



Good-Enough Synthesis

Shaull Almagor¹  and Orna Kupferman² 

¹ Department of Computer Science, Technion, Haifa, Israel

`shaull@cs.technion.ac.il`

² School of Computer Science and Engineering, The Hebrew University,
Jerusalem, Israel

`orna@cs.huji.ac.il`

Abstract. We introduce and study *good-enough synthesis* (GE-synthesis) – a variant of synthesis in which the system is required to satisfy a given specification ψ only when it interacts with an environments for which a satisfying interaction exists. Formally, an input sequence x is *hopeful* if there exists some output sequence y such that the induced computation $x \otimes y$ satisfies ψ , and a system GE-realizes ψ if it generates a computation that satisfies ψ on all hopeful input sequences. GE-synthesis is particularly relevant when the notion of correctness is *multi-valued* (rather than Boolean), and thus we seek systems of the highest possible quality, and when synthesizing *autonomous systems*, which interact with unexpected environments and are often only expected to do their best.

We study GE-synthesis in Boolean and multi-valued settings. In both, we suggest and solve various definitions of GE-synthesis, corresponding to different ways a designer may want to take hopefulness into account. We show that in all variants, GE-synthesis is not computationally harder than traditional synthesis, and can be implemented on top of existing tools. Our algorithms are based on careful combinations of nondeterministic and universal automata. We augment systems that GE-realize their specifications by monitors that provide satisfaction information. In the multi-valued setting, we provide both a worst-case analysis and an expectation-based one, the latter corresponding to an interaction with a stochastic environment.

1 Introduction

Synthesis is the automated construction of a system from its specification: given a specification ψ , typically by a linear temporal logic (LTL) formula over sets I and O of input and output signals, the goal is to construct a finite-state system that satisfies ψ [9, 20]. At each moment in time, the system reads an assignment, generated by the environment, to the signals in I , and responds with an assignment to the signals in O . Thus, with every input sequence, the system associates an output sequence. The system *realizes* ψ if ψ is satisfied in all the interactions of the system, with all environments [5].

S. Almagor—Supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 837327.

O. Kupferman—Supported in part by the Israel Science Foundation, grant No. 2357/19.

© The Author(s) 2020

S. K. Lahiri and C. Wang (Eds.): CAV 2020, LNCS 12225, pp. 541–563, 2020.

https://doi.org/10.1007/978-3-030-53291-8_28

In practice, the requirement to satisfy the specification in all environments is often too strong. Accordingly, it is common to add assumptions on the behavior of the environment. An assumption may be direct, say given by an LTL formula that restricts the set of possible input sequences [8], less direct, say a bound on the size of the environment [13] or other resources it uses, or conceptual, say rationality from the side of the environment, which may have its own objectives [11, 14]. We introduce and study a new type of relaxation of the requirement to satisfy the specification in all environments. The idea behind the relaxation is that if an environment is such that no system can interact with it in a way that satisfies the specification, then we cannot expect our system to succeed. In other words, the system has to satisfy the specification only when it interacts with environments in which this mission is possible. This is particularly relevant when synthesizing *autonomous systems*, which interact with unexpected environments and often replace human behavior, which is only expected to be *good enough* [28], and when the notion of correctness is multi-valued (rather than Boolean), and thus we seek *high-quality* systems.

Before we explain the relaxation formally, let us consider a simple example, and we start with the Boolean setting. Let $I = \{req\}$ and $O = \{grant\}$. Thus, the system receives requests and generates grants. Consider the specification $\psi = \text{GF}(req \wedge grant) \wedge \text{GF}(\neg req \wedge \neg grant)$. Clearly, ψ is not realizable, as an input sequence need not satisfy $\text{GF}req$ or $\text{GF}\neg req$. However, a system that always generates a grant upon (and only upon) a request, *GE-realizes* ψ , in the sense that for every input sequence, if there is some interaction with it with which ψ is satisfied, then our system generates such an interaction.

Formally, we model a system by a strategy $f : (2^I)^+ \rightarrow 2^O$, which given an input sequence $x = i_0 \cdot i_1 \cdot i_2 \cdots \in (2^I)^\omega$, generates an output sequence $f(x) = f(i_0) \cdot f(i_0 \cdot i_1) \cdot f(i_0 \cdot i_1 \cdot i_2) \cdots \in (2^O)^\omega$, inducing the computation $x \otimes f(x) = (i_0 \cup f(i_0)) \cdot (i_1 \cup f(i_0 \cdot i_1)) \cdot (i_2 \cup f(i_0 \cdot i_1 \cdot i_2)) \cdots \in (2^{I \cup O})^\omega$, obtained by “merging” x and $f(x)$. In traditional realizability, a system realizes ψ if ψ is satisfied in all environments. Formally, for all input sequences $x \in (2^I)^\omega$, the computation $x \otimes f(x)$ satisfies ψ . For our new notion, we first define when an input sequence $x \in (2^I)^\omega$ is *hopeful*, namely there is an output sequence $y \in (2^O)^\omega$ such that the computation $x \otimes y$ satisfies ψ . Then, a system *GE-realizes* ψ if ψ is satisfied in all interactions with hopeful input sequences. Formally, for all $x \in (2^I)^\omega$, if x is hopeful, then the computation $x \otimes f(x)$ satisfies ψ .

Since LTL is Boolean, synthesized systems are correct, but there is no reference to their quality. This is a crucial drawback, as designers would be willing to give up manual design only if automated-synthesis algorithms return systems of comparable quality. Addressing this challenge, researchers have developed quantitative specification formalisms. For example, in [4], the input to the synthesis problem includes also Mealy machines that grade different realizing systems. In [1], the specification formalism is the multi-valued logic $\text{LTL}[\mathcal{F}]$, which augments LTL with quality operators. The satisfaction value of an $\text{LTL}[\mathcal{F}]$ formula is a real value in $[0, 1]$, where the higher the value, the higher the quality in which the computation satisfies the specification. The quality operators in \mathcal{F} can

prioritize and weight different scenarios. The synthesis algorithm for $LTL[\mathcal{F}]$ seeks systems with a highest possible satisfaction value. One can consider either a worst-case approach, where the satisfaction value of a system is the satisfaction value of its computation with the lowest satisfaction value [1], or a stochastic approach, where it is the expected satisfaction value, given a distribution of the inputs [2].

Consider, for example, an acceleration controller of an autonomous car. Normally, the car should maintain a relatively constant speed. However, in order to optimize travel time, if a long stretch of road is visible and is identified as low-risk, the car should accelerate. Conversely, if an obstacle or some risk factor is identified, the car should decelerate. Clearly, the car cannot accelerate and decelerate at the same time. We capture this desired behavior with the following $LTL[\mathcal{F}]$ formula over the inputs $\{safe, obs\}$ and outputs $\{acc, dec\}$:

$$\psi = G(safe \rightarrow (acc \oplus_{\frac{2}{3}} Xacc)) \wedge G(obs \rightarrow (dec \oplus_{\frac{2}{3}} Xdec)) \wedge G(\neg(acc \wedge dec)).$$

Thus, in order to get satisfaction value 1, each detection of a safe stretch should be followed by an acceleration during two transactions, with a preference to the first (by the semantics of the weighted average \oplus_λ operator, the satisfaction value of $safe \rightarrow (acc \oplus_{\frac{2}{3}} Xacc)$ is 1 when $safe$ is followed by two acc s, $\frac{2}{3}$ when it is followed by one acc , and $\frac{1}{3}$ if it is followed by one acc with a delay), and each detection of an obstacle should be followed by a deceleration during two transactions, with a (higher) preference to the first. Clearly, ψ is not realizable with satisfaction value 1, as for some input sequences, namely those with simultaneous or successive occurrences of $safe$ and obs , it is impossible to respond with the desired patterns of acceleration or deceleration. Existing frameworks for synthesis cannot handle this challenge. Indeed, we do not want to add an assumption about $safe$ and obs occurring far apart. Rather, we want our autonomous car to behave in an optimal way also in problematic environments, and we want, when we evaluate the quality of a car, to take into an account the challenge posed by the environment. This is exactly what high-quality GE-synthesis does: for each input sequence, it requires the synthesized car to obtain the maximal satisfaction value that is possible for that input sequence.

We show that in the Boolean setting, GE-synthesis can be reduced to synthesis of LTL with quantification of atomic propositions [26]. Essentially, GE-synthesis of ψ amounts to synthesis of $(\exists O.\psi) \rightarrow \psi$. We show that by carefully switching between nondeterministic and universal automata, we can solve the GE-synthesis problem in doubly-exponential time, thus it is not harder than traditional synthesis. Also, our algorithm is *Safraless*, thus no determinization and parity games are needed [15, 17].

A drawback of GE-synthesis is that we do not actually know whether the specification is satisfied. We describe two ways to address this drawback. The first goes beyond providing satisfaction information and enables the designer to partition the specification into a *strong* component, which is guaranteed to be satisfied in all environments, and a *weak* component, which is guaranteed to be satisfied only in hopeful ones. The second way augments GE-realizing systems

by “satisfaction indicators”. For example, we show that when a system is lucky to interact with an environment that generates a prefix of an input sequence such that, when combined with a suitable prefix of an output sequence, the specification becomes realizable, then GE-synthesis guarantees that the system indeed responds with a suitable prefix of an output sequence. Moreover, it is easy to add to the system a monitor that detects such prefixes, thus indicating that the specification is going to be satisfied in all environments. Additional monitors we suggest detect prefixes after which the satisfaction becomes valid or unsatisfiable.

We continue to the quantitative setting. We parameterize hope by a satisfaction value $v \in [0, 1]$ and say that an input sequence $x \in (2^I)^\omega$ is v -hopeful for an LTL[\mathcal{F}] formula ψ if an interaction with it can generate a computation that satisfies ψ with value at least v . Formally, there is an output sequence $y \in (2^O)^\omega$ such that $\llbracket x \otimes y, \psi \rrbracket \geq v$, where for a computation $w \in (2^{I \cup O})^\omega$, we use $\llbracket w, \psi \rrbracket$ to denote the satisfaction value of ψ in w . As we elaborate below, while the basic idea of GE-synthesis, namely “input sequences with a potential to high quality should realize this potential” is as in the Boolean setting, there are several ways to implement this idea.

We start with a worst-case approach. There, a strategy $f : (2^I)^+ \rightarrow 2^O$ GE-realizes an LTL[\mathcal{F}] formula ψ if for all input sequences $x \in (2^I)^\omega$, if x is v -hopeful, then $\llbracket x \otimes f(x), \psi \rrbracket \geq v$. The requirement can be applied to a threshold value or to all values $v \in [0, 1]$. For example, our autonomous car controller has to achieve satisfaction value 1 in roads with no simultaneous or successive occurrences of *safe* and *obs*, and value $\frac{3}{4}$ in roads that violate the latter only with some *obs* followed by *safe*. We then argue that the situation is similar to that of *high-quality assume guarantee synthesis* [3], where richer relations between a quantitative assumption and a quantitative guarantee are of interest. In our case, the assumption is the hopefulness level of the input sequence, namely $\llbracket x, \exists O.\psi \rrbracket$, and the guarantee is the satisfaction value of the specification in the generated computation, namely $\llbracket x \otimes f(x), \psi \rrbracket$. When synthesizing, for example, a robot controller (e.g., vacuum cleaner) in a building, the doors to rooms are controlled by the environment, whereas the movement of the robot by the system. A measure of the performance of the robot has to take into an account both the number of “hopeful rooms”, namely these with an open door – a projection of this number on $[0, 1]$ serves as the assumption, and the number of room cleaned – which induces the guarantee. We assume that the desired relation between the assumption and the guarantee is given by a function $\text{comb} : [0, 1] \times [0, 1] \rightarrow [0, 1]$, which can capture implication, difference, or ratio.

We continue with an analysis of the expected performance of the system. We do so by assuming a stochastic environment, with a known distribution on the input sequences. We introduce and study two measures for high-quality GE-synthesis in a stochastic environment. In the first, termed *expected GE-synthesis*, all input sequences are sampled, yet the satisfaction value in each input sequence takes its hopefulness level into account, for example by a comb function as in the assume-guarantee setting. In the second, termed *conditional expected*

GE-synthesis, only hopeful input sequences are sampled. For both approaches, our synthesis algorithm is based on the high-quality LTL[\mathcal{F}] synthesis algorithm of [2], which is based on an analysis of deterministic automata associated with the different satisfaction values of the LTL[\mathcal{F}] specification. Here too, the complexity stays doubly exponential. In addition, we extend the synthesized systems with guarantees for satisfaction and monitors indicating satisfaction in various satisfaction levels.

2 Preliminaries

Consider two finite sets I and O of input and output signals, respectively. For two words $x = i_0 \cdot i_1 \cdot i_2 \cdots \in (2^I)^\omega$ and $y = o_0 \cdot o_1 \cdot o_2 \cdots \in (2^O)^\omega$, we define $x \otimes y$ as the word in $(2^{I \cup O})^\omega$ obtained by merging x and y . Thus, $x \otimes y = (i_0 \cup o_0) \cdot (i_1 \cup o_1) \cdot (i_2 \cup o_2) \cdots$. The definition is similar for finite x and y of the same length. For a word $w \in (2^{I \cup O})^\omega$, we use $w|_I$ to denote the projection of w on I . In particular, $(x \otimes y)|_I = x$.

A strategy is a function $f : (2^I)^+ \rightarrow 2^O$. Intuitively, f models the interaction of a system that generates in each moment in time a letter in 2^O with an environment that generates letters in 2^I . For an input sequence $x = i_0 \cdot i_1 \cdot i_2 \cdots \in (2^I)^\omega$, we use $f(x)$ to denote the output sequence $f(i_0) \cdot f(i_0 \cdot i_1) \cdot f(i_0 \cdot i_1 \cdot i_2) \cdots \in (2^O)^\omega$. Then, $x \otimes f(x) \in (2^{I \cup O})^\omega$ is the computation of f on x . Note that the environment initiates the interaction, by inputting i_0 . Of special interest are finite-state strategies, induced by finite state transducers. Formally, an I/O-transducer is $\mathcal{T} = \langle I, O, S, s_0, M, \tau \rangle$, where S is a finite set of states, $s_0 \in S$ is an initial state, $M : S \times 2^I \rightarrow S$ is a transition function, and $\tau : S \rightarrow 2^O$ is a labelling function. For $x = i_0 \cdot i_1 \cdot i_2 \cdots \in (2^I)^*$, let $M^*(x)$ be the state in S that \mathcal{T} reaches after reading x . Thus is, $M^*(\epsilon) = s_0$ and for every $j \geq 0$, we have that $M^*(i_0 \cdot i_1 \cdot i_2 \cdots i_j) = M(M^*(i_0 \cdot i_1 \cdot i_2 \cdots i_{j-1}), i_j)$. Then, \mathcal{T} induces the strategy $f_{\mathcal{T}} : (2^I)^+ \rightarrow 2^O$, where for every $x \in (2^I)^+$, we have that $f_{\mathcal{T}}(x) = \tau(M^*(x))$. We use $\mathcal{T}(x)$ and $x \otimes \mathcal{T}(x)$ to denote the output sequence and the computation of \mathcal{T} on x , respectively, and talk about \mathcal{T} realizing a specification, referring to the strategy $f_{\mathcal{T}}$.

We specify on-going behaviors of reactive systems using the linear temporal logic LTL [19]. Formulas of LTL are constructed from a set AP of atomic proposition using the usual Boolean operators and temporal operators like G (“always”), F (“eventually”), X (“next time”), and U (“until”). Each LTL formula ψ defines a language $L(\psi) = \{w : w \models \psi\} \subseteq (2^{AP})^\omega$. We also use automata on infinite words for specifying and reasoning about on-going behaviors. We use automata with different branching modes (nondeterministic, where some run has to be accepting; universal, where all runs have to be accepting; and deterministic, where there is a single run) and different acceptance conditions (Büchi, co-Büchi, and parity). We use the three letter acronyms NBW, UCW, DPW, and DFW, to refer to nondeterministic Büchi, universal co-Büchi, deterministic parity, and deterministic finite word automata, respectively. Given an LTL formula ψ over AP , one can constructs an NBW \mathcal{A}_ψ with at most $2^{O(|\psi|)}$ states

such that $L(\mathcal{A}_\psi) = L(\psi)$ [27]. Constructing an NBW for $\neg\psi$ and then dualizing it, results in a UCW for $L(\psi)$, also with at most $2^{O(|\psi|)}$ states. Determinization [23] then leads to a DPW for $L(\psi)$ with at at most $2^{2^{O(|\psi|)}}$ states and index $2^{O(|\psi|)}$. For full definitions of LTL, automata, and their relation, see [12].

Consider an LTL formula ψ over $I \cup O$. We say that ψ is *realizable* if there is a finite-state strategy $f : (2^I)^+ \rightarrow 2^O$ such that for all $x \in (2^I)^\omega$, we have that $x \otimes f(x) \models \psi$. That is, the computation of f on every input sequence satisfies ψ . We say that a word $x \in (2^I)^\omega$ is *hopeful* for ψ if there is $y \in (2^O)^\omega$ such that $x \otimes y \models \psi$. Then, we say that ψ is *good-enough realizable* (GE-realizable, for short) if there is a finite-state strategy $f : (2^I)^+ \rightarrow 2^O$ such that for every $x \in (2^I)^\omega$ that is hopeful for ψ , we have that $x \otimes f(x) \models \psi$. That is, if there is some output sequence whose combination with x satisfies ψ , then the computation of f on x satisfies ψ . The LTL GE-synthesis problem is then to decide whether a given LTL formula is GE-realizable, and if so, to return a transducer that GE-realizes it. Clearly, every realizable specification is GE-realizable – by the same transducer. We say that ψ is *universally satisfiable* if all input sequences are hopeful for ψ . It is easy to see that for universally satisfiable specifications, realizability and GE-realizability coincide. On the other hand, as demonstrated in Sect. 1, there are specifications that are not realizable and are GE-realizable.

Example 1. Let $I = \{p\}$ and $O = \{q\}$. Consider the specification $\psi = \text{GF}((Xp) \wedge q) \wedge \text{GF}((X\neg p) \wedge \neg q)$. Clearly, ψ is not realizable, as an input sequence $x \in (2^I)^\omega$ is hopeful for ψ iff $x \models \text{GF}p \wedge \text{GF}\neg p$. Since the system has to assign a value to q before it knows the value of Xp , it seems that ψ is also not GE-realizable. As we show below, however, the specification ψ is GE-realizable. Intuitively, it follows from the fact that hopeful input sequences consists of alternating p -blocks and $(\neg p)$ -blocks. Then, by outputting $\neg q$ in p -blocks and outputting q in $(\neg p)$ -blocks, the system guarantees that each last position in a $(\neg p)$ -block satisfies $q \wedge Xp$ and each last position in a p -block satisfies $(\neg q) \wedge Xp$. Formally, ψ is GE-realized by the transducer $\mathcal{T} = \langle \{p\}, \{q\}, \{s_0, s_1\}, s_0, M, \tau \rangle$, where $M(s_0, \emptyset) = M(s_1, \emptyset) = s_0$, $M(s_0, \{p\}) = M(s_1, \{p\}) = s_1$, $\tau(s_0) = \{q\}$, and $\tau(s_1) = \emptyset$. □

3 LTL Good-Enough Synthesis

Recall that a strategy $f : (2^I)^+ \rightarrow 2^O$ GE-realizes an LTL formula ψ if its computations on all hopeful input sequences satisfy ψ . Thus, for every input sequence $x \in (2^I)^\omega$, either $x \otimes y \not\models \psi$ for all $y \in (2^O)^\omega$, or $x \otimes f(x) \models \psi$. The above suggests that algorithms for solving LTL GE-synthesis involve existential and universal quantification over the behavior of output signals. The logic EQLTL extends LTL by allowing existential quantification over atomic propositions [26]. We refer here to the case the atomic propositions are the signals in $I \cup O$, and the signals in O are existentially quantified. Then, an EQLTL formula is of the form $\exists O.\psi$, and a computation $w \in (2^{I \cup O})^\omega$ satisfies $\exists O.\psi$ iff there is $y \in (2^O)^\omega$ such that $w_I \otimes y \models \psi$. Dually, AQLTL extends LTL by allowing universal quantification over atomic propositions. We consider here formulas of the form

$\forall O.\psi$, which are equivalent to $\neg\exists O.\neg\psi$. Indeed, a computation $w \in (2^{I \cup O})^\omega$ satisfies $\forall O.\psi$ iff for all $y \in (2^O)^\omega$, we have that $w|_I \otimes y \models \psi$. Note that in both the existential and universal cases, the O -component of w is ignored. Accordingly, we sometimes interpret EQLTL and AQLTL formulas with respect to input sequences $x \in (2^I)^\omega$. Also note that both EQLTL and AQLTL increase the expressive power of LTL. For example, the EQLTL formula $\exists q.q \wedge X\neg q \wedge G(q \leftrightarrow XXq) \wedge G(q \rightarrow p)$ states that p holds in all even positions of the computation, which cannot be specified in LTL [29].

Theorem 1. *The LTL GE-synthesis problem is 2EXPTIME-complete.*

Proof. We start with the upper bound. Given an LTL formula ψ over $I \cup O$, we describe an algorithm that returns a transducer \mathcal{T} that GE-realizes ψ , or declares that no such transducer exists.

It is not hard to see that \mathcal{T} GE-realizes ψ iff \mathcal{T} realizes $\varphi = \psi \vee \forall O.\neg\psi$. Indeed, an input sequence $x \in (2^I)^\omega$ is hopeful for ψ iff $x \models \exists O.\psi$, and so the specification φ requires all hopeful input sequences to satisfy ψ . A naive construction of an NBW for φ involves a universal projection of the signals in O in an automaton for $\neg\psi$, and results in an NBW that is doubly exponential. In order to circumvent the extra exponent, we construct an NBW $\mathcal{A}_{\neg\varphi}$ for $\neg\varphi$, and then dualize it to get a UCW for φ , as follows.

Let $\mathcal{A}_{\neg\psi}$ be an NBW for $L(\neg\psi)$ and $\mathcal{A}_{\exists O.\psi}$ be an NBW for $L(\exists O.\psi)$. Thus, $\mathcal{A}_{\exists O.\psi}$ is obtained from an NBW \mathcal{A}_ψ for $L(\psi)$ by existentially projecting its transitions on 2^I . In more details, if $\mathcal{A}_\psi = \langle 2^{I \cup O}, Q, Q_0, \delta, \alpha \rangle$, then $\mathcal{A}_{\exists O.\psi} = \langle 2^{I \cup O}, Q, Q_0, \delta', \alpha \rangle$, where for all $q \in Q$ and $i \cup o \in 2^{I \cup O}$, we have $\delta'(q, \sigma) = \bigcup_{o \in 2^O} \{\delta(q, (\sigma \cap I) \cup o)\}$.

Let $\mathcal{A}_{\neg\varphi}$ be an NBW for the intersection of $\mathcal{A}_{\neg\psi}$ and $\mathcal{A}_{\exists O.\psi}$. We can define $\mathcal{A}_{\neg\varphi}$ as the product of $\mathcal{A}_{\neg\psi}$ and $\mathcal{A}_{\exists O.\psi}$, possibly using the generalized Büchi acceptance condition (see Remark 1), thus its size is exponential in ψ . The language of $\mathcal{A}_{\neg\varphi}$ is then $\{w \in (2^{I \cup O})^\omega : w \not\models \psi \text{ and } w \models \exists O.\psi\}$. We then solve usual synthesis for the complementing UCW. Its language is $\{w \in (2^{I \cup O})^\omega : w \models \psi \text{ or } w \models \forall O.\neg\psi\}$, as required. By [17], the synthesis problem for UCW can be solved in EXPTIME, and we are done.

The lower bound follows from the 2EXPTIME-hardness of LTL realizability [22]. The hardness proof there constructs, given a 2EXPTIME Turing machine M , an LTL formula ψ that is realizable iff M accepts the empty tape. Since all input sequences are hopeful for ψ , realizability and GE-realizability coincide, and we are done. □

Note that working with a UCW not only handles the universal quantification for free but also has the advantage of a Safraless synthesis algorithm – no determinization and parity games are needed [15, 17]. Also note that the algorithm we suggest in the proof of Theorem 1 can be generalized to handle specifications that are arbitrary positive Boolean combinations of EQLTL formulas.

Remark 1 [Products and optimizations]. Throughout the paper, we construct products of automata whose state space is $2^{cl(\psi)}$, and states correspond

to maximal consistent subsets of $cl(\psi)$, possibly in the scope of an existential quantifier of O . Accordingly, the product can be minimized to include only consistent pairs. Also, since traditional-synthesis algorithms, in particular the Safraless algorithms we use, can handle automata with *generalized* Büchi and co-Büchi acceptance condition, we need only one copy of the product. \square

Remark 2 [Determinacy of the GE-synthesis game]. Determinacy of games implies that in traditional synthesis, a specification ψ is not I/O -realizable iff $\neg\psi$ is O/I -realizable. This is useful, for example when we want to synthesize a transducer of a bounded size and proceed simultaneously, aiming to synthesize either a system transducer that realizes ψ or an environment transducer that realizes $\neg\psi$ [17]. For GE-synthesis, simple dualization does not hold, but we do have determinacy in the sense that $(\exists O.\psi) \rightarrow \psi$ is not I/O -realizable iff $(\exists O.\psi) \wedge \neg\psi$ is O/I -realizable. Accordingly, ψ is not GE-realizable iff the environment has a strategy that generates, for each output sequence $y \in (2^O)^\omega$, a helpful input sequence $x \in (2^I)^\omega$ such that $x \otimes y \models \neg\psi$. In the full version, we formalize and study this duality further. \square

4 Guarantees in Good-Enough Synthesis

A drawback of GE-synthesis is that we do not actually know whether the specification is satisfied. In this section we describe two ways to address this drawback. The first way goes beyond providing satisfaction information and enables the designer to partition the specification into to a *strong* component, which should be satisfied in all environments, and a *weak* component, which should be satisfied only in hopeful ones. The second way augments GE-realizing transducers by flags, raised to indicate the status of the satisfaction.

4.1 GE-Synthesis with a Guarantee

Recall that GE-realizability is suitable especially in settings where we design a system that has to do its best in all environments. GE-synthesis with a guarantee is suitable in settings where we want to make sure that some components of the specification are satisfied in all environment. Accordingly, a specification is an LTL formula $\psi = \psi_{strong} \wedge \psi_{weak}$. When we GE-synthesize ψ_{weak} with guarantee ψ_{strong} , we seek a transducer \mathcal{T} that realizes ψ_{strong} and GE-realizes ψ_{weak} . Thus, for all input sequences $x \in (2^I)^\omega$, we have that $x \otimes \mathcal{T}(x) \models \psi_{strong}$, and if x is hopeful for ψ_{weak} , then $x \otimes \mathcal{T}(x) \models \psi_{strong}$.

Theorem 2. *The LTL GE-synthesis with guarantee problem is 2EXPTIME-complete.*

Proof. Consider an LTL formula $\psi = \psi_{strong} \wedge \psi_{weak}$ over $I \cup O$. It is not hard to see that a transducer \mathcal{T} GE-realizes ψ_{weak} with guarantee ψ_{strong} iff \mathcal{T} realizes $\varphi = \psi_{strong} \wedge ((\exists O.\psi_{weak}) \rightarrow \psi_{weak})$. We can then construct a UCW \mathcal{A}_φ for $L(\varphi)$ by dualizing an NBW for its negation $\neg\psi_{strong} \vee ((\exists O.\psi_{weak}) \wedge \neg\psi_{weak})$, which

can be constructed using techniques similar to those in the proof of Theorem 1. We then proceed with standard synthesis for \mathcal{A}_φ . Note that the approach is Safrales. Taking an empty (that is, **True**) guarantee, a lower bound follows from the 2EXPTIME-hardness of LTL GE-synthesis. \square

4.2 Flags by a GE-Realizing Transducer

For a language $L \subseteq (2^{I \cup O})^\omega$ and a finite word $w \in (2^{I \cup O})^*$, let $L^w = \{w' \in (2^{I \cup O})^\omega : w \cdot w' \in L\}$. That is, L^w is the language of suffixes of words in L that have w as a prefix. We say that a word $w \in (2^{I \cup O})^*$ is *green for* L if L^w is realizable. Then, a word $x \in (2^I)^*$ is *green for* L if there is $y \in (2^O)^*$ such that $x \otimes y$ is green for L . When a system is lucky to interact with an environment that generates a green input sequence, we want the system to react in a way that generates a green prefix, and then realizes the specification. Formally, we say that a strategy $f : (2^I)^+ \rightarrow 2^O$ *green realizes* L if for every $x \in (2^I)^+$, if x is green for L , then $x \otimes f(x)$ is green for L .^{1,2} We say that a word $w \in (2^{I \cup O})^*$ is *light green for* L if L^w is universally satisfiable, thus all input sequences are hopeful for L^w . A word $x \in (2^I)^*$ is *light green for* L if there is $y \in (2^O)^*$ such that $x \otimes y$ is light green for L . It is not hard to see that for GE-realizable languages, green and light green coincide. Indeed, if L is universally satisfiable and GE-realizable, then L is realizable.

Theorem 3. *GE-realizability is strictly stronger than green realizability.*

Proof. We first prove that every strategy $f : (2^I)^+ \rightarrow 2^O$ that GE-realizes a specification ψ also green realizes ψ . Consider $x \in (2^I)^+$ that is green for ψ . By definition, there is $y \in (2^O)^+$ such that $L^{x \otimes y}$ is realizable. Then, for every $x' \in (2^I)^\omega$, there is $y' \in (2^O)^\omega$ such that $x' \otimes y'$ in $L^{x \otimes y}$. Hence, for every $x' \in (2^I)^\omega$, we have that $x \cdot x'$ is hopeful. Therefore, as f GE-realizes ψ , we have that $(x \cdot x') \otimes f(x \cdot x') \models \psi$. Thus, $x \otimes f(x)$ is green, and so f green realizes ψ .

We continue and describe a specification that is green realizable and not GE-realizable. Let $I = \{p\}$ and $O = \{q\}$. Consider the specification $\psi = \mathbb{G}((Xp) \leftrightarrow q)$. Clearly, ψ is not realizable, as the system has to commit a value for q before a value for Xp is known. Likewise, no word $w \in (2^{I \cup O})^*$ is green for ψ , and so no finite input sequence $x \in (2^I)^*$ is green for ψ . Hence, every strategy (vacuously) green realizes ψ . On the other hand, for every input sequences $x \in (2^I)^\omega$ there is an output sequence $y \in (2^O)^\omega$ such that $x \otimes y \models \psi$. Thus, all input sequences are hopeful for ψ . Thus, synthesis and GE-synthesis coincide for ψ , which is not GE-realizable. \square

Theorem 3 brings with it two good news. The first is that a GE-realizing transducer has the desired property of being also green realizing. The second has

¹ Note that while the definition of green realization does not refer to ϵ directly, we have that ϵ is green iff L is realizable, in which case all $x \in (2^I)^*$ are green.

² While synthesis corresponds to finding a winning strategy for the system, green synthesis can be viewed as a subgame-perfect best-response strategy, where the system does its best in every subgame, even if it loses the overall game.

to do with our goal of providing the user with information about the satisfaction status, in particular raising a green flag whenever a green prefix is detected. By Theorem 3, such a flag indicates that the computation generated by our GE-realizing transducer satisfies the specification. A naive way to detect green prefixes for a specification ψ is to solve the synthesis problem for ψ by solving a game on top of a DPW \mathcal{D}_ψ for ψ . The winning positions in the game are states in \mathcal{D}_ψ . By defining them as accepting states, we can obtain from \mathcal{D}_ψ a DFW for green prefixes. Then, we run this DFW in parallel with the GE-realizing transducer, and raise the green flag whenever a green prefix is detected. This, however, requires a generation of \mathcal{D}_ψ and a solution of parity games. Below we describe a much simpler way, which makes use of the fact that our transducer GE-realizes the specification.

Recall that if L is universally satisfiable and GE-realizable, then L is realizable. Accordingly, given a transducer \mathcal{T} that GE-realizes ψ , we can augment it with green flags by running in parallel a DFW that detects light-green prefixes. As we argue below, constructing such a DFW only requires an application of the subset construction on top of an NBW for the existential projection of ψ on 2^I .

Lemma 1. *Given an LTL formula ψ over $I \cup O$, we can construct a DFA \mathcal{S} of size $2^{2^{O(|\psi|)}}$ such that $L(\mathcal{S}) = \{x \in (2^I)^* : x \text{ is light green for } L(\psi)\}$.*

Proof. Let $\mathcal{A}_\psi = \langle 2^{I \cup O}, Q, \delta, Q_0, \alpha \rangle$ be an NBW for $L(\psi)$, and let $\mathcal{B}_\psi = \langle 2^I, Q, \delta', Q_0, \alpha \rangle$ be its existential projection on 2^I . Thus, for every $q \in Q$ and $i \in 2^I$, we have $\delta'(q, i) = \bigcup_{o \in 2^O} \delta(q, i \cup o)$. We define the DFW $\mathcal{S} = \langle 2^I, 2^Q, M, \{Q_0\}, F \rangle$, where M follows the subset construction of \mathcal{B}_ψ : for every $S \in 2^Q$ and $i \in 2^I$, we have $M(S, i) = \bigcup_{s \in S} \delta'(s, i)$. Then, $F = \{S \in 2^Q : L(\mathcal{B}_\psi^S) = (2^I)^\omega\}$. Observe that \mathcal{S} rejects $x \in (2^I)^*$ iff there is $x' \in (2^I)^\omega$ such that for all $y \in (2^O)^*$ and $y' \in (2^O)^\omega$, no state in $\delta(Q_0, x \otimes y)$ accepts $x' \otimes y'$. Thus, \mathcal{S} rejects x iff x is not light green, and accepts it otherwise. Note that the definition of F involves universality checking, possibly via complementation, yet no determinization is required, and the size of \mathcal{S} is $2^{2^{O(|\psi|)}}$. □

Note that once we reach an accepting state in \mathcal{S} , we can make it an accepting loop. Indeed, once a green prefix is detected, then all prefixes that extend it are green. Accordingly, once the green flag is raised, it stays up. Also note that if an input sequence is not hopeful for ψ , then none of its prefixes is light green for ψ . The converse, however, is not true: an input sequence may be hopeful and still have no light green prefixes. For example, taking $I = \{p\}$, the input sequence $\{p\}^\omega$ is hopeful for $\mathcal{G}p$, yet none of its prefixes is green light, as it can be extended to an input sequence with $\neg p$.

Green flags provide information about satisfaction. Two additional flags of interest are related to safety and co-safety properties:

- A word $w \in (2^{I \cup O})^*$ is *red* for L if $L^w = \emptyset$. A word $x \in (2^I)^*$ is *red* for L if for all $y \in (2^O)^*$, we have that $x \otimes y$ is red for L . Thus, when the environment generates x , then no matter how the system responds, L is not satisfied.

- a word $w \in (2^{I \cup O})^*$ is *blue for L* when $L^w = (2^{I \cup O})^\omega$, and then define a word $x \in (2^I)^*$ as *blue for L* if there is $y \in (2^O)^*$ such that $x \otimes y$ is blue for L . Thus, when the environment generates x , the system can respond in a way that guarantees satisfaction no matter how the interaction continues.

A monitor that detects red and blue prefixes for L can be added to a transducer that GE-realizes L . As has been the case with the monitor for green prefixes, its construction is based on applying the subset construction on an NBW for L [16]. Also, once a red or blue flag is raised, it stays up. In a way analogous to green realizability, we seek a transducer that GE-realizes the specification and generates a red prefix only if all interactions generate a red prefix, and generates a blue prefix whenever this is possible. In the full version, we show that while GE-realization implies *red realization*, it may conflict with *blue realization*.

5 High-Quality Good-Enough Synthesis

GE-synthesis is of special interest when the satisfaction value of the specification is multi-valued, and we want to synthesize high-quality systems. We start by defining the multi-valued logic $LTL[\mathcal{F}]$, which is our multi-valued specification formalism. We then study $LTL[\mathcal{F}]$ GE-synthesis, first in a worst-case approach, where the satisfaction value of a transducer is the satisfaction value of its computation with the lowest satisfaction value, and then in a stochastic approach, where it is the expected satisfaction value, given a distribution of the inputs.

5.1 The Logic $LTL[\mathcal{F}]$

Let AP be a set of Boolean atomic propositions and let $\mathcal{F} \subseteq \{f : [0, 1]^k \rightarrow [0, 1] : k \in \mathbb{N}\}$ be a set of *quality operators*. An $LTL[\mathcal{F}]$ formula is one of the following:

- **True**, **False**, or p , for $p \in AP$.
- $f(\psi_1, \dots, \psi_k)$, $\mathbf{X}\psi_1$, or $\psi_1 \mathbf{U} \psi_2$, for $LTL[\mathcal{F}]$ formulas ψ_1, \dots, ψ_k and a function $f \in \mathcal{F}$.

The semantics of $LTL[\mathcal{F}]$ formulas is defined with respect to infinite computations over AP . For a computation $w = w_0, w_1, \dots \in (2^{AP})^\omega$ and position $j \geq 0$, we use w^j to denote the suffix w_j, w_{j+1}, \dots . The semantics maps a computation w and an $LTL[\mathcal{F}]$ formula ψ to the *satisfaction value* of ψ in w , denoted $\llbracket w, \psi \rrbracket$. The satisfaction value is in $[0, 1]$ and is defined inductively as follows.

- $\llbracket w, \mathbf{True} \rrbracket = 1$ and $\llbracket w, \mathbf{False} \rrbracket = 0$.
- For $p \in AP$, we have that $\llbracket w, p \rrbracket = 1$ if $p \in w_0$, and $\llbracket w, p \rrbracket = 0$ if $p \notin w_0$.
- $\llbracket w, f(\psi_1, \dots, \psi_k) \rrbracket = f(\llbracket w, \psi_1 \rrbracket, \dots, \llbracket w, \psi_k \rrbracket)$.
- $\llbracket w, \mathbf{X}\psi_1 \rrbracket = \llbracket w^1, \psi_1 \rrbracket$.
- $\llbracket w, \psi_1 \mathbf{U} \psi_2 \rrbracket = \max_{i \geq 0} \{ \min \{ \llbracket w^i, \psi_2 \rrbracket, \min_{0 \leq j < i} \llbracket w^j, \psi_1 \rrbracket \} \}$.

The logic LTL can be viewed as $LTL[\mathcal{F}]$ for \mathcal{F} that models the usual Boolean operators. In particular, the only possible satisfaction values are 0 and 1. We abbreviate common functions as described below. Let $x, y, \lambda \in [0, 1]$. Then,

- $\neg x = 1 - x$ • $x \vee y = \max\{x, y\}$ • $x \wedge y = \min\{x, y\}$
- $x \rightarrow y = \max\{1 - x, y\}$ • $\nabla_\lambda x = \lambda \cdot x$ • $x \oplus_\lambda y = \lambda \cdot x + (1 - \lambda) \cdot y$

The realizability problem for $LTL[\mathcal{F}]$ is an optimization problem: For an $LTL[\mathcal{F}]$ specification ψ and a transducer \mathcal{T} , we define the satisfaction value of ψ in \mathcal{T} , denoted $\llbracket \mathcal{T}, \psi \rrbracket$, by $\min\{\llbracket x \otimes \mathcal{T}(x), \psi \rrbracket : x \in (2^I)^\omega\}$, namely the satisfaction value of ψ in the worst-case. Then, the synthesis problem is to find, given ψ , a transducer that maximizes its satisfaction value. Moving to a decision problem, given ψ and a threshold value $v \in [0, 1]$, we say that ψ is *v-realizable* if there exists a transducer \mathcal{T} such that $\llbracket \mathcal{T}, \psi \rrbracket \geq v$, and the synthesis problem is to find, given ψ and v , a transducer \mathcal{T} that *v-realizes* ψ .

For an $LTL[\mathcal{F}]$ formula ψ , let $V(\psi)$ be the set of possible satisfaction values of ψ in arbitrary computations. Thus, $V(\psi) = \{\llbracket w, \psi \rrbracket : w \in (2^{AP})^\omega\}$.

Theorem 4 [1]. *Consider an $LTL[\mathcal{F}]$ formula ψ .*

- $|V(\psi)| \leq 2^{|\psi|}$.
- For every predicate $P \subseteq [0, 1]$, there exists an NBW \mathcal{A}_ψ^P such that $L(\mathcal{A}_\psi^P) = \{w : \llbracket w, \psi \rrbracket \in P\}$. Furthermore, \mathcal{A}_ψ^P has at most $2^{O(|\psi|^2)}$ states [1].

As with LTL, we define the existential and universal extensions $EQLTL[\mathcal{F}]$ and $AQLTL[\mathcal{F}]$ of $LTL[\mathcal{F}]$. Here too, we consider the case $AP = I \cup O$, with the signals in O being quantified. Then, $\llbracket w, \exists O.\psi \rrbracket = \max_{y \in (2^O)^\omega} \{\llbracket w|_I \otimes y, \psi \rrbracket\}$ and $\llbracket w, \forall O.\psi \rrbracket = \min_{y \in (2^O)^\omega} \{\llbracket w|_I \otimes y, \psi \rrbracket\}$.

Remark 3 [**On the semantics of $EQLTL[\mathcal{F}]$**]. It is tempting to interpret an expression like $\llbracket w, \exists O.\psi \rrbracket \leq v$ as “there exists an output sequence y such that $\llbracket w|_I \otimes y, \psi \rrbracket \leq v$ ”. By the semantics of $\exists O.\psi$, however, $\llbracket w, \exists O.\psi \rrbracket \leq v$ actually means that $\max_{y \in (2^O)^\omega} \llbracket w|_I \otimes y, \psi \rrbracket \leq v$. Thus, the correct interpretation is “for all output sequences y , we have that $\llbracket w|_I \otimes y, \psi \rrbracket \leq v$ ”. □

5.2 $LTL[\mathcal{F}]$ GE-Synthesis

For a value $v \in [0, 1]$, we say that x is *v-hopeful for ψ* if there is $y \in (2^O)^\omega$ such that $\llbracket x \otimes y, \psi \rrbracket \geq v$. We study two variants of $LTL[\mathcal{F}]$ GE-synthesis:

- In $LTL[\mathcal{F}]$ *GE-synthesis with a threshold*, the input is an $LTL[\mathcal{F}]$ formula ψ and a value $v \in [0, 1]$, and the goal is to generate a transducer whose computation on every input sequence that is *v-hopeful* has satisfaction value at least v . Formally, a function $f : (2^I)^+ \rightarrow 2^O$ GE-realizes ψ with threshold v if for every $x \in (2^I)^\omega$, if x is *v-hopeful*, then $\llbracket x \otimes f(x), \psi \rrbracket \geq v$.
- In $LTL[\mathcal{F}]$ *GE-synthesis*, the input is an $LTL[\mathcal{F}]$ formula ψ , and the goal is to generate a transducer whose computation on every input sequence has the highest possible satisfaction value for this input sequence. Formally, a function $f : (2^I)^+ \rightarrow 2^O$ GE-realizes ψ if for every $x \in (2^I)^\omega$ and value $v \in [0, 1]$, if x is *v-hopeful*, then $\llbracket x \otimes f(x), \psi \rrbracket \geq v$.

In the Boolean case, the two variants coincide, taking $v = 1$. Indeed, then, for every $x \in (2^I)^\omega$, if x is hopeful, then $x \otimes f(x)$ has to satisfy ψ . We note that GE-realization with a threshold is not monotone, in the sense that decreasing the threshold need not lead to GE-realization. Indeed, the lower is the threshold v , the more input sequences are v -helpful (see Example 2). Accordingly, we do not search for a maximal threshold, and rather may ask about a desired threshold or about GE-synthesis without a threshold.

Solving the GE-synthesis problem, a naive combination of the automata construction of Theorem 4 with the projection technique of Theorem 1, corresponds to an erroneous semantics of EQLTL[\mathcal{F}], as noted in Remark 3. Before describing our construction, it is helpful to state the correct (perhaps less intuitive) interpretation of existential and universal quantification in the quantitative setting:

Lemma 2. *For every LTL[\mathcal{F}] formula ψ and an input sequence $x \in (2^I)^\omega$, we have that $\llbracket x, \exists O.\psi \rrbracket = 1 - \llbracket x, \forall O.\neg\psi \rrbracket$. Accordingly, for every value $v \in [0, 1]$, we have that $\llbracket x, \exists O.\psi \rrbracket < v$ iff $\llbracket x, \forall O.\neg\psi \rrbracket > 1 - v$.*

Proof. By definition, $\llbracket x, \exists O.\psi \rrbracket = \max_{y \in (2^O)^\omega} \llbracket x \otimes y, \psi \rrbracket = 1 - \min_{y \in (2^O)^\omega} 1 - \llbracket x \otimes y, \psi \rrbracket = 1 - \min_{y \in (2^O)^\omega} \llbracket x \otimes y, \neg\psi \rrbracket = 1 - \llbracket x, \forall O.\neg\psi \rrbracket$. Then, $\llbracket x, \exists O.\psi \rrbracket < v$ iff $1 - \llbracket x, \forall O.\neg\psi \rrbracket > 1 - v$ iff $\llbracket x, \forall O.\neg\psi \rrbracket > 1 - v$. \square

Consider an LTL[\mathcal{F}] formula ψ , a value $v \in [0, 1]$, and an input sequence $x \in (2^I)^\omega$. Recall that x is v -hopeful for ψ if there is $y \in (2^O)^\omega$ such that $\llbracket x \otimes y, \psi \rrbracket \geq v$. Equivalently, $\llbracket x, \exists O.\psi \rrbracket \geq v$. Indeed, $\llbracket x, \exists O.\psi \rrbracket = \max_{y \in (2^O)^\omega} \llbracket x \otimes y, \psi \rrbracket$, which is greater or equal to v iff there is $y \in (2^O)^\omega$ such that $\llbracket x \otimes y, \psi \rrbracket \geq v$. Hence, x is not v -hopeful for ψ if $\llbracket x, \exists O.\psi \rrbracket < v$. Equivalently, by Lemma 2, $\llbracket x, \forall O.\neg\psi \rrbracket > 1 - v$. Accordingly, for a strategy $f : (2^I)^+ \rightarrow 2^O$, an input sequence $x \in (2^I)^\omega$, and a value $v \in [0, 1]$, we say that f is v -good for x with respect to ψ , if $\llbracket x \otimes f(x), \psi \rrbracket \geq v$ or $\llbracket x, \forall O.\neg\psi \rrbracket > 1 - v$.

Example 2. Let $I = \{p\}$ and $O = \{q\}$. Consider the LTL[\mathcal{F}] formula $\psi = (\nabla_{\frac{1}{4}}p \vee \nabla_{\frac{1}{2}}q)$. Checking for which values v a strategy f is v -good for x with respect to ψ , we examine whether $\llbracket x \otimes f(x), \nabla_{\frac{1}{4}}p \vee \nabla_{\frac{1}{2}}q \rrbracket \geq v$ or $\llbracket x, \forall q.\neg(\nabla_{\frac{1}{4}}p \vee \nabla_{\frac{1}{2}}q) \rrbracket > 1 - v$. Since ψ refers only to the first position in the computation, it is enough to examine x_0 and $f(x_0)$. For example, if $x_0 = \emptyset$ and $f(x_0) = \emptyset$, then $\llbracket x \otimes f(x), \nabla_{\frac{1}{4}}p \vee \nabla_{\frac{1}{2}}q \rrbracket = 0$, $\llbracket x, \exists q.\nabla_{\frac{1}{4}}p \vee \nabla_{\frac{1}{2}}q \rrbracket = \max\{0, \frac{1}{2}\} = \frac{1}{2}$, and $\llbracket x, \forall q.\neg(\nabla_{\frac{1}{4}}p \vee \nabla_{\frac{1}{2}}q) \rrbracket = \min\{1, 1 - \frac{1}{2}\} = \frac{1}{2}$. Hence, f is v -good for x with respect to ψ if $v = 0$ or $v > \frac{1}{2}$, thus $v \in \{0\} \cup (\frac{1}{2}, 1]$. Similarly, we have the following.

- If $x_0 = \emptyset$ and $f(x_0) = \{q\}$ then f is v -good for x when $v \in [0, 1]$.
- If $x_0 = \{p\}$ and $f(x_0) = \emptyset$ then f is v -good for x when $v \in [0, \frac{1}{4}] \cup (\frac{1}{2}, 1]$.
- If $x_0 = \{p\}$ and $f(x_0) = \{q\}$ then f is v -good for x when $v \in [0, 1]$.

Theorem 5. *The LTL[\mathcal{F}] GE-synthesis with threshold problem is 2EXPTIME-complete.*

Proof. We show we can adjust the upper bound described in the proof of Theorem 1 to the multi-valued setting. Given an $LTL[\mathcal{F}]$ formula ψ over $I \cup O$ and a threshold $v \in [0, 1]$, we describe an algorithm that returns a transducer \mathcal{T} that GE-realizes ψ with threshold v , or declares that no such transducer exists.

By definition, we have that \mathcal{T} GE-realizes ψ with threshold v if for every input sequence x , we have that $f_{\mathcal{T}}$ is v -good for x with respect to ψ . Thus, $\llbracket x \otimes f_{\mathcal{T}}(x), \psi \rrbracket \geq v$ or $\llbracket x, \forall O. \neg \psi \rrbracket > 1 - v$. We construct a UCW whose language is $\{w \in (2^{I \cup O})^\omega : \llbracket w, \psi \rrbracket \geq v \text{ or } \llbracket w, \forall O. \neg \psi \rrbracket > 1 - v\}$.

Let $\mathcal{A}_\psi^{<v}$ be an NBW for $\{w : \llbracket w, \psi \rrbracket < v\}$ and $\mathcal{A}_{\exists O. \psi}^{\geq v}$ be an NBW for $\{w : \llbracket w, \exists O. \psi \rrbracket \geq v\}$. Thus, $\mathcal{A}_{\exists O. \psi}^{\geq v}$ is obtained from an NBW $\mathcal{A}_\psi^{\geq v}$ for $\{w : \llbracket w, \psi \rrbracket \geq v\}$ by existentially projecting its transitions on 2^I . By Theorem 4, both $\mathcal{A}_\psi^{<v}$ and $\mathcal{A}_{\exists O. \psi}^{\geq v}$ are of size exponential in ψ .

Let \mathcal{B}_ψ^v be an NBW for the intersection of $\mathcal{A}_\psi^{<v}$ and $\mathcal{A}_{\exists O. \psi}^{\geq v}$. The language of \mathcal{B}_ψ^v is then $\{w \in (2^{I \cup O})^\omega : \llbracket w, \psi \rrbracket < v \text{ and } \llbracket w, \exists O. \psi \rrbracket \geq v\}$. We then solve usual synthesis for the complementing UCW, whose language is $\{w \in (2^{I \cup O})^\omega : \llbracket w, \psi \rrbracket \geq v \text{ or } \llbracket w, \forall O. \neg \psi \rrbracket > 1 - v\}$, as required. By [17], the synthesis problem for UCW can be solved in EXPTIME.

The lower bound follows from the 2EXPTIME-hardness of LTL GE-realizability. \square

Theorem 6. *The $LTL[\mathcal{F}]$ GE-synthesis problem is 2EXPTIME-complete.*

Proof. We start with the upper bound. Given an $LTL[\mathcal{F}]$ specification ψ over $I \cup O$, we describe an algorithm that returns a transducer \mathcal{T} that GE-realizes ψ or declares that no such transducer exists.

As discussed above, a transducer \mathcal{T} GE-realizes ψ iff for every input sequence $x \in (2^I)^\omega$ and value $v \in [0, 1]$, we have that $f_{\mathcal{T}}$ is v -good for x with respect to ψ . Accordingly, we construct a UCW whose language is $\bigcap_{v \in V(\psi)} \{w \in (2^{I \cup O})^\omega : \llbracket w, \psi \rrbracket \geq v \text{ or } \llbracket w, \forall O. \neg \psi \rrbracket > 1 - v\}$.

For $v \in V(\psi)$, let \mathcal{B}_ψ^v be an NBW for $\{w : \llbracket w, \neg \psi \rrbracket \geq v \text{ and } \llbracket w, \exists O. \psi \rrbracket \geq v\}$, as constructed in the proof of Theorem 5, and let \mathcal{B} be the union of \mathcal{B}_ψ^v for all $v \in V(\psi)$. By Theorem 4, the size of $V(\psi)$ is exponential in ψ , and thus so is the size of \mathcal{B} . We then solve usual synthesis for the complementing UCW, whose language is as required. By [17], the synthesis problem for UCW can be solved in EXPTIME. The lower bound follows from the 2EXPTIME-hardness of LTL GE-realizability. \square

Remark 4 [Tuning hope down]. The quantitative setting allows the designer to tune down “satisfaction by hopelessness”: rather than synthesizing $\psi \vee \forall O. \neg \psi$, we can have a factor λ and synthesize $\psi \vee \nabla_\lambda \forall O. \neg \psi$. In Sect. 5.3 below we study additional ways to refer to hopefulness levels.

5.3 $LTL[\mathcal{F}]$ Assume-Guarantee GE-Synthesis

In Sect. 5.2, we seek a transducer \mathcal{T} such that for a given or for all values $v \in [0, 1]$ and input sequences $x \in (2^I)^\omega$, if $\llbracket x, \exists O. \psi \rrbracket \geq v$ then $\llbracket x \otimes \mathcal{T}(x), \psi \rrbracket \geq v$. In this

section we measure the quality of a transducer \mathcal{T} by analyzing richer relations between $\llbracket x, \exists O.\psi \rrbracket$ and $\llbracket x \otimes \mathcal{T}(x), \psi \rrbracket$. The setting has the flavor of quantitative assume-guarantee synthesis [3]. There, the specification consists of a multi-valued assumption A , which in our case is $\exists O.\psi$, and a multi-valued guarantee G , which in our case is ψ .

There are different ways to analyze the relation between $\llbracket x, \exists O.\psi \rrbracket$ and $\llbracket x \otimes \mathcal{T}(x), \psi \rrbracket$. To this end, we assume that we are given a function $\text{comb} : [0, 1] \times [0, 1] \rightarrow [0, 1]$ that given the satisfaction values of $\exists O.\psi$ and of ψ , outputs a combined satisfaction value. We assume that comb is decreasing in the first component and increasing in the second component. This corresponds to the intuition that a lower satisfaction value of $\exists O.\psi$ and a higher satisfaction value of ψ both yield a higher overall score. Also, since $\llbracket x, \exists O.\psi \rrbracket \geq \llbracket x \otimes \mathcal{T}(x), \psi \rrbracket$ for all $x \in (2^I)^\omega$, we assume that the first component is greater than or equal to the second. Finally, we require comb to be efficiently computed. Some natural comb functions include:

- The quantitative implication function: $\text{comb}(A, G) = \max\{1 - A, G\}$. This captures the quantitative notion of the implication $(\exists O.\psi) \rightarrow \psi$.
- The (negated) difference function: $\text{comb}(A, G) = 1 - (A - G)$. This captures how far the satisfaction value for the given computation is from the best satisfaction value. Since $A \geq G$, the range of the function is indeed $[0, 1]$.
- The ratio function, given by some normalization to $[0, 1]$ of the function $\text{comb}(A, G) = \frac{G}{A}$, which captures the “relative success” with respect to the best possible satisfaction value.

The choice of an appropriate comb function depends on the setting. Implication is in order when harsh environments may outweigh the actual performance of the system. For example, if our specification measures the uptime of a server in a cluster, then environments that cause very frequent power failures render the server unusable, as the overhead of reconnecting it outweighs its usefulness. In such a case, being shut down is better than continuously trying to reconnect, and so we give a higher satisfaction value for the server being down, which depends only on the environment. Then, as demonstrated with the cleaning robot in Sect. 1, the difference and ratio functions are fairly natural when measuring “realization of potential”. We now describe a more detailed example when these measures are in order.

Example 3. Consider a controller for an elevator in an n -floor building. The environment sends to the controller requests, by means of a truth assignment to $I = \{1, \dots, n\}$, indicating the subset of floors in which the elevator is requested. Then, the controller assigns values to $O = \{up, down\}$, directing the elevator to go up, go down, or stay. The satisfaction value of the specification ψ reflects the waiting time of the request with the slowest response: it is 0 when this time is more than $2n$, and is 1 when the slowest request is granted immediately. Sure enough, there is no controller that attains satisfaction value 1 on all input sequences, and so ψ is not realizable with satisfaction value 1. Also, adding assumptions about the behavior of the environment is not of much interest. Using

AG GE-realizability, we can synthesize a controller that behaves in an optimal way. For example, using the difference function, we measure the performance of the controller on an input sequence $x \in (2^I)^\omega$ with respect to the best possible performance on x . Note that such a best performance needs a look-ahead on requests yet to come, which is indeed the satisfaction value of $\exists O.\psi$ in x . Thus, the assumption $\llbracket x, \exists O.\psi \rrbracket$ actually gives us the performance of a good-enough *off-line* controller. Accordingly, using the ratio function, we can synthesize a system with the best *competitive ratio* for an on-line interaction [7]. \square

Given an LTL $[\mathcal{F}]$ formula ψ and a function comb , we define the *GE-AG-realization value* of ψ in a transducer \mathcal{T} by $\min\{\text{comb}(\llbracket x, \exists O.\psi \rrbracket, \llbracket x \otimes \mathcal{T}(x), \psi \rrbracket) : x \in (2^I)^\omega\}$. Then, our goal in *AG GE-realizability* is to find, given an LTL $[\mathcal{F}]$ formula ψ and a function comb , the maximal value $v \in [0, 1]$ such that there exists a transducer \mathcal{T} whose AG GE-realization value of ψ is v . The *AG GE-synthesis* problem is then to find such a transducer.

We start by solving the decision version of AG GE-realizability.

Theorem 7. *The problem of deciding, given an LTL $[\mathcal{F}]$ formula ψ , a function comb , and a threshold $v \in [0, 1]$, whether there exists a transducer \mathcal{T} whose AG GE-realization value of ψ is v , is 2EXPTIME-complete.*

Proof. Recall that $V(\psi)$ is the set of possible satisfaction values of ψ (and hence of $\exists O.\psi$), and that by Theorem 4, we have that $|V(\psi)| \leq 2^{|\psi|}$. Let $G_v = \{\langle v_1, v_2 \rangle \in V(\psi) \times V(\psi) : \text{comb}(v_1, v_2) \geq v\}$. Intuitively, G is the set of satisfaction-value pairs $\langle \llbracket w, \exists O.\psi \rrbracket, \llbracket w, \psi \rrbracket \rangle$ that are allowed to be generated by a transducer whose AG GE-realization value of ψ is at least v . By definition, AG GE-realization of ψ with value v coincides with realization of the language $L_v = \{w \in (2^{I \cup O})^\omega : \text{comb}(\llbracket w, \exists O.\psi \rrbracket, \llbracket w, \psi \rrbracket) \geq v\}$. By the monotonicity assumption on comb , for every $\langle v_1, v_2 \rangle \in G_v$, we have that $\langle v'_1, v'_2 \rangle \in G$ for every $v'_1 \leq v_1$ and $v'_2 \geq v_2$. Hence, we can write $L_v = \bigcup_{\langle v_1, v_2 \rangle \in G_v} \{w \in (2^{I \cup O})^\omega : \llbracket w, \exists O.\psi \rrbracket \leq v_1 \text{ and } \llbracket w, \psi \rrbracket \geq v_2\}$, and proceed to construct an NBW for L_v by taking the union of NBWs \mathcal{A}_{v_1, v_2} for all $\langle v_1, v_2 \rangle \in G_v$, each of which is the product of NBWs $\mathcal{A}_{\exists O.\psi}^{\leq v_1}$ and $\mathcal{A}_\psi^{\geq v_2}$, as in the proof of Theorem 5.

Aiming to proceed Safralessly, we can also construct a UCW for L_v , as follows. First, note that by the monotonicity of comb , for every $\langle v_1, v_2 \rangle \in V(\psi) \times V(\psi)$ we have that $\langle v_1, v_2 \rangle \in G_v$ iff for every $\langle u_1, u_2 \rangle \in V(\psi) \times V(\psi) \setminus G_v$, we have that $v_1 < u_1$ or $v_2 > u_2$. Hence, $L_v = \bigcap_{\langle u_1, u_2 \rangle \in V(\psi) \times V(\psi) \setminus G_v} \{w \in (2^{I \cup O})^\omega : \llbracket w, \exists O.\psi \rrbracket < u_1 \text{ or } \llbracket w, \psi \rrbracket > u_2\}$, and so by dualization we have $(2^{I \cup O})^\omega \setminus L_v = \bigcup_{\langle u_1, u_2 \rangle \in V(\psi) \times V(\psi) \setminus G_v} \{w \in (2^{I \cup O})^\omega : \llbracket w, \exists O.\psi \rrbracket \geq u_1 \text{ and } \llbracket w, \psi \rrbracket \leq u_2\}$. Hence, we can obtain a UCW for L_v by dualizing an NBW that is the union of NBWs \mathcal{A}_{u_1, u_2} , for all $\langle u_1, u_2 \rangle \in V(\psi) \times V(\psi) \setminus G_v$, each of which is the product of NBWs $\mathcal{A}_{\exists O.\psi}^{\geq u_1}$ and $\mathcal{A}_\psi^{\leq u_2}$.

Observe that in all cases, the size of the NBW is $2^{O(|\psi|)}$. Indeed, there are at most $2^{2^{|\psi|}}$ pairs in the union, and, by Theorem 4, the size of the NBW for each pair is $2^{O(|\psi|)}$.

The lower bound follows from the 2EXPTIME-hardness of LTL GE-realizability. \square

By Theorem 4, the number of possible satisfaction values for ψ is at most $2^{|\psi|}$. Thus, the number of possible values for $\text{comb}(A, G)$, where A and G are satisfaction values of ψ , is at most $2^{2^{|\psi|}}$. Using binary search over the image of comb , we can use Theorem 7 to obtain the following.

Corollary 1. *The AG GE-synthesis problem can be solved in doubly-exponential time.*

Remark 5 [GE-synthesis as a special case of AG GE-synthesis]. The two approaches taken in Sect. 5.2 can be captured by an appropriate comb function. Indeed, for GE-synthesis with a threshold, we can use the function comb with $\text{comb}(A, G) = 1$ if $A \geq v \rightarrow G \geq v$, and $\text{comb}(A, G) = 0$ otherwise. For GE-synthesis (without a threshold), we can use the function comb with $\text{comb}(A, G) = 1$ if $A = G$, and $\text{comb}(A, G) = 0$ otherwise (recall that $A \geq G$ by definition). However, the solution described in Sect. 5.2 is simpler than the one described here for the general case. \square

5.4 LTL[\mathcal{F}] GE-Synthesis in Stochastic Environments

The setting of LTL[\mathcal{F}] GE-synthesis studied in Sects. 5.2 and 5.3 takes the different satisfaction values into an account, but is binary, in the sense that a specification is either (possibly AG) GE-realizable, or is not. In particular, in case the specification is not GE-realizable, synthesis algorithms only return “no”. In this section we add a quantitative measure also to the underlying realizability question. We do so by assuming a stochastic environment, with a known distribution on the inputs sequences, and analyzing the expected performance of the system.

For completeness, we remind the reader of some basics of probability theory. For a comprehensive reference see e.g., [25]. Let Σ be a finite alphabet, and let ν be some *probability distribution* over Σ^ω . For example, in the uniform distribution over $(2^I)^\omega$, the probability space is induced by sampling each letter with probability $2^{-|I|}$, corresponding to settings in which each signal in I always holds in probability $\frac{1}{2}$. We assume ν is given by a finite Markov Decision Process (MDP). That is, ν is induced by the distribution of each letter $i \in 2^I$ at each time step, determined by a finite stochastic control process that takes into account also the outputs generated by the system (see [2] for the precise model). A *random variable* is then a function $X : \Sigma^\omega \rightarrow \mathbb{R}$. When X has a finite image V , which is the case in our setting, its *expected value* is $\mathbb{E}[X] = \sum_{v \in V} v \cdot \Pr(X^{-1}(v))$. Intuitively, $\mathbb{E}[X]$ is the “average” value that X attains. Next, consider an *event* $E \subseteq \Sigma^\omega$. The *conditional expectation of X with respect to E* is $\mathbb{E}[X|E] = \frac{\mathbb{E}[\mathbb{1}_E X]}{\Pr(E)}$, where $\mathbb{1}_E X$ is the random variable that assigns $X(w)$ to $w \in E$ and 0 to $w \notin E$. Intuitively, $\mathbb{E}[X|E]$ is the average value that X attains when restricting to words in E , and normalizing according to the probability of E itself.

We continue and review the *high-quality synthesis problem* [2], where the GE variant is not considered. There, the environment is assumed to be stochastic and we care for the expected satisfaction value of an LTL[\mathcal{F}] specification in

the computations of a transducer \mathcal{T} , assuming some given distribution on the inputs sequences. Formally, let $X_{\mathcal{T},\psi} : (2^I)^\omega \rightarrow \mathbb{R}$ be a random variable that assigns each sequence $x \in (2^I)^\omega$ of input signals with $\llbracket \mathcal{T}(x), \psi \rrbracket$. Then, when the sequences in $(2^I)^\omega$ are sampled according to a given distribution ν of $(2^I)^\omega$, we define $\llbracket \mathcal{T}, \psi \rrbracket^\nu = \mathbb{E}[X_{\mathcal{T},\psi}]$. Since ν is fixed, we omit it from the notation and use $\llbracket \mathcal{T}, \psi \rrbracket$ in the following.

Remark 6 [Relating LTL GE-synthesis with stochastic LTL $[\mathcal{F}]$ synthesis] Given an LTL formula ψ , we can view it as an LTL $[\mathcal{F}]$ formula with possible satisfaction values $\{0, 1\}$, apply to it high-quality synthesis *a-la* [2], and find a transducer \mathcal{T} that maximizes $\mathbb{E}[X_{\mathcal{T},\psi}]$. An interesting observation is that if \mathcal{T} GE-realizes ψ , then it also maximizes $\mathbb{E}[X_{\mathcal{T},\psi}]$. Indeed, all input sequences that can contribute to the expected satisfaction value, do so. \square

We introduce and study two measures for high-quality synthesis in a stochastic environment. In the first, termed *expected GE-synthesis*, all input sequences are sampled, yet the satisfaction value in each input sequence takes its hopefulness level into account. In the second, termed *conditional expected GE-synthesis*, only hopeful input sequences are sampled.

We start with expected GE-synthesis. There, instead of associating each sequence $x \in (2^I)^\omega$ with $\llbracket x \otimes \mathcal{T}(x), \psi \rrbracket$, we associate it with $X_{\mathcal{T},\psi}^{\text{comb}} = \text{comb}(\llbracket x, \exists O.\psi \rrbracket, \llbracket x \otimes \mathcal{T}(x), \psi \rrbracket)$, where *comb* is as described in Sect. 5.3, thus capturing the assume-guarantee semantics of quantitative GE-synthesis. Then, we define $\llbracket \mathcal{T}, \psi \rrbracket^{\text{comb}} = \mathbb{E}[X_{\mathcal{T},\psi}^{\text{comb}}]$. For example, taking *comb* as implication, we have $X_{\mathcal{T},\psi}^{\text{comb}} = \max\{\llbracket x \otimes \mathcal{T}(x), \psi \rrbracket, \llbracket x, \forall O.\neg\psi \rrbracket\}$, capturing the semantics of $(\exists O.\psi) \rightarrow \psi$.

Then, in conditional expected GE-synthesis, we consider $\exists O.\psi$ as an environment assumption, and factor it in using conditional expectation, parameterized by a threshold $v \in [0, 1]$. Formally, let $\exists O.\psi \geq v$ denote the event $\{x \in (2^I)^\omega : \llbracket x, \exists O.\psi \rrbracket \geq v\}$. Then, we define $\llbracket \mathcal{T}, \psi \rrbracket^{\text{cond}(v)} = \mathbb{E}[X_{\mathcal{T},\psi} | \exists O.\psi \geq v]$, assuming the event $\exists O.\psi \geq v$ has a strictly positive probability.

In [2], it is shown that the high-quality synthesis problem can be solved in doubly-exponential time, also in the presence of environment assumptions. In the solution, the first step is the translation of the involved formulas to DPWs. In order to extract from [2] the results relevant to us, we describe them by means of *discrete quantitative specifications*, defined as follows. A discrete quantitative specification Ψ over $I \cup O$ is given by means of a sequence $\mathcal{A}_1, \dots, \mathcal{A}_n$ of DPWs, with $(2^{I \cup O})^\omega = L(\mathcal{A}_1) \supseteq L(\mathcal{A}_2) \supseteq \dots \supseteq L(\mathcal{A}_n)$, and sequence $0 \leq v_1 < \dots < v_n \leq 1$ of values. For every $w \in (2^{I \cup O})^\omega$, the satisfaction value of w in Ψ , denoted $\llbracket w, \Psi \rrbracket$, is $\max\{v_i : w \in L(\mathcal{A}_i)\}$. We refer to n as the depth of Ψ .

Theorem 8 ([2]). *Consider a discrete quantitative specification Ψ over $I \cup O$. Let n be its depth and m be the size of the largest DPW in Ψ . For a transducer \mathcal{T} , let $X_{\mathcal{T}}$ be a random variable that assigns a word $x \in (2^I)^\omega$ with $\llbracket x \otimes \mathcal{T}(x), \Psi \rrbracket$.*

1. *We can synthesize a transducer \mathcal{T} that maximizes $\mathbb{E}[X_{\mathcal{T}}]$ in time m^n .*

- Given a DPW \mathcal{B} over 2^I such that $\Pr(L(\mathcal{B})) > 0$, we can synthesize a transducer \mathcal{T} that maximizes $\mathbb{E}[X_{\mathcal{T}}|\mathcal{B}]$ in time $m^n \cdot k$, where k is the size of \mathcal{B} .

We can now state the main results of this section.

Theorem 9. *Consider an LTL $[\mathcal{F}]$ formula ψ .*

- Given a function comb , we can find in doubly-exponential time a transducer that maximizes $\llbracket \mathcal{T}, \psi \rrbracket^{\text{comb}}$.
- Given a threshold $v \in [0, 1]$, we can find in doubly-exponential time a transducer that maximizes $\llbracket \mathcal{T}, \psi \rrbracket^{\text{cond}(v)}$.

Proof. Let $v_1 < v_2 < \dots < v_n$ be the possible satisfaction values of ψ (and hence also of $\exists O.\psi$ and of $\forall O.\psi$). By Theorem 4, we have that $n \leq 2^{|\psi|}$. For each v_i , we can construct a DPW $\mathcal{D}_{\text{comb}(\exists O.\psi, \psi)}^{\geq v_i}$ as in Theorem 7. It is not hard to see that the discrete quantitative specification given by the DPWs $\mathcal{D}_{\text{comb}(\exists O.\psi, \psi)}^{\geq v_i}$ and the values v_i , for $1 \leq i \leq n$, is equal to the specification $\text{comb}(\exists O.\psi, \psi)$. Thus, by Theorem 8 (1), we can find a transducer that maximizes $\mathbb{E}[X_{\mathcal{T}}]$ in time $(2^{2^{O(|\psi|)}})^{2^{|\psi|}} = 2^{2^{O(|\psi|)}}$.

Next, given $v \in [0, 1]$, we can check whether $\Pr(\exists O.\psi > v) > 0$, for example by converting a DPW $\mathcal{D}_{\exists O.\psi}^{\geq v}$ to an MDP, and reasoning about its Ergodic-components. Then, by Theorem 8 (2), we can find a transducer that maximizes $\mathbb{E}[X_{\mathcal{T}}|\exists O.\psi > v]$, in time $(2^{2^{O(|\psi|)}})^{2^{|\psi|}} \cdot 2^{2^{O(\psi)}} = 2^{2^{O(|\psi|)}}$. \square

Corollary 2. *The (possibly conditional) expected GE-synthesis problem for LTL $[\mathcal{F}]$ can be solved in doubly-exponential time.*

5.5 Guarantees in High-Quality GE-Synthesis

As in the Boolean setting, also in the high-quality one we would like to add to a GE-realizing transducer guarantees and indications about the satisfaction level. As we detail below, the quantitative setting offers many possible ways to do so.

High-Quality GE-Synthesis with Guarantees. We consider specifications of the form $\psi = \psi_{strong} \wedge \psi_{weak}$, where essentially, we seek a transducer that realizes ψ_{strong} and (possibly AG) GE-realizes ψ_{weak} . Maximizing the realization value of ψ_{strong} may conflict with maximizing the GE-realization value of ψ_{weak} , and there are different ways to trade-off the two goals. Technically, in the decision-problem variant, we are given two thresholds $v_1, v_2 \in [0, 1]$, and we seek a transducer \mathcal{T} that realizes ψ_{strong} with value at least v_1 , and GE-realizes ψ_{weak} with value at least v_2 . Then, one may start, for example, by maximizing the value v_1 , and then find the maximal value v_2 that may be achieved simultaneously. Alternatively, one may prefer to maximize v_2 , or some other combination of v_1 and v_2 . Also, it is possible to decompose ψ further, to several strong and weak components, each with its desired threshold.

The solutions in the different settings all involve a construction of a UCW $\mathcal{A}_{\psi_{strong}}^{\geq v_1}$, and its product with the automata constructed in the solutions for the

different GE-synthesis variants. We thus have the following. We note that when the solution for ψ_{weak} is Saffraless, we can use a UCW for ψ_{strong} to maintain a Saffraless construction.

Theorem 10. *The problem of LTL[\mathcal{F}] high-quality GE-synthesis with a guarantee can be solved in doubly-exponential time.*

Flags by a High-Quality GE-Realizing Transducer. In the quantitative setting, we parameterized the flags raised by the GE-realizing transducer by values in $[0, 1]$, indicating the announced satisfaction level. Thus, rather than talking about prefixes being green, red, or blue, we talk about them being v -green, v -red, and v -blue, for $v \in [0, 1]$, which essentially means that a satisfaction value of at least v is guaranteed (in green and blue flags) or is impossible (in red ones). We can think of those as “degrees” of green, red, and blue. Below, we formalize this intuition and argue that even an augmentation of a transducer that GE-realizes ψ by flags for all values in $V(\psi)$ leaves the problem in doubly-exponential time.

A quantitative language over $2^{I \cup O}$ is $L : (2^{I \cup O})^\omega \rightarrow [0, 1]$. For a quantitative language L and a word $w \in (2^{I \cup O})^*$, we define L^w as the quantitative language where for all $w' \in (2^{I \cup O})^\omega$, we have $L^w(w') = L(w \cdot w')$. For a value $v \in [0, 1]$, a word $w \in (2^{I \cup O})^*$ is v -green for L if L^w is v -realizable. That is, there is a transducer \mathcal{T} such that $\llbracket \mathcal{T}, L^w \rrbracket \geq v$. A word $x \in (2^I)^*$ is v -green for L if there is $y \in (2^O)^*$ such that $x \otimes y$ is v -green for L . Thus, when the environment generates x , the system can respond in a way that would guarantee v -realizability. Finally, we say that L is *green realizable* if there is a strategy $f : (2^I)^+ \rightarrow 2^O$ that for every threshold v and for every input $x \in (2^I)^+$ that is v -green for L , we have that $x \otimes f(x)$ is v -green for L . It is not hard to see that Theorem 3 carries over to the quantitative setting, thus quantitative optimal realizability is strictly stronger than quantitative green realizability. In particular, if a transducer \mathcal{T} optimally realizes an LTL[\mathcal{F}] formula ψ , then \mathcal{T} also green realizes ψ . In the full version, we describe quantitative definitions also for red and blue prefixes, and describe monitors for the detection of the various types of prefixes.

6 Discussion

We introduced and solved several variants of GE-synthesis. Our complexity results are tight and show that GE-synthesis is not more complex than traditional synthesis. In practice, however, traditional synthesis algorithms do not scale well, and much research is devoted for the development of methods and heuristics for coping with the implementation challenges of synthesis. A natural future research direction is to extend these heuristics and methods for GE-synthesis. We mention here two specific examples.

Efficient synthesis algorithms have been developed for fragments of LTL [21]. Most notable is the *GR(1) fragment* [18], which supports assume-guarantee reasoning, and for which synthesis has an efficient symbolic solution. Adding existential quantification to GR(1) specifications, which is how we handled LTL

GE-synthesis, is not handled by its known algorithms, and is an interesting challenge. The success of SAT-based model-checking have led to the development of SAT-based synthesis algorithms [6], where the synthesis problem is reduced to satisfiability of a QBF formula. The fact the setting already includes quantifiers suggests it can be extended to GE-synthesis. A related effort is *bounded synthesis* algorithms [13,24], where the synthesized systems are assumed to be of a bounded size and can be represented symbolically [10].

References

1. Almagor, S., Boker, U., Kupferman, O.: Formalizing and reasoning about quality. *J. ACM* **63**(3), 24:1–24:56 (2016)
2. Almagor, S., Kupferman, O.: High-quality synthesis against stochastic environments. In: Proceedings of 25th Annual Conference of the European Association for Computer Science Logic, LIPICs, vol. 62, pp. 28:1–28:17 (2016)
3. Almagor, S., Kupferman, O., Ringert, J.O., Velnér, Y.: Quantitative assume guarantee synthesis. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10427, pp. 353–374. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63390-9_19
4. Bloem, R., Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Better quality in synthesis through quantitative objectives. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 140–156. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02658-4_14
5. Bloem, R., Chatterjee, K., Jobstmann, B.: Graph games and reactive synthesis. *Handbook of Model Checking*, pp. 921–962. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-10575-8_27
6. Bloem, R., Egly, U., Klampfl, P., Könighofer, R., Lonsing, F.: Sat-based methods for circuit synthesis. In: Proceedings of 14th International Conference on Formal Methods in Computer-Aided Design, pp. 31–34. IEEE (2014)
7. Borodin, A., El-Yaniv, R.: *Online Computation and Competitive Analysis*. Cambridge University Press, New York (1998)
8. Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Environment assumptions for synthesis. In: van Breugel, F., Chechik, M. (eds.) CONCUR 2008. LNCS, vol. 5201, pp. 147–161. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85361-9_14
9. Church, A.: Logic, arithmetics, and automata. In: Proceedings of International Congress of Mathematicians, vol. 1962, pp. 23–35. Institut Mittag-Leffler (1963)
10. Ehlers, R.: Symbolic bounded synthesis. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 365–379. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14295-6_33
11. Fisman, D., Kupferman, O., Lustig, Y.: Rational synthesis. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 190–204. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12002-2_16
12. Kupferman, O.: Automata theory and model checking. In: Clarke, E., Henzinger, T., Veith, H., Bloem, R. (eds.) *Handbook of Model Checking*, pp. 107–151. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-10575-8_4
13. Kupferman, O., Lustig, Y., Vardi, M.Y., Yannakakis, M.: Temporal synthesis for bounded systems and environments. In: Proceedings of 28th Symposium on Theoretical Aspects of Computer Science, pp. 615–626 (2011)

14. Kupferman, O., Perelli, G., Vardi, M.Y.: Synthesis with rational environments. *Ann. Math. Artif. Intell.* **78**(1), 3–20 (2016). <https://doi.org/10.1007/s10472-016-9508-8>
15. Kupferman, O., Piterman, N., Vardi, M.Y.: Safriless compositional synthesis. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 31–44. Springer, Heidelberg (2006). https://doi.org/10.1007/11817963_6
16. Kupferman, O., Vardi, M.Y.: Model checking of safety properties. *Formal Methods Syst. Des.* **19**(3), 291–314 (2001). <https://doi.org/10.1023/A:1011254632723>
17. Kupferman, O., Vardi, M.Y.: Safriless decision procedures. In: Proceedings of 46th IEEE Symposium on Foundations of Computer Science, pp. 531–540 (2005)
18. Piterman, N., Pnueli, A., Sa'ar, Y.: Synthesis of reactive(1) designs. In: Emerson, E.A., Namjoshi, K.S. (eds.) VMCAI 2006. LNCS, vol. 3855, pp. 364–380. Springer, Heidelberg (2005). https://doi.org/10.1007/11609773_24
19. Pnueli, A.: The temporal semantics of concurrent programs. *Theor. Comput. Sci.* **13**, 45–60 (1981)
20. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: Proceedings of 16th ACM Symposium on Principles of Programming Languages, pp. 179–190 (1989)
21. Alur, R., La Torre, S., Madhusudan, P.: Playing games with boxes and diamonds. In: Amadio, R., Lugiez, D. (eds.) CONCUR 2003. LNCS, vol. 2761, pp. 128–143. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45187-7_8
22. Rosner, R.: Modular synthesis of reactive systems. Ph.D thesis, Weizmann Institute of Science (1992)
23. Safra, S.: On the complexity of ω -automata. In: Proceedings of 29th IEEE Symposium on Foundations of Computer Science, pp. 319–327 (1988)
24. Schewe, S., Finkbeiner, B.: Bounded synthesis. In: Namjoshi, K.S., Yoneda, T., Higashino, T., Okamura, Y. (eds.) ATVA 2007. LNCS, vol. 4762, pp. 474–488. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75596-8_33
25. Sheldon, R.: *A First Course in Probability*. Pearson Education India, Delhi (2002)
26. Sistla, A.P., Vardi, M.Y., Wolper, P.: The complementation problem for Büchi automata with applications to temporal logic. *Theor. Comput. Sci.* **49**, 217–237 (1987)
27. Vardi, M.Y., Wolper, P.: Reasoning about infinite computations. *Inf. Comput.* **115**(1), 1–37 (1994)
28. Winnicott, D.W.: *Playing and Reality*. Penguin, Harmondsworth (1971)
29. Wolper, P.: Temporal logic can be more expressive. In: Proceedings of 22nd IEEE Symposium on Foundations of Computer Science, pp. 340–348 (1981)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

