Qualitative Controller Synthesis for Consumption Markov Decision Processes

František Blahoudek¹, Tomáš Brázdil², Petr Novotný², Melkior Ornik³,
Pranay Thangeda³(✉), and Ufuk Topcu¹

¹ The University of Texas at Austin, Austin, USA
frantisek.blahoudek@gmail.com,
utopcu@utexas.edu

² Masaryk University, Brno, Czech Republic
{xbrazdil, petr.novotny}@fi.muni.cz

³ University of Illinois at Urbana-Champaign, Urbana, USA
{mornik, pranayt2}@illinois.edu

Abstract. Consumption Markov Decision Processes (CMDPs) are probabilistic decision-making models of resource-constrained systems. In a CMDP, the controller possesses a certain amount of a critical resource, such as electric power. Each action of the controller can consume some amount of the resource. Resource replenishment is only possible in special *reload states*, in which the resource level can be reloaded up to the full capacity of the system. The task of the controller is to prevent resource exhaustion, i.e. ensure that the available amount of the resource stays non-negative, while ensuring an additional linear-time property. We study the complexity of strategy synthesis in consumption MDPs with almost-sure Büchi objectives. We show that the problem can be solved in polynomial time. We implement our algorithm and show that it can efficiently solve CMDPs modelling real-world scenarios.

1 Introduction

In the context of formal methods, controller synthesis typically boils down to computing a strategy in an *agent-environment* model, a nondeterministic state-transition model where some of the nondeterministic choices are resolved by the controller and some by an uncontrollable environment. Such models are typically either two-player graph games with an adversarial environment or Markov decision process (MDPs); the latter case being apt for modelling statistically predictable environments. In this paper, we consider controller synthesis for *resource-constrained MDPs*, where the computed controller must ensure, in addition to satisfying some linear-time property, that the system's operation is not compromised by a lack of necessary resources.

This work was partially supported by NASA under Early Stage Innovations grant No. 80NSSC19K0209, and by DARPA under grant No. HR001120C0065. Petr Novotný is supported by the Czech Science Foundation grant No. GJ19-15134Y.

© The Author(s) 2020

S. K. Lahiri and C. Wang (Eds.): CAV 2020, LNCS 12225, pp. 421–447, 2020.

https://doi.org/10.1007/978-3-030-53291-8_22

Resource-Constrained Probabilistic Systems. *Resource-constrained* systems need a supply of some resource (e.g. power) for steady operation: the interruption of the supply can lead to undesirable consequences and has to be avoided. For instance, an autonomous system, e.g. an autonomous electric vehicle (AEV), is not able to draw power directly from an endless source. Instead, it has to rely on an internal storage of the resource, e.g. a battery, which has to be replenished in regular intervals to prevent resource exhaustion. Practical examples of AEVs include driverless cars, drones, or planetary rovers [8]. In these domains, resource failures may cause a costly mission failure and even safety risks. Moreover, the operation of autonomous systems is subject to probabilistic uncertainty [54]. Hence, in this paper, we study the resource-constrained strategy synthesis problem for MDPs.

Models of Resource-Constrained Systems & Limitations of Current Approaches.

There is a substantial body of work in the area of verification of resource-constrained systems [3, 5, 7, 9, 11, 23, 38, 39, 53, 58]. The typical approach is to model them as finite-state systems augmented with an integer-valued counter representing the current *resource level*, i.e. the amount of the resource present in the internal storage. The resource constraint requires that the resource level never drops below zero.¹ In the well-known *energy* model [11, 23], each transition is labelled by an integer, and performing an ℓ -labelled transition results in ℓ being added to the counter. Thus, negative numbers stand for resource consumption while positive ones represent re-charging by the respective amount. Many variants of both MDP and game-based energy models were studied, as detailed in the related work. In particular, [26] considers controller synthesis for energy MDPs with qualitative Büchi and parity objectives. The main limitation of energy-based agent-environment models is that in general, they are not known to admit polynomial-time controller synthesis algorithms. Indeed, already the simplest problem, deciding whether a non-negative energy can be maintained in a two-player energy game, is at least as hard as solving mean-payoff graph games [11]; the complexity of the latter being a well-known open problem [45]. This hardness translates also to MDPs [26], making polynomial-time controller synthesis for energy MDPs impossible without a theoretical breakthrough.

Consumption models, introduced in [14], offer an alternative to energy models. In a consumption model, a non-negative integer, *cap*, represents the maximal amount of the resource the system can hold, e.g. the battery capacity. Each transition is labelled by a non-negative number representing the amount of the resource *consumed* when taking the transition (i.e., taking an ℓ -labelled transition decreases the resource level by ℓ). The resource replenishment is different from the energy approach. The consumption approach relies on the fact that reloads are often *atomic events*, e.g. an AEV plugging into a charging station and waiting to finish the charging cycle. Hence, some states in the consumption model are designated as *reload states*, and whenever the system visits a

¹ In some literature, the level is required to stay positive as opposed to non-negative, but this is only a matter of definition: both approaches are equivalent.

reload state, the resource level is replenished to the full capacity cap . Modelling reloads as atomic events is natural and even advantageous: consumption models typically admit more efficient analysis than energy models [14, 47]. However, consumption models have not yet been considered in the probabilistic setting.

Our Contribution. We study strategy synthesis in consumption MDPs with Büchi objectives. Our main theoretical result is stated in the following theorem.

Theorem 1. *Given a consumption MDP $\mathcal{M}$ with a capacity cap , an initial resource level $0 \leq d \leq cap$, and a set T of accepting states, we can decide, in polynomial time, whether there exists a strategy σ such that when playing according to σ , the following consumption-Büchi objectives are satisfied:*

- *Starting with resource level d , the resource level never² drops below 0.*
- *With probability 1, the system visits some state in T infinitely often.*

Moreover, if such a strategy exists then we can compute, in polynomial time, its polynomial-size representation.

For the sake of clarity, we restrict to proving Theorem 1 for a natural sub-class of MDPs called *decreasing consumption MDPs*, where there are no cycles of zero consumption. The restriction is natural (since in typical resource-constrained systems, each action – even idling – consumes some energy, so zero cycles are unlikely) and greatly simplifies presentation. In addition to the theoretical analysis, we implemented the algorithm behind Theorem 1 and evaluated it on several benchmarks, including a realistic model of an AEV navigating the streets of Manhattan. The experiments show that our algorithm is able to efficiently solve large CMDPs, offering a good scalability.

Significance. Some comments on Theorem 1 are in order. First, all the numbers in the MDP, and in particular the capacity cap , are encoded in binary. Hence, “polynomial time” means time polynomial in the encoding size of the MDP itself and in $\log(cap)$. In particular, a naive “unfolding” of the MDP, i.e. encoding the resource levels between 0 and cap into the states, does not yield a polynomial-time algorithm, but an exponential-time one, since the unfolded MDP has size proportional to cap . We employ a value-iteration-like algorithm to compute minimal energy levels with which one can achieve the consumption-Büchi objectives.

A similar concern applies to the “polynomial-size representation” of the strategy σ . To satisfy a consumption-Büchi objective, σ generally needs to keep track of the current resource level. Hence, under the standard notion of a finite-memory (FM) strategy (which views FM strategies as transducers), σ would require memory proportional to cap , i.e. a memory exponentially large w.r.t. size of the input. However, we show that for each state s we can partition the integer interval $[0, \dots, cap]$ into polynomially many sub-intervals $I_1^s, \dots, I_k^s$ such that, for each $1 \leq j \leq k$, the strategy σ picks the same action whenever the current state is

² In our model, this is equivalent to requiring that with probability 1, the resource level never drops below 0.

s and the current resource level is in I_j^s . As such, the endpoints of the intervals are the only extra knowledge required to represent σ , a representation which we call a *counter selector*. We instrument our main algorithm so as to compute, in polynomial time, a polynomial-size counter selector representing the witness strategy σ .

Finally, we consider linear-time properties encoded by Büchi objectives over the states of the MDP. In essence, we assume that the translation of the specification to the Büchi automaton and its product with the original MDP model of the system were already performed. Probabilistic analysis typically requires the use of deterministic Büchi automata, which cannot express all linear-time properties. However, in this paper we consider qualitative analysis, which can be performed using restricted versions of non-deterministic Büchi automata that are still powerful enough to express all ω -regular languages. Examples of such automata are limit-deterministic Büchi automata [51] or good-for-MDPs automata [41]. Alternatively, consumption MDPs with parity objectives could be reduced to consumption-Büchi MDPs using the standard parity-to-Büchi MDP construction [25, 30, 32, 33]. We abstract from these aspects and focus on the technical core of our problem, solving consumption-Büchi MDPs.

Consequently, to our best knowledge, we present the first polynomial-time algorithm for controller synthesis in resource-constrained MDPs with ω -regular objectives.

Related Work. There is an enormous body of work on energy models. Stemming from the models introduced in [11, 23], the subsequent work covered energy games with various combinations of objectives [10, 12, 13, 18, 20, 21, 27, 48], energy games with multiple resource types [15, 24, 28, 31, 37, 43, 44, 57] or the variants of the above in the MDP [17, 49], infinite-state [1], or partially observable [34] settings. As argued previously, the controller synthesis within these models is at least as hard as solving mean-payoff games. The paper [29] presents polynomial-time algorithms for non-stochastic energy games with special weight structures. Recently, an abstract algebraic perspective on energy models was presented in [22, 35, 36].

Consumption systems were introduced in [14] in the form of consumption games with multiple resource types. Minimizing mean-payoff in automata with consumption constraints was studied in [16].

Our main result requires, as a technical sub-component, solving the *resource-safety* (or just *safety*) problem in consumption MDPs, i.e. computing a strategy which prevents resource exhaustion. The solution to this problem consists (in principle) of a Turing reduction to the problem of minimum cost reachability in two-player games with non-negative costs. The latter problem was studied in [46], with an extension to arbitrary costs considered in [19] (see also [40]). We present our own, conceptually simple, value-iteration-like algorithm for the problem, which is also used in our implementation.

Elements of resource-constrained optimization and minimum-cost reachability are also present in the line of work concerning *energy-utility quantiles* in MDPs [4–7, 42]. In this setting, there is no reloading in the consumption- or

energy-model sense, and the task is typically to minimize the total amount of the resource consumed while maximizing the probability that some other objective is satisfied.

Paper Organization & Outline of Techniques. After the preliminaries (Sect. 2), we present counter selectors in Sect. 3. The next three sections contain the three main steps of our analysis. In Sect. 4, we solve the safety problem in consumption MDPs. The technical core of our approach is presented in Sect. 5, where we solve the problem of *safe positive reachability*: finding a resource-safe strategy which ensures that the set T of accepting states is visited with positive probability. Solving consumption-Büchi MDPs then, in principle, consists of repeatedly applying a strategy for safe positive reachability of T , ensuring that the strategy is “re-started” whenever the attempt to reach T fails. Details are given in Sect. 6. Finally, Sect. 7 presents our experiments. Due to space constraints, most technical proofs were moved to the full version.

2 Preliminaries

We denote by $\mathbb{N}$ the set of all non-negative integers and by $\bar{\mathbb{N}}$ the set $\mathbb{N} \cup \{\infty\}$. Given a set I and a vector $\mathbf{v} \in \bar{\mathbb{N}}^I$ of integers indexed by I , we use $\mathbf{v}(i)$ to denote the i -component of $\mathbf{v}$. We assume familiarity with basic notions of probability theory. In particular, a *probability distribution* on an at most countable set X is a function $f: X \rightarrow [0, 1]$ s.t. $\sum_{x \in X} f(x) = 1$. We use $\mathcal{D}(X)$ to denote the set of all probability distributions on X .

Definition 1 (CMDP). A consumption Markov decision process (CMDP) is a tuple $\mathcal{M} = (S, A, \Delta, C, R, \text{cap})$ where S is a finite set of states, A is a finite set of actions, $\Delta: S \times A \rightarrow \mathcal{D}(S)$ is a total transition function, $C: S \times A \rightarrow \mathbb{N}$ is a total consumption function, $R \subseteq S$ is a set of reload states where the resource can be reloaded, and cap is a resource capacity.

Figure 1 shows a visual representation of an CMDP. We denote by $\mathcal{M}(R')$ for $R' \subseteq S$ the CMDP obtained from $\mathcal{M}$ by changing the set of reloads to R' . For

Distributions in Δ are indicated by gray numbers (we leave out 1 when an action has only one successor), and the cost of an action follows its name in the edge labels. Actions labeled by $a_{1,2}$ represent that Δ and C are defined identically for both actions a_1 and a_2 . The blue background indicates a target set $T = \{s_2\}$, while the double circles represent the reload states.

Fig. 1. CMDP $\mathcal{M} = (\{s_1, s_2, s_3, s_4, s_5\}, \{a_1, a_2\}, \Delta, C, \{s_2, s_5\}, 20)$. Details are given on the right.

$s \in S$ and $a \in A$, we denote by $Succ(s, a)$ the set $\{t \mid \Delta(s, a)(t) > 0\}$. A *path* is a (finite or infinite) state-action sequence $\alpha = s_1 a_1 s_2 a_2 s_3 \dots \in (S \times A)^\omega \cup (S \cdot A)^* \cdot S$ such that $s_{i+1} \in Succ(s_i, a_i)$ for all i . We define $\alpha_i = s_i$ and $Act^i(\alpha) = a_i$. We use $\alpha_{..i}$ for the finite prefix $s_1 a_1 \dots s_i$ of α , $\alpha_{i..}$ for the suffix $s_i a_i \dots$, and $\alpha_{i..j}$ for the infix $s_i a_i \dots s_j$. A finite path is a *cycle* if it starts and ends in the same state and is *simple* if none of its infixes forms a cycle. The *length* of a path α is the number $len(\alpha)$ of actions on α and $len(\alpha) = \infty$ if α is infinite.

A CMDP is *decreasing* if for every cycle $s_1 a_1 s_2 \dots a_{k-1} s_k$ there exists $1 \leq i < k$ such that $C(s_i, a_i) > 0$. Throughout this paper we consider only decreasing CMDPs. The only place where this assumption is used are the proofs of Theorem 4 and Theorem 8.

An infinite path is called a *run*. We typically name runs by variants of the symbol ϱ . The set of all runs in $\mathcal{M}$ is denoted $Runs_{\mathcal{M}}$. A finite path is called *history*. The set of all possible histories of $\mathcal{M}$ is $hist_{\mathcal{M}}$ or simply $hist$. We use $last(\alpha)$ for the last state of α . Let α be a history with $last(\alpha) = s_1$ and $\beta = s_1 a_1 s_2 a_2 \dots$; we define a *joint path* as $\alpha \odot \beta = \alpha a_1 s_2 a_2 \dots$.

A *strategy* for $\mathcal{M}$ is a function $\sigma: hist_{\mathcal{M}} \rightarrow A$ assigning to each history an action to play. A strategy is *memoryless* if $\sigma(\alpha) = \sigma(\beta)$ whenever $last(\alpha) = last(\beta)$. We do not consider randomized strategies in this paper, as they are non-necessary for qualitative ω -regular objectives on finite MDPs [30, 32, 33].

A computation of $\mathcal{M}$ under the control of a given strategy σ from some initial state $s \in S$ creates a path. The path starts with $s_1 = s$. Assume that the current path is α and let $s_i = last(\alpha)$ (we say that $\mathcal{M}$ is currently in s_i). Then the next action on the path is $a_i = \sigma(\alpha)$ and the next state s_{i+1} is chosen randomly according to $\Delta(s_i, a_i)$. Repeating this process *ad infinitum* yields an infinite sample run ϱ . We say that ϱ is σ -*compatible* if it can be produced using this process, and s -*initiated* if it starts in s . We denote the set of all σ -compatible, s -initiated runs by $Comp_{\mathcal{M}}(\sigma, s)$.

We denote by $\mathbb{P}_{\mathcal{M}, s}^\sigma(A)$ the probability that a sample run from $Comp_{\mathcal{M}}(\sigma, s)$ belongs to a given measurable set of runs A . For details on the formal construction of measurable sets of runs as well as the probability measure $\mathbb{P}_{\mathcal{M}, s}^\sigma$ see [2]. Throughout the paper, we drop the $\mathcal{M}$ subscripts in symbols whenever $\mathcal{M}$ is known from the context.

2.1 Resource: Consumption, Levels, and Objectives

We denote by $cap(\mathcal{M})$ the battery capacity in the MDP $\mathcal{M}$. A resource is consumed along paths and can be reloaded in the reload states up to the full capacity. For a path $\alpha = s_1 a_1 s_2 \dots$ we define the consumption of α as $cons(\alpha) = \sum_{i=1}^{len(\alpha)} C(s_i, a_i)$ (since the consumption is non-negative, the sum is always well defined, though possibly diverging). Note that $cons$ does not consider reload states at all. To accurately track the remaining amount of the resource, we use the concept of a *resource level*.

Definition 2 (Resource level). *Let $\mathcal{M}$ be a CMDP with a set of reload states R , let α be a history, and let $0 \leq d \leq cap(\mathcal{M})$ be an integer called initial load.*

Then the energy level after α initialized by d , denoted by $RL_d^M(\alpha)$ or simply as $RL_d(\alpha)$, is defined inductively as follows: for a zero-length history s we have $RL_d^M(s) = d$. For a non-zero-length history $\alpha = \beta a$ we denote $c = C(\text{last}(\beta), a)$, and put

$$RL_d^M(\alpha) = \begin{cases} RL_d^M(\beta) - c & \text{if } \text{last}(\beta) \notin R \text{ and } c \leq RL_d^M(\beta) \neq \perp \\ \text{cap}(\mathcal{M}) - c & \text{if } \text{last}(\beta) \in R \text{ and } c \leq \text{cap}(\mathcal{M}) \text{ and } RL_d^M(\beta) \neq \perp \\ \perp & \text{otherwise} \end{cases}$$

Consider $\mathcal{M}$ from Fig. 1 and the history $\alpha(i) = (s_1 a_2 s_5 a_2)^i s_1$ with i as a parameter. We have $\text{cons}(\alpha(i)) = 3i$ and at the same time, following the inductive definition of $RL_d(\alpha(i))$ we have $RL_d(\alpha(i)) = 19$ for all $i \geq 1$ as the resource is reloaded every time in s_5 . This generalizes into the following. Let α be a history and let $f, l \geq 0$ be the minimal and maximal indices i such that $\alpha_i \in R$, respectively. For $RL_d(\alpha) \neq \perp$, it holds $RL_d(\alpha_{..i}) = d - \text{cons}(\alpha_{..i})$ for all $i \leq f$ and $RL_d(\alpha) = \text{cap}(\mathcal{M}) - \text{cons}(\alpha_{l..})$. Further, for each history α and d such that $e = RL_d(\alpha) \neq \perp$, and each history β suitable for joining with α it holds that $RL_d(\alpha \odot \beta) = RL_e(\beta)$.

A run ϱ is d -safe if and only if the energy level initialized by d is a non-negative number for each finite prefix of ϱ , i.e. if for all $i > 0$ we have $RL_d(\varrho_{..i}) \neq \perp$. We say that a run is safe if it is $\text{cap}(\mathcal{M})$ -safe. The next lemma follows immediately from the definition of an energy level.

Lemma 1. *Let $\varrho = s_1 a_1 s_2 \dots$ be a d -safe run for some d and let α be a history such that $\text{last}(\alpha) = s_1$. Then the run $\alpha \odot \varrho$ is e -safe if $RL_e(\alpha) \geq d$.*

Example 1. Recall the CMDP and the parameterized history $\alpha(i)$ from above. We know that $RL_2(\alpha(i)) = 19$ for all i . Therefore, a strategy that always picks a_2 in s_1 is d -safe in s_1 for all $d \geq 2$. On the other hand, a strategy that always picks a_1 in s_1 is *not* d -safe in s_1 for any $0 \leq d \leq 20 = \text{cap}(\mathcal{M})$ because for all runs ϱ that visit s_3 at least three times before s_2 we have $RL_d(\varrho) = \perp$.

Objectives. An *objective* is a set of runs. The objective $\text{SafeRuns}(d)$ contains exactly d -safe runs. Given a *target set* $T \subseteq S$ and $i \in \mathbb{N}$, we define $\text{Reach}_T^i = \{\varrho \in \text{Runs} \mid \varrho_j \in T \text{ for some } 1 \leq j \leq i + 1\}$ to be the set of all runs that reach some state from T within the first i steps. We put $\text{Reach}_T = \bigcup_{i \in \mathbb{N}} \text{Reach}_T^i$. Finally, the set $\text{Büchi}_T = \{\varrho \in \text{Runs} \mid \varrho_i \in T \text{ for infinitely many } i \in \mathbb{N}\}$.

Problems. We solve three main qualitative problems for CMDPs, namely *safety*, *positive reachability*, and *Büchi*.

Let us fix a state s and a target set of states T . We say that a strategy σ is d -safe in s if $\text{Comp}(\sigma, s) \subseteq \text{SafeRuns}(d)$. We say that σ is T -positive d -safe in s if it is d -safe in s and $\mathbb{P}_s^\sigma(\text{Reach}_T) > 0$, which means that there exists a run in $\text{Comp}(\sigma, s)$ that visits T . Finally, we say that σ is T -Büchi d -safe in a state s if it is d -safe in s and $\mathbb{P}_s^\sigma(\text{Büchi}_T) = 1$.

The vectors Safe , SafePR_T (PR for “positive reachability”), and SafeBüchi_T of type $\overline{\mathbb{N}}^S$ contain, for each $s \in S$, the minimal d such that there exists a strategy

that is d -safe in s , T -positive d -safe in s , and T -Büchi d -safe in s , respectively, and ∞ if no such strategy exists.

The problems we consider for a given CMDP are:

- *Safety*: compute the vector Safe and a strategy that is $\text{Safe}(s)$ -safe in every $s \in S$.
- *Positive reachability*: compute the vector SafePR_T and a strategy that is T -positive $\text{SafePR}_T(s)$ -safe in every state s .
- *Büchi*: compute SafeBüchi_T and a strategy that is T -Büchi $\text{SafeBüchi}_T(s)$ -safe in every state s .

Example 2. Now consider again the d -safe strategy from Example 1 that always picks a_2 ; such a strategy is 2-safe in s_1 , but is not useful if we attempt to eventually reach T . Hence memoryless strategies are not sufficient in our setting. Consider, instead, a strategy σ that picks a_1 in s_1 whenever the current resource level is at least 10 and picks a_2 otherwise. Such a strategy is 2-safe in s_1 and guarantees reaching s_2 with a positive probability: we need at least 10 units of energy to return to s_5 in the case we are unlucky and picking a_1 leads us to s_3 . If we are lucky, a_1 leads us to s_2 by consuming just 5 units of the resource, witnessing that σ is T -positive. As a matter of fact, during *every* revisit of s_5 there is a $\frac{1}{2}$ chance of hitting s_2 during the next try, so σ actually ensures that s_2 is visited with probability 1.

Solving a CMDP is substantially different from solving a consumption 2-player game [14]. Indeed, imagine that in $\mathcal{M}$ from Fig. 1, the outcome of the action a_1 from state s_1 is resolved by an adversarial player. In such a game, the strategy σ does not produce any run that reaches s_2 . In fact, there would be no strategy that guarantees reaching T in a 2-player game like this at all.

The strategy σ from our example uses finite memory to track the resource level exactly. We describe an efficient representation of such strategies in the next section.

3 Counter Strategies

In this section, we define a succinct representation of finite-memory strategies via so called counter selectors. Under the standard definition, a strategy σ is a *finite memory* strategy, if σ can be encoded by a *memory structure*, a type of finite transducer. Formally, a memory structure is a tuple $\mu = (M, \text{next}, \text{up}, m_0)$ where M is a finite set of *memory elements*, $\text{next}: M \times S \rightarrow A$ is a *next action* function, $\text{up}: M \times S \times A \times S \rightarrow M$ is a *memory update* function, and $m_0: S \rightarrow M$ is the *memory initialization function*. The function up can be lifted to a function $\text{up}^*: M \times \text{hist} \rightarrow M$ as follows.

$$\text{up}^*(m, \alpha) = \begin{cases} m & \text{if } \alpha = s \text{ has length } 0 \\ \text{up}(\text{up}^*(m, \beta), \text{last}(\beta), a, t) & \text{if } \alpha = \beta a t \text{ for some } a \in A \text{ and } t \in S \end{cases}$$

The structure μ encodes a strategy σ_μ such that for each history $\alpha = s_1 a_1 s_2 \dots s_n$ we have $\sigma_\mu(\alpha) = \text{next}(\text{up}^*(m_0(s_1), \alpha), s_n)$.

In our setting, strategies need to track energy levels of histories. Let us fix an CMDP $\mathcal{M} = (S, A, \Delta, C, R, \text{cap})$. A non-exhausted energy level is always a number between 0 and $\text{cap}(\mathcal{M})$, which can be represented with a binary-encoded bounded counter. We call strategies with such counters *finite counter (FC) strategies*. An FC strategy selects actions to play according to *selection rules*.

Definition 3 (Selection rule). A selection rule φ for $\mathcal{M}$ is a partial function from the set $\{0, \dots, \text{cap}(\mathcal{M})\}$ to A . Undefined value for some n is indicated by $\varphi(n) = \perp$.

We use $\text{dom}(\varphi) = \{n \in \{0, \dots, \text{cap}(\mathcal{M})\} \mid \varphi(n) \neq \perp\}$ to denote the domain of φ and we use $\text{Rules}_{\mathcal{M}}$ or simply Rules for the set of all selection rules for $\mathcal{M}$. Intuitively, a selection according to rule φ selects the action that corresponds to the largest value from $\text{dom}(\varphi)$ that is not larger than the current energy level. To be more precise, if $\text{dom}(\varphi)$ consists of numbers $n_1 < n_2 < \dots < n_k$, then the action to be selected in a given moment is $\varphi(n_i)$, where n_i is the largest element of $\text{dom}(\varphi)$ which is less than or equal to the current amount of the resource. In other words, $\varphi(n_i)$ is to be selected if the current resource level is in $[n_i, n_{i+1})$ (putting $n_{k+1} = \infty$).

Definition 4 (Counter selector). A counter selector for $\mathcal{M}$ is a function $\Sigma: S \rightarrow \text{Rules}$.

A counter selector itself is not enough to describe a strategy. A strategy needs to keep track of the energy level throughout the path. With a vector $\mathbf{r} \in \{0, \dots, \text{cap}(\mathcal{M})\}^S$ of initial resource levels, each counter selector Σ defines a strategy $\Sigma^{\mathbf{r}}$ that is encoded by the following memory structure $(M, \text{next}, \text{up}, m_0)$ with $a \in A$ being a globally fixed action (for uniqueness). We stipulate that $\perp < n$ for all $n \in \mathbb{N}$.

- $M = \{\perp\} \cup \{0, \dots, \text{cap}(\mathcal{M})\}$.
- Let $m \in M$ be a memory element, let $s \in S$ be a state, let $n \in \text{dom}(\Sigma(s))$ be the largest element of $\text{dom}(\Sigma(s))$ such that $n \leq m$. Then $\text{next}(m, s) = \Sigma(s)(n)$ if n exists, and $\text{next} = a$ otherwise.
- The function up is defined for each $m \in M, a \in A, s, t \in S$ as follows.

$$\text{up}(m, s, a, t) = \begin{cases} m - C(s, a) & \text{if } s \notin R \text{ and } C(s, a) \leq m \neq \perp \\ \text{cap}(\mathcal{M}) - C(s, a) & \text{if } s \in R \text{ and } C(s, a) \leq \text{cap}(\mathcal{M}) \text{ and } m \neq \perp \\ \perp & \text{otherwise.} \end{cases}$$

- The function m_0 is $m_0(s) = \mathbf{r}(s)$.

A strategy σ is a finite counter (FC) strategy if there is a counter selector Σ and a vector $\mathbf{r}$ such that $\sigma = \Sigma^{\mathbf{r}}$. The counter selector can be imagined as a finite-state device that implements σ using $\mathcal{O}(\log(\text{cap}(\mathcal{M})))$ bits of additional memory (counter) used to represent numbers $0, 1, \dots, \text{cap}(\mathcal{M})$. The device uses the counter to keep track of the current resource level, the element $\perp$ representing energy exhaustion. Note that a counter selector can be exponentially more succinct than the corresponding memory structure.

Example 3. Consider again the CMDP $\mathcal{M}$ in Fig. 1 and a counter selector Σ defined as follows: Let φ be a selection rule with $\text{dom}(\varphi) = \{0, 10\}$ such that $\varphi(0) = a_2$ and $\varphi(10) = a_1$. Then let φ' be a selection rule such that $\text{dom}(\varphi') = \{0\}$ and $\varphi'(0) = a_1$. Finally, let Σ be a counter selector such that $\Sigma(s_1) = \varphi$ and $\Sigma(s_i) = \varphi'$ for all $i \neq 1$. Then, for a vector of initial resource levels $\mathbf{r}$, the strategy σ informally described in Example 2 can be formally represented by putting $\sigma = \Sigma^{\mathbf{r}}$. Note that for any $\mathbf{r}$ with $\mathbf{r}(s_1) \geq 2$, $\mathbf{r}(s_2) \geq 0$, $\mathbf{r}(s_3) \geq 5$, $\mathbf{r}(s_4) \geq 4$, and $\mathbf{r}(s_5) \geq 0$ and for any state s of $\mathcal{M}$ the strategy $\Sigma^{\mathbf{r}}$ is $\mathbf{r}(s)$ -safe in s .

4 Safety

In this section, we present an algorithm that computes, for each state, the minimal value d (if it exists) such that there exists a d -safe strategy from that state. We also provide the corresponding strategy. In the remainder of the section we fix an MDP $\mathcal{M}$.

A d -safe run has the following two properties: (i) It consumes at most d units of the resource (energy) before it reaches the first reload state, and (ii) it never consumes more than $\text{cap}(\mathcal{M})$ units of the resource between 2 visits of reload states. To ensure (ii), we need to identify a maximal subset $R' \subseteq R$ of reload states for which there is a strategy σ that, starting in some $r \in R'$, can always reach R' again (within at least one step) using at most $\text{cap}(\mathcal{M})$ resource units. The d -safe strategy we seek can be then assembled from σ and from a strategy that suitably navigates towards R' , which is needed for (i).

In the core of both properties (i) and (ii) lies the problem of *minimum cost reachability*. Hence, in the next subsection, we start with presenting necessary results on this problem.

4.1 Minimum Cost Reachability

The problem of minimum cost reachability with non-negative costs was studied before [46]. Here we present a simple approach to the problem used in our implementation and most of the technical details are available in the full version.

Definition 5. Let $T \subseteq S$ be a set of target states, let $\alpha = s_1 a_1 s_2 \dots$ be a finite or infinite path, and let $1 \leq f$ be the smallest index such that $s_f \in T$. We define consumption of α to T as $\text{ReachCons}_{\mathcal{M},T}(\alpha) = \text{cons}(\alpha_{..f})$ if f exists and we set $\text{ReachCons}_{\mathcal{M},T}(\alpha) = \infty$ otherwise. For a strategy σ and a state $s \in S$ we define $\text{ReachCons}_{\mathcal{M},T}(\sigma, s) = \sup_{\varrho \in \text{Comp}(\sigma, s)} \text{ReachCons}_{\mathcal{M},T}(\varrho)$.

A minimum cost reachability of T from s is a vector defined as

$$\text{MinReach}_{\mathcal{M},T}(s) = \inf \{ \text{ReachCons}_{\mathcal{M},T}(\sigma, s) \mid \sigma \text{ is a strategy for } \mathcal{M} \}.$$

Intuitively, $d = \text{MinReach}_T(s)$ is the minimal initial load with which some strategy can ensure reaching T with consumption at most d , when starting

in s . We say that a strategy σ is optimal for MinReach_T if we have that $\text{MinReach}_T(s) = \text{ReachCons}_T(\sigma, s)$ for all states $s \in S$.

We also define functions $\text{ReachCons}_{\mathcal{M},T}^+$ and the vector $\text{MinReach}_{\mathcal{M},T}^+$ in a similar fashion with one exception: we require the index f from definition of $\text{ReachCons}_{\mathcal{M},T}(\alpha)$ to be strictly larger than 1, which enforces to take at least one step to reach T .

For the rest of this section, fix a target set T and consider the following functional $\mathcal{F}$:

$$\mathcal{F}(\mathbf{v})(s) = \begin{cases} \min_{a \in A} (C(s, a) + \max_{t \in \text{Succ}(s, a)} \mathbf{v}(t)) & s \notin T \\ 0 & s \in T \end{cases}$$

$\mathcal{F}$ is a simple generalization of the standard Bellman functional used for computing shortest paths in graphs. The proof of the following Theorem is rather standard and moved to the full version of the paper.

Theorem 2. *Denote by n the length of the longest simple path in $\mathcal{M}$. Let $\mathbf{x}_T$ be a vector such that $\mathbf{x}_T(s) = 0$ if $s \in T$ and $\mathbf{x}_T(s) = \infty$ otherwise. Then iterating $\mathcal{F}$ on $\mathbf{x}_T$ yields a fixpoint in at most n steps and this fixpoint equals MinReach_T .*

To compute $\text{MinReach}_{\mathcal{M},T}^+$, we construct a new CMDP $\widetilde{\mathcal{M}}$ from $\mathcal{M}$ by adding a copy $\tilde{s}$ of each state $s \in S$ such that dynamics in $\tilde{s}$ is the same as in s ; i.e. for each $a \in A$, $\Delta(\tilde{s}, a) = \Delta(s, a)$ and $C(\tilde{s}, a) = C(s, a)$. We denote the new state set as $\tilde{S}$. We don't change the set of reload states, so $\tilde{s}$ is *never* in T , even if s is. Given the new CMDP $\widetilde{\mathcal{M}}$ and the new state set as $\tilde{S}$, the following lemma is straightforward.

Lemma 2. *Let $\mathcal{M}$ be a CMDP and let $\widetilde{\mathcal{M}}$ be the CMDP constructed as above. Then for each state s of $\mathcal{M}$ it holds $\text{MinReach}_{\mathcal{M},T}^+(s) = \text{MinReach}_{\widetilde{\mathcal{M}},T}(\tilde{s})$.*

4.2 Safely Reaching Reload States

In the following, we use $\text{MinInitCons}_{\mathcal{M}}$ (read *minimal initial consumption*) for the vector $\text{MinReach}_{\mathcal{M},R}^+$ – minimal resource level that ensures we can surely reach a reload state in at least one step. By Lemma 2 and Theorem 2 we can construct $\widetilde{\mathcal{M}}$ and iterate the operator $\mathcal{F}$ for $|S|$ steps to compute $\text{MinInitCons}_{\mathcal{M}}$. Note that S is the state space of $\mathcal{M}$ since introducing the new states into $\widetilde{\mathcal{M}}$ did not increase the length of the maximal simple path. However, we can avoid the construction of $\widetilde{\mathcal{M}}$ and still compute $\text{MinInitCons}_{\mathcal{M}}$ using a *truncated* version of the functional $\mathcal{F}$, which is the approach used in our implementation. We first introduce the following truncation operator:

$$\|\mathbf{x}\|_{\mathcal{M}}(s) = \begin{cases} \mathbf{x}(s) & \text{if } s \notin R, \\ 0 & \text{if } s \in R. \end{cases}$$

Algorithm 1: Algorithm for computing $MinInitCons_{\mathcal{M}}$.**Input:** CMDP $\mathcal{M} = (S, A, \Delta, C, R, cap)$ **Output:** The vector $MinInitCons_{\mathcal{M}}$

```

1 initialize  $\mathbf{x} \in \overline{\mathbb{N}}^S$  to be  $\infty$  in every component;
2 repeat
3    $\mathbf{x}_{old} \leftarrow \mathbf{x}$ ;
4   foreach  $s \in S$  do
5      $c \leftarrow \min_{a \in A} \left\{ C(s, a) + \max_{s' \in Succ(s, a)} \|\mathbf{x}_{old}\|_{\mathcal{M}(s')} \right\}$ ;
6     if  $c < \mathbf{x}(s)$  then
7        $\mathbf{x}(s) \leftarrow c$ ;
8 until  $\mathbf{x}_{old} = \mathbf{x}$ ;
9 return  $\mathbf{x}$ 

```

Then, we define a truncated functional $\mathcal{G}$ as follows:

$$\mathcal{G}(\mathbf{v})(s) = \min_{a \in A} \left(C(s, a) + \max_{s' \in Succ(s, a)} \|\mathbf{v}\|_{\mathcal{M}(s')} \right).$$

The following lemma connects the iteration of $\mathcal{G}$ on $\mathcal{M}$ with the iteration of $\mathcal{F}$ on $\widetilde{\mathcal{M}}$.

Lemma 3. *Let $\infty \in \overline{\mathbb{N}}^S$ be a vectors with all components equal to ∞ . Consider iterating $\mathcal{G}$ on ∞ in $\mathcal{M}$ and $\mathcal{F}$ on $\mathbf{x}_R$ in $\widetilde{\mathcal{M}}$. Then for each $i \geq 0$ and each $s \in R$ we have $\mathcal{G}^i(\infty)(s) = \mathcal{F}^i(\mathbf{x}_R)(\tilde{s})$ and for every $s \in S \setminus R$ we have $\mathcal{G}^i(\infty)(s) = \mathcal{F}^i(\mathbf{x}_R)(s)$.*

Algorithm 1 uses $\mathcal{G}$ to compute the vector $MinInitCons_{\mathcal{M}}$.

Theorem 3. *Algorithm 1 correctly computes the vector $MinInitCons_{\mathcal{M}}$. Moreover, the repeat-loop terminates after at most $|S|$ iterations.*

4.3 Solving the Safety Problem

We want to identify a set $R' \subseteq R$ such that we can reach R' in at least 1 step and with consumption at most $cap = cap(\mathcal{M})$, from each $r \in R'$. This entails identifying the maximal $R' \subseteq R$ such that $MinInitCons_{\mathcal{M}(R')} \leq cap$ for each $r \in R'$. This can be done by initially setting $R' = R$ and iteratively removing states that have $MinInitCons_{\mathcal{M}(R')} > cap$, from R' , as in Algorithm 2.

Theorem 4. *Algorithm 2 computes the vector $Safe_{\mathcal{M}}$ in polynomial time.*

Proof. The algorithm clearly terminates. Computing $MinInitCons_{\mathcal{M}(Rel)}$ on line 5 takes a polynomial number of steps per call due to Theorem 3 and since $\mathcal{M}(Rel)$ has asymptotically the same size as $\mathcal{M}$. Since the repeat loop performs at most $|R|$ iterations, the complexity follows.

Algorithm 2: Computing the vector $\text{Safe}_{\mathcal{M}}$.

Input: CMDP $\mathcal{M}$
Output: The vector $\text{Safe}_{\mathcal{M}}$

```

1  $\text{cap} \leftarrow \text{cap}(\mathcal{M});$ 
2  $\text{Rel} \leftarrow R; \text{ToRemove} \leftarrow \emptyset;$ 
3 repeat
4    $\text{Rel} \leftarrow \text{Rel} \setminus \text{ToRemove};$ 
5    $\mathbf{mic} \leftarrow \text{MinInitCons}_{\mathcal{M}(\text{Rel})};$ 
6    $\text{ToRemove} \leftarrow \{r \in \text{Rel} \mid \mathbf{mic}(r) > \text{cap}\};$ 
7 until  $\text{ToRemove} = \emptyset;$ 
8 foreach  $s \in S$  do
9   if  $\mathbf{mic}(s) > \text{cap}$  then  $\text{out}(s) = \infty;$ 
10  else  $\text{out}(s) = \mathbf{mic}(s);$ 
11 return out
```

As for correctness, we first prove that $\mathbf{out} \leq \text{Safe}_{\mathcal{M}}$. It suffices to prove for each $s \in S$ that upon termination, $\mathbf{mic}(s) \leq \text{Safe}_{\mathcal{M}}(s)$ whenever the latter value is finite. Since $\text{MinInitCons}_{\mathcal{M}'}(s) \leq \text{Safe}_{\mathcal{M}'}(s)$ for each MDP $\mathcal{M}'$ and each its state such that $\text{Safe}_{\mathcal{M}'}(s) < \infty$, it suffices to show that $\text{Safe}_{\mathcal{M}(\text{Rel})} \leq \text{Safe}_{\mathcal{M}}$ is an invariant of the algorithm (as a matter of fact, we prove that $\text{Safe}_{\mathcal{M}(\text{Rel})} = \text{Safe}_{\mathcal{M}}$). To this end, it suffices to show that at every point of execution $\text{Safe}_{\mathcal{M}}(t) = \infty$ for each $t \in R \setminus \text{Rel}$: indeed, if this holds, no strategy that is safe for some state $s \neq t$ can play an action a from s such that $t \in \text{Succ}(s, a)$, so declaring such states non-reloading does not influence the $\text{Safe}_{\mathcal{M}}$ -values. So denote by Rel_i the contents of Rel after the i -th iteration. We prove, by induction on i , that $\text{Safe}_{\mathcal{M}}(s) = \infty$ for all $s \in R \setminus \text{Rel}_i$. For $i = 0$ we have $R = \text{Rel}$, so the statement holds. For $i > 0$, let $s \in R \setminus \text{Rel}_i$, and let σ be any strategy. If some run from $\text{Comp}(\sigma, s)$ visits a state from $R \setminus \text{Rel}_{i-1}$, then σ is not cap -safe, by induction hypothesis. Now assume that all such runs only visit reload states from Rel_{i-1} . Then, since $\text{MinInitCons}_{\mathcal{M}(\text{Rel}_{i-1})}(s) > \text{cap}$, there must be a run $\varrho \in \text{Comp}(\sigma, s)$ with $\text{ReachCons}_{\text{Rel}_{i-1}}^+(\varrho) > \text{cap}$. Assume that ϱ is cap -safe in s . Since we consider only decreasing CMDPs, ϱ must infinitely often visit a reload state (as it cannot get stuck in a zero cycle). Hence, there exists an index $f > 1$ such that $\varrho_f \in \text{Rel}_{i-1}$, and for this f we have $RL_{\text{cap}}(\varrho..f) = \perp$, a contradiction. So again, σ is not safe in s . Since there is no safe strategy from s , we have $\text{Safe}_{\mathcal{M}}(s) = \infty$.

Finally, we need to prove that upon termination, $\mathbf{out} \geq \text{Safe}_{\mathcal{M}}$. Informally, per the definition of $\mathbf{out}$, from every state s we can ensure reaching a state of Rel by consuming at most $\mathbf{out}(s)$ units of the resource. Once in Rel , we can ensure that we can again return to Rel without consuming more than cap units of the resource. Hence, when starting with $\mathbf{out}(s)$ units, we can surely prevent resource exhaustion. $\square$

Definition 6. We call an action a safe in a state s if one of the following conditions holds:

- $s \notin R$ and $C(s, a) + \max_{t \in \text{Succ}(s, a)} \text{Safe}_{\mathcal{M}}(t) \leq \text{Safe}_{\mathcal{M}}(s)$; or
- $s \in R$ and $C(s, a) + \max_{t \in \text{Succ}(s, a)} \text{Safe}_{\mathcal{M}}(t) \leq \text{cap}(\mathcal{M})$.

Note that by the definition of $\text{Safe}_{\mathcal{M}}$, for each state s with $\text{Safe}_{\mathcal{M}}(s) < \infty$ there is always at least one action safe in s . For states s s.t. $\text{Safe}_{\mathcal{M}}(s) = \infty$, we stipulate all actions to be safe in s .

Theorem 5. Any strategy which always selects an action that is safe in the current state is $\text{Safe}_{\mathcal{M}}(s)$ -safe in every state s . In particular, in each consumption MDP $\mathcal{M}$ there is a memoryless strategy σ that is $\text{Safe}_{\mathcal{M}}(s)$ -safe in every state s . Moreover, σ can be computed in polynomial time.

Proof. The first part of the theorem follows directly from Definition 6, Definition 2 (resource levels), and from definition of d -safe runs. The second part is a corollary of Theorem 4 and the fact that in each state, the safe strategy from Definition 6 can fix one such action in each state and thus is memoryless. The complexity follows from Theorem 4. $\square$

Example 4. Consider again the $\mathcal{M}$ from Fig. 1. Algorithm 1 returns, for input $\mathcal{M}$, the vector $\mathbf{mic} = (2, 1, 5, 4, 3)$. Algorithm 2 reuses $\mathbf{mic}$ on line 5 and returns it unchanged. Hence, the vector $\mathbf{mic}$ equals $\text{Safe}_{\mathcal{M}}$. The strategies described in Example 1 witness that $\text{Safe}(s_1) \leq 2$. Here we see that there is no strategy that would be 1-safe in s_1 .

5 Positive Reachability

In this section, we focus on strategies that are safe and such that at least one run they produce visits a given set $T \subseteq S$ of *targets*. The main contribution of this section is Algorithm 3 used to compute such strategies as well as the vector $\text{SafePR}_{\mathcal{M}, T}$ of minimal initial resource levels for which such a strategy exist. As before, for the rest of this section we fix a CMDP $\mathcal{M}$.

We define a function $\text{SPR-Val}_{\mathcal{M}}: S \times A \times \overline{\mathbb{N}}^S \rightarrow \overline{\mathbb{N}}$ (*SPR* for safe positive reachability) s.t. for all $s \in S, a \in A$, and $\mathbf{x} \in \overline{\mathbb{N}}^S$ we have

$$\text{SPR-Val}_{\mathcal{M}}(s, a, \mathbf{x}) = C(s, a) + \min_{t \in \text{Succ}(s, a)} \left\{ \max \{ \mathbf{x}(t), \text{Safe}_{\mathcal{M}}(t') \mid t' \in \text{Succ}(s, a), t' \neq t \} \right\}$$

The max operator considers, for given t , the value $\mathbf{x}(t)$ and the values needed to survive from all possible outcomes of a other than t . Let $v = \text{SPR-Val}_{\mathcal{M}}(s, a, \mathbf{x})$ and t the outcome selected by min. Intuitively, v is the minimal amount of resource needed to reach t with at least $\mathbf{x}(t)$ resource units, or survive if the outcome of a is different from t .

We now define a functional whose fixed point characterizes $\text{SPR-Val}_{\mathcal{M}, T}$. We first define a two-sided version of the truncation operator from the previous section: the operator $\llbracket \cdot \rrbracket_{\mathcal{M}}$ such that

$$\llbracket \mathbf{x} \rrbracket_{\mathcal{M}}(s) = \begin{cases} \infty & \text{if } \mathbf{x}(s) > \text{cap}(\mathcal{M}) \\ \mathbf{x}(s) & \text{if } \mathbf{x}(s) \leq \text{cap}(\mathcal{M}) \text{ and } s \notin R \\ 0 & \text{if } \mathbf{x}(s) \leq \text{cap}(\mathcal{M}) \text{ and } s \in R \end{cases}$$

Using the functions *SPR-Val* and $\llbracket \cdot \rrbracket_{\mathcal{M}}$, we now define an auxiliary operator $\mathcal{A}$ and the main operator $\mathcal{B}$ as follows.

$$\mathcal{A}_{\mathcal{M}}(\mathbf{r})(s) = \begin{cases} \text{Safe}_{\mathcal{M}}(s) & \text{if } s \in T \\ \min_{a \in A} (\text{SPR-Val}_{\mathcal{M}}(s, a, \mathbf{r})) & \text{otherwise;} \end{cases}$$

$$\mathcal{B}_{\mathcal{M}}(\mathbf{r}) = \llbracket \mathcal{A}_{\mathcal{M}}(\mathbf{r}) \rrbracket_{\mathcal{M}}$$

Let SafePR_T^i be the vector such that for a state $s \in S$ the number $d = \text{SafePR}_T^i(s)$ is the minimal number ℓ such that there exists a strategy that is ℓ -safe in s and produces at least one run that visits T within first i steps. Further, we denote by $\mathbf{y}_T$ a vector such that

$$\mathbf{y}_T(s) = \begin{cases} \text{Safe}_{\mathcal{M}}(s) & \text{if } s \in T \\ \infty & \text{if } s \notin T \end{cases}$$

The following lemma can be proved by a rather straightforward but technical induction.

Lemma 4. *Consider the iteration of $\mathcal{B}_{\mathcal{M}}$ on the initial vector $\mathbf{y}_T$. Then for each $i \geq 0$ it holds that $\mathcal{B}_{\mathcal{M}}^i(\mathbf{y}_T) = \text{SafePR}_{\mathcal{M},T}^i$.*

The following lemma says that iterating $\mathcal{B}_{\mathcal{M}}$ reaches a fixed point in a polynomial number of iterations. Intuitively, this is because when trying to reach T , it doesn't make sense to perform a cycle between two visits of a reload state (as this can only increase the resource consumption) and at the same time it doesn't make sense to visit the same reload state twice (since the resource is reloaded to the full capacity upon each visit). The proof is straightforward and is omitted in the interest of brevity. Detailed proofs for Lemma 4 and Lemma 5 are available in the full version of the paper.

Lemma 5. *Let $K = |R| + (|R| + 1) \cdot (|S| - |R| + 1)$. Taking the same initial vector $\mathbf{y}_T$ as in Lemma 4, we have $\mathcal{B}_{\mathcal{M}}^K(\mathbf{y}_T) = \text{SafePR}_{\mathcal{M},T}$.*

The computation of $\text{SafePR}_{\mathcal{M},T}$ and of the associated witness strategy is presented in Algorithm 3.

Example 5. Consider again the CMDP $\mathcal{M}$ from Fig. 1. After one iteration of the loop on line 5, we have $\mathbf{r} = (10, 0, \infty, \infty, \infty)$, as $\mathbf{r}$ is only finite for s_2 before this iteration. In the next iteration, we have $\mathbf{r} = (10, 0, \infty, 12, 0)$. Thus, the next iteration changes the value for s_1 to 2 and in the end, we end up with $\mathbf{r} = (2, 0, 4, 5, 0)$. The iteration with $\mathbf{r}(s_1) = 10$ influences the selector Σ . Note that the computed $\mathbf{r}$ and Σ match those mentioned in Example 3.

Theorem 6. *The Algorithm 3 always terminates after a polynomial number of steps, and upon termination, $\mathbf{r} = \text{SafePR}_{\mathcal{M},T}$.*

Algorithm 3: Positive reachability of T in $\mathcal{M}$ **Input:** CMDP $\mathcal{M}$ with states S , set of target states $T \subseteq S$ **Output:** The vector $\text{SafePR}_{\mathcal{M},T}$, coreresponding rule selector Σ

```

1  $\mathbf{r} \leftarrow \{\infty\}^S$ ;
2 foreach  $s \in S$  s.t.  $\text{Safe}_{\mathcal{M}}(s) < \infty$  do
3    $\Sigma(s)(\text{Safe}_{\mathcal{M}}(s)) \leftarrow$  arbitrary action safe in  $s$ 
4 foreach  $t \in T$  do  $\mathbf{r}(t) \leftarrow \text{Safe}_{\mathcal{M}}(t)$ ;
5 repeat
6    $\mathbf{r}_{old} \leftarrow \mathbf{r}$ ;
7   foreach  $s \in S \setminus T$  do
8      $\mathbf{a}(s) \leftarrow \arg \min_{a \in A} \text{SPR-Val}(s, a, \mathbf{r}_{old})$ ;
9      $\mathbf{r}(s) \leftarrow \min_{a \in A} \text{SPR-Val}(s, a, \mathbf{r}_{old})$ ;
10   $\mathbf{r} \leftarrow \llbracket \mathbf{r} \rrbracket_{\mathcal{M}}$ ;
11  foreach  $s \in S \setminus T$  do
12    if  $\mathbf{r}(s) < \mathbf{r}_{old}(s)$  then
13       $\Sigma(s)(\mathbf{r}(s)) \leftarrow \mathbf{a}(s)$ ;
14 until  $\mathbf{r}_{old} = \mathbf{r}$ ;
15 return  $\mathbf{r}, \Sigma$ 

```

Proof. The repeat loop on lines 1–4 initialize $\mathbf{r}$ to $\mathbf{y}_T$. The repeat loop on lines 5–14 then iterates the operator $\mathcal{B}$. By Lemma 5, the iteration reaches a fixed point in at most K steps, and this fixed point equals $\text{SafePR}_{\mathcal{M},T}$. The complexity bound follows easily, since K is of polynomial magnitude.

The most intricate part of our analysis is extracting a strategy that is T -positive $\text{SafePR}_{\mathcal{M},T}(s)$ -safe in every state s .

Theorem 7. *Let $\mathbf{v} = \text{SafePR}_{\mathcal{M},T}$. Upon termination of Algorithm 3, the computed selector Σ has the property that the finite counter strategy $\Sigma^{\mathbf{v}}$ is, for each state $s \in S$, T -positive $\mathbf{v}(s)$ -safe in s . That is, a polynomial-size finite counter strategy for the positive reachability problem can be computed in polynomial time.*

The rest of this section is devoted to the proof of Theorem 7. The complexity follows from Theorem 6. Indeed, since the algorithm has a polynomial complexity, also the size of Σ is polynomial. The correctness proof is based on the following invariant of the main repeat loop: the finite counter strategy $\pi = \Sigma^{\mathbf{r}}$ has these properties:

- (a) Strategy π is $\text{Safe}_{\mathcal{M}}(s)$ -safe in every state $s \in S$; in particular, we have for $l = \min\{\mathbf{r}(s), \text{cap}(\mathcal{M})\}$ that $RL_l(\alpha) \neq \perp$ for every finite path α produced by π from s .
- (b) For each state $s \in S$ such that $\mathbf{r}(s) \leq \text{cap}(\mathcal{M})$ there exists a π -compatible finite path $\alpha = s_1 a_1 s_2 \dots s_n$ such that $s_1 = s$ and $s_n \in T$ and such that “the resource level with initial load $\mathbf{r}(s)$ never decreases below $\mathbf{r}$ along α ”, which means that for each prefix $\alpha_{..i}$ of α it holds $RL_{\mathbf{r}(s)}(\alpha_{..i}) \geq \mathbf{r}(s_i)$.

The theorem then follows from this invariant (parts (a) and the first half of (b)) and from Theorem 6. We start with the following support invariant, which is easy to prove.

Lemma 6. *The inequality $\mathbf{r} \geq \text{Safe}_{\mathcal{M}}$ is an invariant of the main repeat-loop.*

Proving Part (a) of the Main Invariant. We use the following auxiliary lemma.

Lemma 7. *Assume that Σ is a counter selector such that for all $s \in S$ such that $\text{Safe}(s) < \infty$:*

- (1.) *$\text{Safe}(s) \in \text{dom}(\Sigma(s))$.*
- (2.) *For all $x \in \text{dom}(\Sigma(s))$, for $a = \Sigma(s)(x)$ and for all $t \in \text{Succ}(s, a)$ we have $RL_x(\text{sat}) = d - C(s, a) \geq \text{Safe}(t)$ where $d = x$ for $s \notin R$ and $d = \text{cap}(\mathcal{M})$ otherwise.*

Then for each vector $\mathbf{y} \geq \text{Safe}$ the strategy $\pi = \Sigma^{\mathbf{y}}$ is $\text{Safe}(s)$ -safe in every state s .

Proof. Let s be a state such that $\mathbf{y}(s) < \infty$. It suffices to prove that for every π -compatible finite path α started in s it holds $\perp \neq RL_{\mathbf{y}(s)}(\alpha)$. We actually prove a stronger statement: $\perp \neq RL_{\mathbf{y}(s)}(\alpha) \geq \text{Safe}(\text{last}(\alpha))$. We proceed by induction on the length of α . If $\text{len}(\alpha) = 0$ we have $RL_{\mathbf{y}(s)}(\alpha) = \mathbf{y}(s) \geq \text{Safe}_{\mathcal{M}}(s) \geq 0$. Now let $\alpha = \beta \odot t_1at_2$ for some shorter path β with $\text{last}(\beta) = t_1$ and $a \in A$, $t_1, t_2 \in S$. By induction hypothesis, $l = RL_{\mathbf{y}(s)}(\beta) \geq \text{Safe}_{\mathcal{M}}(t_1)$, from which it follows that $\text{Safe}_{\mathcal{M}}(t_1) < \infty$. Due to (1.), it follows that there exists at least one $x \in \text{dom}(\Sigma(t_1))$ such that $x \leq l$. We select maximal x satisfying the inequality so that $a = \Sigma(t_1)(x)$. We have that $RL_{\mathbf{y}(s)}(\alpha) = RL_l(t_1at_2)$ by definition and from (2.) it follows that $\perp \neq RL_x(t_1at_2) \geq \text{Safe}(t_2) \geq 0$. All together, as $l \geq x$ we have that $RL_{\mathbf{y}(s)}(\alpha) \geq RL_x(t_1at_2) \geq \text{Safe}(t_2) \geq 0$. $\square$

Now we prove the part (a) of the main invariant. We show that throughout the execution of Algorithm 3, Σ satisfies the assumptions of Lemma 7. Property (1.) is ensured by the initialization on line 3. The property (2.) holds upon first entry to the main loop by the definition of a safe action (Definition 6). Now assume that $\Sigma(s)(\mathbf{r}(s))$ is redefined on line 13, and let a be the action $\mathbf{a}(s)$.

We first handle the case when $s \notin R$. Since a was selected on line 8, from the definition of *SPR-Val* we have that there is $t \in \text{Succ}(s, a)$ such that after the loop iteration,

$$\mathbf{r}(s) = C(s, a) + \max\{\mathbf{r}_{\text{old}}(t), \text{Safe}(t') \mid t \neq t' \in \text{Succ}(s, a)\} \geq C(s, a) + \max_{t' \in \text{Succ}(s, a)} \text{Safe}_{\mathcal{M}}(t'), \quad (1)$$

the latter inequality following from Lemma 6. Satisfaction of property (2.) in s then follows immediately from the Eq. (1).

If $s \in R$, then (1) holds before the truncation on line 10, at which point $\mathbf{r}(s) < \text{cap}(\mathcal{M})$. Hence, $\text{cap}(\mathcal{M}) - C(s, a) \geq \max_{t \in \text{Succ}(s, a)} \text{Safe}_{\mathcal{M}}(t)$ as required by (2.). From Lemmas 6 and 7 it follows that $\Sigma^{\mathbf{r}}$ is $\text{Safe}_{\mathcal{M}}(s)$ -safe in every state s . This finishes the proof of part (a) of the invariant.

Proving Part (b) of the Main Invariant. Clearly, (b) holds after initialization. Now assume that an iteration of the main repeat loop was performed. Denote by π_{old} the strategy $\Sigma^{\mathbf{r}_{old}}$ and by π the strategy $\Sigma^{\mathbf{r}}$. Let s be any state such that $\mathbf{r}(s) \leq \text{cap}(\mathcal{M})$. If $\mathbf{r}(s) = \mathbf{r}_{old}(s)$, then we claim that (b) follows directly from the induction hypothesis: indeed, we have that there is an s -initiated π_{old} -compatible path α ending in a target state s.t. the $\mathbf{r}_{old}(s)$ -initiated resource level along α never drops $\mathbf{r}_{old}$, i.e. for each prefix β of α it holds $RL_{\mathbf{r}_{old}(s)}(\beta) \geq \mathbf{r}_{old}(\text{last}(\beta))$. But then β is also π -compatible, since for each state q , $\Sigma(q)$ was only redefined for values smaller than $\mathbf{r}_{old}(q)$.

The case when $\mathbf{r}(s) < \mathbf{r}_{old}(s)$ is treated similarly. As in the proof of part (a), denote by a the action $\mathbf{a}(s)$ assigned on line 13. There must be a state $t \in \text{Succ}(s, a)$ s.t. (1) holds before the truncation on line 10. In particular, for this t it holds $RL_{\mathbf{r}(s)}(sat) \geq \mathbf{r}_{old}(t)$. By induction hypothesis, there is a t -initiated π_{old} -compatible path β ending in T satisfying the conditions in (b). We put $\alpha = sat \odot \beta$. Clearly α is s -initiated and reaches T . Moreover, it is π -compatible. To see this, note that $\Sigma^{\mathbf{r}}(s)(\mathbf{r}(s)) = a$; moreover, the resource level after the first transition is $e(t) = RL_{\mathbf{r}(s)}(sat) \geq \mathbf{r}_{old}(t)$, and due to the assumed properties of β , the $\mathbf{r}_{old}(t)$ -initiated resource level (with initial load $e(t)$) never decreases below $\mathbf{r}_{old}$ along β . Since Σ was only re-defined for values smaller than those given by the vector $\mathbf{r}_{old}$, π mimics π_{old} along β . Since $\mathbf{r} \leq \mathbf{r}_{old}$, we have that along α , the $\mathbf{r}(s)$ -initiated resource level never decreases below $\mathbf{r}$. This finishes the proof of part (b) of the invariant and thus also the proof of Theorem 7. $\square$

6 Büchi

This section proves Theorem 1 which is the main theoretical result of the paper. The proof is broken down into the following steps.

- (1.) We identify a largest set $R' \subseteq R$ of reload states such that from each $r \in R'$ we can reach R' again (in at least one step) while consuming at most cap resource units and restricting ourselves only to strategies that (i) avoid $R \setminus R'$ and (ii) guarantee positive reachability of T in $\mathcal{M}(R')$.
- (2.) We show that $\text{SafeBüchi}_{\mathcal{M}, T} = \text{SafePR}_{\mathcal{M}(R'), T}$ and that the corresponding strategy (computed by Algorithm 3) is also T -Büchi $\text{SafeBüchi}_{\mathcal{M}, T}(s)$ -safe for each $s \in S$.

Algorithm 4: Almost-sure Büchi reachability of T in $\mathcal{M}$.**Input:** CMDP $\mathcal{M} = (S, A, \Delta, C, R, \text{cap})$, target states $T \subseteq S$ **Output:** The largest set $\text{Rel} \subseteq R$ such that $\text{SafePR}_{\mathcal{M}(\text{Rel}), T}(r) \leq \text{cap}$ for all $r \in \text{Rel}$.

```

1  $\text{Rel} \leftarrow R$ ;  $\text{ToRemove} \leftarrow \emptyset$ ;
2 repeat
3    $\text{Rel} \leftarrow \text{Rel} \setminus \text{ToRemove}$ ;
4    $(\text{reach}, \Sigma) \leftarrow \text{SafePR}_{\mathcal{M}(\text{Rel}), T}$ ;
5    $\text{ToRemove} \leftarrow \{r \in \text{Rel} \mid \text{reach}(r) > \text{cap}\}$ ;
6 until  $\text{ToRemove} = \emptyset$ ;
7 return  $\text{reach}, \Sigma$ 

```

Algorithm 4 solves (1.) in a similar fashion as Algorithm 2 handled safety. In each iteration, we declare as non-reloading all states from which positive reachability of T and safety within $\mathcal{M}(\text{Rel})$ cannot be guaranteed. This is repeated until we reach a fixed point. The number of iterations is clearly bounded by $|R|$.

Theorem 8. Let $\mathcal{M} = (S, A, \Delta, C, R, \text{cap})$ be a CMDP and $T \subseteq S$ be a target set. Moreover, let R' be the contents of Rel upon termination of Algorithm 4 for the input $\mathcal{M}$ and T . Finally let $\mathbf{r}$ and Σ be the vector and the selector returned by Algorithm 3 for the input $\mathcal{M}$ and T . Then for every state s , the finite counter strategy $\sigma = \Sigma^{\mathbf{r}}$ is T -Büchi $\mathbf{r}(s)$ -safe in s in both $\mathcal{M}(R')$ and $\mathcal{M}$. Moreover, the vector $\mathbf{r}$ is equal to $\text{SafeBüchi}_{\mathcal{M}, T}$.

Proof. We first show that σ is T -Büchi $\mathbf{r}(s)$ -safe in $\mathcal{M}(R')$ for all $s \in S$ with $\mathbf{r}(s) \leq \text{cap}$. Clearly it is $\mathbf{r}(s)$ -safe, so it remains to prove that T is visited infinitely often with probability 1. We know that upon every visit of a state $r \in R'$, σ guarantees a future visit to T with positive probability. As a matter of fact, since σ is a finite memory strategy, there is $\delta > 0$ such that upon every visit of some $r \in R'$, the probability of a future visit to T is at least δ . As $\mathcal{M}(R')$ is decreasing, every s -initiated σ -compatible run must visit the set R' infinitely many times. Hence, with probability 1 we reach T at least once. The argument can then be repeated from the first point of visit to T to show that with probability 1 we visit T at least twice, three times, etc. *ad infinitum*. By the monotonicity of probability, $\mathbb{P}_{\mathcal{M}, s}^{\sigma}(\text{Büchi}_T) = 1$.

It remains to show that $\mathbf{r} \leq \text{SafeBüchi}_{\mathcal{M}, T}$. Assume that there is a state $s \in S$ and a strategy σ' such that σ' is d -safe in s for some $d < \mathbf{r}(s) = \text{SafePR}_{\mathcal{M}(R'), T}(s)$. We show that this strategy is not T -Büchi d -safe in $\mathcal{M}$. If all σ' -compatible runs reach T , then there must be at least one history α produced by σ' that visits $r \in R \setminus R'$ before reaching T (otherwise $d \geq \mathbf{r}(s)$).

Then either (a) $\text{SafePR}_{\mathcal{M},T}(r) = \infty$, in which case any σ' -compatible extension of α avoids T ; or (b) since $\text{SafePR}_{\mathcal{M}(R'),T}(r) > \text{cap}$, there must be an extension of α that visits, between the visit of r and T , another $r' \in R \setminus R'$ such that $r' \neq r$. We can then repeat the argument, eventually reaching the case (a) or running out of the resource, a contradiction with σ' being d -safe. $\square$

We can finally proceed to prove Theorem 1.

Proof (of Theorem 1). The theorem follows immediately from Theorem 8 since we can (a) compute $\text{SafeBüchi}_{\mathcal{M},T}$ and the corresponding strategy σ_T in polynomial time (see Theorem 7 and Algorithm 4); (b) we can easily check whether $d \geq \text{SafeBüchi}_{\mathcal{M},T}(s)$, if yes, then σ_T is the desired strategy σ ; and (c) represent σ_T in polynomial space as it is a finite counter strategy represented by a polynomial-size counter selector. $\square$

7 Implementation and Case Studies

We implemented the presented algorithms in Python and released it as an open-source tool called *FiMDP* (*Fuel in MDP*) available at <https://github.com/xblahoud/FiMDP>. The docker artifact is available at <https://hub.docker.com/r/xblahoud/fimdp> and can be run without installation via the Binder project [50]. We investigate the practical behavior of our algorithms using two case studies: (1) An autonomous electric vehicle (AEV) routing problem in the streets of Manhattan modeled using realistic traffic and electric car energy consumption data, and (2) a multi-agent grid world model inspired by the Mars Helicopter Scout [8] to be deployed from the planned Mars 2020 rover. The first scenario demonstrates the utility of our algorithm for solving real-world problems [59], while the second scenario studies the algorithm’s scalability limits.

The consumption-Büchi objective can be also solved by a naive approach that encodes the energy constraints in the state space of the MDP, and solves it using techniques for standard MDPs [33]. States of such an MDP are tuples (s, e) where s is a state of the input CMDP and e is the current level of energy. Naturally, all actions that would lead to states with $e < 0$ lead to a special sink state. The standard techniques rely on decomposition of the MDP into maximal end-components (MEC). We implemented the explicit encoding of CMDP into MDP, and the MEC-decomposition algorithm.

All computations presented in the following were performed on a PC with Intel Core i7-8700 3.20 GHz 12 core processor and a RAM of 16 GB running Ubuntu 18.04 LTS. All running times are means from at least 5 runs and the standard deviation was always below 5% among these runs.

7.1 Electric Vehicle Routing

We consider the area in the middle of Manhattan, from 42nd to 116th Street, see Fig. 2. Street intersections and directions of feasible movement form the state and action spaces of the MDP. Intersections in the proximity of real-world fast charging stations [56] represent the set of reload states.

After the AEV picks a direction, it reaches the next intersection in that direction deterministically with a stochastic energy consumption. We base our model of consumption on distributions of vehicle travel times from the area [55] and conversion of velocity and travel times to energy consumption [52]. We discretize the consumption distribution into three possible values (c_1, c_2, c_3) reached with corresponding probabilities (p_1, p_2, p_3) . We then model the transition from one intersection (I_1) to another (I_2) using additional dummy states as explained in Fig. 2.

The corresponding CMDP has 7378 states and 8473 actions. For a fixed set of 100 randomly selected target states, Fig. 3 shows influence of requested capacity on running times for (a) strategy for Büchi objective using CMDP (our approach), and (b) MEC-decomposition for the corresponding explicit MDP. With constant number of states, our algorithm runs rea-

Fig. 2. (Top:) Street network in the considered area. Charging stations are red, one way roads green, and two-way roads blue. (Bottom:) Transition from intersection I_1 to I_2 with stochastic consumption. The small circles are dummy states. (Color figure online)

Fig. 3. Mean computation times for a fixed target set of size 100 and varying capacity: (a) **CMDP** – computing Büchi objective via CMDP, (b) **explicit** – computing MEC decomposition of the explicit MDP, (c) **combined** – (a) and (b) combined for small capacity values.

sonably fast for all capacities and the running time stabilizes for $cap > 95$; this is not the case for the explicit approach where the number of states keeps growing (52747 for $cap = 95$) as well as the running time. The decomposition to MECs is slightly faster than solving Büchi using CMDP for the small capacities (Fig. 3 (c)), but MECs decomposition is only a part of the solution and running the full algorithm for Büchi would most likely diminish this advantage.

7.2 Multi-agent Grid World

We use multi-agent grid world to generate CMDP with huge number of states to study the scalability limits of the proposed algorithms. We model the rover and the helicopter of the Mars 2020 mission with the following realistic considerations: the rover enjoys infinite energy while the helicopter is restricted by batteries recharged at the rover. These two vehicle jointly operate on a mission where the helicopter reaches areas inaccessible to the rover. The outcomes of the helicopter’s actions are deterministic while those of the rover—influenced by terrain dynamics—are stochastic. For a grid world of size n , this system can be naturally modeled as a CMDP with n^4 states. Figure 4 shows the running times of the Büchi objective for growing grid sizes and capacities in CMDP. We observe that the increase in the computational time of CMDP follows the growth in the number of states roughly linearly, and our implementation deals with an MDP with 1.6×10^5 states in no more than seven minutes. The figure also shows the running time for the MEC decomposition of the corresponding explicit MDP when the capacity is 10 and, for certain smaller, computationally feasible grid sizes, when the capacity is 20.

Fig. 4. Mean computation times for varying grid sizes and of size capacities: **(a) CMDP** – computing Büchi objective via CMDP, the gray line shows the corresponding growth in the number of states on separate scale, **(b) explicit** – computing MEC decomposition of the explicit MDP, **(c) combined** – combined computation time for a capacity of 10.

8 Conclusion and Future Work

We presented a first study of consumption Markov decision processes (CMDPs) with qualitative ω -regular objectives. We developed and implemented a polynomial-time algorithm for CMDPs with an objective of probability-1 satisfaction of a given Büchi condition. Possible directions for the future work are extensions to quantitative analysis (e.g. minimizing the expected resource consumption), stochastic games, or partially observable setting.

Acknowledgements. We acknowledge the kind help of Vojtěch Forejt, David Kláška, and Martin Kučera in the discussions leading to this paper.

References

1. Abdulla, P.A., Atig, M.F., Hofman, P., Mayr, R., Kumar, K.N., Totzke, P.: Infinite-state energy games. In: Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic and the 29th Annual ACM/IEEE Symposium on Logic in Computer Science, pp. 7:1–7:10 (2014)
2. Ash, R., Doléans-Dade, C.: Probability and Measure Theory. Harcourt/Academic Press, San Diego (2000)
3. Bacci, G., Bouyer, P., Fahrenberg, U., Larsen, K.G., Markey, N., Reynier, P.-A.: Optimal and robust controller synthesis. In: Havelund, K., Peleska, J., Roscoe, B., de Vink, E. (eds.) FM 2018. LNCS, vol. 10951, pp. 203–221. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-95582-7_12
4. Baier, C., Chrszon, P., Dubslaff, C., Klein, J., Klüppelholz, S.: Energy-utility analysis of probabilistic systems with exogenous coordination. In: de Boer, F., Bonsangue, M., Rutten, J. (eds.) It's All About Coordination. LNCS, vol. 10865, pp. 38–56. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-90089-6_3
5. Baier, C., Daum, M., Dubslaff, C., Klein, J., Klüppelholz, S.: Energy-utility quantiles. In: Badger, J.M., Rozier, K.Y. (eds.) NFM 2014. LNCS, vol. 8430, pp. 285–299. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06200-6_24
6. 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
7. Baier, C., Dubslaff, C., Klüppelholz, S., Leuschner, L.: Energy-utility analysis for resilient systems using probabilistic model checking. In: Ciardo, G., Kindler, E. (eds.) PETRI NETS 2014. LNCS, vol. 8489, pp. 20–39. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07734-5_2
8. Balaram, B., et al.: Mars helicopter technology demonstrator. In: AIAA Atmospheric Flight Mechanics Conference (2018)
9. Boker, U., Henzinger, T.A., Radhakrishna, A.: Battery transition systems. In: 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 595–606 (2014)
10. Bouyer, P., Fahrenberg, U., Larsen, K.G., Markey, N.: Timed automata with observers under energy constraints. In: 13th ACM International Conference on Hybrid Systems: Computation and Control, pp. 61–70. ACM (2010)

11. Bouyer, P., Fahrenberg, U., Larsen, K.G., Markey, N., Srba, J.: Infinite runs in weighted timed automata with energy constraints. In: Cassez, F., Jard, C. (eds.) FORMATS 2008. LNCS, vol. 5215, pp. 33–47. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85778-5_4
12. Bouyer, P., Hofman, P., Markey, N., Randour, M., Zimmermann, M.: Bounding average-energy games. In: Esparza, J., Murawski, A.S. (eds.) FoSSaCS 2017. LNCS, vol. 10203, pp. 179–195. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54458-7_11
13. Bouyer, P., Markey, N., Randour, M., Larsen, K.G., Laursen, S.: Average-energy games. *Acta Informatica* **55**(2), 91–127 (2018)
14. Brázdil, T., Chatterjee, K., Kučera, A., Novotný, P.: Efficient controller synthesis for consumption games with multiple resource types. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 23–38. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31424-7_8
15. Brázdil, T., Jančar, P., Kučera, A.: Reachability games on extended vector addition systems with states. In: Abramsky, S., Gavioille, C., Kirchner, C., Meyer auf der Heide, F., Spirakis, P.G. (eds.) ICALP 2010. LNCS, vol. 6199, pp. 478–489. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14162-1_40
16. Brázdil, T., Kláška, D., Kučera, A., Novotný, P.: Minimizing running costs in consumption systems. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 457–472. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08867-9_30
17. Brázdil, T., Kučera, A., Novotný, P.: Optimizing the expected mean payoff in energy Markov decision processes. In: Artho, C., Legay, A., Peled, D. (eds.) ATVA 2016. LNCS, vol. 9938, pp. 32–49. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46520-3_3
18. Brenguier, R., Cassez, F., Raskin, J.-F.: Energy and mean-payoff timed games. In: 17th International Conference on Hybrid Systems: Computation and Control, pp. 283–292 (2014)
19. Brihaye, T., Geeraerts, G., Haddad, A., Monmege, B.: Pseudopolynomial iterative algorithm to solve total-payoff games and min-cost reachability games. *Acta Informatica* **54**(1), 85–125 (2017)
20. Brim, L., Chaloupka, J., Doyen, L., Gentilini, R., Raskin, J.: Faster algorithms for mean-payoff games. *Form. Methods Syst. Des.* **38**(2), 97–118 (2011)
21. Bruyère, V., Hautem, Q., Randour, M., Raskin, J.-F.: Energy mean-payoff games. In: 30th International Conference on Concurrency Theory, pp. 21:1–21:17 (2019)
22. Cachera, D., Fahrenberg, U., Legay, A.: An ω -algebra for real-time energy problems. *Log. Methods Comput. Sci.* **15**(2) (2019)
23. Chakrabarti, A., de Alfaro, L., Henzinger, T.A., Stoelinga, M.: Resource interfaces. In: Alur, R., Lee, I. (eds.) EMSOFT 2003. LNCS, vol. 2855, pp. 117–133. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45212-6_9
24. Chaloupka, J.: Z-reachability problem for games on 2-dimensional vector addition systems with states is in P. *Fundamenta Informaticae* **123**(1), 15–42 (2013)
25. Chatterjee, K.: Stochastic ω -regular games. Ph.D. thesis, University of California, Berkeley (2007)
26. Chatterjee, K., Doyen, L.: Energy and mean-payoff parity Markov decision processes. In: Murlak, F., Sankowski, P. (eds.) MFCS 2011. LNCS, vol. 6907, pp. 206–218. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22993-0_21
27. Chatterjee, K., Doyen, L.: Energy parity games. *Theor. Comput. Sci.* **458**, 49–60 (2012)

28. Chatterjee, K., Doyen, L., Henzinger, T., Raskin, J.-F.: Generalized mean-payoff and energy games. In: 30th Annual Conference on Foundations of Software Technology and Theoretical Computer Science, pp. 505–516 (2010)
29. Chatterjee, K., Henzinger, M., Krinninger, S., Nanongkai, D.: Polynomial-time algorithms for energy games with special weight structures. In: 20th Annual European Symposium on Algorithms, pp. 301–312 (2012)
30. Chatterjee, K., Jurdziński, M., Henzinger, T.: Quantitative stochastic parity games. In: 15th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 121–130 (2004)
31. Chatterjee, K., Randour, M., Raskin, J.-F.: Strategy synthesis for multi-dimensional quantitative objectives. *Acta informatica* **51**(3–4), 129–163 (2014)
32. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. *J. ACM* **42**(4), 857–907 (1995)
33. de Alfaro, L.: Formal verification of probabilistic systems. Ph.D. thesis, Stanford University (1998)
34. Degorre, A., Doyen, L., Gentilini, R., Raskin, J.-F., Toruńczyk, S.: Energy and mean-payoff games with imperfect information. In: Dawar, A., Veith, H. (eds.) *CSL 2010. LNCS*, vol. 6247, pp. 260–274. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15205-4_22
35. Ésik, Z., Fahrenberg, U., Legay, A., Quaas, K.: An algebraic approach to energy problems I - continuous Kleene ω -algebras. *Acta Cybernetica* **23**(1), 203–228 (2017)
36. Ésik, Z., Fahrenberg, U., Legay, A., Quaas, K.: An algebraic approach to energy problems II - the algebra of energy functions. *Acta Cybernetica* **23**(1), 229–268 (2017)
37. Fahrenberg, U., Juhl, L., Larsen, K.G., Srba, J.: Energy games in multiweighted automata. In: Cerone, A., Pihlajasaari, P. (eds.) *ICTAC 2011. LNCS*, vol. 6916, pp. 95–115. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23283-1_9
38. Fahrenberg, U., Legay, A.: Featured weighted automata. In: 5th International FME Workshop on Formal Methods in Software Engineering, pp. 51–57 (2017)
39. Fijalkow, N., Zimmermann, M.: Cost-parity and cost-Streett games. In: 32nd Annual Conference on Foundations of Software Technology and Theoretical Computer Science, pp. 124–135 (2012)
40. Filiot, E., Gentilini, R., Raskin, J.-F.: Quantitative languages defined by functional automata. In: Koutny, M., Ulidowski, I. (eds.) *CONCUR 2012. LNCS*, vol. 7454, pp. 132–146. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32940-1_11
41. Hahn, E.M., Perez, M., Schewe, S., Somenzi, F., Trivedi, A., Wojtczak, D.: Good-for-MDPs automata for probabilistic analysis and reinforcement learning. *TACAS 2020. LNCS*, vol. 12078, pp. 306–323. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45190-5_17
42. Herrmann, L., Baier, C., Fetzner, C., Klüppelholz, S., Napierkowski, M.: Formal parameter synthesis for energy-utility-optimal fault tolerance. In: Bakhshi, R., Ballarini, P., Barbot, B., Castel-Taleb, H., Remke, A. (eds.) *EPEW 2018. LNCS*, vol. 11178, pp. 78–93. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-02227-3_6
43. Juhl, L., Guldstrand Larsen, K., Raskin, J.-F.: Optimal bounds for multiweighted and parametrised energy games. In: Liu, Z., Woodcock, J., Zhu, H. (eds.) *Theories of Programming and Formal Methods. LNCS*, vol. 8051, pp. 244–255. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39698-4_15

44. Jurdziński, M., Lazić, R., Schmitz, S.: Fixed-dimensional energy games are in pseudo-polynomial time. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015. LNCS, vol. 9135, pp. 260–272. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47666-6_21
45. Jurdziński, M.: Deciding the winner in parity games is in $UP \cap co-UP$. *Inf. Process. Lett.* **68**(3), 119–124 (1998)
46. Khachiyan, L., et al.: On short paths interdiction problems: total and node-wise limited interdiction. *Theory Comput. Syst.* **43**(2), 204–233 (2008)
47. Klačka, D.: Complexity of Consumption Games. Bachelor's thesis, Masaryk University (2014)
48. Larsen, K.G., Laursen, S., Zimmermann, M.: Limit your consumption! Finding bounds in average-energy games. In: 14th International Workshop Quantitative Aspects of Programming Languages and Systems, pp. 1–14 (2016)
49. Mayr, R., Schewe, S., Totzke, P., Wojtczak, D.: MDPs with energy-parity objectives. In: 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, pp. 1–12 (2017)
50. Jupyter, P., et al.: Binder 2.0 - reproducible, interactive, sharable environments for science at scale. In: 17th Python in Science Conference, pp. 113–120 (2018)
51. Sickert, S., Esparza, J., Jaax, S., Křetínský, J.: Limit-deterministic Büchi automata for linear temporal logic. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9780, pp. 312–332. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41540-6_17
52. Straubel, J.B.: Roadster efficiency and range (2008). <https://www.tesla.com/blog/roadster-efficiency-and-range>
53. Sugumar, G., Selvamuthukumar, R., Dragicevic, T., Nyman, U., Larsen, K.G., Blaabjerg, F.: Formal validation of supervisory energy management systems for microgrids. In: 43rd Annual Conference of the IEEE Industrial Electronics Society, pp. 1154–1159 (2017)
54. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press, Cambridge (2018)
55. Uber Movement: Traffic speed data for New York City (2019). <https://movement.uber.com/>
56. United States Department of Energy. Alternative fuels data center (2019). <https://afdc.energy.gov/stations/>
57. Velner, Y., Chatterjee, K., Doyen, L., Henzinger, T.A., Rabinovich, A.M., Raskin, J.: The complexity of multi-mean-payoff and multi-energy games. *Inf. Comput.* **241**, 177–196 (2015)
58. Wognsen, E.R., Hansen, R.R., Larsen, K.G., Koch, P.: Energy-aware scheduling of FIR filter structures using a timed automata model. In: 19th International Symposium on Design and Diagnostics of Electronic Circuits and Systems, pp. 1–6 (2016)
59. Zhang, H., Sheppard, C.J.R., Lipman, T.E., Moura, S.J.: Joint fleet sizing and charging system planning for autonomous electric vehicles. *IEEE Trans. Intell. Transp. Syst.* (2019)

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.