

Weak and Nested Class Memory Automata

Conrad Cotton-Barratt^{1,*}, Andrzej S. Murawski^{2,**}, and C.-H. Luke Ong^{1,***}

¹ Department of Computer Science, University of Oxford, UK
{conrad.cotton-barratt,luke.ong}@cs.ox.ac.uk

² Department of Computer Science, University of Warwick, UK
a.murawski@warwick.ac.uk

Abstract. Automata over infinite alphabets have recently come to be studied extensively as potentially useful tools for solving problems in verification and database theory. One popular model of automata studied is the Class Memory Automata (CMA), for which the emptiness problem is equivalent to Petri Net Reachability. We identify a restriction – which we call weakness – of CMA, and show that they are equivalent to three existing forms of automata over data languages. Further, we show that in the deterministic case they are closed under all Boolean operations, and hence have an EXPSPACE-complete equivalence problem. We also extend CMA to operate over multiple levels of nested data values, and show that while these have undecidable emptiness in general, adding the weakness constraint recovers decidability of emptiness, via reduction to coverability in well-structured transition systems. We also examine connections with existing automata over nested data.

Keywords: Automata, Data Languages, Nested Data

1 Introduction

A data word is a word over a finite alphabet in which every position in the word also has an associated *data value*, from an infinite domain. Data languages provide a useful formalism both for problems in database theory and verification [13, 16, 3]. For example, data words can be used to model a system of a potentially unbounded number of concurrent processes: the data values are used as identifiers for the processes, and the data word then gives an interleaving of the actions of the processes. Having expressive, decidable logics and automata over data languages then allows properties of the modelled system to be checked.

Class memory automata (CMA) [3] are a natural form of automata over data languages. CMA can be thought of as finite state machines extended with the ability, on reading a data value, to remember what state the automaton was in when it last saw that data value. A run of a CMA is accepting if the following two conditions hold: (i) the run ends in a *globally accepting* state; and (ii) each

* Supported by an EPSRC Doctoral Training Grant

** Supported by EPSRC (EP/J019577/1)

*** Partially supported by Merton College Research Fund

data value read in the run was last seen in a *locally accepting* state. If using data values to distinguish semi-autonomous parts of a system, while the first condition can check the system as a whole has behaved correctly, the second of these conditions can be used to check that each part of the system independently behaved correctly. The emptiness problem for CMA is equivalent to Petri net reachability, and while closed under intersection, union, and concatenation, they are not closed under complementation, and do not have a decidable equivalence problem.

We earlier described how data words can be used to model concurrent systems: each process can be identified by a data value, and CMA can then verify properties of the system. What happens when these processes can spawn subprocesses, which themselves can spawn subprocesses, and so on? In these situations the parent-child relationship between processes becomes important, and a single layer of data values cannot capture this; instead we want a notion of *nested* data values, which themselves contain the parent-child relationship. In fact, such nested data values have applications beyond just in concurrent systems: they are prime candidates for modelling many computational situations in which names are used hierarchically. This includes higher-order computation where intermediate functional values are being created and named, and later used by referring to these names. More generally, this feature is characteristic of numerous encodings into the π -calculus [14].

This paper is concerned with finding useful automata models which are expressive enough to decide properties we may wish to verify, as well as having good closure and decidability properties, which make them easy to abstract our queries to. We study a restriction of class memory automata, which we find leads to improved complexity and closure results, at the expense of expressivity. We then extend class memory automata to a nested data setting, and find a decidable class of automata in this setting.

Contributions. In Section 3 we identify a natural restriction of Class Memory Automata, which we call *weak* Class Memory Automata, in which the local-acceptance condition of CMA is dropped. We show that these weak CMA are equivalent to: (i) Class Counting Automata, which were introduced in [12]; (ii) non-reset History Register Automata, introduced in [18]; and (iii) locally prefix-closed Data Automata, introduced in [5].

These automata have an EXPSPACE-complete emptiness problem. The primary advantage of having this equivalent model as a kind of Class Memory Automaton is that there is a natural notion of determinism, and we show that Deterministic Weak CMA are closed under all Boolean operations (and hence have decidable containment and equivalence problems).

In Section 4 we introduce a new notion of nesting for data languages, based on tree-structured datasets. This notion does not commit all letters to be at the same level of nesting and appears promising from the point of view of modelling scenarios with hierarchical name structure, such as concurrent or higher-order computation. We extend Class Memory Automata to operate over these nested datasets, and show that this extension is Turing-powerful in general,

but reintroducing the Weakness constraint recovers decidability. We show how these Nested Data CMA recognise the same string languages as Higher-Order Multicounter Automata, introduced in [2], and also how the weakness constraint corresponds to a natural weakness constraint on these Higher-Order Multicounter Automata. Finally, we show these automata to be equivalent to the Nested Data Automata introduced in [5].

Related Work. Class memory automata are equivalent to data automata (introduced in [4]), though unlike data automata, they admit a notion of determinism. Data automata (and hence class memory automata) were shown in [4] to be equiexpressive with the two-variable fragment of existential monadic second order logic over data words. Temporal logics have also been studied over data words [6], and the introduction of locally prefix-closed data automata and of nested data automata in [5] is motivated by extensions to BD-LTL, a form of LTL over multiple data values introduced in [10].

Fresh register automata [17] are a precursor to the History Register automata [18] which we examine a restriction of in this paper. Class counting automata, which we show to be equivalent to weak CMA, have been extended to be as expressive as CMA by adding resets and counter acceptance conditions [12, 11].

We note that our restriction of class memory automata, which we call weak class memory automata, sound similar to the weak data automata introduced in [9]. However, these are two quite different restrictions, with emptiness problems of different complexities, and the two automata models should not be confused.

In the second part of this paper we examine automata over nested data values. First-order logic over nested data values has been studied in [3], where it was shown that the $<$ predicate quickly led to undecidability, but that only having the $+1$ predicate preserved decidability. They also examined the link between nested data and shuffle expressions. In [5] Decker et al. introduced ND-LTL, extending BD-LTL to nested data values. To show decidability of certain fragments of ND-LTL they extended data automata to run over nested data values, giving the nested data automata we examine in this paper.

Further Work. We would like to understand better whether there is a natural fragment of the π -calculus that corresponds to the new classes of automata. On the logical side, an interesting outstanding question is to characterize languages accepted by our classes of automata with suitable logics.

2 Preliminaries

Let Σ be a finite alphabet, and \mathcal{D} an infinite set of data values. A *data alphabet*, \mathbb{D} , is of the form $\Sigma \times \mathcal{D}$. The set of finite data words over \mathbb{D} is denoted \mathbb{D}^* .

Class Memory Automata and Data Automata. Given a set S , we write S_{\perp} to mean $S \cup \{\perp\}$, where \perp is a distinguished symbol (representing a fresh data value). A *Class Memory Automaton* [3] is a tuple $\langle Q, \Sigma, q_I, \delta, F_L, F_G \rangle$ where Q is a finite set of states, Σ is a finite alphabet, $q_I \in Q$ is the initial state, $F_G \subseteq F_L \subseteq Q$ are sets of globally- and locally-accepting sets (respectively), and δ is the transition map $\delta : Q \times \Sigma \times Q_{\perp} \rightarrow \mathcal{P}(Q)$. The automaton is deterministic

if each set in the image of the transition function is a singleton. A *class memory function* is a map $f : \mathcal{D} \rightarrow Q_\perp$ such that $f(d) \neq \perp$ for only finitely many $d \in \mathcal{D}$. We view f as a record of the history of computation: it holds the state of the automaton after the data value d was last read, where $f(d) = \perp$ means that d is fresh. A configuration of the automaton is a pair (q, f) where $q \in Q$ and f is a class memory function. The initial configuration is (q_0, f_0) where $f_0(d) = \perp$ for every $d \in \mathcal{D}$. Suppose $(a, d) \in \Sigma \times \mathcal{D}$ is the input. The automaton can transition from configuration (q, f) to configuration (q', f') just if $q' \in \delta(q, a, f(d))$ and $f' = f[d \mapsto q']$. A data word w is *accepted* by the automaton just if the automaton can make a sequence of transitions from the initial configuration to a configuration (q, f) where $q \in F_G$ and $f(d) \in F_L \cup \{\perp\}$ for every data value d .

A *Data Automaton* [4] is a pair $(\mathcal{A}, \mathcal{B})$ where \mathcal{A} is a letter-to-letter string transducer with output alphabet Γ , called the Base Automaton, and \mathcal{B} is a NFA with input alphabet Γ , called the Class Automaton. A data word $w = w_1 \dots w_n \in \mathbb{D}^*$ is accepted by the automaton if there is a run of \mathcal{A} on the string-projection of w (to Σ) with output $b_1 \dots b_n$ such that for each maximal set of positions $\{x_1, \dots, x_k\} \subseteq \{1, \dots, n\}$ such that w_{x_1}, \dots, w_{x_k} share the same data value, the word $b_{x_1} \dots b_{x_k}$ is accepted by \mathcal{B} .

CMA and DA are expressively equivalent, with PTIME translation [3]. The emptiness problem for these automata is decidable, and equivalent to Petri Net Reachability [4]. The class of languages recognised by CMA is closed under intersection, union, and concatenation. It is not closed under complementation or Kleene star. Of the above, the class of languages recognised by deterministic CMA is closed only under intersection.

Locally Prefix-Closed Data Automata. A Data Automaton $\mathcal{D} = (\mathcal{A}, \mathcal{B})$ is locally prefix-closed (pDA) [5] if all states in \mathcal{B} are final. The emptiness problem for pDA is EXPSpace-complete [5].

Class Counting Automata. A *bag* over \mathcal{D} is a function $h : \mathcal{D} \rightarrow \mathbb{N}$ such that $h(d) = 0$ for all but finitely many $d \in \mathcal{D}$. Let $C = \{=, \neq, <, >\} \times \mathbb{N}$, which we call the set of constraints. If $c = (\text{op}, e) \in C$ and $n \in \mathbb{N}$ we write $n \models c$ iff $n \text{ op } e$. A *Class Counting Automaton* (CCA) [12] is a tuple $\langle Q, \Sigma, \Delta, q_0, F \rangle$ where Q is a finite set of states, Σ is a finite alphabet, q_0 is the initial state, $F \subseteq Q$ is the set of accepting states, and Δ , the transition relation, is a finite subset of $Q \times \Sigma \times C \times \{\uparrow^+, \downarrow\} \times \mathbb{N} \times Q$. A configuration of a CCA, $\mathcal{C} = \langle Q, \Sigma, \Delta, q_0, F \rangle$, is a pair (q, h) where $q \in Q$ and h is a bag. The initial configuration is (q_0, h_0) where h_0 is the zero function. Given a data word $w = (a_1, d_1)(a_2, d_2) \dots (a_n, d_n)$ a run of w on \mathcal{C} is a sequence of configurations $(q_0, h_0)(q_1, h_1) \dots (q_n, h_n)$ such that for all $0 \leq i < n$ there is a transition (q, a, c, π, m, q') where $q = q_i$, $q' = q_{i+1}$, $a = a_{i+1}$, $h_i(d_{i+1}) \models c$, and $h_{i+1} = h_i[d_{i+1} \mapsto h_i(d_{i+1}) + m]$ if $\pi = \uparrow^+$ or $h_{i+1} = h_i[d_{i+1} \mapsto m]$ if $\pi = \downarrow$. The run is accepting if $q_n \in F$. The emptiness problem for Class Counting Automata was shown to be EXPSpace-complete in [12].

Non-Reset History Register Automata. For a positive integer k write $[k]$ for the set $\{1, 2, \dots, k\}$. Fixing a positive integer m , define the set of labels $\text{Lab} = \mathcal{P}([m])^2$. A non-reset History Register Automaton (nrHRA) of type m

with initially empty assignment is a tuple $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, F \rangle$ where $q_0 \in Q$ is the initial state, $F \subseteq Q$ is the set of final states, and $\delta \subseteq Q \times \Sigma \times \text{Lab} \times Q$. A configuration of \mathcal{A} is a pair (q, H) where $q \in Q$ and $H : [m] \rightarrow \mathcal{P}_{fn}(\mathcal{D})$ where $\mathcal{P}_{fn}(\mathcal{D})$ is the set of finite subsets of \mathcal{D} . We call H an *assignment*, and for $d \in \mathcal{D}$ we write $H^{-1}(d)$ for the set $\{i \in [m] : d \in H(i)\}$. The initial configuration is (q_0, H_0) , where H_0 assigns every integer in $[m]$ to the empty set. When the automaton is in configuration (q, H) , on reading input (a, d) it can transition to configuration (q', H') providing there exists $X \subseteq [m]$ such that $(q, a, (H^{-1}(d), X), d) \in \delta$ and H' is obtained by removing d from $H(i)$ for each i then adding d to each $H(i)$ such that $i \in X$. A run is defined in the usual way, and a run is accepting if it ends in a configuration (q, H) where $q \in F$.

Higher-Order Multicounter Automata. A multiset over a set A is a function $m : A \rightarrow \mathbb{N}$. A level-1 multiset over A is a finite multiset over A . A level- $(k+1)$ multiset over A is a finite multiset of level- k multisets over A . We can visualise this with nested set notation: e.g. $\{\{a, a\}, \{\}, \{\}\}$ represents the level-2 multiset containing one level-1 multiset containing two copies of a , and two empty level-1 multisets. A multiset is *hereditarily empty* if, written in nested set notation, it contains no symbols from A .

Higher-Order Multicounter Automata (HOMCA) were introduced in [2], and their emptiness problem was shown to be Turing-complete at level-2 and above. A level- k multicounter automaton is a tuple $\langle Q, \Sigma, A, \Delta, q_0, F \rangle$ where Q is a finite set of states, Σ is the input alphabet, A is the multiset alphabet, q_0 is the initial state, and F is the set of final states. A configuration is a tuple $(q, m_1, m_2, \dots, m_k)$ where $q \in Q$ and each m_i is either undefined (\perp) or a level- i multiset over A . The initial configuration is $(q_0, \perp, \dots, \perp)$. Δ is the transition relation, and is a subset of $Q \times \Sigma \times ops \times Q$ where ops is the set of possible counter operations. These operations, and meanings, are as follows: (i) *new_i* ($i \leq k$) turns m_i from \perp into the empty level- i multiset; (ii) *inc_a* ($a \in A$) adds a to m_1 ; (iii) *dec_a* ($a \in A$) removes a from m_1 ; (iv) *store_i* ($i < k$) adds m_i to m_{i+1} and sets m_i to \perp ; (v) *load_i* ($i < k$) non-deterministically removes an m from m_{i+1} and turns m_i from \perp to m . This can happen only when $m_1 \dots m_i$ are all \perp . The automaton reads the input word from left to right, updating $m_1 \dots m_k$ as determined by the transitions. A word is accepted by the automaton just if there is a run of the word such that the automaton ends up in configuration (q, m_1, \dots, m_k) where $q \in F$ and each m_i is hereditarily empty.

3 Weak Class Memory Automata

In this section we introduce a restriction of class memory automata, *weak* class memory automata (WCMA), that have improved closure and complexity properties. We show that WCMA correspond to a natural restriction of data automata, locally-prefix closed data automata, as well as two other independent automata models, class counting automata and non-reset history register automata.

Definition 1. A class memory automaton $\langle Q, \Sigma, \Delta, q_0, F_L, F_G \rangle$ is weak if all states are locally accepting (i.e. $F_L = Q$).

When defining a weak CMA (WCMA) we may omit the set of locally accepting states, and just give one set of final states, F .

The emptiness problem for class memory automata is reducible (in fact, equivalent) to emptiness of multicounter automata (MCA) [4, 3]. This reduction works by using counters to store the number of data values last seen in each state. The local-acceptance condition is checked by the zero-test of each counter at the end of a run of an MCA. In the weak CMA case, this check is not necessary, and so emptiness is reducible to emptiness of weak MCA. Just as MCA emptiness is equivalent to Petri net reachability, weak MCA emptiness is equivalent to Petri net coverability.

Example 2. We give an example showing how a very simple Petri net reachability query can be reduced to an emptiness of CMA problem, and the small change required to reduce coverability queries to emptiness of WCMA. The idea is to encode tokens in the Petri net using data values: the location of the token is stored by the class memory function's memory for the data value. Transitions in the Petri net will be simulated by sequences of transitions in the automaton, which change class memory function appropriately. Consider the Petri net shown in Figure 1, with initial marking above and target marking below.




Fig. 1. An example Petri net with initial and target markings

We give the automaton which models this reachability query in Figure 2. The first transitions from the initial state just set up the initial marking. As there is only one token in the initial marking, this just involves reading one fresh data value: this is the transition from s_0 to s_1 below. Once the initial marking has been set up (reaching s_2 below), the automaton can simulate the transitions firing any number of times. Each loop from s_2 back to itself represents one transition in the Petri




Fig. 2. A class memory automaton simulating the Petri net query shown in Figure 1

net firing: the loop above represents t_1 firing, and the loop below represents t_2 firing. For t_1 to fire, no preconditions must be met, and a new data value can be read in state s_3 , thus data values last seen in either of states s_1 and s_3 represent tokens in p_1 . For t_2 to fire, a token must be removed from p_1 , since tokens in p_1 are represented by tokens in either s_1 or s_3 , the first transition in this loop – to s_4 – involves reading a data value last seen in one of these states. Thus data values seen in s_4 represent removed tokens, which we do not use again. Then a

new token is placed in p_2 by reading a fresh data value in s_5 . Once back in s_2 these loops can be taken more, or the fact that a marking covering the target marking has been reached can be checked by reading two data values last seen in s_5 to reach a final state. The only globally accepting state is s_7 , and all states except those which are used to represent tokens – i.e. all except s_1 , s_3 , and s_5 – are locally accepting. The local acceptance condition thus checks that no other tokens remain in the simulated Petri net.

If we were interested in a coverability query, the same automaton, but without the local acceptance condition, would obviously suffice. Thus emptiness of WCMA is equivalent to Petri net coverability, which is EXPSPACE-complete.

We now give the main observation of this section: that weak CMA are equivalent to three independent existing automata models.

Theorem 3. *Weak CMA, locally prefix-closed DA, class counting automata, and non-reset history register automata are all PTIME-equivalent.*

Proof. That Weak CMA and pDA are equivalent is a simple alteration of the proof of equivalence of CMA and DA provided in [3].

Recall that CCA use a “bag”, which essentially gives a counter for each data value. Weak CMA can easily be simulated by CCA by identifying each state with a natural number; then the bag can easily simulate the class memory function, by setting the data value’s counter to the appropriate number when it is read. To simulate a CCA with a WCMA, we first observe that for any CCA, since counter values can only be incremented or reset, there is a natural number, N , above which different counter values are indistinguishable to the automaton. Thus we need only worry about a finite set of values. This means the value for the counter of each data value can be stored in the automaton state, and thereby the class memory function.

In [18] the authors already show that nrHRA can be simulated by CMA. Their construction does not make use of the local-acceptance condition, so the fact that nrHRA can be simulated by WCMA is immediate. In order to simulate a given WCMA with state set $[m]$, one can take a nrHRA of type m , with place i storing the data values last seen in state i .

Data automata, and hence pDA, unlike CMA and WCMA, do not have a natural notion of determinism, nor a natural restriction corresponding to deterministic CMA or WCMA. What about for CCA? We define CCA to be deterministic if for each state q and input letter a , the transitions $(q, a, c, \dots) \in \Delta$ are such that the c ’s partition \mathbb{N} . The same translations given in Theorem 3 also show that deterministic WCMA and deterministic CCA are equivalent. We can ask the same question of non-reset HRA. We find that the natural notion of determinism here is: for each $q \in Q$, $a \in \Sigma$, and $X \subseteq [m]$ there is precisely one $Y \subseteq [m]$ and $q' \in Q$ such that $(q, a, (X, Y), q') \in \delta$. Similarly, the translations discussed above show deterministic WCMA to be equivalent to deterministic nrHRA.

It follows from the results for CCA in [12] that Weak CMA, like normal CMA, are closed under intersection and union, though these closures can easily be shown directly using product constructions (and these constructions preserve

determinism). In fact, Deterministic Weak CMA have even nicer closure properties: the language recognised can be complemented using the same method as for DFA: complementing the final states.

Proposition 4. *Deterministic Weak CMA are closed under all Boolean operations.*

Corollary 5. *The containment and equivalence problems for Deterministic Weak CMA are EXPSPACE-complete.*

4 Nested Data Class Memory Automata

In Section 1 we discussed how data values fail to provide a good model for modelling computations in which names are used hierarchically, such as a system of concurrent processes which can spawn subprocesses. Motivated by these applications, in this section we introduce a notion of nested data values in which the data set has a forest-structure. This is a stylistically different presentation to earlier work on nested data in that [3, 5] require that each position in the words considered have a data value in each of a fixed number of levels. By giving the data set a forest-structure, we can explicitly handle variable levels of nesting within a word. However, we note that there is a natural translation between the two presentations.

Definition 6. *A rooted tree (henceforth, just tree) is a simple directed graph $\langle D, \text{pred} \rangle$, where $\text{pred} : D \rightarrow D$ is the predecessor map defined on every node of the tree except the root, such that every node has a unique path to the root. A node n of a tree has level l just if $\text{pred}^{l-1}(n)$ is the root (thus the root has level 1). A tree has bounded level just if there exists a least $l \geq 1$ such that every node has level no more than l ; we say that such a tree has level l .*

We define a nested dataset $\langle \mathcal{D}, \text{pred} \rangle$ to be a forest of infinitely many trees of level l which is full in the sense that for each data value d of level less than l , d has infinitely many children (i.e. there are infinitely many data values d' s.t. $\text{pred}(d') = d$).

We now extend CMA to nested data by allowing the nested data class memory automaton (NDCMA), on reading a data value d , to access the class memory function's memory of not only d , but each ancestor of d in the nested data set. Once a transition has been made, the class memory function updates the remembered state not only of d , but also of each of its ancestors. Formally:

Definition 7. *Fix a nested data set of level l . A Nested Data CMA of level l is a tuple $\langle Q, \Sigma, \delta, q_0, F_L, F_G \rangle$ where Q is a finite set of states, $q_0 \in Q$ is the initial state, $F_G \subseteq F_L \subseteq Q$ are sets of globally and locally accepting states respectively, and δ is the transition map. δ is given by a union $\delta = \bigcup_{1 \leq i \leq l} \delta_i$ where each δ_i is a function:*

$$\delta_i : Q \times \Sigma \times (\{i\} \times (Q_{\perp})^i) \rightarrow \mathcal{P}(Q)$$

The automaton is deterministic if each set in the image of δ is a singleton; and is weak if $F_L = Q$. A configuration is a pair (q, f) where $q \in Q$, and $f : \mathcal{D} \rightarrow Q_\perp$ is a class memory function (i.e. $f(d) = \perp$ for all but finitely many $d \in \mathcal{D}$). The initial configuration is (q_0, f_0) where f_0 is the class memory function mapping every data value to \perp . A configuration (q, f) is final if $q \in F_G$ and $f(d) \in F_L \cup \{\perp\}$ for all $d \in \mathcal{D}$. The automaton can transition from configuration (q, f) to configuration (q', f') on reading input (a, d) just if d is a level- i data value, $q' \in \delta(q, a, (i, f(\text{pred}^{i-1}(d)), \dots, f(\text{pred}(d)), f(d)))$, and $f' = f[d \mapsto q, \text{pred}(d) \mapsto q, \dots, \text{pred}^{i-1}(d) \mapsto q]$. A run $(q_0, f_0), (q_1, f_1), \dots, (q_n, f_n)$ is accepting if the configuration (q_n, f_n) is final. $w \in L(\mathcal{A})$ if there is an accepting run of \mathcal{A} on w .

It is clear that level-1 NDCMA are equivalent to normal CMA. We know that emptiness of class memory automata is equivalent to reachability in Petri nets; it is natural to ask whether there is any analogous correspondence – to some kind of high-level Petri net – once nested data is used.

Example 8. In Example 2, we showed how CMA (resp. WCMA) can encode Petri net reachability (resp. coverability). A similar technique allows reachability (resp. coverability) of Petri nets with reset arcs to be reduced to emptiness of NDCMA (resp. weak NDCMA). The key idea is to have, for each place in the net, a level-1 data value – essentially as a “bag” holding the tokens for that place. Nested under the level-1 data value, level-2 data values are used to represent tokens just as before. When a reset arc is fired, the corresponding level-1 data value is moved to a “dead” state – from where it and the data values nested under it are not moved again – and a fresh level-1 data value is then used to hold subsequently added tokens to that place.

Theorem 9. *The emptiness problem for NDCMA is undecidable. Emptiness of Weak NDCMA is decidable, but Ackermann-hard.*

Proof. This result follows from Theorem 12 together with results in [5], though we also provide a direct proof.

We show decidability by reduction to a well-structured transition system [8] constructed as follows: a class memory function on a nested data set can be viewed as a labelling of the data set by labels from the set of states. Since we only care about the shape of the class memory function (i.e. up to automorphisms of the nested data set), we can remove the nodes labelled by \perp , and view a class memory function as a finite set of labelled trees. The set of finite forests of finite trees of bounded depth with the order given by $F \leq F'$ iff there is a forest homomorphism from F to F' (where a forest is the natural extension of tree homomorphisms to forests) is a well-quasi-order [7], which provides the basis for the well-structured transition system.

Undecidability for NDCMA and Ackermann-hardness for Weak NDCMA follow from the ideas in Example 8: the reachability (resp. coverability) problem for Petri nets with reset arcs is encodable in NDCMA (resp. Weak NDCMA), and this is undecidable [1] (resp. Ackermann-hard [15]).

Weak Nested Data CMA have similar closure properties to Weak CMA (and this can be shown using the same techniques as for DFA).

Proposition 10. (i) *Weak NDCMA are closed under intersection and union.*
(ii) *Deterministic Weak NDCMA are closed under all Boolean operations.*

Hence, as for weak CMA, the containment and equivalence problems for Deterministic Weak NDCMA are decidable.

4.1 Link with Nested Data Automata

In [5] Decker et al. also examined “Nested Data Automata” (NDA), and showed the locally prefix-closed NDA (pNDA) to have decidable emptiness (via reduction to well-structured transition systems). In fact, these NDA precisely correspond to NDCMA, and again being locally prefix-closed corresponds to weakness. In this section we briefly outline this connection.

Nested Data Automata. ([5]) A k -nested data automaton (k -NDA) is a tuple $(\mathcal{A}, \mathcal{B}_1, \dots, \mathcal{B}_k)$ where $(\mathcal{A}, \mathcal{B}_i)$ is a data automaton for each i . Such automata run on words over the alphabet $\Sigma \times \mathcal{D}^k$, where \mathcal{D} is a (normal, unstructured) dataset. As for normal data automata, the transducer, \mathcal{A} , runs on the string projection of the word, giving output w . Then for each i the class automaton \mathcal{B}_i runs on each subsequence of w corresponding to the positions which agree on the first i data values. The NDA is locally prefix-closed if each $(\mathcal{A}, \mathcal{B}_i)$ is.

Since these NDA are defined on a slightly different presentation of nested data, we provide the following presentation of NDCMA over multiple levels of data.

Definition 11. A Nested Data CMA of level k over the alphabet $\Sigma \times \mathcal{D}^k$ is a tuple $\langle Q, \Sigma, \delta, q_0, F_L, F_G \rangle$ where Q is a finite set of states, $q_0 \in Q$, $F_G \subseteq F_L \subseteq Q$, and $\delta : Q \times \Sigma \times (Q_\perp)^k \rightarrow \mathcal{P}(Q)$ is the transition map.

A configuration is a tuple $(q, f_1, f_2, \dots, f_k)$, where each $f_i : \mathcal{D}^i \rightarrow Q_\perp$ maps an i -tuple of data values to a state in the automaton (or \perp). The initial configuration is $(q_0, f_1^0, \dots, f_k^0)$ where f_i^0 maps every tuple in the domain to \perp . A configuration (q, f_1, \dots, f_k) is final if each f_i maps into $F_L \cup \{\perp\}$. The automaton can transition from configuration (q, f_1, \dots, f_k) to configuration (q', f'_1, \dots, f'_k) on reading input (a, d_1, \dots, d_k) just if $q' \in \delta(q, a, (f_1(d_1), f_2(d_1, d_2), \dots, f_k(d_1, \dots, d_k)))$, and each $f'_i = f_i[(d_1, \dots, d_i) \mapsto q']$.

Using ideas from the proof of equivalence between CMA and DA in [3], we can show the following result:

Theorem 12. *NDCMA (resp. weak NDCMA) and NDA (resp. pNDA) are expressively equivalent, with effective translations.*

4.2 Link with Higher-Order Multicounter Automata

In [2] the authors examined a link between nested data values and shuffle expressions. In doing so, they introduced higher-order multicounter automata

(HOMCA). While not explicitly over nested data values, they are closely related to the ideas involved, and we show that, just as multicounter automata and CMA are equivalent, there is a natural translation between HOMCA and the NDCMA we have introduced. Further, just as the equivalence between MCA and CMA descends to one between weak multicounter automata and weak CMA, we find an equivalence between weak NDCMA and “weak” HOMCA in which the corresponding acceptance condition – hereditary emptiness – is dropped. To show this, we introduce HOMCA’, which add restrictions to the $store_i$ and new_i counter operations analogous to the restriction for the $load_i$ operation. We show that these HOMCA’ are equivalent to HOMCA, and that HOMCA’ are equivalent to NDCMA, with both of these equivalences descending to the weak versions. These equivalences are summarised in Figure 3.

Definition 13. We define weak HOMCA to be just as HOMCA, but without the hereditary-emptiness condition on acceptance, i.e. a run is accepting just if it ends in a final state.

Definition 14. We define HOMCA’ to be the same as HOMCA, except for the following changes to the counter operations: (i) $store_i$ operations are only enabled when $m_1 = m_2 = \dots = m_{i-1} = \perp$; and (ii) new_i operations are only enabled when $m_k \neq \perp, m_{k-1} \neq \perp, \dots, m_{i+1} \neq \perp$ and $m_{i-1} = m_{i-2} = \dots = m_1 = \perp$.

As for HOMCA, we define weak HOMCA’ to be HOMCA’ without the hereditary-emptiness condition.

This means that each reachable configuration (q, m_1, \dots, m_k) is such that there is a unique $0 \leq i \leq k$ such that for all $j \leq i$, $m_j = \perp$ and each $l > i$, $m_l \neq \perp$.

Theorem 15. HOMCA (respectively weak HOMCA) and HOMCA’ (resp. weak HOMCA’) are expressively equivalent, with effective translations between them.

Proof. This requires simulating the HOMCA operations $store_i$ and new_i in HOMCA’: which can be difficult if, for instance, the HOMCA is carrying out a $store_i$ operation when it has a current level- $(i-1)$ multiset in memory. The trick is to move the level- $(i-1)$ multiset across to be nested under a new level- i multiset, and this can be done one element at a time in a “folding-and-unfolding” method. The hereditary emptiness condition checks that the each of these movements was completed, i.e. no element was left unmoved. In the weak case some elements not being moved could not change an accepting run to a non-accepting run, so the fallibility of the moving method does not matter.

Theorem 16. For every (weak) level- k NDCMA, \mathcal{A} , there is a (weak) level- k HOMCA’, \mathcal{A}' , such that $\mathcal{L}(\mathcal{A}')$ is equal to the Σ -projection of $\mathcal{L}(\mathcal{A})$, and vice-versa.




Fig. 3. Translations between HOMCA, HOMCA’, NDCMA, and their weak counterparts

Proof. This proof rests on the strong similarity between the nesting of data values, and the nesting of level- i multisets in level- $(i + 1)$ multisets. For a NDCMA to simulate a HOMCA', we use level- k data values to represent instances of the multiset letters, level- $(k - 1)$ data values to represent level-1 multisets, and so on, up to level-1 data values representing level- $(k - 1)$ multisets. Since each run of a HOMCA' can have at most one level- k multiset, this does not need to be encoded in data values. Conversely, when simulating a NDCMA with a HOMCA', a level- k data value is represented by an instance of an appropriate multiset letter. The letter contains the information on which state the data value was last seen in. Level- $(k - 1)$ data values are represented by level-1 multisets, which also include a multiset letter storing the state that data value was last seen in.

References

1. Araki, T., Kasami, T.: Some decision problems related to the reachability problem for Petri nets. *Theor. Comput. Sci.* 3(1), 85–104 (1976)
2. Björklund, H., Bojanczyk, M.: Shuffle expressions and words with nested data. In: *Proceedings of MFCS. LNCS*, vol. 4708, pp. 750–761 (2007)
3. Björklund, H., Schwentick, T.: On notions of regularity for data languages. *Theor. Comput. Sci.* 411(4-5), 702–715 (2010)
4. Bojanczyk, M., Muscholl, A., Schwentick, T., Segoufin, L., David, C.: Two-variable logic on words with data. In: *LICS*. pp. 7–16. IEEE Computer Society (2006)
5. Decker, N., Habermehl, P., Leucker, M., Thoma, D.: Ordered navigation on multi-attributed data words. In: *CONCUR 2014*. pp. 497–511. Springer (2014)
6. Demri, S., Lazic, R.: LTL with the freeze quantifier and register automata. *ACM Trans. Comput. Log.* 10(3) (2009)
7. Ding, G.: Subgraphs and well-quasi-ordering. *J. Graph Theory* 16, 489–502 (1992)
8. Finkel, A., Schnoebelen, P.: Well-structured transition systems everywhere! *Theor. Comput. Sci.* 256(1-2), 63–92 (2001)
9. Kara, A., Schwentick, T., T.: Feasible automata for two-variable logic with successor on data words. In: *LATA. LNCS*, vol. 7183, pp. 351–362. Springer (2012)
10. Kara, A., Schwentick, T., Zeume, T.: Temporal logics on words with multiple data values. In: *FSTTCS. LIPIcs*, vol. 8, pp. 481–492. Schloss Dagstuhl (2010)
11. Manuel, A.: Counter automata and classical logics for data words. Ph.D. thesis, Institute of Mathematical Sciences, Chennai (2011)
12. Manuel, A., Ramanujam, R.: Counting multiplicity over infinite alphabets. In: *RP. LNCS*, vol. 5797, pp. 141–153. Springer (2009)
13. Neven, F., Schwentick, T., Vianu, V.: Finite state machines for strings over infinite alphabets. *ACM Trans. Comput. Log.* 5(3), 403–435 (2004)
14. Sangiorgi, D.: Expressing Mobility in Process Algebras: First-Order and Higher-Order Paradigms. Ph.D. thesis, University of Edinburgh (1992)
15. Schnoebelen, P.: Verifying lossy channel systems has nonprimitive recursive complexity. *Inf. Process. Lett.* 83(5), 251–261 (2002)
16. Segoufin, L.: Automata and logics for words and trees over an infinite alphabet. In: *CSL. LNCS*, vol. 4207, pp. 41–57. Springer (2006)
17. Tzevelekos, N.: Fresh-register automata. In: *POPL 2011*. pp. 295–306. ACM (2011)
18. Tzevelekos, N., Grigore, R.: History-register automata. In: *FoSSaCS. Lecture Notes in Computer Science*, vol. 7794, pp. 17–33. Springer (2013)