

# Weight Monitoring with Linear Temporal Logic: Complexity and Decidability

Extended Version (2014-06-30) \*

Christel Baier   Joachim Klein   Sascha Klüppelholz   Sascha Wunderlich

Institute for Theoretical Computer Science  
Technische Universität Dresden, Germany

{baier,klein,klueppel,wunder}@tcs.inf.tu-dresden.de

## Abstract

Many important performance and reliability measures can be formalized as the accumulated values of weight functions. In this paper, we introduce an extension of linear time logic including past (LTL) with new operators that impose constraints on the accumulated weight along path fragments. The fragments are characterized by regular conditions formalized by deterministic finite automata (monitor DFA). This new logic covers properties expressible by several recently proposed formalisms. We study the model-checking problem for weighted transition systems, Markov chains and Markov decision processes with rational weights. While the general problem is undecidable, we provide algorithms and sharp complexity bounds for several sublogics that arise by restricting the monitoring DFA.

**Categories and Subject Descriptors** F.4.1 [Mathematical logic and formal languages]: Mathematical logic—Temporal logic

**General Terms** Theory

**Keywords** LTL, MDP, Markov chain, transition system, accumulation, weights, rewards, model checking, finite automata, monitor

## 1. Introduction

In many application scenarios weight accumulation occurs rather naturally. One example is the total win or loss of a share at the stock market over one day. However, not only fixed periods of time are of interest. Formalizing more general time spans is often necessary, for example in between certain events, e.g., when considering the average CPU load within a specific computation phase. Many performance and reliability measures can be formalized using

automata models with weight functions for the states or transitions. Resource requirements (e.g., bandwidth, energy consumption) and other quantitative system properties (e.g., the number of service-level violations) are then formally modeled as accumulated weights of path fragments.

Various models, logics and specification formalisms for weighted structures and accumulated weights have been proposed in the literature. For the analysis of Markovian models, the traditional approaches mainly concentrate on branching time logics with cost- or reward-bounded temporal modalities and state formulas for reasoning about total expected costs or long-run averages, see e.g. [2, 4, 16, 19]. This work has mainly focused on *non-negative* weight functions, called *reward functions*. In the case of models with two or more weight functions, the properties that can be specified in such branching-time logics are mostly Boolean combinations of formulas, each of them referring to a *single* reward function. Although nested state formulas can impose constraints for different weight functions, these logics are not adequate to express, e.g., a reachability constraint with bounds for the accumulated value of two reward functions. Approaches for multi-objective reasoning in Markovian models with nondeterminism (MDPs) [11, 18, 19] focus on the task to synthesize schedulers satisfying multiple constraints on the probabilities of  $\omega$ -regular events or expected (total) accumulated rewards.

There is a recent trend to study logics and algorithms for reasoning about properties on the accumulated values of multiple weight function that might have both positive and negative values. Such weight functions appear naturally when modeling, e.g., the amount of available energy in a battery or the market trend of stocks that might vary (increase or decrease) over time. Conditions on the relation between the accumulated values of different weight functions can be useful, e.g., for the analysis of load balancing algorithms for multi-core systems that might trigger a migration in case the load of one core is two times larger than the average load of the cores. Another example is constraints on the cost/utility ratio. Ratio objectives for weighted MDPs have been studied for example in [29] where the goal is to synthesize a scheduler for a given weighted MDP that maximizes or minimizes the average ratio payoff of two integer weight functions. Quantitative objectives for efficient synthesis have also been proposed in [6, 12]. An extension of linear time and branching time logic with prefix-accumulation assertions, which sum the weights from the start of the computation, is proposed in [7]. Similarly, variants with assertions on the long-run (mean-payoff) accumulation are considered in [7, 17, 27]. Decidability results for special property types  $\psi \wedge \varphi$  where  $\psi$  is a condition on the accumulated value of a single weight function and  $\varphi$  an  $\omega$ -regular path property have been established for various types of

\* The authors are supported by the DFG through the collaborative research centre HAEC (SFB 912), the cluster of excellence cfAED (center for Advancing Electronics Dresden), the Graduiertenkolleg 1763 (QuantLA), and the DFG/NWO-project ROCKS, the ESF young researcher groups IMData (100098198) and SREX (100111037), and the EU-FP-7 grant 295261 (MEALS).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.  
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.  
<http://dx.doi.org/10.1145/2603088.2603162>

weighted game structures, e.g., for energy and mean-payoff games in [1, 9, 10].

We introduce a temporal logical framework based on linear temporal logic (LTL) with the standard future and past temporal modalities and two new operators  $\diamond^A$  and  $\heartsuit^A$ . The latter impose constraints on the accumulated weight of path fragments  $\pi$  satisfying a given regular condition formalized by a deterministic finite automaton (DFA)  $\mathcal{A}$ , possibly under side constraints formalized by nested formulas for the first and the last position of  $\pi$ . Formulas are interpreted over the paths of weighted structures. We consider finite-state Markov decision processes with multiple weight functions, briefly called WMDPs. MDPs are a standard mathematical model with nondeterministic and probabilistic choices. For linear-time specifications, the probabilistic model-checking (PMC) problem for MDPs denotes the task to compute the maximal or minimal probabilities of a specification, where extrema are taken when ranging over all resolutions of the nondeterministic choices. Weighted transition systems (WTSs) and weighted Markov chains (WMCs) can be seen as degenerated WMDPs without nondeterminism or probabilism, respectively.

Four different classes of automata are considered here. Window DFA simply impose a restriction on the length of words. They can be used to formalize, e.g., constraints on the load of a processor in the past few clock cycles or the chart-development of a stock within one day. The fixed window properties studied in [13] for non-probabilistic game structures are expressible in our logic with window DFA. More flexibility is provided by the class of acyclic DFA. They can serve as monitors for the weights accumulated along path fragments whose traces belong to a finite set of words, e.g., formalizing finitely many request-response patterns between processes. Obviously, the full class of DFA is even more expressive and can express conditions on the total weight accumulated along any regular pattern. As a special case we also consider DFA formalizing reachability conditions. These can be used to reason about the cost accumulated until a certain event occurs. The resulting logic is indeed very expressive and covers the core features of several other logics that have been studied in the literature [7, 27]. For a detailed discussion on related formalisms we refer to Section 3.4.

**Contribution.** Besides the presentation of the syntax and semantics of linear temporal logic with weight assertions, our main contribution is to study the impact of different types of weight assertions and classes of DFA on the decidability of the PMC problem and to provide sharp complexity bounds for decidable fragments.

For the class of acyclic (and window) DFA we show in Section 4 that the PMC problem for WMDPs is solvable using a reduction to the standard PMC problem for unweighted MDPs. The reduction is, however, exponential in the size of the acyclic DFA occurring in the given formula. We then study the complexity of several simple patterns of formulas with weight assertions and establish NP- or coNP-hardness for the model-checking problem in WTSs and WMCs.

The border between decidability and undecidability will be addressed in Section 5. For the full class of DFA, we immediately get undecidability of the model-checking problem by the undecidability results presented in [7] for temporal logic with assertions on the weights accumulated along prefixes of infinite paths. We strengthen this result by proving that the model-checking problem is even undecidable for propositional logic with weight assertions obtained by DFA imposing reachability conditions. This is surprising given that [7] proves decidability for Boolean combinations of the CTL-modality  $\exists\heartsuit$  and prefix-accumulation assertions. Furthermore, we discuss the impact of weight functions versus non-negative reward functions, multiple versus single weight functions and the type of weight constraints. Omitted proofs and further technical details can be found in the appendix.

## 2. Preliminaries

Throughout the paper we assume the reader is familiar with temporal logics, automata over finite and infinite words and Markovian models. We provide a brief summary of the relevant concepts for Markov decision processes. Further details can be found, e.g., in [3, 14, 16, 25].

**Markov decision processes (MDPs).** An MDP is a tuple  $\mathcal{M} = (S, Act, P, AP, L)$ , where  $S$  is a finite set of states,  $Act$  a finite set of actions,  $P : S \times Act \times S \rightarrow [0, 1]$ ,  $AP$  is a finite set of atomic propositions and  $L : S \rightarrow 2^{AP}$  a labeling function. We require that the values  $P(s, act, s')$  are rational and  $\sum_{s' \in S} P(s, act, s') \in \{0, 1\}$  for all states  $s \in S$  and actions  $act \in Act$ . The triples  $(s, act, s')$  with  $P(s, act, s') > 0$  are called *steps*. Action  $act$  is said to be *enabled* in state  $s$  if  $P(s, act, s') > 0$  for some state  $s'$ .  $Act(s)$  denotes the set of actions that are enabled in  $s \in S$ . To avoid terminal behaviors, we require that  $Act(s) \neq \emptyset$  for all states  $s$ .

A *pointed* MDP is an MDP with a distinguished initial state  $s_{init}$ . Intuitively, the computation starts in  $s_0 = s_{init}$ . If after  $i$  steps the current state is  $s_i$  then  $\mathcal{M}$  selects nondeterministically one of the enabled actions  $act_i \in Act(s_i)$ , followed by an internal probabilistic choice to move to one of the states  $s_{i+1}$  where  $P(s_i, act_i, s_{i+1})$  is positive.

Paths in an MDP  $\mathcal{M}$  can be seen as sample runs that are obtained in this way. Formally, the paths are finite or infinite sequences where states and actions alternate:

$$\zeta = s_0 act_0 s_1 act_1 \dots \in (S \times Act)^* S \cup (S \times Act)^\omega$$

with  $act_i \in Act(s_i)$  and  $P(s_i, act_i, s_{i+1}) > 0$  for all  $i$ . The trace of  $\zeta$  is obtained by ignoring the actions and taking the projection to the state labels. If  $0 \leq h \leq k$  then  $\zeta[h \dots k]$  denotes the path fragment starting in the  $(h+1)$ -st state and ending in the  $(k+1)$ -st state. Thus, if  $\zeta$  is as above then:

$$\begin{aligned} trace(\zeta) &= L(s_0) L(s_1) L(s_2) \dots \in (2^{AP})^* \cup (2^{AP})^\omega \\ \zeta[h \dots k] &= s_h act_h s_{h+1} act_{h+1} \dots act_{k-1} s_k \end{aligned}$$

In particular,  $\zeta[k]$  is the  $(k+1)$ -st state in  $\zeta$ . We write  $first(\zeta)$  to denote the first state of  $\zeta$ . If  $\pi$  is a finite path, then  $last(\pi)$  denotes its last state and  $|\pi|$ , the length of  $\pi$ , stands for the number of steps that are taken in  $\pi$ .  $IPaths$  and  $FPaths$  stand for the set of all infinite resp. finite paths.

**Schedulers and induced probability space.** Reasoning about probabilities for path properties in MDPs requires the selection of an initial state and the resolution of the nondeterministic choices between the possible transitions. The latter is formalized via *schedulers*, also called policies or adversaries, which take as input a finite path and select an action to be executed. A (deterministic) scheduler is a function  $\mathfrak{S} : FPaths \rightarrow Act$  such that  $\mathfrak{S}(\pi) \in Act(last(\pi))$  for all finite paths  $\pi$ . Given an initial state  $s$ , the behavior of  $\mathcal{M}$  under  $\mathfrak{S}$  is purely probabilistic. Standard concepts of measure and probability theory can be applied to define a sigma-algebra and a probability measure  $\text{Pr}_{\mathcal{M}, s}^{\mathfrak{S}}$  for measurable sets of the infinite paths, also called (*path*) *events* or *path properties*. For further details, we refer to standard text books such as [25].

**Weighted MDPs (WMDP).** A weight function for  $\mathcal{M}$  is a function  $wgt : S \times Act \rightarrow \mathbb{Q}$ . We extend  $wgt$  to a function that assigns to each finite path its accumulated weight.

$$wgt(s_0 act_0 s_1 act_2 \dots act_{n-1} s_n) = \sum_{j=0}^{n-1} wgt(s_j, act_j)$$

The logic introduced in Section 3 will be interpreted over *weighted* MDPs (WMDPs), i.e., tuples  $(S, Act, P, AP, L, wgt)$  consisting of an MDP and a  $d$ -tuple  $wgt = (wgt_1, \dots, wgt_d)$  of weight functions. Non-negative weight functions  $wgt : S \times Act \rightarrow \mathbb{Q}_{\geq 0}$  are called

*reward functions.*  $wgt$  is called *positive* if  $wgt(s, act) > 0$  for all  $s \in S$  and  $act \in Act(s)$ .

**Weighted Markov chains (WMC).** Markov chains (MC) can be seen as special instances of MDPs where the action set is a singleton. Thus, the behavior of MCs is purely probabilistic. The action set will be dropped when talking about Markov chains. We write MCs as tuples  $(S, P, AP, L)$  and  $P(s, s')$  for the transition probabilities. Paths in Markov chains are just state sequences and weight functions are functions of the type  $wgt : S \rightarrow \mathbb{Q}$ . Intuitively,  $wgt(s)$  stands for the costs resp. the reward earned when leaving  $s$ . Thus, the weight of finite paths is given by  $wgt(s_0 s_1 \dots s_n) = wgt(s_0) + wgt(s_1) + \dots + wgt(s_{n-1})$ . The concept of schedulers is irrelevant for MCs and we write  $Pr_{\mathcal{M}, s}$  for the probability measure induced by  $\mathcal{M}$  when  $s$  is viewed as the initial state.

**Weighted transition systems (WTS).** A WMDP where the values of the transition probabilities are 0 or 1 can be seen as a weighted transition system, briefly called WTS. We write  $s \xrightarrow{act} s'$  if  $P(s, act, s') = 1$ .

With abuse of notations, the abbreviations WMDP, WMC and WTS are often used for pointed structures.

**Deterministic finite automata (DFA).** The logic introduced in the next section will use DFA serving as monitors for the accumulated weights in a WMDP. In this context, each DFA is given by a tuple  $\mathcal{A} = (Q, \delta, q_{init}, F)$  where  $Q$  is a finite set of states,  $\delta : Q \times 2^{AP} \rightarrow Q$  a partial transition function,  $q_{init} \in Q$  the initial state and  $F \subseteq Q$  a set of final states. We write  $\mathcal{L}(\mathcal{A})$  for the accepted language.

If  $\phi$  is a propositional formula over AP then  $\mathcal{A}[\dots\phi]$  denotes the minimal DFA where  $\mathcal{L}(\mathcal{A}[\dots\phi])$  consists of all finite words  $A_1 A_2 \dots A_n$  over  $2^{AP}$  with  $A_n \models \phi$ . Similarly,  $\mathcal{A}[\phi\dots]$  and  $\mathcal{A}[\phi_1 \dots \phi_2]$  denote minimal DFA accepting precisely the words  $A_1 A_2 \dots A_n$  with  $A_1 \models \phi$  and the words  $A_1 A_2 \dots A_n$  with  $A_1 \models \phi_1$  and  $A_n \models \phi_2$ .

For  $\ell \in \mathbb{N}$ ,  $\ell \geq 1$ , let  $\mathcal{A}_{\leq \ell}$  and  $\mathcal{A}_{=\ell}$  denote minimal DFA for the languages consisting of all words over the alphabet  $2^{AP}$  of length at most  $\ell+1$  resp. of length precisely  $\ell+1$ .

### 3. LTL with monitored weight assertions

Linear temporal logic with monitored weight assertions extends standard LTL by two new modalities  $\diamond$  and  $\diamond^A$  that impose linear constraints on the accumulated weights along finite paths that meet a regular condition given by a DFA.

#### 3.1 Syntax

A signature Sig for the linear temporal logic with monitored weight constraints consists of finitely many weight symbols  $wgt_1, \dots, wgt_d$ , a finite set AP of atomic propositions and a class AUT consisting of DFA over the alphabet  $2^{AP}$ . We consider here the following classes AUT:

Window:	the class of DFA $\mathcal{A}_{=\ell}$ and $\mathcal{A}_{\leq \ell}$ for $\ell \geq 1$
Acyc:	the class of acyclic DFA
Reach:	the class of DFA of the form $\mathcal{A}[\dots\phi]$
All:	the full class of DFA

For all classes, we require additionally that all DFA  $\mathcal{A} \in \text{AUT}$  are minimal and that the accepted language  $\mathcal{L}(\mathcal{A})$  is nonempty and does not contain the empty word  $\varepsilon$ . The assumption that  $\mathcal{A}$  is minimal and  $\mathcal{L}(\mathcal{A})$  is nonempty implies that all states  $q \in Q$  can reach  $F$ . The requirement that  $\mathcal{A}$  does not accept the empty word  $\varepsilon$  is equivalent to the requirement that  $q_{init} \notin F$ . For each DFA  $\mathcal{A} \in \text{Acyc}$  its language  $\mathcal{L}(\mathcal{A})$  is finite and the length of the longest run is the length of a longest word in  $\mathcal{L}(\mathcal{A})$ . The accepted language of each DFA  $\mathcal{A}[\dots\phi] \in \text{Reach}$  can be seen as a reachability condition  $\diamond\phi$ .

A *basic weight constraint* over Sig is a constraint of the form  $\text{expr} \bowtie c$ . Here,  $\bowtie$  is one of the four comparison operators  $<$ ,  $>$ ,  $\leq$  or  $\geq$ ,  $c \in \mathbb{Q}$  and  $\text{expr}$  is a *weight expression* of the form

$$\text{expr} = \sum_{i=1}^d a_i \cdot wgt_i \quad \text{with coefficients } a_i \in \mathbb{Q}.$$

A *weight constraint* is a Boolean combination of basic weight constraints. The class WC of weight constraints is closed under negation, and so is the class of basic weight constraints as, e.g.,  $\neg(\text{expr} \leq c)$  is equivalent to  $\text{expr} > c$ .

A weight expression is called *simple* if it has the form  $wgt_i$  for some  $i \in \{1, \dots, d\}$ . Simple basic weight constraints have the form  $wgt_i \bowtie c$ . A weight constraint  $\text{constr}$  is said to be simple if all its weight expressions are simple. Obviously, for  $d=1$  all weight constraints are equivalent to simple ones.

The logic  $\text{LTL}[\diamond, \diamond^A : \text{AUT}]$  extends LTL by two new modalities  $\diamond$  and  $\diamond^A$  to formalize constraints on the accumulated weight of path fragments. For the LTL fragment we use standard temporal modalities U (until) and S (since). The previous and next operators are omitted from the basis syntax since they are derivable (see Remark 2). The abstract syntax of  $\text{LTL}[\diamond, \diamond^A : \text{AUT}]$ -formulas is defined as follows:

$$\begin{aligned} \varphi ::= & \text{tt} \mid \mathbf{a} \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \text{S} \varphi_2 \mid \varphi_1 \text{U} \varphi_2 \mid \\ & \diamond^A(\varphi_1; \text{constr}; \varphi_2) \mid \diamond^A(\varphi_1; \text{constr}; \varphi_2) \end{aligned}$$

where  $\mathbf{a} \in \text{AP}$ ,  $\mathcal{A} \in \text{AUT}$ ,  $\text{constr} \in \text{WC}$  and  $\varphi_1$  and  $\varphi_2$  are again  $\text{LTL}[\diamond, \diamond^A : \text{AUT}]$ -formulas. We refer to  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  as (*generalized*) *weight assertion*. We simply write  $\diamond^A \text{constr}$  for  $\diamond^A(\text{tt}; \text{constr}; \text{tt})$  and refer to formulas of this type as *pure weight assertions*. Formulas of the form  $\diamond^A(\text{expr} \bowtie c)$ , are called *basic weight assertions*. Generalized, pure and basic weight assertions using the  $\diamond$ -modality are defined accordingly.

While  $\diamond^A$  is a past operator,  $\diamond^A$  imposes a constraint on the future behavior. Intuitively,  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  asserts that for the current position  $k$  of a path  $\zeta$ , there is a path fragment  $\pi = \zeta[h \dots k]$  ending in the current position accepted by  $\mathcal{A}$  such that  $\pi$  satisfies  $\text{constr}$  and the precondition  $\varphi_1$  holds in the  $h$ -th position of  $\zeta$  as well as the postcondition  $\varphi_2$  in its  $k$ -th position. Similarly,  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  imposes a constraint on the weights that will be accumulated along some path fragment that is accepted by  $\mathcal{A}$ , satisfying the precondition  $\varphi_1$  and the postcondition  $\varphi_2$  in the position of their first and last state respectively. Thus,  $\mathcal{A}$  can be viewed as a monitor that observes the traces of  $\mathcal{M}$ . For example,  $\mathcal{A}$  can be used to formalize requirements on the accumulated weights between sending a request and receiving a response.

**Length of formulas.** The length of an  $\text{LTL}[\diamond, \diamond^A : \text{AUT}]$ -formula  $\varphi$  is defined as the number of occurrences of logical operators  $\neg$ ,  $\wedge$ , S and U plus the length of all generalized weight assertions that appear in  $\varphi$ . The length of  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  or  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  is the number of states in  $\mathcal{A}$  plus the sum of the lengths of the precondition  $\varphi_1$ , the postcondition  $\varphi_2$  and the constraint  $\text{constr}$ . The latter is defined as the sum of the lengths of the binary encodings of the coefficients  $a_1, \dots, a_d$  and the constant  $c$  in the basic weight constraints  $(a_1 \cdot wgt_1 + \dots + a_d \cdot wgt_d) \bowtie c$  of  $\text{constr}$ .

**Derived operators.** As usual we can derive all operators from propositional logic (disjunction  $\vee$ , implication  $\rightarrow$ , etc.). The temporal modalities  $\diamond$  (eventually),  $\square$  (always) and R (release) can be derived as in standard LTL by  $\diamond\varphi \stackrel{\text{def}}{=} \text{tt} \text{U} \varphi$ ,  $\square\varphi \stackrel{\text{def}}{=} \neg\diamond\neg\varphi$  and  $\varphi_1 \text{R} \varphi_2 \stackrel{\text{def}}{=} \neg(\neg\varphi_1 \text{U} \neg\varphi_2)$ .

$(\zeta, k) \models \mathbf{a}$	iff	$\mathbf{a} \in \mathcal{L}(\zeta[k])$	$(\zeta, k) \models \mathbf{tt}$
$(\zeta, k) \models \neg\varphi$	iff	$(\zeta, k) \not\models \varphi$	
$(\zeta, k) \models \varphi_1 \wedge \varphi_2$	iff	$(\zeta, k) \models \varphi_1$ and $(\zeta, k) \models \varphi_2$	
$(\zeta, k) \models \varphi_1 \cup \varphi_2$	iff	there exists $h \geq k$ such that $(\zeta, h) \models \varphi_2$ and $(\zeta, i) \models \varphi_1$ for $k \leq i < h$	
$(\zeta, k) \models \varphi_1 \text{ S } \varphi_2$	iff	there exists $h \leq k$ such that $(\zeta, h) \models \varphi_2$ and $(\zeta, i) \models \varphi_1$ for $k \geq i > h$	
$(\zeta, k) \models \diamond^A(\varphi_1; \text{constr}; \varphi_2)$	iff	there exists $h \leq k$ s.t. $\text{trace}(\zeta[h \dots k]) \in \mathcal{L}(\mathcal{A})$ , $\zeta[h \dots k] \models \text{constr}$ and $(\zeta, h) \models \varphi_1$ and $(\zeta, k) \models \varphi_2$	
$(\zeta, k) \models \boxplus^A(\varphi_1; \text{constr}; \varphi_2)$	iff	there exists $h \geq k$ s.t. $\text{trace}(\zeta[k \dots h]) \in \mathcal{L}(\mathcal{A})$ , $\zeta[k \dots h] \models \text{constr}$ and $(\zeta, k) \models \varphi_1$ and $(\zeta, h) \models \varphi_2$	

**Figure 1.** Semantics of  $\text{LTL}[\diamond, \boxplus : \text{AUT}]$  over infinite paths  $\zeta$  and position  $k \in \mathbb{N}$

We can also ask whether all path fragments accepted by a monitor automaton fulfill some weight constraint instead of just one. The formula

$$\boxminus^A \text{constr} \stackrel{\text{def}}{=} \neg \diamond^A \neg \text{constr}$$

states that each suffix of the current system history accepted by  $\mathcal{A}$  fulfills  $\text{constr}$ . Similarly,

$$\boxplus^A \text{constr} \stackrel{\text{def}}{=} \neg \boxminus^A \neg \text{constr}$$

asserts that each future behavior accepted by  $\mathcal{A}$  satisfies  $\text{constr}$ .

**Sublogics.** We write  $\text{PL}[\diamond : \text{AUT}]$  for propositional logic where the atoms are future pure weight assertions, i.e., formulas of  $\text{PL}[\diamond : \text{AUT}]$  are Boolean combinations of formulas of the type  $\diamond^A \text{constr}$ . Note that the analogous logic  $\text{PL}[\boxplus : \text{AUT}]$  is pointless as it is as expressive as propositional logic.  $\text{LTL}_{\text{simple}}[\diamond, \boxplus : \text{AUT}]$  denotes  $\text{LTL}[\diamond, \boxplus : \text{AUT}]$  restricted to simple weight constraints.

### 3.2 Semantics

Formulas of the logic  $\text{LTL}[\diamond, \boxplus : \text{AUT}]$  can be interpreted over structures consisting of directed graphs with a  $d$ -dimensional weight function and node-labels in AP. Here, we deal with an MDP-semantic of  $\text{LTL}[\diamond, \boxplus : \text{AUT}]$  and interpret formulas over the infinite paths of a WMDP  $\mathcal{M} = (S, \text{Act}, P, \text{AP}, \mathcal{L}, \overline{\text{wgt}})$  with  $\overline{\text{wgt}} = (\text{wgt}_1, \dots, \text{wgt}_d)$  as in Section 2. Weight expressions are evaluated over the finite paths in  $\mathcal{M}$  in the expected way. Given a basic weight constraint  $\text{expr} \bowtie c$  and a finite path  $\pi$  in  $\mathcal{M}$  we define:

$$\pi \models \text{expr} \bowtie c \quad \text{iff} \quad \llbracket \text{expr}, \pi \rrbracket \bowtie c$$

where  $\llbracket \text{expr}, \pi \rrbracket$  denotes the value of the weight expression  $\text{expr}$  when interpreting the weight symbols  $\text{wgt}_i$  with the accumulated weight of  $\pi$  under weight function  $\text{wgt}_i$ :

$$\llbracket \text{expr}, \pi \rrbracket \stackrel{\text{def}}{=} \sum_{i=1}^d a_i \cdot \text{wgt}_i(\pi) \quad \text{for} \quad \text{expr} = \sum_{i=1}^d a_i \cdot \text{wgt}_i$$

The satisfaction relation  $\models$  for finite paths and weight constraints (i.e., Boolean combination of basic weight constraints) is now defined in the obvious way. The interpretation of  $\text{LTL}[\diamond, \boxplus : \text{AUT}]$ -formulas in the WMDP  $\mathcal{M}$  is defined over pairs  $(\zeta, k)$  where  $\zeta = s_0 \text{ act}_0 s_1 \text{ act}_1 s_2 \text{ act}_2 \dots$  is an infinite path in  $\mathcal{M}$  and  $k \in \mathbb{N}$  as shown in Figure 1. Thus,  $(\zeta, k) \models \diamond^A \text{constr}$  iff there exists  $h \leq k$  with  $\text{trace}(\zeta[h \dots k]) \in \mathcal{L}(\mathcal{A})$  and  $\zeta[h \dots k] \models \text{constr}$ . Similarly,  $(\zeta, k) \models \boxminus^A \text{constr}$  iff for each  $h \leq k$  we have:  $\text{trace}(\zeta[h \dots k]) \notin \mathcal{L}(\mathcal{A})$  or  $\zeta[h \dots k] \not\models \text{constr}$ . The semantics of future pure weight assertions is analogous. To reason about the probabilities for properties specified in  $\text{LTL}[\diamond, \boxplus : \text{AUT}]$  we lift the semantics to infinite paths:

$$\zeta \models \varphi \quad \text{iff} \quad (\zeta, 0) \models \varphi$$

**Remark 1 (Distributivity)** Weight constraints can be arbitrary Boolean combinations of basic weight constraints. For disjunctive weight constraints we can rely on the distributivity law for *existential* quantification and disjunction and obtain:

$$\begin{aligned} & \diamond^A(\varphi_1; \text{constr}_1 \vee \text{constr}_2; \varphi_2) \\ & \equiv \diamond^A(\varphi_1; \text{constr}_1; \varphi_2) \vee \diamond^A(\varphi_1; \text{constr}_2; \varphi_2) \end{aligned}$$

where  $\equiv$  denotes the equivalence of formulas. By duality, we get the equivalence of the formulas  $\boxminus^A \text{constr}_1 \wedge \text{constr}_2$  and  $\boxminus^A \text{constr}_1 \wedge \boxminus^A \text{constr}_2$ . The analogous statements hold for  $\boxplus$  and  $\boxminus$ . ■

**Remark 2 (Pre-/postconditions, past vs. future)** The postcondition in past weight assertions and the precondition in future weight assertions impose a constraint on the current position. Hence:

$$\begin{aligned} \diamond^A(\varphi_1; \text{constr}; \varphi_2) & \equiv \diamond^A(\varphi_1; \text{constr}; \mathbf{tt}) \wedge \varphi_2 \\ \boxplus^A(\varphi_1; \text{constr}; \varphi_2) & \equiv \boxplus^A(\mathbf{tt}; \text{constr}; \varphi_2) \wedge \varphi_1 \end{aligned}$$

In combination with an eventually operator, the semantics of  $\diamond$  and  $\boxplus$  coincide. That is,

$$\diamond \diamond^A(\varphi_1; \text{constr}; \varphi_2) \equiv \diamond \boxplus^A(\varphi_1; \text{constr}; \varphi_2).$$

Note that for each infinite path  $\zeta$ :

$$\begin{aligned} \zeta & \models \diamond \diamond^A(\varphi_1; \text{constr}; \varphi_2) \\ \text{iff} \quad \zeta & \models \diamond \boxplus^A(\varphi_1; \text{constr}; \varphi_2) \\ \text{iff} \quad & \text{there exists } h, k \in \mathbb{N} \text{ with } h \leq k \text{ such that} \\ & (1) \quad \text{trace}(\zeta[h \dots k]) \in \mathcal{L}(\mathcal{A}) \\ & (2) \quad \zeta[h \dots k] \models \text{constr} \\ & (3) \quad (\zeta, h) \models \varphi_1 \\ & (4) \quad (\zeta, k) \models \varphi_2 \end{aligned}$$

By duality,  $\boxminus \boxminus^A \text{constr} \equiv \boxminus \boxplus^A \text{constr}$ . ■

**Remark 3 (Window weight assertions, next and previous)** Step-bounded properties can be expressed using the automata class Window. The DFA  $\mathcal{A}_{\leq \ell}$  and  $\mathcal{A}_{=\ell}$  can be seen as monitors that observe paths up to length  $\ell$  or of length precisely  $\ell$ . In what follows,  $\diamond^{\leq \ell}$  and  $\diamond^{=\ell}$  are used as brief notations for  $\diamond^{\mathcal{A}_{\leq \ell}}$  and  $\diamond^{\mathcal{A}_{=\ell}}$  respectively, and called (past) *window weight operators*. Future window weight operators are defined accordingly. The standard next and previous operators are definable using generalized window weight assertions:

$$\bigcirc \varphi \stackrel{\text{def}}{=} \diamond^{=1}(\mathbf{tt}; \text{true}; \varphi)$$

The previous operator  $\ominus \varphi$  is obtained by  $\diamond^{=1}(\varphi; \text{true}; \mathbf{tt})$ . ■

**Remark 4 (Existential vs. universal window weight assertions)** Existential and universal pure future window weight assertions agree

for precise window length, i.e.,  $\diamond^{=\ell} \text{ constr} \equiv \boxplus^{=\ell} \text{ constr}$ , as we have for all path-position pairs  $(\zeta, k)$ :

$$\begin{aligned} (\zeta, k) \models \diamond^{=\ell} \text{ constr} & \text{ iff } (\zeta, k) \models \boxplus^{=\ell} \text{ constr} \\ & \text{ iff } \zeta[k \dots k+\ell] \models \text{ constr} \end{aligned}$$

In contrast, the formulas  $\diamond^{=\ell} \text{ constr}$  and  $\boxplus^{=\ell} \text{ constr}$  are not equivalent since for  $k < \ell$  and each infinite path  $\zeta$  we have  $(\zeta, k) \models \boxplus^{=\ell} \text{ constr}$ , while  $(\zeta, k) \not\models \diamond^{=\ell} \text{ constr}$ .

However, if  $k \geq \ell$  then for each infinite path  $\zeta$ :

$$\begin{aligned} (\zeta, k) \models \diamond^{=\ell} \text{ constr} & \text{ iff } (\zeta, k) \models \boxplus^{=\ell} \text{ constr} \\ & \text{ iff } \zeta[k-\ell \dots k] \models \text{ constr} \end{aligned}$$

In combination with prefix-independent temporal modalities the effect of  $\boxplus$  and  $\diamond$  collapses, e.g.:

$$\begin{aligned} \square \diamond \diamond^{=\ell} \text{ constr} & \equiv \square \diamond \boxplus^{=\ell} \text{ constr} \\ \diamond \square \diamond^{=\ell} \text{ constr} & \equiv \diamond \square \boxplus^{=\ell} \text{ constr} \end{aligned}$$

**Remark 5 (Step- and weight-bounded until and release)** Interpreted over a structure with a single weight function  $\text{wgt}$ , the formula  $\mathbf{a} \mathbf{U}_{\boxplus c}^{\leq \ell} \mathbf{b}$  is a variant of a  $\mathbf{U} \mathbf{b}$  with step bound  $\ell$  and weight constraint  $\text{wgt} \bowtie c$ . Formally, if  $\ell \geq 1$  then  $(\zeta, k) \models \mathbf{a} \mathbf{U}_{\boxplus c}^{\leq \ell} \mathbf{b}$  iff there exists  $k \leq h \leq k + \ell$  with  $\text{wgt}(\zeta[k \dots h]) \bowtie c$  such that  $\mathbf{b} \in \mathbf{L}(\zeta[h])$  and  $\mathbf{a} \in \mathbf{L}(\zeta[i])$  for  $k \leq i < h$ . Let  $\mathcal{A}[\mathbf{a} \mathbf{U}_{\boxplus c}^{\leq \ell} \mathbf{b}] \in \text{Acyc}$  be a DFA for the language consisting of the words  $A_1 A_2 \dots A_n$  over  $2^{\text{AP}}$  such that  $n \leq \ell + 1$ ,  $\mathbf{a} \in A_i$  for  $0 \leq i < n$  and  $\mathbf{b} \in A_n$ . Then:

$$\mathbf{a} \mathbf{U}_{\boxplus c}^{\leq \ell} \mathbf{b} \equiv \diamond^{\mathcal{A}[\mathbf{a} \mathbf{U}_{\boxplus c}^{\leq \ell} \mathbf{b}]} (\text{wgt} \bowtie c)$$

and by duality we get:

$$\mathbf{a} \mathbf{R}_{\boxplus c}^{\leq \ell} \mathbf{b} \equiv \boxplus^{\mathcal{A}[\neg \mathbf{a} \mathbf{U}_{\boxplus c}^{\leq \ell} \neg \mathbf{b}]} (\neg (\text{wgt} \bowtie c)).$$

**Interpretation over WMDP, WMC, WTS.** Our main interest is in reasoning about the probabilities of  $\text{LTL}[\diamond, \diamond : \text{AUT}]$ -specifications  $\varphi$  in weighted Markovian models. If  $(\mathcal{M}, s)$  is a WMC then:

$$\text{Pr}_{\mathcal{M}, s}(\varphi) \stackrel{\text{def}}{=} \text{Pr}_{\mathcal{M}, s} \{ \zeta \in \text{IPaths} : \zeta \models \varphi \}$$

Similarly, if  $(\mathcal{M}, s)$  is a WMDP and  $\mathfrak{S}$  a scheduler for  $\mathcal{M}$  then  $\text{Pr}_{\mathcal{M}, s}^{\mathfrak{S}}(\varphi)$  denotes the probability for  $\varphi$  under scheduler  $\mathfrak{S}$  and for starting state  $s$ . As usual, we define:

$$\text{Pr}_{\mathcal{M}, s}^{\max}(\varphi) \stackrel{\text{def}}{=} \sup_{\mathfrak{S}} \text{Pr}_{\mathcal{M}, s}^{\mathfrak{S}}(\varphi), \quad \text{Pr}_{\mathcal{M}, s}^{\min}(\varphi) \stackrel{\text{def}}{=} \inf_{\mathfrak{S}} \text{Pr}_{\mathcal{M}, s}^{\mathfrak{S}}(\varphi)$$

where  $\mathfrak{S}$  ranges over all schedulers for  $\mathcal{M}$ . When interpreting  $\text{LTL}[\diamond, \diamond : \text{AUT}]$ -formulas over WTS, we use CTL-like notations such as  $s \models \exists \varphi$  to indicate the existence of an infinite path  $\zeta$  starting in state  $s$  with  $\zeta \models \varphi$ .

**Notation 6 (Model-checking problem)** To discuss the complexity and decidability we will study decision variants of the model-checking problem for  $\text{LTL}[\diamond, \diamond : \text{AUT}]$  interpreted over WMDP, WMC or WTS. For WMDP the notion model-checking problem will be used to refer to one of the four problems asking whether (a), (b), (c) or (d) holds, where

$$\begin{aligned} \text{(a)} \quad \text{Pr}_{\mathcal{M}, s_{\text{init}}}^{\max}(\varphi) > 0 & \quad \text{(c)} \quad \text{Pr}_{\mathcal{M}, s_{\text{init}}}^{\min}(\varphi) > 0 \\ \text{(b)} \quad \text{Pr}_{\mathcal{M}, s_{\text{init}}}^{\max}(\varphi) = 1 & \quad \text{(d)} \quad \text{Pr}_{\mathcal{M}, s_{\text{init}}}^{\min}(\varphi) = 1 \end{aligned}$$

(a) and (d) are trivially interreducible since

$$\text{Pr}_{\mathcal{M}, s_{\text{init}}}^{\min}(\varphi) = 1 - \text{Pr}_{\mathcal{M}, s_{\text{init}}}^{\max}(\neg \varphi).$$

The analogous statement holds for problems (b) and (c). We refer to (a) as the *positive* model-checking problem and to (b) as the *almost-sure* model-checking problem. For WMC all four problems collapse from a computational point of view since the concept of schedulers is irrelevant. For WTS, the task of the *existential* model-checking problem is to decide whether  $s_{\text{init}} \models \exists \varphi$  for a given formula  $\varphi$ , while the *universal* model-checking problem asks whether  $s_{\text{init}} \models \forall \varphi$ . ■

**Remark 7 (Integer vs. rational weights)** We introduced weight functions in MDPs as functions of the type  $\text{wgt} : S \times \text{Act} \rightarrow \mathbb{Q}$ . When using  $\text{LTL}[\diamond, \diamond : \text{AUT}]$  as a specification formalism for finite WMDPs, however, integer-valued weight functions are equally expressive. Further details are in the appendix, Remark 22. ■

**Remark 8 (Transformations of weight functions)** Using the idea of [7], each basic weight constraint  $\text{expr} \bowtie c$  where  $\text{expr} = a_1 \cdot \text{wgt}_1 + \dots + a_d \cdot \text{wgt}_d$  can be replaced with a simple basic weight constraint  $\text{wgt} \bowtie c$  where  $\text{wgt}$  is a fresh weight symbol representing a new weight function  $\text{wgt} : S \times \text{Act} \rightarrow \mathbb{Q}$  where  $\text{wgt}(s, \text{act}) = \sum_{i=1}^d a_i \text{wgt}_i(s, \text{act})$ . See appendix, Remark 23. ■

### 3.3 Examples

To illustrate the expressiveness and usefulness of our formalism we provide a number of examples from different domains. In what follows, we use some shorthand notations for weight constraints with obvious meanings. For instance,  $\text{wgt} = c$  is short for  $(\text{wgt} \leq c) \wedge (\text{wgt} \geq c)$  and  $c_1 \leq \text{wgt} \leq c_2$  means  $(\text{wgt} \geq c_1) \wedge (\text{wgt} \leq c_2)$ .

The following formula specifies global bounds on the load, measured in the number of incoming requests to a system within a given monitor  $\mathcal{A}$ , is between  $c_{\min}$  and  $c_{\max}$ .

$$\square \boxplus^{\mathcal{A}[\text{begin} \dots \text{end}]} (c_{\min} \leq \text{load} \leq c_{\max})$$

The property that whenever there is a request to the system, the utility of the system exceeds a certain threshold  $c$  within monitor  $\mathcal{A}$  is asserted by the following formula.

$$\square (\text{request} \rightarrow \boxplus^{\mathcal{A}} (\text{utility} \geq c))$$

Similarly, the formula

$$\square (\text{request} \rightarrow \diamond^{\leq \ell} (\text{utility} \geq c))$$

can be used to specify that after each request a utility value of at least  $c$  is guaranteed within  $\ell$  or fewer steps.

The following example formalizes a resilience property, in which we require that whenever in the last ten rounds of some protocol, the accumulated number of errors a certain component produced exceeded five, the component will receive no load until the replacement procedure (formalized by  $\mathcal{A}$ ) is complete.

$$\diamond^{=10} (\text{error} > 5) \rightarrow \boxplus^{\mathcal{A}} (\text{load} = 0)$$

Using two (or more) different weight functions allows, e.g., to reason about the load balancing of two (or more) subsystems.

$$\square (\diamond^{\mathcal{A}} (|\text{load}_1 - \text{load}_2| \leq c))$$

The formula above states that globally, within a given monitor  $\mathcal{A}$  the load difference must not exceed a certain threshold  $c$ .

With two weight functions one can also express properties related to the tradeoff between cost and utility. E.g., the following formula might state that whenever the system consumes a certain amount of energy  $c_e$  for processing a query then the accumulated utility exceeds some utility threshold  $c_u$ :

$$\square (\diamond^{\mathcal{A}} (\text{energy} \geq c_e) \rightarrow \diamond^{\mathcal{A}} (\text{utility} \geq c_u))$$

Nesting of formulas allows, e.g., expressing properties of the following type. For this, let  $\mathcal{A}_{\text{init}}$  formalize an initialization process and  $\mathcal{A}_{\text{work}}$  a working phase.

Then the formula:

$$\diamond^A \text{init} (\text{tt}; \text{energy} < c_e; \diamond^A \text{work} (\text{utility} \geq c_u))$$

stands for the requirements that there is an initialization process which uses not more than  $c_e$  energy and is followed by a working phase which in turn gains at least utility  $c_u$ .

Assertions on the ratio of two weight functions can be expressed using weight expressions. For example, the following formula expresses that the monitored ratio of utility and energy exceeds some threshold  $c$ :

$$\square (\boxplus^A (\frac{\text{utility}}{\text{energy}} \geq c))$$

where  $\frac{\text{util}}{\text{energy}} \geq c$  is a short form notation for the weight expression  $\text{utility} - c \cdot \text{energy} \geq 0$ .

### 3.4 Variants and related logics

**Average.** Up to now the semantics of weight expressions over finite paths is based on the accumulated weight given by the sum of all state-action pairs along the given finite path. Alternatively one might deal with the average defined by

$$\text{avg}[wgt](\pi) = wgt(\pi)/|\pi|$$

if  $|\pi| > 0$ . Let  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{AUT}]$  denote the extension of  $\text{LTL}[\diamond, \diamond : \text{AUT}]$  where basic weight constraints either have the form  $\text{expr} \bowtie c$  as before or  $\text{avgexpr} \bowtie c$  where

$$\text{avgexpr} = \sum_{i=1}^d a_i \cdot \text{avg}[wgt_i] \quad \text{with } a_i \in \mathbb{Q}$$

is an *average weight expression*. The symbol  $\text{avg}[wgt_i]$  indicates that the weight function represented by  $wgt_i$  will be interpreted by the average weight of finite paths. To ensure that the average weight of all finite paths  $\pi$  with  $\text{trace}(\pi) \in \mathcal{L}(\mathcal{A})$  is well-defined, we impose the side constraint that for all subformulas  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  and  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$ , where  $\text{constr}$  contains an average weight constraint, the DFA  $\mathcal{A}$  does not accept words of length 1. Following the idea described in [7], average weight expressions can be transformed into sum weight expressions of the form  $wgt \bowtie 0$ . This transformation is applicable for our purposes as well. Details on the transformation can be found in the appendix, Remark 24. This yields that the (probabilistic) model-checking problem for  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{AUT}]$ -formulas is reducible to the one for  $\text{LTL}[\diamond, \diamond : \text{AUT}]$ .

Indeed several authors considered logics or specific properties that are in the spirit of or even expressible in  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{AUT}]$  for some automata class AUT.

**Fixed window properties.** The (direct) fixed window properties studied in [13] for non-probabilistic weighted game structures have the form  $\diamond^{\leq \ell}(\text{avg}[wgt] \geq c)$ ,  $\square \diamond^{\leq \ell}(\text{avg}[wgt] \geq c)$  and  $\diamond \square \diamond^{\leq \ell}(\text{avg}[wgt] \geq c)$ . Thus, they are expressible in  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{Window}]$ .

**Temporal logic with prefix accumulation.** The concept of prefix-accumulation assertions as introduced by Boker et al [7] for weighted Kripke structures and branching-time and linear-time temporal logics is very much in the spirit of the logic  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{All}]$ . The differences between weighted Kripke structures and WTSs in our sense are mostly of a syntactic nature. Rephrased for our notations, *prefix-accumulation assertions* as in [7] can be defined as  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{All}]$ -formulas:

$$\text{assert}[\text{constr}] \stackrel{\text{def}}{=} \diamond^A[\text{init}\dots] \text{constr}$$

We suppose here that  $\text{init}$  is an atomic proposition that characterizes the initial state. Thus, with this side assumption,  $\text{LTL}$  with prefix-accumulation assertions as in [7] is a sublogic of  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{All}]$ .

Given a DFA  $\mathcal{A} \in \text{All}$  imposing a regular constraint that is not LTL-definable, we cannot expect to get an LTL formula with prefix-accumulation assertions that is equivalent to  $\diamond^A \text{constr}$ . However, the  $\text{LTL}[\diamond, \diamond : \text{Reach}]$ -formula  $\diamond^A[\dots \diamond]$   $\text{constr}$  is equivalent to the formula  $\diamond(\phi \wedge \text{assert}[\text{constr}])$ . Thus, e.g.,  $\text{LTL}[\diamond, \diamond : \text{Reach}]$  can be seen as a sublogic of LTL with prefix-accumulation assertions. [7] also considers a variant of prefix accumulation ‘‘controlled’’ by some regular expression. This approach, however, departs from the regular conditions imposed by the DFA  $\mathcal{A}$  in generalized or pure weight assertions. The purpose of regular conditions in controlled prefix accumulation as in [7] relies on an alternative definition of the weight of finite paths (where the weights of certain transitions can be ignored), while the operators  $\diamond^A$  and  $\diamond^A$  impose conditions on the standard weight of finite paths satisfying a given regular constraint.

**Mean-payoff, long-run averages.** Several authors studied game structures or logics with mean-payoff objectives. The latter are typically defined as requirements on the limit superior or limit inferior of the accumulated weight along the prefixes of a given infinite path. Such requirements can be formalized in  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{All}]$  by formulas of the form

$$\diamond \square \diamond^A[\text{init}\dots](\text{avgexpr} \bowtie c) \quad \text{or} \quad \diamond \square \diamond^A[\text{init}\dots](\text{avgexpr} \bowtie c)$$

Again, up to some minor syntactic differences, this yields an embedding of the extension of LTL with mean-payoff assertions of [7] into  $\text{LTL}^{\text{avg}}[\diamond, \diamond : \text{All}]$ . [27] proposes a further extension where expressions might be polynomial and might refer to so-called characteristic properties. For the latter, our logic provides no corresponding concept. However, the switch from (linear) weight expressions to polynomial weight expressions would be possible for our framework as well. The decidability results presented for  $\text{LTL}[\diamond, \diamond : \text{Acyc}]$  in Section 4 would not be affected as we just require that weight constraints can be evaluated efficiently over a given  $d$ -tuple of values. Since mean-payoff assertions are prefix-independent properties, the logic presented in [27] is incomparable to our logic concerning expressiveness. Neither [7] nor [27] considers probabilistic structures. WMDPs or weighted game structures with mean-payoff objectives have been considered by several authors, see, e.g., [8, 9]. Extensions of temporal logics with formulas for weight constraints in WMDPs are mostly restricted to branching-time logics such as PRCTL with reward-bounded until and release modalities (see Remark 5) or state conditions on the expected total reward [2, 16, 19]. An exception is the logic introduced in [17] with state formulas asserting that  $\text{Pr}_{\mathcal{M} \otimes \mathcal{A}, s, \text{init}}^{\min}(\psi \rightarrow \varphi) = 1$  where  $\psi = \square \diamond \diamond^{\leq 1}(\text{wgt}_2 > 0)$ ,  $\varphi = \diamond \square \diamond^A[\text{init}\dots](\text{wgt}_1/\text{wgt}_2 > c)$  and  $\mathcal{A}$  is a DFA (so-called experiment) that runs in parallel to the WMDP  $\mathcal{M}$ .  $\text{wgt}_1, \text{wgt}_2$  stand for reward functions in  $\mathcal{A}$  lifted to the product  $\mathcal{M} \otimes \mathcal{A}$ . Intuitively,  $\text{wgt}_2$  counts the number of successful experiments and  $\text{wgt}_1$  the total outcome of successful experiments. This particular concept of experiments is not expressible in our logic, but inspired our work.

## 4. Model checking against

### $\text{LTL}[\diamond, \diamond : \text{Acyc}]$ -specifications

We now address the probabilistic model-checking (PMC) problem for  $\text{LTL}[\diamond, \diamond : \text{Acyc}]$ , where we are given a  $\text{LTL}[\diamond, \diamond : \text{Acyc}]$ -formula  $\varphi$ , a WMDP  $\mathcal{M} = (S, \text{Act}, P, \text{AP}, L, \text{wgt})$  and a state  $s_{\text{init}} \in S$  and the task is to compute  $\text{Pr}_{\mathcal{M}, s}^{\max}(\varphi)$  or  $\text{Pr}_{\mathcal{M}, s}^{\min}(\varphi)$ .

We first present a general approach that relies on a reduction to the task of computing extremal probabilities for LTL formulas in (unweighted) MDPs (Section 4.1). This approach is computationally expensive and relies on a product construction. It inherently uses a refined powerset construction for the automata appearing in subformulas  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  or  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  to store the relevant information on the possible runs in  $\mathcal{A}$  and the weight

for the suffixes of the current history in  $\mathcal{M}$ . We then discuss in Section 4.2 the time complexity of the model-checking problem for LTL $[\diamond, \diamond : \text{Acyc}]$  and sublogics and show that no efficient algorithms can be expected that run in time polynomial in the size of the automata  $\mathcal{A}$ . Efficient model-checking algorithms for special patterns of LTL $[\diamond, \diamond : \text{Acyc}]$ -formulas are presented in Section 4.3.

#### 4.1 Reduction to the LTL-PMC problem

The goal is to provide a reduction from the LTL $[\diamond, \diamond : \text{Acyc}]$ -PMC problem to the LTL-PMC problem. Given an LTL $[\diamond, \diamond : \text{Acyc}]$ -formula  $\varphi$  and a WMDP  $\mathcal{M} = (S, \text{Act}, P, \text{AP}, \text{L}, \text{wgt})$  where  $\text{wgt} = (\text{wgt}_1, \dots, \text{wgt}_d)$ , the idea is to replace all weight assertions  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  and  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  with an until or since formula, while adding information on the possible runs in  $\mathcal{A}$  for the path fragments in  $\mathcal{M}$ . This is done by enhancing each state  $s$  with a partial function  $f$  for each occurring automaton  $\mathcal{A}$ . The function tracks all the states  $q$  the automaton  $\mathcal{A}$  can possibly be in after reading the trace of a path fragment ending in  $s$ , along with a vector  $\bar{w}$  of the accumulated weights along this fragment.

Let  $\mathcal{A}_1, \dots, \mathcal{A}_m$  be the DFA that occur as parameters of weight assertions in  $\varphi$ . Recall that  $\mathcal{A}_i$  are supposed to be minimal acyclic DFA over the alphabet  $2^{\text{AP}}$ . Let  $\ell_i$  be the number of states in a longest run in  $\mathcal{A}_i$ . Then, the length of each word accepted by  $\mathcal{A}_i$  is at most  $\ell_i - 1$ . In particular, the maximal length of a path  $\pi$  in  $\mathcal{M}$  where  $\text{trace}(\pi)$  is accepted by  $\mathcal{A}_i$  is  $\ell_i - 2$ . (Recall that the length  $|\pi|$  of a finite path  $\pi$  is the number of transitions taken in  $\pi$ . Thus,  $\text{trace}(\pi)$  consists of  $|\pi| + 1$  symbols.) We are going to construct an (unweighted) MDP

$$\widetilde{\mathcal{M}} = \text{Monitor}(\mathcal{M}, \mathcal{A}_1, \dots, \mathcal{A}_m) = (\widetilde{S}, \text{Act}, \widetilde{P}, \widetilde{\text{AP}}, \widetilde{\text{L}})$$

whose states have the form  $\tilde{s} = \langle s, f_1, \dots, f_m \rangle$  where  $s \in S$  and  $f_i$  is a partial function from  $\{0, 1, \dots, \ell_i\}$  to pairs  $(q, \bar{w})$  where  $q$  is a state in  $\mathcal{A}_i$  and  $\bar{w} \in \mathbb{Q}^d$  such that  $f_i(k) = \perp$  (undefined) for at least one  $k$ . The set of all these tuples is, of course, infinite. Below we provide the definition of the state space  $\widetilde{S}$  of  $\widetilde{\mathcal{M}}$  which ensures that  $\widetilde{\mathcal{M}}$  has only finitely many states. See Remark 9.

The actions that are enabled in state  $\tilde{s} = \langle s, f_1, \dots, f_m \rangle$  of  $\widetilde{\mathcal{M}}$  are precisely the actions in  $\text{Act}(s)$ . The transition probability function  $\widetilde{P}$  of  $\widetilde{\mathcal{M}}$  is defined as follows. Suppose that  $\text{act} \in \text{Act}(s)$ . Then:

$$\widetilde{P}(\langle s, f_1, \dots, f_m \rangle, \text{act}, \langle s', f'_1, \dots, f'_m \rangle) = P(s, \text{act}, s')$$

where  $f'_i$  is the unique  $(\text{act}, s')$ -successor of  $\langle s, f_i \rangle$  in  $\mathcal{A}_i$  that is defined as follows. Let us now fix some  $i \in \{1, \dots, m\}$  and suppose  $\mathcal{A}_i = (Q, \delta, q_{\text{init}}, F)$ . Then,  $f_i : \{0, 1, \dots, \ell_i\} \rightarrow Q \times \mathbb{Q}^d$  is a partial function such that  $f_i(k) = \perp$  for at least one  $k$ . The  $(\text{act}, s')$ -successor of  $\langle s, f_i \rangle$  in  $\mathcal{A}_i$  is the partial function  $f'_i : \{0, 1, \dots, \ell_i\} \rightarrow Q \times \mathbb{Q}^d$  where for  $0 \leq k \leq \ell_i$ :

- If  $f_i(k) = (q, \bar{w})$  where  $q \in Q$  and  $q \neq q_{\text{init}}$  then  $f'_i(k) = (\delta(q, \text{L}(s')), \bar{w} + \text{wgt}(s, \text{act}))$ .
- If  $k$  is the smallest index such that  $f_i(k) = \perp$  then  $f'_i(k) = (\delta(q_{\text{init}}, \text{L}(s')), \bar{0})$ .
- In all other cases:  $f'_i(k) = \perp$ .

Here, we identify the tuples  $(\perp, \bar{w})$  with  $\perp$ . Furthermore, we define the initial function  $f_i^s$  by  $f_i^s(0) = (\delta(q_{\text{init}}, \text{L}(s)), \bar{0})$  and  $f_i^s(k) = \perp$  for  $k \in \{1, \dots, \ell_i\}$ .

In all other cases,  $\widetilde{P}(\cdot) = 0$ . The state space  $\widetilde{S}$  of  $\widetilde{\mathcal{M}}$  is the smallest set that contains the states  $\tilde{s} \stackrel{\text{def}}{=} \langle s, f_1^s, \dots, f_m^s \rangle$  for all  $s \in S$  and that is closed under the steps induced by the transition probability function  $\widetilde{P}$ .

**Remark 9 (Size of the state space)** The set  $\widetilde{S}$  is indeed finite. Specifically, for  $\ell = \max\{\ell_1, \dots, \ell_m\}$  we have

$$|\widetilde{S}| < |S|^\ell \cdot |\text{Act}|^{\ell-1} \cdot 2^{m \cdot (\ell+1) \cdot \log(\ell+1)}.$$

This bound is obtained by the observation that  $\widetilde{S}$  can be written as the union of sets  $\widetilde{S}_\pi$ , where  $\pi$  is a path fragment of length at most  $\ell - 2$  in  $\mathcal{M}$  and all states in  $\widetilde{S}_\pi$  have the form  $\langle \text{last}(\pi), f_1, \dots, f_m \rangle$  where  $\{f_i(0), \dots, f_i(\ell_i)\} \setminus \{\perp\}$  consists of the pairs  $(\delta(q_{\text{init}}, \text{trace}(\pi')), \text{wgt}(\pi))$  for some prefix  $\pi'$  of  $\pi$ . This yields

$$|\widetilde{S}_\pi| \leq (\ell+1)! < 2^{m \cdot (\ell+1) \cdot \log(\ell+1)}.$$

The factor  $|S|^\ell \cdot |\text{Act}|^{\ell-1}$  is an upper bound for the number of path fragments of length  $\ell - 2$ . ■

The set  $\widetilde{\text{AP}}$  of atomic propositions in  $\widetilde{\mathcal{M}}$  consists of:

- the atomic propositions in AP that appear in  $\varphi$ ,
- fresh symbols  $\text{init}_i(k)$ ,  $\text{run}_i(k)$  and  $\text{goal}_i(k)$  for  $i = 1, \dots, m$  and  $k \in \{0, 1, \dots, \ell_i\}$

The labeling function  $\widetilde{\text{L}} : \widetilde{S} \rightarrow 2^{\widetilde{\text{AP}}}$  is then defined by the following conditions. Let  $\tilde{s} = \langle s, f_1, \dots, f_m \rangle \in \widetilde{S}$ . Then  $\text{AP} \cap \widetilde{\text{L}}(\tilde{s}) = \text{AP} \cap \text{L}(s)$ . For  $i \in \{1, \dots, m\}$  the semantics of  $\diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$  or  $\diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$  will be encoded using the atomic propositions  $\text{init}_i(k)$ ,  $\text{run}_i(k)$  and  $\text{goal}_i(k)$ . The requirements for the labeling function is as follows where we suppose that  $\mathcal{A}_i = (Q, \delta, q_{\text{init}}, F)$ :

$$\text{init}_i(k) \in \widetilde{\text{L}}(\tilde{s}) \quad \text{iff} \quad f_i(k) = (\delta(q_{\text{init}}, \text{L}(s)), \bar{0})$$

$$\text{run}_i(k) \in \widetilde{\text{L}}(\tilde{s}) \quad \text{iff} \quad f_i(k) \neq \perp$$

$$\text{goal}_i(k) \in \widetilde{\text{L}}(\tilde{s}) \quad \text{iff} \quad f_i(k) = (q, \bar{w}) \text{ for some } q \in F \text{ and } \text{constr}[\text{wgt}/\bar{w}]$$

We use  $\text{constr}[\text{wgt}/\bar{w}]$  to denote the variable-free arithmetic condition resulting from  $\text{constr}$  by replacing the weight symbols  $\text{wgt}_k$  in the weight expressions of  $\text{constr}$  with the values  $w_k$  for  $k = 1, \dots, d$ . Thus,  $\text{constr}[\text{wgt}/\bar{w}]$  can be treated as a truth value.

Let  $\tilde{\varphi}$  be the LTL formula that results from  $\varphi$  by replacing the subformulas  $\psi_+ = \diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$  and  $\psi_- = \diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$  with:

$$\tilde{\psi}_+ \stackrel{\text{def}}{=} \bigvee_{0 \leq k \leq \ell_i} (\varphi_1 \wedge \text{init}_i(k) \wedge (\text{run}_i(k) \text{U}(\text{goal}_i(k) \wedge \varphi_2)))$$

$$\tilde{\psi}_- \stackrel{\text{def}}{=} \bigvee_{0 \leq k \leq \ell_i} (\varphi_2 \wedge \text{goal}_i(k) \wedge (\text{run}_i(k) \text{S}(\text{init}_i(k) \wedge \varphi_1)))$$

**Theorem 10 (Soundness)** For each state  $s$  in  $\mathcal{M}$ :

$$\text{Pr}_{\mathcal{M}, s}^{\min}(\varphi) = \text{Pr}_{\widetilde{\mathcal{M}}, \tilde{s}}^{\min}(\tilde{\varphi}) \quad \text{and} \quad \text{Pr}_{\mathcal{M}, s_{\text{init}}}^{\max}(\varphi) = \text{Pr}_{\widetilde{\mathcal{M}}, \tilde{s}}^{\max}(\tilde{\varphi})$$

where  $\tilde{s} = \langle s, f_1^s, \dots, f_m^s \rangle$ .

The proof is presented in appendix D.

Thus, we can rely on well-known model-checking techniques for MDPs and LTL. Most prominent is the automata-based approach that transforms the LTL formula into a deterministic  $\omega$ -automaton  $\mathcal{D}$  and then analyzes the end components of the product of the given MDP and  $\mathcal{D}$  (see, e.g., [3, 15]). The worst-case time complexity of this approach is dominated by the generation of a deterministic automaton for the LTL formula and runs in time polynomial in the size of the MDP and double exponential in the length of the LTL formula.

In our case, the size of the generated MDP  $\widetilde{\mathcal{M}}$  is polynomial in the size of  $\mathcal{M}$ , but (single) exponential in the length of the longest runs in the automata of subformulas  $\diamond^A(\varphi_1; \text{constr}; \varphi_2)$  or

$\diamond^A (\varphi_1; \text{constr}; \varphi_2)$  (see Remark 9). Thus, the time complexity of our algorithm is double exponential, too.

## 4.2 Complexity

We now discuss the complexity of the model-checking problem for LTL $[\diamond, \diamond : \text{Acyc}]$  and its sublogic LTL $[\diamond, \diamond : \text{Window}]$  over WMDPs, WMCs and WTSs (see Notation 6).

The model-checking problem for standard LTL is known to be 2EXPTIME-complete for MDPs and PSPACE-complete for Markov chains and transition systems [15, 26, 28]. Obviously, the lower bounds carry over to any logic that extends LTL. Nondeterministic polynomially space-bounded algorithms for the model-checking problem for LTL $[\diamond, \diamond : \text{Acyc}]$  in WMCs and WTSs arise by adapting the approaches of [26] and [28]. Hence:

**Theorem 11** *The model-checking problem for LTL $[\diamond, \diamond : \text{Acyc}]$  and LTL $[\diamond, \diamond : \text{Window}]$  is 2EXPTIME-complete for WMDPs and PSPACE-complete for WMCs and WTSs.*

With the reduction presented in the previous section, the probabilistic analysis has to be carried out with the MDP  $\widehat{\mathcal{M}} = \text{Monitor}(\mathcal{M}, \dots)$  whose size grows exponentially in the sizes of (more precisely, the length of longest runs in) the automata  $\mathcal{A} \in \text{AUT}$  that appear as parameters of the modalities  $\diamond^A$  and  $\square^A$ . However, we cannot expect much more efficient algorithms for the LTL $[\diamond, \diamond : \text{Acyc}]$ -PMC problem since even simple patterns of PL $[\diamond : \text{Window}]$ -formulas interpreted over WTSs and WMCs can encode NP-hard problems as shown in the proof of the following theorem.

**Theorem 12 (NP/coNP-completeness for WTSs)** *For WTSs the problem “does  $s_{\text{init}} \models \Phi_i$  hold?” is NP-complete for formulas of the type  $\Phi_1, \Phi_2, \Phi_3$  and coNP-complete for  $\Phi_4, \Phi_5$  and  $\Phi_6$  where:*

$$\begin{aligned} \Phi_1 &= \exists \diamond^{\leq \ell} \text{constr} & \Phi_4 &= \forall \diamond^{\leq \ell} \text{constr} \\ \Phi_2 &= \exists \square \diamond^{\leq \ell} \text{constr} & \Phi_5 &= \forall \square \diamond^{\leq \ell} \text{constr} \\ \Phi_3 &= \exists \square \square \diamond^{\leq \ell} \text{constr} & \Phi_6 &= \forall \square \square \diamond^{\leq \ell} \text{constr} \end{aligned}$$

The same holds when  $\diamond^{\leq \ell}$  is replaced with  $\diamond^{\leq \ell}$ .

The hardness results in Theorem 12 can be shown using reductions from the subset sum problem and its complement. See appendix, Theorem 27. The coNP-hardness proof for  $\diamond^{\leq \ell}$  uses two positive reward functions. In the remaining cases, the hardness already holds for a WTS with a single positive integer-valued reward function and when  $\text{constr}$  is a simple basic weight constraint  $\text{wgt} = c$  or its negation. Analogous results are obtained for Markov chains, extending the NP-hardness result of [24] for the quantitative PMC decision problem and reward-bounded reachability:

**Theorem 13 (NP/coNP-completeness for WMCs)** *For WMCs the problems to decide whether*

$$\begin{aligned} \Pr_{\mathcal{M}, s_{\text{init}}}(\diamond^{\leq \ell} \text{constr}) &> 0 \\ \Pr_{\mathcal{M}, s_{\text{init}}}(\diamond \diamond^{\leq \ell} \text{constr}) &> 0 & \Pr_{\mathcal{M}, s_{\text{init}}}(\diamond \diamond^{\leq \ell} \text{constr}) &= 1 \\ \Pr_{\mathcal{M}, s_{\text{init}}}(\square \diamond \diamond^{\leq \ell} \text{constr}) &> 0 & \Pr_{\mathcal{M}, s_{\text{init}}}(\square \diamond \diamond^{\leq \ell} \text{constr}) &= 1 \end{aligned}$$

are NP-complete. NP-hardness even holds with a single positive integer-valued reward function and if  $\text{constr}$  has the form  $\text{wgt} = c$ . The problem “does  $\Pr_{\mathcal{M}, s_{\text{init}}}(\diamond^{\leq \ell} \text{constr}) = 1$  hold?” is coNP-complete. The same holds when  $\diamond^{\leq \ell}$  is replaced with  $\diamond^{\leq \ell}$ .

For the proof, see the appendix, Theorem 29.

So far we presented hardness results for simple patterns of formulas with monitors  $\mathcal{A} \in \text{Window}$ , partly with simple weight assertions

and single positive reward functions. But even for Boolean combinations of simple window weight assertions, the model-checking problem is computationally hard.

**Theorem 14** *For PL $[\diamond : \text{Acyc}]$  and PL $[\diamond : \text{Window}]$ , the positive model-checking problem for WMDPs and WMCs and the existential model-checking problem for WTSs are NP-complete. Hardness already holds for a single positive integer-valued reward function.*

For the proof, see the appendix, Theorem 30.

## 4.3 Special algorithms for selected formula patterns

As a consequence of Theorem 12, the task to compute  $\Pr_{\mathcal{M}, s_{\text{init}}}^{\max}(\diamond \diamond^A \text{constr})$  for WMDPs with a single reward functions is computationally hard if  $\text{constr}$  is a conjunctive weight constraint  $\text{wgt} = c$  (which is  $(\text{wgt} \leq c) \wedge (\text{wgt} \geq c)$ ). However, the analogous problem for simple basic weight constraints  $\text{wgt} \bowtie c$  can be solved efficiently, even for WMDP with a (possibly negative) weight function.

**Proposition 15** *For WMDPs with a single weight function, the problems*

$$\begin{aligned} \text{“does } \Pr_{\mathcal{M}, s_{\text{init}}}^{\max}(\diamond \diamond^A(\text{wgt} \bowtie c)) &> 0 \text{ hold?”} \\ \text{“does } \Pr_{\mathcal{M}, s_{\text{init}}}^{\min}(\square \square^A(\text{wgt} \bowtie c)) &= 1 \text{ hold?”} \end{aligned}$$

are in P.

**Proof.** It suffices to consider the problem to decide whether the probability of  $\diamond \diamond^A(\text{wgt} \bowtie c)$  is positive as:

$$\Pr_{\mathcal{M}, s_{\text{init}}}^{\min}(\square \square^A(\text{wgt} \bowtie c)) = 1 - \Pr_{\mathcal{M}, s_{\text{init}}}^{\max}(\diamond \diamond^A(\text{wgt} \nabla c))$$

We define  $\text{Graph}[\mathcal{M} \otimes \mathcal{A}]$  as a weighted directed graph with the vertices  $\langle s, q \rangle \in S \times Q$  and the following edge relation  $E$ :

$$\begin{aligned} (\langle s, q \rangle, \langle s', q' \rangle) &\in E \\ \text{iff } P(s, \text{act}, s') &> 0 \text{ for some } \text{act} \in \text{Act}(s) \\ \text{and } q' = \delta(q, s') &\neq \perp \end{aligned}$$

The weight of the edge from vertex  $\langle s, q \rangle$  to vertex  $\langle s', q' \rangle$  is

$$\min \{ \text{wgt}(s, \text{act}) : P(s, \text{act}, s') > 0 \}.$$

We may apply standard polynomial-time shortest path algorithms (e.g., the algorithms by Bellman-Ford or Floyd) to compute the lengths  $\ell_{\min}(s)$  and  $\ell_{\max}(s)$  of shortest and longest paths from  $\langle s, q_{\text{init}} \rangle$  to some state  $\langle t, p \rangle$  with  $p \in F$  in  $\text{Graph}[\mathcal{M} \otimes \mathcal{A}]$ . Here, the notion of length is to be understood in terms of accumulated weight. (Longest paths are obtained by shortest path algorithms for the weighted graph that results from  $\text{Graph}[\mathcal{M} \otimes \mathcal{A}]$  by multiplying all weights by  $-1$ .) Here, we deal with  $\ell_{\min}(s) = +\infty$  if no state in  $S \times F$  is reachable from  $\langle s, q_{\text{init}} \rangle$  and  $\ell_{\min}(s) = -\infty$  if some cycle with negative weight is reachable from  $\langle s, q_{\text{init}} \rangle$ . Similarly, we have  $\ell_{\max}(s) \in \mathbb{Q} \cup \{-\infty, +\infty\}$  with the same conditions for  $-\infty$  and  $+\infty$ .

Let  $\triangleleft \in \{\leq, <\}$  and  $\triangleright \in \{\geq, >\}$ . The statement follows from the fact that there is a scheduler  $\mathfrak{S}$  for  $\mathcal{M}$  where  $\Pr_{\mathcal{M}, s_{\text{init}}}^{\mathfrak{S}}(\diamond \diamond^A(\text{wgt} \triangleleft c))$  is positive if and only if there exists a state  $s$  in  $\mathcal{M}$  that is reachable from  $s_{\text{init}}$  with  $\ell_{\min}(s) \triangleleft c$ . Similarly, there is a scheduler  $\mathfrak{S}$  for  $\mathcal{M}$  where  $\Pr_{\mathcal{M}, s_{\text{init}}}^{\mathfrak{S}}(\diamond \diamond^A(\text{wgt} \triangleright c))$  is positive if and only if there exists a state  $s$  in  $\mathcal{M}$  that is reachable from  $s_{\text{init}}$  with  $\ell_{\max}(s) \triangleright c$ . ■

According to Theorem 13, the computation of  $\Pr_{\mathcal{M}, s_{\text{init}}}(\square \square \diamond^A \text{constr})$  is hard in WMCs, even in the 1-dimensional case. The problem becomes considerably simpler for basic weight constraints:



	General	Non-negative weight functions, simple weight constraints
PL[ $\diamond$ : Acyc] PL[ $\diamond$ : Window]	<b>NP-complete</b> (Thm. 14)	
LTL[ $\diamond, \diamond$ : Acyc] LTL[ $\diamond, \diamond$ : Window]	WTS, WMC: WMDP:	<b>PSPACE-complete</b> (Thm. 11) <b>2EXPTIME-complete</b> (Thm. 11)
PL[ $\diamond$ : Reach] PL[ $\diamond$ : All]	<b>undecidable</b> (Thm. 17,20)	<i>decidable</i>
LTL[ $\diamond, \diamond$ : Reach] LTL[ $\diamond, \diamond$ : All]	<i>undecidable</i>	<b>decidable</b> (Thm. 18)

**Table 1.** Decidability and complexity of the model-checking problem.

**Proposition 16** *For WMCs with a single weight function, the probabilities*

$$\Pr_{\mathcal{M}, s_{init}}(\Box \diamond \diamond^A(\text{wgt} \bowtie c))$$

$$\Pr_{\mathcal{M}, s_{init}}(\diamond \Box \Box^A(\text{wgt} \bowtie c))$$

*can be computed in polynomial time.*

The proof relies on the computation of the bottom strongly connected components that are *good* according to the weight constraint  $\text{wgt} \bowtie c$ , i.e. those BSCCs  $T$  for which there is no finite path  $\pi$  in  $T$  with  $\text{trace}(\pi) \in \mathcal{L}(\mathcal{A})$  and  $\text{wgt}(\pi) \not\bowtie c$ . Checking whether a BSCC is good can be done using shortest path algorithms. The proof is presented in the appendix, Proposition 32.

## 5. Unbounded weight assertions

So far, we studied the model-checking problem for LTL[ $\diamond, \diamond$  : Acyc] where the operators  $\diamond^A$  and  $\Box^A$  are parametrized by acyclic DFA, i.e., their accepted languages are finite. Dropping this assumption leads to undecidability. This is an immediate consequence of the undecidability results by [7] for LTL with prefix-accumulation assertions which can be seen as a sublogic of LTL[ $\diamond, \diamond$  : All]; see Section 3.4.

More interesting is the observation that undecidability even holds for the logic PL[ $\diamond$  : Reach], i.e., propositional logic where the atoms are pure weight assertions  $\diamond^{A[\dots\phi]}$  constr where  $\phi$  is an ordinary propositional formula with atoms in AP. Recall that the path conditions specified by the automata  $\mathcal{A}[\dots\phi] \in \text{Reach}$  are reachability constraints  $\diamond\phi$ . Given the fact that Boker et al [7] prove decidability for the branching-time logic obtained by adding the CTL-modality  $\exists\diamond$  and prefix-accumulation assertions to propositional logic, this result appears surprising to us. Using a reduction from the Post correspondence problem, we get:

**Theorem 17 (Undecidability for PL[ $\diamond$  : Reach])** *The model-checking problem for PL[ $\diamond$  : Reach] over WTSs and WMCs is undecidable.*

The proof can be found in the appendix, Theorem 33.

By Theorem 17, there is no chance to design algorithms for the computation of (maximal or minimal) probabilities for the events specified as formulas of PL[ $\diamond$  : Reach] (or more expressive logics such as LTL[ $\diamond, \diamond$  : Reach] and LTL[ $\diamond, \diamond$  : All]) in WMCs and WMDPs. We now discuss the case of non-negative structures where all weight functions are non-negative.

Recall that simple weight constraints are Boolean combinations of simple basic weight constraints  $\text{wgt}_i \bowtie c$  and that LTL<sub>simple</sub>[ $\diamond, \diamond$  : All] is the sublogic of LTL[ $\diamond, \diamond$  : All] where all weight constraints are simple.

**Theorem 18** *The LTL<sub>simple</sub>[ $\diamond, \diamond$  : All]-PMC problem is decidable for non-negative WMDPs.*

The proof can be found in the appendix, Theorem 35. It relies on an adaption of the algorithm presented in Section 4.1 using a threshold technique that avoids the expansion of pairs  $f_i(k) = (q, \bar{w})$  where the values of  $\bar{w}$  are larger than the largest constant in the weight constraints of the given formula. To treat 0-weight cycles we use the fact that as soon as  $f_i(k) = f_i(k')$  then we can release one argument  $k$  or  $k'$  and reuse it for fresh runs. The full proof is in the appendix, Theorem 35.

For  $d = 1$ , all weight constraints are equivalent to simple ones:

**Corollary 19** *The LTL[ $\diamond, \diamond$  : All]-PMC problem is decidable for WMDPs with a single non-negative weight function.*

In the multi-dimensional case the requirement that the basic weight expressions are simple cannot be dropped as we have:

**Theorem 20** *The model-checking problem for PL[ $\diamond$  : Reach] and non-negative WTSs is undecidable. Likewise, the model-checking problem for PL[ $\diamond$  : Reach] and non-negative WMCs is undecidable.*

The full proof is presented in the appendix, Theorem 36.

## 6. Conclusions

We established sharp complexity bounds and investigated the border of decidability for the model-checking problem of our new logics. Our main results are depicted in table 4.2, where we distinguish between the general case with arbitrary (rational) weight functions and weight assertions and the case of simple weight assertions with non-negative weight functions. Note that the second column is only applicable if both of these restrictions apply. The given results refer to the positive model-checking problem for WMCs and WMDPs and the existential one for WTSs. The results for the automata class Acyc also hold for the automata class Window. Similarly, the results for Reach also hold for All. The decidability results in the table written in italic are a direct consequence of some result in boldface.

Even though the stated complexity bounds seem to make a practical application unfeasible, there are many techniques to make LTL model checking for MDPs applicable to real-world scenarios. An evaluation of the methods used in some popular model-checking tools can be found, e.g., in [23] for PRISM, in [21] for MRMC and in [5] for ProbDiVinE.

**Acknowledgements** We are grateful to the operating systems group, especially to Hermann Härtig and Marcus Völp, for providing inspiration for this work through real-world problems. We would also like to thank Marcus Daum, Clemens Dubsclaff, Daniel Krähhmann, Linda Leuschner, Steffen Märcker and David Müller for fruitful discussions within our group.

## References

- [1] P. A. Abdulla, R. Mayr, A. Sangnier, and J. Sproston. Solving parity games on integer vectors. In *24th International Conference on Concurrency Theory (CONCUR)*, volume 8052 of *Lecture Notes in Computer Science*, pages 106–120. Springer, 2013.
- [2] S. Andova, H. Hermanns, and J.-P. Katoen. Discrete-time rewards model-checked. In *First International Workshop on Formal Modeling and Analysis of Timed Systems Workshop (FORMATS)*, volume 2791 of *Lecture Notes in Computer Science*, pages 88–104. Springer, 2003.
- [3] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [4] C. Baier, L. Cloth, B. R. Haverkort, H. Hermanns, and J.-P. Katoen. Performability assessment by model checking of Markov reward models. *Formal Methods in System Design*, 36(1):1–36, 2010.
- [5] J. Barnat, L. Brim, I. Cema, M. Ceska, and J. Tumova. ProbDiVinE: A parallel qualitative ltl model checker. *4th International Conference on Quantitative Evaluation of Systems (QEST)*, pages 215–216, 2007.
- [6] R. Bloem, K. Chatterjee, T. Henzinger, and B. Jobstmann. Better quality in synthesis through quantitative objectives. In *21st International Conference on Computer Aided Verification (CAV)*, volume 5643 of *Lecture Notes in Computer Science*, pages 140–156. Springer, 2009. .
- [7] U. Boker, K. Chatterjee, T. A. Henzinger, and O. Kupferman. Temporal specifications with accumulative values. In *26th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 43–52. IEEE Computer Society, 2011.
- [8] T. Brázdil, V. Brozek, K. Chatterjee, V. Forejt, and A. Kucera. Two views on multiple mean-payoff objectives in Markov decision processes. In *26th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 33–42. IEEE Computer Society, 2011. ISBN 978-0-7695-4412-0.
- [9] K. Chatterjee and L. Doyen. Energy and mean-payoff parity Markov decision processes. In *36th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 6907 of *Lecture Notes in Computer Science*, pages 206–218. Springer, 2011.
- [10] K. Chatterjee and L. Doyen. Energy parity games. *Theoretical Computer Science*, 458:49–60, 2012.
- [11] K. Chatterjee, R. Majumdar, and T. A. Henzinger. Markov decision processes with multiple objectives. In *23rd Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 3884 of *Lecture Notes in Computer Science*, pages 325–336. Springer, 2006. ISBN 978-3-540-32301-3. .
- [12] K. Chatterjee, T. Henzinger, B. Jobstmann, and R. Singh. Measuring and synthesizing systems in probabilistic environments. In *22nd International Conference on Computer Aided Verification (CAV)*, volume 6174 of *Lecture Notes in Computer Science*, pages 380–395. Springer, 2010. .
- [13] K. Chatterjee, L. Doyen, M. Randour, and J.-F. Raskin. Looking at mean-payoff and total-payoff through windows. In *11th International Symposium on Automated Technology for Verification and Analysis (ATVA)*, pages 118–132, 2013.
- [14] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 2000.
- [15] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.
- [16] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, Department of Computer Science, 1997.
- [17] L. de Alfaro. How to specify and verify the long-run average behavior of probabilistic systems. In *13th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 454–465. IEEE Computer Society, 1998.
- [18] K. Etessami, M. Kwiatkowska, M. Vardi, and M. Yannakakis. Multi-objective model checking of Markov decision processes. *Logical Methods in Computer Science*, 4(4), 2008.
- [19] V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker. Automated verification techniques for probabilistic systems. In *11th International School on Formal Methods for the Design of Computer, Communication and Software Systems (SFM)*, volume 6659 of *Lecture Notes in Computer Science*, pages 53–113. Springer, 2011.
- [20] B. Haverkort. *Performance of Computer Communication Systems: A Model-Based Approach*. Wiley, 1998.
- [21] J.-P. Katoen, I. Zapreev, E. Hahn, H. Hermanns, and D. Jansen. The ins and outs of the probabilistic model checker MRMC. *Performance Evaluation*, 68(2):90–104, 2011.
- [22] V. Kulkarni. *Modeling and Analysis of Stochastic Systems*. Chapman & Hall, 1995.
- [23] M. Kwiatkowska, G. Norman, and D. Parker. Advances and challenges of probabilistic model checking. In *48th Annual Allerton Conference on Communication, Control and Computing*, pages 1691–1698. IEEE Press, 2010.
- [24] F. Laroussinie and J. Sproston. Model checking durational probabilistic systems. In *8th International Conference on Foundations of Software Science and Computational Structures (FOSSACS)*, volume 3441 of *Lecture Notes in Computer Science*, pages 140–154. Springer, 2005.
- [25] M. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
- [26] A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logic. *Journal of the ACM*, 32(3):733–749, 1985.
- [27] T. Tomita, S. Hiura, S. Hagihara, and N. Yonezaki. A temporal logic with mean-payoff constraints. In *14th International Conference on Formal Engineering Methods. Formal Methods and Software Engineering (ICFEM)*, volume 7635 of *Lecture Notes in Computer Science*, pages 249–265. Springer, 2012.
- [28] M. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification (preliminary report). In *1st Symposium on Logic in Computer Science (LICS)*, pages 332–344. IEEE Computer Society Press, 1986.
- [29] C. von Essen and B. Jobstmann. Synthesizing systems with optimal average-case behavior for ratio objectives. In *International Workshop on Interactions, Games and Protocols (iWIGP)*, volume 50 of *EPTCS*, pages 17–32, 2011.

## A. Additional notations

We write  $IPaths(s)$  for the set of infinite paths  $\zeta$  with  $first(\zeta) = s$ , while  $FPaths(s)$  denotes the set of finite paths starting in  $s$ .

If  $\zeta$  is an infinite path then  $pref(\zeta, k) = \zeta[0 \dots k]$  is the prefix of  $\zeta$  consisting of the first  $k$  steps, ending in state  $\zeta[k] = s_k$ . Similarly,  $suff(\zeta, k)$  denotes the suffix of  $\zeta$ , starting in  $\zeta[k] = s_k$ .

Given a scheduler  $\mathfrak{S} : FPaths \rightarrow Act$  for an MDP  $\mathcal{M}$  a finite  $\mathfrak{S}$ -path is any path  $\pi = s_0 act_0 \dots act_{n-1} s_n$  in  $\mathcal{M}$  that arises when the nondeterministic choices in  $\mathcal{M}$  are resolved using  $\mathfrak{S}$ , i.e.,  $\mathfrak{S}(pref(\pi, k-1)) = act_k$  for all  $1 \leq k \leq n$ . Infinite  $\mathfrak{S}$ -paths are defined accordingly. Given some scheduler  $\mathfrak{S}$  and state  $s$  (viewed as the initial state), the behavior of  $\mathcal{M}$  under  $\mathfrak{S}$  is purely probabilistic and can be formalized by a tree-like (infinite-state) Markov chain  $\mathcal{M}_s^{\mathfrak{S}}$ . One can think of the states in  $\mathcal{M}_s^{\mathfrak{S}}$  as the finite  $\mathfrak{S}$ -paths starting in state  $s$ , where the probability to move from some finite path  $\pi$  to  $\pi act s'$  is simply  $P(last(\pi), act, s')$ . Using standard concepts of measure and probability theory, a sigma-algebra and a probability measure  $\Pr_{\mathcal{M}_s^{\mathfrak{S}}}$  for measurable sets of the infinite paths in the Markov chain  $\mathcal{M}_s^{\mathfrak{S}}$ , also called (*path*) *events* or *path properties*, is defined and can be transferred to infinite  $\mathfrak{S}$ -paths in  $\mathcal{M}$  starting in  $s$ . For further details, we refer to standard text books such as [20, 22, 25].

**Sub-MDPs, end components.** We use the notion *sub-MDP* of  $\mathcal{M}$  for any pair  $(T, \mathfrak{A})$  where  $T \subseteq S$  and  $\mathfrak{A} : T \rightarrow 2^{Act}$  such that for all  $t \in T$ :

- (1)  $\mathfrak{A}(t) \subseteq Act(t)$  and
- (2) if  $act \in \mathfrak{A}(t)$  and  $P(t, act, t') > 0$  then  $t' \in T$ .

An *end component* [16] of  $\mathcal{M}$  is a sub-MDP  $(T, \mathfrak{A})$  of  $\mathcal{M}$  where  $\mathfrak{A}(t)$  is nonempty for all  $t \in T$  and the underlying directed graph with node set  $T$  and the edge relation  $t \rightarrow t'$  iff  $P(t, act, t') > 0$  for some  $act \in \mathfrak{A}(t)$  is strongly connected. An end component is said to be *maximal* if it is not contained in any other end component.

## B. Remarks on the semantics

**Remark 21 (Universal generalized weight assertions)**

In Section 3 we introduced the operators  $\boxplus$  and  $\boxtimes$  only for *pure* weight constraints. This was simply because  $LTL[\boxplus, \boxtimes : AUT]$  does not directly support the complementation of triples  $(\varphi_1; \text{constr}; \varphi_2)$ . Their semantics is a conjunction “precondition  $\varphi_1$  and weight constraint  $\text{constr}$  and postcondition  $\varphi_2$ ”. Generalized universal weight assertions could have been defined using triples  $[\varphi_1 \circ \text{constr} \circ \varphi_2]$  with a disjunctive meaning “precondition  $\varphi_1$  or weight constraint  $\text{constr}$  or postcondition  $\varphi_2$ ”:

$$\boxplus^A[\varphi_1 \circ \text{constr} \circ \varphi_2] \stackrel{\text{def}}{=} \neg \boxtimes^A(\neg \varphi_1; \neg \text{constr}; \neg \varphi_2)$$

with the following semantics over path-position pairs:

$$(\zeta, k) \models \boxplus^A[\varphi_1 \circ \text{constr} \circ \varphi_2]$$

iff for each  $h \in \mathbb{N}$  with  $h > k$  with  $trace(\zeta[k \dots h]) \in \mathcal{L}(A)$  we have:

$$\begin{aligned} & \zeta[k \dots h] \models \text{constr} \\ \text{or } & (\zeta, k) \models \varphi_1 \\ \text{or } & (\zeta, h) \models \varphi_2 \end{aligned}$$

An analogous definition could have been provided for universal generalized past weight assertions  $\boxplus^A[\varphi_1 \circ \text{constr} \circ \varphi_2]$ . ■

**Remark 22 (Integer vs. rational weights; see Remark 7)** We introduced weight functions in MDPs as functions of the type  $wgt : S \times Act \rightarrow \mathbb{Q}$ . When using  $LTL[\boxplus, \boxtimes : AUT]$  as a specification formalism for (finite) WMDPs, however, integer-valued weight functions are equally expressive. Let us briefly explain why. With  $N \in \mathbb{N}$

being the least common multiple of the denominators of all weights  $wgt_i(s, act)$ , we replace the rational weight function  $wgt_i$  of the given WMDP  $\mathcal{M}$  with the integer weight function  $wgt'_i = N \cdot wgt_i$ . Using  $wgt'_i$  as interpretation of the weight symbol  $wgt_i$  in the given signature  $\text{Sig}$ , we can then replace each basic weight constraint  $\text{expr} \boxtimes c$  with  $\text{expr} \boxtimes N \cdot c$ . Since for each finite path  $\pi$  we have:

$$\begin{aligned} & a_1 \cdot wgt_1(\pi) + \dots + a_d \cdot wgt_d(\pi) \boxtimes c \\ \text{iff } & a_1 \cdot wgt'_1(\pi) + \dots + a_d \cdot wgt'_d(\pi) \boxtimes N \cdot c \end{aligned}$$

this transformation is justified. ■

**Remark 23 (Transformation; see Remark 8)** Instead of the explicit representation of the accumulated weight for  $d$  weight functions one can combine  $wgt_1, \dots, wgt_d$  into one weight function for each basic weight expression. This idea has been suggested by Boker et al. [7] for temporal logics with prefix-accumulation assertions and is applicable here as well. For example, if we are given an  $LTL[\boxplus, \boxtimes : AUT]$ -formula  $\varphi$  that contains a basic weight constraint  $\text{expr} \boxtimes c$  where  $\text{expr} = a_1 \cdot wgt_1 + \dots + a_d \cdot wgt_d$  then one can introduce a new weight function  $wgt : S \times Act \rightarrow \mathbb{Q}$  given by

$$wgt(s, act) = \sum_{i=1}^d a_i \cdot wgt_i(s, act)$$

and replace the original basic weight expression  $\text{expr} \boxtimes c$  in  $\varphi$  with  $wgt \boxtimes c$  where  $wgt$  is a fresh weight symbol. Let  $\varphi'$  be the resulting formula over the extended signature with the weight symbol  $wgt$  and  $\mathcal{M}'$  the WMDP resulting from  $\mathcal{M}$  by adding  $wgt$ . Then:

$$\begin{aligned} \Pr_{\mathcal{M}, s_{init}}^{\min}(\varphi) &= \Pr_{\mathcal{M}', s_{init}}^{\min}(\varphi') \\ \Pr_{\mathcal{M}, s_{init}}^{\max}(\varphi) &= \Pr_{\mathcal{M}', s_{init}}^{\max}(\varphi') \end{aligned}$$

If all weight constraints in  $\varphi$  have the form  $\text{expr} \boxtimes c$  for the same weight expression  $\text{expr}$ , then  $\varphi'$  does not contain any weight constraint for  $wgt_1, \dots, wgt_d$ . Hence, in this case the model-checking problem for  $LTL[\boxplus, \boxtimes : AUT]$ -formulas  $\varphi$  over signatures with multiple weight symbols  $wgt_1, \dots, wgt_d$  is reducible to the model-checking problem for  $LTL[\boxplus, \boxtimes : AUT]$ -formulas over the signature with a single weight symbol  $wgt$ . ■

## C. From average to sum expressions

**Remark 24** As described in [7] for temporal logic with prefix-accumulation assertions, average expressions can be transformed into sum expressions of the form  $wgt \boxtimes 0$ . We briefly recall this transformation where, for simplicity, we suppose that we are given an  $LTL^{avg}[\boxplus, \boxtimes : AUT]$ -formula  $\varphi$  with a single average weight constraint  $\text{avgexpr} \boxtimes c$  (see Section 3.4). We extend the signature by a fresh weight symbol  $wgt' = wgt_{d+1}$ . Given a WMDP

$$\mathcal{M} = (S, Act, P, AP, L, \overline{wgt})$$

where  $\overline{wgt} = (wgt_1, \dots, wgt_d)$  provides the meanings for  $wgt_1, \dots, wgt_d$ , we switch from  $\mathcal{M}$  to the WMDP structure

$$\mathcal{M}' = (S, Act, P, AP, L, \overline{wgt}, wgt')$$

that interpretes  $wgt'$  by

$$wgt' : S \times Act \rightarrow \mathbb{Q}, \quad wgt'(s, act) = \sum_{i=1}^d a_i \cdot wgt_i - c$$

Obviously,  $\mathcal{M}$  and  $\mathcal{M}'$  have the same paths and for each finite path  $\pi$  in  $\mathcal{M}$  we have:

$$\pi \models_{\mathcal{M}} \text{avgexpr} \boxtimes c \quad \text{iff} \quad \pi \models_{\mathcal{M}'} wgt' \boxtimes 0$$

Let  $\varphi'$  be the  $LTL[\boxplus, \boxtimes : AUT]$ -formula that results from  $\varphi$  by replacing the average weight expression  $\text{avgexpr} \boxtimes c$  with  $wgt' \boxtimes 0$ . Then:

$$\begin{aligned}\Pr_{\mathcal{M},s}^{\min}(\varphi) &= \Pr_{\mathcal{M}',s}^{\min}(\varphi') \\ \Pr_{\mathcal{M},s}^{\max}(\varphi) &= \Pr_{\mathcal{M}',s}^{\max}(\varphi')\end{aligned}$$

The same technique is applicable if the given LTL<sup>avg</sup>[ $\diamond, \heartsuit$  : AUT]-formula  $\varphi$  contains two or more average weight constraints. ■

## D. Proof of Theorem 10

**Theorem 25 (Soundness of the reduction; see Thm. 10)** *For each state  $s$  in  $\mathcal{M}$ :*

$$\Pr_{\mathcal{M},s}^{\min}(\varphi) = \Pr_{\widetilde{\mathcal{M}},\bar{s}}^{\min}(\widetilde{\varphi}) \text{ and } \Pr_{\mathcal{M},s_{init}}^{\max}(\varphi) = \Pr_{\widetilde{\mathcal{M}},\bar{s}}^{\max}(\widetilde{\varphi})$$

where  $\bar{s} = \langle s, f_1^s, \dots, f_m^s \rangle$ .

**Proof.**  $\widetilde{\mathcal{M}}$  can be seen as a refinement of  $\mathcal{M}$  where the states in  $\mathcal{M}$  are augmented with information on the potential states of the monitors  $\mathcal{A}_i$  and the accumulated weight of the path fragment observed by the monitors. The behavior of  $\mathcal{M}$  is, however, not affected by synchronizing with  $\mathcal{A}_1, \dots, \mathcal{A}_m$ . Note that precisely the actions  $act \in Act(s)$  are enabled in all states  $\langle s, \bar{f} \rangle$  of  $\widetilde{\mathcal{M}}$  and that  $\widetilde{P}(\langle s, \bar{f} \rangle, act, \langle s', \bar{f}' \rangle)$  equals  $P(s, act, s')$ . Thus, all states  $\langle s, \bar{f} \rangle$  in  $\widetilde{\mathcal{M}}$  can be seen as copies of state  $s$  in  $\mathcal{M}$ . (Formally, state  $s$  in  $\mathcal{M}$  is strongly bisimilar to all states  $\langle s, \bar{f} \rangle$  of  $\widetilde{\mathcal{M}}$  when only the atomic propositions in AP are assumed to be observable.) Hence, for each (finite or infinite) path  $\widetilde{\pi}$  in  $\widetilde{\mathcal{M}}$ , the projection  $\widetilde{\pi}|_{\mathcal{M}}$  is a path in  $\mathcal{M}$  where  $\widetilde{\pi}|_{\mathcal{M}}$  arises from  $\widetilde{\pi}$  by dropping all  $f_i$  from any state  $\langle s, f_1, \dots, f_m \rangle$  in  $\widetilde{\pi}$ . Vice versa, if  $\pi = s_0 act_0 s_1 act_1 \dots$  is a path in  $\mathcal{M}$  with  $s_0 = s_{init}$  then  $\pi$  can be lifted to a path

$$lift(\pi) \stackrel{\text{def}}{=} \langle s_0, \bar{f}_0 \rangle act_0 \langle s_1, \bar{f}_1 \rangle act_1 \dots$$

in  $\widetilde{\mathcal{M}}$  where

$$\bar{f}_0 = \langle f_1^{s_0}, \dots, f_m^{s_0} \rangle \quad \bar{f}_{h+1} = \langle f_{h+1,1}, \dots, f_{h+1,m} \rangle$$

with  $f_{h+1,i}$  the  $(act_h, s_{h+1})$ -successor of  $\langle s_h, f_{h,i} \rangle$ . Moreover, we have:

$$\widetilde{\pi} = lift(\widetilde{\pi}|_{\mathcal{M}})$$

If  $\zeta$  is an infinite path in  $\mathcal{M}$  and  $k \in \mathbb{N}$  and  $i \in \{1, \dots, m\}$  then:

$$\langle \zeta, h \rangle \models \diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$$

iff there exists  $k \in \{0, 1, \dots, l_i\}$  such that

$$lift(\zeta, h) \models \begin{cases} \varphi_1 \wedge \text{init}_i(k) \wedge \\ (\text{run}_i(k) \cup \text{goal}_i(k) \wedge \varphi_2) \end{cases}$$

and

$$\langle \zeta, h \rangle \models \diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$$

iff there exists  $k \in \{0, 1, \dots, l_i\}$  such that

$$lift(\zeta, h) \models \begin{cases} \varphi_2 \wedge \text{goal}_i(k) \wedge \\ (\text{run}_i(k) \text{S}(\text{init}_i(k) \wedge \varphi_1)) \end{cases}$$

This yields:

$$\zeta \models \varphi \text{ iff } lift(\varphi) \models \widetilde{\varphi}$$

Using the fact that exactly the same actions are enabled in state  $s$  of  $\mathcal{M}$  and all its copies  $\langle s, \bar{f} \rangle$  in  $\widetilde{\mathcal{M}}$ , the correspondence between paths in  $\mathcal{M}$  and in  $\widetilde{\mathcal{M}}$  can be lifted to schedulers. Let  $\mathfrak{S}$  be a scheduler for  $\mathcal{M}$ . The induced  $lift(\mathfrak{S})$  is defined as follows. If  $\widetilde{\pi}$  is a finite path in  $\widetilde{\mathcal{M}}$  then:

$$lift(\mathfrak{S})(\widetilde{\pi}) = \mathfrak{S}(\widetilde{\pi}|_{\mathcal{M}})$$

Then, the  $lift(\mathfrak{S})$ -paths starting in  $\bar{s} = \langle s, \bar{f}^s \rangle$  are precisely the liftings of the  $\mathfrak{S}$ -paths starting in  $s$ . This yields:

$$\Pr_{\mathcal{M},s}^{\mathfrak{S}}(\varphi) = \Pr_{\widetilde{\mathcal{M}},\bar{s}}^{lift(\mathfrak{S})}(\widetilde{\varphi})$$

Vice versa, if  $\widetilde{\mathfrak{S}}$  is a scheduler for  $\widetilde{\mathcal{M}}$  then a corresponding scheduler  $\mathfrak{S} = \widetilde{\mathfrak{S}}|_{\mathcal{M}}$  is obtained by:

$$\mathfrak{S}(\pi) = \widetilde{\mathfrak{S}}(lift(\pi))$$

Obviously, we then have  $lift(\mathfrak{S}) = \widetilde{\mathfrak{S}}$  and therefore:

$$\Pr_{\mathcal{M},s}^{\mathfrak{S}}(\varphi) = \Pr_{\widetilde{\mathcal{M}},\bar{s}}^{\widetilde{\mathfrak{S}}}(\widetilde{\varphi})$$

The established scheduler-correspondence yields the claim. ■

## E. A heuristic for the 1-dimensional case

Let  $\mathcal{M}$  be a WMDP with a single weight function ,i.e., for  $d = 1$ , in which case we write  $wgt$  rather than  $wgt_1$ . In this case, a simple threshold heuristic can be integrated in the algorithm for the PMC-problem for the logic LTL[ $\diamond, \heartsuit$  : Acyc] presented in Section 4.1.

**Remark 26 (Threshold heuristic)** Suppose we are given an LTL[ $\diamond, \heartsuit$  : Acyc]-formula with a single basic weight constraint, say the task is to compute the probability for the formula  $\varphi = \diamond \heartsuit^A(wgt < c)$ . We apply a standard shortest-path algorithm to  $Graph[\mathcal{M} \otimes \mathcal{A}]$  as in proposition 15 and use the information on the weights of shortest paths to reduce the state space of  $\mathcal{M}'$ . Given a state  $\langle s, f \rangle$  in  $\widetilde{\mathcal{M}} = Monitor(\mathcal{M}, \mathcal{A})$  and a pair  $f(k) = (q, w)$  where the weights of all (shortest) paths from  $\langle s, q \rangle$  to all states  $\langle t, p \rangle$  with  $p \in F$  are at least  $c - w$ , then the expansion of  $(q, w)$  is irrelevant. That is, for the successors  $\langle s', f' \rangle$  of state  $\langle s, f \rangle$  we can put  $f'(k) = \perp$ . ■

## F. Proofs for the complexity-theoretic results

**Theorem 27 (NP/coNP-completeness for WTSs; see Thm. 12)** *For WTSs the problem “does  $s_{init} \models \Phi_i$  hold?” is NP-complete for formulas of the type  $\Phi_1, \Phi_2, \Phi_3$  and coNP-complete for  $\Phi_4, \Phi_5$  and  $\Phi_6$  where:*

$$\begin{aligned}\Phi_1 &= \exists \diamond^{=\ell} \text{ constr} & \Phi_4 &= \forall \diamond^{=\ell} \text{ constr} \\ \Phi_2 &= \exists \diamond \diamond^{=\ell} \text{ constr} & \Phi_5 &= \forall \diamond \diamond^{=\ell} \text{ constr} \\ \Phi_3 &= \exists \square \diamond \diamond^{=\ell} \text{ constr} & \Phi_6 &= \forall \square \diamond \diamond^{=\ell} \text{ constr}\end{aligned}$$

The same holds when  $\diamond^{=\ell}$  is replaced with  $\diamond^{< \ell}$ .

**Proof.** We first address the formulas of the form  $\Phi_1, \Phi_2$  and  $\Phi_3$ . Membership to NP is easy to prove since we can deal with a naïve guess-and-check approach where we first guess nondeterministically a finite path  $\pi$  in the given WTS  $\mathcal{T}$  with initial state  $s_{init}$  and then check whether  $\pi$  yields a witness for the formula. More precisely:

$\Phi_1$ : We simply guess a path  $\pi$  in  $\mathcal{T}$  from  $s_{init}$  of length  $\ell$  and check whether  $\text{constr}[\overline{wgt}/wgt(\pi)]$ .

$\Phi_2$ : We guess a state  $s_0$  and a path  $\pi$  of length  $\ell$  from  $s_0$  and check whether the following conditions (a) and (b) are satisfied:

- (a)  $s_0$  is reachable from  $s_{init}$  in  $\mathcal{T}$
- (b) The weight constraint holds for  $\overline{wgt}(s_0 s_1 \dots, s_n)$ .

$\Phi_3$ : We guess  $s_0$  and a path  $\pi$  as for part (2) and check conditions (a) and (b) and additionally

- (c)  $s_0$  is reachable from  $last(\pi)$  in  $\mathcal{T}$ .

For  $\Phi_1, \Phi_2$  and  $\Phi_3$  with  $\diamond^{=\ell}$  being replaced with  $\diamond^{< \ell}$ , the same guess-and-check approach is applicable by guessing paths  $\pi$  of length  $|\pi| \leq \ell$ .

The NP-hardness proofs are provided using polynomial reductions from the subset sum problem (SUBSUM). For SUBSUM, we are given a finite sequence  $x_1, \dots, x_k, c$  of positive natural numbers and the task is to decide whether there exists an index-set  $I \subseteq \{1, \dots, k\}$  such that  $\sum_{i \in I} x_i = c$ .

Recall that the simple weight constraint  $\text{wgt} = c$  is a short form notation for  $(\text{wgt} \leq c) \wedge (\text{wgt} \geq c)$ . Similarly, we write  $\text{wgt} \neq c$  shortly for  $(\text{wgt} < c) \vee (\text{wgt} > c)$ .

We now construct a WTS  $\mathcal{T}$  with initial state  $s_{\text{init}}$  of size  $\mathcal{O}(k)$  with a single weight function  $\text{wgt}$  with values  $x_1, \dots, x_k$  such that for  $j \in \{1, 2, 3\}$ :

$$s_{\text{init}} \models \Phi_j \quad \text{iff} \quad \text{SUBSUM is solvable for } x_1, \dots, x_k, c$$

where  $\Phi_1, \Phi_2, \Phi_3$  are as above (or their variants with  $\diamond^{\leq \ell}$ ) with  $\text{constr}$  being the simple weight constraint  $\text{wgt} = c$  and  $\ell = k$ . This yields the NP-hardness statement for  $\Phi_1, \Phi_2$  and  $\Phi_3$ . The same WTS  $\mathcal{T}$  will be used to prove the coNP-hardness statements. For  $j \in \{4, 5, 6\}$  we will have:

$$s_{\text{init}} \not\models \Phi_j \quad \text{iff} \quad \text{SUBSUM is solvable for } x_1, \dots, x_k, c$$

For the  $\diamond^{\leq \ell}$ -variant,  $\text{constr}$  is  $\text{wgt} \neq c$ , while for the  $\diamond^{\leq \ell}$ -variant, we will use a second weight function  $\text{wgt}'$  and use the weight constraint  $(\text{wgt} \neq c) \wedge (\text{wgt}' = c)$ .

*Construction of the WTS.* Let  $\mathcal{T} = (S, \text{Act}, \rightarrow, \text{AP}, L, \text{wgt})$  be the WTS with state space

$$S = \{s_{\text{init}}\} \cup \{yes_1, \dots, yes_k, no_1, \dots, no_k\},$$

action set  $\text{Act} = \{yes_1, \dots, yes_k, no_1, \dots, no_k\}$  and transitions

$$\begin{array}{cc} s_{\text{init}} \xrightarrow{yes_1} yes_1 & s_{\text{init}} \xrightarrow{no_1} no_1 \\ s_i \xrightarrow{yes_{i+1}} yes_{i+1} & s_i \xrightarrow{no_{i+1}} no_{i+1} \end{array}$$

for  $s_i \in \{yes_i, no_i\}$  and  $i < k$ , as well as transitions

$$s_k \xrightarrow{yes_1} yes_1 \quad s_k \xrightarrow{no_1} no_1$$

for  $s_k \in \{yes_k, no_k\}$ . Thus, all incoming transitions of state  $s$  are labeled by the action name  $s$ . Intuitively,  $yes_i$  means “yes,  $i \in I$ ”, while  $no_i$  means “no,  $i \notin I$ ”. This intuitive meaning is formalized by the weight function given by:

$$\text{wgt}(yes_i, yes_i) = x_i \quad \text{for } 1 \leq i \leq k$$

and  $\text{wgt}(\cdot) = 0$  in all other cases. Clearly, the size of  $\mathcal{T}$  is polynomial in the size of the input of SUBSUM.

There is a natural correspondence between the finite paths in  $\mathcal{T}$  up to length  $k$  and the subsets of  $\{1, \dots, k\}$ . Let  $\pi$  be a finite path in  $\mathcal{T}$  of length at most  $k$ . For every  $1 \leq i \leq k$ , there is at most one occurrence of an action indexed by  $i$  in  $\pi$  i.e., at most one occurrence of either  $yes_i$  or  $no_i$ . We define  $I_\pi$  as the following index set:

$$I_\pi = \{i \in \{1, \dots, k\} : yes_i \text{ occurs in } \pi\}$$

Then:

$$\text{wgt}(\pi) = c_I \quad \text{where } c_I \stackrel{\text{def}}{=} \sum_{i \in I_\pi} x_i$$

Given an index-set  $I \subseteq \{1, \dots, k\}$ , let  $\pi_I^s$  be the path in  $\mathcal{T}$  of length  $k$  that starts in state  $s$  constructed according to the choices provided by  $I$ , i.e., if the successors of the current state are  $yes_i$  and  $no_i$  then action  $yes_i$  is chosen if  $i \in I$  and action  $no_i$  is chosen if  $i \notin I$ . Then:

$$\text{wgt}(\pi_I^s) = c_I$$

We obtain an infinite path  $\zeta_I^{s_{\text{init}}}$  from  $I$  with

$$\zeta_I^{s_{\text{init}}} \models \square \diamond^k (\text{wgt} = c_I)$$

as follows. Let  $s_k$  be the last state of  $\pi_I^{s_{\text{init}}}$ . We construct  $\zeta_I^{s_{\text{init}}}$  by starting with  $\pi_I^{s_{\text{init}}}$  and then appending infinitely often  $\pi_I^{s_k}$  while merging the last and first states of the paths. As the weights along

any fragment of  $\zeta_I^{s_{\text{init}}}$  of length  $k$  are a permutation of the weights of  $\pi_I^{s_{\text{init}}} = \zeta_I^{s_{\text{init}}}[0 \dots k]$ , we obtain  $\zeta_I^{s_{\text{init}}} \models \square \diamond^k (\text{wgt} = c_I)$ .

*NP-hardness for  $\Phi_1$ .* We show that  $s_{\text{init}} \models \exists \diamond^k (\text{wgt} = c)$  iff there is some  $I \subseteq \{1, \dots, k\}$  such that  $c_I = c$ .

Suppose first that  $s_{\text{init}} \models \exists \diamond^k (\text{wgt} = c)$ . Then, there exists an infinite path  $\zeta$  starting in  $s_{\text{init}}$  with  $\text{wgt}(\pi) = c$  for  $\pi = \zeta[0 \dots k]$ . Then,  $I_\pi$  is a solution to the given instance of SUBSUM.

Vice versa, let  $I \subseteq \{1, \dots, k\}$  be an index set such that  $c_I = c$ . Then:

$$\zeta_I^{s_{\text{init}}} \models \diamond^k (\text{wgt} = c)$$

Hence,  $s_{\text{init}} \models \exists \diamond^k (\text{wgt} = c)$ .

NP-hardness for the variant of  $\Phi_1$  with  $\diamond^{\leq k}$  follows by the fact that  $s_{\text{init}} \models \exists \diamond^{\leq k} (\text{wgt} = c)$  if and only if  $s_{\text{init}} \models \exists \diamond^k (\text{wgt} = c)$ .

*NP-hardness for  $\Phi_2$  and  $\Phi_3$*  is obtained by showing that:

$$\begin{array}{l} \text{there exists an index set } I \text{ with } c_I = c \\ \text{iff } s_{\text{init}} \models \exists \diamond \diamond^k (\text{wgt} = c) \\ \text{iff } s_{\text{init}} \models \exists \square \diamond \diamond^k (\text{wgt} = c) \end{array}$$

Let  $\zeta$  be an infinite path starting in  $s_{\text{init}}$  that satisfies  $\diamond \diamond^k (\text{wgt} = c)$  or  $\square \diamond \diamond^k (\text{wgt} = c)$ . Then, there exists some (smallest) index  $h$  such that  $\text{wgt}(\pi) = c$  for  $\pi = \zeta[h \dots h+k]$ , yielding a solution  $I_\pi$  for the instance  $x_1, \dots, x_k, c$  for SUBSUM.

Vice versa, given an index set  $I$  with  $c_I = c$ , we consider the infinite path  $\zeta_I^{s_{\text{init}}}$  as constructed above. As  $c = c_I$  and  $\zeta_I^{s_{\text{init}}} \models \square \diamond^k (\text{wgt} = c)$ , we have:

$$\zeta_I^{s_{\text{init}}} \models \diamond \diamond^k (\text{wgt} = c) \quad \wedge \quad \square \diamond \diamond^k (\text{wgt} = c).$$

For the variants of  $\Phi_2$  and  $\Phi_3$  with  $\diamond^{\leq k}$  similar arguments with the modifications as for  $\Phi_1$  apply.

*coNP-hardness for  $\Phi_4, \Phi_5$  and  $\Phi_6$  and  $\diamond^{\leq \ell}$ .* We reduce the complement of SUBSUM to the problem that asks whether

$$\begin{array}{l} s_{\text{init}} \models \forall \diamond^k (\text{wgt} \neq c) \\ s_{\text{init}} \models \forall \diamond \diamond^k (\text{wgt} \neq c) \\ s_{\text{init}} \models \forall \square \diamond \diamond^k (\text{wgt} \neq c). \end{array}$$

If  $s_{\text{init}} \not\models \forall \diamond^k (\text{wgt} \neq c)$ , there is an infinite path  $\zeta$  starting in  $s_{\text{init}}$  with  $\zeta \not\models \diamond^k (\text{wgt} \neq c)$ . Then  $\text{wgt}(\pi) = c$  for  $\pi = \zeta[0 \dots k]$ , yielding a solution  $I_\pi$  for the instance  $x_1, \dots, x_k, c$  of SUBSUM.

If  $s_{\text{init}} \not\models \forall \diamond \diamond^k (\text{wgt} \neq c)$  or  $s_{\text{init}} \not\models \forall \square \diamond \diamond^k (\text{wgt} \neq c)$ , then there is an infinite path  $\zeta$  starting in  $s_{\text{init}}$  and an index  $h$  such that:

$$\text{suff}(\zeta, h) \models \square \neg \diamond^k (\text{wgt} \neq c)$$

Recall that  $\text{suff}(\zeta, h)$  denotes the suffix of  $\zeta$  that starts in the  $(h+1)$ -st state of  $\zeta$ . Then, in particular,  $\text{suff}(\zeta, h) \models \neg \diamond^k (\text{wgt} \neq c)$  and  $\text{wgt}(\pi) = c$  for  $\pi = \zeta[h \dots h+k]$ , yielding a solution  $I_\pi$ .

In the other direction, we show that given a solution  $I$  of the SUBSUM instance,  $\zeta_I^{s_{\text{init}}}$  serves as a counterexample to each of the three formulas. As  $\zeta_I^{s_{\text{init}}} \models \diamond^k (\text{wgt} = c)$  we have  $s_{\text{init}} \not\models \forall \diamond^k (\text{wgt} \neq c)$ . Furthermore, by construction of  $\zeta_I^{s_{\text{init}}}$  we have  $\zeta_I^{s_{\text{init}}} \models \square \diamond^k (\text{wgt} = c)$ , as the weights along any fragment of  $\zeta_I^{s_{\text{init}}}$  of length  $k$  are a permutation of the weights of  $\pi = \zeta_I^{s_{\text{init}}}[0 \dots k]$  and  $\text{wgt}(\pi) = c$ . Thus,  $s_{\text{init}} \not\models \forall \diamond \diamond^k (\text{wgt} \neq c)$  and  $s_{\text{init}} \not\models \forall \square \diamond \diamond^k (\text{wgt} \neq c)$ .

*coNP-hardness for  $\Phi_4, \Phi_5$  and  $\Phi_6$  and  $\diamond^{\leq \ell}$ .* To show coNP-hardness for the variants of  $\Phi_4, \Phi_5$  and  $\Phi_6$  with  $\diamond^{\leq \ell}$ , we introduce a second weight function  $\text{wgt}'$  in  $\mathcal{T}$  with  $\text{wgt}'(s, \text{act}) = 1$  for all

$s \in S$  and  $act \in Act$ . That is,  $wgt'$  simply serves as a step counter. We now reduce the complement of SUBSUM to the problem that asks whether

$$\begin{aligned} s_{init} &\models \forall \diamond^{\leq k} ((wgt \neq c) \wedge (wgt' = k)) \\ s_{init} &\models \forall \diamond \diamond^{\leq k} ((wgt \neq c) \wedge (wgt' = k)) \\ s_{init} &\models \forall \square \diamond \diamond^{\leq k} ((wgt \neq c) \wedge (wgt' = k)) \end{aligned}$$

If  $\zeta$  is an infinite path of  $\mathcal{T}$  then:

$$\begin{aligned} \zeta &\models \diamond^{\leq k} (wgt \neq c) \\ \text{iff } \zeta &\models \diamond^{\leq k} ((wgt \neq c) \wedge (wgt' = k)) \end{aligned}$$

Thus, the arguments presented above for the variant  $\diamond^{\leq \ell}$  can be adapted.  $\blacksquare$

**Remark 28 (Hardness results for WMDP)** When interpreting a WTS  $\mathcal{T}$  as a WMDP then we have  $\Pr_{\mathcal{T}, s_{init}}^{\mathcal{G}}(\varphi) \in \{0, 1\}$  for all schedulers  $\mathcal{G}$  and all LTL $[\diamond, \square : \text{AUT}]$ -formulas  $\varphi$  and:

$$\begin{aligned} \Pr_{\mathcal{T}, s_{init}}^{\max}(\varphi) > 0 \text{ iff } \Pr_{\mathcal{T}, s_{init}}^{\max}(\varphi) = 1 \text{ iff } s_{init} \models \exists \varphi \\ \Pr_{\mathcal{T}, s_{init}}^{\min}(\varphi) > 0 \text{ iff } \Pr_{\mathcal{T}, s_{init}}^{\min}(\varphi) = 1 \text{ iff } s_{init} \models \forall \varphi \end{aligned}$$

Hence, by Theorem 12, the problems to decide whether

$$\Pr_{\mathcal{M}, s_{init}}^{\max}(\varphi) > 0, \text{ or } \Pr_{\mathcal{M}, s_{init}}^{\max}(\varphi) = 1$$

are NP-hard, when  $\varphi$  is an LTL $[\diamond, \square : \text{Window}]$ -formula or an LTL $[\diamond, \square : \text{Acyc}]$ -formula. Similarly, the analogous problems for  $\Pr^{\min}$  rather than  $\Pr^{\max}$  are coNP-hard.  $\blacksquare$

**Theorem 29 (NP/coNP-completeness for WMCs; see Thm. 13)**

For WMCs the problems to decide whether

- (1)  $\Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} \text{ constr}) > 0$
- (2)  $\Pr_{\mathcal{M}, s_{init}}(\diamond \diamond^{\leq \ell} \text{ constr}) > 0$
- (3)  $\Pr_{\mathcal{M}, s_{init}}(\square \diamond \diamond^{\leq \ell} \text{ constr}) > 0$
- (4)  $\Pr_{\mathcal{M}, s_{init}}(\diamond \diamond^{\leq \ell} \text{ constr}) = 1$
- (5)  $\Pr_{\mathcal{M}, s_{init}}(\square \diamond \diamond^{\leq \ell} \text{ constr}) = 1$

are NP-complete. NP-hardness even holds with a single positive integer-valued reward function and if constr has the form  $wgt = c$ . The problem “does  $\Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} \text{ constr}) = 1$  hold?” is coNP-complete. The same holds when  $\diamond^{\leq \ell}$  is replaced with  $\diamond^{\leq \ell}$ .

**Proof.** Similarly to the proof of Theorem 27, we provide a polynomial reduction from SUBSUM, the problem where we are given a finite sequence  $x_1, \dots, x_k, c$  of natural numbers and the task is to decide whether there exists an index-set  $I \subseteq \{1, \dots, k\}$  such that  $\sum_{i \in I} x_i = c$ .

The construction differs marginally in the details from the one in the proof of 27 due to the differences in the syntax and semantics of the weight function for WMCs and WTSs.

Let  $\mathcal{M} = (S, P, wgt)$  be the WMC with state space

$$S = \{s_{init}\} \cup \{yes_1, \dots, yes_k, no_1, \dots, no_k\} \cup \{s_{stop}\}$$

As in the proof of Theorem 27,  $yes_i$  means “yes,  $i \in I$ ”, while  $no_i$  means “no,  $i \notin I$ ”. This intuitive meaning is formalized by the weight function given by:

$$wgt(yes_i) = x_i \quad \text{for } 1 \leq i \leq k$$

and  $wgt(\cdot) = 0$  in all other cases. The transition probabilities are given by:

$$\begin{aligned} P(s_{init}, yes_1) &= P(s_{init}, no_1) = \frac{1}{2} \\ P(yes_k, s_{stop}) &= P(no_k, s_{stop}) = 1 \\ P(yes_i, yes_{i+1}) &= P(yes_i, no_{i+1}) = \frac{1}{2} \\ P(no_i, yes_{i+1}) &= P(no_i, no_{i+1}) = \frac{1}{2} \end{aligned}$$

for  $1 \leq i < k$ . Furthermore, we have  $P(s_{stop}, s_{init}) = 1$  and  $P(\cdot) = 0$  in all other cases. Thus, each path from  $s_{init}$  to  $s_{stop}$  has the form

$$\pi_I = s_{init} s_1 s_2 \dots s_k s_{stop} \text{ where } s_i = \begin{array}{ll} yes_i & : \text{ if } i \in I \\ no_i & : \text{ if } i \notin I \end{array}$$

and  $wgt(\pi_I) = c_I$  where  $c_I = \sum_{i \in I} x_i$ .

Obviously, the size of  $\mathcal{M}$  is polynomial in  $k$  and for  $\ell \stackrel{\text{def}}{=} k+1$  we have:

$$\begin{aligned} \Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} (wgt = c)) &> 0 \\ \text{iff } s_{init} \models \exists \diamond^{\leq \ell} (wgt = c) \\ \text{iff there exists } I \subseteq \{1, \dots, k\} \text{ such that } wgt(\pi_I) &= c \end{aligned}$$

This yields that  $\Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} (wgt = c))$  is positive if and only if there exists a subset  $I$  of  $\{1, \dots, k\}$  such that  $\sum_{i \in I} x_i = c$ . Obviously,

the transformation is polynomial. NP-hardness of problems (2), (3), (4) and (5) follow by the observation that  $\mathcal{M}$  is strongly connected. Hence, almost all paths of  $\mathcal{M}$  contain all finite paths infinitely often as path fragments.

Membership to NP is clear for problems (1), (2), (3), (4) and (5) since we can guess a path  $\pi$  in  $\mathcal{M}$  of length  $\ell$  and check in polynomial time whether it meets the weight constraint imposed by constr. For problem (1), we need  $first(\pi) = s_{init}$ .

The complement of SUBSUM is polynomially reducible to the problem “does  $\Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} \text{ constr}) = 1$  hold?”. Indeed, for the transformation above we have:

$$\begin{aligned} \Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} (wgt \neq c)) &= 1 \\ \text{iff } s_{init} \models \forall \diamond^{\leq \ell} (wgt \neq c) \\ \text{iff there is no } I \subseteq \{1, \dots, k\} \text{ such that } wgt(\pi_I) &= c \end{aligned}$$

This yields that the problem to decide whether  $\diamond^{\leq \ell} \text{ constr}$  holds almost surely in a given WMC is coNP-hard. Membership to coNP is obvious since we can guess a path  $\pi$  from  $s_{init}$  of length  $\ell$  and check whether  $\pi \not\models \text{constr}$ .

For the variant of the problem where  $\diamond^{\leq \ell}$  is replaced by  $\diamond^{\leq \ell}$ , similar arguments as in the proof of Theorem 27 apply, yielding the NP-completeness of (1)-(5). To show the coNP-hardness for the problem “does  $\Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} \text{ constr}) = 1$  hold?”, we introduce a second weight function  $wgt'$  in  $\mathcal{T}$  with  $wgt'(s) = 1$  for all  $s \in S$  and reduce the complement of SUBSUM to

$$\Pr_{\mathcal{M}, s_{init}}(\diamond^{\leq \ell} (wgt = c \wedge wgt' = \ell)) = 1$$

with  $\ell \stackrel{\text{def}}{=} k + 1$ .  $\blacksquare$

**Theorem 30 (See Thm. 14)**

For PL $[\diamond : \text{Acyc}]$  and PL $[\diamond : \text{Window}]$ , the positive model-checking problem for WMDPs and WMCs and the existential model-checking problem for WTSs are NP-complete. Hardness already holds for a single positive integer-valued reward function.

**Proof.** The problem to decide  $\Pr_{\mathcal{M}, s_{init}}^{\max}(\varphi) > 0$  is in NP since it is solvable by the following guess-and-check method. Let  $\ell$  be the maximal number of states in the automata  $\mathcal{A} \in \text{Acyc}$  that appear in subformulas  $\diamond^A \text{ constr}$  of  $\varphi$ . Then, the truth value of  $\varphi$  for a

given infinite path  $\zeta$  only depends on the prefix  $\zeta[0 \dots \ell-2]$ . Note that  $\text{trace}(\zeta[0 \dots \ell-2])$  consists of  $\ell-1$  letters and its run in a DFA consist of (at most)  $\ell$  states. Hence, if  $\pi$  is a finite path of length  $\ell-2$  then either  $\zeta \models \varphi$  for all infinite paths  $\zeta$  with  $\zeta[0 \dots \ell-2] = \pi$  or there is no infinite path  $\zeta$  with  $\zeta \models \varphi$  and  $\zeta[0 \dots \ell-2] = \pi$ .

Thanks to this observation, we can use the following polynomially time-bounded guess-&-check method for deciding the existential model-checking problem for  $\text{PL}[\diamond : \text{Acyc}]$ . We guess non-deterministically a finite path  $\pi$  in  $\mathcal{M}$  with  $\text{first}(\pi) = s_{\text{init}}$  and  $|\pi| = \ell-2$ . Then, we can evaluate deterministically in polynomial time whether  $\pi$  is a prefix of an infinite  $\zeta$  satisfying  $\varphi$ . For this, we can use the fact that if  $\pi$  is a prefix of  $\zeta$  then  $\zeta \models \diamond^A \text{constr}$  if and only if there is a prefix  $\pi'$  of  $\pi$  with  $\text{trace}(\pi') \in \mathcal{L}(\mathcal{A})$  and  $\pi' \models \text{constr}$ .

We now prove the NP-hardness for WTSs and  $\text{PL}[\diamond : \text{Acyc}]$ . In Theorem 27 we saw that the problem to decide where

$$s_{\text{init}} \models \exists \diamond^{=k}(\text{wgt} = c)$$

in a WTS with a single integer-valued weight function is NP-hard. Clearly, we have:

$$\begin{aligned} s_{\text{init}} &\models \exists \diamond^{=k}(\text{wgt} = c) \\ \text{iff } s_{\text{init}} &\models \exists \diamond^{=k}((\text{wgt} \leq c) \wedge (\text{wgt} \geq c)) \\ \text{iff } s_{\text{init}} &\models \exists((\diamond^{=k}(\text{wgt} \leq c) \wedge (\diamond^{=k} \text{wgt} \geq c)) \end{aligned}$$

The formula  $\diamond^{=k}(\text{wgt} \leq c) \wedge (\diamond^{=k} \text{wgt} \geq c)$  is a formula of  $\text{PL}[\diamond : \text{Acyc}]$  and  $\text{PL}[\diamond : \text{Window}]$ . This yields the NP-hardness of  $\text{PL}[\diamond : \text{Acyc}]$  and  $\text{PL}[\diamond : \text{Window}]$ . The hardness of the positive model-checking problem for WMCs is obtained by the same arguments applied to case (1) of Theorem 29. ■

**Remark 31** Thanks to Remark 2, for the cases  $\Phi_2, \Phi_3, \Phi_5$  and  $\Phi_6$  of Theorem 12 and parts (2), (3), (4) and (5) of Theorem 13 we can replace  $\diamond$  with  $\diamond$ . This, however, does not hold for the cases  $\Phi_1$  and  $\Phi_4$  of Theorem 12 and part (1) of Theorem 13 since  $\zeta \models \diamond^A \text{constr}$  iff  $(\zeta, 0) \models \diamond^A \text{constr}$  iff the word consisting of the symbol  $L(\zeta[0])$  is not accepted by  $\mathcal{A}$  or  $\text{constr}[\overline{\text{wgt}}/\overline{0}]$ . Thus, the model-checking problem for the formulas of the type  $\diamond^A \text{constr}$  is trivial, even in WMDPs. ■

**Proposition 32 (See Proposition 16)** *For WMCs with a single weight function, the probabilities  $\Pr_{\mathcal{M}, s_{\text{init}}}(\square \diamond \diamond^A(\text{wgt} \bowtie c))$  and  $\Pr_{\mathcal{M}, s_{\text{init}}}(\diamond \square \boxplus^A(\text{wgt} \bowtie c))$  can be computed in polynomial time.*

**Proof.** In what follows, let:

$$\begin{aligned} \varphi_{\diamond \square \boxplus} &= \diamond \square \boxplus^A(\text{wgt} \bowtie c) \\ \varphi_{\diamond \square \diamond} &= \square \diamond \diamond^A(\text{wgt} \bowtie c) \end{aligned}$$

To compute  $\Pr_{\mathcal{M}, s_{\text{init}}}(\varphi_{\diamond \square \boxplus})$  and  $\Pr_{\mathcal{M}, s_{\text{init}}}(\varphi_{\diamond \square \diamond})$  where we can rely on the fact that almost surely a bottom strongly connected component (BSCC)  $T$  will be reached and all states visited infinitely often. Given that the language of  $\mathcal{A}$  is finite there is an upper bound  $\ell \in \mathbb{N}$  on the length of the words in  $\mathcal{L}(\mathcal{A})$  (namely the length of a longest run in  $\mathcal{A}$  from  $q_{\text{init}}$  to some final state). Obviously, the same bound  $\ell$  applies to the maximal length of path fragments  $\pi$  such that  $\text{trace}(\pi) \in \mathcal{L}(\mathcal{A})$ . We say that a BSCC  $T$  of  $\mathcal{M}$  is *good* for  $\varphi_{\diamond \square \boxplus}$  if there is no finite path  $\pi$  in  $T$  with  $\text{trace}(\pi) \in \mathcal{L}(\mathcal{A})$  and  $\text{wgt}(\pi) \not\bowtie c$ . We then have:

- (a) If  $T$  is good for  $\varphi_{\diamond \square \boxplus}$  then  $\Pr_t(\square \boxplus^A(\text{wgt} \bowtie c)) = 1$  for all states  $t \in T$ .
- (b) If  $T$  is not good for  $\varphi_{\diamond \square \boxplus}$  then  $\Pr_t(\square \boxplus^A(\text{wgt} \bowtie c)) = 0$  for all states  $t \in T$ .

Part (a) is obvious. Part (b) follows from the fact that each finite paths in  $T$  appears infinitely often in almost all paths starting in some state  $t \in T$ . As a consequence,

$$\Pr_{\mathcal{M}, s_{\text{init}}}(\diamond \square \boxplus^A(\text{wgt} \bowtie c)) = \Pr_{\mathcal{M}, s_{\text{init}}}(\diamond \text{goodBSCC})$$

and

$$\begin{aligned} \Pr_{\mathcal{M}, s_{\text{init}}}(\square \diamond \diamond^A(\text{wgt} \bowtie c)) \\ = 1 - \Pr_{\mathcal{M}, s_{\text{init}}}(\diamond \square \boxplus^A \neg(\text{wgt} \bowtie c)) \end{aligned}$$

where *goodBSCC* is the union of all BSCC that are good for  $\varphi_{\diamond \square \boxplus}$ .

Checking whether a BSCC is good can be done using shortest path algorithms either in  $\text{Graph}[\mathcal{M} \otimes \mathcal{A}]$  or in the modified weighted graph obtained from  $\text{Graph}[\mathcal{M} \otimes \mathcal{A}]$  by multiplying all weights with  $-1$ , depending on whether  $\bowtie c$  is an upper or a lower bound. ■

## G. Unbounded weight assertions

**Theorem 33 (Undecidability for  $\text{PL}[\diamond : \text{Reach}]$ ; see Thm. 17)** *The model-checking problem for  $\text{PL}[\diamond : \text{Reach}]$  over WTSs and WMCs is undecidable.*

**Proof.** We provide a reduction from the Post correspondence problem (PCP) to the problem where we are given a WTS  $\mathcal{T}$  with initial state  $s_{\text{init}}$  and a  $\text{PL}[\diamond : \text{Reach}]$ -formula  $\varphi$  and the task is to check whether  $s_{\text{init}} \models \exists \varphi$ . As PCP is known to be undecidable, this yields the proof for the undecidability of the model-checking problem for  $\text{PL}[\diamond : \text{Reach}]$  over WTSs.

The input of the PCP is a sequence of  $k \geq 2$  pairs

$$(u_0, v_0), (u_1, v_1), \dots, (u_{k-1}, v_{k-1})$$

consisting of finite, nonempty words over some alphabet  $\Sigma$ . The task is to check whether there exists a sequence  $i_1, i_2, \dots, i_n$  of indices  $i_h \in \{0, 1, \dots, k-1\}$  with  $n \geq 2$  such that:

$$u_{i_1} u_{i_2} \dots u_{i_n} = v_{i_1} v_{i_2} \dots v_{i_n}$$

We construct a WTS  $\mathcal{T}$  that results from the parallel composition

$$\mathcal{T} = \text{Words} \parallel \text{Gadget}[u] \parallel \text{Gadget}[v]$$

of an unweighted transition system *Words* that serves to “generate” pairs of finite words over  $\Sigma$  of the form

$$(u_{i_1} u_{i_2} \dots u_{i_n}, v_{j_1} v_{j_2} \dots v_{j_m})$$

$$\text{where } u_{i_1} u_{i_2} \dots u_{i_n} = v_{j_1} v_{j_2} \dots v_{j_m}$$

and two WTSs *Gadget*[ $u$ ] and *Gadget*[ $v$ ] with the “main” weight functions *useq* and *vseq*, respectively, where the current values of the accumulated weight have the intuitive meaning:

$$\begin{aligned} \text{useq} &\hat{=} i_n \cdot k^{n-1} + i_{n-1} \cdot k^{n-2} + \dots + i_2 \cdot k + i_1 \\ \text{vseq} &\hat{=} j_m \cdot k^{m-1} + j_{m-1} \cdot k^{m-2} + \dots + j_2 \cdot k + j_1 \end{aligned}$$

Then,  $\text{useq} = \text{vseq}$  if and only if the index sequences  $i_1, i_2, \dots, i_n$  and  $j_1, j_2, \dots, j_m$  agree. In addition to the weight function *useq*, the WTS *Gadget*[ $u$ ] has a weight function *ufactor* representing the factor  $k^n$  of the next index  $i = i_{n+1}$  and an auxiliary weight function *ucounter* that serves to realize the assignments:

$$\text{useq} := \text{useq} + \text{ufactor} \cdot i, \quad \text{ufactor} := k \cdot \text{ufactor}$$

*Gadget*[ $v$ ] has analogous additional weight functions *vfactor* and *vcounter*.

**Signature.** Altogether  $\mathcal{T}$  has six weight functions that will be used as meanings for the following weight symbols in the signature for  $\text{PL}[\diamond : \text{Reach}]$ :

$$\text{useq}, \text{vseq}, \text{ufactor}, \text{vfactor}, \text{ucounter}, \text{vcounter}$$

The set AP of atomic propositions will be the names of the local states of the two gadgets  $Gadget[u]$  and  $Gadget[v]$  plus an atomic proposition goal.

**Transition system for the word-pairs.** Let  $\ell_u$  and  $\ell_v$  be the maximal length of the  $u$ -words resp.  $v$ -words in the given input sequence for PCP. That is:

$$\begin{aligned}\ell_u &= \max\{|u_0|, |u_1|, \dots, |u_{k-1}|\} \\ \ell_v &= \max\{|v_0|, |v_1|, \dots, |v_{k-1}|\}\end{aligned}$$

The transition system  $Words$  has the state space

$$W = \Sigma^{\leq \ell_u} \times \Sigma^{\leq \ell_v} \times \{normal, wait\} \cup \{init, fail\}$$

where  $\Sigma^{\leq \ell}$  denotes the set of words over  $\Sigma$  of length at most  $\ell$ . The components  $normal$  and  $wait$  are called modes of  $Words$ . In what follows, we use simplifying notations for the states in normal mode by dropping the component  $normal$  from the states  $(\alpha_1 \alpha_2 \dots \alpha_\nu, \beta_1 \beta_2 \dots \beta_\mu, normal)$ . That is, states in normal mode will be written simply as pairs  $(\alpha_1 \alpha_2 \dots \alpha_\nu, \beta_1 \beta_2 \dots \beta_\mu)$  of words over  $\Sigma$  with  $\nu \leq \ell_u$  and  $\mu \leq \ell_v$ . Intuitively, if  $(\alpha_1 \alpha_2 \dots \alpha_\nu, \beta_1 \beta_2 \dots \beta_\mu)$  is a state in  $Words$  then  $\alpha_1 \alpha_2 \dots \alpha_\nu$  is a prefix of the next generated subword  $u_{i_{n+1}}$  of the  $u$ -string and  $\beta_1 \beta_2 \dots \beta_\mu$  a prefix of the next generated subword  $v_{j_{m+1}}$  of the  $v$ -string. States in waiting mode have the same meaning, but generating the next symbol is disabled, since  $Words$  is waiting for a synchronization with one of the gadgets. There are two special states in  $Words$ , the initial state  $w_{init} = init$  and the failure state  $fail$ , which indicates that  $Words$  failed to generate a word pair  $(u_{i_1} \dots, u_{i_n}, v_{j_1} \dots, v_{j_m})$ .

The action set of  $Words$  consists of the symbols in  $\Sigma$  (which represents generating the next symbol in the current  $u$ -string and  $v$ -string), the actions  $ustart_i, vstart_i, ustop_i, vstop_i$  for  $0 \leq i < k$  that are used to synchronize with the gadgets and a special action  $\tau$  for the fail-state.

The transition relation of  $Words$  is defined by the following SOS-rules, where  $\varepsilon$  denotes the empty word:

$$\begin{array}{c} \frac{\sigma \in \Sigma}{init \xrightarrow{\sigma} (\sigma, \sigma)} \\ \\ \frac{\nu < \ell_u, \quad \mu < \ell_v, \quad \sigma \in \Sigma}{(\alpha_1 \dots \alpha_\nu, \beta_1 \dots \beta_\mu) \xrightarrow{\sigma} (\alpha_1 \dots \alpha_\nu \sigma, \beta_1 \dots \beta_\mu \sigma)} \\ \\ \frac{\nu = \ell_u, \quad \sigma \in \Sigma}{(\alpha_1 \dots \alpha_\nu, \beta_1 \dots \beta_\mu) \xrightarrow{\sigma} fail} \\ \\ \frac{\mu = \ell_v, \quad \sigma \in \Sigma}{(\alpha_1 \dots \alpha_\nu, \beta_1 \dots \beta_\mu) \xrightarrow{\sigma} fail} \\ \\ \frac{\alpha_1 \dots \alpha_\nu = u_i, \quad 0 \leq i < k}{(\alpha_1 \dots \alpha_\nu, \beta_1 \dots \beta_\mu) \xrightarrow{ustart_i} (\varepsilon, \beta_1 \dots \beta_\mu, wait)} \\ \\ \frac{\beta_1 \dots \beta_\mu = v_j, \quad 0 \leq j < k}{(\alpha_1 \dots \alpha_\nu, \beta_1 \dots \beta_\mu) \xrightarrow{vstart_j} (\alpha_1 \dots \alpha_\nu, \varepsilon, wait)} \end{array}$$

and the following axioms:

$$\begin{aligned}(\varepsilon, \beta_1 \dots \beta_\mu, wait) &\xrightarrow{ustop_i} (\varepsilon, \beta_1 \dots \beta_\mu) \\ (\alpha_1 \dots \alpha_\nu, \varepsilon, wait) &\xrightarrow{vstop_j} (\alpha_1 \dots \alpha_\nu, \varepsilon) \\ fail &\xrightarrow{\tau} fail\end{aligned}$$

**The WTSs for the gadgets.** We now explain the structure of  $Gadget[u]$ . All states of  $Gadget[u]$  have the form  $g_*^u$  where  $*$  stands for some subscript (index). In what follows, we simply write  $g_*$  instead of  $g_*^u$ . The starting state is  $g_{init}$ . The action set of  $Gadget[u]$  consists of the actions:

- $ustart_i$  and  $ustop_i$  to be synchronized with  $Words$ ,
- a special action symbol  $silent = silent^u$  for silent transitions that have no effect on the weights, i.e.,  $wgt(g_*, silent) = 0$  for all weight functions  $wgt$  and states  $g_*$  of  $Gadget[u]$
- action  $start = start^u$  used for the initial assignment  $ufactor := 1$
- action  $reset = reset^u$  used to clear the counter which corresponds to the assignment  $ucounter := 0$
- actions  $add_i = add_i^u$  for  $i \in \{0, 1, \dots, k-1\}$  used to realize the assignment  $useq := useq + ufactor \cdot i$
- actions  $incr = incr^u$ ,  $clear = clear^u$  and  $swap = swap^u$  used to realize the assignment  $ufactor := k \cdot ufactor$

We simply write  $g \rightarrow g'$  rather than  $g \xrightarrow{silent} g'$ . All actions except for  $ustart_i$  and  $ustop_i$  are internal actions, not to be synchronized with  $Words$  or  $Gadget[v]$ . To distinguish the internal actions of  $Gadget[u]$  from those in  $Gadget[v]$ , the names of the internal actions are parametrized by  $u$ . For better readability, we drop this parametrization in the following description of the behavior of  $Gadget[u]$ .  $Gadget[u]$  operates in several modes:

**Mode 0: Initialization.** The initialization phase consists of the assignment  $ufactor := 1$ , realized by the transition

$$g_{init} \xrightarrow{start} g_1$$

and by setting  $ufactor(g_{init}, start) = 1$ .

**Mode 1: Synchronize with the TS for the word-pairs.** The ‘‘computation’’ of  $Gadget[u]$  is invoked by synchronizing with  $Words$  over one of the actions  $ustart_i$  where  $i \in \{1, \dots, k\}$  by means of the transition:

$$g_1 \xrightarrow{ustart_i} g_{2,i}$$

**Mode 2: Reset counter.** The assignment  $ucounter := 0$  is realized by the transitions

$$\begin{aligned}g_{2,i} &\rightarrow g_{2,i,loop}, \\ g_{2,i,loop} &\xrightarrow{reset} g_{2,i,loop}, \\ g_{2,i,loop} &\rightarrow g_{3,i}\end{aligned}$$

where  $ucounter(g_{2,i,loop}, reset) = -1$  with the side-constraint:

$$\psi_{2,i}^u = \boxplus^{A[\dots g_{3,i}]}(ucounter = 0)$$

**Mode 3: Assignment for the encoding of the index sequence for the  $u$ -word.** We mimic the effect of the assignment  $useq := useq + ufactor \cdot i$  by the transitions

$$\begin{aligned}g_{3,i} &\rightarrow g_{3,i,loop} \\ g_{3,i,loop} &\xrightarrow{add_i} g_{3,i,loop} \\ g_{3,i,loop} &\rightarrow g_{4,i}\end{aligned}$$

where  $ucounter(g_{3,i,loop}, add_i) = 1$  and  $useq(g_{3,i,loop}, add_i) = i$  with the side-constraint:

$$\psi_{3,i}^u = \boxplus^{A[\dots g_{4,i}]}(ucounter = ufactor)$$

**Mode 4: Reset counter.** As in mode 2, the purpose of mode 4 is to realize the assignment  $ucounter := 0$  by means of the



transitions

$$\begin{aligned} g_{4,i} &\longrightarrow g_{4,i,loop} \\ g_{4,i,loop} &\xrightarrow{\text{reset}} g_{4,i,loop} \\ g_{4,i,loop} &\longrightarrow g_{5,i} \end{aligned}$$

where  $ucounter(g_{4,i,loop}, \text{reset}) = -1$  with the side-constraint:

$$\psi_{4,i}^u = \boxplus^{\mathcal{A}[\dots g_{5,i}]}(\text{ucounter} = 0)$$

**Mode 5: Shift.** The assignment  $ufactor := k \cdot ufactor$  has the same effect as the following sequence of instructions:

1. increment  $ucounter$  until  $ucounter = k \cdot ufactor$ ;
2. clear  $ufactor$  by setting  $ufactor := 0$ ;
3. swap the values of  $ucounter$  and  $ufactor$

The third instruction (swap) can be realized by performing simultaneously the assignments  $ucounter := ucounter - 1$  and  $ufactor := ufactor + 1$  until  $ucounter = 0$ . In  $Gadget[u]$ , the three instructions are simulated by the transitions:

$$\begin{aligned} g_{5,i} &\longrightarrow g_{5,i,loop} \\ g_{5,i,loop} &\xrightarrow{\text{incr}} g_{5,i,loop} \\ g_{5,i,loop} &\longrightarrow g_{6,i} \\ \\ g_{6,i} &\longrightarrow g_{6,i,loop} \\ g_{6,i,loop} &\xrightarrow{\text{clear}} g_{6,i,loop} \\ g_{6,i,loop} &\longrightarrow g_{7,i} \\ \\ g_{7,i} &\longrightarrow g_{7,i,loop} \\ g_{7,i,loop} &\xrightarrow{\text{swap}} g_{7,i,loop} \\ g_{7,i,loop} &\longrightarrow g_{8,i} \end{aligned}$$

where

$$\begin{aligned} ucounter(g_{5,i,loop}, \text{incr}) &= 1 \\ ufactor(g_{6,i,loop}, \text{clear}) &= -1 \\ ucounter(g_{7,i,loop}, \text{swap}) &= -1 \\ ufactor(g_{7,i,loop}, \text{swap}) &= 1 \end{aligned}$$

with the side-constraints:

$$\begin{aligned} \psi_{5,i}^u &= \boxplus^{\mathcal{A}[\dots g_{6,i}]}(\text{ucounter} = k \cdot ufactor) \\ \psi_{6,i}^u &= \boxplus^{\mathcal{A}[\dots g_{7,i}]}(\text{ufactor} = 0) \\ \psi_{7,i}^u &= \boxplus^{\mathcal{A}[\dots g_{8,i}]}(\text{ucounter} = 0) \end{aligned}$$

**Mode 6: Return.** In its last mode, the gadget moves back to mode 1 by the transition  $q_{8,i} \xrightarrow{\text{ustop}_i} q_1$ .

So far, we defined the values of the weight functions only in selected cases. In all remaining cases, the values of the weight functions are 0. The WTS  $Gadget[v]$  for the  $v$ -words is defined analogously.

**The composite WTS.** The WTS  $\mathcal{T}$  arises by the parallel composition of the transition system  $Words$  and the two gadgets  $Gadget[u]$  and  $Gadget[v]$ . Thus, the states in  $\mathcal{T}$  have the form

$$s = \langle w, g_x^u, g_y^v \rangle$$

where  $w \in W$  is a state of  $Words$ ,  $g_x^u$  is a state of  $Gadget[u]$  and  $g_y^v$  a state of  $Gadget[v]$ . The initial state of  $\mathcal{T}$  is:

$$s_{init} = \langle \text{init}, g_{init}^u, g_{init}^v \rangle$$

In addition,  $\mathcal{T}$  contains states of the form  $\langle \text{fail}, g_x^u, g_y^v \rangle$ .

The action set of  $\mathcal{T}$  is the union of the action sets of  $Words$  and the two gadgets. The actions  $ustart_i, ustop_i$  with  $i \in \{0, 1, \dots, k-1\}$  are performed synchronously in  $Words$  and

$Gadget[u]$ . Similarly,  $Words$  synchronizes with  $Gadget[v]$  over the actions  $vstart_j$  and  $vstop_j$  for  $j \in \{0, 1, \dots, k-1\}$ . All other actions of the three systems are executed in an interleaved way.

As stated above,  $\mathcal{T}$  is equipped with six weight functions  $ucounter, vcounter, ufactor, vfactor, useq$  and  $vseq$  that arise by the liftings of the corresponding weight functions in the two gadgets. Suppose  $s = \langle w, g_x^u, g_y^v \rangle$  is a state of  $\mathcal{T}$ , then for example:

- If  $g_x^u = g_{init}^u$  then  $ufactor(s, start^u) = 1$  and  $wgt(s, start^u) = 0$  for the other five weight functions.
- If  $g_x^u = g_{7,i,loop}^u$  then  $ufactor(s, swap^u) = 1$  and  $ucounter(s, swap^u) = -1$ . For the other four weight functions, we have  $wgt(s, swap^u) = 0$ .

The set AP of atomic propositions of  $\mathcal{T}$  are the names of the states in the gadgets plus the special proposition goal that characterizes the states where the  $Words$ -component is  $(\varepsilon, \varepsilon)$ . The labeling function  $L$  is defined by the following conditions:

$$\begin{aligned} g_x^u \in L(s) &\text{ iff } s \text{ has the form } \langle w, g_x^u, g_y^v \rangle \\ g_y^v \in L(s) &\text{ iff } s \text{ has the form } \langle w, g_x^u, g_y^v \rangle \\ \text{goal} \in L(s) &\text{ iff } s \text{ has the form } \langle (\varepsilon, \varepsilon), g_1^u, g_1^v \rangle \end{aligned}$$

**Formula.** The PL $[\Diamond : \text{Reach}]$ -formula  $\varphi$  is defined as follows:

$$\varphi = \boxplus^{\mathcal{A}[\dots \text{goal}]}(\text{useq} = \text{vseq}) \wedge \psi^u \wedge \psi^v$$

where for  $\xi \in \{u, v\}$ :

$$\begin{aligned} \psi^\xi &= \bigwedge_{0 \leq i < k} \left( \psi_{2,i}^\xi \wedge \psi_{3,i}^\xi \wedge \psi_{4,i}^\xi \right. \\ &\quad \left. \wedge \psi_{5,i}^\xi \wedge \psi_{6,i}^\xi \wedge \psi_{7,i}^\xi \right) \end{aligned}$$

**Soundness.** We now have to show that  $s_{init} \models \exists \varphi$  if and only if the given instance  $(u_0, v_0), (u_1, v_1), \dots, (u_{k-1}, v_{k-1})$  for PCP has a solution.

“ $\implies$ ”: Let the infinite path  $\zeta = s_0 \text{ act}_0 s_1 \text{ act}_1 \dots$  be a witness for  $s_{init} \models \exists \varphi$ , i.e.,  $s_0 = s_{init}$  and  $\zeta \models \varphi$ . We have to show that there is then a corresponding sequence  $i_1, \dots, i_n$  such that  $u_{i_1} u_{i_2} \dots u_{i_n} = v_{i_1} v_{i_2} \dots v_{i_n}$ .

As  $\zeta$  satisfies  $\varphi$ , in particular  $\zeta \models \boxplus^{\mathcal{A}[\dots \text{goal}]}(\text{useq} = \text{vseq})$ . Thus there exists a (smallest) index  $n$  such that  $s_n \models \text{goal}$  and  $useq(\zeta[0 \dots n]) = vseq(\zeta[0 \dots n])$ . Let  $\pi = \zeta[0 \dots n]$ . W.l.o.g., we assume that the first two actions in  $\pi$  are the initializing actions of  $Gadget[u]$  and  $Gadget[v]$ , i.e.,  $act_0 = start_u$  and  $act_1 = start_v$  and thus  $s_2 = \langle \text{init}, g_1^u, g_1^v \rangle$ . If this is not the case, we can move these actions to the beginning of  $\zeta$  without affecting the satisfaction of  $\varphi$ .

Due to the structure of  $\mathcal{T}$ , we can partition the sequence of actions  $act_2 \dots act_n$  into subsequences, each consisting of a finite number of symbol actions  $\sigma \in \Sigma$  of  $Words$  and followed by a sequence  $ustart_i \dots ustop_i$  consisting of  $Gadget[u]$ -actions for some  $i$  or a sequence  $vstart_j \dots vstop_j$  consisting of  $Gadget[v]$ -actions for some  $j$ , or both. As the  $ustart_i$  and  $vstart_j$  actions switch  $Words$  into the *wait* mode, interleaving is suspended until the next occurrence of  $ustop_i$  ( $vstop_j$ , respectively). As the local state of  $Words$  of the goal state has the form  $(\varepsilon, \varepsilon)$ , at least one  $ustart_i \dots ustop_i$  sequence and one  $vstart_j \dots vstop_j$  sequence has to occur, as this is the only way for the words to become empty. In particular, the last action in  $\pi$  has either the form  $ustop_i$  or  $vstop_j$ .

Let  $w$  be the sequence of  $\Sigma$ -actions in  $\pi$  and let  $i_1, \dots, i_m$  be the sequence of indices  $i$  of the  $ustop_i$  actions in  $\pi$ . Then, by construction,  $w = u_{i_1} u_{i_2} \dots u_{i_m}$ . Similarly, let  $j_1, \dots, j_{m'}$  be the sequence of indices  $j$  of the  $vstop_j$  actions in  $\pi$ . Then again, by

construction,  $w = v_{j_1} v_{j_2} \dots v_{j_{m'}}$ . It remains to show that  $m = m'$  and  $i_1 = j_1, i_2 = j_2, \dots, i_m = j_{m'}$ .

Let  $\pi_n$  be the prefix of  $\pi$  up to the  $n$ -th occurrence of an *ustop*-action. As above, the index of this action is denoted by  $i_n$ . We consider the evaluation of the accumulated weights along these prefixes, as enforced by satisfaction of  $\psi^u$ . For  $n = 1$ , we have  $useq(\pi_1) = i_1$  and  $ufactor(\pi_1) = k$ , while for  $n > 1$  we have

$$\begin{aligned} useq(\pi_{n+1}) &= useq(\pi_n) + i_{n+1} \cdot k^n \\ ufactor(\pi_{n+1}) &= k \cdot ufactor(\pi_n) = k^{n+1} \end{aligned}$$

Applying the same arguments to *vseq* and *vfactor*, we obtain

$$\begin{aligned} useq(\pi) &= i_m \cdot k^{m-1} + i_{m-1} \cdot k^{m-2} + \dots + i_2 \cdot k + i_1 \\ vseq(\pi) &= j_{m'} \cdot k^{m'-1} + j_{m'-1} \cdot k^{m'-2} + \dots + j_2 \cdot k + j_1 \end{aligned}$$

On the other hand, by the construction of  $\pi$ , we have  $useq(\pi) = vseq(\pi)$ . As all the  $i$  and  $j$  are elements of  $\{0, \dots, k-1\}$  it follows that  $m = m'$  and  $i_1 = j_1, i_2 = j_2, \dots, i_m = j_m$ . Thus,  $i_1, \dots, i_m$  is a solution for the given PCP instance.

“ $\Leftarrow$ ”: Let  $i_1, \dots, i_m$  be a solution of the given PCP instance, with  $w = u_{i_1} u_{i_2} \dots u_{i_m}$ . We construct an initial, infinite path  $\zeta$  of  $\mathcal{T}$  such that  $\zeta \models \varphi$ , thus serving as a witness for  $s_{init} \models \exists \varphi$ . As the sequence of actions uniquely determines the successor states in  $\mathcal{T}$ , we only specify the sequence of actions  $act_0 act_1 \dots$  of  $\zeta$ . The initial actions  $act_0 = start^u$  and  $act_1 = start^v$  initialize *Gadget*[ $u$ ] and *Gadget*[ $v$ ], respectively. Let  $act_0 \dots act_n$  be the sequence of actions already chosen and let  $\pi_n = s_0 act_0 \dots act_n s_n$  be the corresponding path in  $\mathcal{T}$ . Action  $act_{n+1}$  is then chosen as follows:

1. If  $act_n = ustart_i$  for some  $i$ , choose the sequence of *Gadget*[ $u$ ]-actions ending with  $ustop_i$  that lead back to  $g_1$ , while ensuring that  $\psi^u$  is satisfied for the resulting path.
2. Else, if  $act_n = vstart_j$  for some  $j$ , choose the sequence of *Gadget*[ $v$ ]-actions ending with  $vstop_j$  that lead back to  $g_1$ , while ensuring that  $\psi^v$  is satisfied for the resulting path.
3. Else, if the sequence of  $\Sigma$ -actions in  $act_0 \dots act_n$  equals  $u_{i_1} \dots u_{i_{m'}}$  for some  $m'$ , chose  $ustart_{i_{m'}}$ .
4. Else, if the sequence of  $\Sigma$ -actions in  $act_0 \dots act_n$  equals  $v_{i_1} \dots v_{i_{m'}}$  for some  $m'$ , chose  $vstart_{i_{m'}}$ .
5. Else, let  $n$  be the number of  $\Sigma$ -actions in  $act_0 \dots act_n$ . If  $|w| > n$ , choose the  $n+1$ -th symbol of  $w$ .
6. Else, finish.

This construction terminates, as the length of  $w$  is finite. Let  $\pi$  be the constructed finite path. By construction, the last state of  $\pi$  is labeled with *goal*, i.e.,  $s_n = \langle (\varepsilon, \varepsilon), g_1^u, g_1^v \rangle$  and  $useq(\pi) = vseq(\pi)$ . We can extend  $\pi$  to an infinite path  $\zeta$  by choosing some arbitrary action  $\sigma \in \Sigma$  as long as the *fail* state has not been reached. From then on, we always choose the  $\tau$  action. By construction,  $\zeta \models \diamond^{A[\dots goal]}(useq = vseq)$ . Additionally,  $\zeta \models \psi^u$ , as during construction of  $\pi$  care is taken to ensure satisfaction. After  $\pi$ , none of the states/labels monitored by  $\psi^u$  are reached and thus  $\psi^u$  is satisfied as well for all prefixes of  $\zeta$ . The same arguments apply to the satisfaction of  $\psi^v$ . Thus,  $\zeta \models \varphi$  and  $s_{init} \models \exists \varphi$ .

**Proof sketch for the undecidability over WMCs.** To show the undecidability of the qualitative model-checking problem for  $PL[\diamond : Reach]$  and WMCs, we can use fairly the same construction as above and resolve all nondeterministic choices in the constructed WTS  $\mathcal{T}$  by uniform distributions. Some care is needed to ensure that the weights can be attached to states (rather than state-action pairs). This problem is of technical nature only, since for each transition  $s \xrightarrow{act} s'$  in  $\mathcal{T}$  we can introduce an auxiliary state  $s_{act}$  with  $wgt_{\mathcal{M}}(s_{act}) = wgt_{\mathcal{T}}(s, act)$  and  $P_{\mathcal{M}}(s_{act}, s') = 1$ . For the resulting Markov chain  $\mathcal{M}$  we then have:  $\Pr_{\mathcal{M}, s_{init}}(\varphi) > 0$  if and

only if the given instance  $(u_0, v_0), (u_1, v_1), \dots, (u_{k-1}, v_{k-1})$  for PCP is solvable. ■

**Remark 34** The model-checking problem for the logic  $PL[\diamond : All]$  is trivially decidable, since Boolean combinations of  $\diamond$ -formulas can be seen as conditions on the labeling of the initial state. However, the model-checking problem for the logic  $PL[\diamond : InitReach]$  consisting of Boolean combinations of formulas of the form  $\diamond \diamond^{A[\phi_1 \dots \phi_2]} constr$  is undecidable, too. Here, if  $\phi_1$  and  $\phi_2$  are propositional formulas with atoms in AP then  $A[\phi_1 \dots \phi_2]$  denotes the minimal DFA for the language consisting of all finite words over the alphabet  $2^{AP}$  starting with a symbol A where  $A \models \phi_1$  and ending with a symbol B where  $B \models \phi_2$ . Let  $A[\phi \dots true]$  be a short form notation for  $A[\phi \dots true]$ .

Indeed, in the proof of Theorem 33 for WTSs and WMCs, we can replace  $PL[\diamond : Reach]$  with  $PL[\diamond : InitReach]$ . We then simply replace

- the formulas  $\psi_{h,i}^\xi = \boxplus^{A[\dots g]} constr$  with

$$\psi_{h,i}^\xi = \boxplus^{A[g_{init} \dots]} constr$$

- the formula  $\diamond^{A[\dots goal]}(useq = vseq)$  with

$$\diamond \diamond^{A[g_{init} \dots goal]}(useq = vseq)$$

The construction of the WTS and WMC from the given instance of the PCP remains unchanged. ■

A WMDP  $\mathcal{M} = (S, Act, P, AP, L, wgt_1, \dots, wgt_d)$  is called *non-negative* if  $wgt_i(s, act) \geq 0$  for all state-action pairs  $(s, \alpha)$  and  $i = 1, \dots, d$ .

**Theorem 35 (See Thm. 18)**

The LTL<sub>simple</sub> $[\diamond, \diamond : All]$ -PMC problem is decidable for non-negative WMDPs.

**Proof.** Let  $\varphi$  be an LTL<sub>simple</sub> $[\diamond, \diamond : All]$ -formula,

$$\mathcal{M} = (S, Act, P, AP, L, wgt_1, \dots, wgt_d)$$

a non-negative WMDP and  $s_{init} \in S$ . By Remark 7, we can safely assume that  $wgt_i(s, act) \in \mathbb{N}$  for all state-action pairs.

We now describe how to modify the construction as in Section 4.1 to provide a reduction to the LTL-PMC problem. Again, the states of the constructed MDP  $\tilde{\mathcal{M}}$  have the form

$$\tilde{s} = \langle s, f_1, \dots, f_m \rangle$$

where  $s$  is a state in  $\mathcal{M}$  and with one function  $f_i$  per automaton  $\mathcal{A}_i$  occurring in some subformula of  $\varphi$  of the form  $\diamond^{A_i}(\varphi_1; constr; \varphi_2)$  or  $\diamond^{A_i}(\varphi_1; constr; \varphi_2)$ . As  $\varphi$  is an LTL<sub>simple</sub> $[\diamond, \diamond : All]$ -formula, the *constr* are boolean combinations of simple weight constraints  $wgt_j \bowtie c$ . For each weight function  $wgt_j$ , let  $c_j^{\max}$  be the largest  $c$  that occurs in any of the weight constraints  $wgt_j \bowtie c$  referencing  $wgt_j$ . Any value above this threshold  $c_j^{\max}$  can not be distinguished by the weight constraints referencing  $wgt_j$  in  $\varphi$  and can thus be abstracted by a single symbol  $\top$ . Let

$$W_j = \{0, 1, \dots, c_j^{\max}\} \cup \{\top\}$$

be the weight values relevant in  $\varphi$  for weight function  $wgt_j$  and let

$$QW_i = Q_i \times W_1 \times \dots \times W_d$$

be the set of tuples  $(q, \bar{w})$  encoding a state of  $\mathcal{A}_i$  and accumulated values for the weight functions  $wgt_1, \dots, wgt_d$ .

In contrast to the construction in Section 4.1, we can no longer rely on the fact that the runs in the DFA have a finite length. Instead, we *merge* runs that are in the same state  $q$  and have the same accumulated weights  $\bar{w}$ . Formally, the function  $f_i$  corresponding to

$\mathcal{A}_i$  is now a partial function of the type

$$f_i: \{0, 1, \dots, |QW_i| + 2\} \rightarrow QW_i$$

with  $f(k) = \perp$  for at least one  $k$ .

The successors  $f'_i(k)$  for  $(act, s')$  are computed as before, with the following differences. Any accumulated value above  $c_j^{\max}$  is abstracted to  $\top$ . Furthermore, if  $f_i(k) = f_i(h)$  for some  $h < k$ , then  $f'_i(k) = \perp$ , i.e., among all runs with the same state and accumulated weights, only the one with the smallest index  $h$  remains. The other runs are merged with the run tracked by  $f_i(h)$ . This ensures that there is always at least one “free” index  $k$  with  $f_i(k) = \perp$ .

Formally, the  $(act, s')$ -successor of  $\langle s, f_i \rangle$  in  $\mathcal{A}_i$  is obtained for  $k \in \{0, 1, \dots, |QW_i| + 2\}$  by the following procedure:

- If  $f_i(k) = (q, \bar{w})$  and  $f_i(h) = f_i(h)$  for some  $h < k$ , then  $f'_i(k) = \perp$ .
- Else, if  $f_i(k) = (q, \bar{w})$  where  $q \in Q$  and  $q \neq q_{init}$  then  $f'_i(k) = (\delta(q, L(s')), \bar{w} + \overline{wgt}(s, act))$ . Here, if the result of the addition for element  $w_j$  is larger than the corresponding threshold  $c_j^{\max}$  then the result is replaced by  $\top$ .
- If  $k$  is the smallest index such that  $f_i(k) = \perp$  then  $f'_i(k) = (\delta(q_{init}, L(s')), \bar{0})$ .
- In all other cases:  $f'_i(k) = \perp$ .

The initial functions are obtained analogously to the construction in Section 4.1.

To keep track of the merges, we introduce additional atomic propositions  $\text{merge}_i(k, h)$  with the intuitive meaning “run  $k$  is merged with run  $h$  in  $f_i$ ”. Thus, a state  $\tilde{s} = \langle s, f_1, \dots, f_m \rangle$  is labeled by atomic proposition  $\text{merge}_i(k, h)$  iff  $f_i(k) = f_i(h) \neq \perp$  and there is no  $j < h$  such that  $f_i(j) = f_i(h)$ .

The condition for labeling function  $L$  with respect to the atomic proposition  $\text{goal}_i$  is the same as in Section 4.1 with the obvious evaluation of  $\text{wgt}_j \bowtie c$  according to the value of the corresponding element in  $f_i$ . In contrast to the normal construction, we now not only have to ensure a sequence from  $\text{init}$  via  $\text{run}$  to  $\text{goal}$ , but have to deal with the possibility of multiple merges along the way. Let  $M_i \subseteq \{0, \dots, |QW_i| + 2\}^+$  be the set of nonempty sequences consisting of arguments of  $f_i$  such that:

$$k_n < k_{n-1} < \dots < k_1 \text{ for all } \sigma = k_1 \dots k_n \in M_i$$

Intuitively, each pair of subsequent symbols in such a word  $\sigma$  corresponds to a merge from  $k_h$  to  $k_{h+1}$ .

For each merge sequence  $\sigma = k_1 \dots k_n \in M_i$ , we define an LTL formula with cascades of until formulas, that tracks a run in  $f_i$  along the merges specified by  $\sigma$ .

$$\varphi_i(k_1 \dots k_n) \stackrel{\text{def}}{=} \text{run}_i(k_1) \text{U} (\text{merge}_i(k_1, k_2) \wedge \bigcirc \varphi_i(k_2 \dots k_n))$$

and

$$\varphi_i(k_n) = \text{run}_i(k_n) \text{U} \text{goal}_i(k_n)$$

We then replace each subformula of the form  $\psi_+ = \diamond^{A_i} \text{constr}$  by the formula

$$\tilde{\psi}_+ = \bigvee_{k_1 \dots k_n \in M_i} \text{init}_i(k_1) \wedge \varphi_i(k_1 \dots k_n)$$

For  $\psi_+ = \diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$  we additionally replace in  $\tilde{\psi}_+$  occurrences of  $\text{init}_i(k)$  with  $\varphi_1 \wedge \text{init}_i(k)$  and  $\text{goal}_i(k)$  with  $\text{goal}_i(k) \wedge \varphi_2$ . For the subformulas of the form  $\psi_- = \diamond^{A_i}(\varphi_1; \text{constr}; \varphi_2)$ , we adapt the construction of  $\psi_+$  in the obvious way, analogous to the construction in Section 4.1. ■

**Theorem 36 (See Thm. 20)** *The model-checking problem for  $\text{PL}[\diamond : \text{Reach}]$  and non-negative WTSs is undecidable. Likewise,*

*the model-checking problem for  $\text{PL}[\diamond : \text{Reach}]$  and non-negative WMCs is undecidable.*

**Proof.** We provide a reduction from the model-checking problems for  $\text{PL}[\diamond : \text{Reach}]$  studied in Theorem 17. Let

$$\mathcal{T} = (S, Act, \longrightarrow, AP, L, \text{wgt}_1, \dots, \text{wgt}_d)$$

be a WTS with weight functions  $\text{wgt}_i : S \times Act \rightarrow \mathbb{Q}$  and distinguished state  $s_{init} \in S$  and let  $\varphi$  be a  $\text{PL}[\diamond : \text{Reach}]$  formula. We define a new non-negative WTS:

$$\tilde{\mathcal{T}} = (S, Act, \longrightarrow, AP, L, \text{wgt}_1^+, \dots, \text{wgt}_d^+, \text{wgt}_1^-, \dots, \text{wgt}_d^-)$$

where  $\text{wgt}_i^+, \text{wgt}_i^- : S \times Act \rightarrow \mathbb{Q}_{\geq 0}$  are defined as follows.

- If  $\text{wgt}_i(s, act) > 0$  then  $\text{wgt}_i^+(s, act) = \text{wgt}_i(s, act)$ .
- If  $\text{wgt}_i(s, act) < 0$  then  $\text{wgt}_i^-(s, act) = -\text{wgt}_i(s, act)$ .
- In all other cases:  
 $\text{wgt}_i^+(s, act) = 0$  and  $\text{wgt}_i^-(s, act) = 0$ .

We do an analogous transformation on the formula-level. That is, we modify the given signature by replacing the weight symbols  $\text{wgt}_1, \dots, \text{wgt}_d$  with  $2d$  weight symbols  $\text{wgt}_1^+, \dots, \text{wgt}_d^+, \text{wgt}_1^-, \dots, \text{wgt}_d^-$ . Let  $\tilde{\varphi}$  be the  $\text{PL}[\diamond : \text{Reach}]$  formula over the modified signature that arises from  $\varphi$  by replacing each weight expression  $\sum_{i=1}^d a_i \cdot \text{wgt}_i$  in  $\varphi$  with

$$\sum_{i=1}^d a_i \cdot \text{wgt}_i^+ - \sum_{i=1}^d a_i \cdot \text{wgt}_i^-$$

Obviously, we have:

$$s_{init} \models_{\mathcal{T}} \exists \varphi \text{ iff } s_{init} \models_{\tilde{\mathcal{T}}} \exists \tilde{\varphi}$$

The reduction for WMCs is analogous. ■