On the length and depth of temporal formulae distinguishing non-bisimilar transition systems

Valentin Goranko
Department of Philosophy
Stockholm University, Sweden
Email: valentin.goranko@philosophy.su.se

Louwe B. Kuijer
Department of Computer Science
University of Liverpool, UK
Laboratoire Lorrain de Recherche
en Informatique et ses Applications, CNRS, France
Email: lbkuijer@liverpool.ac.uk

***Abstract*—We investigate the minimal length and nesting depth of temporal formulae that distinguish two given non-bisimilar finite pointed transition systems. We show that such formula can always be constructed in length at most exponential in the combined number of states of both transition systems, and give an example with exponential lower bound, for several common temporal languages. We then show that by using renamings of subformulae or explicit assignments the length of the distinguishing formula can always be reduced to one that is bounded above by a cubic polynomial on the combined size of both transition systems. This is also a bound for the size obtained by using DAG representation of formulae. We also prove that the minimal nesting depth for such formula is less than the joint size of the two state spaces and obtain some tight upper bounds.**

***Keywords*-non-bisimilar transition systems; temporal logics; size of distinguishing formula;**

I. Introduction

Modal and temporal languages are suitable for describing properties of transition systems that are invariant under behavioural equivalence, i.e. bisimulations. Furthermore, formulae in sufficiently expressive languages can describe any finite transition system up to bisimulation equivalence, thus distinguishing it from any other non-equivalent transition system; in particular, see [3] for characteristic formulae in CTL and [9] where such formulae are constructed in the EF fragment of CTL. The length of such characteristic formulae, however, typically grows exponentially in the size of the transition system.

In this paper we address the questions of the minimal length and nesting depth of a formula in a given modal or temporal logic that distinguishes between two pointed transition systems, in the sense of being true in one and false in the other. We begin with the basic modal logic ML and then indicate how the results extend to more expressive temporal languages, such as the extension TL with past operators, and the computation tree logics CTL and CTL*. It is somewhat surprising that, despite that most answers are as expected, to our current knowledge they have apparently not been explicitly proved and published yet, so this paper is also intended to fill some gaps in the literature. In any case, the methods we apply are widely known, using well established and explored in the literature links between modal equivalence, bisimulations, bisimulation games, and characteristic formulae; see e.g. [3], [10], [9], [8], [5].

To put the main problem studied here in perspective, we first note a well known fact, that the basic modal logic, with language $\mathcal{L}_{\mathrm{ML}}$, has the so called *small model property* (see e.g. [2], [8], [5], usually phrased as follows:

Every satisfiable formula $\varphi \in \mathcal{L}_{\mathrm{ML}}$ is satisfied in a pointed transition system $(\mathcal{M}, w)$ where $\mathcal{M}$ has a size at most exponential in the length of φ.

For our purpose, we propose a somewhat different, but equivalent, formulation:

Every contingent[1] formula $\varphi \in \mathcal{L}_{\mathrm{ML}}$ distinguishes between two pointed transition systems $(\mathcal{M}_1, w_1)$ and $(\mathcal{M}_2, w_2)$, both of size at most exponential in the length of φ.

We now state a dually analogous *small formula property*[2]:

Every two pointed transition systems $(\mathcal{M}_1, w_1)$ and $(\mathcal{M}_2, w_2)$ that are not modally equivalent are distinguishable by a formula $\varphi \in \mathcal{L}_{\mathrm{ML}}$ with a length that is at most exponential (more precisely, n^n, i.e. $2^{n \log n}$) in the combined number of states n of $\mathcal{M}_1$ and $\mathcal{M}_2$.

In this paper, we show inter alia that, unsurprisingly, the usual modal and temporal logics have the small formula property. Furthermore, we show that the exponential upper bound given by the small formula property is almost tight, viz. there are sequences $(\mathcal{M}_1^n, w_1^n)$ and $(\mathcal{M}_2^n, w_2^n)$ (in fact, even with $\mathcal{M}_1^n = \mathcal{M}_2^n$ for each n) of non-equivalent finite pointed transition systems such that the smallest modal formula that distinguishes between $(\mathcal{M}_1^n, w_1^n)$ and $(\mathcal{M}_2^n, w_2^n)$ is exponentially large with respect to either of their sizes $|\mathcal{M}_1^n|$ and $|\mathcal{M}_2^n|$ (though there is still a logarithmic gap between the exponents in the two bounds). This exponential lower bound persists when the expressiveness of the language is extended with past operators (basic temporal logic), or even to the full computation tree logic CTL*. Further, we show that using renaming of subformulae by fresh propositional variables (an idea probably used first by Scott in the early 1960s for FOL and by Tseitin (1968) for propositional logic;

[1] i.e., neither valid, nor unsatisfiable

[2] Even though it is, arguably, more of a *large* formula property.

see [4] for more details and further references) or explicit assignments (like those used in various logics for program correctness, e.g., 'local assignments' in PDL [11], see also the 'public assignments' in [6]) to such variables, the length of the distinguishing formula can always be reduced to one bounded above by a cubic polynomial on the combined size of both transition systems. This reduction in size is due to the fact that only so many different subformulae occur in a distinguishing formula of minimal length, so this polynomial size is readily attained by using representation of formulae by means of directed acyclic graphs (DAGs) rather than strings of symbols.

Another important parameter of a distinguishing formula is the nesting depth of modal/temporal operators in it. This parameter has a very natural interpretation, at least for the basic modal and temporal languages, viz. it is equal to the minimal number of rounds needed for Spoiler to win the respective bisimulation game starting from the initial configuration defined by the pointed transition systems, by following a winning strategy (that can be extracted from any distinguishing formula). It turns out that the minimal nesting depth is closely related to the number of iterations in the computation of the largest bisimulation between the two pointed transition systems as a greatest fixed point of the respective monotone operator encoding the one-step back-and-forth property. Using that observation, we prove that the minimal nesting depth for such distinguishing formula is less than the total number of states in the two transition systems, and obtain tight upper bounds.

Here we only consider finite pointed transition systems, on which modal equivalence coincides with bisimulation equivalence (aka, bisimilarity), so hereafter we will reason interchangeably about bisimilar or modally equivalent transition systems.

II. Preliminaries

We assume that the reader is familiar with the modal and temporal logics considered here, as well as with Kripke models, called here (interpreted) transition systems, bisimulations between them and the standard semantics of ML, CTL, and CTL* in interpreted transition systems (see e.g. [7], [1] or [5]). Still, we provide some basic preliminaries, mainly for the sake of terminological and notational self-containment.

Let $\mathcal{P}$ be a fixed *finite* set of atomic propositions (sometimes also regarded as propositional variables). We will be denoting the standard modal operators $\Diamond$ and $\Box$ by EX and AX instead, to represent the basic modal language $\mathcal{L}_{\mathrm{ML}}$ as a fragment of CTL. The formulae of $\mathcal{L}_{\mathrm{ML}}$ over $\mathcal{P}$ are defined as usual:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid (\varphi \wedge \varphi) \mid \mathsf{EX}\varphi \mid \mathsf{AX}\varphi$$

where $p \in \mathcal{P}$. We assume more primitive operators than necessary, for convenience of computing lengths of formulae and use $\bigwedge$ and $\bigvee$ in the usual way as abbreviations.

We also consider the following extensions of $\mathcal{L}_{\mathrm{ML}}$:

- the basic tense language $\mathcal{L}_{TL}$, obtained by adding past operators, here denoted as inverse modalities EY and AY.
- the computation tree logic CTL, obtained from $\mathcal{L}_{\mathrm{ML}}$ by adding the operators EG, AG, EU, AU, while regarding EF and AF as definable in terms of EU and AU.
- the full computation tree logic CTL*, obtained from $\mathcal{L}_{\mathrm{ML}}$ by adding the temporal operators G, U, and the path quantifiers E and A, while regarding F as definable.

For any $\varphi \in$ CTL, the *length* $|\varphi|$, being the number of primitive symbols occurring in it, is explicitly defined as follows:
$|p| = 1$;
$|O\varphi| = |\varphi|+1$ for $O \in \{\neg, \mathsf{EX}, \mathsf{AX}, \mathsf{EY}, \mathsf{AY}, \mathsf{EF}, \mathsf{EG}, \mathsf{AF}, \mathsf{AG}\}$;
$|(\varphi_1 O \varphi_2)| = |\varphi_1| + |\varphi_2| + 3$ for each $O \in \{\wedge, \vee, \mathsf{EU}, \mathsf{AU}\}$.

A small adjustment is needed for formulae of CTL*, which we will not address here. Note that we define the length of a formula literally, as the number of occurring symbols. We could have alternatively defined the size of a formula by the size of its syntax tree, but doing so would not have made an essential difference in the results. In section V-C, however, we discuss how the results change when a more succinct, DAG-based representation is adopted.

Respectively, the *nesting depth* $\mathsf{nd}(\varphi)$ is defined for any $\varphi \in$ CTL, as follows:
$\mathsf{nd}(p) = 0$; $\mathsf{nd}(\neg\varphi) = \mathsf{nd}(\varphi)$,
$\mathsf{nd}(\varphi_1 \wedge \varphi_2) = \mathsf{nd}(\varphi_1 \vee \varphi_2) = \max(\mathsf{nd}(\varphi_1), \mathsf{nd}(\varphi_2))$;
$\mathsf{nd}(O\varphi) = \mathsf{nd}(\varphi) + 1$
for $O \in \{\mathsf{EX}, \mathsf{AX}, \mathsf{EY}, \mathsf{AY}, \mathsf{EF}, \mathsf{EG}, \mathsf{AF}, \mathsf{AG}\}$;
$\mathsf{nd}(\mathsf{E}(\varphi_1 \mathsf{U} \varphi_2)) = \mathsf{nd}(\mathsf{A}(\varphi_1 \mathsf{U} \varphi_2)) = \max(\mathsf{nd}(\varphi_1), \mathsf{nd}(\varphi_2)) + 1$.

An *interpreted transition system* (aka *Kripke model*) is a triple $\mathcal{M} = (W, R, V)$, where W is a set of states, $R \subseteq W \times W$ is a transition relation[3], and $V : \mathcal{P} \to 2^W$ is a valuation function. Its *size*[4] $|\mathcal{M}|$ is defined as

$$|\mathcal{M}| = \mathrm{card}(W) + \mathrm{card}(R) + \sum_{p \in \mathcal{P}} |V(p)|,$$

where $\mathrm{card}(X)$ is the number of elements in the set X.

Here, we omit 'interpreted' and simply write 'transition system'. A transition system is *finite* if its state space (and therefore, its size) is finite.

Remark. Our definition of transition systems and measure of their size use *valuation functions* (from atomic propositions to sets of states), while often these are defined by means of *labelling functions* (from states to sets of atomic propositions). However, since we only consider here finite sets of atomic propositions and finite transition systems, both measures yield the same sizes.

A *pointed transitions system* is a pair $(\mathcal{M}, w)$ with $\mathcal{M}$ a transition system and w a state in it.

All logics considered here have a well-known standard semantics in transition systems. For relevant background, see e.g. [7], [2], [8], [1], or [5].

[3]We consider here only systems with one transition relation, but the results generalise easily to any labelled transitions systems

[4]We are using the same notation for length of a formula and size of a model, but that should not cause any confusion.

Let $\mathcal{M} = \langle W, R, V\rangle$, $\mathcal{M}' = \langle W', R', V'\rangle$ be transitions systems and let $w \in W, w' \in W'$. We say that w and w' are *propositionally equivalent*, denoted $w \simeq w'$, if $w \in V(p)$ iff $w' \in V'(p)$ for each atomic proposition p.

The property of a relation $\beta \subseteq W \times W'$ to be a *k-bisimulation between the pointed transitions systems* $(\mathcal{M}, w)$ *and* $(\mathcal{M}', w')$, denoted $(\mathcal{M}, w) \overset{\beta}{\rightleftarrows}_k (\mathcal{M}', w')$, is defined inductively on $k \in \mathbb{N}$ as follows:

$(\mathbf{B_0})$ $(\mathcal{M}, w) \overset{\beta}{\rightleftarrows}_0 (\mathcal{M}', w')$ iff $w \beta w'$ and $w \simeq w'$.

$(\mathbf{B_{k+1}})$ $(\mathcal{M}, w) \overset{\beta}{\rightleftarrows}_{k+1} (\mathcal{M}', w')$ iff $(\mathcal{M}, w) \overset{\beta}{\rightleftarrows}_0 (\mathcal{M}', w')$ and the following conditions hold:

Forth: if wRu for some $u \in W$, then there is $u' \in W'$ such that $w'R'u'$ and $(\mathcal{M}, u) \overset{\beta}{\rightleftarrows}_k (\mathcal{M}', u')$.

Back: If $w'R'u'$ for some $u' \in W'$ then there is $u \in W$ such that wRu and $(\mathcal{M}, u) \overset{\beta}{\rightleftarrows}_k (\mathcal{M}', u')$.

Clearly, every k-bisimulation is also an m-bisimulation for every $m < k$. We say that $(\mathcal{M}, w)$ and $(\mathcal{M}', w')$ are *k-bisimular*, or *k-bisimulation equivalent*, denoted $(\mathcal{M}, w) \rightleftarrows_k (\mathcal{M}', w')$, if there is a k-bisimulation β between them. When β is a k-bisimulation for every $k \in \mathbb{N}$, we call it a *finite bisimulation* and say that $(\mathcal{M}, w)$ and $(\mathcal{M}', w')$ are *finitely bisimilar*. Since we only consider finite transitions systems in this paper, and it is well known (from Hennessy-Milner's theorem, see e.g. [8] and [5]), or [1]) that on them finite bisimulation coincides with (unbounded) bisimulation, hereafter we will omit 'finite' and will simply talk about bisimilar (resp. non-bisimilar) transitions systems. We will denote the claim that $(\mathcal{M}, w)$ and $(\mathcal{M}', w')$ are (finitely) bisimilar by $(\mathcal{M}, w) \rightleftarrows (\mathcal{M}', w')$.

Bisimilarity between pointed transitions systems can be characterised in terms of existence of winning strategy for the proponent (Duplicator) in the respective *bisimulation games* between them, defined as follows. The bisimulation game on transitions systems $\mathcal{M}_1 = (W_1, R_1, V_1)$ and $\mathcal{M}_2 = (W_2, R_2, V_2)$ is played by two players **I** (*Spoiler*) and **II** (*Duplicator*), with two tokens, one in $\mathcal{M}_1$ and one in $\mathcal{M}_2$, to mark the 'current state' in each structure. A *configuration* in the game is a pair of pointed transitions systems $(\mathcal{M}_1, s_1; \mathcal{M}_2, s_2)$, where the distinguished points are the current positions of the two tokens. The game starts from *initial configuration* and is played in *rounds*. In each round Spoiler selects a token and moves it forward along a transition in the respective structure, to a successor state. Then Duplicator responds by similarly moving forward the token in the other structure along a transition with the same label. During the game Duplicator loses if she cannot respond correctly to the move of Spoiler, or if the two token positions in the resulting new configuration do not match on some atomic proposition. On the other hand, Spoiler loses during the game if he cannot make a move in the current round because both tokens are in states without successors.

The bisimulation game can be played for a pre-determined number of rounds, or indefinitely. The *n-round bisimulation game* terminates after n rounds, or earlier if either player loses during one of these rounds. If the n-th round is completed without violating the atom equivalence in any configuration, Duplicator wins the game. Respectively, if Duplicator can play the *unbounded bisimulation game* forever, without loosing at any round, she wins the game.

Duplicator has a *winning strategy* in a given bisimulation game if she has responses to any challenges of Spoiler that guarantee her to win the game. A winning strategy of Spoiler is defined likewise.

The following claims relate these games and bisimulations.

1) Duplicator has a winning strategy in the n-round bisimulation game with initial configuration $(\mathcal{M}_1, s_1; \mathcal{M}_2, s_2)$ if and only if $(\mathcal{M}_1, s_1) \rightleftarrows_n (\mathcal{M}_2, s_2)$.
2) Duplicator has a winning strategy in the unbounded bisimulation game with initial configuration $(\mathcal{M}_1, s_1; \mathcal{M}_2, s_2)$ if and only if $(\mathcal{M}_1, s_1) \rightleftarrows (\mathcal{M}_2, s_2)$.

Finally, bisimulations are closely related to logical equivalence in the modal and temporal logics mentioned above. First, the truth of every CTL*-formula is invariant with respect to (finite) bisimulations. More precisely, the truth of every ML-formula of modal depth at most k is invariant with respect to k-bisimulations. Furthermore, two finite pointed transition systems are k-bisimilar if an only if they satisfy the same ML-formulae of modal depth at most k; hence, they are bisimilar if an only if they satisfy the same ML-formulae.

For further background on the relationships between bisimulations, games, and modal equivalence, see e.g. [3], [10], [9], [8] and [5].

III. The small formula property for modal logic

Theorem 1. *For every pair of non-bisimilar finite pointed transition systems $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$, where $\mathcal{M}_1 = (W_1, R_1, V_1)$ and $\mathcal{M}_2 = (W_2, R_2, V_2)$, there is a modal formula φ such that $\mathcal{M}_1, u_1 \models \varphi$, $\mathcal{M}_2, u_2 \not\models \varphi$ and $|\varphi| < n^n$, where $n = \mathrm{card}(W_1) + \mathrm{card}(W_2)$.*

Proof: Let $G_i \subseteq W_1 \times W_2$, for $0 \leq i \leq n$, be the set of pairs (w_1, w_2) that are distinguishable by a modal formula of depth i, but not by a formula of lower depth. We show in Section VI that $n = \mathrm{card}(W_1) + \mathrm{card}(W_2)$ is large enough to ensure that $G_n = \varnothing$, i.e. every non-bisimilar pair (w_1, w_2) in $W_1 \times W_2$ is distinguishable by a modal formula of depth $i < n$. Let m be the largest index such that $G_m \neq \varnothing$. Thus, $m < n$.

For any $(w_1, w_2) \in G_i$, let $\varphi_{(w_1,w_2)}$ be a minimal length formula of depth i that holds on w_1 but not on w_2. Now, for $0 \leq i \leq m$, let $S_i := \max_{(w_1,w_2) \in G_j \wedge j \leq i} |\varphi_{(w_1,w_2)}|$.

If $(w_1, w_2) \in G_0$, then there is a propositional variable, or its negation, that holds on w_1 but not on w_2. So, $S_0 \leq 2$. Now, consider any $(w_1, w_2) \in G_i$ with $i > 0$. By assumption, w_1 and w_2 are distinguishable by a formula of depth i. This implies that either

(a) there is a successor w_1' of w_1 that is distinguishable from every successor w_2' of w_2 by a formula of depth at most $i-1$, or

(b) there is a successor w_2' of w_2 that is distinguishable from every successor w_1' of w_1 by a formula of depth at most $i-1$.

In the first case, the formula $\mathsf{EX} \bigwedge_{w_2 R_2 w_2'} \varphi_{(w_1', w_2')}$ distinguishes between w_1 and w_2; in the second case $\mathsf{AX} \bigvee_{w_1 R_1 w_1'} \varphi_{(w_1' w_2')}$ does. Every such successor pair is in G_j for some $j < i$, hence $|\varphi_{(w_1', w_2')}| \leq S_{i-1}$ and each of w_1, w_2 has less than n successors, so in either case we have[5] $|\varphi_{(w_1,w_2)}| < n(S_{i-1}+3)$. Thus,

$$S_i < n(S_{i-1}+3).$$

Putting $a_i = S_i + \frac{3n}{n-1}$ we obtain $a_i < n a_{i-1}$. Therefore, $S_i < a_i < n^i a_0 \leq n^i(2+\frac{3n}{n-1})$. Assuming $n > 5$, we eventually get $S_i < 6n^i$, hence:

$$|\varphi| \leq S_m < 6n^m \leq n^n.$$

For $n \leq 5$ the same upper bound is easily verified directly. ∎

The upper bound stated in Theorem 1 seems rather crude and we conjecture that a more refined calculation can produce an upper bound for $|\varphi|$ of $2^{O(n)}$.

IV. Shortest distinguishing formulae of exponential length

Now, we show that there are cases of non-bisimilar transition systems where the smallest distinguishing formulae are of length exponential in the size of each of the transition systems (coinciding in our example). We first prove that for the case of ML and then adapt the argument for the extensions.

Theorem 2. *The sequence $\{\mathcal{M}_k \mid k \in \mathbb{N}\}$ of finite transition systems defined recursively in Figures 1 and 2 is such that, for all $k \in \mathbb{N}$:*

1) *the pointed transition systems $(\mathcal{M}_k, w_k)$ and $(\mathcal{M}_k, v_k)$ are not bisimilar,*
2) *every formula $\varphi \in \mathcal{L}_{\mathrm{ML}}$ that distinguishes between $(\mathcal{M}_k, w_k)$ and $(\mathcal{M}_k, v_k)$ has a length exponential in the size of $\mathcal{M}_k$. More precisely, the length of every such formula φ satisfies the following lower bound:*

$$|\varphi| \geq 9 \cdot 2^{\frac{|\mathcal{M}_k|-5}{23}} - 8.$$

Proof: For $k \in \mathbb{N}$, let $\mathcal{M}_k$ be as shown in Figure 2. Note that for every $i \leq \min(k_1, k_2)$ it holds that $(\mathcal{M}_{k_1}, w_i)$ is bisimilar to $(\mathcal{M}_{k_2}, w_i)$ (since the generated submodels are the same), and likewise for $v_i, s_i, t_i, u_i, x_i, y_i$ and z_i. By bisimulation invariance, every state that occurs in both $\mathcal{M}_{k_1}$ and $\mathcal{M}_{k_2}$ satisfies exactly the same formulae in both transition systems. We therefore omit mention of the transition systems, and say simply that a formula is true at a given state.

We prove by induction on $k \in \mathbb{N}$ that every formula $\varphi \in \mathcal{L}_{\mathrm{ML}}$ that distinguishes between w_k and v_k is of length at least $a_k = 9 \cdot 2^k - 8$, and that there is at least one formula φ_k of that length which is true at w_k and false at v_k.

[5]The added 3 accounts for the number of conjunction/disjunction and parentheses symbols.

Fig. 1: The transition system $\mathcal{M}_0$.

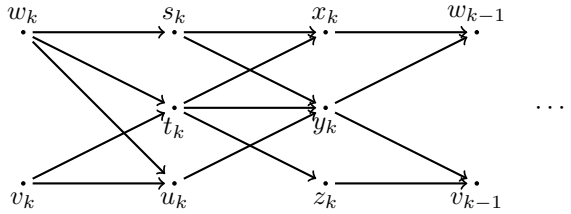

Fig. 2: The transition systems $\mathcal{M}_k$, for $k \in \mathbb{N}_{>0}$.

As base case, suppose $k = 0$. Every formula is of length at least $1 = 9 \cdot 2^0 - 8$, and such formula is $\varphi_0 = p$.

Suppose therefore, as induction hypothesis, that the claim holds for $k-1$, i.e. there is a formula φ_{k-1} of length $9 \cdot 2^{k-1} - 8$ that distinguishes between w_{k-1} and v_{k-1}, and this φ_{k-1} is length-minimal. It is then straightforward to verify that the formula

$$\varphi_k = \mathsf{EX}(\mathsf{AXEX}\varphi_{k-1} \wedge \mathsf{EXAX}\varphi_{k-1})$$

is true at w_k and false at v_k, and its length is $2(9 \cdot 2^{k-1} - 8) + 8 = 9 \cdot 2^k - 8$.

Now, we will show the length-minimality of φ_k.

First, note that the only way to distinguish between any two states in $\{x_k, y_k, z_k\}$ is by which of w_{k-1} and v_{k-1} are accessible from each of them. A distinguishing formula must therefore contain a subformula of the form $\mathsf{AX}\chi$ or $\mathsf{EX}\chi$ where χ distinguishes between w_{k-1} and v_{k-1}. By the induction hypothesis, it follows that any formula ξ that distinguishes between any two of $\{x_k, y_k, z_k\}$ is of length at least $(9 \cdot 2^{k-1} - 8) + 1 = 9 \cdot 2^{k-1} - 7$.

Now, let φ be any minimal length formula such that $w_k \models \varphi$ and $v_k \not\models \varphi$. If a Boolean combination of formulae distinguishes between two states, then so does at least one of the component formulae. It therefore follows from the minimality of φ that the main connective of φ is either AX or EX. Furthermore, every successor of v_k is also a successor of w_k, so the main connective of φ cannot be AX. So $\varphi = \mathsf{EX}\psi$, where $s_k \models \psi$, $t_k \not\models \psi$ and $u_k \not\models \psi$. The formula ψ is a Boolean combination of formulae ξ_j, where each ξ_j has AX or EX as main connective. Since we are after the lower bound of length, and because $|\neg\xi_j| = |\xi_j| + 1$, we can assume without loss of generality that all ξ_j occur positively in ψ.

Because ψ holds in s_k but not in t_k or u_k and every ξ_j occurs positively, there must be some ξ_1 that holds in s_k but not in t_k, and some ξ_2 that holds in s_k but not in u_k. Note that the set of successors of u_k is a proper subset of the set

of successors of s_k, which is a proper subset of the set of successors of t_k. Therefore, ξ_1 must be a AX-formula, while ξ_2 must be a EX-formula. Let $\xi_1 = \mathsf{AX}\zeta_1$ and $\xi_2 = \mathsf{EX}\zeta_2$. Furthermore, both ζ_1 and ζ_2 distinguish between at least two states from the set $\{x_k, y_k, z_k\}$. As shown above, this implies that ζ_1 and ζ_2 are both of length at least $9{\cdot}2^{k-1}{-}7$. This means that the length of ψ is at least the length of $(\mathsf{AX}\zeta_1 \wedge \mathsf{EX}\zeta_2)$, i.e. $2(9 \cdot 2^{k-1} - 7) + 5 = 9 \cdot 2^k - 9$. Therefore, the formula $\varphi = \mathsf{EX}\psi$ is of length at least $9 \cdot 2^k - 8$. This completes the induction.

Finally, note that the transition system $\mathcal{M}_0$ has size 5, and $|\mathcal{M}_k| = |\mathcal{M}_{k-1}| + 23$, hence $|\mathcal{M}_k| = 23k + 5$, whence the claim of the theorem follows easily. ∎

The proof of Theorem 2 can be easily adapted to prove exponential lower bounds for TL, CTL and CTL*.

First, let us consider the case of TL. The only way to distinguish between any two states in the same "column" of $\mathcal{M}_k$ (i.e. between v_i and w_i, between any two of $\{s_i, t_i, u_i\}$ or between any two of $\{x_i, y_i, z_i\}$) is by in which ways the unique p state can be reached. This p state is in the future, so the shortest formula distinguishing between two states in the same column does not contain any EY or AY operators. It follows that the shortest TL formula distinguishing between v_k and w_k is an ML formula and therefore of length exponentially bounded below as in Theorem 1.

Now, consider CTL. Take any path σ from either w_k or v_k. Then there is a path σ' from the other state that differs from σ only in the first two states, i.e. $\sigma(i) = \sigma'(i)$ for all $i > 1$. So, if any CTL state formula of the type $\mathsf{EG}\varphi$, $\mathsf{AF}\varphi$, $\mathsf{E}(\varphi\mathsf{U}\psi)$ or $\mathsf{A}(\varphi\mathsf{U}\psi)$ distinguishes between σ and σ', it must do so based on the first two states. But then either φ or ψ or $\mathsf{EX}\psi$ distinguishes between these states, too. Thus, the extra operators that CTL has over modal logic do not make it easier to distinguish between the worlds.

Essentially the same argument works for state formulae of the full CTL*, too, by considering a few more cases.

The lower bound for the distinguishing formulae obtained in Theorem 2 is most likely not exact in terms of the exponent and coefficients, but any further improvement of these would not be substantial. Therefore, there is a certain gap between the exponents in this lower bound and the upper bound established in Theorem 1 (even though the latter uses the combined number of states, whereas the former refers to the combined size), which we leave open.

V. Distinguishing formulae of polynomial length

Now, we will show that there are simple variations of the framework considered so far that enable constructing distinguishing formulae of polynomial length, even in ML. We discuss two such variations. Both variations use the fact that while the smallest distinguishing formula of ML (or any of the extensions we consider) may be of exponential length, it contains only polynomially many non-equivalent subformulae. By introducing abbreviations for these, we can obtain a polynomial length description of the distinguishing formula.

Before introducing these variations, we need a few more preliminaries. We add *explicit assignments* to modal logic, producing the language $\mathcal{L}_{\mathrm{ML+A}}$. The explicit assignment operator is denoted as $[p := \varphi]$, where $p \in \mathcal{P}$ and $\varphi \in \mathcal{L}_{\mathrm{ML+A}}$. The length of the formula $[p := \varphi_1]\,\varphi_2$ is defined by

$$|[p := \varphi_1]\,\varphi_2| = |\varphi_1| + |\varphi_2| + 1,$$

and the semantics is given as follows:

$$\mathcal{M}, w \models [p := \varphi]\psi \text{ iff } \mathcal{M}[p := \varphi], w \models \psi,$$

where $\mathcal{M}[p := \varphi] = (W, R, V[p := \varphi])$ and
$V[p := \varphi](p) := \{w \in W \mid \mathcal{M}, w \models \varphi\}$, and
$V[p := \varphi](q) := V(q)$ for all $q \in \mathcal{P} \setminus \{p\}$.
Thus, $[p := \varphi]$ assigns to p the extension $\|\varphi\|_{\mathcal{M}}$ of φ in $\mathcal{M}$.

A. Building distinguishing formulae of polynomial length by means of renaming

The first variation is based on the idea of adding fresh propositional variables to the language and the transition systems, and using them to rename the distinguishing subformulae on the fly.

Consider a temporal formula φ of modal depth m, and a subformula ψ. Take a fresh (not occurring in φ) variable p_ψ.

Now, for each $k = 0, 1, \ldots$ let

$$\Gamma^k(\psi, p_\psi) := (\psi \leftrightarrow p_\psi) \wedge \mathsf{AX}(\psi \leftrightarrow p_\psi) \wedge \ldots \wedge \mathsf{AX}^k(\psi \leftrightarrow p_\psi).$$

Further, let $\varphi(p_\psi/\psi)$ be the result of the uniform substitution of all occurrences of ψ in φ by p_ψ. Now, we define the formula

$$\varphi[p_\psi \Leftarrow \psi] := \Gamma^m(\psi, p_\psi) \wedge \varphi(p_\psi/\psi).$$

Proposition 3. *For every transition system $\mathcal{M} = (W, R, V)$, $w \in W$, formula φ of modal depth m not containing the variable p_ψ, and a subformula ψ, the following are equivalent:*

1) $\mathcal{M}, w \models \varphi$.
2) $\mathcal{M}[p_\psi := \psi], w \models \varphi[p_\psi \Leftarrow \psi]$.

The proof is by straightforward induction on φ. Now, using renaming as above on the fly, we can reduce the length of distinguishing formulae down to cubic in the joint size of the two transition systems, as follows. Consider the procedure described in the proof of Theorem 1, applied to two non-bisimilar pointed transition systems $(\mathcal{M}_1, v_1)$ and $(\mathcal{M}_2, v_2)$.

At every step of the construction when a new distinguishing formula $\varphi_{(w_1,w_2)}$ is obtained which is not a variable itself, introduce a fresh variable $p_{\varphi_{(w_1,w_2)}}$, not in the language of the current transition systems $\widehat{\mathcal{M}}_1$ and $\widehat{\mathcal{M}}_2$, and expand these to $\widehat{\mathcal{M}}_1[p_{\varphi_{(w_1,w_2)}} := \varphi_{(w_1,w_2)}]$ and $\widehat{\mathcal{M}}_2[p_{\varphi_{(w_1,w_2)}} := \varphi_{(w_1,w_2)}]$ respectively. Thereafter, wherever $\varphi_{(w_1,w_2)}$ is used further, replace it by $p_{\varphi_{(w_1,w_2)}}$. Eventually, take as a distinguishing formula for the resulting pointed transition systems $\widehat{\mathcal{M}}_1, v_1$ and $\widehat{\mathcal{M}}_2, v_2$:

$$\widehat{\varphi}_{(v_1,v_2)} := p_{\varphi_{(v_1,v_2)}} \wedge \bigwedge \Gamma^m(\varphi_{(w_1,w_2)}, p_{\varphi_{(w_1,w_2)}})$$

where m is the modal depth of $\varphi_{(v_1,v_2)}$ (as noted earlier, $m < \max(|\mathcal{M}_1|, |\mathcal{M}_2|)$) and the conjunction is over all formulae $\varphi_{(w_1,w_2)}$ generated during the construction of $\varphi_{(v_1,v_2)}$.

We leave out the easy details of proving that the formula $\widehat{\varphi}_{(v_1,v_2)}$ is, indeed, distinguishing for $\widehat{\mathcal{M}}_1, v_1$ and $\widehat{\mathcal{M}}_2, v_2$, and thereby encoding the distinction between the original pointed transition systems. Note that, the length of $\widehat{\varphi}_{(v_1,v_2)}$ is roughly bounded above by $O(m^3)$.

For example, the procedure outlined above produces the following distinguishing formulae for the pointed transition systems $(\mathcal{M}_k, w_k)$ and $(\mathcal{M}_k, v_k)$ from Theorem 2 as follows:
$\widehat{\varphi}_{(w_0,v_0)} = \varphi_{(w_0,v_0)} = p$, $\varphi_{(x_1,z_1)} = \mathsf{EX}p$,
$\varphi_{(x_1,y_1)} = \mathsf{AX}p$, $\varphi_{(s_1,t_1)} = \mathsf{AX}p_{\mathsf{EX}p}$,
$\varphi_{(s_1,u_1)} = \mathsf{EX}p_{\mathsf{AX}p}$, $\varphi_{(w_1,v_1)} = \mathsf{EX}(p_{\mathsf{AX}p_{\mathsf{EX}p}} \wedge p_{\mathsf{EX}p_{\mathsf{AX}p}})$.
Now,
$\widehat{\varphi}_{(w_1,v_1)} = p_{\varphi_{(w_1,v_1)}} \wedge \Gamma^3(\mathsf{EX}p, p_{\mathsf{EX}p}) \wedge \Gamma^3(\mathsf{AX}p, p_{\mathsf{AX}p}) \wedge \Gamma^3(\mathsf{AX}p_{\mathsf{EX}p}, p_{\mathsf{AX}p_{\mathsf{EX}p}}) \wedge \Gamma^3(\mathsf{EX}p_{\mathsf{AX}p}, p_{\mathsf{EX}p_{\mathsf{AX}p}}) \wedge \Gamma^3(\mathsf{EX}(p_{\mathsf{AX}p_{\mathsf{EX}p}} \wedge p_{\mathsf{EX}p_{\mathsf{AX}p}}), p_{\varphi_{(w_1,v_1)}})$.

Indeed, the formula above is much longer than φ_1 defined earlier, but its length grows only polynomially fast.

B. Building distinguishing formulae of polynomial length by means of explicit assignments

The alternative approach is to use explicit assignment operators in order to declare within the formula the required renamings. As above, we define a new variable $p_{(w_1,w_2)}$ for each formula $\varphi_{(w_1,w_2)}$. However, instead of forcing $p_{(w_1,w_2)}$ to have the same extension as $\varphi_{(w_1,w_2)}$ by using $\Gamma^m(\varphi_{(w_1,w_2)}, p_{(w_1,w_2)})$ as defined above, we ensure that by using the explicit assignment $[p_{(w_1,w_2)} := \varphi_{(w_1,w_2)}]$.

As in the proof of Theorem 1, let G_i be the set of pairs (w_1, w_2) that are distinguishable by a formula of depth i but not by a formula of lesser depth. Recall that for every $(w_1, w_2) \in G_i$, either $\mathsf{EX} \bigwedge_{w_2 R_2 w_2'} \varphi_{(w_1',w_2')}$ or $\mathsf{AX} \bigvee_{w_1 R_1 w_1'} \varphi_{(w_1',w_2')}$ distinguishes between $(\mathcal{M}_1, w_1)$ and $(\mathcal{M}_2, w_2)$. In the first case, let $\psi_{(w_1,w_2)} := \mathsf{EX} \bigwedge_{w_2 R_2 w_2'} p_{(w_1',w_2')}$, in the second case let $\psi_{(w_1,w_2)} := \mathsf{AX} \bigvee_{w_1 R_1 w_1'} p_{(w_1',w_2')}$.

Now, order $\bigcup_{i=0}^{n} G_i = \{(x_0, y_0), \cdots, (x_k, y_k)\}$ in such a way that (x_j, y_j) comes before $(x_{j'}, y_{j'})$ if there is some i such that $(x_j, y_j) \in G_i$ and $(x_{j'}, y_{j'}) \notin G_i$. Finally, for any distinguishable $(w_1, w_2) \in W_1 \times W_2$ let

$$\chi_{(w_1,w_2)} := [p_{(x_0,y_0)} := \psi_{(x_0,y_0)}] \cdots [p_{(x_k,y_k)} := \psi_{(x_k,y_k)}]\psi_{(w_1,w_2)}.$$

This $\chi_{(w_1,w_2)}$ is equivalent to $\varphi_{(w_1,w_2)}$, so it distinguishes between w_1 and w_2. It contains at most $\mathrm{card}(W_1 \times W_2)$ subformulae $\psi_{(x_j,y_j)}$ and each such formula is of length at most $\max(\mathrm{card}(W_1), \mathrm{card}(W_2))$. So $\chi_{(w_1,w_2)}$ is of length at most cubic in the size of $\mathcal{M}_1$ and $\mathcal{M}_2$.

C. Polynomial size distinguishing formulae in DAG format

Both proposals above for producing distinguishing formulae of polynomial size hinge on the observation that the shortest such distinguishing formulae contain only polynomially many non-equivalent subformulae. This observation can be put to work more explicitly by treating formulae not as strings of symbols, but as represented by directed acyclic graphs (DAG), with nodes labelled by subformulae of the formula at the root and arcs representing the subformula relation (see e.g. [5]). Thus, the DAG-based representation of the shortest distinguishing formulae only involves polynomially many nodes and can therefore be exponentially more succinct than the string representations of these formulae. Therefore, the polynomial size can be achieved automatically, by adopting the more succinct representation.

D. Finding distinguishing formulae in polynomial time

One important consequence of the observations and results from this section is that, the shortest distinguishing formula can be constructed and represented in polynomial space by using renaming, or explicit assignments, or DAGs, as a data structure for their representation.

Furthermore, a quick look at the way we construct the distinguishing formulae in the proofs of Theorem 1 and in Sections V-A and V-B shows that we can do so in polynomial time as well. For example, Algorithm 1, which is inspired by the proof of Theorem 1 and the assignments used in Section V-B, computes a distinguishing formula for every non-bisimilar pair $(w_1, w_2) \in W_1 \times W_2$, and it runs in $O(n^5)$ time, where $n = \mathrm{card}(W_1) + \mathrm{card}(W_2)$.

Do note that, in order for the $O(n^5)$ bound to hold, it is critical that we treat occurrences of $f(w_1', w_2')$ as atoms; i.e. when we write $\mathsf{EX} \bigwedge_{w_2' \in R_2(w_2)} f(w_1', w_2')$ or $\mathsf{AX} \bigvee_{w_1' \in R_1(w_1)} f(w_1', w_2')$ we do not add a copy of each $f(w_1', w_2')$ but merely a reference.

VI. The minimal nesting depth of distinguishing formulae

Lastly, we analyse the question of the minimal nesting depth of distinguishing formulae between given non-bisimilar pointed transition systems $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$. It is easy to see that minimal nesting depth equals 0 if $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$ are not 0-bisimilar, else equals $n + 1$ where n is the unique number such that $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$ are n-bisimilar but not $(n + 1)$-bisimilar. Equivalently, this is the minimal number of rounds needed for Spoiler to win the bisimulation game from the respective initial configuration. Here we obtain some tight upper bounds for that parameter.

To begin with, given a pair of non-bisimilar finite pointed transition systems $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$, with $\mathcal{M}_1 = (W_1, R_1, V_1)$ and $\mathcal{M}_2 = (W_2, R_2, V_2)$, an obvious upper bound for the smallest nesting depth of a distinguishing formula between $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$ is $\mathrm{card}(W_1) \times \mathrm{card}(W_2)$. Indeed, if Spoiler has a winning strategy for the bisimulation game between $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$, then Spoiler has such a winning strategy that avoids repeating configurations (for, if Duplicator can force one repetition, then she can force repetitions forever).

Algorithm 1 Diff$(\mathcal{M}_1, \mathcal{M}_2)$

Input: transition systems $\mathcal{M}_1 = (W_1, R_1, V_1)$ and $\mathcal{M}_2 = (W_2, R_2, V_2)$.
Output: $f : W_1 \times W_2 \to \mathcal{L}_{\mathrm{ML}} \cup \{\mathrm{NULL}\}$ such that for every $w_1, w_2 \in W_1 \times W_2$: (a) $f(w_1, w_2) = \mathrm{NULL}$ if and only if $(\mathcal{M}_1, w_1) \rightleftarrows (\mathcal{M}_2, w_2)$ and (b) $\mathcal{M}_1, w_1 \models f(w_1, w_2)$ and $\mathcal{M}_2, w_2 \not\models f(w_1, w_2)$ if $(\mathcal{M}_1, w_1) \not\sim (\mathcal{M}_2, w_2)$.

Initialize $f(w_1, w_2) = \mathrm{NULL}$ for all $(w_1, w_2) \in W_1 \times W_2$
For all $(w_1, w_2) \in W_1 \times W_2$ do
  For all $p \in \mathcal{P}_{\mathcal{M}_1} \cup \mathcal{P}_{\mathcal{M}_2}$ do
    If $w_1 \in V_1(p)$ and $w_2 \notin V_2(p)$ then set $f(w_1, w_2) = p$
    If $w_1 \notin V_1(p)$ and $w_2 \in V_2(p)$ then set $f(w_1, w_2) = \neg p$
  od
od
For $1 \leq i \leq \mathrm{card}(W_1) + \mathrm{card}(W_2)$ do
  For all (w_1, w_2) such that $f(w_1, w_2) == \mathrm{NULL}$ do
    For all $w_1' \in R_1(w_1)$ do
      If for all $w_2' \in R_2(w_2) : f(w_1', w_2') \neq \mathrm{NULL}$
      then set $f(w_1, w_2) = \mathsf{EX} \bigwedge_{w_2' \in R_2(w_2)} f(w_1', w_2')$
    od
    For all and $w_2' \in R_2(w_1)$ do
      If for all $w_1' \in R_1(w_1) : f(w_1', w_2') \neq \mathrm{NULL}$
      then set $f(w_1, w_2) = \mathsf{AX} \bigvee_{w_1' \in R_1(w_1)} f(w_1', w_2')$
    od
  od
od
Return f

This upper bound, however, turns out to be very crude. A much better bound is obtained if we observe that the minimal such nesting depth is closely related to the number of iterations in the computation of the largest bisimulation between the two pointed transition systems as a greatest fixed point of the respective monotone operator encoding the one-step back-and-forth property. We only sketch the relevant definitions and claims here, and refer the reader to [8, Section 3.5], or partly to [5, Section 3.4] for further details and proofs.

First, we note that the largest bisimulation relation between two given pointed transition systems $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$, is uniquely defined as the union of all bisimulation relations between them, and equivalently as the greatest fixed point of a monotone operator on $W_1 \times W_2$ defined as follows.

Given a relation $X \subseteq W_1 \times W_2$ and a pair $(s_1, s_2) \in W_1 \times W_2$, we say that (s_1, s_2) has the *back-and-forth property with respect to* X, denoted $BF((s_1, s_2), X)$, iff Spoiler has a single round strategy in the bisimulation game between $\mathcal{M}_1$ and $\mathcal{M}_2$ to lead from the configuration (s_1, s_2) to a configuration $(r_1, r_2) \in X$; that is, if the respective Back and Forth conditions are satisfied with respect to the pair (s_1, s_2) and the relation X.[6] Now, consider the following operator $F = F_{(\mathcal{M}_1, \mathcal{M}_2)}$ on subsets $X \subseteq W_1 \times W_2$:

$$F(X) := \big\{(s_1, s_2) \in X \mid BF((s_1, s_2), X)\big\}.$$

Clearly, F is monotone in the sense that $X \subseteq Y$ implies that $F(X) \subseteq F(Y)$. Therefore, by the Knaster-Tarski Theorem, F has a (unique) greatest fixpoint in restriction to any subset of $X \subseteq W_1 \times W_2$. We are interested in the greatest fixpoint of F that respects propositional equivalence, so by default we will apply F to the set

$$X_{\simeq} := \{(s_1, s_2) \in W_1 \times W_2 \mid s_1 \simeq s_2\}$$

(Recall that $\simeq$ denotes propositional equivalence, i.e. satisfying the same atomic propositions.)

The iterations of the application of F computing that greatest fixpoint, computed as $X_0 = X_{\simeq}$ and $X_{n+1} = F(X_n)$ for $n \geq 0$, eventually stabilise with value $\nu X.F(X)$ which gives the largest bisimulation $\beta(\mathcal{M}_1, \mathcal{M}_2)$ between $\mathcal{M}_1$ and $\mathcal{M}_2$. Thus, $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$ are bisimilar iff $(u_1, u_2) \in \beta(\mathcal{M}_1, \mathcal{M}_2)$.

Now, the following result (often used without being explicitly stated and proved) will eventually yield the tight upper bounds for the smallest nesting depth of distinguishing formulae.

Proposition 4. *Given the finite transition systems $\mathcal{M}_1$ and $\mathcal{M}_2$ with sets of states resp. W_1 and W_2, the greatest fixpoint $\nu X.F(X)$ of the operator F is reached within a number of iterations bounded above by $\mathrm{card}(W_1)+\mathrm{card}(W_2)-m$, where m is the number of different labels[7] of states in $\mathcal{M}_1 \cup \mathcal{M}_2$.*

Proof: Suppose, without loss of generality, that $\mathcal{M}_1$ and $\mathcal{M}_2$ have disjoint sets of states and let $W = W_1 \cup W_2$. Now, consider the operator F, defined above, as applied in the (disjoint) union $\mathcal{M}_1 \cup \mathcal{M}_2$. Then note that the starting set $X_0 = X_{\simeq}$, defined above, is an equivalence relation in W. It is easy to show by induction on n that every iteration $F^n(X_0)$ is an equivalence relation in W. Since every equivalence relation in W can be identified with the partition in W that it generates, it follows that every iteration step of the computation of $\nu X.F(X)$ corresponds to a refinement of the previous partition of W. Next, note that the number of clusters in the partition corresponding to X_0 equals the number of different labels of states in $\mathcal{M}_1 \cup \mathcal{M}_2$, that is m, and that every refinement before stabilisation strictly increases the number of clusters in the current partition. Therefore, the maximal number of refining iterations is bounded above by $\mathrm{card}(W_1) + \mathrm{card}(W_2) - m$, whence the claim. ∎

What remains to be seen is how the number of iterations in the computation of $\nu X.F(X)$ relates to the depth of distinguishing formulae between non-bisimilar transition systems. The next proposition gives the answer.

Proposition 5. *For every $n \in \mathbb{N}$:*

$$F^n(X_{\simeq}) = \big\{(s_1, s_2) \in W_1 \times W_2 \mid \mathcal{M}_1, s_1 \rightleftarrows_n \mathcal{M}_2, s_2\big\}.$$

[6] Note that the Back and Forth conditions for a bisimulation relation β say that each pairs in β has the *back*-and-*forth* property with respect to β itself.

[7] The label of a state is the set of atomic propositions true at that state.

The proof is by straightforward induction on $n \in \mathbb{N}$, using directly the definition of the operator F.

Finally, here is the argument that relates the results above. If the pointed transition systems $(\mathcal{M}_1, s_1)$ and $(\mathcal{M}_2, s_2)$ are not bisimilar, then either they are not 0-bisimular (i.e., the labels of s_1 and s_2 differ) – in which case there is a distinguishing formula of depth 0 – or there is a number $n \in \mathbb{N}$ such that they are n-bisimular, but not $(n+1)$-bisimular. Then, by Proposition 5, (s_1, s_2) is in $F^n(X_\simeq)$ but not in $F^{n+1}(X_\simeq)$, which is only possible if the greatest fixpoint of the operator F occurs after more than n iterations, hence $n \leq \mathrm{card}(W_1) + \mathrm{card}(W_2) - m$, where m is the number of different labels of states in $\mathcal{M}_1 \cup \mathcal{M}_2$. This gives us an upper bound for the minimal nesting depth of a distinguishing formula between $(\mathcal{M}_1, s_1)$ and $(\mathcal{M}_2, s_2)$, thus proving the following.

Theorem 6. *For every pair of finite non-bisimilar pointed transition systems $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$ with respective sets of states W_1 and W_2 there is a modal formula φ such that $\mathcal{M}_1, u_1 \models \varphi$, $\mathcal{M}_2, u_2 \not\models \varphi$ and $\mathsf{nd}(\varphi) \leq \mathrm{card}(W_1) + \mathrm{card}(W_2) - m$, where m is the number of different labels of states in $\mathcal{M}_1 \cup \mathcal{M}_2$.*

In every pair of transition systems there is at least one state label, so we have the following corollary.

Corollary 7. *For every pair of finite non-bisimilar pointed transitions systems $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$ with respective sets of states W_1 and W_2, there is a formula φ such that $\mathcal{M}_1, u_1 \models \varphi$, $\mathcal{M}_2, u_2 \not\models \varphi$ and $\mathsf{nd}(\varphi) \leq \mathrm{card}(W_1) + \mathrm{card}(W_2) - 1$.*

A natural constraint in transition systems is seriality, i.e. that every state has a successor. Note that any two serial transition systems are bisimilar unless there are at least two different state labels, so we also have the following.

Corollary 8. *For every pair of finite non-bisimilar pointed serial transitions systems $(\mathcal{M}_1, u_1)$ and $(\mathcal{M}_2, u_2)$ with respective sets of states W_1 and W_2, there is a formula φ such that $\mathcal{M}_1, u_1 \models \varphi$, $\mathcal{M}_2, u_2 \not\models \varphi$ and $\mathsf{nd}(\varphi) \leq \mathrm{card}(W_1) + \mathrm{card}(W_2) - 2$.*

The bounds given in both these corollaries are tight. In order to see that the bound $\mathrm{card}(W_1)+\mathrm{card}(W_2)-1$ is tight, let $\mathcal{M}_1$ and $\mathcal{M}_2$ be as shown in Figure 3. The lowest depth formula that distinguishes between $(\mathcal{M}_1, s_1)$ and $(\mathcal{M}_2, t)$ is $\mathsf{AX}^n\bot$, which is of depth $n = \mathrm{card}(W_1) + \mathrm{card}(W_2) - 1$.

Fig. 3: Pointed transition systems $(\mathcal{M}_1, s_1)$ and $(\mathcal{M}_2, t)$.

(a) $\mathcal{M}_1$

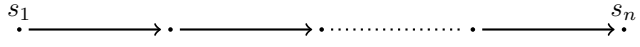

(b) $\mathcal{M}_2$

In order to see that the bound of $\mathrm{card}(W_1) + \mathrm{card}(W_2) - 2$ for serial transition systems is tight, let $\mathcal{M}_3$ and $\mathcal{M}_4$ be as in Figure 4. We have $\mathrm{card}(W_1) = n$ and $\mathrm{card}(W_2) = n+1$ Furthermore, the lowest depth formula that distinguishes between $(\mathcal{M}_3, s_1)$ and $(\mathcal{M}_4, t_1)$ is $\mathsf{EX}^{2n-1}p$, which is of depth $2n - 1 = \mathrm{card}(W_1) + \mathrm{card}(W_2) - 2$.

Fig. 4: Pointed transition systems $(\mathcal{M}_3, s_1)$ and $(\mathcal{M}_4, t)$.

(a) $\mathcal{M}_3$

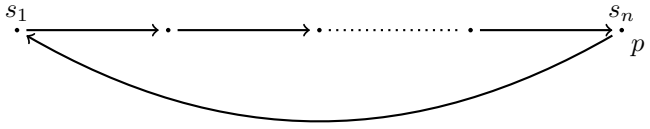

(b) $\mathcal{M}_4$

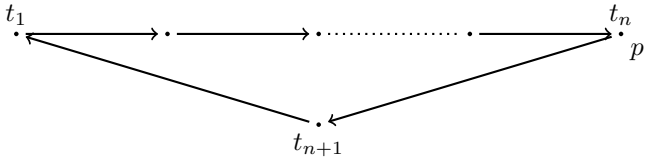

Lastly, we note that the minimal nesting depth of distinguishing formulae in temporal logics where the accessibility relation is transitive, or that contain (as primitive or definable) universal, master, or reachability modality, depends substantially on structural details and specific additional assumptions, so we leave it out here. For instance, note that using the reachability modality EF, the smallest distinguishing formula for the example in Figure 3 is $\mathsf{EF}\,\mathsf{AX}\bot$, of depth 2, while the smallest such distinguishing formula for the example in Figure 4 is $\mathsf{EF}(p \wedge \mathsf{AX}^n p)$, of depth $n + 1$. Thus, the bounds for the nesting depth established for formulae in the basic modal logic are generally not optimal for stronger languages under suitable assumptions.

VII. Concluding remarks

In summary, here we have showed that the smallest formula in either of the basic modal logic ML, its extension with past operators TL, and the computation tree logics CTL and CTL*, distinguishing between two non-bisimilar pointed transition systems is of size at most exponential (more precisely, n^n) in the combined number of states n of the transition systems. Furthermore, we have showed an example with exponential lower bound. We have also showed that the lowest nesting depth of a formula in basic modal logic that distinguishes between two non-bisimilar pointed transition systems is bounded above by $\mathrm{card}(W_1)+\mathrm{card}(W_2)-1$, where W_1 and W_2 are the domains of the transition systems. For serial transition systems, we have obtained the sharper bound of $\mathrm{card}(W_1)+\mathrm{card}(W_2)-2$. Both these bounds are tight.

Most of the facts and results used here are widely known, some almost folklore, but we have not found published references where they are explicitly stated and proved, so we have done that here *inter alia*, thus probably filling some gaps in the literature.

The present work leaves a few still open questions, of which we mention again two.

1) As noted in Section IV, there is a certain gap between the upper bound $n^n = 2^{n \log n}$ established in Theorem 1 and the single exponential lower bound obtained in Theorem 2. We conjecture that the upper bound can be reduced to $2^{O(n)}$.

2) We have not explored yet the precise bounds for the length and nesting depth of distinguishing formulae in more expressive languages, most notably the μ-calculus.

Acknowledgments. We thank the anonymous reviewers for some useful comments and several corrections.

REFERENCES

[1] C. Baier and J.P. Katoen. *Principles of Model Checking*. MIT Press, 2008.

[2] P. Blackburn, M. de Rijke, and V. Venema. *Modal Logic*. Cambridge Univ. Press, 2001.

[3] M. Browne, E. Clarke, and O. Grümberg. Characterizing finite Kripke structures in propositional temporal logic. *Theoretical Computer Science*, 59:115–131, 1988.

[4] Th. Boy de la Tour. An optimality result for clause form translation. *J. of Symb. Computation*, 14:283–301, 1992.

[5] S. Demri, V. Goranko, and M. Lange. *Temporal Logics in Computer Science*. Cambridge Univ. Press, 2016.

[6] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. Dynamic epistemic logic with assignment. In *Proc. of AAMAS 2005*, pages 141–148, 2005.

[7] E.A. Emerson. Temporal and modal logics. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics, pages 995–1072. MIT Press, 1990.

[8] V. Goranko and M. Otto. Model theory of modal logic. In *Handbook of Modal Logic*, pages 249–330. Elsevier, 2007.

[9] Petr Jancar, Antonín Kucera, and Richard Mayr. Deciding bisimulation-like equivalences with finite-state processes. *Theoretical Computer Science*, 258(1-2):409–433, 2001.

[10] C. Stirling. Bisimulation, modal logic and model checking games. *Logic Journal of the IGPL*, 7(1):103–124, 1999.

[11] Michael L. Tiomkin and Johann A. Makowsky. Propositional dynamic logic with local assignments. *Theor. Comput. Sci.*, 36:71–87, 1985.