classes for some $r, t \in \mathbb{N}$. Any such equivalence class is an intersection of languages as in Lemma 4.2. Now the result follows from the fact that termination-detecting Muller MPAs are closed under union and intersection. \square

The number of states of the termination-detecting Muller MPA \mathcal{A} constructed from a given EMSO $^\infty(Ag, C)$ -formula φ is elementary in the size of the formula φ : The radius r from Theorem 3.3 is bounded by $3^{|\varphi|}$. Similarly, the number t is bounded by $|\varphi| \cdot r$ or infinite. We only remark that the number of states of the MPA from Prop. 4.1 is also elementary in r and t . But it is not clear whether \mathcal{A} can be constructed in elementary time. For this, one has to verify that one can compute easily the equivalence classes of $\stackrel{\infty}{\simeq}_{r,t}$ that contribute to the language of φ (in the proof above).

4.2 Staiger-Wagner MPAs and Logic

As far as finite counting is concerned, termination-detecting Staiger-Wagner MPAs can do as much as termination-detecting Muller MPAs. Similarly to Lemma 4.2, we can show the following:

Lemma 4.4. *Let $r, t \in \mathbb{N}$ and S be some r -sphere in some MSC around a single vertex. There exist termination-detecting Staiger-Wagner MPAs that recognize the sets of MSCs M with $|M|_S = t$ and $t < |M|_S$, respectively.*

Theorem 4.5. *Let L be a set of MSCs. Then, the following are equivalent:*

1. *there exists a termination-detecting Staiger-Wagner MPA \mathcal{A} such that $L = L(\mathcal{A})$.*
2. *there exists an EMSO(Ag, C) sentence φ such that $L = \{M \mid M \models \varphi\}$.*

Proof. The proof is similar to that of Theorem 4.3. The only difference in the transformation of an automaton into a formula concerns the acceptance condition. For the other transformation, we use Hanf's theorem [11] instead of Theorem 3.3 and Lemma 4.4 instead of Lemma 4.2. \square

5 Conclusion

We conclude summarizing the main results of this paper:

Theorem 5.1. *Let L be a set of MSCs. Then, the following are equivalent:*

1. *there exists a termination-detecting Muller MPA that accepts L .*
2. *there exists a Muller MPA that accepts L .*
3. *there exists a Büchi MPA that accepts L .*
4. *there exists a sentence $\varphi \in \text{EMSO}^\infty(\text{Ag}, C)$ such that $L = \{M \mid M \models \varphi\}$.*

References

1. B. Bollig. *Automata and Logics for Message Sequence Charts*. PhD thesis, Department of Computer Science, RWTH Aachen, Germany, May 2005.
2. B. Bollig and D. Kuske. Distributed Muller automata and logics. Research Report LSV-06-11, Laboratoire Spécification et Vérification, ENS Cachan, France, May 2006.
3. B. Bollig and M. Leucker. Message-passing automata are expressively equivalent to EMSO logic. *Theoretical Computer Science*, 358(2-3):150–172, 2006.
4. D. Brand and P. Zafiropolo. On communicating finite-state machines. *Journal of the ACM*, 30(2), 1983.
5. J. Büchi. Weak second order logic and finite automata. *Z. Math. Logik, Grundlag. Math.*, 5:66–62, 1960.
6. M. Droste, P. Gastin, and D. Kuske. Asynchronous cellular automata for pomsets. *Theoretical Computer Science*, 247(1-2):1–38, 2000.
7. W. Ebinger and A. Muscholl. Logical definability on infinite traces. *Theoretical Computer Science*, 154:67–84, 1996.
8. C. C. Elgot. Decision problems of finite automata design and related arithmetics. *Trans. Amer. Math. Soc.*, 98:21–52, 1961.
9. P. Gastin and A. Petit. Asynchronous cellular automata for infinite traces. In *Proceedings of ICALP 1992*, volume 623 of *LNCS*, pages 583–594. Springer, 1992.
10. B. Genest, D. Kuske, and A. Muscholl. A Kleene theorem and model checking algorithms for existentially bounded communicating automata. *Information and Computation*, 204(6):920–956, 2006.
11. W. Hanf. Model-theoretic methods in the study of elementary logic. In J. W. Addison, L. Henkin, and A. Tarski, editors, *The Theory of Models*. North-Holland, Amsterdam, 1965.
12. J. G. Henriksen, M. Mukund, K. Narayan Kumar, M. Sohoni, and P. S. Thiagarajan. A theory of regular MSC languages. *Information and Computation*, 202(1):1–38, 2005.
13. ITU-TS Recommendation Z.120: Message Sequence Chart 1999 (MSC99), 1999.
14. D. Kuske. Regular sets of infinite message sequence charts. *Information and Computation*, 187:80–109, 2003.
15. L. Libkin. *Elements of Finite Model Theory*. Springer, Berlin, 2004.
16. A. Muscholl. On the complementation of asynchronous cellular Büchi automata. *Theoretical Computer Science*, 169(2):123–145, 1996.
17. W. Thomas. On logical definability of trace languages. In *Proc. of Workshop of the ESPRIT Basic Research Action No 3166: Algebraic and Syntactic Methods in Computer Science (ASMICS), Kochel am See, Germany (1989)*, Report TUM-I9002, TU Munich, 1990.
18. W. Thomas. Languages, automata and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, pages 389–455. Springer, Berlin, 1997.