



HAL
open science

Temporal Synthesis for Bounded Systems and Environments

Orna Kupferman, Yoad Lustig, Moshe Vardi, Mihalis Yannakakis

► **To cite this version:**

Orna Kupferman, Yoad Lustig, Moshe Vardi, Mihalis Yannakakis. Temporal Synthesis for Bounded Systems and Environments. Symposium on Theoretical Aspects of Computer Science (STACS2011), Mar 2011, Dortmund, Germany. pp.615-626. hal-00573612

HAL Id: hal-00573612

<https://hal.science/hal-00573612>

Submitted on 5 Mar 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Temporal Synthesis for Bounded Systems and Environments*

Orna Kupferman¹, Yoad Lustig², Moshe Y. Vardi², and Mihalis Yannakakis³

- 1 Hebrew University
- 2 Rice University
- 3 Columbia University

Abstract

Temporal synthesis is the automated construction of a system from its temporal specification. It is by now realized that requiring the synthesized system to satisfy the specifications against all possible environments may be too demanding, and, dually, allowing all systems may be not demanding enough. In this work we study *bounded temporal synthesis*, in which bounds on the sizes of the state space of the system and the environment are additional parameters to the synthesis problem. This study is motivated by the fact that such bounds may indeed change the answer to the synthesis problem, as well as the theoretical and computational aspects of the synthesis problem. In particular, a finer analysis of synthesis, which takes system and environment sizes into account, yields deeper insight into the quantificational structure of the synthesis problem and the relationship between strong synthesis – there exists a system such that for all environments, the specification holds, and weak synthesis – for all environments there exists a system such that the specification holds.

We first show that unlike the unbounded setting, where determinacy of regular games implies that strong and weak synthesis coincide, these notions do not coincide in the bounded setting. We then turn to study the complexity of deciding strong and weak synthesis. We show that bounding the size of the system or both the system and the environment, turns the synthesis problem into a search problem, and one cannot expect to do better than brute-force search. In particular, the synthesis problem for bounded systems and environment is Σ_2^P -complete (in terms of the bounds, for a specification given by a deterministic automaton). We also show that while bounding the environment may lead to the synthesis of specifications that are otherwise unrealizable, such relaxation of the problem comes at a high price from a complexity-theoretic point of view.

Digital Object Identifier 10.4230/LIPIcs.STACS.2011.615

1 Introduction

Temporal synthesis is the automated construction of a system from its temporal specification. The basic idea is simple and appealing: instead of developing a system and verifying that it satisfies its specification, we would like to have an automated procedure that, given a specification, constructs a system that is correct by construction. The first formulation of synthesis goes back to Church [4]; the modern approach was initiated by Pnueli and Rosner, who introduced LTL (linear temporal logic) synthesis [22]. The *LTL synthesis problem* receives as input a specification given by means of an LTL formula and outputs a reactive system modeled by a finite-state transducer satisfying the given specification — if such exists.

In the specification to the system, it is important to distinguish between output signals, controlled by the system, and input signals, controlled by the environment. A system should satisfy its specification against all possible environments. Therefore, the quantification structure on input and output

* Supported in part by NSF grants CCF-0613889, CCF-0728882, CCF-0728736, CCF-1017955, and CNS 1049862, by BSF grant 9800096, and by gift from Intel. Part of this work was done while Moshe Y. Vardi was on sabbatical at the Hebrew University in Jerusalem, supported by a Forchheimer Visiting Professorship.



signals is different. Output signals are existentially quantified while input signals are universally quantified [22]. It is by now realized that requiring the synthesized system to satisfy the specification against all possible environments is often too demanding. Dually, allowing all possible systems is perhaps not demanding enough. This issue is traditionally approached by adding assumptions on the system and/or the environment, which are modeled as part of the specification (c.f., [3]).¹

In this work we study *bounded temporal synthesis*, in which assumptions on the system and its environment are given by means of bounds on the sizes of their state space. Thus, in addition to a specification ψ , the input to the bounded synthesis problem contains two parameters $n, m \geq 1$, and a specification ψ is (n, m) -*realizable*, if there is a transducer \mathcal{T} with n states such that for all transducers \mathcal{T}' with m states, the computation $\mathcal{T} \parallel \mathcal{T}'$ – generated by the interaction of \mathcal{T} with \mathcal{T}' , satisfies ψ . Note that traditional synthesis corresponds to the case $n = m = \infty$.² Also note that by setting only one of n or m to ∞ , we can consider a setting in which only one of the components is bounded. In particular, [25] studies the setting in which only the system is bounded. We note that the need to bound the environment is of interest in several other paradigms in computer science. For example, in cryptography, one studies the security of a given cryptosystem with respect to attackers with bounded (typically polynomial) computational power [18], and in the analysis of on-line algorithms one sometimes care for the competitive ratio of a given on-line algorithm with respect to requests issued by a bounded adversary [1]. Even closer to the work here is the study of bounded rationality in games, where bounds are placed on the power of the players. In particular, having the players be automata with a bounded number of states is a natural way of doing this. As shown in [20], such bounds affect the kind of equilibria one gets, and gives in fact a way of getting around some of the problematic cases of equilibria, (e.g., in the Prisoner’s Dilemma [24]).

It is not hard to see that bounding the size of the system or its environment may indeed change the answer to the synthesis problem. Clearly, already in a setting with no interaction, bounding the size of a system may prevent it from satisfying some specifications. In the presence of interaction, bounding the size of the environment both restricts the possible behaviors of the environment and enables the system to “learn” the environment. For example, knowing that the environment has a single state implies that the input to the system is fixed, thus a specification like “if p holds in the present, then p holds always”, for an input signal p , is realizable against environments with a single state, while it is clearly not realizable in general.

Traditional temporal synthesis is *determined*, in the sense that for every specification ψ , either there is a system that realizes ψ , or there is an environment that realizes $\neg\psi$. Note that not having a system that realizes ψ only means that for every system \mathcal{T} , there is an environment \mathcal{T}' such that the computation $\mathcal{T} \parallel \mathcal{T}'$ does not satisfy ψ . This by itself does not imply that there is an environment \mathcal{T}' such that for all systems \mathcal{T} , the computation $\mathcal{T} \parallel \mathcal{T}'$ satisfies $\neg\psi$. However, by determinacy of Borel games [17], we know that the lack of \mathcal{T} that realizes ψ does imply the existence of \mathcal{T}' that realizes $\neg\psi$. We show that determinacy no longer holds in the bounded setting. In particular, for every $k \geq 1$ there is a specification ψ_k such that ψ_k cannot be realized against environments of size k , nor can $\neg\psi_k$ be realized by an environment of size k .

The observation about determinacy, which uses the theory of *checking sequences* for transducers [16], yields deeper insight into the quantificational structure of the synthesis problem and the relationship between two possible definitions of synthesis in the bounded setting: *strong synthesis*, where there is a system \mathcal{T} such that for all environments \mathcal{T}' , the computation $\mathcal{T} \parallel \mathcal{T}'$ satisfies the

¹ A different, more conceptual, way to restrict the range of environments with respect to which the system has to satisfy the specification is to assume that the environment has objectives of its own, and is therefore *rational*, rather than hostile [9].

² In fact, by the small model property, already to the case n and m are doubly exponential in the length of $|\psi|$ [7, 22].

specification, and *weak synthesis*, where for all environments \mathcal{T}' there exists a system \mathcal{T} such that the computation $\mathcal{T} \parallel \mathcal{T}'$ satisfies the specification. In contrast, in the unbounded setting, strong and weak synthesis coincide.

We study the complexity of strong and weak synthesis. Recall that LTL synthesis is 2EXPTIME-complete [22]. This high complexity, along with technical challenges in the implementation of synthesis algorithms [12, 30], have led researchers to develop synthesis algorithms for fragments of LTL or alternative specification formalisms [15, 21]. A different approach for coping with the complexity of LTL synthesis, tightly related to our work here, is to restrict attention to systems of a bounded size [25]. As argued in [25], bounding the size of the system enables a reduction of the synthesis problem to the SAT problem, and also leads to the decidability of synthesis of distributed systems. For the bounded setting, researchers were also able to come up with a symbolic implementation [6, 8].³

Recall that the bounded synthesis problem has three parameters: ψ , n , and m . We would like to study the complexity in terms of each of n and m . One standard way to analyze the complexity in terms of one parameter is to fix the other parameters. This standard way does not work in our setting, as fixing ψ implies fixing also n and m . Indeed, by [7, 22], if ψ is realizable, then it is also realizable by a transducer with doubly-exponentially many states. Also, by determinacy, ψ is not realizable if the environment can realize $\neg\psi$ by a transducer that is doubly-exponential in the length of ψ . Accordingly, we have to neutralize the dominance of ψ in the complexity analysis in a different way. We do so by assuming that the temporal specification is given, instead of as an LTL formula, as a deterministic Büchi automaton. For such specifications, the unbounded realizability problem amounts to checking the nonemptiness of a deterministic Büchi automaton, and can therefore be solved in quadratic time [28]. In Section 4, we justify the choice of deterministic Büchi automata further.

We first show that bounding the size of the system enables an easy reduction from the synthesis problem to the model-checking problem. At the same time, one cannot expect to do better than brute-force search in synthesizing bounded-size systems. Formally, we show that deciding whether a specification \mathcal{A} given by means of a deterministic Büchi automaton is realizable by a system with n states is NP-complete, for n given in unary. The proof of NP-hardness is technically easy and is similar to known NP-hardness proofs in the context of formal methods [5, 13]. Still, it justifies the reduction to SAT given in [25] without a lower bound, and it sets the stage to the much more challenging lower bound, for the case both the system and the environment are bounded: we show that deciding whether a specification \mathcal{A} is realizable by a system with n states against all environment with m states is Σ_2^P -complete, for n and m given in unary. Thus, brute-force search is the best we can do also here, and one cannot expect to do better than model checking the interaction of all systems with n states with all environments with m states.

We also show that while bounding the environment may indeed lead to the synthesis of specifications that are otherwise unrealizable, such relaxation of the problem comes at a high price from a complexity-theoretic point of view. As pointed above, synthesis with respect to specifications given by deterministic Büchi automata is quadratic. Adding a bound on the size of the environment seems to add a cost that is doubly exponential in that size. In fact, we show that even model checking against bounded environments is apparently harder than standard model checking. Finally, we formalize the intuition of the system being able to learn the bounded environment by showing that for absolute liveness properties, which are insensitive to additions of prefixes [26], weak and strong realizability

³ As mentioned above, [25] also studies synthesis of bounded systems. The contributions here and in [25] are, however, different. In [25], the focus is on practical algorithm for the setting of a bounded system. Here, the focus is on the theoretical aspects of the problem and its complexity, and rather than studying bounded systems, we consider bounds on the system, the environment, and both.

coincide, and the complexity of the problem is only exponential in the bound for the environment.

Due to the lack of space, some proofs are omitted, and can be found in the full version in the authors' home pages.

2 Preliminaries

Consider finite sets I and O of input and output signals, respectively. We model finite-state reactive systems with inputs in I and outputs in O by *transducers* (I/O -transducers, when I and O are not clear from the context). A transducer is a finite graph with a designated start state, where the edges are labeled by letters in 2^I and the states are labeled by letters in 2^O . Formally, a transducer is a tuple $\mathcal{T} = \langle I, O, S, s_{in}, \eta, L \rangle$, where I and O are the sets of input and output signals, S is a finite set of states, $s_{in} \in S$ is an initial state, $\eta : S \times 2^I \rightarrow S$ is a deterministic transition function, and $L : S \rightarrow 2^O$ is a labeling function. We extend η to words in $(2^I)^*$ in the straightforward way. Thus, $\eta : (2^I)^* \rightarrow S$ is such that $\eta(\varepsilon) = s_{in}$, and for $x \in (2^I)^*$ and $i \in 2^I$, we have $\eta(x \cdot i) = \eta(\eta(x), i)$. Each transducer \mathcal{T} induces a *strategy* $f_{\mathcal{T}} : (2^I)^* \rightarrow 2^O$ where for all $w \in (2^I)^*$, we have $f_{\mathcal{T}}(w) = \tau(L(w))$. Thus, $f_{\mathcal{T}}(w)$ is the letter that \mathcal{T} outputs after reading the sequence w of input letters. A transducer with at most k states is referred to as a *k-transducer*.

Consider an infinite sequence $w = i_0, i_1, i_2, i_3, \dots \in (2^I)^\omega$ of input letters. The computation of \mathcal{T} on w , denoted $\mathcal{T}(w)$, is $\rho = (o_0 \cup i_0), (o_1 \cup i_1), (o_2 \cup i_2), \dots \in (2^{I \cup O})^\omega$ such that for all $j \geq 0$, we have $o_j = f_{\mathcal{T}}(i_0 \cdot i_1 \cdots i_{j-1})$. Note that, in particular, $o_0 = f_{\mathcal{T}}(\varepsilon)$. Thus, the mode of interaction we assume is that the transducer initiates the interaction with the environment by outputting $f_{\mathcal{T}}(\varepsilon)$, the environment then responds with i_0 (making $f_{\mathcal{T}}(\varepsilon) \cup i_0$ the set of signals that are valid in the first time unit), then the transducer responds with $f_{\mathcal{T}}(i_0)$, the environment with i_1 , and so on. In order to emphasize the fact that the transducer moves first, we sometimes refer to ρ as a *1-computation* of \mathcal{T} . Also, we sometimes refer to ρ as w/y , for $y = o_0, o_1, o_2, \dots$.

One could also consider a dual type of interaction, in which the environment moves first. Then, a *2-computation* of \mathcal{T} is $\rho = (o_0 \cup i_0), (o_1 \cup i_1), (o_2 \cup i_2), \dots$ where for all $j \geq 0$, we have $o_j = f_{\mathcal{T}}(i_0 \cdot i_1 \cdots i_j)$. In particular, $o_0 = f_{\mathcal{T}}(i_0)$. Note that when two transducers interact with each other, one of them initiates the interaction and moves first, thus its computations are 1-computations, and the second moves second, and its computations are 2-computations. We say that $\rho \in (2^{I \cup O})^\omega$ is a computation of \mathcal{T} if there is $w \in (2^I)^\omega$ such that $\rho = \mathcal{T}(w)$. Finally, we sometimes refer also to finite sequences of input letters and the finite computations of \mathcal{T} on them.

We specify on-going behaviors of I/O -transducers by means of LTL formulas over the set $I \cup O$ of atomic propositions, or automata on infinite words over the alphabet $2^{I \cup O}$. For a specification ψ over $I \cup O$ and a "who moves first" flag $b \in \{1, 2\}$, we use $real_{I,O,b}(\psi)$ to indicate that there is an I/O transducer \mathcal{T} such that all the b -computations of \mathcal{T} satisfy ψ . Given a specification ψ over the sets I and O of input and output signals, the *realizability problem* for ψ is to decide whether $real_{I,O,1}(\psi)$ [22]. For $b \in \{1, 2\}$ let \tilde{b} dualize b (that is, $\tilde{b} = 3 - b$). By determinacy of Borel games [17], we have the following (see also [10]).

► **Theorem 1.** *For every specification ψ , precisely one of $real_{I,O,b}(\psi)$ or $real_{O,I,\tilde{b}}(\neg\psi)$ holds.*

3 Strong and Weak realizability

The traditional definition of realizability requires ψ to be satisfied in all the computations of \mathcal{T} , ignoring the ability of the environment to feasibly generate the sequences of input letters that induce these computations. In this work, we are interested in realizability with components of a bounded size. In order to define this setting, let us first formalize the interaction between two transducers. For an I/O -transducer \mathcal{T} that induces a strategy $f_{\mathcal{T}} : (2^I)^* \rightarrow 2^O$ and an O/I -transducer \mathcal{T}' that induces a

strategy $f_{\mathcal{T}'} : (2^O)^* \rightarrow 2^I$, the single 1-computation of $\mathcal{T} \parallel \mathcal{T}'$ is $\rho = (o_0 \cup i_0), (o_1 \cup i_1), (o_2 \cup i_2), \dots$ where for all $j \geq 0$, we have $o_j = f_{\mathcal{T}}(i_0 \cdot i_1 \cdots i_{j-1})$ and $i_j = f_{\mathcal{T}'}(o_0 \cdot o_1 \cdots o_j)$. The single 2-computation of $\mathcal{T} \parallel \mathcal{T}'$ is $\rho = (o_0 \cup i_0), (o_1 \cup i_1), (o_2 \cup i_2), \dots$ where for all $j \geq 0$, we have $i_j = f_{\mathcal{T}'}(o_0 \cdot o_1 \cdots o_{j-1})$ and $o_j = f_{\mathcal{T}}(i_0 \cdot i_1 \cdots i_j)$.

Recall that a specification ψ is realizable if there is an I/O -transducer \mathcal{T} such that all the 1-computations of \mathcal{T} satisfy ψ . It is known that if a transducer \mathcal{T} does have a computation that violates ψ , then \mathcal{T} also has a computation that violates ψ and is induced by an input sequence that is generated by a transducer [2]. Accordingly, ψ is realizable iff there is an I/O -transducer \mathcal{T} such that for all O/I transducers \mathcal{T}' , the 1-computation of $\mathcal{T} \parallel \mathcal{T}'$ satisfies ψ . In Definition 2 below, we call this “strong realizability” and introduce also a weaker type of realizability.

► **Definition 2.** [*Strong and Weak Realizability*] Consider a specification ψ over $I \cup O$, a first-move flag $b \in \{1, 2\}$, and $m, n \in \mathcal{N} \cup \{\infty\}$.

- We say that ψ is *strongly* (I, O, b) -realizable with respect to systems with n states and environments with m states, denoted $s_real_{I,O,b}(\psi, n, m)$, if there is an I/O -transducer \mathcal{T} with at most n states such that for every O/I -transducer \mathcal{T}' with at most m states, the b -computation of $\mathcal{T} \parallel \mathcal{T}'$ satisfies ψ .
- We say that ψ is *weakly* (I, O, b) -realizable with respect to systems with n states and environments with m states, denoted $w_real_{I,O,b}(\psi, n, m)$, if for every O/I -transducer \mathcal{T}' with at most m states, there is an I/O -transducer \mathcal{T} with at most n states such that the b -computation of $\mathcal{T} \parallel \mathcal{T}'$ satisfies ψ .

Strong realizability means that the system can defeat all environment’s strategies under consideration. Weak realizability means that the environment does not have a strategy to defeat all the strategies of the system. When strong realizability is impossible, a designer may settle for a weak one. Also, weak realizability of $\neg\psi$ by the environment explains why ψ is not realizable by the system. Formally, we have the following.

► **Lemma 3.** $w_real_{I,O,b}(\psi, n, m)$ iff not $s_real_{O,I,\bar{b}}(\neg\psi, m, n)$.

As discussed above, $real_{I,O,b}(\psi)$ iff $s_real_{I,O,b}(\psi, \infty, \infty)$. In fact, as we state below, weak and strong realizability coincide in the unbounded setting.

► **Theorem 4.** For every specification ψ , we have that $w_real_{I,O,b}(\psi, \infty, \infty)$ iff $s_real_{I,O,b}(\psi, \infty, \infty)$.

Proof. By definition, $s_real_{I,O,b}(\psi, \infty, \infty)$ implies $w_real_{I,O,b}(\psi, \infty, \infty)$. For the other direction, assume that $s_real_{I,O,b}(\psi, \infty, \infty)$ does not hold. Then, by Theorem 1, we have that $s_real_{O,I,\bar{b}}(\neg\psi, \infty, \infty)$. Thus, by Lemma 3, we have that $w_real_{I,O,b}(\psi, \infty, \infty)$ does not hold, and we are done. ◀

► **Example 5.** Let $I = \{p\}$, $O = \{q\}$, and $\psi = G(q \leftrightarrow Xp)$. Note that the specification requires the next value of the input signal p to depend on the current value of the output signal q . Since the system has no control on the input signals, we have that not $s_real_{I,O,1}(\psi, \infty, \infty)$, and in fact even not $s_real_{I,O,1}(\psi, \infty, 1)$. To formally prove this, we use Lemma 3 and show that $w_real_{O,I,2}(\neg\psi, 1, \infty)$. Note that $\neg\psi = F((q \wedge X\neg p) \vee (\neg q \wedge Xp))$. Now, for every I/O -transducer \mathcal{T} , if \mathcal{T} outputs q in its initial state, the O/I -1-transducer \mathcal{T}' that always outputs $\neg p$ is such that the unique 2-computation of $\mathcal{T}' \parallel \mathcal{T}$ satisfies $\neg\psi$. Similarly, if \mathcal{T} outputs $\neg q$ in its initial state, the corresponding O/I -1-transducer \mathcal{T}' always outputs p .

On the other hand, $w_real_{I,O,1}(\psi, \infty, 1)$. Indeed, there are two possible O/I -1-transducers: \mathcal{T}'_1 that always outputs p , and \mathcal{T}'_2 that always outputs $\neg p$. For \mathcal{T}'_1 , the I/O -transducer \mathcal{T} that always outputs q is such that the unique 1-computation of $\mathcal{T} \parallel \mathcal{T}'_1$ is $\{p, q\}^\omega$. For \mathcal{T}'_2 , the I/O -transducer \mathcal{T} that always outputs $\neg q$ is such that the unique 1-computation of $\mathcal{T} \parallel \mathcal{T}'_2$ is \emptyset^ω . Since both satisfy ψ , we are done.

Example 5 shows that strong and weak realizability do not coincide in the bounded setting. We now generalize the example to all bounds, and show that it implies that the games that correspond to bounded synthesis are not determined. We first need some notations.

For two I/O -transducers \mathcal{T}_1 and \mathcal{T}_2 , we say that \mathcal{T}_1 and \mathcal{T}_2 are *equivalent*, denoted $\mathcal{T}_1 \equiv \mathcal{T}_2$, if $\mathcal{T}_1(w) = \mathcal{T}_2(w)$ for every word $w \in (2^I)^*$; otherwise, \mathcal{T}_1 and \mathcal{T}_2 are inequivalent. If \mathcal{T} is a transducer and s an state of \mathcal{T} , then denote by \mathcal{T}/s the transducer that is the same as \mathcal{T} except that it has s as the start state. We say that \mathcal{T} is *minimized* if there is no transducer \mathcal{T}' such that $\mathcal{T} \equiv \mathcal{T}'$ and \mathcal{T}' has strictly fewer states than \mathcal{T} . We say that two transducers $\mathcal{T}_1, \mathcal{T}_2$ are *structurally equivalent*, denoted $\mathcal{T}_1 \sim \mathcal{T}_2$, if (i) for every state s_1 of \mathcal{T}_1 there is a state s_2 of \mathcal{T}_2 such that $\mathcal{T}_1/s_1 \equiv \mathcal{T}_2/s_2$, and (ii) conversely, for every state s_2 of \mathcal{T}_2 there is a state s_1 of \mathcal{T}_1 such that $\mathcal{T}_1/s_1 \equiv \mathcal{T}_2/s_2$. Note that two minimized transducers are structurally equivalent iff they are isomorphic (but their start states do not need to map to each other in the isomorphism) [16].

Consider an I/O -transducer \mathcal{T} with k states. A *checking sequence* for \mathcal{T} is a word $w \in (2^I)^*$ such that for all transducers \mathcal{T}' with at most k states, if $\mathcal{T}'(w) = \mathcal{T}(w)$ then $\mathcal{T}' \sim \mathcal{T}$. A transducer is *strongly connected* if every state can reach every other state. It is known that every strongly connected k -state transducer \mathcal{T} has a checking sequence (of length at most exponential in k), and furthermore, if \mathcal{T} is minimized, then it has one of length polynomial in k ; see [16] for background and an overview on checking and other test sequences of transducers.

For a word $w \in (2^I)^*$, let $\theta(w)$ be an LTL formula over I that holds in exactly all computations in $(2^{I \cup O})^\omega$ whose prefix agrees with w on the signals in I . Thus, if $w = i_0, i_1, \dots, i_l$, then $\theta(w) = \bigwedge_{0 \leq j \leq l} X^j((\bigwedge_{x \in i_j} x) \wedge (\bigwedge_{x \notin i_j} \neg x))$. We define $\theta(y)$ similarly, for $y \in (2^O)^*$.

We now use checking sequences in order to show that bounding the sizes of the components, strong and weak synthesis no longer coincide, and determinacy fails.

► **Theorem 6.** *For every $k \geq 1$ there is an LTL formula ψ_k such that $w_real_{I,O,1}(\psi_k, \infty, k)$ but not $s_real_{I,O,1}(\psi_k, \infty, k)$; equivalently, neither $s_real_{I,O,1}(\psi_k, \infty, k)$ nor $s_real_{I,O,2}(\neg\psi_k, k, \infty)$.*

Proof. We give first a high level idea of the approach. We construct an LTL formula ψ_k that is weakly realizable, and is “almost strongly realizable”, in the sense that the system can succeed in ensuring the formula by using one of two strategies (transducers): the first strategy works for all environment k -transducers except for those that are isomorphic to a specific one \mathcal{T}'_1 , and the second strategy works for all environment k -transducers except for those isomorphic to a second one \mathcal{T}'_2 . However, no system strategy works for both $\mathcal{T}'_1, \mathcal{T}'_2$, and thus if the system does not know the environment transducer, then it cannot ensure ψ_k , i.e. ψ_k is not strongly realizable. In our construction, we pick suitable transducers \mathcal{T}'_1 and \mathcal{T}'_2 and construct from them a formula ψ_k that has the desired properties. Note that the presence of a bound on the size of the transducers is critical for this to be possible (and is obviously essential for the construction), i.e. there is no such formula ψ_∞ , since weak and strong realizability coincide in the unbounded case.

We proceed now with the details of the construction. For an O/I transducer \mathcal{T}'_i with k states that is minimized and strongly connected, let $y_i \in (2^O)^*$ be a checking sequence for \mathcal{T}'_i , let y_i/w_i be the 2-computation of \mathcal{T}'_i on y_i , and let $\varphi_i = \theta(y_i) \wedge \neg\theta(w_i)$. Note that $s_real_{I,O,1}(\varphi_i, \infty, k-1)$. To see why, consider an I/O -transducer \mathcal{T} that ignores the input and outputs y_i . Since y_i is a checking sequence for \mathcal{T}'_i , and \mathcal{T}'_i has k states and is minimized, the sequence y_i/w_i cannot be generated in an interaction with an O/I -transducer with at most k states that is not isomorphic with \mathcal{T}'_i . Hence, for every transducer \mathcal{T}' with $k-1$ states, the 1-computation of $\mathcal{T} \parallel \mathcal{T}'$ satisfies both $\theta(y_i)$ and $\neg\theta(w_i)$. On the other hand, note also that $\neg s_real_{I,O,1}(\varphi_i, \infty, k)$ and even $\neg w_real_{I,O,1}(\varphi_i, \infty, k)$. Indeed, the O/I -transducer \mathcal{T}'_i is such that for all I/O -transducers \mathcal{T} , if the 1-computation of $\mathcal{T} \parallel \mathcal{T}'_i$ satisfies $\theta(y_i)$, then it also satisfies $\theta(w_i)$.

We would like to use φ_i in order to construct a specification ψ_k as required. Let \mathcal{T}'_1 and \mathcal{T}'_2 be two nonisomorphic minimized strongly connected O/I transducers with k states; thus $\mathcal{T}'_1, \mathcal{T}'_2$ are not

structurally equivalent (since they are minimized). Let o_1 and o_2 be two different letters in 2^O , let s_1 be the state that \mathcal{T}'_1 reaches when it reads o_1 from its initial state, and let s_2 be the state that \mathcal{T}'_2 reaches when it reads o_2 from its initial state. Finally, let y_1 be a checking sequence for the transducer \mathcal{T}'_1/s_1 (i.e., \mathcal{T}'_1 with initial state s_1) with corresponding output w_1 , and let y_2 be a checking sequence for \mathcal{T}'_2/s_2 with corresponding output w_2 . We define $\psi_k = (o_1 \wedge X(\theta(y_1) \wedge \neg\theta(w_1))) \vee (o_2 \wedge X(\theta(y_2) \wedge \neg\theta(w_2)))$. In the full version we prove that $w_real_{I,O,1}(\psi_k, \infty, k)$ but $\neg s_real_{I,O,1}(\psi_k, \infty, k)$. ◀

4 Bounded Systems

We start by studying the strong-realizability problem for bounded systems. This problem is first studied in [25]. The motivation there is to deal with the computational complexity of temporal synthesis: the complexity of deciding whether $real_{I,O,b}(\psi)$ holds is known to be 2EXPTIME-complete [22]. Furthermore, for each $n > 0$ one can construct an LTL formula φ_n of length $O(n)$ such that the smallest transducer \mathcal{T} realizing φ_n has at least 2^{2^n} states [23]. In practice, therefore, we may want to bound the size of the systems under consideration, motivating us to ask whether $s_real_{I,O,b}(\psi, k, \infty)$ holds, instead of asking whether $real_{I,O,b}(\psi)$ holds.

Note that the strong realizability problem for bounded systems has two parameters: a specification ψ and a bound k . Here we would like to understand the complexity with respect to k . Thus, we would like to “neutralize” here the effect of ψ on the complexity of checking whether $s_real_{I,O,b}(\psi, k, \infty)$ holds. Fixing ψ would not help us, as, by the small model property, it induces a fixed bound on the size of a realizing transducer for ψ , if one exists. The complexity of the unbounded synthesis problem for LTL follows from the need to translate the LTL specification to a deterministic automaton on infinite words. Such a translation involves a doubly exponential blow-up [14], and is the source of the complexity of the synthesis problem. Indeed, for specifications given by means of nondeterministic or deterministic Büchi automata, the synthesis problem is complete in EXPTIME and PTIME, respectively [28].

Accordingly, we neutralize the dominance of the specification ψ by considering, instead of a specification given by an LTL formula, a specification given by a deterministic Büchi automaton \mathcal{A} over the alphabet $2^{I \cup O}$. We denote classes of automata by acronyms in $\{D, N\} \times \{F, B\} \times \{W, T\}$. The first letter stands for the branching mode of the automaton (deterministic or nondeterministic); the second letter stands for the acceptance-condition type (finite words or Büchi); the third letter stands for the object over which the automaton runs (words or trees). For example, NFW stands for nondeterministic automata on finite words, and DBT stands for deterministic Büchi tree automata. We note that while DBWs are less expressive than NBWs, we still work with DBW rather than, say, deterministic parity word automata. The reason is that the nonemptiness problem for deterministic parity tree automata is not known to be polynomial, and we want to emphasize the fact that the hardness results we are going to prove are not due to the automaton and are due to k ; the upper bounds we are going to present for DBWs are valid also for specifications given by a deterministic parity word automata.

Working with specifications that are automata, it is convenient to talk about alphabets Σ_I and Σ_O , where the transducers are Σ_I/Σ_O -transducers, in which the transitions are labeled by letters in Σ_I and the states are labeled by letters in Σ_O (or dually, are Σ_O/Σ_I -transducers). The alphabet of the specification DBW is then $\Sigma_O \times \Sigma_I$ when we specify 1-computations, and is $\Sigma_I \times \Sigma_O$ when we specify 2-computations.

▶ **Theorem 7.** *Deciding $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, k, \infty)$, for $b \in \{1, 2\}$, is NP-complete.*

The proof of Theorem 7 is given in the full version. The lower bound is by a simple reduction from the Vertex Cover problem: the states of the system correspond to the cover, and whenever the

environment gives an edge of the graph to the system, the system should respond with the vertex that covers it.

By Lemma 3, we have that $w_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$ iff not $s_real_{\Sigma_O, \Sigma_I, \bar{b}}(\mathcal{A}, k, \infty)$. Theorem 7 then immediately implies the following.

► **Corollary 8.** *Deciding $w_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$, for $b \in \{1, 2\}$, is co-NP-complete.*

The implication of Theorem 7 and Corollary 8 is that deciding strong realizability with bounded systems (resp., weak realizability of bounded environments) amounts to search for a bounded system (resp., environment) that satisfies (resp., falsifies) the specification. Thus, what we have is essentially exhaustive search combined with model checking. The (co)-NP lower bounds tell us that there is no way of getting around the need to do an exhaustive search. Note that Corollary 8 already refers to the setting in which the environment is bounded, and it studies weak realizability there. In Section 5 below we study strong realizability in this setting.

5 Bounded Environments

We now turn to study the case of strong realizability in a setting of bounded environments (or, dually, weak realizability in a setting of bounded systems). In fact, bounding the environment is of interest already in the context of model-checking. For a system modeled by an I/O -transducer \mathcal{T} , a specification \mathcal{A} given by a DBW, and $k \geq 1$, we say that \mathcal{T} satisfies \mathcal{A} with respect to k -environments if for all O/I - k -transducers \mathcal{T}' , the computation $\mathcal{T} \parallel \mathcal{T}'$ satisfies \mathcal{A} .

► **Theorem 9.** *Given an I/O -transducer \mathcal{T} , a DBW \mathcal{A} over the alphabet $2^{I \cup O}$, and $k \geq 1$, the problem of deciding whether \mathcal{T} satisfies \mathcal{A} with respect to k -environments is co-NP-complete.*

The proof of Theorem 9 is given in the full version. The lower bound is by a reduction from the complement of the Hamiltonian Circle Problem [11].

We now turn to study the strong realizability problem when both the system and the environment are bounded.

► **Theorem 10.** *Deciding $s_real_{\Sigma_I, \Sigma_O, 1}(\mathcal{A}, n, m)$ and $s_real_{\Sigma_I, \Sigma_O, 2}(\mathcal{A}, n, m)$, for a given DBW \mathcal{A} over alphabets Σ_I, Σ_O , and positive integers n, m in unary, is Σ_2^P -complete.*

Proof. We prove the claim for $b = 1$. The argument for $b = 2$ is analogous. The strong realizability property $s_real_{\Sigma_I, \Sigma_O, 1}(\mathcal{A}, n, m)$ holds iff there exists a Σ_I/Σ_O transducer \mathcal{T} with at most n states such that for every Σ_O/Σ_I transducer \mathcal{T}' with at most m states, the 1-execution of $\mathcal{T} \parallel \mathcal{T}'$ satisfies the DBW specification \mathcal{A} . Membership in Σ_2^P follows from the fact that the sizes of the existentially and universally quantified transducers \mathcal{T} and \mathcal{T}' respectively are polynomially bounded in the size of the input, and the fact that we can check in polynomial time whether the 1-execution of $\mathcal{T} \parallel \mathcal{T}'$ satisfies \mathcal{A} for given $\mathcal{T}, \mathcal{T}'$, and \mathcal{A} [27].

To prove the hardness we reduce from the problem of deciding the truth of a formula of the form $\exists x \forall y \Phi(x, y)$, where x and y are vectors of Boolean variables and Φ is a formula in disjunctive normal form [11]. Let $x = (x_1, \dots, x_k)$, $y = (y_1, \dots, y_k)$ (we assume without loss of generality that x, y have the same number k of variables), and let $\Phi = C_1 \vee C_2 \vee \dots \vee C_p$, where each term C_i is a conjunction of literals (variables in x or y or their negations). We construct an instance (\mathcal{A}, n, m) of the strong realizability problem such that $s_real_{\Sigma_I, \Sigma_O, 1}(\mathcal{A}, n, m)$ iff $\exists x \forall y \Phi(x, y)$.

We describe the reduction in detail in the full version. Here we describe the general idea. The integers n, m are both set to $2p + 1 + k$. The input and output alphabets have size $2p + 1 + 2k$: The input alphabet is $\Sigma_I = \{d_0\} \cup \{d_i, d'_i \mid 1 \leq i \leq p\} \cup \{y_j, \bar{y}_j \mid 1 \leq j \leq k\}$, and the output alphabet is $\Sigma_O = \{c_0\} \cup \{c_i, c'_i \mid 1 \leq i \leq p\} \cup \{x_j, \bar{x}_j \mid 1 \leq j \leq k\}$.

Recall that the alphabet of the DBW \mathcal{A} is $\Sigma_O \times \Sigma_I$ and it specifies the expected behavior of the system with respect to all input sequences. Thus, \mathcal{A} should not limit the behavior of the environment. Since, however, the interaction between the system and the environment would be of interest only for some behaviors of the environment, we are going to describe how \mathcal{A} prescribe an expected format of interaction for both the system and the environment, and we assume that if, during the interaction, the system deviates from this format then \mathcal{A} moves to rejecting sink, and if the system follows the format but the environment deviates, then \mathcal{A} accepts.

The DBW \mathcal{A} prescribes the expected format in two phases. The goal of Phase 1 is to force the system and the environment to commit to a truth assignments for the variables in x and y , respectively, and to set aside states that output $c_1, c'_1, \dots, c_p, c'_p$, and $d_1, d'_1, \dots, d_p, d'_p$, respectively. Thus, if the system follows the prescribed format in Phase 1, then it outputs $2p + 1 + k$ different symbols during this phase: c_0, c_i, c'_i for $i = 1, \dots, p$, and either x_j or \bar{x}_j for $j = 1, \dots, k$. Since the bound n on its number of states is $2p + 1 + k$, this implies that each state of the system has as output one of the letters above. In particular, this means that for each variable x_j , the system has exactly one corresponding state that outputs x_j or \bar{x}_j , which corresponds to assigning value true or false, respectively, to the variable x_j . Similarly, if the environment follows the prescribed format in Phase 1, and then it has exactly $2p + 1 + k$ states, corresponding to the n letters generated d_0, d_i, d'_i for $i = 1, \dots, p$, and either y_j or \bar{y}_j for $j = 1, \dots, k$, and they induce a truth assignment for the y variables.

During Phase 2, the DBW \mathcal{A} checks that the assignment to which the system and environment commit in Phase 1 satisfies Φ . Phase 2 consists of p stages, one for each term C_i of Φ . In the stage for C_i , the DBW \mathcal{A} goes over the $2k$ variables and checks whether the assignment for them satisfies C_i . Recall that C_i is a conjunction. When \mathcal{A} detects that the assignment contradicts a requirement imposed by C_i , it moves to the phase for C_{i+1} , or rejects, if $i = p$. When \mathcal{A} concludes that the assignment satisfies C_i , it accepts. ◀

We can now turn to the problem of synthesis with bounded environments. We first need some definitions. For a word $w/y \in (\Sigma_I \times \Sigma_O)^*$ and $k \geq 1$, we say that w/y is k -generable if there is a Σ_I/Σ_O -transducer \mathcal{T} with at most k states such that $\mathcal{T}(w) = w/y$. Let $L_k \subseteq (\Sigma_I \times \Sigma_O)^*$ be the set of k -generable words.

► **Lemma 11.** *There is a DFW with $2^{k^{O(k)}}$ states that recognizes L_k .*

Proof. We prove that there is an NFW with $k^{O(k)}$ states that recognizes L_k . The NFW guesses a Σ_I/Σ_O -transducer \mathcal{T} with at most k states and then simulates it, checking that the y track is indeed the output of the w track. The number of Σ_I/Σ_O -transducers \mathcal{T} with exactly k states is $k!|\Sigma_O|^k k^{k|\Sigma_I|}$, and the number of transducers with at most k states, is no more than k times this expression. ◀

We can now solve the strong realizability problem for bounded environments.

► **Theorem 12.** *Deciding $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$, for $b \in \{1, 2\}$ is in 2EXPTIME.*

Proof. We prove the claim for $b = 1$. The argument for $b = 2$ is analogous.

We recall first how ones decides $s_real_{\Sigma_I, \Sigma_O, 1}(\mathcal{A}, \infty, \infty)$. The key idea is to consider the following game between two players, called System and Environment. In each round, System first chooses a letter in Σ_I and Environment then chooses a letter in Σ_O . System wins if the play, which is the infinite word in $(\Sigma_I \times \Sigma_O)^\omega$ that results from the choices of System and Environment is accepted by \mathcal{A} . Checking if System wins the game is equivalent to testing nonemptiness of Büchi tree automata and can be done in quadratic time [29]. Furthermore, if System wins the game, then there is a strategy that depends only on the states of \mathcal{A} . Thus, if $\mathcal{A} = \langle \Sigma_I \times \Sigma_O, Q, q_0, \rho, F \rangle$, then the nonemptiness algorithm yields a function $L : Q \rightarrow \Sigma_O$, which means that \mathcal{A} is strongly realized by the transducer $\mathcal{T} = \langle \Sigma_I, \Sigma_O, Q, q_0, \eta, L \rangle$, where $\eta(q, \sigma) = \rho(q, \langle \sigma, L(q) \rangle)$.

In deciding $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$, System only has to win against environments with at most k states. Thus, if a prefix of the play is not k -generable then System wins immediately, since no transducer with at most k states can generate such a play. Let \mathcal{D}_k be a DFW that accepts exactly all words that are not k -generable. By Lemma 11, such a DFW with $2^{2^{O(k)}}$ states exists.

Then, $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$ holds iff $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A} \times \mathcal{D}_k, \infty, \infty)$, where $\mathcal{A} \times \mathcal{D}_k$ is the product of \mathcal{A} and \mathcal{D}_k , which accepts a word $w \in (\Sigma_I \times \Sigma_O)^\omega$ if w is accepted by \mathcal{A} or is not k -generable. It follows that if $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$, then \mathcal{A} is realized by a transducer whose state space is that of $\mathcal{A} \times \mathcal{D}_k$. ◀

Theorem 12 tells us that we can solve strong realizability against bounded transducers, but at a cost that is doubly exponential in k . Thus, on the one hand, we expect more specifications to be realizable when we bound the size of the adversaries, but, on the other hand, deciding such realizability comes at a considerable cost. The question is whether this cost is unavoidable. To prove this, we would have to show that deciding $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$ is 2EXPTIME-hard.

OPEN QUESTION: Is deciding $s_real_{\Sigma_I, \Sigma_O, b}(\mathcal{A}, \infty, k)$ 2EXPTIME-hard?

As in the case of bounded systems, we can show that the strong realizability problem for bounded environment is at least NP-hard. The problem, however, seems to be much harder. As supporting evidence for the hardness of the problem, we show that recognizing k -generable words by an automaton requires doubly-exponential size.

For a word $x = x_1 \dots x_n \in \Sigma_I^*$, a *combination lock* transducer for x is $\mathcal{T}_x = \langle \Sigma_I, \Sigma_O, \{0, \dots, n\}, 0, \tau \rangle$, where $\Sigma_O = \{0, 1\}$, and for all $0 \leq j < n$, we have $\rho(j, x_{j+1}) = j + 1$ and $\rho(j, \sigma) = 0$, for all $\sigma \neq x_{j+1}$. Also, $\rho(n, \sigma) = n$, for all $\sigma \in \Sigma_I$. Finally, $\tau(j) = 0$ for all $0 \leq j < n$ and $\tau(n) = 1$. Thus, \mathcal{T}_x outputs 0 in all states but n . It can read x from its initial state, in which case it reaches the state n , where it outputs 1. When a violation of x is detected, \mathcal{T}_x goes back to its initial state.

► **Theorem 13.** *A DFW that recognizes L_k has at least 2^{2^k} states.*

Proof. Recall that two words $x_1, x_2 \in (\Sigma_I \times \Sigma_O)^*$ are L_k -equivalent, in the Myhill-Nerode sense, iff for all $z \in (\Sigma_I \times \Sigma_O)^*$, we have that $x_1 \cdot z \in L_k$ iff $x_2 \cdot z \in L_k$. The number of states in a minimal DFW for L_k is the number of equivalence classes of the L_k -equivalence relation.

Let $\Sigma_I = \{a, b, \#\}$ and $\Sigma_O = \{0, 1\}$. For a word $x \in (a + b)^k$, let \mathcal{T}_x be the transducer obtained from the combination lock $\{a, b\}/\{0, 1\}$ -transducer for x by adding $\#$ -transitions from all states to the initial state. Thus, $\#$ is an input reset symbol that resets to the initial state from all states. For every subset P of $(a + b)^k$, let $w(P)$ be the Σ_I/Σ_O word in which the input part consists of the words in P in some order, say lexicographic, with each word preceded by a reset, and the output part is all 0. We claim that all the words $w(P)$, for different subsets P , are not L_k -equivalent. Let P and Q be two different subsets. Assume, without loss of generality, that $x \in Q \setminus P$.

Note that the transducer \mathcal{T}_x is strongly connected, and it is easy to see also that it is minimized. Let w be a checking sequence for \mathcal{T}_x , and let $\mathcal{T}_x(w) = w/y$. That is, the only k -state transducer that can generate w/y (including k -transducers that are not combination locks) is \mathcal{T}_x starting from some state. If we append to $w(P)$ the word $\#/0 \cdot w/y$, then the resulting word is in L_k , as \mathcal{T}_x generates it. On the other hand, if we append the word $\#/0 \cdot w/y$ to $w(Q)$, then the resulting word is not in L_k . Indeed, because of the w/y portion, the only k -transducer that could possibly generate the resulting word is \mathcal{T}_x , but since $x \in Q$, the transducer \mathcal{T}_x cannot generate $w(Q)$, as it would output 1 after reading $\#x$.⁴ ◀

⁴ A small variant of the argument holds also for binary input alphabet, i.e., without the additional reset symbol. Restrict to combination locks where the first and the last symbol of the combination word is a , and replace the reset symbol by b^k . Then, b^k acts like a reset for these machines.

5.1 Absolute liveness properties

A specification ψ is of an *absolute liveness* property if for every computation π , we have that $\pi \models \psi$ iff $\pi \models F\psi$ [26]. In this section we show that for absolute liveness properties, weak and strong realizability against bounded environments coincide, and the complexity of deciding realizability is only exponential in the bound. Intuitively, it follows from the fact that the system can take its time to learn the environment with which it interacts, and then follow a strategy against that environment.

We formalize this intuition by means of the so-called *machine identification problem*: Given $k \geq 1$, we say that a word w is a *k-identifier* for every two O/I - k -transducers $\mathcal{T}'_1, \mathcal{T}'_2$, if the two transducers produce the same output sequence y in response to w , and if s_1, s_2 are their states respectively after processing w , then $\mathcal{T}'_1/s_1 \equiv \mathcal{T}'_2/s_2$. In other words, observing the response of the environment to w , identifies uniquely up to equivalence the part of the environment transducer that is reachable from the final state *after* w . The machine identification problem was formulated and solved by Moore in his classical paper [19]. It is shown there that for every $k \geq 1$, there is a k -identifier of length exponential in k , and it can be constructed in exponential time. (The word w is essentially a homing sequence of the disjoint union of all k -transducers.)

► **Theorem 14.** *Let ψ be an absolute liveness specification. Then*

1. $w_real_{I,O,1}(\psi, \infty, k)$ iff $s_real_{I,O,1}(\psi, \infty, k)$.
2. If ψ is given as a DBW, we can decide $s_real_{I,O,1}(\psi, \infty, k)$ in time polynomial in ψ and exponential in k .

Proof. Clearly $s_real_{I,O,1}(\psi, \infty, k)$ implies $w_real_{I,O,1}(\psi, \infty, k)$. For the other direction, assume that $w_real_{I,O,1}(\psi, \infty, k)$ holds. Then, for each O/I - k -transducer \mathcal{T}' there is an I/O -transducer \mathcal{T} that guarantees that ψ holds. We need, however, one I/O -transducer \mathcal{T} that can guarantee ψ against all O/I - k -transducers. What \mathcal{T} can do is first output a k -identifier sequence w . After observing the response y of the k -transducer \mathcal{T}' of the environment, we can construct a transducer \mathcal{T}'' with at most k states that is equivalent to \mathcal{T}'/s where s is the current state of \mathcal{T}' after w . Then, using weak realizability, \mathcal{T} can simulate the I/O -transducer that wins against \mathcal{T}'' . This proves the first claim.

The second claim follows from the fact that, for every given k -transducer \mathcal{T}' for the environment, we can determine in polynomial time whether there is a system that satisfies a DBW specification ψ with environment \mathcal{T}' , hence $w_real_{I,O,1}(\psi, \infty, k)$ (and $s_real_{I,O,1}(\psi, \infty, k)$) can be decided in time polynomial in ψ and exponential in k . If $s_real_{I,O,1}(\psi, \infty, k)$ holds, then a transducer \mathcal{T} for the system that realizes ψ can be constructed also in time polynomial in ψ and exponential in k , as explained above. ◀

References

- 1 A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- 2 J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite-state strategies. *Trans. AMS*, 138:295–311, 1969.
- 3 K. Chatterjee, T. Henzinger, and B. Jobstmann. Environment assumptions for synthesis. In *Proc. 19th CONCUR*, LNCS 5201, pages 147–161, 2008.
- 4 A. Church. Logic, arithmetics, and automata. In *Proc. Int. Congress of Mathematicians, 1962*, pages 23–35. Institut Mittag-Leffler, 1963.
- 5 E.M. Clarke, O. Grumberg, K.L. McMillan, and X. Zhao. Efficient generation of counterexamples and witnesses in symbolic model checking. In *Proc. 32st DAC*, pages 427–432, 1995.
- 6 R. Ehlers. Symbolic bounded synthesis. In *Proc. 22nd CAV*, LNCS 6174, pages 365–379, 2010.

- 7 E.A. Emerson. Automata, tableaux, and temporal logics. In *Proc. Workshop on Logic of Programs*, LNCS 193, pages 79–87. Springer, 1985.
- 8 E. Filiot, N. Jin, and J.-F. Raskin. An antichain algorithm for LTL realizability. In *Proc. 21st CAV*, LNCS 5643, pages 263–277, 2009.
- 9 D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th TACAS*, LNCS 6015, pages 190–204. Springer, 2010.
- 10 M. Faella and M. Napoli and M. Parente. Graded Alternating-Time Temporal Logic In *Proc. 10th LPAR*, LNCS 6355, pages 192–211. Springer 2010.
- 11 M. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. Freeman and Co., 1979.
- 12 A. Harding, M. Ryan, and P. Schobbens. A new algorithm for strategy synthesis in LTL games. In *Proc. 11th TACAS*, LNCS 3440, pages 477–492. Springer, 2005.
- 13 O. Kupferman and S. Sheinvald-Faragy. Finding shortest witnesses to the nonemptiness of automata on infinite words. In *Proc. 17th CONCUR*, LNCS 4137, pages 492–508. Springer, 2006.
- 14 O. Kupferman and M.Y. Vardi. From linear time to branching time. *ACM Transactions on Computational Logic*, 6(2):273–294, 2005.
- 15 O. Kupferman and M.Y. Vardi. Synthesis of trigger properties. In *Proc. 16th LPAR*, LNCS 6355, pages 312–331. Springer, 2010.
- 16 D. Lee and M. Yannakakis. Principles and methods of testing finite state machines - a survey. *Proc. IEEE*, 84:1089–1123, 1996.
- 17 D.A. Martin. Borel determinacy. *Annals of Mathematics*, 65:363–371, 1975.
- 18 A. Menezes, P.C. van Oorschot, and S.A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996.
- 19 E.F. Moore. Gedanken-experiments on sequential machines. In *Automata studies*, volume 34 of *Ann. Math. Studies*, pages 129–153. Princeton Univ. Press, 1956.
- 20 C.H. Papadimitriou and M. Yannakakis. On complexity as bounded rationality (extended abstract). In *Proc. 26th STOC*, pages 726–733, 1994.
- 21 N. Piterman, A. Pnueli, and Y. Saar. Synthesis of reactive(1) designs. In *Proc. 7th VMCAI*, LNCS 3855, pages 364–380. Springer, 2006.
- 22 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.
- 23 R. Rosner. *Modular Synthesis of Reactive Systems*. PhD thesis, Weizmann Inst. of Science, 1992.
- 24 A. Rubinstein. Finite automata play the repeated prisoner’s dilemma. *J. Economic Theory*, 39(1):83–96, 1986.
- 25 S. Schewe and B. Finkbeiner. Bounded synthesis. In *Proc. 5th ATVA*, LNCS 4762, pages 474–488. Springer, 2007.
- 26 A.P. Sistla. Safety, liveness and fairness in temporal logic. *Formal Aspects of Computing*, 6:495–511, 1994.
- 27 M.Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Logics for Concurrency: Structure versus Automata*, LNCS 1043, pages 238–266. Springer, 1996.
- 28 M.Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proc. 1st LICS*, pages 332–344, 1986.
- 29 M.Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logics of programs. *Journal of Computer and Systems Science*, 32(2):182–221, 1986.
- 30 G. Wang, A. Mishchenko, R. Brayton, and A. Sangiovanni-Vincentelli. Synthesizing FSMs according to co-Büchi properties. Technical report, UC Berkeley, 2005.