

TWO VARIABLE VS. LINEAR TEMPORAL LOGIC IN MODEL CHECKING AND GAMES*

MICHAEL BENEDIKT, RASTISLAV LENHARDT, AND JAMES WORRELL

Department of Computer Science, University of Oxford, United Kingdom
e-mail address: {michael.benedikt, rastislav.lenhardt, jbw}@cs.ox.ac.uk

ABSTRACT. Model checking linear-time properties expressed in first-order logic has non-elementary complexity, and thus various restricted logical languages are employed. In this paper we consider two such restricted specification logics, linear temporal logic (LTL) and two-variable first-order logic (FO^2). LTL is more expressive but FO^2 can be more succinct, and hence it is not clear which should be easier to verify. We take a comprehensive look at the issue, giving a comparison of verification problems for FO^2 , LTL, and various sublogics thereof across a wide range of models. In particular, we look at unary temporal logic (UTL), a subset of LTL that is expressively equivalent to FO^2 ; we also consider the stutter-free fragment of FO^2 , obtained by omitting the successor relation, and the expressively equivalent fragment of UTL, obtained by omitting the next and previous connectives.

We give three logic-to-automata translations which can be used to give upper bounds for FO^2 and UTL and various sublogics. We apply these to get new bounds for both non-deterministic systems (hierarchical and recursive state machines, games) and for probabilistic systems (Markov chains, recursive Markov chains, and Markov decision processes). We couple these with matching lower-bound arguments.

Next, we look at combining FO^2 verification techniques with those for LTL. We present here a language that subsumes both FO^2 and LTL, and inherits the model checking properties of both languages. Our results give both a unified approach to understanding the behaviour of FO^2 and LTL, along with a nearly comprehensive picture of the complexity of verification for these logics and their sublogics.

1. INTRODUCTION

The complexity of verification problems clearly depends on the specification language for describing properties. Arguably the most important such language is *Linear Temporal Logic* (LTL). LTL has a simple syntax, one can verify LTL properties over Kripke structures in polynomial space, and one can check satisfiability also in polynomial space. Moreover, Kamp [Kam68] has shown that LTL has the same expressiveness as first-order logic over words. For example, the first-order property “after we are born, we live until we die”:

$$\forall x (\text{born}(x) \rightarrow \exists y \geq x \text{ die}(y) \wedge \forall z (x \leq z < y \rightarrow \text{live}(z)))$$

2012 ACM CCS: [Theory of computation]: Logic—Verification by model checking; Formal languages and automata theory; Models of computation—Abstract machines.

Key words and phrases: Finite Model Theory, Verification, Automata.

* This includes material presented in Concur 2011 and QEST 2012 extended abstracts.

is expressed in LTL by the formula $\Box(\text{born} \rightarrow \text{live } \mathcal{U} \text{ die})$.

In contrast with LTL, model checking first-order queries has non-elementary complexity [Sto74]—thus LTL could be thought of as a tractable syntactic fragment of FO. Another approach to obtaining tractability within first-order logic is by maintaining first-order syntax, but restricting to two-variable formulas. The resulting specification language FO^2 has also been shown to have dramatically lower complexity than full first-order logic. In particular, Etessami, Vardi and Wilke [EVW02] showed that satisfiability for FO^2 is NEXPTIME-complete and that FO^2 is strictly less expressive than FO (and thus less expressive than LTL also). Indeed, [EVW02] shows that FO^2 has the same expressive power as *Unary Temporal Logic* (UTL): the fragment of LTL with only the unary operators “previous”, “next”, “sometime in the past”, “sometime in the future”. Consider the example above. We have shown that it can be expressed in LTL, but it is easy to show that it cannot be expressed in UTL, and therefore cannot be expressed in FO^2 .

Although FO^2 is less expressive than LTL, there are some properties that are significantly easier to express in FO^2 than in LTL. Consider the property that two n -bit identifiers agree:

$$\exists x \exists y (x < y \wedge \bigwedge_{1 \leq i \leq n} b_i(x) \leftrightarrow b_i(y)).$$

It is easy to show that there is an exponential blow-up in transforming the above FO^2 formula into an equivalent LTL formula. We thus have three languages UTL, LTL and FO^2 , with UTL and FO^2 equally expressive, LTL more expressive, and with FO^2 incomparable in succinctness with LTL.

Are verification tasks easier to perform in LTL, or in FO^2 ? This is the main question we address in this paper. There are well-known examples of problems that are easier in LTL than in FO^2 : in particular satisfiability, which is PSPACE-complete for LTL and NEXPTIME-complete for FO^2 [EVW02]. We will show that there are also tasks where FO^2 is more tractable than LTL.

Our main contribution is a uniform approach to the verification of FO^2 via automata. We show that translations to the appropriate automata can give optimal bounds for verification of FO^2 on both non-deterministic and probabilistic structures. We also show that such translations allow us to understand the verification of the fragment of FO^2 formed by removing the successor relation from the signature, denoted $\text{FO}^2[<]$. It turns out, somewhat surprisingly, that for this fragment we can get the same complexity upper bounds for verification as for the simplest temporal logic— $\text{TL}[\diamond, \heartsuit]$. For our translations from $\text{FO}^2[<]$ to automata, we make use of a key result from Weis [Wei11], showing that models of $\text{FO}^2[<]$ formulas realise only a polynomial number of types. We extend this “few types” result from finite to infinite words and use it to characterise the structure of automata for $\text{FO}^2[<]$.

The outcome of our translations is a comprehensive analysis of the complexity of FO^2 and UTL verification problems, together with those for the respective stutter-free fragments $\text{FO}^2[<]$ and $\text{TL}[\diamond, \heartsuit]$. We begin with model checking problems for Kripke structures and for recursive state machines (RSMs), which we compare to known results for LTL on these models. We then turn to two-player games, considering the complexity of the problem of determining which player has a strategy to ensure that a given formula is satisfied. We then move from non-deterministic systems to probabilistic systems. We start with Markov chains and recursive Markov chains, the analogs of Kripke structures and RSMs in the

probabilistic case. Finally we consider one-player stochastic games, looking at the question of whether the player can devise a strategy that is winning with a given probability.

Towards the end of the paper, we consider extensions of FO^2 , and in particular how FO^2 verification techniques can be combined with those for Linear Temporal Logic (LTL). We present here a language that we denote $\text{FO}^2[\text{LTL}]$, subsuming both FO^2 and LTL. We show that the complexity of verification problems for $\text{FO}^2[\text{LTL}]$ can be attacked by our automata-theoretic methods, and indeed reduces to verification of FO^2 and LTL individually. As a result we show that the worst-case complexity of probabilistic verification, as well as non-deterministic verification, for $\text{FO}^2[\text{LTL}]$ is (roughly speaking) the maximum of the complexity for FO^2 and LTL.

This paper expands on results presented in two conference papers, [BLW11, BLW12].

Organization: Section 2 contains preliminaries, while Section 3 gives fundamental results on the model theory of FO^2 and its relation to UTL that will be used in the remainder of the paper. Section 4 presents the logic-to-automata translations used in our upper bounds. The first is a translation of a given UTL formula to a large disjoint union of Büchi automata with certain structural restrictions. This can also be used to give a translation from a given FO^2 formula to an (still larger) union of Büchi automata. The second does something similar for $\text{FO}^2[<]$ formulas. The last translation maps $\text{FO}^2[<]$ and FO^2 formulas to deterministic parity automata, which is useful for certain problems involving games.

Section 6 gives upper and lower bounds for non-deterministic systems, while Section 7 is concerned with probabilistic systems. In Section 8 we consider model checking of $\text{FO}^2[\text{LTL}]$, which subsumes both FO^2 and LTL, and finally in Section 9 we consider the impact of extending all the previous logics with *let definitions*.

2. LOGIC, AUTOMATA AND COMPLEXITY CLASSES

We consider a first-order signature with set of unary predicates $\mathcal{P} = \{P_1, \dots, P_m\}$ and binary predicates $<$ (less than) and suc (successor). Fixing two distinct variables x and y , we denote by FO^2 the set of first-order formulas over the above signature involving only the variables x and y . We denote by $\text{FO}^2[<]$ the sublogic in which the binary predicate suc is not used. We write $\varphi(x)$ for a formula in which only the variable x occurs free.

In this paper we are interested in interpretations of FO^2 on infinite words. An ω -word $u = u_0u_1\dots$ over the powerset alphabet $\Sigma = 2^{\mathcal{P}}$ represents a first-order structure extending $\langle \mathbb{N}, <, \text{suc} \rangle$, in which predicate P_i is interpreted by the set $\{n \in \mathbb{N} : P_i \in u_n\}$ and the binary predicates $<$ and suc have the obvious interpretations.

We also consider Linear Temporal Logic LTL on ω -words. The formulas of LTL are built from atomic propositions using Boolean connectives and the temporal operators \bigcirc (*next*), \ominus (*previously*), \diamond (*eventually*), \diamondsuit (*sometime in the past*), \mathcal{U} (*until*), and \mathcal{S} (*since*). Formally, LTL is defined by the following grammar:

$$\varphi ::= P_i \mid \varphi \wedge \varphi \mid \neg \varphi \mid \varphi \mathcal{U} \varphi \mid \varphi \mathcal{S} \varphi \mid \diamond \varphi \mid \diamondsuit \varphi \mid \bigcirc \varphi \mid \ominus \varphi,$$

where P_0, P_1, \dots are propositional variables. Unary temporal logic (UTL) denotes the subset without \mathcal{U} and \mathcal{S} , while $\text{TL}[\diamond, \diamondsuit]$ denotes the *stutter-free* subset of UTL without \bigcirc and \ominus . We use $\square \varphi$ as an abbreviation for $\neg \diamond \neg \varphi$.

Let (u, i) be the suffix $u_i u_{i+1} \dots$ of ω -word u . We define the semantics of LTL inductively on the structure of the formulas as follows:

- (1) $(u, i) \models P_k$ iff atomic prop. P_k holds at position i of u

- (2) $(u, i) \models \varphi_1 \wedge \varphi_2$ iff $(u, i) \models \varphi_1$ and $(u, i) \models \varphi_2$
- (3) $(u, i) \models \neg\varphi$ iff it is not the case that $(u, i) \models \varphi$
- (4) $(u, i) \models \bigcirc\varphi$ iff $(u, i+1) \models \varphi$
- (5) $(u, i) \models \ominus\varphi$ iff $(u, i-1) \models \varphi$
- (6) $(u, i) \models \varphi_1 \mathcal{U} \varphi_2$ iff $\exists j \geq i$ s.t. $(u, j) \models \varphi_2$ and $\forall k, i \leq k < j$ we have $(u, k) \models \varphi_1$
- (7) $(u, i) \models \varphi_1 \mathcal{S} \varphi_2$ iff $\exists j \leq i$ s.t. $(u, j) \models \varphi_2$ and $\forall k, j < k \leq i$ we have $(u, k) \models \varphi_1$
- (8) $(u, i) \models \blacklozenge\varphi$ iff $(u, i) \models \text{true } \mathcal{U} \varphi$
- (9) $(u, i) \models \blacklozenge\varphi$ iff $(u, i) \models \text{true } \mathcal{S} \varphi$

It is well known that over ω -words LTL has the same expressiveness as first-order logic, and UTL has the same expressiveness as FO^2 . Moreover, while FO^2 is less expressive than LTL, it can be exponentially more succinct [EVW02] – for concrete examples of these facts, see the introduction.




Figure 1: Expressiveness Diagram

We can combine the succinctness of FO^2 and the expressiveness of LTL by extending the former with the temporal operators \mathcal{U} and \mathcal{S} (applied to formulas with at most one free variable). We call the resulting logic $\text{FO}^2[\text{LTL}]$. The syntax of $\text{FO}^2[\text{LTL}]$ divides formulas into two syntactic classes: *temporal formulas* and *first-order formulas*. Temporal formulas are given by the grammar

$$\varphi ::= P_i \mid \varphi \wedge \varphi \mid \neg\varphi \mid \varphi \mathcal{U} \varphi \mid \varphi \mathcal{S} \varphi \mid \psi,$$

where P_i is an atomic proposition and ψ is a first-order formula with one free variable. First-order formulas are given by the grammar

$$\psi ::= \varphi(x) \mid x < y \mid \text{suc}(x, y) \mid \psi \wedge \psi \mid \neg\psi \mid \exists x \psi,$$

where φ is a temporal formula. Here the first-order formula $\varphi(x)$ asserts that the temporal formula φ holds at position x . The temporal operators \bigcirc , \ominus , \blacklozenge and \blacklozenge can all be introduced as derived operators. An example of $\text{FO}^2[\text{LTL}]$ formula is:

$$b_0 \mathcal{U} (\exists y (y < x \wedge \bigwedge_{1 \leq i \leq n} b_i(x) \leftrightarrow b_i(y))).$$

The relative expressiveness of the logics defined thus far is illustrated in Figure 1.

Finally, we consider an extension of $\text{FO}^2[\text{LTL}]$ with *let definitions*. We inductively define the formulas and the unary predicate subformulas that occur *free* in such a formula. The atomic formulas of $\text{FO}^2[\text{LTL}]_{\text{Let}}$ are as in $\text{FO}^2[\text{LTL}]$, with the formula $P(x)$ occurring freely in itself. The constructors include all those of $\text{FO}^2[\text{LTL}]$, with the set of free subformula occurrences being preserved by all of these constructors.

There is one new formula constructor of the form:

$$\varphi ::= \text{Let } P_i(x) \text{ be } \varphi_1(x) \text{ in } \varphi_2$$

where P_i is a unary predicate, $\varphi_1(x)$ is an $\text{FO}^2[\text{LTL}]_{\text{Let}}$ formula in which x is the only free variable and no occurrence of predicate P_i is free, and φ_2 is an arbitrary $\text{FO}^2[\text{LTL}]_{\text{Let}}$ formula. A subformula $P_j(z)$ occurs freely in $\varphi(x)$ iff it occurs freely in $\varphi_1(x)$ or it occurs freely in φ_2 and the predicate is not P_i .

The semantics of $\text{FO}^2[\text{LTL}]_{\text{Let}}$ is defined via a translation function T to $\text{FO}^2[\text{LTL}]$, with the only non-trivial rule being:

$$T(\text{Let } P_i(x) \text{ be } \varphi_1(x) \text{ in } \varphi_2) ::= \\ T(\varphi_2[P_i(x) \mapsto T(\varphi_1), P_i(y) \mapsto T(\varphi_1)[x \mapsto y]])$$

where $T(\varphi_1)[x \mapsto y]$ denotes the formula obtained by substituting variable y for all free occurrences of x in $T(\varphi_1)$, and $T(\varphi_2[P_i(x) \mapsto T(\varphi_1), P_i(y) \mapsto T(\varphi_1)[x \mapsto y]])$ denotes substitution of any free occurrence of the form $P_i(x)$ in $T(\varphi_1)$ and every occurrence of $P_i(y)$ by $T(\varphi_1)[x \mapsto y]$. We let UTL_{Let} be the extension of UTL by the operator above, and similarly define $\text{TL}[\diamond, \heartsuit]_{\text{Let}}$, $\text{FO}^2[\prec]_{\text{Let}}$, etc.

For φ a temporal logic formula or an FO^2 formula with one free variable, we denote by $L(\varphi)$ the set $\{w \in \Sigma^\omega : (w, 0) \models \varphi\}$ of infinite words that satisfy φ at the initial position. The quantifier depth of an FO^2 formula φ is denoted $qdp(\varphi)$ and the operator depth of a UTL formula φ is denoted $odp(\varphi)$. In either case the length of the formula is denoted $|\varphi|$.

The notion of a subformula of an $\text{FO}^2[\text{LTL}]$ formula is defined as usual. For an $\text{FO}^2[\text{LTL}]_{\text{Let}}$ formula φ , let $\text{sub}(\varphi)$ denote the set of subformulas of the equivalent $\text{FO}^2[\text{LTL}]$ formula $T(\varphi)$, where T is the translation function defined above.

Lemma 2.1. *Given an LTL_{Let} formula φ , $|\text{sub}(\varphi)|$ is linear in $|\varphi|$.*

Proof. Notice that if $\varphi = \text{Let } P_i(x) \text{ be } \varphi_1(x) \text{ in } \varphi_2(x)$, then $|\text{sub}(\varphi)| \leq |\text{sub}(\varphi_1)| + |\text{sub}(\varphi_2)|$. Then by structural induction it holds that for a LTL_{Let} -formula φ , $\text{sub}(\varphi)$ has size at most $|\varphi|$. \square

Büchi Automata. Our results will be obtained via transforming formulas to automata that accept ω -words. We will be most concerned with *generalised Büchi automata* (GBA). A GBA A is a tuple $(\Sigma, S, S_0, \Delta, \mathcal{F})$ with alphabet Σ , set of states S , set of initial states $S_0 \subseteq S$, transition function Δ and set of sets of final states \mathcal{F} . The accepting condition is that for each $F \in \mathcal{F}$ there is a state $s \in F$ which is visited infinitely often. We can have labels either on states or on transitions, but both models are equivalent. For more details, see [VW86]. We will consider two important classes of Büchi automata: the automaton A is said to be *deterministic in the limit* if all states reachable from accepting states are deterministic; A is *unambiguous* if for each state s each word is accepted along at most one run that starts at s .

Deterministic Parity Automata. For some model checking problems, we will need to work with deterministic automata. In particular, we will use deterministic parity automata. A deterministic parity automaton A is a tuple $(\Sigma, S, s_0, \Delta, Pr)$ with alphabet Σ , set of states S , an initial state $s_0 \in S$, transition function Δ and a priority function Pr mapping each state to a natural number. The transition function Δ maps each state and symbol of the alphabet exactly to one new state. A run of such an automaton on input ω -word induces an infinite sequence of priorities. The acceptance condition is that the highest infinitely often occurring priority in this sequence is even.

Complexity Classes. Our complexity bounds involve *counting classes*. $\#P$ is the class of functions f for which there is a non-deterministic polynomial-time Turing Machine T such that $f(x)$ is the number of accepting computation paths of T on input x . A

complete problem for $\#P$ is $\#SAT$, the problem of counting the number of satisfying assignments of a given boolean formula. We will be considering computations of probabilities, not integers, so our problems will technically not be in $\#P$; but some of them will have representations computable in the related class $FP^{\#P}$, and will be $\#P$ -hard. For brevity, we will sometimes abuse notation by saying that such probability computation problems are $\#P$ -complete. The class of functions $\#EXP$ is defined analogously to $\#P$, except with T a non-deterministic exponential-time machine. We will deal with a decision version of $\#EXP$, $PEXP$, the set of problems solvable by nondeterministic Turing machine in exponential time, where the acceptance condition is that more than a half of computation paths accept [BFT98].

Notation: In our complexity bounds, we will often write *poly* to denote a fixed but arbitrary polynomial.

3. FO^2 MODEL THEORY AND SUCCINCTNESS

We now discuss the model theory of FO^2 , summarizing and slightly extending the material presented in Etessami, Vardi, and Wilke [EVW02] and in Weis and Immerman [WI09].

Recall that we will consider strings over alphabet $\Sigma = 2^{\mathcal{P}}$, where \mathcal{P} is the set of unary predicates appearing in the input $FO^2[<]$ formula. We start by recalling the small-model property of FO^2 that underlies the NEXPTIME satisfiability result of Etessami, Vardi, and Wilke [EVW02], it is also implicit in Theorem 6.2 of [WI09].

The *domain* of a word $u \in \Sigma^* \cup \Sigma^\omega$ is the set $\text{dom}(u) = \{i \in \mathbb{N} : 0 \leq i < |u|\}$ of positions in u . The *range* of u is the set $\text{ran}(u) = \{u_i : i \in \text{dom}(u)\}$ of letters occurring in u . Write also $\text{inf}(u)$ for the set of letters that occur infinitely often in u .

Given a finite or infinite word $u \in \Sigma^* \cup \Sigma^\omega$, a position $i \in \text{dom}(u)$, and $k \in \mathbb{N}$, we define the *k-type of u at position i* to be the set of $FO^2[<]$ formulas

$$\tau_k(u, i) = \{\varphi(x) : \text{qdp}(\varphi) = k \text{ and } (u, i) \models \varphi\}.$$

Given $u, v \in \Sigma^* \cup \Sigma^\omega$ and positions $i \in \text{dom}(u)$ and $j \in \text{dom}(v)$, write $(u, i) \sim_k (v, j)$ if and only if $\tau_k(u, i) = \tau_k(v, j)$. Furthermore, we write $u \sim_k v$ for two strings $u, v \in \Sigma^* \cup \Sigma^\omega$ if for all $FO^2[<]$ -formulas $\varphi(x)$ of quantifier depth at most k we have $(u, 0) \models \varphi$ iff $(v, 0) \models \varphi$.

The small model property of [EVW02] can then be stated as follows:

Proposition 3.1 ([EVW02]). *Let $\Sigma = 2^{\mathcal{P}}$. Then (i) For any string $u \in \Sigma^*$ and positive integer k there exists $v \in \Sigma^*$ such that $u \sim_k v$ and $|v| \in 2^{O(|\mathcal{P}|k)}$; (ii) for any infinite string $u \in \Sigma^\omega$ and positive integer k there are finite strings v and w , with $|v|, |w| \in 2^{O(|\mathcal{P}|k)}$, such that $u \sim_k vw^\omega$.*

For completeness, we give a constructive proof of Proposition 3.1, which will be used in one of our translations of FO^2 to automata. This is Lemma 3.9 at the end of this section. For this it is convenient to use the following inductive characterisation of \sim_k , which is proven in [EVW02] by a straightforward induction:

Proposition 3.2 ([EVW02]). *Let $u, v \in \Sigma^* \cup \Sigma^\omega$. Then $\tau_k(u, i) = \tau_k(v, j)$ if and only if (i) $u_i = v_j$, (ii) $\{\tau_{k-1}(u, i') : i' < i\} = \{\tau_{k-1}(v, j') : j' < j\}$, and (iii) $\{\tau_{k-1}(u, i') : i' > i\} = \{\tau_{k-1}(v, j') : j' > j\}$.*

The next proposition states that we can collapse any two positions in a string that have the same k -type without affecting the k -type of the string.

Proposition 3.3 ([EVW02]). *Let $u \in \Sigma^* \cup \Sigma^\omega$ and let $i < j$ be such that $(u, i) \sim_k (u, j)$. Writing $u = u_1 \dots u_j u'$, we have $u \sim_k u_1 \dots u_i u'$.*

From these two propositions it follows that every finite string is equivalent under \sim_k to a string of length exponential in k and $|\mathcal{P}|$.

Proposition 3.4. *Given a nonnegative integer k , for all strings $u \in \Sigma^*$ there exists a string $v \in \Sigma^*$ such that $u \sim_k v$ and $|v|$ is bounded by $2^{O(|\mathcal{P}|^k)}$.*

Proof. We prove by induction on k that the set $\{\tau_k(u, i) : i \in \text{dom}(u)\}$ of k -types occurring along u has size at most $|\Sigma|(2|\Sigma| + 2)^k$.

The base case $k = 0$ is clear.

For the induction step, assume that the number of $(k - 1)$ -types occurring along u is at most $|\Sigma|(2|\Sigma| + 2)^{k-1}$. Define a *boundary point* in u to be the position of the first or last occurrence of a given $(k - 1)$ -type. Then there are at most $2|\Sigma|(2|\Sigma| + 2)^{k-1}$ boundary points. But by Proposition 3.2 the k -type at a given position i in u is determined by u_i , the set of boundary points strictly less than i , and the set of boundary points strictly greater than i . Thus the number of k -types along u is at most

$$(|\Sigma| + 1)2|\Sigma|(2|\Sigma| + 2)^{k-1} = |\Sigma|(2|\Sigma| + 2)^k. \quad (3.1)$$

By Proposition 3.3, given any string v in which there are two distinct positions with the same k -type there exists a shorter string w with $v \sim_k w$. From the bound (3.1) on the number of boundary points, we conclude that there exists a string v such that $u \sim_k v$ and $|v| \leq |\Sigma|(2|\Sigma| + 2)^k \in 2^{O(|\mathcal{P}|^k)}$. \square

The relation \sim_k is also easy to compute:

Proposition 3.5. *Given $u, v \in \Sigma^*$ of length at most h we can compute whether $u \sim_k v$ in time at most $h2^{O(|\mathcal{P}|^k)}$.*

Proof. For $m = 0, 1, \dots, k$ we successively pass along u labelling each position i with its m -type $\tau_m(u, i)$. Each rank m requires two passes: we pass leftward through u computing the set of $(m - 1)$ -types to the left of each position, then we pass rightward computing the set of $(m - 1)$ -types to the right of each position. This requires $2k$ passes, with each pass taking time linear in h and at most quadratic in the number of k -types that occur along u . The bound now follows using the estimate of the number of types given in Proposition 3.4. \square

Combining Propositions 3.4 and 3.5 we get:

Corollary 3.6. *Given k there exists a set $\text{Rep}_k(\Sigma) \subseteq \Sigma^*$ of representative strings such that each $v \in \text{Rep}_k(\Sigma)$ has $|v| \leq |\Sigma|(2|\Sigma| + 2)^k$ and for each string $u \in \Sigma^*$ there exists a unique $v \in \text{Rep}_k(\Sigma)$ such that $u \sim_k v$. Moreover $\text{Rep}_k(\Sigma)$ can be computed from k in time $2^{2^{O(|\mathcal{P}|^k)}}$.*

The following result is classical, and can be proven using games.

Proposition 3.7. *Given $u, v \in \Sigma^*$ and $u', v' \in \Sigma^\omega$, for all k if $u \sim_k v$ and $u' \sim_k v'$ then $uu' \sim_k vv'$.*

From the above we infer that the equivalence class of an infinite string under \sim_k is determined by a prefix of the string and the set of letters appearing infinitely often within it.

Proposition 3.8. *Fix $k \in \mathbb{N}$. Given $u = u_0u_1 \dots \in \Sigma^\omega$, there exists $N \in \mathbb{N}$ such that for all $n \geq N$ and any word $w \in \Sigma^\omega$ with $\text{inf}(w) = \text{ran}(w) = \text{inf}(u)$ it holds that $u \sim_k u_0u_1 \dots u_nw$.*

Proof. Define a strictly increasing sequence of integers $n_0 < n_1 < \dots < n_k$ inductively as follows.

Let n_0 be such that for all $i \geq n_0$ letter u_i occurs infinitely often in u . For $0 < s \leq k$ let n_s be such that $\text{ran}(u_{n_{s-1}} \dots u_{n_s}) = \text{inf}(u)$. Now define $N := n_k$.

Let $n \geq N$ and let $v := u_0u_1 \dots u_nw$ for some w such that $\text{inf}(w) = \text{ran}(w) = \text{inf}(u)$. We claim that for all $0 \leq s \leq k$:

- (1) if $i \leq n_s$ then $\tau_s(u, i) = \tau_s(v, i)$;
- (2) if $i, j > n_s$ then $\tau_s(u, i) = \tau_s(v, j)$ if $u_i = v_j$.

This claim entails the proposition. We prove the claim by induction on s . The base case $s = 0$ is obvious.

The induction step for Clause 1 is as follows. Suppose that $i \leq n_s$; we must show that $\tau_s(u, i) = \tau_s(v, i)$. Certainly $u_i = v_i$ since u and v agree in the first N letters. Similarly for all $j < i$ we have $\tau_{s-1}(u, j) = \tau_{s-1}(v, j)$ by Parts 1 and 2 of the induction hypothesis. Now for all $j > i$ there exists $j' > i$ such that $u_j = v_{j'}$ and hence by Part 2 of the induction hypothesis $\tau_{s-1}(u, j) = \tau_{s-1}(v, j')$. We conclude that $\tau_s(u, i) = \tau_s(v, i)$ by Proposition 3.2.

The induction step for Clause 2 is as follows. Suppose that $i, j > n_s$ and $u_i = v_j$; we must show that $\tau_s(u, i) = \tau_s(v, j)$. We will again use Proposition 3.2. Certainly for all $i' > i$ there exists $j' > j$ such that $u_{i'} = v_{j'}$ and hence $\tau_{s-1}(u, i') = \tau_{s-1}(v, j')$. Now let $i' < i$. If $i' \leq n_s$ then $i' < j$, $u_{i'} = v_{i'}$ and hence $\tau_{s-1}(u, i') = \tau_{s-1}(v, i')$. Otherwise suppose $n_s < i' < i$. By definition of n_s there exists j' , $n_{s-1} < j' \leq n_s$ such that $u_{i'} = v_{j'}$. Then $\tau_{s-1}(u, i') = \tau_{s-1}(u, j')$ by Clause 2 of the induction hypothesis. \square

Combining Proposition 3.7 and Proposition 3.8, we complete the proof of Proposition 3.1, giving a slight strengthening of the conclusion for infinite words.

Lemma 3.9. *For any string $u \in \Sigma^\omega$ and positive integer k there exists $v \in \Sigma^*$ with $|v| \in 2^{O(|P|^k)}$ such that $v \sim_k u'$ for infinitely many prefixes u' of u , and $u \sim_k vw^\omega$, where w is a list of the letters occurring infinitely often in u .*

3.1. FO² and temporal logic. We now examine the relationship between FO² and UTL. Again we will be summarizing previous results while adding some new ones about the complexity of translation.

As mentioned previously, Etesami, Vardi and Wilke [EVW02] have studied the expressiveness and complexity of FO² on words. They show that FO² has the same expressiveness as unary temporal logic UTL, giving a linear translation of UTL into FO², and an exponential translation in the reverse direction.

Lemma 3.10 ([EVW02]). *Every FO² formula $\varphi(x)$ can be converted to an equivalent UTL formula φ' with $|\varphi'| \in 2^{O(|\varphi|(qdp(\varphi)+1))}$ and $odp(\varphi') \leq 2qdp(\varphi)$. The translation runs in time polynomial in the size of the output.*

With regard to complexity, [EVW02] shows that satisfiability for FO² over finite words or ω -words is NEXP-complete. The NEXP upper bound follows immediately from their “small model” theorem (see Proposition 3.1 stated earlier). NEXP-hardness is by reduction

from a tiling problem. This reduction requires either the use of the successor predicate, or consideration of models where an arbitrary Boolean combination of predicates can hold, that is, they consider words over an alphabet of the form $\Sigma = 2^{\{P_1, P_2, \dots, P_n\}}$.

The NEXP-hardness result for $\text{FO}^2[<]$ does not carry over from satisfiability to model checking since the collection of alphabet symbols that can appear in a word generated by the system being checked is bounded by the size of the system. However the complexity of model checking is polynomially related to the complexity of satisfiability when the latter is measured as a function of both formula size and alphabet size. Hence in the rest of the section we will deal with words over alphabet $\Sigma = \{P_0, P_1, \dots, P_n\}$, i.e., in which a unique proposition holds in each position. We call this the *unary alphabet restriction*.

One obvious approach to obtaining upper bounds for model checking $\text{FO}^2[<]$ would be to give a polynomial translation to $\text{TL}[\diamond, \diamond]$, and use logic-to-automata translation for $\text{TL}[\diamond, \diamond]$. Without the unary alphabet restriction an exponential blow-up in translating from $\text{FO}^2[<]$ to $\text{TL}[\diamond, \diamond]$ was shown necessary by Etessami, Vardi, and Wilke:

Proposition 3.11 ([EVW02]). *There is a sequence $(\psi_n)_{n \geq 1}$ of $\text{FO}^2[<]$ sentences over $\{P_1, P_2, \dots, P_n\}$ of size $O(n^2)$ such that the shortest temporal logic formula equivalent to ψ_n has size $2^{\Omega(n)}$.*

The sequence given in [EVW02] to prove the above theorem is

$$\psi_n = \forall x \forall y \left(\bigwedge_{i < n} (P_i(x) \leftrightarrow P_i(y)) \rightarrow (P_n(x) \leftrightarrow P_n(y)) \right).$$

In particular, their proof does not apply under the unary alphabet restriction. However below we show that the exponential blow-up is necessary even in this restricted setting. Our proof is indirect; it uses the following result about extensions of FO^2 with let definitions:

Lemma 3.12. *There is a sequence $(\varphi_n)_{n \geq 1}$ of $\text{FO}^2[<]_{\text{Let}}$ sentences mentioning predicates $\{P_1, P_2, \dots, P_n\}$ such that the shortest model of φ_n under the unary alphabet restriction has size $2^{\Omega(|\varphi_n|)}$.*

Proof. We define φ_n as follows.

$$\begin{aligned} \varphi_n = & \text{Let } R_1(x) \text{ be } \exists y (y \leq x \wedge P_1(y)) \text{ in} \\ & \text{Let } R_2(x) \text{ be } \exists y (y \leq x \wedge P_2(y) \wedge (R_1(x) \leftrightarrow R_1(y))) \text{ in} \\ & \dots \\ & \text{Let } R_n(x) \text{ be } \exists y \left(y \leq x \wedge P_n(y) \wedge \bigwedge_{k=1}^{n-1} (R_k(x) \leftrightarrow R_k(y)) \right) \\ & \text{in } \forall x \bigwedge_{i=1}^n \exists y \left((\neg(R_i(x) \leftrightarrow R_i(y)) \wedge \bigwedge_{j \neq i} (R_j(x) \leftrightarrow R_j(y))) \right) \end{aligned}$$

The body of the nested sequence of let definitions states that for all x and for all $1 \leq i \leq n$ there exists y such that the vector of formulas $(R_1(x), R_2(x), \dots, R_n(x))$ has the same truth value as the vector $(R_1(y), R_2(y), \dots, R_n(y))$ in all but position i . Hence the vector $(R_1(x), R_2(x), \dots, R_n(x))$ must take all 2^n possible truth values as x ranges over all positions in the word, i.e., any model of φ_n must have length at least 2^n .

We now claim that φ_n is satisfiable. To show this, recursively define a sequence of words $w^{(k)}$ over alphabet $\Sigma = \{P_0, P_1, \dots, P_n\}$ by $w^{(0)} = \varepsilon$ and $w^{(k+1)} = w^{(k)} P_{n-k} w^{(k)}$, where $0 \leq$

$k < n$. Finally write $w = w_n P_0$. Then the vector of truth values $(R_1(x), R_2(x), \dots, R_n(x))$ counts down from $2^n - 1$ to 0 in binary as one moves along w . \square

In contrast, we show that basic temporal logic enhanced with let definitions has the small model property:

Lemma 3.13. *There is a polynomial poly such that every satisfiable $\text{TL}[\diamond, \diamond]_{\text{Let}}$ formula φ has a model of size $\text{poly}(|\varphi|)$.*

Proof. In [EVW02, Section 5], Etessami, Vardi, and Wilke prove a small model property for $\text{TL}[\diamond, \diamond]$, which follows the same lines as the one given for FO^2 , but with polynomial rather than exponential bounds on sizes. Instead of using types based on quantifier-rank, the notion of type is based on the nesting of modalities; they thus look at modal k -type, where k is the nesting of modalities in φ . It was shown how to collapse infinite ω -words in order to get "smaller" ω -words with essentially the same type structure. Then in Lemma 4 of [EVW02] it is shown that for each $u \in \Sigma^\omega$ there are strings v, w such that the type of u at position 0 is equal to the type of vw^ω at position 0 and the length of both v and w is less than $(t + 1)^2$, where t is number of types occurring along u (that is, a polynomial version to Proposition 3.1).

The type is determined by the predicate and the combination of temporal subformulas of φ holding at the given position. Each temporal subformula, i.e. subformula which starts with \diamond or \diamond , can change its truth value at most once along the infinite word. Therefore there are at most polynomially many (in $|\Sigma|$ and in number of temporal subformulas of φ) different combinations and so also types along u .

Lemma 2.1 tells us that number of temporal subformulas of φ is linear in $|\varphi|$, and therefore the number of types t occurring along any word is polynomial in $|\varphi|$. Thus applying the above-mentioned type-collapsing argument of [EVW02] we conclude that there is a polynomial size model of φ . \square

The small model property for $\text{TL}[\diamond, \diamond]_{\text{Let}}$ will allow the lifting of NP model-checking results to this language. Most relevant to our discussion of succinctness, it can be combined with the previous result to show that $\text{FO}^2[<]$ is succinct with respect to $\text{TL}[\diamond, \diamond]$:

Proposition 3.14. *Even assuming the unary alphabet restriction, there is no polynomial translation from $\text{FO}^2[<]$ formulas to equivalent $\text{TL}[\diamond, \diamond]$ -formulas.*

Proof. Proof by contradiction. Assuming there were such a polynomial translation, we could apply it locally to the body of every let definition in an $\text{FO}^2[<]_{\text{Let}}$ formula. This would allow us to translate an $\text{FO}^2[<]_{\text{Let}}$ formula to a $\text{TL}[\diamond, \diamond]_{\text{Let}}$ formula of polynomial size. Therefore it would follow from Lemma 3.13 that every $\text{FO}^2[<]_{\text{Let}}$ formula that is satisfiable has a polynomial sized model, which is a contradiction of Lemma 3.12. \square

Proposition 3.14 shows that we cannot obtain better bounds for $\text{FO}^2[<]$ merely by translation to $\text{TL}[\diamond, \diamond]$. Weis [Wei11] showed an NP-bound on satisfiability of $\text{FO}^2[<]$ under the unary alphabet restriction (compared to NEXP-completeness of satisfiability in the general case). His approach is to show that models realise only polynomially many types. We will later show that the approach of Weis can be extended to obtain complexity bounds for model checking $\text{FO}^2[<]$ that are as low as one could hope, i.e., that match the complexity bounds for the simplest temporal logic, $\text{TL}[\diamond, \diamond]$. We do so by building sufficiently small unambiguous Büchi automata for $\text{FO}^2[<]$ formulas.

4. TRANSLATIONS

This section contains a key contribution of this paper—three logic-to-automata translations for UTL, FO^2 , and $\text{FO}^2[<]$. We will later use these translations to obtain upper complexity bounds for model checking both non-deterministic and probabilistic systems. As we will show, for most of the problems it is sufficient to translate a given formula to an unambiguous Büchi automaton. Our first translation produces such an automaton from a given UTL formula. This is then lifted to full FO^2 via a standard syntactic transformation from FO^2 to UTL. Our second translation goes directly from the stuffer-free fragment $\text{FO}^2[<]$ to unambiguous Büchi automata, and is used to obtain optimal bounds for this fragment. Our third translation constructs a deterministic parity automaton from an FO^2 formula. Having a deterministic automaton is necessary for solving two-player games and quantitative model checking of Markov decision processes.

4.1. Translation I: From UTL to unambiguous Büchi automata. We begin with a translation that takes UTL formulas to Büchi automata. Combining this with the standard syntactic transformation of FO^2 to UTL, we obtain a translation from FO^2 to Büchi automata.

Recall from the preliminaries section that a Büchi automaton A is said to be *deterministic in the limit* if all accepting states and their descendants are deterministic, and that A is *unambiguous* if each word has at most one accepting run.

We will aim at the following result:

Theorem 4.1. *Let φ be a UTL formula over set of propositions \mathcal{P} with operator depth n with respect to \bigcirc and \bigodot . Given an alphabet $\Sigma \subseteq 2^{\mathcal{P}}$, there is a family of at most $2^{|\varphi|^2}$ Büchi automata $\{A_i\}_{i \in I}$ such that (i) $\{w \in \Sigma^\omega : w \models \varphi\}$ is the disjoint union of the languages $L(A_i)$; (ii) A_i has at most $O(|\varphi||\Sigma|^{n+1})$ states; (iii) A_i is unambiguous and deterministic in the limit; (iv) there is a polynomial-time procedure that outputs A_i given input φ and index $i \in I$.*

We first outline the construction of the family $\{A_i\}$. Let φ be a formula of $\text{TL}[\diamond, \heartsuit]$ over set of atomic propositions \mathcal{P} . Following Wolper’s construction [Wol01], define $cl(\varphi)$, the *closure* of φ , to consist of all subformulas of φ (including φ) and their negations, where we identify $\neg\neg\psi$ with ψ . Furthermore, say that $\mathbf{s} \subseteq cl(\varphi)$ is a *subformula type* if (i) for each formula $\psi \in cl(\varphi)$ precisely one of ψ and $\neg\psi$ is a member of \mathbf{s} ; (ii) $\psi \in \mathbf{s}$ implies $\diamond\psi, \heartsuit\psi \in \mathbf{s}$; (iii) $\psi_1 \wedge \psi_2 \in \mathbf{s}$ iff $\psi_1 \in \mathbf{s}$ and $\psi_2 \in \mathbf{s}$. Given subformula types \mathbf{s} and \mathbf{t} , write $\mathbf{s} \sim \mathbf{t}$ if \mathbf{s} and \mathbf{t} agree on all formulas whose outermost connective is a temporal operator, i.e., for all formulas ψ we have $\diamond\psi \in \mathbf{s}$ iff $\diamond\psi \in \mathbf{t}$, and $\heartsuit\psi \in \mathbf{s}$ iff $\heartsuit\psi \in \mathbf{t}$. Note that these types are different from the types based on modal depth considered before.

Fix an alphabet $\Sigma \subseteq 2^{\mathcal{P}}$ and write tp_φ^Σ for the set of subformula types $\mathbf{s} \subseteq cl(\varphi)$ with $\mathbf{s} \cap P \in \Sigma$. In subsequent applications Σ will arise as the set of propositional labels in a structure to be model checked. Following [Wol01] we define a generalised Büchi automaton $A_\varphi^\Sigma = (\Sigma, S, S_0, \Delta, \lambda, \mathcal{F})$ such that $L(A_\varphi^\Sigma) = \{w \in \Sigma^\omega : (w, 0) \models \varphi\}$. The set of states is $S = tp_\varphi^\Sigma$, with the set S_0 of initial states comprising those $\mathbf{s} \in tp_\varphi^\Sigma$ such that (i) $\varphi \in \mathbf{s}$ and (ii) $\heartsuit\psi \in \mathbf{s}$ if and only if $\psi \in \mathbf{s}$ for any formula ψ . The state labelling function $\lambda : S \rightarrow \Sigma$ is defined by $\lambda(\mathbf{s}) = \mathbf{s} \cap P$. The transition relation Δ consists of those pairs (\mathbf{s}, \mathbf{t}) such that

- (i) $\heartsuit\psi \in \mathbf{t}$ iff either $\psi \in \mathbf{t}$ or $\heartsuit\psi \in \mathbf{s}$;
- (ii) $\diamond\psi \in \mathbf{s}$ and $\psi \notin \mathbf{s}$ implies $\diamond\psi \in \mathbf{t}$;

(iii) $\neg\Diamond\psi \in \mathbf{s}$ implies $\neg\Diamond\psi \in \mathbf{t}$.

The collection of accepting sets is $\mathcal{F} = \{F_{\Diamond\psi} : \Diamond\psi \in cl(\varphi)\}$, where $F_{\Diamond\psi} = \{\mathbf{s} : \psi \in \mathbf{s} \text{ or } \Diamond\psi \notin \mathbf{s}\}$.

A run of A_{φ}^{Σ} on a word $u \in \Sigma^{\omega}$ yields a function $f : \mathbb{N} \rightarrow 2^{cl(\varphi)}$. Moreover it can be shown that if the run is accepting then for all formulas $\psi \in cl(\varphi)$, $\psi \in f(i) \Rightarrow (u, i) \models \psi$ [Wol01, Lemma 2]. But since $f(i)$ contains each subformula or its negation, we have $\psi \in f(i)$ if and only if $(u, i) \models \psi$ for all $\psi \in cl(\varphi)$. We conclude that A_{φ}^{Σ} is unambiguous and accepts the language $L(\varphi)$. The following lemma summarises some structural properties of the automaton A_{φ}^{Σ} .

Lemma 4.2. *Consider the automaton A_{φ}^{Σ} as a directed graph with set of vertices S and set of edges Δ . Then (i) states \mathbf{s} and \mathbf{t} are in the same strongly connected component iff $\mathbf{s} \sim \mathbf{t}$; (ii) each strongly connected component has size at most $|\Sigma|$; (iii) the dag of strongly connected components has depth at most $|\varphi|$ and outdegree at most $2^{|\varphi|}$; (iv) A_{φ}^{Σ} is deterministic within each strongly connected component, i.e., given transitions (\mathbf{s}, \mathbf{t}) and (\mathbf{s}, \mathbf{u}) with \mathbf{s}, \mathbf{t} and \mathbf{u} in the same strongly connected component, we have $\mathbf{t} = \mathbf{u}$ if and only if $\lambda(\mathbf{t}) = \lambda(\mathbf{u})$.*

Proof. (i) If $\mathbf{s} \sim \mathbf{t}$ then by definition of the transition relation Δ we have that $(\mathbf{s}, \mathbf{t}) \in \Delta$. Thus \mathbf{s} and \mathbf{t} are in the same connected component. Conversely, suppose that \mathbf{s} and \mathbf{t} are in the same connected component. By clauses (i) and (iii) in the definition of the transition relation Δ we have that $\Diamond\psi \in \mathbf{s}$ iff $\Diamond\psi \in \mathbf{t}$ and likewise $\neg\Diamond\psi \in \mathbf{s}$ iff $\neg\Diamond\psi \in \mathbf{t}$. But for each formula $\psi \in cl(\varphi)$ either \mathbf{s} contains ψ or its negation, and similarly for \mathbf{t} ; it follows that $\mathbf{s} \sim \mathbf{t}$.

(ii) If $\mathbf{s} \sim \mathbf{t}$, then $\mathbf{s} = \mathbf{t}$ if and only if $\lambda(\mathbf{s}) = \lambda(\mathbf{t})$. Thus the number of states in an SCC is at most the number $|\Sigma|$ of labels.

(iii) Suppose that $(\mathbf{s}, \mathbf{t}) \in \Delta$ is an edge connecting two distinct SCC's, i.e., $\mathbf{s} \not\sim \mathbf{t}$. Then there is a subformula $\Diamond\psi \in \mathbf{s}$ such that $\neg\Diamond\psi \in \mathbf{t}$. Note that $\neg\Diamond\psi$ lies in all states reachable from \mathbf{t} under Δ . Since there are at most $|\varphi|$ such subformulas, we conclude that the depth of the DAG of SCC's is at most $|\varphi|$.

(iv) This follows immediately from (i). \square

We proceed to the proof of Theorem 4.1.

Proof. We first treat the case $n = 0$, i.e., φ does not mention \bigcirc or \ominus .

Let $A_{\varphi}^{\Sigma} = (\Sigma, S, S_0, \Delta, \lambda, \mathcal{F})$ be the automaton corresponding to φ , as defined above. For each path $\pi = C_0, C_1, \dots, C_k$ of SCC's in the SCC dag of A_{φ}^{Σ} we define a sub-automaton A_{π} as follows. A_{π} has set of states $S_{\pi} = C_0 \cup C_1 \cup \dots \cup C_k$; its set of initial states is $S_0 \cap S_{\pi}$; its transition relation is $\Delta_{\pi} = \Delta \cap (S_{\pi} \times S_{\pi})$, i.e., the transition relation of A_{φ}^{Σ} restricted to S_{π} ; its collection of accepting states is $\mathcal{F}_{\pi} = \{F \cap C_k : F \in \mathcal{F}\}$.

It follows from observations (ii) and (iii) in Lemma 4.2 that A_{π} has at most $|\varphi||\Sigma|$ states, and from observation (iii) that there are at most $2^{|\varphi|^2}$ such automata. Since A_{φ}^{Σ} is unambiguous, each accepting run of A_{φ}^{Σ} yields an accepting run of A_{π} for a unique path π , and so the $L(A_{\pi})$ partition $L(A_{\varphi}^{\Sigma})$.

Finally, A_{π} is deterministic in the limit since all accepting states lie in a bottom strongly connected component, and all states in such a component are deterministic by Lemma 4.2(iv). If we convert A_{π} from a generalised Büchi automaton to an equivalent

Büchi automaton (using the construction from [Wol01]), then the resulting automaton remains unambiguous and deterministic in the limit. This transformation touches only the bottom strongly connected component of A_π , whose size will become at most quadratic.

This completes the proof in case $n = 0$. The general case can be handled by reduction to this case. A UTL formula φ can be transformed to a normal form such that all next-time \bigcirc and last-time \ominus operators are pushed inside the other Boolean and temporal operators. Now the formula can be regarded as a TL $[\diamond, \diamond]$ formula φ' over an extended set of propositions $\{\bigcirc^i P, \ominus^i P : 0 \leq i \leq n, P \in \mathcal{P}\}$. Applying the case $n = 0$ to φ' we obtain a family of automata $\{A'_i\}$ over alphabet $\Sigma' = 2^{\mathcal{P}'}$ such that $L(A_{\varphi'}^{\Sigma'}) = \bigcup_i L(A'_i)$, A'_i is unambiguous and deterministic in the limit, and A'_i has at most $O(|\varphi'| |\Sigma'|) = O(|\varphi| |\Sigma|^n)$ states.

Now we can construct a deterministic transducer T with $|\Sigma|^n$ states that transforms (in the natural way) an ω -word over alphabet Σ into an ω -word over alphabet Σ' . Such a machine can be made deterministic by having T produce its output n positions behind the input. To do this we maintain an n -place buffer in the states of T , which requires $|\Sigma|^n$ states.

We construct automaton A_i over alphabet Σ by composing A'_i with T , i.e., by synchronising the output of T with the input of A'_i . The number of states of the composition is the product of the number of states of A'_i and T which are consistent with respect to their label in Σ' . Thus the product has at most $O(|\varphi| |\Sigma|^{n+1})$ states.

This completes the proof of Theorem 4.1. \square

From Theorem 4.1 we can get a translation of FO^2 to automata with bounds as stated below:

Theorem 4.3. *Given an FO^2 formula φ , there is a collection of $2^{2^{\text{poly}(|\varphi|)}}$ generalised Büchi automata A_i , each of size at most $2^{\text{poly}(|\varphi|)}$ such that the languages they accept partition the language $\{w \in \Sigma^\omega : w \models \varphi\}$. Moreover, each automaton A_i is unambiguous and can be constructed by a non-deterministic Turing machine in polynomial time in its size.*

Proof. First we apply Lemma 3.10 to translate the FO^2 formula φ to an equivalent UTL formula φ' . We then apply Theorem 4.1 to φ' , noting that the size of φ' is exponential in the size of φ , while the operator depth of φ' is polynomial in the quantifier depth of φ . Finally, we apply Theorem 4.1 to φ' . \square

4.2. Translation II: From $\text{FO}^2[<]$ to unambiguous Büchi automata. The previous translation via UTL will be useful for giving bounds on verifying both UTL and FO^2 . However it does not give insight into the sublanguage $\text{FO}^2[<]$. We will thus give another translation specific to this fragment. The main idea for getting upper bounds on verification problems for $\text{FO}^2[<]$ will be to show that for any $\text{FO}^2[<]$ formula φ , the number of one-variable subformula types realised along a finite or infinite word is polynomial in the size of φ . Informally these subformula types are the collections of one-variable subformulas of φ that might hold at a given position. Note that the types we consider here are collections of $\text{FO}^2[<]$ formulas, not temporal logic formulas as in the last section. Also note the contrast with the k -types of Proposition 3.1, which consider all formulas of a given quantifier rank.

Recall that the *domain* of a word $u \in \Sigma^* \cup \Sigma^\omega$ is the set $\text{dom}(u) = \{i \in \mathbb{N} : 0 \leq i < |u|\}$ of positions in u . Given an $\text{FO}^2[<]$ -formula φ , let $\text{cl}(\varphi)$ denote the set of all subformulas of φ with at most one free variable (including atomic predicates). Given a finite or infinite

word $u \in \Sigma^* \cup \Sigma^\omega$, a position $i \in \text{dom}(u)$, we define the *subformula type of u at position i* to be the set of $\text{FO}^2[<]$ formulas

$$\tau(u, i) = \{\psi : \psi \in \text{cl}(\varphi) \text{ and } (u, i) \models \psi\}.$$

We have omitted φ in this notation since it will be fixed for the remainder of the proof.

Few Types Property for $\text{FO}^2[<]$. We will base our result on the following theorem of Weis [Wei11], showing that $\text{FO}^2[<]$ formulas divide a finite word into a small number of segments based on subformula type:

Proposition 4.4 ([Wei11]). *Let φ be an $\text{FO}^2[<]$ -formula. A string $u \in \Sigma^*$ can be written $u = v_1 \dots v_n$, where $v_i \in \Sigma^*$, n is polynomial in $|\varphi|$ and $|\Sigma|$, and for any two positions i, j lying in the same factor v_k having the same symbol, $\tau(u, i) = \tau(u, j)$.*

We will need an extension of this result to infinite words:

Proposition 4.5. *Let φ be an $\text{FO}^2[<]$ -formula. A string $u \in \Sigma^\omega$ can be written $u = v_1 \dots v_n$, where $v_k \in \Sigma^*$ for $k < n$ and $v_n \in \Sigma^\omega$, n is polynomial in $|\varphi|$ and $|\Sigma|$, and for any two positions i, j lying within the same factor and having the same symbol we have $\tau(u, i) = \tau(u, j)$.*

Proof. We note that for any $u \in \Sigma^\omega$, from some position onwards, the subformula type is determined only by the current symbol. In fact, the proof of Proposition 3.8 shows that we have $u = vw$ for some prefix $v \in \Sigma^*$ of u and $w \in \Sigma^\omega$ such that for any two positions i, j of vw such that $i, j > |v|$ having the same symbol $\tau(vw, i) = \tau(vw, j)$.

Given an infinite u , we can take a finite prefix v as above and apply Proposition 4.4 to it, adding on the infinite interval w as one additional member of the partition. Now if i and j are in the final partition, then agreement on the same symbol determines the entire set of formulas, and hence we are done. Otherwise, fix any two positions $i, j \leq |v|$ in u with the symbol $a \in \Sigma$ holding at both i and j , and lying in the same partition within v . We claim that the subformula types $\tau(u, i)$ and $\tau(u, j)$ contain the same set of formulas. An atomic predicate $\psi \in \text{cl}(\varphi)$ holds at position i iff it holds at j by assumption, since there is only one symbol true at each position. Positions i and j then by assumption satisfy the same subformula type within v . But using the hypothesis on v we can easily see inductively that a subformula holds on a position within v iff it holds at that position within vw . \square

We now present a result showing that the few subformula types property can be used to get a better translation to automata:

Theorem 4.6. *Assume the unary alphabet restriction. Then given an $\text{FO}^2[<]$ formula φ , there is a collection of $2^{\text{poly}(|\varphi|, |\Sigma|)}$ generalised Büchi automata A_i (each of polynomial size in $|\varphi|$ and $|\Sigma|$) such that the languages they accept are disjoint and the union of these languages is exactly $\{w \in \Sigma^\omega : w \models \varphi\}$. Moreover, each automaton A_i is unambiguous and deterministic in the limit and can be constructed by a non-deterministic Turing machine in polynomial-time.*

Proof. We say that $\tau \subseteq \text{cl}(\varphi)$ is a *subformula pre-type* if: (i) if $\varphi_1 \wedge \varphi_2 \in \text{cl}(\varphi)$, then $\varphi_1 \wedge \varphi_2 \in \tau$ iff $\varphi_1 \in \tau$ and $\varphi_2 \in \tau$; (ii) if $\varphi_1 \vee \varphi_2 \in \text{cl}(\varphi)$, then $\varphi_1 \vee \varphi_2 \in \tau$ iff $\varphi_1 \in \tau$ or $\varphi_2 \in \tau$; (iii) if $\neg\psi \in \text{cl}(\varphi)$, then $\neg\psi \in \tau$ iff $\psi \notin \tau$.

This notion is similar to the notion of “subformula type of a node” used in the prior results, except that a collection of formulas satisfying the above property may not be consistent, since the semantics of existential quantifiers is not taken into account.

In general the formulas in a (subformula) pre-type τ can have either x or y as free variables. We write $\tau(x)$ for the subformula pre-type obtained by interchanging x and y in all formulas in τ with y as free variable. Thus all formulas in $\tau(x)$ have free variable x . We similarly define $\tau(y)$.

An *order formula* is an atomic formula

$$\alpha ::= x < y \mid y < x \mid x = y.$$

Given $m, n \in \mathbb{N}$ let $\alpha_{m,n}$ denote the unique order formula satisfied by the valuation $x, y \mapsto m, n$.

Given a pair of pre-types τ_1, τ_2 , an order formula α , and a subformula θ of φ , we write $\tau_1(x), \tau_2(y), \alpha \models \theta(x, y)$ to denote that when θ is transformed by replacing top-level subformulas by their truth values as specified by $\tau_1(x)$, $\tau_2(y)$, or α , then the resulting Boolean combination evaluates to true. Note that this implies that if word w and positions i, j satisfy $\tau_1(x) \cup \tau_2(y) \cup \{\alpha\}$, then they also satisfy θ .

A *closure labelling* is a function $f : \mathbb{N} \rightarrow 2^{cl(\varphi)}$ such that

- (1) $f(n)$ is a pre-type for each $n \in \mathbb{N}$ and
- (2) for each $n \in \mathbb{N}$, if $\exists y \theta \in cl(\varphi)$ then $\exists y \theta \in f(n)$ iff there exists $m \in \mathbb{N}$ such that $f(n)(x), f(m)(y), \alpha_{n,m} \models \theta$.

It is easy to see that an ω -word $w : \mathbb{N} \rightarrow \Sigma$ has a unique extension to a closure labelling $f : \mathbb{N} \rightarrow 2^{cl(\varphi)}$. Namely, f is defined by $f(n) = \{\psi \in cl(\varphi) : w, n \models \psi\}$.

We now define a generalised Büchi automaton A_φ corresponding to φ .

Definition 4.7. The alphabet of A_φ is Σ , and the other components of A_φ are as follows:

States. The states of A_φ are tuples $(\mathbf{s}, \tau, \mathbf{t})$, where $\tau \subseteq cl(\varphi)$ is a pre-type and $\mathbf{s}, \mathbf{t} \subseteq 2^{cl(\varphi)}$ are sets of pre-types of size at most $p(|\varphi|, |\Sigma|)$, where p is the polynomial from Proposition 4.5, such that the following *consistency condition* holds: for each formula $\exists y \theta \in \tau$ we have that either $\tau(x), \tau(y), x = y \models \theta$, $\tau(x), \tau'(y), x < y \models \theta$ for some $\tau' \in \mathbf{t}$, or $\tau'(y), \tau(x), y < x \models \theta$ for some $\tau' \in \mathbf{s}$. (This condition corresponds to the second clause in the definition of closure labelling.) Informally, a state consists of an assertion about the subformula pre-types seen in the past, the current subformula pre-type, and the subformula pre-types to be seen in the future.

Initial State. A state $(\mathbf{s}, \tau, \mathbf{t})$ is initial if $\mathbf{s} = \emptyset$ and $\varphi \in \tau$.

Accepting States. There is a set of accepting states F_τ for each pre-type τ . We have $(\mathbf{s}, \tau', \mathbf{t}) \in F_\tau$ if and only if $\tau = \tau'$ or $\tau \notin \mathbf{t}$.

Transitions. For each $a \in \Sigma$ there is an a -labelled transition from $(\mathbf{s}, \tau, \mathbf{t})$ to $(\mathbf{s}', \tau', \mathbf{t}')$ iff (i) for the unique proposition $P_i(x)$ in τ , $P_i = a$; (ii) $\mathbf{s}' = \mathbf{s} \cup \{\tau\}$; (iii) $\tau' \in \mathbf{t}$; (iv) either $\mathbf{t}' = \mathbf{t}$ or $\mathbf{t}' = \mathbf{t} \setminus \{\tau'\}$.

The following proposition, whose proof follows straightforwardly from Proposition 4.5, shows that the automaton captures the formula:

Proposition 4.8. *If $(\mathbf{s}_0, \tau_0, \mathbf{t}_0), (\mathbf{s}_1, \tau_1, \mathbf{t}_1), (\mathbf{s}_2, \tau_2, \mathbf{t}_2), \dots$ is an accepting run of A_φ , then the function $f : \mathbb{N} \rightarrow 2^{cl(\varphi)}$ defined by $f(n) = \tau_n$ is a closure labelling. Moreover every closure labelling f such that $\varphi \in f(0)$ arises from a run of A_φ in this manner.*

We now analyze the automaton A_φ . Because of the polynomial restriction on the number of pre-types, the automaton has at most exponentially many states. But by Proposition 4.5, any accepting run goes through only polynomially many states. For every path π in the DAG of strongly-connected components, we take the subautomaton A_π of A_φ obtained

by restricting to the components in this path. We claim that this is the required decomposition of A_φ . Note that an NP machine can construct these restrictions by iteratively making choices of successor components that are strictly lower in the DAG. Clearly the automata corresponding to distinct paths accept disjoint languages, since they correspond to different collections of pre-types holding in the word. One can show that for any word satisfying the formula, the unique accepting run is the one in which the state at a position corresponds to the pre-types seen before the positions, the pre-type seen at the position, and the pre-types seen after the position. In particular, this shows that each automaton is unambiguous. Finally, because the only nondeterministic choice is whether to leave an SCC or not, upon reaching the bottom SCC the automaton is deterministic—hence each automaton is deterministic in the limit. Thus this decomposition witnesses Theorem 4.6. \square

The above translation of $\text{FO}^2[\prec]$ formulas to unambiguous Büchi automata can be extended to handle formulas with successor, i.e., the full logic FO^2 , at the same time removing the unary alphabet restriction. Given an FO^2 formula φ over set of predicates \mathcal{P} , we can consider an “equivalent” $\text{FO}^2[\prec]$ formula φ' over a set of new predicates $2^{|\varphi||\mathcal{P}|}$. Intuitively each predicate in \mathcal{P}' specifies the truth values of all predicates in \mathcal{P} in a neighbourhood of radius $|\varphi|$ around the current position. Applying Theorem 4.6 to φ' we obtain a collection of double-exponentially many automata A_i , each of size exponential in φ and Σ . Thus, we get a weaker version of Theorem 4.3 of the previous subsection, in which the size bound on the component automata has an exponential dependence on the alphabet as well as the formula size.

4.3. Translation III: From FO^2 to deterministic parity automata. While the previous translations are useful for relating FO^2 to unambiguous automata, for some problems it is useful to have deterministic automata. We now give a translation of FO^2 formulas to “small” deterministic parity automata. We give the translation first for the fragment $\text{FO}^2[\prec]$ without successor and show later how to handle the full logic. Specifically, we will show:

Theorem 4.9. *Given an $\text{FO}^2[\prec]$ formula φ over set of predicates \mathcal{P} with quantifier depth k , there exists a deterministic parity automaton \mathcal{A}_φ accepting the language $L(\varphi)$ such that \mathcal{A}_φ has $2^{2^{O(|\mathcal{P}|k)}}$ states, $2^{O(|\mathcal{P}|)}$ priorities, and can be computed from φ in time $|\varphi|^{O(1)} \cdot 2^{2^{O(|\mathcal{P}|k)}}$.*

The definition of the automaton \mathcal{A}_φ in Theorem 4.9 relies on the small-model property, as stated in Proposition 3.1. By Lemma 3.9, to know whether $u \in \Sigma^\omega$ satisfies an $\text{FO}^2[\prec]$ -formula of quantifier depth k it suffices to know some k -type such that infinitely many prefixes of u have that type, as well as which letters occur infinitely often in u . We will translate φ to a deterministic parity automaton \mathcal{A}_φ that detects this information. As \mathcal{A}_φ reads an input string u it stores a representative of the k -type of the prefix read so far. By Proposition 3.1(i) the number of such representatives is bounded by $2^{2^{O(|\mathcal{P}|k)}}$. Applying Lemma 3.9, we use a parity acceptance condition to determine whether u satisfies φ , based on which representatives and input letters occur infinitely often.

We are now ready to formally define \mathcal{A}_φ . To this end, define the *last appearance record* of a finite string $u = u_0 \dots u_n \in \Sigma^*$ to be the substring $\text{LAR}(u) := u_{i_1} u_{i_2} \dots u_{i_m}$ such that for all $k \in \text{dom}(u)$ there exists a unique $i_j \geq k$ such that $u_{i_j} = u_k$. Thus we obtain $\text{LAR}(u)$ from u by keeping only the last occurrence of each symbol from u . Write $\text{LAR}(\Sigma)$ for the

set $\{\text{LAR}(u) : u \in \Sigma^*\}$ of all possible last appearance records. Recall also the set of strings $\text{Rep}_k(\Sigma)$ from Corollary 3.6 that represent the different k -types of strings in Σ^* .

Definition 4.10. Let $\varphi(x)$ be an $\text{FO}^2[<]$ -formula of quantifier depth k . We define a deterministic parity automaton \mathcal{A}_φ as follows.

- \mathcal{A}_φ has set of states $\text{Rep}_k(\Sigma) \times \text{LAR}(\Sigma) \times \{0, 1, \dots, |\Sigma|\}$.
- The initial state is $(\varepsilon, \varepsilon, 0)$.
- The transition function maps a state (s, ℓ, i) , where $\ell = \ell_1 \dots \ell_j$, and input letter $a \in \Sigma$ to the unique state (t, ℓ', j') such that $sa \sim_k t$, $\ell' = \text{LAR}(a\ell)$, $j' = 0$ if a does not occur in ℓ and otherwise $\ell_{j'} = a$.
- The set of priorities is $0, 1, \dots, 2|\Sigma| + 1$.
- The priority of state (s, ℓ, i) where $\ell = \ell_1 \ell_2 \dots \ell_j$ is given by

$$\text{pr}(s, \ell, i) = \begin{cases} 2 \cdot |\ell_i \dots \ell_j| & \text{if } (s(\ell_i \dots \ell_j)^\omega, 0) \models \varphi \\ 2 \cdot |\ell_i \dots \ell_j| + 1 & \text{otherwise.} \end{cases}$$

It follows from Proposition 3.7 that in a run of \mathcal{A}_φ on a finite word $u = u_0 u_1 \dots u_n \in \Sigma^*$ the last state (s, ℓ, i) is such that s has the same k -type as u . Also we note that ℓ is the LAR of u and i is the position in the previous LAR of u_n .

The following two results prove Theorem 4.9:

Proposition 4.11. $L(\mathcal{A}_\varphi) = \{u \in \Sigma^\omega : (u, 0) \models \varphi\}$.

Proof. Let $u \in \Sigma^\omega$ and let N be as in Proposition 3.8. Suppose that the highest infinitely often occurring priority in a run of \mathcal{A}_φ on u is even. Then there exists $n \geq N$ such that \mathcal{A}_φ is in state (s, ℓ, i) after reading $u_0 u_1 \dots u_n$, where $\ell = \ell_1 \ell_2 \dots \ell_j$, $\{\ell_i, \dots, \ell_j\} = \text{inf}(u)$ and $(s(\ell_i \dots \ell_j)^\omega, 0) \models \varphi$. Now

$$\begin{aligned} u &\sim_k u_0 u_1 \dots u_n (\ell_i \dots \ell_j)^\omega && \text{by Proposition 3.8} \\ &\sim_k s(\ell_i \dots \ell_j)^\omega && \text{by Proposition 3.7.} \end{aligned}$$

We conclude that $(u, 0) \models \varphi$.

Similarly we can show that if the highest infinitely often occurring priority in a run of \mathcal{A}_φ on u is odd then $(u, 0) \not\models \varphi$. \square

Proposition 4.12. *If φ over set of monadic predicates \mathcal{P} has quantifier depth k , then \mathcal{A}_φ has number of states at most $2^{2^{O(|\mathcal{P}|^k)}}$ and can be computed from φ in time $|\varphi|^{O(1)} \cdot 2^{2^{O(|\mathcal{P}|^k)}}$.*

Proof. The set of states $\text{Rep}_k(\Sigma)$ has size at most $2^{2^{O(|\mathcal{P}|^k)}}$ and can be constructed in time at most $2^{2^{O(|\mathcal{P}|^k)}}$ by Corollary 3.6. We can establish the existence of a transition between any pair of states of \mathcal{A}_φ in time at most $2^{O(|\mathcal{P}|^k)}$ by Proposition 3.5. Finally we can compute the priority of a state (s, ℓ, i) by model checking φ on a lasso of length at most $2^{O(|\mathcal{P}|^k)}$, which can be done in time $|\varphi|^{O(1)} \cdot 2^{O(|\mathcal{P}|^k)}$. \square

Extension to FO^2 with successor. We now extend to successor using the same approach as in the proof of Theorem 4.1. By Lemma 3.10, given an FO^2 formula φ of quantifier depth k there is an equivalent UTL formula φ' of at most exponential size and operator depth at most $2k$. Moreover, φ' can be transformed to a normal form such that all next-time \bigcirc and last-time \ominus operators are pushed inside the other operators. Again, we consider φ' also as a $\text{TL}[\diamond, \heartsuit]$ -formula over an extended set of predicates $\mathcal{P}' = \{\bigcirc^i P_j, \ominus^i P_j \mid P_j \in \mathcal{P}, i \leq k\}$.

By a straightforward transformation we get an equivalent $\text{FO}^2[<]$ formula φ'' over P' . Overall, this transformation creates exponentially larger formulas, but the quantifier depth is only doubled and the set of predicates is quadratic. Applying Theorem 4.9 for φ'' over set of predicates \mathcal{P}' gives:

Theorem 4.13. *Given an FO^2 formula φ with quantifier depth k , there is a deterministic parity automaton having $2^{2^{O(k^2|\mathcal{P}|)}}$ states and $2^{O(k|\mathcal{P}|)}$ priorities that accepts the language $L(\varphi)$.*

5. MODELS CONSIDERED

Next we collect together definitions of the various different types of state machine that we consider in this paper. For non-deterministic machines we will be interested in the existence of an accepting path through the machine that satisfies a formula, while for probabilistic models we want to know the probability of such paths.

Kripke Structures, Hierarchical and Recursive State Machines. Our most basic model of non-deterministic computation is a Kripke structure, which is just a graph with an additional set of nodes (the initial states), and a labelling of nodes with a subset of a collection of propositions. The behavior represented by such a structure is the set of paths through the graph, where paths can be seen as ω -words.

We will look also at more expressive and succinct structures for representing behaviours. A recursive state machine (RSM) \mathcal{M} over a set of propositions \mathcal{P} is given by a tuple (M_1, \dots, M_k) where each component state machine $M_i = (N_i \cup B_i, Y_i, X_i, En_i, Ex_i, \delta_i)$ contains

- a set N_i of *nodes* and a disjoint set B_i of *boxes*;
- an indexing function $Y_i : B_i \mapsto \{1, \dots, k\}$ that assigns to every box an index of one of the component machines, M_1, \dots, M_k ;
- a labelling function $X_i : N_i \mapsto 2^{\mathcal{P}}$;
- a set of *entry nodes* $En_i \subseteq N_i$ and a set of *exit nodes* $Ex_i \subseteq N_i$;
- a *transition relation* δ_i , where transitions are of the form (u, v) where the source u is either a node of N_i , or a pair (b, x) , where b is a box in B_i and x is an exit node in Ex_j for $j = Y_i(b)$. We require that the destination v be either a node in N_i or a pair (b, e) , where b is a box in B_i and e is an entry node in En_j for $j = Y_i(b)$.

Informally, an RSM represents behaviors that can transition through a box into the entry node of the machine called by the box, and can transition via an exit node back to the calling box, as with function calls. The semantics can be found in [ABE⁺05]. A hierarchical state machine (HSM) is an RSM in which the dependency relation between boxes is acyclic. HSMs have the same expressiveness as flat state machines, but can be exponentially more succinct.

Markov Chains. The basic probabilistic model corresponding to a Kripke structure is a (labelled) *Markov chain*, specified as $\mathcal{M} = (\Sigma, X, V, E, M, \rho)$, consisting of an *alphabet* Σ , a set X of *states*; a *valuation* $V : X \rightarrow \Sigma$; a set $E \subseteq X \times X$ of *edges*; a *transition probability* M_{xy} for each pair of states $(x, y) \in E$ such that for each state x , $\sum_y M_{xy} = 1$; an *initial probability distribution* ρ on the set of states X .

A Markov chain defines a probability distribution on trajectories—paths through the chain. Given a language $L \subseteq \Sigma^\omega$, we denote by $P_{\mathcal{M}}(L)$ the probability of the set of

trajectories of \mathcal{M} whose image under V lies in L . We consider the complexity of the following model checking problem: Given a Markov chain \mathcal{M} and an LTL- or FO^2 -formula φ , calculate $P_{\mathcal{M}}(L(\varphi))$. There is a decision version of this problem that asks whether this probability exceeds a given rational threshold.

Recursive Markov Chains. Recursive Markov chains (RMCs) are the analog of RSMs in the probabilistic context. They are defined as RSMs, except that the transition relation consists of triples $(u, p_{u,v}, v)$ where u and v are as with RSMs, and the $p_{u,v}$ are non-negative reals with $\sum_v p_{u,v} = 1$ or 0 for every u . As with Markov chains, these define a probability distribution on trajectories, but now trajectories are paths which must obey the box-entry/box-exit discipline of an RSM. The semantics of an RMC can be found in [EY05]. A hierarchical Markov chain (HMC) is the probabilistic analog of an HSM, that is, an RMC in which the calling graph is acyclic. An HMC can be converted to an ordinary Markov chain via unfolding, possibly incurring an exponential blow-up. An example of an RMC is shown in Figure 2.




Figure 2: A sample Recursive Markov Chain

Markov Decision Processes. We will also deal with verification problems related to control of a probabilistic process by a scheduler. A *Markov decision process (MDP)* $\mathcal{M} = (\Sigma, X, N, R, V, E, M, \rho)$ consists of an *alphabet* Σ , a set X of *states*, which is partitioned into a set N of *non-deterministic states* and a set R of *randomising states*; a *valuation* $V : X \rightarrow \Sigma$, a set $E \subseteq X \times X$ of *edges*, a *transition probability* M_{xy} for each pair of states $(x, y) \in E$, $x \in R$ such that $\sum_y M_{xy} = 1$; an *initial probability distribution* ρ . This model is considered in [CY95] under the name *Concurrent Markov chain*.

We can view non-deterministic states as being controlled by the scheduler, which given a trajectory leading to a non-deterministic state s chooses a transition out of s . There are two basic qualitative model checking problems: the *universal problem* (\forall) asks that a given formula be satisfied with probability one for all schedulers; the *existential problem* (\exists) asks that the formula be satisfied with probability one for some scheduler. The latter corresponds to the problem of designing a system that behaves correctly in a probabilistic environment. In the *quantitative model checking problem*, we ask for the maximal probability for the formula to be satisfied on a given MDP when the scheduler chooses optimal moves in the non-deterministic states.

Two-player Games. A *two-player game* $G = (\Sigma, X, X_1, X_2, V, E, x_0)$ consists of an *alphabet* Σ ; a set X of *states*, which is partitioned into a set X_1 of states controlled by *Player I* and a set X_2 controlled by *Player II*; a set of $E \subseteq X \times X$ of *transitions*; a *valuation* $V : X \rightarrow \Sigma$; an *initial state* x_0 .

The game starts in the initial state and then the player who controls the current state, taking into account the whole history of the game, chooses one of the possible transitions. The verification problem of interest is whether Player I has a strategy such that for all infinite plays the induced infinite word $u \in \Sigma^\omega$ satisfies a given formula φ .

Stochastic Two-player Games. A *Stochastic two-player game* ($2\frac{1}{2}$ -player game) $G = (X, X_1, X_2, R, V, E, M, p_0)$ consists of a set X of *states*, which is partitioned into a set X_1 of states controlled by *the first player*, a set X_2 controlled by *the second player* and a set R of *randomising states*; a *valuation* $V : X \rightarrow \Sigma$; a set of $E \subseteq X \times X$ of *transitions*, a *transition probability* M_{xy} for each pair of states $(x, y) \in E$, $x \in R$ such that $\sum_y M_{xy} = 1$; an *initial probability distribution* ρ . See Figure 3 for an example.

The *universal* (\forall) qualitative model checking problem asks if the first player can enforce that the infinite word u , induced by the path through the game, satisfies φ with probability one.




Figure 3: A sample Stochastic Two-player Game. Diamonds are states of the first player, squares are states of the second player and circles represent randomising states.

6. VERIFYING NON-DETERMINISTIC SYSTEMS

Model checking for traditional Kripke Structures is fairly well-understood. All of our logics subsume propositional logic, and the model checking problems we deal with generalise propositional satisfiability—hence they are all NP-hard. LTL and UTL are both PSPACE-complete [SC82], while $(\text{TL}[\diamond, \heartsuit])$ is NP-complete.

Translation I shows how to convert an FO^2 formula to a union of exponential sized automata. A NEXPTIME algorithm can guess such an automaton, take its product with a given Kripke Structure, and then determine non-emptiness of the resulting product. Coupled with the hardness argument in [EVW02], this gives an alternative proof of the result of Etessami, Vardi, and Wilke:

Theorem 6.1. [EVW02] *FO^2 model-checking is complete for NEXPTIME.*

Below we extend these results to give a comparison of the complexity of model checking for recursive state machines and two-player games, applying all of the translations in the previous section.

6.1. Recursive State Machines. Using Translation II, we show that $\text{FO}^2[<]$ model checking can be done as efficiently as for $\text{TL}[\diamond, \heartsuit]$ on non-deterministic systems, and in particular for RSMs.

Proposition 6.2. *Model checking $\text{FO}^2[<]$ properties on Kripke structures, hierarchical and recursive state machines is in NP.*

Proof. We give the upper bound for RSMs only, since the other classes are special cases. We describe an NP algorithm that checks satisfiability of an $\text{FO}^2[<]$ sentence φ on the language of RSM \mathcal{M} . Model checking the structure involves only combinations of propositions occurring in the structure, and hence by expanding out these combinations explicitly, we can assume that the unary alphabet restriction holds. Thus we can apply Translation II, from $\text{FO}^2[<]$ to Büchi Automata, Theorem 4.6. It suffices to check that one of the automata A_i produced by the translation accepts a word produced by \mathcal{M} . We can thus guess such an A_i and can then check intersection of A_i with \mathcal{M} in polynomial time, by forming the product and checking that we can reach an accepting bottom strongly connected component. This reachability analysis can be done efficiently using the “summary edge construction”—see, e.g., [ABE⁺05]. \square

In the same way, we can obtain the result for model checking full FO^2 on RSMs, but now using the FO^2 to automata translation in Translation 1, Theorem 4.3. Again we guess an automata A_i , which is now of exponential size. Thus we have:

Proposition 6.3. *FO^2 model checking of RSMs can be done in NEXPTIME.*

This result matches the known result for ordinary Kripke structures.

6.2. Two-player games with FO^2 winning condition. Two-player games are known to be in 2EXPTIME for LTL [PR89]. We now show that the same is true for FO^2 , making use of Translation III in the previous section, which translates to deterministic parity automata. We also utilise the fact that a parity game with n vertices, m edges and d priorities can be solved in time $O(dmn^d)$ [Jur00].

From these two results we easily conclude the 2EXPTIME upper bound:

Proposition 6.4. *Two-player games with FO^2 winning conditions are solvable in 2EXPTIME.*

Proof. Using Theorem 4.13, we construct in 2EXPTIME a deterministic parity automaton for the FO^2 formula φ with doubly exponentially many states and at most exponentially many priorities. By taking the product of this automaton with the graph of the game, we get a parity game with doubly exponentially many states but only exponentially many priorities. (In fact if we define the automaton over an alphabet $\Sigma \subseteq 2^P$ containing only sets of propositions that occur as labels of states in the game, then polynomially many priorities suffice.) We can then determine the winner in double exponential time, again applying the $O(dmn^d)$ bound for solving games of [Jur00] mentioned above. \square

Combining this with the result by Alur, La Torre, and Madhusudan, who showed that two-player games are 2EXPTIME-hard [ATM03] already for the simplest $\text{TL}[\diamond, \diamond]$, along with the fact that we can convert UTL formula to FO^2 formula in polynomial time, we get 2EXPTIME-completeness:

Corollary 6.5. *Deciding two-player games with FO^2 winning conditions is complete for 2EXPTIME.*

The table below summarises both the known results and the results from this paper (in bold) concerning non-deterministic systems. All bounds are tight.

	$\text{TL}[\diamond, \diamond]$	UTL	$\text{FO}^2[\langle]$	FO^2	LTL
Kripke Structure	NP	PSPACE	NP	NEXP	PSPACE
HSM	NP	PSPACE	NP	NEXP	PSPACE
RSM	NP	EXP	NP	NEXP	EXP
Two pl. games	2EXP	2EXP	2EXP	2EXP	2EXP

The PSPACE bound for model checking LTL on HSMs follows by expanding the HSMs to ‘flat’ Kripke structures and recalling that model checking LTL on Kripke structures can be done in space polynomial in the logarithm of the model size. Additionally, the complexity of model checking UTL and LTL on RSMs is EXPTIME-complete [BEM97], and model checking $\text{TL}[\diamond, \diamond]$ on RSMs is NP-complete [LTP07].

7. VERIFYING PROBABILISTIC SYSTEMS

We now turn to probabilistic systems. Here we will make use of two key properties of the automata produced by the first two translations—unambiguity and determinism in the limit. We will need two lemmas, which show that the complexity bounds for model checking unambiguous Büchi automata on various probabilistic systems are the same as the bounds for deterministic Büchi automata on these systems. First, following [CSS03], we note the following property of unambiguous automata:

Lemma 7.1. *Given a Markov chain $\mathcal{M} = (\Sigma, X, V, E, M, \rho)$ and a generalised Büchi automaton $A = (\Sigma, S, S_0, \Delta, \lambda, \mathcal{F})$ that is unambiguous, $P_{\mathcal{M}}(L(A))$ can be computed in time polynomial in \mathcal{M} and A .*

Proof. We define a directed graph $\mathcal{M} \otimes A$ representing the synchronised product of \mathcal{M} and A . The vertices of $\mathcal{M} \otimes A$ are pairs $(x, s) \in X \times S$ with matching propositional labels, i.e., such that $V(x) = \lambda(s)$; the set of directed edges is $\{((x, s), (y, t)) : (x, y) \in E \text{ and } (s, t) \in \Delta\}$. We say that a strongly connected component (SCC) of $\mathcal{M} \otimes A$ is *accepting* if (i) for each set of accepting states $F \in \mathcal{F}$ it contains a pair (x, s) with $s \in F$ and (ii) for each pair (x, s) and each transition $(x, y) \in E$, there exists $(s, t) \in \Delta$ such that (y, t) is in the same SCC as (x, s) . This guarantees that we can stay in the SCC and visit each of its states infinitely often.

Let $L(A, s)$ denote the set of words accepted by A starting in state s . For each vertex (x, s) of $\mathcal{M} \otimes A$ we have a variable $\xi_{x,s}$ representing the probability $P_{\mathcal{M},x}(L(A, s))$ of all runs of \mathcal{M} starting in state x that are in $L(A, s)$. These probabilities can be computed as

the unique solution of the following linear system of equations:

$$\begin{aligned} \xi_{x,s} &= 1 && (x, s) \text{ in an accepting SCC} \\ \xi_{x,s} &= 0 && (x, s) \text{ in a non-accepting SCC} \\ \xi_{x,s} &= \sum_{(s,t) \in \Delta} \sum_{y:V(y)=\lambda(t)} M_{xy} \cdot \xi_{y,t} && \text{otherwise.} \end{aligned}$$

The correctness of the third equation follows from the following calculation:

$$\begin{aligned} P_{\mathcal{M},x}(L(A, s)) &= P_{\mathcal{M},x}\left(\bigcup_{(s,t) \in \Delta} \lambda(s) \cdot L(A, t)\right) \\ &= \sum_{(s,t) \in \Delta} P_{\mathcal{M},x}(\lambda(s) \cdot L(A, t)) \quad (\text{since } A \text{ is unambiguous}) \\ &= \sum_{(s,t) \in \Delta} \sum_{y:V(y)=\lambda(t)} M_{xy} \cdot P_{\mathcal{M},y}(L(A, y)). \quad \square \end{aligned}$$

For an RMC \mathcal{M} , we can compute reachability probabilities $q_{(u,ex)}$ of exiting a component M_i starting at state $u \in V_i$ going to exit $ex \in Ex_i$. Etessami and Yannakakis [EY05] show that these probabilities are the unique solution of a system of non-linear equations which can be found in polynomial space using a decision procedure for the existential theory of the reals. Following [EY05] for every vertex $u \in V_i$ we let $ne(u) = 1 - \sum_{ex \in Ex_i} q_{(u,ex)}$ be the probability that a trajectory beginning from node u never exits the component M_i of u . Etessami and Yannakakis [YE05] also show that one can check properties specified by deterministic Büchi automata in PSPACE, while for non-deterministic Büchi automata they give a bound of EXPSPACE. Thus the prior results would give a bound of EXPSPACE for UTL and 2EXPSPACE for FO². We will improve upon both these bounds. We observe that the technique of [YE05] can be used to check properties specified by non-deterministic Büchi automata that are unambiguous in the same complexity as deterministic ones. This will then allow us to apply our logic-to-automata translations.

Proposition 7.2. *Given an unambiguous Büchi automaton A and a RMC \mathcal{M} , we can compute the probability that A accepts a trajectory of \mathcal{M} in PSPACE.*

Proof. Let A be an unambiguous Büchi automaton with set of states Q , transition function Δ and labelling function λ . Let \mathcal{M} be an RMC with valuation V . We define a product RMC $\mathcal{M} \otimes A$ with component and call structure coming from \mathcal{M} whose states are pairs (x, s) , with x a state of \mathcal{M} and s a state of A such that $V(x) = \lambda(s)$ (i.e., x and s have the same label). Such a pair (x, s) is accepting if s is an accepting state of A . A run through the product chain is accepting if at least one of the accepting states is visited infinitely often. Note that a path through \mathcal{M} may expand to several runs in $\mathcal{M} \otimes A$ since A is non-deterministic.

For each i , for each vertex $x \in V_i$, exit $ex \in Ex_i$ and states $s, t \in Q$ we define $p(x, s \rightarrow ex, t)$ to be the probability that a trajectory in RMC \mathcal{M} that begins from a configuration with state x and some non-empty context (i.e. not at top-level) expands to an accepting run in $\mathcal{M} \otimes A$ from (x, s) to (ex, t) .

Just as in the case of deterministic automata, we can compute $p(x, s \rightarrow ex, t)$ as the solution of the following system of non-linear equations:

If $x \in V_i$ is not entrance of the box we have:

$$p(x, s \rightarrow ex, t) = \sum_{x':(x, M_{xx'}, x') \in \delta_i} M_{xx'} \sum_{s':(s, s') \in \Delta \wedge \lambda(s') = V(x')} p(x', s' \rightarrow ex, t)$$

If $x \in V_i$ is entrance of the box $b \in B_i$ then we include the equations:

$$p(x, s \rightarrow ex, t) = \sum_{j, s' \in Q} p((b, en), s \rightarrow (b, ex_j), s') p((b, ex_j), s' \rightarrow ex, t)$$

where $p((b, en), s \rightarrow (b, ex_j), s') = p(en_{Y_i(b)}, s \rightarrow ex_j, s')$ and $ex_j \in Ex_{Y_i(b)}$.

The justification for these equations is as follows. Since A is unambiguous, each trajectory of \mathcal{M} expands to at most one accepting run of $\mathcal{M} \otimes A$. Thus in summing over automaton states s' in the two equations above we are summing probabilities over disjoint events which correctly gives us the probability of the union of these events.

We now explain how these probabilities can be used to compute the probability of acceptance. We assume without loss of generality that the transition function of A is total.

We construct a finite-state *summary chain* for the product $\mathcal{M} \otimes A$ exactly as in the case of deterministic automata [YE05]. For each component M_i of \mathcal{M} , vertex x of M_i , exit $ex \in Ex_i$ and for each pair of states s, t of A the probability to transition from (x, s) to (ex, t) in the summary chain is calculated from $p(x, s \rightarrow ex, t)$ after adjusting for probability $ne(x)$ that \mathcal{M} never exits M_i starting at vertex x . Note that since automaton A is non-blocking, the probability of never exiting the current component of $\mathcal{M} \otimes A$ starting at (x, s) is the same as $ne(x)$ (the probability of never exiting the current component from vertex x in the RMC \mathcal{M} alone).

To summarise, we first compute reachability probabilities $q_{(u, ex)}$ and probabilities $ne(u)$ for the RMC \mathcal{M} . Then we consider the product $\mathcal{M} \otimes A$ and solve a system of non-linear equations to compute the probabilities of summary transitions $p(x, s \rightarrow ex, t)$. From these data we build the summary chain, identify accepting SCCs and compute the resulting probabilities in the same way as in [YE05]. All these steps can be expressed as a formula and its truth value can be decided using existential theory of the reals in PSPACE. \square

7.1. Markov chains. We are now ready to prove a new bound for the model checking problem on our most basic probabilistic system, Markov chains. Courcoubetis and Yannakakis [CY95] showed that one can determine if an LTL formula holds with non-zero probability in a Markov chain in PSPACE. This gives a PSPACE upper bound for $TL[\diamond, \diamond]$ and an EXPSpace upper bound for FO^2 . We will show how to get better bounds, even in the quantitative case, using the logic-to-automata translations.

Proposition 7.3. *Model checking $TL[\diamond, \diamond]$ or $FO^2[<]$ on Markov chains is in $\#P$.*

Proof. Let φ be a $TL[\diamond, \diamond]$ or $FO^2[<]$ formula and \mathcal{M} a Markov chain. Using Theorem 4.1 in case of $TL[\diamond, \diamond]$ and Theorem 4.6 in case of $FO^2[<]$, we have that for formula φ there is a family $\{A_i\}$ comprising at most $2^{poly(|\varphi|, |\Sigma|)}$ unambiguous generalised Büchi automata, whose languages partition $\{w \in \Sigma^\omega : w \models \varphi\}$. Moreover, each A_i has at most $|\varphi| \cdot |\Sigma|$ states and can be generated in polynomial time from φ and index i . By Lemma 7.1 we can further compute the probability p_i of \mathcal{M} satisfying A_i in polynomial time in the sizes of \mathcal{M} and A_i . Since each p_i is computable in polynomial time we can determine $\sum_i p_i$ in $\#P$. \square

Proposition 7.4. *The threshold problem for model checking FO^2 on Markov chains is in PEXP.*

Proof. The result follows by the same argument as in Proposition 7.3, as we are essentially in the same situation, but now by Theorem 4.3 we have a collection of doubly-exponentially many automata, each of exponential size. \square

7.2. Hierarchical and Recursive Markov chains. Similarly, we get the following results for recursive Markov chains (and in particular for hierarchical Markov chains):

Proposition 7.5. *The probability of a $\text{TL}[\diamond, \diamond]$ or $\text{FO}^2[<]$ formula holding on a recursive Markov chain can be computed in PSPACE.*

Proof. By Theorem 4.1 in case of $\text{TL}[\diamond, \diamond]$ and by Theorem 4.6 in case of $\text{FO}^2[<]$, we can convert a formula φ into an equivalent disjoint union of exponentially many unambiguous automata of polynomial size (in $|\varphi|$ and $|\Sigma|$) and the RMC. Using polynomial space we can generate each automaton, calculate the probability that the RMC generates an accepting trajectory by Proposition 7.2, and sum these probabilities for each automaton. \square

Corollary 7.6. *The probability of a $\text{TL}[\diamond, \diamond]$ or $\text{FO}^2[<]$ formula holding on a hierarchical Markov chain can be computed in PSPACE.*

Proposition 7.7. *The probability of an FO^2 formula holding on an RMC can be computed in EXPSPACE.*

Proof. The result follows by the same argument as in Proposition 7.5, but now by Theorem 4.3 we have family of doubly exponentially many automata each of exponential size, with a non-deterministic exponential time algorithm for building each automaton. Therefore applying Proposition 7.2 we immediately obtain upper bounds for FO^2 . \square

For an ordinary Markov chain, calculating the probability of an LTL formula can be done in PSPACE [Yan10], while we have seen previously that we can calculate the probability of an FO^2 formula in PEXP. One can achieve the same bounds for LTL and FO^2 on hierarchical Markov chains. In each case we expand the HMC into an ordinary Markov chain and then use the model checking algorithm for a Markov chain. This does not impact the complexity, since the space complexity is only polylog in the size of the machine for LTL and the time complexity is only polynomial in the machine size for FO^2 . We thus get:

Proposition 7.8. *The probability of a FO^2 formula holding on a HMC can be computed in PEXP, while for an LTL formula it can be computed in PSPACE.*

7.3. Markov decision processes. Courcoubetis and Yannakakis [CY95] have shown that the maximal probability with which a scheduler can achieve an UTL objective on an MDP can be computed in 2EXPTIME. It follows from results of [ATM03] that even the qualitative problem of determining whether every scheduler achieves probability 1 is 2EXPTIME-hard. Combining the 2EXPTIME upper bound with the exponential translation from FO^2 to UTL [EVW02] yields a 3EXPTIME bound for FO^2 . Below we see that using our FO^2 -to-automaton construction we are able to improve this bound to 2EXPTIME.

We begin with universal formulation of qualitative model checking MDPs. To deal with MDP's, we will make use of determinism in the limit.

Proposition 7.9. *Determining whether for all schedulers a $\text{FO}^2[\langle \rangle]$ -formula φ holds almost surely on a Markov decision process \mathcal{M} is co-NP-complete.*

Proof. The corresponding complement problem asks whether there exists a scheduler σ such that the probability of $\neg\varphi$ is greater than 0. For this problem, there is an NP algorithm, as we now explain. In Courcoubetis and Yannakakis [CY95], there is a polynomial time algorithm for qualitative model checking deterministic Büchi automata on MDPs. As noted there, the algorithm applies to automata that are deterministic in the limit as well. Therefore we can just guess a particular automaton A_i from the family of automata corresponding to $\neg\varphi$, as described in Theorem 4.6. The theorem guarantees that this automaton will be deterministic in the limit.

It is easy to see that the co-NP is tight, even for $\text{TL}[\diamond, \heartsuit]$, since qualitative model checking for MDPs generalises validity for both $\text{TL}[\diamond, \heartsuit]$ formulas, which is co-NP hard. \square

Proposition 7.10. *Determining whether for all schedulers a UTL-formula φ holds almost surely on a Markov decision process \mathcal{M} is in EXPTIME. For FO^2 the problem is in co-NEXPTIME.*

Proof. The result for FO^2 follows along the lines of the proof of Proposition 7.9, but now we guess an automaton A_i of exponential size (using Theorem 4.3).

Similarly, for UTL we can use Theorem 4.1. We still have exponential sized automata A_i , but only exponentially many of them, so we can iterate over all of them, which gives us a single exponential algorithm. \square

Note that here the FO^2 problem is *easier* than the corresponding LTL problem, which is known to be 2EXPTIME-complete.

For the existential case of the qualitative model-checking problem, an upper bound of 2EXPTIME for all of our languages will follow from the quantitative case below. On the other hand the arguments from [ATM03] can be adapted to get a 2EXPTIME lower bound (see Proposition 7.18) even for qualitative model-checking $\text{TL}[\diamond, \heartsuit]$ in the existential case. Hence we have:

Proposition 7.11. *Determining if there is a scheduler that enforces a formula with probability one is 2EXPTIME-complete for each of $\text{TL}[\diamond, \heartsuit]$, UTL, LTL, $\text{FO}^2[\langle \rangle]$ and FO^2 .*

We now turn to the quantitative case. We apply the translation from FO^2 to deterministic parity automata from Subsection 4.3, along with the result that the value of a Markov decision process with parity winning objective can be computed in polynomial time [CH12]. Using Theorem 4.13 we immediately get bounds for FO^2 that match the known bounds for LTL:

Proposition 7.12. *We can compute the maximum probability of an FO^2 formula φ over all schedulers on a Markov decision processes \mathcal{M} in 2EXPTIME.*

7.4. Stochastic two-player games with FO^2 winning condition. We can reduce the qualitative case of stochastic two-player games to the case of ordinary two-player games using the following result of Chatterjee, Jurdzinski and Henzinger:

Proposition 7.13 ([CJH03]). *Every (universal) qualitative simple stochastic parity game with n vertices, m edges and d priorities can be translated to a simple parity game with*

the same set of priorities, with $O(dn)$ vertices and $O(d(m+n))$ edges, and hence it can be solved in time $O(d(m+n)(nd)^{d/2})$.

Now combining the reduction with our results for two-player games, we ascertain the complexity of stochastic two-player games:

Corollary 7.14. *The universal qualitative model checking problem for Stochastic two-player games ($2\frac{1}{2}$ -player game) with FO^2 winning condition is 2EXPTIME-complete.*

Proof. Hardness follows from 2EXPTIME-hardness for two-player games with FO^2 winning conditions. Membership is a consequence of the above reduction and our bounds for two-player games (see Proposition 6.4 and Proposition 7.13). \square

7.5. Lower bounds. We can get corresponding tight lower bounds for most of the probabilistic model checking problems.

Proposition 7.15. *The quantitative model checking problem for a $\text{TL}[\diamond, \diamond]$ formula ψ on a Markov chain \mathcal{M} is #P-hard.*

Proof. The proof is by reduction from #SAT. Let φ be a propositional formula over literals a_1, a_2, \dots, a_n . We construct a Markov chain \mathcal{M} such that each trajectory generated by \mathcal{M} corresponds to an assignment of truth values to literals a_1, \dots, a_n , with each of the 2^n possible truth assignments arising with equal probability. We also construct a $\text{TL}[\diamond, \diamond]$ formula ψ such that only trajectories of \mathcal{M} that encode satisfying valuations contribute to the probability $P_{\mathcal{M}}(L(\psi))$. Therefore the number of satisfying assignments of the original propositional formula φ is $2^n P_{\mathcal{M}}(L(\psi))$.

See Figure 4 for a depiction of the Markov chain \mathcal{M} in case $n = 3$. All probabilities equal $1/2$, except those on transitions leading to the final vertex f . A path going through vertex a_i corresponds to assigning true to the literal a_i and a path through a'_i to an assignment of false. We construct the $\text{TL}[\diamond, \diamond]$ formula ψ corresponding to the propositional formula φ by replacing each positive literal a_i in φ with $\diamond a_i$ and each negative literal $\neg a_i$ in φ with $\diamond a'_i$.




Figure 4: Markov chain \mathcal{M} for $n = 3$

Recalling the upper bound from Proposition 7.3, we conclude that the quantitative model checking problem for $\text{TL}[\diamond, \diamond]$ on Markov chains is #P-complete. \square

Proposition 7.16. *The quantitative model checking problem for FO^2 on Markov chains is PEXP-hard.*

Proof. PEXP-hardness is by reduction from the problem of whether a strict majority of computation paths of a given non-deterministic EXPTIME Turing machine T on a given input I are accepting. The Markov chain generates a uniform distribution over strings of the appropriate length, and the formula checks whether a given string encodes an accepting computation of \mathcal{M} . The ability of FO^2 to check validity of such a string has already been exploited in the NEXPTIME-hardness proof for FO^2 satisfiability in [EVW02]. The details of this approach can be found in the proof of Proposition 9.3.

Combining with the upper bound from Proposition 7.4, the quantitative model checking problem for FO^2 on Markov chains is PEXP-complete. \square

Turning to lower bounds for MDPs, note that co-NEXPTIME-hardness for FO^2 is inherited from the lower bound for Markov chains. On the other hand, we can show that the EXPTIME bound for UTL is tight:

Proposition 7.17. *Determining whether for all schedulers a UTL-formula φ holds almost surely on a Markov decision process \mathcal{M} is EXPTIME-hard.*

Proof. The argument is based on the idea of Courcoubetis and Yannakakis for lower bounds in the LTL case. We reduce the acceptance problem for an alternating PSPACE Turing machine to the problem of whether there is a scheduler that enforces that a UTL formula φ holds with positive probability. Thus we reduce to the complement of the problem of interest.

Consider an alternating PSPACE Turing machine T with input I . Without loss of generality we assume that each configuration of T has exactly two successors and that T uses space at most n on an input I of length n . Then we can encode a branch of the computation tree of T as a finite string in which each configuration is represented by a consecutive block of $n+1$ letters: one bit to represent the choice to branch left or right, and n letters to represent the configuration. Let $L_{T(I)}$ be the language of infinite strings, each of which is an infinite concatenation of finite strings that encode accepting computations. It is standard that one can write a UTL formula φ that captures $L_{T(I)}$.

Next we describe the MDP \mathcal{M} . Intuitively the goal of the scheduler is to choose a path through \mathcal{M} so as to generate a word in $L_{T(I)}$. A high-level depiction of \mathcal{M} is given in Figure 5. The boxes *init-conf* and *next-conf* contain gadgets that are used by the scheduler to generate the initial configuration and all successive configurations of T as strings of length n . The number of such strings is exponential in n , but clearly the gadgets can be constructed using only linearly many states. After producing an existential configuration of the Turing Machine, the scheduler sends control to the state *sch*, where it decides whether T should branch left or right. After generating a universal configuration, an honest scheduler sends control to *pro*, the only randomising state in \mathcal{M} , where the branching direction T is selected uniformly at random. When the scheduler has successfully generated an accepting computation it visits *acc*, which is the only accepting state of \mathcal{M} , and the simulation starts over again from the beginning. Only those computations that visit *acc* infinitely often and in which the scheduler behaves honestly satisfy φ .

We claim that there exists a scheduler such that $P_{\mathcal{M}}(L(\varphi)) > 0$ if and only if T accepts its input.

If the Turing Machine T accepts its input, then the scheduler can simply follow the strategy from the alternating computation of T . Regardless of the choice made by the probabilistic opponent, the scheduler can always go to an accepting vertex with probability 1. Therefore even if we repeat the whole simulation, for this scheduler $P_M(L(\varphi)) = 1$, which is greater than 0 as required.

The infinite repetition is important in the second case, when the Turing Machine T rejects its input. If the process ran only once, it could happen that in the probabilistic choice, only one option would lead to a rejecting state, but it would not be chosen if the probabilistic opponent of the scheduler were unlucky. Therefore we repeat this process infinitely many times and thus guarantee that with probability 1 we will reach the rejecting vertex and then stay there forever, i.e. $P_M(L(\varphi))$ will be 0 as required.

Combining with the upper bound from Proposition 7.10, determining whether for all schedulers a UTL-formula holds with probability one on a Markov decision process is EXPTIME-complete. □




Figure 5: Sketch of the Markov decision process \mathcal{M}

The above was a lower bound for checking whether all schedulers enforce the property with probability 1. We now show a tight lower-bound for the existence of a probability one scheduler:

Proposition 7.18. *Given a Markov decision process and a $TL[\diamond, \heartsuit]$ formula, determining whether the formula holds with probability one for some scheduler is 2EXPTIME-hard.*

Proof. The proof is an adaptation of the 2EXPTIME-hardness proof of Alur *et al.* for model checking $TL[\diamond, \heartsuit]$ formulas on two-player games in [ATM03]. The proof there is based on a reduction from the membership problem for an alternating exponential-space Turing machine, where a game graph and a $TL[\diamond, \heartsuit]$ formula are constructed such that the Turing machine accepts the given input if and only if the existential player has a winning strategy in the game.

We can adapt the proof by assigning the existential vertices of the game graph to a scheduler and assigning the universal vertices from the game graph to the probabilistic

player (by setting the uniform outgoing probabilities from these vertices). When the Turing machine accepts its input we are guaranteed that there is a corresponding scheduler that leads to acceptance with the probability 1. On the other hand, if the Turing machine does not accept its input then after some finite number of transitions in the Markov decision process, either the scheduler “cheats” (does not follow the Turing machine transition function or cell numbering) or we get to a rejecting state. In both cases, the probability of acceptance is less than 1. \square

Table 7.5 summarises the known results and the results from this paper (in bold) on probabilistic systems. An asterisk indicates bounds that are not known to be tight. Note that for the more complex verification problems, from strategy synthesis for MDPs onwards, all problems are 2EXP-complete. Intuitively the complexity of the model overwhelms the difference in the respective logics. Similarly, we see that in the stutter-free case the extra succinctness of $\text{FO}^2[\prec]$ comes at “no cost” over $\text{TL}[\diamond, \diamond]$ — at least, for the complexity classes we consider, and where we can establish tight bounds, the respective columns are identical.

	$\text{TL}[\diamond, \diamond]$	UTL	$\text{FO}^2[\prec]$	FO^2	LTL
Markov chain	#P	PSPACE	#P	PEXP	PSPACE
HMC	PSPACE*	PSPACE	PSPACE*	PEXP	PSPACE
RMC	PSPACE*	EXPSpace*	PSPACE*	EXPSpace*	EXPSpace*
MDP (\forall)	co-NP	EXP	co-NP	co-NEXP	2EXP
MDP (\exists)	2EXP	2EXP	2EXP	2EXP	2EXP
MDP (quant)	2EXP	2EXP	2EXP	2EXP	2EXP
$2\frac{1}{2}$ -game (\forall)	2EXP	2EXP	2EXP	2EXP	2EXP

8. MODEL CHECKING $\text{FO}^2[\text{LTL}]$

We now turn to combining FO^2 with automata-based techniques for LTL, examining verification of the hybrid language $\text{FO}^2[\text{LTL}]$. As was done with FO^2 , we first show that we can translate $\text{FO}^2[\text{LTL}]$ into temporal logic with exponential blow-up in the size of the formula, giving a simple upper bound. While for FO^2 the translation was to unary temporal logic, in this case we have a translation to LTL_{Let} .

We can look at every $\text{FO}^2[\text{LTL}]$ formula as being rewritable using let definitions such that every let definition involves either a pure FO^2 formula or a pure LTL operator. We get this form by introducing a let definition for every subformula with one free variable. For example, rewriting the formula $\varphi = ((\exists y (\text{succ}(x, y) \wedge P_1(x))) \mathcal{U} P_0)(x)$ with *let definitions* yields

$$\begin{aligned} \varphi_{\text{Let}} = & \text{Let } R_0(x) \text{ be } P_0(x) \text{ in} \\ & \text{Let } R_1(x) \text{ be } P_1(x) \text{ in} \\ & \text{Let } R_2(x) \text{ be } \exists y (\text{succ}(x, y) \wedge R_1(x)) \text{ in} \\ & (R_2 \mathcal{U} R_0)(x) \end{aligned}$$

Note that although the above uses a combination of FO^2 and LTL, each individual definition is either “pure FO^2 ”, or “pure LTL”, and we can apply the translation of FO^2 to UTL in Lemma 3.10 to each FO^2 definition. This gives the following result:

Lemma 8.1. *Given an $\text{FO}^2[\text{LTL}]$ formula φ , we can convert it to an equivalent LTL_{Let} formula ψ such that $|\psi| = O(2^{|\varphi|^2})$.*

We could then translate the let definitions away for LTL, to get an ordinary LTL formula—thus showing that $\text{FO}^2[\text{LTL}]$ and LTL have the same expressiveness. However, there is no need to perform this second transformation to get a bound on the complexity of model checking. Let definitions do not increase complexity for model checking LTL, since non-deterministic Büchi automata for LTL and LTL_{Let} have the same asymptotic size:

Lemma 8.2. *Given an LTL_{Let} formula φ , there is an unambiguous Büchi automaton A with at most $O(2^{|\varphi|^2})$ states accepting exactly the language $\{w \in \Sigma^\omega : w \models \varphi\}$. Moreover this automaton can be constructed in polynomial time in its size.*

This follows from the fact that the number of subformulas of LTL_{Let} formulas is linear in the formula size (Lemma 2.1) and from the following result of Couvreur et al:

Lemma 8.3 ([CSS03]). *Given an LTL formula φ , there is an unambiguous Büchi automaton A with at most $O(|\Sigma| |\text{sub}(\varphi)| 2^{|\text{sub}(\varphi)|})$ states accepting exactly the language $\{w : w \in \Sigma^\omega \wedge w \models \varphi\}$. Moreover this automaton can be constructed in polynomial time in its size.*

As a corollary of Lemmas 8.1 and 8.2 we see that we can convert from an $\text{FO}^2[\text{LTL}]$ formula to an unambiguous Büchi automaton in doubly exponential time, giving a doubly-exponential bound on the complexity of model-checking. However, just as in the previous section, we show that we can do better by direct analysis than via this translation approach.

We begin by looking at the translation given in Lemma 8.1 from a different perspective. Let us extend the set of atomic propositions \mathcal{P} and alphabet $\Sigma = 2^{\mathcal{P}}$ by adding new atomic propositions \mathcal{R} for every predicate created in that translation. Thus we have an extended alphabet $\Sigma' = 2^{\mathcal{P} \cup \mathcal{R}}$. There is an obvious restriction mapping taking an infinite word w' over Σ' to a word over Σ , simply by ignoring all propositions in \mathcal{R} ; we denote this by $\text{restrict}(w', \Sigma)$.

Lemma 8.4. *Given an $\text{FO}^2[\text{LTL}]$ formula φ alphabet Σ , there is an FO^2 formula φ_F and an LTL formula φ_L over Σ' having the following two properties for all $w \in \Sigma^\omega$: (i) if $w \models \varphi$ then there is a unique extension w' of w such that $w' \models \varphi_L \wedge \varphi_F$; (ii) if $w \not\models \varphi$ then there is no extension to w' such that $w' \models \varphi_L \wedge \varphi_F$. Moreover, $|\varphi_L|, |\varphi_F| = O(|\varphi|^2)$*

Proof. We use the translation in Lemma 8.1, but consider it simply returning the collection of let definitions. Corresponding to each definition is a conjunct stating that R_i holds iff φ_i holds. We now examine the form of this conjunct.

Each φ_i is either a basic two-variable formula or an LTL atomic formula. If φ_i is in LTL then the iff can be expressed again in LTL: $\Box(R_i \leftrightarrow \varphi_i)$. If φ_i is in FO^2 then the iff above can be expressed as $\forall x.(R_i(x) \leftrightarrow \varphi_i(x))$. We can simply let φ_F be the FO^2 conjuncts and φ_L be the LTL conjuncts to obtain the desired conclusion.

The upper bounds for lengths $|\varphi_L|$ and $|\varphi_F|$ follow from the fact that $k \leq |\varphi|$ and $|\varphi_i| \leq |\varphi|$. \square

For the formula from the example at the beginning of this section we get following formulas φ_L and φ_F over Σ' :

$$\begin{aligned} \varphi_L &= (R_2 \mathcal{U} R_0)(x) \wedge \Box(R_0(x) \leftrightarrow P_0(x)) \wedge \\ &\quad \Box(R_1(x) \leftrightarrow P_1(x)) \\ \varphi_F &= \forall x.(R_2(x) \leftrightarrow \exists y.(\text{succ}(x, y) \wedge R_1(x))) \end{aligned}$$

8.1. Combining automata constructions for FO^2 and LTL. Given $\text{FO}^2[\text{LTL}]$ formula φ , we can apply Lemma 8.4 to obtain an equisatisfiable formula $\varphi_L \wedge \varphi_F$, where φ_L is an LTL formula and φ_F is an FO^2 formula over the extended alphabet Σ' . Now we can build a Büchi automaton B_L for φ_L using the construction from Lemma 8.3, as well as a collection of $2^{2^{\text{poly}(|\varphi_F|)}}$ Büchi automata B_{F_i} for φ_F , using Theorem 4.3.

For each i we build a product automaton $A_i = B_L \otimes B_{F_i}$ synchronising on the truth values of the newly introduced atomic propositions R_i . We claim that each product automaton A_i is unambiguous, the languages they accept are disjoint, and their union is exactly $\{w \in \Sigma^\omega : w \models \varphi\}$. This follows from the fact that each word over Σ has only one extension to a word over Σ' for which B_L accepts, along with the fact that the languages accepted by the B_{F_i} are disjoint.

After producing the synchronised cross product, we can restrict the input alphabet back to Σ , because the values of all newly introduced atomic propositions $p_i \in \Sigma' \setminus \Sigma$ are fully determined by the truth values of atomic predicates P_i and the relations defined by φ .

Therefore we get the following theorem:

Theorem 8.5. *$\text{FO}^2[\text{LTL}]$ formula φ , there is a collection of doubly exponentially many (in $|\varphi|$) generalized Büchi automata A_i , each of exponential size in $|\varphi|$, such that the languages they accept are disjoint and the union of these languages is exactly $\{w \in \Sigma^\omega : w \models \varphi\}$. Moreover, each automaton A_i is unambiguous and can be constructed by a non-deterministic Turing machine in polynomial time in its size.*

This translation will now allow us to read off bounds for many $\text{FO}^2[\text{LTL}]$ verification problems.

8.2. Model Checking $\text{FO}^2[\text{LTL}]$. Comparing Theorem 4.3 with Theorem 8.5, we can easily see that automata for FO^2 in isolation and $\text{FO}^2[\text{LTL}]$ have the same asymptotic size. We can therefore use all automata-based bounds on verification results for FO^2 , provided that they rely only on unambiguity of the resulting automata. This allows us to replace FO^2 with $\text{FO}^2[\text{LTL}]$ in the results of the previous sections, giving the following:

Proposition 8.6. *Model checking $\text{FO}^2[\text{LTL}]$ properties on Kripke structures, hierarchical and recursive state machines is in the complexity class NEXP.*

Proposition 8.7. *The threshold problem for model checking $\text{FO}^2[\text{LTL}]$ on both Markov chains and hierarchical Markov chains is in PEXP.*

Proposition 8.8. *The probability of an $\text{FO}^2[\text{LTL}]$ formula holding on a recursive Markov chain can be computed in EXPSPACE.*

Now let us consider model checking Markov decision processes. Recall that in the proof of the corresponding bound for FO^2 , Theorem 7.10, we relied on the fact that the automata are deterministic in the limit. Thus our translation for $\text{FO}^2[\text{LTL}]$ does not give us the same bounds as for FO^2 . And indeed, the corresponding bound for checking whether all schedulers achieve probability 1 is worse for LTL in this case, namely doubly-exponential. We will show that we can achieve the same bound as for LTL.

Proposition 8.9. *Determining whether for all schedulers an $\text{FO}^2[\text{LTL}]$ -formula φ holds on a Markov decision process with probability one is in the complexity class 2EXPTIME.*

Proof. We will decide the corresponding complement problem which asks whether there exists a scheduler σ such that the probability satisfying $\neg\varphi$ is greater than 0. By applying the translation from Theorem 8.5, we get a collection of doubly-exponentially many automata, each of exponential size. We can go through all these automata and check if the probability is greater than 0 for one of them. For each automaton, we make a call to the exponential time algorithm for qualitative model checking Büchi automata on MDPs from Courcoubetis and Yannakakis [CY95]. \square

The following table summarises the results for $\text{FO}^2[\text{LTL}]$ from this paper (in bold) concerning both non-deterministic and probabilistic systems in the context of results for FO^2 and LTL alone. An asterisk indicates bounds that are not known to be tight. The table shows that for the models considered in this paper the complexity of verifying $\text{FO}^2[\text{LTL}]$ is the maximum of the respective complexities of FO^2 and LTL.

	$\text{FO}^2[\text{LTL}]$	FO^2	LTL
Kripke structure	NEXP	NEXP	PSPACE
HSM	NEXP	NEXP	PSPACE
RSM	NEXP	NEXP	EXPTIME
Markov chain	PEXP	PEXP	PSPACE
HMC	PEXP	PEXP	PSPACE
RMC	EXPSPACE*	EXPSPACE*	EXPSPACE*
MDP (\forall)	2EXP	co-NEXP	2EXP

9. THE IMPACT OF LET DEFINITIONS ON MODEL CHECKING

In the process of examining two-variable logics and their extensions, we have utilized results on logics extended with Let definitions. We now return to considering the impact of Let for several temporal logics. First, we note that model checking $\text{TL}[\diamond, \diamond]_{\text{Let}}$, UTL_{Let} and LTL_{Let} properties on both non-deterministic (Kripke structures, HSMs, RSMs) and probabilistic systems (Markov chains, HMCs, RMCs, MDPs (\forall)) has similar computational complexity as for the corresponding logics without let definitions. We get these results by simply substituting let definitions to obtain formulas in the base logic, and then analyze the complexity of model-checking the resulting formulas.

In the case of LTL_{Let} , we have already noted that the size of the automaton for LTL is exponential only in the number of subformulas (see, e.g. Couvreur et. al. [CSS03])—this leads to Lemma 8.2. Similarly, for $\text{TL}[\diamond, \diamond]_{\text{Let}}$ and UTL_{Let} , we get the corresponding automata of the same asymptotic size as for $\text{TL}[\diamond, \diamond]$ and UTL respectively, because their size depends on the number of subformulas and the operator depth and not directly on the size of the formula (see translation in Subsection 4.1).

In the case of FO_{Let}^2 , we can use Lemma 3.10 to translate the formula to UTL_{Let} and then use the result above that the sizes of the automata for UTL and UTL_{Let} formulas of the same length are asymptotically equal. Moreover, since LTL_{Let} and FO_{Let}^2 have unambiguous Büchi automata of equal asymptotic size as for LTL and FO^2 respectively, we can combine them in the same way as in the proof of Theorem 8.5 to get the same complexity upper bounds for model checking $\text{FO}^2[\text{LTL}]_{\text{Let}}$ as for $\text{FO}^2[\text{LTL}]$. Thus we have:

Proposition 9.1. *For LTL, FO^2 , UTL, $\text{TL}[\diamond, \diamond]$, and $\text{FO}^2[\text{LTL}]$, all the upper bounds previously shown hold also in the presence of Let definitions.*

Finally, we will show that, in contrast to the cases above, the complexity of model checking $\text{FO}^2[<]_{\text{Let}}$ is exponentially worse than that of $\text{FO}^2[<]$ on both non-deterministic and probabilistic systems. Thus this is the only logic we have considered where the introduction of let definitions makes a difference in the computational complexity of model checking. The following two theorems show the lower bounds on the complexity of model checking $\text{FO}^2[<]_{\text{Let}}$, which match exactly the upper bounds for FO^2_{Let} (compare with Proposition 6.2).

Proposition 9.2. *The satisfiability of a $\text{FO}^2[<]_{\text{Let}}$ formula under the unary alphabet restriction is NEXP-hard.*

Proof. The proof is by reduction from the halting problem of a non-deterministic EXPTIME Turing machine T on a given input I . Let Γ and Q be respectively the tape alphabet and set of control states of T . We consider infinite strings over alphabet

$$\Sigma := (\{P_0, P_1, \dots, P_{2n-1}\} \times \{\Gamma \cup (\Gamma \times Q)\}) \cup \{\#\}.$$

An infinite word $u \in \Sigma^\omega$ encodes a computation of T as follows. Each configuration is encoded in a block of contiguous letters in u , with successive configurations arranged in successive blocks. Each such block comprises 2^n symbols denoting the contents of each tape cell in the configuration. A symbol encoding a tape cell consists of: a letter from $\Gamma \cup (\Gamma \times Q)$ to denote the contents of the tape cell and whether the read head of the Turing Machine is currently on the cell (and if so, the current control state of T), and a predicate P_i denoting the address of the tape cell and the configuration number. Here we use the power of Let definitions to transform the sequence of $2n$ predicates to values of $2n$ -bit counter (see the proof of Lemma 3.12), which represent the address of configuration and tape cell. Having thus encoded a computation of T in a finite prefix of u we require that the remaining infinite tail of u be the string $\#^\omega$.

We can use short $\text{FO}^2[<]_{\text{Let}}$ formulas to identify the position in the string representing the previous or next position of the tape cell in the same configuration. We can also use such formulas to identify the same position of the tape cell in the previous or next configuration. Thus we can easily check if the tape symbols are consistent with the transition function of T . Finally, we ensure T is in the accepting state in the last configuration. \square

Proposition 9.3. *The decision problem of whether a Markov chain \mathcal{M} satisfies an $\text{FO}^2[<]_{\text{Let}}$ -formula φ with probability greater than $1/2$ is PEXP-hard.*

Proof. The proof is by reduction from the problem of whether a strict majority of computation paths of a given non-deterministic EXPTIME Turing machine T on a given input I are accepting. Without loss of generality we can assume that any non-halting configuration of T has exactly two successors and that all computations of T on input I make exactly 2^n steps, where n is the length of I .

The basic idea, following the proof of NEXPTIME-hardness of satisfiability for $\text{FO}^2[<]_{\text{Let}}$, is to encode computations of T as strings. We can define an $\text{FO}^2[<]_{\text{Let}}$ formula that is satisfied by a word $u \in \Sigma^\omega$ precisely when u encodes a legitimate computation of T on input I according to the encoding scheme used in Proposition 9.2. Indeed, the definition is just as described in the proof of NEXPTIME-hardness for $\text{FO}^2[<]_{\text{Let}}$ satisfiability in Proposition 9.2.

The Markov Chain \mathcal{M} in our reduction is constructed from two copies of a component \mathcal{M}' . The definition of \mathcal{M}' is very simple; it consists of a directed clique augmented with a single sink state. In detail, there is a state s_σ for each letter $\sigma \in \Sigma$; $s_\#$ is a sink that makes a transition to itself with probability 1; the next-state distribution from s_σ , $\sigma \neq \#$, is given by a uniform distribution over all states; finally, the label of state s_σ is σ .

The Markov chain \mathcal{M} consists of two disjoint copies \mathcal{M}_{left} and \mathcal{M}_{right} of \mathcal{M}' that are identical except that their states are distinguished by propositions P_{left} and P_{right} . The initial state of \mathcal{M} is a uniform distribution over all states.

We can partition Σ^ω into three sets N , A and R , respectively comprising those strings that don't encode computations of T on input I , those strings that encode accepting computations, and those strings that encode rejecting computations. Moreover each of these sets is definable in $\text{FO}^2[<]_{\text{Let}}$ by formulas φ_N , φ_A and φ_R respectively.

We define the formula φ by

$$\varphi := ((\forall x P_{left}(x)) \wedge (\varphi_N \vee \varphi_A)) \vee ((\forall x P_{right}(x)) \wedge \varphi_A).$$

To complete the reduction, we claim that $P_{\mathcal{M}}(L(\varphi)) > 1/2$ if and only if a strict majority of the computations of Turing Machine T on input I are accepting. To see this, observe that if \mathcal{M} produces a trajectory from $N \subseteq \Sigma^\omega$ then that trajectory is equally likely to have come from \mathcal{M}_{left} or \mathcal{M}_{right} . Using this we can see that $P_{\mathcal{M}}(L(\varphi))$ is $(P_{\mathcal{M}}(A) + P_{\mathcal{M}}(N))/2 + P_{\mathcal{M}}(A)/2$. Thus $P_{\mathcal{M}}(L(\varphi)) > 1/2$ iff $2P_{\mathcal{M}}(A) > 1 - P_{\mathcal{M}}(N)$. From this we see that $P_{\mathcal{M}}(L(\varphi)) > 1/2$ if and only if $|A| > |R|$, as required. \square

The table below summarises the results for the selected logics. An asterisk indicates bounds that are not known to be tight.

	$\text{TL}[\diamond, \diamond]_{\text{Let}}$	$\text{FO}^2[<]_{\text{Let}}$	FO^2_{Let}	$\text{FO}^2[\text{LTL}]_{\text{Let}}$
Kripke structure	NP	NEXP	NEXP	NEXP
HSM	NP	NEXP	NEXP	NEXP
RSM	NP	NEXP	NEXP	NEXP
Markov chain	$\#P$	PEXP	PEXP	PEXP
HMC	PSPACE*	PEXP	PEXP	PEXP
RMC	PSPACE*	EXSPACE*	EXSPACE*	EXSPACE*
MDP (\forall)	co-NP	co-NEXP	co-NEXP	2EXP

10. CONCLUSIONS AND ONGOING WORK

In this paper we have compared the complexity of verifying properties in the two best-known elementary fragments of monadic first-order logic on words: LTL and FO^2 . We provided several different logic-to-automaton constructions that are useful for verification of FO^2 . One translations allows us to understand the complexity of verifying full FO^2 via analysis of unary temporal logic; a second is useful for the sublanguage of FO^2 with only the linear-ordering; the third is useful for getting deterministic automata, which is needed for obtaining bounds for certain game-related problems. We have shown that these translations put together allow us to understand the complexity of verification and synthesis problems for both non-deterministic and probabilistic models transition systems, including those arising from hierarchical and recursive state machines.

While LTL is more expressive than FO^2 , FO^2 can be exponentially more succinct. We have shown that the effect of these opposing factors on the complexity of model checking depends on the model, e.g., FO^2 has higher complexity on Markov chains while LTL has higher complexity on MDPs. By contrast, in the stutter-free case the extra succinctness of $\text{FO}^2[<]$ comes for free—all verification problems have the same complexity as for $\text{TL}[\diamond, \diamond]$. For the most structured models e.g., two-player games and quantitative verification of MDPs, the complexity of the model dominates any difference in the logics.

We are currently examining the succinctness of Let definitions when added to each of our logics. A number of succinctness results can be found in this work, but we have left open the succinctness of Let in certain situations, e.g., for the logic $\text{FO}^2[\text{LTL}]$. Finally, we are investigating the extension of the techniques introduced here from words to trees.

Acknowledgments: M. Benedikt is supported in part by EPSRC grants EP/G004021/1 and EP/H017690/1. Worrell is supported in part by EPSRC grant EP/G069727/1.

REFERENCES

- [ABE⁺05] R. Alur, M. Benedikt, K. Etessami, P. Godefroid, T. Reps, and M. Yannakakis. Analysis of recursive state machines. *ACM Trans. Program. Lang. Syst.*, 27:786–818, July 2005.
- [ATM03] R. Alur, S. La Torre, and P. Madhusudan. Playing games with boxes and diamonds. In *CONCUR*, pages 127–141, 2003.
- [BEM97] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR*, pages 135–150, 1997.
- [BFT98] H. Buhrman, L. Fortnow, and T. Thierauf. Nonrelativizing separations. In *Int’l Conference on Computational Complexity*, pages 8–12, 1998.
- [BLW11] M. Benedikt, R. Lenhardt, and J. Worrell. Two variable vs. linear temporal logic in model checking and games. In *CONCUR*, pages 497–511, 2011.
- [BLW12] M. Benedikt, R. Lenhardt, and J. Worrell. Verification of two variable logic revisited. In *QEST*, 2012.
- [CH12] K. Chatterjee and T. A. Henzinger. A survey of stochastic ω -regular games. *J. Comput. Syst. Sci.*, pages 394–413, 2012.
- [CJH03] K. Chatterjee, M. Jurdzinski, and T. A. Henzinger. Simple stochastic parity games. In Matthias Baaz and Johann A. Makowsky, editors, *CSL*, volume 2803 of *Lecture Notes in Computer Science*, pages 100–113. Springer, 2003.
- [CSS03] J.-M. Couvreur, N. Saheb, and G. Sutre. An optimal automata approach to LTL model checking of probabilistic systems. In *LPAR*, pages 361–375. 2003.
- [CY95] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *J. ACM*, 42(4):857–907, 1995.
- [EVW02] K. Etessami, M. Y. Vardi, and T. Wilke. First-order logic with two variables and unary temporal logic. *Inf. and Comp.*, 179(2):279–295, 2002.
- [EY05] K. Etessami and M. Yannakakis. Recursive Markov Chains, stochastic grammars, and monotone systems of nonlinear equations. In *STACS*, pages 340–352, 2005.
- [Jur00] M. Jurdzinski. Small progress measures for solving parity games. In *STACS*, pages 290–301, 2000.
- [Kam68] H. W. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, UCLA, 1968.
- [LTP07] S. La Torre and G. Parlato. On the complexity of LTL model-checking of recursive state machines. In *ICALP*, pages 937–948, 2007.
- [PR89] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989.
- [SC82] A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. In *STOC*, pages 159–168, 1982.
- [Sto74] L. J. Stockmeyer. *The Complexity of Decision Problems in Automata Theory and Logic*. PhD thesis, MIT, Cambridge, Massachusetts, USA, 1974.
- [VW86] M. Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *LICS*, pages 332–344, June 1986.

- [Wei11] P. Weis. *Expressiveness and Succinctness of First-Order Logic on Finite Words*. PhD thesis, University of Massachusetts, 2011.
- [WI09] P. Weis and N. Immerman. Structure theorem and strict alternation hierarchy for FO^2 on words. *LMCS*, 5(3), 2009.
- [Wol01] P. Wolper. Constructing automata from temporal logic formulas: A tutorial. In *European Educational Forum: School on Formal Methods and Performance Analysis*, pages 261–277, 2001.
- [Yan10] M. Yannakakis. Personal communication, 2010.
- [YE05] M. Yannakakis and K. Etessami. Checking LTL properties of Recursive Markov Chains. In *QEST*, pages 155–165, 2005.