



HAL
open science

The Sensing Cost of Monitoring and Synthesis

Shaull Almagor, Denis Kuperberg, Orna Kupferman

► **To cite this version:**

Shaull Almagor, Denis Kuperberg, Orna Kupferman. The Sensing Cost of Monitoring and Synthesis. FSTTCS, 2015, Bangalore, India. hal-02070826

HAL Id: hal-02070826

<https://hal.science/hal-02070826>

Submitted on 18 Mar 2019

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

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

The Sensing Cost of Monitoring and Synthesis

Shaull Almagor¹, Denis Kuperberg², and Orna Kupferman¹

1 The Hebrew University, Israel.

2 IRIT/Onera, Toulouse.

Abstract

In [2], we introduced *sensing* as a new complexity measure for the complexity of regular languages. Intuitively, the sensing cost quantifies the detail in which a random input word has to be read by a deterministic automaton in order to decide its membership in the language. In this paper, we consider sensing in two principal applications of deterministic automata. The first is *monitoring*: we are given a computation in an on-line manner, and we have to decide whether it satisfies the specification. The second is *synthesis*: we are given a sequence of inputs in an on-line manner and we have to generate a sequence of outputs so that the resulting computation satisfies the specification. In the first, our goal is to design a monitor that handles all computations and minimizes the expected average number of sensors used in the monitoring process. In the second, our goal is to design a transducer that realizes the specification for all input sequences and minimizes the expected average number of sensors used for reading the inputs.

We argue that the two applications require new and different frameworks for reasoning about sensing, and develop such frameworks. We focus on safety languages. We show that for monitoring, minimal sensing is attained by a monitor based on the minimal deterministic automaton for the language. For synthesis, however, the setting is more challenging: minimizing the sensing may require exponentially bigger transducers, and the problem of synthesizing a minimally-sensing transducer is EXPTIME-complete even for safety specifications given by deterministic automata.

1998 ACM Subject Classification F.4.3 : Formal Languages, B.8.2 Performance Analysis and Design Aids, F.1.1 Models of Computation

Keywords and phrases Automata, regular languages, ω -regular languages, complexity, sensing, minimization.

1 Introduction

Studying the complexity of a formal language, there are several complexity measures to consider. When the language is given by means of a Turing Machine, the traditional measures are time and space requirements. Theoretical interest as well as practical considerations have motivated additional measures, such as randomness (the number of random bits required for the execution) [11] or communication complexity (number and length of messages required) [10]. For ω -regular languages, given by means of finite-state automata, the classical complexity measure is the size of a minimal deterministic automaton that recognizes the language.

In [2], we introduced and studied a new complexity measure, namely the *sensing cost* of the language. Intuitively, the sensing cost of a language measures the detail with which a random input word needs to be read in order to decide membership in the language. Sensing has been studied in several other CS contexts. In theoretical CS, in methodologies such as PCP and property testing, we are allowed to sample or query only part of the input [8]. In more practical applications, mathematical tools in signal processing are used to reconstruct information based on compressed sensing [3], and in the context of data streaming, one cannot store in memory the entire input, and therefore has to approximate its properties according to partial “sketches” [12].

Our study in [2] considered regular and ω -regular languages, where sensing is defined as follows. Consider a deterministic automaton \mathcal{A} over an alphabet 2^P , for a finite set P of *signals*. For a state



© Shaull Almagor, Denis Kuperberg, and Orna Kupferman;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

q of \mathcal{A} , we say that a signal $p \in P$ is *sensed* in q if at least one transition taken from q depends on the truth value of p . The *sensing cost* of q is the number of signals it senses, and the sensing cost of a run is the average sensing cost of states visited along the run. We extend the definition to automata by assuming a given distribution of the inputs. The sensing cost of a language with respect to this distribution is then the infimum sensing cost of an automaton for the language. For simplicity, we focus on the uniform distribution, and we refer to the sensing cost of an automaton without parameterizing it by a distribution. As detailed in Remark 1, all our results can be extended to a setting with a parameterized distribution.

In [2], we showed that computing the sensing cost of a language can be done in polynomial time. We further showed that while in finite words the minimal sensing cost is always attained, this is not the case for infinite words. For example, recognizing the language L over $2^{\{p\}}$ of all words with infinitely many p 's, one can give up sensing of p for unboundedly-long intervals, thus the sensing cost of L is 0, yet every deterministic automaton \mathcal{A} that recognizes L must sense p infinitely often, causing the sensing cost of \mathcal{A} to be strictly greater than 0.

In the context of formal methods, sensing has two appealing applications. The first is *monitoring*: we are given a computation and we have to decide whether it satisfies a specification. When the computations are over 2^P , we want to design a monitor that minimizes the expected average number of sensors used in the monitoring process. Monitoring is especially useful when reasoning about *safety* specifications [7]. There, every computation that violates the specification has a bad prefix – one all whose extensions are not in L . Hence, as long as the computation is a prefix of some word in L , the monitor continues to sense and examine the computation. Once a bad prefix is detected, the monitor declares an error and no further sensing is required. The second application is *synthesis*. Here, the set P of signals is partitioned into sets I and O of input and output signals, respectively. We are given a specification L over the alphabet $2^{I \cup O}$, and our goal is to construct an *I/O transducer* that realizes L . That is, for every sequence of assignments to the input signals, the transducer generates a sequence of assignments to the output signals so that the obtained computation is in L [13]. Our goal is to construct a transducer that minimizes the expected average number of sensors (of input signals) that are used along the interaction.

The definition of sensing cost in [2] falls short in the above two applications. For the first, the definition in [2] does not distinguish between words in the language and words not in the language, whereas in monitoring we care only for words in the language. In particular, according to the definition in [2], the sensing cost of a safety language is always 0. For the second, the definition in [2] considers automata and does not partition P into I and O , whereas synthesis refers to *I/O-transducers*. Moreover, unlike automata, correct transducers generate only computations in the language, and they need not generate all words in the language – only these that ensure receptiveness with respect to all sequences of inputs.

In this work we study sensing in the context of monitoring and synthesis. We suggest definitions that capture the intuition of “required number of sensors” in these settings and solve the problems of generating monitors and transducers that minimize sensing. For both settings, we focus on safety languages.

Consider, for example, a traffic monitor that has access to various sensors on roads and whose goal is to detect accidents. Once a road accident is detected, an alarm is raised to the proper authorities and the monitoring is stopped until the accident has been taken care of. The monitor can read the speed of cars along the roads, as well as the state of traffic lights. An accident is detected when some cars do not move even-though no traffic light is stopping them. Sensing the speed of every car and checking every traffic light requires huge sensing. Our goal is to find a monitor that minimizes the required sensing and still detects all accidents. In the synthesis setting, our goal is extended to designing a transducer that controls the traffic lights according to the speed of the traffic in each

direction, and satisfies some specification (say, give priority to slow traffic), while minimizing the sensing of cars.

We can now describe our model and results. Let us start with monitoring. Recall that the definition of sensing in [2] assumes a uniform probability on the assignments to the signals, whereas in monitoring we want to consider instead more intricate probability spaces – ones that restrict attention to words in the language. As we show, there is more than one way to define such probability spaces, each leading to a different measure. We study two such measures. In the first, we sample a word randomly, letter by letter, according to a given distribution, allowing only letters that do not generate bad prefixes. In the second, we construct a sample space directly on the words in the language. We show that in both definitions, we can compute the sensing cost of the language in polynomial time, and that the minimal sensing cost is attained by a minimal-size automaton. Thus, luckily enough, even though different ways in which a computation may be given in an online manner calls for two definitions of sensing cost, the design of a minimally-sensing monitor is the same in the two definitions.

Next, we proceed to study sensing for synthesis. The main challenge there is that we no longer need to consider all words in the language. Also, giving up sensing has a flavor of synthesis with incomplete information [9]: the transducer has to realize the specification no matter what the incomplete information is. This introduces a new degree of freedom, which requires different techniques than those used in [2]. In particular, while a minimal-size transducer for a safety language can be defined on top of the state space of a minimal-size deterministic automaton for the language, this is not the case when we seek minimally-sensing transducers. This is different also from the results in [2] and even these in the monitoring setting, where a minimally-sensing automaton or monitor for a safety language coincides with the minimal-size automaton for it. In fact, we show that a minimally-sensing transducer for a safety language might be exponentially bigger than a minimal-size automaton for the language. Consequently, the problems of computing the minimal sensing cost and finding a minimally-sensing transducer are EXPTIME-complete even for specifications given by means of deterministic safety automata. On the positive side, a transducer that attains the minimal sensing cost always exists for safety specifications. For general ω -regular specifications, even decidability of computing the optimal sensing cost remains open.

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

2 Preliminaries

Automata and Transducers A *deterministic automaton on infinite words* is $\mathcal{A} = \langle \Sigma, Q, q_0, \delta, \alpha \rangle$, where Q is a finite set of states, $q_0 \in Q$ is an initial state, $\delta : Q \times \Sigma \rightarrow Q$ is a partial transition function, and α is an acceptance condition. We sometimes refer to δ as a relation $\Delta \subseteq Q \times \Sigma \times Q$, with $\langle q, \sigma, q' \rangle \in \Delta$ iff $\delta(q, \sigma) = q'$. A run of \mathcal{A} on a word $w = \sigma_1 \cdot \sigma_2 \cdots \in \Sigma^\omega$ is a sequence of states q_0, q_1, \dots such that $q_{i+1} = \delta(q_i, \sigma_{i+1})$ for all $i \geq 0$. Note that since δ is deterministic and partial, \mathcal{A} has at most one run on a word. A run is accepting if it satisfies the acceptance condition. A word $w \in \Sigma^\omega$ is accepted by \mathcal{A} if \mathcal{A} has an accepting run on w . The language of \mathcal{A} , denoted $L(\mathcal{A})$, is the set of words that \mathcal{A} accepts. We denote by \mathcal{A}^q the automaton \mathcal{A} with the initial state set to q .

In a deterministic *looping* automaton (DLW), every run is accepting. Thus, a word is accepted if there is a run of the automaton on it.¹ Since every run is accepting, we omit the acceptance condition

¹ For readers familiar with the Büchi acceptance condition, a looping automaton is a special case of Büchi with $\alpha = Q$.

and write $\mathcal{A} = \langle \Sigma, Q, q_0, \delta \rangle$.

For finite sets I and O of input and output signals, respectively, an I/O transducer is $\mathcal{T} = \langle I, O, Q, q_0, \delta, \rho \rangle$, where Q is a finite set of states, $q_0 \in Q$ is an initial state, $\delta : Q \times 2^I \rightarrow Q$ is a total transition function, and $\rho : Q \rightarrow 2^O$ is a labeling function on the states. The run of \mathcal{T} on a word $w = i_0 \cdot i_1 \cdots \in (2^I)^\omega$ is the sequence of states q_0, q_1, \dots such that $q_{k+1} = \delta(q_k, i_k)$ for all $k \geq 0$. The *output* of \mathcal{T} on w is then $o_1, o_2, \dots \in (2^O)^\omega$ where $o_k = \rho(q_k)$ for all $k \geq 1$. Note that the first output assignment is that of q_1 , and we do not consider $\rho(q_0)$. This reflects the fact that the environment initiates the interaction. The *computation* of \mathcal{T} on w is then $\mathcal{T}(w) = i_0 \cup o_1, i_1 \cup o_2, \dots \in (2^{I \cup O})^\omega$.

Note that the structure of each I/O -transducer \mathcal{T} induces a DLW $\mathcal{A}_{\mathcal{T}}$ over the alphabet 2^I with a total transition relation. Thus, the language of the DLW is $(2^I)^\omega$, reflecting the receptiveness of \mathcal{T} .

Safety Languages Consider a language $L \subseteq \Sigma^\omega$. A finite word $x \in \Sigma^*$ is a *bad prefix* for L if for every $y \in \Sigma^\omega$, we have that $x \cdot y \notin L$. That is, x is a bad prefix if all its extensions are words not in L . The language L is then a *safety language* if every word not in L has a bad prefix. For a language L , let $\text{pref}(L) = \{x \in \Sigma^* : \text{there exists } y \in \Sigma^\omega \text{ such that } x \cdot y \in L\}$ be the set of prefixes of words in L . Note that each word in Σ^* is either in $\text{pref}(L)$ or is a bad prefix for L . Since the set $\text{pref}(L)$ for a safety language L is *fusion closed* (that is, a word is in L iff all its prefixes are in $\text{pref}(L)$), an ω -regular language is safety iff it can be recognized by a DLW [15].

Consider a safety language L over sets I and O of input and output signals. We say that L is *I/O -realizable* if there exists an I/O transducer \mathcal{T} all whose computations are in L . Thus, for every $w \in (2^I)^\omega$, we have that $\mathcal{T}(w) \in L$. We then say that \mathcal{T} *I/O -realizes* L . When I and O are clear from the context, we omit them. The *synthesis* problem gets as input a safety language L over $I \cup O$, say by means of a DLW, and returns an I/O -transducer that realizes L or declares that L is not I/O -realizable.

Sensing In [2], we defined regular sensing as a measure for the number of sensors that need to be operated in order to recognize a regular language. We study languages over an alphabet $\Sigma = 2^P$, for a finite set P of signals. A letter $\sigma \in \Sigma$ corresponds to a truth assignment to the signals, and sensing a signal amounts to knowing its assignment. Describing sets of letters in Σ , it is convenient to use Boolean assertions over P . For example, when $P = \{a, b\}$, the assertion $\neg b$ stands for the set $\{\emptyset, \{a\}\}$ of two letters.

For completeness, we bring here the definitions from [2]. Consider a language L and a deterministic automaton $\mathcal{A} = \langle 2^P, Q, q_0, \delta, \alpha \rangle$ such that $L(\mathcal{A}) = L$. We assume that δ is total. For a state $q \in Q$ and a signal $p \in P$, we say that p is *sensed in* q if there exists a set $S \subseteq P$ such that $\delta(q, S \setminus \{p\}) \neq \delta(q, S \cup \{p\})$. Intuitively, a signal is sensed in q if knowing its value may affect the destination of at least one transition from q . We use $\text{sensed}(q)$ to denote the set of signals sensed in q . The *sensing cost* of a state $q \in Q$ is $\text{scost}(q) = |\text{sensed}(q)|$.²

For a finite run $r = q_1, \dots, q_m$ of \mathcal{A} , we define the sensing cost of r , denoted $\text{scost}(r)$, as $\frac{1}{m} \sum_{i=0}^{m-1} \text{scost}(q_i)$. That is, $\text{scost}(r)$ is the average number of sensors that \mathcal{A} uses during r . Now, for a finite word w , we define the sensing cost of w in \mathcal{A} , denoted $\text{scost}_{\mathcal{A}}(w)$, as the sensing cost of the run of \mathcal{A} on w . Finally, the sensing cost of \mathcal{A} is the expected sensing cost of words of length that tends to infinity, where we assume that the letters in Σ are uniformly distributed (see Remark 1 below). Thus, $\text{scost}(\mathcal{A}) = \lim_{m \rightarrow \infty} |\Sigma|^{-m} \sum_{w \in \Sigma^m} \text{scost}_{\mathcal{A}}(w)$.

² We note that, alternatively, one could define the *sensing level* of states, with $\text{slevel}(q) = \frac{\text{scost}(q)}{|P|}$. Then, for all states q , we have that $\text{slevel}(q) \in [0, 1]$. All our results hold also for this definition, simply by dividing the sensing cost by $|P|$.

Note that the definition applies to automata on both finite and infinite words, and it corresponds to the closed setting: the automaton gets as input words over 2^P and uses sensors in order to monitor the input words and decide their membership in L . We define the *sensing cost* of a language L to be the minimal cost of an automaton for L . A-priori, the minimal cost might not be attained by a single automaton, thus we define $scost(L) = \inf \{scost(\mathcal{A}) : \mathcal{A} \text{ is an automaton for } L\}$.

► **Remark 1** (On the choice of uniform distribution). The choice of a uniform distribution on the letters in Σ may be unrealistic in practice. Indeed, in real scenarios, the distribution on the truth assignments to the underlying signals may be complicated. Generally, such a distribution can be given by a Markov chain (in monitoring) or by an MDP (in synthesis). As it turns out, adjusting our setting and algorithms to handle such distributions involves only a small technical elaboration, orthogonal to the technical challenges that exists already in a uniform distribution.

Accordingly, throughout the paper we assume a uniform distribution on the truth assignments to the signals. In Appendix B we describe how our setting and algorithms are extended to the general case. ◀

The definition of sensing in [2] essentially considers the sensing required in the Ergodic SCC of a deterministic automaton for the language. Since in safety languages, the Ergodic SCCs are accepting or rejecting sinks, which require no sensing, we have the following, which implies that the definition in [2] is not too informative for safety languages.

► **Lemma 2.** *For every safety language $L \subseteq \Sigma^\omega$, we have $scost(L) = 0$.*

Markov Chains and Decision Processes A Markov chain $\mathcal{M} = \langle S, P \rangle$ consists of a finite state space S and a stochastic transition matrix $P : S \times S \rightarrow [0, 1]$. That is, for all $s \in S$, we have $\sum_{s' \in S} P(s, s') = 1$. Given an initial state s_0 , consider the vector v^0 in which $v^0(s_0) = 1$ and $v^0(s) = 0$ for every $s \neq s_0$. The *limiting distribution* of \mathcal{M} is $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{m=0}^n v^0 P^m$. The limiting distribution satisfies $\pi P = \pi$, and can be computed in polynomial time [6].

A Markov decision process (MDP) is $\mathcal{M} = \langle S, s_0, (A_s)_{s \in S}, P, cost \rangle$ where S is a finite set of states, $s_0 \in S$ is an initial state, A_s is a finite set of actions that are available in state $s \in S$. Let $A = \bigcup_{s \in S} A_s$. Then, $P : S \times A \times S \rightarrow [0, 1]$ is a partial transition probability function, defining for every two states $s, s' \in S$ and action $a \in A_s$, the probability of moving from s to s' when action a is taken. Accordingly, $\sum_{s' \in S} P(s, a, s') = 1$. Finally, $cost : S \times A \rightarrow \mathbb{N}$ is a partial cost function, assigning each state s and action $a \in A_s$, the cost of taking action a in state s .

An MDP can be thought of as a game between a player who chooses the actions and nature, which acts stochastically according to the transition probabilities.

A *policy* for an MDP \mathcal{M} is a function $f : S^* \times S \rightarrow A$ that outputs an action given the history of the states, such that for s_0, \dots, s_n we have $f(s_0, \dots, s_n) \in A_{s_n}$. Policies correspond to the strategies of the player. The *cost* of a policy f is the expected average cost of a random walk in \mathcal{M} in which the player proceeds according to f . Formally, for $m \in \mathbb{N}$ and for a sequence of states $\tau = s_0, \dots, s_{m-1}$, we define $P_f(\tau) = \prod_{i=1}^{m-1} P(s_{i-1}, f(s_0 \dots s_{i-1}), s_i)$. Then, $cost_m(f, \tau) = \frac{1}{m} \sum_{i=1}^m cost(s_i, f(s_1 \dots s_i))$ and we define the cost of f as $cost(f) = \liminf_{m \rightarrow \infty} \frac{1}{m} \sum_{\tau: |\tau|=m} cost_m(f, \tau) \cdot P_f(\tau)$.

A policy is *memoryless* if it depends only on the current state. We can describe a memoryless policy by $f : S \rightarrow A$. A memoryless policy f induces a Markov chain $\mathcal{M}^f = \langle S, P_f \rangle$ with $P_f(s, s') = P(s, f(s), s')$. Let π be the limiting distribution of \mathcal{M}^f . It is not hard to prove that $cost(f) = \sum_{s \in S} \pi_s cost(s, f(s))$. Let $cost(\mathcal{M}) = \inf \{cost(f) : f \text{ is a policy for } \mathcal{M}\}$. That is, $cost(\mathcal{M})$ is the expected cost of a game played on \mathcal{M} in which the player uses an optimal policy.

► **Theorem 3.** *Consider an MDP \mathcal{M} . Then, $cost(\mathcal{M})$ can be attained by a memoryless policy, which can be computed in polynomial time.*

3 Monitoring


As described in Section 2, the definition of sensing in [2] takes into account all words in $(2^P)^\omega$, regardless their membership in the language. In monitoring, we restrict attention to words in the language, as once a violation is detected, no further sensing is required. In particular, in safety languages, violation amounts to a detection of a bad prefix, and indeed safety languages are the prominent class of languages for which monitoring is used [7].

As it turns out, however, there are many approaches to define the corresponding probability space. We suggest here two. Let \mathcal{A} be a DLW and let $L = L(\mathcal{A})$.

1. **[Letter-based]** At each step, we uniformly draw a “safe” letter – one with which we are still generating a word in $\text{pref}(L)$, thereby iteratively generating a random word in L .
2. **[Word-based]** At the beginning, we uniformly draw a word in L .

We denote the sensing cost of \mathcal{A} in the letter- and word-based approaches $\text{lcost}(\mathcal{A})$ and $\text{wcost}(\mathcal{A})$, respectively. The two definitions yield two different probability measures on L , as demonstrated in Example 4 below.

► **Example 4.** Let $P = \{a\}$ and consider the safety language $L = a^\omega + (\neg a) \cdot (\text{True})^\omega$. That is, if the first letter is $\{a\}$, then the suffix should be $\{a\}^\omega$, and if the first letter is \emptyset , then all suffixes result in a word in L . Consider the DLW \mathcal{A} for L in Figure 1.



■ **Figure 1** A DLW for $a^\omega + (\neg a) \cdot (\text{True})^\omega$.

In the letter-based definition, we initially draw a letter from $2^{\{a\}}$ uniformly, i.e., either a or $\neg a$ w.p. $\frac{1}{2}$. If we draw $\neg a$, then we move to q_1 and stay there forever. If we draw a , then we move to q_2 and stay there forever. Since $\text{scost}(q_1) = 0$ and $\text{scost}(q_2) = 1$, and we reach q_1 and q_2 w.p. $\frac{1}{2}$, we get $\text{lcost}(\mathcal{A}) = \frac{1}{2}$.

In the word-based definition, we assign a uniform probability to the words in L . In this case, almost all words are not a^ω , and thus the probability of a^ω is 0. This means that we will get to q_1 w.p. 1, and thus $\text{wcost}(\mathcal{A}) = 0$. ◀

As a more realistic example, recall our traffic monitor in Section 1. There, the behavior of the cars is the random input, and the two approaches can be understood as follows. In the letter-based approach, we assume that the drivers do their best to avoid accidents regardless of the history of the traffic and the traffic lights so far. Thus, after every safe prefix, we assume that the next input is also safe. In the word-based approach, we assume that the city is planned well enough to avoid accidents. Thus, we a-priori set the distribution to safe traffic behaviors according to their likelihood.

We now define the two approaches formally.

The Letter-Based Approach Consider a DLW $\mathcal{A} = \langle \Sigma, Q, \delta, q_0 \rangle$. For a state $q \in Q$, let $\text{avail}(q)$ be the set of letters available in q , namely letters that do not cause \mathcal{A} to get stuck. Formally, $\text{avail}(q) = \{\sigma \in \Sigma : \delta(q, \sigma) \text{ is defined}\}$. We model the drawing of available letters by the Markov chain $\mathcal{M}_{\mathcal{A}} = \langle Q, P \rangle$, where the probability of a transition from state q to state q' in $\mathcal{M}_{\mathcal{A}}$ is $P(q, q') = \frac{|\{\sigma \in \Sigma : \delta(q, \sigma) = q'\}|}{|\text{avail}(q)|}$. Let π be the limiting distribution of $\mathcal{M}_{\mathcal{A}}$. We define $\text{lcost}(\mathcal{A}) = \sum_{q \in Q} \pi(q) \cdot \text{scost}(q)$.

Since computing the limiting distribution can be done in polynomial time, we have the following.

► **Theorem 5.** *Given a DLW \mathcal{A} , the sensing cost $lcost(\mathcal{A})$ can be calculated in polynomial time.*

The Word-Based Approach Consider a DLW $\mathcal{A} = \langle 2^P, Q, q_0, \delta \rangle$ recognizing a non-empty safety language L . From [2], we have $scost(\mathcal{A}) = \lim_{n \rightarrow \infty} \frac{1}{|\Sigma|^n} \sum_{u \in \Sigma^n} scost_{\mathcal{A}}(u)$, which is proven to coincide with $\mathbb{E}[scost_{\mathcal{A}}(u)]$ where \mathbb{E} is the expectation with respect to the standard measure on Σ^ω . Our goal here is to replace this standard measure with one that restricts attention to words in L . Thus, we define $wcost(\mathcal{A}) = \mathbb{E}[scost(u) \mid u \in L]$. For $n \geq 0$, let $pref(L, n)$ be the set of prefixes of L of length n . Formally, $pref(L, n) = pref(L) \cap \Sigma^n$. As in the case of the standard measure, the expectation-based definition coincides with one that is based on a limiting process: $wcost(\mathcal{A}) = \lim_{n \rightarrow \infty} \frac{1}{|pref(L, n)|} \sum_{u \in pref(L, n)} scost_{\mathcal{A}}(u)$. Thus, the expressions for $scost$ and $wcost$ are similar, except that in the expectation-based definition we add conditional probability, restricting attention to words in L , and in the limiting process we replace Σ^n by $pref(L, n)$.

Note that the term $\frac{1}{|pref(L, n)|}$ is always defined, as L is a non-empty safety language. In particular, the expectation is well defined even if L has measure 0 in Σ^ω .

► **Theorem 6.** *Given a DLW \mathcal{A} , we can compute $wcost(\mathcal{A})$ in polynomial time.*

Proof. We will use here formal power series on one variable z , a classical tool for graph and automata combinatorics. They can be thought of as polynomials of infinite degree.

For states $p, q \in Q$ and for $n \in \mathbb{N}$, let $\#paths(p, q, n)$ denote the number of paths (each one labeled by a distinct word) of length n from p to q in \mathcal{A} . We define the generating functions: $C_{p,q}(z) = \sum_{n \in \mathbb{N}} \#paths(p, q, n)z^n$ and $F_q(z) = C_{q_0,q}(z) \sum_{p \in Q} C_{q,p}(z)$. Let $[z^n]F_q(z)$ be the coefficient of z^n in $F_q(z)$. By the definition of $C_{q_0,q}$, we get

$$[z^n]F_q(z) = \sum_{k=0}^n \#paths(q_0, q, k) \sum_{p \in Q} \#paths(q, p, n-k).$$

Therefore, $[z^n]F_q(z)$ is the total number of times the state q is used when listing all paths of length n from q_0 .

Thus, we have $\sum_{u \in pref(L, n)} scost(u) = \frac{1}{n} \sum_{q \in Q} scost(q) [z^n]F_q(z)$. Finally, let $S(z) = \sum_{p \in P} C_{q_0,p}(z)$. Then, $wcost(\mathcal{A}) = \lim_{n \rightarrow \infty} \frac{1}{n \cdot [z^n]S(z)} \sum_{q \in Q} scost(q) [z^n]F_q(z)$. In Appendix A.1 we use techniques from [4] and [14] to compute the latter limit in polynomial time, by asymptotic estimations of the coefficients, thus concluding the proof. ◀

Sensing cost of languages

For a safety language L , we define $lcost(L) = \inf\{lcost(\mathcal{A}) : \mathcal{A} \text{ is a DLW for } L\}$, and similarly for $wcost(L)$. Different DLWs for a language L may have different sensing costs. We show that the minimal sensing cost in both approaches is attained at the minimal-size DLW. We first need some definitions and notations.

Consider a safety language $L \subseteq \Sigma^\omega$. For two finite words u_1 and u_2 , we say that u_1 and u_2 are *right L -indistinguishable*, denoted $u_1 \sim_L u_2$, if for every $z \in \Sigma^\omega$, we have that $u_1 \cdot z \in L$ iff $u_2 \cdot z \in L$. Thus, \sim_L is the Myhill-Nerode right congruence used for minimizing DFAs. For $u \in \Sigma^*$, let $[u]$ denote the equivalence class of u in \sim_L and let $\langle L \rangle$ denote the set of all equivalence classes. Each class $[u] \in \langle L \rangle$ is associated with the *residual language* $u^{-1}L = \{w : uw \in L\}$. Note that for safety languages, there is at most one class $[u]$, namely the class of bad prefixes, such that $u^{-1}L = \emptyset$. We denote this class $[\perp]$. When $L \neq \emptyset$ is a regular safety language, the set $\langle L \rangle$ is finite, and induces the *residual automaton* of L , defined by $\mathcal{R}_L = \langle \Sigma, \langle L \rangle \setminus \{[\perp]\}, \delta_L, [\epsilon] \rangle$, with $\delta_L([u], a) = [u \cdot a]$ for all $[u] \in \langle L \rangle \setminus \{[\perp]\}$ and $a \in \Sigma$ such that $[u \cdot a] \neq [\perp]$. The automaton \mathcal{R}_L is well defined and is the unique minimal-size DLW for L .

Consider a DLW $\mathcal{A} = \langle \Sigma, Q, q_0, \delta \rangle$ such that $L(\mathcal{A}) = L$. For a state $s = [u] \in \langle L \rangle \setminus \{[\perp]\}$, we associate with s a set $states(\mathcal{A}, s) = \{q \in Q : L(\mathcal{A}^q) = u^{-1}L\}$. That is, $states(\mathcal{A}, s) \subseteq Q$ contains exactly all state that \mathcal{A} can be in after reading a word that leads \mathcal{R}_L to $[u]$.

The following claims are simple exercises.

► **Proposition 3.1.** Consider a safety language L and a DLW \mathcal{A} for it.

1. The set $\{states(\mathcal{A}, s) : s \in \langle L \rangle \setminus \{[\perp]\}\}$ forms a partition of the states of \mathcal{A} .
2. For every state $s \in \langle L \rangle \setminus \{[\perp]\}$ of \mathcal{R}_L , letter $\sigma \in \Sigma$, and state $q \in states(\mathcal{A}, s)$, we have $\delta(q, \sigma) \in states(\mathcal{A}, \delta_L(s, \sigma))$.

► **Lemma 7.** Consider a safety language $L \subseteq \Sigma^\omega$. For every DLW \mathcal{A} with $L(\mathcal{A}) = L$, we have that $lcost(\mathcal{A}) \geq lcost(\mathcal{R}_L)$ and $wcost(\mathcal{A}) \geq wcost(\mathcal{R}_L)$


Proof. We outline the key points in the proof for $lcost$. The arguments for $wcost$ are similar. For a detailed proof see Appendix A.2.

Recall that the states of \mathcal{R}_L are $\langle L \rangle \setminus \{[\perp]\}$. We start by showing that for every $s \in \langle L \rangle \setminus \{[\perp]\}$ and for every $q \in states(\mathcal{A}, s)$ we have that $scost(q) \geq scost(s)$. Next, we consider the Markov chains $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{R}_L}$. Using Proposition 3.1 we show that if π and τ are the limiting distributions of $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{R}_L}$ respectively, then for every $s \in \langle L \rangle \setminus \{[\perp]\}$ we have that $\tau(s) = \sum_{q \in states(\mathcal{A}, s)} \pi(q)$. Finally, since Q is partitioned by $\{states(\mathcal{A}, s)\}_s$ we conclude that $lcost(\mathcal{A}) \geq lcost(\mathcal{R}_L)$. ◀

Lemma 7 and Theorems 5 and 6 allow us to conclude with the following.

► **Theorem 8.** Given a DLW \mathcal{A} , we can compute $lcost(L(\mathcal{A}))$ and $wcost(L(\mathcal{A}))$ in polynomial time.

► **Example 9.** Consider the DLW \mathcal{A} over the alphabet $2^{\{a,b\}}$ appearing in Figure 2.



► **Figure 2** A DLW for $(\neg a \vee \neg b)^\omega + (\neg a \vee \neg b)^* \cdot (a \wedge b) \cdot a^\omega$.

Clearly, \mathcal{A} is a minimal automaton for $L = (\neg a \vee \neg b)^\omega + (\neg a \vee \neg b)^* \cdot (a \wedge b) \cdot a^\omega$. By Lemma 7, we can calculate the sensing cost of \mathcal{A} in order to find the sensing cost of L .

Clearly, $scost(q_0) = 2$ and $scost(q_1) = 1$. We start by computing $lcost(\mathcal{A})$. The corresponding Markov chain $\mathcal{M}_{\mathcal{A}}$ has only one ergodic component $\{q_1\}$, so we obtain $lcost(\mathcal{A}) = scost(q_1) = 1$.

The computation of $wcost(\mathcal{A})$ is more intricate. In Appendix A.3 we show that $wcost(\mathcal{A}) = 2$. We remark that unlike in the other versions of sensing cost, transient components can play a role in $wcost$. In particular, If the self-loop on q_0 has been labeled by two rather than three letters, then we would have gotten $wcost(\mathcal{A}) = \frac{3}{2}$. ◀

4 Synthesis

In the setting of synthesis, the signals in P are partitioned into sets I and O of input and output signals. An I/O -transducer \mathcal{T} senses only input signals and we define its sensing cost as the sensing cost of the DLW $\mathcal{A}_{\mathcal{T}}$ it induces.


We define the I/O -sensing cost of a realizable specification $L \in (2^{I \cup O})^\omega$ as the minimal cost of an I/O -transducer that realizes L . Thus, $scost_{I/O}(L) = \inf\{scost(\mathcal{T}) : \mathcal{T} \text{ is an } I/O\text{-transducer}\}$

that realizes L }. In this section we consider the problem of finding a minimally-sensing I/O -transducer that realizes L .


The realizability problem for DLW specifications can be solved in polynomial time. Indeed, given a DLW \mathcal{A} , we can view \mathcal{A} as a game between a system, which controls the outputs, and an environment, which controls the inputs. We look for a strategy for the system that never reaches an undefined transition. This amounts to solving a turn-based safety game, which can be done in polynomial time.

When sensing is introduced, it is not enough for the system to win this game, as it now has to win while minimizing the sensing cost. Intuitively, not sensing some inputs introduces incomplete information to the game: once the system gives up sensing, it may not know the state in which the game is and knows instead only a set of states in which the game may be. In particular, unlike usual realizability, a strategy that minimizes the sensing need not use the state space of the DLW. We start with an example illustrating this.


► **Example 10.** Consider the DLW \mathcal{A} appearing in Figure 3. The DLW is over $I = \{p, q\}$ and $O = \{a\}$. A realizing transducer over the structure of \mathcal{A} (see \mathcal{T}_1 in Figure 4) senses p and q , responds with a if $p \wedge q$ was sensed and responds with $\neg a$ if $\neg p \wedge \neg q$ was sensed. In case other inputs are sensed, the response is arbitrary (denoted $*$ in the figure). As \mathcal{T}_1 demonstrates, every transducer that is based on the structure of \mathcal{A} senses two input signals (both p and q) every second step, thus its sensing cost is 1. As demonstrated by the transducer \mathcal{T}_2 in Figure 5, it is possible to realize \mathcal{A} with sensing cost of $\frac{1}{2}$ by only sensing p every second step. Indeed, knowing the value of p is enough in order to determine the output. Note that \mathcal{T}_2 may output sometimes a and sometimes $\neg a$ after reading assignments that causes \mathcal{A} to reach q_3 . Such a behavior cannot be exhibited by a transducer with the state-structure of \mathcal{A} . ◀



■ **Figure 3** The DLW \mathcal{A} in Example 10.



■ **Figure 4** The transducer \mathcal{T}_1 for \mathcal{A} .



■ **Figure 5** The transducer \mathcal{T}_2 for \mathcal{A} .

Solving games with incomplete information is typically done by some kind of a subset - construction, which involves an exponential blow up. Unlike usual games with incomplete information, here the strategy of the system should not only take care of the realizability but also decides which input signals should be sensed, where the goal is to obtain a minimally sensing transducer. In order to address these multiple objectives, we first construct an MDP in which the possible policies are all winning for the system, and corresponds to different choices of sensing. An optimal policy in this MDP then induces a minimally-sensing transducer.

► **Theorem 11.** Consider a DLW \mathcal{A} over $2^{I \cup O}$. If \mathcal{A} is realizable, then there exists an MDP \mathcal{M} in which an optimal strategy corresponds to a minimally-sensing I/O -transducer that realizes \mathcal{A} . The MDP \mathcal{M} has size exponential in $|\mathcal{A}|$ and can be computed in time exponential in $|\mathcal{A}|$.

Proof. Consider a DLW $\mathcal{A} = \langle 2^{I \cup O}, Q, q_0, \delta \rangle$. We obtain from \mathcal{A} an MDP $\mathcal{M} = \langle \mathcal{S}, \text{START}, A, P, \text{cost} \rangle$, where $\mathcal{S} = (2^Q \times \{0, 1, \perp\}^I) \cup \{\text{START}\}$, and $A = 2^I \times 2^O$. Intuitively, when \mathcal{M} is in state $\langle S, \ell \rangle$, for $S \subseteq Q$ and $\ell : I \rightarrow \{0, 1, \perp\}$, then \mathcal{A} can be in every state in S , and for each input signal $b \in I$, we have that either b is true ($\ell(b) = 1$), b is false ($\ell(b) = 0$), or b is not sensed ($\ell(b) = \perp$). The action $\langle o, i \rangle$ means that we now output o and in the next state we will sense only inputs in i . For $\star \in \{\perp, 0, 1\}$, we define $\ell_\star = \{b \in I : \ell(b) = \star\}$.

We define the actions so that an action $\langle o, i \rangle$ is available in state $\langle S, \ell \rangle$ if for every $q \in S$ and $i' \subseteq \ell_\perp$, we have that $\delta(q, \ell_1 \cup i' \cup o)$ is defined. That is, an action is available if its o component does not cause \mathcal{A} to get stuck no matter what the assignment to the signals that are not sensed is.

The transition probabilities are defined as follows. Consider a state $\langle S, \ell \rangle$, and an available action $\langle o, i \rangle$. Let $S' = \bigcup_{q \in S} \bigcup_{i' \subseteq \ell_\perp} \{\delta(q, \ell_1 \cup i' \cup o)\}$. Recall that by taking action $\langle o, i \rangle$, we decide that in the next state we will only sense signals in i . For $i \subseteq I$, we say that an assignment $\ell' : I \rightarrow \{0, 1, \perp\}$ senses i if $\ell'_1 \cup \ell'_0 = i$. Note that there are $2^{|i|}$ assignments that sense i . Accordingly, we have $P(\langle S, \ell \rangle, \langle o, i \rangle, \langle S', \ell' \rangle) = 2^{-|i|}$ for every $\ell' : I \rightarrow \{0, 1, \perp\}$ that senses i . That is, a transition from $\langle S, \ell \rangle$ with $\langle o, i \rangle$ goes to the set of all possible successors of S under inputs that are consistent with ℓ and the output assignment o , and the ℓ' component is selected with uniform distribution among all assignments that sense i . The cost function depends on the number of signals we sense, thus $\text{cost}(\langle S, \ell \rangle) = |\ell_1 \cup \ell_0|$.

Finally, in the state START we only choose an initial set of input signals to sense. Thus, for every ℓ such that $\ell_1 \cup \ell_0$, we have $P(\text{START}, \langle o, i \rangle, \langle \{q_0\}, \ell \rangle) = 2^{-|i|}$. Note that START is not reachable from any state in \mathcal{M} , and thus its cost is irrelevant. We arbitrarily set $\text{cost}(\text{START}) = 0$.

In Appendix A.4 we prove that $\text{cost}(\mathcal{M}) = \text{scost}_{I,O}(\mathcal{A})$ and that a minimal-cost policy f in \mathcal{M} induces a minimally-sensing I/O -transducer that realizes \mathcal{A} . Intuitively, we prove this by showing a correspondence between transducers and policies, such that the sensing cost of a transducer \mathcal{T} equals the value of the policy it corresponds to in \mathcal{M} .

Finally, we observe that the size of \mathcal{M} is single exponential in the size of \mathcal{A} , and that we can construct \mathcal{M} in time exponential in the size of \mathcal{A} . ◀

► **Theorem 12.** *A minimally-sensing transducer for a realizable DLW \mathcal{A} has size tightly exponential in $|\mathcal{A}|$.*

Proof. The upper bound follows from Theorem 3 applied to the MDP constructed in Theorem 11.

For the lower bound, we describe a family of realizable DLWs $\mathcal{A}_1, \mathcal{A}_2, \dots$ such that for all $k \geq 1$, the DLW \mathcal{A}_k has $1 + \sum_{i=1}^k p_i$ states, yet a minimally-sensing transducer for it requires at least $\prod_{i=1}^k p_i$ states, where p_1, p_2, \dots are prime numbers. Intuitively, \mathcal{A}_k is constructed as follows. In the initial state q_{reset} , the inputs signals determine a number $1 \leq i \leq k$, and \mathcal{A}_k moves to component i , which consists of a cycle of length p_i . In every state j in component i , the output signals must acknowledge that \mathcal{A}_k is in state $0 \leq j < p_i$ of component i . Furthermore, we force a sensing of 1 in every state except for q_{reset} by requiring a signal to be acknowledged in every step. Finally, we can go back to q_{reset} only with a special output signal, which can be outputted only in state 0 of an i component.

Thus, a realizing transducer essentially only chooses which signals to read in q_{reset} . We show that 0 bits can be read, but in that case we need $\prod_{i=1}^k p_i$ states. Indeed, the transducer needs to keep track of the location in all the i components simultaneously, which means keeping track of the modulo from each p_i . Since every combination of such modulus is possible, the transducer needs $\prod_{i=1}^k p_i$ states. In Appendix A.5 we formalize this intuition. ◀

We now turn to study the complexity of the problem of finding a minimally-sensing transducer. By the construction in Theorem 11 and the polynomial time algorithm from Theorem 3, we have the following.

► **Theorem 13.** *Consider a realizable DLW \mathcal{A} over $2^{I \cup O}$. We can calculate $\text{cost}_{I,O}(\mathcal{A})$ and return a minimally-sensing I/O -transducer that realizes \mathcal{A} in time exponential in $|\mathcal{A}|$.*

In order to complete the picture, we consider the corresponding decision problem. Given a DLW \mathcal{A} over $2^{I \cup O}$ and a threshold γ , the sensing problem in the open setting is to decide whether $\text{cost}_{I,O}(\mathcal{A}) < \gamma$.

► **Theorem 14.** *The sensing problem in the open setting is EXPTIME-complete.*

Proof. The upper bound follows from Theorem 13. For the lower bound, we show that the problem is EXPTIME hard even for a fixed γ . Given a DLW specification \mathcal{A} over $2^{I \cup O}$, we show that it is EXPTIME-hard to decide whether there exists a transducer \mathcal{T} that realizes \mathcal{A} with $\text{scost}(\mathcal{T}) < 1$. We show a reduction from the problem of deciding the nonemptiness of an intersection of finite deterministic tree automata proved to be EXPTIME-hard in [5]. The idea is similar to that of Theorem 12, where a reset state is used to select an object, and a transducer can ignore the inputs in this state by using a response which is acceptable in every possible selected object.

A deterministic automaton on finite trees (DFT) is $\mathcal{U} = \langle \Sigma, Q, \delta, q_0, F \rangle$, where Σ is a finite alphabet, Q is a finite set of states, $q_0 \in Q$ is an initial state, $\delta : Q \times \Sigma \rightarrow Q \times Q$ is a transition function, and $F \subseteq Q$ is a set of accepting states. We refer to the left and right components of δ as δ_{\triangleleft} and δ_{\triangleright} . For example, when $\delta(q, \sigma) = \langle q_l, q_r \rangle$, we write $\delta_{\triangleleft}(q, \sigma) = q_l$. An DFT runs on Σ -trees. A (binary) Σ -tree is $T = \langle \tau, \ell \rangle$ where $\tau \subseteq \{\triangleleft, \triangleright\}^*$ is prefix-closed: for every $x \cdot \sigma \in \tau$ it holds that $x \in \tau$, and $\ell : \tau \rightarrow \Sigma$ is a labeling function. For simplicity, we require that for every $x \in \tau$, either $\{x\triangleleft, x\triangleright\} \subseteq \tau$, or $\{x\triangleleft, x\triangleright\} \cap \tau = \emptyset$, in which case x is a leaf. Given a tree $T = \langle \tau, \ell \rangle$, the run of \mathcal{U} on T is a Q -tree $\langle \tau, \ell' \rangle$ where $\ell'(\epsilon) = q_0$, and for every $x \in \tau$ such that x is not a leaf, we have $\delta(\ell'(x), \ell(x)) = \langle \ell'(x\triangleleft), \ell'(x\triangleright) \rangle$. A run is *accepting* if every leaf is labeled by an accepting state. A Σ -tree T is accepted by \mathcal{U} if the run of \mathcal{U} on T is accepting.

The nonempty-intersection problem gets as input DFTs $\mathcal{U}_1, \dots, \mathcal{U}_n$, and decides whether their intersection is nonempty, that is $\bigcap_{t=1}^n L(\mathcal{U}_t) \neq \emptyset$. Given $\mathcal{U}_1, \dots, \mathcal{U}_n$, we construct a specification DLW \mathcal{A} such that $\bigcap_{t=1}^n L(\mathcal{U}_t) \neq \emptyset$ iff $\text{scost}(\mathcal{A}) < 1$. We assume w.l.o.g. that $L(\mathcal{U}_t) \neq \emptyset$ for all $1 \leq t \leq n$.

We construct \mathcal{A} as follows. Initially, the inputs specify an index $1 \leq t \leq n$. Then, the transducer should respond with a tree in $L(\mathcal{U}_t)$. This is done by challenging the transducer with a branch in the tree, until some reset input signal is true, and the process repeats. Now, if $\bigcap_{t=1}^n L(\mathcal{U}_t) \neq \emptyset$, the transducer can ignore the input signals that specify the index t and just repeatedly output a tree in the intersection. On the other hand, if $\bigcap_{t=1}^n L(\mathcal{U}_t) = \emptyset$, the transducer must sense some information about the specified index.³

We now formalize this intuition. For $1 \leq t \leq n$, let $\mathcal{U}_t = \langle 2^J, Q^t, \delta^t, q_0^t, F^t \rangle$. Note that we assume w.l.o.g. that the alphabet of all the DFTs is 2^J . We construct a specification DLW $\mathcal{A} = \langle 2^{I \cup O}, Q, q_0, \delta \rangle$ as follows. The set of states of \mathcal{A} is $Q = \bigcup_{t=1}^n Q^t \cup \{\text{RESET}\}$. Assume w.l.o.g. that $n = 2^k$ for some $k \in \mathbb{N}$. We define $I = \{b_1, \dots, b_k\} \cup \{dI\}$ and $O = J \cup \{dO, e\}$. The input signal dI and the output signal dO denote the direction of branching in the tree. For clarity, in an input letter $i \in I$ we write $i(dI) = \triangleleft$ (and $i(dI) = \triangleright$) to indicate that $dI \notin i$ (and $dI \in i$). We use a similar notation for dO .

We define the transition function as follows. In state RESET, we view the inputs b_1, \dots, b_k as a binary encoding of a number $t \in \{1, \dots, n\}$. Then, $\delta(\text{RESET}, t) = q_0^t$. Next, consider a state

³ Note that since a tree in the intersection of DFTs may be exponentially bigger than the DFTs, the lower bound here also suggests an alternative lower bound to the exponential size of a minimally-sensed transducer, now with a polynomial set of signals (as opposed to the proof of Theorem 12).

$q \in Q^t$, and consider letters $i \subseteq I$ and $o \subseteq O$. We define δ as follows:

$$\delta(q, i \cup o) = \begin{cases} \text{RESET} & q \in F \wedge e \in o \wedge o(dO) = i(dI) \\ \delta_{\triangleleft}^t(q, o \cap J) & e \notin o \wedge o(dO) = i(dI) = \triangleleft \\ \delta_{\triangleright}^t(q, o \cap J) & e \notin o \wedge o(dO) = i(dI) = \triangleright \end{cases}$$

Note that $\delta(q, i \cup o)$ is undefined when $o(dO) \neq i(dI)$ or when $q \notin F$ and $e \in o$. Intuitively, in state RESET, an index $1 \leq t \leq n$ is chosen. From then on, in a state $q \in Q^t$, we simulate the run of \mathcal{U}_t on the left or right branch of the tree, depending on the signal dI . The next letter is outputted in o , and additionally, we require that dO matches dI .

We claim that $\text{scost}(\mathcal{A}) < 1$ iff $\bigcap_{t=1}^n L(\mathcal{U}_t) \neq \emptyset$. In the first direction, assume that $\bigcap_{t=1}^n L(\mathcal{U}_t) \neq \emptyset$, and let T be a tree such that $T \in \bigcap_{t=1}^n L(\mathcal{U}_t) \neq \emptyset$. Consider the following transducer \mathcal{T} : in the state RESET it does not sense any inputs, and then it outputs a branch of T according to the signal dI , while always acknowledging the dI bit with the correct dO . When the end of the branch is reached, it outputs e . Since T is accepted by every DFT U^t , it follows that \mathcal{T} realizes \mathcal{A} . Moreover, let l be the longest branch in T , then every l steps at most, \mathcal{T} visits a state corresponding to RESET, in which it senses nothing. Thus, \mathcal{T} senses 1 for at most l steps, and then 0. It follows that $\text{scost}(\mathcal{T}) \leq \frac{l}{l+1} = 1 - \frac{1}{l+1} < 1$.

Conversely, observe that in every state $q \in Q \setminus \{\text{RESET}\}$, a realizing transducer must sense at least 1 signal, namely dI . Thus, the only way to get sensing cost of less than 1 is to visit RESET infinitely often (in fact, with bounded sparsity), and to sense 0 in RESET. However, sensing 0 in RESET means that the next state could be the initial state of any of the n DFTs. Moreover, visiting RESET again means that at some point e was outputted in an accepting state of one of the DFTs. Thus, the transducer outputs a tree that is accepted in every DFT, so $\bigcap_{t=1}^n L(\mathcal{U}_t) \neq \emptyset$.

Finally, observe that the reduction is clearly polynomial, and thus we conclude that deciding whether $\text{scost}(\mathcal{A}) < 1$ is EXPTIME-hard. \blacktriangleleft

5 Discussion and Future Research

Sensing is a basic measure of the complexity of monitoring and synthesis. In monitoring safety properties, the definition of sensing presented in [2] is not informative, as it gives sensing cost 0 to properties that are satisfied with probability 0. We argue that in the context of monitoring, the definition of sensing cost should consider only computations that satisfy the property, and we study the complexity of computing the sensing cost of a property in the new definition. We distinguish between two approaches to define a probabilistic measure with respect to the set of computations that satisfy a property. We show that while computing the sensing cost according to the new definitions is technically more complicated than in [2], the minimal sensing is still attained by a minimal-size automaton, and it can still be computed in polynomial time.

In synthesis, we introduce a new degree of freedom, namely choosing the outputs when realizing a specification. We study the complexity of finding a minimal-sensing transducer for safety specifications. We show that the minimal-sensing transducer is not necessarily minimal in size. Moreover, interestingly, unlike the case of traditional synthesis, a minimal-sensing transducer need not even correspond to a strategy embodied in the specification deterministic automaton. On the positive side, we show that a minimal-sensing transducer always exists (for a realizable safety specification) and that its size is at most exponential in the size of the minimal-size transducer. We also provide matching lower bounds.

We now turn to discuss some future directions for research.

Non-safety properties We focus on safety properties. The study in [2] completes the monitoring picture for all other ω -regular properties. We plan to continue the study of synthesis of ω -regular

properties. An immediate complication in this setting is that a finite minimal-sensing transducer does not always exist. Indeed, even in the monitoring setting studied in [2], a minimal-sensing automaton does not always exist. Even the decidability of computing the optimal cost remains open.

A trade-off between sensing and quality Reducing the sensing cost of a transducer can often be achieved by *delaying* the sensing of some letter, thus sensing it less often. This, however, means that eventualities may take longer to be fulfilled, resulting in transducers of lower quality [1]. We plan to formalize and study the trade-off between the sensing and quality and relate it to the trade-offs between size and sensing, as well as between size and quality.

Acknowledgment We thank Elie de Panafieu for helpful discussions.

References

- 1 S. Almagor, U. Boker, and O. Kupferman. Discounting in LTL. In *Proc. 20th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science. Springer, 2014.
- 2 Shaull Almagor, Denis Kuperberg, and Orna Kupferman. Regular sensing. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India*, pages 161–173, 2014.
- 3 D.L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52:1289–1306, 2006.
- 4 Philippe Flajolet and Robert Sedgewick. Analytic combinatorics: functional equations, rational and algebraic functions. 2001.
- 5 Jean Goubault. Rigid e-unifiability is dextime-complete. In *Proceedings of the Ninth Annual Symposium on Logic in Computer Science (LICS '94), Paris, France, July 4-7, 1994*, pages 498–506, 1994.
- 6 C. Grinstead and J. Laurie Snell. Chapter 11: markov chains. In *Introduction to Probability*. American Mathematical Society, 1997.
- 7 K. Havelund and G. Rosu. Efficient monitoring of safety properties. *Software Tools for Technology Transfer*, 6(2):18–173, 2004.
- 8 G. Kindler. *Property Testing, PCP, and Juntas*. PhD thesis, Tel Aviv University University, 2002.
- 9 O. Kupferman and M.Y. Vardi. Synthesis with incomplete information. In *Advances in Temporal Logic*, pages 109–127. Kluwer Academic Publishers, 2000.
- 10 E. Kushilevitz and N. Nisan. *Communication complexity*. Cambridge University Press, 1997.
- 11 C. Mauduit and A. Sárköz. On finite pseudorandom binary sequences. i. measure of pseudorandomness, the legendre symbol. *Acta Arith.*, 82(4):365–377, 1997.
- 12 S. Muthukrishnan. Theory of data stream computing: where to go. In *Proc. 30th Symposium on Principles of Database Systems*, pages 317–319, 2011.
- 13 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.
- 14 Marie-Françoise Roy and Aviva Szpirglas. Complexity of the computation on real algebraic numbers. *J. Symb. Comput.*, 10(1):39–52, 1990.
- 15 A.P. Sistla. Safety, liveness and fairness in temporal logic. *Formal Aspects of Computing*, 6:495–511, 1994.

A Appendix

A.1 Proof of Theorem 6

By [4], for every $p, q \in Q$, we can compute in polynomial time (using standard algorithms on matrices) rational expressions for $C_{p,q}(z)$. The base case is for computing coefficients in the same irreducible aperiodic SCC represented by a matrix M : it suffices to compute the matrix $R(z) = (Id - zM)^{-1}$, its coefficient (p, q) is $C_{p,q}(z)$. For instance if $\{p\}$ is a SCC in \mathcal{A} with a self-loop labeled by k letters, then $C_{p,p}(z) = \frac{1}{1-kz}$. Other $C_{p,q}$ are then computed from this base case via standard operations on rational functions. In particular, from [4] there is a period $d \leq |Q|$ such that for every $i \in \{0, \dots, d-1\}$ and for all rational functions $Q(z) \in \{S(z)\} \cup \bigcup_{q \in Q} \{F_q(z)\}$ considered here, we can compute in polynomial time γ, k , and λ such that $[z^{nd+i}]Q(z) \sim \gamma(nd+i)^k \lambda^{nd+i}$ (where for functions $f, g : \mathbb{N} \rightarrow \mathbb{R}$ we have $f(n) \sim g(n)$ iff $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 1$). We remind the formula that allows us to do so. Let $Q(z) = \frac{A(z)+B(z)(1-\frac{z}{r})^{-j}}{z^i}$ be a rational function of convergence radius r , where $i, j \in \mathbb{N}$, $A(z)$ and $B(z)$ have convergence radius strictly greater than r , and $B(r) \neq 0$. Then we have

$$[z^n]Q(z) \sim \frac{B(r)}{(j-1)! \cdot r^i} n^{j-1} (1/r)^n.$$

Notice that r will in general be a real algebraic number, that will be represented in our algorithm as a root of a polynomial with integer coefficients. Standard operations as sum, product, and comparisons on algebraic numbers (represented by polynomials) can be done in polynomial time, using techniques as described in [14].

Therefore, we can compute asymptotic equivalents of the form $\alpha \cdot n^k \cdot \lambda^n$ with α, λ real algebraic and $k \in \mathbb{N}$ [4], for both $[z^n]S(z)$ and $\frac{1}{n} \sum_{q \in Q} scost(q)[z^n]F_q(z)$, performing an averaging operation if $d > 1$. Finally, we can compute the wanted limit thanks to these asymptotic equivalents, thereby achieving the polynomial-time computation of $wcost(\mathcal{A})$.

A.2 Detailed proof of Lemma 7

We start with $lcost$. Consider a finite word $u \in \Sigma^*$ that is not a bad prefix for L . After reading u , the DLW \mathcal{R}_L reaches the state $[u]$ and the DLW \mathcal{A} reaches a state q with $L(\mathcal{A}^q) = u^{-1}L$. Indeed, otherwise we can point to a word with prefix u that is accepted only in one of the DLWs. We claim that for every state $q \in Q$ such that $L(\mathcal{A}^q) = u^{-1}L$, it holds that $sensed([u]) \subseteq sensed(q)$. To see this, consider a signal $p \in sensed([u])$. By definition, there exists a set $S \subseteq P$ and words u_1 and u_2 such that $([u], S \setminus \{p\}, [u_1]) \in \Delta_L$, $([u], S \cup \{p\}, [u_2]) \in \Delta_L$, yet $[u_1] \neq [u_2]$. By the definition of \mathcal{R}_L , there exists $z \in (2^P)^*$ such that, w.l.o.g, $z \in u_1^{-1}L \setminus u_2^{-1}L$. Hence, as $L(\mathcal{A}^q) = u^{-1}L$, we have that \mathcal{A}^q accepts $(S \setminus \{p\}) \cdot z$ and rejects $(S \cup \{p\}) \cdot z$. Let $\delta_{\mathcal{A}}$ be the transition function of \mathcal{A} . By the above, $\delta_{\mathcal{A}}(q, S \setminus \{p\}) \neq \delta_{\mathcal{A}}(q, S \cup \{p\})$. Therefore, $p \in sensed(q)$, and we are done. Now, $sensed([u]) \subseteq sensed(q)$ implies that $scost(q) \geq scost([u])$. Since our assumption on q is only that $L(\mathcal{A}^q) = u^{-1}L$, we get that $q \in states(\mathcal{A}, [u])$. Thus, we conclude that for every $s \in \langle L \rangle \setminus \{\perp\}$ and for every $q \in states(\mathcal{A}, s)$ we have that $scost(q) \geq scost(s)$.

Next, consider the Markov chains $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{R}_L}$, and let P and R be their respective transition matrices. We index the rows and columns of P (resp. R) by Q (resp. $\langle L \rangle \setminus \{\perp\}$). Let $v^0 \in [0, 1]^Q$ be the initial vector for \mathcal{A} , thus $v^0(q_0) = 1$ and $v^0(q) = 0$ for $q \neq q_0$. Similarly, let $u^0([\epsilon]) = 1$ and $u^0([u]) = 0$ for all $[u] \neq [\epsilon]$. For $m \in \mathbb{N}$, let $v^m = v^0 P^m$ and $u^m = u^0 R^m$.

We claim that for every $s \in \langle L \rangle \setminus \{\perp\}$ it holds that $u^m(s) = \sum_{q \in states(\mathcal{A}, s)} v^m(q)$.

The proof of the claim proceeds by an induction on m . For $m = 0$, we have $q_0 \in [\epsilon]$, and the claim follows trivially. Assume correctness for m , we prove the claim for $m + 1$.

Consider states $s, s' \in \langle L \rangle \setminus \{\perp\}$. For every state $q \in \text{states}(\mathcal{A}, s)$ it holds that

$$R_{s,s'} = \frac{|\{\sigma : \delta_L(s, \sigma) = s'\}|}{\text{dom}(s)} = \frac{|\{\sigma : \delta(q, \sigma) = q' \in \text{states}(\mathcal{A}, s')\}|}{\text{dom}(q)} = \sum_{q' \in \text{states}(\mathcal{A}, s')} P_{q,q'} \quad (1)$$

where the second equality follows from Observation 3.1 and by observing that $\text{dom}(q) = \text{dom}(s)$. The latter holds since if $\sigma \in \text{dom}(q)$, then there exists $q' \in Q$ such that $q' = \delta(q, \sigma)$, so $q' \in \text{states}(\delta_L(s, \sigma))$ and $\sigma \in \text{dom}(s)$, so $\text{dom}(q) \subseteq \text{dom}(s)$, and conversely - if $\sigma \in \text{dom}(s)$ then by Proposition 3.1 we have that $\sigma \in \text{dom}(q)$, so $\text{dom}(s) \subseteq \text{dom}(q)$.

Now, for every state $s' \in \langle L \rangle \setminus \{\perp\}$, we have that

$$\begin{aligned} u^{m+1}(s') &= \sum_{s \in \langle L \rangle \setminus \{\perp\}} R_{s,s'} u^m(s) \\ &= \sum_{s \in \langle L \rangle \setminus \{\perp\}} R_{s,s'} \sum_{q \in \text{states}(\mathcal{A}, s)} v^m(q) && \text{(Induction Hypothesis)} \\ &= \sum_{s \in \langle L \rangle \setminus \{\perp\}} \sum_{q \in \text{states}(\mathcal{A}, s)} R_{s,s'} v^m(q) \\ &= \sum_{s \in \langle L \rangle \setminus \{\perp\}} \sum_{q \in \text{states}(\mathcal{A}, s)} \sum_{q' \in \text{states}(\mathcal{A}, s')} P_{q,q'} v^m(q) && \text{(Equation (1))} \\ &= \sum_{q' \in \text{states}(\mathcal{A}, s')} \sum_{s \in \langle L \rangle \setminus \{\perp\}} \sum_{q \in \text{states}(\mathcal{A}, s)} P_{q,q'} v^m(q) \\ &= \sum_{q' \in \text{states}(\mathcal{A}, s')} \sum_{q \in Q} P_{q,q'} v^m(q) && \text{(Proposition 3.1)} \\ &= \sum_{q' \in \text{states}(\mathcal{A}, s')} v^{m+1}(q') \end{aligned}$$

and the induction is complete.

Now, let π and τ be the limiting distributions of $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{R}_L}$ respectively, then $\pi = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{m=1}^n v^0 P^n$ and $\tau = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{m=1}^n u^0 R^n$, and by the above we have that $\tau(s) = \sum_{q \in \text{states}(\mathcal{A}, s)} \pi(q)$.

Since $\text{scost}(q) \geq \text{scost}(s)$ for every $q \in \text{states}(\mathcal{A}, s)$, we conclude that

$$\begin{aligned} \text{lcost}(\mathcal{A}) &= \sum_{q \in Q} \pi(q) \text{scost}(q) = \sum_{s \in \langle L \rangle \setminus \{\perp\}} \sum_{q \in \text{states}(\mathcal{A}, s)} \pi(q) \text{scost}(q) \\ &\geq \sum_{s \in \langle L \rangle \setminus \{\perp\}} \sum_{q \in \text{states}(\mathcal{A}, s)} \pi(q) \text{scost}(s) = \sum_{s \in \langle L \rangle \setminus \{\perp\}} \text{scost}(s) \sum_{q \in \text{states}(\mathcal{A}, s)} \pi(q) \\ &= \sum_{s \in \langle L \rangle \setminus \{\perp\}} \text{scost}(s) \tau(s) = \text{lcost}(\mathcal{R}_L). \end{aligned}$$

We proceed to $w\text{cost}$. Following similar arguments as above, we see that for every finite word w , we have $\text{scost}_{\mathcal{A}}(w) \geq \text{scost}_{\mathcal{R}_L}(w)$. Therefore, for any $n \geq 0$ we have $\sum_{w \in \text{pref}(L) \cap \Sigma^n} \text{scost}_{\mathcal{A}}(w) \geq \sum_{w \in \text{pref}(L) \cap \Sigma^n} \text{scost}_{\mathcal{R}_L}(w)$, and finally $w\text{cost}(\mathcal{A}) \geq w\text{cost}(\mathcal{R}_L)$. This shows that $w\text{cost}(\mathcal{R}_L) = w\text{cost}(L)$.

A.3 Computing $w\text{cost}(\mathcal{A})$ in Example 9

First, note that $\neg a \vee \neg b$ corresponds to 3 letters, $a \wedge b$ to 1 letter, and a to 2 letters.

We have $C_{q_0, q_0}(z) = \frac{1}{1-3z}$, $C_{q_1, q_1}(z) = \frac{1}{1-2z}$, and $C_{q_0, q_1}(z) = C_{q_0, q_0}(z)C_{q_1, q_1}(z) = \frac{1}{(1-3z)(1-2z)}$, whereas $C_{q_1, q_0}(z) = 0$.

Moreover, we have

$$\begin{aligned} S(z) &= C_{q_0, q_0}(z) + C_{q_0, q_1}(z) = \frac{2 - 5z}{(1 - 3z)(1 - 2z)} \\ F_{q_0}(z) &= C_{q_0, q_0}(z)(C_{q_0, q_0}(z) + C_{q_0, q_1}(z)) = \frac{2 - 5z}{(1 - 3z)^2(1 - 2z)} \\ F_{q_1}(z) &= C_{q_0, q_1}(z)C_{q_1, q_1}(z) = \frac{1}{(1 - 3z)(1 - 2z)} \end{aligned}$$

Using standard algorithms on rational functions, we get $[z^n]S(z) \sim \frac{2-5/3}{1-2/3}3^n = 3^n$, $[z^n]F_{q_0}(z) \sim n3^n$ and $[z^n]F_{q_1}(z) \sim 3^{n+1}$. We finally obtain $wcost(\mathcal{A}) = \lim_{n \rightarrow \infty} \frac{2 \cdot n3^n + 1 \cdot 3^{n+1}}{n3^n} = 2$.

Note that $wcost$, unlike $scost$ and $lcost$, allows to take into a consideration the cost of transient components when a long word in L is likely to spend time in them.

If the self-loop on q_0 has been labeled by two letters, say by b , rather than by three, then q_0 and q_1 would have participated equally and we would have gotten $wcost(\mathcal{A}) = 3/2$.

A.4 Proof of Theorem 11

We claim that $cost(\mathcal{M}) = scost_{I,O}(\mathcal{A})$, and that a minimal-cost policy f in \mathcal{M} induces a minimally-sensing I/O -transducer that realizes \mathcal{A} .

Consider a memoryless minimal-cost policy f . We construct from f a transducer $\mathcal{T} = \langle I, O, \mathcal{S}, \text{START}, \mu, \rho \rangle$, where μ and ρ are defined as follows. For sets $i_1 \subseteq i \subseteq I$, we say that an assignment $\ell : I \rightarrow \{0, 1, \perp\}$ the (unique) i -sensed i_1 -true assignment if for every signal $b \in I$, we have that $\ell(b)$ is \perp if $b \notin i$, is 1 if $b \in i_1$, and is 0 if $b \in i \setminus i_1$.

Let $f(\text{START}) = \langle o_0, i_0 \rangle$. We arbitrarily⁴ set $\rho(\text{START})$ to o_0 . For input $i \in 2^I$, we set $\mu(\text{START}, i) = \langle \{q_0\}, \ell_0 \rangle$, for the i_0 -sensed i -true assignment ℓ_0 .

Next, consider a state $\langle S, \ell \rangle$ and an input $i \in 2^I$. Let $\langle o, i' \rangle = f(\langle S, \ell \rangle)$. We define $\rho(\langle S, \ell \rangle) = o$ and $\mu(\langle S, \ell \rangle, i) = \langle S', \ell' \rangle$, where $S' = \bigcup_{q \in S} \bigcup_{i_1 \subseteq i \subseteq \perp} \delta(q, \ell_1 \cup i_1 \cup o)$ and ℓ' is the i' -sensed i -true assignment.

We claim that $scost(\mathcal{T}) = cost(f)$. To see this, let \mathcal{M}' be the Markov Chain obtained from \mathcal{M} by fixing the action in each state according to f , and let \mathcal{T}' be the Markov chain obtained from \mathcal{T} by assigning uniform distributions to the input signals. It is easy to see that the Markov chains \mathcal{M}' and \mathcal{T}' are identical (with the exception of START , which, as we mentioned, does not affect the cost). Thus, $cost(f) = scost(\mathcal{T})$. Since $cost(\mathcal{M}) = cost(f)$, we can conclude that there is a transducer \mathcal{T} that realizes \mathcal{A} and for which $scost(\mathcal{T}) = cost(\mathcal{M})$. Thus, $scost_{I,O}(\mathcal{A}) \leq cost(\mathcal{M})$.

Conversely, consider a transducer \mathcal{T} for \mathcal{A} . By following the set of sensed input signals and the output at each state, \mathcal{T} induces a (possibly non-memoryless) policy f in \mathcal{M} . Moreover, as above, $cost(f) = scost(\mathcal{T})$. Thus, $scost(\mathcal{T}) \geq cost(\mathcal{M})$. Since this holds for all transducers \mathcal{T} , it follows that $scost_{I,O}(\mathcal{A}) \geq cost(\mathcal{M})$.

A.5 Proof of Theorem 12

We define $\mathcal{A}_k = \{2^{I \cup O}, Q, q_{\text{reset}}, \delta\}$, where

- $I = \{i_1, \dots, i_{\lceil \log k \rceil}\} \cup \{dI\}$, and we view 2^I as $\{1, \dots, k\} \times \{dI, \neg dI\}$. Then, $O = \bigcup_{i=1}^k \{o_{i,1}, \dots, o_{i, \lceil \log p_i \rceil}\} \cup \{dO, e\}$, and we view it as $(\prod_{i=1}^k \{0, \dots, p_i - 1\}) \times \{dO, \neg dO\} \times \{e, \neg e\}$. Thus, a letter $\sigma \in 2^{I \cup O}$ is a pair $\sigma = \langle \sigma_I, \sigma_O \rangle$ with $\sigma_I = \langle m, dI \rangle$ where $1 \leq m \leq k$ and $dI \in \{0, 1\}$,

⁴ Since the output in START is ignored, this is indeed arbitrary.

and with $\sigma_O = \langle r_1, \dots, r_k, dO, e \rangle$ with $0 \leq r_i < p_i - 1$ for all $1 \leq i \leq k$, $dO \in \{0, 1\}$, and $e \in \{0, 1\}$.

- $Q = \{q_{\text{reset}}\} \cup \bigcup_{i=1}^k \{q_{i,0}, \dots, q_{i,p_i-1}\}$.
- The transition function δ is defined as follows. Consider a letter $\sigma = \langle \langle m, dI \rangle, \langle r_1, \dots, r_k, dO, e \rangle \rangle$. In state q_{reset} , we have $\delta(q_{\text{reset}}, \sigma) = q_{m,0}$. For a state $q_{i,j}$, we have

$$\delta(q_{i,j}, \sigma) = \begin{cases} q_{i,(j+1) \bmod p_i} & e = 0 \wedge r_i = j \wedge dI = dO \\ q_{\text{reset}} & j = 0 \wedge e = 1 \wedge r_i = j \wedge dI = dO \end{cases}$$

Note that $\delta(q_{i,j}, \sigma)$ is undefined in all other cases, thus when $dI \neq dO$, when $j \neq 0$ and $e = 1$, and when $r_i \neq j$.

Intuitively, in order to take a transition from $q_{i,j}$, r_i has to match j , and $dI = dO$. Providing that, if $e = 0$ then the transition progresses along the cycle in the i component, and if $e = 1$ and the state is $q_{i,0}$, then the run moves to q_{reset} .

Consider a transducer \mathcal{T} with $\text{scost}(\mathcal{T}) < 1$ that realizes \mathcal{A}_k . We show that \mathcal{T} must have at least $\prod_{i=1}^k p_i$ states. Observe that whenever \mathcal{A}_k is in a state that is different from q_{reset} , the transducer \mathcal{T} must sense at least dI in order to match dO . Thus, the only state that \mathcal{A}_k visits and in which the sensing cost can be lower than 1 (i.e. 0) is q_{reset} . Thus, \mathcal{T} senses 0 in q_{reset} , and moreover, \mathcal{A}_k has to visit q_{reset} infinitely often in order for the 0 sensing to reduce the sensing cost. Sensing 0 in q_{reset} means that in the next step, \mathcal{A}_k can be in $q_{i,0}$ for every $1 \leq i \leq k$. Since \mathcal{T} has to output $r_i = j$ in every $q_{i,j}$, then \mathcal{T} has to keep track of j , which runs from 0 to $p_i - 1$. Since \mathcal{T} does not “know” the value of i , then \mathcal{T} has to keep track of every reachable combination of r_1, \dots, r_k with $0 \leq r_i \leq p_i - 1$ for every $1 \leq i \leq k$. From the Chinese remainder theorem, every such possible combination of r_1, \dots, r_k is reachable. Thus, \mathcal{T} needs at least $\prod_{i=1}^k p_i$ states.

Finally, note that there is a transducer \mathcal{T} that realizes \mathcal{A}_k . Indeed, \mathcal{T} has $\prod_{i=1}^k p_i$ states and outputs e , thus causing \mathcal{A}_k to return to q_{reset} , every $\prod_{i=1}^k p_i$ steps.

B Non-uniform Distributions

In this section we consider non-uniform distributions, and show that incorporating them in the model does not involve any significant changes in our techniques. We start by defining the model.

In the closed setting, we are given a *labeled Markov chain* $\mathcal{M} = \langle S, P, \ell \rangle$ where ℓ is a labeling function $\ell : S \rightarrow \Sigma$. The probability of a letter $l \in \Sigma$ in state s is $P_s(l) = \sum_{s' \in S: \ell(s')=l} P(s, s')$. We assume that the probability has full support. That is, for every $s \in S$ and $l \in \Sigma$ it holds that $P_s(l) > 0$, and that all probabilities are rational. We lift the probability to words as follows. A sequence of states $\pi = s_1, \dots, s_n$ induces the word $w_\pi = \ell(s_1) \cdot \dots \cdot \ell(s_n)$. The probability of a word w from a state s is then $P_s(w) = \sum_{\pi: w=w_\pi} P_s(\pi)$, where $P_s(\pi) = \prod_{1 \leq i \leq n} P(s_{i-1}, s_i)$, where we set $s_0 = s$ and $\pi = s_1, \dots, s_n$. The distribution on Σ^* is lifted to a distribution on Σ^ω based on cylinder sets, in the usual manner.

In the open setting, the specification is over $2^{I \cup O}$. We are given an MDP $\mathcal{M} = \langle S, s_0, (A_s)_{s \in S}, P, \text{cost} \rangle$ and a labeling function $\ell : S \rightarrow 2^{AP}$, where $A_s = 2^O$ for every $s \in S$. Then, for a sequence of outputs $\rho \in (2^O)^*$ (which corresponds to a strategy for the MDP), the induced Markov chain defined the probability space on $(2^I)^*$, similarly to the above.

In the following we explain how to adapt the various results of the paper to the setting of non-uniform inputs.

B.1 The letter based approach - adapting Theorem 5

In the letter based approach, we still sample letters, restricting to those that keep us within the language. However, instead of sampling them uniformly, we need to take into account the MC

describing the distribution. Thus, we are given a DLW $\mathcal{A} = \langle \Sigma, Q, \delta, q_0 \rangle$ and a labeled MC $\mathcal{M} = \langle S, P, \ell \rangle$ as above. Instead of constructing $\mathcal{M}_{\mathcal{A}}$ as in Section 3, we construct $\mathcal{M}_{\mathcal{A}}$ as follows. We start with the product of \mathcal{A} and \mathcal{M} . That is, the state space is $Q \times S$, there is a transition between $\langle q, s \rangle$ and $\langle q', s' \rangle$ if $\delta(q, \sigma) = q'$ in \mathcal{A} , and the probability of the transition is determined according to \mathcal{M} . Since δ is only a partial function, we re-normalize the transition probabilities. Computing the sensing cost $lcost(\mathcal{A})$ under \mathcal{M} proceeds by finding a limiting distribution in this chain, similarly to Theorem 5.

B.2 The letter based approach - adapting Theorem 6

Computing $wcost(\mathcal{A})$ is done similarly to the case of a uniform distribution, with the difference being the construction of the generating function $C_{p,q}(z)$. Instead of having the coefficient of z^n being $\#paths(p, q, n)$, we need to account for the probabilities of the different paths. To do so, we consider again the product of \mathcal{A} and the states S of the MC describing the distribution. Consider a state $\langle q, s \rangle$. Since the transition probabilities in $Q \times S$ are assumed to be rational, we assume that these probabilities are all multiples of some common denominator p . Then, when computing the number of paths from $\langle p, s \rangle$ to $\langle q, s' \rangle$, we count each path according to its multiple of p . The rest of the computation is the same as the proof of Theorem 6.

B.3 The open setting - adapting Theorem 11

Consider a realizable DLW specification \mathcal{A} over $2^{I \cup O}$, and an MDP $\mathcal{M} = \langle S, s_0, (A_s)_{s \in S}, P, cost \rangle$ describing the input distribution. Instead of directly constructing an MDP from \mathcal{A} (using a uniform distribution), we construct the MDP from \mathcal{A} using the probabilities defined by \mathcal{M} , as follows. The state space of the MDP is $\mathcal{S} = (2^Q \times \{0, 1, \perp\}^I \times S) \cup \{\text{START}\}$. Next, Consider a state $\langle R, \ell, s \rangle$ and an action $\langle o, i \rangle$. For an available transition to state $\langle R', \ell', s' \rangle$ the transition probability is defined as the probability to read ℓ'_1 in s' . That is, $P_s(\ell'_1)$. The rest of the proof follows the same lines as that of Theorem 11. We note that the size of a memoryless strategy may now depend (polynomially) on the size of the distribution MDP.