



# Anytime Guarantees for Reachability in Uncountable Markov Decision Processes

Kush Grover  

Technische Universität München, Germany

Jan Křetínský  

Technische Universität München, Germany

Tobias Meggendorfer   

Institute of Science and Technology Austria, Wien, Austria

Maximilian Weininger  

Technische Universität München, Germany

---

## Abstract

---

We consider the problem of approximating the reachability probabilities in Markov decision processes (MDP) with uncountable (continuous) state and action spaces. While there are algorithms that, for special classes of such MDP, provide a sequence of approximations converging to the true value in the limit, our aim is to obtain an algorithm with guarantees on the precision of the approximation.

As this problem is undecidable in general, assumptions on the MDP are necessary. Our main contribution is to identify sufficient assumptions that are as weak as possible, thus approaching the “boundary” of which systems can be correctly and reliably analyzed. To this end, we also argue why each of our assumptions is necessary for algorithms based on processing finitely many observations.

We present two solution variants. The first one provides converging lower bounds under weaker assumptions than typical ones from previous works concerned with guarantees. The second one then utilizes stronger assumptions to additionally provide converging upper bounds. Altogether, we obtain an *anytime* algorithm, i.e. yielding a sequence of approximants with known and iteratively improving precision, converging to the true value in the limit. Besides, due to the generality of our assumptions, our algorithms are very general templates, readily allowing for various heuristics from literature in contrast to, e.g., a specific discretization algorithm. Our theoretical contribution thus paves the way for future practical improvements without sacrificing correctness guarantees.

**2012 ACM Subject Classification** Mathematics of computing → Markov processes; Mathematics of computing → Continuous mathematics; Computing methodologies → Continuous models

**Keywords and phrases** Uncountable system, Markov decision process, discrete-time Markov control process, probabilistic verification, anytime guarantee

**Digital Object Identifier** 10.4230/LIPIcs.CONCUR.2022.11

**Related Version** *Full Version*: <https://arxiv.org/abs/2008.04824> [17]

**Funding** *Kush Grover*: The author has been supported by the DFG research training group GRK 2428 *ConVeY*.

*Maximilian Weininger*: The author has been partially supported by DFG projects 383882557 *Statistical Unbounded Verification (SUV)* and 427755713 *Group-By Objectives in Probabilistic Verification (GOPro)*.

## 1 Introduction

The standard formalism for modelling systems with both non-deterministic and probabilistic behaviour are Markov decision processes (MDP) [43]. In the context of many applications such as cyber-physical systems, states and actions are used to model real-valued phenomena like position or throttle. Consequently, the state space and the action space may be uncountably



© Kush Grover, Jan Křetínský, Tobias Meggendorfer, and Maximilian Weininger; licensed under Creative Commons License CC-BY 4.0

33rd International Conference on Concurrency Theory (CONCUR 2022).

Editors: Bartek Klin, Slawomir Lasota, and Anca Muscholl; Article No. 11; pp. 11:1–11:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

infinite. For example, the intervals  $[a, b] \times [c, d] \subseteq \mathbb{R}^2$  can model a safe area for a robot to move in or a set of available control inputs such as acceleration and steering angle. This gives rise to MDP with uncountable state- and action-spaces (sometimes called controlled discrete-time Markov process [51, 52] or discrete-time Markov control process [11, 28]), with applications ranging from modelling a Mars rover [10, 24], over water reservoir control [36] and warehouse storage management [38], to energy control [51], and many more [41].

Although systems modelled by MDP are often safety-critical, the analysis of uncountable systems is so complex that practical approaches for verification and controller synthesis are usually based on “best effort” learning techniques, for example *reinforcement learning*. While efficient in practice, these methods guarantee, even in the best case, convergence to the true result only in the limit, e.g. [40], or for increasingly precise discretization, e.g. [51, 32]. In line with the tradition of learning and to make the analysis more feasible, the typical objectives considered for MDP are either finite-horizon [37, 3] or discounted properties [18, 53, 25], together with restrictive assumptions. Note that when it comes to approximation, discounted properties effectively are finite-horizon. In contrast, ensuring safety of a reactive system or a certain probability to satisfy its mission goals requires an *unbounded* horizon and reduces to optimizing the reachability probabilities. Moreover, the safety-critical context requires *reliable* bounds on the probability, not an approximation with *unknown* precision.

In this paper, we provide the first provably correct anytime algorithm for (unbounded) reachability in uncountable MDP. As an *anytime* algorithm, it can at every step of the execution return correct lower and upper bounds on the true value. Moreover, these bounds gradually converge to the true value, allowing approximation up to an arbitrary precision. Since the problem is undecidable, the core of our contribution is identifying sufficient conditions on the uncountable MDP to allow for approximation.

Our *primary goal* is to provide *conditions as weak as possible*, thereby pushing towards the boundary of which systems can be analyzed provably correctly. To this end, we do not rely on any particular representation of the system. Nonetheless, for classical scenarios, and, in particular, for finite MDP, our conditions are mostly satisfied trivially.

Our *secondary goal* is to derive the respective algorithms as an *extension of value iteration* (VI) [29, 43], while *avoiding drawbacks of discretization*-based approaches. VI is a de facto standard method for numerical analysis of finite MDP, in particular with reachability objectives, regarded as practically efficient and allowing for heuristics avoiding the exploration of the complete state space, e.g. [9]. Interestingly, even for finite MDP, anytime VI algorithms with precision guarantees are quite recent [9, 19, 4, 44, 22]. Previous to that, the most used model checkers could return arbitrarily wrong results [19]. Providing VI with precision guarantees for general uncountable MDP is thus worthwhile on its own. Finally, while discretization is conceptually simple, we prefer to provide a solution that avoids the need to introduce arbitrary boundaries through gridding the whole state space and, moreover, instead utilizes information from one “cell” of the grid in other places, too.

To summarize, while algorithmic aspects form an important motivation, our primary contribution is theoretical: an explicit and complete set of generic assumptions allowing for guarantees, disregarding practical efficiency at this point. Consequently, while our approach lays foundations for further, more tailored approaches, it is not to be seen as a competitor to the existing practical, best-effort techniques, as these aim for a completely different goal.

**Our Contribution.** In this work, we provide the following:

**Section 3:** A set of assumptions that allow for computing converging *lower* bounds on the reachability probability in MDP with uncountable state and action spaces. We discuss in detail why they are weaker than usual, necessary, and applicable to typically considered systems. With these assumptions, we extend the standard (convergent but precision-ignorant) VI to this general setting.

**Section 4:** An additional set of assumptions that yield the first *anytime* algorithm, i.e. with provable bounds on the precision/error of the result, converging to 0. We combine the preceding algorithm with the technique of *bounded real-time dynamic programming (BRTDP)* [39] and provide also converging *upper* bounds on the reachability probability.

**Section 5:** A discussion of theoretical extensions and practical applications.

**Related work.** For detailed *theoretical* treatment of reachability and related problems on uncountable MDP, see e.g. [52, 11]. Reachability on uncountable MDP generalizes numerous problems known to be undecidable. For example, we can encode the halting problem of (probabilistic) Turing machines by encoding the tape content as real value. Similarly, almost-sure termination of probabilistic programs (undecidable [33]) is a special case of reachability on general uncountable MDP (see e.g. [16]). As precise reachability analysis is undecidable even for non-stochastic linear hybrid systems [26], many works turn their attention to more relaxed notions such as  $\delta$ -reachability, e.g. [48], and/or employ many assumptions.

In order to obtain precision bounds, we assume that the *value* function, mapping states to their reachability probability, is *Lipschitz continuous* (and that we know the Lipschitz constant). This is slightly *weaker* than the classical approach of assuming Lipschitz continuity of the *transition* function (and knowledge of the constant), e.g. [2, 49]. In particular, these assumptions (i) imply our assumption (as we show in [17, App. B.2.1]) and (ii) are used even in the simpler settings of finite-horizon and discounted reward scenarios [5, 2, 49, 51] or even more restricted settings to obtain practical efficiency, e.g. [35]. In contrast to our approach, they are not anytime algorithms and require treatment of the whole state space.

To provide context, we outline how continuity is used (explicitly or implicitly) in related work and mention their respective results. Firstly, [25, 47] assume Lipschitz continuity, *but not* explicit knowledge of the constant. In essence, these approaches solve the problem by successively increasing internal parameters. The parameters then eventually cross a bound implied by the Lipschitz constant, yielding an “eventual correctness”. In particular, they provide “convergence in the limit” or “probably approximately correct” results, but no bounds on the error or the convergence rate; these would depend on knowledge of the constant.

Secondly, [18, 40, 2, 49, 51] (and our work) assume Lipschitz continuity *and* knowledge of the constant. Relying on the constant being provided externally, these works derive guarantees. Previously, the guarantees given are weaker than our convergent anytime bounds: Either convergence in the limit [40] or a bound on a discretization error, relativized to sub-optimal strategies [18] or bounded horizon [2, 49, 51].

Several of the above mentioned works employ *discretization* [18, 2, 49, 51]. This method is quite general, but obtaining any bounds on the error requires continuity assumptions [1]. Further, there are works that use other assumptions: [23, 24] use *reinforcement learning* methods to tackle reachability and more general problems, without any continuity assumption. However, they do not provide any guarantees. See [53] for a detailed exposition of similar approaches. Assuming an abstraction is given, *abstraction and bisimulation* approaches, e.g. [21, 20], provide guarantees, but only on the lower bounds. With significant assumptions on the system’s structure, *symbolic* approaches [37, 54, 45, 14] may even obtain exact solutions.

## 2 Preliminaries

In this section, we recall basics of probabilistic systems and set up the notation. As usual,  $\mathbb{N}$  and  $\mathbb{R}$  refer to the (positive) natural numbers and real numbers, respectively. For a set  $S$ ,  $\mathbb{1}_S$  denotes its *characteristic function*, i.e.  $\mathbb{1}_S(x) = 1$  if  $x \in S$  and 0 otherwise. We write  $S^*$  and  $S^\omega$  to refer to the set of finite and infinite sequences comprising elements of  $S$ , respectively.

We assume familiarity with basic notions of measure theory, e.g. *measurable set* or *measurable function*, as well as probability theory, e.g. *probability spaces* and *measures* [8]. For a measure space  $X$  with sigma-algebra  $\Sigma_X$ ,  $\Pi(X)$  denotes the set of all probability measures on  $X$ . For a measure  $\mu \in \Pi(X)$ , we write  $\mu(Y) = \int \mathbb{1}_Y d\mu$  to denote the mass of a measurable set  $Y \in \Sigma_X$  (also called *event*). For two probability measures  $\mu$  and  $\nu$ , the *total variation distance* is defined as  $\delta_{TV}(\mu, \nu) := 2 \cdot \sup_{Y \in \Sigma_X} |\mu(Y) - \nu(Y)|$ . Some event happens *almost surely* (a.s.) w.r.t. some measure  $\mu$  if it happens with probability 1. We write  $\text{supp}(\mu)$  to denote the *support* of the probability measure  $\mu$ .

► **Remark 1.** It is surprisingly difficult to give a well-defined notion of support for measures in general. Intuitively,  $\text{supp}(\mu)$  describes the “smallest” set which  $\mu$  assigns a value of 1. However, this is not well-defined for general measures. We discuss these issues and a proper definition in [17, App. E]. Throughout this work, similar subtle issues related to measure theory arise. For the sake of readability, these are mostly delegated to footnotes or the appendix of the full version [17], and readers may safely skip over these points.

We work with *Markov decision processes* (MDP) [43], a widely used model to capture both non-determinism and probability. We consider uncountable state and action spaces.

► **Definition 2.** A (continuous-space, discrete-time) Markov decision process (MDP) is a tuple  $\mathcal{M} = (S, \text{Act}, \text{Av}, \Delta)$ , where  $S$  is a compact set of states (with topology  $\mathcal{T}_S$  and Borel  $\sigma$ -algebra  $\Sigma_S = \mathfrak{B}(\mathcal{T}_S)$ ),  $\text{Act}$  is a compact set of actions (with topology  $\mathcal{T}_{\text{Act}}$  and Borel  $\sigma$ -algebra  $\Sigma_{\text{Act}} = \mathfrak{B}(\mathcal{T}_{\text{Act}})$ ),  $\text{Av}: S \rightarrow \Sigma_{\text{Act}} \setminus \{\emptyset\}$  assigns to every state a non-empty, measurable, and compact set of available actions, and  $\Delta: S \times \text{Act} \rightarrow \Pi(S)$  is a transition function that for each state  $s$  and (available) action  $a \in \text{Av}(s)$  yields a probability measure over successor states (i.e. a Markov Kernel). An MDP is called *finite* if  $|S| < \infty$  and  $|\text{Act}| < \infty$ .

See [43, Sec. 2.3] and [6, Chp. 9] for a more detailed discussion on the technical considerations arising from uncountable state and action spaces. Note that we assume the set of available actions to be non-empty. This means that the system can never get “stuck” in a degenerate state without successors. *Markov chains* are a special case of MDP where  $|\text{Av}(s)| = 1$  for all  $s \in S$ , i.e. a completely probabilistic system without any non-determinism. Our presented methods thus are directly applicable to Markov chains as well.

Given a measure  $\mu \in \Pi(X)$  and a measurable function  $f: X \rightarrow \mathbb{R}$  mapping elements of a set  $X$  to real numbers, we write  $\mu\langle f \rangle := \int f(x) d\mu(x)$  to denote the integral of  $f$  with respect to  $\mu$ . For example,  $\Delta(s, a)\langle f \rangle$  denotes the expected value  $\mathbb{E}_{s' \sim \Delta(s, a)} f(s')$  of  $f: S \rightarrow \mathbb{R}$  over the successors of  $s$  under action  $a$ . Moreover, abusing notation, for some set of state  $S' \subseteq S$  and function  $\text{Av}': S' \rightarrow \text{Act}$ , we write  $S' \times \text{Av}' = \{(s, a) \mid s \in S', a \in \text{Av}'(s)\}$  to denote the set of state-action pairs with states from  $S'$  under  $\text{Av}'$ .

An *infinite path* in an MDP is some infinite sequence  $\rho = s_1 a_1 s_2 a_2 \dots \in (S \times \text{Av})^\omega$ , such that for every  $i \in \mathbb{N}$  we have  $s_{i+1} \in \text{supp}(\Delta(s_i, a_i))$ . A *finite path* (or *history*)  $\varrho = s_1 a_1 s_2 a_2 \dots s_n \in (S \times \text{Av})^* \times S$  is a non-empty, finite prefix of an infinite path of length  $|\varrho| = n$ , ending in state  $s_n$ , denoted by  $\text{last}(\varrho)$ . We use  $\rho(i)$  and  $\varrho(i)$  to refer to the  $i$ -th state in an (in)finite path. We refer to the set of finite (infinite) paths of an MDP  $\mathcal{M}$  by  $\text{FPaths}_{\mathcal{M}}$  ( $\text{Paths}_{\mathcal{M}}$ ). Analogously, we write  $\text{FPaths}_{\mathcal{M}, s}$  ( $\text{Paths}_{\mathcal{M}, s}$ ) for all (in)finite paths starting in  $s$ .

In order to obtain a probability measure, we first need to eliminate the non-determinism. This is done by a so-called *strategy* (also called *policy*, *controller*, or *scheduler*). A strategy on an MDP  $\mathcal{M} = (S, Act, Av, \Delta)$  is a function  $\pi: \text{FPaths}_{\mathcal{M}} \rightarrow \Pi(Act)$ , s.t.  $\text{supp}(\pi(\varrho)) \subseteq Av(\text{last}(\varrho))$ . The set of all strategies is denoted by  $\Pi_{\mathcal{M}}$ . Intuitively, a strategy is a “recipe” describing which step to take in the current state, given the evolution of the system so far.

Given an MDP  $\mathcal{M}$ , a strategy  $\pi \in \Pi_{\mathcal{M}}$ , and an initial state  $s_0$ , we obtain a measure on the set of infinite paths  $\text{Paths}_{\mathcal{M}}$ , which we denote as  $\text{Pr}_{\mathcal{M}, s_0}^{\pi}$ . See [43, Sec. 2] for further details. Thus, given a measurable set  $A \subseteq \text{Paths}_{\mathcal{M}}$ , we can define its maximal probability starting from state  $s_0$  under any strategy by  $\text{Pr}_{\mathcal{M}, s_0}^{\text{sup}}[A] := \sup_{\pi \in \Pi_{\mathcal{M}}} \text{Pr}_{\mathcal{M}, s_0}^{\pi}[A]$ . Depending on the structure of  $A$  it may be the case that no optimal strategy exists and we have to resort to the supremum instead of the maximum. This may already arise for finite MDP, see [12].

For an MDP  $\mathcal{M} = (S, Act, Av, \Delta)$  and a set of *target states*  $T \subseteq S$ , (*unbounded*) *reachability* refers to the set  $\diamond T = \{\rho \in \text{Paths}_{\mathcal{M}} \mid \exists i \in \mathbb{N}. \rho(i) \in T\}$ , i.e. all paths which eventually reach  $T$ . The set  $\diamond T$  is measurable if  $T$  is measurable [51, Sec. 3.1], [52, Sec. 2].

Now, it is straightforward to define the *maximal reachability problem* of a given set of states. Given an MDP  $\mathcal{M}$ , target set  $T$ , and state  $s_0$ , we are interested in computing the maximal probability of eventually reaching  $T$ , starting in state  $s_0$ . Formally, we want to compute the *value* of the state  $s_0$ , defined as  $\mathcal{V}(s_0) := \text{Pr}_{\mathcal{M}, s_0}^{\text{sup}}[\diamond T] = \sup_{\pi \in \Pi_{\mathcal{M}}} \text{Pr}_{\mathcal{M}, s_0}^{\pi}[\diamond T]$ . This state value function satisfies a straightforward fixed point equation, namely

$$\mathcal{V}(s) = 1 \quad \text{if } s \in T \quad \mathcal{V}(s) = \sup_{a \in Av(s)} \Delta(s, a) \langle \mathcal{V} \rangle \quad \text{otherwise.} \quad (1)$$

Moreover,  $\mathcal{V}$  is the *smallest* fixed point of this equation [6, Prop. 9.8, 9.10], [52, Thm. 3]. In our approach, we also deal with values of state-action pairs  $(s, a) \in S \times Av$ , where  $\mathcal{V}(s, a) := \Delta(s, a) \langle \mathcal{V} \rangle$ . Intuitively, this represents the value achieved by choosing action  $a$  in state  $s$  and then moving optimally. Clearly, we have that  $\mathcal{V}(s) = \sup_{a \in Av(s)} \mathcal{V}(s, a)$ . See [15, Sec. 4] for a discussion of reachability on finite MDP and [52] for the general case.

In this work, we are interested in *approximate solutions* due to the following two reasons. Firstly, obtaining precise solutions for MDP is difficult already under strict assumptions and undecidable in our general setting.<sup>(1)</sup> We thus resort to approximation, allowing for much lighter assumptions. Secondly, by considering approximation we are able to apply many different optimization techniques, potentially leading to algorithms which are able to handle real-world systems, which are out of reach for precise algorithms even for finite MDP [9].

We are interested in two types of approximations. Firstly, we consider approximating the value function in the limit, without knowledge about how close we are to the true value. This is captured by a semi-decision procedure for queries of the form  $\text{Pr}_{\mathcal{M}, s}^{\text{sup}}[\diamond T] > \xi$  for a threshold  $\xi \in [0, 1]$ . We call this problem **ApproxLower**. Secondly, we consider the variant where we are given a precision requirement  $\varepsilon > 0$  and obtain  $\varepsilon$ -optimal values  $(l, u)$ , i.e. values with  $\mathcal{V}(s_0) \in [l, u]$  and  $0 \leq u - l < \varepsilon$ . We refer to this variant as **ApproxBounds**.

### 3 Converging Lower Bounds

In this section, we present the first set of assumptions, enabling us to compute *converging lower bounds* on the true value, solving the **ApproxLower** problem. In Section 3.1, we discuss each assumption in detail and argue on an intuitive level why it is necessary by means of

<sup>(1)</sup>For example, one can encode the tape of a Turing machine into the binary representation of a real number and reduce the halting problem to a reachability query.

counterexamples. With the assumptions in place, in Section 3.2 we then present our first algorithm, also introducing several ideas we employ again in the following section.

Our assumptions and algorithms are motivated by *value iteration* (VI) [29], which we briefly outline. In a nutshell, VI boils down to repeatedly applying an iteration operator to a value vector  $v_n$ . For example, the canonical value iteration for reachability on finite MDP starts with  $v_0(s) = 1$  for all  $s \in T$  and 0 otherwise and then iterates

$$v_{n+1}(s) = \max_{a \in \text{Act}(s)} \sum_{s' \in S} \Delta(s, a, s') \cdot v_n(s') \quad (2)$$

for all  $s \notin T$ . The vector  $v_n$  converges monotonically from below to the true value for all states. We mention two important points. Firstly, the iteration can be applied “asynchronously”. Instead of updating all states in every iteration, we can pick a single state and only update its value. The values  $v_n$  still converge to the correct value as long as all states are updated infinitely often. Secondly, instead of storing a value per state, we can store a value for each state-action pair and obtain the state value as the maximum of these values. Both points are a technical detail for finite MDP, however they play an essential role in our uncountable variant. See [17, App. A.1] for more details on VI for finite MDP.

In the uncountable variant of Equation (2),  $v$  is a function,  $\text{Act}(s)$  is potentially uncountable, and the sum is replaced by integration. As in this setting the problem is undecidable, naturally we have to employ some assumptions. Our goal is to sufficiently imitate the essence of Equation (2), obtaining convergence without being overly restrictive. In particular, we want to (i) represent (an approximation of)  $v_n$  using finite memory, (ii) safely approximate the maximum and integration, and (iii) select appropriate points to update  $v_n$ .

### 3.1 Assumptions

Before discussing each assumption in detail, we first put them into context. As we argue in the following, most of our assumptions typically hold implicitly. Still, by stating even basic computability assumptions in a form as weak as possible, we avoid “hidden” assumptions, e.g. by assuming that the state space is a subset of  $\mathbb{R}^d$ . Two of our assumptions are more restrictive, namely **Assumption C: Value Lipschitz Continuity** (Section 3.1.3) and, introduced later, **Assumption D: Absorption** (Section 4.1.2). However, they are also often used in related works, as we detail in the respective sections. Moreover, in light of previous results, the necessity of restrictive assumptions is to be expected: Computing bounds is hard or even undecidable already for very restricted classes. Aside from the discussion in the introduction, we additionally mention two further cases. In the setting of probabilistic programs (which are a very special case of uncountable MDP), deciding almost sure termination for a fixed initial state (which is a severely restricted subclass of reachability on uncountable MDP without non-determinism) is an actively researched topic with recent advances, see e.g. [30, 31], and shown to be  $\Pi_2^0$ -complete [33], i.e. highly undecidable. In [27] and the references therein, the authors present (un-)decidability results for hybrid automata, which are a special case of uncountable MDP *without any stochastic dynamics* (flow transitions can be modelled as actions indicating the delay). As such, it is to be expected that the general class of models we consider has to be pruned very strictly in order to hope for any decidability results.

► **Remark 3.** As already mentioned, we want to provide assumptions which are as general as possible. Importantly, we avoid (unnecessarily) assuming any particular representation of the system. Our motivation is to ultimately identify the boundary of what is necessary to derive guarantees. While our assumptions are motivated by VI and built around Equation (2), we

note that being able to represent the state values and evaluate (some aspect of) the transition dynamics intuitively are a necessity for *any* method dealing with such systems. We do not claim that our framework of assumptions is the only way to approach the problem, instead we provide arguments why it is a sensible way to do so.

### 3.1.1 A: Basic Assumptions (Asm. A1-A4)

We first present a set of basic computability assumptions (**A1-A4**). These are essential, since for uncountable systems even the simplest computations are intractable without any assumptions. More specifically, such systems cannot be given explicitly (due to their infinite size), but instead have to be described symbolically by, e.g., differential equations. Thus, we necessarily require some notion of computability and structural properties for each part of this symbolic description. And indeed, each assumption essentially corresponds to one part of the MDP description (**Metric Space** to  $S \times Act$ , **Maximum Approximation** to  $Av$ , **Transition Approximation** to  $\Delta$ , and **Target Computability** to  $T$ ). They are weak and hold on practically all commonly considered systems (see [17, App. B.1]). In particular, finite MDP and discrete components are trivially subsumed by considering the discrete metric.

**A1: Metric Space**  $S$  and  $Act$  are metric spaces with (computable) metrics  $d_S$  and  $d_{Act}$ , respectively, and  $d_\times$  is a compatible<sup>(2)</sup> metric on the space of state-action pairs  $S \times Av$ ,

**A2: Maximum Approximation** For each state  $s$  and computable Lipschitz  $f : Av(s) \rightarrow [0, 1]$ , the value  $\max_{a \in Av(s)} f(a)$  can be under-approximated to arbitrary precision.

**A3: Transition Approximation** For each state-action pair  $(s, a)$  and Lipschitz  $g : S \rightarrow [0, 1]$  which can be under-approximated to arbitrary precision, the successor expectation  $\Delta(s, a)\langle g \rangle$  can be under-approximated to arbitrary precision.

**A4: Target Computability** The target set  $T$  is decidable, i.e. we are given a computable predicate which, given a state  $s$ , decides whether  $s \in T$ .

We denote the approximations for **A2** and **A3** by  $\text{APPROX}_{\leq}$ , i.e. given a pair  $(s, a)$  and functions  $f, g$  as in the assumptions, we write (abusing notation)  $\text{APPROX}_{\leq}(\max_{a \in Av(s)} f(a), \varepsilon)$  and  $\text{APPROX}_{\leq}(\Delta(s, a)\langle g \rangle, \varepsilon)$  for approximation of the respective values up to precision  $\varepsilon$ , i.e.  $0 \leq \max_{a \in Av(s)} f(a) - \text{APPROX}_{\leq}(\max_{a \in Av(s)} f(a), \varepsilon) \leq \varepsilon$  and analogous for  $\Delta(s, a)\langle g \rangle$ . Note that **A2** and **A3** are satisfied if we can sample densely in  $Av(s)$  and approximate  $\Delta(s, a)$ .

### 3.1.2 B: Sampling (Asm. B.VI)

As there are uncountably many states, we are unable to explicitly update all of them at once and instead update values asynchronously. Moreover, as there may also be uncountably many actions, we instead store and update the values of state-action pairs. Together, we need to pick state-action pairs to update. We delegate this choice to a selection mechanism `GETPAIR`, an oracle for state-action pairs. We allow for `GETPAIR` to be “stateful”, i.e. the sampled state-action pair may depend on previously returned pairs. This is required in, for example, round-robin or simulation-based approaches. We only require a basic notion of fairness in order to guarantee that we do not miss out on any information. Note the additional identifier **.VI** (*value iteration*) on the assumption name; later on, a similar, but weaker variant (**B.BRTDP**) is introduced.

<sup>(2)</sup>For two pairs  $(s, a)$  and  $(s, a')$  we have that  $k \cdot d_{Act}(a, a') \leq d_\times((s, a), (s, a')) \leq K \cdot d_{Act}(a, a')$  for some constants  $k, K \geq 0$ , analogous for  $d_S$ , achieved by, e.g.  $d_\times((s, a), (s', a')) := d_S(s, s') + d_{Act}(a, a')$ .

**B.VI: State-Action Sampling** Let  $S^\diamond = \{\text{last}(\varrho) \mid \varrho \in \text{FPaths}_{\mathcal{M},s}\}$  the set of all reachable states. Then, for any  $\varepsilon > 0$ ,  $s \in S^\diamond$ , and  $a \in \text{Av}(s)$  we have that GETPAIR eventually yields a pair  $(s', a')$  with  $d_\times((s, a), (s', a')) < \varepsilon$  and  $\delta_{TV}(\Delta(s, a), \Delta(s', a')) < \varepsilon$  a.s.<sup>(3)</sup>

Essentially, this means that GETPAIR provides a way to “exhaustively” generate all behaviours of the system up to a precision of  $\varepsilon$ . This fairness assumption is easily satisfied under usual conditions. For example, if  $S \times \text{Av}$  is a bounded subset of  $\mathbb{R}^d$ , we can randomly sample points in that space or consider increasingly dense grids. Alternatively, if we can sample from the set of actions and from the distributions of  $\Delta$ , GETPAIR can be implemented by sampling paths of random length, following random actions. Note that we can view the procedure as a “template”: Instead of requiring a concrete method to acquire pairs to update, we leave this open for generality; we discuss implications of this in Sections 5.1 and 5.3.

The requirement on total variation may seem unnecessary, especially given that we will also assume continuity. However, otherwise we could, for example, miss out on solitary actions which are the “witnesses” for a state’s value: suppose that  $\text{Av}(s) = [0, 1]$  and  $\Delta(s, 0)$  moves to the goal, while  $\Delta(s, a)$  just loops back to  $s$ . Only selecting actions close to  $a = 0$  w.r.t. the product metric is not sufficient to observe that we can move to the goal. Note that this would not be necessary if we assumed continuity of the transition function – selecting “nearby” actions then also yields “similar” behaviour.

### 3.1.3 C: Lipschitz Continuity

Finally, we present our already advertised continuity assumption. For simplicity, we give it in its strict form and discuss relaxations later in Section 5.2. Intuitively, Lipschitz continuity allows us to extrapolate the behaviour of the system from a single state to its surroundings.

**C: Value Lipschitz Continuity** The value functions  $\mathcal{V}(s)$  and  $\mathcal{V}(s, a)$  are Lipschitz continuous with *known* constants  $C_S$  and  $C_\times$ , i.e. for all  $s, s' \in S$  and  $a \in \text{Av}(s), a' \in \text{Av}(s')$  we have


$$|\mathcal{V}(s) - \mathcal{V}(s')| \leq C_S \cdot d_S(s, s') \quad |\mathcal{V}(s, a) - \mathcal{V}(s', a')| \leq C_\times \cdot d_\times((s, a), (s', a'))$$

This requirement may seem quite restrictive at first glance. Indeed, it is the only one in this section to not usually hold on “standard” systems. However, in order to obtain any kind of (provably correct) bounds, some notion of continuity is elementary, since otherwise we cannot safely extrapolate from finitely many observations to an uncountable set. The immediately arising questions are (i) why *Lipschitz* continuity is necessary compared to, e.g., regular or uniform continuity, and (ii) why *knowledge of the Lipschitz constant* is required. For the first point, note that we want to be able to extrapolate from values assigned to a single state to its immediate surroundings. While continuity means that the values in the surroundings do not “jump”, it does not give us any way of bounding the rate of change, and this rate may grow arbitrarily (for example, consider the continuous but not Lipschitz function  $\sin(\frac{1}{x})$  for  $x > 0$ ). So, also relating to the second point, without knowledge of the Lipschitz constant, regular continuity and Lipschitz continuity are (mostly) equivalent from a computational perspective: The function does not have discontinuities, but we cannot safely estimate the rate of change in general. To illustrate this point further, we give an intuitive example.

---

<sup>(3)</sup>Technically, it is sufficient to satisfy this property on any subset of  $S^\diamond$  which only differs from it up to measure 0. More precisely, we only require that this assumption holds for  $S^\diamond = \text{supp}(\text{Pr}_{\mathcal{M},s}^{\text{sup}})$ , i.e. the set of all reachable paths with non-zero measure. We omit this rather technical notion and the discussion it entails in order to avoid distracting from the central results of this work.





■ **Figure 1** The value function of Example 4, showing that knowledge of the constant is important.

► **Example 4.** We construct an MDP with a periodic, Lipschitz continuous value function, as illustrated in Figure 1 and formally defined below. Intuitively, for a given period width  $w$  (e.g. 0.25) and a periodic function  $f$  (e.g. a triangle function), a state  $s$  between 0 and  $w$  moves to a target or sink with probability  $f(s)$ . All larger states  $s \geq w$  transition to  $s - w$  with probability 1. The value function thus is periodic and Lipschitz continuous, see Figure 1 for a possible value function and [17, App. B.2.3] for a formal definition.

For a finite number of samples, we can choose  $f$  and  $w$  such that all samples achieve a value of 1. Nevertheless, we cannot conclude anything about states we have not sampled yet: *Without knowledge of the constant, we cannot extrapolate from samples.*

We note the underlying connection to the Nyquist-Shannon sampling theorem [46, Thm. 1]. Intuitively, the theorem states that, for a function that contains no frequencies higher than  $W$ , it is completely determined by giving its ordinates at a series of points spaced  $0.5 \cdot W$  apart. If we know the Lipschitz constant, this gives us a way of bounding the “frequency” of the value function, and thus allows us to determine it by sampling a finite number of points. On the other hand, without the Lipschitz constant, we do not know the frequency and cannot judge whether we are “undersampling”.

Since we do not assume any particular representation of the transition system, we cannot derive such constants in general. Instead, these would need to be obtained by, e.g., domain knowledge, or tailored algorithms. As in previous approaches [18, 40, 2, 49, 51], we thus resort to assuming that we are given this constant, offloading this (highly non-trivial) step. Recall that Lipschitz continuity of the transition function implies Lipschitz continuity of the value function (see [17, App. B.2.1]), but can potentially be checked more easily.


### 3.2 Assumptions Applied: Value Iteration Algorithm

Before we present our new algorithm, we explain how our assumptions allow us to lift VI to the uncountable domain. Contrary to the finite state setting, we are unable to store precise values for each state explicitly, since there are uncountably many states. Hence, the algorithm exploits the Lipschitz-continuity of the value function as follows. Assume that we know that the value of a state  $s$  is bounded from below by a value  $l$ , i.e.  $\mathcal{V}(s) \geq l$ . Then, by Lipschitz-continuity of  $\mathcal{V}$ , we know that the value of a state  $s'$  is bounded by  $l - d_S(s, s') \cdot C_S$ . More generally, if we are given a finite set of states  $\text{Sampled}$  with correct lower bounds  $\widehat{\mathcal{L}}: \text{Sampled} \rightarrow [0, 1]$ , we can safely extend these values to the whole state space by

$$\mathcal{L}(s) := \max_{s' \in \text{Sampled}} \left( \widehat{\mathcal{L}}(s') - C_S \cdot d_S(s, s') \right).$$

Since  $\mathcal{V}(s) \geq \widehat{\mathcal{L}}(s)$  for all  $s \in \text{Sampled}$ , we have  $\mathcal{V}(s) \geq \mathcal{L}(s)$  for all  $s \in S$ , i.e.  $\mathcal{L}(\cdot)$  is a valid lower bound. We thus obtain a lower bound for all of the uncountably many states, described symbolically as a combination of finitely many samples. See Figure 2 for an illustration.

This is sufficient to deal with Markov chains, but for MDPs we additionally need to take care of the (potentially uncountably many) actions. Recall that value iteration updates state values with the maximum over available actions,  $v_{n+1}(s) = \max_{a \in Av(s)} \Delta(s, a)(v_n)$ .



■ **Figure 2** Example of the function extension on the set  $[0, 2]$  with a Lipschitz constant of  $C_S = 1$ . Dots represent stored values in  $\widehat{L}$ , while the solid line represents the extrapolated function  $L$ . Note that it is possible to have  $\widehat{L}(s) < L(s)$ , as seen in the graph.

■ **Algorithm 1** The Value Iteration (VI) Algorithm for MDPs with general state- and action-spaces.

**Input:** `ApproxLower` query with threshold  $\xi$ , satisfying **A1–A4**, **B.VI** and **C**.

**Output:** `yes`, if  $\mathcal{V}(s_0) > \xi$ .

---

```

1: Sampled  $\leftarrow \emptyset, t \leftarrow 1$  ▷ Initialize
2: while APPROX≤(L(s0), PRECISION(t)) ≤ ξ do
3:    $(s, a) \leftarrow \text{GETPAIR}$  ▷ Sample state-action pair
4:   if  $s \in T$  then  $\widehat{L}(s, \cdot) \leftarrow 1$  ▷ Handle target states
5:   else  $\widehat{L}(s, a) \leftarrow \text{APPROX}_{\leq}(\Delta(s, a)\langle L \rangle, \text{PRECISION}(t))$  ▷ Update  $\widehat{L}$ 
6:   Sampled  $\leftarrow \text{Sampled} \cup \{(s, a)\}, t \leftarrow t + 1$ 
7: return yes
    
```

---

This is straightforward to compute when there are only finitely many actions, but in the uncountable case obtaining  $L(s) = \sup_{a \in Av(s)} L(s, a)$  is much more involved. We apply the idea of Lipschitz continuity again, storing values for a set `Sampled` of state-action pairs instead of only states. We bound the value of every state-action pair by

$$L(s, a) := \max_{(s', a') \in \text{Sampled}} \left( \widehat{L}(s', a') - d_x((s, a), (s', a')) \cdot C_x \right) \quad (3)$$

Observe that  $L(s, a)$  is computable and Lipschitz-continuous as well, so by **Maximum Approximation** we can approximate the bound of any state, i.e.  $L(s) = \max_{a \in Av(s)} L(s, a)$ , based on such a finite set of values assigned to state-action pairs. (Recall that  $Av(s)$  is compact and  $L(s, a)$  continuous, hence the maximum is attained.) Consequently, we can also under-approximate  $\Delta(s, a)\langle L \rangle$  by **Transition Approximation**. To avoid clutter, we omit the following two special cases in the definition of  $L(s, a)$ : Firstly, if `Sampled` =  $\emptyset$ , we naturally set  $L(s, a) = 0$ . Secondly, if all pairs  $(s', a')$  are too far away for a sensible estimate, i.e. if Equation (3) was yielding  $L(s, a) < 0$ , we also set  $L(s, a)$  to 0.

We present VI for MDPs with general state- and action-spaces in Algorithm 1. It depends on `PRECISION(t)`, a sequence of precisions converging to zero in the limit, e.g. `PRECISION(t) =  $\frac{1}{t}$` . The algorithm executes the main loop until the current approximation of the lower bound of the initial state  $L(s_0) = \max_{a \in Av(s_0)} L(s_0, a)$  exceeds the given threshold  $\xi$ . Inside the loop, the algorithm updates state-action pairs yielded by `GETPAIR`. For target states, the lower bound is set to 1. Otherwise, we set the bound of the selected pair to an approximation of the expected value of  $L$  under the corresponding transition. Here is the crucial difference to VI in the finite setting: Instead of using Equation (2), we have to use Equation (3) and `APPROX≤`, the approximations that exist by assumption, see Section 3.1.1. Since `PRECISION(t)` converges to zero, the approximations eventually get arbitrarily fine. The procedure `PRECISION(t)` may be adapted heuristically in order to speed up computation. For example, it may be beneficial to only approximate up to 0.01 precision at first to quickly get a rough overview. We show that Algorithm 1 is correct, i.e. the stored values (i) are lower bounds and (ii) converge to the true values in [17, App. E.1]. Here, we only provide a sketch, illustrating the main steps.

► **Theorem 5.** *Algorithm 1 is correct under Assumptions A1–A4, B.VI, and C, i.e. it outputs yes iff  $\mathcal{V}(s) > \xi$ .*

**Proof sketch.** First, we show that  $L_t(s) \leq L_{t+1}(s) \leq \mathcal{V}(s)$  by simple induction on the step. Initially, we have  $L_1(s) = 0$ , obviously satisfying the condition. The updates in Lines 4 and 5 both keep correctness, i.e.  $L_{t+1}(s) \leq \mathcal{V}(s)$ , proving the claim.

Since  $L_t$  is monotone as argued above, its limit for  $t \rightarrow \infty$  is well defined, denoted by  $L_\infty$ . By **State-Action Sampling**, the set of accumulation points of  $s_t$  contains all reachable states  $S^\diamond$ . We then prove that  $L_\infty$  satisfies the fixed point equation Equation (1). For this, we use the second part of the assumption on GETPAIR, namely that for every  $(s, a) \in S^\diamond \times Av$  we get a converging subsequence  $(s_{t_k}, a_{t_k})$  where additionally  $\Delta(s_{t_k}, a_{t_k})$  converges to  $\Delta(s, a)$  in total variation. Intuitively, since infinitely many updates occur infinitely close to  $(s, a)$ , its limit lower bound  $L_\infty(s, a)$  agrees with the limit of the updates values  $\lim_{k \rightarrow \infty} \Delta(s_{t_k}, a_{t_k}) \langle L_{t_k} \rangle$ . Since  $L_\infty$  satisfies the fixed point equation and is less or equal to the value function  $\mathcal{V}$ , we get the result, since  $\mathcal{V}$  is the smallest fixed point. ◀

## 4 Converging Upper Bounds

In this section, we present the second set of assumptions, allowing us to additionally compute converging *upper* bounds. With both lower and upper bounds, we can quantify the progress of the algorithm and, in particular, terminate the computation once the bounds are sufficiently close. Therefore, instead of only providing a semi-decision procedure for reachability, this algorithm is able to determine the maximal reachability probability up to a given precision. Thus, we obtain the first algorithm able to handle such general systems with guarantees on its result. We again present our assumptions together with a discussion of their necessity (Section 4.1), and then introduce the subsequent algorithm and prove its correctness (Section 4.2). As expected, obtaining this additional information also requires additional assumptions. On the other hand, quite surprisingly, we can use the additional information of upper bounds to actually *speed up* the computation, as discussed in Section 5.3.

As before, our approach is inspired by algorithms for finite MDP, in this case by *Bounded Real-Time Dynamic Programming* (BRTDP) [39, 9]. BRTDP uses the same update equations as VI, but iterates both lower and upper bounds. A major contribution of [9] was to solve the long standing open problem of how to deal with *end components*. These parts of the state space prevent convergence of the upper bounds by introducing additional fixpoints of Equation (1). We direct the interested reader to [17, App. A.2] for further details on BRTDP and insights on the issue of end components. In the uncountable setting, these issues arise as well alongside several other, related problems, which we discuss in Section 4.1.2.

### 4.1 Assumptions

The basic assumptions A1–A4 as well as **Lipschitz continuity** (Assumption C) remain unchanged. For **Maximum Approximation** (A2) and **Transition Approximation** (A3), we additionally require that we are able to *over*-approximate the respective results. The respective assumptions are denoted by A5 and A6, respectively, and both over-approximations by APPROX<sub>≥</sub>. Further, we only require a weakened variant of **State-Action Sampling**, now called Assumption **B.BRTDP** instead of Assumption **B.VI**. Finally, there is the new Assumption **D** called **Absorption**, addressing the aforementioned issue of end components.

### 4.1.1 B: Weaker Sampling (Asm. B.BRTDP)

We again assume a GETPAIR oracle, but, perhaps surprisingly, with *weaker* assumptions. Instead of requiring it to return “all” actions, we only require it to yield “optimal” actions, respective to a given state-action value function. We first introduce some notation. Intuitively, we want GETPAIR to yield actions which are optimal with respect to the *upper bounds* computed by the algorithm. However, these upper bounds potentially change after each update. Thus, assume that  $f_n: S \times Av \rightarrow [0, 1]$  is an arbitrary sequence of computable, Lipschitz continuous, (point-wise) monotone decreasing functions, assigning a value to each state-action pair, and set  $\mathcal{F} = (f_1, f_2, \dots)$ . For each state  $s \in S$ , set

$$Av_{\mathcal{F}}(s) := \{a \in Av(s) \mid \forall \varepsilon > 0. \forall N \in \mathbb{N}. \exists n > N. \max_{a' \in Av(s)} f_n(s, a') - f_n(s, a) < \varepsilon\},$$

i.e. actions that infinitely often achieve values arbitrarily close to the optimum of  $f_n$ . Let  $S_{\mathcal{F}}^{\diamond} = \{last(\varrho) \mid \varrho \in \text{FPaths}_{\mathcal{M}, s_0} \cap (S \times Av_{\mathcal{F}})^* \times S\}$  be the set of all states reachable using these optimal actions.<sup>(4)</sup> Essentially, we require that GETPAIR samples densely in  $S_{\mathcal{F}}^{\diamond} \times Av_{\mathcal{F}}$ .

**B.BRTDP: State-Action Sampling** For any  $\varepsilon > 0$ ,  $\mathcal{F}$  as above,  $s \in S_{\mathcal{F}}^{\diamond}$ , and  $a \in Av_{\mathcal{F}}(s)$  we have that GETPAIR a.s. eventually yields a pair  $(s', a')$  with  $d_{\times}((s, a), (s', a')) < \varepsilon$  and  $\delta_{TV}(\Delta(s, a), \Delta(s', a')) < \varepsilon$ .

While this new variant may seem much more involved, it is *weaker* than its previous variant, since  $Av_{\mathcal{F}}(s) \subseteq Av(s)$  for each  $s \in S$  and thus also  $S_{\mathcal{F}}^{\diamond} \subseteq S^{\diamond}$ . As such, it also allows for more practical optimizations, which we briefly discuss in Section 5.3.

### 4.1.2 D: Absorption

We present our most specific assumption. While it is *not needed* for correctness, we require it for convergence of the upper bounds to the value and thus for termination of the algorithm.

**D: Absorption** There exists a *known and decidable* set  $R$  (called *sink*) such that  $\mathcal{V}(s) = 0$  for all  $s \in R$ . Moreover, for any  $s \in S$  and strategy  $\pi$  we have  $\Pr_{\mathcal{M}, s}^{\pi}[\diamond(T \cup R)] = 1$ .

Intuitively, the assumption requires that for all strategies, the system will eventually reach a target or a goal state; in other words: It is not possible to avoid both target and sink infinitely long. Variants of this assumption are used in numerous settings: On MDP, it is similar to the contraction assumption, e.g. [6, Chp. 4]; in stochastic game theory (a two-player extension of MDP) it is called *stopping*, e.g. [13]; and, using terms from the theory of the stochastic shortest path problem, we require all strategies to be *proper*, see e.g. [7].

This assumption already is important in the finite setting: There, **Absorption** is equivalent to the absence of *end components*, which introduce multiple solutions of Equation (1). Then, a VI algorithm computing upper bounds can be “stuck” at a greater fixpoint than the value and thus does not converge [9, 19]. *Any* procedure using value iteration thus either needs to exclude such cases or detect and treat them. Aside from end components, which are the only issue in the finite setting, uncountable systems may feature other complex behaviour, such as Zeno-like approaching the target closer and closer without reaching it.

Unfortunately, even just detecting these problems already is difficult. For the mentioned, restricted setting of probabilistic programs, almost sure termination is  $\Pi_0^2$ -complete [33]. Yet, universal termination with goal set  $T \cup R$  is exactly what we require for **Absorption**. So, already on a restricted setting (together with a *given* guess for  $R$ ), we cannot decide whether the assumption holds, let alone treat the underlying problems. Thus, we decide to exclude this issue and delegate treatment to specialized approaches.

<sup>(4)</sup>As in Section 3, we simplify the definition of  $S_{\mathcal{F}}^{\diamond}$  slightly in order to avoid technical details.

■ **Algorithm 2** The BRTDP algorithm for MDPs with general state- and action-spaces.

**Input:** ApproxBounds query with precision  $\varepsilon$ , satisfying **A1–A6**, **B.BRTDP**, **C** and **D**.

**Output:**  $\varepsilon$ -optimal values  $(l, u)$ .

---

```

1: Sampled  $\leftarrow \emptyset, t \leftarrow 1$  ▷ Initialize
2: while  $\text{APPROX}_{\geq}(U(s_0), \text{PRECISION}(t)) - \text{APPROX}_{\leq}(L(s_0), \text{PRECISION}(t)) \geq \varepsilon$  do
3:    $s, a \leftarrow \text{GETPAIR}$  ▷ Sample stat-action pair
4:   if  $s \in T$  then  $\hat{L}(s, \cdot) \leftarrow 1$  ▷ Handle special cases
5:   else if  $s \in R$  then  $\hat{U}(s, \cdot) \leftarrow 0$ 
6:   else ▷ Update upper and lower bounds
7:      $\hat{U}(s, a) \leftarrow \text{APPROX}_{\geq}(\Delta(s, a)\langle U \rangle, \text{PRECISION}(t))$ 
8:      $\hat{L}(s, a) \leftarrow \text{APPROX}_{\leq}(\Delta(s, a)\langle L \rangle, \text{PRECISION}(t))$ 
9:   Sampled  $\leftarrow \text{Sampled} \cup \{(s, a)\}, t \leftarrow t + 1$ 
10: return  $(L(s_0), U(s_0))$ 

```

---

In summary, while this assumption is indeed restrictive, it is the key point that allows us to obtain convergent upper bounds and thus an anytime algorithm. As argued above, an assumption of this kind seems to be *necessary* to obtain such an algorithm in this generality.

► **Remark 6.** These problems do not occur when considering *finite horizon* or *discounted* properties, which are frequently used in practice. For details on treating finite horizon objectives, see [17, App. C.1]. Discounted reachability with a factor of  $\gamma < 1$  is equivalent to normal reachability where at each step the system moves into a sink state with probability  $(1 - \gamma)$ . **Absorption** is trivially satisfied and our methods are directly applicable.

## 4.2 Assumptions Applied: The Convergent Anytime Algorithm

With our assumptions in place, we are ready to present our adaptation of BRTDP to the uncountable setting. Compared to VI, we now also store upper bounds, again using Lipschitz-continuity to extrapolate the stored values. In particular, together with the definitions of Equation (3) we additionally set

$$U(s, a) = \min_{(s', a') \in \text{Sampled}} \left( \hat{U}(s', a') + d_{\times}((s, a), (s', a')) \cdot C_{\times} \right).$$

We also set  $U(s, a) = 1$  if either  $\text{Sampled} = \emptyset$  or the above equation would yield  $U(s, a) > 1$ .

We present BRTDP in Algorithm 2. It is structurally similar to BRTDP in the finite setting (see [17, App. A.2]). The major difference is given by the storage tables  $\hat{U}$  and  $\hat{L}$  used to compute the current bounds  $U$  and  $L$ , again exploiting Lipschitz continuity. As before, the central idea is to repeatedly update state-action pairs given GETPAIR. If GETPAIR yields a state of the terminal sets  $T$  and  $R$ , we update the stored values directly. Otherwise, we back-propagate the value of the selected pair by computing the expected value under this transition. Moreover, we again require that  $\text{PRECISION}(t)$  converges to zero. Note that the algorithm can easily be supplied with a-priori knowledge by initializing the upper and lower bounds to non-trivial values. Moreover, in contrast to VI, this algorithm is an *anytime* algorithm, i.e. it can at any time provide an approximate solution together with its precision.

Despite the algorithm being structurally similar to the finite variant of [9], the proof of correctness unsurprisingly is more intricate due to the uncountable sets. We again provide both a simplified proof sketch here and the full technical proof in [17, App. E.2].

► **Theorem 7.** *Algorithm 2 is correct under Assumptions A1–A6, B.BRTDP, C and D, and terminates with probability 1.*

**Proof sketch.** We again obtain monotonicity of the bounds, i.e.  $L_t(s, a) \leq L_{t+1}(s, a) \leq V(s, a) \leq U_{t+1}(s, a) \leq U_t(s, a)$  by induction on  $t$ , using completely analogous arguments.

By monotonicity, we also obtain well defined limits  $U_\infty$  and  $L_\infty$ . Further, we define the difference function  $\text{Diff}_t(s, a) = U_t(s, a) - L_t(s, a)$  together with its state based counterpart  $\text{Diff}_t(s)$  and its limit  $\text{Diff}_\infty(s)$ . We show that  $\text{Diff}_\infty(s_0) = 0$ , proving convergence. To this end, similar to the previous proof, we prove that  $\text{Diff}_\infty$  satisfies a fixed point equation on  $S_+^\diamond$  (see **B.BRTDP**), namely  $\text{Diff}_\infty(s) = \Delta(s, a(s)) \langle \text{Diff}_\infty \rangle$  where  $a(s)$  is a specially chosen “optimal” action for each state satisfying  $\text{Diff}_\infty(s, a(s)) = \text{Diff}_\infty(s)$ . Now, set  $\text{Diff}_* = \max_{s \in S_+^\diamond} \text{Diff}_\infty(s)$  the maximal difference on  $S_+^\diamond$  and let  $S_*^\diamond$  be the set of witnesses obtaining  $\text{Diff}_*$ . Then,  $\Delta(s, a(s), S_*^\diamond) = 1$ : If a part of the transition’s probability mass would move to a region with smaller difference, an appropriate update of a pair close to  $(s, a(s))$  would reduce its difference. Hence, the set of states  $S_*^\diamond$  is a “stable” subset of the system when following the actions  $a(s)$ . By **Absorption**, we eventually have to reach either the target  $T$  or the sink  $R$  starting from any state in  $S_*^\diamond$ . Since  $\text{Diff}_\infty(s) = 0$  for all (sampled) states in  $T \cup R$  and  $\text{Diff}_\infty$  satisfies the fixed point equation, we get that  $\text{Diff}_\infty(s) = 0$  for all states  $S_*^\diamond$  and consequently  $\text{Diff}_\infty(s_0) = 0$ . ◀

## 5 Discussion

### 5.1 Relation to Algorithms for Finite Systems and Discretization

Our algorithm directly generalizes the classical value iteration as well as BRTDP for finite MDP by an appropriate choice of GETPAIR. In value iteration, it proceeds in round-robin fashion, enumerating all state-action pairs. Note that the algorithm immediately uses the results of previous updates, corresponding to the *Gauß-Seidel variant* of VI; to exactly obtain synchronous value iteration, we would have to slightly modify the structure for saving the values. In BRTDP, GETPAIR simulates paths through the MDP and we update only those states encountered during the simulation.

Approaches based on discretization through, e.g., grids with increasing precision, essentially reduce the uncountable state space to a finite one. This is also encompassed by GETPAIR, e.g. by selecting the grid points in round robin or randomized fashion. However, our algorithm has the following key advantages when compared to classical discretization. Firstly, it avoids the need to grid the whole state space (typically into cells of regular sizes). Secondly, in discretization, updating the value of one cell does not directly affect the value in other cells; in contrast in our algorithm, knowledge about a state fluently propagates to other areas (by using Equation (3)) without being hindered by (arbitrarily chosen) cell boundaries.

### 5.2 Extensions

We outline possible extensions and augmentations of our approach to showcase its versatility.

**Discontinuities.** Our Lipschitz assumption **C** actually is slightly stronger than required. We first give an example of a system exhibiting discontinuities and then describe how our approach can be modified to deal with it. More details are in [17, App. C.2].

► **Example 8.** Consider a robot navigating a terrain with cliffs, where falling down a cliff immediately makes it impossible to reach the target. There, states which are barely on the edge may still reach the goal with significant probability, while a small step to the side results in falling down the cliff and zero probability of reaching the goal.

To solve this example, one could model the cliff as a steep but continuous slope, which would make our approach still possible. Unfortunately, this might not be very practical, since the Lipschitz constant then is quite large.

However, if we know of discontinuities, e.g. the location of cliffs in the terrain the robot navigates, both our algorithms can be extended as follows: Instead of requiring  $\mathcal{V}$  to be continuous on the whole domain, we may assume that we are given a (finite, decidable) partitioning of the state set  $S$  into several sets  $S_i$ . We allow the value function to be discontinuous along the boundaries of  $S_i$  (the cliffs), as long as it remains Lipschitz-continuous inside each  $S_i$ . We only need to slightly modify the assumption on GETPAIR by requiring that for any state-action pair  $(s, a)$  with  $s \in S_i$  we eventually get a nearby, similarly behaving state-action pair  $(s', a')$  of the same region, i.e.  $s' \in S_i$ . While computing the bounds of a particular state-action pair, e.g.  $U(s, a)$ , we first determine which partition  $S_i$  the state  $s$  belongs to and then only consider the stored values of states inside the region  $S_i$ .

**Linear Temporal Logic.** In [9], the authors extend BRTDP to LTL queries [42]. Several difficulties arise in the uncountable setting. For example, in order to prove *liveness* conditions, we need to solve the *repeated reachability* problem, i.e. whether a particular set of states is reached infinitely often. This is difficult even for restricted classes of uncountable systems, and impossible in the general case. In particular, [9] relies on analysing end components, which we already identified as an unresolved problem. We provide further insight in [17, App. C.3]. Nevertheless, there is a straightforward extension of our approach to the subclass of *reach-avoid* problems [50] (or *constrained reachability* [52]), see [17, App. C.4].

### 5.3 Implementation and Heuristics

For completeness, we implemented a prototype of our BRTDP algorithm to demonstrate its effectiveness. See [17, App. D] for details and an evaluation on both a one- and two-dimensional navigation model. Our implementation is barely optimized, with no delegation to high-performance libraries. Yet, these non-trivial models are solved in reasonable time. However, since we aim for assumptions that are as general as possible, one cannot expect our generic approach perform on par with highly optimized tools. Our prototype serves as a proof-of-concept and does not aim to be competitive with specialized approaches. We highlight again that the goal of our paper is *not* to be practically efficient in a particular, restricted setting, but rather to provide general assumptions and theoretical algorithms applicable to all kinds of uncountable systems.

Aside from several possible optimizations concerning the concrete implementation, we suggest two more general directions for heuristics:

**Adaptive Lipschitz constants.** As an example, suppose that a robot is navigating mostly flat land close to its home, but more hilly terrain further away. The flat land has a smaller Lipschitz constant than the hilly terrain, and thus here we can infer tighter bounds. More generally, given a partitioning of the state space and local Lipschitz constants for every subset, we use this local knowledge when computing  $\hat{L}$  and  $\hat{U}$  instead of using the global Lipschitz constant, which is the maximum of all local ones. See [17, App. C.2] for details.

**GetPair-heuristics.** In Section 3.1.2, we mentioned two simple implementations of GETPAIR. Firstly, we can discretize both state and action space, yielding each state-action pair in the discretization for a finite number of iterations, choosing a finer discretization constant, and repeating the process until convergence. Assuming that we can sample all state-action pairs in the discretization, this method eventually samples arbitrarily close to *any* state-action pair in  $S \times Av$  and thus trivially satisfies the sampling assumption. This intuitively corresponds to executing interval iteration [19] on the (increasingly refined) discretized systems. Note that this approach completely disregards the reachability probability of certain states and invests the same computational effort for all of them. In particular, it invests the same amount of computational effort into regions which are only reached with probability  $10^{-100}$  as in regions around the initial state  $s_0$ .

Thus, a second approach is to sample a path through the system at random, following random actions. This approach updates states roughly proportional to the probability of being reached, which already in the finite setting yields dramatic speed-ups [34].

However, we can also use further information provided by the algorithm, namely the upper bounds. As mentioned in [9], following “promising” actions with a large upper bound proves to be beneficial, since actions with small upper bound likely are suboptimal. To extend this idea to the general domain, we need to apply a bit of care. In particular, it might be difficult to select exactly from the optimal set of actions, since already  $\arg \max_{a \in Av(s)} U(s, a)$  might be very difficult to compute. Yet, it is sufficient to choose some constant  $\xi > 0$  and over-approximate the set of  $\xi$ -optimal actions in a given state, randomly selecting from this set. This over-approximation can easily be performed by, for example, randomly sampling the set of available actions  $Av(s)$  until we encounter an action close to the optimum (which can approximate due to our assumptions). By generating paths only using these actions, we combine the previous idea of focussing on “important” states (in terms of reachability) with an additional focus on “promising” states (in terms of upper bounds). This way, the algorithm *learns* from its experiences, using it as a guidance for future explorations.

More generally, we can easily apply more sophisticated learning approaches by interleaving it with one of the above methods. For example, by following the learning approach with probability  $\nu$  and a “safe” method with probability  $1 - \nu$  we still obtain a safe heuristic, since the assumption only requires limit behaviour. As such, we can combine our approach with existing, learning based algorithm by following their suggested heuristic and interleave it with some sampling runs guided by the above ideas. In other words, this means that the learning algorithm can focus on finding a reasonable solution quickly, which is then subsequently *verified* by our approach, potentially *improving* the solution in areas where the learner is performing suboptimally. On top, the (guaranteed) bounds identified by our algorithm can be used as feedback to the learning algorithm, creating a positive feedback loop, where both components improve each other’s behaviour and performance.

## 6 Conclusion

In this work, we have presented the first anytime algorithm to tackle the reachability problem for MDP with uncountable state- and action-spaces, giving both correctness and termination guarantees under general assumptions. The experimental evaluation of our prototype implementation shows both promising results and room for improvements.

On the theoretical side, we conjecture that **Assumption D: Absorption** can be weakened if we complement it with an automatic procedure that finds and treats problematic parts of the state space of a certain kind, similar to the collapsing approach on finite MDP



[19, 9]. Note that as the general problem is undecidable, some form of **Absorption** will remain necessary. On the practical side, we aim for a more sophisticated tool, applying our theoretical foundation to the full range of MDP, including discrete discontinuities. Moreover, we want to combine the tool with existing ways of identifying the Lipschitz constant.

---

## References

---

- 1 Alessandro Abate, Saurabh Amin, Maria Prandini, John Lygeros, and Shankar Sastry. Computational approaches to reachability analysis of stochastic hybrid systems. In *HSCC*, volume 4416 of *Lecture Notes in Computer Science*, pages 4–17. Springer, 2007. doi:10.1007/978-3-540-71493-4\_4.
- 2 Alessandro Abate, Joost-Pieter Katoen, John Lygeros, and Maria Prandini. Approximate model checking of stochastic hybrid systems. *Eur. J. Control*, 16(6):624–641, 2010. doi:10.3166/ejc.16.624-641.
- 3 Alessandro Abate, Maria Prandini, John Lygeros, and Shankar Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008. doi:10.1016/j.automatica.2008.03.027.
- 4 Christel Baier, Joachim Klein, Linda Leuschner, David Parker, and Sascha Wunderlich. Ensuring the reliability of your model checker: Interval iteration for Markov decision processes. In *CAV (1)*, volume 10426 of *Lecture Notes in Computer Science*, pages 160–180. Springer, 2017.
- 5 Dimitri Bertsekas. Convergence of discretization procedures in dynamic programming. *IEEE Transactions on Automatic Control*, 20(3):415–419, 1975.
- 6 Dimitri P Bertsekas and Steven Shreve. Stochastic optimal control: the discrete-time case, 1978.
- 7 Dimitri P. Bertsekas and John N. Tsitsiklis. An analysis of stochastic shortest path problems. *Math. Oper. Res.*, 16(3):580–595, 1991. doi:10.1287/moor.16.3.580.
- 8 Patrick Billingsley. *Probability and Measure*, volume 939. John Wiley & Sons, 2012.
- 9 Tomáš Brázdil, Krishnendu Chatterjee, Martin Chmelik, Vojtech Forejt, Jan Křetínský, Marta Z. Kwiatkowska, David Parker, and Mateusz Ujma. Verification of Markov decision processes using learning algorithms. In *ATVA*, volume 8837 of *Lecture Notes in Computer Science*, pages 98–114. Springer, 2014. doi:10.1007/978-3-319-11936-6\_8.
- 10 John L. Bresina, Richard Dearden, Nicolas Meuleau, Sailesh Ramakrishnan, David E. Smith, and Richard Washington. Planning under continuous time and resource uncertainty: A challenge for AI. *CoRR*, abs/1301.0559, 2013. arXiv:1301.0559.
- 11 Debasish Chatterjee, Eugenio Cinquemani, and John Lygeros. Maximizing the probability of attaining a target prior to extinction. *Nonlinear Analysis: Hybrid Systems*, 5(2):367–381, 2011.
- 12 Krishnendu Chatterjee, Zuzana Křetínská, and Jan Křetínský. Unifying two views on multiple mean-payoff objectives in Markov decision processes. *Logical Methods in Computer Science*, 13(2), 2017. doi:10.23638/LMCS-13(2:15)2017.
- 13 Anne Condon. The complexity of stochastic games. *Inf. Comput.*, 96(2):203–224, 1992. doi:10.1016/0890-5401(92)90048-K.
- 14 Zhengzhu Feng, Richard Dearden, Nicolas Meuleau, and Richard Washington. Dynamic programming for structured continuous Markov decision problems. In *UAI*, pages 154–161. AUAI Press, 2004. URL: [https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article\\_id=1102&proceeding\\_id=20](https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=1102&proceeding_id=20).
- 15 Vojtech Forejt, Marta Z. Kwiatkowska, Gethin Norman, and David Parker. Automated verification techniques for probabilistic systems. In *SFM*, volume 6659 of *Lecture Notes in Computer Science*, pages 53–113. Springer, 2011. doi:10.1007/978-3-642-21455-4\_3.

- 16 Hongfei Fu and Krishnendu Chatterjee. Termination of nondeterministic probabilistic programs. In *VMCAI*, volume 11388 of *Lecture Notes in Computer Science*, pages 468–490. Springer, 2019. doi:10.1007/978-3-030-11245-5\_22.
- 17 Kush Grover, Jan Kretínský, Tobias Meggendorfer, and Maximilian Weininger. Anytime guarantees for reachability in uncountable markov decision processes. *CoRR*, abs/2008.04824, 2020. arXiv:2008.04824.
- 18 Carlos Guestrin, Milos Hauskrecht, and Branislav Kveton. Solving factored MDPs with continuous and discrete variables. In *UAI*, pages 235–242. AUAI Press, 2004. URL: [https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article\\_id=1113&proceeding\\_id=20](https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=1113&proceeding_id=20).
- 19 Serge Haddad and Benjamin Monmege. Interval iteration algorithm for MDPs and IMDPs. *Theor. Comput. Sci.*, 735:111–131, 2018. doi:10.1016/j.tcs.2016.12.003.
- 20 Sofie Haesaert, Sadegh Soudjani, and Alessandro Abate. Temporal logic control of general Markov decision processes by approximate policy refinement. In *ADHS*, volume 51(16) of *IFAC-PapersOnLine*, pages 73–78. Elsevier, 2018. doi:10.1016/j.ifacol.2018.08.013.
- 21 Sofie Haesaert, Sadegh Esmaeil Zadeh Soudjani, and Alessandro Abate. Verification of general Markov decision processes by approximate similarity relations and policy refinement. *SIAM J. Control and Optimization*, 55(4):2333–2367, 2017. doi:10.1137/16M1079397.
- 22 Arnd Hartmanns and Benjamin Lucien Kaminski. Optimistic value iteration. In *CAV (2)*, volume 12225 of *Lecture Notes in Computer Science*, pages 488–511. Springer, 2020.
- 23 Mohammadhosein Hasanbeig, Alessandro Abate, and Daniel Kroening. Logically-constrained neural fitted q-iteration. In *AAMAS*, pages 2012–2014. International Foundation for Autonomous Agents and Multiagent Systems, 2019. URL: <http://dl.acm.org/citation.cfm?id=3331994>.
- 24 Mohammadhosein Hasanbeig, Alessandro Abate, and Daniel Kroening. Certified reinforcement learning with logic guidance. *CoRR*, abs/1902.00778, 2019. arXiv:1902.00778.
- 25 William B. Haskell, Rahul Jain, Hiteshi Sharma, and Pengqian Yu. A universal empirical dynamic programming algorithm for continuous state MDPs. *IEEE Trans. Automat. Contr.*, 65(1):115–129, 2020. doi:10.1109/TAC.2019.2907414.
- 26 Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. What’s decidable about hybrid automata? In *STOC*, pages 373–382. ACM, 1995.
- 27 Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. What’s decidable about hybrid automata? *J. Comput. Syst. Sci.*, 57(1):94–124, 1998. doi:10.1006/jcss.1998.1581.
- 28 Onésimo Hernández-Lerma and Jean B Lasserre. *Discrete-time Markov control processes: basic optimality criteria*, volume 30. Springer Science & Business Media, 2012.
- 29 Ronald A Howard. *Dynamic programming and Markov processes*. John Wiley, 1960.
- 30 Mingzhang Huang, Hongfei Fu, and Krishnendu Chatterjee. New approaches for almost-sure termination of probabilistic programs. In *Program. Lang. and Sys.*, volume 11275 of *Lecture Notes in Computer Science*, pages 181–201. Springer, 2018. doi:10.1007/978-3-030-02768-1\_11.
- 31 Mingzhang Huang, Hongfei Fu, Krishnendu Chatterjee, and Amir Kafshdar Goharshady. Modular verification for almost-sure termination of probabilistic programs. *Proc. ACM Program. Lang.*, 3(OOPSLA):129:1–129:29, 2019. doi:10.1145/3360555.
- 32 Manfred Jaeger, Peter Gjøøl Jensen, Kim Guldstrand Larsen, Axel Legay, Sean Sedwards, and Jakob Haahr Taankvist. Teaching stratego to play ball: Optimal synthesis for continuous space MDPs. In *ATVA*, volume 11781 of *Lecture Notes in Computer Science*, pages 81–97. Springer, 2019. doi:10.1007/978-3-030-31784-3\_5.
- 33 Benjamin Lucien Kaminski and Joost-Pieter Katoen. On the hardness of almost-sure termination. In *MFCS*, volume 9234 of *Lecture Notes in Computer Science*, pages 307–318. Springer, 2015. doi:10.1007/978-3-662-48057-1\_24.
- 34 Jan Kretínský and Tobias Meggendorfer. Of cores: A partial-exploration framework for Markov decision processes. In *CONCUR*, volume 140 of *LIPICs*, pages 5:1–5:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.CONCUR.2019.5.

- 35 Ratan Lal and Pavithra Prabhakar. Bounded verification of reachability of probabilistic hybrid systems. In *QEST*, volume 11024 of *Lecture Notes in Computer Science*, pages 240–256. Springer, 2018. doi:10.1007/978-3-319-99154-2\_15.
- 36 Bernard F Lamond and Abdeslem Boukhtouta. Water reservoir applications of Markov decision processes. In *Handbook of Markov decision processes*, pages 537–558. Springer, 2002.
- 37 Lihong Li and Michael L. Littman. Lazy approximation for solving continuous finite-horizon MDPs. In *AAAI*, pages 1175–1180. AAAI Press / The MIT Press, 2005. URL: <http://www.aaai.org/Library/AAAI/2005/aaai05-186.php>.
- 38 Masoud Mahootchi. *Storage system management using reinforcement learning techniques and nonlinear models*. PhD thesis, University of Waterloo, 2009.
- 39 H. Brendan McMahan, Maxim Likhachev, and Geoffrey J. Gordon. Bounded real-time dynamic programming: RTDP with monotone upper bounds and performance guarantees. In *ICML*, volume 119 of *ACM International Conference Proceeding Series*, pages 569–576. ACM, 2005. doi:10.1145/1102351.1102423.
- 40 Francisco S. Melo, Sean P. Meyn, and M. Isabel Ribeiro. An analysis of reinforcement learning with function approximation. In *ICML*, volume 307 of *ACM International Conference Proceeding Series*, pages 664–671. ACM, 2008. doi:10.1145/1390156.1390240.
- 41 Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006.
- 42 Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 46–57. IEEE Computer Society, 1977. doi:10.1109/SFCS.1977.32.
- 43 Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, 1994. doi:10.1002/9780470316887.
- 44 Tim Quatmann and Joost-Pieter Katoen. Sound value iteration. In *CAV (1)*, volume 10981 of *Lecture Notes in Computer Science*, pages 643–661. Springer, 2018.
- 45 Scott Sanner, Karina Valdivia Delgado, and Leliane Nunes de Barros. Symbolic dynamic programming for discrete and continuous state MDPs. In *UAI*, pages 643–652. AUAI Press, 2011. URL: [https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article\\_id=2223&proceeding\\_id=27](https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=2223&proceeding_id=27).
- 46 Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.
- 47 Hiteshi Sharma, Mehdi Jafarnia-Jahromi, and Rahul Jain. Approximate relative value learning for average-reward continuous state MDPs. In *UAI*, page 341. AUAI Press, 2019. URL: <http://auai.org/uai2019/proceedings/papers/341.pdf>.
- 48 Fedor Shmarov and Paolo Zuliani. Probreach: verified probabilistic delta-reachability for stochastic hybrid systems. In *HSCC*, pages 134–139. ACM, 2015.
- 49 Sadegh Esmail Zadeh Soudjani and Alessandro Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *QEST*, pages 59–68. IEEE Computer Society, 2011. doi:10.1109/QEST.2011.16.
- 50 Sean Summers and John Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961, 2010. doi:10.1016/j.automatica.2010.08.006.
- 51 Ilya Tkachev, Alexandru Mereacre, Joost-Pieter Katoen, and Alessandro Abate. Quantitative automata-based controller synthesis for non-autonomous stochastic hybrid systems. In *HSCC*, pages 293–302. ACM, 2013. doi:10.1145/2461328.2461373.
- 52 Ilya Tkachev, Alexandru Mereacre, Joost-Pieter Katoen, and Alessandro Abate. Quantitative model-checking of controlled discrete-time Markov processes. *Inf. Comput.*, 253:1–35, 2017. doi:10.1016/j.ic.2016.11.006.
- 53 Hado van Hasselt. Reinforcement learning in continuous state and action spaces. In *Reinforcement Learning*, volume 12 of *Adaptation, Learning, and Optimization*, pages 207–251. Springer, 2012. doi:10.1007/978-3-642-27645-3\_7.

## 11:20 Anytime Guarantees for Reachability in Uncountable Markov Decision Processes

- 54 Luis Gustavo Rocha Vianna, Scott Sanner, and Leliane Nunes de Barros. Continuous real time dynamic programming for discrete and continuous state MDPs. In *2014 Brazilian Conference on Intelligent Systems, BRACIS 2014, Sao Paulo, Brazil, October 18-22, 2014*, pages 134–139. IEEE Computer Society, 2014. doi:10.1109/BRACIS.2014.34.