Benedikt Bollig and Dietrich Kuske

Distributed Muller Automata and Logics

Research Report LSV-06-11

May 2006

Laboratoire Spécification et Vérification

CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

Ecole Normale Supérieure de Cachan
61, avenue du Président Wilson
94235 Cachan Cedex France

Distributed Muller Automata and Logics

Benedikt Bollig[1] and Dietrich Kuske[2]

[1] LSV, CNRS UMR 8643 & ENS de Cachan
61 Av. du Président Wilson, F-94235 Cachan Cedex, France
bollig@lsv.ens-cachan.fr
[2] Institut für Informatik, Universität Leipzig
Augustusplatz 10–11, D-04109 Leipzig, Germany
kuske@informatik.uni-leipzig.de

Abstract. We consider Muller asynchronous cellular automata running on infinite dags over distributed alphabets. 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 an extension of the classical Ehrenfeucht-Fraïssé game to cope with infinite structures and the new first-order quantifier. As a byproduct, we obtain a logical characterization of unbounded Muller message-passing automata running on infinite message sequence charts.

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. Not only yield those connections beautiful theories; from a system developer's 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 so that an automatic translation from logic to automata is a desirable feature in the development process of reliable software.

The probably most famous connection between automata theory and classical logic has been established by Büchi and Elgot who showed that the language of a finite automaton can be expressed by a formula from monadic second-order (MSO) logic and that, conversely, any MSO formula can be effectively transformed into an equivalent finite automaton [3, 5]. When modeling distributed systems, finite automata, as a rather sequential model, are often no longer adequate to represent real-life implementations. More general models employ several components (which might be finite automata on their part) that are equipped with a communication mechanism: they might be connected by channels and capable of sending and receiving messages, or they might communicate by executing certain actions simultaneously. Asynchronous automata and asynchronous cellular automata over the partial-order based domain of Mazurkiewicz traces belong to the latter communication paradigm [18]. Actually, they also enjoy a logical characterization in terms of MSO logic [16], which also applies to message passing automata over message sequence charts (MSCs) as a representative of the channel-based communication paradigm, provided channels have bounded capacity [7, 9, 11].

For the modeling of reactive systems such as telecommunication systems, medical devices, and many more, which are often intended to run forever, one is interested in infinite behaviors. Indeed, Büchi showed that, over infinite words, MSO logic is still as expressive as finite automata that are equipped with a natural acceptance condition for infinite strings requiring that at least one final state is visited infinitely often. Such an acceptance condition comes in many flavors, and variations thereof give rise to Büchi, Muller, Rabin, and Streett automata (see [15] for an overview). They are all equivalent in the nondeterministic case. Restricting to deterministic automata, however, the Büchi condition suffers from loss of expressiveness, whereas the other three types retain their expressive power. The very same results apply to the settings of asynchronous (cellular) automata over infinite Mazurkiewicz traces, which might be equipped with a Büchi, Muller, Rabin, or Streett acceptance condition as well [6, 14]. More recently, Kuske extended the classical model of a message passing automaton by a Muller acceptance condition to make it capable of accepting infinite MSCs. He succeeded in characterizing that model in terms of an MSO logic over MSCs when the model is supposed to have bounded channel capacity [11].

In all cases mentioned above, one single (infinite) behavior was modeled as a partial order $(V, \leq)$ with a countable set V of events and a partial-order relation $\leq \subseteq V \times V$ to restrict the order of execution of dependent events (for example, of those that are executed by one and the same sequential process). Accordingly, all the logics that have been considered to characterize sets of infinite behaviors make use of (or can at least encode) the transitive closure predicate $x \leq y$, which formalizes that, in the partial order associated with a behavior, an event y is causally ordered after an event x. While such a predicate is still feasible by automata in the domains of words, Mazurkiewicz traces, and channel-bounded MSC languages, it was shown in [2] that, in the framework of unbounded MSC languages, it allows to define system behaviors that have no automata-theoretic counterpart. Moreover, even the first-order fragment of MSO logic over words employing $\leq$ as atomic predicate, involves a nonelementary transformation from formulas into finite automata. More precisely, there is no elementary bound that restricts the size of the automata arising from first-order formulas featuring $\leq$.

Following the approach of [2, 10, 13], we model one single execution of a system as a directed acyclic graph (V, E) without autoconcurrency (dag for short) rather than a partial order and employ the existential fragment of MSO (EMSO) logic with the atomic predicate xEy to express that x and y are connected with an edge. Nevertheless, this general class of dags allows us to model words, traces, and MSCs, as well as combinations of several communication paradigms. However, the logical formalism of EMSO formulas equipped with the atomic predicate xEy instead of $x \leq y$ appears to be weak as it is generally not capable of defining the transitive closure of the edge relation. Actually, an acceptance condition of a finite-state device such as "infinitely many final states are visited" cannot be expressed as an EMSO formula without transitive-closure predicate. 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 $\varphi(x)$. As we deal with structures of bounded degree (which would not be the case if we employed the transitive closure of the edge relation), we can exploit the close connection of Ehrenfeucht-Fraïssé games and locally threshold testable languages [12, 17].

Consider, for example, infinite words over the alphabet $\{a, b\}$ and consider two first-order logics: the one, $\mathrm{FO}(\sigma_{\lessdot})$, allows us to access the immediate successor y of a position x by means of a formula $x \lessdot y$ (xEy and $Succ(x, y)$ are other common notations); instead, the logic $\mathrm{FO}(\sigma_{\leq})$ features the predicate $x \leq y$. In $\mathrm{FO}(\sigma_{\lessdot})$, the set of words in which b occurs infinitely often, cannot be defined. However, this is an easy exercise in $\mathrm{FO}(\sigma_{\leq})$: $\forall x\exists y(x \leq y \wedge R_b(y))$ (with the suggested meaning of $R_b(y)$). Obviously, $\mathrm{FO}(\sigma_{\leq})$ is strictly more expressive than $\mathrm{FO}(\sigma_{\lessdot})$. If we equip both logics with the possibility to use existential second-order quantifications at the very beginning of a formula, we obtain the logics $\mathrm{EMSO}(\sigma_{\lessdot})$ and $\mathrm{EMSO}(\sigma_{\leq})$. Still, $\mathrm{EMSO}(\sigma_{\leq})$, is strictly more expressive than $\mathrm{EMSO}(\sigma_{\lessdot})$, as the latter cannot express the set of words with infinitely many b's either. We will see that, however, if we equip $\mathrm{EMSO}(\sigma_{\lessdot})$ with our quantifier $\exists^\infty$, the resulting logic, $\mathrm{EMSO}^\infty(\sigma_{\lessdot})$, is as powerful as $\mathrm{EMSO}(\sigma_{\leq})$. Relative to the domains of infinite words, it is even as powerful as full MSO logic. Nevertheless, the transformation of $\mathrm{EMSO}^\infty(\sigma_{\lessdot})$-formulas into automata is elementary.

As a model of a system implementation, we consider Muller asynchronous cellular automata over dags, which subsume the models of Büchi automata, Muller message passing automata over (infinite) MSCs, and Büchi asynchronous (cellular) automata over (infinite) traces, and we let them twinkle in a unifying framework. To establish a logical characterization of asynchronous cellular automata over arbitrary dags in terms of $\mathrm{EMSO}^\infty(\sigma_E)$ logic (i.e., EMSO logic over graphs enriched by the new first-order quantifier), we accordingly extend the classical theory of Ehrenfeucht-Fraïssé games and locally threshold testable languages. We argue that our model provides a unifying framework to characterize many well-known models of distributed systems over infinite behaviors. In particular, it yields a logical characterization of message passing automata over infinite MSCs in the general case of unbounded channels. Finally, in each case, we gain an effective transformation procedure from formulas to automata of elementary size.

2 Structures, Logic, and the Ehrenfeucht-Fraïssé Game

2.1 Structures and Monadic Second-Order Logic

In this paper, we consider 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 $\mathfrak{B}$.

The *Gaifman graph* $G(\mathfrak{A})$ 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 $(a_1, \ldots, a_n) \in P$ such that $a, b \in \{a_1, a_2, \ldots, a_n\}$. We will speak of the degree of a in $\mathfrak{A}$ whenever we actually mean the degree of a in the Gaifman graph $\mathfrak{A}$. 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 $\overline{a} = (a_1, \ldots, a_n) \in A^n$ and $b \in A$, we write $d(\overline{a}, b) = \min\{d(a_1, b), \ldots, d(a_n, b)\}$. Let $r \in \mathbb{N}$ and $\overline{c}$ denote the sequence of constants in the structure $\mathfrak{A}$. The *r-sphere of* $\mathfrak{A}$ is the substructure of $\mathfrak{A}$ generated by the universe $\{b \in A \mid d_{\mathfrak{A}}(\overline{c}, b) \leq r\}$ (note that, if $\mathfrak{A}$ does not contain any constants, this set is empty and the sphere is the empty structure). For a tuple $\overline{a}$ of elements in $\mathfrak{A}$, the *r-sphere of* $\mathfrak{A}$ *around* $\overline{a}$ is the r-sphere of the extension $(\mathfrak{A}, \overline{a})$ of $\mathfrak{A}$ by constants $\overline{a}$.

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

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

where $n \in \mathbb{N}$, $P \in \sigma$ is an n-ary predicate symbol, $x_1, x_2, \ldots \in \mathrm{Var}$, $c \in \sigma$ is a constant symbol and $X \in \mathrm{VAR}$.

Let $\mathfrak{A}$ be a σ-structure, $\varphi(x_1, \ldots, x_m, X_1, \ldots, X_n) \in \mathrm{MSO}^\infty$ a formula, and $\overline{a} = (a_1, \ldots, a_m) \in A^m$ and $\overline{A} = (A_1, \ldots, A_n) \in (2^A)^n$ tuples of elements and subsets of A. Then the *satisfaction* relation $\mathfrak{A} \models \varphi(\overline{a}, \overline{A})$ is defined as usual such that

- $\mathfrak{A} \models \exists^\infty y \varphi(\overline{a}, \overline{A})$ iff $\mathfrak{A} \models \varphi(a, \overline{a}, \overline{A})$ for infinitely many $a \in A$.

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

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

In the following, we sometimes omit the reference to σ and write, for example, FO instead of $\mathrm{FO}(\sigma)$. For any logic $\mathcal{L}$, we denote by $\mathrm{w}\mathcal{L}$ the *weak* variant of $\mathcal{L}$, which syntactically coincides with $\mathcal{L}$ but restricts the interpretation of second-order variables to finite sets.

Example 2.1. There are two ways to consider an ω-word $u = a_0 a_1 a_2 \ldots$ with $a_i \in \Sigma$ for some alphabet Σ as a relational structure. Both consider the set $\mathbb{N}$ of natural numbers as universe and contain the relations $P_a = \{i \in \mathbb{N} \mid a_i = a\}$ for $a \in \Sigma$. The difference lies in the way the universe $\mathbb{N}$ is structured: first, one can deal with the linear order $\leq$. This results in a structure $(\mathbb{N}, \leq, (P_a)_{a\in\Sigma})$ over a signature $\sigma_{\leq}$ with unary predicates P_a and one binary predicate $\leq$. Alternatively, one can deal with the successor relation $\lessdot$ resulting in a relational structure $(\mathbb{N}, \lessdot, (P_a)_{a\in\Sigma})$ over a signature $\sigma_{\lessdot}$ with unary predicates P_a and one binary predicate $\lessdot$. This gives a whole lot of logics to speak about infinite words, ranging from $\mathrm{FO}(\sigma_{\lessdot})$ to $\mathrm{MSO}^{\infty}(\sigma_{\leq})$. For any such logic $\mathcal{L}$, let $\mathbb{L}(\mathcal{L})$ denote the set of ω-languages that can be defined by a sentence from $\mathcal{L}$. Then we obtain the strict hierarchy depicted in Fig. 1 (provided Σ contains at least two letters). There, the topmost class $\mathbb{L}(\mathrm{EMSO}^{\infty}(\sigma_{\lessdot}))$ equals $\mathbb{L}(\mathcal{L})$ for any of the (non first-order) logics that are not mentioned. For the proof, see the appendix.

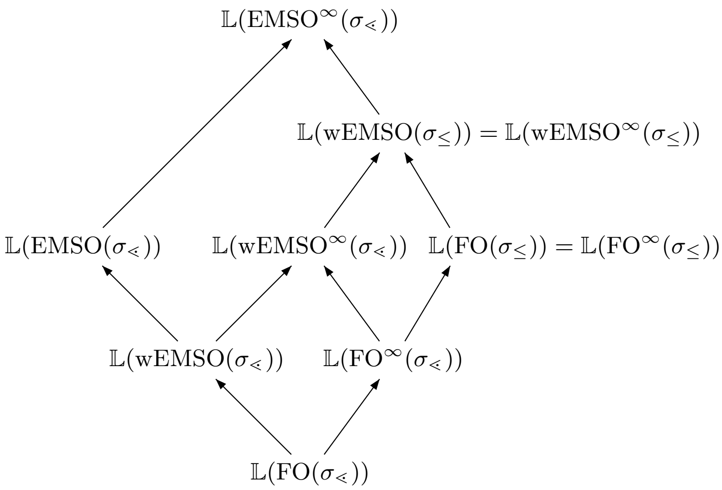

Fig. 1. The hierarchy of logics over infinite words

The *quantifier-rank* $\mathrm{qr}(\varphi)$ of a formula φ in FO^{∞} is the number of nested quantifiers in φ. More precisely, $\mathrm{qr}(\varphi) = 0$ if φ is atomic, $\mathrm{qr}(\neg\varphi) = \mathrm{qr}(\varphi)$, $\mathrm{qr}(\varphi \vee \psi) = \max\{\mathrm{qr}(\varphi), \mathrm{qr}(\psi)\}$, and $\mathrm{qr}(\exists x\varphi) = \mathrm{qr}(\exists^{\infty} x\varphi) = \mathrm{qr}(\varphi) + 1$. For $k \in \mathbb{N}$, we denote by $\mathrm{FO}^{\infty}[k]$ the set of first-order formulas of quantifier rank at most k. Any formula $\varphi \in \mathrm{FO}^{\infty}[k+1]$ is logically equivalent to a Boolean combination of formulas of the form $\exists^{\infty} x\psi$ and $\exists x\psi$ with $\psi \in \mathrm{FO}^{\infty}[k]$. Hence, by induction on k, one shows that $\mathrm{FO}^{\infty}[k]$ is a finite set up to logical equivalence.

Let $k, m \in \mathbb{N}$, $\mathfrak{A}$ be a σ-structure, and $\overline{a}$ an m-ary vector of elements of $\mathfrak{A}$. The *rank-k m-type* of $\overline{a}$ in $\mathfrak{A}$ comprises those FO^{∞} formulas of quantifier rank at most k that hold true for $\overline{a}$: $\mathit{type}_k(\mathfrak{A}, \overline{a}) := \{\varphi \in \mathrm{FO}^{\infty}[k] \mid \mathfrak{A} \models \varphi(\overline{a})\}$. Since, up to logical equivalence, $\mathrm{FO}^{\infty}[k]$ contains only finitely many formulas with free variables $x_1, x_2, \ldots, x_m$, the number of different rank-k m-types is finite.

Moreover, for any *rank-k* m-type T, there is a formula $\alpha_T(\overline{x}) \in \mathrm{FO}^\infty[k]$ (with $\overline{x}$ an m-tuple) such that, for every σ-structure $\mathfrak{A}$ and $\overline{a} \in A^m$, $\mathfrak{A} \models \alpha_T(\overline{a})$ iff $type_k(\mathfrak{A}, \overline{a}) = T$.

2.2 Ehrenfeucht-Fraïssé Games

The classical Ehrenfeucht-Fraïssé 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

$$\{(c^{\mathfrak{A}}, c^{\mathfrak{B}}) \mid c \in \sigma \text{ constant symbol}\}$$

is a partial isomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. If $k > 0$, the game proceeds as follows:

(1) The spoiler chooses either $\mathfrak{A}$ or $\mathfrak{B}$.
(2) The spoiler chooses an element of the structure selected in step (1) (i.e., either $a \in A$ or $b \in B$).
(3) The duplicator chooses an element in the other structure (i.e., either $b \in B$ or $a \in A$).
(4) 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 $\mathfrak{A} \equiv_k \mathfrak{B}$ if duplicator can force the play started in $(\mathfrak{A}, \mathfrak{B}, k)$ into a winning position. The existence of such a winning strategy describes precisely those properties that can be expressed using formulas of $\mathrm{FO}[k]$. See e.g. [12, 17] for the proof of the following theorem.

Theorem 2.2 (Ehrenfeucht-Fraïssé). *Let $\mathfrak{A}$ and $\mathfrak{B}$ be σ-structures and $k \in \mathbb{N}$. Then, $\mathfrak{A}$ and $\mathfrak{B}$ agree on* $\mathrm{FO}[k]$ *iff* $\mathfrak{A} \equiv_k \mathfrak{B}$.

We now propose an extension of the above game that will, in the same spirit, capture the expressive power of the extended logic $\mathrm{FO}^\infty[k]$. The only difference is that the game now proceeds from $(\mathfrak{A}, \mathfrak{B}, k)$ as follows (with $k > 0$):

(1) The spoiler chooses either $\mathfrak{A}$ or $\mathfrak{B}$ and chooses to proceed with (2) or (2').
(2) The spoiler chooses an element of the structure selected in step (1) (i.e., either $a \in A$ or $b \in B$).
(3) The duplicator chooses an element in the other structure (i.e., either $b \in B$ or $a \in A$).
(4) The game proceeds with $((\mathfrak{A}, a), (\mathfrak{B}, b), k - 1)$.
(2') The spoiler chooses an infinite subset Z of the structure selected in step (1) (i.e., either $Z \subseteq A$ or $Z \subseteq B$).
(3') The 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') The 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 such play started in $(\mathfrak{A}, \mathfrak{B}, k)$ into a winning position.

Theorem 2.3. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be σ-structures and let $k \in \mathbb{N}$. Then, $\mathfrak{A}$ and $\mathfrak{B}$ agree on $\mathrm{FO}^\infty[k]$ iff $\mathfrak{A} \equiv_k^\infty \mathfrak{B}$.*

Proof. 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 $\mathrm{FO}^\infty[k+1]$. We consider the case that, in step (2'), duplicator 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 $T = \mathit{type}_k(\mathfrak{A}, a)$ 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 $B' = \{b \in B \mid T = \mathit{type}_k(\mathfrak{B}, b)\}$ is infinite. For any $(a, b) \in A' \times B'$, we have $\mathit{type}_k(\mathfrak{A}, a) = \mathit{type}_k(\mathfrak{B}, b)$, i.e., $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ agree on $\mathrm{FO}^\infty[k]$. Hence, 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 some $a \in A$ in step (2), similar arguments allow duplicator to choose $b \in B$ with $(\mathfrak{A}, a) \equiv_k^\infty (\mathfrak{B}, b)$. For symmetry reasons, this implies that 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 $\exists^\infty x \varphi$, it suffices to consider these two cases. Suppose the set $Z = \{a \in A \mid \mathfrak{A} \models \varphi(a)\}$ is infinite. Spoiler can, in step (2'), choose this set Z. 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 $\mathrm{FO}^\infty[k]$. 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. □

2.3 Threshold Equivalence

In the context of structures of bounded degree (e.g., with respect to classes of graphs 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 [12, 17].

Definition 2.4 (Finite Threshold Equivalence). *Let $r, t \in \mathbb{N}$. Given σ-structures $\mathfrak{A}$ and $\mathfrak{B}$, we write $\mathfrak{A} \leftrightarrows_{r,t} \mathfrak{B}$ if, for any isomorphism type τ of an r-sphere around a single element,*

- *$|\mathfrak{A}|_\tau = |\mathfrak{B}|_\tau$ (where $|\mathfrak{A}|_\tau$ denotes the number of occurrences of τ in $\mathfrak{A}$) or*
- *both $t \le |\mathfrak{A}|_\tau$ and $t \le |\mathfrak{B}|_\tau$.*

In other words, $\leftrightarrows_{r,t}$ distinguishes structures on the basis of the number of realizations of r-spheres up to some threshold t.

Theorem 2.5 ("Hanf's Theorem" [8]). *For any $k, l \ge 0$, there are $r, t \ge 0$ such that, for any σ-structures $\mathfrak{A}$ and $\mathfrak{B}$ of degree at most l,*

$$\mathfrak{A} \leftrightarrows_{r,t} \mathfrak{B} \quad \textit{implies} \quad \mathfrak{A} \equiv_k \mathfrak{B} \ .$$

Thus, $\mathfrak{A}$ and $\mathfrak{B}$ agree on FO$[k]$ whenever $\mathfrak{A} \leftrightarrows_{r,t} \mathfrak{B}$. We now want to have a similar result for the logic FO$^{\infty}$.

Definition 2.6 (Threshold Equivalence). *Let $r, t \in \mathbb{N}$. Given σ-structures $\mathfrak{A}$ and $\mathfrak{B}$, we write $\mathfrak{A} \leftrightarrows^{\infty}_{r,t} \mathfrak{B}$ if, for any isomorphism type τ of an r-sphere of a single element,*

- $|\mathfrak{A}|_\tau = |\mathfrak{B}|_\tau$ *or*
- *both* $t \leq |\mathfrak{A}|_\tau < \infty$ *and* $t \leq |\mathfrak{B}|_\tau < \infty$.

Note the difference between $\leftrightarrows_{r,t}$ and $\leftrightarrows^{\infty}_{r,t}$: The former does not distinguish between "many" and "infinitely many" realizations of a sphere. The latter identifies all natural numbers $t, t+1, \ldots$, but makes a difference between them and infinity.

Theorem 2.7. *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,*

$$\mathfrak{A} \leftrightarrows^{\infty}_{r,t} \mathfrak{B} \quad \textit{implies} \quad \mathfrak{A} \equiv^{\infty}_{k} \mathfrak{B} \ .$$

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 $i \cdot c$ where c is the maximal size of an r_i-sphere whose elements have degree at most l.

We first show that for any $k \in \mathbb{N}$, if one round of the extended Ehrenfeucht-Fraïssé game is played from $(\mathfrak{A}, \mathfrak{B}, k+1)$ with $\mathfrak{A} \leftrightarrows^{\infty}_{r_{k+1},t_{k+1}} \mathfrak{B}$, then duplicator can enforce $(\mathfrak{A}, a) \leftrightarrows^{\infty}_{r_k,t_k} (\mathfrak{B}, b)$ for the following game position $((\mathfrak{A}, a), (\mathfrak{B}, b), k)$.

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}-Sph$(\mathfrak{A}, a) \cong r_{k+1}$-Sph$(\mathfrak{A}, a') =: S$. Since r_{k+1}-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 $\mathfrak{A} \leftrightarrows^{\infty}_{r_{k+1},t_{k+1}} \mathfrak{B}$, we obtain the existence of infinitely many $b \in B$ with r_{k+1}-Sph$(\mathfrak{B}, b) \cong S$. Let B' denote this infinite set. Then duplicator chooses these two sets A' and B' in step (3'). Now let $a \in A'$ and $b \in B'$ be chosen by spoiler in step (4'). We want to show that $(\mathfrak{A}, a) \leftrightarrows^{\infty}_{r_k,t_k} (\mathfrak{B}, b)$. For this, let $a' \in A$ be arbitrary. Then r_k-Sph$(\mathfrak{A}, a, a')$ can be determined as follows (for the details see, e.g., [12, Sect. 4.2])

1. if $d(a, a') > 2r_k + 1$, then it is completely determined by r_{k+1}-Sph$(\mathfrak{A}, a)$ and r_{k+1}-Sph$(\mathfrak{A}, a')$
2. if $d(a, a') \leq 2r_k + 1$, then it equals the r_k-sphere in r_{k+1}-Sph$(\mathfrak{A}, a)$ around a'.

These observations imply $(\mathfrak{A}, a) \leftrightarrows^{\infty}_{r_k,t_k} (\mathfrak{B}, b)$.

Thus, duplicator can force any play from $(\mathfrak{A}, \mathfrak{B}, k)$ with $\mathfrak{A} \leftrightarrows^{\infty}_{r_k,t_k} \mathfrak{B}$ into a game position $(\mathfrak{A}', \mathfrak{B}', 0)$ with $\mathfrak{A}' \leftrightarrows^{\infty}_{r_0,t_0} \mathfrak{B}'$. Note that r_0-Sph$(\mathfrak{A}')$ equals the restriction of $\mathfrak{A}'$ to its constants and that $\mathfrak{A} \equiv^{\infty}_{r_0} \mathfrak{B}$ holds if and only if these restrictions are isomorphic. Thus, duplicator can force the play into a winning position. □

3 Asynchronous Cellular Automata with Types

We fix a nonempty finite set Ag of *agents*, a *distributed alphabet* $\widetilde{\Sigma}$, which is a tuple $(\Sigma_i)_{i \in Ag}$ of (not necessarily disjoint) alphabets Σ_i, and an alphabet C. In the following, let Σ stand for $\bigcup_{i \in Ag} \Sigma_i$, the set of *actions*. Its elements will label the nodes of a dag, while the elements of C will label its edges, providing some control information. Elements from Σ_i are understood to be actions that are performed by agent i. So let, for $a \in \Sigma$, $loc(a) := \{i \in Ag \mid a \in \Sigma_i\}$ denote the set of agents that are involved in a.

3.1 Infinite Dags over Distributed Alphabets

We now introduce the models representing the behavior of a system of communicating agents. In doing so, we combine the standard models of [4] and [10].

Definition 3.1 ($(\widetilde{\Sigma}, C)$-Dag). *A* $(\widetilde{\Sigma}, C)$-dag *is a structure* $(V, \{E_\ell\}_{\ell \in C}, \lambda)$ *with* $E_\ell \subseteq V \times V$ *mutually disjoint binary* edge relations *on* V*, and* $\lambda : V \to \Sigma$ *a* labeling function *such that the following holds with* $E = \bigcup_{\ell \in C} E_\ell$.

- *E is irreflexive and E^* is a partial order on V,*
- *for any $v \in V$, $\{u \in V \mid (u, v) \in E^*\}$ is a finite set,*
- *for any $i \in Ag$, $\lambda^{-1}(\Sigma_i)$ is totally ordered by E^*, and*
- *for any $\ell \in C$ and any $(u, v), (u', v') \in E_\ell$ with $\lambda(u) = \lambda(u')$ and $\lambda(v) = \lambda(v')$, we have $(u, u') \in E^*$ iff $(v, v') \in E^*$.*

The idea is that this models a distributed execution. The set V consists of the events of this execution, the mapping λ assigns the action executed by an event, and the relations E_ℓ describe the communication structure.

We conclude that, in a $(\widetilde{\Sigma}, C)$-dag $(V, \{E_\ell\}_{\ell \in C}, \lambda)$, for any $u \in V$, $\ell \in C$, and $a \in \Sigma$, there is at most one vertex $v \in V$ such that both $uE_\ell v$ ($vE_\ell u$) and $\lambda(v) = a$.

For a $(\widetilde{\Sigma}, C)$-dag $(V, \{E_\ell\}_{\ell \in C}, \lambda)$, we will use the following abbreviations:

- $E = \bigcup_{\ell \in C} E_\ell$ is the set of edges, $\leq\, = E^*$ is the transitive and reflexive closure of E, and $<\, = E^+$ is the irreflexive part of $\leq$ (it is irreflexive since E is required to be irreflexive)
- for any agent $i \in Ag$, let $V_i = \lambda^{-1}(\Sigma_i)$ denote the set of nodes with label in Σ_i. A node $u \in V$ is Σ_i*-maximal* if $u \in V_i$ and there is no $v \in V_i$ such that $u < v$.

3.2 The Operational Behavior of Infinite Dags

Let Q be a nonempty finite set of *states*. A Q*-extended* $(\widetilde{\Sigma}, C)$*-dag* is a structure $(V, \{E_\ell\}_{\ell \in C}, \lambda, \rho)$ such that $(V, \{E_\ell\}_{\ell \in C}, \lambda)$ is a $(\widetilde{\Sigma}, C)$-dag and $\rho : V \to Q$ is a mapping. While, as explained above, the mapping λ assigns basic actions to the events of a distributed execution, the function ρ describes the local state assumed after executing an event. This state will depend on the states and actions of the events immediately preceding the current one. To define this, the following notions will turn out to be useful: Let $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda, \rho)$ be some Q-extended $(\widetilde{\Sigma}, C)$-dag and $v \in V$ be some node.

- The *read domain* $\mathrm{Read}(v)$ of the node v is the set of pairs $(b, \ell) \in \Sigma \times C$ such that there exists a node $u \in V$ with $\lambda(u) = b$ and $(u, v) \in E_\ell$. We call this (unique) node u the (b, ℓ)*-predecessor of* v and denote it by (b, ℓ)-$\mathrm{pred}(v)$.
- The *write domain* $\mathrm{Write}(v)$ is the set of pairs $(b, \ell) \in \Sigma \times C$ such that there exists a node $w \in V$ with $\lambda(w) = b$ and $(v, w) \in E_\ell$.
- If $\iota \in Q$ is some designated element, then the *transition taken at* v is the tuple $\mathit{trans}^\iota_{\mathcal{D}}(v) = (\overline{q}, \lambda(v), \rho(v))$ with $\overline{q} \in Q^{\Sigma \times C}$ where, for any $(b, \ell) \in \Sigma \times C$,

$$\overline{q}[(b, \ell)] = \begin{cases} \rho((b, \ell)\text{-pred}(v)) & \text{if } (b, \ell) \in \mathrm{Read}(v) \\ \iota & \text{otherwise} \end{cases}$$

Now we are ready to define distributed automata whose behavior will be described by $(\widetilde{\Sigma}, C)$-dags. Later, they will be equipped with an acceptance condition.

Definition 3.2 (Asynchronous Cellular Automaton with Types). *An* asynchronous cellular automaton with types *or* ACAT *over the pair* $(\widetilde{\Sigma}, C)$ *is a structure* $\mathcal{A} = (Q, \iota, \Delta, T)$ *where*

- Q *is the nonempty finite set of* states*,*
- $\iota \in Q$ *is some* initial state*,*
- $\Delta \subseteq Q^{\Sigma \times C} \times \Sigma \times Q$ *is the set of* transitions*, and*
- $T : (\Sigma \times Q) \to 2^{\Sigma \times C}$ *is the* type function*.*

We call the ACAT $\mathcal{A}$ deterministic *if, for any* $\overline{q} \in Q^{\Sigma \times C}$ *and* $a \in \Sigma$*, there is at most one state* $q \in Q$ *such that* $(\overline{q}, a, q) \in \Delta$*.*

Definition 3.3. *Let* $\mathfrak{A} = (Q, \iota, \Delta, T)$ *be an ACAT over* $(\widetilde{\Sigma}, C)$ *and let* $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda)$ *be a* $(\widetilde{\Sigma}, C)$*-dag. A run of* $\mathcal{A}$ *on* $\mathcal{D}$ *is a mapping* $\rho : V \to Q$ *such that, for any* $u \in V$*,*

- *the transition taken at* u *is legal, i.e.,* $\mathit{trans}^\iota_{(\mathcal{D},\rho)}(u) \in \Delta$*, and*
- *the communication request for* u *is satisfied, i.e.,* $T(\lambda(u), \rho(u)) \subseteq \mathrm{Write}(u)$*.*

Even without any notion of acceptance, these automata can be used to compute the sphere around any node of a dag. This feature, described formally in the following proposition, is the key connection between these automata and the logical characterization of first-order expressible properties as described in the previous section.

Proposition 3.4 (cf. [1]). *Let* $r \in \mathbb{N}$*. There are an ACAT* $\mathcal{A}_r = (Q, \iota, \Delta, T)$ *over* $(\widetilde{\Sigma}, C)$ *and a mapping* η *from* Q *into the set of spheres of radius* r *such that for any* $(\widetilde{\Sigma}, C)$*-dag* $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda)$*,*

– *there exists a run of $\mathcal{A}_r$ on $\mathcal{D}$, and*
– *for any run ρ of $\mathcal{A}_r$ on $\mathcal{D}$ and any $u \in V$, we have $\eta(\rho(u)) = r$-$\mathrm{Sph}(\mathcal{D}, u)$.*

We will now extend our automata model by some acceptance modes that originate from the work on automata on infinite words. Since all except the Staiger-Wagner acceptance depend on those states that appear infinitely often in a run, we first give the following definitions. Let $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda, \rho)$ be some Q-extended $(\widetilde{\Sigma}, C)$-dag (we think of Q as the set of states of some ACAT and of ρ as a run on $(V, \{E_\ell\}_{\ell \in C}, \lambda)$). For $\iota \in Q$, we define functions $\mathit{final}^\iota_{\mathcal{D}} : Ag \to 2^Q$ and $\mathit{efinal}^\iota_{\mathcal{D}} : Ag \to 2^Q \times \{\infty, \overline{\infty}\}$ as follows (with $i \in Ag$):

$$\mathit{final}^\iota_{\mathcal{D}}[i] = \begin{cases} \{q \mid \forall u \in V_i \; \exists v \in V_i : u \leq v \text{ and } q = \rho(v)\} & \text{if } V_i \neq \emptyset \\ \{\iota\} & \text{otherwise} \end{cases}$$

$$\mathit{efinal}^\iota_{\mathcal{D}}[i] = \begin{cases} (\mathit{final}^\iota_{\mathcal{D}}[i], \overline{\infty}) & \text{if } V_i \text{ is finite} \\ (\mathit{final}^\iota_{\mathcal{D}}[i], \infty) & \text{otherwise} \end{cases}$$

If V_i is finite, then $\mathit{final}^\iota_{\mathcal{D}}[i]$ describes the state assumed at the Σ_i-maximal event (which is the initial state ι if V_i is even empty). If V_i is infinite, then $\mathit{final}^\iota_{\mathcal{D}}[i]$ is the set of states assumed infinitely often. If $\mathit{final}^\iota_{\mathcal{D}}[i]$ is a singleton, we do not know whether V_i is finite or not – this additional information is present in $\mathit{efinal}^\iota_{\mathcal{D}}[i]$.

Definition 3.5. *A* Staiger-Wagner or Muller ACAT *over $(\widetilde{\Sigma}, C)$ is a tuple $\mathcal{A} = (Q, \iota, \Delta, T, \mathcal{F})$ such that (Q, ι, Δ, T) is an ACAT and $\mathcal{F} \subseteq (2^Q)^{Ag}$.*

Now let ρ be some run of $\mathcal{A}$ on the $(\widetilde{\Sigma}, C)$-dag $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda)$.

(1) If $\mathcal{A}$ is a Staiger-Wagner ACAT, then the run ρ is accepting if we have $(\rho^{-1}(V_i))_{i \in Ag} \in \mathcal{F}$.
(2) If $\mathcal{A}$ is a Muller ACAT, then the run ρ is accepting in case $\mathit{final}^\iota_{(\mathcal{D},\rho)} \in \mathcal{F}$.

Definition 3.6. *A* generalized Muller ACAT *is a tuple $\mathfrak{A} = (Q, \iota, \Delta, T, \mathcal{F})$ such that (Q, ι, Δ, T) is an ACAT and $\mathcal{F} \subseteq (2^Q \times \{\overline{\infty}, \infty\})^{Ag}$.*

Now let ρ be some run of $\mathfrak{A}$ on the $(\widetilde{\Sigma}, C)$-dag $\mathcal{D}$. Then ρ is accepting if $\mathit{efinal}^\iota_{(\mathcal{D},\rho)} \in \mathcal{F}$.

If $\mathcal{A}$ is some of these ACAT, then the *language $L(\mathcal{A})$ accepted by $\mathcal{A}$* is the set of those $(\widetilde{\Sigma}, C)$-dags that admit an accepting run of $\mathcal{A}$.

4 The Accepting Power of ACATs

We will deal with the accepting power of ACATs in the context of $(\widetilde{\Sigma}, C)$-dags and of message sequence charts.

4.1 ACATs over $(\widetilde{\Sigma}, C)$-dags

In this section, we consider the class of all $(\widetilde{\Sigma}, C)$-dags. We will relate the expressive power of all types of ACATs and the extended logic. As far as the logic is concerned, $(\widetilde{\Sigma}, C)$-dags can be seen as relational structures whose signature contains binary relations E_ℓ for $\ell \in C$ and unary relations R_a for $a \in \Sigma$. As expected, we will write the formula $R_a(x)$ as $\lambda(x) = a$. Furthermore, we write $\mathrm{EMSO}^\infty(\widetilde{\Sigma}, C)$ for $\mathrm{EMSO}^\infty((E_\ell)_{\ell \in C}, (R_a)_{a \in \Sigma})$, $\mathrm{FO}^\infty(\widetilde{\Sigma}, C)$ etc. are to be understood similarly.

Generalized Muller ACATs

The following closure properties of generalized Muller ACATs will turn out to be useful:

Proposition 4.1. *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be generalized Muller automata over $(\widetilde{\Sigma}, C)$. There are generalized Muller automata $\mathcal{A}$ and $\mathcal{B}$ such that $L(\mathcal{A}) = L(\mathcal{A}_1) \cup L(\mathcal{A}_2)$ and $L(\mathcal{B}) = L(\mathcal{A}_1) \cap L(\mathcal{A}_2)$.*

Proof. Suppose $\mathcal{A}_1 = (Q_1, \iota_1, \Delta_1, T_1, \mathcal{F}_1)$ and $\mathcal{A}_2 = (Q_2, \iota_2, \Delta_2, T_2, \mathcal{F}_2)$. Without loss of generality, we assume that Q_1 and Q_2 are disjoint. In the following, we only give the construction of the automata and leave the actual proof of correctness to the interested reader.

To recognize $L(\mathcal{A}_1) \cup L(\mathcal{A}_2)$, $\mathcal{A} = (Q, \iota, \Delta, T, \mathcal{F})$ simulates either $\mathcal{A}_1$ or $\mathcal{A}_2$. Hence its set of states equals $Q = Q_1 \uplus Q_2 \uplus \{\iota\}$. To define the set of transitions, we need the partial functions $\mathrm{n}_j : Q \to Q_j$ with $\mathrm{n}_j(\iota) = \iota_j$, $\mathrm{n}_j(q) = q$ for $q \in Q_j$, and $\mathrm{n}_j(q)$ undefined for $q \in Q_{3-j}$. Then we set

$$\begin{aligned}
\Delta &= \Delta_1 \cup \Delta_2 \\
&\quad \cup \{(\bar{q}, a, q) \in Q^{\Sigma \times C} \times \Sigma \times Q \mid (\mathrm{n}_1 \circ \bar{q}, a, q) \in \Delta_1 \text{ or } (\mathrm{n}_2 \circ \bar{q}, a, q) \in \Delta_2\} \\
\mathcal{F} &= \mathcal{F}_1 \cup \mathcal{F}_2 \\
&\quad \cup \{(F'_i, \theta_i)_{i \in Ag} \mid \exists j \in \{1,2\} \exists F_i \subseteq Q_j : (F_i, \theta_i)_{i \in Ag} \in \mathcal{F}_j, \\
&\qquad\qquad F_i \setminus \{\iota_j\} = F'_i \setminus \{\iota, \iota_j\}, \\
&\qquad\qquad \iota_j \in F_i \iff F'_i \cap \{\iota, \iota_j\} \neq \emptyset\}
\end{aligned}$$

$$T(a, q) = \begin{cases} \emptyset & \text{if } q = \iota \\ T_j(a, q) & \text{if } q \in Q_j \end{cases}$$

The generalized Muller ACAT $\mathcal{B} = (Q', \iota', \Delta', T', \mathcal{F}')$, to recognize $L(\mathcal{A}_1) \cap L(\mathcal{A}_2)$, connects $\mathcal{A}_1$ and $\mathcal{A}_2$ in parallel: we set $Q' = Q_1 \times Q_2$, $\iota' = (\iota_1, \iota_2)$, $T'(a, (q_1, q_2)) = T_1(a, q_1) \cup T_2(a, q_2)$ for any $a \in \Sigma$, $q_1 \in Q_1$, and $q_2 \in Q_2$. Moreover, for states $q_i^{(b,\ell)} \in Q_i$ (where $(b, \ell) \in \Sigma \times C$ and $i \in \{1, 2\}$), $a \in \Sigma$, and $(q_1, q_2) \in Q'$, $((q_1^{(b,\ell)}, q_2^{(b,\ell)})_{(b,\ell) \in \Sigma \times C}, a, (q_1, q_2)) \in \Delta'$ iff both $((q_1^{(b,\ell)})_{(b,\ell) \in \Sigma \times C}, a, q_1) \in \Delta_1$ and $((q_2^{(b,\ell)})_{(b,\ell) \in \Sigma \times C}, a, q_2) \in \Delta_2$. Finally, a mapping $\overline{q} : Ag \to 2^{Q'} \times \{\overline{\infty}, \infty\}$ is contained in $\mathcal{F}'$ iff there are $\overline{q}_1 : Ag \to 2^{Q_1} \times \{\overline{\infty}, \infty\} \in \mathcal{F}_1$ and $\overline{q}_2 : Ag \to 2^{Q_2} \times \{\overline{\infty}, \infty\} \in \mathcal{F}_2$ such that, for any agent $i \in Ag$,

– $\pi_2(\overline{q}[i]) = \pi_2(\overline{q}_1[i]) = \pi_2(\overline{q}_2[i])$ and
– $\pi_1(\overline{q}_1[i]) = \{q \in Q_1 \mid \exists q' \in Q_2 : (q, q') \in \pi_1(\overline{q}[i])\}$ and
$\pi_1(\overline{q}_2[i]) = \{q \in Q_2 \mid \exists q' \in Q_1 : (q', q) \in \pi_2(\overline{q}[i])\}$. □

Lemma 4.2. *Let $r \in \mathbb{N}$, $t \in \mathbb{N} \cup \{\infty\}$ and S be some r-sphere in some $(\widetilde{\Sigma}, C)$-dag around a single vertex. There exists a generalized Muller ACAT $\mathcal{A}$ such that $L(\mathcal{A})$ is the set of $(\widetilde{\Sigma}, C)$-dags $\mathcal{D}$ satisfying*

$$|\mathcal{D}|_S = t \text{ or } t < |\mathcal{D}|_S < \infty \ .$$

Proof. Let $\mathcal{A}_r = (Q, \iota, \Delta, T)$ and η denote the ACAT and the mapping from Prop. 3.4. Let $i \in Ag$ and, furthermore, suppose that the node in the center of S is labeled by $a \in \Sigma_i$.

If $t = \infty$, then this ACAT is just extended by a generalized Muller condition $\mathcal{F}$, namely the set of tuples $(F_j, \theta_j)_{j \in Ag}$ such that $\theta_i = \infty$ and $F_i \subseteq \eta^{-1}(S)$. Now let $\mathcal{D}$ be a $(\widetilde{\Sigma}, C)$-dag and ρ some run of the generalized Muller ACAT $(\mathcal{A}_r, \mathcal{F})$ on $\mathcal{D}$. Then ρ is accepting if and only if there are infinitely many nodes u such that $\eta(\rho(u)) = S$. Now, by Prop. 3.4, this is equivalent to $|\mathcal{D}|_S = \infty = t$. Since any $(\widetilde{\Sigma}, C)$-dag admits some run of $\mathcal{A}_r$, the result follows.

From now on, let $t \in \mathbb{N}$. Then we extend the states of $\mathcal{A}_r$ by a counter that counts up to t, i.e., $Q' = Q \times \{0, 1, \ldots, t\}$. The new initial state will be $(\iota, 0)$. Now let, for $(b, \ell) \in \Sigma \times C$, $(q_{(b,\ell)}, n_{(b,\ell)}) \in Q'$, $k \in Ag$, $a \in \Sigma_k$, and $(q, n) \in Q'$. Then $((q_{(b,\ell)}, n_{(b,\ell)})_{(b,\ell) \in \Sigma \times C}, a, (q, n))$ is a transition of Δ' iff $((q_{(b,\ell)})_{(b,\ell) \in \Sigma \times C}, a, q) \in \Delta$ is a transition of the ACAT $\mathcal{A}_r$, and

$$n = \begin{cases} \max\{n_{(b,\ell)} \mid (b, \ell) \in \Sigma \times C\} & \text{if } \eta(q) \neq S \\ \min\{t, 1 + \max\{n_{(b,\ell)} \mid (b, \ell) \in \Sigma \times C\}\} & \text{if } \eta(q) = S \end{cases}$$

Furthermore, $T'(a, (q, n)) = T(a, q)$ for any $a \in \Sigma$ and $(q, n) \in Q'$. Finally, let $\mathcal{F}'$ denote the set of pairs $(F_j, \theta_j)_{j \in Ag}$ such that

– $\theta_i = \overline{\infty}$ implies $F_i \subseteq Q \times \{t\}$ and
– $\theta_i = \infty$ implies $F_i \subseteq (Q \setminus \eta^{-1}(S)) \times \{t\}$.

Then consider the generalized Muller ACAT $\mathcal{A} = (Q', \Delta', (\iota, 0), T')$.

Any run of $\mathcal{A}$ first simulates a run of $\mathcal{A}_r$ and, secondly, counts the occurrences of the r-sphere S. This counting is correct since all the occurrences of the r-sphere are linearly ordered by the transitive closure of the edge relation. If the number of occurrences exceeds t, the automaton does not count up. □

Theorem 4.3. *Let L be a set of $(\widetilde{\Sigma}, C)$-dags. Then the following are equivalent*

1. *there exists a generalized Muller ACAT $\mathcal{A}$ such that $L = L(\mathcal{A})$*
2. *there exists an* $\mathrm{EMSO}^\infty(\widetilde{\Sigma}, C)$ *sentence φ such that L is the set of $(\widetilde{\Sigma}, C)$-dags $\mathcal{D}$ with $\mathcal{D} \models \varphi$.*

Proof. The construction of an $\text{EMSO}^\infty(\widetilde{\Sigma}, C)$ sentence from a given generalized Muller ACAT follows similar instances of that problem, e.g., [4]. Basically, the second order variables $(X_q)_{q \in Q}$ encode an assignment of states to vertices. The first-order body of the formula then expresses that this assignment is a run. In addition, we have to take care of the accepting condition. Any such condition $\overline{q} \in (2^Q \times \{\overline{\infty}, \infty\})^{Ag}$ is translated into a conjunction of the following conjuncts for $i \in Ag$

- if $\overline{q}[i] = (F, \overline{\infty})$ with F not a singleton, or $\overline{q}[i] = (\emptyset, \infty)$, then the whole conjunction is false (e.g., $\exists x (x \neq x)$).
- if $\overline{q}[i] = (F, \infty)$, then we have a conjunct $\bigwedge_{q \in F} \exists^\infty x (x \in X_q \wedge \lambda(x) \in \Sigma_i) \wedge \bigwedge_{q \in Q \setminus F} \neg \exists^\infty x (x \in X_q \wedge \lambda(x) \in \Sigma_i)$
- if $\overline{q}[i] = (\{q\}, \overline{\infty})$, then we have a conjunct

$$\exists X \exists x \begin{pmatrix} x \in X_q \wedge \lambda(x) \in \Sigma_i \wedge x \in X \\ \wedge \forall y, z (y \in X \wedge (y, z) \in E \rightarrow z \in X) \\ \wedge \forall y (y \in X \wedge \lambda(y) \in \Sigma_i \rightarrow x = y) \end{pmatrix}$$

  (supplemented by $\dots \vee \forall x \neg \lambda(x) \in \Sigma_i$ if $q = \iota$). The above formula expresses that X contains at least those nodes y with $x \leq y$. Hence it ensures that x is the maximal element of V_i.

Now suppose $\varphi = \exists X_1 \exists X_2 \dots \exists X_n \, \psi$ to be an $\text{EMSO}^\infty(\widetilde{\Sigma}, C)$-sentence with $\psi \in \text{FO}^\infty(\widetilde{\Sigma}, C)$. For $i \in Ag$, let $\Gamma_i = \Sigma_i \times 2^n$ (where 2^n denotes the powerset of $\{1, 2, \dots, n\}$) and $\widetilde{\Gamma} = (\Gamma_i)_{i \in Ag}$. Let ψ' be obtained from ψ by the following replacements

- any occurrence of $\lambda(x) = a$ with $\bigvee_{M \in 2^n} \lambda(x) = (a, M)$
- any occurrence of $x \in X_i$ with $\bigvee_{a \in \Sigma, M \in 2^n} \lambda(x) = (a, M \cup \{i\})$.

Let H be the set of $(\widetilde{\Gamma}, C)$-dags that satisfy ψ'. Then, by Theorem 2.7, H is a finite union of $\leftrightarrows^\infty_{r,t}$-equivalence classes for some $r, t \in \mathbb{N}$. By Proposition 4.1 the intersection and the union of sets accepted by generalized Muller ACATs can be accepted by a generalized Muller ACAT. Hence, by Lemma 4.2, there is a generalized Muller ACAT $\mathcal{A}'$ that accepts H.

Any $(\widetilde{\Gamma}, C)$-dag (V, E, λ) with $\lambda : V \to \Gamma$ can be seen as a 2^n-extended $(\widetilde{\Sigma}, C)$-dag $(V, E, \pi_1 \circ \lambda, \pi_2 \circ \lambda)$ (where π_i is the ith projection from $\Sigma \times 2^n$). To check whether some $(\widetilde{\Sigma}, C)$-dag $\mathcal{D}$ satisfies φ, a generalized Muller ACAT $\mathcal{A}$ will therefore guess for any node the additional labeling from the set 2^n and then simulate $\mathcal{A}'$. $\square$

The number of states of the generalized Muller ACAT $\mathcal{A}$ constructed from an $\text{EMSO}^\infty(\widetilde{\Sigma}, C)$-formula φ is elementary in the size of the formula φ. The radius r from Theorem 2.7 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 ACAT from Prop. 3.4 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 $\leftrightarrows^\infty_{r,t}$ that contribute to H (in the proof above).

Staiger-Wagner ACATs

The following lemma describes the counting power of Staiger-Wagner ACATs: as far as finite counting is concerned, they can do as much as generalized Muller ACATs can do.

Lemma 4.4. *Let $r, t \in \mathbb{N}$ and S be some r-sphere in some $(\widetilde{\Sigma}, C)$-dag around a single vertex. There exists a Staiger-Wagner ACAT $\mathcal{A}$ such that $L(\mathcal{A})$ is the set of $(\widetilde{\Sigma}, C)$-dags $\mathcal{D}$ satisfying*

$$|\mathcal{D}|_S = t \text{ or } t < |\mathcal{D}|_S \ .$$

Proof. The proof differs only slightly from that of Lemma 4.2: one again starts from the ACAT $\mathcal{A}_r$ and the mapping η, extends the states by a counter, and defines the transition relation and the initial states as there. But the acceptance condition $\mathcal{F} \subseteq (2^{Q'})^{Ag}$ now contains all functions $\overline{q} : Ag \to 2^{Q'}$ such that there exists $q \in Q$ with $(q, t) \in \overline{q}[i]$. □

Theorem 4.5. *Let L be a set of $(\widetilde{\Sigma}, C)$-dags. Then the following are equivalent*

1. *there exists a Staiger-Wagner ACAT $\mathcal{A}$ such that $L = L(\mathcal{A})$*
2. *there exists an $\mathrm{EMSO}(\widetilde{\Sigma}, C)$ sentence φ such that L is the set of $(\widetilde{\Sigma}, C)$-dags $\mathcal{D}$ with $\mathcal{D} \models \varphi$.*

Proof. The proof is similar to the proof of Theorem 4.3. The only difference in the transformation of an automaton into a formula concerns the acceptance condition, which, this time, is given as (a set of) function(s) $\overline{q} : Ag \to 2^Q$. It is expressed as a conjunction of the following conjunct for any $i \in Ag$:

$$\bigwedge_{q \in \overline{q}[i]} \exists x (x \in X_q \wedge \lambda(x) \in \Sigma_i) \wedge \bigwedge_{q \in Q \setminus \overline{q}[i]} \forall x (\lambda(x) \in \Sigma_i \to x \notin X_q)$$

For the other transformation, one uses Theorem 2.5 instead of Theorem 2.7 and Lemma 4.4 instead of 4.2. □

Muller ACATs

Recall that, in a generalized Muller ACAT, 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 acceptance condition. To solve this problem, we first show that a Muller ACAT can determine whether a particular agent performs finitely or infinitely many actions.

Lemma 4.6. *Let $k \in Ag$. There exist Muller ACATs $\mathcal{A}$ and $\mathcal{B}$ such that, for any $(\widetilde{\Sigma}, C)$-dag $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda)$, we have*

- *$\mathcal{D} \in L(\mathcal{A})$ iff V_k is infinite*

– $\mathcal{D} \in L(\mathcal{B})$ *iff* V_k *is finite*

Proof. Let $Q = \{0, 1, 2, 3\}$. The state 0 will be the initial state; it signals in particular that no action from Σ_k has been executed. The state 3 signals that some Σ_k-action has been executed but that process k will not move again. Finally, the states 1 and 2 denote that some Σ_k-action has been executed and that some will be executed in the future. To ensure this, let $\overline{q} : \Sigma \times C \to Q$, $a \in \Sigma$, and $q \in Q$. Then $(\overline{q}, a, q) \in \Delta$ if

– $q \geq \max\{\overline{q}[b, \ell] \mid (b, \ell) \in \Sigma \times C\}$ and $a \notin \Sigma_k$ or
– $q \geq 1$, $3 \notin \{\overline{q}[b, \ell] \mid (b, \ell) \in \Sigma \times C\}$ and $a \in \Sigma_k$.

Furthermore, we set

– $\mathcal{F}_\mathcal{A} = \{\overline{q} : Ag \to 2^Q \mid \overline{q}[k] = \{1, 2\}\}$
– $\mathcal{F}_\mathcal{B} = \{\overline{q} : Ag \to 2^Q \mid \{1, 2\} \cap \overline{q}[k] = \emptyset\}$

Finally, set $\mathcal{A} = (Q, 0, \Delta, T, \mathcal{F}_\mathcal{A})$ and $\mathcal{B} = (Q, 0, \Delta, T, \mathcal{F}_\mathcal{B})$ with $T(a, q) = \emptyset$ for all $q \in Q$ and $a \in \Sigma$.

We first show the statement of the lemma for $\mathcal{A}$. If $\mathcal{D}$ is accepted by $\mathcal{A}$, then there exists a run ρ of $\mathcal{A}$ on $\mathcal{D}$ with $\mathit{final}^0_{(\mathcal{D},\rho)}[k]$ not a singleton. Hence V_k is infinite. Conversely, let V_k be infinite. Then define $\rho : V \to Q$ such that

1. $\rho^{-1}(1) \cap V_k$ and $\rho^{-1}(2) \cap V_k$ are infinite
2. $\rho(u) = 1$ for all $u \in V \setminus V_k$.

Then one can check that ρ is a run of $\mathcal{A}$ on $\mathcal{D}$ such that $\mathit{final}^0_{(\mathcal{D},\rho)}[k] = \{1, 2\}$. Hence ρ is accepting.

Next we show the statement of the lemma for $\mathcal{B}$. First suppose that V_k is finite. The define $\rho : V \to Q$ by

$$\rho(u) = \begin{cases} 3 & \text{if } u \geq v \text{ for all } v \in V_k \\ 2 & \text{otherwise} \end{cases}$$

As above, this is a run of $\mathcal{B}$ on $\mathcal{D}$. It satisfies $\mathit{final}^0_{(\mathcal{D},\rho)}[k] = \{3\}$ if $V_k \neq \emptyset$ and $\mathit{final}^0_{(\mathcal{D},\rho)}[k] = \{0\}$ otherwise. Thus, ρ is accepting. Conversely let ρ be an accepting run of $\mathcal{B}$ on $\mathcal{D}$ and suppose, towards a contradiction, that V_k is infinite. Then $\rho(u) \geq 1$ for all $u \in V_k$. Since ρ is accepting, there therefore exists a node $u \in V_k$ with $\rho(u) = 3$. But then, for any node $v \in V$ with $u \leq v$, we have $\rho(v) = 3$. Since V_k is infinite, there exists $w \in V_k$ with $u < w$. Hence there exists $u \leq v E_\ell w$ for some $\ell \in C$ with $\rho(v) = 3$. Hence the transition taken at w does not belong to Δ, i.e., ρ is not a run. □

The following can be shown similarly to Proposition 4.1.

Proposition 4.7. *Let* $\mathcal{A}_1$ *and* $\mathcal{A}_2$ *be Muller automata over* $(\widetilde{\Sigma}, C)$. *There are Muller automata* $\mathcal{A}$ *and* $\mathcal{B}$ *such that* $L(\mathcal{A}) = L(\mathcal{A}_1) \cup L(\mathcal{A}_2)$ *and* $L(\mathcal{B}) = L(\mathcal{A}_1) \cap L(\mathcal{A}_2)$.

Theorem 4.8. *Let L be a set of $(\widetilde{\Sigma}, C)$-dags. Then the following are equivalent*

1. *there exists a generalized Muller ACAT $\mathcal{A}$ such that $L = L(\mathcal{A})$*
2. *there exists a Muller ACAT $\mathcal{A}'$ such that $L = L(\mathcal{A}')$*

Proof. Suppose $\mathcal{A} = (Q, \iota, \Delta, T, \mathcal{F})$ is a Muller ACAT. Denote by $\pi_1 : 2^Q \times \{\infty, \overline{\infty}\} \to 2^Q$ the projection onto the first component. Then let $\mathcal{F}'$ comprise all mappings $\overline{q} : Ag \to 2^Q \times \{\infty, \overline{\infty}\}$ with $\pi_1 \circ \overline{q} \in \mathcal{F}$. This defines a generalized Muller ACAT $\mathcal{A}' = (Q, \iota, \Delta, T, \mathcal{F}')$ that certainly accepts the same language as $\mathcal{A}$ does.

For the other implication, let $\mathcal{A}' = (Q, \iota, \Delta, T, \mathcal{F}')$ be some generalized Muller ACAT. For $\overline{q} \in \mathcal{F}'$, consider the Muller ACAT $\mathcal{A}_{\overline{q}} = (Q, \iota, \Delta, \{\pi_1 \circ \overline{q}\})$. Then the language of the generalized Muller ACAT $\mathcal{A}' = (Q, \iota, \Delta, T, \{\overline{q}\})$ is an intersection of $L(\mathcal{A}_{\overline{q}})$ with some sets of the form $\{\mathcal{D} \mid V_k \text{ is infinite}\}$ and $\{\mathcal{D} \mid V_k \text{ is finite}\}$. Since they can be accepted by a Muller ACAT, the result follows from Proposition 4.7. □

4.2 Message Sequence Charts and Message-Passing Automata

If distributed components synchronize by exchanging messages rather than by executing actions simultaneously, we deal with a channel system. The behavior of such a system can be described by a set of *message sequence charts* (MSCs), a subclass of dags, which is defined over a suitable distributed alphabet. We fix a nonempty finite set $\mathcal{M}$ of *messages*. The messages from $\mathcal{M}$ might be exchanged between the agents by performing send and receive actions. Accordingly, for an agent $i \in Ag$, we set Σ_i to be $\{i!^m j, i?^m j \mid j \in Ag \setminus \{i\},\ m \in \mathcal{M}\}$, the set of *(communication) actions* of agent i. The action $i!^m j$ is to be read as "i sends the message m to j", while $j?^m i$ is the complementary action of receiving a message m sent from i to j. Let Σ stand for the union of the (disjoint) sets Σ_i and set $\widetilde{\Sigma}$ to be the distributed alphabet $(\Sigma_i)_{i \in Ag}$. Furthermore, let $C = Ag \uplus \{\mathrm{m}\}$. A *message sequence chart over Ag* (or MSC) is a $(\widetilde{\Sigma}, C)$-dag $(V, \{E_\ell\}_{\ell \in C}, \lambda)$ such that

1. for any $i \in Ag$, E_i is the cover relation of $\leq_i = \leq \cap (V_i \times V_i)$,
2. for any $(u, v) \in E_{\mathrm{m}}$, there exist $m \in \mathcal{M}$ and $i, j \in Ag$ distinct with $\lambda(u) = i!^m j$ and $\lambda(v) = j?^m i$, and
3. for any $u \in V$, there is $v \in V$ satisfying $(u, v) \in E_{\mathrm{m}}$ or $(v, u) \in E_{\mathrm{m}}$.

Let $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda)$ be some MSC. Suppose $m, m' \in \mathcal{M}$ and, moreover, $(u, v), (u', v') \in E_{\mathrm{m}}$ with $\lambda(u) = i!^m j$, $\lambda(v) = j?^m i$, $\lambda(u') = i!^{m'} j$, and $\lambda(v') = j?^{m'} i$. Assume $m = m'$. Since $u, u' \in V_i$, u and u' are ordered by $\leq$. We may assume $u \leq u'$. Then, by the very definition of a $(\widetilde{\Sigma}, C)$-dag, we have $v \leq v'$. Hence the messages are exchanged in a fifo manner, i.e., messages cannot overtake other messages of the same type. If, otherwise, $m \neq m'$, we may have the situation that both $u < u'$ and $v' < v$, i.e., messages sent from i to j of distinct type m and m' may overtake each other. Thus, we may deal with a non-fifo or a fifo setting, depending on whether we consider automata and logics relative to the class of all MSCs or relative to the class of pure fifo MSCs, respectively. In particular, the following results hold in both the pure fifo and the non-fifo setting.

In [11], Muller message-passing automata were considered. In our setting, they are best defined as follows:

Definition 4.9. *A* message-passing automaton *(MPA, for short) is an ACAT* (Q, ι, Δ, T) *such that* $T(a, q) = \emptyset$ *for all* $q \in Q$ *and* $a \in \Sigma$*. A* (generalized) Muller MPA *is a (generalized) Muller ACAT* $(Q, \iota, \Delta, T, \mathcal{F})$ *such that* (Q, ι, Δ, T) *is an MPA.*

Thus, differently from an ACAT, an MPA cannot make communication requests. For $(\widetilde{\Sigma}, C)$-dags, this assumption severely restricts the expressive power (see [1]). We will show that, in the realm of MSCs, we can safely require it.

Proposition 4.10. *Let* $\mathcal{A} = (Q, \iota, \Delta, T, \mathcal{F})$ *be some generalized Muller ACAT. Then there exists a generalized Muller MPA* $\mathcal{A}' = (Q', \iota', \Delta', T', \mathcal{F}')$ *that accepts the same set of MSCs.*

Proof. Let $Q' = Q \times (\Sigma \cup \{0, 1\})$. The idea behind this is as follows: First, the first component of the elements of Q will be used to simulate the automaton $\mathcal{A}$. Now suppose, in an accepting run of $\mathcal{A}'$, the node u labeled in Σ_i assumes the state (q, a). Then this shall mean

- that u is the last event of agent i if $a = 1$
- that the next action of agent i is a if $a \in \Sigma$ (in particular, the node u will never assume a state of the form (q, a) with $a \in \Sigma_j$ for some $j \neq i$)
- no state $(q, 0)$ will ever be assumed (0 serves as label for the initial state).

Having this idea in mind, we define the transition relation Δ' as follows: let $(q_{(c,\ell)}, b_{(c,\ell)}) \in Q'$ for $c \in \Sigma$ and $\ell \in C$, let $k \in Ag$, $a \in \Sigma_k$, and $(q, b) \in Q'$. Then set $((q_{(c,\ell)}, b_{(c,\ell)})_{(c,\ell) \in \Sigma \times C}, a, (q, b)) \in \Delta'$ iff

- $b_{(c,k)} \in \{0, a\}$ for all $c \in \Sigma_k$
- $((q_{(c,\ell)})_{(c,\ell) \in \Sigma \times C}, a, q) \in \Delta$,
- $T(a, q) \subseteq \{(\bar{a}, \mathrm{m})\} \cup (\Sigma_k \times \{k\})$ where $\bar{a}$ and a are complementary communication actions such that
  - $(\bar{a}, \mathrm{m}) \in T(a, q)$ implies that a is a send action
  - for all $c \in \Sigma_k$, $(c, k) \in T(a, q)$ implies $b = c$
- $b \neq 0$

Let $\pi_1 : Q' \to Q$ be the projection to the first component and define $\pi_1 : 2^{Q'} \to 2^Q$ by $\pi_1(M) = \{q \mid (q, a) \in M\}$. Then we define the acceptance condition $\mathcal{F}'$ as the set of mappings $\overline{q}' : Ag \to 2^{Q'} \times \{\overline{\infty}, \infty\}$ with $\overline{q}'[i] = (M_i, \theta_i)$ such that

- there exists $\overline{q} \in \mathcal{F}$ with $\overline{q}[i] = (\pi_1(M_i), \theta_i)$ for all $i \in Ag$,
- for all $i \in Ag$, $M_i \subseteq Q \times \{1\}$ if $\theta_i = \overline{\infty}$, and
- for all $i \in Ag$, $M_i \cap (Q \times \{1\}) = \emptyset$ if $\theta_i = \infty$.

Now consider the generalized Muller MPA $\mathcal{A}' = (Q', \iota', \Delta', T', \mathcal{F}')$ with $\iota' = (\iota, 0)$ and $T(a,q) = \emptyset$ for all $q \in Q'$ and $a \in \Sigma$.

It remains to be shown that these two generalized Muller ACATs accept the same set of MSCs.

(1) Suppose ρ' is some accepting run of $\mathcal{A}'$ on the MSC $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda)$. Then let $\rho = \pi_1 \circ \rho' : V \to Q$. We show that ρ is an accepting run of $\mathcal{A}$ on $\mathcal{D}$. First note that any transition taken by ρ belongs to Δ since this was the case for ρ' and the transitions in Δ' are "extensions" of transitions in Δ. What is more involved is the proof that any communication request is fulfilled. To this aim, let $u \in V$ be some node labeled by $a \in \Sigma_k$ and suppose $\rho'(u) = (q, b)$. If $b = 1$, then $T(a,q) = \emptyset$ or $T(a,q) = \{(\bar{a}, \mathrm{m})\}$ depending on whether a is a receive action or not. Since any send has a matching receive, all communication requests of (q,a) are satisfied. Now suppose $b \neq 1$ which implies (by the definition of transitions) $b \in \Sigma_k$. Then $T(a,q) = \{(b,k)\}$ or $T(a,q) = \{(b,k), (\bar{a}, \mathrm{m})\}$ depending on whether a is a receive action or not. If u was maximal in V_k, then $\mathit{efinal}^{\iota'}_{(\mathcal{D},\rho')} = (\{(q,b)\}, \overline{\infty})$ which is impossible because of $b \neq 1$. Hence there exists $v \in V_k$ with $(u,v) \in E_k$. Since the transition taken at v is legal, we obtain $\lambda(v) = b$. Thus, all communication requests of (q,a) are satisfied. Hence, ρ is indeed a run of $\mathcal{A}$ on $\mathcal{D}$. It is immediate to show that it is accepting by the first requirement in the definition of $\mathcal{F}'$.

(2) Now let ρ be some accepting run of $\mathcal{A}$ on the MSC $\mathcal{D} = (V, \{E_\ell\}_{\ell \in C}, \lambda)$. For $v \in V_i$, define $\rho'(v)$ as follows

- if v is the maximal element of V_i, then set $\rho'(v) = (\rho(v), 1)$
- otherwise there exists $w \in V_i$ with $(v,w) \in E_i$. Then set $\rho'(v) = (\rho(v), \lambda(w))$.

Let $((q_{(c,\ell)}, b_{(c,\ell)})_{(c,\ell) \in \Sigma \times C}, a, (q,b))$ be the transition taken at v with $v \in V_k$. We have to show that it is legal, i.e., satisfies the four stipulations above. To verify the first, let $c \in \Sigma_k$. If there is $u \in V_k$ with $\lambda(u) = c$ and $(u,v) \in E_k$, then $b_{(c,k)} = \lambda(v) = a$ by the definition of ρ'. If there is no such u, then $b_{(c,k)} = 0$ which ensures the first condition in the definition of Δ'. The second one is immediate since the first entry in $\rho'(u)$ equals $\rho(u)$ for any $u \in V$ and since ρ is a run of $\mathcal{A}$. Since the fourth condition is immediate by the definition of ρ', it remains to verify the third one. Since ρ is a run of $\mathcal{A}$, we have $T(a,q) \subseteq \mathrm{Write}(v)$. Since $\mathcal{D}$ is an MSC, the write domain $\mathrm{Write}(v)$ is contained in $\{(\bar{a}, \mathrm{m})\} \cup (\Sigma_k \times \{k\})$. Now suppose $(\bar{a}, \mathrm{m}) \in T(a,q)$. Then a is a send event since $(\bar{a}, \mathrm{m})$ belongs to the write domain of v. Next let $(c,k) \in T(a,q)$. Since it also belongs to the write domain of v, there is $w \in V$ with $(v,w) \in E_k$ and $\lambda(w) = c$. But this implies $b = c$ by the definition of ρ'. Thus, we showed that ρ' is a run of $\mathcal{A}'$. Finally, it is not hard to believe that it is accepting. □

Theorem 4.11. *Let L be a set of MSCs. Then the following are equivalent*

(1) there exists a generalized Muller MPA that accepts L relative to the set of all MSCs.
(2) there exists a Muller MPA that accepts L relative to the set of all MSCs.
(3) there exists a sentence $\varphi \in \mathrm{EMSO}^\infty$ such that L comprises precisely those MSCs that satisfy φ

(4) there exists a generalized Muller ACAT that accepts L relative to the set of all MSCs.

Proof. Note that in the proofs of Lemma 4.6 and Theorem 4.8, we did not introduce nontrivial communication requests. Hence the implication (1)⇒(2) can be shown in the same way. The implications (2)⇒(3)⇒(4) are immediate consequences of Theorems 4.3 and 4.8. Finally, the implication (4)⇒(1) is the above proposition. □

References

1. B. Bollig. On the expressiveness of asynchronous cellular automata. In *Proceedings of the 15th International Symposium on Fundamentals of Computation Theory (FCT 2005)*, volume 3623 of *Lecture Notes in Computer Science*, pages 528–539. Springer, 2005.
2. B. Bollig and M. Leucker. Message-passing automata are expressively equivalent to EMSO logic. *Theoretical Computer Science*, to appear, 2005. A preliminary version appeared in *Proceedings of CONCUR 2004*, volume 3170 of Lecture Notes in Computer Science, pages 146–160. Springer, 2004.
3. J. Büchi. Weak second order logic and finite automata. *Z. Math. Logik, Grundlag. Math.*, 5:66–62, 1960.
4. M. Droste, P. Gastin, and D. Kuske. Asynchronous cellular automata for pomsets. *Theoretical Computer Science*, 247(1-2):1–38, 2000.
5. C. C. Elgot. Decision problems of finite automata design and related arithmetics. *Trans. Amer. Math. Soc.*, 98:21–52, 1961.
6. P. Gastin and A. Petit. Asynchronous cellular automata for infinite traces. In *Proceedings of the 19th International Colloquium on Automata, Languages and Programming (ICALP 1992)*, volume 623 of *Lecture Notes in Computer Science*, pages 583–594. Springer, 1992.
7. B. Genest, D. Kuske, and A. Muscholl. A Kleene theorem for a class of communicating automata with effective algorithms. In *Proceedings of the 8th International Conference on Developments in Language Theory (DLT 2004)*, volume 3340 of *Lecture Notes in Computer Science*, pages 30–48. Springer, 2004.
8. 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.
9. 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.
10. D. Kuske. Emptiness is decidable for asynchronous cellular machines. In *Proceedings of the 11th International Conference on Concurrency Theory (CONCUR 2000)*, volume 1877 of *Lecture Notes in Computer Science*, pages 536–551. Springer, 2000.
11. D. Kuske. Regular sets of infinite message sequence charts. *Information and Computation*, 187:80–109, 2003.
12. L. Libkin. *Elements of Finite Model Theory*. Springer, Berlin, 2004.
13. P. Madhusudan. Reasoning about sequential and branching behaviours of message sequence graphs. In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001)*, volume 2076 of *Lecture Notes in Computer Science*, pages 809–820. Springer, 2001.

14. A. Muscholl. On the complementation of asynchronous cellular büchi automata. *Theoretical Computer Science*, 169(2):123–145, 1996.
15. W. Thomas. Automata on infinite objects. In *Handbook of theoretical computer science (vol. B): formal models and semantics*, pages 133–191. MIT Press, 1990.
16. W. Thomas. On logical definability of trace languages. In *Proceedings of a 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, Technical University of Munich, pages 172–182, 1990.
17. W. Thomas. Languages, automata and logic. In A. Salomaa and G. Rozenberg, editors, *Handbook of Formal Languages*, volume 3, Beyond Words, pages 389–455. Springer, Berlin, 1997.
18. W. Zielonka. Notes on finite asynchronous automata. *R.A.I.R.O. — Informatique Théorique et Applications*, 21:99–135, 1987.

Appendix

Proof details for Example 2.1

First note that $\exists^\infty x : \varphi$ is equivalent to $\forall y \exists x (y \leq x \wedge \varphi)$. This allows to eliminate the infinity quantifier from all logics over the signature $\sigma_\leq$. In particular, we get $\mathbb{L}(\mathrm{FO}(\sigma_\leq)) = \mathbb{L}(\mathrm{FO}^\infty(\sigma_\leq))$ and $\mathbb{L}(\mathrm{wEMSO}(\sigma_\leq)) = \mathbb{L}(\mathrm{wEMSO}^\infty(\sigma_\leq))$. Furthermore, $x \lessdot z$ is equivalent to $x < z \wedge \forall y (x \leq y < z \to x = y)$. To establish $\mathbb{L}(\mathrm{wEMSO}(\sigma_\lessdot)) \subseteq \mathbb{L}(\mathrm{EMSO}(\sigma_\lessdot))$, let $\varphi = \exists X_1 \exists X_2 \ldots \exists X_n : \psi$ be a formula from $\mathrm{wEMSO}(\sigma_\lessdot)$ where ψ does not contain any set quantifiers. In φ, replace ψ with

$$\exists Y : (Y \neq \emptyset \wedge \forall x, y (x \in Y \wedge x \lessdot y \to y \in Y) \wedge \bigwedge_{1 \leq i \leq n} Y \cap X_i = \emptyset \wedge \psi)$$

and consider the resulting formula φ' as $\mathrm{EMSO}(\sigma_\lessdot)$-formula. The requirement on Y to be an upwards closed nonempty set of natural numbers that does not intersect any of the sets X_i forces these sets to be finite. Hence a word satisfies the $\mathrm{wEMSO}(\sigma_\lessdot)$-formula φ iff it satisfies the $\mathrm{EMSO}(\sigma_\lessdot)$-formula φ'. From this, the remaining inclusions follow directly form the definitions since the logics involved are fragments of each other. It therefore remains to show strictness and incomparability in the hierarchy. Consider Fig. 2: for any dashed edge from a class $\mathcal{C}_1$ to a class $\mathcal{C}_2$, we will provide a language contained in $\mathcal{C}_1$ but not in $\mathcal{C}_2$ (for an edge labeled n, see Claim n). It is easy to see that this suffices to establish the required strictness and incomparability result. Let $\Sigma = \{a, b\}$.

Claim 1 $L_1 = \{\alpha b a^\omega \mid \alpha \in \Sigma^* \text{ has odd length}\} \in \mathbb{L}(\mathrm{wEMSO}(\sigma_\lessdot)) \setminus \mathbb{L}(\mathrm{FO}(\sigma_\leq))$.

Proof. A corresponding $\mathrm{wEMSO}(\sigma_\lessdot)$-formula requests finite sets X and Y such that, up to some b-labeled position $y \in Y$, X contains any odd and Y contains any even position. Moreover, if a position is labeled by a b, then it must be contained in $X \cup Y$. On the other hand, given $k \in \mathbb{N}$, the duplicator has a k-round winning strategy in the standard Ehrenfeucht-Fraïssé game on a^{2^k} and a^{2^k+1} (both considered to be equipped with the relation $\leq$, see [17, Example 4.3]).

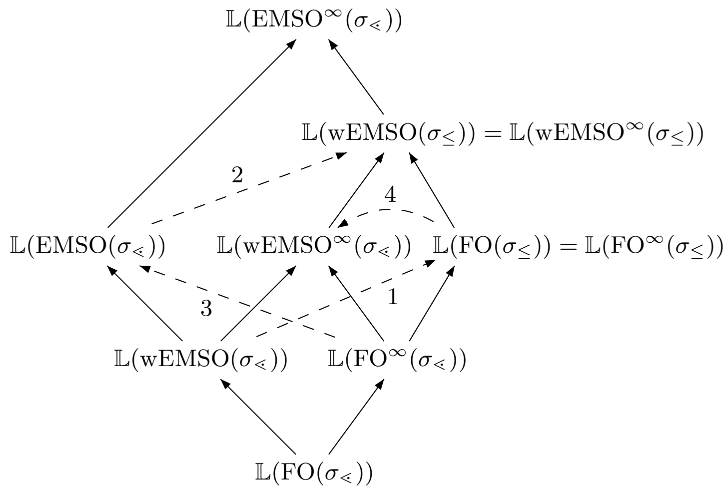

Fig. 2. Strictness and incomparability in the hierarchy of logics over infinite words

Duplicator has a trivial winning strategy on ba^ω and ba^ω. These two winning strategies can be combined in the obvious way to yield a winning strategy on $a^{2^k}ba^\omega$ and $a^{2^k+1}ba^\omega$ so that $a^{2^k}ba^\omega \equiv_k a^{2^k+1}ba^\omega$. Hence L_1 is not in $\mathbb{L}(\mathrm{FO}(\sigma_\leq))$ by Theorem 2.2. $\square$

Let $\alpha = a_1a_2a_3\ldots \in \Sigma^\omega$ be a word and $X_1, X_2, \ldots, X_n \subseteq \mathbb{N}$ sets of positions. We encode all this into a single word $\beta = b_1b_2b_3\ldots$ over $\Sigma \times \{0,1\}^n$ in the standard way: $b_i = (a_i, x_i^1, x_i^2, \ldots, x_i^n)$ with $x_i^j = 1$ iff $i \in X_j$. To simplify notations further, for two words u and v of the same length, we write $\binom{u}{v}$ to denote the word whose ith letter is the pair consisting of the ith letters of u and v.

Claim 2 $L_2 = \{a_1a_2a_3\ldots \in \Sigma^\omega \mid \forall i : a_i = b \to i \text{ is odd}\} \in \mathbb{L}(\mathrm{EMSO}(\sigma_\lessdot)) \setminus \mathbb{L}(\mathrm{wEMSO}(\sigma_\leq))$.

Proof. Containment in $\mathbb{L}(\mathrm{EMSO}(\sigma_\lessdot))$ is shown similarly to the first claim: a corresponding formula requests a set of positions X containing the first position and then every second position so that X comprises all b-labeled positions. On the other hand, let $n, k \geq 1$ and suppose there is a sentence $\varphi = \exists X_1 \exists X_2 \ldots \exists X_n : \psi \in \mathrm{wEMSO}(\sigma_\leq)$ with $\psi \in \mathrm{FO}(\sigma_\leq)[k]$ such that $L(\varphi) = L_2$. Clearly, $(a^{2^k}b)^\omega \in L_2$ and therefore $(a^{2^k}b) \models \varphi$. Hence, there are $m \in \mathbb{N}$ and $\bar{v}_1, \ldots, \bar{v}_m \in (\{0,1\}^n)^{2^k+1}$ such that

$$\alpha = \binom{a^{2^k}b}{\bar{v}_1}\binom{a^{2^k}b}{\bar{v}_2}\cdots\binom{a^{2^k}b}{\bar{v}_m}\binom{a^{2^k}b}{\bar{0}}\binom{a^{2^k}b}{\bar{0}}^\omega \models \psi$$

where $\bar{0} \in (0,0,\ldots,0)^\infty$ stands for the word of appropriate length. Now set

$$\alpha' = \binom{a^{2^k}b}{\bar{v}_1}\binom{a^{2^k}b}{\bar{v}_2}\cdots\binom{a^{2^k}b}{\bar{v}_m}\binom{a^{2^k+1}b}{\bar{0}}\binom{a^{2^k}b}{\bar{0}}^\omega$$

and suppose that both α and α' are equipped with the predicate $\leq$. Note that the first m blocks as well as the final blocks of α and α' are identical. Furthermore, duplicator has a winning strategy for the standard Ehrenfeucht-Fraïssé game played on blocks number $m+1$ in α and α' (see [17, Example 4.3]). Since the remaining blocks are identical, duplicator can extend this strategy to a winning strategy on α and α'. Hence, by Theorem 2.2, $\alpha' \models \psi$. Therefore, its first row satisfies φ, i.e., belongs to $L(\varphi) \setminus L_2$, a contradiction. □

Claim 3 $L_3 = (\Sigma^* b)^\omega \in \mathbb{L}(\mathrm{FO}^\infty(\sigma_\lessdot)) \setminus \mathbb{L}(\mathrm{EMSO}(\sigma_\lessdot))$

Proof. Obviously, L_3 is $\mathrm{FO}^\infty(\sigma_\lessdot)$-definable. Again, let $n, k \geq 1$ and suppose there is a sentence $\varphi = \exists X_1 \exists X_2 \ldots \exists X_n : \psi \in \mathrm{EMSO}(\sigma_\lessdot)$ with first-order kernel $\psi \in \mathrm{FO}(\sigma_\lessdot)[k]$ satisfying $L(\varphi) = L_3$. There are $r, t \geq 0$ such that, for any two words $\alpha, \alpha' \in (\Sigma \times \{0,1\}^n)^\omega$, $\alpha \leftrightarrows_{r,t} \alpha'$ implies $\alpha \equiv_k \alpha'$ (Theorem 2.5). Moreover, for any $N \in \mathbb{N}$, there are $\bar{v}_1, \bar{v}_2, \ldots \in (\{0,1\}^n)^{N+1}$ such that

$$\alpha = \binom{a^N b}{\bar{v}_1}\binom{a^N b}{\bar{v}_2}\cdots \models \psi \ .$$

Now let $m, N_1, N_2, N_3 > 0$ with $N = N_1 + N_2 + N_3$. Then there are words $\bar{u}_i \in (\{0,1\}^n)^*$ of length N_i such that $\bar{v}_{m+1} = \bar{u}_1\bar{u}_2\bar{u}_3$. Set

$$\alpha' = \binom{a^N b}{\bar{v}_1}\binom{a^N b}{\bar{v}_2}\cdots\binom{a^N b}{\bar{v}_m}\binom{a^{N_1}}{\bar{u}_1}\binom{a^{N_2}}{\bar{u}_2}^\omega \ .$$

The numbers N, N_1, N_2, N_3, and m can be chosen such that $\alpha \leftrightarrows_{r,t} \alpha'$ and therefore $\alpha' \models \psi$. Hence its first row belongs to $L(\varphi) \setminus L_3$, a contradiction. □

Claim 4 $L_4 = (a^* b a^* c)^\omega \in \mathbb{L}(\mathrm{FO}(\sigma_\leq)) \setminus \mathbb{L}(\mathrm{wEMSO}^\infty(\sigma_\lessdot))$.

Proof. Clearly, L_4 is in $\mathbb{L}(\mathrm{FO}(\Sigma_\leq))$. Assume it is contained in $\mathbb{L}(\mathrm{wEMSO}^\infty(\sigma_\lessdot))$ as well. Let $n, k \geq 1$ and suppose there is a sentence $\varphi = \exists X_1 \ldots \exists X_n : \psi \in \mathrm{wEMSO}^\infty(\sigma_\lessdot)$ with $\psi \in \mathrm{FO}^\infty(\sigma_\lessdot)[k]$ satisfying $L(\varphi) = L_4$. There are $r, t \geq 0$ such that, for any two words $\alpha, \alpha' \in (\Sigma \times \{0,1\}^n)^\omega$, $\alpha \leftrightarrows^\infty_{r,t} \alpha'$ implies $\alpha \equiv_k \alpha'$ (Theorem 2.7). Moreover, for any natural number $N \in \mathbb{N}$, there are $m \in \mathbb{N}$ and $\bar{v}_1, \ldots, \bar{v}_m \in (\{0,1\}^n)^{N+1}$ such that

$$\alpha = \binom{a^N b}{\bar{v}_1}\binom{a^N c}{\bar{v}_2}\cdots\binom{a^N b}{\bar{v}_{2m-1}}\binom{a^N c}{\bar{v}_{2m}}\left[\binom{a^N b}{\bar{0}}\binom{a^N c}{\bar{0}}\right]^\omega \models \psi \ .$$

Now let

$$\alpha' = \binom{a^N b}{\bar{v}_1}\binom{a^N c}{\bar{v}_2}\cdots\binom{a^N b}{\bar{v}_{2m-1}}\binom{a^N c}{\bar{v}_{2m}}\binom{a^N c}{\bar{0}}\left[\binom{a^N b}{\bar{0}}\binom{a^N c}{\bar{0}}\right]^\omega \models \psi \ .$$

If N is large enough, then we obtain $\alpha \leftrightarrows^\infty_{r,t} \alpha'$ and therefore $\alpha' \models \psi$. Hence the first row of α' belongs to $L(\varphi) \setminus L_4$, a contradiction. □