Logical Methods in Computer Science
Volume 21, Issue 3, 2025, pp. 12:1–12:26
https://lmcs.episciences.org/

Submitted Nov. 14, 2024
Published Jul. 29, 2025

SYNTHESIS WITH PRIVACY AGAINST AN OBSERVER

ORNA KUPFERMAN, OFER LESHKOWITZ, AND NAMMA SHAMASH HALEVY

School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel

Abstract. We study automatic *synthesis* of systems that interact with their environment and maintain *privacy* against an observer to the interaction. The system and the environment interact via sets I and O of input and output signals. The input to the synthesis problem contains, in addition to a specification, also a list of *secrets*, a function $\mathsf{cost} : I \cup O \to \mathbb{N}$, which maps each signal to the cost of hiding it, and a bound $b \in \mathbb{N}$ on the budget that the system may use for hiding of signals. The desired output is an (I/O)-transducer $\mathcal{T}$ and a set $\mathcal{H} \subseteq I \cup O$ of signals that respects the bound on the budget, thus $\sum_{s \in \mathcal{H}} \mathsf{cost}(s) \leq b$, such that for every possible interaction of $\mathcal{T}$, the generated computation satisfies the specification, yet an observer from which the signals in $\mathcal{H}$ are hidden, cannot evaluate the secrets.

We first show that the complexity of the problem is 2EXPTIME-complete for specifications and secrets in LTL, thus it is not harder than synthesis with no privacy requirements. We then analyze the complexity of the problem more carefully, isolating the two aspects that do not exist in traditional synthesis, namely the need to hide the value of the secrets and the need to choose the set $\mathcal{H}$. We do this by studying settings in which traditional synthesis can be solved in polynomial time – when the specification formalism is deterministic automata and when the system is closed, and show that each of the two aspects involves an exponential blow-up in the complexity. We continue and study *bounded synthesis with privacy*, where the input also includes a bound on the size of the synthesized transducer, as well as a variant of the problem in which the observer has *knowledge*, either about the specification or about the system, which can be helpful in evaluating the secrets. In addition, we study *certified privacy*, where the synthesis algorithm also provides to the user a certification about the secrets being hidden.

1. Introduction

Synthesis is the automated construction of correct systems from their specifications [BCJ18]. While synthesized systems are correct, there is no guarantee about their *quality*. Since designers will be willing to give up manual design only after being convinced that the automatic process replacing it generates systems of comparable quality, it is extremely important to develop and study quality measures for automatically-synthesized systems. An

Key words and phrases: Synthesis, Privacy, LTL, Automata, Complexity.

* A preliminary version of this paper appears in the Proceedings of the 27th Foundations of Software Science and Computation Structures, FoSSaCS 2024.

This research is supported by the Israel Science Foundation, Grant 2357/19, and the European Research Council, Advanced Grant ADVANSYNT.

DOI:10.46298/LMCS-21(3:12)2025

important quality measure is *privacy*: making sure that the system and its environment do not reveal information they prefer to keep private. Privacy is a vivid research area in Theoretical Computer Science. There, the notion of *differential privacy* is used for formalizing when an algorithm maintains privacy. Essentially, an algorithm is differentially private if by observing its output, one cannot tell if a particular individual's information is used in the computation [DN03, DMNS16]. Another related notion is *obfuscation* in system development, where we aim to develop systems whose internal operation is hidden [BGI+12, GGH+16]. Obfuscation has been mainly studied in the context of software, where it has exciting connections with cryptography [BGI+12, GGH+16].

In the setting of automated synthesis in formal methods, a very basic notion of privacy has been studied by means of synthesis with *incomplete information* [Rei84, KV00, CDHR06]. There, the system should satisfy its specification eventhough it only has a partial view of the environment. Lifting differential privacy to formal methods, researchers have introduced the temporal logic *HyperLTL*, which extends LTL with explicit trace quantification [CFK+14]. Such a quantification can relate computations that differ only in non-observable elements, and can be used for specifying that computations with the same observable input have the same observable output. The synthesis problem of HyperLTL is undecidable, yet is decidable for the fragment with a single existential quantifier, which can specify interesting properties [FHL+20b]. In [KL22], the authors suggested a general framework for automated synthesis of privacy-preserving reactive systems. In their framework, the input to the synthesis problem includes, in addition to the specification, also *secrets*. During its interaction with the environment, the system may keep private some of the assignments to the output signals, and it directs the environment which assignments to the input signals it should keep private. Consequently, the satisfaction value of the specification and secrets may become unknown. The goal is to synthesize a system that satisfies the specification yet keeps the value of the secrets unknown. Finally, lifting obfuscation to formal methods, researchers have studied the synthesis of obfuscation policies for temporal specifications. In [WRR+18], an obfuscation mechanism is based on edit functions that alter the output of the system, aiming to make it impossible for an observer to distinguish between secret and non-secret behaviors. In [DDM10], the goal is to synthesize a control function that directs the user which actions to disable, so that the observed sequence of actions would not disclose a secret behavior.

In this paper we continue to study privacy-preserving reactive synthesis. As in [KL22], our setting is based on augmenting the specification with secrets whose satisfaction value should remain unknown. Unlike [KL22], the system and the environment have complete information about the assignments to the input and output signals, and the goal is to hide the secrets from a third party, and to do so by hiding the assignment to some of the signals throughout the interaction. As an example, consider a system that directs a robot patrolling a warehouse storage. Typical specifications for the system require it to direct the robot so that it eventually reaches the shelves of requested items, it never runs out of energy, etc. An observer to the interaction between the system and the robot may infer properties we may want to keep private, like dependencies between customers and shelves visited, locations of battery docking stations, etc. If we want to prevent the observer from inferring these properties (a.k.a,. the secrets), we have to hide the interaction from it. Different effort should be made in order to hide different components of the interaction (alarm sound, content of shelves, etc.). Our framework synthesizes a system that realizes the specification without the secrets being revealed, subject to restrictions on hiding of signals. As another example, consider a scheduler that should grant access to a joint resource. The scheduler

should maintain mutual exclusion (grants are not given to different users simultaneously) and non-starvation (all requests are granted), while hiding details like waiting time or priority to specific users. In Examples 2.1 and 2.2, we describe in detail the application of our framework for the synthesis of such a scheduler, as well as its application in the synthesis of a robot that paints parts of manufactured pieces. The robot should satisfy some requirements about the generated pattern of colors while hiding other features of the pattern.

Formally, we consider a reactive system that interacts with its environments via sets I and O of input and output signals. At each moment in time, the system reads a truth assignment, generated by the environment, to the signals in I, and it generates a truth assignment to the signals in O. The interaction between the system and its environment generates a *computation*. The system *realizes* a specification φ if all its computations satisfy φ [PR89]. We introduce and study the problem of *synthesis with privacy in the presence of an observer*. Given a specification φ, and secrets $\psi_1, \ldots, \psi_k$ over $I \cup O$, our goal is to return, in addition to a system that realizes the specification φ, also a set $\mathcal{H} \subseteq I \cup O$ of *hidden signals*, such that the satisfaction value of the secrets $\psi_1, \ldots, \psi_k$ is unknown to an observer that does not know the truth values of the signals in $\mathcal{H}$. Thus, secrets are evaluated according to a *three-valued semantics*. The use of secrets enables us to hide *behaviors*, rather than just signals. [1] Obviously, hiding all signals guarantees that the satisfaction value of every secret is unknown. Hiding of signals, however, is not always possible or involves some cost. We formalize this by adding to the setting a function $\mathsf{cost} : I \cup O \to \mathbb{N}$, which maps each signal to the cost of hiding its value, and a bound $b \in \mathbb{N}$ on the budget that the system may use for hiding of signals. The set $\mathcal{H}$ of hidden signals has to respect the bound, thus $\sum_{s \in \mathcal{H}} \mathsf{cost}(s) \leq b$.

In some cases, it is desirable to hide the truth value of a secret only when some condition holds. For example, we may require to hide the content of selves only in some sections of the warehouse. We extend our framework to *conditional secrets*: pairs of the form $\langle \theta, \psi \rangle$, where the satisfaction value of the secret ψ should be hidden from the observer only when the trigger θ holds. In particular, when $\theta = \psi$, we require to hide the secret only when it holds. For example, we may require to hide an unfair scheduling policy only when it is applied. Note that a conditional secret $\langle \theta, \psi \rangle$ is not equivalent to a secret $\theta \rightarrow \psi$ or $\theta \rightarrow \neg\psi$, and that the synthesized system may violate the trigger, circumventing the need to hide the secret. For example, by synthesizing a fair scheduler, the designer circumvents the need to hide an unfair policy.

We show that synthesis with privacy is 2EXPTIME-complete for specifications and secrets in LTL. Essentially, once the set $\mathcal{H}$ of hidden signals is determined, we can compose an automaton for the specification with automata that verify, for each secret, that the assignments to the signals in $(I \cup O) \setminus \mathcal{H}$ can be completed both in a way that satisfies the secret and in a way that does not satisfy it. A similar algorithm works for conditional secrets.

While the complexity of our algorithm is not higher than that of LTL synthesis with no privacy, it would be misleading to conclude that handling of privacy involves no increase in the complexity. The 2EXPTIME complexity follows from the need to translate LTL specifications to deterministic automata on infinite words. Such a translation involves a doubly-exponential blow-up [KV05a, KR10], which possibly dominates other computational tasks of the algorithm. In particular, two aspects of synthesis with privacy that do not exist

[1] Hiding of signals is a special case of our framework. Specifically, hiding of a signal p can be done with the secrets Fp and $F\neg p$.

in usual synthesis are a need to go over all possible choices of signals to hide, and a need to go over all assignments to the hidden signals.

Our main technical contribution is a finer complexity analysis of the problem, which reveals that each of the two aspects above involves an exponential complexity: the first in the number of signals and the second in the size of the secret. We start with the need to go over all assignments of hidden signals and show that even when the specification is $\mathtt{T}$, the set $\mathcal{H}$ of hidden signals is given, and there is only one secret, given by a deterministic Büchi automaton, synthesis with privacy is EXPTIME-complete. This is exponentially higher than synthesis of deterministic Büchi automata, which can be solved in polynomial time. We continue to the need to go over all possible choices of $\mathcal{H}$. For that, we focus on the closed setting, namely when $I = \varnothing$, and the case the specification and secrets are given by deterministic automata. We show that while synthesis with privacy can be then solved in polynomial time for a given set $\mathcal{H}$, it is NP-complete when $\mathcal{H}$ is not given, even when the function cost is uniform.

We continue and study three variants of the problem: *bounded synthesis*, *certified privacy*, and *knowledgeable observer*, briefly described below. One way to cope with the 2EXPTIME complexity of LTL synthesis, which is carried over to a doubly-exponential lower bound on the size of the generated system [Ros92], is bounded synthesis. There, the input to the problem includes also a bound on the size of the system [SF07, Ehl10, KLVY11]. In a setting with no privacy, the bound reduces the complexity of LTL synthesis to PSPACE, as one can go over all candidate systems. We study bounded synthesis with privacy and show that privacy makes the problem much harder: it is EXPSPACE-complete when the specification and secrets are given by LTL formulas, and is PSPACE-complete when they are given by deterministic parity (or Büchi) automata.

The second variant we consider is one in which in addition to hiding the secret, the system has to generate a *certificate for privacy*. The certificate is an assignment to the hidden signals that differs from the assignment generated during the interaction and with which the secret has a different satisfaction value. This way, the user gets a witness computation that proves privacy is preserved. We examine the difference from the standard setting and specifically show that the need to generate certificates in an online manner makes specification and secrets harder to realize. Then, we provide a solution for synthesis with certified privacy and show that it is 2EXPTIME-complete and can be solved Safralessly (i.e., without determinization of a nondeterministic Büchi automaton) [KV05b].

Finally, recall that a system keeps a secret ψ private if an observer cannot reveal the truth value of ψ: every observable computation can be completed both to a computation that satisfies ψ and to a computation that does not satisfy ψ. We study a setting in which the observer has additional knowledge. First, we consider the setting in which the observer knows the specification φ of the system. Consequently, the observer knows that only completions that satisfy φ should be taken into account. If, for example, $\varphi \rightarrow \psi$, then ψ cannot be kept private. We describe an algorithm for this variant of the problem and analyze the way knowledge of the specification affects the complexity. In particular, we show that the problem becomes EXPTIME-complete even when the specification is given by a deterministic Büchi automaton and the secrets are of a fixed size. Then, we consider the even more demanding setting in which the observer knows the structure of the system. Consequently, the observer knows the set of possible computations of the system and thus may restrict to them when trying to evaluate the secret. We describe a procedure for checking whether a

given (I/O)-transducer $\mathcal{T}$ hides a given secret from an observer that knows $\mathcal{T}$. We leave the synthesis problem for this variant open and briefly explain why we suspect it is undecidable.

2. Preliminaries

2.1. **Synthesis.** For a finite nonempty alphabet Σ, an infinite *word* $w = \sigma_0 \cdot \sigma_1 \cdot \ldots \in \Sigma^\omega$ is an infinite sequence of letters from Σ. A *language* $L \subseteq \Sigma^\omega$ is a set of infinite words.

Let I and O be disjoint finite sets of input and output signals, respectively. We consider the alphabet $2^{I \cup O}$ of truth assignments to the signals in $I \cup O$. Then, a languages $L \subseteq (2^{I \cup O})^\omega$ can be viewed as a *specification*, and the *truth value* of L in a computation $w \in (2^{I \cup O})^\omega$ is $\mathtt{T}$ if $w \in L$, and is $\mathtt{F}$ otherwise.

An (I/O)*-transducer* is a tuple $\mathcal{T} = \langle I, O, S, s_0, \eta, \tau \rangle$, where S is a finite set of states, $s_0 \in S$ is an initial state, $\eta : S \times 2^I \to S$ is a transition function, and $\tau : S \to 2^O$ is a labeling function. We extend the transition function η to words in $(2^I)^*$ in the expected way, thus $\eta^* : S \times (2^I)^* \to S$ is such that for all $s \in S$, $x_I \in (2^I)^*$, and $i \in 2^I$, we have that $\eta^*(s, \epsilon) = s$, and $\eta^*(s, x_I \cdot i) = \eta(\eta^*(s, x_I), i)$. For a word $w_I = i_0 \cdot i_1 \cdot i_2 \cdot \ldots \in (2^I)^\omega$, we define the *computation of* $\mathcal{T}$ *on* w_I to be the word $\mathcal{T}(w_I) = (i_0 \cup o_0) \cdot (i_1 \cup o_1) \cdot \ldots \in (2^{I \cup O})^\omega$, where for all $j \geq 0$, we have that $o_j = \tau(\eta^*(s_0, i_0 \cdots i_j))$. The *language* of $\mathcal{T}$, denoted $L(\mathcal{T})$, is the set of computations of $\mathcal{T}$, that is $L(\mathcal{T}) = \{\mathcal{T}(w_I) : w_I \in (2^I)^\omega\}$.

We say that $\mathcal{T}$ *realizes* a language $L \subseteq (2^{I \cup O})^\omega$ if $L(\mathcal{T}) \subseteq L$. We say that a language $L \subseteq (2^{I \cup O})^\omega$ is *realizable* if there is an (I/O)-transducer that realizes it. In the *synthesis* problem, we are given a specification language $L \subseteq (2^{I \cup O})^\omega$ and we have to return an (I/O)-transducer that realizes L or decide that L is not realizable. The language L is given by an automaton over the alphabet $2^{I \cup O}$ or a temporal logic formula over $I \cup O$ (see definitions in Section 2.4).

2.2. **Synthesis with privacy.** In the *synthesis with privacy* problem, we are given, in addition to the specification language $L_\varphi \subseteq (2^{I \cup O})^\omega$, also a *secret* $L_\psi \subseteq (2^{I \cup O})^\omega$, which defines a behavior that we want to hide from an observer[2]. Thus, we seek an (I/O)-transducer that realizes L_φ without revealing the truth value of L_ψ in the generated computations. Keeping the truth value of L_ψ secret is done by hiding the truth value of some signals in $I \cup O$. Before we define synthesis with privacy formally, we first need some notations.

Consider a set $\mathcal{H} \subseteq I \cup O$ of *hidden signals*. Let $\mathcal{V} = (I \cup O) \setminus \mathcal{H}$ denote the set of *visible signals*. For an assignment $\sigma \in 2^{I \cup O}$, let $\sigma|_\mathcal{V} \in 2^\mathcal{V}$ and $\sigma|_\mathcal{H}$ be the restriction of σ to the visible and hidden signals respectively. Thus, $\sigma = \sigma|_\mathcal{V} \cup \sigma|_\mathcal{H}$. Then, for an infinite word $w = \sigma_0 \cdot \sigma_1 \cdot \sigma_2 \cdots \in (2^{I \cup O})^\omega$, we denote by $w|_\mathcal{V}$ and $w|_\mathcal{H}$, the restriction of w to $\mathcal{V}$ and $\mathcal{H}$ respectively. Thus, $w|_\mathcal{V} = \sigma_0|_\mathcal{V} \cdot \sigma_1|_\mathcal{V} \cdot \sigma_2|_\mathcal{V} \cdots \in (2^\mathcal{V})^\omega$ and $w|_\mathcal{H} = \sigma_0|_\mathcal{V} \cdot \sigma_1|_\mathcal{V} \cdot h_2 \cdots \in (2^\mathcal{H})^\omega$. Then, let $\mathsf{noise}_\mathcal{H}(w)$ be the set of all computations that differ from w in assignments to the signals in $\mathcal{H}$. Thus, when the signals in $\mathcal{H}$ are hidden, then an observer of a computation $w \in (2^{I \cup O})^\omega$ only knows that the computation is one from $\mathsf{noise}_\mathcal{H}(w)$. Formally, $\mathsf{noise}_\mathcal{H}(w) = \{w' \in (2^{I \cup O})^\omega : w'|_\mathcal{V} = w|_\mathcal{V}\}$.

Consider a specification $L_\varphi \subseteq (2^{I \cup O})^\omega$ and a secret $L_\psi \subseteq (2^{I \cup O})^\omega$. For a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, we say that an (I/O)-transducer $\mathcal{T}$ $\mathcal{H}$*-hides* L_ψ if for all words $w_I \in (2^I)^\omega$, the truth value of the secret L_ψ in the computation $\mathcal{T}(w_I)$ cannot be deduced from $\mathcal{T}(w_I)|_\mathcal{V}$.

[2]See Remark 2.3 for an extension of the setting to multiple and conditional secrets.

Formally, for every $w_I \in (2^I)^\omega$, there exist two computations $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I))$, such that $w^+ \in L_\psi$ and $w^- \notin L_\psi$. We say that $\mathcal{T}$ *realizes* $\langle L_\varphi, L_\psi, \mathcal{H}\rangle$ *with privacy* if $\mathcal{T}$ realizes L_φ and $\mathcal{H}$-hides L_ψ. We say that $\langle L_\varphi, L_\psi, \mathcal{H}\rangle$ is *realizable with privacy* if there exists an (I/O)-transducer that realizes $\langle L_\varphi, L_\psi, \mathcal{H}\rangle$ with privacy.

Clearly, hiding is monotone with respect to $\mathcal{H}$, in the sense that the larger $\mathcal{H}$ is, the more likely it is for an (I/O)-transducer $\mathcal{T}$ to $\mathcal{H}$-hide L_ψ. Indeed, if $\mathcal{T}$ $\mathcal{H}$-hides L_ψ, then $\mathcal{T}$ $\mathcal{H}'$-hides L_ψ for all $\mathcal{H}'$ with $\mathcal{H} \subseteq \mathcal{H}'$. In particular, taking $\mathcal{H} = I \cup O$, we can hide all non-trivial secrets. Hiding of signals, however, is not always possible, and may sometimes involve a cost. Formally, we consider a *hiding cost function* $\mathsf{cost} : I \cup O \to \mathbb{N}$, which maps each signal to the cost of hiding it, and a *hiding budget* $b \in \mathbb{N}$, which bounds the cost that the system may use for hiding of signals. The cost of hiding a set $\mathcal{H} \subseteq I \cup O$ of signals is then $\mathsf{cost}(\mathcal{H}) = \sum_{p\in\mathcal{H}} \mathsf{cost}(p)$, and we say that $\mathcal{H}$ *respects* b if $\mathsf{cost}(\mathcal{H}) \leq b$. Note that if $\mathsf{cost}(p) > b$, for $p \in I \cup O$, then p cannot be hidden. Also, when $\mathsf{cost}(p) = 1$ for all $p \in I \cup O$, we say that cost is *uniform*. Note that then, b bounds the number of signals we may hide.

Now, we say that $\langle L_\varphi, L_\psi, \mathsf{cost}, b\rangle$ is *realizable with privacy* if there exists a set $\mathcal{H} \subseteq I \cup O$ such that $\mathcal{H}$ respects b and $\langle L_\varphi, L_\psi, \mathcal{H}\rangle$ is realizable with privacy. Finally, in the *synthesis with privacy* problem, we are given L_φ, L_ψ, cost, and b, and we have to return a set $\mathcal{H} \subseteq I \cup O$ that $\mathcal{H}$ respects b and an (I/O)-transducer $\mathcal{T}$ that realizes $\langle L_\varphi, L_\psi, \mathcal{H}\rangle$ with privacy, or determine that $\langle L_\varphi, L_\psi, \mathsf{cost}, b\rangle$ is not realizable with privacy.

2.3. **Multiple and conditional secrets.** In this section we discuss two natural extensions of our setting. First, often we need to hide form the observer more than one secret. We extend the definition of synthesis with privacy to a set of secrets $S = \{L_{\psi_1}, L_{\psi_2}, \ldots, L_{\psi_k}\}$ in the natural way. Thus, an (I/O)-transducer $\mathcal{T}$ realizes $\langle \varphi, S, \mathcal{H}\rangle$ with privacy if it realizes φ and $\mathcal{H}$-hides L_{ψ_i}, for all $i \in [k]$. Note that

Then, a *conditional secret* is a pair $\langle L_\theta, L_\psi\rangle$, consisting of a *trigger* and a *secret*. The truth value of the secret should be unknown only in computations that satisfy the trigger. Formally, for a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, we say that an (I/O)-transducer $\mathcal{T}$ $\mathcal{H}$*-hides* $\langle L_\theta, L_\psi\rangle$ if for all input sequences $w_I \in (2^I)^\omega$ such that $\mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I)) \subseteq L_\theta$, the truth value of L_ψ in the computation $\mathcal{T}(w_I)$ cannot be deduced from $\mathcal{T}(w_I)|_{\mathcal{V}}$, thus there exist two computations $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I))$, such that $w^+ \in L_\psi$ and $w^- \notin L_\psi$. A useful special case of conditional secrets is when the trigger and the secret coincide, and so we have to hide the truth value of the secret only if there are computations where the value of secret is $\mathtt{T}$. Formally, $\mathcal{T}$ $\mathcal{H}$-hides $\langle L_\psi, L_\psi\rangle$ if for all input sequences $w_I \in (2^I)^\omega$, there exists a computation $w^- \in \mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I))$ such that $w^- \notin L_\psi$.

Note that unlike a collection of specifications, which can be conjuncted, hiding a set of secrets is not equivalent to hiding their conjunction. Likewise, hiding a conditional secret is not equivalent to hiding the implication of the secret by the trigger. Thus, the two variants require an extension of the solution for the case of a single or unconditional secret. In Remark 3.5, we describe such an extension.

2.4. **Automata and LTL.** An *automaton* on infinite words is $\mathcal{A} = \langle \Sigma, Q, q_0, \delta, \alpha\rangle$, where Σ is an alphabet, Q is a finite set of *states*, $q_0 \in Q$ is an *initial state*, $\delta : Q \times \Sigma \to 2^Q$ is a *transition function*, and α is an *acceptance condition*, to be defined below. For states $q, s \in Q$ and a letter $\sigma \in \Sigma$, we say that s is a σ-successor of q if $s \in \delta(q, \sigma)$. Note that we do not require the transition function to be *total*. That is, we allow that $\delta(q, \sigma) = \varnothing$.

If $|\delta(q,\sigma)| \leq 1$ for every state $q \in Q$ and letter $\sigma \in \Sigma$, then $\mathcal{A}$ is *deterministic*. For a deterministic automaton $\mathcal{A}$ we view δ as a function $\delta : Q \times \Sigma \to Q \cup \{\bot\}$, where $\bot$ is a distinguished symbol, and instead of writing $\delta(q,\sigma) = \{q\}$ and $\delta(q,\sigma) = \varnothing$, we write $\delta(q,\sigma) = q$ and $\delta(q,\sigma) = \bot$, respectively.

A *run* of $\mathcal{A}$ on $w = \sigma_0 \cdot \sigma_1 \cdots \in \Sigma^\omega$ is an infinite sequence of states $r = r_0 \cdot r_1 \cdot r_2 \cdot \ldots \in Q^\omega$, such that $r_0 = q_0$, and for all $i \geq 0$, we have that $r_{i+1} \in \delta(r_i, \sigma_i)$. The acceptance condition α determines which runs are "good". We consider here the *Büchi*, *co-Büchi*, *generalized Büchi* and *parity* acceptance conditions. All conditions refer to the set $inf(r) \subseteq Q$ of states that r traverses infinitely often. Formally, $inf(r) = \{q \in Q : q = r_i \text{ for infinitely many } i\text{'s}\}$. In generalized Büchi the acceptance condition is of the form $\alpha = \{\alpha_1, \alpha_2, \ldots, \alpha_k\}$, for $k \geq 1$ and sets $\alpha_i \subseteq Q$. In a generalized Büchi automaton, a run r is accepting if for all $1 \leq i \leq k$, we have that $inf(r) \cap \alpha_i \neq \varnothing$. Thus, r visits each of the sets in α infinitely often. Büchi automata is a special case of its generalized form with $k = 1$. That is, a run r is accepting with respect to the Büchi condition $\alpha \subseteq Q$, if $inf(r) \cap \alpha \neq \varnothing$. Dually, in co-Büchi automata, a run r is accepting if $inf(r) \cap \alpha = \varnothing$. Finally, in a parity automaton, the acceptance condition $\alpha : Q \to \{1, ..., k\}$, for some $k \geq 1$, maps states to ranks, and a run r is accepting if the maximal rank of a state in $inf(r)$ is even. Formally, $\max_{q \in inf(r)}\{\alpha(q)\}$ is even. A run that is not accepting is *rejecting*. We refer to the number k in α as the *index* of the automaton. A word w is accepted by $\mathcal{A}$ if there is an accepting run of $\mathcal{A}$ on w. The language of $\mathcal{A}$, denoted $L(\mathcal{A})$, is the set of words that $\mathcal{A}$ accepts. Two automata are *equivalent* if their languages are equivalent.

We denote the different classes of automata by three-letter acronyms in {D,N} × {B,C,GB,P} × {W}. The first letter stands for the branching mode of the automaton (deterministic, nondeterministic); the second for the acceptance condition type (Büchi, co-Büchi, generalized Büchi, or parity); and the third indicates we consider automata on words. For example, NBWs are nondeterministic Büchi word automata.

LTL is a linear temporal logic used for specifying on-going behaviors of reactive systems [Pnu81]. Specifying the behavior of (I/O)-transducers, formulas of LTL are defined over the set $I \cup O$ of signals using the usual Boolean operators and the temporal operators G ("always") and F ("eventually"), X ("next time") and U ("until"). The semantics of LTL is defined with respect to infinite computations in $(2^{I \cup O})^\omega$. Thus, each LTL formula φ over $I \cup O$ induces a language $L_\varphi \subseteq (2^{I \cup O})^\omega$ of all computations that satisfy φ.

Recall that the input to the problem of synthesis with privacy includes languages L_φ and L_ψ. We sometimes replace L_φ and L_ψ in the different notations with automata or LTL formulas that describe them, thus talk about realizability with privacy of $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H} \rangle$ or $\langle \varphi, \psi, \mathcal{H} \rangle$, for automata $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$, or LTL formulas φ and ψ.

Example 2.1. Consider a scheduler that serves two users and grant them with access to a joint resource. The scheduler can be viewed as an open system with $I = \{\mathsf{req}_1, \mathsf{req}_2\}$, with req_i ($i \in \{1, 2\}$) standing for a request form User i, and $O = \{\mathsf{grant}_1, \mathsf{grant}_2\}$, with grant_i standing for a grant to User i. The system should satisfy mutual exclusion and non-starvation. Formally, the specification for the system is $\varphi_1 \wedge \varphi_2 \wedge \varphi_3$, for $\varphi_1 = G((\neg\mathsf{grant}_1) \vee (\neg\mathsf{grant}_2))$, $\varphi_2 = G(\mathsf{req}_1 \to F\mathsf{grant}_1)$, and $\varphi_3 = G(\mathsf{req}_2 \to F\mathsf{grant}_2)$.

We may want to hide from an observer of the interaction the exact policy scheduling of the system. For example,[3] the secret $\psi_1 = ((\neg\mathsf{grant}_1)W\mathsf{req}_1) \wedge G(\mathsf{grant}_1 \to X((\neg\mathsf{grant}_1)W\mathsf{req}_1))$ reveals whether the system gives User 1 grants only after requests that have not been granted

[3]The LTL operator W is "weak Until", thus $p_1 W p_2 = (p_1 U p_2) \vee G p_1$.

yet. Indeed, ψ_1 specifies that once a grant to User 1 is given, no more grants are given to her, unless a new request from her arrives. A similar secret can be specified for User 2. Note that in order to hide ψ_1, it is sufficient to hide only one of the signals req_1 or grant_1. In fact, this is true even when the observer knows the specification for the system. Then, the secret $\psi_2 = G((\mathsf{req}_1 \rightarrow \mathsf{grant}_1 \vee X\mathsf{grant}_1) \wedge (\mathsf{req}_2 \rightarrow \mathsf{grant}_2 \vee X\mathsf{grant}_2))$ reveals whether delays in grants are limited to one cycle. Here, unlike with ψ_2, it is not sufficient hiding only a single request or even both. Indeed, some policies disclose the satisfaction value of ψ_2 even when requests are hidden. For example, a system that simply alternates between grants, thus outputs $\{\mathsf{grant}_1\}, \{\mathsf{grant}_2\}, \{\mathsf{grant}_1\}, \{\mathsf{grant}_2\}, \ldots$, satisfies the specification and clearly satisfies ψ_2 regardless of the users' requests.

Consider now the secret $\psi_3 = FG(\mathsf{req}_1 \rightarrow \mathsf{grant}_1)$, which asserts that eventually, the requests of User 1 are always granted immediately. A system that satisfies ψ_3 is unfair to User 2. Aiming to hide this unfair behavior, we can use the conditional secret $\langle \psi_3, \psi_3 \rangle$, which requires a system that satisfies ψ_3 to hide its satisfaction.

Some computations that satisfy ψ_3, however, may still be fair to User 2. For example, if ψ_3 is satisfied vacuously or if only finitely many requests are sent from User 2, then the behavior specified in ψ_3 is fair, and we need not hide it. Accordingly, we can strengthen the trigger ψ_3 and restrict further the computations in which the satisfaction value of ψ_3 should be hidden. Formally, we replace the trigger ψ_3 by a trigger $\psi_3 \wedge \theta$, for a behavior θ in which a scheduling policy that satisfies ψ_3 is not fair (and hence, need to be hidden).

Let us consider possible behaviors θ for the conditional secret $\langle \psi_3 \wedge \theta, \psi_3 \rangle$. As discussed above, behaviors that make ψ_3 unfair are $GF\mathsf{req}_1$, implying that ψ_3 is not satisfied vacuously, and $GF\mathsf{req}_2$, implying that the immediate grants to User 1 are not due to no requests from User 2. Taking $\theta = (GF\mathsf{req}_1) \wedge (GF\mathsf{req}_2)$ results in a more precise conditional secret.

The trigger θ can be made more precise: taking $\theta = GF(\mathsf{req}_1 \wedge \mathsf{req}_2)$ still guarantees no vacuous satisfaction and also asserts that immediate grants to User 1 are given even when the requests of User 1 arrive together with those of User 2. In fact, $\theta = GF(\mathsf{req}_2 \wedge (\neg\mathsf{grant}_2)U\mathsf{req}_1))$ is even more precise, as it excludes the possibility that the requests of User 1 arrive before those of User 2. Note that the secret can be made less restrictive too, for example with $\psi_3' = FG(\mathsf{req}_1 \rightarrow ((\neg\mathsf{grant}_2)U\mathsf{grant}_1))$, which specifies that eventually, grants to User 1 are always given before grants to User 2. □

Example 2.2. As a different example, consider a paint robot that paints parts of manufactured pieces. The set O includes 3 signals c_0, c_1, c_2 that encode 8 colors. The encoding corresponds to the composition of the color from paint in three different containers. For example, color 101 stands for the robot mixing paints from containers 0 and 2. The observer does not see the generated pattern, but, unless we hide it, may see the arm of the robot when it reaches a container. Accordingly, hiding of signals in O involve different costs.

The user instructs the robot whether to stay with the current color or change it, thus $I = \{\mathsf{change}\}$. We seek a system that directs the robot which color to choose, in a way that satisfies requirements about the generated pattern. For example, in addition to the requirement to respect the changing instructions ($\xi_{respect}$), the specification φ may require the pattern to start with color 000, and if there are infinitely many changes, then all colors are used (ξ_{all}), yet color 000 repeats between each two colors (ξ_{repeat}). Formally,

- $\xi_{respect} = G((X\neg\mathsf{change}) \leftrightarrow ((c_0 \leftrightarrow Xc_0) \wedge (c_1 \leftrightarrow Xc_1) \wedge (c_2 \leftrightarrow Xc_2)))$,
- $\xi_{all} = GF(\bar{c_0} \wedge \bar{c_1} \wedge \bar{c_2}) \wedge GF(\bar{c_0} \wedge \bar{c_1} \wedge c_2) \wedge \cdots \wedge GF(c_0 \wedge c_1 \wedge c_2)$,
- $\xi_{repeat} = G((c_0 \vee c_1 \vee c_2) \rightarrow X(\mathsf{change} \rightarrow (\bar{c_0} \wedge \bar{c_1} \wedge \bar{c_2})))$, and

- $\varphi = (\bar{c_0} \wedge \bar{c_1} \wedge \bar{c_2}) \wedge \xi_{respect} \wedge ((GF\mathsf{change}) \rightarrow (\xi_{all} \wedge \xi_{repeat}))$.

We may want to hide from an observer certain patterns that the robot may produce. For example, the fact color 111 is used only after color 110 (with color 000 between them), the fact there are colors other than 000 that repeat without a color different from 000 between them, and more. Note that not all the signals in O need to be hidden, and that the choice of signals to hide depends on the secrets as well as the cost of hiding. □

3. Solving Synthesis with Privacy

In this section we describe a solution to the problem of synthesis with privacy for LTL specifications and show that it is TIME-complete, thus not harder than LTL synthesis. The solution is based on replacing the specification by one that guarantees the hiding of the secret. For this, we need the following two constructions.

Lemma 3.1. *Consider a nondeterministic automaton $\mathcal{A} = \langle 2^{I \cup O}, Q, q_0, \delta, \alpha \rangle$. Given a set $\mathcal{H} \subseteq I \cup O$, there is a transition function $\delta^{\mathcal{H}} : Q \times 2^{I \cup O} \rightarrow 2^Q$ such that the nondeterministic automaton $\mathcal{A}^{\mathcal{H}} = \langle 2^{I \cup O}, Q, q_0, \delta^{\mathcal{H}}, \alpha \rangle$ is such that $L(\mathcal{A}^{\mathcal{H}}) = \mathsf{noise}_{\mathcal{H}}(L(\mathcal{A}))$.*

Proof. Intuitively, the transition function $\delta^{\mathcal{H}}$ increases the nondeterminism of δ by guessing an assignment to the signals in $\mathcal{H}$. Formally, for $q \in Q$ and $\sigma \in 2^{I \cup O}$, we define $\delta^{\mathcal{H}}(q, \sigma) = \bigcup_{\sigma' \in \mathsf{noise}_{\mathcal{H}}(\sigma)} \delta(q, \sigma')$. It is easy to see that a word w' is accepted by $\mathcal{A}^{\mathcal{H}}$ iff there is a word w accepted by $\mathcal{A}$ such that $w' \in \mathsf{noise}_{\mathcal{H}}(w)$. □

Note that while $\mathcal{A}^{\mathcal{H}}$ maintains the state space and acceptance condition of $\mathcal{A}$, it does not preserve determinism. Indeed, unless $\mathcal{H} = \varnothing$, we have that $\mathcal{A}^{\mathcal{H}}$ is nondeterministic even when $\mathcal{A}$ is deterministic. Next, in Lemma 3.2 we construct automata that accept computations that satisfy the specification and hide the secret when a given set of signals is hidden.

Lemma 3.2. *Consider two disjoint finite sets I and O, a subset $\mathcal{H} \subseteq I \cup O$, and ω-regular languages L_φ and L_ψ over the alphabet $2^{I \cup O}$. There exists a DPW $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ with alphabet $2^{I \cup O}$ that accepts a computation $w \in (2^{I \cup O})^\omega$ iff $w \in L_\varphi$ and there exist $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(w)$ such that $w^+ \in L_\psi$ and $w^- \notin L_\psi$.*

1. *If L_φ and L_ψ are given by LTL formulas φ and ψ, then $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ has $2^{2^{O(|\varphi|+|\psi|)}}$ states and index $2^{O(|\varphi|+|\psi|)}$.*
2. *If L_φ and L_ψ are given by DPWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$ with n_φ and n_ψ states, and of indices k_φ and k_ψ, then $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ has $2^{O(n_\varphi \cdot k_\varphi \cdot (n_\psi \cdot k_\psi)^2 \log(n_\varphi \cdot k_\varphi \cdot n_\psi \cdot k_\psi))}$ states and index $O(n_\varphi \cdot k_\varphi \cdot (n_\psi \cdot k_\psi)^2)$.*

Proof. We start with the case L_φ and L_ψ are given by LTL formulas φ and ψ. Let $\mathcal{A}_\varphi$, $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$ be NGBWs for L_φ, L_ψ, and $L_{\neg\psi}$. By [VW94], such NGBWs exist, and are of size exponential in the corresponding LTL formulas. Let $\mathcal{A}^{\mathcal{H}}_\psi$ and $\mathcal{A}^{\mathcal{H}}_{\neg\psi}$ be the NGBWs for $\mathsf{noise}_{\mathcal{H}}(L(\mathcal{A}_\psi))$ and $\mathsf{noise}_{\mathcal{H}}(L(\mathcal{A}\neg\psi))$, respectively, constructed as in Lemma 3.1.

Now, let $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$ be an NGBW for the intersection of the three automata $\mathcal{A}_\varphi$, $\mathcal{A}^{\mathcal{H}}_\psi$, and $\mathcal{A}^{\mathcal{H}}_{\neg\psi}$. The NGBW $\mathcal{N}$ can be easily defined on top of the product of the three automata, and hence is of size $2^{O(|\varphi|+|\psi|)}$ and index $O(|\varphi| + |\psi|)$. Observe that indeed, a word $w \in (2^{I \cup O})^\omega$ is accepted by $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$ iff $w \models \varphi$ and there exist $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(w)$ such that $w^+ \models \psi$ and

$w^- \not\models \psi$. By [Saf88, Pit06], determinizing $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$ results in a DPW $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ with $2^{2^{O(|\varphi|+|\psi|)}}$ states and index $2^{O(|\varphi|+|\psi|)}$, and we are done.

We continue with the case L_φ and L_ψ are given by DPWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$. We first obtain from $\mathcal{D}_\psi$ two NBWs, $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$ for L_ψ and $L_{\neg\psi} = (2^{I\cup O})^\omega \setminus L_\psi$ respectively, and also we translate $\mathcal{D}_\varphi$ into an equivalent NBW $\mathcal{A}_\varphi$. Note that the NBWs $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$ can be defined with $O(n_\psi \cdot k_\psi)$ states, and that $\mathcal{A}_\varphi$ can be defined with $O(n_\varphi \cdot k_\varphi)$ states. We then obtain the NBWs $\mathcal{A}^{\mathcal{H}}_\psi$ and $\mathcal{A}^{\mathcal{H}}_{\neg\psi}$ by applying the construction in Lemma 3.1 on $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$, respectively. We then define an NBW of size $O(n_\varphi \cdot k_\varphi \cdot (n_\psi \cdot k_\psi)^2)$ for the intersection of the three NBWs $\mathcal{A}_\varphi$, $\mathcal{A}^{\mathcal{H}}_\psi$, and $\mathcal{A}^{\mathcal{H}}_{\neg\psi}$, and finally determinize it into a DPW $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ with $2^{O(n_\varphi \cdot k_\varphi \cdot (n_\psi k_\psi)^2 \log(n_\varphi \cdot k_\varphi \cdot n_\psi k_\psi))}$ states and index $O(n_\varphi \cdot k_\varphi \cdot (n_\psi k_\psi)^2)$. □

Remark 3.3 (The size of $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ for specifications and secrets given by DBWs)**.** The exponential dependency of $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ in the DPW $\mathcal{D}_\varphi$ in the construction in Lemma 3.2 follows from the exponential blow up in DPW intersection [Bok18]. When L_φ is given by a DBW $\mathcal{D}_\varphi$, we can first construct a DPW for the intersection of $\mathcal{A}^{\mathcal{H}}_\psi$ and $\mathcal{A}^{\mathcal{H}}_{\neg\psi}$, and only then take its intersection with $\mathcal{D}_\varphi$. This results in a DPW $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ of size exponential in $\mathcal{D}_\psi$, but only polynomial in $\mathcal{D}_\varphi$.

We can now solve synthesis with privacy for LTL formulas.

Theorem 3.4 (Synthesis with privacy, LTL)**.** *Given two disjoint finite sets I and O, LTL formulas φ and ψ over $I \cup O$, a cost function $\mathsf{cost} : I \cup O \to \mathbb{N}$, and a budget $b \in \mathbb{N}$, deciding whether $\langle \varphi, \psi, \mathsf{cost}, b\rangle$ is realizable with privacy is 2EXPTIME-complete.*

Proof. We start with the upper bound. Given φ, ψ, cost, and b, we go over all $\mathcal{H} \subseteq I \cup O$ such that $\mathsf{cost}(\mathcal{H}) \leq b$, construct the DPW $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ defined in Lemma 3.2, and check whether $L(\mathcal{D}^{\mathcal{H}}_{\varphi,\psi})$ is realizable. Since realizability of a DPW with n states and index k can be solved in time at most $O(n^k)$ [CJK+17], the 2EXPTIME upper bound follows from $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ having $2^{2^{O(|\varphi|+|\psi|)}}$ states and index $2^{O(|\varphi|+|\psi|)}$.

For the lower bound, we describe a reduction from LTL synthesis with no privacy. Note that adding to a specification φ a secret `T` or `F` does not work, as an observer knows its satisfaction value. It is easy, however, to add a secret that is independent of the specification. Specifically, given a specification φ over $I \cup O$, let $O' = O \cup \{p\}$, where p is a fresh signal not in $I \cup O$. Consider the secret $\psi = p$ and a cost function with $\mathsf{cost}(p) = 0$. Clearly, an (I/O)-transducer $\mathcal{T}$ realizes φ iff the (I/O')-transducer $\mathcal{T}'$ that agrees with $\mathcal{T}$ and always assigns `F` to p, realizes φ and $\{p\}$-hides ψ. Conversely, for an I/O'-transducer $\mathcal{T}'$, let $\mathcal{T}$ be the (I/O)-transducer obtained from $\mathcal{T}$ by ignoring the assignments to p. Clearly, $\mathcal{T}'$ $\{p\}$-hides ψ. In addition, as φ does not refer to p, we have that $\mathcal{T}'$ realizes φ iff $\mathcal{T}$ realizes φ. Thus, φ is realizable iff $\langle \varphi, \psi, \mathsf{cost}, 0\rangle$ is realizable with privacy. □

Remark 3.5 (Solving privacy with multiple and conditional secrets)**.** Recall that for a set of secrets $S = \{\psi_1, \psi_2, \ldots, \psi_k\}$, an (I/O)-transducer $\mathcal{T}$ realizes $\langle \varphi, S, \mathcal{H}\rangle$ with privacy if it realizes φ and $\mathcal{H}$-hides ψ_i, for all $i \in [k]$. It is easy to extend Theorem 3.4 to the setting of multiple secrets by replacing the DPW $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ by a DPW obtained by determinizing the product of $\mathcal{A}_\varphi$ with automata $\mathcal{A}^{\mathcal{H}}_{\psi_i}$ and $\mathcal{A}^{\mathcal{H}}_{\neg\psi_i}$, for all $1 \leq i \leq k$.

As for conditional secrets, recall that a computation should satisfy the specification, and from the point of view of an observer, either the trigger is not triggered, thus $\pi \in L(\mathcal{A}^{\mathcal{H}}_{\neg\theta})$,

or the secret is hidden, thus $\pi \in L(\mathcal{A}^{\mathcal{H}}_{\psi}) \cap L(\mathcal{A}^{\mathcal{H}}_{\neg\psi})$. Accordingly, we need to construct a deterministic automaton for $L(\mathcal{A}_{\varphi}) \cap (L(\mathcal{A}^{\mathcal{H}}_{\neg\theta}) \cup L(\mathcal{A}^{\mathcal{H}}_{\psi}) \cup L(\mathcal{A}^{\mathcal{H}}_{\neg\psi}))$. This can done by determinizing an NBW that is defined on top of the product of $\mathcal{A}_{\varphi}, \mathcal{A}^{\mathcal{H}}_{\neg\theta}, \mathcal{A}^{\mathcal{H}}_{\psi}$, and $\mathcal{A}^{\mathcal{H}}_{\neg\psi}$.

While the complexity of our algorithm is not higher than that of LTL synthesis with no privacy, it would be misleading to state that handling of privacy involves no increase in the complexity. Indeed, the algorithm involved two components whose complexity may have been dominated by the doubly exponential translation of the LTL formulas to deterministic automata:

(1) A need to go over all candidate sets $\mathcal{H} \subseteq I \cup O$.
(2) A need to check that the generated transducer $\mathcal{H}$-hides the secret.

In the next two sections, we isolate these two components of synthesis with privacy and show that each of them involves an exponential complexity: the first in the number of signals and the second in the size of the secret.

3.1. **Hiding secrets is hard.** The synthesis problem for DBWs can be solved in polynomial time. Indeed, the problem can be reduced to solving a Büchi game played on top of the specification automaton. In this section we show that synthesis with privacy is EXPTIME hard even for a given set $\mathcal{H}$ of hidden signals (in fact, even a singleton set $\mathcal{H} \subseteq I$), a trivial specification, and a secret given by a DBW.

We start by showing that $\mathcal{H}$-hiding is hard even for secrets given by DBWs.

Theorem 3.6. *Given two disjoint finite sets I and O, a DBW $\mathcal{D}_{\psi}$ over $2^{I \cup O}$, and a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, deciding whether there exists an (I/O)-transducer that $\mathcal{H}$-hides $\mathcal{D}_{\psi}$ is EXPTIME-hard. The problem is EXPTIME-hard already when $\mathcal{H} \subseteq I$.*

Proof. We describe a polynomial-time reduction from NBW realizability, which is EXPTIME-hard [Rab72, HKKM02]. Given an NBW $\mathcal{A}$ over $2^{I \cup O}$, we define a set of signals $\mathcal{H}$ and a DBW $\mathcal{D}_{\psi}$ over $2^{I \cup O \cup \mathcal{H}}$, such that $L(\mathcal{A})$ is realizable iff there exists an $((I \cup \mathcal{H})/O)$-transducer that $\mathcal{H}$-hides $L(\mathcal{D}_{\psi})$.

Let $\mathcal{A} = \langle 2^{I \cup O}, Q, q_0, \delta, \alpha \rangle$. W.l.o.g, we assume that $\mathcal{A}$ has a single initial state and that every word in $(2^{I \cup O})^{\omega}$ has at least one rejecting run in $\mathcal{A}$. The latter can be achieved, for example, by adding a nondeterministic transition from the initial state to a rejecting sink upon any assignment $i \cup o \in 2^{I \cup O}$. Let $\mathcal{H}$ be a set of signals that encode Q. Thus, each assignment $s \in 2^{\mathcal{H}}$ is associated with a single state in Q. We refer to a letter in $2^{I \cup O \cup \mathcal{H}}$ as a pair $\langle \sigma, q \rangle \in 2^{I \cup O} \times Q$, and we view a word in $(2^{I \cup O \cup \mathcal{H}})^{\omega}$ as the combination $w \oplus r$, of a word $w \in (2^{I \cup O})^{\omega}$ with a word $r \in Q^{\omega}$. Formally, for $w = \sigma_0 \cdot \sigma_1 \cdots \in (2^{I \cup O})^{\omega}$ and $r = r_1 \cdot r_2 \cdots \in Q^{\omega}$, let $w \oplus r = \langle \sigma_0, r_1 \rangle \cdot \langle \sigma_1, r_2 \rangle \cdots \in (2^{I \cup O \cup \mathcal{H}})^{\omega}$. Then, we define $\mathcal{D}_{\psi}$ so that $L(\mathcal{D}_{\psi}) = \{w \oplus r \in (2^{I \cup O \cup \mathcal{H}})^{\omega} : \text{the sequence } q_0 \cdot r \text{ is an accepting run of } \mathcal{A} \text{ on } w\}$. Note that since every word in $(2^{I \cup O})^{\omega}$ has at least one rejecting run in $\mathcal{A}$, then every word $w \in L(\mathcal{A})$ has at least one word $r^+ \in Q^{\omega}$ such that $w \oplus r^+ \in L(\mathcal{D}_{\psi})$ and at least one word $r^- \in Q^{\omega}$ such that $w \oplus r^- \notin L(\mathcal{D}_{\psi})$.

Formally, $\mathcal{D}_{\psi} = \langle 2^{I \cup O \cup \mathcal{H}}, Q, q_0, \delta', \alpha \rangle$ has the same state space and acceptance condition as $\mathcal{A}$, and it uses the Q-component of each letter in order to resolve the nondeterministic choices in $\mathcal{A}$. Thus, the transitions function $\delta' : Q \times 2^{I \cup O \cup \mathcal{H}} \to Q$ is defined as follows. For every state $q \in Q$ and letter $\langle \sigma, s \rangle \in 2^{I \cup O \cup \mathcal{H}}$, we have that $\delta'(q, \langle \sigma, s \rangle) = s$ if $s \in \delta(q, \sigma)$, and otherwise $\delta'(q, \langle \sigma, s \rangle) = \bot$.

We prove that indeed $L(\mathcal{D}_\psi)$ accepts exactly all words $w \oplus r$ such that $q_0 \cdot r$ is an accepting run of $\mathcal{A}$ on w. By definition of δ', it holds that r' is a run of $\mathcal{D}_\psi$ over $w \oplus r$ iff $r' = q_0 \cdot r$, and $q_0 \cdot r$ is a run of $\mathcal{A}$ over w. Hence, a run $r' = q_0 \cdot r$ of $\mathcal{D}_\psi$ over $w \oplus r$ is accepting, iff $\mathit{inf}(r') \cap \alpha \neq \varnothing$, iff $r' = q_0 \cdot r$ is an accepting run of $\mathcal{A}$ over w, and we are done. □

Note that in the proof of Theorem 3.6, we could have defined $\mathcal{H}$ so that it resolves the nondeterminism in $\mathcal{A}$ in a more concise way. In particular, if we assume that the nondeterminism degree in $\mathcal{A}$ is at most 2, then a set $\mathcal{H}$ of size 1 can resolve the nondeterminizm of δ. Hence, as NBW synthesis is EXPTIME-hard already for NBWs with branching degree 2 (this follows from the fact that a bigger branching degree can be decomposed along several transitions), EXPTIME hardness holds already when hiding a single input signal.

Theorem 3.7 (Synthesis with privacy, DPWs)**.** *Given two disjoint finite sets I and O, DPWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$ over $2^{I \cup O}$, and a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, deciding whether $\langle \mathcal{D}_\varphi, \mathcal{D}_\psi, \mathcal{H} \rangle$ is realizable with privacy is EXPTIME-complete. Moreover, hardness holds already when the specification is trivial and the secret is given by a DBW.*

Proof. For the upper bound, we solve the synthesis problem for the DPW $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ defined in Lemma 3.2. As specified there, the size of $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$ is exponential in the size and index of both $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$, and its index is polynomial in the size and index of $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$. Membership in EXPTIME then follows from the complexity of the synthesis problem for DPWs [BCJ18].

For the lower bound, fix a DBW $\mathcal{D}_\mathtt{T}$ such that $L(\mathcal{D}_\mathtt{T}) = (2^{I \cup O})^\omega$. Then, it is easy to see that $\langle \mathcal{D}_\mathtt{T}, \mathcal{D}_\psi, \mathcal{H} \rangle$ is realizable with privacy iff there is an (I/O)-transducer that $\mathcal{H}$-hides $\mathcal{D}_\psi$. Thus, hardness in EXPTIME follows from Theorem 3.6. □

3.2. **Searching for a set of signals to hide is hard.** Another component in the algorithm that is dominated by the doubly-exponential translation of LTL to DPWs is the need to go over all subsets of $I \cup O$ in a search for the set $\mathcal{H}$ of signals to hide. Trying to isolate the influence of this search, it is not enough to consider specifications and secrets that are given by DBWs, as synthesis with privacy is EXPTIME-hard already for a given set $\mathcal{H}$, and so again, the complexity of the search is dominated by the complexity of the synthesis problem. Fixing the size of the secret, which is the source of the exponential complexity, does not work either, as it also fixes the number of signals that we may need to hide. We address this challenge by moving to an even simpler setting for the problem, namely synthesis with privacy of a *closed* system. We are going to show that in this setting, the search for $\mathcal{H}$ is the only non-polynomial component in the algorithm.

In the closed setting, all signals are controlled by the system, namely $I = \varnothing$. Consequently, each transducer has a single computation, and realizability coincides with satisfiability. In particular, for $I = \varnothing$, we have that $\langle L_\varphi, L_\psi, \mathcal{H} \rangle$ is realizable with privacy iff there exists a word $w \in L_\varphi$, for which there exist two words $w^+, w^- \in \mathsf{noise}_\mathcal{H}(w)$ such that $w^+ \in L_\psi$ and $w^- \notin L_\psi$. We show that while synthesis with privacy in the closed setting can be solved in polynomial time for a given set $\mathcal{H}$ of hidden signals, it is NP-complete when $\mathcal{H}$ is not given, even when the function *cost* is uniform.

We start with the case $\mathcal{H}$ is given.

Theorem 3.8. *Given a finite set O of output signals, a set $\mathcal{H} \subseteq O$ of hidden signals, and DPWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$ over 2^O, deciding whether $\langle \mathcal{D}_\varphi, \mathcal{D}_\psi, \mathcal{H} \rangle$ is realizable with privacy can be done in polynomial time.*

Proof. First, we complement $\mathcal{D}_\psi$, which results in a DPW $\mathcal{D}_{\neg\psi}$ of the same size, and of index $k+1$, where k is the index of $\mathcal{D}_\psi$. Then, we translate $\mathcal{D}_\varphi$, $\mathcal{D}_\psi$ and $\mathcal{D}_{\neg\psi}$ into equivalent NBWs $\mathcal{A}_\varphi$, $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$, respectively. All three NBWs can be defined in size that is polynomial in their deterministic DPW counterpart. Let $\mathcal{A}_\psi^{\mathcal{H}}$ and $\mathcal{A}_{\neg\psi}^{\mathcal{H}}$ be NBWs obtained by applying the construction in Lemma 3.1 on $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$, respectively. By Lemma 3.1, the NBWs $\mathcal{A}_\psi^{\mathcal{H}}$ and $\mathcal{A}_{\neg\psi}^{\mathcal{H}}$ have the same number of states as $\mathcal{A}_\psi$ and $\mathcal{A}_{\neg\psi}$ respectively. Let $\mathcal{N}$ be an NBW for the intersection $L(\mathcal{A}_\varphi) \cap L(\mathcal{A}_\psi^{\mathcal{H}}) \cap L(\mathcal{A}_{\neg\psi}^{\mathcal{H}})$. Note that $\mathcal{N}$ can be defined with size that is polynomial in $\mathcal{A}_\varphi$, $\mathcal{A}_\psi^{\mathcal{H}}$ and $\mathcal{A}_{\neg\psi}^{\mathcal{H}}$. Moreover, $\mathcal{N}$ accepts a word w iff $w \in L(\mathcal{A}_\varphi)$, and there exist two words $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(w)$ such that $w^+ \in L(\mathcal{A}_\psi)$ and $w^- \notin L(\mathcal{A}_\psi)$. Thus, realizability with privacy of $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H}\rangle$ can be reduced to the nonemptiness of $\mathcal{N}$, which can be decided in polynomial time. □

We continue to the case $\mathcal{H}$ should be searched.

Theorem 3.9. *Given a finite set of output signals O, DPWs $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$ over 2^O, a hiding cost function $\mathsf{cost} : O \to \mathbb{N}$, and a budget $b \in \mathbb{N}$, deciding whether $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathsf{cost}, b\rangle$ is realizable with privacy is NP-complete. Moreover, hardness holds already when the specification and secret are given by DBWs.*

Proof. For the upper bound, a nondeterministic Turing machine can guess a set $\mathcal{H} \subseteq O$, check whether $\mathsf{cost}(\mathcal{H}) \leq b$, and, by Theorem 3.8, check in polynomial time whether $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H}\rangle$ is realizable with privacy.

For the lower bound, we describe a polynomial-time reduction from the *vertex-cover* problem. In this problem, we are given an undirected graph $G = \langle V, E\rangle$ and $k \geq 1$, and have to decide whether there is a set $S \subseteq V$ such that $|S| \leq k$ and for every edge $\{v, u\} \in E$, we have that $E \cap S \neq \varnothing$. Given an undirected graph $G = \langle V, E\rangle$, with $E = \{e_1, e_2, \ldots, e_m\}$, we consider a closed setting with $O = V$ and construct DBWs $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$ over the alphabet 2^V such that for all $\mathcal{H} \subseteq V$, it holds that $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H}\rangle$ is realizable with privacy iff $\mathcal{H}$ is a vertex cover of G. Accordingly, there is a vertex cover of size k in G iff $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathsf{cost}, k\rangle$ is realizable with privacy for the uniform cost function that assigns 1 to all signals in O.

We define $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$ over the alphabet 2^V as follows. The DBW $\mathcal{A}_\varphi$ is a 2-state DBW that accepts the single word $\varnothing^\omega$. The DBW $\mathcal{A}_\psi = \langle 2^V, Q, q_1, \delta, \alpha\rangle$ for the secret is defined as follows. The set of states is $Q = \{q_1, q_2, \ldots, q_{m+1}\}$. The set of accepting states is $\alpha = \{q_{m+1}\}$. The transition function δ is defined for all $S \subseteq O$ and $i \leq m$ as follows. If $S \cap e_i \neq \varnothing$, then $\delta(q_i, S) = q_{i+1}$, and otherwise, $\delta(q_i, S) = \bot$. Finally, $\delta(q_{m+1}, S) = q_{m+1}$, for all $S \subseteq V$.

Intuitively, words in $L(\mathcal{A}_\psi)$ encode vertex covers of G. Indeed, if $w = S_1 \cdot S_2 \cdot S_2 \cdot \ldots \in L(\mathcal{A}_\psi)$, then for all $i \leq m$, we have that $S_i \cap e_i \neq \varnothing$. Thus, if for all $i \leq m$ we set $v_i \in V$ to be some vertex in $S_i \cap e_i$, then $\{v_i, \ldots, v_m\}$ is a vertex cover of G.

Formally, we prove that $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H}\rangle$ is realizable with privacy iff $\mathcal{H}$ is a vertex cover of G. Consider first a vertex cover $\mathcal{H}$ of G. Let $\mathcal{T}$ be a transducer that generates the computation $\varnothing^\omega$ and hides the signals in $\mathcal{H}$. In every step of the computation, the observer sees that the signals in $V \setminus \mathcal{H}$ are assigned `F`. Assigning `F` to all the signals in $\mathcal{H}$ results in a word w^- that is not in $L(\mathcal{A}_\psi)$. On the other hand, since $\mathcal{H}$ is a vertex cover, i.e., it intersects every $e \in E$, assigning `T` to all the signals in $\mathcal{H}$ results in a word w^+ that is in $L(\mathcal{A}_\psi)$. Thus, $\mathcal{T}$ realizes $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H}\rangle$ with privacy.

For the other direction, consider a set $\mathcal{H} \subseteq V$ that is not a vertex cover of G. Consider a transducer $\mathcal{T}$ that realizes $L(\mathcal{A}_\varphi)$ and hides the set $\mathcal{H}$ of signals. Since $\mathcal{T}$ realizes $L(\mathcal{A}_\varphi)$, it

generates the computation $\varnothing^\omega$. Since $\mathcal{H}$ is not a vertex cover of G, there is $0 \leq i \leq m$ such that $\mathcal{H} \cap \{e_i\} = \varnothing$. Then, at every step of the computation, the observer sees the values of the signals in e_i. In particular, at the i-th step of the computation, the observer sees that the signals in e_i are assigned `F`, which reveals that the computation is not in $L(\mathcal{A}_\psi)$. Thus, $\mathcal{T}$ does not realize $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H} \rangle$ with privacy, and we are done. □

4. Bounded Synthesis with Privacy

In the general synthesis problem, there is no bound on the size of the generated system. It is not hard to see that if a system that realizes the specification exists, then there is also one whose size is bounded by the size of a deterministic automaton for the specification. For the case of LTL specifications, this gives a doubly-exponential bound on the size of the generated transducer, which is known to be tight [Ros92]. In [SF07], the authors suggested to study *bounded synthesis*, where the input to the problem includes also a bound on the size of the system. The bound not only guarantees the generation of a small system, if it exists, but also reduces the complexity of the synthesis problem and gives rise to a symbolic implementation and further extensions [Ehl10, KLVY11]. In particular, for LTL, it is easy to see that bounded synthesis can be solved in PSPACE, as one can go over and model-check all candidate systems. For specifications in DPW, the bound actually increases the complexity, as going over all candidates results in an algorithm in NP.

In this section we study bounded synthesis with privacy. As in traditional synthesis, the hope is to both reduce the complexity of the problem and to end up with smaller systems. In addition to a specification L_φ, a secret L_ψ, and a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, we are given a bound $n \in \mathbb{N}$, represented in unary, and we are asked to construct an (I/O)-transducer with at most n states that realizes $\langle L_\varphi, L_\psi, \mathcal{H} \rangle$ with privacy, or to determine that no such transducer exists. As in the unbounded case, we can define the problem also with respect to a hiding cost function and a budget.

4.1. Hiding secrets by a bounded system is hard.

We first show that hiding secrets in a bounded setting is hard. In fact, the complexity of hiding goes beyond the complexity of bounded synthesis with no privacy already in the case the specification and secrets are given by LTL formulas. In the case of DBWs and DPWs, hiding is also more complex than bounded synthesis without privacy, but the difference is not significant.

The key idea in both results is similar to the one in the proof of Theorem 3.6. There, we reduce realizability of NBWs to hiding of secrets given by DBWs. Essentially, we use the hidden signals to imitate nondeterminism. Here, with a bound on the size of the system, we cannot reduce from realizability, as the problem has the flavor of model checking many candidates. Accordingly, we reduce from universality, either in the form of LTL formulas with universally-quantified atomic propositions, or in the form of language-universality for NBWs.

Theorem 4.1. *Given two disjoint finite sets I and O, an LTL formula ψ over $2^{I \cup O}$, a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, and a bound $n \geq 1$, given in unary, deciding whether there exists an (I/O)-transducer of size at most n that $\mathcal{H}$-hides ψ is EXPSPACE-hard. The problem is EXPSPACE-hard already when $n = 1$.*

Proof. We describe a reduction from satisfiability of AQLTL, namely LTL with universally-quantified atomic propositions, which is known to be EXPSPACE-complete [SVW87]. An AQLTL formula is of the form $\forall p_1, p_2, \ldots, p_k \theta$, where θ is an LTL formula over the set of atomic propositions AP, and $A = \{p_1, p_2, \ldots, p_k\} \subseteq AP$. The satisfiability problem is to determine whether there is $w \in (2^{AP})^\omega$ such that for all $w' \in \mathsf{noise}_A(w)$, we have that $w' \models \theta$.

Consider an AQLTL formula $\forall p_1, p_2, \ldots, p_k \theta$. Let $A = \{p_1, p_2, \ldots, p_k\} \subseteq AP$. Let $I = AP \setminus A$ and $O = \mathcal{H} = A \cup \{p\}$, where p is a fresh signal not in AP. Also, let $\psi = \theta \vee p$. We claim that $\forall p_1, p_2, \ldots, p_k \theta$ is not satisfiable iff there is a single state (I/O)-transducer that $\mathcal{H}$-hides ψ.

Assume first that there is a single state (I/O)-transducer $\mathcal{T}$ that $\mathcal{H}$-hides ψ. Then, for every $w_I \in (2^{AP \setminus A})^\omega$, there is $w' \in \mathsf{noise}_A(\mathcal{T}(w_I))$ such that $w' \not\models \psi$. In particular $w' \not\models \theta$. Thus, $\forall p_1, p_2, \ldots, p_k \theta$ is not satisfiable.

Assume now that $\forall p_1, p_2, \ldots, p_k \theta$ is not satisfiable. Thus, for every $w \in (2^{AP})^\omega$, there is $w' \in \mathsf{noise}_A(w)$ such that $w' \not\models \theta$. Also, as $p \in \mathcal{H}$, there are $w^+ \in \mathsf{noise}_{\{p\}}(w')$ and $w^- \in \mathsf{noise}_{\{p\}}(w')$ such that $w^+ \models p$ and $w^- \not\models p$. Note that $w^+, w^- \in \mathsf{noise}_A(w)$ and that $w^+ \models \psi$ whereas $w^- \not\models \psi$. It follows that all (I/O)-transducers $\mathcal{H}$-hides ψ, and we are done. □

Theorem 4.2. *Given two disjoint finite sets I and O, a DBW $\mathcal{D}_\psi$ over $2^{I \cup O}$, a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, and a bound $n \geq 1$, represented in unary, deciding whether there exists an (I/O)-transducer of size at most n that $\mathcal{H}$-hides $\mathcal{D}_\psi$ is PSPACE-hard. The problem is PSPACE-hard already when $n = 1$.*

Proof. We describe a reduction from the NBW universality problem, which is PSPACE-hard [SVW87]. Let $\mathcal{N} = \langle 2^I, Q, q_0, \delta, \alpha \rangle$ be an NBW and w.l.o.g assume that for all $w \in (2^I)^\omega$, there is a rejecting run of $\mathcal{N}$ over w. Let $O = \mathcal{H}$ be a set of output signals that encode the states in Q. For convenience, we abuse notation and refer to elements in 2^O as states $q \in Q$. Let $\mathcal{D}_\psi = \langle 2^{I \cup O}, Q, q_0, \delta', \alpha \rangle$ be the DBW over the alphabet $2^{I \cup O}$, where the deterministic transition function $\delta' : Q \times 2^{I \cup O} \to Q$ is defined for all $q \in Q$, $o \in 2^O$ and $i \in 2^I$, by $\delta'(q, i \cup o) = o$ if $o \in \delta(q, i)$, and $\delta'(q, i \cup o) = \bot$ otherwise. Clearly, the size of $\mathcal{D}_\psi$ is linear in that of $\mathcal{N}$. Also, by its definition, we have that $\mathcal{D}_\psi$ accepts a word $\pi \in (2^{I \cup O})^\omega$ iff $r = q_0 \cdot \mathsf{hide}_I(\pi)$ is an accepting run of $\mathcal{N}$ over the word $w_I = \mathsf{hide}_{\mathcal{H}}(\pi)$.

Let $o \in 2^O$ be some arbitrary output assignment, and consider the single state (I/O)-transducer $\mathcal{T}_o = \langle I, O, \{s_0\}, s_0, \eta, \tau \rangle$, where $\eta(s_0, i) = s_0$ for all $i \in 2^I$ and $\tau(s_0) = o$. We claim that $\mathcal{N}$ is universal iff $\mathcal{T}_o$ $\mathcal{H}$-hides $\mathcal{D}_\psi$.

First, as every word has a rejecting run in $\mathcal{N}$, then $\mathcal{N}$ is universal iff for all $w_I \in (2^I)^\omega$, there exist $\pi^+, \pi^- \in \mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I))$ such that $r^+ = q_0 \cdot \mathsf{hide}_I(\pi^+)$ is an accepting run of $\mathcal{N}$ over w_I, and $r^- = q_0 \cdot \mathsf{hide}_I(\pi^-)$ is a rejecting run of $\mathcal{N}$ over w_I. Thus, $\pi^+ \in L(\mathcal{D}_\psi)$ and $\pi^- \notin L(\mathcal{D}_\psi)$. It follows that $\mathcal{N}$ is universal iff $\mathcal{T}_o$ $\mathcal{H}$-hides $\mathcal{D}_\psi$. Since $o \in 2^O$ is arbitrary, we conclude that $\mathcal{N}$ is universal iff there exists a single state (I/O)-transducer $\mathcal{T}$ that $\mathcal{H}$-hides $\mathcal{D}_\psi$. □

4.2. **Solving bounded synthesis with privacy.** We can now present the tight complexity for bounded synthesis with privacy for both types of specification formalisms. For the upper bounds, we construct an NGBW $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$ that accepts exactly all words that satisfy φ and hide ψ, and search for an (I/O)-transducer of size n whose language is contained in that of $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$.

Theorem 4.3 (Bounded synthesis with privacy)**.** *Given two disjoint finite sets I and O, specification L_φ and secret L_ψ over $I \cup O$, a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, and a bound $n \in \mathbb{N}$, represented in unary, deciding whether there is an (I/O)-transducer with at most n states that realizes $\langle L_\varphi, L_\psi, \mathcal{H}\rangle$ with privacy is PSPACE-complete for L_φ and L_ψ given by DPWs, and is EXPSPACE-complete for L_φ and L_ψ given by LTL formulas. Hardness in PSPACE holds already for DBWs.*

Proof. We start with the upper bound for the case L_φ and L_ψ are given by DPWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$. As described in Lemma 3.2, given $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$, it is possible to construct an NGBW $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$ that accepts exactly all words $w \in L(\mathcal{D}_\varphi)$ for which there exists $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(w)$ with $w^+ \in L(\mathcal{D}_\psi)$ and $w^- \notin L(\mathcal{D}_\psi)$. Moreover, $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$ is of constant index $k = 3$, and of size polynomial in the size and index of $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$.

Once we have constructed $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$, we can go through all (I/O)-transducers of size n and to check if for one of them it holds that all its computations satisfy $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$. This reduces to checking the containment of the language of each (I/O)-transducer of size n in that of $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$, which can be done in PSPACE [SVW87].

When we start with three NBWs $\mathcal{A}_\varphi$, $\mathcal{A}_\psi$, and $\mathcal{A}_{\neg\psi}$, for L_φ, L_ψ, and $L_{\neg\psi}$, the construction of $\mathcal{N}^{\mathcal{H}}_{\varphi,\psi}$ is polynomial. Hence, an EXPSPACE upper bound for LTL follows from the exponential translation of LTL formulas to NBWs.

For the lower bounds, note that by taking $L_\varphi = (2^{I\cup O})^\omega$ (that is, $\mathtt{T}$ for an LTL formula and the universal $\mathcal{D}_{\mathtt{T}}$ for DBW), we have that there is a transducer with n states that realizes $\langle L_\varphi, L_\psi, \mathcal{H}\rangle$ with privacy iff there is a transducer with n states that $\mathcal{H}$-hides L_ψ. Hence, the lower bounds follow from Theorems 4.1 and 4.2. □

5. Privacy with Certificates

In this section we study the problem of synthesizing systems that not only maintain privacy against an observer in the usual sense, but also generate for the user a *certificate for privacy*. This certificate is a computation that agrees with the real computation on the visible signals, yet it disagrees with it on the hidden signals in a way that induces a different satisfaction value for the secret. Accordingly, the certificate witnesses that the secret is indeed hidden from the observer.

Recall that a secret ψ is $\mathcal{H}$-hidden in a computation $w \in (2^{I\cup O})^\omega$ iff there exist two computations $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(w)$, such that $w^+ \models \psi$ and $w^- \models \neg\psi$. Clearly, if ψ is indeed $\mathcal{H}$-hidden in w, then we can always use w as one of the witnesses w^+ or w^-, according to whether ψ is satisfied or not in w. Hence, ψ is $\mathcal{H}$-hidden in w iff there exists $w' \in \mathsf{noise}_{\mathcal{H}}(w)$ such that ($w \models \psi$ iff $w' \models \neg\psi$) holds. Moreover, as all words in $\mathsf{noise}_{\mathcal{H}}(w)$ agree with w on $\mathcal{V}$, the witness word w' is identified by its restriction $w'|_{\mathcal{H}}$ to the signals in $\mathcal{H}$. Accordingly, we may simply require the existence of a word $w_{\mathcal{H}} \in (2^{\mathcal{H}})^\omega$, such that the witness word $w' \in \mathsf{noise}_{\mathcal{H}}(w)$ for ψ being $\mathcal{H}$-hidden in w is such that $w'|_{\mathcal{H}} = w_{\mathcal{H}}$. Thus, $w_{\mathcal{H}}$ describes a sequence of assignments to the signals in $\mathcal{H}$ with which the value of the secret ψ is different from its value in w. We call $w_{\mathcal{H}}$ a *certificate for the privacy of ψ in w*. When ψ and w are clear from the context we simply refer to $w_{\mathcal{H}}$ as a *certificate for privacy*.

We can thus reformulate the notion of privacy against an observer as follows. An (I/O)-transducer $\mathcal{T}$ is said to $\mathcal{H}$-hide ψ if for all $w_I \in (2^I)^\omega$, there exists a certificate $w_{\mathcal{H}} \in (2^{\mathcal{H}})^\omega$ for ψ in $\mathcal{T}(w_I)$. Note that a transducer $\mathcal{T}$ that $\mathcal{H}$-hides ψ does not have to provide a

certificate for ψ. An observer that tries to reveal the value of ψ is not going to succeed, and the user is guaranteed that such a certificate exists, but there is no requirement for $\mathcal{T}$ to generate it. Adding a certification may, however, be of interest: it assures the user that ψ is indeed hidden. In this section we study the problem of synthesizing a system that not only realizes $\langle \varphi, \psi, \mathcal{H} \rangle$ in the usual sense, but also generates for every input word an appropriate certificate. As we discuss below, adding certificates makes specifications and secrets more difficult to realize. That is, there are specification and secrets that become unrealizable with privacy once requiring the generation of certificates. Essentially, this follows from the fact that the certificate $w_{\mathcal{H}}$ has to be generated in an on-line manner.

Before we formally define the setting of *certifiable privacy*, note that the complication added to the problem when requiring the online generation of a certificate is analogous to the difficulty in usual synthesis (with no privacy considerations) that is added on top of *universal satisfiability*. In universal satisfiability we ask whether a specification is such that for all $w_I = i_0 \cdot i_1 \cdot i_2 \cdots \in (2^I)^\omega$, there exists $w_O = o_0 \cdot o_1 \cdot o_2 \cdots \in (2^O)^\omega$, such that $w_I \oplus w_O = (i_0 \cup o_0) \cdot (i_1 \cup o_1) \cdot (i_2 \cup o_2) \cdots$ satisfies the specification. Then, in the synthesis problem, the inputs in w_I are read one after the other, and we require that the output assignments in w_O are generated online. Privacy (with no certificates) has the flavor of universal satisfiability with respect to the secret, while the setting with certification is closer to realizability as it is no longer enough that a certificate exists, instead we need to construct one on-line. It is hence not surprising that generating certificates makes the task of preserving privacy more difficult to realize. Indeed, we already know this is the case for realizability and universal satisfiability: There are specifications that are universally satisfiable yet not realizable. A simple example for such a specification is $\varphi = o \leftrightarrow Xi$, which is clearly universally satisfiable, but cannot be realized as the system needs to predict the value of the input i in order to generate the correct value for o in the present. In Example 5.1 we use similar formula to demonstrate the gap between synthesis with privacy, with and without certificates.

We continue and formally define the setting for synthesis with *certified privacy*. A *certifying transducer* $\mathcal{C} = \langle I, O, \mathcal{H}, S, s_0, \rho, \tau, c \rangle$ is defined similarly to a standard (I/O)-transducer but with a distinguished set of hidden signals $\mathcal{H}$ and an additional labeling function $c : S \to 2^{\mathcal{H}}$. For every input word $w_I = i_0 \cdot i_1 \cdot i_2 \cdots \in (2^I)^\omega$, the run $r(w_I) = s_0 \cdot s_1 \cdot s_2 \cdots \in S^\omega$, and the computation $\mathcal{C}(w_I) = (i_0 \cup \tau(s_0)) \cdot (i_1 \cup \tau(s_1)) \cdot (i_2 \cup \tau(s_2)) \cdots \in (2^{I \cup O})^\omega$, are defined as in a standard (I/O)-transducer. Then, the *certificate* $c(w_I) \in (2^{\mathcal{H}})^\omega$, for $\mathcal{C}(w_I)$, is defined to be $c(w_I) = c(s_0) \cdot c(s_1) \cdot c(s_2) \cdots$. Finally, the *alternative computation* $\mathcal{C}'(w_I) \in (2^{I \cup O})^\omega$, is the unique computation that agrees with $\mathcal{C}(w_I)$ on the visible signals and with $c(w_I)$ on the hidden signals. Formally, $\mathcal{C}'(w_I)|_{\mathcal{V}} = \mathcal{C}(w_I)|_V$ and $\mathcal{C}'(w_I)|_{\mathcal{H}} = c(w_I)$.

Given LTL formulas φ and ψ over $I \cup O$, and a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, we say that $\mathcal{C}$ realizes $\langle \varphi, \psi, \mathcal{H} \rangle$ with *certified privacy*, if $\mathcal{C}$ realizes φ and for all $w_I \in (2^I)^\omega$, it holds that $c(w_I)$ is a certificate for ψ being $\mathcal{H}$-hidden in $\mathcal{C}(w_I)$. I.e., for all $w_I \in (2^I)^\omega$ it holds that $\mathcal{C}(w_I) \models \varphi$ and $(\mathcal{C}(w_I) \models \psi$ iff $\mathcal{C}'(w_I) \models \neg\psi)$.

Finally, in the *synthesis with certified privacy problem*, we are given $\langle \varphi, \psi, \mathcal{H} \rangle$, and we need to construct a certifying transducer $\mathcal{C}$ that realizes $\langle \varphi, \psi, \mathcal{H} \rangle$ with certified privacy, or determine that no such transducer exists.

Example 5.1. Let $I = \{i\}$, $\mathcal{H} = O = \{o\}$, and $\psi = G(o \leftrightarrow Xi)$. As we explain below, ψ is $\mathcal{H}$-hidden in every (I/O)-transducer, yet there exists no certifying transducer that realizes $\langle \mathtt{T}, \psi, \mathcal{H} \rangle$ with certified privacy.

Consider a word $w = (i_0 \cup o_0) \cdot (i_1 \cup o_1) \cdot (i_2 \cup o_2) \cdots \in (2^I)^\omega$. Clearly, $w \models \psi$ iff for all $j \geq 0$ we have that $o_j = \emptyset$ iff $i_{j+1} = \emptyset$. This implies that ψ is $\mathcal{H}$-hidden in (I/O)-transducers. Indeed, consider a transducer $\mathcal{T}$, then for every input word $w_I \in (2^I)^\omega$ there exist a unique word $w^+ \in \mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I))$ that satisfies the aforementioned criterion, and thus satisfy ψ, while all other words in $\mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I))$ violate ψ. Note that then it follows that a certifying transducer $\mathcal{C}$ should manage to generate the exact unique w^+ either as the real computation $\mathcal{C}(w_I)$ or as the alternative computation $\mathcal{C}'(w_I)$. We show that this is impossible. Specifically, that the environment can force both computations to violate ψ.

Consider a certifying transducer $\mathcal{C}$. We prove that there exists an input word $w_I \in (2^I)^\omega$ such that both $\mathcal{C}(w_I)$ and $\mathcal{C}'(w_I)$ do not satisfy ψ. The word w_I is defined so that the criterion $i_{j+1} = \emptyset$ iff $o_j = \emptyset$ is violated in $\mathcal{C}(w_I)$ at time $j = 0$ and in $\mathcal{C}'(w_I)$ at time $j = 1$. Thus, after two time steps both the real computation and the alternative computation violate the secret ψ, implying that $\mathcal{C}$ does not realize $\langle \mathtt{T}, \psi, \mathcal{H} \rangle$ with certified privacy.

First, we define $i_0 = \emptyset$. Then, let $o_0 \in 2^O$ and $o'_0 \in 2^O$ be the first real output of $\mathcal{C}$ and the first certificate output of $\mathcal{C}$, when reading the first input i_0. I.e., if $s_1 = \rho(s_0, i_0)$, then $o_0 = \tau(s_1)$ and $o'_0 = c(s_1)$. We choose $i_1 \in 2^I$ so that $i_1 = \emptyset$ iff $o_0 \neq \emptyset$ holds. Observe that now the computation of $\mathcal{C}$ already violates ψ. Let now o_1 and o'_1 be the real and certificate outputs of $\mathcal{C}$ after reading the prefix $i_0 \cdot i_1$. Similar to the previous choice of i_1, we now choose i_2 so that $i_2 = \emptyset$ iff $o'_1 \neq \emptyset$. As before, we can see that now the alternative computation generated by $\mathcal{C}$ is guaranteed to violate ψ as well. Thus, no matter how we extend $i_0 \cdot i_1 \cdot i_2$ into an infinite word w_I, it holds that both $\mathcal{C}(w_I)$ and $\mathcal{C}'(w_I)$ violate ψ. In particular, if we let $w_I = i_0 \cdot i_1 \cdot i_2 \cdot \emptyset^\omega$, then $\mathcal{C}(w_I) \models \neg\psi$ and $\mathcal{C}'(w_I) \models \neg\psi$. That is, $\mathcal{C}$ does not certify that ψ is $\mathcal{H}$-hidden and we are done. □

5.1. **Solving certified privacy.** We continue to the solution of the synthesis problem with certified privacy. As discussed above, the need to generate the certificate on-line makes some specifications and secrets non-realizable. On the positive side, it enables the use of universal co-Büchi automata, which enable a Safraless synthesis procedure [KV05b]. Thus, the same way universal satisfiability is used as a heuristic to solve the realizability problem [KSS11], our procedure here can be viewed as a heuristic to the problem of synthesis with privacy (without a certificate).

Lemma 5.2. *Consider two disjoint finite sets I and O, a subset $\mathcal{H} \subseteq I \cup O$, and LTL formula φ and ψ over the signals $I \cup O$. There exists a UCW $\mathcal{U}^{\mathcal{H}}_{\varphi,\psi}$ with alphabet $2^{I \cup O} \times 2^{\mathcal{H}}$ that accepts all words $w \oplus w'_{\mathcal{H}}$ where $w \models \varphi$ and $w'_{\mathcal{H}}$ is a certificate for ψ being $\mathcal{H}$-hidden in w. Moreover, the UCW $\mathcal{U}^{\mathcal{H}}_{\varphi,\psi}$ is of size $2^{O(|\varphi|+|\psi|)}$.*

Proof. Let $\mathcal{H}'$ be a copy of the signals in $\mathcal{H}$, and let ψ' be the LTL formula obtained from ψ when replacing every signal in $\mathcal{H}$ with its copy in $\mathcal{H}'$. Consider the LTL formula $\Theta = \varphi \wedge (\psi \leftrightarrow \neg\psi')$ above $I \cup O \cup \mathcal{H}'$. Since letters in $2^{I \cup O} \times 2^{\mathcal{H}}$ correspond to assignments to $I \cup O \cup \mathcal{H}'$, it follows immediately from the definition of Θ that $\mathcal{U}^{\mathcal{H}}_{\varphi,\psi}$ can be defined as a UCW for Θ. The latter can be constructed by constructing an NBW for $\neg\Theta$ and then dualizing into the required UCW. By [VW94], such an NBW can be defined with size $2^{O(|\Theta|)} = 2^{O(|\varphi|+|\psi|)}$, and so we are done. □

Theorem 5.3 (LTL Synthesis with proved privacy)**.** *Given two disjoint finite sets I and O, LTL formulas φ and ψ over $I \cup O$, and a set of hidden signals $\mathcal{H} \subseteq I \cup O$, deciding whether*

$\langle \varphi, \psi, \mathcal{H} \rangle$ is realizable with certified privacy is 2EXPTIME-complete, and can be decided using a Safraless procedure.

Proof. The proof is very similar to the proof of Theorem 3.4 with the only difference that the upper bound goes through UCW synthesis. By Lemma 5.2 there exists a UCW $\mathcal{U}^{\mathcal{H}}_{\varphi,\psi}$ of exponential size that recognizes all words $w \oplus w'_{\mathcal{H}} \in (2^{I \cup O} \times 2^{\mathcal{H}})^\omega$ that satisfy $\Theta = \varphi \wedge (\psi \leftrightarrow \neg\psi')$ as defined in the proof of Lemma 5.2. The first conjunct asserts that w satisfy the specification φ. Then, the conjunct $\psi \leftrightarrow \neg\psi'$ asserts that w disagrees with the alternative computation $w|_{\mathcal{V}} \oplus w'_{\mathcal{H}}$ on the secret ψ, and thus $w'_{\mathcal{H}}$ is a certificate for ψ being $\mathcal{H}$-hidden in w. We then consider the set $O' = O \cup \mathcal{H}'$ as a new set of output signals, and reduce the problem of synthesizing a system that realizes $\langle \varphi, \psi, \mathcal{H} \rangle$ with certified privacy to the UCW synthesis of an (I/O')-transducer that realizes $\mathcal{U}^{\mathcal{H}}_{\varphi,\psi}$ in the traditional sense. Note that as elaborated in the proof of Theorem 5.3, assignments to O' correspond to pairs of assignments, one to O and another to $\mathcal{H}$, and so the alphabet $2^{I \cup O} \times 2^{\mathcal{H}}$ of $\mathcal{U}^{\mathcal{H}}_{\varphi,\psi}$ is compatible with an (I/O')-transducer. In [KV05b] the authors show that UCW synthesis can be solved in exponential time while avoiding Safra's determinization procedure (in fact it avoids determinization completely), and so we obtain the required upper bound.

Finally, the correspondence between (I/O')-transducers that realize $\mathcal{U}^{\mathcal{H}}_{\varphi,\psi}$ and certifying transducers that realize $\langle \varphi, \psi, \mathcal{H} \rangle$ with certified privacy is straight forward. Indeed, a labeling function $\tau' : S \to 2^{O'}$ of an (I/O')-transducer, correspond to pair of labeling functions $\langle \tau, c \rangle$ of type $\tau : S \to 2^O$ and $c : S \to 2^{\mathcal{H}}$. Indeed, the natural bijection between $2^{O'}$ and $2^O \times 2^{\mathcal{H}}$ induces the needed decomposition of τ' to a pair $\langle \tau, c \rangle$. Formally, τ is simply the restriction of τ' to the signals in O and c is the restriction of τ' to the signals in $\mathcal{H}'$, where every signal $h \in \mathcal{H}$ is identified with its copy $h' \in \mathcal{H}'$. Thus, $c(h) = \tau(h')$ for all $h \in \mathcal{H}$. □

6. When the Observer Knows the Specification or the Transducer

In this section we study two settings in which the observer has some knowledge that may help her to evaluate the secret. In the first setting, the observer knows the specification, and in the second, she knows the transducer.

6.1. **An observer that knows the secret.**

Since all the computations of the system satisfies the specification, an observer that knows the specification φ knows that only computations that satisfy φ should be taken into account when she tries to evaluate the secret. If, for example, $\varphi \to \psi$, then ψ cannot be kept private in a setting in which the observer knows φ. Indeed, the fact φ is realized by the system reveals that ψ is satisfied. Formally, we say that $\mathcal{T}$ *realizes $\langle \varphi, \psi, \mathcal{H} \rangle$ with privacy under the knowledge of the specification* if $\mathcal{T}$ realizes φ, and for every $w_I \in (2^I)^\omega$, there exist $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(\mathcal{T}(w_I)) \cap L_\varphi$ such that $w^+ \models \psi$ and $w^- \not\models \psi$. Thus, the satisfaction of the secret ψ in a computation $\mathcal{T}(w_I)$ cannot be deduced from the observable computation $\mathsf{hide}_{\mathcal{H}}(\mathcal{T}(w_I))$ even when the observer knows that φ is satisfied in $\mathcal{T}(w_I)$. The adjustment for the definition of the problem with respect to a hiding cost function and a budget is similar.

We start by showing the analogue of Lemma 3.2 for the setting in which the observer knows the specification. The construction is similar to that of Lemma 3.2, except that now, the construction of the DPW $\mathcal{D}^{\mathcal{H}}_{\psi|\varphi}$ involves an existential projection on $\mathcal{H}$ also in the

automaton for the specification. Accordingly, the size of the DPW is exponential in both the specification and the secret even in the case they are given by DBWs.

Lemma 6.1. *Consider two disjoint finite sets I and O, a subset $\mathcal{H} \subseteq I \cup O$, and regular languages L_φ and L_ψ over the alphabet $2^{I \cup O}$. There exists a DPW $\mathcal{D}^{\mathcal{H}}_{\psi|\varphi}$ with alphabet $2^{I \cup O}$ that accepts a computation $w \in (2^{I \cup O})^\omega$ iff $w \in L_\varphi$ and there exist $w^+, w^- \in \mathsf{noise}_{\mathcal{H}}(w)$ such that $w^+ \in L_\varphi \cap L_\psi$ and $w^- \in L_\varphi \setminus L_\psi$.*

(1) *If L_φ and L_ψ are given by LTL formulas φ and ψ, then $\mathcal{D}^{\mathcal{H}}_{\psi|\varphi}$ has $2^{2^{O(|\varphi|+|\psi|)}}$ states and index $2^{O(|\varphi|+|\psi|)}$.*
(2) *If L_φ and L_ψ are given by DPWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$ with n_φ and n_ψ states, and of indices k_φ and k_ψ, then $\mathcal{D}^{\mathcal{H}}_{\psi|\varphi}$ has $2^{O((n_\varphi \cdot k_\varphi)^3 \cdot (n_\psi \cdot k_\psi)^2 \log(n_\varphi \cdot k_\varphi \cdot n_\psi \cdot k_\psi))}$ states and index $O((n_\varphi \cdot k_\varphi)^3 \cdot (n_\psi \cdot k_\psi)^2)$.*

Lemma 6.1 implies that all the asymptotic upper bounds described in Section 3 are valid also in a setting with an observer that knows the specification. Also, as the lower bounds in Theorems 3.4 and 3.7 involve secrets that are independent of the specification, they are valid for this setting too. Two issues require a consideration:

(1) The need to search for $\mathcal{H}$: the NP-hardness proof in Theorem 3.9 is no longer valid, as there, $\varphi \to \neg\psi$, and so the satisfaction value of the secret is revealed in a setting with an observer that knows the specification.
(2) The construction in Lemma 6.1 results in an algorithm that is exponential also in the specification, even when given by a DBW. On the other hand, the EXPTIME-hardness proof in Theorem 3.6 does not imply an exponential lower bound in the specification.

Below we address the two issues, providing lower bounds for a setting in which the observer knows the specification. Matching upper bounds follow the same considerations in Theorems 3.7 and 3.9, where $\mathcal{D}^{\mathcal{H}}_{\psi|\varphi}$ replaces $\mathcal{D}^{\mathcal{H}}_{\varphi,\psi}$. We start with a variant of Theorem 3.9, showing NP-hardness also in the setting of a knowledgeable observer. As mentioned above, the lower bound in the proof of Theorem 3.9 does not work when the observer knows the specification, yet, can easily be modified to work for the case of a knowledgeable observer.

Theorem 6.2. *Given a set O of output signals, a cost function $\mathsf{cost} : O \to \mathbb{N}$, a hiding budget $b \in \mathbb{N}$, and DBWs $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$ over 2^O, deciding whether there is $\mathcal{H} \subseteq O$, with $\mathsf{cost}(\mathcal{H}) \leq b$, such that $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H} \rangle$ is realizable with privacy under knowledge of the specification is NP-hard. Moreover, hardness holds already when cost is uniform.*

Proof. We describe a polynomial-time reduction from the vertex cover problem. Consider an undirected graph $G = \langle V, E \rangle$, with $E = \{e_0, e_2, \ldots, e_m\}$. We construct DBWs $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$ over the alphabet 2^V, such that for all $\mathcal{H} \subseteq V$, it holds that $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H} \rangle$ is realizable with privacy under knowledge of the specification iff $\mathcal{H}$ is a vertex cover of G. Let $L^1_\varphi = \{\varnothing^\omega\}$ and L^2_φ be the set of all words $w = \sigma_0 \cdot \sigma_1 \cdot \sigma_2 \cdots \in (2^V)^\omega$ such that $\sigma_i \cap e_i \neq \varnothing$, for all $0 \leq i \leq m$. Now, let $L_\varphi = L^1_\varphi \cup L^2_\varphi$, and $L_\psi = L^2_\varphi$. It is not hard to see that there are DBWs $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$ with $O(|E|)$ states that recognize L_φ and L_ψ, respectively, and can be constructed in polynomial time. We argue that $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H} \rangle$ is realizable with privacy under knowledge of the specification iff $\mathcal{H}$ is a vertex cover of G. Thus, G has a vertex cover of size b iff $\langle L_\varphi, L_\psi, \mathsf{cost}, b \rangle$ is realizable with privacy under the knowledge of the specification, for the uniform cost function that assign 1 to all $v \in V$.

For the first direction, consider a vertex cover $\mathcal{H}$ of G. Let $\mathcal{T}$ be a transducer that generates the computation $\varnothing^\omega$ and hides the signals in $\mathcal{H}$. In every step of the computation, the observer sees that the signals in $V \setminus \mathcal{H}$ are assigned F. Assigning F to the signals in $\mathcal{H}$ results in a word w^- that is in $L^1_\varphi \setminus L^2_\varphi = L_\varphi \setminus L_\psi$. On the other hand, as $\mathcal{H}$ is a vertex cover, and so it intersects every edge $e \in E$, then assigning T to the signals in $\mathcal{H}$ results in a word w^+ that is in $L^2_\varphi = L_\varphi \cap L_\psi$. Hence, the membership of the computation in L_ψ is not known to the observer, even when she knows it is in L_φ. That is, $\mathcal{T}$ raelizes $\langle \mathcal{A}_\varphi, \mathcal{A}_\psi, \mathcal{H}\rangle$ with privacy under knowledge of the specification.

In the other direction, consider a set $\mathcal{H} \subseteq V$ that is not a vertex cover of G. Consider a transducer $\mathcal{T}$ that realizes L_φ and hides the signals in $\mathcal{H}$. Let $w = \sigma_0 \cdot \sigma_1 \cdot \sigma_2 \cdots \in (2^V)^\omega$ be the computation generated by the transducer. Since $\mathcal{H}$ is not a vertex cover of G, there is $0 \leq i \leq m$ such that $\mathcal{H} \cap \{e_i\} = \varnothing$. Let $\kappa_i = \mathsf{hide}_\mathcal{H}(\sigma_i)$. We claim that the membership of w to L_ψ can be revealed from κ_i when the observer knows that $w \in L_\varphi$. Indeed, if $\kappa_i \cap e_i \neq \varnothing$, then $\sigma_i \neq \varnothing$ and so $w \neq \varnothing^\omega$. Hence, $w \in L_\varphi \setminus L^1_\varphi = L_\psi$. Also, if $\kappa_i \cap e_i = \varnothing$, then sa $\mathcal{H} \cap e_i = \varnothing$, we have that $\sigma_i \cap e_i = \varnothing$. Thus, $w \in L_\varphi \setminus L^2_\varphi = L_\varphi \setminus L_\psi$. □

We continue to the second issue, proving that synthesis with privacy under knowledge of the specification is EXPTIME-hard even for specifications in DBWs and secrets of a fixed size. Note that synthesis with privacy (without knowledge of the specification) can be solved in PTIME in this case (see Remark 3.3). The proof is similar to that of Theorem 3.6, except that here the lower bound needs the secret to be a of a fixed size, making the specification more complex. It follows that the exponential blow-up in $\mathcal{D}_\varphi$, which exists in Lemma 6.1 cannot be avoided even when $\mathcal{D}_\varphi$ is a deterministic Büchi (rather than parity) automaton and $\mathcal{D}_\psi$ is of a fixed size.

Theorem 6.3. *Given two disjoint finite sets I and O, DBWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$ over $2^{I\cup O}$, and a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, deciding whether $\langle \mathcal{D}_\varphi, \mathcal{D}_\psi, \mathcal{H}\rangle$ is realizable with privacy under knowledge of the specification is EXPTIME-hard already when $\mathcal{D}_\psi$ is of fixed size.*

Proof. We describe a polynomial-time reduction from NBW realizability: given an NBW $\mathcal{A}$ over $2^{I\cup O}$, we construct a new set of signals $\mathcal{H}$, and DBWs $\mathcal{D}_\varphi$ and $\mathcal{D}_\psi$ over the alphabet $2^{I\cup O\cup\mathcal{H}}$, such that $\mathcal{D}_\psi$ has only two states and $L(\mathcal{A})$ is realizable iff $\langle \mathcal{D}_\varphi, \mathcal{D}_\psi, \mathcal{H}\rangle$ is realizable with privacy under knowledge of the specification.

As in the proof of Theorem 3.6, the signals in $\mathcal{H}$ encode the states of $\mathcal{A}$, and thus we refer to words in $(2^{I\cup O\cup\mathcal{H}})^\omega$ as a combination $w \oplus r \in (2^{I\cup O} \times Q)^\omega$, where $w \in (2^{I\cup O})^\omega$ and $r \in Q^\omega$. Also, we assume that $\mathcal{A}$ has a single initial state and that every word in $(2^{I\cup O})^\omega$ has at least one rejecting run in $\mathcal{A}$. Consider a word $w \oplus r \in (2^{I\cup O} \times Q)^\omega$. The key difference between the proof of Theorem 3.6 and the one here, is that now the specification requires the Q-component r to be a legal run of $\mathcal{A}$ on the $2^{I\cup O}$-component w. Then the secret only requires r to be accepting. Thus, $\mathcal{D}_\psi$ only needs two states to detect infinitely many visits in accepting states of $\mathcal{A}$.[4]

We prove that $L(\mathcal{A})$ is realizable iff $\langle \mathcal{D}_\varphi, \mathcal{D}_\psi, \mathcal{H}\rangle$ is realizable with privacy under knowledge of the specification. Indeed, as in the proof of Theorem 3.6, a transducer $\mathcal{T}$ realizes $L(\mathcal{A})$ iff $\mathcal{T}$ realizes $\langle \mathcal{D}_\varphi, \mathcal{D}_\psi, \mathcal{H}\rangle$ with privacy under knowledge of the specification. □

[4]Note that while $\mathcal{D}_\psi$ only has two states, its representation requires an encoding of Q. If one wants to get rid of this, it is possible to add a new signal p_{acc} to $\mathcal{H}$, and use the specification in order to require p_{acc} to hold exactly when the encoded state is accepting. That way, $\mathcal{D}_\psi$ has two states and its transitions depend only on p_{acc}.

6.2. **An observer that knows the transducer.** Recall that an observer that knows the specification can evaluate the secret only with respect to computations that satisfy the specification. In fact, the observer can do better, and restricts the search to computations that are generated by an (I/O)-transducer that realizes the specification. In this section we show that such a restriction is indeed stronger, and study an even stronger setting, in which the observer not only knows that computations are generated by an (I/O)-transducer that realizes the specification, but actually knows the (I/O)-transducer $\mathcal{T}$. Note that even though the observer knows all the components of $\mathcal{T}$, she still may not be able to evaluate the secret, as she sees only the signals not in $\mathcal{H}$. Note also that since $\mathcal{T}$ realizes the specification φ, then $L(\mathcal{T}) \subseteq L_\varphi$. As an observer that knows $\mathcal{T}$ also knows $L(\mathcal{T})$, the latter implies that knowing $\mathcal{T}$ is more helpful for the observer than knowing the specification.

We first demonstrate that an observer that takes into an account the fact that the computations are generated by an (I/O)-transducer that realizes φ can reveal secrets that are not revealed by an observer that only takes into account the fact that all the computations satisfy φ. In order to see the difference between the two types of observers, consider the case $I = \mathcal{H} = \{p_1, p_2\}$, $O = \{q\}$, $\varphi = (q \leftrightarrow p_1) \vee G p_2$, and $\psi = p_1$. An observer that only knows that all the computations satisfy φ does not know which of its two disjuncts is satisfied, and thus, even though she observes q, the value of p_1 stays secret. Formally, a transducer that realizes ϕ $\mathcal{H}$-hides ψ from the observer, even if the observer knows that φ is satisfied. Indeed, for every observable computation $\kappa \in 2^{\{q\}}$, there is a computation $w^+ \in \mathsf{noise}_\mathcal{H}(\kappa)$ that satisfies $p_1 \wedge G p_2$ and a computation $w^- \in \mathsf{noise}_\mathcal{H}(\kappa)$ that satisfies $(\neg p_1) \wedge G p_2$. Hence, $\langle \varphi, \psi, \mathcal{H} \rangle$ is realizable with privacy even when the observer knows the specification.

On the other hand, an observer that knows that the computations are generated by a transducer $\mathcal{T}$ that realizes φ can do more. To see this, note that the specification φ is *open equivalent* to the formula $q \leftrightarrow p_1$: for every transducer $\mathcal{T}$, we have that $\mathcal{T}$ realizes φ iff $\mathcal{T}$ realizes $q \leftrightarrow p_1$. Indeed, if $\mathcal{T}$ does not realize $q \leftrightarrow p_1$, then φ is not satisfied in computations that do not satisfy $G p_2$, which is the case for almost all the computations of $\mathcal{T}$. Accordingly, an observer that knows that the computations are generated by a transducer that realizes φ also knows that they are generated by a transducer that realizes $q \leftrightarrow p_1$, and so she can learn the value of the secret p_1 by observing the value of q.[5]

We continue to an observer that knows $\mathcal{T}$. Formally, we say that $\mathcal{T}$ $\mathcal{H}$-hides ψ from an observer that knows $\mathcal{T}$ if for all $w_I \in (2^I)^\omega$, there is $w'_I \in (2^I)^\omega$ such that $\mathsf{hide}_\mathcal{H}(\mathcal{T}(w_I)) = \mathsf{hide}_\mathcal{H}(\mathcal{T}(w'_I))$ and the two computations $\mathcal{T}(w_I)$ and $\mathcal{T}(w'_I)$ do not agree on ψ.

Note that $\mathcal{T}$ does not $\mathcal{H}$-hide ψ from an observer that knows $\mathcal{T}$ if there is $w_I \in (2^I)^\omega$ such that $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T}) \subseteq \psi$ or $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T}) \subseteq \neg\psi$. Indeed, if such a word w_I exists, then the set $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T})$, namely the set of all the computations that are possible interactions of $\mathcal{T}$ with w_I from the point of view of the observer, reveals the satisfaction value of the secret ψ. Technically, rather than considering the set $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I))$ of computations, we consider its restriction $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T})$ to computations generated by $\mathcal{T}$. As discussed above, since $L(\mathcal{T}) \subseteq L_\varphi$, such a restriction is stronger than the one induced by knowing φ.

We consider the problem of checking whether $\mathcal{T}$ $\mathcal{H}$-hides ψ from an observer that knows $\mathcal{T}$. The input to the problem is a transducer $\mathcal{T}$, an LTL formula ψ, and a set $\mathcal{H} \subseteq I \cup O$, and the mission is to decide whether $\mathcal{T}$ $\mathcal{H}$-hides ψ from an observer that knows $\mathcal{T}$.

[5]Note that the example uses the fact that the specification φ is *inherently vacuous*; namely that it is open equivalent to a simpler formula [GBJV08, FKSV08]. The question whether there are also examples in which the two types of an observer differ for a specification that is not inherently vacuous is left open.

Theorem 6.4. *Given an (I/O)-transducer $\mathcal{T}$, an LTL formula ψ, and a set $\mathcal{H} \subseteq I \cup O$ of hidden signals, checking whether $\mathcal{T}$ $\mathcal{H}$-hides ψ from an observer that knows $\mathcal{T}$ can be solved in EXPSPACE.*

Proof. Consider an LTL formula ψ, a set $\mathcal{H} \subseteq I \cup O$, and an (I/O)-transducer $\mathcal{T} = \langle I, O, S, s_0, \mu, \tau\rangle$. Let $\mathcal{U}_\psi = \langle 2^{I\cup O}, Q, q_0, \delta, \alpha\rangle$ be a UCW for ψ. There is such a UCW of size exponential in $|\psi|$ [VW94]. We construct two UCWs $\mathcal{U}^+$ and $\mathcal{U}^-$ over the alphabet 2^I, such that $L(\mathcal{U}^+) = \{w_I : \mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T}) \subseteq \psi\}$ and $L(\mathcal{U}^-) = \{w_I : \mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T}) \subseteq \neg\psi\}$. Then, $\mathcal{T}$ $\mathcal{H}$-hides ψ from an observer that knows $\mathcal{T}$ iff $L(\mathcal{U}^+) \neq \varnothing$ or $L(\mathcal{U}^-) \neq \varnothing$.

We now describe the UCW $\mathcal{U}^+$. Reading a word $w_I \in 2^I$, the UCW $\mathcal{U}^+$ traces every word w'_I with $\mathcal{T}(w'_I) \in \mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T})$. The run corresponding to w'_I is accepting iff $\mathcal{T}(w'_I) \in \psi$. Thus, all of the runs on w_I are accepting iff $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T}) \subseteq \psi$. The state space of $\mathcal{U}^+$ is $S \times S \times Q$. The first S-component of each state is used to simulate the run of $\mathcal{T}$ on w_I. The second component$S$-component is used to simulate a run of $\mathcal{T}$ on a word in $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I))$. That is, after reading a word $x \in (2^I)^*$, each run of $\mathcal{U}^+$ reaches a state of the form $\langle \mu^*(s_0, x), s_2, q\rangle$, where $s_2 = \mu^*(s_0, x')$ for some word $x' \in 2^I$ with $\mathcal{T}(x') \in \mathsf{noise}_\mathcal{H}(\mathcal{T}(x))$. The Q-component of each state of $\mathcal{U}^+$ simulates the runs of $\mathcal{U}_\psi$ on words in $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I))$. That is, $q \in \delta^*(q_0, \mathcal{T}(x'))$ for some word $x' \in 2^I$ with $\mathcal{T}(x') \in \mathsf{noise}_\mathcal{H}(\mathcal{T}(x))$.

Formally, $\mathcal{U}^+ = \langle 2^I, S \times S \times Q, \langle s_0, q_0\rangle, \delta', \alpha'\rangle$, where for all $\langle s_1, s_2, q\rangle \in S \times S \times Q$ and $i \in 2^I$, we have that $\langle s_1, s_2, q\rangle \in \alpha'$ iff $q \in \alpha$, and $\delta'(\langle s_1, s_2, q\rangle, i) = \{\langle s'_1, s'_2, q'\rangle : s'_1 = \mu(s_1, i) \wedge \exists i' \in \mathsf{noise}_\mathcal{H}(i)$ s.t. $(s'_2 = \mu(s_2, i') \wedge \tau(s'_2) \in \mathsf{noise}_\mathcal{H}(\tau(s'_1)) \wedge q' \in \delta(q, i' \cup \tau(s'_2)))\}$.

We prove that the language of $\mathcal{U}^+$ is exactly all words w_I with $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T}) \subseteq \psi$. For the first direction, consider a word $w_I = i_0 \cdot i_1 \cdot \ldots \in L(\mathcal{U}^+)$. That is, all runs of $\mathcal{U}^+$ on w_I are accepting. Let $w' = (i'_0 \cup o'_0) \cdot (i'_1 \cup o'_1) \cdot \ldots \in \mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T})$. We show that $w' \in \psi$. Let $q_0 \cdot q_1 \cdot \ldots$ be a run of $\mathcal{U}_\psi$ on w'. We show that this run is accepting. Let $s_0 \cdot s_1 \cdot \ldots$ be the run of $\mathcal{T}$ on w_I, and let $s'_0 \cdot s'_1 \cdot \ldots$ be the run of $\mathcal{T}$ on $w'|_I = i'_0 \cdot i'_1 \cdot \ldots$. Since $w' \in L(\mathcal{T})$, we have that $o'_j = \tau(s'_{j+1})$, for all $j \geq 0$. Also, since $w' \in \mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I))$, we have that $i'_j \in \mathsf{noise}_\mathcal{H}(i_j)$ and $o'_j \in \mathsf{noise}_\mathcal{H}(\tau(s_{j+1}))$, for all $j \geq 0$. Thus, by definition of $\mathcal{U}^+$, we have that $\langle s_0, s'_0, q_0\rangle \cdot \langle s_1, s'_1, q_1\rangle \cdot \ldots$ is a run of $\mathcal{U}^+$ on w_I. Since all runs of $\mathcal{U}^+$ on w_I are accepting, we get that $q_j \in \alpha$ for finitely many j-s. Thus, the run $q_0 \cdot q_1 \cdot \ldots$ is accepting as needed.

For the second direction, consider a word $w_I = i_0 \cdot i_1 \cdot \ldots \in (2^I)^\omega$ with $\mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T}) \subseteq \psi$. We show that $w_I \in L(\mathcal{U}^+)$. Let $r = \langle s_0, s'_0, q_0\rangle \cdot \langle s_1, s'_1, q_1\rangle$ be a run of $\mathcal{U}^+$ on w_I. We show that this run is accepting. First, by the definition of $\mathcal{U}^+$, we have that $s_0 \cdot s_1 \cdot \ldots$ is the run of $\mathcal{T}$ on w_I. For all $j \geq 0$, let o_j be $\tau(s_{j_1})$. By the definition of $\mathcal{U}^+$, there is a word $w' = (i'_0 \cup o'_0) \cdot (i'_1 \cup o'_1) \cdot \ldots$ with $i'_j \in \mathsf{noise}_\mathcal{H}(i_j)$, $o'_j = \tau(s'_{j+1})$ and $o'_j \in \mathsf{noise}_\mathcal{H}(o_j)$ for all $j \geq 0$, for which $q_{j+1} \in \delta(q_j, i'_j \cup o'_j)$. Since $o'_j = \tau(s'_{j+1})$ for all $j \geq 0$, we get that $w' \in L(\mathcal{T})$. Since $i'_j \in \mathsf{noise}_\mathcal{H}(i_j)$ and $o'_j \in \mathsf{noise}_\mathcal{H}(o_j)$ for all $j \geq 0$, we get that $w' \in \mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I))$. Overall, $w' \in \mathsf{noise}_\mathcal{H}(\mathcal{T}(w_I)) \cap L(\mathcal{T})$, and by that $w' \in \psi$. Moreover, since $q_{j+1} \in \delta(q_j, i'_j \cup o'_j)$ for all $j \geq 0$, we get that $q_o \cdot q_1 \cdot \ldots$ is a run of $\mathcal{U}_\psi$ on w', and thus accepting. That is, $q_j \in \alpha$ for only finitely many j-s, and so r is accepting as needed.

We construct the UCW $\mathcal{U}^-$ in a similar way, using $\mathcal{U}_{\neg\psi}$ instead of $\mathcal{U}_\psi$. Recall that the observer wins iff $L(\mathcal{U}^+) \neq \varnothing$ or $L(\mathcal{U}^-) \neq \varnothing$, which can be decided by checking nonemptiness twice, once for $\mathcal{U}^+$ and once for $\mathcal{U}^-$. Since universality of UCWs can be solved in PSPACE [SVW87], and the UCWs $\mathcal{U}^+$ and $\mathcal{U}^-$ are of size exponential in the size of ψ, the induced algorithm is in EXPSPACE. □

Readers familiar with HyperLTL [CFK+14, FHL+20b], might notice that $\mathcal{H}$-hiding from an observer that knows $\mathcal{T}$ has the flavor of a hyper property. Indeed, a different approach to decide whether $\mathcal{T}$ $\mathcal{H}$-hides a secret from an observer that knows $\mathcal{T}$, is by model checking of hyper properties: Given an LTL formula ψ, consider the hyper LTL formula $f = \forall\pi\exists\pi'(\bigwedge_{v\in V} v_\pi \leftrightarrow v_{\pi'}) \wedge (\psi(\pi) \leftrightarrow \neg\psi(\pi'))$, where $\psi(\pi)$ and $\psi(\pi')$ are the LTL formulas obtained from ψ by replacing each signal p by the path variables p_π and $p_{\pi'}$, respectively. It is easy to see that $\mathcal{T}$ $\mathcal{H}$-hides ψ from an observer that knows $\mathcal{T}$, if and only if $\mathcal{T} \models f$. Indeed, $\mathcal{T} \models f$ if and only if for every computation $\pi \in L(\mathcal{T})$, there is a computation $\pi' \in L(\mathcal{T})$ that agrees with π on the visible signals, but does not agree with π on the secret ψ. That is, our problem can be reduced to model checking of HyperLTL formula, which is decidable [FHL20a].

The formulation of $\mathcal{H}$-hiding from an observer that knows $\mathcal{T}$ as a HyperLTL formula of the form $\forall\pi\exists\pi' f(\pi, \pi')$ suggests that synthesis with privacy against an observer that knowns the transducer is undecidable, as is synthesis of HyperLTL formulas of this form [CFK+14].

7. Directions for Future Research

We suggested a framework for the synthesis of systems that satisfy their specifications while keeping some behaviors secret. Behaviors are kept secret from an observer by hiding the truth value of some input and output signals, subject to budget restrictions: each signal has a hiding cost, and there is a bound on the total hiding cost. Our framework captures settings in which the choice and cost of hiding are fixed throughout the computation. For example, settings with signals that cannot be hidden (e.g., alarm sound, or the temperature outside), signals that can be hidden throughout the computation with some effort (e.g., hand movement of a robot), or signals that are anyway hidden (e.g., values of internal control variables). Our main technical contribution are lower bounds for the complexity of the different aspects of privacy: the need to choose the hidden signals, and the need to hide the secret behaviors. We show that both aspects involve an exponential blow up in the complexity of synthesis without privacy.

The exponential lower bounds apply already in the relatively simple cost mechanism we study. Below we discuss possible extensions of this mechanism. In settings with a *dynamic hiding of signals*, we do not fix a set $\mathcal{H} \subseteq I \cup O$ of hidden signals. Instead, the output of the synthesis algorithm contains a transducer that describes not only the assignments to the output signals but also the choice of input and output signals that are hidden in the next cycle of the interaction. Thus, signals may be hidden only in segments of the interaction – segments that depend on the history of the interaction so far. For example, we may hide information about a string that is being typed only after a request for a password. In addition, the cost function need not be fixed and may depend on the history of the interaction too. For example, hiding the location of a robot may be cheap in certain sections of the warehouse and expensive in others. Solving synthesis with privacy in a setting with such dynamic hiding and pricing of signals involves automata over the alphabet $3^{I\cup O}$, reflecting the ability of signals to get an "unknown" truth value in parts of the computation. Moreover, as the cost is not known in advance (even when the cost of hiding signals is fixed), several mechanisms for bounding the budget are possible (energy, mean-payoff, etc. [BBFR13, CD10]).

References

[BBFR13] A. Bohy, V. Bruyère, E. Filiot, and J-F. Raskin. Synthesis from LTL specifications with mean-payoff objectives. In *Proc. 19th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 7795 of *Lecture Notes in Computer Science*, pages 169–184. Springer, 2013.

[BCJ18] R. Bloem, K. Chatterjee, and B. Jobstmann. Graph games and reactive synthesis. In *Handbook of Model Checking.*, pages 921–962. Springer, 2018.

[BGI+12] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S.P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, 2012.

[Bok18] U. Boker. Why these automata types? In *Proc. 22nd Int. Conf. on Logic for Programming Artificial Intelligence and Reasoning*, volume 57 of *EPiC Series in Computing*, pages 143–163, 2018.

[CD10] K. Chatterjee and L. Doyen. Energy parity games. In *Proc. 37th Int. Colloq. on Automata, Languages, and Programming*, pages 599–610, 2010.

[CDHR06] K. Chatterjee, L. Doyen, T. A. Henzinger, and J-F. Raskin. Algorithms for ω-regular games with imperfect information. In *Proc. 15th Annual Conf. of the European Association for Computer Science Logic*, volume 4207 of *Lecture Notes in Computer Science*, pages 287–302, 2006.

[CFK+14] M.R. Clarkson, B. Finkbeiner, M. Koleini, K.K. Micinski, M.N. Rabe, and C. Sánchez. Temporal logics for hyperproperties. In *3rd International Conference on Principles of Security and Trust*, volume 8414 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2014.

[CJK+17] C.S. Calude, S. Jain, B. Khoussainov, W. Li, and F. Stephan. Deciding parity games in quasipolynomial time. In *Proc. 49th ACM Symp. on Theory of Computing*, pages 252–263, 2017.

[DDM10] J. Dubreil, Ph. Darondeau, and H. Marchand. Supervisory control for opacity. *IEEE Transactions on Automatic Control*, 55(5):1089–1100, 2010.

[DMNS16] C. Dwork, F. McSherry, K. Nissim, and A.D. Smith. Calibrating noise to sensitivity in private data analysis. *J. Priv. Confidentiality*, 7(3):17–51, 2016.

[DN03] I. Dinur and K. Nissim. Revealing information while preserving privacy. In *Proceedings of the 22nd ACM Symposium on Principles of Database Systems*, pages 202–210. ACM, 2003.

[Ehl10] R. Ehlers. Symbolic bounded synthesis. In *Proc. 22nd Int. Conf. on Computer Aided Verification*, volume 6174 of *Lecture Notes in Computer Science*, pages 365–379. Springer, 2010.

[FHL20a] B. Finkbeiner, C. Hahn, and P. Lukert. Synthesis from hyperproperties. *Acta Informatica*, 57:137–163, 2020.

[FHL+20b] B. Finkbeiner, C. Hahn, P. Lukert, M. Stenger, and L. Tentrup. Synthesis from hyperproperties. *Acta Informatica*, 57(1-2):137–163, 2020.

[FKSV08] D. Fisman, O. Kupferman, S. Sheinvald, and M.Y. Vardi. A framework for inherent vacuity. In *4th International Haifa Verification Conference*, volume 5394 of *Lecture Notes in Computer Science*, pages 7–22. Springer, 2008.

[GBJV08] K. Greimel, R. Bloem, B. Jobstmann, and M. Vardi. Open implication. In *Proc. 35th Int. Colloq. on Automata, Languages, and Programming*, volume 5126 of *Lecture Notes in Computer Science*, pages 361–372. Springer, 2008.

[GGH+16] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM J. Comput.*, 45(3):882–929, 2016.

[HKKM02] T.A. Henzinger, S.C. Krishnan, O. Kupferman, and F.Y.C. Mang. Synthesis of uninitialized systems. In *Proc. 29th Int. Colloq. on Automata, Languages, and Programming*, volume 2380 of *Lecture Notes in Computer Science*, pages 644–656. Springer, 2002.

[KL22] O. Kupferman and O. Leshkowitz. Synthesis of privacy-preserving systems. In *Proc. 42nd Conf. on Foundations of Software Technology and Theoretical Computer Science*, volume 250 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 42:1–42:23, 2022.

[KLVY11] O. Kupferman, Y. Lustig, M.Y. Vardi, and M. Yannakakis. Temporal synthesis for bounded systems and environments. In *Proc. 28th Symp. on Theoretical Aspects of Computer Science*, pages 615–626, 2011.

[KR10] O. Kupferman and A. Rosenberg. The blow-up in translating LTL to deterministic automata. In *Proc. 6th Workshop on Model Checking and Artificial Intelligence*, volume 6572 of *Lecture Notes in Artificial Intelligence*, pages 85–94. Springer, 2010.

[KSS11] O. Kupferman, D. Sadigh, and S.A. Seshia. Synthesis with clairvoyance. In *7th International Haifa Verification Conference*, volume 7261 of *Lecture Notes in Computer Science*, pages 5–19. Springer, 2011.

[KV00] O. Kupferman and M.Y. Vardi. Synthesis with incomplete informatio. In *Advances in Temporal Logic*, pages 109–127. Kluwer Academic Publishers, 2000.

[KV05a] O. Kupferman and M.Y. Vardi. From linear time to branching time. *ACM Transactions on Computational Logic*, 6(2):273–294, 2005.

[KV05b] O. Kupferman and M.Y. Vardi. Safraless decision procedures. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 531–540, 2005.

[Pit06] N. Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. In *Proc. 21st ACM/IEEE Symp. on Logic in Computer Science*, pages 255–264. IEEE press, 2006.

[Pnu81] A. Pnueli. The temporal semantics of concurrent programs. *Theoretical Computer Science*, 13:45–60, 1981.

[PR89] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th ACM Symp. on Principles of Programming Languages*, pages 179–190, 1989.

[Rab72] M.O. Rabin. Automata on infinite objects and Church's problem. *Amer. Mathematical Society*, 1972.

[Rei84] J.H. Reif. The complexity of two-player games of incomplete information. *Journal of Computer and Systems Science*, 29:274–301, 1984.

[Ros92] R. Rosner. *Modular Synthesis of Reactive Systems*. PhD thesis, Weizmann Institute of Science, 1992.

[Saf88] S. Safra. On the complexity of ω-automata. In *Proc. 29th IEEE Symp. on Foundations of Computer Science*, pages 319–327, 1988.

[SF07] S. Schewe and B. Finkbeiner. Bounded synthesis. In *5th Int. Symp. on Automated Technology for Verification and Analysis*, volume 4762 of *Lecture Notes in Computer Science*, pages 474–488. Springer, 2007.

[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.

[VW94] M.Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994.

[WRR+18] Y. Wu, V. Raman, B.C. Rawlings, S. Lafortune, and S.A. Seshia. Synthesis of obfuscation policies to ensure privacy and utility. *Journal of Automated Reasoning*, 60(1):107–131, 2018.