Partial and Conditional Expectations in Markov Decision Processes with Integer Weights

Jakob Piribauer^(✉) and Christel Baier

Technische Universität Dresden, Dresden, Germany
{jakob.piribauer, christel.baier}@tu-dresden.de

Abstract. The paper addresses two variants of the stochastic shortest path problem (“optimize the accumulated weight until reaching a goal state”) in Markov decision processes (MDPs) with integer weights. The first variant optimizes partial expected accumulated weights, where paths not leading to a goal state are assigned weight 0, while the second variant considers conditional expected accumulated weights, where the probability mass is redistributed to paths reaching the goal. Both variants constitute useful approaches to the analysis of systems without guarantees on the occurrence of an event of interest (reaching a goal state), but have only been studied in structures with non-negative weights. Our main results are as follows. There are polynomial-time algorithms to check the finiteness of the supremum of the partial or conditional expectations in MDPs with arbitrary integer weights. If finite, then optimal weight-based deterministic schedulers exist. In contrast to the setting of non-negative weights, optimal schedulers can need infinite memory and their value can be irrational. However, the optimal value can be approximated up to an absolute error of ϵ in time exponential in the size of the MDP and polynomial in $\log(1/\epsilon)$.

1 Introduction

Stochastic shortest path (SSP) problems generalize the shortest path problem on graphs with weighted edges. The SSP problem is formalized using finite state Markov decision processes (MDPs), which are a prominent model combining probabilistic and nondeterministic choices. In each state of an MDP, one is allowed to choose nondeterministically from a set of actions, each of them is augmented with probability distributions over the successor states and a weight (cost or reward). The SSP problem asks for a policy to choose actions (here called a scheduler) maximizing or minimizing the expected accumulated weight until reaching a goal state. In the classical setting, one seeks an optimal *proper* scheduler where proper means that a goal state is reached almost surely. Polynomial-time solutions exist exploiting the fact that optimal memoryless deterministic

The authors are supported by the DFG through the Research Training Group QuantLA (GRK 1763), the DFG-project BA-1679/11-1, the Collaborative Research Center HAEC (SFB 912), and the cluster of excellence CeTI.

© The Author(s) 2019

M. Bojańczyk and A. Simpson (Eds.): FOSSACS 2019, LNCS 11425, pp. 436–452, 2019.

https://doi.org/10.1007/978-3-030-17127-8_25

schedulers exist (provided the optimal value is finite) and can be computed using linear programming techniques, possibly in combination with model transformations (see [1, 5, 10]). The restriction to proper schedulers, however, is often too restrictive. First, there are models that have no proper scheduler. Second, even if proper schedulers exist, the expectation of the accumulated weight of schedulers missing the goal with a positive probability should be taken into account as well. Important such applications include the semantics of probabilistic programs (see e.g. [4, 7, 12, 14, 16]) where no guarantee for almost sure termination can be given and the analysis of program properties at termination time gives rise to stochastic shortest (longest) path problems in which the goal (halting configuration) is not reached almost surely. Other examples are the fault-tolerance analysis (e.g., expected costs of repair mechanisms) in selected error scenarios that can appear with some positive, but small probability or the trade-off analysis with conjunctions of utility and cost constraints that are achievable with positive probability, but not almost surely (see e.g. [2]).

This motivates the switch to variants of classical SSP problems where the restriction to proper schedulers is relaxed. One option (e.g., considered in [8]) is to seek a scheduler optimizing the expectation of the random variable that assigns weight 0 to all paths not reaching the goal and the accumulated weight of the shortest prefix reaching the goal to all other paths. We refer to this expectation as *partial expectation*. Second, we consider the *conditional expectation* of the accumulated weight until reaching the goal under the condition that the goal is reached. In general, partial expectations describe situations in which some reward (positive and negative) is accumulated but only retrieved if a certain goal is met. In particular, partial expectations can be an appropriate replacement for the classical expected weight before reaching the goal if we want to include schedulers which miss the goal with some (possibly very small) probability. In contrast to conditional expectations, the resulting scheduler still has an incentive to reach the goal with a high probability, while schedulers maximizing the conditional expectation might reach the goal with a very small positive probability.

Previous work on partial or conditional expected accumulated weights was restricted to the case of non-negative weights. More precisely, partial expectations have been studied in the setting of stochastic multiplayer games with non-negative weights [8]. Conditional expectations in MDPs with non-negative weights have been addressed in [3]. In both cases, optimal values are achieved by weight-based deterministic schedulers that depend on the current state and the weight that has been accumulated so far, while memoryless schedulers are not sufficient. Both [8] and [3] prove the existence of a *saturation point* for the accumulated weight from which on optimal schedulers behave memoryless and maximize the probability to reach a goal state. This yields exponential-time algorithms for computing optimal schedulers using an iterative linear programming approach. Moreover, [3] proves that the threshold problem for conditional expectations (“does there exist a scheduler $\mathfrak{S}$ such that the conditional expectation under $\mathfrak{S}$ exceeds a given threshold?”) is PSPACE-hard even for acyclic MDPs.

The purpose of the paper is to study partial and conditional expected accumulated weights for MDPs with integer weights. The switch from non-negative to integer weights indeed causes several additional difficulties. We start with the following observation. While optimal partial or conditional expectations in non-negative MDPs are rational, they can be irrational in the general setting:

Fig. 1. Enabled actions are denoted by Greek letters and the weight associated to the action is stated after the bar. Probabilistic choices are marked by a bold arc and transition probabilities are denoted next to the arrows.

Example 1. Consider the MDP $\mathcal{M}$ depicted on the left in Fig. 1. In the initial state s_{init} , two actions are enabled. Action τ leads to $goal$ with probability 1 and weight 0. Action σ leads to the states s and t with probability $1/2$ from where we will return to s_{init} with weight -2 or $+1$, respectively. The scheduler choosing τ immediately leads to an expected weight of 0 and is optimal among schedulers reaching the goal almost surely. As long as we choose σ in s_{init} , the accumulated weight follows an asymmetric random walk increasing by 1 or decreasing by 2 with probability $1/2$ before we return to s_{init} . It is well known that the probability to ever reach accumulated weight $+1$ in this random walk is $1/\Phi$ where $\Phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. Likewise, ever reaching accumulated weight n has probability $1/\Phi^n$ for all $n \in \mathbb{N}$. Consider the scheduler $\mathfrak{S}_k$ choosing τ as soon as the accumulated weight reaches k in s_{init} . Its partial expectation is k/Φ^k as the paths which never reach weight k are assigned weight 0. The maximum is reached at $k = 2$. In Sect. 4, we prove that there are optimal schedulers whose decisions only depend on the current state and the weight accumulated so far. With this result we can conclude that the maximal partial expectation is indeed $2/\Phi^2$, an irrational number.

The conditional expectation of $\mathfrak{S}_k$ in $\mathcal{M}$ is k as $\mathfrak{S}_k$ reaches the goal with accumulated weight k if it reaches the goal. So, the conditional expectation is not bounded. If we add a new initial state making sure that the goal is reached with positive probability as in the MDP $\mathcal{N}$, we can obtain an irrational maximal conditional expectation as well: The scheduler $\mathfrak{T}_k$ choosing τ in c as soon as the weight reaches k has conditional expectation $\frac{k/2\Phi^k}{1/2+1/2\Phi^k}$. The maximum is obtained for $k = 3$; the maximal conditional expectation is $\frac{3/\Phi^3}{1+1/\Phi^3} = \frac{3}{3+\sqrt{5}}$.

Moreover, while the proposed algorithms of [3, 8] crucially rely on the monotonicity of the accumulated weights along the prefixes of paths, the accumulated

weights of prefixes of path can oscillate when there are positive and negative weights. As we will see later, this implies that the existence of saturation points is no longer ensured and optimal schedulers might require infinite memory (more precisely, a counter for the accumulated weight). These observations provide evidence why linear-programming techniques as used in the case of non-negative MDPs [3, 8] cannot be expected to be applicable for the general setting.

Contributions. We study the problem of maximizing the partial and conditional expected accumulated weight in MDPs with integer weights. Our first result is that the finiteness of the supremum of partial and conditional expectations in MDPs with integer weights can be checked in polynomial time (Sect. 3). For both variants we show that there are optimal weight-based deterministic schedulers if the supremum is finite (Sect. 4). Although the suprema might be irrational and optimal schedulers might need infinite memory, the suprema can be ϵ -approximated in time exponential in the size of the MDP and polynomial in $\log(1/\epsilon)$ (Sect. 5). By duality of maximal and minimal expectations, analogous results hold for the problem of minimizing the partial or conditional expected accumulated weight. (Note that we can multiply all weights by -1 and then apply the results for maximal partial resp. conditional expectations.)

Related Work. Closest to our contribution is the above mentioned work on partial expected accumulated weights in stochastic multiplayer games with non-negative weights in [8] and on computation schemes for maximal conditional expected accumulated weights in non-negative MDPs [3]. Conditional expected termination time in probabilistic push-down automata has been studied in [11], which can be seen as analogous considerations for a class of infinite-state Markov chains with non-negative weights. The recent work on notions of conditional value at risk in MDPs [15] also studies conditional expectations, but the considered random variables are limit averages and a notion of (non-accumulated) weight-bounded reachability.

2 Preliminaries

We give basic definitions and present our notation. More details can be found in textbooks, e.g. [18].

Notations for Markov Decision Processes. A *Markov decision process* (MDP) is a tuple $\mathcal{M} = (S, Act, P, s_{init}, wgt)$ where S is a finite set of states, Act a finite set of actions, $s_{init} \in S$ the initial state, $P : S \times Act \times S \rightarrow [0, 1] \cap \mathbb{Q}$ is the transition probability function and $wgt : S \times Act \rightarrow \mathbb{Z}$ the weight function. We require that $\sum_{t \in S} P(s, \alpha, t) \in \{0, 1\}$ for all $(s, \alpha) \in S \times Act$. We write $Act(s)$ for the set of actions that are enabled in s , i.e., $\alpha \in Act(s)$ iff $\sum_{t \in S} P(s, \alpha, t) = 1$. We assume that $Act(s)$ is non-empty for all s and that all states are reachable from s_{init} . We call a state absorbing if the only enabled action leads to the state itself with probability 1 and weight 0. The paths of $\mathcal{M}$ are finite or infinite sequences $s_0 \alpha_0 s_1 \alpha_1 s_2 \alpha_2 \dots$ where states and actions alternate such that $P(s_i, \alpha_i, s_{i+1}) > 0$ for all $i \geq 0$. If $\pi = s_0 \alpha_0 s_1 \alpha_1 \dots \alpha_{k-1} s_k$ is finite, then

$wgt(\pi) = wgt(s_0, \alpha_0) + \dots + wgt(s_{k-1}, \alpha_{k-1})$ denotes the accumulated weight of π , $P(\pi) = P(s_0, \alpha_0, s_1) \dots P(s_{k-1}, \alpha_{k-1}, s_k)$ its probability, and $last(\pi) = s_k$ its last state. The *size* of $\mathcal{M}$, denoted $size(\mathcal{M})$, is the sum of the number of states plus the total sum of the logarithmic lengths of the non-zero probability values $P(s, \alpha, s')$ as fractions of co-prime integers and the weight values $wgt(s, \alpha)$.

Scheduler. A (*history-dependent, randomized*) scheduler for $\mathcal{M}$ is a function $\mathfrak{S}$ that assigns to each finite path π a probability distribution over $Act(last(\pi))$. $\mathfrak{S}$ is called *memoryless* if $\mathfrak{S}(\pi) = \mathfrak{S}(\pi')$ for all finite paths π, π' with $last(\pi) = last(\pi')$, in which case $\mathfrak{S}$ can be viewed as a function that assigns to each state s a distribution over $Act(s)$. $\mathfrak{S}$ is called *deterministic* if $\mathfrak{S}(\pi)$ is a Dirac distribution for each path π , in which case $\mathfrak{S}$ can be viewed as a function that assigns an action to each finite path π . Scheduler $\mathfrak{S}$ is said to be *weight-based* if $\mathfrak{S}(\pi) = \mathfrak{S}(\pi')$ for all finite paths π, π' with $wgt(\pi) = wgt(\pi')$ and $last(\pi) = last(\pi')$. Thus, deterministic weight-based schedulers can be viewed as functions that assign actions to state-weight-pairs. By $HR^{\mathcal{M}}$ we denote the class of all schedulers, by $WR^{\mathcal{M}}$ the class of weight-based schedulers, by $WD^{\mathcal{M}}$ the class of weight-based, deterministic schedulers, and by $MD^{\mathcal{M}}$ the class of memoryless deterministic schedulers. Given a scheduler $\mathfrak{S}$, $\varsigma = s_0 \alpha_0 s_1 \alpha_1 \dots$ is a $\mathfrak{S}$ -path iff ς is a path and $\mathfrak{S}(s_0 \alpha_0 s_1 \alpha_1 \dots \alpha_{k-1} s_k)(\alpha_k) > 0$ for all $k \geq 0$.

Probability Measure. We write $\Pr_{\mathcal{M},s}^{\mathfrak{S}}$ or briefly $\Pr_s^{\mathfrak{S}}$ to denote the probability measure induced by $\mathfrak{S}$ and s . For details, see [18]. We will use LTL-like formulas to denote measurable sets of paths and also write $\Diamond(wgt \bowtie x)$ to describe the set of infinite paths having a prefix π with $wgt(\pi) \bowtie x$ for $x \in \mathbb{Z}$ and $\bowtie \in \{<, \leq, =, \geq, >\}$. Given a measurable set ψ of infinite paths, we define $\Pr_{\mathcal{M},s}^{\min}(\psi) = \inf_{\mathfrak{S}} \Pr_{\mathcal{M},s}^{\mathfrak{S}}(\psi)$ and $\Pr_{\mathcal{M},s}^{\max}(\psi) = \sup_{\mathfrak{S}} \Pr_{\mathcal{M},s}^{\mathfrak{S}}(\psi)$ where $\mathfrak{S}$ ranges over all schedulers for $\mathcal{M}$. Throughout the paper, we suppose that the given MDP has a designated state *goal*. Then, $p_s^{\max}$ and $p_s^{\min}$ denote the maximal resp. minimal probability of reaching *goal* from s . That is, $p_s^{\max} = \sup_{\mathfrak{S}} \Pr_s^{\mathfrak{S}}(\Diamond goal)$ and $p_s^{\min} = \inf_{\mathfrak{S}} \Pr_s^{\mathfrak{S}}(\Diamond goal)$. Let $Act^{\max}(s) = \{\alpha \in Act(s) \mid \sum_{t \in S} P(s, \alpha, t) \cdot p_t^{\max} = p_s^{\max}\}$, and $Act^{\min}(s) = \{\alpha \in Act(s) \mid \sum_{t \in S} P(s, \alpha, t) \cdot p_t^{\min} = p_s^{\min}\}$.

Mean Payoff. A well-known measure for the long-run behavior of a scheduler $\mathfrak{S}$ in an MDP $\mathcal{M}$ is the *mean payoff*. Intuitively, the mean payoff is the amount of weight accumulated per step on average in the long run. Formally, we define the mean payoff as the following random variable on infinite paths $\zeta = s_0 \alpha_0 s_1 \alpha_1 \dots$: $MP(\zeta) := \liminf_{k \rightarrow \infty} \frac{\sum_{i=0}^k wgt(s_i, \alpha_i)}{k+1}$. The mean payoff of the scheduler $\mathfrak{S}$ starting in s_{init} is then defined as the expected value $\mathbb{E}_{s_{init}}^{\mathfrak{S}}(MP)$. The maximal mean payoff is the supremum over all schedulers which is equal to the maximum over all MD -schedulers: $\mathbb{E}_{s_{init}}^{\max}(MP) = \max_{\mathfrak{S} \in MD} \mathbb{E}_{s_{init}}^{\mathfrak{S}}(MP)$. In strongly connected MDPs, the maximal mean payoff does not depend on the initial state.

End Components, MEC-Quotient. An *end component* of $\mathcal{M}$ is a strongly connected sub-MDP. End components can be formalized as pairs $\mathcal{E} = (E, \mathfrak{A})$ where E is a nonempty subset of S and $\mathfrak{A}$ a function that assigns to each state $s \in E$ a nonempty subset of $Act(s)$ such that the graph induced by $\mathcal{E}$ is strongly

connected. $\mathcal{E}$ is called *maximal* if there is no end component $\mathcal{E}' = (E', \mathfrak{A}')$ with $\mathcal{E} \neq \mathcal{E}'$, $E \subseteq E'$ and $\mathfrak{A}(s) \subseteq \mathfrak{A}'(s)$ for all $s \in E$. The *MEC-quotient* of an MDP $\mathcal{M}$ is the MDP $MEC(\mathcal{M})$ arising from $\mathcal{M}$ by collapsing all states that belong to the same maximal end component $\mathcal{E}$ to a state $s_{\mathcal{E}}$. All actions enabled in some state in $\mathcal{E}$ not belonging to $\mathcal{E}$ are enabled in $s_{\mathcal{E}}$. Details and the formal construction can be found in [9]. We call an end component $\mathcal{E}$ *positively weight-divergent* if there is a scheduler $\mathfrak{S}$ for $\mathcal{E}$ such that $\Pr_{\mathcal{E},s}^{\mathfrak{S}}(\Diamond(wgt \geq n)) = 1$ for all $s \in \mathcal{E}$ and $n \in \mathbb{N}$. In [1], it is shown that the existence of positively weight-divergent end components can be decided in polynomial time.

3 Partial and Conditional Expectations in MDPs

We define *partial* and *conditional expectations* in MDPs. We extend the definition of [8] by introducing partial expectations with *bias* which are closely related to conditional expectations. Afterwards, we sketch the computation of maximal partial expectations in MDPs with non-negative weights and in Markov chains.

Partial and Conditional Expectation. In the sequel, let $\mathcal{M}$ be an MDP with a designated absorbing goal state *goal*. Furthermore, we collapse all states from which *goal* is not reachable to one absorbing state *fail*. Let $b \in \mathbb{R}$. We define the random variable $\oplus^b goal$ on infinite paths ζ by

$$\oplus^b goal(\zeta) = \begin{cases} wgt(\zeta) + b & \text{if } \zeta \models \Diamond goal, \\ 0 & \text{if } \zeta \not\models \Diamond goal. \end{cases}$$

We call the expectation of this random variable under a scheduler $\mathfrak{S}$ the *partial expectation with bias b* of $\mathfrak{S}$ and write $PE_{\mathcal{M},s_{init}}^{\mathfrak{S}}[b] := \mathbb{E}_{\mathcal{M},s_{init}}^{\mathfrak{S}}(\oplus^b goal)$ as well as $PE_{\mathcal{M},s_{init}}^{\sup}[b] := \sup_{\mathfrak{S} \in HR^{\mathcal{M}}} PE_{\mathcal{M},s_{init}}^{\mathfrak{S}}[b]$. If $b = 0$, we sometimes drop the argument b ; if $\mathcal{M}$ is clear from the context, we drop the subscript. In order to maximize the partial expectation, intuitively one has to find the right balance between reaching *goal* with high probability and accumulating a high positive amount of weight before reaching *goal*. The bias can be used to shift this balance by additionally rewarding or penalizing a high probability to reach *goal*.

The *conditional expectation* of $\mathfrak{S}$ is defined as the expectation of $\oplus^0 goal$ under the condition that *goal* is reached. It is defined if $\Pr_{\mathcal{M},s_{init}}^{\mathfrak{S}}(\Diamond goal) > 0$. We write $CE_{\mathcal{M},s_{init}}^{\mathfrak{S}} := \mathbb{E}_{\mathcal{M},s_{init}}^{\mathfrak{S}}(\oplus^0 goal | \Diamond goal)$ and $CE_{\mathcal{M},s_{init}}^{\sup} = \sup_{\mathfrak{S}} CE_{\mathcal{M},s_{init}}^{\mathfrak{S}}$ where the supremum is taken over all schedulers $\mathfrak{S}$ with $\Pr_{\mathcal{M},s_{init}}^{\mathfrak{S}}(\Diamond goal) > 0$. We can express the conditional expectation as $CE_{\mathcal{M},s_{init}}^{\mathfrak{S}} = PE_{\mathcal{M},s_{init}}^{\mathfrak{S}} / \Pr_{\mathcal{M},s_{init}}^{\mathfrak{S}}(\Diamond goal)$. The following proposition establishes a close connection between conditional expectations and partial expectations with bias.

Proposition 2. *Let $\mathcal{M}$ be an MDP, $\mathfrak{S}$ a scheduler with $\Pr_{s_{init}}^{\mathfrak{S}}(\Diamond goal) > 0$, $\theta \in \mathbb{Q}$, and $\bowtie \in \{<, \leq, \geq, >\}$. Then we have $PE_{s_{init}}^{\mathfrak{S}}[-\theta] \bowtie 0$ iff $CE_{s_{init}}^{\mathfrak{S}} \bowtie \theta$. Further, if $\Pr_{s_{init}}^{\min}(\Diamond goal) > 0$, then $PE_{s_{init}}^{\sup}[-\theta] \bowtie 0$ iff $CE_{s_{init}}^{\sup} \bowtie \theta$.*

Proof. The first claim follows from $PE_{s_{init}}^{\mathfrak{S}}[-\theta] = PE_{s_{init}}^{\mathfrak{S}}[0] - \Pr_{s_{init}}^{\mathfrak{S}}(\Diamond goal) \cdot \theta$. The second claim follows by quantification over all schedulers. $\square$

In [3], it is shown that deciding whether $CE_{s_{init}}^{\sup} \bowtie \theta$ for $\bowtie \in \{<, \leq, \geq, >\}$ and $\theta \in \mathbb{Q}$ is PSPACE-hard even for acyclic MDPs. We conclude:

Corollary 3. *Given an MDP $\mathcal{M}$, $\bowtie \in \{<, \leq, \geq, >\}$, and $\theta \in \mathbb{Q}$, deciding whether $PE_{\mathcal{M}, s_{init}}^{\sup} \bowtie \theta$ is PSPACE-hard.*

Finiteness. We present criteria for the finiteness of $PE_{s_{init}}^{\sup}[b]$ and $CE_{s_{init}}^{\sup}$. Detailed proofs can be found in Appendix A.1 of [17]. By slightly modifying the construction from [1] which removes end components only containing 0-weight cycles, we obtain the following result.

Proposition 4. *Let $\mathcal{M}$ be an MDP which does not contain positively weight-divergent end components and let $b \in \mathbb{Q}$. Then there is a polynomial time transformation to an MDP $\mathcal{N}$ containing all states from $\mathcal{M}$ and possibly an additional absorbing state fail such that*

- *all end components of $\mathcal{N}$ have negative maximal expected mean payoff,*
- *for any scheduler $\mathfrak{S}$ for $\mathcal{M}$ there is a scheduler $\mathfrak{S}'$ for $\mathcal{N}$ with $\Pr_{\mathcal{M}, s}^{\mathfrak{S}}(\Diamond goal) = \Pr_{\mathcal{N}, s}^{\mathfrak{S}'}(\Diamond goal)$ and $PE_{\mathcal{M}, s}^{\mathfrak{S}}[b] = PE_{\mathcal{N}, s}^{\mathfrak{S}'}[b]$ for any state s in $\mathcal{M}$, and vice versa.*

Hence, we can restrict ourselves to MDPs in which all end components have negative maximal expected mean payoff if there are no positively weight divergent end components. The following result is now analogous to the result in [1] for the classical SSP problem.

Proposition 5. *Let $\mathcal{M}$ be an MDP and $b \in \mathbb{R}$ arbitrary. The optimal partial expectation $PE_{s_{init}}^{\sup}[b]$ is finite if and only if there are no positively weight-divergent end components in $\mathcal{M}$.*

To obtain an analogous result for conditional expectations, we observe that the finiteness of the maximal partial expectation is necessary for the finiteness of the maximal conditional expectation. However, this is not sufficient. In [3], a *critical scheduler* is defined as a scheduler $\mathfrak{S}$ for which there is a path containing a positive cycle and for which $\Pr_{s_{init}}^{\mathfrak{S}}(\Diamond goal) = 0$. Given a critical scheduler, it is easy to construct a sequence of schedulers with unbounded conditional expectation (see Appendix A.1 of [17] and [3]). On the other hand, if $\Pr_{\mathcal{M}, s_{init}}^{\min}(\Diamond goal) > 0$, then $CE_{s_{init}}^{\sup}$ is finite if and only if $PE_{s_{init}}^{\sup}$ is finite. We will show how we can restrict ourselves to this case if there are no critical schedulers:

So, let $\mathcal{M}$ be an MDP with $\Pr_{\mathcal{M}, s_{init}}^{\min}(\Diamond goal) = 0$ and suppose there are no critical schedulers for $\mathcal{M}$. Let S_0 be the set of all states reachable from s_{init} while only choosing actions in $Act^{\min}$. As there are no critical schedulers, $(S_0, Act^{\min})$ does not contain positive cycles. So, there is a finite maximal weight w_s among paths leading from s_{init} to s in S_0 . Consider the following MDP $\mathcal{N}$: It contains the MDP $\mathcal{M}$ and a new initial state t_{init} . For each $s \in S_0$ and each $\alpha \in Act(s) \setminus Act^{\min}(s)$, $\mathcal{N}$ also contains a new state $t_{s, \alpha}$ which is reachable from

t_{init} via an action $\beta_{s,\alpha}$ with weight w_s and probability 1. In $t_{s,\alpha}$, only action α with the same probability distribution over successors and the same weight as in s is enabled. So in $\mathcal{N}$, one has to decide immediately in which state to leave S_0 and one accumulates the maximal weight which can be accumulated in $\mathcal{M}$ to reach this state in S_0 . In this way, we ensure that $\Pr_{\mathcal{N},t_{init}}^{\min}(\Diamond goal) > 0$.

Proposition 6. *The constructed MDP $\mathcal{N}$ satisfies $CE_{\mathcal{N},t_{init}}^{\sup} = CE_{\mathcal{M},s_{init}}^{\sup}$.*

We can rely on this reduction to an MDP in which *goal* is reached with positive probability for ϵ -approximations and the exact computation of the optimal conditional expectation. In particular, the values w_s for $s \in S_0$ are easy to compute by classical shortest path algorithms on weighted graphs. Furthermore, we can now decide the finiteness of the maximal conditional expectation.

Proposition 7. *For an arbitrary MDP $\mathcal{M}$, $CE_{\mathcal{M},s_{init}}^{\sup}$ is finite if and only if there are no positively weight-divergent end components and no critical schedulers.*

Partial and Conditional Expectations in Markov Chains. Markov chains with integer weights can be seen as MDPs with only one action α enabled in every state. Consequently, there is only one scheduler for a Markov chain. Hence, we drop the superscripts in $p^{\max}$ and $PE^{\sup}$.

Proposition 8. *The partial and conditional expectation in a Markov chain $\mathcal{C}$ are computable in polynomial time.*

Proof. Let α be the only action available in $\mathcal{C}$. Assume that all states from which *goal* is not reachable have been collapsed to an absorbing state *fail*. Then $PE_{\mathcal{C},s_{init}}$ is the value of $x_{s_{init}}$ in the unique solution to the following system of linear equations with one variable x_s for each state s :

$$\begin{aligned} x_{goal} &= x_{fail} = 0, \\ x_s &= wgt(s, \alpha) \cdot p_s + \sum_t P(s, \alpha, t) \cdot x_t \text{ for } s \in S \setminus \{goal, fail\}. \end{aligned}$$

The existence of a unique solution follows from the fact that $\{goal\}$ and $\{fail\}$ are the only end components (see [18]). It is straight-forward to check that $(PE_{\mathcal{C},s})_{s \in S}$ is this unique solution. The conditional expectation is obtained from the partial expectation by dividing by the probability $p_{s_{init}}$ to reach the goal. $\square$

This result can be seen as a special case of the following result. Restricting ourselves to schedulers which reach the goal with maximal or minimal probability in an MDP without positively weight-divergent end components, linear programming allows us to compute the following two memoryless deterministic schedulers (see [3, 8]).

Proposition 9. *Let $\mathcal{M}$ be an MDP without positively weight-divergent end components. There is a scheduler $\mathfrak{Max} \in MD^{\mathcal{M}}$ such that for each $s \in S$ we have $\Pr_s^{\mathfrak{Max}}(\Diamond \text{goal}) = p_s^{\max}$ and $PE_s^{\mathfrak{Max}} = \sup_{\mathfrak{S}} PE_s^{\mathfrak{S}}$ where the supremum is taken over all schedulers $\mathfrak{S}$ with $\Pr_s^{\mathfrak{S}}(\Diamond \text{goal}) = p_s^{\max}$. Similarly, there is a scheduler $\mathfrak{Min} \in MD^{\mathcal{M}}$ maximizing the partial expectation among all schedulers reaching the goal with minimal probability. Both these schedulers and their partial expectations are computable in polynomial time.*

These schedulers will play a crucial role for the approximation of the maximal partial expectation and the exact computation of maximal partial expectations in MDPs with non-negative weights.

Partial Expectations in MDPs with Non-negative Weights. In [8], the computation of maximal partial expectations in stochastic multiplayer games with non-negative weights is presented. We adapt this approach to MDPs with non-negative weights. A key result is the existence of a *saturation point*, a bound on the accumulated weight above which optimal schedulers do not need memory.

In the sequel, let $R \in \mathbb{Q}$ be arbitrary, let $\mathcal{M}$ be an MDP with non-negative weights, $PE_{s_{init}}^{\sup} < \infty$, and assume that end components have negative maximal mean payoff (see Proposition 4). A saturation point for bias R is a natural number $\mathfrak{p}$ such that there is a scheduler $\mathfrak{S}$ with $PE_{s_{init}}^{\mathfrak{S}}[R] = PE_{s_{init}}^{\sup}[R]$ which is memoryless and deterministic as soon as the accumulated weight reaches $\mathfrak{p}$. I.e. for any two paths π and π' , with $\text{last}(\pi) = \text{last}(\pi')$ and $\text{wgt}(\pi), \text{wgt}(\pi') > \mathfrak{p}$, $\mathfrak{S}(\pi) = \mathfrak{S}(\pi')$.

Transferring the idea behind the saturation point for conditional expectations given in [3], we provide the following saturation point which can be considerably smaller than the saturation point given in [8] in stochastic multiplayer games. Detailed proofs to this section are given in Appendix A.2 of [17].

Proposition 10. *We define $p_{s,\alpha}^{\max} := \sum_{t \in S} P(s, \alpha, t) \cdot p_t^{\max}$ and $PE_{s,\alpha}^{\mathfrak{Max}} := p_{s,\alpha}^{\max} \cdot \text{wgt}(s, \alpha) + \sum_{t \in S} P(s, \alpha, t) \cdot PE_t^{\mathfrak{Max}}$. Then,*

$$\mathfrak{p}_R := \sup \left\{ \frac{PE_{s,\alpha}^{\mathfrak{Max}} - PE_s^{\mathfrak{Max}}}{p_s^{\max} - p_{s,\alpha}^{\max}} \mid s \in S, \alpha \in \text{Act}(s) \setminus \text{Act}^{\max}(s) \right\} - R$$

is an upper saturation point for bias R in $\mathcal{M}$.

The saturation point $\mathfrak{p}_R$ is chosen such that, as soon as the accumulated weight exceeds $\mathfrak{p}_R$, the scheduler $\mathfrak{Max}$ is better than any scheduler deviating from $\mathfrak{Max}$ for only one step. So, the proposition states that $\mathfrak{Max}$ is then also better than any other scheduler.

As all values involved in the computation can be determined by linear programming, the saturation point $\mathfrak{p}_R$ is computable in polynomial time. This also means that the logarithmic length of $\mathfrak{p}_R$ is polynomial in the size of $\mathcal{M}$ and hence $\mathfrak{p}_R$ itself is at most exponential in the size of $\mathcal{M}$.

Proposition 11. *Let $R \in \mathbb{Q}$ and let B_R be the least integer greater or equal to $\mathfrak{p}_R + \max_{s \in S, \alpha \in \text{Act}(s)} \text{wgt}(s, \alpha)$ and let $S' := S \setminus \{\text{goal}, \text{fail}\}$. The values*

$(PE_{s_{init}}^{\sup}[r+R])_{s \in S', 0 \leq r \leq B_R}$ form the unique solution to the following linear program in the variables $(x_{s,r})_{s \in S', 0 \leq r \leq B_R}$ (r ranges over integers):

Minimize $\sum_{s \in S', 0 \leq r \leq B_R} x_{s,r}$ under the following constraints:

$$\text{For } r \geq \mathfrak{p}_R : x_{s,r} = p_s^{\max} \cdot (r+R) + E_s^{\mathfrak{Max}},$$

for $r < \mathfrak{p}_R$ and $\alpha \in \text{Act}(s)$:

$$x_{s,r} \geq P(s, \alpha, \text{goal}) \cdot (r+R + \text{wgt}(s, \alpha)) + \sum_{t \in S'} P(s, \alpha, t) \cdot x_{t, r + \text{wgt}(s, \alpha)}.$$

From a solution x to the linear program, we can easily extract an optimal weight-based deterministic scheduler. This scheduler only needs finite memory because the accumulated weight increases monotonically along paths and as soon as the saturation point is reached $\mathfrak{Max}$ provides the optimal decisions. As B_R is exponential in the size of $\mathcal{M}$, the computation of the optimal partial expectation via this linear program runs in time exponential in the size of $\mathcal{M}$.

4 Existence of Optimal Schedulers

We prove that there are optimal weight-based deterministic schedulers for partial and conditional expectations. After showing that, if finite, $PE_{s_{init}}^{\sup}$ is equal to $\sup_{\mathfrak{S} \in WD^{\mathcal{M}}} PE_{s_{init}}^{\mathfrak{S}}$, we take an analytic approach to show that there is a weight-based deterministic scheduler maximizing the partial expectation. We define a metric on $WD^{\mathcal{M}}$ turning it into a compact space. Then, we prove that the function assigning the partial expectation to schedulers is upper semi-continuous. We conclude that there is a weight-based deterministic scheduler obtaining the maximum. Proofs to this section can be found in Appendix B of [17].

Proposition 12. *Let $\mathcal{M}$ be an MDP with $PE_{s_{init}}^{\sup} < \infty$. Then we have $PE_{s_{init}}^{\sup} = \sup_{\mathfrak{S} \in WD^{\mathcal{M}}} PE_{s_{init}}^{\mathfrak{S}}$.*

Proof sketch. We can assume that all end components have negative maximal expected mean payoff (see Proposition 4). Given a scheduler $\mathfrak{S} \in HR^{\mathcal{M}}$, we take the expected number of times $\theta_{s,w}$ that s is visited with accumulated weight w under $\mathfrak{S}$ for each state-weight pair (s, w) , and the expected number of times $\theta_{s,w,\alpha}$ that $\mathfrak{S}$ then chooses α . These values are finite due to the negative maximal mean payoff in end components. We define the scheduler $\mathfrak{T} \in WR^{\mathcal{M}}$ choosing α in s with probability $\theta_{s,w,\alpha}/\theta_{s,w}$ when weight w has been accumulated. Then, we show by standard arguments that we can replace all probability distributions that $\mathfrak{T}$ chooses by Dirac distributions to obtain a scheduler $\mathfrak{T}' \in WD^{\mathcal{M}}$ such that $PE_{s_{init}}^{\mathfrak{T}'} \geq PE_{s_{init}}^{\mathfrak{S}}$. $\square$

It remains to show that the supremum is obtained by a weight-based deterministic scheduler. Given an MDP $\mathcal{M}$ with arbitrary integer weights, we define the following metric $d^{\mathcal{M}}$ on the set of weight-based deterministic schedulers, i.e. on the set of functions from $S \times \mathbb{Z} \rightarrow \text{Act}$: For two such schedulers $\mathfrak{S}$ and

$\mathfrak{T}$, we let $d^{\mathcal{M}}(\mathfrak{S}, \mathfrak{T}) := 2^{-R}$ where R is the greatest natural number such that $\mathfrak{S} \upharpoonright S \times \{-(R-1), \dots, R-1\} = \mathfrak{T} \upharpoonright S \times \{-(R-1), \dots, R-1\}$ or ∞ if there is no greatest such natural number.

Lemma 13. *The metric space $(Act^{S \times \mathbb{Z}}, d^{\mathcal{M}})$ is compact.*

Having defined this compact space of schedulers, we can rely on the analytic notion of upper semi-continuity.

Lemma 14 (Upper Semi-Continuity of Partial Expectations). *If $PE_{s_{init}}^{\sup}$ is finite in $\mathcal{M}$, then the function $PE : (WD, d^{WD}) \rightarrow (\mathbb{R}_{\infty}, d^{euclid})$ assigning $PE_{s_{init}}^{\mathfrak{S}}$ to a weight-based deterministic scheduler $\mathfrak{S}$ is upper semi-continuous.*

The technical proof of this lemma can be found in Appendix B of [17]. We arrive at the main result of this section.

Theorem 15 (Existence of Optimal Schedulers for Partial Expectations). *If $PE_{s_{init}}^{\sup}$ is finite in an MDP $\mathcal{M}$, then there is a weight-based deterministic scheduler $\mathfrak{S}$ with $PE_{s_{init}}^{\sup} = PE_{s_{init}}^{\mathfrak{S}}$.*

Proof. If $PE_{s_{init}}^{\sup}$ is finite, then the map $PE : (WD, d^{WD}) \rightarrow (\mathbb{R}_{\infty}, d^{euclid})$ is upper semi-continuous. So, this map has a maximum because (WD, d^{WD}) is a compact metric space. $\square$

Corollary 16 (Existence of Optimal Schedulers for Conditional Expectations). *If $CE_{s_{init}}^{\sup}$ is finite in an MDP $\mathcal{M}$, then there is a weight-based deterministic scheduler $\mathfrak{S}$ with $CE_{s_{init}}^{\sup} = CE_{s_{init}}^{\mathfrak{S}}$.*

Proof. By Proposition 6, we can assume that $\Pr_{s_{init}}^{\min}(\diamond \text{goal}) > 0$. We know that $PE_{s_{init}}^{\sup}[-CE_{s_{init}}^{\sup}] = 0$ and that there is a weight-based deterministic scheduler $\mathfrak{S}$ with $PE_{s_{init}}^{\mathfrak{S}}[-CE_{s_{init}}^{\sup}] = 0$. By Proposition 2, $\mathfrak{S}$ maximizes the conditional expectation as it reaches *goal* with positive probability. $\square$

Fig. 2. All non-trivial transition probabilities are $1/2$. In the MDP $\mathcal{M}$, the optimal choice to maximize the partial expectation in t depends on the parity of the accumulated weight.

In MDPs with non-negative weights, the optimal decision in a state s only depends on s as soon as the accumulated weight exceeds a saturation point. In MDPs with arbitrary integer weights, it is possible that the optimal choice of action does not become stable for increasing values of accumulated weight as we see in the following example.

Example 17. Let us first consider the MDP $\mathcal{N}$ depicted in Fig. 2. Let π be a path reaching t for the first time with accumulated weight r . Consider a scheduler which chooses β for the first k times and then α . In this situation, the partial expectation from this point on is:

$$\frac{1}{2^{k+1}}(r-k) + \sum_{i=1}^k \frac{1}{2^i}(r-i) = \frac{1}{2^{k+1}} + \sum_{i=1}^{k+1} \frac{1}{2^i}(r-i) = \frac{k-r+4}{2^{k+1}} + r-2.$$

For $r \geq 2$, this partial expectation has its unique maximum for the choice $k = r-2$. This already shows that an optimal scheduler needs infinite memory. No matter how much weight r has been accumulated when reaching t , the optimal scheduler has to count the $r-2$ times it chooses β .

Furthermore, we can transfer the optimal scheduler for the MDP $\mathcal{N}$ to the MDP $\mathcal{M}$. In state t , we have to make a nondeterministic choice between two action leading to the states q_0 and q_1 , respectively. In both of these states, action β is enabled which behaves like the same action in the MDP $\mathcal{N}$ except that it moves between the two states if *goal* is not reached. So, the action α is only enabled every other step. As in $\mathcal{N}$, we want to choose α after choosing β $r-2$ times if we arrived in t with accumulated weight $r \geq 2$. So, the choice in t depends on the parity of r : For $r = 1$ or r even, we choose δ . For odd $r \geq 3$, we choose γ . This shows that the optimal scheduler in the MDP $\mathcal{M}$ needs specific information about the accumulated weight, in this case the parity, no matter how much weight has been accumulated.

In the example, the optimal scheduler has a periodic behavior when fixing a state and looking at optimal decisions for increasing values of accumulated weight. The question whether an optimal scheduler always has such a periodic behavior remains open.

5 Approximation

As the optimal values for partial and conditional expectation can be irrational, there is no hope to compute these values by linear programming as in the case of non-negative weights. In this section, we show how we can nevertheless approximate the values. The main result is the following.

Theorem 18. *Let $\mathcal{M}$ be an MDP with $PE_{\mathcal{M},s_{init}}^{\sup} < \infty$ and $\epsilon > 0$. The maximal partial expectation $PE_{\mathcal{M},s_{init}}^{\sup}$ can be approximated up to an absolute error of ϵ in time exponential in the size of $\mathcal{M}$ and polynomial in $\log(1/\epsilon)$. If further, $CE_{\mathcal{M},s_{init}}^{\sup} < \infty$, also $CE_{\mathcal{M},s_{init}}^{\sup}$ can be approximated up to an absolute error of ϵ in time exponential in the size of $\mathcal{M}$ and polynomial in $\log(1/\epsilon)$.*

We first prove that upper bounds for $PE_{\mathcal{M}, s_{init}}^{\sup}$ and $CE_{\mathcal{M}, s_{init}}^{\sup}$ can be computed in polynomial time. Then, we show that there are ϵ -optimal schedulers for the partial expectation which become memoryless as soon as the accumulated weight leaves a sufficiently large weight window around 0. We compute the optimal partial expectation of such a scheduler by linear programming. The result can then be extended to conditional expectations.

Upper Bounds. Let $\mathcal{M}$ be an MDP in which all end components have negative maximal mean payoff. Let δ be the minimal non-zero transition probability in $\mathcal{M}$ and $W := \max_{s \in S, \alpha \in Act(s)} |wgt(s, \alpha)|$. Moving through the MEC-quotient, the probability to reach an accumulated weight of $|S| \cdot W$ is bounded by $1 - \delta^{|S|}$ as *goal* or *fail* is reached within S steps with probability at least $1 - \delta^{|S|}$. It remains to show similar bounds inside an end component.

We will use the characterization of the maximal mean payoff in terms of super-harmonic vectors due to Hordijk and Kallenberg [13] to define a super-martingale controlling the growth of the accumulated weight in an end component under any scheduler. As the value vector for the maximal mean payoff in an end component is constant and negative in our case, the results of [13] yield:

Proposition 19 (Hordijk, Kallenberg). *Let $\mathcal{E} = (S, Act)$ be an end component with maximal mean payoff $-t$ for some $t > 0$. Then there is a vector $(u_s)_{s \in S}$ such that $-t + u_s \geq wgt(s, \alpha) + \sum_{s' \in S} P(s, \alpha, s') \cdot u_{s'}$.*

Furthermore, let v be the vector $(-t, \dots, -t)$ in $\mathbb{R}^S$. Then, (v, u) is the solution to a linear program with $2|S|$ variables, $2|S||Act|$ inequalities, and coefficients formed from the transition probabilities and weights in $\mathcal{E}$.

We will call the vector u a *super-potential* because the expected accumulated weight after i steps is at most $u_s - \min_{t \in S} u_t - i \cdot t$ when starting in state s . Let $\mathfrak{S}$ be a scheduler for $\mathcal{E}$ starting in some state s . We define the following random variables on $\mathfrak{S}$ -runs in $\mathcal{E}$: let $s(i) \in S$ be the state after i steps, let $\alpha(i)$ be the action chosen after i steps, let $w(i)$ be the accumulated weight after i steps, and let $\pi(i)$ be the history, i.e. the finite path after i steps.

Lemma 20. *The sequence $m(i) := w(i) + u_{s(i)}$ satisfies $\mathbb{E}(m(i+1)|\pi(0), \dots, \pi(i)) \leq m(i) - t$ for all i .*¹

Proof. By Proposition 19, $\mathbb{E}(m(i+1)|\pi(0), \dots, \pi(i)) - m(i) = wgt(s(i), \mathfrak{S}(\pi(i))) + \sum_{s' \in S} P(s(i), \mathfrak{S}(\pi(i)), s') \cdot u_{s'} - u_{s(i)} \leq -t$. $\square$

We are going to apply the following theorem by Blackwell [6].

Theorem 21 (Blackwell [6]). *Let $X_1, X_2, \dots$ be random variables, and let $S_n := \sum_{k=1}^n X_k$. Assume that $|X_i| \leq 1$ for all i and that there is a $u > 0$ such that $\mathbb{E}(X_{n+1}|X_1, \dots, X_n) \leq -u$. Then, $\Pr(\sup_{n \in \mathbb{N}} S_n \geq t) \leq \left(\frac{1-u}{1+u}\right)^t$.*

¹ This means that $m(i) + i \cdot t$ is a super-martingale with respect to the history $\pi(i)$.

We denote $\max_{s' \in S} u_{s'} - \min_{s' \in S} u_{s'}$ by $\|u\|$. Observe that $|m(i+1) - m(i)| \leq \|u\| + W =: c_{\mathcal{E}}$. We can rescale the sequence $m(i)$ by defining $m'(i) := (m(i) - m(0))/c_{\mathcal{E}}$. This ensures that $m'(0) = 0$, $|m'(i+1) - m'(i)| \leq 1$ and $\mathbb{E}(m'(i+1)|m'(0), \dots, m'(i)) \leq -t/c_{\mathcal{E}}$ for all i . In this way, we arrive at the following conclusion, putting $\lambda_{\mathcal{E}} := \frac{1-t/c_{\mathcal{E}}}{1+t/c_{\mathcal{E}}}$.

Corollary 22. *For any scheduler $\mathfrak{S}$ and any starting state s in $\mathcal{E}$, we have $\Pr_s^{\mathfrak{S}}(\Diamond \text{wgt} \geq (k+1) \cdot c_{\mathcal{E}}) \leq \lambda_{\mathcal{E}}^k$.*

Proof. By Theorem 21, $\Pr_s^{\mathfrak{S}}(\Diamond \text{wgt} \geq (k+1) \cdot c_{\mathcal{E}}) \leq \Pr_s^{\mathfrak{S}}(\Diamond \text{wgt} \geq \|u\| + k \cdot c_{\mathcal{E}}) \leq \Pr_s^{\mathfrak{S}}(\exists i : m(i) - m(0) \geq k \cdot c_{\mathcal{E}}) = \Pr_s^{\mathfrak{S}}(\sup_{i \in \mathbb{N}} m'(i) \geq k) \leq \left(\frac{1-t/c_{\mathcal{E}}}{1+t/c_{\mathcal{E}}}\right)^k$. $\square$

Let MEC be the set of maximal end components in $\mathcal{M}$. For each $\mathcal{E} \in MEC$, let $\lambda_{\mathcal{E}}$ and $c_{\mathcal{E}}$ be as in Corollary 22. Define $\lambda_{\mathcal{M}} := 1 - (\delta^{|S|} \cdot \prod_{\mathcal{E} \in MEC} (1 - \lambda_{\mathcal{E}}))$, and $c_{\mathcal{M}} := |S| \cdot W + \sum_{\mathcal{E} \in MEC} c_{\mathcal{E}}$. Then an accumulated weight of $c_{\mathcal{M}}$ cannot be reached with a probability greater than $\lambda_{\mathcal{M}}$ because reaching accumulated weight $c_{\mathcal{M}}$ would require reaching weight $c_{\mathcal{E}}$ in some end component $\mathcal{E}$ or reaching weight $|S| \cdot W$ in the MEC-quotient and $1 - \lambda_{\mathcal{M}}$ is a lower bound on the probability that none of this happens (under any scheduler).

Proposition 23. *Let $\mathcal{M}$ be an MDP with $PE_{s_{init}}^{\sup} < \infty$. There is an upper bound PE^{ub} for the partial expectation in $\mathcal{M}$ computable in polynomial time.*

Proof. In any end component $\mathcal{E}$, the maximal mean payoff $-t$ and the superpotential u are computable in polynomial time. Hence, $c_{\mathcal{E}}$ and $\lambda_{\mathcal{E}}$, and in turn also $c_{\mathcal{M}}$ and $\lambda_{\mathcal{M}}$ are also computable in polynomial time. When we reach accumulated weight $c_{\mathcal{M}}$ for the first time, the actual accumulated weight is at most $c_{\mathcal{M}} + W$. So, we conclude that $\Pr_s^{\max}(\Diamond \text{wgt} \geq k \cdot (c_{\mathcal{M}} + W)) \leq \lambda_{\mathcal{M}}^k$. The partial expectation can now be bounded by $\sum_{k=0}^{\infty} (k+1) \cdot (c_{\mathcal{M}} + W) \cdot \lambda_{\mathcal{M}}^k = \frac{c_{\mathcal{M}} + W}{(1 - \lambda_{\mathcal{M}})^2}$. $\square$

Corollary 24. *Let $\mathcal{M}$ be an MDP with $CE_{\mathcal{M}, s_{init}}^{\sup} < \infty$. There is an upper bound CE^{ub} for the conditional expectation in $\mathcal{M}$ computable in polynomial time.*

Proof. By Proposition 6, we can construct an MDP $\mathcal{N}$ in which *goal* is reached with probability $q > 0$ in polynomial time with $CE_{\mathcal{M}, s_{init}}^{\sup} = CE_{\mathcal{N}, s_{init}}^{\sup}$. Now, $CE^{ub} := PE^{ub}/q$ is an upper bound for the conditional expectation in $\mathcal{M}$. $\square$

Approximating Optimal Partial Expectations. The idea for the approximation is to assume that the partial expectation is $PE_{s_{init}}^{\max} + w \cdot p_s^{\max}$ if a high weight w has been accumulated in state s . Similarly, for small weights w' , we use the value $PE_{s_{init}}^{\min} + w' \cdot p_s^{\min}$. We will first provide a lower “saturation point” making sure that only actions minimizing the probability to reach the goal are used by an optimal scheduler as soon as the accumulated weight drops below this saturation point. For the proofs to this section, see Appendix C.1 of [17].

Proposition 25. *Let $\mathcal{M}$ be an MDP with $PE_{s_{init}}^{\sup} < \infty$. Let $s \in S$ and let $q_s := \frac{PE_s^{ub} - PE_s^{\mathfrak{Min}}}{p_s^{\min} - \min_{\alpha \notin Act^{\min}(s)} p_{s,\alpha}^{\min}}$. Then any weight-based deterministic scheduler $\mathfrak{S}$ maximizing the partial expectation in $\mathcal{M}$ satisfies $\mathfrak{S}(s, w) \in Act^{\min}(s)$ if $w \leq q_s$.*

Let $q := \min_{s \in S} q_s$ and let $D := PE^{ub} - \min\{PE_s^{\mathfrak{Max}}, PE_s^{\mathfrak{Min}} | s \in S\}$. Given $\epsilon > 0$, we define $R_\epsilon^+ := (c_{\mathcal{M}} + W) \cdot \left\lceil \frac{\log(2D) + \log(1/\epsilon)}{\log(1/\lambda_{\mathcal{M}})} \right\rceil$ and $R_\epsilon^- := q - R_\epsilon^+$.

Theorem 26. *There is a weight-based deterministic scheduler $\mathfrak{S}$ such that the scheduler $\mathfrak{T}$ defined by*

$$\mathfrak{T}(\pi) = \begin{cases} \mathfrak{S}(\pi) & \text{if any prefix } \pi' \text{ of } \pi \text{ satisfies } R_\epsilon^- \leq \text{wgt}(\pi') \leq R_\epsilon^+, \\ \mathfrak{Max}(\pi) & \text{if the shortest prefix } \pi' \text{ of } \pi \text{ with } \text{wgt}(\pi') \notin [R_\epsilon^-, R_\epsilon^+] \\ & \text{satisfies } \text{wgt}(\pi') > R_\epsilon^+, \\ \mathfrak{Min}(\pi) & \text{otherwise,} \end{cases}$$

satisfies $PE_{s_{init}}^{\mathfrak{T}} \geq PE_{s_{init}}^{\sup} - \epsilon$.

This result now allows us to compute an ϵ -approximation and an ϵ -optimal scheduler with finite memory by linear programming, similar to the case of non-negative weights, in a linear program with $R_\epsilon^+ + R_\epsilon^-$ many variables and $|Act|$ -times as many inequalities.

Corollary 27. *$PE_{s_{init}}^{\sup}$ can be approximated up to an absolute error of ϵ in time exponential in the size of $\mathcal{M}$ and polynomial in $\log(1/\epsilon)$.*

If the logarithmic length of $\theta \in \mathbb{Q}$ is polynomial in the size of $\mathcal{M}$, we can also approximate $PE_{s_{init}}^{\sup}[\theta]$ up to an absolute error of ϵ in time exponential in the size of $\mathcal{M}$ and polynomial in $\log(1/\epsilon)$: We can add a new initial state s with a transition to s_{init} with weight θ and approximate $PE_s^{\sup}$ in the new MDP.

Transfer to Conditional Expectations. Let $\mathcal{M}$ be an MDP with $CE_{s_{init}}^{\sup} < \infty$ and $\epsilon > 0$. By Proposition 6, we can assume that $\Pr_{\mathcal{M}, s_{init}}^{\min}(\Diamond \text{goal}) =: p$ is positive. Clearly, $CE_{s_{init}}^{\sup} \in [CE_{s_{init}}^{\mathfrak{Max}}, CE^{ub}]$. We perform a binary search to approximate $CE_{s_{init}}^{\sup}$: We put $A_0 := CE_{s_{init}}^{\mathfrak{Max}}$ and $B_0 := CE^{ub}$. Given A_i and B_i , let $\theta_i := (A_i + B_i)/2$. Then, we approximate $PE_{s_{init}}^{\sup}[-\theta_i]$ up to an absolute error of $p \cdot \epsilon$. Let E_i be the value of this approximation. If $E_i \in [-2p \cdot \epsilon, 2p \cdot \epsilon]$, terminate and return θ_i as the approximation for $CE_{s_{init}}^{\sup}$. If $E_i < -2p \cdot \epsilon$, put $A_{i+1} := A_i$ and $B_{i+1} := \theta_i$, and repeat. If $E_i > 2p \cdot \epsilon$, put $A_{i+1} := \theta_i$ and $B_{i+1} := B_i$, and repeat.

Proposition 28. *The procedure terminates after at most $\lceil \log((A_0 - B_0)/(p \cdot \epsilon)) \rceil$ iterations and returns a 3ϵ -approximation of $CE_{s_{init}}^{\sup}$ in time exponential in the size of $\mathcal{M}$ and polynomial in $\log(1/\epsilon)$.*

The proof can be found in Appendix C.2 of [17]. This finishes the proof of Theorem 18.

6 Conclusion

Compared to the setting of non-negative weights, the optimization of partial and conditional expectations faces substantial new difficulties in the setting of integer weights. The optimal values can be irrational showing that the linear programming approaches from the setting of non-negative weights cannot be applied for the computation of optimal values. We showed that this approach can nevertheless be adapted for approximation algorithms. Further, we were able to show that there are optimal weight-based deterministic schedulers. These schedulers, however, can require infinite memory and it remains open whether we can further restrict the class of schedulers necessary for the optimization. In examples, we have seen that optimal schedulers can switch periodically between actions they choose for increasing values of accumulated weight. Further insights on the behavior of optimal schedulers would be helpful to address threshold problems (“Is $PE_{s_{init}}^{\sup} \geq \theta$?”).

References

1. Baier, C., Bertrand, N., Dubslaff, C., Gburek, D., Sankur, O.: Stochastic shortest paths and weight-bounded properties in Markov decision processes. In: Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), pp. 86–94. ACM (2018)
2. Baier, C., Dubslaff, C., Klein, J., Klüppelholz, S., Wunderlich, S.: Probabilistic model checking for energy-utility analysis. In: van Breugel, F., Kashefi, E., Palamidessi, C., Rutten, J. (eds.) Horizons of the Mind. A Tribute to Prakash Panangaden. LNCS, vol. 8464, pp. 96–123. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06880-0_5
3. Baier, C., Klein, J., Klüppelholz, S., Wunderlich, S.: Maximizing the conditional expected reward for reaching the goal. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10206, pp. 269–285. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54580-5_16
4. Barthe, G., Espitau, T., Ferrer Fioriti, L.M., Hsu, J.: Synthesizing probabilistic invariants via Doob’s decomposition. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9779, pp. 43–61. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41528-4_3
5. Bertsekas, D.P., Tsitsiklis, J.N.: An analysis of stochastic shortest path problems. Math. Oper. Res. **16**(3), 580–595 (1991)
6. Blackwell, D.: On optimal systems. Ann. Math. Stat. **25**, 394–397 (1954)
7. Chatterjee, K., Fu, H., Goharshady, A.K.: Termination analysis of probabilistic programs through Positivstellensatz’s. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9779, pp. 3–22. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41528-4_1
8. Chen, T., Forejt, V., Kwiatkowska, M., Parker, D., Simaitis, A.: Automatic verification of competitive stochastic systems. Formal Methods Syst. Des. **43**(1), 61–92 (2013)
9. Ciesinski, F., Baier, C., Größer, M., Klein, J.: Reduction techniques for model checking Markov decision processes. In: Proceedings of the Fifth International Conference on Quantitative Evaluation of Systems (QEST), pp. 45–54. IEEE (2008)

10. de Alfaro, L.: Computing minimum and maximum reachability times in probabilistic systems. In: Baeten, J.C.M., Mauw, S. (eds.) CONCUR 1999. LNCS, vol. 1664, pp. 66–81. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48320-9_7
11. Esparza, J., Kucera, A., Mayr, R.: Quantitative analysis of probabilistic pushdown automata: expectations and variances. In: Proceedings of the 20th Annual IEEE Symposium on Logic in Computer Science (LICS), pp. 117–126. IEEE (2005)
12. Gretz, F., Katoen, J.-P., McIver, A.: Operational versus weakest pre-expectation semantics for the probabilistic guarded command language. *Perform. Eval.* **73**, 110–132 (2014)
13. Hordijk, A., Kallenberg, L.: Linear programming and Markov decision chains. *Management Sci.* **25**(4), 352–362 (1979)
14. Katoen, J.-P., Gretz, F., Jansen, N., Kaminski, B.L., Olmedo, F.: Understanding probabilistic programs. In: Meyer, R., Platzer, A., Wehrheim, H. (eds.) *Correct System Design*. LNCS, vol. 9360, pp. 15–32. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23506-6_4
15. Kretínský, J., Meggendorfer, T.: Conditional value-at-risk for reachability and mean payoff in Markov decision processes. In: Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), pp. 609–618. ACM (2018)
16. Olmedo, F., Gretz, F., Jansen, N., Kaminski, B.L., Katoen, J.-P., Mciver, A.: Conditioning in probabilistic programming. *ACM Trans. Program. Lang. Syst. (TOPLAS)* **40**(1), 4:1–4:50 (2018)
17. Piribauer, J., Baier, C.: Partial and conditional expectations in Markov decision processes with integer weights (extended version). [arXiv:1902.04538](https://arxiv.org/abs/1902.04538) (2019)
18. Puterman, M.L.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York (1994)

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

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