



Synthesizing Systems with Optimal Average-Case Behavior for Ratio Objectives

Christian von Essen, Barbara Jobstmann

Verimag Research Report n° TR-2010-19

November 2, 2010

Reports are downloadable at the following address

<http://www-verimag.imag.fr>

Unité Mixte de Recherche 5104 CNRS - INPG - UJF

Centre Equation
2, avenue de VIGNATE
F-38610 GIERES
tel : +33 456 52 03 40
fax : +33 456 52 03 50
<http://www-verimag.imag.fr>



Synthesizing Systems with Optimal Average-Case Behavior for Ratio Objectives

Christian von Essen, Barbara Jobstmann

November 2, 2010

Abstract

We show how to automatically construct a system that satisfies a given logical specification and has an optimal average behavior with respect to a specification with fractional costs.

When synthesizing a system from a logical specification, it is often the case that several different systems satisfy the specification. In this case, it is usually not easy for the user to state formally which system she prefers. Prior work proposed to rank the correct systems by adding a quantitative aspect to the specification. A desired preference relation can be expressed with (i) a quantitative language, which is a function assigning a value to every possible behavior of a system, and (ii) an environment model defining the desired optimization criteria of the system, e.g., worst-case or average-case optimal.

In this paper, we show how to synthesize a system that is optimal for (i) a quantitative language given by an automaton with a fractional cost function, and (ii) an environment model given by a labeled Markov decision process. The objective of the system is to minimize the expected (fractional) costs. The solution is based on a reduction to Markov Decision Processes with extended-fractional cost functions which do not require that the costs in the denominator are strictly positive. We find an optimal strategy for these using a fractional linear program.

Keywords: Quantitative synthesis, MDP, fractional objective

Reviewers:

Notes: The research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7] under grant agreement no 257414 (ASCENS).

How to cite this report:

```
@techreport {TR-2010-19,  
  title = {Synthesizing Systems with Optimal Average-Case Behavior for Ratio Objectives},  
  author = {Christian von Essen, Barbara Jobstmann},  
  institution = {{Verimag} Research Report},  
  number = {TR-2010-19},  
  year = {}  
}
```

1 Introduction

Quantitative analysis techniques are usually used to measure quantitative properties of systems, such as timing, performance, or reliability (cf. [BHHK10, HKNP06, BBD⁺02]). We use quantitative reasoning in the classically Boolean contexts of verification and synthesis because they allow us to distinguish systems with respect to “soft constraints” like robustness [BGHJ09] or default behavior [BCHJ09]. This is particularly helpful in synthesis, where a system is automatically derived from a specification, because quantitative specifications allow us to guide the synthesis tool towards a desired implementation.

In this paper we show how quantitative specifications based on fractional objectives can be used to guide the synthesis process. In particular, we present a technique to synthesize a system with an average-case behavior that satisfies a logical specification and optimizes a quantitative objective given by a fractional objective.

The synthesis problem can be seen as a game between two players: the system and the environment (the context in which the system operates). The system has a fixed set of interface variables with a finite domain to interact with its environment. The variables are partitioned into a set of input and output variables. The environment can modify the set of input variables. For instance, an input variable can indicate the arrival of some packet on a router on a given port or the request of a client to use a shared resource. Each assignment to the input variables is a possible move of the environment in the synthesis game. The system reacts to the behavior of the environment by changing the value of the output variables. An assignment to the output variables is called an action of the system and describes a possible move of the system in the synthesis game. E.g., the system can grant a shared resource to Client *C* by setting a corresponding output variable. Environment and system change their variables in turns. In every step, first the system makes modification to the output variables, then the environment changes the input variables. The sequence of variable evaluations built up by this interplay is evaluated with respect to a specification. A logical (or qualitative) specification maps every sequence to 1 or 0, indicating whether the sequence satisfies the specification or not. For example, a sequence of evaluations in which the system grants a shared resource to two clients at the same time is mapped to 0 if the specification requires mutual exclusive access to this resource. The aim of the system in the synthesis game is to satisfy the specification independent of the choices of the environment. There might be several systems that can achieve this goal for a given specification. Therefore, Bloem et al. [BCHJ09] proposed to add a quantitative specification in order to rank the correct systems. A quantitative specification maps every sequence of variable evaluations to a value indicating how desirable this behavior is. A system can be seen as a set of behaviors. We can use, for example, the worst or the average value over all behaviors to assign a value to a system. Then we can ask for a system that optimizes this value, i.e., a system that achieves a better value than another system. Taking the worst value over the possible behaviors corresponds to assuming that the system is in an adversary environment. The average value is computed with respect to a probabilistic model of the environment [CHJS10]. In the average-case synthesis game, the environment player is replaced by a probabilistic player that is playing according to the probabilistic environment model.

In this paper, we present the first average-case synthesis algorithm for specifications that evaluate a behavior of the system with respect to the ratio of two cost functions [BCHJ09]. This ratio objective allows us, e.g., to ask for a system that optimizes the ratio between requests and acknowledgments in a server-client system. For the average-case analysis, we present a new environment model, which is based on Markov decision processes and generalizes the one in [CHJS10]. We solve the average-case synthesis problem with ratio objective by reduction to Markov decision processes with extended-fractional cost functions, which are fractional cost functions that do not require that the costs in the denominator are strictly positive. For irreducible Markov Decision Processes with extended-fractional cost functions, we present a solution based on linear programming. We also extended our algorithm to non-irreducible MDPs but due to space limitations we only present the irreducible case here. Our implementation can handle arbitrary (finite-state) MDPs.

Related Work. Researchers have considered a number of formalisms for quantitative specifications [ADMW09, CCH⁺05, CdAHS03, CdAF⁺06, CDH08, dA98, dAHM03, DG07, DKR08, KL07] but most of them (except for [BGHJ09]) do not consider long-run ratio objectives. In [BGHJ09], the environment is assumed

to be adversary, while we assume a probabilistic environment model. Regarding the environment model, there have been several notions of metrics for probabilistic systems and games proposed in the literature [dAMRS07, DGJP04]. The metrics measure the distance of two systems with respect to all temporal properties expressible in a logic, whereas we (like [CHJS10]) uses the quantitative specification to compare systems wrt the property of interest. In contrast to [CHJS10], we use ratio objectives and a more general environment model. Our environment model is the same as the one used for control and synthesis in the presence of uncertainty (cf. [BGL⁺04, CY90, BdA95]). However, in this context usually only qualitative specifications are considered. MDPs with long-run average objectives are well studied. The books [FV96, Put94] present a detailed analysis of this topic. Cyrus Derman [Der62] studied MDPs with a fractional objective. This work differs in two aspects from ours: first, Derman requires that the payoff of the cost function of the denominator is always strictly positive and second, the objective functions used in [Der62] is already given in terms of the expected cost of the first cost function to the expected cost of the second cost functions and not in terms of a single trace. Finally, we would like to note that the two choices we have in a quantitative synthesis problem, namely the choice of the quantitative language and the choice of environment model are the same two choices that appear in weighted automata and max-plus algebras (cf. [DKV09, GP97, CG79]).

2 Preliminaries

Words, Qualitative and Quantitative Languages. Given a finite alphabet Σ , a *word* $w = w_0w_1\dots$ is a finite or infinite sequence of elements of Σ . We use w_i to denote the $(i + 1)$ -th element in the sequence. If w is finite, then $|w|$ denotes the length of w , otherwise $|w|$ is infinity. We denote the empty word by ϵ , i.e., $|\epsilon| = 0$. We use Σ^* and Σ^ω to denote the set of finite and infinite words, respectively. Given a finite word $w \in \Sigma^*$ and a finite or infinite word $v \in \Sigma^* \cup \Sigma^\omega$, we write wv for the concatenation of w and v . A *qualitative language* φ is a function $\varphi : \Sigma^\omega \rightarrow \mathbb{B}$ mapping every infinite word to true or false. Intuitively, a qualitative language partitions the set of words into a set of good and a set of bad traces. A *quantitative language* [CDH08] ψ is a function $\psi : \Sigma^\omega \rightarrow \mathbb{R}^+ \cup \{\infty\}$ associating to each infinite word a value from the extended non-negative reals.

Specifications and automata with cost functions. An *automaton* is a tuple $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$, where Σ is a finite *alphabet*, Q is a finite set of *states*, $q_0 \in Q$ is an *initial state*, $\delta : Q \times \Sigma \rightarrow Q$ is the *transition function*, and $F \subseteq Q$ is a set of *safe states*. We assume without loss of generality that $Q \setminus F$ is closed under δ , i.e., $\forall s \in Q \setminus F, \forall a \in \Sigma : \delta(s, a) \in Q \setminus F$. We use $\delta^* : S \times L^* \rightarrow S$ to denote the closure of δ over finite words. Formally, given a word $w = w_0 \dots w_n \in \Sigma^*$, δ^* is defined inductively as $\delta^*(q, \epsilon) = q$, and $\delta^*(q, w) = \delta(\delta^*(q, w_0 \dots w_{n-1}), w_n)$. We use $|\mathcal{A}|$ to denote the size of the automaton.

The *run* ρ of \mathcal{A} on an infinite word $w = w_0w_1w_2\dots \in \Sigma^\omega$ is an infinite sequence of states $q_0q_1q_2\dots$ such that q_0 is the initial state of \mathcal{A} and $\forall i \geq 0 : \delta(q_i, w_i) = q_{i+1}$ holds. The run ρ is called *accepting* if for all $i \geq 0$, $q_i \in F$. A word w is accepting if the corresponding run is accepting. The *language* of \mathcal{A} , denoted by $\mathcal{L}_{\mathcal{A}}$, is the qualitative language $\mathcal{L}_{\mathcal{A}} : \Sigma^\omega \rightarrow \mathbb{B}$ mapping all accepting words to 1 and non-accepting words to 0, i.e., $\mathcal{L}_{\mathcal{A}}$ is the characteristic function of the set of all accepting words of \mathcal{A} .

Given an automaton $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$, a *cost function* $c : Q \times \Sigma \rightarrow \mathbb{N}$ is a function that maps every transition in \mathcal{A} to a non-negative integer. We use automata with cost functions and *objective functions* to define quantitative languages (or properties). Intuitively, the objective function tells us how to summarize the costs along a run. Given an automaton \mathcal{A} and two cost functions c_1, c_2 , the *ratio objective* [BGHJ09] computes the ratio between the costs seen along a run of \mathcal{A} on a word $w = w_0w_1w_2\dots \in \Sigma^\omega$:

$$\mathcal{R}(w) := \lim_{m \rightarrow \infty} \liminf_{l \rightarrow \infty} \frac{\sum_{i=m}^l c_1(\delta^*(q_0, w_0 \dots w_i), w_{i+1})}{1 + \sum_{i=m}^l c_2(\delta^*(q_0, w_0 \dots w_i), w_{i+1})} \quad (1)$$

The ratio objective is a generalization of the long-run average objective (also known as mean-payoff objective, cf. [ZP96]). We use $\mathcal{R}_{\frac{c_1}{c_2}}^{\mathcal{A}}$ to denote the quantitative language defined by \mathcal{A} , c_1 , c_2 , and the ratio objective function. If \mathcal{A} , c_1 , or c_2 are clear from the context, we drop them.

Intuitively, \mathcal{R} computes the long-run ratio between the costs accumulated along a run. The first limit allows us to ignore a finite prefix of the run, which ensures that we only consider the long-run behavior. The 1 in the denominator avoids division by 0, if the accumulated costs are 0 and has not effect if the accumulated costs are infinite. We need the limit inferior here because the sequence of the limit might not converge. Imagine a sequence $q^1 q'^2 q^4 q'^8 q^{16} \dots$ with $c_1(q) = 0$, $c_2(q) = 1$, $c_1(q') = 1$ and $c_2(q') = 0$. Then the value will alternate between 0 and 1 and hence the sequence will not converge. The limit inferior of this sequence is 0.

Finite-state system and Correctness A *finite-state system* $\mathcal{S} = (S, L, s_0, A, \delta, \tau)$ consists of the automaton $\mathcal{A} = (S, L, s_0, \delta, S)$, an *output (or action) alphabet* A , and an *output function* $\tau : S \rightarrow A$ assigning to each state of the system a letter from the output alphabet. The alphabet of the automaton L is called the *input alphabet* of the system. Given an input word w , the *run of the system \mathcal{S} on the word w* is simply the run of \mathcal{A} on the word w . For every word w over the input alphabet, the system produces a word over the joint input/output alphabet. We use $\mathcal{O}_{\mathcal{S}}$ to denote the function mapping input words to the joint input/output word, i.e., given an input word $w = w_0 w_1 \dots \in L^\omega$, $\mathcal{O}_{\mathcal{S}}(w)$ is the sequence of tuples $(l_0, a_0)(l_1, a_1) \dots \in (L \times A)^\omega$ such that (i) $l_i = w_i$ for all $i \geq 0$, (ii) $a_0 = \tau(s_0)$, and (iii) for all $i > 0$, $a_i = \tau(\delta^*(s_0, w_0 \dots w_{i-1}))$ holds.

Given a system \mathcal{S} with input alphabet L and output alphabet A , and an automaton \mathcal{A} with alphabet $\Sigma = L \times A$, we say that the system \mathcal{S} *satisfies the specification \mathcal{A}* , denoted $\mathcal{S} \models \mathcal{A}$, if for all input words, the joint input/output word produced by the system \mathcal{S} is accepted by the automaton \mathcal{A} , i.e., $\forall w \in L^\omega : (\mathcal{L}_{\mathcal{A}} \circ \mathcal{O}_{\mathcal{S}})(w) = 1$, where \circ denotes the function composition operator.

Probability space. We use the standard definitions of probability spaces. A *probability space* is given by a tuple $\mathcal{P} := (\Omega, \mathcal{F}, \mu)$, where Ω is the set of *outcomes or samples*, $\mathcal{F} \subseteq 2^\Omega$ is the σ -algebra defining the set of *measurable events*, and $\mu \in \mathcal{F} \rightarrow [0, 1]$ is a *probability measure* assigning a probability to each event such that $\mu(\Omega) = 1$ and for each countable set $E_1, E_2, \dots \in \mathcal{F}$ of events we have $\mu(\bigcup E_i) = \sum \mu(E_i)$. Recall that, since \mathcal{F} is a σ -algebra, it satisfies the following three conditions: (i) $\emptyset \in \mathcal{F}$, (ii) $E \in \mathcal{F}$ implies $\Omega \setminus E \in \mathcal{F}$ for any event E , and (iii) the union of any countable set $E_1, E_2, \dots \in \mathcal{F}$ is also in \mathcal{F} , i.e., $\bigcup E_i \in \mathcal{F}$. Given a measurable function $f : \mathcal{F} \rightarrow \mathbb{R} \cup \{+\infty, -\infty\}$, we use $\mathbb{E}_{\mathcal{P}}[f]$ to denote the expected value of f under μ , i.e.,

$$\mathbb{E}_{\mathcal{P}}[f] = \int_{\Omega} f d\mu \quad (2)$$

If \mathcal{P} is clear from the context we drop the subscript or replace it with the structure that defines \mathcal{P} . The integral used here is the Lebesgue Integral, which is commonly used to define the expected value of a random variable. Note that the expected value is always defined if the function f maps only to values in $\mathbb{R}^+ \cup \{\infty\}$.

Markov chains and Markov decision processes (MDP). Let $\mathcal{D}(S) := \{p : S \rightarrow [0, 1] \mid \sum_{s \in S} p(s) = 1\}$ be the *set of probability distributions* over a set S .

A *Markov decision process* is a tuple $\mathcal{M} = (S, s_0, A, \tilde{A}, p)$, where S is a finite set of *states*, $s_0 \in S$ is an *initial state*, A is the set of *actions*, $\tilde{A} : S \rightarrow 2^A$ is the *enabled action function* defining for each state s the set of enabled actions in s , and $p : S \times A \rightarrow \mathcal{D}(S)$ is a probabilistic *transition function*. For technical convenience we assume that every state has at least one enabled action, i.e., $\forall s \in S : |\tilde{A}(s)| \geq 1$. If $|\tilde{A}(s)| = 1$ for all states $s \in S$, then \mathcal{M} is called a *Markov chain (MC)*. In this case, we omit A and \tilde{A} from the definition of \mathcal{M} . Given a Markov chain \mathcal{M} , we say that \mathcal{M} is *irreducible* if every state can be reached from any other.

An *L-labeled Markov decision process* is a tuple $\mathcal{M} = (S, s_0, A, \tilde{A}, p, \lambda)$, where $(S, s_0, A, \tilde{A}, p)$ is a Markov decision process and $\lambda : S \rightarrow L$ is a labeling function such that \mathcal{M} is deterministic with respect to λ , i.e., $\forall s, a, s', s''$ if $p(s, a)(s') > 0$ and $p(s, a)(s'') > 0$, then $\lambda(s') \neq \lambda(s'')$. Since we use L -labeled Markov decision process to represent the behavior of the environment, we require that in every state all actions are enabled, i.e., $\forall s \in S : \tilde{A}(s) = A$.

Sample runs and strategies A (sample) run ρ of \mathcal{M} is an infinite sequence of tuples $(s_0, a_0)(s_1, a_1) \cdots \in (S \times A)^\omega$ of states and actions such that for all $i \geq 0$, (i) $a_i \in \tilde{A}(s_i)$ and (ii) $p(s_i, a_i)(s_{i+1}) > 0$. We use Ω to denote the set of all runs, and Ω_s for the set of runs starting at state s . A finite run of \mathcal{M} is a prefix of some infinite run. To avoid confusion, we use v to refer to a finite run. Given a finite run v , the set $\gamma(v) := \{\rho \in \Omega \mid \exists \rho' \in \Omega : \rho = v\rho'\}$ of all possible infinite extensions of v is called the *cylinder set* of v . We use the usual extension of $\gamma(\cdot)$ to sets of finite words.

A strategy is a function $\pi : (S \times A)^*S \rightarrow \mathcal{D}(A)$ that assigns a probability distribution to all finite sequences in $(S \times A)^*S$. A strategy must refer only to enabled actions, i.e., for all sequences $w \in (S \times A)^*$, states $s \in S$, and actions $a \in A$, if $\pi(ws)(a) > 0$, then action a has to be enabled in s , i.e., $a \in \tilde{A}(s)$. Strategies that do not use randomization are called *pure*. A strategy π is *pure* if for all finite sequences $w \in (S \times A)^*$ and for all states $s \in S$, there is an action $a \in A$ such that $\pi(ws)(a) = 1$. A *memoryless* strategy is independent of the history of the run, i.e., for all $w, w' \in (S \times A)^*$ and for all $s \in S$, $\pi(ws) = \pi(w's)$ holds. A memoryless strategy can be represented as function $\pi : S \rightarrow \mathcal{D}(A)$. A strategy is *positional* if it is pure and memoryless. In this case, it can be represented by a function $\pi : S \rightarrow A$ mapping states to actions. An MDP $\mathcal{M} = (S, s_0, A, \tilde{A}, p)$ together with a positional strategy $\pi : S \rightarrow A$ defines the Markov chain $\mathcal{M}^\pi = (S, s_0, A, \tilde{A}_\pi, p)$, in which only the actions prescribed in the strategy π are enabled, i.e., $\tilde{A}_\pi(s) = \{\pi(s)\}$. Note that every finite-state system \mathcal{S} with input alphabet S and output alphabet A denotes a strategy for \mathcal{M} .

Induced probability space, objective function, and optimal strategies. An MDP $\mathcal{M} = (S, s_0, A, \tilde{A}, p)$ together with a strategy π and a state $s \in S$ induces a probability space $\mathcal{P}_{\mathcal{M},s}^\pi = (\Omega_{\mathcal{M},s}^\pi, \mathcal{F}_{\mathcal{M},s}^\pi, \mu_{\mathcal{M},s}^\pi)$ over the cylinder sets of the runs. Hence, $\Omega_{\mathcal{M},s}^\pi = S^\omega$. The probability measure of a cylinder set is the probability that the MDP starts from state s and follows the common prefix under the strategy π . By convention $\mathcal{P}_{\mathcal{M}}^\pi := \mathcal{P}_{\mathcal{M},s_0}^\pi$. If \mathcal{M} is a Markov chain, then π is fixed (since there is only one available action in every state), and we simply write $\mathcal{P}_{\mathcal{M}}$.

An *objective function* of \mathcal{M} is a measurable function $f : (S \times A)^\omega \rightarrow \mathbb{R}^+ \cup \{\infty\}$ that maps runs of \mathcal{M} to values in $\mathbb{R}^+ \cup \{\infty\}$. We use $\mathbb{E}_{\mathcal{M},s}^\pi[f]$ to denote the expected value of f wrt the probability space induced by the MDP \mathcal{M} , a strategy π , and a state s .

We are interested in a strategy that has the least expected value for a given state. Given an MDP \mathcal{M} and a state s , a strategy π is called *optimal* for objective f and state s if $\mathbb{E}_{\mathcal{M},s}^\pi[f] = \min_{\pi'} \mathbb{E}_{\mathcal{M},s}^{\pi'}[f]$, where π' ranges over all possible strategies.

Given an MDP $\mathcal{M} = (S, s_0, A, \tilde{A}, p)$ and two cost function $c_1 : S \times A \rightarrow \mathbb{N}$ and $c_2 : S \times A \rightarrow \mathbb{N}$, the *ratio or fractional payoff value* is the function $\mathcal{R} : (S \times A)^\omega \rightarrow \mathbb{R}^+ \cup \{\infty\}$ mapping every run ρ to a value in $\mathbb{R}^+ \cup \{\infty\}$ as follows:

$$\mathcal{R}_{\frac{c_1}{c_2}}(\rho) := \lim_{m \rightarrow \infty} \liminf_{l \rightarrow \infty} \frac{\sum_{i=m}^l c_1(\rho_i)}{1 + \sum_{i=m}^l c_2(\rho_i)} \quad (3)$$

We drop the subscript $\frac{c_1}{c_2}$ if c_1 and c_2 are clear from the context.

3 Synthesis with Ratio Objective in Probabilistic Environments

In this section, we first present a variant of the quantitative synthesis problem introduced in [BCHJ09]. Then, we show how to solve the synthesis problem with a safety and a ratio specifications in a probabilistic environment described by an MDP.

The quantitative synthesis problem with probabilistic environments asks to construct a finite-state system \mathcal{S} that satisfies a qualitative specification and optimizes a quantitative specification under the given environment. The specifications are qualitative and quantitative languages over letters in $(L \times A)$, where L and A are the input and output alphabet of \mathcal{S} , respectively.

In order to compute the average behavior of a system, we need a model of the environment. In [CHJS10], the environment model is a probability space $\mathcal{P} = (L^\omega, \mathcal{F}, \mu)$ over the input words L^ω of the system defined by a finite L -labeled Markov chain. This model assumes that the behavior of the environment is independent of the behavior of the system. This restricts the modeling possibilities. It is impossible to model a client-server system, in which a client increases the probability of sending a request

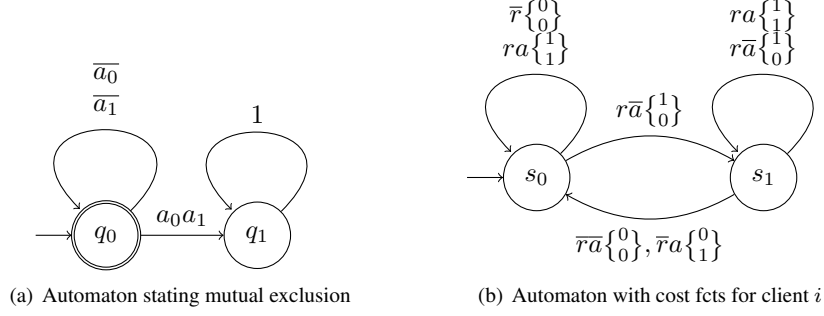


Figure 1: Specifications for the client-server example

if it has not been served in the previous step. Therefore, *our environment model* is a function f_e that maps every system $f_s : L^* \rightarrow A$ to a probability space $\mathcal{P} = (L^\omega, \mathcal{F}, \mu)$ over the input words L^ω . Note that every finite-state system defines such a system function f_s but not vice versa. To describe a particular environment model f_e , we use a finite L -labeled Markov decision process. Once we have an environment model, we can define what it means for a system to satisfy a specification under a given environment.

Definition 1 (Satisfaction). *Given a finite-state system \mathcal{S} with alphabets L and A , a qualitative specification φ over alphabet $L \times A$, and an environment model f_e , we say that \mathcal{S} satisfies φ under f_e (written $\mathcal{S} \models_{f_e} \varphi$ ¹) iff \mathcal{S} satisfies φ with probability 1, i.e.,*

$$\mathbb{E}_{f_e(\mathcal{S})}[\varphi \circ \mathcal{O}_{\mathcal{S}}] = 1.$$

Next, we define the value of a system with respect to a specification under an environment model and what it means for a system to optimize a specification. Then, we are ready to define the quantitative synthesis problem.

Definition 2 (Value of a system). *Given a finite-state system \mathcal{S} with alphabets L and A , a qualitative (φ) and a quantitative specification (ψ) over alphabet $L \times A$, and an environment model f_e , the value of \mathcal{S} with respect to φ and ψ under f_e is defined as the expected value of the function $\psi \circ \mathcal{O}_{\mathcal{S}}$ in the probability space $f_e(\mathcal{S})$, if \mathcal{S} satisfies φ , and ∞ otherwise. Formally,*

$$\text{Value}_{\varphi\psi}^{f_e}(\mathcal{S}) := \begin{cases} \mathbb{E}_{f_e(\mathcal{S})}[\psi \circ \mathcal{O}_{\mathcal{S}}] & \text{if } \mathcal{S} \models_{f_e} \varphi, \\ \infty & \text{otherwise.} \end{cases}$$

If φ is the set of all words, then we write $\text{Value}_{\psi}^{f_e}(\mathcal{S})$. Furthermore, we say \mathcal{S} optimizes ψ wrt f_e , if $\text{Value}_{\psi}^{f_e}(\mathcal{S}) \leq \text{Value}_{\psi}^{f_e}(\mathcal{S}')$ for all systems \mathcal{S}' .

Definition 3 (Quantitative realizability and synthesis problem). *Given a qualitative specification φ and a quantitative specification ψ over the alphabets $L \times A$ and an environment model f_e , the realizability problem asks to decide if there exists a finite-state system \mathcal{S} with alphabets L and A such that $\text{Value}_{\varphi\psi}^{f_e}(\mathcal{S}) \neq \infty$. The synthesis problem asks to construct a finite-state system \mathcal{S} (if it exists) s. t.*

1. $\text{Value}_{\varphi\psi}^{f_e}(\mathcal{S}) \neq \infty$ and
2. \mathcal{S} optimizes ψ wrt f_e .

In the following, we give an example of a quantitative synthesis problem.

¹Note that $\mathcal{S} \models_{f_e} \varphi$ and $\mathcal{S} \models \varphi$ coincide if (i) φ is prefix-closed (which is the case for the specifications, we consider here), and (ii) $f_e(\mathcal{S})$ assigns, for every finite word $w \in L^*$, a positive probability to the set of infinite words wL^ω .

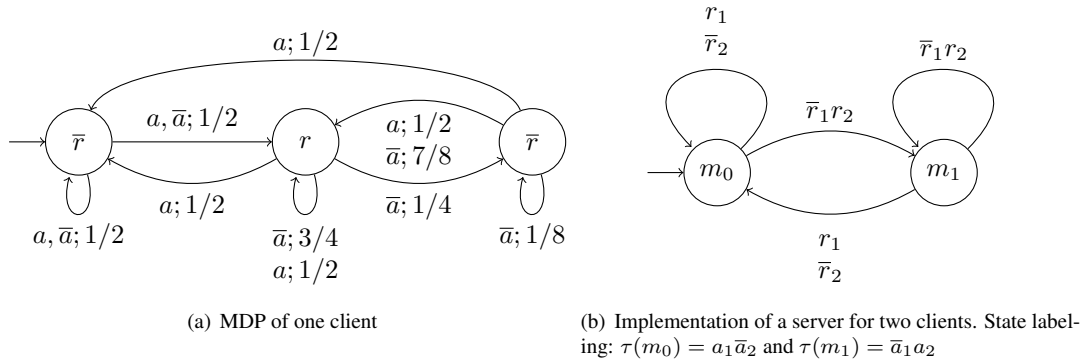


Figure 2: Specifications and implementation for the client-server example

Server-client example. Consider a server-client system with two clients and one server. Each server-client interface consists of two variables r_i (request) and a_i (acknowledge). Client i sends a request by setting r_i to 1. The server acknowledges the request by setting a_i to 1. We require that the server does not acknowledge both clients at the same time. Hence, our qualitative specification demands mutual exclusion. Figure 1(a) shows an automaton stating the mutual exclusion property for a_1 and a_2 . Edges are labeled with evaluations of a_1 and a_2 , e.g., \bar{a}_1a_2 states that $a_1 = 0$ and $a_2 = 1$. States drawn with a double circle are safe states. Among all systems satisfying the mutual exclusion property, we ask for a system that minimizes the average ratio between requests and useful acknowledges. An acknowledge is only useful if it is sent as a response to a request. To express this property, we can give a quantitative language defined by an automaton with two cost functions (c_1, c_2) and the ratio objective (Eqn. 1). Figure 1(b) shows an automaton labeled with tuples representing the two cost functions c_1 and c_2 for one client. The first component of the tuples represents cost function c_1 , the second component defines cost function c_2 . The cost function c_1 is 1, whenever we see a request. The cost function c_2 is 1, when we see a “useful” acknowledge, which is an acknowledge that matches an unacknowledged request. E.g., every acknowledge in state s_1 is useful, since the last request has not been acknowledged yet. In state s_0 only acknowledgments that answer a direct request are useful and get cost 1 (in the second component). This corresponds to a server with a buffer that can hold exactly one request. State s_1 says that there is a request in the buffer. If the machine is in this state and its request is granted, then it can clear its buffer. If another request comes while it is in this state and if it does not get an acknowledgment in the same round, then it has to drop one request.

Assume we know the expected behavior of the clients. E.g., in every step, Client 1 is expected to send a request with probability 0.5 independent of the acknowledgments. Client 2 changes its behavior based on the acknowledgments. We can describe the behavior of Client 2 by the labeled MDP shown in Figure 2(a). In the beginning the chance of getting a request from this client is 0.5. Once it has sent a request, i.e., it is in state r , the probability of sending a request again is very high until at least one acknowledgment is given. This is modeled by action \bar{a} at state r having a probability of $3/4$ to get into state r again, and a probability of $1/4$ to not send a request in the next step. In this case, we move to the right \bar{r} state. In this state, the probability of receiving a request from this client in the next step is even $7/8$. This means, that if this client does not receive an acknowledgment after having sent a request, then the possibility of receiving another request from this client in the next two steps is $1 - 1/4 * 1/8 = 31/32$.

Consider the finite-state system \mathcal{S} shown in Figure 2(b). It is an implementation of a server for two clients. The system has two state m_0 and m_1 labeled with $a_1\bar{a}_2$ and \bar{a}_1a_2 , respectively. We can compute the value of \mathcal{S} using the following two lemmas (Lem. 1, Lem. 2).

Lemma 1. Given (i) a finite-state system \mathcal{S} with alphabets L and A , (ii) an automaton \mathcal{A} with alphabet $L \times A$, and (iii) a L -labeled MDP \mathcal{M} defining an environment model for \mathcal{S} , there exists a Markov chain \mathcal{M}_c and two cost functions c_1 and c_2 such that

$$\mathcal{S} \models_{\mathcal{M}} \mathcal{L}_{\mathcal{A}} \stackrel{\text{Def. 1}}{\iff} \mathbb{E}_{\mathcal{M}}^{\mathcal{S}}[\mathcal{L}_{\mathcal{A}} \circ \mathcal{O}_{\mathcal{S}}] = 1 \iff \mathbb{E}_{\mathcal{M}_c}[R_{\frac{c_1}{c_2}}] = 0$$

Proof idea: The Markov chain \mathcal{M}_c is constructed by taking the synchronous product of \mathcal{S} , \mathcal{A} , and \mathcal{M} . In every state $(s, q, m) \in (S_S \times Q_A \times S_M)$, we take the action $a \in A$ given by the labeling function of the system $\tau(s)$ and move to a successor state for every input labels $l \in L$ such that there exists a state m' in the MDP \mathcal{M} with $\lambda(m') = l$ and $p(m, a)(m') > 0$. The corresponding successor states of the system and the automaton-state components are $s' = \delta_S(s, l)$ and $q' = \delta_A(q, (l, a))$. The probability distribution of \mathcal{M}_c is taken from the \mathcal{M} -component. The two cost functions are defined as follows: for state (s, q, m) and an action a , then $c_1((s, q, m), a) = 1$ and $c_2((s, q, m), a) = 0$, if q is a safe state in \mathcal{A} , otherwise $c_1((s, q, m), a) = 0$ and $c_2((s, q, m), a) = 1$. Intuitively, since the non-safe states of \mathcal{A} are (by definition) closed under δ_A and all actions in this set have the same cost, they all have the same value, namely ∞ , so does every state from which there is a positive probability to reach this set.²

Lemma 2. *Given (i) a finite-state system \mathcal{S} with alphabets L and A , (ii) an automaton \mathcal{A} with alphabet $L \times A$ with two cost functions c_1 and c_2 , and (iii) a L -labeled MDP \mathcal{M} defining an environment model for \mathcal{S} , there exists a Markov chain \mathcal{M}_c and two cost functions d_1 and d_2 such that*

$$\text{Value}_{\mathcal{R}_{\frac{c_1}{c_2}}}^{\mathcal{M}}(\mathcal{S}) \stackrel{\text{Def. 2}}{=} \mathbb{E}_{\mathcal{M}}^{\mathcal{S}}[\mathcal{R}_{\frac{c_1}{c_2}} \circ \mathcal{O}_{\mathcal{S}}] = \mathbb{E}_{\mathcal{M}_c}[\mathcal{R}_{\frac{d_1}{d_2}}]$$

Proof idea: The construction is the same as the one for Lem. 1 except for the cost functions. The cost functions are simply copied from the component referring to the automaton, e.g., given a state $(s, q, m) \in (S_S \times Q_A \times S_M)$ and an action $a \in A$, $d_1((s, q, m), a) = c_1(q)$ and $d_2((s, q, m), a) = c_2(q)$.

In Section 4, we show how to compute an optimal value for MDP with ratio objectives in polynomial time. Since Markov chains with ratio objectives are a special case of MDP with ratio objectives, we can first use Lem. 1 to check if $\mathcal{S} \models_{\mathcal{M}} \mathcal{L}_{\mathcal{A}}$. If the check succeeds, we then use Lem. 2 to compute the value $\text{Value}_{\mathcal{R}_{\frac{c_1}{c_2}}}^{\mathcal{M}}(\mathcal{S})$. This algorithm leads to the following theorem.

Theorem 1 (System value wrt Safety and Ratio specifications). *Given a finite-state system \mathcal{S} with alphabets L and A , an automaton \mathcal{A} with alphabet $L \times A$ defining a qualitative language, an automaton \mathcal{B} with alphabet $L \times A$ and two cost functions c_1 and c_2 defining a quantitative language, and a L -labeled MDP \mathcal{M} defining an environment model, we can compute value of \mathcal{S} with respect to $\mathcal{L}_{\mathcal{A}}$ and $\mathcal{R}_{\frac{c_1}{c_2}}$ under $\mathcal{P}_{\mathcal{M}}^{\mathcal{S}}$ in time polynomial in the maximum of $|\mathcal{S}| \cdot |\mathcal{A}| \cdot |\mathcal{M}|$ and $|\mathcal{S}| \cdot |\mathcal{B}| \cdot |c_1| \cdot |c_2| \cdot |\mathcal{M}|$.*

In order to construct an optimal system with respect to a Safety and a Ratio specification, we modify Lem. 1 and Lem. 2 slightly and construct an MDP instead of two Markov chains.

Lemma 3. *Given an automaton \mathcal{A} with alphabet $L \times A$ defining a qualitative language, an automaton \mathcal{B} with alphabet $L \times A$ and two cost functions c_1 and c_2 defining a quantitative language, and a L -labeled MDP \mathcal{M} defining an environment model, we can construct an \mathcal{M}' and two cost functions d_1 and d_2 such that for all systems \mathcal{S} ,*

$$\text{Value}_{\mathcal{L}_{\mathcal{A}} \mathcal{R}_{\frac{c_1}{c_2}}}^{\mathcal{M}}(\mathcal{S}) = \text{Value}_{\mathcal{R}_{\frac{d_1}{d_2}}}^{\mathcal{M}'}(\mathcal{S}) = \mathbb{E}_{\mathcal{M}'}^{\mathcal{S}}[\mathcal{R}_{\frac{d_1}{d_2}}].$$

Theorem 2 (Optimal System wrt Safety and Ratio specifications). *Given an automaton \mathcal{A} with alphabet $L \times A$ defining a qualitative language, an automaton \mathcal{B} with alphabet $L \times A$ and two cost functions c_1 and c_2 defining a quantitative language, and a L -labeled MDP \mathcal{M} defining an environment model, we can compute value of \mathcal{S} with respect to $\mathcal{L}_{\mathcal{A}}$ and $\mathcal{R}_{\frac{c_1}{c_2}}$ under $\mathcal{P}_{\mathcal{M}}^{\mathcal{S}}$ in time polynomial in $|\mathcal{S}| \cdot |\mathcal{A}| \cdot |\mathcal{B}| \cdot |c_1| \cdot |c_2| \cdot |\mathcal{M}|$.*

4 Calculating the best strategy

In this section we will first outline a proof showing that for every MDP there is a positional optimal strategy for our payoff function. To this end, we argue how the proof given by [Gim07] can be adapted to our case. After that we will show how we can calculate an optimal positional strategy.

²Note that instead of an MDP with ratio objective, we could have also set up a two-player safety game here.

4.1 Positional strategies suffice

In [Gim07], Gimbert proved that in an MDP any payoff function mapping to \mathbb{R} that is submixing and prefix independent admits optimal positional strategies. Since our payoff function \mathcal{R} may also take the value ∞ , we cannot apply the result immediately. However, since \mathcal{R} maps only to non-negative values and the set of measurable functions is closed under addition, multiplication, limit inferior and superior and division, provided that the divisor is not equal to 0, the expected value of \mathcal{R} is always defined and theory presented in [Gim07] also applies in this case. Furthermore, to adapt the proof of [Gim07] to minimize the payoff function instead of maximizing it, one only needs to inverse the used inequalities and replace max by min. What remains to show is that \mathcal{R} fulfills the following two properties.

Lemma 4 (\mathcal{R} is submixing and prefix independent). *Let $\mathcal{M} = (S, A, \tilde{A}, p)$ be a MDP and ρ be a run.*

1. *For every $i \geq 0$ the prefix of ρ up to i does not matter, i.e., $\mathcal{R}(\rho) = \mathcal{R}(\rho_i \rho_{i+1} \dots)$.*
2. *For every sequence of non-empty words $u_0, v_0, u_1, v_1 \dots \in (A \times S)^+$ such that $\rho = u_0 v_0 u_1 v_1 \dots$ we have that the payoff of the sequence is greater than or equal to the minimal payoff of sequences $u_0 u_1 \dots$ and $v_0 v_1 \dots$, i.e., $\mathcal{R}(\rho) \geq \min\{\mathcal{R}(u_0 u_1 \dots), \mathcal{R}(v_0 v_1 \dots)\}$.*

Proof. The first property follows immediately from the first limit in the definition of \mathcal{R} .

For the second property we partition \mathbb{N} into U and V such that U contains the indexes of the parts of ρ that belong to a u_k for some $k \in \mathbb{N}$ and such that V contains the other indexes. Formally, we define $U := \bigcup_{i \in \mathbb{N}} U_i$ where $U_0 := \{k \in \mathbb{N} \mid 0 \leq k < |u_0|\}$ and $U_i := \{\max(U_{i-1}) + |v_{i-1}| + k \mid 1 \leq k \leq |u_i|\}$. Let $V := \mathbb{N} \setminus U$ be the other indexes.

Now we look at the payoff from m to l for some $m \leq l \in \mathbb{N}$, i.e. $\mathcal{R}_m^l := (\sum_{i=m \dots l} c_1(\rho_i)) / (1 + \sum_{i=m \dots l} c_2(\rho_i))$. We can divide the sums into two parts, the one belonging to U and the one belonging to V and we get

$$\mathcal{R}_m^l = \frac{\left(\sum_{i \in \{m \dots l\} \cap U} c_1(\rho_i) \right) + \left(\sum_{i \in \{m \dots l\} \cap V} c_1(\rho_i) \right)}{1 + \left(\sum_{i \in \{m \dots l\} \cap U} c_2(\rho_i) \right) + \left(\sum_{i \in \{m \dots l\} \cap V} c_2(\rho_i) \right)}$$

We now define the different sub-sums as $u_1 := \sum_{i \in \{m \dots l\} \cap U} c_1(\rho_i)$, $u_2 := \sum_{i \in \{m \dots l\} \cap U} c_2(\rho_i)$, $v_1 := \sum_{i \in \{m \dots l\} \cap V} c_1(\rho_i)$ and $v_2 := \sum_{i \in \{m \dots l\} \cap V} c_2(\rho_i)$. Then we receive

$$\mathcal{R}_m^l = \frac{u_1 + v_1}{1 + u_2 + v_2}$$

We will now show

$$\mathcal{R}_m^l \geq \min \left\{ \frac{u_1}{u_2 + 1}, \frac{v_1}{v_2 + 1} \right\}$$

Without loss of generality we can assume $u_1/(u_2 + 1) \geq v_1/(v_2 + 1)$, then we have to show that

$$\frac{u_1 + v_1}{1 + u_2 + v_2} \geq \frac{v_1}{v_2 + 1}.$$

This holds if and only if $(u_1 + v_1)(1 + v_2) = u_1 + v_1 + u_1 v_2 + v_1 v_2 \geq v_1(1 + u_2 + v_2) = v_1 + v_1 u_2 + v_1 v_2$ holds. By subtracting v_1 and $v_1 v_2$ from both sides we receive $u_1 + u_1 v_2 = u_1(1 + v_2) \geq u_2 v_1$. If u_2 is equal to 0 then this holds because u_1 and v_2 are greater than or equal to 0. Otherwise, this holds if and only if $u_1/u_2 \geq v_1/(1 + v_2)$ holds. This follows from the assumption $u_1/(u_2 + 1) \geq v_1/(v_2 + 1)$ because $u_1/u_2 \geq u_1/(u_2 + 1)$ holds. The claim follows because we have shown this for any pair of m and l . \square

Theorem 3 (There is always a positional optimal strategy). *For each MDP with the ratio payoff function, there is a positional optimal strategy.*

Proof. See [Gim07] \square

4.2 Reduction of MDP to a Linear Fractional Program

In this section, we show how to calculate a positional optimal strategy for an MDP with ratio objective by reducing the problem to a fractional linear programming problem. A fractional linear programming problem is similar to a linear programming problem, but the function that one wants to optimize is the fraction of two linear functions. A fractional linear programming problem can be reduced to a series of conventional linear programming problems to calculate the optimal value.

We present the reduction only for MDPs for which every positional strategy induces a irreducible Markov chain because the resulting linear program is easier to understand. A similar reduction exists for general MDPs.

Our reduction uses the fact that an MDP with a positional strategy induces a Markov chain and that the runs of a Markov chain have a special property akin to the law of large numbers, which we can use to calculate the expected value.

Definition 4 (Random variables of MCs). *Let $p^n(l)$ be the probability of being in state l at step n and let $p^*(l) := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^n p^i(l)$ be the Cesaro limit of p^n . Let further ν_l^n denote the number of visits to state l up to time n .*

We have the following lemma describing the long-run behavior of Markov Chains [Tij03, Nor03].

Lemma 5 (Expected number of visits of a state and well-behaved runs). *For every infinite run of a Markov Chain the fraction of visits to a specific state l equals $p^*(l)$ almost surely, i.e., $P(\lim_{l \rightarrow \infty} \frac{\nu_l^s}{l} = p^*(s)) = 1$. We call the set of runs that have this property well-behaved.*

When we calculate the expected payoff, we only need to consider well-behaved words as shown in the following lemma.

Lemma 6. *Let N denote the set of runs that are not well-behaved. Then*

$$\mathbb{E}[\mathcal{R}] = \int_{\Omega_M \setminus N} \mathcal{R} d\mu_M$$

Proof. The probability measure of the set of well-behaved words is 1. Hence the probability measure of the complement of this set, i.e., N , has to be 0. Sets like these are called *null sets*. A classical result says that *null sets* do not need to be considered for the Lebesgue integral. \square

For a well-behaved run, i.e., for every run that we need to consider when calculating the expected value, we can calculate the payoff in the following way.

Lemma 7 (Calculating the payoff of a well-behaved run). *Let ρ be a well-behaved run of a irreducible Markov chain. Denote by $\pi : S \rightarrow A$ the only action available at a state. Then*

$$\mathcal{R}(\rho) = \frac{\sum_{s \in S} p^*(s) c_1(s, \pi(s))}{\lim_{l \rightarrow \infty} \frac{1}{l} + \sum_{s \in S} p^*(s) c_2(s, \pi(s))}$$

Proof. By definition of \mathcal{R} we have

$$\mathcal{R}(\rho) = \lim_{m \rightarrow \infty} \liminf_{l \rightarrow \infty} \frac{\sum_{i=l}^m c_1(\rho_i)}{1 + \sum_{i=l}^m c_2(\rho_i)}$$

Since the Markov chain is irreducible, it does not matter whether we start calculating at time 0 or at a later point. Hence

$$\mathcal{R}(\rho) = \liminf_{l \rightarrow \infty} \frac{\sum_{i=0}^l c_1(\rho_i)}{1 + \sum_{i=0}^l c_2(\rho_i)}$$

We can calculate the sums in a different way: we take the sum over the states and count how often we visit one state, i.e.,

$$\frac{\sum_{i=0}^l c_1(\rho_i)}{1 + \sum_{i=0}^l c_2(\rho_i)} = \frac{\sum_{s \in S} c_1(s, \pi(s)) \nu_s^l}{1 + \sum_{s \in S} c_2(s, \pi(s)) \nu_s^l} = \frac{\sum_{s \in S} c_1(s, \pi(s)) (\nu_s^l / l)}{1/l + \sum_{s \in S} c_2(s, \pi(s)) (\nu_s^l / l)}$$

Now we take \lim instead of \liminf . We will see later that the sequence converges for $l \rightarrow \infty$ and hence \lim and \liminf have the same value. Because both sides of the fraction are finite values we can safely draw the limit into the fraction, i.e.,

$$\begin{aligned} (\dagger) \lim_{l \rightarrow \infty} \left(\frac{\sum_{s \in S} c_1(s, \pi(s)) (\nu_s^l / l)}{1/l + \sum_{s \in S} c_2(s, \pi(s)) (\nu_s^l / l)} \right) &= \frac{\lim_{l \rightarrow \infty} (\sum_{s \in S} c_1(s, \pi(s)) (\nu_s^l / l))}{\lim_{l \rightarrow \infty} (1/l + \sum_{s \in S} c_2(s, \pi(s)) (\nu_s^l / l))} \\ &= \frac{\sum_{s \in S} c_1(s, \pi(s)) \lim_{l \rightarrow \infty} (\nu_s^l / l)}{\lim_{l \rightarrow \infty} (1/l) + \sum_{s \in S} c_2(s, \pi(s)) \lim_{l \rightarrow \infty} (\nu_s^l / l)} \end{aligned}$$

Finally, by the definition of well-behaved runs we have $\lim_{l \rightarrow \infty} \frac{\nu_i^l}{l} = p^*(l)$. Hence

$$\frac{\sum_{s \in S} c_1(s, \pi(s)) \lim_{l \rightarrow \infty} (\nu_s^l / l)}{\lim_{l \rightarrow \infty} (1/l) + \sum_{s \in S} c_2(s, \pi(s)) \lim_{l \rightarrow \infty} (\nu_s^l / l)} = \frac{\sum_{s \in S} c_1(s, \pi(s)) p^*(l)}{\lim_{l \rightarrow \infty} (1/l) + \sum_{s \in S} c_2(s, \pi(s)) p^*(l)}$$

The limit diverges to ∞ if and only if the second costs are all equal to zero and at least one first cost is not. In this case the original definition of \mathcal{R} diverges and hence \mathcal{R} and the last expression are the same. Otherwise the last expression converges, hence \dagger converges, ergo \liminf and \lim of this sequence are the same. \square

Note that the previous lemma implies that the value of a well-behaved run is independent of the actual run. In other words, on the set of well-behaved runs of an irreducible Markov chain the payoff function is constant. Ergo the expected value of such a Markov chain is equal to the payoff of any of its well-behaved runs.

Theorem 4 (Expected payoff of a MDP and a strategy). *Let \mathcal{M} be a MDP and π be a strategy. Let further p^* denote the Cesaro limit of p^n of the induced Markov chain. Then*

$$\mathbb{E}_{\mathcal{M}}^{\pi}[\mathcal{R}] = \frac{\sum_{s \in S} c_1(s, \pi(s)) p^*(l)}{\lim_{l \rightarrow \infty} (1/l) + \sum_{s \in S} c_2(s, \pi(s)) p^*(l)}$$

Proof. This follows from the previous lemma and the fact that \mathcal{R} is constant on any well-behaved run. \square

Note that this means that an expected value is ∞ if and only if the second cost of every action in the Markov chain is 0 and there is at least one first cost that is not.

Using this lemma, we are now able to transform the MDP into a fractional linear program. This is done in the same way as is done for the expected average payoff case (cf. [Put94]). We define variables $x(s, a)$ for every state $s \in S$ and every available action $a \in \tilde{A}(s)$. This variable intuitively corresponds to the probability of being in state s and choosing action a at any time. Then we have, for example $p^*(s) = \sum_{a \in \tilde{A}(s)} x(s, a)$. We need to restrict this set of variables. First of all, we always have to be in some state and choose some action, i.e., the sum over all $x(s, a)$ has to be one. The second set of restrictions ensures that we have a stationary distribution, i.e., the sum of the probabilities of going out of (i.e., being in) a state is equal to the sum of the probabilities of moving into this state.

Definition 5 (Linear program for MDP). *Let \mathcal{M} be an MDP such that every Markov chain induced by any strategy is recurrent and irreducible and such that it contains at least one non-zero second cost. Then we define the following linear program for it.*

$$\text{Minimize } \frac{\sum_{s \in S} \sum_{a \in \tilde{A}(s)} x(s, a) c_1(s, a)}{\sum_{s \in S} \sum_{a \in \tilde{A}(s)} x(s, a) c_2(s, \pi(s, a))} \quad (4)$$

subject to

$$\sum_{s \in S} \sum_{a \in \tilde{A}(s)} x(s, a) = 1 \quad (5)$$

$$\sum_{a \in \tilde{A}(s)} x(s, a) = \sum_{s' \in S} \sum_{a \in \tilde{A}(s')} x(s', a) p(s', a)(s) \quad \forall s \in S \quad (6)$$

There is a one to one correspondence between positional strategies and feasible solutions to the linear program³. See [Put94] for a detailed analysis of this in the expected average reward case, and how this can

³A feasible solution is one that fulfills the linear equations that every solution is subject to.

be extended to the case of general MDPs.

Once we have calculated a solution of the linear program, we can calculate the strategy as follows.

Definition 6 (Strategy from solution of linear program). *Let $x(s, a)$ be the solutions to the linear program. Then we define the strategy as follows.*

$$\pi(s) = a \Leftrightarrow x(s, a) > 1$$

Note that this definition is well defined because for each s there is at most one action a such that $x(s, a) > 0$ because of the bijection between feasible solutions and strategies and because the optimal strategy is always positional.

4.3 From LFP to LP

Since solvers to linear fractional programs are not common and there are good free solvers to linear programs, we presented a method of converting a linear fractional program to a sequence of linear programs that calculate the solution. This algorithm is due to [IM56]. Let $f(x)$ denote the value of Eqn. 4 under variable assignment x .

Input: feasible solution x_0 , MDP \mathcal{M}

Output: Variable assignment, optimal solution

$n \leftarrow 0$

repeat

$g \leftarrow f(x_n)$

$n \leftarrow n + 1$

 Solve

$$\text{Minimize } \sum_{s \in S} \sum_{a \in \bar{A}(s)} x_n(s, a) c_s^1 - g \sum_{s \in S} \sum_{a \in \bar{A}(s)} x_n(s, a) c_s^2$$

 subject to Eqn. 5 and Eqn. 6.

until $f(x_n) \neq f(x_{n+1})$;

return $x_n, f(x_n)$

4.4 Preliminary Implementation

We have developed a tool that can handle arbitrary (finite) MDPs with ratio objectives based on the approach presented in this paper. Our tool is implemented in Scala and uses the *GNU Linear Programming Kit* to solve the resulting linear programs.

We made some initial experiments using the server-client example from Section 3. In the case of two clients we have a MDP with 24 states and 288 edges. The resulting machine has 8 states and behaves as follows: if it receives two requests at the same time, then it decides to serve Client 2 first, i.e., the client with the behavior described in Figure 2(a). If Client 2 stops requesting, then the machine serves the other client. Intuitively, this is the best strategy because answering Client 1 first would result, with a high probability, in a dropped request of Client 2. Even if there is an outstanding request for Client 1, i.e., if a grant would earn a point to the system, then the system sends an acknowledgment to Client 2, if Client 2 is requesting, even if this means a dropped request for Client 1. In other words, if Client 2 sends a request, then Client 2 gets an acknowledgment. The expected value is $1.2 = 12/10$. This means that, out of 12 requests, 10 can be served, which means 83.3%.

5 Conclusions and Future Work

We have presented a technique to automatically synthesize system that satisfy a qualitative specification and optimize a quantitative specification under a given environment model. Our technique can handle qualitative specifications given by an automaton with a set of safe states, and quantitative specifications defined by an automaton with ratio objective.

Currently, we are working on a better representation of the input specifications. In particular, we are aiming for a symbolic representation that would allow us to use a combined symbolic and explicit approach, which has shown to be very effective for MDP with long-run average objective [?]. Furthermore, we are extending the presented approach to qualitative specifications described by arbitrary ω -regular specifications.

References

- [ADMW09] Rajeev Alur, Aldric Degorre, Oded Maler, and Gera Weiss. On omega-languages defined by mean-payoff conditions. In Luca de Alfaro, editor, *FOSSACS*, volume 5504 of *Lecture Notes in Computer Science*, pages 333–347. Springer, 2009. 1
- [BBD⁺02] G. Behrmann, J. Bengtsson, A. David, K. G. Larsen, P. Pettersson, and W. Yi. UPPAAL implementation secrets. In *Formal Techniques in Real-Time and Fault Tolerant Systems*, 2002. 1
- [BCHJ09] Roderick Bloem, Krishnendu Chatterjee, Thomas A. Henzinger, and Barbara Jobstmann. Better quality in synthesis through quantitative objectives. In Ahmed Bouajjani and Oded Maler, editors, *CAV*, volume 5643 of *Lecture Notes in Computer Science*, pages 140–156. Springer, 2009. 1, 3
- [BdA95] Andrea Bianco and Luca de Alfaro. Model checking of probabilistic and nondeterministic systems. In P. S. Thiagarajan, editor, *FSTTCS*, volume 1026 of *Lecture Notes in Computer Science*, pages 499–513. Springer, 1995. 1
- [BGHJ09] Roderick Bloem, Karin Greimel, Thomas A. Henzinger, and Barbara Jobstmann. Synthesizing robust systems. In *FMCAD*, pages 85–92. IEEE, 2009. 1, 1, 2
- [BGL⁺04] C. Baier, M. Größer, M. Leucker, B. Bollig, and F. Ciesinski. Controller synthesis for probabilistic systems. In *IFIP TCS*, pages 493–506, 2004. 1
- [BHHK10] Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Performance evaluation and model checking join forces. *Commun. ACM*, 53(9):76–85, 2010. 1
- [CCH⁺05] Arindam Chakrabarti, Krishnendu Chatterjee, Thomas A. Henzinger, Orna Kupferman, and Rupak Majumdar. Verifying quantitative properties using bound functions. In Dominique Borriore and Wolfgang J. Paul, editors, *CHARME*, volume 3725 of *Lecture Notes in Computer Science*, pages 50–64. Springer, 2005. 1
- [CdAF⁺06] Krishnendu Chatterjee, Luca de Alfaro, Marco Faella, Thomas A. Henzinger, Rupak Majumdar, and Mariëlle Stoelinga. Compositional quantitative reasoning. In *QEST*, pages 179–188. IEEE Computer Society, 2006. 1
- [CdAHS03] Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In Rajeev Alur and Insup Lee, editors, *EMSOFT*, volume 2855 of *Lecture Notes in Computer Science*, pages 117–133. Springer, 2003. 1
- [CDH08] Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. In Michael Kaminski and Simone Martini, editors, *CSL*, volume 5213 of *Lecture Notes in Computer Science*, pages 385–400. Springer, 2008. 1, 2
- [CG79] R. A. Cuninghame-Green. Minimax algebra. In *Lecture Notes in Economics and Mathematical Systems*, volume 166. Springer-Verlag, 1979. 1

- [CHJS10] Krishnendu Chatterjee, Thomas A. Henzinger, Barbara Jobstmann, and Rohit Singh. Measuring and synthesizing systems in probabilistic environments. In Tayssir Touili, Byron Cook, and Paul Jackson, editors, *CAV*, volume 6174 of *Lecture Notes in Computer Science*, pages 380–395. Springer, 2010. 1, 1, 3
- [CY90] Costas Courcoubetis and Mihalis Yannakakis. Markov decision processes and regular events (extended abstract). In Mike Paterson, editor, *ICALP*, volume 443 of *Lecture Notes in Computer Science*, pages 336–349. Springer, 1990. 1
- [dA98] Luca de Alfaro. Stochastic transition systems. In Davide Sangiorgi and Robert de Simone, editors, *CONCUR*, volume 1466 of *Lecture Notes in Computer Science*, pages 423–438. Springer, 1998. 1
- [dAHM03] Luca de Alfaro, Thomas A. Henzinger, and Rupak Majumdar. Discounting the future in systems theory. In Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow, and Gerhard J. Woeginger, editors, *ICALP*, volume 2719 of *Lecture Notes in Computer Science*, pages 1022–1037. Springer, 2003. 1
- [dAMRS07] Luca de Alfaro, Rupak Majumdar, Vishwanath Raman, and Mariëlle Stoelinga. Game relations and metrics. In *LICS*, pages 99–108. IEEE Computer Society, 2007. 1
- [Der62] C. Derman. On sequential decisions and markov chains. *Management Science*, 9(1):16–24, Oct. 1962. 1
- [DG07] M. Droste and P. Gastin. Weighted automata and weighted logics. *Theoretical Computer Science*, 380:69–86, 2007. 1
- [DGJP04] Josee Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labelled markov processes. *Theor. Comput. Sci.*, 318(3):323–354, 2004. 1
- [DKR08] Manfred Droste, Werner Kuich, and George Rahonis. Multi-valued mso logics over words and trees. *Fundam. Inform.*, 84(3-4):305–327, 2008. 1
- [DKV09] M. Droste, W. Kuich, and H. Vogler. *Handbook of Weighted Automata*. Springer Publishing Company, Incorporated, 2009. 1
- [FV96] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1996. 1
- [Gim07] Hugo Gimbert. Pure stationary optimal strategies in markov decision processes. In Wolfgang Thomas and Pascal Weil, editors, *STACS*, volume 4393 of *Lecture Notes in Computer Science*, pages 200–211. Springer, 2007. 4, 4.1, 3
- [GP97] Stephane Gaubert and Max Plus. Methods and applications of (max, +) linear algebra. In Rüdiger Reischuk and Michel Morvan, editors, *STACS*, volume 1200 of *Lecture Notes in Computer Science*, pages 261–282. Springer, 1997. 1
- [HKNP06] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A tool for automatic verification of probabilistic systems. In *TACAS*, 2006. 1
- [IM56] J. R. Isbell and W. H. Marlow. Attrition games. *Naval Research Logistics Quarterly*, 3:71–94, 1956. 4.3
- [KL07] Orna Kupferman and Yoad Lustig. Lattice automata. In Byron Cook and Andreas Podelski, editors, *VMCAI*, volume 4349 of *Lecture Notes in Computer Science*, pages 199–213. Springer, 2007. 1
- [Nor03] J.R. Norris. *Markov Chains*. Cambridge University Press, 2003. 4.2
- [Put94] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience, April 1994. 1, 4.2, 4.2

- [Tij03] H. C. Tijms. *A First Course in Stochastic Models*. Chichester: Wiley, 2003. [4.2](#)
- [ZP96] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theor. Comput. Sci.*, 158(1&2):343–359, 1996. [2](#)