A Temporal Logic with Mean-Payoff Constraints

Takashi Tomita[1], Shin Hiura[2], Shigeki Hagihara[1], and Naoki Yonezaki[1]

[1] Tokyo Institute of Technology, Tokyo, Japan
{tomita,hagihara,yonezaki}@fmx.cs.titech.ac.jp
[2] NS Solutions Corporation, Tokyo, Japan

Abstract. In the quantitative verification and synthesis of reactive systems, the states or transitions of a system are associated with payoffs, and a quantitative property of a behavior of the system is often characterized by the mean payoff for the behavior. This paper proposes an extension of LTL that describes mean-payoff constraints. For each step of a behavior of a system, the payment depends on a system transition and a temporal property of the behavior. A mean-payoff constraint is a threshold condition for the limit supremum or limit infimum of the mean payoffs of a behavior. This extension allows us to describe specifications reflecting qualitative and quantitative requirements on long-run average of costs and the frequencies of satisfaction of temporal properties. Moreover, we develop an algorithm for the emptiness problems of multi-dimensional payoff automata with Büchi acceptance conditions and multi-threshold mean-payoff acceptance conditions. The emptiness problems are decided by solving linear constraint satisfaction problems, and the decision problems of our logic are reduced to the emptiness problems. Consequently, we obtain exponential-time algorithms for the model- and satisfiability-checking of the logic. Some optimization problems of the logic can also be reduced to linear programming problems.

Keywords: LTL, automata, mean payoff, formal verification, decision problems, specification optimization, linear programming.

1 Introduction

Research on the formal verification and synthesis of reactive systems has focused on the qualitative properties of behaviors (e.g., "undesirable properties never hold" and "some properties hold infinitely often"). Linear Temporal Logic [19] (LTL), which is a subset of the class of ω-regular languages (i.e., languages recognized by finite-state automata such as Büchi automata and Rabin automata), is widely used to describe such properties. For LTL specifications, several model- and realizability- [18] checkers (e.g., SPIN [21] and Acacia+ [1], respectively) have been provided.

Alternatively, as an approach for describing quantitative properties, quantitative languages [15,17,12,2,11] have recently been proposed. A quantitative language is a function that gives a value in a certain ordered range to each word.[1]

[1] It is a Boolean language if the range is Boolean.

T. Aoki and K. Tagushi (Eds.): ICFEM 2012, LNCS 7635, pp. 249–265, 2012.

In the models of these languages, a payoff (or weight/cost/reward) is associated with transitions or states. Some quantitative attributes of a system behavior (e.g., the long-run average cost and the frequency of being in unexpected states) can be characterized as certain values pertaining to the mean payoff of the behavior. In quantitative synthesis [14,7,13,10], a program or strategy is optimized for such a value in the ordered range.

Alur et al. proposed a multi-threshold mean-payoff language [2], as a tractable Boolean language for describing quantitative aspects of behaviors. This language is recognized by a payoff automaton with a *multi-threshold mean-payoff acceptance condition*. A payoff is a real vector associated with a transition of the automaton. It accepts a word over a run satisfying the mean-payoff acceptance condition, given by a Boolean combination of *threshold conditions* (i.e., inequalities relating a constant threshold and the maximum or minimum value of the interval of a certain coordinate projection of accumulation points of mean payoffs of the run). The closure property under Boolean operations and the decidability of the emptiness problem for the language have been proved in [2]. However, the languages are incompatible with ω-regular languages, and cannot capture qualitative fairness, such as "a certain property holds infinitely often". Boker et al. proposed $\mathrm{LTL}^{\mathrm{lim}}$ [8], which is an extension of LTL with *path-accumulation assertions* (mean-payoff assertions). In a manner analogous to the multi-threshold mean-payoff languages, a path-accumulation assertion $\mathsf{LimSupAvg}(v) \geq c$ (resp. $\mathsf{LimInfAvg}(v) \geq c$) is a threshold condition; i.e., an inequality relating a constant c and the limit supremum (resp., limit infimum) of mean payoffs for v, where v is a numeric variable whose value depends on the state of a system. They also presented a model-checking algorithm for $\mathrm{LTL}^{\mathrm{lim}}$ against quantitative Kripke structures (in other words, multi-dimensional weighted transition systems). In this algorithm, model-checking is modified to the emptiness problem in [2], considering the Büchi condition reflecting an LTL portion of a specification. Consequently, $\mathrm{LTL}^{\mathrm{lim}}$ allows us to check whether a system satisfies a specification which reflects both qualitative and quantitative requirements. However, mean-payoff assertions are almost meaningless for satisfiability-checking, because either a combination of assertions is inconsistent according to algebraic rules or there exists a trivial variable assignment for which the assertions are true.

This paper is aimed to develop a temporal logic that can describe both qualitative and quantitative properties, and can be used as a verifiable specification language for realizability-checking and synthesis. We propose $\mathrm{LTL}^{\mathrm{mp}}$, which is an extension of $\mathrm{LTL}^{\mathrm{lim}}$ with a payment for satisfying temporal properties. In this logic, for each step of a behavior of a system, the payoff depends not only on a system transition but also on a temporal property of the behavior. Concretely, a payment t consists of free variables $v_1, \ldots, v_n$ (for associating with the transitions of a system), *characteristic variables* $\mathbf{1}_{\varphi_1}, \ldots, \mathbf{1}_{\varphi_m}$ for formulae $\varphi_1, \ldots, \varphi_m$ in the logic (i.e., each $\mathbf{1}_{\varphi_i} = 1$ if φ_i holds at the time, and otherwise $\mathbf{1}_{\varphi_i} = 0$), and algebraic operations. The mean-payoff formula has a form $\overline{\mathsf{MP}}(t) \sim c$ ($\equiv \mathsf{LimSupAvg}(t) \sim c$) or $\underline{\mathsf{MP}}(t) \sim c$ ($\equiv \mathsf{LimInfAvg}(t) \sim c$) for a payment t and $\sim \in \{<, >, \leq, \geq\}$. $\mathrm{LTL}^{\mathrm{mp}}$ can represent the quantitative properties; e.g., "the frequency of satisfying φ is bounded below by 0.1" is represented by $\underline{\mathsf{MP}}(\mathbf{1}_{\varphi}) > 0.1$, and "the long-run average cost is bounded above by 3" is expressed by $\overline{\mathsf{MP}}(6 \cdot \mathbf{1}_{\neg on \wedge \mathbf{X} on} + 4 \cdot \mathbf{1}_{on} + 5 \cdot \mathbf{1}_{on \wedge \mathbf{X} \neg on}) < 3$ if the operating cost is 4 and additional costs for booting and shutdown are 6 and 5, respectively. In addition, we can check the satisfiability of specifications with such meaningful mean-payoff constraints that have no free variable.

We reduce the decision problems of this logic to the emptiness problems of payoff automata Büchi conditions and with multi-threshold mean-payoff conditions. This type of emptiness problem can be also decided by a part of the algorithm in [8]. However, the complexity of that algorithm is roughly estimated to be exponential with respect to the size of the state space of the automaton. Therefore, we develop an algorithm for the emptiness problems of the automata, by reducing these problems to *linear constraint satisfaction problems* (LCSPs). In terms of LCSPs, the difference between the two algorithms is explained as follows: in their algorithm, the solution region is computed explicitly for finding some solutions, whereas our algorithm captures the region implicitly via linear constraints, and then finds the solutions. With this reduction, the emptiness problem of an automaton is decidable in polynomial time for the state space of the automaton. As a result, we obtain exponential-time algorithms for the model- and satisfiability-checking of the logic.

An additional advantage of this reduction is that some optimization problems concerning $\mathrm{LTL}^{\mathrm{mp}}$ specifications can be solved via *linear programming* (LP) techniques, which are widely used and well-studied optimization methods. For example, maximization/minimization problems for the limit supremum $\overline{\mathsf{MP}}(t)$ (or limit infimum $\underline{\mathsf{MP}}(t)$) of the mean payoff for a payment t, which is subject to a specification described in $\mathrm{LTL}^{\mathrm{mp}}$, are reduced to LP problems. Consequently, we can analyze performance limitations under specifications. We conjecture that this specification optimization method can be applied to realizability-checking as well as optimal synthesis for specifications described in the logic.

Related Work. [12,2,11] introduced quantitative languages focusing on mean-payoff properties. The multi-threshold mean-payoff language [2] and $\mathrm{LTL}^{\mathrm{lim}}$ [8] have been proposed as Boolean languages for describing mean-payoff properties. A multi-threshold mean-payoff language can represent threshold mean-payoff properties and some qualitative properties. $\mathrm{LTL}^{\mathrm{lim}}$ is an LTL extension with threshold mean-payoff assertions for payoffs associated with transitions of a model. $\mathrm{LTL}^{\mathrm{lim}}$ can be used as a specification language for model-checking. However, the mean-payoff assertions are almost meaningless for satisfiability-checking. This paper introduces $\mathrm{LTL}^{\mathrm{mp}}$, which is an extension of $\mathrm{LTL}^{\mathrm{lim}}$ with payments for satisfying temporal properties. $\mathrm{LTL}^{\mathrm{mp}}$ can represent quantitative properties which are meaningful for satisfiability-checking.

In existing methods [14,7,13,10] for the quantitative synthesis, a program (resp., strategy) is synthesized from a partial program or deterministic automaton (resp., Markov decision process or game). A probabilistic environment is often assumed [14,13,10], and a synthesized program (or strategy) is optimal in the average case. The notion of probability is also introduced in quantitative verification. Probabilistic temporal logics [16,4,5] (and their reward extensions [6,3]) are often used as specification languages, and some probabilistic model-checking tools (e.g., PRISM [20]) have been provided. However, the decidability of their satisfiability problems is an open question.[2] This paper provides an optimization method of $\mathrm{LTL}^{\mathrm{mp}}$ specifications, and we conjecture that our approach to the specification optimization can be applied to optimal synthesis for temporal logic specifications in which quantitative properties are described.

Previously, we introduced a probabilistic temporal logic, with a frequency operator that can describe quantitative linear-time properties pertaining only to conditional frequencies of satisfaction of temporal properties [22]. By contrast, $\mathrm{LTL}^{\mathrm{mp}}$ is a non-probabilistic linear-time logic with mean-payoff formulae. A payment for a mean-payoff formula can be flexibly described. Therefore, the mean-payoff formulae can be used to represent linear-time properties pertaining not only to conditional frequencies, but also to other types of frequencies, such as long-run average costs. (However, the semantics of the mean-payoff formulae are incompatible with those of the frequency operator.)

Organization of the Paper. In Section 2, we introduce the syntax and semantics of $\mathrm{LTL}^{\mathrm{mp}}$, which is an extension of $\mathrm{LTL}^{\mathrm{lim}}$ with payments for satisfying temporal properties. In Section 3, we provide definitions and related notions of payoff automata that accept words over runs satisfying both Büchi conditions and multi-threshold mean-payoff conditions. In addition, we develop an algorithm for the emptiness problems of the automata, in which the problems are reduced to LCSPs. In Section 4, we show how to construct an automaton that recognizes a given $\mathrm{LTL}^{\mathrm{mp}}$ formula, and how to reduce the decision problems of $\mathrm{LTL}^{\mathrm{mp}}$ to the emptiness problems of the automata. We also show that some optimization problems of $\mathrm{LTL}^{\mathrm{mp}}$ specifications can be solved by LP methods. Our conclusions are stated in Section 5.

2 LTL with Mean-Payoff Constraints

In this section, we introduce the syntax and semantics of $\mathrm{LTL}^{\mathrm{mp}}$, which is an extension of $\mathrm{LTL}^{\mathrm{lim}}$ [8] with payments for satisfying temporal properties. In $\mathrm{LTL}^{\mathrm{lim}}$, an assertion has the form either $\mathsf{LimSupAvg}(v) \sim c$ or $\mathsf{LimInfAvg}(v) \sim c$ for a variable v associated with transitions of the system. In comparison, in $\mathrm{LTL}^{\mathrm{mp}}$, a payment for each step of a behavior of a system depends not only on a transition of the system, but also on a temporal property of the behavior. An assertion in $\mathrm{LTL}^{\mathrm{mp}}$ has the form either $\overline{\mathsf{MP}}(t) \sim c$ ($\equiv \mathsf{LimSupAvg}(t) \sim c$) or $\underline{\mathsf{MP}}(t) \sim c$ ($\equiv \mathsf{LimInfAvg}(t) \sim c$), for a payment t consisting of free variables for associating with transitions of the system, characteristic variables associated with temporal properties of the behavior, and algebraic operations.

[2] For the qualitative fragment of Probabilistic CTL [16], the satisfiability problem is decidable [9].

First, we define the syntax of LTL$^{\text{mp}}$. In the following discussion, we fix the set AP of atomic propositions.

Definition 1 (Syntax). *LTL^{mp} over a set V of free variables is defined inductively as follows:*

$$\begin{aligned} \varphi &::= p \mid \overline{\mathsf{MP}}(t) \sim c \mid \underline{\mathsf{MP}}(t) \sim c \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi\mathbf{U}\varphi \\ t &::= v \mid \mathbf{1}_\varphi \mid t + t \mid -t \mid t \cdot t \mid c \cdot t \end{aligned}$$

where $p \in AP$, $v \in V$, $\sim\ \in \{<, >, \leq, \geq\}$ and $c \in \mathbb{R}$.

The operators $\mathbf{X}$ and $\mathbf{U}$ are standard temporal operators representing "next" and "until", respectively. Intuitively, $\mathbf{X}\varphi$ means that "φ holds in the *next* step", and $\varphi_1\mathbf{U}\varphi_2$ means that "φ_2 holds eventually and φ_1 holds *until* then". A payment t consists of free variables $v_1, \ldots, v_n \in V$, characteristic variables $\mathbf{1}_{\varphi_1}, \ldots, \mathbf{1}_{\varphi_m}$ for formulae $\varphi_1, \ldots, \varphi_m$, and algebraic operators ($+, -$ and $\cdot$). The major difference between LTL$^{\text{mp}}$ and LTL$^{\text{lim}}$ is the existence of characteristic variables. A characteristic variable $\mathbf{1}_\varphi$ for a formula φ represents a payment for satisfying the property φ; i.e., $\mathbf{1}_\varphi = 1$ if φ holds at the given time, and otherwise $\mathbf{1}_\varphi = 0$. The satisfaction of φ at a given time depends on a temporal property of the present and future. In this sense, a characteristic variable is bounded. A free variable v is used for associating with transitions of a system, and an LTL$^{\text{mp}}$ formula is a *sentence* if it has no free variable. Intuitively, $\overline{\mathsf{MP}}(t)$ and $\underline{\mathsf{MP}}(t)$ give the limit supremum and limit infimum, respectively, of the mean payoff for t. The formulae $\overline{\mathsf{MP}}(t) \sim c$ and $\underline{\mathsf{MP}}(t) \sim c$ are called *mean-payoff formulae*, and are *simple* if t is constructed without characteristic variables for mean-payoff formulae.

We allow common abbreviations of normal logical symbols ($\texttt{tt} \equiv \varphi \vee \neg\varphi$ and $\texttt{ff} \equiv \neg\texttt{tt}$), and connectives ($\varphi_1 \wedge \varphi_2 \equiv \neg(\neg\varphi_1 \vee \neg\varphi_2)$, $\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$ and $\varphi_1 \leftrightarrow \varphi_2 \equiv (\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_2 \rightarrow \varphi_1)$), and standard temporal operators ($\Diamond\varphi \equiv \texttt{tt}\mathbf{U}\varphi$ and $\Box\varphi \equiv \neg\Diamond\neg\varphi$). Intuitively, $\Diamond\varphi$ (resp., $\Box\varphi$) means that "φ eventually (resp., always) holds". We also use c instead of $c \cdot \mathbf{1}_{\texttt{tt}}$, for short.

LTL$^{\text{mp}}$ can represent a combination of qualitative properties described in classical LTL and quantitative properties given by mean-payoff formulae. We present some simple examples of quantitative properties.

Example 1 (Conditional frequency). A mean-payoff formula for the payment $t = (c_1 \cdot \mathbf{1}_{\varphi_1} - c_2 \cdot \mathbf{1}_{\neg\varphi_1}) \cdot \mathbf{1}_{\varphi_2}$ can represent a property pertaining to the conditional frequency of satisfaction of φ_1 under the condition φ_2, where $c_1, c_2 > 0$. Our previous work [22] focused on the conditional frequencies of satisfying temporal properties and introduced a new binary temporal operator to describe only this type of property. For $\varphi_1 = \mathbf{X}\mathit{response}$ and $\varphi_2 = \mathit{request}$, the formula $\underline{\mathsf{MP}}(t) > 0$ means that "the occurrence frequency of *requests* is not negligible (i.e., $\underline{\mathsf{MP}}(\mathbf{1}_{\mathit{request}}) > 0$) and the limit infimum of the conditional frequency of *responding* to *requests* in the next step is greater than $\frac{c_2}{c_1+c_2}$".

Example 2 (Long-run average costs). Usually, a cost is associated with an event, which has a corresponding proposition. A property of the long-run average of event-based costs is expressed as a mean-payoff formula for a payment $t = \sum c_i \cdot \mathbf{1}_{p_i}$, where p_i is a proposition representing the occurrence of an event e_i and c_i is the cost for the event e_i. For example, $\overline{\mathsf{MP}}(t) \leq 5$ means that "the long-run average of costs obeying t is bounded above by 5". In addition, switching costs for p_i are described by the characteristic variables $\mathbf{1}_{p_i \wedge \mathbf{X} \neg p_i}$ and $\mathbf{1}_{\neg p_i \wedge \mathbf{X} p_i}$.

Next we define the semantics of $\mathrm{LTL^{mp}}$.

Definition 2 (Semantics). *For an infinite word $\sigma = a_0 a_1 \cdots \in (2^{AP})^\omega$, an LTL^{mp} formula φ over a set V of free variables, and an assignment $\alpha : V \to \mathbb{R}^\omega$, the satisfaction relation $\models$ is defined inductively as follows:*

$$\begin{aligned}
\sigma, \alpha, i \models p &\Leftrightarrow p \in a_i, \\
\sigma, \alpha, i \models \neg\varphi &\Leftrightarrow \sigma, \alpha, i \not\models \varphi, \\
\sigma, \alpha, i \models \varphi_1 \vee \varphi_2 &\Leftrightarrow \sigma, \alpha, i \models \varphi_1 \ \textit{or}\ \sigma, \alpha, i \models \varphi_2, \\
\sigma, \alpha, i \models \mathbf{X}\varphi &\Leftrightarrow \sigma, \alpha, i+1 \models \varphi, \\
\sigma, \alpha, i \models \varphi_1 \mathbf{U} \varphi_2 &\Leftrightarrow \exists j \geq i.(\sigma, \alpha, j \models \varphi_2 \ \textit{and}\ \forall k \in [i, j).\sigma, \alpha, k \models \varphi_1), \\
\sigma, \alpha, i \models \overline{\mathsf{MP}}(t) \sim c &\Leftrightarrow \limsup_{n \to \infty} \tfrac{1}{n+1} \cdot \textstyle\sum_{m=0}^{n} [\![t]\!]_\sigma^\alpha(i+m) \sim c, \\
\sigma, \alpha, i \models \underline{\mathsf{MP}}(t) \sim c &\Leftrightarrow \liminf_{n \to \infty} \tfrac{1}{n+1} \cdot \textstyle\sum_{m=0}^{n} [\![t]\!]_\sigma^\alpha(i+m) \sim c,
\end{aligned}$$

$$\begin{aligned}
[\![v]\!]_\sigma^\alpha(i) &= \alpha(v)[i] \ \textit{for}\ v \in V, & [\![\mathbf{1}_\varphi]\!]_\sigma^\alpha(i) &= \begin{cases} 1 & \textit{if}\ \sigma, \alpha, i \models \varphi, \\ 0 & \textit{otherwise}, \end{cases} \\
[\![t_1 + t_2]\!]_\sigma^\alpha(i) &= [\![t_1]\!]_\sigma^\alpha(i) + [\![t_2]\!]_\sigma^\alpha(i), & [\![-t]\!]_\sigma^\alpha(i) &= -[\![t]\!]_\sigma^\alpha(i), \\
[\![t_1 \cdot t_2]\!]_\sigma^\alpha(i) &= [\![t_1]\!]_\sigma^\alpha(i) \cdot [\![t_2]\!]_\sigma^\alpha(i), & [\![c \cdot t]\!]_\sigma^\alpha(i) &= c \cdot [\![t]\!]_\sigma^\alpha(i),
\end{aligned}$$

where, for an infinite sequence $x = x_0 x_1 \cdots \in \mathbb{R}^\omega$ of real numbers, we denote by $x[i]$ the i-th element of x.

We omit i and/or α from $\sigma, \alpha, i \models \varphi$ if $i = 0$ and/or $V = \emptyset$.

The semantics of mean-payoff formulae are expressed by the limit supremum or limit infimum, and hence, for any word and assignment, the truth-value of a mean-payoff formula is either always true or always false. In a manner analogous to $\mathrm{LTL^{lim}}$, a formula φ with a mean-payoff subformula ψ is equivalent to a formula $(\varphi[\psi/\mathtt{tt}] \wedge \psi) \vee (\varphi[\psi/\mathtt{ff}] \wedge \neg\psi)$. Furthermore, any payment over $\mathrm{LTL^{mp}}$ can be represented in the form $\sum(c_i \cdot \mathbf{1}_{\varphi_i} \cdot \prod v_{ij})$. Therefore, we can restrict the syntax of $\mathrm{LTL^{mp}}$, without loss of generality, to the form $\bigvee(\varphi_i \wedge \bigwedge \psi_{ij})$, where each φ_i is a classical LTL formula (not necessarily conjunctive), each ψ_{ij} is a simple mean-payoff formula, and each payment for ψ_{ij} is of the form $\sum(c_{ijk} \cdot \mathbf{1}_{\varphi_{ijk}} \cdot \prod v_{ijkl})$. We call such a form a *mean-payoff normal form* (MPNF). An $\mathrm{LTL^{mp}}$ formula φ with n mean-payoff formulae can be transformed, at worst, into an equivalent MPNF formula with 2^n disjuncts, where each distinct has one LTL formula φ_i ($|\varphi_i| \leq |\varphi|$) and n simple mean-payoff formulae.

3 Multi-threshold Mean-Payoff Büchi Automata

In [8], model-checking for an $\mathrm{LTL^{lim}}$ formula is modified to the emptiness problem of a multi-dimensional payoff automaton with a multi-threshold mean-payoff condition [2], considering the Büchi condition reflecting the LTL portion of the formula. In this paper, we define payoff automata with both Büchi conditions and multi-threshold mean-payoff conditions. Such automata are called *multi-threshold mean-payoff Büchi automata* (MTMPBAs). In Subsection 3.1, we introduce definitions and concepts related to the automata. The decision problems of LTL^{mp} can be reduced to the emptiness problems of the automata, and it can be solved via the part of the algorithm in [8], but with a high complexity. In Subsection 3.2, we develop an algorithm for solving the emptiness problem, using a different approach with lower complexity than that of [8].

3.1 Definitions

In this subsection, we introduce the definitions of the payoff systems and MTMPBAs, together with some concepts related to them.

A payoff system is a multi-dimensional weighted transition system. It is used as a model in quantitative verification.

Definition 3. *A d-dimensional payoff system PS is a tuple $\langle Q, \Sigma, \Delta, q_0, \mathbf{w} \rangle$, where Q is a finite set of states, Σ is a finite alphabet, $\Delta \subseteq Q \times \Sigma \times Q$ is a transition relation, $q_0 \in Q$ is an initial state, and $\mathbf{w} : \Delta \to \mathbb{R}^d$ is a weight function that maps each transition to a d-dimensional real vector. We denote by $\mathbf{w}[i]$ the i-th coordinate function of $\mathbf{w}$; i.e., $\mathbf{w}(\delta) = \langle \mathbf{w}[1](\delta), \ldots, \mathbf{w}[d](\delta) \rangle$.*

For a transition $\delta = \langle q, a, q' \rangle \in \Delta$, we denote by $pre(\delta)$ the pre-state q, $post(\delta)$ the post-state q', and $letter(\delta)$ the letter a. A finite *run* r on Δ is a finite sequence $\delta_0 \cdots \delta_n \in \Delta^*$ of transitions such that $post(\delta_i) = pre(\delta_{i+1})$ for $0 \leq i < n$. A finite *word* σ $(= word(r))$ over a finite run $r = \delta_0 \cdots \delta_n$ is a finite sequence $letter(\delta_0) \cdots letter(\delta_n) \in \Sigma^*$ of letters. A (d-dimensional) finite *trace* τ is a finite sequence of (d-dimensional) real vectors. We denote by $payoff_{\mathbf{w}}(r)$ the trace $\mathbf{w}(\delta_0) \cdots \mathbf{w}(\delta_n)$ of payoffs, and by $mp_{\mathbf{w}}(r)$ the trace $\mathbf{w}(\delta_0) \cdots (\frac{1}{n+1} \sum_{i=0}^{n} \mathbf{w}(\delta_i))$ of mean payoffs, over a finite run $r = \delta_0 \cdots \delta_n$ for a d-dimensional weighted function $\mathbf{w}$. Infinite runs, words and traces are defined in a manner analogous to the finite case. We denote by $run(\Delta)$ the set of finite or infinite runs on Δ, and by $run(PS)$ the set of infinite runs starting from the initial state q_0 and belonging to $run(\Delta)$. A finite run $r = \delta_0 \cdots \delta_n \in run(\Delta)$ is *cyclic* if $pre(\delta_0) = post(\delta_n)$. A state q is *reachable* from q' on Δ if $q = q'$ or there exists a finite run $\delta_0 \cdots \delta_n \in run(\Delta)$ such that $pre(\delta_0) = post(\delta_n)$. A subgraph $\langle Q', \Delta' \rangle$ is a *strongly connected component* (SCC) on PS if $\Delta' \subseteq \Delta \cap Q' \times \Sigma \times Q'$, and for any two states in Q', one is reachable from the other on Δ'.

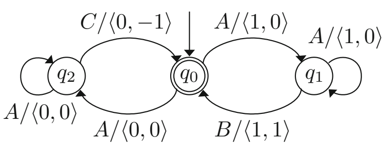

Fig. 1. Example 3

Example 3. Consider the payoff system $PS = \langle Q, 2^{\{p_0,p_1\}}, \Delta, q_0, w \rangle$, where $Q = \{q_0, q_1, q_2\}$, $A = \emptyset$, $B = \{p_0\}$, $C = \{p_1\}$, $\Delta = \{\delta_1, \ldots, \delta_6\}$, $\delta_1 = \langle q_0, A, q_1 \rangle$, $\delta_2 = \langle q_1, A, q_1 \rangle$, $\delta_3 = \langle q_1, B, q_0 \rangle$, $\delta_4 = \langle q_0, A, q_2 \rangle$, $\delta_5 = \langle q_2, A, q_2 \rangle$, $\delta_6 = \langle q_2, C, q_0 \rangle$, $\mathbf{w}[1](\delta) = 1$ if $\delta \in \{\delta_1, \delta_2, \delta_3\}$ and otherwise $\mathbf{w}[1](\delta) = 0$, and $\mathbf{w}[2](\delta_3) = 1$, $\mathbf{w}[2](\delta_6) = -1$ and $\mathbf{w}[2](\delta) = 0$ if $\delta \in \{\delta_1, \delta_2, \delta_4, \delta_5\}$ (Fig. 1). Consider runs $r_1 = (\delta_1\delta_3)^1\delta_4\delta_6(\delta_1\delta_3)^2\delta_4\delta_6(\delta_1\delta_3)^3\delta_4\delta_6\dots$ and $r_2 = (\delta_1\delta_3)(\delta_4\delta_5^{2^2-2}\delta_6)(\delta_1\delta_2^{2^3-2}\delta_3)(\delta_4\delta_5^{2^4-2}\delta_6)\dots$. Then, the trace of payoffs over r_1 is $(\langle 1,0\rangle\langle 1,1\rangle)^1\langle 0,0\rangle\langle 0,\text{-}1\rangle(\langle 1,0\rangle\langle 1,1\rangle)^2\langle 0,0\rangle\langle 0,\text{-}1\rangle(\langle 1,0\rangle\langle 1,1\rangle)^3\langle 0,0\rangle\langle 0,\text{-}1\rangle\dots$, and the trace of mean payoffs over r_1 converges to the point $\langle 1, 1/2\rangle$. The trace of payoffs over r_2 is $\langle 1,0\rangle^{2-1}\langle 1,1\rangle\langle 0,0\rangle^{2^2-1}\langle 0,\text{-}1\rangle\langle 1,0\rangle^{2^3-1}\langle 1,1\rangle\langle 0,0\rangle^{2^4-1}\langle 0,\text{-}1\rangle\dots$, and the trace of mean payoffs over r_2 has the set of accumulation points[3] $\{\langle x, 0\rangle | 1/3 \le x \le 2/3\}$.

Next, we define an MTMPBA which is a payoff system with two acceptance conditions F and G on Büchi fairness and mean payoffs, respectively. We capture a quantitative attribute of a run r via the set of accumulation points of the trace $mp_{\mathbf{w}}(r)$ of mean payoffs over r. Then, a mean-payoff acceptance condition G is given by a Boolean combination of the threshold conditions for the maximum or minimum value of the i-th projection of the set of accumulation points.

Definition 4. *An MTMPBA $\mathcal{A}$ is a tuple $\langle Q, \Sigma, \Delta, q_0, \mathbf{w}, F, G\rangle$ (or $\langle PS, F, G\rangle$ for a payoff system $PS = \langle Q, \Sigma, \Delta, q_0, \mathbf{w}\rangle$), where*

- *$F \subseteq Q$ is a* Büchi acceptance condition *given by a set of final states,*
- *$G : 2^{\mathbb{R}^d} \to$ `Bool` is a* multi-threshold mean-payoff acceptance condition *such that $G(X)$ is a Boolean combination of* threshold conditions *of the form either $\max \pi_i(X) \sim c$ or $\min \pi_i(X) \sim c$ for $\sim \in \{<, >, \le, \ge\}$, $c \in \mathbb{R}$ and the i-th projection π_i.*

The concepts of MTMPBAs are defined in a manner analogous to those of payoff systems. We denote by $Acc(\tau)$ the set of accumulation points of a trace τ. Note that, for an infinite run $r \in \text{run}(\mathcal{A})$, the maximum (resp., minimum) of the set $\pi_i(Acc(mp_{\mathbf{w}}(r)))$ is equal to the limit supremum (resp., limit infimum) of the trace $mp_{\mathbf{w}[i]}(r)$. A threshold condition is *universal* if it has the form either $\max \pi_i(\cdot) < c$, $\max \pi_i(\cdot) \le c$, $\min \pi_i(\cdot) > c$, or $\min \pi_i(\cdot) \ge c$; i.e., it asserts that "all" accumulation points of the i-th coordinate trace of mean payoffs over a run satisfy the inequality. Otherwise, it is *existential*; i.e., it asserts that "some" of the accumulation points satisfy the inequality.

An infinite run $r \in run(\mathcal{A})$ is *accepted* by $\mathcal{A}$ if both the Büchi acceptance condition F (i.e., a certain state $q \in F$ occurs infinitely often on r) and the mean-payoff acceptance condition $G(Acc(mp_{\mathbf{w}}(r)))$ hold. An infinite word $\sigma \in \Sigma^\omega$ is *accepted* by $\mathcal{A}$ if there exists a run r such that $\sigma = word(r)$ and r is accepted by $\mathcal{A}$ (i.e., $\mathcal{A}$ is an existential MTMPBA in a strict sense). A language $L \subseteq \Sigma^\omega$ (resp., an LTL^{mp} sentence φ) is *recognized* by $\mathcal{A}$ if, for all $\sigma \in \Sigma^\omega$, σ is accepted by $\mathcal{A} \Leftrightarrow \sigma \in L$ (resp., $\sigma \models \varphi$). A language recognized by an MTMPBA with $\Delta : Q \times \Sigma \to Q$ and $F = Q$ is called a multi-threshold mean-payoff language [2]. Therefore, the class of languages recognized by MTMPBAs is the superclass of ω-regular languages and of multi-threshold mean-payoff languages.

[3] A point $\mathbf{x} \in \mathbb{R}^d$ is an accumulation point of a trace $\mathbf{x}_0\mathbf{x}_1 \dots \in (\mathbb{R}^d)^\omega$ if, for any open set containing $\mathbf{x}$, there are infinitely many indices $i_1, i_2, \dots$ such that $\mathbf{x}_{i_1}, \mathbf{x}_{i_2}, \dots$ belong to the open set.

Example 4. Consider the MTMPBAs $\mathcal{A}_1 = \langle PS, \{q_0\}, \min \pi_1(\cdot) \geq 1/2 \wedge \min \pi_2(\cdot) \geq 0\rangle$ and $\mathcal{A}_2 = \langle PS, \{q_0\}, \max \pi_1(\cdot) > 1/2 \wedge \min \pi_2(\cdot) < 0\rangle$, where PS is the payoff system of Example 3. Both runs r_1 and r_2 in Example 3 satisfy the Büchi condition $\{q_0\}$. The traces of mean payoffs over r_1 and r_2 have the respective sets $\{\langle 1, 1/2\rangle\}$ and $\{\langle x, 0\rangle | 1/3 \leq x \leq 2/3\}$ of accumulation points. Thus $\mathcal{A}_1$ accepts r_1, but rejects r_2, and $\mathcal{A}_2$ rejects both r_1 and r_2.

Regarding the closure properties of the class of languages recognized by MTMPBAs, the following theorem holds. (The proof is omitted from this paper.)

Theorem 5. *The class of languages recognized by MTMPBAs is closed under union and intersection.*

3.2 Emptiness Problems

An algorithm for the emptiness problems of multi-threshold mean-payoff languages has been proposed in [2]. An algorithm for the emptiness problems of MTMPBAs has also been proposed as a part of a procedure for the model-checking of $\mathrm{LTL}^{\mathrm{lim}}$ [8], and is based on the algorithm of [2]. The decision problems of $\mathrm{LTL}^{\mathrm{mp}}$ can be reduced to the emptiness problems of MTMPBAs (see Section 4), and hence can be decided by the algorithm of [8]. However, the complexity of that algorithm is exponential with respect to the size of the state space of the automaton.

In this paper, we reduce the emptiness problems of MTMPBAs to *linear constraint satisfaction problems* (LCSPs), which can be solved by *linear programming* (LP) methods. For an MTMPBA, the existence of an accepting run can be inferred from the existence of some sets of cyclic runs. Then, the solution of each LCSP is associated with a set of cyclic runs, and a set of solutions indicates the existence of an accepting run on the automaton.

Lemma 6. *Let* $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \mathbf{w}, F, G\rangle$ *be a* d*-dimensional MTMPBA, where* $G(\cdot) = \bigwedge_{1 \leq i \leq d} \min \pi_i(\cdot) \sim_i 0$. *The following statements are equivalent.*

- *There exists an accepting run on* $\mathcal{A}$.
- *There exists a reachable (and maximal) SCC* $\langle Q', \Delta'\rangle$ *on* $\mathcal{A}$ *such that (i)* $F \cap Q' \neq \emptyset$ *and (ii) there exists a non-negative solution* $\mathbf{x}$ *for linear constraints (1)-(4), and the following conditions also hold:*
  - *(ii-a) for each existential threshold condition of the form* $\min \pi_i(\cdot) \leq 0$, *there exists a non-negative solution* $\mathbf{x}$ *for linear constraints (1)-(4) and (5),*
  - *(ii-b) for each existential threshold condition of the form* $\min \pi_i(\cdot) < 0$, *there exists a non-negative solution* $\mathbf{x}$ *for linear constraints (1)-(4) and (6),*

where $\mathbf{x}$ *is a* $|\Delta'|$*-dimensional vector,* x_δ *is the element of the vector* $\mathbf{x}$ *associated with* $\delta \in \Delta'$ *and the linear constraints are:*

$$\textstyle\sum_{\delta \in \Delta'} x_\delta \geq 1, \tag{1}$$

$$\textstyle\sum_{\delta \in \Delta' \text{ s.t. } post(\delta) = q} x_\delta = \sum_{\delta \in \Delta' \text{ s.t. } pre(\delta) = q} x_\delta \quad \text{for each } q \in Q', \tag{2}$$

$$\textstyle\sum_{\delta \in \Delta'} \mathbf{w}[j](\delta) \cdot x_\delta \geq 0 \quad \text{for each } j \text{ such that } \sim_j \text{ is } \geq, \tag{3}$$

$$\sum_{\delta\in\Delta'} \mathbf{w}[j](\delta)\cdot x_\delta \;\geq\; 1 \quad \textit{for each } j \textit{ such that } \sim_j \textit{ is } >, \tag{4}$$
$$\sum_{\delta\in\Delta'} \mathbf{w}[i](\delta)\cdot x_\delta \;\leq\; 0, \tag{5}$$
$$\sum_{\delta\in\Delta'} \mathbf{w}[i](\delta)\cdot x_\delta \;\leq\; -1. \tag{6}$$

Proof. Let n be the number of existential threshold conditions in G, and fix a reachable SCC $S = \langle Q', \Delta'\rangle$ on $\mathcal{A}$.

First, consider a solution $\mathbf{x}$ for the linear constraints (1) and (2). If $\mathbf{x}$ is an integer vector, each variable x_δ can be interpreted as the number of occurrences of the transition δ on runs. With this interpretation, $\mathbf{x}$ implies the existence of m cyclic finite runs $r_1,\ldots,r_m \in \mathit{run}(\Delta')$. This is because the linear constraint (1) implies the existence of runs with positive length, and the linear constraint (2) implies that, for each state, the number of incoming transitions is equal to the number of outgoing transitions. Here, we shall denote by $\mathit{WM}_{\mathbf{x}}(\mathbf{w})$ the weighted mean $(\sum_{\delta\in\Delta'} x_\delta\cdot\mathbf{w}(\delta))/\sum_{\delta\in\Delta'} x_\delta$ of $\mathbf{w}$ with respect to $\mathbf{x}$ (in this sense, $\mathbf{x}$ and $\mathbf{w}$ are "weight" and "data" vectors, respectively). If $m = 1$, there exists a trivial run $r_0(r_1)^\omega \in \mathit{run}(\mathcal{A})$, since S is reachable. The trace $\mathit{mp}_{\mathbf{w}}(r_0(r_1)^\omega)$ of mean payoffs over this run converges on $\mathit{WM}_{\mathbf{x}}(\mathbf{w})$. It is equal to the mean payoff of r_1, and is independent of the prefix r_0. Otherwise, there exists a larger cyclic finite run of the form $r_1r_1'\cdots r_mr_m' \in \mathit{run}(\Delta')$, since S is a SCC. Then, we can obtain a run $r_0(r_1r_1'\cdots r_mr_m')((r_1)^2r_1'\cdots(r_m)^2r_m')\cdots \in \mathit{run}(\mathcal{A})$. The trace of mean payoffs over the run also converges on $\mathit{WM}_{\mathbf{x}}(\mathbf{w})$ (i.e., the mean payoffs of $r_1,\ldots,r_m$). With this type of LCSP, given a solution $\mathbf{x}$ and a constant $c \geq 1$, the scalar product $c\cdot\mathbf{x}$ is also a solution. Therefore, even if $\mathbf{x}$ is a real vector, there still exists a run in $\mathit{run}(\mathcal{A})$ such that the ratio of the occurrence of transitions on r asymptotically approaches that of $\mathbf{x}$.

Next, consider a solution $\mathbf{x}$ of the linear constraints (1), (2) and $\sum_{\delta\in\Delta'}\mathbf{w}[k](\delta)\cdot x_\delta \geq 0$ (resp., $\sum_{\delta\in\Delta'}\mathbf{w}[k](\delta)\cdot x_\delta \geq 1$, $\sum_{\delta\in\Delta'}\mathbf{w}[k](\delta)\cdot x_\delta \leq 0$, and $\sum_{\delta\in\Delta'}\mathbf{w}[k](\delta)\cdot x_\delta \leq -1$). In a manner analogous to the first case, if a solution exists, there exists a run $r \in \mathit{run}(\mathcal{A})$ such that $\min\pi_k(\mathit{Acc}(\mathit{mp}_{\mathbf{w}}(r))) \sim_k 0$ holds, where $\sim_k$ is $\geq$ (resp., $>$, $\leq$ and $<$). This is because the k-th coordinate $\mathit{WM}_{\mathbf{x}}(\mathbf{w}[k])$ of the weighted mean with respect to $\mathbf{x}$ is greater than or equal to 0 (resp., greater than 0, less than or equal to 0, and less than 0). Otherwise, there is no run satisfying the threshold condition $\min\pi_k(\cdot) \sim_k 0$ on S. Hence, there exists a solution $\mathbf{x}$ of linear constraints (1)-(4) (and either (5) for $\min\pi_i(\cdot) \leq 0$ or (6) for $\min\pi_i(\cdot) < 0$) iff there exists a run $r_{\mathbf{x}} \in \mathit{run}(\mathcal{A})$ such that $r_{\mathbf{x}}$ satisfies all universal threshold conditions in G (and either $\min\pi_i(\cdot) \leq 0$ or $\min\pi_i(\cdot) < 0$).

Accordingly, if $n = 0$, the condition (ii) holds iff there exists a run satisfying G. Otherwise, the condition (ii) holds iff there exist runs $r_{\mathbf{x}_{\theta_1}},\ldots,r_{\mathbf{x}_{\theta_n}} \in \mathit{run}(\mathcal{A})$ corresponding to solutions $\mathbf{x}_{\theta_1},\ldots,\mathbf{x}_{\theta_n}$ for existential threshold conditions $\theta_1,\ldots,\theta_n$ in G. The trace of mean payoffs over $r_{\mathbf{x}_{\theta_k}}$ converges on the point $\mathit{WM}_{\mathbf{x}_{\theta_k}}(\mathbf{w})$, and $G(\{\,\mathit{WM}_{\mathbf{x}_{\theta_1}}(\mathbf{w}),\ldots,\mathit{WM}_{\mathbf{x}_{\theta_n}}(\mathbf{w})\})$ holds. This is because each of the runs satisfies all of the universal threshold conditions in G, and each of the existential threshold conditions is satisfied at least by one of the runs. Therefore, we can construct a run such that the trace of mean payoffs over the run comes arbitrarily close to every accumulation point $\mathit{WM}_{\mathbf{x}_{\theta_k}}(\mathbf{w})$ infinitely often. Consequently, the condition (ii) holds iff there exists a run satisfying G.

In addition, if such a run exists and condition (i) holds, there exists a run such that both F and G hold [8]. Hence, there exists an accepting run on $\mathcal{A}$ iff there exists a SCC on $\mathcal{A}$ satisfying the conditions (i) and (ii). □

Note that we can assume, without loss of generality, that (a) each coordinate is referred to by just one threshold condition, since the duplication of the coordinates of a weight function $\mathbf{w}$ does not change the recognizing language, and (b) a threshold condition has the form $\min \pi_i(\cdot) \sim 0$, since any threshold condition can be represented in this form via an affine transformation of $\mathbf{w}$.

Therefore, the emptiness problems of MTMPBAs can be reduced to LCSPs.

Theorem 7. *The emptiness problem of an MTMPBA is decidable in exponential time.*

Proof. Let $\mathcal{A} = \langle Q, \Sigma, \Delta, q_0, \mathbf{w}, F, G \rangle$ be an MTMPBA, G_i the i-th disjunct of a DNF formula $\bigvee G_i$ equivalent to G, and $\mathbf{w}_i$ the affine transformation of $\mathbf{w}$ for G_i, where each coordinate is referred to by just one threshold condition in G_i, and each G_i has the form $\bigwedge \min \pi_j(\cdot) \sim 0$.

The language recognized by $\mathcal{A}$ is empty iff the language recognized by the MTMPBA $\mathcal{A}_i = \langle Q, \Sigma, \Delta, q_0, \mathbf{w}_i, F, G_i \rangle$ is empty for all G_i. For G_i and reachable (and maximal) SCC S_{ik} on $\mathcal{A}_i$, each LCSP in Lemma 6 can be solved by LP methods in polynomial time for $|S_{ik}|$ and $|G|$. We must solve $\mathcal{O}(|G|)$ LCSPs to decide whether S_{ik} satisfies the condition (i) and (ii) in Lemma 6. Therefore, the emptiness problem of $\mathcal{A}_i$ can be solved in polynomial time for $|Q|$ and $|G|$.

In general, the number of disjuncts of a negation-free DNF formula equivalent to G is exponential in $|G|$. Hence, the complexity of the emptiness problem is polynomial in $|Q|$ and exponential in $|G|$. □

The algorithm of [8] computes explicitly a counterpart to the solution region for the linear constraints (1) and (2). The complexity is linear in the number of simple cyclic finite runs on an automaton, and that number is roughly estimated to be exponential in $|Q|$ [8]. In comparison, our algorithm captures the region implicitly, and solutions can be found in polynomial time for $|Q|$.

Example 5. Consider the MTMPBA $\mathcal{A}_2' = \langle Q, 2^{\{p_0,p_1\}}, \Delta, q_0, \mathbf{w}', \{q_0\}, \min \pi_1(\cdot) < 0 \wedge \min \pi_2(\cdot) < 0 \rangle$, obtained by affine transformation of the MTMPBA $\mathcal{A}_2 = \langle Q, 2^{\{p_0,p_1\}}, \Delta, q_0, \mathbf{w}, \{q_0\}, \max \pi_1(\cdot) > 1/2 \wedge \min \pi_2(\cdot) < 0 \rangle$ of Example 4, where $\mathbf{w}'[1](\delta) = -\mathbf{w}[1](\delta) + 1/2$ and $\mathbf{w}'[2](\delta) = \mathbf{w}[2](\delta)$ for $\delta \in \Delta$. Trivially, the SCC $\langle Q, \Delta \rangle$ is reachable and maximal, and has a final state q_0. $\mathcal{A}_2'$ has two existential threshold conditions, and hence we must solve two LCSPs to decide whether $\mathcal{A}_2'$ (and also $\mathcal{A}_2$) is empty. For a non-negative vector $\langle x_{\delta_1}, \ldots, x_{\delta_6} \rangle$, the linear constraints are: (1) $\sum_{k=1}^{6} x_{\delta_k} \geq 1$, (2) $x_{\delta_1} = x_{\delta_3}$, $x_{\delta_1} + x_{\delta_4} = x_{\delta_3} + x_{\delta_6}$ and $x_{\delta_4} = x_{\delta_6}$, (6-1) $-\frac{1}{2}(x_{\delta_1} + x_{\delta_2} + x_{\delta_3}) + \frac{1}{2}(x_{\delta_4} + x_{\delta_5} + x_{\delta_6}) \leq -1$ for $\min \pi_1(\cdot) < 0$, and (6-2) $x_{\delta_3} - x_{\delta_6} \leq -1$ for $\min \pi_2(\cdot) < 0$. As a result, $\langle 1, 0, 1, 0, 0, 0 \rangle$ and $\langle 0, 0, 0, 1, 0, 1 \rangle$ (for example) turn out to be solutions for {(1), (2) and (6-1)} and {(1), (2) and (6-2)}, respectively. These vectors are indicative of the sets $\{\delta_1\delta_3\}$ and $\{\delta_4\delta_6\}$ of single cyclic run. Therefore, we can construct an accepting run on $\mathcal{A}_2$; e.g., $(\delta_1\delta_3)^{2^2}(\delta_4\delta_6)^{2^{2^2}}(\delta_1\delta_3)^{2^{2^3}}(\delta_4\delta_6)^{2^{2^4}} \ldots$.

4 Decision and Optimization Problems of LTL$^{\mathrm{mp}}$

In this section, we present algorithms for the decision and optimization problems of LTL$^{\mathrm{mp}}$.

In a manner analogous to the decision problems of classical LTL, we can reduce the decision problems of LTL$^{\mathrm{mp}}$ to the emptiness problems of automata, which recognize LTL$^{\mathrm{mp}}$ formulae. First, we show how to construct such an automaton from a given LTL$^{\mathrm{mp}}$ sentence.

Lemma 8. *For an LTL^{mp} sentence φ, there exists an MTMPBA recognizing φ.*

Proof. Consider a *future-independent* payment t of the form $\sum c_i \cdot \mathbf{1}_{\psi_i}$ (i.e., temporal operators do not appear in every ψ_i). We can easily translate t into an alphabetic weight function $\mathbf{w}_t(\delta) = \sum_{letter(\delta) \text{ satisfies } \psi_i} c_i$, for $\delta \in \Delta$. Thus, the MTMPBA $\mathcal{A} = \langle \{q_0\}, 2^{AP}, \{q_0\} \times 2^{AP} \times \{q_0\}, q_0, \mathbf{w}_t, \{q_0\}, \min \pi_1(\cdot) > c\rangle$ recognizes the formula $\underline{\mathsf{MP}}(t) > c$. For a simple mean-payoff formula for such a payment in another form, we can construct a recognizing MTMPBA in a similar way. Therefore, we can construct an MTMPBA recognizing a given LTL$^{\mathrm{mp}}$ sentence φ if φ has no future-dependent variable, since any LTL$^{\mathrm{mp}}$ formula can be represented in MPNF and Theorem 5 holds. The size of an MPNF formula equivalent to φ is at worst linear in $|\varphi|$ and exponential in the number n of mean-payoff formulae in φ. Hence, the size of the resulting automaton is at worst exponential in $|\varphi|$ and n.

Next, consider an LTL$^{\mathrm{mp}}$ sentence φ with m future-dependent characteristic variables $\mathbf{1}_{\psi_1}, \ldots, \mathbf{1}_{\psi_m}$. Then, we can obtain another LTL$^{\mathrm{mp}}$ sentence φ', which has fresh predictive propositions $p_1, \ldots, p_m$ as follows:

$$\varphi' = \varphi[\psi_1, \ldots, \psi_m / p_1, \ldots, p_m] \wedge \bigwedge_{1 \le j \le m} \Box(p_j \leftrightarrow \psi_j).$$

This sentence φ' has no future-dependent variable, and preserves the behavioral characteristics represented by φ. Therefore, we can obtain an MTMPBA $\mathcal{A}_\varphi$ that recognizes φ, by eliminating $p_1, \ldots, p_m$ from an MTMPBA $\mathcal{A}_{\varphi'}$ that recognizes φ'. The size of the resulting automaton is at worst exponential in also m. □

Example 6. Consider the following LTL$^{\mathrm{mp}}$ formulae $\varphi_1, \ldots, \varphi_4$: $\varphi_1 = \neg p_0 \wedge \neg p_1$, $\varphi_2 = \Box((\neg p_0 \vee \neg p_1) \wedge ((p_0 \vee p_1) \rightarrow \mathbf{X}(\neg p_0 \wedge \neg p_1)))$, $\varphi_3 = \Box\Diamond(p_0 \vee p_1)$ and $\varphi_4 = \underline{\mathsf{MP}}(\mathbf{1}_{\psi_1}) \ge 1/2 \wedge \underline{\mathsf{MP}}(\mathbf{1}_{p_0} - \mathbf{1}_{p_1}) \ge 0$, where $\psi_1 = (\neg p_0 \wedge \neg p_1)\mathbf{U} p_1$. The MTMPBA $\mathcal{A}_1$ in Example 4 recognizes the LTL$^{\mathrm{mp}}$ sentence $\bigwedge_{1 \le i \le 4} \varphi_i$, and $\mathcal{A}_1$ or an MTMPBA equivalent to it can easily be obtained from the sentence. Intuitively, φ_1 represents the outgoing transitions $\langle q_0, A, q_1\rangle$ and $\langle q_0, A, q_2\rangle$ from the initial state q_0 of $\mathcal{A}_1$, φ_2 represents the transition relation of $\mathcal{A}_1$, and φ_3 and φ_4 represent the Büchi and mean-payoff acceptance conditions of $\mathcal{A}_1$, respectively. The nondeterminism of the transitions $\langle q_0, A, q_1\rangle$ and $\langle q_0, A, q_2\rangle$ on $\mathcal{A}_1$ is caused by the future-dependent payment $\mathbf{1}_{\psi_1}$.[4]

[4] Even if we consider *multi-threshold mean-payoff "Rabin" automata*, there is no deterministic automaton that recognizes the sentence $\bigwedge_{1 \le i \le 4} \varphi_i$. (The proof of this fact is omitted from this paper.) However, some future-dependent payments (e.g., $\mathbf{1}_{\mathbf{X}p}$ for $p \in AP$) do not exhibit this result.

In a manner analogous to the classical LTL model-checking, the model-checking of an $\mathrm{LTL}^{\mathrm{mp}}$ formula φ against a payoff system PS can be reduced to the emptiness problem of a synchronized product of PS and an automaton recognizing $\neg\varphi$, considering the proper variable assignment. In this paper, we define the satisfaction relation between a payoff system and an $\mathrm{LTL}^{\mathrm{mp}}$ formula as follows.

Definition 9 **($PS \models \varphi$).** *Let $PS = \langle Q, 2^{AP}, \Delta, q_0, \mathbf{w}\rangle$ be a d-dimensional payoff system, and let φ be an LTL^{mp} formula over a set V of free variables, where the free variables are indexed and $|V| = d$. In addition, we assume that the i-th free variable v_i is associated with the i-th coordinate of $\mathbf{w}$; i.e., we employ the assignment $\alpha_{\mathbf{w},r}$ such that $\alpha_{\mathbf{w},r}(v_i) = payoff_{\mathbf{w}[i]}(r)$ for an infinite run r on PS.*

We define the satisfaction relation $PS \models \varphi$ by $word(r), \alpha_{\mathbf{w},r} \models \varphi$ for all $r \in run(PS)$.

A little trick is required to assign traces of payoffs over a run of PS to free variables in φ on the synchronized product. Then, we reduce the model-checking to the emptiness problem.

Theorem 10. *The model-checking for an LTL^{mp} specification against a payoff system is decidable in exponential time.*

Proof. Let $PS = \langle Q, 2^{AP}, \Delta, q_0, \mathbf{w}\rangle$ be a payoff system, φ an $\mathrm{LTL}^{\mathrm{mp}}$ formula over a set V of free variables, φ_i the i-th disjunct of an MPNF formula $\bigvee \varphi_i$ equivalent to $\neg\varphi$, ψ_{ij} the j-th simple mean-payoff formula in φ_i, $\sum(c_{ijk} \cdot \mathbf{1}_{\varphi_{ijk}} \cdot \prod v_{ijkl})$ the payment for ψ_{ij}, $\mathbf{w}[ijkl]$ a coordinate function of $\mathbf{w}$ associated with the free variable v_{ijkl}, and n the number of mean-payoff formulae in φ.

$PS \models \varphi$ iff the language L_{PS,φ_i} is empty for all φ_i, where L_{PS,φ_i} is a set of words over runs $r \in run(PS)$ such that $word(r)$ satisfies φ_i under the assignment of each trace $payoff_{\mathbf{w}[ijkl]}(r)$ of payoffs over r to the corresponding free variable v_{ijkl} in φ_i. Therefore, we construct MTMPBAs recognizing such languages, and decide whether $PS \models \varphi$ by checking the emptiness of them.

Then, we show how to construct such MTMPBAs. First, we construct an MTMPBA $\mathcal{A}_{\varphi_i'} = \langle Q_i, 2^{AP}, \Delta_i, q_{i0}, \mathbf{w}_i, F_i, G_i\rangle$ that recognize $\varphi_i' = \varphi_i[v/0$ for all free variable $v \in V]$. We assume that the j-th coordinate function $\mathbf{w}_i[j]$ of $\mathbf{w}_i$ is associated with the payment for ψ_{ij}, and predictive propositions for future-dependent characteristic variables are still annotated on the automaton. Next, we construct a synchronized product $PS \otimes \mathcal{A}_{\varphi_i} = \langle Q \times Q_i, 2^{AP}, \Delta_i', \langle q_0, q_{i0}\rangle, \mathbf{w}_i', Q \times F_i, G_i\rangle$ of PS and $\mathcal{A}_{\varphi_i}$, considering the proper variable assignment as follows:

$$\Delta_i' = \{\langle\langle q_1, q_1'\rangle, a, \langle q_2, q_2'\rangle\rangle | \langle q_1, a, q_2\rangle \in \Delta, \langle q_1', a, q_2'\rangle \in \Delta_i\},$$

$$\mathbf{w}_i'[j](\langle\langle q_1, q_1'\rangle, a, \langle q_2, q_2'\rangle\rangle) = \mathbf{w}_i[j](\langle q_1', a, q_2'\rangle) + \textstyle\sum_{a \text{ satisfies } \varphi_{ijk}}(c_{ijk} \cdot \prod \mathbf{w}[ijkl](\langle q_1, a, q_2\rangle)).$$

We use annotated predictive propositions to check whether the letter a satisfies φ_{ijk} if φ_{ijk} has temporal operators. The automaton $PS \otimes \mathcal{A}_{\varphi_i}$ recognizes L_{PS,φ_i}.

The automaton $PS \otimes \mathcal{A}_{\varphi_i}$ has $|Q| \cdot |Q_i|$ states and a conjunctive mean-payoff acceptance condition G_i, where $|Q_i| = \mathcal{O}(2^{|\varphi|})$ and $|G_i| = \mathcal{O}(n)$. The emptiness problem for $PS \otimes \mathcal{A}_{\varphi_i}$ can be solved in polynomial time for $|Q| \cdot |Q_i|$ and $|G_i|$, since G_i is conjunctive (Theorem 7). The number of disjuncts of an MPNF formula equivalent to $\neg\varphi$ is exponential in n, and hence the complexity of the model-checking is polynomial in $|Q|$ and exponential in $|\varphi|$ and n. □

Our algorithm can accomplish the model-checking of LTL$^{\rm mp}$ with much less complexity than the algorithm of [8], which is roughly estimated to be exponential in $|Q|$, doubly exponential in $|\varphi|$, and triply exponential in n.

In a manner analogous to the classical LTL satisfiability problem, we can reduce the satisfiability problem for an LTL$^{\rm mp}$ sentence φ to the non-emptiness problem of an automaton recognizing φ.

Theorem 11. *The satisfiability problem of an LTL^{mp} sentence is decidable in exponential time.*

Proof. An LTL$^{\rm mp}$ sentence φ with n mean-payoff formulae is satisfiable iff an MTMPBA that recognizes φ is not empty. By Theorem 7 and Lemma 8, the satisfiability problem is decidable in exponential time for $|\varphi|$ and n. □

We can eventually reduce the satisfiability problem of LTL$^{\rm mp}$ to LCSPs, which can be solved by LP methods. Therefore, some optimization problems of LTL$^{\rm mp}$ can be also solved by LP methods.

Theorem 12. *The maximization/minimization problem for a mean-payoff objective ($\overline{\mathsf{MP}}(t)$ or $\underline{\mathsf{MP}}(t)$), which is subject to an LTL^{mp} sentence, is solvable in exponential time.*

Proof. Let θ be an objective ($\overline{\mathsf{MP}}(t)$ or $\underline{\mathsf{MP}}(t)$), φ an LTL$^{\rm mp}$ sentence, φ_i the i-th disjunct MPNF formula $\bigvee \varphi_i$ equivalent to φ, and n the number of mean-payoff formulae in φ.

The optimal value for θ, which is subject to φ, can be obtained as the optimal value in the set of values $opt_\theta(\varphi_i)$, where $opt_\theta(\varphi_i)$ is the optimal value for θ, which is subject to φ_i.

Such $opt_\theta(\varphi_i)$ can be found by using an MTMPBA $\mathcal{A}^t_{\varphi_i}$, which recognizes φ_i and has an additional coordinate associated with t. For a disjunct φ_i, $\mathcal{A}^t_{\varphi_i}$ can be obtained by a construction similar to that used in Lemma 8. Let w_t be the weight function of this additional coordinate, and let G_i be a mean-payoff acceptance condition of $\mathcal{A}^t_{\varphi_i}$, where G_i has the form $\bigwedge \min \pi_j(\cdot) \sim 0$ and m existential threshold conditions $p_{i1}, \ldots, p_{im}$ $(m \leq |G_i| \leq n)$.

Then, $opt_\theta(\varphi_i)$ is obtained as the optimal value in a set of values $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ for reachable SCC S_{ik} on $\mathcal{A}^t_{\varphi_i}$, where S_{ik} satisfies the condition (i) and (ii) in Lemma 6 and $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ is the optimal value for θ, subject to φ_i on S_{ik}.

If S_{ik} satisfies the condition (i) and (ii), we obtain $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ as follows:

– If (a) $m = 0$ or (b) the problem is the maximization of $\overline{\mathsf{MP}}(t)$ or the minimization of $\underline{\mathsf{MP}}(t)$, $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ is the maximized/minimized weighted mean $WM_{\mathbf{x}}(w_t)$ of w_t with respect to the solution $\mathbf{x}$, where $\mathbf{x}$ is subject to all universal threshold conditions in G_i on S_{ik} (i.e., the linear constraints (1)-(4) in Lemma 6).

– Otherwise, $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ is the minimum/maximum value in the set of the maximized/minimized weighted means $WM_{\mathbf{x}_{ik1}}(w_t), \ldots, WM_{\mathbf{x}_{ikm}}(w_t)$ of w_t with respect to the solutions $\mathbf{x}_{ik1}, \ldots, \mathbf{x}_{ikm}$, where each $\mathbf{x}_{ikl}$ is subject to all universal threshold conditions in G_i and p_{il} on S_{ik} (i.e., the linear constraints (1)-(4), and either (5) or (6) depending on p_{il} in Lemma 6).

If (a) holds, $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ is trivially optimal for φ_i on S_{ik}. Otherwise, note that each p_{il} asserts that "some" accumulation points of mean payoffs over a run satisfy the inequality. Therefore, if (b) holds, $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ is obtained independently from $p_{i1}, \ldots, p_{im}$. Otherwise, $opt_\theta(\mathcal{A}^t_{\varphi_i}, S_{ik})$ is given by the minimum/maximum value in the set of the maximized/minimized weighted means $WM_{\mathbf{x}_{ik1}}(w_t), \ldots, WM_{\mathbf{x}_{ikm}}(w_t)$ for $p_{i1}, \ldots, p_{im}$. Each optimization problem for the weighted mean is a linear fractional programming problem (LFPP), which can be solved by LP methods in polynomial time for $|S_{ik}|$ and n.

The basic flow of the optimization on $\mathcal{A}^t_{\varphi_i}$ is similar to that of the emptiness problem of an MTMPBA $\mathcal{A}_{\varphi_i}$ recognizing φ_i. Instead of each LCSP in the emptiness problem, we solve the corresponding LFPP in the optimization. Hence, the complexity of the optimization problem is exponential in $|\varphi|$ and n. □

Therefore, we can analyze performance limitations under given qualitative and quantitative specifications described in LTL$^{\mathrm{mp}}$. LP and related techniques can be also applied to other optimization problems (e.g., maximization/minimization problems for the limit supremum or limit infimum of the ratio of the mean payoffs for two payments) and multi-objective optimization problems, as in [14]. We conjecture that this LP-based approach for specification optimization can effectively be applied to optimal synthesis for temporal logic specifications.

5 Conclusions and Future Work

In this paper, we introduced LTL$^{\mathrm{mp}}$, which is an extension of LTL with mean-payoff formulae. A mean-payoff formula is a threshold condition for the limit supremum or limit infimum of the mean payoffs pertaining to a given payment. This extension allows us to describe specifications that reflect qualitative and quantitative requirements on long-run average costs and frequencies of satisfying temporal properties. Moreover, we introduced multi-threshold mean-payoff Büchi automata (MTMPBAs), which are payoff automata with Büchi acceptance conditions and multi-threshold mean-payoff acceptance conditions. Then, we developed an algorithm for solving the emptiness problems of MTMPBAs, by reducing the problems to linear constraint satisfaction problems. The decision problems of the logic can be reduced to the emptiness problems, and hence we obtained exponential-time algorithms for model- and satisfiability-checking of the logic. An additional advantage of the reduction is that some optimization problems for specifications described in the logic can be solved by linear programming methods in exponential time.

Future work will be devoted to a detailed analysis of the determinizability of automata that recognize sentences described in mean-payoff extensions of LTL and to developing the realizability-checking and quantitative synthesis methods of the extensions.

References

1. Acacia+, http://lit2.ulb.ac.be/acaciaplus/
2. Alur, R., Degorre, A., Maler, O., Weiss, G.: On Omega-Languages Defined by Mean-Payoff Conditions. In: de Alfaro, L. (ed.) FOSSACS 2009. LNCS, vol. 5504, pp. 333–347. Springer, Heidelberg (2009)
3. Andova, S., Hermanns, H., Katoen, J.P.: Discrete-Time Rewards Model-Checked. In: Larsen, K.G., Niebert, P. (eds.) FORMATS 2003. LNCS, vol. 2791, pp. 88–104. Springer, Heidelberg (2004)
4. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.: Verifying Continuous Time Markov Chains. In: Alur, R., Henzinger, T.A. (eds.) CAV 1996. LNCS, vol. 1102, pp. 269–276. Springer, Heidelberg (1996)
5. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time markov chains. IEEE Transactions on Software Engineering 29(6), 524–541 (2003)
6. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.-P.: On the Logical Characterisation of Performability Properties. In: Welzl, E., Montanari, U., Rolim, J.D.P. (eds.) ICALP 2000. LNCS, vol. 1853, pp. 780–792. Springer, Heidelberg (2000)
7. 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)
8. Boker, U., Chatterjee, K., Henzinger, T., Kupferman, O.: Temporal specifications with accumulative values. In: LICS 2011, pp. 43–52 (2011)
9. Brázdil, T., Forejt, V., Kretínský, J., Kucera, A.: The satisfiability problem for probabilistic ctl. In: LICS 2008, pp. 391–402 (2008)
10. Černý, P., Chatterjee, K., Henzinger, T.A., Radhakrishna, A., Singh, R.: Quantitative Synthesis for Concurrent Programs. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 243–259. Springer, Heidelberg (2011)
11. Chatterjee, K., Doyen, L., Edelsbrunner, H., Henzinger, T.A., Rannou, P.: Mean-Payoff Automaton Expressions. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 269–283. Springer, Heidelberg (2010)
12. Chatterjee, K., Doyen, L., Henzinger, T.A.: Quantitative Languages. In: Kaminski, M., Martini, S. (eds.) CSL 2008. LNCS, vol. 5213, pp. 385–400. Springer, Heidelberg (2008)
13. Chatterjee, K., Henzinger, T.A., Jobstmann, B., Singh, R.: Measuring and Synthesizing Systems in Probabilistic Environments. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 380–395. Springer, Heidelberg (2010)
14. Chatterjee, K., Majumdar, R., Henzinger, T.A.: Markov Decision Processes with Multiple Objectives. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 325–336. Springer, Heidelberg (2006)
15. Droste, M., Gastin, P.: Weighted automata and weighted logics. Theoretical Computer Science 380(1-2), 69–86 (2007)
16. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. Formal Aspects of Computing 6(5), 512–535 (1994)

17. Kupferman, O., Lustig, Y.: Lattice Automata. In: Cook, B., Podelski, A. (eds.) VMCAI 2007. LNCS, vol. 4349, pp. 199–213. Springer, Heidelberg (2007)
18. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: POPL 1989, pp. 179–190 (1989)
19. Pnueli, A.: The temporal logic of programs. In: FOCS 1977, pp. 46–57 (1977)
20. PRISM, `http://www.prismmodelchecker.org/`
21. SPIN, `http://spinroot.com/spin/`
22. Tomita, T., Hagihara, S., Yonezaki, N.: A probabilistic temporal logic with frequency operators and its model checking. In: INFINITY 2011. EPTCS, vol. 73, pp. 79–93 (2011)