Asynchronous cellular automata for pomsets without auto-concurrency*

Manfred Droste
Institut für Algebra
Technische Universität Dresden
D-01062 Dresden
droste@math.tu-dresden.de

Paul Gastin
LITP, IBP
Université Paris 7
2, place Jussieu
F-75251 Paris Cedex 05
Paul.Gastin@litp.ibp.fr

Abstract

This paper extends to pomsets without auto-concurrency the fundamental notion of asynchronous cellular automata (ACA) which was originally introduced for traces by Zielonka. We generalize to pomsets the notion of asynchronous mapping introduced by Zielonka and we show how to construct a deterministic ACA from an asynchronous mapping. Our main result generalizes Büchi's theorem for a class of pomsets without auto-concurrency which satisfy a natural axiom. This axiom ensures that an asynchronous cellular automaton works on the pomset as a concurrent read owner write machine. More precisely, we prove the equivalence between non deterministic ACA, deterministic ACA and monadic second order logic for this class of pomsets.

1 Introduction

In a distributed system, some events may occur concurrently, meaning that they may occur in any order or simultaneously or even that their executions may overlap. This is the case in particular when two events use independent resources. On the other hand, some events may causally depend on each other. For instance, the receiving of a message must follow its sending. Therefore, a distributed behavior may be abstracted as a pomset, that is a set

*This research was partly carried out during a stay of the second author in Dresden.

of events together with a partial order which describes causal dependencies of events and with a labeling function. In this paper, we mainly deal with pomsets without auto-concurrency : concurrent events must have different labels. These pomsets are called semi-words in [12, 3]. For studies how general pomsets can be used to represent parallel processes and how they can be composed, we refer the reader e.g. to [11, 7].

There are several ways to describe the behaviors of a system. For instance, logic formulas are suited for specification purposes. Depending on the properties we have to express, we can use various logics such as the temporal logic, the first order logic or the (monadic) second order logic. On the other hand, transition systems are often used to give more operational descriptions. On this paper, we will concentrate on these two kinds of descriptions of systems.

When dealing with distributed systems, it is natural to look for transition systems which faithfully reflect the concurrency. For instance, Petri nets are a widely studied class of such transition systems. Asynchronous cellular automata (ACA) form another fundamental class of transition systems with built-in concurrency. They were introduced for traces by Zielonka [14, 15]. Mazurkiewicz introduced traces in order to describe the behaviors of one-safe Petri nets [8, 9]. A trace is a pomset where the partial order is dictated by a static dependence relation over the actions of the system.

The primary aim of this work is to generalize the notion of ACA so that they can work on pomsets without auto-concurrency. There are several possible definitions for such ACA. In Section 3, we give a natural definition of ACA which intuitively run over the Hasse diagrams of pomsets. We also discuss possible alternative definitions.

Asynchronous mappings prove to be a basic tool to construct ACA for traces [15, 2]. In Section 4, we give a definition of asynchronous mappings for general pomsets. We prove that a pomset language recognized by an asynchronous mapping can be accepted by a *deterministic* ACA.

The rest of this paper is devoted to the equivalence between ACA and monadic second order (MSO) logic for pomsets. We prove in Section 5 that from a (non deterministic) ACA one can construct a MSO formula which defines precisely the pomset language accepted by the automaton. We do not know whether the converse holds in general. In Section 6, we prove the converse for the special subclass of pomsets for which the ACA works as a concurrent read owner write (CROW) machine. These pomsets are called CROW-pomsets. More precisely, from a MSO formula we construct a *deterministic* ACA which accepts precisely the CROW-pomsets defined by the formula. Therefore, for CROW-pomsets, we have the equivalence between MSO logic, deterministic ACA and non deterministic ACA. This result is crucial since it opens the way of model checking for distributed

systems whose behaviors are described as CROW-pomsets.

Note that the CROW assumption for pomsets is satisfied for a large class of systems. For instance, a stably concurrent automaton $\mathcal{A}$ [6] defines dynamic dependencies between actions of a system, whence it generalizes the static dependencies on actions used for traces. With the computation sequences of a stably concurrent automaton $\mathcal{A}$, one can associate dependence orders [1]. These dependence orders are pomsets which generalize traces but are special cases of pomsets without auto-concurrency. Under the assumption that the automaton $\mathcal{A}$ is stably concurrent and forwardly weakly preserves dependency, these dependence orders satisfy the CROW axiom. Therefore, as a consequence of our results, we obtain a new, Zielonka-type characterization of the recognizable languages of concurrency monoids studied in [6].

We do not give all proofs in this extended abstract. They will appear in a full version of this paper.

2 Preliminaries

2.1 Pomsets

Let Σ be an alphabet. A pomset over Σ is (an isomorphism class of) a finite labeled partial order $t = (V, \leq, \lambda)$ where V is a finite set of vertices, $\leq$ is the partial order on V and $\lambda : V \longrightarrow \Sigma$ is the labeling function.

Let $s = (V_s, \leq_s, \lambda_s)$ and $t = (V_t, \leq_t, \lambda_t)$ be two pomsets. We say that s is a prefix of t, denoted by $s \preceq t$ if V_s is a downward closed subset of V_t (i.e. $(u \leq_t v \wedge v \in V_s) \Longrightarrow u \in V_s$ for all $u, v \in V_t$) and $\leq_s$ and λ_s are the restrictions of $\leq_t$ and λ_t to V_s. The prefix order relation is a partial order on the set of pomsets. In the following, we will identify a downward closed subset of vertices with the corresponding prefix of the pomset. Let $s_1 = (V_{s_1}, \leq_{s_1}, \lambda_{s_1})$ and $s_2 = (V_{s_2}, \leq_{s_2}, \lambda_{s_2})$ be two prefixes of a pomset $t = (V_t, \leq_t, \lambda_t)$. Then, $V_{s_1} \cup V_{s_2}$ is downward closed and the corresponding prefix of t will be denoted by $s_1 \cup s_2$.

Let $t = (V, \leq, \lambda)$ be a pomset. The downward closure of a vertex v is denoted by $\downarrow v = \{u \in V \mid u \leq v\}$. The strict downward closure of a vertex v is denoted by $\Downarrow v = \downarrow v \setminus \{v\}$. Since $\downarrow v$ and $\Downarrow v$ are downward closed subsets of V, we will identify these sets to the corresponding prefixes of t.

Let $\Sigma_1, \ldots, \Sigma_n$ be pairwise disjoint alphabets and let $\Sigma = \Sigma_1 \dot{\cup} \cdots \dot{\cup} \Sigma_n$. Intuitively, we can view $[n] = \{1, \ldots, n\}$ as a set of labels of sequential processes and $\Sigma_1, \ldots, \Sigma_n$ as the sets of actions of these sequential processes. Let $p : \Sigma \longrightarrow [n]$ be the mapping which associates with each letter $a \in \Sigma$ the process $p(a) \in [n]$ which executes the letter a, i.e. $a \in \Sigma_{p(a)}$.

Let $t = (V, \leq, \lambda)$ be a pomset. We say that a vertex v covers a vertex u, denoted by $u \lessdot v$, if $u < v$ and there is no vertex w such that $u < w < v$. We say that two vertices $u, v \in V$ are concurrent, denoted by $u \parallel v$, if $u \not\leq v$ and $u \not\geq v$. We may see the covering relation as the description of the interactions between the processes. More precisely, we consider that an event $v \in V$ reads the states of the processes $p \circ \lambda(\{u \mid u \lessdot v\})$ and writes in the process $p \circ \lambda(v)$. We will not allow concurrent writes, therefore two concurrent events $u \parallel v$ must write in different process $p \circ \lambda(u) \neq p \circ \lambda(v)$. This leads to the following definition.

A $(\Sigma_1, \ldots, \Sigma_n)$-pomset is a pomset $t = (V, \leq, \lambda)$ for which $\lambda^{-1}(\Sigma_i)$ is totally ordered for all $1 \leq i \leq n$. The set of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets will be denoted by $\mathbb{P}(\Sigma_1, \ldots, \Sigma_n)$. Note that with this notation the set $\mathbb{P}(\Sigma)$ is the set of words over Σ. Another special case which will come into play later is when the sets $\Sigma_1, \ldots, \Sigma_n$ are all singletons.

For $A \subseteq \Sigma$, we denote by $\partial_A(t) = {\downarrow}\lambda^{-1}(A)$ the least prefix of a pomset t which contains all letters from A. Note that $\partial_A(t) = \bigcup_{v \mid \lambda(v) \in A} {\downarrow}v$. For $a \in \Sigma$ and $i \in [n]$, we will use the following simplified notations : $\partial_a(t) = \partial_{\{a\}}(t)$ and $\partial_i(t) = \partial_{\Sigma_i}(t)$.

Note that if $\lambda^{-1}(A)$ is totally ordered then $\partial_A(t)$ is either empty or has exactly one maximal vertex. In particular, if t is a $(\Sigma_1, \ldots, \Sigma_n)$-pomset then $\partial_i(t)$ is either empty or has exactly one maximal vertex.

2.2 Traces

We recall now basic definitions for Mazurkiewicz traces which will be needed in this paper. The reader is referred to [5] for a general presentation of traces. Our presentation is not the classical one but is more adequate in our context. Basically, a trace is a pomset which satisfies additional requirements. We consider a reflexive and symmetric relation $\Delta \subseteq \Sigma \times \Sigma$ over the alphabet. Intuitively, two dependent actions must be ordered while two independent actions may occur concurrently. More precisely, a trace over a dependence alphabet (Σ, Δ) is a pomset $t = (V, \leq, \lambda)$ such that for all $u, v \in V$

1. $(\lambda(u), \lambda(v)) \in \Delta \Longrightarrow u \not\parallel v$
2. $u \lessdot v \Longrightarrow (\lambda(u), \lambda(v)) \in \Delta$.

The set of traces over (Σ, Δ) will be denoted by $\mathbb{M}(\Sigma, \Delta)$.

We will now define recognizable trace languages. A trace automaton is a quadruple $\mathcal{A} = (Q, T, I, F)$ where Q is a finite set of states, $I \subseteq Q$ is the set of initial states, $F \subseteq Q$ is the set of final states and $T \subseteq Q \times \Sigma \times Q$ is the set of transitions which verifies the diamond property : for all $(a, b) \in \Sigma \times \Sigma \setminus \Delta$

and $q, q', q'' \in Q$, if $(q, a, q') \in T$ and $(q', b, q'') \in T$ then there exists some $\bar{q}' \in Q$ such that $(q, b, \bar{q}') \in T$ and $(\bar{q}', a, q'') \in T$. A word $w = a_1 \cdots a_n \in \Sigma^*$ is accepted by $\mathcal{A}$ if there is a run $q_0, a_1, q_1, \ldots, a_n, q_n$ such that $q_0 \in I$, $q_n \in F$ and $(q_{i-1}, a_i, q_i) \in T$ for all $1 \leq i \leq n$. Note that a linearization of a pomset may be identified with a word of Σ^*. A trace $t \in \mathbb{M}(\Sigma, \Delta)$ is accepted by $\mathcal{A}$ if some linear extension of t is accepted by $\mathcal{A}$. Note that, thanks to the diamond property of a trace automaton, if some linear extension of a trace is accepted by $\mathcal{A}$ then all linear extensions of t are accepted by $\mathcal{A}$. A trace language $L \subseteq \mathbb{M}(\Sigma, \Delta)$ is recognizable if it is the set of traces accepted by some trace automaton.

3 Asynchronous Cellular Automata

Definition 3.1 *A $(\Sigma_1, \ldots, \Sigma_n)$-ACA (asynchronous cellular automata) is a tuple*

$$\mathcal{A} = ((Q_i)_{i \in [n]}, (\delta_{a,J})_{a \in \Sigma, J \subseteq [n]}, F)$$

where

1. *for all $i \in [n]$, Q_i is a finite set of local states for process i,*
2. *for all $a \in \Sigma$ and $J \subseteq [n]$, $\delta_{a,J} : \prod_{i \in J} Q_i \longrightarrow \mathcal{P}(Q_{p(a)})$ is a transition function,*
3. *and $F \subseteq \bigcup_{J \subseteq [n]} \prod_{i \in J} Q_i$ is the set of accepting states.*

The automaton is deterministic if all the transition functions are deterministic, i.e. if $|\delta_{a,J}((q_i)_{i \in J})| \leq 1$ for all $a \in \Sigma$, $J \subseteq [n]$ and $(q_i)_{i \in J} \in \prod_{i \in J} Q_i$.

In order to explain how a $(\Sigma_1, \ldots, \Sigma_n)$-ACA accepts a $(\Sigma_1, \ldots, \Sigma_n)$-pomset, we need to introduce first some new notations. Let $t = (V, \leq, \lambda)$ be a $(\Sigma_1, \ldots, \Sigma_n)$-pomset and let $v \in V$ be a vertex in t. We define the write domain of v by $W(v) = p \circ \lambda(v)$ and the read domain of v as the set of processes of vertices covered by v : $R(v) = p \circ \lambda(\{u \in V \mid u \prec\!\!\!\cdot\, v\})$. A process i is maximal in t if $\lambda(v) \in \Sigma_i$ for some maximal vertex v of t. The set of maximal processes of t is denoted by $\max(t)$.

Intuitively, in order to perform an event v, the ACA reads the states of the processes in $R(v)$ and according to its transition functions determines the new state of the process $W(v)$. At the end of its run, the ACA collects the states of the maximal processes to decide whether the run is accepted or rejected. The formal definition is given below.

Note that if $i \in R(v)$ then $\partial_i(\Downarrow v)$ is non empty, whence has only one maximal vertex which will be identified with $\partial_i(\Downarrow v)$ in Definition 3.2. Similarly, if $i \in \max(t)$ then $\partial_i(t)$ is non empty and will also be identified with its maximum vertex.

Definition 3.2 *A run of a $(\Sigma_1, \ldots, \Sigma_n)$-ACA $\mathcal{A}$ over a $(\Sigma_1, \ldots, \Sigma_n)$-pomset $t = (V, \leq, \lambda)$ is a mapping $r : V \longrightarrow \bigcup_{i\in[n]} Q_i$ such that for all $v \in V$,*

$$r(v) \in \delta_{\lambda(v),R(v)}\left(r(\partial_i(\Downarrow v))_{i\in R(v)}\right)$$

The run r is accepted if its final maximal state $f(r) = r(\partial_i(t))_{i\in\max(t)}$ is in the accepting set F.

A $(\Sigma_1, \ldots, \Sigma_n)$-pomset t is accepted by $\mathcal{A}$ if there is some accepting run of $\mathcal{A}$ on t. The set of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets accepted by $\mathcal{A}$ is denoted by $L(\mathcal{A})$.

The covering of pomsets by the chains formed by the fixed sequential processes is crucial in the definition of asynchronous cellular automata. It allows to use a fixed number of local states and to determine the read and write domains of the actions using the labeling and the covering relation. The weakest covering is when each Σ_i is a singleton. In this case we have a set of local states per letter as in the asynchronous cellular automata for traces [15, 2]. Note that, even with this trivial covering, our definition is not the same as that of Zielonka for traces. Mainly, in our definition, a run of the ACA is over the Hasse diagram of the pomset whereas with Zielonka's ACA for traces, a run is in fact over the dependence graph of the trace. A dependence graph is an intermediary representation of a trace between its Hasse diagram and its partial order. This intermediary representation is possible thanks to the existence of a static dependence relation over actions. More precisely, our definition of ACA for pomsets and that of Zielonka for traces differ in three respects. First, Zielonka's definition uses a global initial state which in our case is coded in the transition functions. Second, our transition functions only read the states of the processes covered by the current action whereas in Zielonka's definition a fixed set of processes is read even if the last executions of some of these processes are far below the current action. Third, we only read the states of maximal processes to determine whether a run is successful whereas in Zielonka's definition the states of all processes are collected globally to decide acceptance.

Remark 3.3 As discussed below, one could give several alternative definitions of accepting runs. One of these variants corresponds precisely to

asynchronous cellular automata for traces. In fact, these differences are not crucial and we will see in Sections 4 and 6 that under some assumptions they are equivalent. In particular, this is true for traces.

1. First, one can change the set of processes read by transition functions. In the definition above, a transition only reads the processes covered by the current vertex. On the contrary, one can allow to read all processes which occur in the past of the current vertex by using $R_{\text{occur}}(v) = p \circ \lambda(\Downarrow v)$ instead of $R(v)$. For traces, we use the trivial covering of Σ by singletons and we identify processes with letters of Σ. Then, we consider a fixed dependence (symmetric and reflexive) relation D_{trace} on Σ and we use $D_{\text{trace}}(v)$ instead of $R(v)$. This static approach enforces the use of an initial state $\perp$. Indeed, for some $a \in D_{\text{trace}}(v)$, the prefix $\partial_a(\Downarrow v)$ may be empty and thus cannot be identified with a vertex of the trace. Thus we define $r(1) = \perp$ where 1 denotes the empty trace.

2. Second, one can change the acceptance condition. In the definition above, we only read the final state of the maximal processes. One can read the final state of all occurring processes by using $f_{\text{occur}}(r) = r(\partial_i(t))_{i \in p \circ \lambda(V)}$ instead of $f(r)$. For traces, we read the final state of all processes by using $f_{\text{trace}}(r) = r(\partial_a(t))_{a \in \Sigma}$ instead of $f(r)$. Again we need the initial state $\perp$ and the convention $r(1) = \perp$.

4 Asynchronous mapping

The domain of an asynchronous mapping must be a prefix closed subset of pomsets, that is a subset $\mathbb{Q}$ of pomsets such that for all pairs of pomsets s and t, if $s \preceq t$ and $t \in \mathbb{Q}$ then $s \in \mathbb{Q}$. For instance, $\mathbb{P}(\Sigma_1, \ldots, \Sigma_n)$ and $\mathbb{M}(\Sigma, \Delta)$ are prefix closed sets of pomsets.

Definition 4.1 *Let $\mathbb{Q}$ be a prefix closed set of pomsets and let S be a finite set. A mapping $\sigma : \mathbb{Q} \longrightarrow S$ is asynchronous if for all $t = (V, \leq, \lambda) \in \mathbb{Q}$,*

1. *for all vertices $v \in V$, the value $\sigma(\downarrow v)$ is uniquely determined by $\sigma(\Downarrow v)$ and $\lambda(v)$.*

2. *for all prefixes t_1, t_2 of t, the value $\sigma(t_1 \cup t_2)$ is uniquely determined by $\sigma(t_1)$ and $\sigma(t_2)$.*

A language $L \subseteq \mathbb{Q}$ is recognized by an asynchronous mapping $\sigma : \mathbb{Q} \longrightarrow S$ if $L = \sigma^{-1}(\sigma(L))$.

Remark 4.2 The definition of asynchronous mappings is valid for arbitrary pomsets and does not depend on the existence of sequential processes $(\Sigma_1, \ldots, \Sigma_n)$ in the pomsets. The definition of asynchronous cellular automata requires at least that each set of vertices labeled with the same letter forms a chain. This simple and weak property is not even required for asynchronous mappings.

Proposition 4.3 *Let $\mathbb{Q} \subseteq \mathbb{P}(\Sigma_1, \ldots, \Sigma_n)$ be a prefix closed set of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets. Let $L \subseteq \mathbb{Q}$ be a language of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets recognized by some asynchronous mapping $\sigma : \mathbb{Q} \longrightarrow S$. Then there exists a deterministic $(\Sigma_1, \ldots, \Sigma_n)$-asynchronous cellular automaton $\mathcal{A}$ such that $L(\mathcal{A}) \cap \mathbb{Q} = L$.*

Proof: (sketch) This proof follows the same ideas as the corresponding one for traces. We only recall here the main steps of this proof. Assume that $\sigma : \mathbb{Q} \longrightarrow S$ recognizes the language $L \subseteq \mathbb{Q}$. We define a deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA $\mathcal{A}$ as follows.

1. For all $i \in [n]$, let $Q_i = S$,
2. For all $a \in \Sigma$, $J \subseteq [n]$ and $(q_i)_{i \in J} \in S^J$, let

$$\begin{aligned} \delta_{a,J}\left((q_i)_{i \in J}\right) \;=\; & \{\sigma(t) \mid t \in \mathbb{Q}, t = {\downarrow}v \text{ for some } v \text{ such that} \\ & \lambda(v) = a, R(v) = J \text{ and } \sigma(\partial_i({\Downarrow}v)) = q_i \forall i \in J\} \end{aligned}$$

3. $F = \{\sigma(\partial_i(t))_{i \in \max(t)} \mid t \in L\}$.

Claim 1: $\mathcal{A}$ is deterministic. (Follows from σ asynchronous).

Claim 2: Let $t = (V, \leq, \lambda) \in \mathbb{Q}$ be a $(\Sigma_1, \ldots, \Sigma_n)$-pomset. Then, the mapping $r : V \longrightarrow \bigcup_{i \in [n]} Q_i$ defined by $r(v) = \sigma({\downarrow}v)$ is the run of $\mathcal{A}$ on t. (Follows from the definition of the transition functions).

Claim 3: $L(\mathcal{A}) \cap \mathbb{Q} = L$ (Follows from σ asynchronous and the definition of F). □

Prop. 4.3 holds also for the alternative definitions of accepting runs discussed in Rem. 3.3. The proofs remain essentially the same. One only has to change accordingly the definitions of the transition functions and of the final set.

Note that, for trace languages, the converse of Proposition 4.3 is also true and all alternative definitions of ACA are equivalent. Indeed, it is easy

to show that a trace language accepted by an ACA is a recognizable trace language, whatever alternative chosen for accepting runs. Moreover, if L is a recognizable trace language, the existence of an asynchronous mapping which recognizes L was proven in [2]. Finally, as mentioned above, from this asynchronous mapping one can easily get an ACA which accepts L, whatever alternative chosen for accepting runs.

The equivalence between alternative definitions of accepting runs will be extended to a more general class of pomsets in Section 6. On the other hand, in the general setting of $\mathbb{P}(\Sigma_1, \ldots, \Sigma_n)$, the converse of Proposition 4.3 is still open.

5 From ACA to MSOL

In this section, we will define monadic second order (MSO) formulas and their interpretations over pomsets. We will then prove that for all ACA $\mathcal{A}$ (deterministic or not), there exists a MSO formula which defines the language accepted by $\mathcal{A}$.

Let Σ be a finite alphabet. The MSO language over Σ that we consider consists of the unary predicates $(P_a)_{a\in\Sigma}$, a binary predicate R, first order variables $x, y, z, \ldots$, monadic second order variables $X, Y, Z, \ldots$, boolean connectives $\neg, \vee, \wedge, \longrightarrow, \longleftrightarrow$ and quantifiers $\exists, \forall$. A sentence is a formula without free variables. For instance, the following formulas are first order and MSO sentences respectively.

$$\begin{aligned}
\varphi_1 &:: \exists x(P_a(x) \wedge \forall y(R(x,y) \longrightarrow \neg P_b(y))) \\
\varphi_2 &:: \exists X \exists Y(\forall x(x \in X \vee x \in Y) \wedge \exists x\ x \in X \wedge \exists y\ y \in Y \\
&\qquad \wedge \forall x \forall y(x \in X \wedge y \in Y \longrightarrow \neg R(x,y) \wedge \neg R(y,x))
\end{aligned}$$

A pomset $t = (V, \leq, \lambda)$ can be seen as an interpretation of this MSO language as follows: the domain is the set V of vertices, i.e. first order variables range over vertices and MSO variables range over sets of vertices; for all $a \in \Sigma$ $P_a(x)$ means $\lambda(x) = a$ and $R(x,y)$ means $x \leq y$. We say that a pomset t satisfies a sentence φ, denoted by $t \models \varphi$, when φ is true for the interpretation defined by t. The set of pomsets which satisfy a sentence φ is denoted by $L(\varphi)$. For instance, $L(\varphi_1)$ is the set of pomsets which have a vertex labeled by a with no vertex labeled by b above and $L(\varphi_2)$ is the set of non connected pomsets.

In order to make the formulas more readable, we will write $x \leq y$ for $R(x,y)$ and $\lambda(x) = a$ for $P_a(x)$. For instance, the formula φ_1 will be written $\exists x(\lambda(x) = a \wedge \forall y(x \leq y \longrightarrow \lambda(y) \neq b))$. Moreover, we will use several

abbreviations which can be easily translated in our MSO language. For instance, we will write

$$\begin{aligned}
x < y \quad &\text{for} \quad x \leq y \wedge \neg y \leq x \\
x \prec y \quad &\text{for} \quad x < y \wedge \neg \exists z (x < z \wedge z < y) \\
\lambda(x) \in A \quad &\text{for} \quad \bigvee_{a \in A} \lambda(x) = a \\
p \circ \lambda(x) = p \circ \lambda(y) \quad &\text{for} \quad \bigvee_{1 \leq i \leq n} (\lambda(x) \in \Sigma_i \wedge \lambda(y) \in \Sigma_i) \\
X \cap Y = \emptyset \quad &\text{for} \quad \neg \exists x (x \in X \wedge x \in Y)
\end{aligned}$$

Note that the language defined by a formula can contain pomsets with auto-concurrency (concurrent vertices with the same label). We do not need to put restrictions on the pomsets defined by a formula because all restrictions we need can be described by MSO formulas. For instance the set $\mathbb{P}(\Sigma_1, \ldots, \Sigma_n)$ of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets is defined by the formula

$$\varphi_{(\Sigma_1,\ldots,\Sigma_n)} :: \forall x \forall y (p \circ \lambda(x) = p \circ \lambda(y) \longrightarrow (x \leq y \vee y \leq x)).$$

We are now ready to state

Theorem 5.1 *Let $\Sigma = \Sigma_1 \dot{\cup} \cdots \dot{\cup} \Sigma_n$ and let $\mathcal{A}$ be a possibly non deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA. There exists a MSO formula φ over Σ such that*

$$L(\varphi) = L(\mathcal{A}).$$

Proof: (sketch)

Let $\mathcal{A} = ((Q_i)_{i \in [n]}, (\delta_{a,J})_{a \in \Sigma, J \subseteq [n]}, F)$ be a $(\Sigma_1, \ldots, \Sigma_n)$-ACA. We will construct a MSO formula which will be satisfied exactly by those $(\Sigma_1, \ldots, \Sigma_n)$-pomsets accepted by $\mathcal{A}$. Let k be the number of states in $\bigcup_{i \in [n]} Q_i$. We may assume that $\bigcup_{i \in [n]} Q_i = [k] = \{1, \ldots, k\}$. The following formula claims the existence of an accepting run of the automaton.

$$\psi :: \exists X_1 \ldots \exists X_k \left(\text{partition}(X_1, \ldots, X_k) \wedge \left(\forall x \ \text{transition}(x) \right) \wedge \text{accepted} \right)$$

We will now explain this formula and give the sub-formulas *partition*, *transition* and *accepted*. A run over a pomset $t = (V, \leq, \lambda)$ is coded by the MSO variables $X_1, \ldots, X_k$. More precisely, X_i stands for the set of vertices mapped on the state i by the run. The formula partition$(X_1, \ldots, X_k)$ makes

sure that the MSO variables $X_1, \ldots, X_k$ do describe a mapping from V to $\bigcup_{i\in[n]} Q_i$.

$$\text{partition}(X_1, \ldots, X_k) :: \left(\forall x \bigvee_{i\in[k]} x \in X_i\right) \wedge \left(\bigwedge_{1\le i\neq j\le k} X_i \cap X_j = \emptyset\right)$$

Then, we have to claim that this labeling of vertices by states agrees with the transition functions of the automaton.

$$\text{transition}(x) :: \bigvee \Bigg(\lambda(x) = a \wedge x \in X_q \wedge \forall y\ (y \prec x \longrightarrow p \circ \lambda(y) \in J) \\ \wedge \bigwedge_{i\in J} \exists y\ (y \prec x \wedge p \circ \lambda(y) = i \wedge y \in X_{q_i})\Bigg)$$

where the disjunction ranges over all letters $a \in \Sigma$, subsets $J \subseteq [n]$, states $q \in Q_{p(a)}$ and tuples $(q_i)_{i\in J} \in \prod_{i\in J} Q_i$ such that $q \in \delta_{a,J}((q_i)_{i\in J})$.

It remains to state that the run reaches a final state of the automaton.

$$\text{accepted} :: \bigvee_{(f_i)_{i\in J}\in F} \Bigg(\forall x \Big((\neg\exists y\ x < y) \longrightarrow p \circ \lambda(x) \in J\Big) \\ \wedge \bigwedge_{i\in J} \exists x \Big((\neg\exists y\ x < y) \wedge p \circ \lambda(x) = i \wedge x \in X_{f_i}\Big)\Bigg)$$

In fact, the formula ψ describes an accepting run of the automaton only for $(\Sigma_1, \ldots, \Sigma_n)$-pomsets. Therefore, we need in addition the formula $\varphi_{(\Sigma_1,\ldots,\Sigma_n)}$ described above. Finally, the theorem follows from the

Claim:

$$L(\varphi_{(\Sigma_1,\ldots,\Sigma_n)} \wedge \psi) = L(\mathcal{A})$$

□

Remark 5.2 Clearly, Theorem 5.1 holds also for the alternative definitions of accepting runs described in Remark 3.3. The proof is essentially the same. One only has to change the formulas *transition* and *accepted* accordingly.

6 From MSOL to deterministic ACA

The converse of Theorem 5.1 is an open problem. In this section, we prove that this converse holds for the special subclass of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets for which the ACA works as a Concurrent Read Owner Write (CROW) machine. More precisely, we consider n processes whose sets of actions are $\Sigma_1, \ldots, \Sigma_n$

respectively. Each process has a memory which can be read by all actions but can be written by its own actions only (Owner Write). We allow concurrent reads of memories but no concurrent writes. As quoted in Section 2.1, this restriction is already enforced by the very definition of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets. Without further restrictions, two concurrent events may respectively read from and write to the same location. This is the case when there exist two concurrent events $y \parallel z$ such that z writes in the memory of some process i ($p \circ \lambda(z) = i$) and y reads the memory of this process i ($p \circ \lambda(x) = i$ for some $x \prec\!\!\!\cdot\, y$). This motivates the following definition.

Definition 6.1 *A $(\Sigma_1, \ldots, \Sigma_n)$-pomset $t = (V, \leq, \lambda)$ satisfies the Concurrent Read Owner Write (CROW) axiom if $\forall x, y, z \in V$,*

$$x \prec\!\!\!\cdot\, y \wedge x < z \wedge y \parallel z \quad \Longrightarrow \quad p \circ \lambda(x) \neq p \circ \lambda(z) \tag{1}$$

The set of $(\Sigma_1, \ldots, \Sigma_n)$-pomsets which satisfy the CROW axiom is denoted by $\mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$.

Theorem 6.2 *Let $\Sigma = \Sigma_1 \dot\cup \cdots \dot\cup \Sigma_n$ and let φ be a MSO formula over Σ. There exists a deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA $\mathcal{A}$ such that*

$$L(\varphi) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) = L(\mathcal{A}) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n).$$

In order to prove this theorem, one can use an induction on the structure of the formula. Disjunction and existential quantification are easily dealt with when non deterministic ACA are allowed. On the other hand, complement is easy for deterministic ACA. Whence the core of such an approach is the determinization of ACA. For this problem, starting from a non deterministic ACA $\mathcal{A}$, one can directly construct an asynchronous mapping which accepts the language $L(\mathcal{A})$ and then use Proposition 4.3. This construction is similar to that of [10] and uses the asynchronous time stamping of Zielonka [15, 2] but the proofs are more involved. In particular, it is known that for traces Zielonka's ν mapping is asynchronous by itself [2, 4] but this is not the case for $\mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$-pomsets. Here we give a simpler proof which uses Zielonka's theorem. For this, we first map CROW-pomsets into traces by simply changing the labeling.

Let $\Sigma' = \Sigma \times \mathcal{P}([n])$ be a new set of labels and for all $i \in [n]$, let $\Sigma'_i = \Sigma_i \times \mathcal{P}([n])$ be the associated new processes. We define an embedding g from $\mathbb{P}(\Sigma_1, \ldots, \Sigma_n)$ into $\mathbb{P}(\Sigma'_1, \ldots, \Sigma'_n)$ by $g(V, \leq, \lambda) = (V, \leq, \lambda')$ where for all $v \in V, \lambda'(v) = (\lambda(v), R(v))$. Note that g is well defined, since for all $i \in [n]$, $\lambda'^{-1}(\Sigma'_i)$ is totally ordered. Let Δ' be the dependence relation defined on Σ' by

$$\Delta' = \{((a, A), (b, B)) \mid p(a) = p(b) \vee p(a) \in B \vee p(b) \in A\}.$$

Proposition 6.3

$$\mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) = g^{-1}(\mathbb{M}(\Sigma', \Delta'))$$

Proof: We first prove that $\mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) \subseteq g^{-1}(\mathbb{M}(\Sigma', \Delta'))$. Let $t = (V, \leq, \lambda) \in \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$ and let $g(t) = (V, \leq, \lambda')$. Let $x, y \in V$ and assume that $x \prec\!\!\!\cdot\, y$. Then, $p \circ \lambda(x) \in R(y)$ and it follows $(\lambda'(x), \lambda'(y)) \in \Delta'$. Now, let $y, z \in V$ and assume that $(\lambda'(y), \lambda'(z)) \in \Delta'$. If $p \circ \lambda(y) = p \circ \lambda(z)$ then $y \not\parallel z$ since t is a $(\Sigma_1, \ldots, \Sigma_n)$-pomset. Otherwise, we have for instance $p \circ \lambda(z) \in R(y)$. Hence, there exists $x \in V$ such that $x \prec\!\!\!\cdot\, y$ and $p \circ \lambda(x) = p \circ \lambda(z)$. Therefore, x and z must be ordered. Since $x < z$ would contradict the CROW-axiom, it follows $z \leq x$, whence $z < y$. Therefore, $g(t) \in \mathbb{M}(\Sigma', \Delta')$ (see Section 2.2) and it follows $t \in g^{-1}(\mathbb{M}(\Sigma', \Delta'))$.

Conversely, let $t = (V, \leq, \lambda) \in g^{-1}(\mathbb{M}(\Sigma', \Delta'))$ and let $g(t) = (V, \leq, \lambda')$. Let $x, y, z \in V$ be such that $x \prec\!\!\!\cdot\, y \wedge x < z \wedge y \parallel z$. By definition, $p \circ \lambda(x) \in R(y)$ and $(\lambda'(y), \lambda'(z)) \notin \Delta'$. Therefore, $p \circ \lambda(z) \notin R(y)$ and it follows $p \circ \lambda(x) \neq p \circ \lambda(z)$. □

Proposition 6.4 *Let φ be a MSO formula over Σ. There exists a MSO formula φ' over Σ' such that*

$$L(\varphi) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) = g^{-1}(L(\varphi') \cap \mathbb{M}(\Sigma', \Delta'))$$

Proof: (sketch) Let φ be a MSO formula over Σ. Let φ' be the MSO formula over Σ' obtained from φ by substituting to atomic formulas of the form $x \in P_a$ the disjunction $\bigvee_{J \subseteq [n]} x \in P_{(a,J)}$:

$$\varphi' = \varphi \left[x \in P_a \;\Big/ \bigvee_{J \subseteq [n]} x \in P_{(a,J)} \right].$$

The proposition follows easily using Proposition 6.3. □

Proposition 6.5 *Let $\mathcal{A}'$ be a (deterministic) $(\Sigma'_1, \ldots, \Sigma'_n)$-ACA. There exists a (deterministic) $(\Sigma_1, \ldots, \Sigma_n)$-ACA $\mathcal{A}$ such that $L(\mathcal{A}) = g^{-1}(L(\mathcal{A}'))$.*

Proof: (sketch) Let $\mathcal{A}' = ((Q_i)_{i \in [n]}, (\delta'_{a',J})_{a' \in \Sigma', J \subseteq [n]}, F)$ be a $(\Sigma'_1, \ldots, \Sigma'_n)$-ACA. For all $a \in \Sigma$ and $J \subseteq [n]$, let $\delta_{a,J} = \delta'_{(a,J),J}$. We claim that $\mathcal{A} = ((Q_i)_{i \in [n]}, (\delta_{a,J})_{a \in \Sigma, J \subseteq [n]}, F)$ is the required $(\Sigma_1, \ldots, \Sigma_n)$-ACA. Note first that if $\mathcal{A}'$ is deterministic then so is $\mathcal{A}$. Now, let $t = (V, \leq, \lambda) \in \mathbb{P}(\Sigma_1, \ldots, \Sigma_n)$ and

let $g(t) = (V, \leq, \lambda')$. For all $v \in V$, we have $R'(v) = p \circ \lambda'(\{u \mid u \prec v\}) = p \circ \lambda(\{u \mid u \prec v\}) = R(v)$. Therefore $\delta'_{\lambda'(v),R'(v)} = \delta'_{(\lambda(v),R(v)),R(v)} = \delta_{\lambda(v),R(v)}$.

Let $r : V \longrightarrow \bigcup_{i\in[n]} Q_i$. It follows that r is an accepting run of $\mathcal{A}'$ on $g(t)$ if and only if r is an accepting run of $\mathcal{A}$ on t, that is,

$$g(L(\mathcal{A})) \subseteq L(\mathcal{A}') \quad \text{and} \quad g^{-1}(L(\mathcal{A}')) \subseteq L(\mathcal{A}).$$

The proposition follows. □

Proof of Theorem 6.2: Let φ be a MSO formula over Σ. By Proposition 6.4, there exists a MSO formula φ' over Σ' such that

$$L(\varphi) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) = g^{-1}(L(\varphi') \cap \mathbb{M}(\Sigma', \Delta')).$$

The language $L(\varphi') \cap \mathbb{M}(\Sigma', \Delta')$ is a recognizable trace language [13]. Hence [15, 2], there exists an asynchronous mapping σ from $\mathbb{M}(\Sigma', \Delta')$ into a finite set which recognizes $L(\varphi') \cap \mathbb{M}(\Sigma', \Delta')$. By Proposition 4.3, there exists a deterministic $(\Sigma'_1, \ldots, \Sigma'_n)$-ACA $\mathcal{A}'$ such that $L(\mathcal{A}') \cap \mathbb{M}(\Sigma', \Delta') = L(\varphi') \cap \mathbb{M}(\Sigma', \Delta')$. It follows by Proposition 6.5 that there exists a deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA $\mathcal{A}$ such that $L(\mathcal{A}) = g^{-1}(L(\mathcal{A}'))$. Finally, applying Proposition 6.3 we obtain

$$\begin{aligned} L(\mathcal{A}) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) &= g^{-1}(L(\mathcal{A}')) \cap g^{-1}(\mathbb{M}(\Sigma', \Delta')) \\ &= g^{-1}(L(\mathcal{A}') \cap \mathbb{M}(\Sigma', \Delta')) \\ &= g^{-1}(L(\varphi') \cap \mathbb{M}(\Sigma', \Delta')) \\ &= L(\varphi) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) \end{aligned}$$

□

As a corollary of Theorems 5.1 and 6.2 we obtain that MSO formulas, non deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA and deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA have the same expressive power for $\mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$-pomsets.

Theorem 6.6 *Let $L \subseteq \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$. The following are equivalent:*

1. *L is definable by a MSO formula,*
2. *there exists a non deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA $\mathcal{A}$ such that*
$$L = L(\mathcal{A}) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$$
3. *there exists a deterministic $(\Sigma_1, \ldots, \Sigma_n)$-ACA $\mathcal{A}$ such that*
$$L = L(\mathcal{A}) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$$

Remark 6.7 Finally, we come back to the alternative definitions of accepting runs discussed in Remark 3.3. As a corollary of previous results, we obtain that these alternatives are equivalent for CROW-pomsets. The question whether the same equivalence holds for general pomsets is still open.

7 Conclusion

We believe that the CROW axiom is really natural if one sees an ACA as an abstract representation of a parallel machine where each process writes in its own memory and can read the memories of other processes. In this case, we have proved that deterministic ACA are closed under the boolean operations and have the same expressive power as the MSOL. Since emptiness for ACA is decidable when restricted to $\mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$ pomsets, we can use ACA to perform model checking : to check whether an implementation (an ACA) $\mathcal{A}$ satisfies a specification (a MSO formula) φ, one computes the deterministic ACA $\mathcal{B}$ such that $L(\varphi) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n) = L(\mathcal{B}) \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$ and then checks for emptiness $L(\mathcal{A}) \cap \overline{L(\mathcal{B})} \cap \mathbb{CROW}(\Sigma_1, \ldots, \Sigma_n)$.

On the other hand, ACA may be seen as an abstract representation of distributed systems communicating asynchronously. In this setting, the covering relation $x \prec\!\!\!\cdot\; y$ can denote a message sent by x to y. If the communication is asynchronous, the process $p(x)$ may perform some actions concurrently to the reception (y) of the message sent (x). In this case, the CROW axiom is not natural anymore. Therefore, an important open problem is to know whether the properties proved in this paper for CROW-pomsets also hold in a more general setting.

Finally, we would like to point out that both the definition of asynchronous mappings and of MSOL can be extended without any change to pomsets with auto-concurrency. It would be very interesting to find a general model of ACA which allows auto-concurrency. Note that it is possible with our ACA to cope with bounded auto-concurrency by suitably relabeling auto-concurrent vertices.

References

[1] F. Bracho, M. Droste, and D. Kuske. Dependence orders for computations of concurrent automata. In *Proceedings of STACS'95*, number 900 in Lecture Notes in Computer Science, pages 467–478. Springer Verlag, 1995. Full version to appear in Theoret. Comp. Science.

[2] R. Cori, Y. Métivier, and W. Zielonka. Asynchronous mappings and asynchronous cellular automata. *Information and Computation*, 106:159–202, 1993.

[3] V. Diekert. A partial trace semantics for petri nets. *Theoretical Computer Science*, 113:87–105, 1994. Special issue of ICWLC 92, Kyoto (Japan).

[4] V. Diekert and A. Muscholl. Construction of asynchronous automata. In G. Rozenberg and V. Diekert, editors, *Book of Traces*, pages 249–267. World Scientific, Singapore, 1995.

[5] V. Diekert and G. Rozenberg, editors. *Book of Traces*. World Scientific, Singapore, 1995.

[6] M. Droste and D. Kuske. Logical definability of recognizable and aperiodic languages in concurrency monoids. In *Proceedings of CSL'95*, Lecture Notes in Computer Science. Springer Verlag, 1996. to appear.

[7] J.L. Gischer. The equational theory of pomsets. *Theoretical Computer Science*, 61:199–224, 1988.

[8] A. Mazurkiewicz. Concurrent program schemes and their interpretations. Tech. rep. DAIMI PB 78, Aarhus University, 1977.

[9] A. Mazurkiewicz. Trace theory. In W. Brauer et al., editors, *Advances in Petri Nets'86*, number 255 in Lecture Notes in Computer Science, pages 279–324. Springer Verlag, 1987.

[10] A. Muscholl. On the complementation of Büchi asynchronous cellular automata. In S. Abiteboul and E. Shamir, editors, *Proceedings of the 21st International Colloquium on Automata, Languages and Programming (ICALP'94)*, number 820 in Lecture Notes in Computer Science. Springer Verlag, 1994.

[11] V.R. Pratt. Modelling concurrency with partial orders. *J. of Parallel Programming*, 15:33–71, 1987.

[12] P.H. Starke. Processes in petri nets. *EIK*, 17:389–416, 1981.

[13] Wolfgang Thomas. On logical definability of trace languages. In V. Diekert, editor, *Proceedings of a workshop of the ESPRIT BRA No 3166: Algebraic and Syntactic Methods in Computer Science (ASMICS), 1989*, Report TUM-I9002, Technical University of Munich, pages 172–182, 1990.

[14] W. Zielonka. Notes on finite asynchronous automata. *R.A.I.R.O. — Informatique Théorique et Applications*, 21:99–135, 1987.

[15] W. Zielonka. Safe executions of recognizable trace languages by asynchronous automata. In A. R. Meyer et al., editors, *Proceedings of the Symposium on Logical Foundations of Computer Science (Logic at*

Botik'89), number 363 in Lecture Notes in Computer Science, pages 278–289. Springer Verlag, 1989.