

Open Systems in Reactive Environments: Control and Synthesis*

Orna Kupferman¹, P. Madhusudan^{2**}, P.S. Thiagarajan³, and Moshe Y. Vardi^{4***}

¹ Hebrew University orna@cs.huji.ac.il

² The Institute of Mathematical Sciences madhu@imsc.ernet.in

³ Chennai Mathematical Institute pst@smi.ernet.in

⁴ Rice University vardi@cs.rice.edu

Abstract. We study the problems of synthesizing open systems as well as controllers for them. The key aspect of our model is that it caters to *reactive environments*, which can disable different sets of responses when reacting with the system. We deal with specifications given as formulas in CTL* and its sub-logic CTL. We show that both these problems, with specifications in CTL (CTL*), are 2EXPTIME-complete (resp. 3EXPTIME-complete). Thus, in a sense, reactive environments constitute a provably harder setting for the synthesis of open systems and controllers for them.

1 Introduction

A *closed* system is a system whose behavior is completely determined by its state. An *open* system is one that interacts with its environment and whose behavior crucially depends on this interaction [HP85]. In an open system, we are required to distinguish between output signals (generated by the system) over which we have control and input signals (generated by the environment) over which we have no control. Given a specification of an open system, say as a temporal logic formula, satisfiability of the formula only guarantees that we can synthesize a system for *some* environment, whereas we need to synthesize a system that meets the specification for *all* environments.

To make this intuition more precise, suppose we are given finite sets I and O of input and output signals. A *program* can be viewed as a *strategy* $f : (2^I)^* \rightarrow 2^O$ that maps finite sequences of input signal sets into an output signal set. When f interacts with an environment that generates infinite input sequences what results is an infinite computation over $2^{I \cup O}$. Though f is deterministic, it produces a computation tree. The branches of the tree correspond to external non-determinism caused by the different possible inputs. One can now specify a

* For a full version of this extended abstract, see Technical Report TCS-2000-03, Chennai Mathematical Institute, available at www.smi.ernet.in

** On leave visiting Informatik VII, RWTH-Aachen, Germany

*** Supported in part by NSF grant CCR-9700061, and by a grant from the Intel Corporation.

property of an open system by a linear or branching temporal logic formula (over $I \cup O$). Unlike linear temporal logics, in branching temporal logics one can specify possibility requirements such as “every input sequence can be extended so that the output signal v eventually becomes true” (cf. [DTV99]). This is achieved via existential and universal quantification provided in branching temporal logics [Lam80, Eme90]. In this paper, we concentrate on branching temporal logics.

The *realizability problem* for a branching temporal logic is to determine, given a branching-time specification φ , whether there exists a program $f : (2^I)^* \rightarrow 2^O$ whose computation tree satisfies φ [ALW89, PR89]. Realizing φ boils down to synthesizing such a function f . An important aspect of the computation tree associated with f is that it has a fixed branching degree $|2^I|$. It reflects the assumption that at each stage, all possible subsets of I are provided by the environment. Such environments are referred to as *maximal environments*. Intuitively, these are static environments in terms of the branching possibilities they contribute to the associated computation trees. Equivalently, as we have noted already, each program has just one computation tree capturing its behavior in the presence of a maximal environment. In a more general setting, however, we have to consider environments that are, in turn, open systems. We term such environments *reactive*. They might offer different subsets of 2^I as input possibilities at different stages in the computation.

As an illustration, consider $I = \{r_1, \dots, r_n\}$ and $O = \{t_1, \dots, t_n\}$ where I represents n different types of resources and O represents n different types of tasks with the understanding that, at each stage, the system needs to receive r_i from the environment in order to execute t_i . In the case of the maximal environment, the specification “it is always possible to reach a stage where t_i is executed” ($AGEF(t_i)$ in CTL parlance) is realizable. This is so because at each stage in the computation, the maximal environment presents all possible combinations of the resources. In the case of the reactive environment, the above specification is *not* realizable; there could be an environment driven by an open system that can produce only a finite number of the resource r_i along any computation. In the resulting computation tree, each path eventually reaches a node in which the environment stop offering r_i . From then on, t_i cannot be executed.

So, a reactive environment associates a *set* of computation trees with a program. Consequently, in the presence of reactive environments, the realizability problem must seek a program *all* of whose computation trees satisfy the specification.

A closely related problem concerns the controllability of open systems. Here we are given an open system, often called a *plant* in this context, typically consisting of system states and environment states. We can control the moves made from the system states (i.e. disable some of the possible moves), but not the moves made from the environment states. Given a branching-time specification φ , the *control problem* is to come up with a strategy for controlling the moves made from the system states so that the resulting computation tree satisfies φ . Here again, assuming a reactive environment requires the controller to func-

tion correctly no matter how the environment disables some of its moves; thus correctness should be checked with respect to a whole set of computation trees.

We study control problems for both CTL* and CTL specifications. We show that the realizability problem can be reduced to the control problem. We prove that both these problems are 3EXPTIME-complete and 2EXPTIME-complete for CTL* and CTL, respectively. As established in [KV99a,KV99b], these problems are 2EXPTIME-complete and EXPTIME-complete for maximal environments, respectively. In this sense, reactive environments make it more difficult to realize open systems and synthesize controllers for them.

In the literature, a variety of realizability and control problems have been studied for open systems. These studies have been mainly in linear-time settings where the nature of the environment (maximal or reactive) does not play a role (this is since a maximal environment is the most challenging one for the system). In branching-time settings, the emphasis has been mainly on the realizability problem (often referred to as the *synthesis problem*) in the presence of maximal environments. For the linear time case, the literature goes back to a closely related realizability problem due to Church [Chu63], which was solved in [BL69] but more elegantly dealt with later using tree automata [Rab72,PR89]. In branching-time settings, the realizability problem for CTL* and CTL has been solved for maximal environments as cited above and more recently also for the μ -calculus [KV00].

As for controller synthesis, as mentioned above, most of the settings considered in the literature are linear time ones and often involve dealing with incomplete information [KG95,KS95,PR89,Var95]. As for branching-time settings, synthesis of memoryless controllers for settings with maximal environments is studied in [Ant95]. Also, for maximal environments, the control problem can be transformed (by flipping the role of the system and the environment) into *module-checking* problems solved in [KV96,KV97]. Yet another work, but in a different framework is [MT98] where the specification is also modeled by a transition system and the correctness criterion is that there should be a behavior preserving simulation from the plant to the specification. This has been later extended to the case of bisimulation in [MT00].

2 The Problem Setting

We assume familiarity with the branching temporal logic CTL* and its sublogic CTL (see [Eme90] for details). Here we just fix the notation for (Kripke) structures, which provide the semantics. A (*Kripke*) *structure* is a tuple $S = \langle AP, W, R, w_0, L \rangle$, where AP is the set of atomic propositions, W is a set of states, $R \subseteq W \times W$ is a transition relation that must be total (i.e., for every $w \in W$ there exists $w' \in W$ such that $R(w, w')$), w_0 is an initial state, and $L : W \rightarrow 2^{AP}$ maps each state to a set of atomic propositions true in this state. For w and w' with $R(w, w')$, we say that w' is a successor of w . A *path* of S is an infinite sequence $\pi = w^0, w^1, \dots$ of states such that for every $i \geq 0$, we have $R(w^i, w^{i+1})$. The suffix w^i, w^{i+1}, \dots of π is denoted by π^i . We use $w \models \varphi$ to

indicate that a state formula φ holds at state w , and we use $\pi \models \varphi$ to indicate that a path formula φ holds at path π (assuming a structure S). We say that S is a model of φ to mean that $w_0 \models \varphi$.

We model a *plant* as $P = \langle AP, W_s, W_e, R, w_0, L \rangle$, where AP, R, w_0 , and L are as in a structure with $W = W_s \cup W_e$. Here W_s is a set of *system states* and W_e is a set of *environment states*. Throughout what follows we are concerned only with *finite* plants; AP and W are both finite sets. We also assume that $W_s \cap W_e = \emptyset$. The *size* of the plant is $|P| = |W| + |R|$.

Given a finite set \mathcal{Y} , an \mathcal{Y} -*tree* is a nonempty set $T \subseteq \mathcal{Y}^*$ such that, if $v.c \in T$ with $v \in \mathcal{Y}^*$ and $c \in \mathcal{Y}$, then $v \in T$. When \mathcal{Y} is clear from the context we call T a tree. The elements of T are called *nodes* and the empty word ε is the *root* of T . For every $v \in T$, the set $\text{succ}_T(v) = \{v.c \mid c \in \mathcal{Y} \text{ and } v.c \in T\}$ is the set of children (successors) of v . Where T is clear from the context, we drop the subscript T and write $\text{succ}(v)$. In what follows, every tree T we consider is assumed to satisfy $\text{succ}(v) \neq \emptyset$ for every v in T . We associate a direction $\text{dir}(v) \in \mathcal{Y}$ with each node $v \in T$. A designated element $c_0 \in \mathcal{Y}$ is the direction of ε . For each non-root node $v.c$ with $c \in \mathcal{Y}$, we set $\text{dir}(v.c) = c$. A *path* π of a tree is a minimal set $\pi \subseteq T$ such that $\varepsilon \in \pi$ and for each $v \in \pi$, we have $|\pi \cap \text{succ}(v)| = 1$. Finally, given a set Σ , a Σ -labeled \mathcal{Y} -tree is a pair (T, V) where T is a \mathcal{Y} -tree and $V : T \rightarrow \Sigma$ is a labeling function. Of special interest to us are 2^{AP} -labeled trees. We call such trees *computation trees* and we sometimes interpret CTL* formulas with respect to them. Formally, a computation tree (T, V) satisfies a CTL* formula φ if φ is satisfied in the infinite-state structure $\langle AP, T, R_T, \varepsilon, V \rangle$, where $R_T(v, v.c)$ iff $v.c \in \text{succ}_T(v)$.

Let $P = \langle AP, W_s, W_e, R, w_0, L \rangle$ be a plant. Then P can be unwound into a W -tree T_P in the obvious manner; thus $\varepsilon \in T_P$, and we set $\text{dir}(\varepsilon) = w_0$, and for all $v \in T_P$, we have that $v.c \in T_P$ iff $R(\text{dir}(v), c)$. The tree T_P induces the 2^{AP} -labelled tree (T_P, V_P) where for each $v \in T_P$, we have $V_P(v) = L(\text{dir}(v))$. A *restriction* (T, V) of (T_P, V_P) is a tree $T \subseteq T_P$ such that V is V_P restricted to T . We denote by $T_P^s = \{v \mid v \in T_P \text{ and } \text{dir}(v) \in W_s\}$ the set of nodes of T_P that correspond to system states, and by $T_P^e = T_P \setminus T_P^s$ the set of states that correspond to environment states.

Consider a plant P and a restriction (T, V) of (T_P, V_P) . We can think of (T, V) as a computation tree generated by an interaction of the system with its environment: each position in this interaction corresponds to a node $v \in T$. When $v \in T_P^s$, the system chooses a nonempty subset of $\text{succ}_{T_P}(v)$. This subset corresponds to successors of $\text{dir}(v)$ to which the system enables the transition from $\text{dir}(v)$. Similarly, when $v \in T_P^e$, the environment chooses a nonempty subset of $\text{succ}_{T_P}(v)$ corresponding to the successors of $\text{dir}(v)$ to which transitions are enabled. Note we can have v and v' with $\text{dir}(v) = \text{dir}(v')$ and still $\text{succ}_T(v) \neq \text{succ}_T(v')$. Indeed, the decisions made by the system and the environment depend not only on the current state of the interaction (that is, $\text{dir}(v)$), but also in the entire interaction between the system and the environment so far (that is, v). Intuitively, given P and a CTL* formula φ , the control problem is to come up

with a strategy for the system so that no matter how the environment responds, the resulting restriction of (T_P, V_P) satisfies φ .

We now make this intuition formal. A *strategy* for the system is a function g that assigns to each $v \in T_P^s$, a non-empty subset of $\text{succ}_{T_P}(v)$. A *g -respecting execution* of P is a restriction (T, V) of (T_P, V_P) such that for every $v \in T \cap T_P^s$, we have $\text{succ}_T(v) = g(v)$. Note that a strategy g determines only the successors of nodes in T_P that correspond to system states. Thus, P may have many g -respecting executions, each corresponding to a different strategy of the environment. Given a CTL* formula φ , the strategy g is *winning* for φ if all the g -respecting executions satisfy φ . The *control problem* is then to decide, given a plant P and a CTL* specification φ , whether the system has a strategy winning for φ , denoted $\text{controllable}(P, \varphi)$.

The realizability problem for programs with reactive environments can be tackled along very similar lines: we consider a program f interacting with its environment via two finite sets I and O of input and output signals respectively. We can view f as a strategy $f : (2^I)^* \rightarrow 2^O$ that maps finite sequences of input signal sets into an output signal set. The interaction starts by the program generating the output $f(\varepsilon)$. The environment replies with some $i_1 \subseteq I$. In general, $f(i_1, \dots, i_j)$, is the response of f for the input sequence i_1, \dots, i_j . This (infinite) interaction can be represented by a computation tree. The branches of the tree correspond to external non-determinism caused by different input signal sets chosen by the environment. Thus f can be viewed as the full $2^{I \cup O}$ -labeled 2^I -tree (T_f, V_f) with $T_f = (2^I)^*$ and $V_f(v) = \text{dir}(v) \cup f(v)$ for each $v \in T_f$. Given a CTL* formula φ , the *realizability problem* is to find a strategy f so that φ is satisfied in f no matter how the environment disables (in a non-blocking manner) its possible responses at different stages. Formally, let $f : (2^I)^* \rightarrow 2^O$ be a strategy and let (T, V) be a $2^{I \cup O}$ -labeled 2^I -tree. We say that (T, V) is an *f -respecting execution* iff $V(v) = f(v) \cup \text{dir}(v)$ for each $v \in T$. (In other words, (T, V) is a restriction of the computation tree (T_f, V_f)). We say that f *realizes* φ if every f -respecting execution satisfies φ . Also, φ is *realizable* iff there is a program f that realizes φ .

Using a universal plant, which embodies all the possible assignments to I and O , the program-synthesis problem can be reduced to supervisory control.

Lemma 1. *Let φ be a CTL* (CTL) formula over $AP = I \cup O$. We can effectively construct a finite plant P and a CTL* (resp. CTL) formula φ' such that $|P| = O((2^{AP})^2)$, $|\varphi'| = O(|\varphi| + 2^{AP})$, and φ is realizable iff $\text{controllable}(P, \varphi')$.*

Proof. We sketch the reduction for the case of CTL. A similar procedure can be developed for CTL*. We define $P = (AP', W_s, W_e, R, w_0, L)$ as follows:

- $AP' = AP \cup \{p_e\}$ with $p_e \notin AP$. (The role of p_e will become clear soon).
- $W_e = 2^I \times 2^O$, and $W_s = \{w_0\} \cup 2^I$ with $w_0 \notin 2^I \cup W_e$
- $R = R_0 \cup R_1 \cup R_2$ where $R_0 = \{w_0\} \times (\{\emptyset\} \times 2^O)$, $R_1 = W_e \times 2^I$, and $R_2 = \{(X, (X, Y)) \mid X \subseteq I \text{ and } Y \subseteq O\}$.
- $L((X, Y)) = X \cup Y \cup \{p_e\}$ for each $(X, Y) \in W_e$ and $L(w) = \emptyset$ for each $w \in W_s$.

Next we construct the CTL formula φ' over AP' by setting $\varphi' = \varphi'_1 \wedge \varphi'_2$. The role of φ'_1 is to ensure that the truthhood of φ matters only at the states in W_e . The conjunct φ'_1 is the formula $EX(\|\varphi\|)$ where $\|\varphi\|$ is given inductively via: $\|p\| = p$ for $p \in AP$, $\|\neg\varphi\| = \neg\|\varphi\|$ and $\|\varphi_1 \vee \varphi_2\| = \|\varphi_1\| \vee \|\varphi_2\|$. Further, $\|EX\varphi\| = EXEX(\|\varphi\|)$ and $\|E(\varphi_1 U \varphi_2)\| = E((p_e \rightarrow \|\varphi_1\|) U (p_e \wedge \|\varphi_2\|))$. Finally, $\|A(\varphi_1 U \varphi_2)\|$ is defined by replacing E by A in the clause for $\|E(\varphi_1 U \varphi_2)\|$.

The conjunct φ'_2 ensures that the system chooses only one move at states in W_s (since the 2^O labelling required must be unique). It is given by $\varphi'_2 = AG(\neg p_e \rightarrow (\bigwedge_{z \in O} (EXz \rightarrow AXz)))$. It is easy to check that P and φ' have the required properties. \square

3 Upper bounds

Since our decision procedure uses tree automata, we first introduce Büchi and parity tree automata (see [Tho97] for more details). Tree automata run on Σ -labeled Υ -trees. We assume that the set Υ of directions is ordered. For an Υ -tree T and a node $v \in T$, we use $o_succ(v)$ to denote an ordered list of $succ(v)$. Let $|\Upsilon| = k$, and let $[k] = \{1, \dots, k\}$. A *nondeterministic tree automaton* over Σ -labeled Υ -trees is $A = \langle \Sigma, Q, \delta, Q_0, \mathcal{F} \rangle$ where Σ is a finite alphabet, Q is a finite set of states, $\delta : Q \times \Sigma \times [k] \rightarrow 2^{Q^*}$ is a transition function that maps a state, a letter, and a branching degree $i \in [k]$ to a set of i -tuples of states, $Q_0 \subseteq Q$ is the set of initial states, and \mathcal{F} is an acceptance condition that depends on the kind of automata we consider. For *Büchi automata*, \mathcal{F} is a subset of Q and for *parity automata*, \mathcal{F} is a function $\mathcal{F} : Q \rightarrow \{0, \dots, h\}$ for some $h \in \mathbb{N}$ (the set $\{0, \dots, h\}$ is called the set of *colors*).

Let (T, V) be a Σ -labeled tree. A *run* of A over (T, V) is a Q -labeled tree (T, ρ) such that $\rho(\varepsilon) \in Q_0$, and for all $v \in T$ with $o_succ(v) = \langle x_1, \dots, x_l \rangle$, we have $\langle \rho(x_1), \dots, \rho(x_l) \rangle \in \delta(q, V(x), l)$. For a path π of T , let $Inf_\rho(\pi)$ be the set of states that appear as the labels of infinitely many nodes on π . If A is a Büchi automaton, then ρ is *accepting* if for every path π of T , $Inf_\rho(\pi) \cap \mathcal{F} \neq \emptyset$. For parity automata, ρ is *accepting* if for every path π of T , $\min(\mathcal{F}(Inf_\rho(\pi)))$ is an even number. A Σ -labeled tree (T, V) is accepted by A iff there is an accepting run of A over (T, V) . The language accepted by A is the set of all Σ -labeled trees accepted by A . We say that A is *universal* iff it accepts all Σ -labeled Υ -trees.

There are well-known connections relating branching temporal logics and tree automata. For our purposes, it suffices to know the following.

Theorem 1.

- (1) [VW86] *Given a CTL formula φ over AP and a finite set Υ , we can construct a nondeterministic Büchi tree automaton $A_{\Upsilon, \varphi}$ with $2^{O(|\varphi|)}$ states that accepts exactly the set of 2^{AP} -labeled Υ -trees that satisfy φ .*
- (2) [EJ88, Saf88, Tho97] *Given a CTL* formula φ over AP and a set Υ , one can construct a nondeterministic parity tree automaton $A_{\Upsilon, \varphi}$ with $2^{2^{O(|\varphi|)}}$ states and $2^{O(|\varphi|)}$ colors that accepts exactly the set of 2^{AP} -labeled Υ -trees that satisfy φ . \square*

In the control problem, we are given a plant $P = \langle AP, W_s, W_e, R, w_0, L \rangle$ and a CTL (or CTL*) formula φ over AP , and we have to decide whether there is a strategy g for the system so that all the g -respecting executions of P satisfy φ .

Recall that a strategy g for the system assigns to each $v \in T_P^s$ a nonempty subset of $\text{succ}(v)$. We can associate with g a $\{\perp, \top, d\}$ -labeled W -tree (T_P, V_g) , where for every $v \in T_P$, the following hold:

- If $v \in T_P^s$, then the children of v that are members of $g(v)$ are labeled by \top , and the children of v that are not members of $g(v)$ are labeled by \perp .
- If $v \in T_P^e$, then all the children of v are labeled by d .

Intuitively, nodes $v.c$ are labeled by \top if g enables the transitions from $\text{dir}(v)$ to c (given that the execution so far has traversed v), they are labeled by \perp if g disables the transition from $\text{dir}(v)$ to c , and they are labeled by d if $\text{dir}(v)$ is an environment state, where the system has no control about the transition from $\text{dir}(v)$ to c being enabled. We call the tree (T_P, V_g) the *strategy tree* of g . Note that not every $\{\perp, \top, d\}$ -labeled W -tree (T_P, V) is a strategy tree. Indeed, in order to be a strategy tree, V should label all the successors of nodes corresponding to environment states by d , and it should label all the successors of nodes corresponding to system states by either \top or \perp , with at least one successor being labeled by \top . In fact, given a plant P with state space W , there is a nondeterministic Büchi automaton A_{stra} over $\{\perp, \top, d\}$ -labeled W -trees, such that A_{stra} has $|W|$ states and it accepts exactly all the strategy trees of P .

Next, given a formula φ we first construct a tree automaton A such that A accepts a strategy tree (T_P, V_g) iff there is a g -respecting execution of P that does not satisfy φ . More precisely, we have:

Theorem 2. *Given a plant P with state space W and a branching-time formula φ , we can construct a nondeterministic tree automaton A over $\{\perp, \top, d\}$ -labeled W -trees such that the following hold:*

1. *A accepts a strategy tree (T_P, V_g) iff there is a g -respecting execution of P that does not satisfy φ .*
2. *If φ is a CTL formula, then A is a Büchi automaton with $|W| \cdot 2^{O(|\varphi|)}$ states.*
3. *If φ is a CTL* formula, then A is a parity automaton with $|W| \cdot 2^{2^{O(|\varphi|)}}$ states and $2^{O(|\varphi|)}$ colors*

Proof. Let $P = \langle AP, W_s, W_e, R, w_0, L \rangle$, and let $A_{W, \neg\varphi} = \langle 2^{AP}, Q, \delta, Q_0, \mathcal{F} \rangle$ be the automaton that accepts exactly all 2^{AP} -labeled W -trees that satisfy $\neg\varphi$, as described in Theorem 1. We define $A = \langle \{\perp, \top, d\}, Q', \delta, Q'_0, \mathcal{F}' \rangle$ as follows.

- $Q' = (W \times Q \times \{\top, \perp\}) \cup \{q_{\text{acc}}\}$. The state q_{acc} is an accepting sink. Consider a state $(w, q, m) \in W \times Q \times \{\top, \perp\}$. The last component m is the *mode* of the state. When $m = \top$, it means that the transition to the current node is enabled (by either the system or the environment). When $m = \perp$, it means that the transition to the current node is disabled.

When A is at a state (w, q, \top) as it reads a node v , it means that $\text{dir}(v) = w$, and that v has to participate in the g -respecting execution. Hence, A can

read \top or d , but not \perp . If v is indeed labeled by \top or d , the automaton A guesses a subset of successors of w of some size $l \geq 1$. It then moves to states corresponding to the successors of w and q , with an appropriate update of the mode (\top for the successors in the guessed subset, \perp for the rest).

When A is in a state (w, q, \perp) and it reads a node v , it means that $\text{dir}(v) = w$ and that v does not take part in a g -respecting execution. Then, A expects to read \perp or d , in which case it goes to the accepting sink.

- $Q'_0 = \{w_0\} \times Q_0 \times \{\top\}$.
- The transition function δ' is defined, for all $w \in W$, $q \in Q$, and $l = |\text{succ}_{T_P}(w)|$ as follows.
 - $\delta((w, q, \top), \perp, l) = \delta((w, q, \perp), \top, l) = \emptyset$
 - If $x \in \{\perp, d\}$, then $\delta'((w, q, \perp), x, l) = \{q_{acc}, \dots, q_{acc}\}$,
 - If $x \in \{\top, d\}$, then $\delta'((w, q, \top), x, l)$ is defined as follows. Let $o_succ(w) = \langle w_1, \dots, w_l \rangle$ and let $Y = \{w_{y_1}, \dots, w_{y_n}\}$ be a nonempty subset (of size n) of $\text{succ}(w)$. Then, $\delta'((w, q, \top), x, l)$ contains all tuples $\langle (w_1, s_1, m_1), \dots, (w_l, s_l, m_l) \rangle$ such that there is $\langle q_1, \dots, q_n \rangle \in \delta(q, L(w), n)$ and for all $1 \leq i \leq l$, the following hold:
 - * If $w_i \in Y$, namely $w_i = w_{y_j}$ for some $1 \leq j \leq n$, then $s_i = q_j$ and $m_i = \top$.
 - * If $w_i \notin Y$, then $w_i = w$ (in fact, w_i can be an arbitrary state) and $m_i = \perp$.

Intuitively, δ' propagates the requirements imposed by $\delta(q, L(w), n)$ among the successors of w to which the transition from w is enabled.

Note that δ' is independent of w being a system or an environment state.

The type of w is taken into consideration only in the definition of A_{stra} .

- The final states are inherited from the formula automaton. Thus, if φ is in CTL, then $\mathcal{F} = (W \times \mathcal{F} \times \{\top, \perp\}) \cup \{q_{acc}\}$. If φ is in CTL*, let $\mathcal{F} : Q \rightarrow \{0, \dots, h\}$. Then, $\mathcal{F}' : Q' \rightarrow \{0, \dots, h\}$ is such that $\mathcal{F}(q_{acc}) = 0$ and for all $w \in W$, $q \in Q$, and $m \in \{\perp, \top\}$, we have $\mathcal{F}'((w, q, m)) = \mathcal{F}(q)$. \square

It is left to check whether the language of A_{stra} is contained in the language of A . Since tree automata are closed under complementation [Rab69, Tho97], we can complement A , get an automaton \tilde{A} , and then check the nonemptiness of the intersection of A_{stra} with \tilde{A} . Hence the following theorem.

Theorem 3. *Given a plant P and a formula φ in CTL, the control problem for φ is in 2EXPTIME. More precisely, it can be solved in time $O(\exp(|P|^2 \cdot 2^{O(|\varphi|)}))$.¹ For φ in CTL*, the problem is in 3EXPTIME. More precisely, it can be solved in time $O(\exp(|P|^2 \cdot 2^{2^{O(|\varphi|)}}))$.*

Proof. For the complexity of this procedure, it is easy to see that if φ is in CTL, the automaton A has a state-space size of $O(|P| \cdot 2^{O(|\varphi|)})$. Though A runs on k -ary trees (where k depends on P), it can be complemented as easily as automata on binary trees — the complemented automaton \tilde{A} (as well as its intersection with A_{stra}) is a parity automaton with $O(\exp(|P| \cdot 2^{O(|\varphi|)}))$ states

¹ $\exp(x)$ stands for $2^{O(x)}$

and $O(|P| \cdot 2^{O(|\varphi|)})$ colors ([Tho97]). Since emptiness of parity tree automata can be done in time polynomial in the state-space and exponential in the number of colors [EJ88,PR89], we can check emptiness of this intersection in time $O(\exp(|P|^2 \cdot 2^{O(|\varphi|)}))$. For CTL* specifications, the analysis is similar except that the complexity contributed by the formula increases by one exponential. \square

By [Rab69], if there is indeed a strategy that is winning for the system, then the automaton that is the product of A_{stra} and the complement of the automaton constructed on in Theorem 2 accepts it and when we test it for emptiness, we can get a *regular* tree accepted by the automaton. This then provides a finite-memory winning strategy that can be realized as a finite-state controller for the system. We note that our upper bounds apply also to the realizability problem.

4 Lower bounds

For two 2^{AP} -labeled trees (T, V) and (T', V') , and a set $Q = \{q_1, \dots, q_k\} \subseteq AP$, we say that (T, V) and (T', V') are *Q-different* if they agree on everything except possibly the labels of the propositions in Q . Formally, $T = T'$ and for all $x \in T$, we have $V(x) \setminus Q = V'(x) \setminus Q$. The logic AQCTL* extends CTL* by universal quantification on atomic propositions: if ψ is a CTL* formula and q_1, \dots, q_k are atomic propositions, then $\forall q_1, \dots, q_k \psi$ is an AQCTL* formula. The semantics of $\forall q_1, \dots, q_k \psi$ is given by $S \models \forall q_1, \dots, q_k \psi$ iff for all trees (T, V) such that (T, V) and the unwinding (T_S, V_S) of S are $\{q_1, \dots, q_k\}$ -different, $(T, V) \models \psi$. The logics AQLTL and AQCTL are defined similarly as the extensions of LTL and CTL with universal quantification on atomic propositions.

The following theorem is taken from [SVW87]. We describe here the full proof, as our lower-bound proofs are based on it. We assume familiarity with LTL [Eme90].

Theorem 4. [SVW87] *The satisfiability problem for AQLTL is EXPSPACE-hard.*

Proof. We reduce the problem of checking whether an exponential-space deterministic Turing machine T accepts an input word x . That is, given T and x , we construct an AQLTL formula $\forall q \varphi$ such that T accepts x iff $\forall q \varphi$ is satisfiable. Below we describe the formula φ informally. The formal description of φ and of the function *next* we use below are given in the full version.

Let $T = \langle \Gamma, Q, \rightarrow, q_0, F \rangle$, where Γ is the alphabet, Q is the set of states, $\rightarrow \subseteq Q \times \Gamma \times Q \times \Gamma \times \{L, R\}$ is the transition relation (we use $(q, a) \rightarrow (q', b, \Delta)$ to indicate that when T is in state q and it reads the input a in the current tape cell, it moves to state q' , writes b in the current tape cell, and its reading head moves one cell to the left or to the right, according to Δ), q_0 is the initial state, and $F \subseteq Q$ is the set of accepting states. Let $n = a \cdot |x|$, for some constant a , be such that the working tape of T has 2^n cells. We encode a configuration of T by a word $\gamma_1 \gamma_2 \dots (q, \gamma_i) \dots \gamma_{2^n}$. The meaning of such a configuration is that the j^{th}

cell of T , for $1 \leq j \leq 2^n$, is labeled γ_j , the reading head points on cell i , and T is in state q . We now encode a computation of T by a sequence of configurations.

Let $\Sigma = \Gamma \cup (Q \times \Gamma)$. We can encode letters in Σ by a set $AP(T) = \{p_1, \dots, p_m\}$ (with $m = \lceil \log |\Sigma| \rceil$) of atomic propositions. We define our formulas over the set $AP = AP(T) \cup \{b, c, d, e, q\}$ of atomic propositions. The task of the last five atoms will be explained shortly. Since T is fixed, so is Σ , and hence so is the size of AP .

Consider an infinite sequence π over 2^{AP} . For an atomic proposition $p \in AP$ and a node u in π , we use $p(u)$ to denote the truth value of p at u . That is, $p(u)$ is 1 if p holds at u and is 0 if p does not hold at u . We divide the sequence π to blocks of length n . Every such block corresponds to a single tape cell of the machine T . Consider a block u_1, \dots, u_n that corresponds to a cell ρ . We use the node u_1 to encode the content of cell ρ . Thus, the bit vector $p_1(u_1), \dots, p_m(u_1)$ encodes the letter (in $\Gamma \cup (Q \times \Gamma)$) that corresponds to cell ρ . We use the atomic proposition b to mark the beginning of the block; that is, b should hold on u_1 and fail on u_2, \dots, u_n . Recall that the letter with which cell ρ is labeled is encoded at the node u_1 of the block u_1, \dots, u_n that corresponds to ρ . Why then do we need a block of length n to encode a single letter? The block also encodes the location of the cell ρ on the tape. Since T is an exponential-space Turing machine, this location is a number between 0 and $2^n - 1$. Encoding the location eliminates the need for exponentially many X operators when we attempt to relate two successive configurations. Encoding is done by the atomic proposition c , called *counter*. Let $c(u_n), \dots, c(u_1)$ encode the location of ρ . A sequence of 2^n blocks corresponds to 2^n cells and encodes a configuration of T . The value of the counters along this sequence goes from 0 to $2^n - 1$, and then start again from 0. This can be enforced by a conjunct in φ which is only $O(n)$ -long by using the proposition d as a carry-bit. The atomic proposition e marks the last block of a configuration, that is, e holds in a node u_1 of a block u_1, \dots, u_n iff c holds on all nodes in the block.

Let $\sigma_1 \dots \sigma_{2^n}, \sigma'_1 \dots \sigma'_{2^n}$ be two successive configurations of T . For each triple $\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle$ with $1 \leq i \leq 2^n$ (taking σ_{2^n+1} to be σ'_1 and σ_0 to be the label of the last letter in the configuration before $\sigma_1 \dots \sigma_{2^n}$, or some special label when $\sigma_1 \dots \sigma_{2^n}$ is the initial configuration), we know, by the deterministic transition relation of T , what σ'_i should be. Let $next(\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle)$ denote our expectation for σ'_i .

Consistency with *next* gives us a necessary condition for a word to encode a legal computation. In addition, the computation should start with the initial configuration and end in a final configuration (with $q \in F$) — these properties can easily be enforced by φ .

The difficult part in the reduction is in guaranteeing that the sequence of configurations is indeed consistent with *next*. To enforce this, we have to relate σ_{i-1}, σ_i , and σ_{i+1} with σ'_i for any i in any two successive configurations $\sigma_1 \dots \sigma_{2^n}, \sigma'_1 \dots \sigma'_{2^n}$. One natural way to do so is by a conjunction of formulas like “whenever we meet a cell with counter $i - 1$ and the labeling of the next three cells forms the triple $\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle$, then the next time we meet a cell

with counter i , this cell is labeled $next(\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle)$. The problem is that as i can take any value from 1 to 2^n , there are exponentially many such conjuncts. This is where the universal quantification of the AQLTL comes into the picture. It enables us to relate $\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle$ with σ'_i , for all i .

To understand how this is done, consider the atomic proposition q , and assume that the following hold. **(1)** q is true at precisely two points, both are points in which a block starts, **(2)** there is exactly one point between them (possibly in exactly one of them) in which e holds (thus, the two points are in successive configurations), and **(3)** the value of the counter at the blocks starting at the two points is the same. Then, it should be true that **(4)** if the labels of the three blocks starting one block before the first q are σ_{i-1} , σ_i , and σ_{i+1} , then the block starting at the second q is labeled by $next(\sigma_{i-1}, \sigma_i, \sigma_{i+1})$.

The formula φ contains the conjunct $((\mathbf{(1)} \wedge \mathbf{(2)} \wedge \mathbf{(3)}) \rightarrow \mathbf{(4)})$. and ensures that only computations consistent with $next$ are satisfied by φ . Hence, $\forall q\varphi$ is satisfiable iff there is an accepting computation of T on x . \square

We now show that AQCTL is also strong enough to describe an exponential-space Turing machine with a formula of polynomial length. Moreover, since CTL has both universal and existential path quantification, AQCTL can describe an alternating exponential-space Turing machine, implying a 2EXPTIME lower bound for its satisfiability problem [CKS81].

Theorem 5. *The satisfiability problem for AQCTL is 2EXPTIME-hard.*

Proof. We do a reduction from the problem whether an exponential-space alternating Turing machine T accepts an input word x . That is, given T and x , we construct an AQCTL formula $\forall q\psi$ such that T accepts x iff $\forall q\psi$ is satisfiable.

Let $T = \langle \Gamma, Q_u, Q_e, \mapsto, q_0, F \rangle$, where the sets Q_u and Q_e of states are disjoint, and contain the universal and the existential states, respectively. We denote their union (the set of all states) by Q . Our model of alternation prescribes that $\mapsto \subseteq Q \times \Gamma \times Q \times \Gamma \times \{L, R\}$ has a binary branching degree. When a universal or an existential state of T branches into two states, we distinguish between the left and the right branches. Accordingly, we use $(q, a) \mapsto \langle (q_l, b_l, \Delta_l), (q_r, b_r, \Delta_r) \rangle$ to indicate that when T is in state $q \in Q_u \cup Q_e$ reading input symbol a , it branches to the left with (q_l, b_l, Δ_l) and to the right with (q_r, b_r, Δ_r) . (Note that the directions left and right here have nothing to do with the movement direction of the head; these are determined by Δ_l and Δ_r .)

For a configuration c of T , let $succ_l(c)$ and $succ_r(c)$ be the successors of c when applying to it the left and right choices in \mapsto , respectively. Given an input x , a computation tree of T on x is a tree in which each node corresponds to a configuration of T . The root of the tree corresponds to the initial configuration. A node that corresponds to a universal configuration c has two successors, corresponding to $succ_l(c)$ and $succ_r(c)$. A node that corresponds to an existential configuration c has a single successor, corresponding to either $succ_l(c)$ or $succ_r(c)$. The tree is an accepting computation tree if all its branches reach an accepting configuration.

The formula ψ will describe accepting trees. As in the linear case, we encode a configuration of T by a sequence $\gamma_1\gamma_2 \dots (q, \gamma_i) \dots \gamma_{2^n}$, and we use a block of length n to describe each letter $\sigma_i \in \Gamma \cup (Q \times \Gamma)$ in the sequence. The construction of ψ is similar to the construction described for φ in the linear case. As in the linear case, the atomic propositions c and d are used to count, b is used to mark the beginning of blocks, and e is used to mark the last letter in a configuration. The formulas which ensure the correct behaviour of c, d, b and e can be suitably prefixed with the universal path quantifier to ensure they behave correctly along all paths of the tree.

The difficult part is to check that the $succ_l$ and $succ_r$ relations are maintained. For that, we add two atomic propositions, e_E and e_U , that refine the proposition e — exactly one of them hold whenever e holds and e_E (e_U) holds iff the configuration which has just ended is existential (universal).

In addition, we use an atomic proposition l to indicate whether the nodes belong to a left or a right successor. For clarity, we denote $\neg l$ by r . Formally, ψ contains the conjunct $AG(l \rightarrow (AlUe)) \wedge AG(r \rightarrow (ArUe))$. Since a universal configuration c has both $succ_l(c)$ and $succ_r(c)$ as successors, and an existential c has only one of them, ψ also contains the conjunct $AG(e_E \rightarrow (EXl \vee EXr)) \wedge AG(e_U \rightarrow (EXl \wedge EXr))$.

We can now use universal quantification over atomic propositions in order to check consistency with $succ_l$ and $succ_r$. Note that $succ_l(c)$ and $succ_r(c)$ are uniquely defined. Thus, we can define functions, $next_l$ and $next_r$, analogous to function $next$ of the linear case. Given a sequence $\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle$ of letters in c , the function $next_l(\langle \sigma_{i-1}, \sigma_i, \sigma_{i+1} \rangle)$ returns the expectation for the i 'th letter in $succ_l(c)$. We denote this letter by σ_i^l , and similarly for $next_r$ and σ_i^r .

In the linear case, we considered assignments to q in which q holds at exactly two points in the computation. Here, we look at assignments where q holds at exactly two points in each branch. The first point is a node where a block of σ_i starts, and the second point is a node where a block of σ_i^l or σ_i^r starts (note that each assignment to q may check consistency with $succ$ along different branches)²

The formula can also ensure that in every branch with two occurrences of q , there is exactly one node between them in which e holds (thus, the two nodes are in successive configurations) and that the value of the counter at the blocks starting at the two points is the same. If the q -labelling satisfies these properties, then φ will demand that if the three blocks starting before the first q along a path are $\sigma_{i-1}, \sigma_i, \sigma_{i+1}$, then the blocks starting at the second q must be labeled by $next_l(\sigma_{i-1}, \sigma_i, \sigma_{i+1})$ (if the second q belongs to the left branch) or $next_r(\sigma_{i-1}, \sigma_i, \sigma_{i+1})$ (if it belongs to the right branch).

One can then show that ψ is satisfied only in a computation tree consistent with $succ_l$ and $succ_r$. Hence, $\forall q\psi$ is satisfiable iff there is an accepting computation tree of T on x . \square

² It is convenient to think of a satisfying tree for $\forall q\psi$ as a tree that has branching degree 1 everywhere except for nodes labeled by e_U , where the branching degree is 2. Our reduction, however, makes no assumption about such a structure.

The satisfiability problem for CTL* is exponentially harder than the one for CTL. We now show that this computational difference is preserved when we look at the extensions of these logics with universal quantification over atomic propositions.

Theorem 6. *The satisfiability problem for AQCTL* is 3EXPTIME-hard.*

Proof. We do a reduction from the problem whether a doubly-exponential-space alternating Turing machine T accepts an input word x . That is, given T and x , we construct an AQCTL* formula $\forall q\psi$ such that T accepts x iff ψ is satisfiable.

In [VS85], the satisfiability problem of CTL* is proved to be 2EXPTIME-hard by a reduction from an exponential-space alternating Turing machine. Below we explain how universal quantification can be used to “stretch” the length of the tape that a polynomial CTL* formula can describe by another exponential. As in the proof of Theorem 4, the formula in [VS85] maintains an n -bit counter, and each cell of T ’s tape corresponds to a block of length n .

In order to point on the letters σ_i and σ'_i simultaneously (that is, the letters that the atomic proposition q point on in the proof of Theorem 4), [VS85] adds to each node of the tree a branch such that nodes that belong to the original tree are labeled by some atomic proposition p , and nodes that belong to the added branches are not labeled by p . Every path in the tree has a single location where the atom p stops being true. [VS85] uses this location in order to point on σ' and in order to compare the values of the n -bit counter in the current point (where σ is located) and in the point in the computation where p stops being true.

On top of the method in [VS85], we use the universal quantification in order to maintain a 2^n -bit counter and thus count to 2^{2^n} . Typically, each bit of our 2^n -bit counter is kept in a block of length n , which maintains the index of the bit (a number between 0 to $2^n - 1$). For example, when $n = 3$, the counter looks as follows.

```

000 001 010 011 100 101 110 111  000 001 010 011 100 101 110 111 ← n-bit counter
  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  1 ← 2^n-bit counter
000 001 010 011 100 101 110 111  000 001 010 011 100 101 110 111 ...
  0  0  0  0  0  0  1  0  0  0  0  0  0  0  1  1 ...

```

To check that the 2^n -bit counter proceeds properly, we use a universally quantified proposition q and we check that if q holds at exactly two points (say, last points in a block of the n -bit counter), with the same value to the n -bit counter, and with only one block between them in which the n -bit counter has value 1^n , then the bit of the 2^n -bit counter that is maintained at the block of the second q is updated properly (we also need to relate and update carry bits, but the idea is the same). \square

Note that the number of atomic propositions in ψ in the proofs of both Theorems 5 and 6 is fixed. Note also that if ψ is satisfiable, then it is also satisfied in a tree of a fixed branching degree (a careful analysis can show that for CTL the sufficient branching degree is 2, and for CTL* it is 3).

The logic EAQCTL* extends AQCTL* by adding existential quantification on atomic propositions: if $\forall q_1, \dots, q_k \psi$ is an AQCTL* formula and p_1, \dots, p_m are

atomic propositions, then $\exists p_1, \dots, p_m \forall q_1, \dots, q_k \psi$ is an EAQCTL* formula. The semantics is given by $S \models \exists p_1, \dots, p_m \forall q_1, \dots, q_k \psi$ iff there is a tree (T, V) such that (T_S, V_S) and (T, V) are $\{p_1, \dots, p_m\}$ -different and $(T, V) \models \forall q_1, \dots, q_k \psi$. The logic EAQCTL is the subset of EAQCTL* corresponding to CTL. It is not difficult now to show that the following Theorem follows from Theorems 5 and 6.

Theorem 7. *The model-checking problems for EAQCTL and EAQCTL* are 2EXPTIME-hard and 3EXPTIME-hard in the size of the specification, respectively.* \square

Intuitively, the model-checking problem for EAQCTL* asks whether we can find an assignment to the propositions that are existentially quantified so that no matter how we assign values to the propositions that are universally quantified, the formula is satisfied. Recall that in the control problem we ask a similar question, namely whether there a strategy for the system so that no matter which strategy the environment uses, the formula is satisfied. In this spirit, it is not difficult to make the relation between existential and universal quantification over atomic propositions and supervisory control formal. The technique is similar to the relation between existential quantification and the module-checking problem, as described in [KV96]. Consequently, we can show :

Theorem 8. *Given an EAQCTL* formula $\exists p_1, \dots, p_m \forall q_1, \dots, q_k \psi$, and a structure S , there is a plant P and a CTL* formula ψ' such that $|P| = O((1+k+m) \cdot |S|)$, $|\psi'| = O(|S| + |\psi|)$, and $S \models \exists p_1, \dots, p_m \forall q_1, \dots, q_k \psi$ iff controllable(P, ψ').* \square

Since the number of atomic propositions in the formulas used in the reductions in Theorems 5 and 6 is fixed, and since in the case P is fixed the size of ψ' in Theorem 8 is $O(|\psi|)$, we can conclude with the following.

Theorem 9. *The control problems for CTL and CTL* are 2EXPTIME-hard and 3EXPTIME-hard in the size of the specification, respectively.* \square

We note that these lower bounds apply also to the realizability problem.

References

- [ALW89] M. Abadi, L. Lamport, and P. Wolper. Realizable and unrealizable concurrent program specifications. In *Proc. 16th ICALP*, LNCS 372, 1989.
- [Ant95] M. Antoniotti. *Synthesis and verification of discrete controllers for robotics and manufacturing devices with temporal logic and the Control-D system*. PhD thesis, New York University, New York, 1995.
- [BL69] J.R. Büchi and L.H.G. Landweber. Solving sequential conditions by finite-state strategies. *Trans. AMS*, 138:295–311, 1969.
- [Chu63] A. Church. Logic, arithmetics, and automata. In *Proc. International Congress of Mathematicians, 1962*, pages 23–35. institut Mittag-Leffler, 1963.
- [CKS81] A.K. Chandra, D.C. Kozen, and L.J. Stockmeyer. Alternation. *Journal of the Association for Computing Machinery*, 28(1):114–133, January 1981.

- [DTV99] M. Daniele, P. Traverso, and M.Y. Vardi. Strong cyclic planning revisited. In *5th European Conference on Planning*, pages 34–46, 1999.
- [EJ88] E.A. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *Proc. 29th FOCS*, pages 328–337, White Plains, October 1988.
- [Eme90] E.A. Emerson. Temporal and modal logic. *Handbook of Theoretical Computer Science*, pages 997–1072, 1990.
- [HP85] D. Harel and A. Pnueli. On the development of reactive systems. In *Logics and Models of Concurrent Systems*, volume F-13 of *NATO Advanced Summer Institutes*, pages 477–498. Springer-Verlag, 1985.
- [KG95] R. Kumar and V.K. Garg. *Modeling and control of logical discrete event systems*. Kluwer Academic Publishers, 1995.
- [KS95] R. Kumar and M.A. Shayman. Supervisory control of nondeterministic systems under partial observation and decentralization. *SIAM Journal of Control and Optimization*, 1995.
- [KV96] O. Kupferman and M.Y. Vardi. Module checking. In *Proc. 8th CAV*, LNCS 1102, pages 75–86, 1996.
- [KV97] O. Kupferman and M.Y. Vardi. Module checking revisited. In *Proc. 9th CAV*, LNCS 1254, pages 36–47, 1997.
- [KV99a] O. Kupferman and M.Y. Vardi. Church’s problem revisited. *The Bulletin of Symbolic Logic*, 5(2):245 – 263, June 1999.
- [KV99b] O. Kupferman and M.Y. Vardi. Robust satisfaction. In *Proc. 10th CONCUR*, LNCS 1664, pages 383–398, 1999.
- [KV00] O. Kupferman and M.Y. Vardi. μ -calculus synthesis. In *Proc. 25th MFCS*.
- [Lam80] L. Lamport. Sometimes is sometimes “not never” - on the temporal logic of programs. In *Proc. 7th POPL* pages 174–185, January 1980.
- [MT98] P. Madhusudan and P. S. Thiagarajan. Controllers for discrete event systems via morphisms. In *Proc 9th CONCUR*, LNCS 146, pages 18–33, 1998.
- [MT00] P. Madhusudan and P. S. Thiagarajan. Branching time controllers for discrete event systems. *To appear in CONCUR’98 Special Issue, Theoretical Computer Science*, 2000.
- [PR89] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.
- [Rab69] M.O. Rabin. Decidability of second order theories and automata on infinite trees. *Transaction of the AMS*, 141:1–35, 1969.
- [Rab72] M.O. Rabin. Automata on infinite objects and Church’s problem. *Amer. Mathematical Society*, 1972.
- [Saf88] S. Safra. On the complexity of ω -automata. In *Proc. 29th IEEE Symposium on Foundations of Computer Science*, pages 319–327, October 1988.
- [SVW87] A.P. Sistla, M.Y. Vardi, and P. Wolper. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49:217–237, 1987.
- [Tho97] W. Thomas. Languages, automata, and logic. *Handbook of Formal Language Theory*, III:389–455, 1997.
- [Var95] M.Y. Vardi. An automata-theoretic approach to fair realizability and synthesis. In *Proc. 7th CAV*, LNCS 939, pages 267–292, 1995.
- [VS85] M.Y. Vardi and L. Stockmeyer. Improved upper and lower bounds for modal logics of programs. In *Proc 17th STOC*, pages 240–251, 1985.
- [VW86] M.Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logics of programs. *Journal of Computer and System Science*, 32(2):182–221, 1986.