Representing Regular Languages of Infinite Words Using Mod 2 Multiplicity Automata

Dana Angluin[1], Timos Antonopoulos[1](✉), Dana Fisman[2], and Nevin George[1]

[1] Yale University, New Haven, CT, USA
timos.antonopoulos@yale.edu
[2] Ben-Gurion University, Beer-Sheva, Israel

Abstract. We explore the suitability of mod 2 multiplicity automata (M2MAs) as a representation for regular languages of infinite words. M2MAs are a deterministic representation that is known to be learnable in polynomial time with membership and equivalence queries, in contrast to many other representations. Another advantage of M2MAs compared to non-deterministic automata is that their equivalence can be decided in polynomial time and complementation incurs only an additive constant size increase. Because learning time is parameterized by the size of the representation, particular attention is focused on the relative succinctness of alternate representations, in particular, LTL formulas and Büchi automata of the types: deterministic, non-deterministic and strongly unambiguous. We supplement the theoretical results of worst case upper and lower bounds with experimental results computed for randomly generated automata and specific families of LTL formulas.

Keywords: Multiplicity Automata · Regular Omega Languages · Büchi Automata · Linear Temporal Logic · Conciseness

1 Introduction

Regular languages of infinite words (or ω-words) play an important role in verification of reactive systems. The question of whether a system S satisfies a specification given by a temporal logic formula φ can be reduced to the question of whether $L(S) \cap L(\neg\varphi)$ is empty, where $L(S)$ is the set of ω-words representing the computation paths of the system S and $L(\neg\varphi)$ is the set of ω-words representing computations that violate φ. Automata are a useful machinery for performing operations on languages such as complementation and intersection, and for deciding properties such as emptiness and equivalence. Many verification tools are implemented using reductions to automata [20].

Regular ω-languages can be represented using various types of automata (e.g. Büchi, Rabin, Parity, etc.). Different automata types differ in their succinctness and in the complexity of performing operations of interest. Non-deterministic Büchi automata (NBAs) are one of the most popular acceptor types for regular ω-languages, mainly due to their simplicity, succinctness, and good complexity

for the emptiness problem. An issue with Büchi automata is that their deterministic version (DBAs) is strictly less expressive: while NBAs accept all regular ω-languages, DBAs recognize only a strict subset thereof. Another issue is that complementation of NBAs is hard; it has a $2^{\Omega(n \log n)}$ lower bound (where n is the number of states) [16]. This motivated the introduction of *complete unambiguous Büchi automata* (CUBA) by Carton and Michel who showed that every regular ω-language can be represented by a CUBA, i.e. there is a way to limit the non-determinism without losing expressiveness [8]. Bousquet and Löding proposed *strongly unambiguous Büchi automata* (SUBA), a slight relaxation of CUBA for which they have shown that equivalence can be decided in polynomial time [6].

The SUBA model was also shown useful in terms of learnability of regular ω-languages — Angluin, Antonopoulos and Fisman have shown that SUBAs are polynomially predictable using membership queries (while NBAs, under plausible cryptographic assumptions, are not) [1]. Their proof makes use of a model of automata called *Mod 2 Multiplicity Automata* (M2MA). Informally, *multiplicity automata* are an algebraic variant of automata that compute functions from finite words to a field $\mathcal{K}$ [4,5], and *M2MAs* are multiplicity automata that work over the field $GF(2) = \{0, 1\}$ where sum and product are computed modulo 2.

In this paper we look at questions concerning the adequacy of M2MAs for representing regular ω-languages. We note that M2MAs operate on finite words, and their use for representing regular ω-languages follows a reduction, by Calbrix, Nivat and Podelski from a regular ω-language L to a regular language $(L)_\$$ of finite words [7]. We thus start by reviewing the succinctness of M2MAs with respect to automata on finite words, particularly of types non-deterministic (NFAs), deterministic (DFAs), and unambiguous (UFAs). We show that M2MAs are more succinct than DFAs and UFAs, whereas with respect to NFAs there are in the worst case exponential gaps in going from M2MAs to NFAs and vice versa.

We also study the complexity of performing basic operations on M2MAs; complementation can be done with an additive constant increase in size, and union and intersection with the product of sizes. There is a known cubic algorithm to minimize a weighted automaton [10,19], which applies to an M2MA and also implies cubic procedures for determining emptiness and equivalence.

We then investigate the succinctness of M2MAs in representing regular ω-languages, by comparing translations from *linear temporal logic* (LTL) formulas and Büchi automata (deterministic, non-deterministic and strongly unambiguous) into M2MAs, DFAs, UFAs, SUBAs and NBAs (where the former three use the $(L)_\$$ representation). The results are summarized in Fig. 3.

To complement the theoretical bounds, we implemented procedures to transform SUBAs to UFAs and UFAs to M2MAs, and to minimize and learn M2MAs, and report estimates of the average size increases in transforming random SUBAs, DBAs, and NBAs to M2MAs. We also determine the minimum dimensions of M2MAs and minimum sizes of DFAs for a few members of three specific families of LTL formulas and compare them with the respective ω-automaton sizes.

2   Preliminaries

For nonnegative integers k and ℓ, $[k..\ell]$ is the set of nonnegative integers n such that $k \leq n \leq \ell$. Given a finite alphabet Σ, Σ^* is the set of finite words over Σ. The length of a word x is $|x|$ and the empty word is ε. $\Sigma^n = \{x \in \Sigma^* \mid |x| = n\}$. The reverse of a word x is x^r. A language L is any subset of Σ^*. The reverse of L, denoted L^r, is $\{x^r \mid x \in L\}$. The *Hankel matrix* of a language L is the infinite matrix whose rows and columns are indexed by elements of Σ^*, where the entry for row x and column y is 1 if $xy \in L$ and 0 if $xy \notin L$.

The set of infinite words (or ω-words) over Σ is the set of all maps from the positive integers to Σ and is denoted Σ^ω. An ω-language is any subset of Σ^ω. For a finite or infinite word w, $w[i]$ denotes the symbol at position i, with indices starting at 1. Concatenation of a finite word x with a finite or infinite word y is denoted xy. The word x is a prefix of xy and the word y is a suffix of xy. The suffix of w starting at position i is denoted $w[i :]$. If $x \in \Sigma^*$ and k is a nonnegative integer, x^k denotes the concatenation of k copies of x, and x^ω denotes the concatenation of x with itself infinitely many times. An ω-word is ultimately periodic if it can be written in the form $u(v)^\omega$ for $u, v \in \Sigma^*$ with $|v| > 0$. If A_1 and A_2 are sets and $S \subseteq A_1 \times A_2$, then we define the projection $\pi_1(S) = \{a_1 \mid (\exists a_2)(a_1, a_2) \in S\}$ and analogously for the projection π_2.

2.1   NFAs, UFAs, DFAs, NBAs, UBAs, SUBAs, and DBAs

A (nondeterministic) finite-state automaton A is a tuple $(\Sigma, Q, I, \Delta, F)$ consisting of a finite alphabet Σ, a finite set Q of states, a set $I \subseteq Q$ of initial states, a set $F \subseteq Q$ of final states, and a transition relation $\Delta \subseteq Q \times \Sigma \times Q$. The transition relation Δ is deterministic if for every state $q \in Q$ and every symbol $\sigma \in \Sigma$, there is at most one state $q' \in Q$ such that $(q, \sigma, q') \in \Delta$. The size of a finite-state automaton is $|Q|$.

For a word w, a run of A on w is a sequence of states $q_0, q_1, \ldots$ such that for each i that indexes a symbol in w, $(q_{i-1}, w[i], q_i) \in \Delta$. Thus, for $w \in \Sigma^*$ a run on w is a sequence of length $|w| + 1$, and for $w \in \Sigma^\omega$, a run on w is an infinite sequence of states. A run on w is *initial* if $q_0 \in I$. A finite run is *final* if $q_{|w|} \in F$, and an infinite run is *final* if there are infinitely many values of i for which $q_i \in F$. Acceptors of languages and ω-languages may be defined using finite-state automata, as follows. In each case, the language of words accepted by an acceptor A is denoted $L(A)$.

A *nondeterministic finite acceptor* (NFA) is a finite-state automaton A that accepts a word $w \in \Sigma^*$ if there exists a run of A on w that is both initial and final. An NFA A is an *unambiguous finite acceptor* (UFA) if for every word $w \in L(A)$ there is exactly one run of A on w that is initial and final. An NFA A is a *deterministic finite acceptor* (DFA) if there is exactly one initial state ($|I| = 1$) and the transition relation Δ is deterministic. The languages over Σ that are accepted by NFAs, UFAs, or DFAs is precisely the regular languages over Σ.

A *nondeterministic Büchi acceptor* (NBA) is a finite-state automaton A that accepts a word $w \in \Sigma^\omega$ if there exists a run of A on w that is both initial and

final. An NBA is an *unambiguous Büchi acceptor* (UBA) if for every $w \in L(A)$, there exists exactly one run of A on w that is initial and final. Bousquet and Löding [6] introduced the concept of a *strongly unambiguous Büchi acceptor* (SUBA), which is an NBA such that for every $w \in \Sigma^\omega$, there is at most one final run of the acceptor on w — note that the condition of being initial is dropped. Thus, every SUBA is a UBA. The ω-languages over Σ that are accepted by NBAs, UBAs, or SUBAs are precisely the regular ω-languages. An NBA is a *deterministic Büchi acceptor* (DBA) if there is exactly one initial state ($|I| = 1$) and the transition relation Δ is deterministic. Every DBA is a UBA, but is not necessarily a SUBA. The ω-languages that are accepted by DBAs are a proper subclass of the class of all regular ω-languages.

For Büchi acceptors, we also consider a generalized version, GNBA, in which the acceptance condition is specified not by a single set of final states, but by a collection $\mathcal{F}$ of sets of final states. For a GNBA, a run $q_0, q_1, \ldots$ is *final* iff for each $F \in \mathcal{F}$, there exist infinitely many indices i such that $q_i \in F$. Applying this generalization to a SUBA yields a GSUBA. There is a standard translation of a GNBA of size n with k sets of final states into an NBA of size kn, in which there are k copies of the GNBA automaton. However, applying this construction to a GSUBA does not in general yield a SUBA.

2.2   LTL formulas

The syntax of *linear temporal logic* (LTL) [18] over a set AP of atomic propositions is given by the following grammar $\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \bigcirc\varphi \mid (\varphi_1 \mathcal{U} \varphi_2)$ where $p \in AP$ is an atomic proposition.

The semantics of LTL relates ω-words over 2^{AP} to formulas as shown on the right (recall that indexing of words starts at 1). Additional Boolean and temporal connectives are defined in the usual way. In particular $\top$ (*true*) is defined as $p \vee \neg p$, $\Diamond\varphi$ (*eventually* φ) is defined as $(\top \mathcal{U} \varphi)$ and $\Box\varphi$ (*always* φ) is defined as $\neg\Diamond(\neg\varphi)$.

$$
\begin{aligned}
w &\models p & \text{iff} \quad & p \in w[1] \\
w &\models \neg\varphi & \text{iff} \quad & w \not\models \varphi \\
w &\models \varphi_1 \wedge \varphi_2 & \text{iff} \quad & w \models \varphi_1 \text{ and } w \models \varphi_2 \\
w &\models \bigcirc\varphi & \text{iff} \quad & w[2:] \models \varphi \\
w &\models (\varphi_1 \mathcal{U} \varphi_2) & \text{iff} \quad & \exists j.\ w[j:] \models \varphi_2 \text{ and} \\
& & & \forall i < j.\ w[i:] \models \varphi_1
\end{aligned}
$$

The ω-language of an LTL formula φ, denoted $L(\varphi)$, is the set of ω-words for which it is true. The *size* of an LTL formula φ is the number of distinct subformulas it contains. Every LTL formula represents a regular ω-language (see Section 5). However, not every regular ω-language can be represented by an LTL formula; in particular, the regular ω-languages that can be represented by LTL formulas are noncounting [9].

2.3   M2MAs

A *multiplicity automaton* represents a function mapping Σ^* to elements of a field $\mathcal{K}$. We focus on the case where $\mathcal{K} = \{0, 1\}$ and product and sum are computed modulo 2. A *mod 2 multiplicity acceptor* (M2MA) of dimension d is a tuple $A = (\Sigma, v_I, \{\mu_\sigma\}_{\sigma \in \Sigma}, v_F)$, where Σ is the input alphabet, $v_I \in \mathcal{K}^d$ is the initial

vector, $v_F \in \mathcal{K}^d$ is the final vector, and for each $\sigma \in \Sigma$, μ_σ is a $d \times d$ transition matrix over $\mathcal{K}$, that is, an element of $\mathcal{K}^{d \times d}$.

The vectors v_I and v_F are interpreted as $d \times 1$ column vectors. The transpose operation is denoted by $\top$, and the inner product of two column vectors $v, w \in \mathcal{K}^d$ is denoted $v^\top w$.

To define $L(A)$ we inductively define the matrix μ_x for all $x \in \Sigma^*$. If $x = \varepsilon$, then μ_x is the $d \times d$ identity matrix. If $x = \sigma y$ for some $\sigma \in \Sigma$ and $y \in \Sigma^*$ then $\mu_x = \mu_\sigma \mu_y$. The function $f_A : \Sigma^* \to \mathcal{K}$ computed by A is defined by $f_A(x) = v_I^\top \mu_x v_F$.   A word x is *accepted* by A if $f_A(x) = 1$.

We refer to column vectors $v \in \mathcal{K}^d$ as *states* or *co-states* of A. A state v is *reachable* iff there exists a word $x \in \Sigma^*$ such that $v = (v_I^\top \mu_x)^\top$. A co-state w is *co-reachable* iff there exists a word $x \in \Sigma^*$ such that $w = \mu_x v_F$. For any state v, $L_v(A)$ denotes the language of words accepted by A with its initial vector replaced by v.

We assume standard results from finite dimensional vector spaces. If U is a vector space of dimension k over the field $\{0, 1\}$ then $|U| = 2^k$. If U is a vector subspace of the vector space V, then the *orthogonal complement* of U is the set $U^\perp = \{v \mid v^\top u = 0 \ \forall u \in U\}$, $U^\perp$ is a vector subspace of V which is disjoint from U except for the zero vector, and the dimensions of U and $U^\perp$ sum to the dimension of V.

The following simple lemmas relate M2MAs to UFAs and DFAs, and show that M2MAs accept exactly the regular languages.

Lemma 1. *[Beimel et al. [4]] Let $L \subseteq \Sigma^*$. If L is accepted by a UFA of size n, it is also accepted by an M2MA of dimension n.*

Lemma 2. *Let $L \subseteq \Sigma^*$. If L is accepted by an M2MA of dimension d with R reachable states, then L is also accepted by a DFA of R states. Clearly, $R \leq 2^d$.*

Beimel et al. [4] have shown that there is a polynomial time algorithm to learn an unknown M2MA using equivalence and membership queries.

2.4   Size lower bounds for DFAs, M2MAs and NFAs

Given a language $L \subseteq \Sigma^*$, we define an *observation table* for L as an $\ell \times m$ matrix T of 0's and 1's where each row i is associated with a finite word x_i and each column j is associated with a finite word y_j, and the entry $T_{i,j}$ is 1 if and only if $x_i y_j \in L$. This terminology is derived from its use in algorithms to learn DFAs. An observation table for L is thus a finite submatrix of its Hankel matrix.

Certain properties of observation tables for a language L yield lower bounds on acceptors recognizing L. Recall that the rank of a matrix is the number of linearly independent rows (or columns) it contains.

Lemma 3. *Let T be an observation table for the regular language L with rows associated with finite words x_i for $i = [1..\ell]$ and columns associated with finite words y_j for $j \in [1..m]$. Assume T has n distinct rows and rank d over the field $\{0, 1\}$. Then any DFA to accept L must have at least n states, and any M2MA to accept L must have dimension at least d.*

Proof. Let D be a DFA accepting L. If the rows for x_i and x_k are distinct, then there is a column j on which they differ, that is, $x_i y_j \in L$ iff $x_k y_j \notin L$. Thus, the states of D reached from the initial state on the words x_i and x_k must be different and D has at least n states.

Let M be an M2MA accepting L. Following the argument of Beimel et al. [4], the observation table is a submatrix of the Hankel matrix of the language L, and its rank (modulo 2) is a lower bound for the rank (modulo 2) of the Hankel matrix, which is a lower bound for the size of any M2MA accepting L.   □

For lower bounds for NFAs, we use the concept of covering the observation table by 1-monochromatic rectangles. If R and C are subsets of the indices of the rows and columns (respectively) of a matrix M, then the (R, C)-*rectangle* of M is the matrix obtained from M by deleting those rows whose indices are not in R and those columns whose indices are not in C. The (R, C)-rectangle of a matrix M is v-*monochromatic* iff all of its entries are equal to the value v.

Let M be a matrix of 0 and 1 values. A 1-*rectangle cover* of M is a set $\{(R_s, C_s) \mid s \in [1..t]\}$, of 1-monochromatic rectangles (R_s, C_s) of M such that for every i and j, if $M_{i,j} = 1$ then there exists some $s \in [1..t]$ such that $i \in R_s$ and $j \in C_s$. A *minimum* 1-*rectangle cover* of M is a 1-rectangle cover of M of minimum possible cardinality t.

Lemma 4. *Let T be an $\ell \times m$ observation table for the regular language L. Any NFA M recognizing L must have at least as many states as the cardinality of the minimum 1-rectangle cover of T.*

This is implied by Theorem 5.2.4.10 and Exercise 5.2.5.14 of Hromkovič [12]. For completeness we provide a simple direct proof.

Proof. Let the strings indexing the rows of T be x_i for $i \in [1..\ell]$ and the strings indexing the columns of T be y_j for $j \in [1..m]$. For each state q of M, let R_q be the set of all $i \in [1..\ell]$ such that x_i reaches q from an initial state of M, and let C_q be the set of all $j \in [1..m]$ such that y_j reaches a final state of M from q.

Clearly (R_q, C_q) must be a 1-monochromatic rectangle of T, because if $i \in R_q$ and $j \in C_q$ then $x_i y_j$ is accepted by M and the entry $T_{i,j}$ must be 1. Also, if $T_{i,j} = 1$, then $x_i y_j$ must be accepted by M, so there must exist a state q of M such that x_i reaches q from an initial state of M and y_j reaches a final state of M from q, that is, $i \in R_q$ and $j \in C_q$. Thus, the rectangles (R_q, C_q) for all states q of M form a 1-rectangle covering of T, and the number of states of M is greater than or equal to the cardinality of the minimum 1-rectangle covering of T.   □

Corollary 1. *If L is a regular language with an $n \times n$ observation table T that has exactly one 1 in every row and column, then any DFA, M2MA, or NFA to recognize L must have at least n states.*

As an example of the use of these results, let L be the regular language over $\{a, b, c\}$ consisting of those strings that do not contain any occurrences of the substrings ba or cb, with the observation table for L in Fig. 1. There are 4 different rows, so any DFA to accept L must have at least 4 states. The mod 2 rank of the table is 3 (the first three rows are a row basis) so any M2MA accepting L must have dimension

	ε	a	b
ε	1	1	1
b	1	0	1
c	1	1	0
ba	0	0	0

Fig. 1: Observation table with rank 3.

at least 3. The observation table with rows c and b, and columns a and b is the 2×2 identity matrix, so any NFA to accept L must have at least 2 states. In fact, there is a DFA of 4 states, an M2MA of dimension 3, and an NFA of 2 states accepting L, so for this example, the lower bounds are tight.

3   M2MAs as representations of regular languages

We consider the computational cost and size implications of some common operations and decision questions using M2MAs to represent regular languages.

3.1   M2MAs: procedures for operations and properties

Reverse Given an M2MA A accepting a regular language L, an M2MA A^r accepting the reverse language L^r may be obtained from A by exchanging the initial and final vectors, and transposing each of the transition matrices. Thus, the minimum dimension of an M2MA accepting L is equal to the minimum dimension of an M2MA accepting L^r. Reversing is similarly easy for UFAs and NFAs, but may incur an exponential increase in size for a DFA.

Sum If for $i = 1, 2$, M_i is a multiplicity automaton of dimension d_i computing the function $f_i : \Sigma^* \to \mathcal{K}$, then the sum $f_1 + f_2$ is computed by a multiplicity automaton M of dimension $d_1 + d_2$ constructed as the direct product of M_1 and M_2 as follows. State vectors of M are the concatenation of state vectors of M_1 and M_2, including the initial and final vectors. For each $\sigma \in \Sigma$, the transition matrix μ_σ is a $(d_1 + d_2) \times (d_1 + d_2)$ matrix obtained by putting $(\mu_1)_\sigma$ in the upper left, $(\mu_2)_\sigma$ in the lower right, and setting the remaining entries to 0. This ensures that the state updates of M_1 and M_2 are done in parallel for each symbol, and the output is the sum of the outputs for M_1 and M_2.

Boolean operations For M2MAs, complementation follows directly from the sum construction. If A is an M2MA of dimension d and C is the M2MA of dimension 1 that outputs 1 on every string, then the sum construction with M and C yields an M2MA of dimension $d + 1$ that accepts the regular language $\Sigma^* \setminus L(A)$. For DFAs, complementation is size-preserving, while for NFAs, complementation may incur an exponential increase in size.

Given M2MAs A_i of dimension d_i for $i = 1, 2$, the intersection language $L(A_1) \cap L(A_2)$ is accepted by an M2MA of dimension $d_1 \cdot d_2$ obtained from A_1

and A_2 using the Kronecker product of matrices.[3] Union can then be obtained from complementation and intersection.

Minimization, Equivalence, and Emptiness Sakarovitch [10,19] describes a cubic-time algorithm to minimize a weighted automaton with weights from a skew field, which has the following corollary.

Corollary 2 (of Theorem 5.20 in [10]). *Given an M2MA A of dimension d, an M2MA A' of the minimum possible dimension accepting $L(A)$ may be found in time $O(|\Sigma|d^3)$.*

An M2MA recognizes the empty language iff it has dimension 0 when minimized, and the equivalence of two M2MAs may be tested by determining if their sum is the empty language.

3.2   Conciseness comparisons for regular languages

We summarize known results comparing the conciseness of M2MAs with that of DFAs, UFAs and NFAs as representations of regular languages in Fig. 2. The entry for row A and column B is "$-$" if the representation A is an instance of the representation B, otherwise, starting with a machine of size n in the representation A, how large must an equivalent machine in the representation B be in the worst case? The entry $2^{\Theta(n)}$ means that there is a lower bound of 2^{cn} and an upper bound of 2^{dn} for positive constants c and d.

We briefly explain the entries in the table. A DFA is also a UFA and an NFA, and a UFA is also an NFA. A DFA or UFA of size n can be converted to an equivalent M2MA of dimension n (Lemma 1). The subset construction to determinize an NFA of size n yields a DFA (and therefore also a UFA or M2MA) of size at most 2^n. An M2MA of dimension n can be converted to a DFA (or UFA or NFA) of size

	DFA	UFA	NFA	M2MA
DFA	$-$	$-$	$-$	n
UFA	$2^{\Theta(n)}$	$-$	$-$	n
NFA	$2^{\Theta(n)}$	$2^{\Theta(n)}$	$-$	$2^{\Theta(n)}$
M2MA	$2^{\Theta(n)}$	$2^{\Theta(n)}$	$2^{\Theta(n)}$	$-$

Fig. 2: Worst-case size bounds for representations of regular languages.

at most 2^n (Lemma 2). The language $B_n = \Sigma^* \cdot 1 \cdot \Sigma^n$, for $\Sigma = \{0,1\}$, consisting of binary strings with a 1 located $n+1$ symbols before the end is accepted by a UFA of size $n+2$ (and therefore also an NFA of size $n+2$ and an M2MA of dimension $n+2$), but requires at least 2^{n+1} states for any DFA that accepts it.

For the problem of converting an NFA to an M2MA, Kaznatcheev and Panangaden [13] consider the language $L_n = \Sigma^* \left((0\Sigma^{n-1}1) + (1\Sigma^{n-1}0) \right) \Sigma^*$ for $\Sigma = \{0,1\}$, and show that L_n is recognized by an NFA of $2n+2$ states, but that any M2MA to recognize L_n must have dimension at least 2^n. By Lemma 1, this lower bound applies also to UFAs.

For the problem of converting an M2MA to an NFA, Kaznatcheev and Panangaden [13] give a family of languages $\{L_n\}$ such that L_n is recognized by an

---

[3] If A is an $m \times n$ matrix and B is a $p \times q$ matrix, then the Kronecker product $A \otimes B$ is the $pm \times qn$ block matrix, with blocks of size B, where the block-matrix at position (i, j) is $a_{ij}B$ [17, Def 1.2.1].

M2MA of dimension $n + 2$, and prove that any NFA to recognize L_n must have at least $2^{n/2} - 2$ states. Here we provide a simpler proof of a stronger lower bound. Let L_n be the language recognized by the M2MA given in Fig. 1 of the paper of Kaznatcheev and Panangaden. This M2MA accepts a word iff the number of indices i such that both $w[i]$ and $w[i + n]$ is 1, is odd.

Lemma 5. *Any NFA to recognize L_n must have at least 2^{n-1} states.*

Proof. The language L_n has an observation table T_n of dimension $2^n \times 2^n$, in which the rows and columns are indexed by strings $x, y \in \{0,1\}^n$. We view strings in $\{0,1\}^n$ as vectors of length n over the field $\{0,1\}$, so that the entry corresponding to the pair (x, y) is the inner product of the vectors x and y, that is $x^\top y$. Note that the inner product $x^\top y$ is 1 iff the number of indices i such that both $xy[i]$ and $xy[i + n]$ is 1, is odd. The lower bound of $2^n - 1$ then follows from Lemma 4 and the following Lemma. $\qquad\square$

Lemma 6. *The minimum 1-rectangle covering of the observation table T_n just defined has cardinality $2^n - 1$.*

Proof. For the upper bound it suffices to consider a 1-rectangle covering of T_n consisting of pairs (R, C) where R is the singleton index of a nonzero row and C consists of the indices of the occurrences of 1 in that row.

If $x \in \{0,1\}^n$ is the zero vector, then $x^\top y$ is 0 for all vectors y; otherwise, $x^\top y = 1$ for exactly half the vectors y, that is, for 2^{n-1} columns of T_n. Hence, T_n contains exactly $2^{n-1}(2^n - 1)$ entries of value 1. We now show that any 1-monochromatic rectangle (R, C) of T_n has at most 2^{n-1} entries of 1, which shows that a minimum 1-rectangle covering of T_n must have cardinality at least $2^n - 1$.

Let (R, C) be any 1-monochromatic rectangle of T_n. Let U be the vector subspace spanned by the vectors x corresponding to indices in R, and let B be a basis for U whose indices are drawn from R. Let $k = |B|$, so that $|U| = 2^k$. Every element of U is a sum of elements of B, but a sum of an even number of elements of B will be 0 in all the columns with indices in C, so R can contain the indices of at most half the elements of U, that is, $|R| \leq 2^{k-1}$.

Let $S = \{v \mid u^\top v = 1 \ \forall u \in B\}$, the set of vectors whose inner product with all elements of B is 1; clearly, $|C| \leq |S|$. We use inclusion/exclusion to find the cardinality of S, as follows.

$$\begin{aligned}
|S| &= 2^n - |\bigcup_{u \in B} \{v \mid u^\top v = 0\}| \\
&= 2^n - |\bigcup_{C \subseteq B} C^\perp| \\
&= 2^n - k2^{n-1} + \binom{k}{2} 2^{n-2} - \ldots (-1)^k 2^{n-k} \\
&= 2^n \cdot (1 - \frac{1}{2})^k \\
&= 2^{n-k}
\end{aligned}$$

Thus, $|C| \leq 2^{n-k}$. Then $|R \times C| \leq 2^{k-1} \cdot 2^{n-k} = 2^{n-1}$, concluding the proof. $\quad\square$

4   Representing regular omega-languages using regular languages

In the preliminaries we discussed NBAs, SUBAs and DBAs, and LTL formulas as representations of regular ω-languages. Here we explain that M2MAs and other automata over finite words can also be used to represent regular ω-languages.

A regular ω-language is uniquely determined by the set of ultimately periodic ω-words it contains. Let L be a regular ω-language and let $\$$ be a symbol not in the alphabet of L. To represent the set of ultimately periodic words in L, Calbrix, Nivat and Podelski [7] introduced the related language of finite words $L_\$ = \{u\$v \mid u(v)^\omega \in L\}$ and proved that it is regular.

Thus a regular ω-language L can be represented by an acceptor for the regular language $L_\$$, for example, a DFA, UFA, NFA or M2MA. The representation of $L_\$$ by an M2MA was used by Angluin, Antonopoulos, and Fisman [1] in showing that regular ω-languages are polynomially predictable with membership queries as a function of the size of the smallest SUBA accepting the language.

We note that if for $i = 1, 2$, A_i is an M2MA of dimension d_i accepting $(L_i)_\$$ for the regular ω-language L_i, then there is an M2MA of dimension $d_1 \cdot d_2$ accepting $(L_1 \cap L_2)_\$$, and an M2MA of dimension $d_1 + 3$ accepting $(\Sigma^\omega \setminus L_1)_\$$. The former follows by the intersection result for M2MAs, and the latter follows by the sum result applied to A_1 and the dimension 3 M2MA that accepts the set $\{u\$v \mid u \in \Sigma^*, v \in \Sigma^+\}$.

5   Conciseness comparisons for regular omega-languages

We present known and new results comparing the conciseness of M2MAs with that of several other representations of regular ω-languages, summarized in Fig. 3. The entry for row A and column B gives upper (above) and lower (below) bounds on the worst case increase in size for a representation of type A of size or dimension n to an equivalent representation of type B. The entry is "$-$" if a representation of type A is an instance of a representation of type B. The entries for the columns for DFA, UFA, M2MA, and NFA are for the language $L_\$$. An arrow indicates that the (lower or upper) bound is derived from a related (lower or upper) bound in the table. For example, the upper bound for the row DBA and columns UFA, M2MA and NFA are derived from the upper bound for the row DBA and column DFA. We now discuss the entries.

5.1   Size increases for LTL formulas

Upper bounds
There is a "classic" algorithm, described by Baier and Katoen [3, Chapter 5], to translate an LTL formula of size n into a GNBA of size 2^n with at most n sets of final states, which then yields an NBA of size at most $n2^n$. This shows that every LTL formula represents a regular ω-language, and gives an upper bound for translating an LTL formula to an NBA. Another algorithm to translate LTL formulas into NBAs is given by Gerth, Peled, Vardi and Wolper [11].

	DFA	UFA	M2MA	NFA	(G)SUBA	NBA
LTL	$2^{2^{O(n)}}$ via UFA $\rightarrow$	$2^{O(n)}$ Cor.3 $\rightarrow$	$\leftarrow$ $2^{\Omega(n)}$ Thm. 2	$\leftarrow$ $2^{\Omega(n)}$ Thm. 2	$(2^n, n)$ Prop. 1 $\rightarrow$	$n2^n$ [3] $2^{\Omega(n)}$ [3]
DBA	$n+n3^{n^2}$ [14] $2^{\Omega(n \log n)}$ [2]	$\leftarrow$ $\rightarrow$	$\leftarrow$ $2^{\Omega(n)}$ Thm. 3	$\leftarrow$ $2^{\Omega(n)}$ Thm. 3	$\downarrow$ $2^{\Omega(n)}$ [6]	$-$ $-$
NBA	$2^n + 2^n3^{n^2}$ [14] $\uparrow$	$\leftarrow$ $\uparrow$	$\leftarrow$ $\uparrow$	$n+n3^{n^2}$ [14] $\uparrow$	$(12n)^n$ [8] $\uparrow$	$-$ $-$
SUBA	$\uparrow$ $2^{\Omega(n)}$ [1]	$2n^2 + n$ [6] $\rightarrow$	$\leftarrow$ $2n^2 - n + 2$ Thm. 4	$\leftarrow$ $2n^2 - n + 2$ Thm. 4	$-$ $-$	$-$ $-$

Fig. 3: Worst-case size bounds for representations of regular ω-languages.

Concerning the classic translation algorithm, Bousquet and Löding [6] give a brief argument and state that "Hence the automaton that is constructed in this standard way is strongly unambiguous." Wilke [21] states that "Every temporal formula with n subformulas can be translated into an equivalent backwards deterministic generalized Büchi automaton with at most 2^n states and as many Büchi sets as there are subformulas with leading temporal operator F (eventually) or U (until)." To clarify these earlier statements, we reformulate them in our terminology. This gives an upper bound for transforming an LTL formula to a GSUBA.

Proposition 1. *Let ϕ be an LTL formula of size n with temporal operators next and until, with m until subformulas. Applying the classic translation algorithm to ϕ yields a GSUBA of size 2^n with m sets of final states.*

Proof. Baier and Katoen [3] show that the algorithm yields a GNBA M of the given size in which each state corresponds to an assignment of true or false to every subformula of ϕ. Moreover, if the ω-word w is accepted from a state q, then q assigns true to each subformula ψ of ϕ iff ψ is true for w. Hence there is at most one state of M from which the ω-word w is accepted, and thus M is also GSUBA. $\qquad\square$

To get an upper bound for translation of LTL formulas to UFAs, M2MAs, and NFAs, we would like to use the property of being strongly unambiguous. However, if the resulting GSUBA has more than one set of final states, transforming it in the usual way into an NBA does not in general yield a SUBA. Instead, we generalize to GSUBAs the method of Bousquet and Löding [6] for transforming a SUBA accepting L into a UFA accepting $L_\$$.

Theorem 1. *There is an algorithm to transform a GSUBA of size n with m sets of final states accepting L into a UFA of size $2^m n^2 + n$ accepting $L_\$$. It runs in time polynomial in n and 2^m.*

Proof. Let L be accepted by the GSUBA $M = (\Sigma, Q, I, \Delta, \mathcal{F})$ with $n = |Q|$ and $m = |\mathcal{F}|$. We index the elements of $\mathcal{F}$ as F_i for $i \in [1..m]$. Bousquet and Löding [6, Lemma 1] show that $u(v)^\omega$ is accepted by a SUBA iff there exists a state q reachable from an initial state on reading u, such that on the word v there is a computation path that loops from q back to q while passing through an accepting state. For the GSUBA M, the condition is that the computation path that loops from q back to q must pass through at least one state from each F_i for $i \in [1..m]$.

We define an NFA $M' = (\Sigma', Q', I', \Delta', F')$ as follows. The alphabet is $\Sigma' = \Sigma \cup \{\$\}$. The state set is $Q' = Q \cup Q_1$, where $Q_1 = \{(q_1, q_2, S) \mid q_1, q_2 \in Q, S \subseteq [1..m]\}$. The initial states are $I' = I$. The transition relation is $\Delta' = \Delta \cup \Delta_1 \cup \Delta_2$, where Δ_1 is the set of all triples $((q_1, q_2, S), \sigma, (q_1', q_2', S'))$ such that $q_1' = q_1$, $(q_2, \sigma, q_2') \in \Delta$, and $S' = S \cup T$, where $T = \{i \in [1..m] \mid q_2' \in F_i\}$. And Δ_2 contains all triples $(q, \$, (q, q, \emptyset))$ such that $q \in Q$. The set of final states F' is the set of triples (q_1, q_2, S) such that $S = [1..m]$ and $q_1 = q_2$.

Then M' has $2^m n^2 + n$ states, and can be constructed in time polynomial in n and 2^m given the GSUBA M. On an input $u\$v$, the NFA M' behaves like M on the word u, reaching some state q. Then on the symbol $\$$, M' transitions to the state $(q, q, \emptyset)$, recording the state q reached after reading u. As M' continues reading v, the first component remembers q while the second component transitions as in M. The third component, S, records the set of indices of those final sets F_i that have been visited in the processing of v. The input $u\$v$ is accepted by M' iff there is a state q of M reachable from a state of I on input u such that there exists a computation path in M on input v from q to q that visits at least one state in F_i for every $i \in [1..m]$. Thus M' accepts $L_\$$. (Note that the set S generalizes the single bit used in Bousquet and Löding's construction.)

To see that M' is a UFA, we note that if there are two different accepting computations in M' for $u\$v$, then these may be used to construct two different accepting computations in M for $u(v)^\omega$, contradicting the fact that M is a GSUBA. $\square$

The entry in Fig. 3 for row LTL and column UFA is then justified by the following.

Corollary 3. *Let ϕ be an LTL formula of size n with temporal operators next and until, with m until subformulas. Then there is a UFA of size $2^{2n+m} + 2^n$ to accept $L(\phi)_\$$.*

For transforming LTL to DFA, we have only the doubly-exponential bound for transforming an LTL formula to a UFA and the UFA to DFA.

Lower bounds
We first generalize Lemma 3 to DBAs and Lemma 4 to NBAs. An *observation*

table for an ω-language L is a matrix $T \in \{0,1\}^{\ell \times m}$ with rows indexed by finite words x_i for $i \in [1..\ell]$ and columns indexed by ω-words y_j for $j \in [1..m]$ such that $T_{i,j} = 1$ iff $x_i y_j \in L$. Then we have the following, proved analogously to Lemma 3 and Lemma 4.

Lemma 7. *Let T be an observation table for the ω-language L. If T has n distinct rows, then any DBA accepting L has at least n states.*

Lemma 8. *Let T be an observation table for the ω-language L. If the minimum 1-cover of T has cardinality n, then any NBA to recognize L has at least n states.*

Baier and Katoen [3, Theorem 5.4.2] give a lower bound for a family of LTL formulas ϕ_n of size poly(n) for which equivalent NBAs must have at least 2^n states. Below we give a simplified and slightly strengthened version of their lower bound, which also applies to M2MAs or NFAs for $L_\$$.

Theorem 2. *For every positive integer n there exists an LTL formula ψ_n of size at most $2n+6$ such that any NBA accepting $L(\psi_n)$ must have size at least 2^n. Any NFA or M2MA accepting $L(\psi_n)_\$$ must have size or dimension at least 2^n.*

Proof. Let p be a propositional variable. For any positive integer n we define the LTL formula $\quad \psi_n = \Box(p \to \bigcirc^n(p)) \land (\neg p \to \bigcirc^n(\neg p))$. We use $\bigcirc^n$ to represent the composition of $\bigcirc$ with itself n times, so $\bigcirc^3(p)$ abbreviates $\bigcirc(\bigcirc(\bigcirc(p)))$. The formula ψ_n has size $2n+6$. Let the symbols 0 and 1 represent the assignment of false and true to p. Then $L(\psi_n)$ is the language of ω-words w over $\{0,1\}$ such that for some $x \in \Sigma^n$, $w = x^\omega$.

For $L(\psi_n)_\$$, let $x_1, x_2, \ldots, x_{2^n}$ be any total ordering of all the elements of $\{0,1\}^n$, and consider the observation table T with rows corresponding to x_i and columns corresponding to $\$x_i$ for $i \in [1..2^n]$. Clearly, there is exactly one 1 in row x_i, in the column $\$x_i$, so this observation table is the $2^n \times 2^n$ identity matrix, which has rank 2^n, and any NFA or M2MA accepting $L(\psi_n)_\$$ must have size at least 2^n by Corollary 1.

For the lower bound on NBAs, we observe that if we instead index the columns of T with $(x_i)^\omega$, it becomes an observation table for the ω-language $L(\psi_n)$, and remains the $2^n \times 2^n$ identity matrix, which implies that any NBA accepting $L(\phi_n)$ must have at least 2^n states, by Lemma 8. $\qquad\square$

5.2 Size increases for DBAs, NBAs, SUBAs

Upper bounds

For an NBA of n states accepting L, Calbrix, Nivat and Podelski [7] show that there is a DFA of $2^n + 2^{2n^2+n}$ states to accept $L_\$$. Kuperberg, Pinault and Pous [14] give a more concise construction that yields for $L_\$$ an NFA of size $n + n3^{n^2}$ and a DFA of size $2^n + 2^n 3^{n^2}$. For the conversion of an NBA of n states to a SUBA, Carton and Michel provide the upper bound of $(12n)^n$ [8].

Starting with a DBA instead of an NBA, the NFA construction of Kuperberg, Pinault and Pous is fully deterministic, so the upper bound of $n + n3^{n^2}$ holds for transforming a DBA into a DFA. Bousquet and Löding [6] show that a SUBA of n states accepting the ω-language L may be transformed into a UFA of $2n^2 + n$ states accepting $L_\$$.

Lower bounds

For transforming a DBA for L into a DFA for $L_\$$, Angluin and Fisman [2] prove that for every n there is a DBA of $n + 2$ states accepting a language L such that no DFA of fewer than $n!$ states accepts $L_\$$. For transforming a DBA into a UFA, M2MA or NFA, we prove the following result.

Theorem 3. *For every even positive integer n there is an ω-language L_n that is accepted by a DBA of $n + 5$ states such that any UFA, NFA or M2MA to accept $(L_n)_\$$ must have size or dimension at least $\binom{n}{n/2}$, which is $\sim 2^n/\sqrt{\pi n/2}$.*

Proof (Sketch). The proof uses a modification of the DBAs in the construction by Angluin and Fisman [2]. Here we sketch the main idea and give an example. Let $n = 2k$ for some nonnegative integer k, let $\Sigma_{2k} = \{\sigma_1, \ldots, \sigma_{2k}\}$ and let Σ be $\Sigma_{2k} \cup \{0, L, E, F\}$. Consider the regular ω-language defined by the ω-regular expression $\left(\cup_{\sigma \in \Sigma \setminus \{0\}} (\sigma \cdot (\Sigma \setminus \{\sigma\})^* \cdot \sigma) \right)^\omega$, which is accepted by a DBA with $2k + 5$ states. Given two subsets C and D of Σ_{2k}, each of size k, we define words u_C and v_D such that $(u_C \cdot v_D)^\omega$ is in the language if and only if $C = D$. The main idea behind the construction is that v_D forces each symbol σ_D in $\Sigma_{2k} \setminus D$ to be followed by the character 0. Thus, if the string preceding (and including) an occurrence of such a symbol σ_D is described by the (unambiguous) regular expression $\left(\bigcup_{\sigma \in \Sigma \setminus \{0\}} \sigma \cdot (\Sigma \setminus \{\sigma\})^* \cdot \sigma \right)^*$, then the symbol 0 that follows cannot be properly consumed, resulting in the ω-word being not in the language. We construct the words u_C and v_D in such a way that this can happen if and only if such a symbol $\sigma_D \in \Sigma_{2k} \setminus D$ is also in C. Since C and D are subsets of Σ_{2k}, each of size k, this happens exactly when $C \neq D$. There is therefore an observation table with rows indexed by $\$u_C$ for all subsets C of size k and whose columns are indexed by v_D for all subsets D of size k, and where each entry, corresponding to row and column subsets C and D respectively, is 1 if and only if $C = D$. By Corollary 1, the result follows.   □

Example. Let $\Sigma_{2k} = \{1, 2, 3, 4\}$, let Σ be $\Sigma_{2k} \cup \{0, L, E, F\}$, let $C = \{2, 3\}$, and let $D = \{2, 4\}$. Then u_C, v_C and v_D are defined on the right. Then $(u_C \cdot v_C)^\omega$ is in the language, whereas $(u_C \cdot v_D)^\omega$ is not (since $C \neq D$).

$$u_C = F \cdot 2 \cdot F \cdot 2 \cdot 3 \cdot 2 \cdot 3 \cdot L \cdot 3$$
$$v_C = L \cdot E \cdot 1 \cdot 0 \cdot 4 \cdot 0 \cdot E$$
$$v_D = L \cdot E \cdot 1 \cdot 0 \cdot 3 \cdot 0 \cdot E$$

For the lower bound on transforming a DBA into a SUBA, Bousquet and Löding [6] show that for every positive integer n there exists an ω-language that is accepted by a DBA with $n + 1$ states, and cannot be accepted by a SUBA with fewer than 2^{n-1} states.

For transforming a SUBA into a DFA, Angluin, Antonopoulos and Fisman [1, Theorem 5] give a family of ω-languages such that L_n is accepted by a SUBA of size $4n + 5$, but any DFA to accept $(L_n)_\$$ or its reverse must have size at least 2^n. For transforming a SUBA into a UFA, M2MA or NFA, we prove the following asymptotically tight lower bound.

Theorem 4. *For every positive integer m greater than 3, there is an ω-language L that is accepted by a SUBA with m states, but no M2MA of dimension less than $2m^2 - m + 2$ or NFA or UFA of size less than $2m^2 - m + 2$ accepts $(L)_\$$.*

Proof (Sketch). For every $n \in \mathbb{N}$ we define L_n to be the regular ω-language over $\Sigma = \{a, b, c\}$ given by the expression $((cc \cdot b^n)^* \cdot aa \cdot b^n)^\omega$. This language is accepted by a SUBA S_n, with $m = n + 3$ states. We construct a specific observation table M for the language $(L_n)_\$$. We then show that any 1-rectangle cover of M is of size at least $2m^2 - m + 2$, which implies by Lemma 4 that the number of states of any NFA (or UFA) for the language $(L_n)_\$$ is at least $2m^2 - m + 2$. We further show that the rank of M is $2m^2 - m + 2$, and by Lemma 3, obtain that the dimension of any M2MA for this language is also at least $2m^2 - m + 2$.   □

6   Empirical results

We report typical size increases in going from a random SUBA, DBA or NBA acceptor for a regular ω-language L to a minimized M2MA (and DFA, in the case of a SUBA) for $L_\$$. We also report computed sizes of minimized M2MAs and DFAs for $L(\phi_n)_\$$ for members of particular families $\{\phi_n\}$ of LTL formulas. Code is available in the GitHub repository:
https://github.com/nevingeorge/Learning_Automata.
    For the generation of random SUBAs, DBAs or NBAs, our procedure is as follows. Given parameters n, f, and t we generate a transition relation on n states (random reverse-deterministic for a SUBA, random deterministic for a DBA, and all possible transitions for an NBA), select f of the n states at random to be final, and randomly remove t of the transitions. The resulting transition relation is trimmed to remove non-live states and their transitions. The trimmed acceptor may have fewer than n states.
    If the goal is a SUBA, using the criterion of Wilke [21], we check that there do not exist two different states q_1 and q_2 and a nonempty finite word v such that for $i = 1, 2$, there is a loop on v from q_i to q_i that passes through a final state. If the acceptor fails this test, it is rejected, and the procedure is repeated until a SUBA is successfully generated.

6.1   SUBAs to minimized M2MAs and DFAs

For random SUBAs to minimized M2MAs, we first generate a random SUBA with $\Sigma = \{a, b, c\}$, $n \in \{5, 10, 15\}$, $t \in \{[1, 5], [2, 10], [18, 22]\}$ (resp.), and $f = 2$ or $f = 3$ with equal probability. We then convert it into a UFA using the algorithm of Bousquet and Löding [6], and minimize the equivalent M2MA.
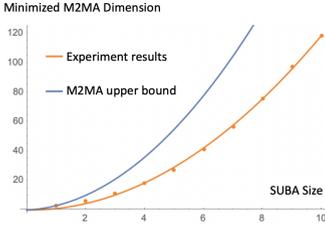
Fig. 4: Random SUBAs to minimized M2MAs
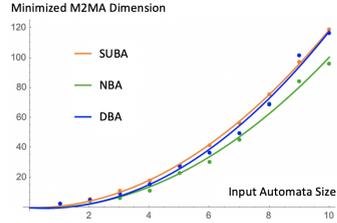
Fig. 5: Random SUBAs, NBAs, and DBAs to minimized M2MAs
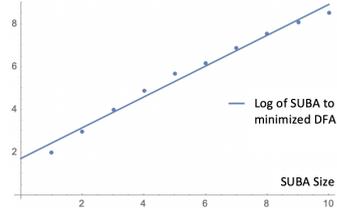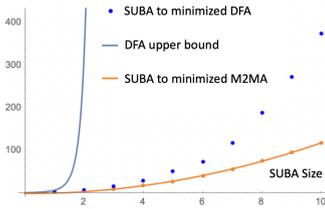
Fig. 6: Random SUBAs to minimized DFAs

We performed the above process on approximately $220,000$ randomly generated SUBAs.

Fig. 4 is a plot of the average minimized M2MA dimension for each trimmed SUBA size from 1 to 10. Upon performing quadratic regression, we obtain the orange curve $1.212n^2 - .2248n$, and the blue curve is the theoretical upper bound of $2n^2 + n$ given in Fig. 3. The quadratic fit has a R^2 of 0.9996 while a linear fit has a R^2 of 0.9370, suggesting that the growth is indeed quadratic. This curve satisfies the theoretical upper bound of $2n^2 + n$, and suggests that the lower bound of $\Omega(n^2)$ holds on average.

For random SUBAs to minimized DFAs, we also calculated the number of reachable states of each minimized M2MA. This is the number of states in the equivalent minimized DFA, by a property of the minimization algorithm of Corollary 2. From Fig. 3, the lower bound in going from a SUBA to a DFA is $2^{\Omega(n)}$, and the upper bound is $2^n + 2^n 3^{n^2}$.

In the left graph in Fig. 6, the blue data points representing the results of the SUBA to DFA experiment grow much more sharply than the results of the SUBA to M2MA experiment, so it is clear that a SUBA can be represented more concisely as an M2MA than as a DFA on average. Upon taking the log (base 2), we obtain a roughly linear fit as seen in the right graph with equation $.7196n + 1.738$ and a R^2 of .9841, suggesting that on average the growth is exponential. The standard deviation and range of converted DFA sizes was large for this conversion, making it difficult to make firm claims about the growth. However, the data suggests that the exponential lower bound likely holds on average, and that in general the upper bound of $2^n + 2^n 3^{n^2}$ is a severe overestimate.

6.2   NBAs and DBAs to minimized M2MAs

For NBAs and DBAs, a minimized M2MA is computed using the M2MA learning algorithm of Beimel et al. [4], which makes membership and equivalence queries to the NBA or DBA. Instead of exact equivalence queries, we use approximate equivalence queries, implemented by testing membership agreement on a sample of randomly generated ultimately periodic words. Thus, the dimension of the learned M2MA may be an underestimate of the true minimum dimension of an M2MA for $L_\$$.

For the NBA/DBA to M2MA experiments, we generated approximately 1000 random NBAs/DBAs with $\Sigma = \{a, b, c\}$, $n \in \{5, \ldots, 10\}$, $t \in [0, n]$ for DBAs and t in ranges within $[90, 680]$ for NBAs, and $f = 2$ or $f = 3$ with equal probability. For the approximate equivalence queries, we tested 1000 random ultimately periodic words of length at most 25. The results of the experiments can be seen in Fig. 5. The fitted NBA and DBA curves are quadratic with equations $1.096n^2 - .8947n$ and $1.318n^2 - 1.392n$, respectively. The quadratic fits for the NBA and DBA results have a R^2 of .9954 and .9961, respectively, while linear fits have a R^2 of .9227 and .9118, respectively. These experiments have limitations: the use of approximate equivalence queries, the small sample size (because of the time requirements of the learning algorithm), and the large standard deviation and range of converted M2MA sizes. However, the results from all three conversions are very similar, suggesting that in these conditions, SUBAs, NBAs, and DBAs don't vary significantly on average with respect to their equivalent M2MA representations.

6.3   LTL formulas to minimized M2MAs

Random LTL formulas seem not to provide much insight, so we consider specific families of LTL formulas: bounded request/grant formulas and two families based on the hierarchy of Manna and Pnueli [15], namely obligation and reactivity formulas. Empirically, for each of the first few members of each family we calculate the minimum dimension of an M2MA and the minimum size of a DFA accepting the corresponding $L_\$$ language, and use the online tool provided by the Spot website (https://spot.lrde.epita.fr/) to find an ω-language acceptor for the corresponding L. (Omitted Spot entries exceeded the limit on calculation time.)

The canonical request/grant formula is of the form $\Box(p \to \Diamond(q))$, which asserts that whenever a request (p) is made, it is eventually granted (q). In the bounded version, a number of steps n is specified, and the assertion is that the request is granted within n steps. Thus, for each natural number n, we have a formula $R_n = \Box(p \to (q \vee \bigcirc(q) \vee \bigcirc^2(q) \vee \ldots \vee \bigcirc^n(q)))$. The table in Fig. 7a gives the resulting sizes and dimensions for n from 0 to 5. It is reasonable to conjecture $n + 1$ for the size of a DBA, $n^2 + 3n + 3$ for the minimum dimension of an M2MA, and $2n^2 + 3n + 4$ for the minimum size of a DFA representing R_n.

The family of obligation formulas we consider is: $F_n = \wedge_{i=1}^n (\Box p_i \vee \Diamond q_i)$. Using conjunction and minimization, we calculate the minimum dimension M2MA (and minimum size DFA) for $L_\$$ for these formulas for n up to 5. The table in Fig. 7b

n	DBA	M2MA	DFA
0	1	3	4
1	2	7	9
2	3	13	18
3	4	21	31
4	5	31	48
5	6	43	69

(a) R_n sizes.

n	DBA	M2MA	DFA
1	3	7	9
2	9	19	23
3	27	55	63
4	81	163	179
5	–	487	519

(b) F_n sizes.

n	GNBA	M2MA	DFA
1	(4,1)	5	6
2	(10, 2)	11	12
3	(28, 3)	29	30
4	–	83	84
5	–	245	246

(c) G_n sizes.

Fig. 7: Size or dimension of acceptors for families of LTL formulas.

shows the results. It is reasonable to conjecture 3^n for the size of a DBA, $2 \cdot 3^n + 1$ for the minimum dimension of an M2MA, and $2 \cdot 3^n + 2^n + 1$ for the minimum size of a DFA to represent F_n.

The family of reactivity formulas we consider is: $G_n = \wedge_{i=1}^n (\Box \Diamond p_i \vee \Diamond \Box q_i)$. We proceed as for the obligation formulas, with the results shown in the table in Fig. 7c. Note that these formulas cannot be represented by DBAs, but are instead represented by GNBAs, which may have multiple sets of final states. For example, the entry $(10, 2)$ indicates a GNBA with 10 states and 2 sets of final states. A reasonable conjecture in this case is $(3^n + 1, n)$ for the size of a GNBA, $3^n + 2$ for the minimum dimension of an M2MA, and $3^n + 3$ for the minimum size of a DFA representing G_n.

In these cases, the minimum dimension of an M2MA (and size of a DFA) appears to grow at most as a polynomial in the size of an ω-language acceptor, quadratically for the bounded request/grant family, and linearly for the obligation and reactivity families.

7   Summary and conclusions

We provide a survey of size relations of M2MAs as a representation of regular languages and regular ω-languages, as well as empirical results for several of these relations. New theoretical results include an improvement of the lower bound for transforming an M2MA to an NFA, an upper bound of $2^{O(n)}$ for the translation of an LTL formula of size n to a UFA, NFA, or M2MA, a lower bound of $2^{\Omega(n)}$ for the translation of a DBA of n states to an M2MA or NFA, and an asymptotically optimal lower bound of $2n^2 - n + 2$ for the translation of a SUBA of n states to an M2MA or NFA.

M2MAs have many advantages as a representation for regular ω-languages: determinism, succinct complementation, and polynomial time algorithms for minimization, equivalence testing, and learning with membership and equivalence queries. M2MAs are as succinct as DFAs, sometimes exponentially more so, and deserve further study.

References

1. Angluin, D., Antonopoulos, T., Fisman, D.: Strongly unambiguous Büchi automata are polynomially predictable with membership queries. In: 28th EACSL Annual Conference on Computer Science Logic, CSL. pp. 8:1–8:17 (2020)
2. Angluin, D., Fisman, D.: Learning regular omega languages. Theor. Comput. Sci. **650**, 57–72 (2016)
3. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press (2008)
4. Beimel, A., Bergadano, F., Bshouty, N.H., Kushilevitz, E., Varricchio, S.: Learning functions represented as multiplicity automata. J. ACM **47**(3), 506–530 (May 2000)
5. Bergadano, F., Varricchio, S.: Learning behaviors of automata from multiplicity and equivalence queries. SIAM J. Comput. **25**(6), 1268–1280 (1996)
6. Bousquet, N., Löding, C.: Equivalence and inclusion problem for strongly unambiguous Büchi automata. In: Language and Automata Theory and Applications, 4th International Conference, LATA. Proceedings. pp. 118–129 (2010). https://doi.org/10.1007/978-3-642-13089-2_10
7. Calbrix, H., Nivat, M., Podelski, A.: Ultimately periodic words of rational w-languages. In: Proceedings of the 9th International Conference on Mathematical Foundations of Programming Semantics. pp. 554–566. Springer-Verlag (1994)
8. Carton, O., Michel, M.: Unambiguous Büchi automata. Theor. Comput. Sci. **297**(1-3), 37–81 (2003). https://doi.org/10.1016/S0304-3975(02)00618-7
9. Diekert, V., Gastin, P.: First-order definable languages. In: Logic and Automata: History and Perspectives [in Honor of Wolfgang Thomas]. pp. 261–306 (2008)
10. Droste, M., Kuich, W., Vogler, H. (eds.): Handbook of Weighted Automata, chap. 4: Rational and Recognizable Series, by Jaques Sakarovitch, pp. 105–174. Springer-Verlag Berlin Heidelberg (2009)
11. Gerth, R., Peled, D., Vardi, M., Wolper, P.: Simple on-the-fly automatic verification of linear temporal logic. In: Protocol Specification, Testing and Verification XV. PSTV 1995. Springer (1996). https://doi.org/10.1007/978-0-387-34892-6_1
12. Hromkovič, J.: Communication Complexity and Parallel Computing. Springer-Verlag Berlin Heidelberg (1997), (There is also 2013 edition.)
13. Kaznatcheev, A., Panangaden, P.: Weighted automata are compact and actively learnable. Information Processing Letters **171** (2021), (The authors were apparently unaware of prior results on learning multiplicity automata by Beimel et al. and others.)
14. Kuperberg, D., Pinault, L., Pous, D.: Coinductive algorithms for Büchi automata. In: Developments in Language Theory - 23rd International Conference, DLT Proceedings. pp. 206–220 (2019)
15. Manna, Z., Pnueli, A.: A hierarchy of temporal properties (invited paper, 1989). In: Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing. p. 377–410. PODC '90, Association for Computing Machinery (1990). https://doi.org/10.1145/93385.93442
16. Michel, M.: Complementation is much more difficult with automata on infinite words. In: Manuscript, CNET (1988)
17. Moser, B.K.: Linear algebra and related introductory topics. In: Linear Models, A Mean Model Approach, A volume in Probability and Mathematical Statistics. pp. 1–22 (1996)
18. Pnueli, A.: The temporal logic of programs. In: FOCS. pp. 46–57 (1977)
19. Sakarovitch, J.: Elements of Automata Theory. Cambridge University Press, USA (2009)

20. Vardi, M.Y.: An automata-theoretic approach to linear temporal logic. In: Logics for Concurrency - Structure versus Automata (8th Banff Higher Order Workshop, Banff, Canada, August 27 - September 3, 1995, Proceedings). pp. 238–266 (1995). https://doi.org/10.1007/3-540-60915-6_6

21. Wilke, T.: ω-automata. CoRR **abs/1609.03062** (2016), http://arxiv.org/abs/1609.03062