

Synthesis with Rational Environments

Orna Kupferman¹, Giuseppe Perelli²(✉), and Moshe Y. Vardi³

¹ The Hebrew University, Jerusalem, Israel

orna@cs.huji.ac.il

² University of Naples “Federico II”, Naples, Italy

giuseppe.perelli@unina.it

³ Rice University, Houston, USA

vardi@cs.rice.edu

Abstract. *Synthesis* is the automated construction of a system from its specification. The system has to satisfy its specification in all possible environments. The environment often consists of agents that have objectives of their own. Thus, it makes sense to soften the universal quantification on the behavior of the environment and take the objectives of its underlying agents into an account. Fisman et al. introduced *rational synthesis*: the problem of synthesis in the context of rational agents. The input to the problem consists of temporal-logic formulas specifying the objectives of the system and the agents that constitute the environment, and a solution concept (e.g., Nash equilibrium). The output is a profile of strategies, for the system and the agents, such that the objective of the system is satisfied in the computation that is the outcome of the strategies, and the profile is stable according to the solution concept; that is, the agents that constitute the environment have no incentive to deviate from the strategies suggested to them.

In this paper we continue to study rational synthesis. First, we suggest an alternative definition to rational synthesis, in which the agents are rational but not cooperative. In the non-cooperative setting, one cannot assume that the agents that constitute the environment take into account the strategies suggested to them. Accordingly, the output is a strategy for the system only, and the objective of the system has to be satisfied in all the compositions that are the outcome of a stable profile in which the system follows this strategy. We show that rational synthesis in this setting is 2EXPTIME-COMPLETE, thus it is not more complex than traditional synthesis or cooperative rational synthesis. Second, we study a richer specification formalism, where the objectives of the system and the agents are not Boolean but quantitative. In this setting, the goal of the system and the agents is to maximize their outcome. The quantitative setting significantly extends the scope of rational synthesis, making the game-theoretic approach much more relevant.

O. Kupferman—The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement 278410, by the Israel Science Foundation (grant 1229/10), and by the US-Israel Binational Science Foundation (grant 2010431).

G. Perelli—This research was partially done while visiting Rice University.

M.Y. Vardi—NSF Expeditions in Computing project “ExCAPE: Expeditions in Computer Augmented Program Engineering”.

1 Introduction

Synthesis is the automated construction of a system from its specification. The basic idea is simple and appealing: instead of developing a system and verifying that it adheres to its specification, we would like to have an automated procedure that, given a specification, constructs a system that is correct by construction. The first formulation of synthesis goes back to Church [Chu63]; the modern approach to synthesis was initiated by Pnueli and Rosner, who introduced LTL (linear temporal logic) synthesis [PR89]. The LTL *synthesis problem* receives as input a specification in LTL and outputs a reactive system modeled by a finite-state transducer satisfying the given specification — if such exists. It is important to distinguish between input signals, assigned by the environment, and output signals, assigned by the system. A system should be able to cope with all values of the input signals, while setting the output signals to desired values [PR89]. Therefore, the quantification structure on input and output signals is different. Input signals are universally quantified while output signals are existentially quantified.

Modern systems often interact with other systems. For example, the clients interacting with a server are by themselves distinct entities (which we call agents) and are many times implemented by systems. In the traditional approach to synthesis, the way in which the environment is composed of its underlying agents is abstracted. In particular, the agents can be seen as if their only objective is to conspire to fail the system. Hence the term “hostile environment” that is traditionally used in the context of synthesis. In real life, however, many times agents have goals of their own, other than to fail the system. The approach taken in the field of algorithmic game theory [NRTV07] is to assume that agents interacting with a computational system are *rational*, i.e., agents act to achieve their own goals. Assuming agents rationality is a restriction on the agents behavior and is therefore equivalent to softening the universal quantification on the environment.¹ Thus, the following question arises: can system synthesizers capitalize on the rationality and goals of agents interacting with the system?

In [FKL10], Fisman et al. positively answered this question by introducing and studying *rational synthesis*. The input to the rational-synthesis problem consists of LTL formulas specifying the objectives of the system and the agents that constitute the environment, and a solution concept, e.g., dominant strategies, Nash Equilibria, and the like. The atomic propositions over which the objectives are defined are partitioned among the system and the agents, so that each of them controls a subset of the propositions. The desired output is a strategy profile such that the objective of the system is satisfied in the computation that is the outcome of the profile, and the agents that constitute the environment have no incentive to deviate from the strategies suggested to them (formally, the profile is an equilibrium with respect to the solution concept). Fisman et al. showed

¹ Early work on synthesis has realized that the universal quantification on the behaviors of the environment is often too restrictive. The way to address this point, however, has been by adding assumptions on the environment, which can be part of the specification (cf., [CHJ08]).

that there are specifications that cannot be realized in a hostile environment but are realizable in a rational environment. Moreover, the rational-synthesis problem for LTL and common solution concepts used in game theory can be solved in 2EXPTIME thus its complexity coincides with that of usual synthesis.

In this paper we continue the study of rational synthesis. We present the following three contributions. First, we suggest an alternative definition to rational synthesis, in which the agents are rational but not cooperative. Second, we study a richer specification formalism, where the objectives of the system and the agents are not Boolean but quantitative. Third, we show that all these variants of the rational synthesis problems can be reduced to model checking in fragments of *Strategy Logic* [CHP07]. Before we describe our contributions in more detail, let us highlight a different way to consider rational synthesis and our contribution here. *Mechanism design* is a field in game theory and economics studying the design of games whose outcome (assuming agents rationality) achieves some goal [NR01, NRTV07]. The outcome of traditional games depends on the final position of the game. In contrast, the systems we reason about maintain an *on-going interaction* with their environment, and we reason about their behavior by referring not to their final state (in fact, we consider non-terminating systems, with no final state) but rather to the *language* of computations that they generate. Rational synthesis can be viewed as a variant of mechanism design in which the game is induced by the objective of the system, and the objectives of both the system and the agents refer to their on-going interaction and are specified by temporal-logic formulas. Our contributions here correspond to the classic setting assumed in mechanism design: the agents need not be cooperative, and the outcome is not Boolean.

We argue that the definition of rational synthesis in [FKL10] is *cooperative*, in the sense that the agents that constitute the environment are assumed to follow the strategy profile suggested to them (as long as it is in an equilibrium). Here, we consider also a *non-cooperative* setting, in which the agents that constitute the environment may follow any strategy profile that is in an equilibrium, and not necessarily the one suggested to them by the synthesis algorithm. In many scenarios, the cooperative setting is indeed too optimistic, as the system cannot assume that the environment, even if it is rational, would follow a suggested strategy, rather than a strategy that is as good for it. Moreover, sometimes there is no way to communicate with the environment and suggest a strategy for it. From a technical point of view, we show that the non-cooperative setting requires reasoning about *all* possible equilibria, yet, despite this more sophisticated reasoning, it stays 2EXPTIME-COMPLETE. We achieve the upper bound by reducing rational synthesis to the model-checking problem for Strategy Logic (SL, for short). SL is a specification formalism that allows to explicitly quantify over strategies in games as first-order objects [CHP07]. While the model-checking problem for strategy logic is in general non-elementary, we show that it is possible to express rational synthesis in the restricted *Nested-Goal* fragment of SL, introduced in [MMPV14], which leads to the desired complexity. It is important to observe the following difference between the cooperative and the non-cooperative settings. In the cooperative one, we synthesize strategies for

all agents, with the assumption that the agent that corresponds to the system always follows his suggested strategy and the agents that constitute the environment decide in a rational manner whether to follow their strategies. On the other hand, in the non-cooperative setting, we synthesize a strategy only for the agent that corresponds to the system, and we assume that the agents that constitute the environment are rational, thus the suggested strategy has to win against all rational behaviors of the environment.

Our second contribution addresses a weakness of the classical synthesis problem, a weakness that is more apparent in the rational setting. In classical synthesis, the specification is Boolean and describes the expected behavior of the system. In many applications, systems can satisfy their specifications at different levels of quality. Thus, synthesis with respect to Boolean specifications does not address designers needs. This latter problem is a real obstacle, as designers would be willing to give up manual design only after being convinced that the automatic procedure that replaces it generates systems of comparable quality. In the last years we see a lot of effort on developing formalisms that would enable the specification of such quality measures [BCHJ09, ABK13].

Classical applications of game theory consider games with quantitative payoffs. In the Boolean setting, we assumed that the payoff of an agent is, say, 1, if its objective is satisfied, and is 0 otherwise. In particular, this means that agents whose objectives are not satisfied have no incentive to follow any strategy, even if the profile satisfies the solution concept. In real-life, rational objectives are rarely Boolean. Thus, even beyond our goal of synthesizing systems of high quality, the extension of the synthesis problem to the rational setting calls also for an extension to a quantitative setting. Unfortunately, the full quantitative setting is undecidable already in the context of model checking [Hen10]. In [FKL10], Fisman et al. extended cooperative rational synthesis to objectives in the multi-valued logic LLTL where specifications take truth values from a finite lattice.

We introduce here a new quantitative specification formalism, termed *Objective LTL*, (OLTL, for short). We first define the logic, and then study its rational synthesis. Essentially, an OLTL specification is a pair $\theta = \langle \Psi, f \rangle$, where $\Psi = \langle \psi_1, \psi_2, \dots, \psi_m \rangle$ is a tuple of LTL formulas and $f : \{0, 1\}^m \rightarrow \mathbb{Z}$ is a *reward function*, mapping Boolean vectors of length m to an integer. A computation η then maps θ to a reward in the expected way, according to the subset of formulas that are satisfied in η . In the rational synthesis problem for OLTL the input consists of OLTL specifications for the system and the other agents, and the goal of the system is to maximize its reward with respect to environments that are in an equilibrium. Again, we distinguish between a cooperative and a non-cooperative setting. Note that the notion of an equilibrium in the quantitative setting is much more interesting, as it means that all agents in the environment cannot expect to increase their payoffs. We show that the quantitative setting is not more complex than the non-quantitative one, thus quantitative rational synthesis is complete for 2EXPTIME in both the cooperative and non-cooperative settings.

2 Preliminaries

2.1 Games

A *concurrent game structure* (CGS, for short) [AHK02] is a tuple $\mathcal{G} \triangleq \langle \Phi, \Omega, (A_i)_{i \in \Omega}, S, \lambda, \tau, s_0 \rangle$, where Φ and $\Omega = \{\alpha_0, \dots, \alpha_k\}$ are finite sets of *atomic propositions* and *agents*, A_i are disjoint sets of *actions*, one for each agent α_i , S is a set of *states*, $s_0 \in S$ is a designated *initial state*, and $\lambda : S \rightarrow 2^\Phi$ is a *labeling function* that maps each state to the set of atomic propositions true in that state. By $A \triangleq \bigcup_{i \in \Omega} A_i$ we denote the union of all possible action for all the agents. Let $D \triangleq A_0 \times \dots \times A_k$ be the set of *decisions*, i.e., $(k + 1)$ -tuples of actions representing the choices of an action for each agent. Then, $\tau : S \times D \rightarrow S$ is a deterministic *transition function* mapping a pair of a state and a decision to a state.

A *path* in a CGS \mathcal{G} is an infinite sequence of states $\eta = \eta_0 \cdot \eta_1 \cdot \dots \in S^\omega$ that agrees with the transition function, i.e., such that, for all $i \in \mathbb{N}$, there exists a decision $\mathbf{d} \in D$ such that $\eta_{i+1} = \tau(\eta_i, \mathbf{d})$. A *track* in a CGS \mathcal{G} is a prefix ρ of a path η , also denoted by $\eta_{\leq n}$, for a suitable $n \in \mathbb{N}$. A track ρ is *non-trivial* if $|\rho| > 0$, i.e., $\rho \neq \varepsilon$.

We use $\text{Pth} \subseteq S^\omega$ and $\text{Trk} \subseteq S^+$ to denote the set of all paths and non-trivial tracks, respectively. Also, for a given $s \in S$, we use Pth^s and Trk^s to denote the subsets of paths and tracks starting from $s \in S$. Intuitively, the game starts in the state s_0 and, at each step, each agent selects an action in its set. The game then deterministically proceeds to the next state according to the corresponding decision. Thus, the outcome of a CGS is a path, regulated by individual actions of the agents.

A *strategy* for Agent α_i is a tool used to decide which decision to take at each phase of the game. Formally, it is a function $\pi_i : \text{Trk} \rightarrow A_i$ that maps each non-trivial track to a possible action of Agent α_i . By Π_i we denote the set of all possible strategies for agent α_i . A *strategy profile* is a $(k + 1)$ -tuple $\mathbf{P} = \langle \pi_0, \dots, \pi_k \rangle \in \Pi_0 \times \dots \times \Pi_k$ that assigns a strategy to each agent. We denote by $\mathcal{P} \triangleq \Pi_0 \times \dots \times \Pi_k$ the set of all possible strategy profiles. For a strategy profile \mathbf{P} and a state s , we use $\eta = \text{play}(\mathbf{P}, s)$ to denote the path that is the outcome of a game that starts in s and agrees with \mathbf{P} , i.e., for all $i \in \mathbb{N}$, it holds that $\eta_{i+1} = \tau(\eta_i, \mathbf{d}_i)$, where $\mathbf{d}_i = (\pi_0(\eta_{\leq i}), \dots, \pi_k(\eta_{\leq i}))$. By $\text{play}(\mathbf{P}) = \text{play}(\mathbf{P}, s_0)$ we denote the unique path starting from s_0 obtained from \mathbf{P} .

We model reactive systems by deterministic transducers. A *transducer* is a tuple $\mathcal{T} = \langle \text{I}, \text{O}, \text{S}, s_0, \delta, \text{L} \rangle$, where I is a set of input signals assigned by the environment, O is a set of output signals, assigned by the system, S is a set of states, s_0 is an initial state, $\delta : \text{S} \times 2^{\text{I}} \rightarrow \text{S}$ is a transition function, and $\text{L} : \text{S} \rightarrow 2^{\text{O}}$ is a labeling function. When the system is in state $s \in \text{S}$ and it reads an input assignment $\sigma \in 2^{\text{I}}$, it changes its state to $s' = \delta(s, \sigma)$ where it outputs the assignment $\text{L}(s')$. Given a sequence $\varrho = \sigma_1, \sigma_2, \sigma_3, \dots \in (2^{\text{I}})^\omega$ of inputs, the *execution* of \mathcal{T} on ϱ is the sequence of states s_0, s_1, s_2, \dots such that for all $j \geq 0$, we have $s_{j+1} = \delta(s_j, \sigma_j)$. The *computation* $\eta \in (2^{\text{I}} \times 2^{\text{O}})^\omega$ of S on ϱ is then $\langle \text{L}(s_0), \sigma_1 \rangle, \langle \text{L}(s_1), \sigma_2 \rangle, \langle \text{L}(s_2), \sigma_3 \rangle, \dots$

2.2 Strategy Logic

Strategy Logic [CHP07,MMV10,MMPV12] (SL, for short) is a logic that allows to quantify over strategies in games as explicit first-order objects. Intuitively, such quantification, together with a syntactic operator called *binding*, allows us to restrict attention to restricted classes of strategy profiles, determining a subset of paths, in which a temporal specification is desired to be satisfied. Since nesting of quantifications and bindings is possible, such temporal specifications can be recursively formulated by an SL subsentence. From a syntactic point of view, SL is an extension of LTL with strategy variables $\text{Var}_0, \dots, \text{Var}_k$ for the agents, existential ($\langle\langle x_i \rangle\rangle$) and universal ($\llbracket x_i \rrbracket$) strategy quantifiers, and a binding operator of the form (α_i, x_i) that couples an agent α_i with one of its variables $x_i \in \text{Var}_i$.

We first introduce some technical notation. For a tuple $t = s(t_0, \dots, t_k)$, by $t[i \leftarrow d]$ we denote the tuple obtained from t by replacing the i -th component with d . We use \vec{x} as an abbreviation for the tuple $(x_0, \dots, x_k) \in \text{Var}_0 \times \dots \times \text{Var}_k$. By $\langle\langle \vec{x} \rangle\rangle = \langle\langle x_0 \rangle\rangle \dots \langle\langle x_k \rangle\rangle$, $\llbracket \vec{x} \rrbracket = \llbracket x_0 \rrbracket \dots \llbracket x_k \rrbracket$, and $b(\vec{x}) = (\alpha_0, x_0) \dots (\alpha_k, x_k)$ we denote the existential and universal quantification, and the binding of all the agents to the strategy profile variable \vec{x} , respectively. Finally, by $b(\vec{x}_{-i}, y_i) = (\alpha_0, x_0) \dots (\alpha_i, y_i) \dots (\alpha_k, x_k)$ we denote the changing of binding for Agent α_i from the strategy variable x_i to the strategy variable y_i in the global binding $b(\vec{x})$.

Here we define and use a slight variant of the *Nested-Goal* fragment of SL, namely $\text{SL}[\text{NG}]$, introduced in [MMPV14]. Formulas in $\text{SL}[\text{NG}]$ are defined with respect to a set Φ of atomic proposition, a set Ω of agents, and sets Var_i of strategy variables for Agent $\alpha_i \in \Omega$. The set of $\text{SL}[\text{NG}]$ formulas is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \varphi \mathbf{R}\varphi \mid \langle\langle x_i \rangle\rangle\varphi \mid \llbracket x_i \rrbracket\varphi \mid b(\vec{x})\varphi,$$

where, $p \in \Phi$ is an atomic proposition, $x_i \in \text{Var}_i$ is a variable, and $\vec{x} \in \text{Var}_0 \times \dots \times \text{Var}_k$ is a tuple of variables, one for each agent.

The LTL part has the classical meaning. The formula $\langle\langle x_i \rangle\rangle\varphi$ states that there exists a strategy for Agent α_i such that the formula φ holds. The formula $\llbracket x_i \rrbracket\varphi$ states that, for all possible strategies for Agent α_i , the formula φ holds. Finally, the formula $b(\vec{x})\varphi$ states that the formula φ holds under the assumption that the agents in Ω adhere to the strategy evaluation of the variable x_i coupled in $b(\vec{x})$.

As an example, $\langle\langle x_0 \rangle\rangle \llbracket x_1 \rrbracket b(\vec{x})(p\mathbf{U}q) \vee \llbracket y_0 \rrbracket \langle\langle y_1 \rangle\rangle b(\vec{y})(\mathbf{G}Fp \wedge \mathbf{G}\neg q)$ is an $\text{SL}[\text{NG}]$ formula stating that either the system α_0 has a strategy x_0 to enforce $p\mathbf{U}q$ or, for all possible behaviors y_0 , the environment has a strategy y_1 to enforce both $\mathbf{G}Fp$ and $\mathbf{G}\neg q$.

We denote by $\text{free}(\varphi)$ the set of strategy variables occurring in φ but not in a scope of a quantifier. A formula φ is *closed* if $\text{free}(\varphi) = \emptyset$.

Similarly to the case of first order logic, an important concept that characterizes the syntax of SL is the one of *alternation depth* of quantifiers, i.e., the maximum number of quantifier switches $\langle\langle x_i \rangle\rangle \llbracket x_j \rrbracket \langle\langle x_j \rangle\rangle$, in the formula.

A precise formalization of the concepts of alternation depth can be found in [MMV10, MMPV12].

Now, in order to define the semantics of SL, we use the auxiliary concept of assignment. Let $\text{Var} = \bigcup_{i=0}^k \text{Var}_i$ be a set of variables for the agents in Ω , an *assignment* is a function $\chi : \text{Var} \cup \Omega \rightarrow \Pi$ mapping variables and agents to a relevant strategy, i.e., for all $\alpha_i \in \Omega$ and $x_i \in \text{Var}_i$, we have that $\chi(\alpha_i), \chi(x_i) \in \Pi_i$. Let $\text{Asg} \triangleq \Pi^{\text{Var} \cup \Omega}$ denote the set of all assignments. For an assignment χ and elements $l \in \text{Var} \cup \Omega$, we use $\chi[l \mapsto \pi] \in \text{Asg}$ to denote the new assignment that returns π on l and the value of χ on the other ones, i.e., $\chi[l \mapsto \pi](l) \triangleq \pi$ and $\chi[l \mapsto \pi](l') \triangleq \chi(l')$, for all $l' \in (\text{Var} \cup \Omega) \setminus \{l\}$. By $\text{play}(\chi, s)$ we denote the path $\text{play}(\mathcal{P}, s)$, for the strategy profile \mathcal{P} that is compatible with χ .

We now describe when a given game \mathcal{G} and a given assignment χ satisfy an SL formula φ , where $\text{dom}(\chi)^2 = \text{free}(\varphi) \cup \Omega$. We use $\mathcal{G}, \chi, s \models \varphi$ to indicate that the path $\text{play}(\chi, s)$ satisfies φ over the CGS \mathcal{G} . For φ in LTL, the semantics is as usual [MP92]. For the other operators, the semantics is as follows.

1. $\mathcal{G}, \chi, s \models \langle\langle x_i \rangle\rangle \varphi$ if there exists a strategy π_i for α_i such that $\mathcal{G}, \chi[x_i \mapsto \pi_i], s \models \varphi$;
2. $\mathcal{G}, \chi, s \models [[x_i]] \varphi$ if, for all strategies π_i for α_i , it holds that $\mathcal{G}, \chi[x_i \mapsto \pi_i], s \models \varphi$;
3. $\mathcal{G}, \chi, s \models b(\vec{x}) \varphi$ if it holds that $\mathcal{G}, \chi[\alpha_0 \mapsto x_0] \dots [\alpha_k \mapsto x_k], s \models \varphi$.

Finally, we say that \mathcal{G} satisfies φ , and write $\mathcal{G} \models \varphi$, if there exists an assignment χ such that $\mathcal{G}, \chi, s_0 \models \varphi$.

Intuitively, at Items 1 and 2, we evaluate the existential and universal quantifiers over a variable x_i by associating with it a suitable strategy. At Item 3 we commit the agents to use the strategy contained in the tuple variable \vec{x} .

Theorem 1. [MMPV14] *The model-checking problem for SL[NG] is $(d + 1)$ EXPTIME with d being the alternation depth of the specification.*

2.3 Rational Synthesis

We define two variants of rational synthesis. The first, *cooperative rational synthesis*, was introduced in [FKL10]. The second, *non-cooperative rational synthesis*, is new.

We work with the following model: the world consists of a *system* and an environment composed of k agents: $\alpha_1, \dots, \alpha_k$. For uniformity, we refer to the system as Agent α_0 . We assume that Agent α_i controls a set X_i of propositions, and the different sets are pairwise disjoint. At each point in time, each agent sets his propositions to certain values. Let $X = \bigcup_{0 \leq i \leq k} X_i$, and $X_{-i} = X \setminus X_i$. Each agent α_i (including the system) has an objective φ_i , specified as an LTL formula over X .

This setting induces the CGS $\mathcal{G} = \langle \Phi, \Omega, (A_i)_{i \in \Omega}, S, \lambda, \tau, s_0 \rangle$ defined as follows. The set of agents $\Omega = \{\alpha_0, \alpha_1, \dots, \alpha_k\}$ consists of the system and the agents that constitute the environment. The actions of Agent α_i are the possible assignments to its variables. Thus, $A_i = 2^{X_i}$. We use A and A_{-i} to

² By $\text{dom}(f)$ we denote the domain of the function f .

denote the sets 2^X and 2^{X-i} , respectively. The nodes of the game record the current assignment to the variables. Hence, $S = A$, and for all $s \in S$ and $\langle \sigma_0, \dots, \sigma_k \rangle \in A_0 \times A_1 \times \dots \times A_k$, we have $\delta(s, \sigma_0, \dots, \sigma_k) = \langle \sigma_0, \dots, \sigma_k \rangle$.

A strategy for the system is a function $\pi_0 : \text{Trk} \rightarrow A_0$. In the standard synthesis problem, we say that π_0 realizes φ_0 if, no matter which strategies the agents composing the environment follow, all the paths in which the system follows π_0 satisfy φ_0 . In rational synthesis, on instead, we assume that the agents that constitute the environment are rational, which soften the universal quantification on the behavior of the environment.

Recall that the rational-synthesis problem gets a solution concept as a parameter. As discussed in Sect. 1, the fact that a strategy profile is a solution with respect to the concept guarantees that it is not worthwhile for the agents constituting the environment to deviate from the strategies assigned to them. Several solution concepts are studied and motivated in game theory. Here, we focus on the concepts of *dominant strategy* and *Nash equilibrium*, defined below.

The common setting in game theory is that the objective for each agent is to maximize his *payoff* – a real number that is a function of the outcome of the game. We use $\text{payoff}_i : \text{Pth} \rightarrow \mathbb{R}$ to denote the payoff function of Agent α_i . That is, payoff_i assigns to each possible path η a real number $\text{payoff}_i(\eta)$ expressing the payoff of α_i on η . For a strategy profile P , we use $\text{payoff}_i(P)$ to abbreviate $\text{payoff}_i(\text{play}(P, s_0))$. In the case of an LTL goal ψ_i , we define $\text{payoff}_i(\eta) = 1$ if $\eta \models \psi_i$ and $\text{payoff}_i(\eta) = 0$, otherwise.

The simplest and most appealing solution concept is dominant-strategies solution [OR94]. A *dominant strategy* is a strategy that an agent can never lose by adhering to, regardless of the strategies of the other agents. Therefore, if there is a profile of strategies $P = \langle \pi_0, \dots, \pi_k \rangle$ in which all strategies π_i are dominant, then no agent has an incentive to deviate from the strategy assigned to him in P . Formally, P is a *dominant strategy profile* if for every $1 \leq i \leq k$ and for every (other) profile P' , we have that $\text{payoff}_i(P') \leq \text{payoff}_i(P'[i \leftarrow \pi_i])$.

As an example, consider the game in Fig. 1(a), played by three agents, Alice, Bob, and Charlie, whose actions are $\{a_1, a_2\}$, $\{b_1, b_2\}$, and $\{c_1, c_2\}$, respectively. The arrows are labeled with the possible action of the agents. Each agent wants to visit a state marked with his initial letter, infinitely often. In this game, the strategy for Bob of always choosing b_2 on his node 2 is dominant, while all the possible strategies for Charlie are dominant. On the other hand, Alice has no dominant strategies, since her goal essentially depends on the strategies adopted by the other agents. In several games, it can happen that agents have not any dominant strategy. For this reason, one would consider also other kind of solution concepts.

Another well known solution concept is Nash equilibrium [OR94]. A strategy profile is a *Nash equilibrium* if no agent has an incentive to deviate from his strategy in P provided that the other agents adhere to the strategies assigned to them in P . Formally, P is a *Nash equilibrium profile* if for every $1 \leq i \leq k$ and for every (other) strategy π'_i for agent α_i , we have that $\text{payoff}_i(P[i \leftarrow \pi'_i]) \leq \text{payoff}_i(P)$. An important advantage of Nash equilibrium is that it is more likely




Fig. 1. A game.

to exist than an equilibrium of dominant strategies [OR94]³. A weakness of Nash equilibrium is that it is not nearly as stable as a dominant-strategy solution: if one of the other agents deviates from his assigned strategy, nothing is guaranteed.

For the case of repeated-turn games like infinite games, a suitable refinement of Nash Equilibria is the *Subgame perfect Nash-equilibrium* [Sel75] (SPE, for short). A strategy profile $P = \langle \pi_0, \dots, \pi_k \rangle$ is an SPE if for every possible history of the game, no agent α_i has an incentive to deviate from her strategy π_i , assuming that the other agents follow their strategies in P . Intuitively, an SPE requires the existence of a Nash Equilibrium for each subgame starting from a randomly generated finite path of the original one. In [FKL10], the authors have studied cooperative rational synthesis also for the solution concept of SPE. To do this, the synthesis algorithm in [FKL10] was extended to consider all possible histories of the game. In SL such a path property can be expressed combining strategy quantifiers with temporal operators. Indeed, the formula $\varphi = \llbracket \bar{x} \rrbracket b(\bar{x})G \psi(\bar{y})$, with $\text{free}(\varphi) = \bar{y}$, states that, for all possible profile strategies the agents can follow, the game always reaches a position in which the formula $\psi(\bar{y})$ holds. Thus, for all possible paths that can be generated by agents, the property holds. By replacing $\psi(\bar{y})$ with the above formula, we then obtain a formula that represents SPEs. Hence, the cooperative and non-cooperative synthesis problems can be asserted in SL also for SPE, and our results hold also for this solution concept.

In rational synthesis, we control the strategy of the system and assume that the agents that constitute the environment are rational. Consider a strategy profile $P = \langle \pi_0, \dots, \pi_k \rangle$ and a solution concept γ (that is, γ is either “dominant strategies” or “Nash equilibrium”). We say that P is *correct* if $\text{play}(P)$ satisfies φ_0 . We say that P is in a π_0 -fixed γ -equilibrium if the agents composing the environment have no incentive to deviate from their strategies according to the solution concept γ , assuming that the system continues to follow π_0 . Thus, P is a π_0 -fixed dominant-strategy equilibrium if for every $1 \leq i \leq k$ and for every (other) profile P' in which Agent 0 follows π_0 , we have that $\text{payoff}_i(P') \leq \text{payoff}_i(P'[i \leftarrow \pi_i])$. Note that for the case of Nash equilibrium, adding the requirement that P is π_0 -fixed does not change the definition of an equilibrium.

³ In particular, all k -agent turn-based games with ω -regular objectives have Nash equilibrium [CMJ04].

In the context of objectives in LTL, we assume the following simple payoffs. If the objective φ_i of Agent α_i holds, then his payoff is 1, and if φ_i does not hold, then the payoff of Agent i is 0. Accordingly, $\mathbf{P} = \langle \pi_0, \dots, \pi_k \rangle$ is in a dominant-strategy equilibrium if for every $1 \leq i \leq k$ and profile $\mathbf{P}' = \langle \pi_0', \dots, \pi_k' \rangle$ with $\pi_0' = \pi_0$, if $\text{play}(\mathbf{P}') \models \varphi_i$, then $\text{play}(\mathbf{P}'[i \leftarrow \pi_i]) \models \varphi_i$. Also, \mathbf{P} is in a Nash-equilibrium if for every $1 \leq i \leq k$ and strategy π_i' , if $\text{play}(\mathbf{P}[i \leftarrow \pi_i']) \models \varphi_i$, then $\text{play}(\mathbf{P}) \models \varphi_i$.

Definition 1 (Rational synthesis). *The input to the rational-strategy problem is a set X of atomic propositions, partitioned into X_0, \dots, X_k , LTL formulas $\varphi_0, \dots, \varphi_k$, describing the objectives of the system and the agents composing the environment, and a solution concept γ . We distinguish between two variants of the problem:*

1. In Cooperative rational synthesis [FKL10], the desired output is a strategy profile \mathbf{P} such that $\text{play}(\mathbf{P})$ satisfies φ_0 and \mathbf{P} is a π_0 -fixed γ -equilibrium.
2. In Non-cooperative rational synthesis, the desired output is a strategy π_0 for the system such that for every strategy profile \mathbf{P} that includes π_0 and is a π_0 -fixed γ -equilibrium, we have that $\text{play}(\mathbf{P})$ satisfies φ_0 .

Thus, in the cooperative variant of [FKL10], we assume that once we suggest to the agents in the environment strategies that are in a γ -equilibrium, they will adhere to the suggested strategies. In the non-cooperative variant we introduce here, the agents may follow any strategy profile that is in a γ -equilibrium, and thus we require the outcome of all these profiles to satisfy φ_0 . It is shown in [FKL10] that the cooperative rational synthesis problem is 2EXPTIME-COMPLETE.

Note that the input to the rational synthesis problem may not have a solution, so when we solve the rational-synthesis problem, we first solve the *rational-realizability* problem, which asks if a solution exists. As with classical synthesis, the fact that SL model-checking algorithms can be easily modified to return a regular witness for the involved strategies in case an existentially quantified strategy exists, makes the realizability and synthesis problems strongly related.

Example 1. Consider a file-sharing network with the system and an environment consisting of two agents. The system controls the signal \mathbf{d}_1 and \mathbf{d}_2 (Agent α_1 and α_2 can download, respectively) and it makes sure that an agent can download only when the other agent uploads. The system's objective is that both agents will upload infinitely often. Agent α_1 controls the signal \mathbf{u}_1 (Agent α_1 uploads), and similarly for Agent α_2 and \mathbf{u}_2 . The goal of both agents is to download infinitely often.

Formally, the set of atomic propositions is $X = \{\mathbf{d}_1, \mathbf{d}_2, \mathbf{u}_1, \mathbf{u}_2\}$, partitioned into $X_0 = \{\mathbf{d}_1, \mathbf{d}_2\}$, $X_1 = \{\mathbf{u}_1\}$, and $X_2 = \{\mathbf{u}_2\}$. The objectives of the system and the environment are as follows.

- $\varphi_0 = \mathbf{G} (\neg \mathbf{u}_1 \rightarrow \neg \mathbf{d}_2) \wedge \mathbf{G} (\neg \mathbf{u}_2 \rightarrow \neg \mathbf{d}_1) \wedge \mathbf{G} \mathbf{F} \mathbf{u}_1 \wedge \mathbf{G} \mathbf{F} \mathbf{u}_2$,
- $\varphi_1 = \mathbf{G} \mathbf{F} \mathbf{d}_1$,
- $\varphi_2 = \mathbf{G} \mathbf{F} \mathbf{d}_2$.

First, note that in standard synthesis, φ_0 is not realizable, as a hostile environment needs not upload. In the cooperative setting, the system can suggest to both agents the following TIT FOR TAT strategy: upload at the first time step, and from that point onward upload iff the other agent uploads. The system itself follows a strategy π_0 according to which it enables downloads whenever possible (that is, \mathbf{d}_2 is valid whenever Agent α_1 uploads, and \mathbf{d}_1 is valid whenever Agent α_2 uploads). It is not hard to see that the above three strategies are all dominant. Indeed, all the three objectives are satisfied. Thus, the triple of strategies is a solution for the cooperative setting, for both solution concepts.

What about the non-cooperative setting? Consider the above strategy π_0 of the system, and consider strategies for the agents that never upload. The tuple of the three strategies is in a π_0 -fixed Nash equilibrium.

This ensures strategies for the environment to be dominant. Indeed, if Agent α_2 changes her strategy, φ_1 is still satisfied and vice-versa. Indeed, as long as Agent α_2 sticks to his strategy, Agent α_1 has no incentive to change his strategy, and similarly for Agent α_2 . Thus, π_0 is not a solution to the non-cooperative rational synthesis problem for the solution concept of Nash equilibrium. On the other hand, we claim that π_0 is a solution to the non-cooperative rational synthesis problem for the solution concept of dominant strategies. For showing this, we argue that if π_1 and π_2 are dominant strategies for α_1 and α_2 , then φ_0 is satisfied in the path that is the outcome of the profile $\mathbf{P} = \langle \pi_0, \pi_1, \pi_2 \rangle$. To see this, consider such a path $\eta = \text{play}(\pi_0, \pi_1, \pi_2)$. We necessarily have that $\eta \models \varphi_1 \wedge \varphi_2$. Indeed, otherwise π_1 and π_2 would not be dominant, as we know that with the strategies described above, α_1 and α_2 can satisfy their objectives. Now, since $\eta \models \varphi_1 \wedge \varphi_2$, we know that \mathbf{u}_2 and \mathbf{u}_1 hold infinitely often in η . Also, it is not hard to see that the formulas $\mathbf{G}(\neg \mathbf{u}_1 \rightarrow \neg \mathbf{d}_2)$ and $\mathbf{G}(\neg \mathbf{u}_2 \rightarrow \neg \mathbf{d}_1)$ are always satisfied in the context of π_0 , no matter how the other agents behave. It follows that $\eta \models \varphi_0$, thus π_0 is a solution of the non-cooperative rational synthesis problem for dominant strategies.

3 Qualitative Rational Synthesis

In this section we study cooperative and non-cooperative rational synthesis and show that they can be reduced to the model-checking problem for $\text{SL}[\text{NG}]$. The cooperative and non-cooperative rational synthesis problems for several solution concepts can be stated in $\text{SL}[\text{NG}]$.

We first show how to state that a given strategy profile $\vec{y} = (y_0, \dots, y_k)$ is in a y_0 -fixed γ -equilibrium. For $\alpha_i \in \Omega$, let φ_i be the objective of Agent α_i . For a solution concept γ and a strategy profile $\vec{y} = (y_0, \dots, y_k)$, the formula $\varphi^\gamma(\vec{y})$, expressing that the profile \vec{y} is a y_0 -fixed γ -equilibrium, is defined as follows.

- For the solution concept of dominant strategies, we define:

$$\varphi^\gamma(\vec{y}) := \llbracket \vec{z} \rrbracket \bigwedge_{i=1}^k (b(\vec{z})\varphi_i \rightarrow b(\vec{z}_{-i}, y_i)\varphi_i).$$
- For the solution concept of Nash equilibrium, we define:

$$\varphi^\gamma(\vec{y}) := \llbracket \vec{z} \rrbracket \bigwedge_{i=1}^k (b(\vec{y}_{-i}, z_i)\varphi_i \rightarrow b(\vec{y})\varphi_i).$$
- For the solution concept of Subgame Perfect Equilibrium, we define:

$$\varphi^\gamma(\vec{y}) := \llbracket \vec{x} \rrbracket b(\vec{x}_{-0}, y_0)\mathbf{F} \bigwedge_{i=1}^k \llbracket z_i \rrbracket b(\vec{y}_{-i}, z_i)\varphi_i \rightarrow b(\vec{y})\varphi_i.$$

We can now state the existence of a solution to the cooperative and non-cooperative rational-synthesis problem, respectively, with input $\varphi_0, \dots, \varphi_k$ by the closed formulas:

1. $\varphi_{cRS}^\gamma := \langle\langle y_0 \rangle\rangle \langle\langle y_1 \rangle\rangle \dots \langle\langle y_k \rangle\rangle (\varphi^\gamma(\vec{y}) \wedge \varphi_0)$;
2. $\varphi_{noncRS}^\gamma := \langle\langle y_0 \rangle\rangle [\![y_1]\!] \dots [\![y_k]\!] (\varphi^\gamma(\vec{y}) \rightarrow \varphi_0)$.

Indeed, the formula 1 specifies the existence of a strategy profile $\mathbf{P} = \langle \pi_0, \dots, \pi_k \rangle$ that is π_0 -fixed γ -equilibrium and such that the outcome satisfies φ_0 . On the other hand, the formula 2 specifies the existence of a strategy π_0 for the system such that the outcome of all profiles that are in a π_0 -fixed γ -equilibrium satisfy φ_0 .

As shown above, all the solution concepts we are taking into account can be specified in $\text{SL}[\text{NG}]$ with formulas whose length is polynomial in the number of the agents and in which the alternation depth of the quantification is 1. Hence we can apply the known complexity results for $\text{SL}[\text{NG}]$:

Theorem 2 (Cooperative and non-cooperative rational-synthesis complexity). *The cooperative and non-cooperative rational-synthesis problems in the qualitative setting are 2EXP TIME -complete.*

Proof. Consider an input $\varphi_0, \dots, \varphi_k$, X , and γ to the cooperative or non-cooperative rational-synthesis problem. As explained in Sect. 2.3, the input induces a game \mathcal{G} with nodes in 2^X . As detailed above, there is a solution to the cooperative (resp., non-cooperative) problem iff the $\text{SL}[\text{NG}]$ formula φ_{cRS}^γ (resp., φ_{noncRS}^γ) is satisfied in \mathcal{G} . The upper bound then follows from the fact that the model checking problem for $\text{SL}[\text{NG}]$ formulas of alternation depth 1 is in $2\text{EXP TIME-COMplete}$ in the size of the formula (Cf., Theorem 1). Moreover, the model-checking algorithm can return finite-state transducers that model strategies that are existentially quantified.

For the lower bound, it is easy to see that the classical LTL synthesis problem is a special case of the cooperative and non-cooperative rational synthesis problem. Indeed, $\varphi(I, O)$ is realizable against a hostile environment iff the solution to the non-cooperative rational synthesis problem for a system that has an objective φ and controls I and an environment that consists of a single agent that controls O and has an objective *True*, is positive.

4 Quantitative Rational Synthesis

As discussed in Sect. 1, a weakness of classical synthesis algorithms is the fact that specifications are Boolean and miss a reference to the quality of the satisfaction. Applications of game theory consider games with quantitative payoffs. Thus, even more than the classical setting, the rational setting calls for an extension of the synthesis problem to a quantitative setting. In this section we introduce *Objective LTL*, a quantitative extension of LTL, and study an extension of the rational-synthesis problem for specifications in Objective LTL. As opposed to other multi-valued specification formalisms used in the context of synthesis of high-quality systems [BCHJ09, ABK13], Objective LTL uses the syntax and semantics of LTL and only augments the specification with a reward function that enables a succinct and convenient prioritization of sub-specifications.

4.1 The Quantitative Setting

Objective LTL (OLTL, for short) is an extension of LTL in which specifications consist of sets of LTL formulas weighted by functions. Formally, an OLTL *specification* over a set X of atomic propositions is a pair $\theta = \langle \Psi, f \rangle$, where $\Psi = \langle \psi_1, \psi_2, \dots, \psi_m \rangle$ is a tuple of LTL formulas over X and $f : \{0, 1\}^m \rightarrow \mathbb{Z}$ is a *reward function*, mapping Boolean vectors of length m to integers. We assume that f is given by a polynomial. We use $|\Psi|$ to denote $\sum_{i=1}^m |\psi_i|$. For a computation $\eta \in (2^X)^\omega$, the *signature of Ψ in η* , denoted $\text{sig}(\Psi, \eta)$, is a vector in $\{0, 1\}^m$ indicating which formulas in Ψ are satisfied in η . Thus, $\text{sig}(\Psi, \eta) = \langle v_1, v_2, \dots, v_m \rangle$ is such that for all $1 \leq i \leq m$, we have that $v_i = 1$ if $\eta \models \psi_i$ and $v_i = 0$ if $\eta \not\models \psi_i$. The *value of θ in η* , denoted $\text{val}(\theta, \eta)$ is then $f(\text{sig}(\Psi, \eta))$. Thus, the interpretation of an OLTL specification is quantitative. Intuitively, $\text{val}(\theta, \eta)$ indicates the level of satisfaction of the LTL formulas in Ψ in η , as determined by the priorities induced by f . We note that the input of weighted LTL formulas studied in [CY98] is a special case of Objective LTL.

Example 2. Consider a system with m buffers, of capacities c_1, \dots, c_m . Let full_i , for $1 \leq i \leq m$, indicate that buffer i is full. The OLTL specification $\theta = \langle \Psi, f \rangle$, with $\Psi = \langle F \text{full}_1, F \text{full}_2, \dots, F \text{full}_m \rangle$ and $f(v) = c_1 \cdot v_1 + \dots + c_m \cdot v_m$ enables us to give different satisfaction values to the objective of filling a buffer. Note that $\text{val}(\langle \Psi, f \rangle, \eta)$ in a computation η is the sum of capacities of all filled buffers.

In the quantitative setting, the objective of Agent α_i is given by means of an OLTL specification $\theta_i = \langle \Psi_i, f_i \rangle$, specifications describe the payoffs to the agents, and the objective of each agent (including the system) is to maximize his payoff. For a strategy profile P , the payoff for Agent α_i in P is simply $\text{val}(\theta_i, \text{play}(P))$.

In the quantitative setting, rational synthesis is an optimization problem. Here, in order to easily solve it, we provide a decision version by making use of a threshold. It is clear that the optimization version can be solved by searching for the best threshold in the decision one.

Definition 2 (Quantitative rational synthesis). *The input to the quantitative rational-strategy problem is a set X of atomic propositions, partitioned into X_0, \dots, X_k , OLTL specifications $\theta_0, \theta_1, \dots, \theta_k$, with $\theta_i = \langle \Psi_i, f_i \rangle$, and a solution concept γ . We distinguish between two variants of the problem:*

1. In cooperative quantitative rational synthesis, the desired output for a given threshold $t \in \mathbb{N}$ is a strategy profile P such that $\text{payoff}_0(P) \geq t$ and P is in a π_0 -fixed γ -equilibrium.
2. In non-cooperative quantitative rational synthesis, the desired output for a given threshold $t \in \mathbb{N}$ is a strategy π_0 for the system such that, for each strategy profile P including π_0 and being in a π_0 -fixed γ -equilibrium, we have that $\text{payoff}_0(P) \geq t$.

Now, we introduce some auxiliary formula that helps us to formulate also the quantitative rational-synthesis problem in SL[NG].

For a tuple $\Psi = \langle \psi_1, \dots, \psi_m \rangle$ of LTL formulas and a signature $v = \{v_1, \dots, v_m\} \in \{0, 1\}^m$, let $\text{mask}(\Psi, v)$ be an LTL formula that characterizes computations η for which $\text{sig}(\Psi, \eta) = v$. Thus, $\text{mask}(\Psi, v) = (\bigwedge_{i:v_i=0} \neg\psi_i) \wedge (\bigwedge_{i:v_i=1} \psi_i)$.

We adjust the SL formulas $\Phi^\gamma(\vec{y})$ described in Sect. 3 to the quantitative setting. Recall that $\Phi^\gamma(\vec{y})$ holds iff the strategy profile assigned to \vec{y} is in a π_0 -fixed γ -equilibrium. There, the formula is a conjunction over all agents in $\{\alpha_1, \dots, \alpha_k\}$, stating that Agent α_i does not have an incentive to change his strategy. In our quantitative setting, this means that the payoff of Agent α_i in an alternative profile is not bigger than his payoff in \vec{y} . For two strategy profiles, assigned to \vec{y} and \vec{y}' , an SL formula that states that Agent α_i has no incentive that the profile would change from \vec{y} to \vec{y}' can state that the signature of Ψ in play(\vec{y}') results in a payoff to Agent α_i that is smaller than his current payoff. Formally, we have that:

$$\Phi_i^{eq}(\vec{y}, \vec{y}') = \bigvee_{v \in \{1, \dots, m_i\}: f_i(v) \leq \text{payoff}_i(\vec{y})} b(\vec{y}') \text{mask}(\Psi_i, v).$$

We can now adjust $\Phi^\gamma(\vec{y})$ for all the cases of solution concepts we are taking into account.

- For the solution concept of dominant strategies, we define:
 $\Phi^\gamma(\vec{y}) := \bigwedge_{i \in \{1, \dots, n\}} \llbracket \vec{z} \rrbracket \Phi_i^{eq}(\vec{y}, (\vec{z}[\alpha_0 \leftarrow y_0]));$
- For the solution concept of Nash equilibrium, we define:
 $\Phi^\gamma(\vec{y}) := \bigwedge_{i \in \{1, \dots, n\}} \llbracket \vec{z} \rrbracket \Phi_i^{eq}(\vec{y}, (\vec{y}[\alpha_i \leftarrow z_i]));$
- For the solution concept of Subgame Perfect Equilibrium, we define:
 $\varphi^\gamma(\vec{y}) := \llbracket \vec{x} \rrbracket b(\vec{x}[\alpha_0 \leftarrow y_0]) F \bigwedge_{i \in \{1, \dots, n\}} \llbracket \vec{z} \rrbracket \Phi_i^{eq}(\vec{y}, (\vec{z}[\alpha_0 \leftarrow y_0])).$

Once we adjust $\Phi^\gamma(\vec{y})$ to the quantitative setting, we can use the same SL formula used in the non-quantitative setting to state the existence of a solution to the rational synthesis problem. We have the following:

- $\Phi_{RS}^\gamma := \langle\langle y_0 \rangle\rangle \langle\langle y_1 \rangle\rangle \dots \langle\langle y_k \rangle\rangle (\Phi^\gamma(\vec{y}) \wedge \varphi_0);$
- $\Phi_{nonRS}^\gamma := \langle\langle y_0 \rangle\rangle \llbracket y_1 \rrbracket \dots \llbracket y_k \rrbracket (\Phi^\gamma(\vec{y}) \rightarrow \varphi_0).$

Theorem 3. *The cooperative and non-cooperative quantitative rational-synthesis problems are 2EXPTIME-COMplete.*

Proof. We can reduce the problems to the model-checking problem of the SL formulas Φ_{RS}^γ and Φ_{nonRS}^γ , respectively. We should, however, take care when analyzing the complexity of the procedure, as the formulas $\Phi_i^{eq}(\vec{y}, \vec{y}')$, which participate in Φ_{RS}^γ and Φ_{nonRS}^γ involve a disjunction over vectors in $\{0, 1\}^{m_i}$, resulting in Φ_{nonRS}^γ of an exponential length.

While the above prevents us from using the doubly exponential known bound on SL model checking for formulas of alternation depth 1 as is, it is not difficult to observe that the run time of the model-checking algorithm in [MMPV14], when applied to Φ_{nonRS}^γ , is only doubly exponential. The reason is the fact that the

inner exponent paid in the algorithm for SL model checking is due to the blow-up in the translation of the formula to a nondeterministic Büchi automaton over words (NBW, for short). In this translation, the exponentially many disjuncts are dominated by the exponential translation of the innermost LTL to NBW. Thus, the running time of the algorithm is doubly exponential, and it can return the witnessing strategies.

Hardness in 2EXPTIME follows easily from hardness in the non-quantitative setting.

5 Discussion

The understanding that synthesis corresponds to a game in which the objective of each player is to satisfy his specification calls for a mutual adoption of ideas between formal methods and game theory. In *rational synthesis*, introduced in [FKL10], synthesis is defined in a way that takes into an account the rationality of the agents that constitute the environment and involves an assumption that an agent cooperates with a strategy in which his objective is satisfied. Here, we extend the idea and consider also *non-cooperative* rational synthesis, in which agents need not cooperate with suggested strategies and may prefer different strategies that are at least as beneficial for them.

Many variants of the classical synthesis problem has been studied. It is interesting to examine the combination of the rational setting with the different variants. To start, the cooperative and non-cooperative settings can be combined into a framework in which one team of agents is competing with another team of agents, where each team is internally cooperative, but the two teams are non-cooperative. Furthermore, we plan to study rational synthesis with incomplete information. In particular, we plan to study rational synthesis with *incomplete information* [KV99], where agents can view only a subset of the signals that other agents output, and rational *stochastic* synthesis [CMJ04], which models the unpredictability of nature and involves stochastic agents that assign values to their output signals by means of a distribution function. Beyond a formulation of the richer settings, one needs a corresponding extension of strategy logic and its decision problems.

As discussed in Sect. 1, classical applications of game theory consider games with quantitative payoffs. We added a quantitative layer to LTL by introducing Objective-LTL and studying its rational synthesis. In recent years, researchers have developed more refined quantitative temporal logics, which enable a formal reasoning about the quality of systems. In particular, we plan to study rational synthesis for the multi-valued logics LTL[F] [ABK13], which enables a prioritization of different satisfaction possibilities, and LTL[D] [ABK14], in which discounting is used in order to reduce the satisfaction value of specifications whose eventualities are delayed. The rational synthesis problem for these logics induce a game with much richer, possibly infinitely many, profiles, making the search for a stable solution much more challenging.

References

- [ABK13] Almagor, S., Boker, U., Kupferman, O.: Formalizing and reasoning about quality. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (eds.) ICALP 2013, Part II. LNCS, vol. 7966, pp. 15–27. Springer, Heidelberg (2013)
- [ABK14] Almagor, S., Boker, U., Kupferman, O.: Discounting in LTL. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 424–439. Springer, Heidelberg (2014)
- [AHK02] Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *JACM* **49**(5), 672–713 (2002)
- [BCHJ09] 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)
- [CHJ08] 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)
- [CHP07] Chatterjee, K., Henzinger, T.A., Piterman, N.: Strategy Logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 59–73. Springer, Heidelberg (2007)
- [Chu63] Church, A.: Logic, arithmetics, and automata. In: Proceedings of the International Congress of Mathematicians, pp. 23–35. Institut Mittag-Leffler (1963)
- [CMJ04] Chatterjee, K., Majumdar, R., Jurdziński, M.: On nash equilibria in stochastic games. In: Marcinkowski, J., Tarlecki, A. (eds.) CSL 2004. LNCS, vol. 3210, pp. 26–40. Springer, Heidelberg (2004)
- [CY98] Courcoubetis, C., Yannakakis, M.: Markov decision processes and regular events. *IEEE Trans. Autom. Control* **43**(10), 1399–1418 (1998)
- [FKL10] 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)
- [Hen10] Henzinger, T.A.: From boolean to quantitative notions of correctness. In: POPL 2010, pp. 157–158. ACM (2010)
- [KV99] Kupferman, O., Vardi, M.Y.: Church’s problem revisited. *Bull. Symbolic Logic* **5**(2), 245–263 (1999)
- [MMPV12] Mogavero, F., Murano, A., Perelli, G., Vardi, M.Y.: What makes ATL* decidable? A decidable fragment of strategy logic. In: Koutny, M., Ulidowski, I. (eds.) CONCUR 2012. LNCS, vol. 7454, pp. 193–208. Springer, Heidelberg (2012)
- [MMPV14] Mogavero, F., Murano, A., Perelli, G., Vardi, M.Y.: Reasoning about strategies on the model-checking problem. *TOCL* **15**(4), 1–47 (2014). doi:[10.1145/2631917](https://doi.org/10.1145/2631917)
- [MMV10] Mogavero, F., Murano, A., Vardi, M.Y.: Reasoning about strategies. In: FSTTCS 2010, LIPIcs 8, pp. 133–144 (2010)
- [MP92] Manna, Z., Pnueli, A.: *The Temporal Logic of Reactive and Concurrent Systems - Specification*. Springer, Heidelberg (1992)
- [NR01] Nisan, N., Ronen, A.: Algorithmic mechanism design. *Games Econ. Behav.* **35**(1–2), 166–196 (2001)

- [NRTV07] Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
- [OR94] Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
- [PR89] Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: *POPL 1989*, pp. 179–190 (1989)
- [Sel75] Selten, R.: Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory* 4(1), 25–55 (1975)