Chapter 5

The expressive power of ACAs

This chapter deals with the question which properties can be expressed by a Σ-ACA. By Corollary 3.3.6, the expressible properties are at least recursive. On the other hand, Example 2.1.6 shows that not all recursive sets of Σ-dags are recognizable. The situation is similar to that of finite sequential automata and sets of words: Any language that is accepted by a finite sequential automaton is recursive, but the converse is false. In this setting, several answers are known to the question which properties can be checked by a finite sequential automaton: Kleene showed that these are precisely the rational properties. By the Myhill-Nerode Theorem, a property can be checked by a finite sequential automaton if its syntactic monoid is finite. Furthermore, Büchi and Elgot [Büc60, Elg61] showed that a property of words can be checked by a finite automaton if and only if it can be expressed in the monadic second order logic. This relation between a model of a computational device (finite sequential automata) and monadic second order logic is a paradigmatic result. It has been extended in several directions, e.g. to infinite words [Büc60], to trees [Rab69], to finite [Tho90b] and to real [EM93, Ebi94] traces, and to computations of concurrent automata [DK96, DK98]. This relation does clearly not hold for Σ-ACMs in general: Example 2.1.5 provides a word language that can be accepted by a Σ-ACM (that is even monoton and effective), but not by a finite sequential automaton. Hence, this set of Σ-dags cannot be axiomatized in monadic second order logic. Therefore, we examine whether there is such a close relation between Σ-ACAs and MSO.

It is shown that any recognizable set can be axiomatized by a sentence of the monadic second order logic. Since the converse is not true (cf. Example 2.1.6), we then restrict furthermore to so called (Σ, k)-dags and show that a set of (Σ, k)-dags is recognizable (relative to the set of all (Σ, k)-dags and even relative to the set of all Σ-dags) iff it can be monadically axiomatized. But it is necessary to allow nondeterminism in the automata since the expressive power of deterministic Σ-ACAs is shown to be strictly weaker.

5.1 From ACAs to MSO

In this section, we will prove that for any ACAs $\mathcal{A}$, there exists a monadic sentence which axiomatizes the language accepted by $\mathcal{A}$. The proof of this result follows [DG96] (see also [DGK00]). There, the restricted case of Σ-dags that are Hasse-diagrams was dealt with. The only difference between this former result and the result we are going to prove now is the following: The monadic second order logic considered in [DG96] makes statements on partial orders and not on dags. Since the partial order $E^\star$ can be expressed by a monadic formula over dags, this is no difference as far as the expressive power is concerned. But one needs more quantifier alternations which is the reason why in our setting the following theorem states only the existence of a monadic sentence which might not be existential.

Theorem 5.1.1 *Let $\mathcal{A}$ be a possibly nondeterministic Σ-ACA. There exists a monadic sentence φ over Σ such that*

$$L(\mathcal{A}) = \{t \in \mathbb{D} \mid t \models \varphi\}.$$

Proof. Let $\mathcal{A} = ((Q_a)_{a\in\Sigma}, (\delta_{a,J})_{a\in\Sigma,J\subseteq\Sigma}, F)$ be a Σ-ACA. We will construct a monadic sentence which will be satisfied exactly by the Σ-dags that are accepted by $\mathcal{A}$. Let k be the number of states in $\bigcup_{a\in\Sigma} Q_a$. We may assume that $\bigcup_{a\in\Sigma} Q_a = [k] = \{1, \ldots, k\}$. The following sentence ψ claims the existence of a successful run of the automaton.

$$\psi = \exists X_1 \ldots \exists X_k \; \big(\mathrm{partition}(X_1, \ldots, X_k) \wedge \big(\forall x \; \mathrm{trans}(x)\big) \wedge \mathrm{accepted}\big)$$

We will now explain this sentence and give the sub-formulas *partition*, *trans* and *accepted*. A run over a Σ-dag $t = (V, E, \lambda)$ is coded by the set-variables $X_1, \ldots, X_k$. More precisely, X_i stands for the set of vertices mapped to the state i by the run. The formula partition$(X_1, \ldots, X_k)$ ensures that the set-variables $X_1, \ldots, X_k$ describe a mapping from V to $\bigcup_{a\in\Sigma} Q_a$:

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

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

$$\mathrm{trans}(x) = \bigvee_{q\in\delta_{a,J}((q_b)_{b\in J})} \Big(\lambda(x) = a \wedge x \in X_q \wedge \forall y \; ((y, x) \in E \rightarrow \lambda(y) \in J) \wedge \bigwedge_{b\in J} \exists y \; ((y, x) \in E \wedge \lambda(y) = b \wedge y \in X_{q_i})\Big)$$

where the disjunction ranges over all letters $a \in \Sigma$, states $q \in Q_a$, subsets $J \subseteq \Sigma$ and tuples $(q_b)_{b\in J} \in \prod_{b\in J} Q_b$ such that $q \in \delta_{a,J}((q_b)_{b\in J})$.

It remains to state that the run reaches a final state of the automaton. Let *accepted* denote the disjunction of the following sentence for $(f_b)_{b\in J} \in F$:

$$\Big(\forall x\ \big(\lambda(x) \in J\big) \wedge \bigwedge_{b\in J} \exists x\big((\neg\exists y\ (\lambda(x) = \lambda(y) \wedge x < y)) \wedge \lambda(x) = b \wedge x \in X_{f_b}\big)\Big).$$

Since the formula ψ describes an accepting run of the automaton for Σ-dags, we get the statement of the theorem. □

Note that the proof of the theorem above makes use of the finiteness of the sets of local states Q_a in a Σ-ACA. The first language from Example 2.1.5 shows that this finiteness is necessary for the theorem to hold: The language given there can be accepted by a monotone ACM but it is not regular and therefore not monadically axiomatizable. Furthermore, Example 2.1.6 shows that the converse of the theorem does not hold: There, we presented a language that is elementarily axiomatizable, but not acceptable by a monotone ACM and can therefore in particular not be accepted by a Σ-ACA.

5.2 (Σ, k)-dags

Theorem 4.1.10 in particular implies that the set of recognizable Σ-dag-languages is not closed under complementation. Hence, there are monadically axiomatizable languages that cannot be accepted by any Σ-ACA. This section is devoted to the class of (Σ, k)-dags that we introduce next where the expressive power of Σ-ACAs and MSO coincide. The results presented here were originally shown for Hasse-diagrams in [Kus98]. Here, the presentation follows [DGK00] and is in addition extended to (Σ, k)-dags.

5.2.1 k-chain coverings

Let $t = (V, E, \lambda)$ be a Σ-dag. Furthermore, let k be a positive integer and $C_\ell \subseteq V$ for $1 \leq \ell \leq k$. We call the tuple $(C_1, C_2, \ldots, C_k)$ a *k-chain covering* of t if

1. C_ℓ is a chain with respect to the partial order $E^\star$ for $\ell = 1, 2, \ldots, k$,
2. $V = \bigcup_{\ell\in[k]} C_\ell$ and
3. for any $(x, y) \in E$, there exists $\ell \in [k]$ with $x, y \in C_\ell$ and there is no element of C_ℓ properly between x and y (i.e. $xE^\star zE^+y$ and $z \in C_\ell$ imply $x = z$).

The Σ-dag t is a *(Σ, k)-dag* if it has a k-chain covering. Let $\mathbb{D}_k$ denote the set of all (Σ, k)-dags.

Example 2.1.6 (continued) Consider the first Σ-dag in Figure 2.4. It can be covered by the chains $C_i = \{a_1, a_2, \ldots, a_i, b_i, b_{i+1}, b_{i+2}, \ldots, b_8\}$ for $1 \leq i \leq 8$. Hence it is a $(\Sigma, 8)$-dag. The reader may check that it is not possible to cover it by fewer chains, i.e. that it is not a (Σ, k)-dag for $k < 8$. Recall that the set L cannot be accepted by a Σ-ACA. Later (Theorem 5.2.10) we will see that the reason for this is that L is not contained in $\mathbb{D}_k$ for any $k \in \mathbb{N}$.

Example 5.2.1 Let (Σ, D) be some trace alphabet and $(V, \leq, \lambda) \in \mathbb{M}(\Sigma, D)$. Then $(V, \leq, \lambda)$ is a pomset without autoconcurrency. Hence the Hasse-diagram $t = \mathrm{Ha}(V, \leq, \lambda)$ of this trace is a Σ-dag. Even more, it is a (Σ, k)-dag with $k = |D|$: For $(a, b) \in D$, let $C_{a,b} = \lambda^{-1}(a) \cup \lambda^{-1}(b) \subseteq V$. Since a and b are dependent, this set is a chain. Now let $x, y \in V$ with $x \longrightarrow\!\!\!< y$. Then $\lambda(x)$ and $\lambda(y)$ are dependent, i.e. x and y belong to some chain $C_{a,b}$ with $(a, b) \in D$.

Suppose (V, E) is the Hasse-diagram of the partially ordered set $(V, \leq)$. Then, by Dilworth's Theorem [Dil50], the width of $(V, \leq)$ and the number of chains necessary to cover the poset are closely related: The poset $(V, \leq)$ has width at most k iff there are k linearly ordered sets $(C_i, \leq_i)$ such that $(V, \leq)$ is the union of these chains, i.e., such that $V = \bigcup_{1 \leq i \leq k} C_i$ and $\leq = (\bigcup_{1 \leq i \leq k} \leq_i)^\star$. This is similar to our definition of a k-chain-covering of (V, E), the difference being that we require in addition that the edges of (V, E) are covered by the chains C_i. The set of (Σ, k)-dags can nevertheless be described by the width of a partial order whose elements are the edges of the Σ-dag: Let (V, E, λ) be a Σ-dag such that $(V, E^\star)$ has a minimal element and V contains at least 2 nodes. On the set of edges E, define a partial order $\sqsubseteq$ by $(x, y) \sqsubset (x', y')$ iff $(y, x') \in E^\star$. The width of the partial order $(E, \sqsubseteq)$ is called the *chainwidth of* (V, E, λ).

Lemma 5.2.2 *Let $t = (V, E, \lambda)$ be a Σ-dag such that $(V, E^\star)$ has a minimal element and V contains at least 2 nodes. Then $t \in \mathbb{D}_k$ iff the chainwidth of t is at most k.*

Proof. Let $t \in \mathbb{D}_k$ and suppose that its chainwidth is at least $k+1$. Then there exists an antichain $\{(x_i, y_i) \in E \mid 1 \leq i \leq k+1\}$ in $(E, \sqsubseteq)$. Furthermore, there exists a k-chain-covering $(C_\ell)_{1 \leq \ell \leq k}$ of t. Hence, for any i with $1 \leq i \leq k+1$, there is one chain C_ℓ with $x_i, y_i \in C_\ell$. This implies that there are $i < j$ with $x_i, y_i, x_j, y_j \in C_\ell$ for some ℓ, i.e., $\{x_i, y_i, x_j, y_j\}$ is linearly ordered by $E^\star$. Now $y_i E^\star x_j$ or $y_j E^\star x_i$ follows, contradicting our assumption that the edges (x_i, y_i) for $1 \leq i \leq k+1$ form an antichain.

Conversely, let the chainwidth of t be at most k. Then the partially ordered set $(E, \sqsubseteq)$ can be covered by k chains D_ℓ. Let $C_\ell = \{x, y \mid (x, y) \in D_\ell\}$. Then $(C_\ell)_{1 \leq \ell \leq k}$ is a k-chain-covering of t, i.e., $t \in \mathbb{D}_k$. □

For later use (Chapter 6), we also introduce the spine of a Σ-dag. This notion is a generalization from [HR95] where it is defined for Hasse-diagrams. Let $t = (V, E, \lambda)$ be a Σ-dag. We define a new edge relation $E' \supseteq E$ by $(x, y) \in E'$ iff $(x, y) \in E$ or

1. $(x, y) \in E^\star$,
2. for any $z \in V$ we have $(x, z) \in E \Rightarrow zE^\star y$, and
3. for any $z \in V$ we have $(z, y) \in E \Rightarrow xE^\star z$.

Thus, $(x, y) \in E'$ if either $(x, y) \in E$ or $xE^\star y$ and any upper neighbor of x (any lower neighbor of y) is below y (above x, respectively). The directed graph (V, E') is the *spine* spine(t) *of* t. By unc(spine(t)), we denote the maximal size of a set in V that is totally unconnected in spine(t).

To show that for any (Σ, k)-dag t the size of totally unconnected sets in spine(t) is bounded, we will use the following result by Ramsey (cf. [Cam94] for the general formulation):

Ramsey's Theorem [Ram30] *Let c, r be positive integers. Then there is a positive integer $R_r(c)$ such that for any mapping d of the two-element subsets of $[R_r(c)]$ into $[c]$ there exists an r-element subset $A \subseteq [R_r(c)]$ such that we have $d(B) = d(C)$ for any two-element subsets B and C of A.*

Lemma 5.2.3 *Any Σ-dag t in $\mathbb{D}_k$ satisfies* $\text{unc}(\text{spine}(t)) < |\Sigma| \cdot R_2(k+1)$. *Conversely, any Σ-dag t with* $\text{unc}(\text{spine}(t)) \leq m$ *is an element of $\mathbb{D}_k$ with* $k = 2(m+1)|\Sigma|^2 - 1$.

Proof. First suppose $t = (V, E, \lambda) \in \mathbb{D}_k$ and assume, by contradiction, that $\text{unc}(\text{spine}(t)) \geq |\Sigma| \cdot R_2(k+2)$. Since the width of $(V, E^\star)$ is at most $|\Sigma|$, there is a set $A \subseteq V$ with at least $R_2(k+2)$ elements that is totally unconnected in spine(t) and linearly ordered in $(V, E^\star)$. For $x, x' \in A$ with xE^+x' set

$$h(\{x, x'\}) = \begin{cases} 0 & \text{if there exists } y \in V \text{ with } (x, y) \in E \text{ and } (y, x') \notin E^\star \\ 1 & \text{otherwise.} \end{cases}$$

By Ramsey's Theorem, there are $x_1, x_2, \ldots, x_{k+2} \in A$ such that $h(\{x_i, x_j\})$ is constant and $x_i E^+ x_j$ for $i < j$. If $h(\{x_i, x_j\}) = 0$, there exist $y_i \in V$ with $(x_i, y_i) \in E$ and $(y_i, x_{i+1}) \notin E^\star$ for $1 \leq i < k+2$. Hence the edges (x_i, y_i) with $1 \leq i \leq k+1$ form an antichain in $(E, \sqsubseteq)$ of size $k+1$ implying that t has chainwidth larger than k. In case $h(\{x_i, x_j\}) = 1$, we can argue dually and, again, obtain that t has chainwidth larger than k. But this contradicts our assumption $t \in \mathbb{D}_k$ together with the preceding lemma.

Now let $\mathrm{unc}(\mathrm{spine}(t)) \leq m$ and $k = 2(m+1)|\Sigma|^2 - 1$ and suppose $t \notin \mathbb{D}_k$. Then $(E, \sqsubseteq)$ contains an antichain of size $k+1$. Since any element $x \in V$ has at most $|\Sigma|$ upper neighbors in (V, E), there is an antichain $\{(x_i, y_i) \in E \mid 1 \leq i \leq 2(m+1)|\Sigma|\}$ with $x_i \neq x_j$ for $i < j$. Since furthermore $(V, E^\star)$ has width at most $|\Sigma|$, we can assume that $(x_i, x_j) \in E^+$ for $1 \leq i < j \leq 2(m+1)$.

Let $A = \{x_i \mid 1 \leq i \leq 2(m+1)\}$. We define a set $B \subseteq A$ as follows: For any maximal sequence $1 \leq i_1 < i_2 < \ldots i_n \leq 2(m+1)$, let $x_{i_{2j+1}}$ belong to B (i.e., the elements at an odd position belong to B and those at even positions do not). Then B contains at least $m+1$ elements. Let $x_i, x_j \in B$ with $i < j$. Then $(y_i, x_j) \notin E^\star$ since $\{(x_i, y_i) \in E \mid 1 \leq i \leq m+1\}$ is an antichain in $(E, \sqsubseteq)$. But this implies that $\{x_i \mid 1 \leq i \leq m+1\}$ is totally unconnected in $\mathrm{spine}(t)$, contradicting our assumption $\mathrm{unc}(\mathrm{spine}(t)) \leq m$. □

Our last alternative characterization of the Σ-dags that admit a k-chain-covering is in terms of k-chain-mappings that we define next:

Definition 5.2.4 Let $t = (V, E, \lambda)$ be a Σ-dag, $k \in \mathbb{N}$ and $\Lambda : V \to (2^{[k]} \setminus \{\emptyset\})$. The function Λ is a *k-chain mapping* if

(1) for all minimal vertices $x, y \in V$, if $x \neq y$ then $\Lambda(x) \cap \Lambda(y) = \emptyset$,

(2) for all non minimal vertices $y \in V$ and $\ell \in \Lambda(y)$, there exists $x \in V$ with $(x, y) \in E$ and $\ell \in \Lambda(x)$,

(3) for all vertices $x \in V$ that are not maximal and for all $\ell \in \Lambda(x)$, the set $\{y \in V \mid (x, y) \in E, \ell \in \Lambda(y)\}$ is empty or has a least element, and

(4) for all $(x, y) \in E$, there is $\ell \in \Lambda(x) \cap \Lambda(y)$ such that for any $z \in V$ with xE^+zE^+y it holds $\ell \notin \Lambda(z)$.

The following lemma relates k-chain mappings and k-chain coverings thereby justifying the name k-chain mapping.

Lemma 5.2.5 *Let $t = (V, E, \lambda)$ be a Σ-dag. Then $t \in \mathbb{D}_k$ iff there exists a k-chain mapping. In particular, if Λ is a k-chain mapping of t and $\ell \in [k]$, then the set $\Lambda^{-1}(\ell) = \{x \in V \mid \ell \in \Lambda(x)\}$ is a chain with respect to $E^\star$ and $(\Lambda^{-1}(\ell))_{\ell \in [k]}$ is a k-chain covering. Conversely, if $(C_\ell)_{\ell \in [k]}$ is a maximal k-chain covering, then $\Lambda(x) = \{\ell \in [k] \mid x \in C_\ell\}$ defines a k-chain mapping.*

Proof. Let $t \in \mathbb{D}_k$. Then there exists a k-chain covering $(C_\ell)_{\ell \in [k]}$ of t. We may assume that the chain covering $(C_\ell)_{\ell \in [k]}$ is maximal with respect to the componentwise inclusion (i.e. incorporating any vertex newly into one of the chains C_ℓ destroys its property to be a k-chain covering). Now define $\Lambda(x) := \{\ell \in [k] \mid x \in C_\ell\}$. Then $\Lambda : V \to (2^{[k]} \setminus \{\emptyset\})$ since $V = \bigcup_{\ell \in [k]} C_\ell$. Since C_ℓ

is a chain for each $\ell \in [k]$, any two different minimal elements of t belong to disjoint sets of chains. Hence the first property of Definition 5.2.4 is satisfied. Now let $y \in V$ be non minimal and $\ell \in \Lambda(y)$. Since the k-chain covering $(C_\ell)_{\ell\in[k]}$ is maximal, there exists $x \in V$ with $(x, y) \in E$ and $\ell \in \Lambda(x)$. Hence, the second requirement is satisfied. The targets in C_ℓ of edges that originate in a nonmaximal vertex x are linearly ordered. Hence this set admits a least element as required by the third condition. If $(x, y) \in E$, there exists $\ell \in [k]$ such that $x, y \in C_\ell$ and no element of C_ℓ is properly between x and y. Hence $\ell \in \Lambda(x) \cap \Lambda(y)$ and for any z properly between x and y we have $\ell \notin \Lambda(z)$. Thus, we proved the last statement of Definition 5.2.4.

Conversely, let Λ be a k-chain mapping of the Σ-dag t. For $\ell \in [k]$, define $C_\ell := \{x \in V \mid \ell \in \Lambda(x)\}$. Since $\Lambda(x) \neq \emptyset$ for all $x \in V$, we get $V = \bigcup_{\ell\in[k]} C_\ell$. By the last property for Λ, for any $(x, y) \in E$ there exists $\ell \in [k]$ with $x, y \in C_\ell$ such that no element of C_ℓ lies properly between x and y. It remains to show that C_ℓ is a chain for any ℓ: Let $x, y \in C_\ell$. By the second property of Λ, there exists a sequence $x_0, x_1, \dots x_m = x$ of elements of C_ℓ with x_0 minimal in $(V, E^\star)$ and $(x_i, x_{i+1}) \in E$. We can even assume that x_{i+1} is the least element of C_ℓ above x_i such that $(x_i, x_{i+1}) \in E$. Similarly, there exist elements $y_0, y_1, \dots y_n = y$ of C_ℓ with y_0 minimal in $(V, E^\star)$ and $(y_j, y_{j+1}) \in E$ such that y_{j+1} is the least element of C_ℓ above y_j with $(y_j, y_{j+1}) \in E$. Now let $m \leq n$. By the first property of Λ, $x_0 = y_0$. Let $0 \leq i < m$ be such that $x_i = y_i$. This element is the source of edges going to x_{i+1} and to y_{i+1}. Since we chose x_{i+1} and y_{i+1} minimal in C_ℓ above $x_i = y_i$ with $(x_i, x_{i+1}) \in E$ and $(y_i, y_{i+1}) \in E$, we obtain $x_{i+1} = y_{i+1}$. This shows that $(x, y) \in E^\star$, i.e. C_ℓ is a chain. □

5.2.2 k-chain coverings are recognizable

Next we construct an ACA $\mathcal{A}_k$ that "produces" k-chain mappings (i.e., any successful run of $\mathcal{A}_k$ corresponds to a k-chain mapping and vice versa, Lemma 5.2.6). This implies immediately that $\mathbb{D}_k$ is recognizable relative to $\mathbb{D}$. In addition, we will use the produced k-chain mapping to relabel a (Σ, k)-dag into a trace. This latter result will enable us to use the theory of Mazurkiewicz traces to show that any monadically axiomatizable set of (Σ, k)-dags is recognizable relative to $\mathbb{D}_k$.

We start with the definition of the automaton $\mathcal{A}_k$: Recall that $\text{part}([k], \Sigma)$ is the set of *partial* functions g from $[k]$ to Σ with $\text{dom}(g) \neq \emptyset$. We write $\text{part}(k, \Sigma)$ for this set $\text{part}([k], \Sigma)$. For a partial function $f \in \text{part}(k, \Sigma)$, we first define an ACA $\mathcal{A}_k(f)$ whose local states are partial functions in $\text{part}(k, \Sigma)$. Intuitively, a node x of some (Σ, k)-dag t will be labeled by the partial function g in some run of $\mathcal{A}_k(f)$ if $\text{dom}(g)$ is the set of chains C_ℓ going through x and for all $\ell \in \text{dom}(g)$, $g(\ell)$ is the next action for the chain ℓ. The partial mapping f is in some sense the initial state of the automaton $\mathcal{A}_k(f)$: $f(\ell) = a$ iff the chain ℓ starts with an action a. As we will see, runs of this automaton correspond to k-chain mappings.

More precisely, the ACA $\mathcal{A}_k(f)$ is defined as follows: The set of local states (common for all processes) is $Q = \text{part}(k, \Sigma)$. For $a \in \Sigma$, let $\delta_{a,\emptyset}$ consist of all nonempty partial functions $g \in Q$ with $\text{dom}(g) = f^{-1}(a)$. For $\emptyset \neq J \subseteq \Sigma$ and $g_b \in Q$ for $b \in J$, we let $\delta_{a,J}((g_b)_{b\in J})$ be the set of all nonempty partial functions $g \in Q$ such that

1. for $b \in J$ there exists $\ell \in \text{dom}(g)$ with $g_b(\ell) = a$ and

2. for $\ell \in \text{dom}(g)$ there exists $b \in J$ with $g_b(\ell) = a$.

Finally, all tuples of states are accepting. Let $\mathcal{A}_k$ denote the disjoint union of the automata $\mathcal{A}_k(f)$ for all partial functions $f \in \text{part}(k, \Sigma)$. Note that not all runs of $\mathcal{A}_k$ are successful, only those that lie completely inside $\mathcal{A}_k(f)$ for some $f \in \text{part}(k, \Sigma)$ are. This can be easily checked by considering the final global state.

The following lemma shows that the k-chain mappings Λ on a Σ-dag t coincide precisely with the mappings $\text{dom} \circ r : V \to 2^{[k]}$ where r is a successful run of the automaton $\mathcal{A}_k$ constructed above.

Lemma 5.2.6 *For $k \in \mathbb{N}$ and $t = (V, E, \lambda) \in \mathbb{D}$, we have:*

1. *for any successful run r of $\mathcal{A}_k$ on t, the mapping* $\text{dom} \circ r : V \to 2^{[k]} \setminus \{\emptyset\}$ *is a k-chain mapping.*

2. *For any k-chain mapping Λ on t, there exists a successful run r of $\mathcal{A}_k$ on t such that $\Lambda = \text{dom} \circ r$.*

Proof. 1. Let $r : V \to \text{part}(k, \Sigma)$ be a successful run of $\mathcal{A}_k$ on t and let $\Lambda = \text{dom} \circ r$. There exists a partial function $f \in \text{part}(k, \Sigma)$ such that r is a run of $\mathcal{A}_k(f)$. Now let $x, y \in V$ be minimal and distinct. Then $r(x) \in \delta_{\lambda(x),\emptyset}$, and therefore $\text{dom} \circ r(x) = f^{-1}(\lambda(x))$. Similarly, $\text{dom} \circ r(y) = f^{-1}(\lambda(y))$. Since x and y are incomparable with respect to $E^\star$, $\lambda(x) \neq \lambda(y)$. Hence $\Lambda(x)$ and $\Lambda(y)$ are disjoint. Thus we showed the first condition of Definition 5.2.4.

Now, let $x \in V$ be non minimal. For $b \in \text{R}(x)$, there exists a unique vertex $x_b \in V$ with $(x_b, x) \in E$ and $\lambda(x_b) = b$. Let also $g_b = r(x_b)$ and $g = r(x)$. Since r satisfies the run condition of $\mathcal{A}(f)$ at x, we have $g \in \delta_{\lambda(x),\text{R}(x)}((g_b)_{b\in\text{R}(x)})$. Now we deduce that for all $\ell \in \Lambda(x) = \text{dom}(g)$, there exists $b \in \text{R}(x)$ with $\ell \in \text{dom}(g_b) = \Lambda(x_b)$ showing Definition 5.2.4 (2). Next, we show Definition 5.2.4 (4) for the edge (x_b, x): Since $g \in \delta_{\lambda(x),\text{R}(x)}((g_b)_{b\in\text{R}(x)})$, there is $\ell \in \text{dom}(g) \cap \text{dom}(g_b) = \Lambda(x) \cap \Lambda(x_b)$ such that $r(x_b)(\ell) = \lambda(x)$. Now assume $x_b E z E^\star x$ with $\ell \in \Lambda(x)$. Then $x_b = \partial_b(x)$. Since r is a run of $\mathcal{A}_k(f)$, we obtain $\lambda(z) = r(x_b)(\ell) = \lambda(x)$. This shows that z and x are targets of edges that originate in x_b, and that they carry the same label $\lambda(x)$. Hence they are equal, i.e. there is no element z properly between x_b and x such that $\ell \in \Lambda(z)$. Thus, Definition 5.2.4 (4) holds. To show Definition 5.2.4(3), let $x \in V$ and $\ell \in \Lambda(x)$ such that the set

$\{y \in V \mid (x,y) \in E$ and $\ell \in \Lambda(x)\}$ is not empty. Note that this set is a subset of the chain C_ℓ. Hence it has a least element.

2. Assume now that Λ is a k-chain mapping. We will construct a successful run r of $\mathcal{A}_k$ such that $\mathrm{dom} \circ r = \Lambda$. Let $x \in V$. Indeed, the domain of the partial function $r(x) \in \mathrm{part}(k,\Sigma)$ will be $\Lambda(x)$. Now, for all $\ell \in \mathrm{dom}(r(x)) = \Lambda(x)$, if there is $y \in V$ with $(x,y) \in E$ and $\ell \in \Lambda(y)$, then, by Definition 5.2.4 (3), there is a least such y. If such a vertex y exists then we set $r(x)(\ell) = \lambda(y)$ and otherwise we set $r(x)(\ell) = a$ for some $a \in \Sigma$ (in this last case, we can give any value since it will never be used).

Let $f \in \mathrm{part}(k,\Sigma)$ be the partial function defined by $\ell \in \mathrm{dom}(f)$ iff there exists a minimal vertex $x \in V$ with $\ell \in \Lambda(x)$ and in this case we set $f(\ell) = \lambda(x)$. Note that f is well-defined thanks to Definition 5.2.4 (1).

We show that indeed r is a run of $\mathcal{A}_k(f)$: Clearly, if $x \in V$ is minimal then we have $\mathrm{dom}(r(x)) = \Lambda(x) = f^{-1}(\lambda(x))$ as required by the initial transitions of $\mathcal{A}_k(f)$.

Now, let $x \in V$ be non minimal. For all $b \in \mathrm{R}(x)$, let $(x_b, x) \in E$ be such that $\lambda(x_b) = b$. We will show that $r(x) \in \delta_{\lambda(x),\mathrm{R}(x)}(r(x_b)_{b\in\mathrm{R}(x)})$. First, for all $b \in \mathrm{R}(x)$, by Definition 5.2.4 (4), there exists $\ell \in \Lambda(x) \cap \Lambda(x_b) = \mathrm{dom}(r(x)) \cap \mathrm{dom}(r(x_b))$ such that no element z with $\ell \in \Lambda(z)$ lies properly between x_b and x. By the construction of $r(x_b)$, it follows that $r(x_b)(\ell) = \lambda(x)$. Second, for $\ell \in \mathrm{dom}(r(x)) = \Lambda(x)$, there exists $b \in \mathrm{R}(x)$ with $\ell \in \Lambda(x_b) = \mathrm{dom}(r(x_b))$ by Definition 5.2.4 (2). By definition of $r(x_b)$, it follows that $r(x_b)(\ell) = \lambda(x)$. Thus we have shown that r is a run of $\mathcal{A}_k(f)$ which concludes the proof. □

As an immediate consequence of the lemma above and Lemma 5.2.5, we obtain that $\mathbb{D}_k$ is recognizable relative to $\mathbb{D}$:

Corollary 5.2.7 *For $k \in \mathbb{N}$, we have $L(\mathcal{A}_k) = \mathbb{D}_k$.* □

Now we define a trace alphabet (Γ, D) as follows: Let

$$\Gamma := \Sigma \times (2^{[k]} \setminus \{\emptyset\}).$$

The dependence relation D is defined by

$$D = \{((a,M),(b,N)) \mid M \cap N \neq \emptyset \text{ or } a = b\}.$$

This binary relation on Γ is obviously reflexive and symmetric. Thus (Γ, D) is indeed a dependence alphabet. Let $\mathbb{M}(\Gamma, D)$ denote the trace monoid over (Γ, D). Now let $t = (V, \leq, \lambda_\Gamma)$ be a trace over (Γ, D). From this trace, we define a Σ-dag as follows: For $x, y \in V$, let $(x,y) \in E$ iff there exists $\ell \in \pi_2 \circ \lambda_\Gamma(x) \cap \pi_2 \circ \lambda_\Gamma(y)$ such that

$$x = \max\{w < y \mid \ell \in \pi_2 \circ \lambda_\Gamma(w)\}.$$

Now let $\Pi(V, \leq, \lambda) = (V, E, \pi_1 \circ \lambda_\Gamma)$.

For an arbitrary trace $t \in \mathbb{M}(\Gamma, D)$, $\Pi(t)$ is a directed acyclic graph whose vertices are labeled by elements from Σ. Let $\mathbb{M}'$ denote the set of all traces $t \in \mathbb{M}(\Gamma, D)$ such that $\Pi(t) \in \mathbb{D}_k$, i.e. that are mapped to a (Σ, k)-dag by the mapping Π. Note that the relation E defined above is elementarily definable in $(V, \leq, \lambda_\Gamma)$. Since in addition the set of (Σ, k)-dags is monadically axiomatizable relative to all Σ-labeled dags, the set $\mathbb{M}'$ is axiomatizable relative to $\mathbb{M}(\Gamma, D)$.

Next, we define the "inverse" of Π: Let $t = (V, E, \lambda)$ be a (Σ, k)-dag. Then there exists a maximal k-chain covering $(C_\ell)_{\ell \in [k]}$. For $y \in V$, define

$$\lambda_\Gamma(y) := (\lambda(y), \{\ell \in [k] \mid y \in C_\ell\}).$$

The following lemma in particular implies that any (Σ, k)-dag is the image under Π of some trace from $\mathbb{M}'$, i.e. $\Pi(\mathbb{M}') = \mathbb{D}_k$.

Lemma 5.2.8 *Let $t = (V, E, \lambda)$ be a (Σ, k)-dag and let $(C_i)_{i \in [k]}$ be a maximal k-chain covering of t. Let $\lambda_\Gamma(x) = (\lambda(x), \{\ell \in [k] \mid x \in C_\ell\})$ for $x \in V$. Then $\Pi(V, E^\star, \lambda_\Gamma) = t$ and $(V, E^\star, \lambda_\Gamma) \in \mathbb{M}'$.*

Proof. Let $\leq$ denote the partial order $E^\star$. First we show that $(V, \leq, \lambda_\Gamma)$ is a trace from $\mathbb{M}(\Gamma, D)$: Let $x, y \in V$ with $x \longrightarrow\!\!\!< y$ (with respect to the partial order $\leq$). Then $(x, y) \in E$. Since $(C_i)_{i \in [k]}$ is a k-chain covering, there exists $\ell \in [k]$ with $x, y \in C_\ell$. Hence $\ell \in \pi_2 \circ \lambda_\Gamma(x) \cap \pi_2 \circ \lambda_\Gamma(y)$ implying $(\lambda_\Gamma(x), \lambda_\Gamma(y)) \in D$. Now let $x, y \in V$ be incomparable. Since C_i is a chain with respect to $\leq$ for $1 \leq i \leq k$, we get $\emptyset = \pi_2 \circ \lambda_\Gamma(x) \cap \pi_2 \circ \lambda_\Gamma(y)$. Since (V, E, λ) is a Σ-dag, x and y carry different labels from Σ. Hence we showed $(\lambda_\Gamma(x), \lambda_\Gamma(y)) \notin D$ which concludes the proof that $(V, \leq, \lambda_\Gamma)$ is a trace.

Now let $\Pi(V, E^\star, \lambda_\Gamma) = (V, E', \lambda')$. Then, $\lambda' = \pi_1 \circ \lambda_\Gamma = \lambda$. It remains to show $E = E'$. So let $(x, y) \in E$. Since $(C_\ell)_{\ell \in [k]}$ is a k-chain covering, there exists $\ell \in [k]$ such that $x, y \in C_\ell$ and no $z \in C_\ell$ lies properly between x and y. Hence $x = \max\{w \in C_\ell \mid w < y\}$ implying $x = \max\{w < y \mid \ell \in \pi_2 \circ \lambda_\Gamma(w)\}$. Hence $(x, y) \in E'$.

If, conversely, $(x, y) \in E'$, then there exists $\ell \in \pi_2 \circ \lambda_\Gamma(y)$ such that $x = \max\{w < y \mid \ell \in \pi_2 \circ \lambda_\Gamma(w)\}$. Since $\pi_2 \circ \lambda_\Gamma(x) = \{\ell \in [k] \mid x \in C_\ell\}$, we obtain $x \in C_\ell$. In addition, $x < y$ implies xE^+y. By contradiction, assume $(x, y) \notin E$. Then there exists $z \in V$ with xE^+zE^+y. Since there is no element of C_ℓ properly between x and y, the set $C_\ell \cup \{z\}$ is a chain with respect to $E^\star$. Since $(x, y) \notin E$, the tuple $(C_1, \ldots, C_{\ell-1}, C_\ell \cup \{z\}, C_{\ell+1}, \ldots, C_k)$ is a k-chain covering contradicting our assumption that $(C_\ell)_{\ell \in [k]}$ is maximal. Thus, we showed $(x, y) \in E$. $\square$

5.2.3 Monadic second order logic

Lemma 5.2.9 *Let φ be a sentence of the monadic second order logic over the alphabet Σ. Then there exists a sentence ψ of the monadic second order logic over the alphabet Γ using the binary relation $\leq$ such that*

$$\{\Pi(s) \mid s \in \mathbb{M}(\Gamma, D) \text{ and } s \models \psi\} = \{t \in \mathbb{D}_k \mid t \models \varphi\}.$$

Proof. The sentence φ contains atomic formulas of the form $\lambda(x) = a$ for $a \in \Sigma$ and of the form $(x, y) \in E$. Replace any occurrence of an atomic formula $\lambda(x) = a$ by $\bigvee_{A\in\Gamma,\pi_1(A)=a} \lambda_\Gamma(x) = A$. There is a monadic formula η using the relation $\leq$ and the mapping λ_Γ that states for any two vertices x, y in a trace $t \in \mathbb{M}(\Gamma, D)$ that there exists $\ell \in [k]$ such that $x = \max\{w < y \mid \ell \in \pi_2 \circ \lambda_\Gamma(w)\}$. Replace any subformula of φ of the form $(x, y) \in E$ by $\eta(x, y)$. The result of these replacements is denoted by $\overline{\varphi}$. Note that $\overline{\varphi}$ is a sentence of the monadic second order logic over the alphabet Γ using the relation $\leq$. Now let $s \in \mathbb{M}'$. Then it is easily seen that $s \models \overline{\varphi}$ iff $\Pi(s) \models \varphi$. Furthermore, there is a monadic second order sentence μ axiomatizing $\mathbb{M}'$ relative to $\mathbb{M}(\Gamma, D)$. Thus, we have the required equality for $\psi = \eta \wedge \overline{\varphi}$. □

Before showing that any monadically axiomatizable set of (Σ, k)-dags can be accepted by an ACA, we have to introduce a variant of asynchronous cellular automata. This variant is meant to work on traces from $\mathbb{M}(\Gamma, D)$. Differently from ACAs considered so far, not every letter of Γ has its own sequential process, but some of the processes are collected into one new sequential component. This collection is given by a partition of Γ into dependence cliques: So let $\Gamma_i \subseteq \Gamma$ for $i \in [n]$ be mutually disjoint sets satisfying $\Gamma_i \times \Gamma_i \subseteq D$ (i.e. the letters from Γ_i are mutually dependent). A trace-ACA over $(\Gamma_i)_{i\in[n]}$ is a tuple $\mathcal{A} = ((Q_i)_{i\in[n]}, (\delta_{a,J})_{a\in\Gamma,J\subseteq[n]}, F)$ where

- Q_i is a finite set of local states for process i,
- $\delta_{a,J} : \prod_{j\in J} \to 2^{Q_i}$ is a local transition function with $a \in \Gamma_i$, and
- $F \subseteq \bigcup_{\emptyset\neq J\subseteq[n]} \prod_{j\in J} Q_j$ is a set of final states.

As remarked earlier, these automata will run on traces from $\mathbb{M}(\Gamma, D)$, more precisely, on the Hasse-diagram of a trace. The only difference in the definition of a run for trace-ACAs is that the transition $\delta_{a,J}$ writes into the process i with $a \in \Gamma_i$. Thus the formal definition is an obvious variation of that from page 13. Therefore, we omit it here.

Theorem 5.2.10 *Let φ be a monadic sentence over the alphabet Σ and let $k \in \mathbb{N}$. Then there exists a Σ-ACA $\mathcal{A}$ such that $L(\mathcal{A}) = \{t \in \mathbb{D}_k \mid t \models \varphi\}$.*

Proof. By Lemma 5.2.9, there is a monadically axiomatizable set $L \subseteq \mathbb{M}(\Gamma, D)$ such that $\Pi(L) = \{t \in \mathbb{D}_k \mid t \models \varphi\}$. Hence by [Tho90b, EM96], the set L is recognizable in $\mathbb{M}(\Gamma, D)$.

For $a \in \Sigma$, let $\Gamma_a = \{(a, M) \in \Gamma\}$ denote the set of letters from Γ whose first component equals a. Then Γ_a is a dependence clique in (Γ, D) and the sets Γ_a are mutually disjoint and cover Γ. By an immediate variant of Zielonka's result [Zie87] (cf. also [CMZ93, Die90]), there exists a trace-ACA

$$\mathcal{A}_\varphi = ((Q_a^\varphi)_{a\in\Sigma}, (\delta^\varphi_{(a,M),J})_{(a,M)\in\Gamma, J\subseteq\Sigma}, F^\varphi)$$

over $(\Gamma_a)_{a\in\Sigma})$ that accepts L relative to $\mathbb{M}(\Gamma, D)$.

Furthermore, let $\mathcal{A}_k = ((Q_a)_{a\in\Sigma}, (\delta_{a,J}), F)$ be the ACA constructed above that accepts the set of all (Σ, k)-dags. We define a Σ-ACA $\mathcal{A}' = ((Q'_a)_{a\in\Sigma}, (\delta'_{a,J}), F')$ over the alphabet Σ as follows: $Q'_a = Q_a \times Q_a^\varphi$ and a tuple $(g_b, q_b)_{b\in J}$ belongs to F' iff $(g_b)_{b\in J} \in F$ and $(q_b)_{b\in J} \in F^\varphi$. To define the transition functions, let $\delta'_{a,J}((g_b, q_b)_{b\in J})$ be the set of all pairs (g, q) satisfying $g \in \delta_{a,J}((g_b)_{b\in J})$ and $q \in \delta^\varphi_{(a,M),J}((q_b)_{b\in J})$ with $M = \mathrm{dom}(g)$. Note that a run of the Σ-ACA $\mathcal{A}'$ "contains" a run of $\mathcal{A}_k$. This run "relabels" the (Σ, k)-dag t in consideration into some trace $s \in \Pi^{-1}(t)$ (see Lemmas 5.2.6 and 5.2.8). The trace s is in fact the actual input of the trace-ACA $\mathcal{A}_\varphi$. Therefore, the (Σ, k)-dag t is accepted by $\mathcal{A}'$ iff $s \in \Pi^{-1}(t)$ is accepted by $\mathcal{A}_\varphi$, that is, iff $t \models \varphi$. □

Recall that the above proof rests on the Σ-ACA $\mathcal{A}_k$. This automaton guesses a k-chain-mapping, i.e., it is nondeterministic. Hence, this theorem leaves the question open whether deterministic ACAs suffice to capture the expressive power of monadic second order logic relative to the set of (Σ, k)-dags. The following proposition states that this is not the case, i.e. in particular, that the nondeterministic Σ-ACAs stricly exceed the deterministic ACAs in expressive power relative to (Σ, k)-dags:

Proposition 5.2.11 *Let $k \in \mathbb{N}$ with $k > 1$ and let the alphabet Σ contain at least two letters. Then there exists a set of (Σ, k)-dags that is monadically axiomatizable relative to $\mathbb{D}_k$, but not acceptable by any deterministic Σ-ACA.*

Proof. It suffices to prove the statement for $k = 2$, and $\Sigma = \{a, b\}$. So, let L consist of all (Σ, k)-dags (V, E, λ) over Σ that have a largest (with respect to $E^\star$) vertex. This language is trivially axiomatizable in MSO relative to $\mathbb{D}_k$.

We show that there is no deterministic Σ-ACA $\mathcal{A}$ accepting among the (Σ, k)-dags all those that have a largest vertex: By contradiction, assume $\mathcal{A}$ is such a Σ-ACA. Let $\ell = |Q_a| + 2$ and consider the (Σ, k)-dag $t = (V, E, \lambda)$ with vertex set $V = \{a_i \mid i = 1, 2, \ldots, \ell\} \cup \{b_1\}$, $a_1 E a_2 \ldots E a_\ell E b_1$ and with the canonical labeling λ with $\lambda(a_i) = a$ and $\lambda(b_1) = b$. Then $t \in L$. Hence there is a successful run r

of $\mathcal{A}$ on t. Since $\ell > |Q_a| + 1$, there are $i < j < \ell$ such that $r(a_i) = r(a_j)$. Now consider the (Σ, k)-dags t_1 and t_2 with $V_1 = V_2 = \{a_\ell \mid \ell = 1, 2, \ldots, j\} \cup \{b_1\}$ and the canonical labeling. The edge relations are defined by $a_1 E_1 a_2 E_1 a_3 \ldots E_1 a_j E_1 b_1$ (i.e. $E_1^\star$ is a linear ordering with largest element b_1) and $a_1 E_2 a_2 E_2 a_3 \ldots E_2 a_j$ and $a_i E_2 b_1$ (i.e. in $E_2^\star$, the a-labeled elements are linearly ordered, but the maximal element b_1 covers a_i and is not the largest element of $(V_2, E_2^\star)$). Since $t_1 \in L$, there is a successful run r_1 of $\mathcal{A}$ on t_1. Since $\mathcal{A}$ is deterministic, we have $r_1(a_\ell) = r(a_\ell)$ for $\ell \leq j$. This implies $r_1(a_i) = r_1(a_j)$ since the equality holds for the run r. Hence r_1 is a run on t_2, too. The global final state of r_1 considered on t_1 equals that of r_1 considered on t_2. Hence t_2 is accepted by $\mathcal{A}$, contradicting our assumption since t_2 does not have any largest vertex. □

Thus, differently from traces, for (Σ, k)-dags the deterministic ACAs are strictly weaker in expressive power than monadic second order logic.

Our methods in particular imply that the monadic theory of $\mathbb{D}_k$ is decidable for any $k \in \mathbb{N}$: Let φ be a monadic sentence. Using Lemma 5.2.9, we can build a monadic sentence ψ that axiomatizes a preimage under Π of the models of φ in $\mathbb{D}_k$. Hence $\neg\psi$ is a tautology iff $\neg\varphi$ is. Since the monadic theory of traces is decidable [EM96], the result follows. There is another, more direct way to prove this decidability: Given $k \in \mathbb{N}$, one can bound the pathwidth (cf. [Bod98] for an overview) of the dags in $\mathbb{D}_k$ by some n. Since $\mathbb{D}_k$ is monadically axiomatizable, and since the monadic theory of the dags of pathwidth at most n is decidable [Cou90], the decidability follows. Anyway, using Theorem 5.1.1, one obtains the following result:

Corollary 5.2.12 *There exist algorithms that solve the following decision problems:*
input: *an alphabet Σ, $k \in \mathbb{N}$ and a Σ-ACA $\mathcal{A}$.*
output: *Is $L(\mathcal{A}) \cap \mathbb{D}_k$ empty?*
Is $L(\mathcal{A})$ contained in $\mathbb{D}_k$?
Does $L(\mathcal{A}_1) \cap \mathbb{D}_k = L(\mathcal{A}_2) \cap \mathbb{D}_k$?

Recall that by Proposition 5.2.11 the expressive power of deterministic Σ-ACAs does not capture that of monadic second order logic relative to $\mathbb{D}_k$. Hence, we get in particular that nondeterministic ACAs are strictly more powerful than deterministic ACAs within the class $\mathbb{D}_k$ for $k \geq 2$ and the same holds for the set of all Σ-dags (which we already knew from Theorem 4.1.10). In this latter case, the set of Σ-ACAs that have an equivalent deterministic Σ-ACA is not recursive (Theorem 4.1.10). It is an open question whether this holds for the class of (Σ, k)-dags, too, i.e. whether there is an algorithm that given a Σ-ACA $\mathcal{A}$ and a positive integer k decides whether there exists a deterministic Σ-ACA $\mathcal{A}_d$ such that $L(\mathcal{A}) \cap \mathbb{D}_k = L(\mathcal{A}_d) \cap \mathbb{D}_k$.