

Finite Optimal Control for Time-Bounded Reachability in CTMDPs and Continuous-Time Markov Games [★]

Markus Rabe¹ and Sven Schewe²

¹ Universität des Saarlandes

² University of Liverpool

Abstract. We establish the existence of optimal scheduling strategies for time-bounded reachability in continuous-time Markov decision processes, and of co-optimal strategies for continuous-time Markov games. Furthermore, we show that optimal control does not only exist, but has a surprisingly simple structure: The optimal schedulers from our proofs are deterministic and timed-positional, and the bounded time can be divided into a finite number of intervals, in which the optimal strategies are positional. That is, we demonstrate the existence of *finite* optimal control. Finally, we show that these pleasant properties of Markov decision processes extend to the more general class of continuous-time Markov games, and that both early and late schedulers show this behaviour.

1 Introduction

Continuous-time Markov decision processes (CTMDPs) are a widely used framework for dependability analysis and for modelling the control of manufacturing processes [12, 6], because they combine real-time aspects with probabilistic behaviour and non-deterministic choices. CTMDPs can also be viewed as a framework that unifies different stochastic model types [14, 12, 9, 7, 10].

While CTMDPs allow for analysing worst-case and best-case scenarios, they fall short of the demands that arise in many real control problems, as they disregard the different nature that non-determinism can have depending on its source: Some sources of non-determinism are supportive, while others are hostile, and in a realistic control scenario, we face both types of non-determinism at the same time: Supportive non-determinism can be used to model the influence of a controller on the evolution of a system, while hostile non-determinism can capture abstraction or unknown environments. We therefore consider a natural extension of CTMDPs: Continuous-time Markov games (CTMGs) that have two players with opposing objectives [5].

[★] This work was partly supported by the German Research Foundation (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS) and by the Engineering and Physical Science Research Council (EPSRC) through grant EP/H046623/1 “Synthesis and Verification in Markov Game Structures”.




Fig. 1. A CTMDP and the reachability probabilities for all positional schedulers with time bound $t_{\max} = 1$. Time $t' = t_{\max} - \frac{1}{2} \log(2)$ is the optimal time to switch to action b .

The analysis of CTMDPs and CTMGs requires to resolve the non-deterministic choices by means of a scheduler (which consists of a pair of strategies in the case of CTMGs), and typically tries to optimise a given objective function.

In this paper, we study the *time-bounded reachability problem*, which recently enjoyed much attention [3, 16, 10, 11, 5]. Time-bounded reachability in CTMDPs is the standard control problem to construct a scheduler that controls the Markov decision process such that the probability of reaching a goal region within a given time bound is maximised (or minimised), and to determine the value. For CTMGs, time-bounded reachability reduces to finding a Nash equilibrium, that is, a pair of strategies for the players, such that each strategy is optimal for the chosen strategy of her opponent.

While continuous-time Markov games are a young field of research [5, 13], Markov decision processes have been studied for decades [8, 4].

Optimal control in CTMDPs clearly depends on the observational power we allow our schedulers to have when observing a run. In the literature, various classes of schedulers with different restrictions on what they can observe [15, 10, 5, 13] are considered. We focus on the most general class of schedulers, schedulers that can fully observe the system state and may change their decisions at any point in time (*late schedulers*, cf. [4, 10]). To be able to translate our results to the more widespread class of schedulers that fix their decisions when entering a location (*early schedulers*), we introduce discrete locations that allow for a translation from early to late schedulers (see Appendix D).

Due to their practical importance, time-bounded reachability for continuous-time Markov models has been studied intensively [5, 4, 11, 3, 10, 1, 2, 16]. However, most previous research focussed on *approximating* optimal control. (The existence of optimal control is currently only known for the artificial class of *time-abstract* schedulers [5, 13], which assume that the scheduler has no access whatsoever to a clock.) While an efficient approximation is of interest to a practitioner, being unable to determine whether or not optimal control *exists* is very dissatisfying from a scientific point of view.

Contributions. This paper has three main contributions: First, we extend the common model of CTMDPs by adding discrete locations, which are passed in 0 time. This generalisation of the model is mainly motivated by avoiding the discussion about the appropriate scheduler class. In particular, the widespread class of schedulers that fix their actions when entering a location can be encoded by a simple mapping.

The second contribution of this paper is the answer to an intriguing research question that remained unresolved for half a century: We show that optimal control of CTMDPs exists for time-bounded reachability and safety objectives. Moreover, we show that optimal control can always be *finite*.

Our third contribution is to lift these results to continuous-time Markov games.

Pursuing a different research question, we exploit proof techniques that differ from those frequently used in the analysis of CTMDPs. Our proofs build mainly on topological arguments: The proof that demonstrates the existence of measurable optimal schedulers, for example, shows that we can fix the decisions of an optimal scheduler successively on closures of open sets (yielding only measurable sets), and the lift to finiteness uses local optimality of positional schedulers in open left and right environments of arbitrary points of times and the compactness of the bounded time interval.

Structure of the Paper. We follow a slightly unorthodox order of proofs for a mathematical paper: we start with a special case in Section 3 and generalise the results later. Besides keeping the proofs simple, this approach is chosen because the simplest case, CTMDPs, is the classical case, and we assume that a wider audience is interested in results for these structures. In the following section, we strengthen this result by demonstrating that optimal control does not only exist, but can be found among schedulers with finitely many switching points and positional strategies between them. In Section 5, we lift this result to single player games (thus extending it to other scheduler classes like those which fix their decision when entering a location, cf. Appendix D). In the final section, we generalise the existence theorem for finite optimal control to finite co-optimal strategies for general continuous-time Markov games.

2 Preliminaries

A continuous-time Markov game is a tuple $(L, L_d, L_c, L_r, L_s, G, Act, \mathbf{R}, \mathbf{P}, \nu)$, consisting of

- a finite set L of locations, which is partitioned into
 - a set L_d of *discrete* locations and a set L_c of *continuous* locations, and
 - sets L_r and L_s of locations owned by a *reachability* and a *safety* player,
- a dedicated set $G \subseteq L$ of *goal locations*,
- a finite set Act of actions,
- a rate matrix $\mathbf{R} : (L_c \times Act \times L) \rightarrow \mathbb{Q}_{\geq 0}$,
- a discrete transition matrix $\mathbf{P} : (L \times Act \times L) \rightarrow \mathbb{Q}_{\geq 0} \cap [0, 1]$, and

- an initial distribution $\nu \in \text{Dist}(L)$,

that satisfies the following side-conditions: For all continuous locations $l \in L_c$, there must be an action $a \in \text{Act}$ such that $\mathbf{R}(l, a, L) := \sum_{l' \in L} \mathbf{R}(l, a, l') > 0$; we call such actions *enabled*. For actions enabled in continuous locations, we require $\mathbf{P}(l, a, l') = \frac{\mathbf{R}(l, a, l')}{\mathbf{R}(l, a, L)}$, and we require $\mathbf{P}(l, a, l') = 0$ for the remaining actions. For discrete locations, we require that either $\mathbf{P}(l, a, l') = 0$ holds for all $l' \in L$, or that $\sum_{l' \in L} \mathbf{P}(l, a, l') = 1$ holds true. Like in the continuous case, we call the latter actions *enabled* and require the existence of at least one enabled action for each discrete location $l \in L_d$.

The idea behind discrete-time locations is that they execute immediately. We therefore do not permit cycles of only discrete-time locations (counting every positive rate of any action as a transition). This restriction is stronger than it needs to be, but it simplifies our proofs, and the simpler model is sufficient for our means.

We assume that the goal region is absorbing, that is $\mathbf{P}(l, a, l') = 0$ holds for all $l \in G$ and $l' \notin G$. See Section 7 for the extension to non-absorbing goal regions.

Intuitively, it is the objective of the reachability player to maximise the probability to reach the goal region in a predefined time t_0 , while it is the objective of the safety player to minimise this probability. (Hence, it is a zero-sum game.)

We are particularly interested in (traditional) CTMDPs. They are single player CTMGs, where either all positions belong to the reachability player ($L = L_r$), or to the safety player ($L = L_s$), without discrete locations ($L_d = \emptyset$ and $L_c = L$).

Paths. A *timed path* π in a CTMG \mathcal{M} is a finite sequence in $L \times (\text{Act} \times \mathbb{R}_{\geq 0} \times L)^* = \text{Paths}(\mathcal{M})$. We write

$$l_0 \xrightarrow{a_0, t_0} l_1 \xrightarrow{a_1, t_1} \dots \xrightarrow{a_{n-1}, t_{n-1}} l_n$$

for a sequence π , and we require $0 \leq t_{i-1} \leq t_i \leq t_{\max}$ for all $i < n$, where t_{\max} is the time bound for our time-bounded reachability probability. (We are not interested in the behaviour of the system after t_{\max} .) The t_i denote the system's time when the action a_i is selected and a discrete transition from l_i to l_{i+1} takes place. Concatenation of paths π, π' will be written as $\pi \circ \pi'$ if the last location of π is the first location of π' and the points of time are ordered correctly. We call a timed path a *complete* timed path when we want to stress that this path describes a complete system run, not to be extended by further transitions.

Schedulers and Strategies. The nondeterminism in the system needs to be resolved by a scheduler which maps paths to decisions. The power of schedulers is determined by their ability to observe and distinguish paths, and thus by their domain. In this paper, we consider the following common scheduler classes:

- *Timed history-dependent* (TH) schedulers $\text{Paths}(\mathcal{M}) \times \mathbb{R}_{\geq 0} \rightarrow D$
that map timed paths and the remaining time to decisions.

- *Timed positional* (TP) schedulers $L \times \mathbb{R}_{\geq 0} \rightarrow D$
that map locations and the remaining time to decisions.
- *Positional* (P) or memoryless schedulers $L \rightarrow D$
that map locations to decisions.

Decisions D are either randomised (R), in which case $D = \text{Dist}(\text{Act})$ is the set of distributions over enabled actions, or are restricted to deterministic (D) choices, that is $D = \text{Act}$. Where it is necessary to distinguish randomised and deterministic versions we will add a postfix to the scheduler class, for example THD and THR.

Strategies. In case of CTMGs, a scheduler consists of the two participating players' strategies, which can be seen as functions $\text{Paths}(\mathcal{M}_p) \times \mathbb{R}_{\geq 0} \rightarrow D$, where $\text{Paths}(\mathcal{M}_p)$ denotes, for $p \in \{r, s\}$, the paths ending on the position of the reachability or safety player, respectively. As for general schedulers, we can introduce restrictions on what players are able to observe.

Discrete locations. The main motivation to introduce discrete locations was to avoid the discussion whether a scheduler has to fix his decision, as to which action it chooses, upon entering a location, or whether such a decision can be revoked while staying in the location. For example, the general measurable schedulers discussed in [15] have only indirect access to the remaining time (through the timed path), and therefore have to decide upon entrance of a location which action they want to perform. Our definition builds on fully-timed schedulers (cf. [4]) that were recently rediscovered and formalised by Neuhäüßer et al. [10], which may revoke their decision after they enter a location. (As a side result, we lift Neuhäüßer's restriction to local uniformity.) The discrete locations now allow to encode making the decision upon entering a continuous location l by mapping the decision to a discrete location that is 'guarding the entry' to a family of continuous locations, one for each action enabled in l . (See Appendix D for details.)

Cylindrical Schedulers. While it is common to refer to TH schedulers as a class, the truth is that there is no straightforward way to define a measure for the time-bounded reachability probability for the complete class (cf. [15]). We therefore turn to a natural subset that can be used as a building block for a powerful yet measurable sub-class of TH schedulers, which is based on cylindrical abstractions of paths.

Let \mathcal{J} be a finite partition of the interval $[0, t_{\max}]$ into intervals $I_0 = [0, t_0]$ and $I_i = (t_{i-1}, t_i]$ for $i = 1, \dots, n$ with $t_0 \geq 0$ and $t_i > t_{i-1}$ for $i = 1, \dots, n$, where $t_n = t_{\max}$ is the time-bound from the problem definition. Then we denote with $[t]_{\mathcal{J}}$ the interval $I_i \in \mathcal{J}$ that contains t , called the *\mathcal{J} -cylindrification of t* , and we denote with $[\pi]_{\mathcal{J}} = l_0 \xrightarrow{a_0, [t'_0]_{\mathcal{J}}} l_1 \xrightarrow{a_1, [t'_1]_{\mathcal{J}}} \dots \xrightarrow{a_{n-1}, [t'_{n-1}]_{\mathcal{J}}} l_n$ the *\mathcal{J} -cylindrification of the timed path $\pi = l_0 \xrightarrow{a_0, t'_0} l_1 \xrightarrow{a_1, t'_1} \dots \xrightarrow{a_{n-1}, t'_{n-1}} l_n$* .

We call a TH scheduler *\mathcal{J} -cylindrical* if its decisions depend only on the *cylindrification* $[\pi]_{\mathcal{J}}$ and $[t]_{\mathcal{J}}$ of π and t , respectively, and *cylindrical* if it is *\mathcal{J} -cylindrical* for some finite partition \mathcal{J} of the interval $I = [0, t_{\max}]$.

Cylindrical Sets and Probability Space. For a given finite partition \mathcal{J} of the interval $[0, t_{\max}]$, an \mathcal{J} -cylindrical set of timed paths is the set of timed paths with the same \mathcal{J} -cylindrification, and we call a finite partition \mathcal{J}' of $[0, t_{\max}]$ a *refinement* of \mathcal{J} if every interval in \mathcal{J} is the union of intervals in \mathcal{J}' .

For an \mathcal{J} -cylindrical scheduler \mathcal{S} and an \mathcal{J}' -cylindrical set of finite timed paths, where \mathcal{J}' is a refinement³ of \mathcal{J} , the likelihood that a complete path is from this cylindrical set is easy to define: Within each interval of \mathcal{J} , the likelihood that a CTMDP \mathcal{M} with scheduler \mathcal{S} behaves in accordance with the \mathcal{J}' -cylindrical set can—assuming compliance in all previous intervals—be checked like for a finite Markov chain.

The probability p_{I_i} to comply with the i -th segment of the partition \mathcal{J}' of $[0, t_{\max}]$ is the product of three multiplicands ($p_{I_i} = p_1^{I_i} \cdot p_2^{I_i} \cdot p_3^{I_i}$):

1. the probability $p_1^{I_i}$ that the actions are *chosen* in accordance with the \mathcal{J}' -cylindrical set of timed paths (which is either 0 or 1 for discrete schedulers, and the product of the likelihood of the individual decisions for randomised schedulers),
2. the probability $p_2^{I_i}$ that the transitions are *taken* in accordance with the \mathcal{J}' -cylindrical set of timed paths, provided the respective actions are chosen, which is simply the product over the individual probabilities $\mathbf{P}(l_i, a_i, l_{i+1})$ in this sequence of the \mathcal{J}' -cylindrical set of timed paths, and
3. the probability $p_3^{I_i}$ that the *right number of steps* is made in this sequence of the \mathcal{J}' -cylindrical set of timed paths.

The latter probability $p_3^{I_i}$ is 0 if the last location is a discrete location, as the system would leave this location at the same point in time in which it was entered. Otherwise, it is the difference $p_3^{I_i} = p_4^{I_i} - p_5^{I_i}$ between the likelihood that at least the correct number of $n \geq 0$ transitions starting in continuous locations are made ($p_4^{I_i}$), and the likelihood that at least $n + 1$ transitions starting in continuous locations are made ($p_5^{I_i}$) in the relevant sequence of the timed path.

Let $\bar{l}_0, \bar{l}_1, \dots, \bar{l}_n$ be the n continuous locations (named in the required order of appearance; note that n might be 0, and that the same location can occur multiple times), and let $\lambda_0, \lambda_1, \dots, \lambda_n$ be the transition rate one would observe at the respective \bar{l}_i . For deterministic schedulers, this transition rate is simply $\lambda_i = \mathbf{R}(\bar{l}_i, a_i, L)$, where a_i is the decision from \mathcal{S} at the respective position in a timed path and in I_i . For a randomised scheduler \mathcal{S} , it is the respective expected transition rate $\lambda_i = \sum_{a \in Act(\bar{l}_i)} p_a \mathbf{R}(\bar{l}_i, a, L)$, where p_a is the likelihood that \mathcal{S} makes the decision a at the respective position in a timed path and in I_i . Note that the locations and transition rates are fixed.

The likelihood to get a path of length $\geq n$ is then $\int_{(\tau_0, \dots, \tau_{n-1}) \in \Phi_{n,i}} \prod_{k=0}^{n-1} \lambda_k e^{-\lambda_k \tau_k} d\tau_k$ for $\Phi_{n,i} = \{(\tau_0, \dots, \tau_{n-1}) \in [0, t_{\max}]^n \mid \sum_{j=0}^{n-1} \tau_j \leq t_i - t_{i-1}\}$ for $n > 0$, and 1 for $n = 0$. Likewise, the likelihood to get a path of length $\geq n + 1$ is $\int_{(\tau_0, \dots, \tau_n) \in \Phi_{n+1,i}^i} \prod_{k=0}^n \lambda_k e^{-\lambda_k \tau_k} d\tau_k$ for $\Phi_{n+1,i}^i =$

³ The restriction to partitions \mathcal{J}' that refine \mathcal{J} is purely technical, because for arbitrary \mathcal{J}' we can simply use a partition \mathcal{J}'' that refines both \mathcal{J} and \mathcal{J}' , and reconstruct every \mathcal{J}' -cylindrical set as a finite union of \mathcal{J}'' -cylindrical sets.

$\{(\tau_0, \dots, \tau_n) \in [0, t_{\max}]^{n+1} \mid \sum_{j=0}^n \tau_j \leq t_i - t_{i-1}\}$. (Recall that t_i and t_{i-1} are the upper and lower endpoints of the interval I_i .)

The likelihood that a complete timed path is in the \mathcal{J}' -cylindrical set of timed paths for \mathcal{S} is the product $\prod_{I \in \mathcal{J}'} p_I$ over the individual p_{I_i} .

Probability Space. Having defined a measure for the likelihood that, for a given cylindrical scheduler, a complete timed path is in a particular cylindrical set, we define the likelihood that it is in a finite union of disjoint sets of cylindrical paths as the sum over the likelihood for the individual cylindrical sets.

This primitive probability measure for primitive schedulers can be lifted in two steps by a standard space completion, going from this primitive measures to quotient classes of Cauchy sequences of such measures:

1. In a first step, we complete the space of measurable sets of complete timed paths from finite unions of cylindrical sets of timed paths to Cauchy sequences of finite unions of cylindrical sets of timed paths. We define the required difference measure of two sets of timed paths (each a finite disjoint union of cylindrical sets) as the measure of the symmetrical difference of the two sets. This set can obviously be represented as a finite disjoint union of cylindrical sets of timed paths, and we can use our primitive measure to define this difference measure.
2. Having lifted the measure to this completed space of paths, we lift the set of measurable *schedulers* in a second step from cylindrical schedulers to Cauchy sequences of cylindrical schedulers. (The difference measure between two cylindrical schedulers is the likelihood that two schedulers act observably different.)

More details of these standard constructions can be found in Appendix A.2.

Time-Bounded Reachability Probability. For a given CTMG $(L, L_d, L_c, L_r, L_s, G, Act, \mathbf{R}, \mathbf{P}, \nu)$ and a given measurable scheduler \mathcal{S} that resolves the non-determinism, we use the following notations for the probabilities:

- $Pr_{\mathcal{S}}^{\mathcal{M}}(l, t)$ is the probability of reaching the goal region G within time t when starting in location l ,
- $Pr_{\mathcal{S}}^{\mathcal{M}}(t) = \sum_{l \in L} \nu(l) Pr_{\mathcal{S}}^{\mathcal{M}}(l, t)$ denotes the probability of reaching the goal region G within time t .

As usual, the supremum of the time-bounded reachability probability over a particular scheduler class is called the time-bounded reachability of \mathcal{M} for this scheduler class.

3 Optimal Scheduling in CTMDPs

In this section, we demonstrate the existence of optimal schedulers in traditional CTMDPs. Before turning to the proof, let us first consider what happens if time

runs out, that is, at time t_{\max} , and then develop an intuition what an optimal scheduling policy should look like.

If we are still in time ($t \leq t_{\max}$) and we are in a goal location, then we reach a goal location in time with probability 1; and if time has run out ($t = t_{\max}$) and we are not in a goal location, then we reach a goal location in time with probability 0. For ease of notation, we also fix the probability of reaching the goal location in time to 0 for all points in time strictly *after* t_{\max} . For a measurable TPR scheduler \mathcal{S} , we would get:

- $Pr_{\mathcal{S}}^{\mathcal{M}}(l, t) = 1$ holds for all goal locations $l \in G$ and all $t \leq t_{\max}$,
- $Pr_{\mathcal{S}}^{\mathcal{M}}(l, t_{\max}) = 0$ holds for all non-goal locations $l \notin G$, and
- $Pr_{\mathcal{S}}^{\mathcal{M}}(l, t) = 0$ holds for all locations $l \in L$ and all $t > t_{\max}$.

A scheduler \mathcal{S} can, in every point in time, choose from distributions over successor locations. Such a choice should be optimal, if the expected gain in the probability of reaching the goal location is maximised.

This gain has two aspects: first, the probability of reaching the goal location *provided a transition is taken*, and second, the likelihood of *taking a transition*. Both are multiplicands in the defining differential equations, assuming a cylindrical TPD scheduler \mathcal{S}

$$-\dot{Pr}_{\mathcal{S}}^{\mathcal{M}}(l, t) = \sum_{l' \in L} \mathbf{R}(l, \mathcal{S}(l, t), l') \cdot (Pr_{\mathcal{S}}^{\mathcal{M}}(l', t) - Pr_{\mathcal{S}}^{\mathcal{M}}(l, t)).$$

(We skip the simple generalisation to measurable TPD scheduler because it is not required in the following proofs.)

The reachability probability of any scheduler is therefore intuitively dominated by the function f_{\max} and dominates the function f_{\min} defined by the following equations:

$$-\dot{f}_{\text{opt}}(l, t) = \text{opt}_{a \in \text{Act}(l)} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\text{opt}}(l', t) - f_{\text{opt}}(l, t)) \text{ for } t \in [0, t_{\max}],$$

where $\text{opt} \in \{\min, \max\}$. This intuitive result is not hard to prove⁴.

Lemma 1. *The reachability probability of any measurable THR scheduler is dominated by the function f_{\max} and dominates the function f_{\min} .*

Proof Idea: To proof this claim for f_{\max} , assume that there is a scheduler that provides a better time-bounded reachability probability $Pr_{\mathcal{S}}^{\mathcal{M}}(l, t) > f_{\max}(l, t)$

⁴ The systems of non-linear ordinary differential equations used in this paper are all quite obvious, and the challenge is to prove that they can be *taken* and not merely approximated. An approximative argument for these ODE's goes back to Bellman [4], but he uses a less powerful set of schedulers, and only proves that f_{\max} and f_{\min} can be approximated from below and above, respectively, claiming that the other direction is obvious. After starting with a similar claim, we were urged to include a full proof.

for some location $l \in L$ and time $t \in [0, t_{\max}]$ (in particular for $t = 0$), and hence improves over $f_{\max}(l, t)$ at this position by at least 3ε for some $\varepsilon > 0$.

\mathcal{S} is a Cauchy sequence of cylindrical schedulers. Therefore we can sacrifice one ε and get an ε -close cylindrical scheduler from this sequence, which is still at least 2ε better than f_{\max} at position (l, t) .

As the measure for this cylindrical scheduler is a Cauchy sequence of measures for sequences with a bounded number of discrete transitions, we can sacrifice another ε to sharpen the requirement for the scheduler to reach the goal region in time *and* with at most n_ε steps for an appropriate bound $n_\varepsilon \in \mathbb{N}$, still maintaining an ε advantage over f_{\max} . Hence, we can compare with a finite structure, and use an inductive argument to show for paths π of shrinking length that end in any location $l' \in L$ that $f_{\max}(l', t) \leq Pr_{\mathcal{S}}^{\mathcal{M}}(\pi, t)$ holds true. \square

The full proof is moved to Appendix B.2.

Theorem 1. *For a CTMDP, there is a measurable TPD scheduler \mathcal{S} optimal for maximum time-bounded reachability in the class of measurable THR scheduler.*

Proof. We construct a measurable scheduler \mathcal{S} that always chooses an action a that maximises $\sum_{l' \in L} \mathbf{R}(l, \mathcal{S}(l, t), l') \cdot (Pr_{\mathcal{S}}^{\mathcal{M}}(l', t) - Pr_{\mathcal{S}}^{\mathcal{M}}(l, t))$; by Lemma 1, this guarantees that $\sum_{l \in L} \nu(l) Pr_{\mathcal{S}}^{\mathcal{M}}(l, 0) = \sum_{l \in L} \nu(l) f_{\max}(l, 0) = \sup_{\mathcal{S} \in THR} Pr_{\mathcal{S}}^{\mathcal{M}}(t_{\max})$ holds true.

To construct the scheduler decisions for a location l for a measurable scheduler \mathcal{S} , we partition $[0, t_{\max}]$ into measurable sets $\{D_a \mid a \in Act(l)\}$, such that \mathcal{S} only makes decisions that maximise $\sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{\mathcal{S}}^{\mathcal{M}}(l', t) - Pr_{\mathcal{S}}^{\mathcal{M}}(l, t))$. (For positions outside of $[0, t_{\max}]$, the behaviour of the scheduler does not matter. $\mathcal{S}(l, t)$ can therefore be fixed to any constant decision $a \in Act(l)$ for all $l \in L$ and $t \notin [0, t_{\max}]$.)

We start with fixing an arbitrary order \succ on the actions in $Act(l)$ and introduce, for each point $t \in [0, t_{\max}]$, an order \succ_t on the actions determined by the value of $\sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\max}(l', t) - f_{\max}(l, t))$, using \succ as a tie-breaker.

Along the order of \succ , we construct, starting with the minimal element, for each action $a \in Act(l)$:

1. Open sets O_a that contain the positions where our scheduler does *not* make a decision⁵ $a' \prec a$.
(We choose $O_a = [0, t_{\max}]$ for the action a that is minimal with respect to \succ .)
2. A set T_a that is open in O_a and contains the points in time in O_a where a is maximal with respect to \succ_t .
3. A set $D_a = \overline{T_a} \cap O_a$ which is the closure of T_a in O_a .
If a is not maximal, we set $O_{a'} = O_a \setminus D_a$ for the successor a' of a with respect to \succ .

⁵ Note that, for all $a' \prec a$, the points $T_{a'}$ in time where the scheduler does make the decision a' have been fixed earlier by this construction.

To complete the proof, we have to show that the scheduler \mathcal{S} , which chooses a for all $t \in D_a$, makes only decisions that maximise the gain, that is $\sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\max}(l', t) - f_{\max}(l, t))$ (and hence that $f_{\max} = Pr_{\mathcal{S}}^M$ holds true), and we have to show that the resulting scheduler is measurable. As an important lemma on the way, we have to demonstrate the claimed openness of the O_a 's and T_a 's in the compact Euclidean space $[0, t_{\max}]$.

This openness is provided by a simple inductive argument: First, the complete space $[0, t_{\max}]$ is open in itself.

Let us assume that a is maximal w.r.t. \succ_t for a t in the open set O_a . Then the following holds: $\sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\max}(l', t) - f_{\max}(l, t))$ is *strictly* greater for a compared to the respective value of all other actions $a' \succ a$ (because \succ serves as tie-breaker), and hence this holds for some ε -environment of t that is contained in the open set O_a . For the actions $a' \prec a$ in this ε -environment of t , the respective value also cannot be *strictly* greater compared to a , because otherwise one of these actions had been selected before.

Note that this argument provides optimality of the choices in T_a as well as openness of T_a . The optimality for the choices at the fringe of T_a (and hence the extension of the optimality argument to D_a) is a consequence of the continuity of f_{\max} .

As every open and closed set in \mathbb{R} is (Lebesgue) measurable, O_a (or, to be precise, $O_a \cap (0, t_{\max})$) and $\overline{T_a}$, and hence their intersection $D_a = \overline{T_a} \cap O_a$, are measurable.

Our construction therefore provides us with a *measurable* scheduler, which is optimal, deterministic, and timed positional. \square

By simply replacing maximisation by minimisation, sup by inf, and max by min, we can rewrite the proof to yield a similar theorem for the minimisation of time-bounded reachability, or likewise, for the maximisation of time-bounded safety.

Theorem 2. *For a CTMDP, there is a measurable TPD scheduler \mathcal{S} optimal for minimum time-bounded reachability in the class of measurable THR scheduler.*

4 Finite Optimal Control

In this section we show that, once the existence of an optimal scheduler is established, we can refine this result to the existence of a *cylindrical* optimal TPD scheduler, that is, a scheduler that changes only finitely many times between different positional strategies. This is as close as we can hope to get to implementability as optimal points for policy switching are—like in the example from Figure 1—almost inevitably irrational.

Our proof of Theorem 1 makes a purely topological existence claim, and therefore does not imply that a finite number of switching points suffices. In principle, this could mean that the required switching points have one or more limit points, and an unbounded number of switches is required to optimise time-bounded reachability. $x \cdot \sin(x^{-1})$ (cf. Figure 2) is an example for a continuous




Fig. 2. The plot on the left shows the function $x \cdot \sin(x^{-1})$ as an example for the limit point problem. The plot on the right: Theorem 1 would not exclude that the intersection points of derivatives (the ‘loss’) of two positional schedulers have a limit point.

function for which the codomain of 0 has a limit point at 0, and the right curve of Figure 2 shows the derivations for a positional scheduler (black) and a potential comparison with a gain function such that their intersections have a limit point.

To exclude such limit points, and hence to prove the existence of an optimal scheduler with a finite number of switching points, we re-visit the differential equations that define the reachability probability, but this time to answer a different question: Can we use the true values in some point of time to locally find an optimal strategy for an ε -environment? If yes, then we could exploit the compactness of $[0, t_{\max}]$: We could, for all points in time $t \in [0, t_{\max}]$, fix a decision that is optimal in an ε -environment of t . This would provide an open set with a positional optimal strategy around each $t \in [0, t_{\max}]$, and hence an open coverage of a compact set, which would imply a final coverage with segments of positional optimal strategies.

While this is the case for most points, this is not necessarily the case at our switching points. In the remainder of this section, we therefore show something similar: For every point $t \in [0, t_{\max}]$ in time, there is a positional strategy that is optimal in a left ε -environment of t (that is, in a set $(t - \varepsilon, t] \cap [0, t_{\max}]$), and one that is optimal in a right ε -environment of t . Hence, we get an open coverage of strategies with at most one switching point, and thus obtain a strategy with a finite number of switching points.

Theorem 3. *For every CTMDP, there is a cylindrical TPD scheduler S optimal for maximum time-bounded reachability in the class of measurable THR scheduler.*

Proof. We have seen that the true optimal reachability probability is defined by a system of differential equations. In this proof we consider the effect of starting with the ‘correct’ values for a time $t \in [0, t_{\max}]$, but *locally fix a positional strategy* for a small left or right ε -environment of t . That is, we consider only schedulers that keep their decision constant for a (sufficiently) small time ε before or after t .

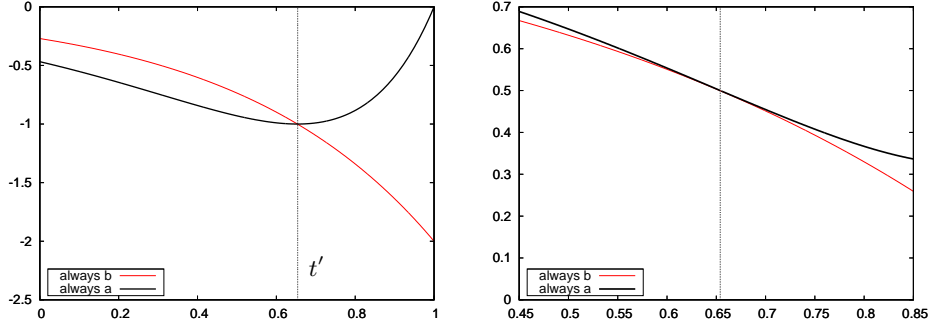


Fig. 3. The gain functions of the competing stationary strategies of Figure 1. To the left: Developed gains $\dot{Pr}_A^{S_a/S_b}(t)$ from $t = t_{\max} = 1$ in order to find the only switching point $t' = t_{\max} - \frac{1}{2} \log(2)$. To the right: Developed values $Pr_A^{S_a/S_b}(t)$ (not gains) from t' in both directions, to show that action a is better for all $t < t'$ whereas b is better for all $t > t'$. This construction is also used to determine the existence of ϵ -environments with a stable strategy around every point t .

Given a CTMDP \mathcal{M} , we consider the differential equations that describe the development near the support point $f_{\max}(l, t)$ for each location l under a positional strategy D :

$$-\dot{Pr}_l^D(\tau) = \sum_{l' \in L} \mathbf{R}(l, a_l, l') \cdot (Pr_{l'}^D(\tau) - Pr_l^D(\tau)),$$

where a_l is the action chosen at l by D (see Figure 3 for an example).

Different to the development of the true probability, the development of these linear differential equations provides us with smooth functions. This provides us with more powerful techniques when comparing two locally positional strategies: Each deterministic scheduler defines a system $\dot{y} = Ay$ of ordinary homogeneous linear differential equations with constant coefficients.

As a result, the solutions $Pr_l^D(\tau)$ of these differential equations—and hence their differences $Pr_l^{D'}(\tau) - Pr_l^D(\tau)$ —can be written as finite sums $\sum_{i=1}^n P_i(\tau)e^{\lambda_i \tau}$, where P_i is a polynomial and the λ_i may be complex. Consequently, these functions are holomorphic.

Using the identity theorem for holomorphic functions, t can only be a limit point of the set of 0 points of $Pr_l^{D'}(\tau) - Pr_l^D(\tau)$ if $Pr_l^{D'}(\tau)$ and $Pr_l^D(\tau)$ are identical on an ϵ -environment of t . The same applies to their derivations: $\dot{Pr}_l^{D'}(\tau) - \dot{Pr}_l^D(\tau)$ either has no limit point in t , or $\dot{Pr}_l^{D'}(\tau)$ and $\dot{Pr}_l^D(\tau)$ are identical on an ϵ -environment of t .

For the remainder of the proof, we fix, for a given time t , a sufficiently small $\epsilon > 0$ such that, for each pair of schedulers D and D' and every location $l \in L$, $\dot{Pr}_l^{D'}(\tau) - \dot{Pr}_l^D(\tau)$ is either < 0 , $= 0$, or > 0 on the complete interval $L_\epsilon^t = (t - \epsilon, t) \cap [0, t_{\max}] \ni \tau$, and, possibly with different sign, for the complete interval $R_\epsilon^t = (t, t + \epsilon) \cap [0, t_{\max}] \ni \tau$.

We argue the case for the left ε -environment L_ε^t . In the ‘>’ case for a location l , we say that D is *l-better* than D' . We call D *preferable* over D' if D' is not *l-better* than D for any location l , and *better* than D' if D is preferable over D' and *l-better* for some $l \in L$.

If D' is *l-better* than D in exactly a non-empty set $L_b \subset L$ of locations, then we can obviously use D' to construct a strategy D'' that is better than D by switching to the strategies of D' in exactly the locations L_b .

Since we choose our strategies from a finite domain—the deterministic positional schedulers—this can happen only finitely many times. Hence we can stepwise *strictly* improve a strategy, until we have constructed a strategy D_{\max} preferable over all others.

By the definition of being preferable over all other strategies, D_{\max} satisfies

$$-Pr_l^{D_{\max}}(\tau) = \max_{a \in Act(l)} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{l'}^{D_{\max}}(\tau) - Pr_l^{D_{\max}}(\tau))$$

for all $\tau \in L_\varepsilon^t$ and all $l \in L$.

We can use the same method for the right ε -environment R_ε^t , and pick the decision for t arbitrarily; we use the decision from the respective left ε environment.

Now we have fixed, for an ε -environment of an arbitrary $t \in [0, t_{\max}]$, an optimal scheduler with at most one switching point. As this is possible for all points in $[0, t_{\max}]$, the sets $I_\varepsilon^t = L_\varepsilon^t \cup R_\varepsilon^t$ define an open cover of $[0, t_{\max}]$. Using the compactness of $[0, t_{\max}]$, we infer a finite sub-cover, which establishes the existence of a strategy with a finite number of switching points. \square

The proof for the minimisation of time-bounded reachability (or maximisation of time-bounded safety) runs accordingly.

Theorem 4. *For every CTMDP, there is a cylindrical TPD scheduler \mathcal{S} optimal for minimal time-bounded reachability in the class of measurable THR scheduler.*

5 Discrete Locations

In this section, we treat the mildly more general case of single player CTMGs, which are traditional CTMDPs plus discrete locations. We reduce the problem of finding optimal measurable schedulers for CTMGs first to *simple CTMGs*, CTMGs whose discrete locations have no incoming transitions from continuous locations. (They hence can only occur initially at time 0.) The extension from CTMDPs to simple CTMGs is trivial.

Lemma 2. *For a simple single player CTMG with only a reachability (or only a safety) player, there is an optimal deterministic scheduler with finitely many switching points.*

Proof. By the definition of simple single player games, the likelihood of reaching the goal location from any continuous location and any point in time is independent of the discrete locations and their transitions. For continuous locations, we can therefore simply reuse the results from the Theorems 3 and 4.

We can only be in discrete locations at time 0, and for every continuous location l there is a fixed time-bounded reachability probability described by $f_{\text{opt}}(l, 0)$. We can show that there is a timed-positional (even a positional) deterministic optimal choice for the discrete locations at time $t = 0$ by induction over the maximal distance to continuous locations: If all successors have been evaluated, we can fix an optimal timed-positional choice. We can therefore use discrete positions with maximal distance 1 as induction basis, and then apply an induction step from positions with distance $\leq n$ to positions with distance $n + 1$. \square

Rebuilding a single player CTMG \mathcal{G} to a *simple* single player CTMG \mathcal{G}_s can be done in a straight forward manner; it suffices to pool all transitions taken between two continuous locations. To construct the resulting simple CTMG \mathcal{G}_s , we add new continuous locations for each possible time abstract path from continuous locations of the CTMG \mathcal{G} , and we add the respective actions: For continuous locations $l_c, l'_c \in L_c$ and discrete locations $l_1^d, \dots, l_n^d \in L_d$ a timed path $l_c \xrightarrow{a_0, t} l_1^d \xrightarrow{a_1, t} l_2^d \dots l_n^d \xrightarrow{a_n, t} l'_c$ translates to $l_c \xrightarrow{\mathbf{a}, t} \underline{l_1^d \xrightarrow{a_0} l_1^d \xrightarrow{a_1} l_2^d \dots l_n^d \xrightarrow{a_n} l'_c}$, where the underlined part is a new continuous location. (For simplicity, we also translate a timed path $l_c \xrightarrow{a, t} l'_c$ to $l_c \xrightarrow{\mathbf{a}, t} \underline{l'_c}$.)

The new *actions* of the resulting simple single player CTMG encode the sequences of actions of \mathcal{G} that a scheduler could make in the current location plus in all possible sequences of discrete locations, until the next continuous location is reached. (Note that this set is finite, and that the scheduler makes all of these transitions at the same point of time.) If \mathbf{a} encodes choices that depend only on the position (but not on this local history), \mathbf{a} is called positional. For continuous locations, all old actions are deleted, and all new continuous locations that end in a location $l_c \in L_c$ get the same outgoing transitions as l_c . The rate matrix is chosen accordingly.

Adding the information about the path to locations allows to reconstruct the timed history in the single player CTMG from a history in the constructed simple CTMG.

Theorem 5. *For a single player CTMG \mathcal{G} with only a reachability (or only a safety) player, there is an optimal deterministic scheduler with finitely many switching points.*

Proof. First, every scheduler \mathcal{S} for \mathcal{G} can be naturally translated into a scheduler of \mathcal{S}_s of \mathcal{G}_s , because every timed-path in \mathcal{G}_s defines a timed-path in \mathcal{G} ; the resulting time-bounded reachability probability coincides.

Let us consider a cylindrical optimal deterministic scheduler \mathcal{S}_{opt} for the simple Markov game, and the function f_{opt} defined by it. For the actions \mathbf{a} \mathcal{S}_{opt} chooses, we can, for each interval in which \mathcal{S}_{opt} is positional, use an inductive

argument similar to the one from the proof of Lemma 2 to show that we can choose a *positional* action \mathbf{a}' instead. The resulting cylindrical deterministic scheduler $\mathcal{S}'_{\text{opt}}$ defines the same f_{opt} (same differential equations).

Clearly, $f_{\text{opt}}(l_c, t) = f_{\text{opt}}(\dots \xrightarrow{a} l_c, t)$ holds true. We use this observation to change $\mathcal{S}'_{\text{opt}}$ to $\mathcal{S}''_{\text{opt}}$ by choosing the action that $\mathcal{S}'_{\text{opt}}$ chooses for l_c for all locations $\dots \xrightarrow{a} l_c$ and at each point of time. The resulting scheduler $\mathcal{S}''_{\text{opt}}$ is still cylindrical and deterministic, and defines the same f_{opt} (same differential equations).

$\mathcal{S}''_{\text{opt}}$ is also the mapping of a cylindrical optimal deterministic scheduler for \mathcal{G} . \square

6 Continuous-Time Markov Games

In this section, we lift our results from single player to general continuous-time Markov games. In general continuous-time Markov games, we are faced with two players with opposing objectives: A reachability player trying to maximise the time-bounded reachability probability, and a safety player trying to minimise it—we consider a 0-sum game.

Thus, all we need to do for lifting our results to games is to show that the quest for optimal strategies for single player games discussed in the previous section can be generalised to a quest for co-optimal strategies—that is, for Nash equilibria—in general games. To demonstrate this, it essentially suffices to show that it is not important whether we first fix the strategy for the reachability player and then the one for the safety player in a strategy refinement loop, or vice versa.

Let us first assume CTMGs without discrete locations.

Lemma 3. *Using the ε -environments I_ε^t from the proof of Theorem 3, we can construct a Nash equilibrium that provides co-optimal deterministic strategies for both players, such that the co-optimal strategies contain at most one strategy switch on I_ε^t .*

Proof. We describe the technique to find a constant co-optimal strategy on the right ε -environment $R_\varepsilon^t = (t, t + \varepsilon) \cap [0, t_{\max}]$ of t .

We write a constant strategy as $D = S + R$ that is composed of the actions chosen by the safety player on L_s , and the actions chosen by the reachability player on L_r . For this simple structure, we introduce a strategy improvement technique on the finite domain of deterministic choices for the respective player.

For a fixed strategy S of the safety player, we can find an optimal counter strategy $R(S)$ of the reachability player by applying the technique described in Theorem 3. (For equivalent strategies, we make an arbitrary but fixed choice.)

We call the resulting vector $(-Pr_{S+R(S)}^{\mathcal{M}}(l, t + \frac{1}{2}\varepsilon) \mid l \in L)$ the *quality vector* of S . Now, we choose an arbitrary \bar{S} for which this vector is minimal. (Note that there could, potentially, be multiple incomparable minimal elements.)

We now show that the following holds for \bar{S} and all $\tau \in R_\varepsilon^t$:

$$-Pr_{\bar{S}+R(\bar{S})}^{\mathcal{M}}(l, \tau) = \max_{a \in Act(l)} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{\bar{S}+R(\bar{S})}^{\mathcal{M}}(l', \tau) - Pr_{\bar{S}+R(\bar{S})}^{\mathcal{M}}(l, \tau))$$

for all $l \in L_r$, and

$$-\dot{P}r_{\overline{S}+R(\overline{S})}^{\mathcal{M}}(l, \tau) = \min_{a \in \text{Act}(l)} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{\overline{S}+R(\overline{S})}^{\mathcal{M}}(l', \tau) - Pr_{\overline{S}+R(\overline{S})}^{\mathcal{M}}(l, \tau))$$

for all $l \in L_s$. (Note that the order between the derivation is maintained on the complete right ε -environment R_ε^t .)

The first of these claims is a trivial consequence from the proof of Theorem 3. (The result is, for example, the same if we had a single player CTMDP that, in the locations L_s of the safety player, has only one possible action: the one chosen by \overline{S} .)

Let us assume that the second claim does not hold. Then we choose a particular $l \in L_s$ where it is violated. Let us consider a slightly changed setting, in which the choices in l are restricted to two actions, the action a_1 chosen by \overline{S} , and the minimising action a_2 . Among these two, one maximises, and one minimises

$$-\dot{P}r_{\overline{S}+R(\overline{S})}^{\mathcal{M}}(l, \tau) = \min_{a \in \{a_1, a_2\}} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{\overline{S}+R(\overline{S})}^{\mathcal{M}}(l', \tau) - Pr_{\overline{S}+R(\overline{S})}^{\mathcal{M}}(l, \tau)).$$

Let us fix all other choices of S , and allow the reachability player to choose among a_1 and a_2 (we ‘pass control’ to the other player). As shown in Theorem 3, she will select an action that produces the well defined set of max equations for the resulting single player game. Hence, choosing a_1 and keeping all other choices from $R(\overline{S})$ is the optimal choice for the reachability player in this setting (as the max equations are satisfied, while they are dissatisfied for a_2).

Consequently, the quality vector for \overline{S} is strictly greater than the one for the adjusted strategy. That is, assuming that choosing an arbitrary maximal element does not lead to a satisfaction of the min and max equations leads to a contradiction.

We can argue symmetrically for the left ε -environment. Note that the satisfaction of the min and max equations implies that it does not matter if we change the rôle of the safety and reachability player in our argumentation. \square

This lemma can easily be extended to construct simple co-optimal strategies:

Theorem 6. *For CTMGs without discrete locations, there are cylindrical deterministic timed-positional co-optimal strategies for the reachability and the safety player.*

Proof. First, Lemma 3 provides us with an open coverage of co-optimal strategies that switch at most once, and we can build a strategy that switches at most finitely many times from a finite sub-cover of the open space $[0, t_{\max}]$. This strategy is everywhere locally co-optimal, and forms a Nash equilibrium:

It is straight forward to cut the interval $[0, t_{\max}]$ into a finite set of sub-intervals $[0, t_0], (t_0, t_1], \dots, (t_{n-1}, t_n]$ with $t_n = t_{\max}$, such that the strategy for the safety player is constant in all of these intervals. We can use the construction from Theorem 3 (note that the proof of Theorem 3 does not use that the

differential equations are initialised to 0 or 1 at t_{\max}) to construct an optimal strategy for the reachability player: We can first solve the problem for the interval $[t_{n-1}, t_n]$, then for the interval $[t_{n-2}, t_{n-1}]$ using $f_{\text{opt}}(l, t_{n-1})$ as initialisation, and so forth. A similar argument can be made for the other player.

This provides us with the same differential equations, namely:

$$-\dot{f}_{\text{opt}}(l, t) = \max_{a \in \text{Act}(l)} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\text{opt}}(l', t) - f_{\text{opt}}(l, t))$$

for $t \in [0, t_{\max}]$ and $l \in L_r$, and

$$-\dot{f}_{\text{opt}}(l, t) = \min_{a \in \text{Act}(l)} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\text{opt}}(l', t) - f_{\text{opt}}(l, t))$$

for $t \in [0, t_{\max}]$ and $l \in L_s$.

Note that all Nash equilibria need to satisfy these equations (with the exception of 0 sets, of course), because otherwise one of the players could improve her strategy. \square

The extension of these results to the full class of CTMGs is straight forward: We would first reprove Theorem 5 in the style of the proof of Theorem 3 (which requires to establish the Theorem in the first place). The only extension is that we additionally get an equation $Pr_l^D(\tau) = \sum_{l' \in L} \mathbf{P}(l, a_l, l') \cdot Pr_{l'}^D(\tau)$ for every discrete location l . The details are moved to Appendix C.

Theorem 7. *For continuous-time Markov Games, there are cylindrical deterministic timed-positional co-optimal strategies for the reachability and the safety player.*

As a small side result, these differential equations show us that we can, for each continuous location $l_c \in L_c$ and every action $a \in \text{Act}(l_c)$, add arbitrary values to $\mathbf{R}(l_c, a, l_c)$ without changing the bounded reachability probability for every pair of schedulers. (Only if we change $\mathbf{R}(l_c, a, l_c)$ to 0 we have to make sure that a is not removed from $\text{Act}(l_c)$.) In particular, this implies that we can locally and globally uniformise a continuous-time Markov game if this eases its computational analysis. (Cf. [10] for the simpler case of CTMDPs.)

7 Variances

In this section, we discuss the impact of small changes in the setting, namely the impact of infinitely many states or actions, and the impact of introducing a non-absorbing goal region.

Infinitely Many States. If we allow for infinitely many states, optimal solutions may require infinitely many switching points. To see this, it suffices to use one copy of the CTMDP from Figure 1, but with rates i and $2i$ for the i -th copy, and assign an initial probability distribution that assigns a weight of 2^{-i} to the




Fig. 4. An example CTMDP with infinitely many actions.

initial state A_i of the i -th copy. (If one prefers to consider only systems with bounded rates, one can choose rates $1 + \frac{1}{i}$ and $2 + \frac{2}{i}$.) The switching points are then different for every copy, and an optimal strategy has to select the correct switching point for every copy.

Infinitely Many Actions. If we allow for infinitely many actions, there is not even an optimal strategy if we restrict our focus to CTMDPs with two locations, an initial location and an absorbing goal location. For the CTMDP of Figure 4 with the natural numbers \mathbb{N} as actions and rate $\lambda_i = 2 - \frac{1}{i}$ for the action $i \in \mathbb{N}$ if we have a reachability player and $\lambda_i = \frac{1}{i}$ if we have a safety player, every strategy \mathcal{S} can be improved over by a strategy \mathcal{S}' that always chooses the successor $i + 1$ when of the action i chosen by \mathcal{S} .

Reachability at t_{\max} . If we drop the assumption that the goal region is absorbing, one might be interested in the marginally more general problem to be (not to be) in the goal region at time t_{\max} for the reachability player (safety player, respectively). For this generalisation, no substantial changes need to be made: It suffices to replace

$$f_{\text{opt}}(l, t) = Pr_{\mathcal{S}}^{\mathcal{M}}(l, t) = 1 \quad \text{for all goal locations } l \in G \text{ and all } t \leq t_{\max}$$

by

$$f_{\text{opt}}(l, t_{\max}) = Pr_{\mathcal{S}}^{\mathcal{M}}(l, t_{\max}) = 1 \quad \text{for all goal locations } l \in G.$$

(In order to be flexible with respect to this condition, the $-f_{\text{opt}}(l, t)$ are defined for goal locations as well. Note that, when all goal locations are absorbing, the value of $-f_{\text{opt}}(l, t)$ is 0 and $f_{\text{opt}}(l, t)$ is 1 for all goal locations $l \in G$ and all $t \in [0, t_{\max}]$.)

References

1. Adnan Aziz, Kumud Sanwal, Vigyan Singhal, and Robert Brayton. Model-checking continuous-time Markov chains. *Transactions on Computational Logic*, 1(1):162–170, 2000.
2. C. Baier, J.-P. Katoen, and H. Hermanns. Approximate Symbolic Model Checking of Continuous-Time Markov Chains. In *Proceedings of CONCUR'99*, volume 1664 of *Lecture Notes in Computer Science*, pages 146–161, 1999.

3. Christel Baier, Holger Hermanns, Joost-Pieter Katoen, and Boudewijn R. Haverkort. Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes. *Theoretical Computer Science*, 345(1):2–26, 2005.
4. Richard Bellman. *Dynamic Programming*. Princeton University Press, 1957.
5. Tomas Brazdil, Vojtech Forejt, Jan Krčal, Jan Kretinsky, and Antonin Kucera. Continuous-time stochastic games with time-bounded reachability. In *Proceedings of FSTTCS'09*, Leibniz International Proceedings in Informatics (LIPIcs), pages 61–72, 2009.
6. Eugene A. Feinberg. Continuous Time Discounted Jump Markov Decision Processes: A Discrete-Event Approach. *Mathematics of Operations Research*, 29(3):492–524, 2004.
7. H. Hermanns. *Interactive Markov Chains and the Quest for Quantified Quality. LNCS 2428*. Springer-Verlag, 2002.
8. Prasadarao Kakumanu. Continuously Discounted Markov Decision Model with Countable State and Action Space. *The Annals of Mathematical Statistics*, 42(3):919–926, 1971.
9. M. A. Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. Modelling with Generalized Stochastic Petri Nets. *SIGMETRICS Performance Evaluation Review*, 26(2):2, 1998.
10. Martin R. Neuhäüßer, Mariëlle Stoelinga, and Joost-Pieter Katoen. Delayed Non-determinism in Continuous-Time Markov Decision Processes. In *Proceedings of FOSSACS '09*, pages 364–379, 2009.
11. Martin R. Neuhäüßer and Lijun Zhang. Time-Bounded Reachability in Continuous-Time Markov Decision Processes. Technical report, 2009.
12. Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience, April 1994.
13. Markus Rabe and Sven Schewe. Optimal Time-Abstract Schedulers for CTMDPs and Markov Games. In *Proceedings of QAPL'10 (accepted)*, 2010.
14. William H. Sanders and John F. Meyer. Reduced Base Model Construction Methods for Stochastic Activity Networks. In *Proceedings of PNPM'89*, pages 74–84, 1989.
15. Nicolás Wolovick and Sven Jöhr. A Characterization of Meaningful Schedulers for Continuous-Time Markov Decision Processes. In *Proceedings of FORMATS'06*, pages 352–367, 2006.
16. L. Zhang, H. Hermanns, E. M. Hahn, and B. Wachter. Time-bounded model checking of infinite-state continuous-time Markov chains. In *Proceedings of ACSD'08*, pages 98–107, 2008.

Appendix

As to be expected by the topic, the paper is based in large parts on measure theory. While the techniques are standard and straight forward for the experts in the field we provide a short introduction to the ideas exploited; while we do not use any technique beyond the standard curriculum of a math degree, we assume that some recap of the ideas behind the completion of metric spaces and its application in our measures in Section A. However, Section A is but a short introduction to the ideas, and cannot serve as a self contained introductory to the techniques.

Section B contains a short recap on the differential equations that describe f_{\min} and f_{\max} , and, more generally, the development of $Pr_S^M(l, t)$ in Subsection B.1, and a proof that f_{\max} truly establishes an upper bound on the performance of any measurable scheduler in Subsection B.2, which constitutes a proof of Lemma 1.

A Completion of Metric Spaces

A metric space is called complete if every Cauchy sequence in it converges. A Cauchy sequence in a metric space (M, d) is a sequence $s : \mathbb{N} \rightarrow M$ such that the following holds:

$$\forall \varepsilon > 0 \exists n \in \mathbb{N} \forall l, m > n. d(s(l), s(m)) < \varepsilon$$

Intuitively one could say that a Cauchy sequence converges, but not necessarily to a point within the space. For example, a sequence of rational numbers that converges to $\sqrt{2}$ is a converging sequence in the real numbers, but not in the rationals—as the limit point is outside of the carrier set—but it is still a Cauchy sequence.

The basic technique to complete an incomplete metric space (M, d) is to use the Cauchy sequences of this space as the new carrier set M' , and define a distance function d' between two Cauchy sequences $s, s' \in M'$ of M to be $d'(s, s') = \lim_{n \rightarrow \infty} d((s(n), s'(n)))$. Now, (M', d') is not yet a metric space, because two different Cauchy sequences—for example the constant 0 sequence and the sequence $s(n) = \frac{1}{2^n}$ of the rationals—can have distance 0.

Technically, one therefore defines equivalence classes of Cauchy sequences that have distance 0 with respect to d' as the new carrier set M'' of a metric space (M'', d'') , where the distance function d'' is defined by using d' on representatives of the quotient classes of Cauchy sequences. This also complies with the intuition: a Cauchy sequence in M is meant to represent its limit point (which is not necessarily in M), and hence two Cauchy sequences with the same limit point should be identified.

The resulting metric space (M'', d'') is complete by construction. The simplest example of such a completion is the completion of the rational numbers

into the real numbers. And on this level, a straight forward effect of completion can be easily explained: To end up with (M'', d'') (or a space isomorphic to it), we can start with any **dense** subset S of M'' .

A subset $S \subset M''$ is dense in M'' if, for every point $m \in M''$ and every $\varepsilon > 0$, there is a point $s \in S$ with $d''(m, s) < \varepsilon$. Looking at the definition, one immediately sees the connection to Cauchy sequences: One could intuitively say that $S \subset M''$ is dense in M'' if, for every point $m \in M''$, there is a Cauchy sequence s with limit point m .

Hence, it does not matter which dense set we use as a starting point. Of course, it works to use M'' , in the example of the real numbers, we can start with the real numbers themselves without gaining anything by applying the completion twice, we can start with the transcendent numbers, the non-transcendent numbers, or, more down to earth, with finite decimal fractions. Note that a subset is dense in M'' if it is dense in some S that is dense in M'' .

A.1 Application in Measure Theory

Another famous application of this completion technique is the completion of Riemann integrable functions to Lebesgue integrable functions. The difference metrics between two functions is the Riemann integral over the absolute value of their difference. (Strictly speaking, this does again not form a metric space, and we again have to use the quotient class of functions with difference 0.)

Riemann integrable functions do, for example, allow only for bounded functions, but there are other problems as well; for example, we cannot integrate over the characteristic function of the rational numbers. (The wikipedia article to Riemann integrable functions is nice and gives a good overview on the weaknesses.)

Using the completion technique defined above, one can, for example, integrate over the characteristic function of the rationals by enumerating them, that is, by defining a surjection $s : \mathbb{N} \rightarrow \mathbb{Q}$, and choose f_i to be the Riemann integrable function that is 1 at the mapping $s(\{1, 2, \dots, i\})$ of the initial sequence of length i of the naturals. Clearly, the limit of the sequence f_1, f_2, f_3, \dots is the characteristic function, the Riemann integral over all f_i is 0 (no matter over which interval we integrate) the sequence is a Cauchy sequence.

The completion of the space of Riemann integrable functions (which essentially establishes the Lebesgue integrable functions) is space of all Cauchy sequences of Riemann integrable functions (or again representatives of the quotient classes of equivalent Cauchy sequences).

Again, it does not make a difference if we start the completion with the Riemann integrable functions, or with weaker concepts, as long as they are dense in the resulting space, or indeed in the Riemann integrable functions. A well known example for such a class is the class of block functions, where the value changes only in finitely many positions.

A.2 Application in Our Measure

The definition of the measures for continuous-time Markov chains, games (with fixed strategies), and decision processes (with a fixed scheduler) works in exactly this way: For Markov chains, one defines the probability measure for simple disjoint (or almost disjoint) sets, for which the probability is simple to determine, and such that one can define the probability of reaching the goal region in time for cylindrical sets.

When considering games and decision processes without fixed scheduling policies, however, we have two layers of completions—one layer *for a given* cylindrical schedulers, and one *on* cylindrical schedulers.

Measure for a given cylindrical schedulers. In the case of games and decision processes, one starts to define it for a particularly friendly and easy to handle class of strategies or schedulers, respectively, such that the techniques from Markov chains can be extended with minor adjustments. Our cylindrical schedulers—which in our paper also represent pairs of strategies in games—with only finitely many switching points are an example of such an extension.

The measure *for* a given cylindrical scheduler is a mild extension of the techniques for Markov chains. The building blocks of the probability measure define the probability on *cylindrical sets*, and they are a straight forward extension of the probability measure for continuous time Markov chains to continuous time Markov decision processes (and games) with a fixed scheduler of this type. However, they only describe the likelihood for cylindrical sets, and without completing the space we can but use finite sums over disjoint sets.

To obtain the time bounded reachability probability, we use Cauchy sequences of such sets that converge against all sets of paths on which a goal region is reached. A representative of this equivalence class would be a sequence P_1, P_2, P_3, \dots or sets of paths, where P_i contains the cylindrical sets of length up to i in which the goal region is reached. This is a Cauchy sequence (cf. the argument in Subsection B.2), and the limit contains all finite paths upon which the goal region is reached.

Measure for a given cylindrical schedulers. While we have established a measure for a given cylindrical scheduler, the class of cylindrical schedulers is not particularly strong, and, like with Riemann integrable functions, we need to strengthen the class of schedulers we allow for. In order to exploit the aforementioned completion technique, we have to create a suitable metric space on them, and in order to introduce such a metric space, we need a measure for the difference between strategies. Such a measure reflects the likelihood that two different strategies ever lead to different actions. For example, for *deterministic* schedulers D and E for a CTMDP \mathcal{M} we define a difference scheduler $\delta_{\{D,E\}}$ that uses the actions of D/E on every history, on which they coincide, and a fresh action (a_D, a_E) if D chose a_D and E chose $a_E \neq a_D$ upon this history. (Note that the difference measure uses a slightly adjusted CTMDP \mathcal{M}' ; it also has a fresh goal location.)

The new action (a_D, a_E) leads to a fresh continuous goal location g ; the old goal locations do not remain goal locations, the new goal region contains only g .

For a continuous location l , we fix

- $\mathbf{R}(l, (a_D, a_E), g) = \mathbf{R}(l, a_D, L) + \mathbf{R}(l, a_E, L)$ and
- $\mathbf{R}(l, (a_D, a_E), l') = 0$ for all locations $l' \neq g$

for the new actions, and maintain the entries to in the rate matrix for the old actions. (To be formally correct, we provide g with only one enabled action a —with $\mathbf{R}(g, a, l') = 0$ for all locations $l' \neq g$ —that $\delta_{\{D, E\}}$ selects in g upon any history.)

For discrete locations, we would fix $\mathbf{P}(l, (a_D, a_E), g) = 1$ (and $\mathbf{P}(l, (a_D, a_E), l') = 0$ for $l' \neq g$) for all new actions, and maintain the entries for the old actions in \mathbf{P} .

The distance between two schedulers is then defined as the likelihood that g is reached in the adjusted CTMDP within the given time bound when using the *cylindrical* scheduler $\delta_{\{D, E\}}$.

The metric space on cylindrical schedulers. This distance function almost defines a metric space on the cylindrical schedulers we used as a basic building block; symmetry, triangle inequation and non-negativity obviously hold. However, we again have to resort to a carrier set of quotient classes of schedulers with distance 0, in order to satisfy $d(x, y) \Leftrightarrow x = y$. (One can think of different actions on unreachable paths.)

Effect of the metric. The distance is defined in a way that guarantees that the absolute value of the difference between the time bounded reachability probabilities for D and E is bounded by the distance between these schedulers (or representatives of their quotient class): To see this, it suffices to look at the scheduler $\delta_{\{D, E\}}$ in \mathcal{M} with two adjusted goal regions:

1. If we use the old goal region plus g as our new goal region, then obviously the time bounded reachability is better than the time bounded reachability of D and E in \mathcal{M} .
2. If we use the old goal region (but not plus g — g rather becomes a non-accepting sink—as a new goal region, then the time bounded reachability is worse than the time bounded reachability of D and E in \mathcal{M} .

The difference between the two cases, however, is exactly the *distance* between D and E , which therefore in particular is an upper bound on the difference between their time bounded reachability probability.

Completion. The time bounded reachability of a Cauchy sequence can therefore be defined as the limit of the time bounded reachability of the elements of the sequence, which is guaranteed to exist (and to be unique) by the previous argument.

Extension to randomised schedulers. If the schedulers D and E from above are randomised, then we would define the distribution chosen by $\delta_{\{D,E\}}$ in line with the definition from above. We first fix an arbitrary global order on all actions used as tie breaker in our construction.

- If, in a particular history that ends in some location l , D chooses an action a with probability p_D^a and E chooses a with probability p_E^a then $\delta_{\{D,E\}}$ chooses a with probability $\min\{p_D, p_E\}$
- If the probabilities assigned by this rule sum up to one for $\delta_{\{D,E\}}$ on this history is thus defined.
- Otherwise, if l is continuous, we choose an action a_D among the actions with $p_D^a > p_E^a$ that maximises $\mathbf{R}(l, a, L)$, using the fixed global order as a tie breaker. For discrete locations, we use the tie breaker only.
- We then accordingly choose an action a_E among the actions with $p_E^a > p_D^a$ that maximises $\mathbf{R}(l, a, L)$, using the fixed global order as a tie breaker. For discrete locations, we again use the tie breaker only.
- We assign the remaining probability weight to the decision (a_D, a_E) .

All arguments from above extend to this case.

Measurable schedulers. The term *measurable scheduler* refers to the limit scheduler of such a Cauchy sequence. To restrict the attention to this class of schedulers is simply a requirement caused by the definition of the measure: Only for such schedulers the time bounded reachability probability is defined.

However, we do not think this is a real drawback, as the set of measurable schedulers far outreaches the power of everything one might have in mind when talking about schedulers (like choosing a on the points in time defined by the Cantor set), and its not easy to describe a non-measurable scheduler in the first place.

B Optimal reachability probability

B.1 Differential Equations

The differential equations defining f_{opt} are simply the differential equations in place when a strategy is locally constant. This holds almost everywhere (everywhere but in a 0-set of positions) in case of the *cylindrical* schedulers that are the basic building blocks in the incomplete space that we have completed by considering Cauchy sequences of cylindrical schedulers.

Hence, for every cylindrical scheduler we can partition the interval $[0, t_{\max}]$ into a finite set of intervals $I_0, I_1, I_2, \dots, I_n$ as described in the preliminaries.

Within such an interval,

$$-\dot{Pr}_{\mathcal{S}}(\pi, t) = \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot \left(Pr_{\mathcal{S}}(\pi \xrightarrow{a,t} l', t) - Pr_{\mathcal{S}}(\pi, t) \right) \quad \text{for } t \in I_i$$

holds for discrete schedulers, where π is a timed path that ends in l , a is the deterministic choice the scheduler makes in I_i on this history, and $\pi \xrightarrow{a,t} l'$ is its extension. For randomised schedulers,

$$-Pr_{\mathcal{S}}(\pi, t) = \sum_{a \in Act} h(a) \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot \left(Pr_{\mathcal{S}}(\pi \xrightarrow{a,t} l', t) - Pr_{\mathcal{S}}(\pi, t) \right) \quad \text{for } t \in I_i$$

holds, where π is a timed path that ends in l , $h(a)$ is the likelihood that the cylindrical scheduler makes the decision a in I_i on this history, and $\pi \xrightarrow{a,t} l'$ is its extension.

To initialise the potentially infinite set of differential equations, we have the following initialisations:

- $Pr_{\mathcal{S}}(\pi, t) = 1$ holds for all timed histories π that contain (and hence end up in) locations $l \in G$ in the goal region and all $t \leq t_{\max}$,
- $Pr_{\mathcal{S}}(\pi, t_{\max}) = 0$ holds for all timed histories π that contain only non-goal locations $l \notin G$, and
- $Pr_{\mathcal{S}}(\pi, t) = 0$ holds for all locations $l \in L$ and all $t > t_{\max}$.

Additionally, we have to consider what happens at the intersection t_i of the fringes of I_i and I_{i+1} for $0 \leq i < n$. But obviously, we can simply first solve the differential equations for I_n , then use the values of $f(\pi, t_{n-1})$ as initialisations for the interval I_{n-1} , and so forth.

Remark: For timed positional deterministic schedulers we get

$$-Pr_{\mathcal{S}}(l, t) = \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{\mathcal{S}}(l', t) - Pr_{\mathcal{S}}(l, t)) \quad \text{for } t \in I_i, \text{ and}$$

$$-Pr_{\mathcal{S}}(l, t) = \sum_{a \in Act} h(a) \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{\mathcal{S}}(l', t) - f(l, t)) \quad \text{for } t \in I_i$$

for timed positional randomised schedulers. In both cases, the initialisation reads

- $Pr_{\mathcal{S}}(l, t) = 1$ holds for all goal locations $l \in G$ and all $t \leq t_{\max}$,
- $Pr_{\mathcal{S}}(l, t_{\max}) = 0$ holds for all non-goal locations $l \notin G$, and
- $Pr_{\mathcal{S}}(l, t) = 0$ holds for all locations $l \in L$ and all $t > t_{\max}$.

Obviously, these differential equations can also be used in the limit.

B.2 Timed positional schedulers suffice for optimal time-bounded reachability

In this subsection we sketch a proof of a variant of Lemma 1; we demonstrate the following claim for arbitrary $t_{\max} \geq 0$:

Lemma 4. *For a CTMDP \mathcal{M} with only continuous locations, $Pr_{\mathcal{S}}^{\mathcal{M}}(l, t) \leq f_{\max}(l, t)$ holds for every scheduler \mathcal{S} , every location l , and every $t \in [0, t_{\max}]$.*

In the proof, we assume a scheduler that provides an 3ε better result, and then sacrifice one ε to transfer to cylindrical schedulers (going back to the simpler incomplete space of cylindrical schedulers, but with completed reachability measure), and then sacrificing a second ε to discard long histories from consideration (going back to the simple space of finite sums over cylindrical sets).

As a result, we can do the comparison in a simple finite structure.

Proof. Let us assume that the claim is incorrect. Then, there is a CTMDP \mathcal{M} with location l_0 and a scheduler $\mathcal{S}_{3\varepsilon}$ for \mathcal{M} such that the time bounded reachability probability is at least 3ε higher for some $\varepsilon > 0$ and $t_0 \in [0, t_{\max}]$. That is, $Pr_{\mathcal{S}_{3\varepsilon}}^{\mathcal{M}}(l_0, t_0) - f_{\max}(l_0, t_0) > 3\varepsilon$

Let us fix appropriate \mathcal{M} , l_0 , and $\mathcal{S}_{3\varepsilon}$. (Note that $\mathcal{S}_{3\varepsilon}$ does not have to be timed positional or deterministic.)

Recall that $\mathcal{S}_{3\varepsilon}$ is the limit point of a Cauchy sequence of cylindrical schedulers. Hence, almost all of these cylindrical schedulers have distance $< \varepsilon$ to $\mathcal{S}_{3\varepsilon}$.

Let us fix such a cylindrical scheduler $\mathcal{S}_{2\varepsilon}$ with distance $< \varepsilon$ to $\mathcal{S}_{3\varepsilon}$. The time bounded reachability probability of $\mathcal{S}_{2\varepsilon}$ is still at least 2ε higher compared to f_{\max} . That is, $Pr_{\mathcal{S}_{2\varepsilon}}^{\mathcal{M}}(l_0, t_0) - f_{\max}(l_0, t_0) > 2\varepsilon$ holds true.

For $\mathcal{S}_{2\varepsilon}$, we now consider a tightened form of time bounded reachability, where we additionally require that the goal region is to be reached within n_ε steps. We choose n_ε big enough that the likelihood of seeing more than n_ε discrete events is less than ε . We call this time bounded n_ε reachability.

Remark: We can estimate n_ε by taking the maximal transition rate $\lambda_{\max} = \max\{\mathbf{R}(l, a, L) \mid l \in L_c, a \in Act\}$, and choose n_ε big enough that the likelihood of having more than n_ε transitions was smaller than ε even if all transitions had transition rate λ_{\max} . As the number of steps is Poisson distributed in this case, a suitable n_ε is easy to find.

The adjustment to time bounded n_ε reachability leads to a small change in the initialisation of the differential equations: For timed histories π of length $> n_\varepsilon$ that do not contain a location $l \in G$ in the goal region within the first n_ε steps, we use $f(\pi, t) = 0$ (even if it contains a goal region after more than n_ε steps) for all $t \in [0, t_{\max}]$. As the probability measure of all timed histories π of length $> n_\varepsilon$ is $< \varepsilon$, time bounded n_ε reachability for $\mathcal{S}_{2\varepsilon}$ is still at least ε higher than the value for f_{\max} .

Let us use $f(\pi, t)$ to express the time bounded n_ε reachability for $\mathcal{S}_{2\varepsilon}$ on a path π at time t . Then this claim can be phrased as $f(l_0, t_0) - f_{\max}(l_0, t_0) > \varepsilon$.

We have now reached a finite structure, and can easily show that this leads to a contradiction: We provide an inductive argument which even demonstrates that $f_{\max}(l, t) \geq f(\pi, t)$ holds for all π that end in l and all $t \in [0, t_{\max}]$.

As a basis for our induction, this obviously holds for all timed histories longer than n_ε : in this case, $f_{\max}(l, t) = 1$ or $f(\pi, t) = 0$ holds true (where the

or is not exclusive).

For our induction step, let us assume we have demonstrated the claim for all histories of length $> n$. Let us, for a timed history π of length n that ends in l and some point $t \in [0, t_{\max}]$ assume that $f_{\max}(l, t) \leq f(\pi, t)$.

For $l \in G$ the initialisation conditions immediately lead to the contradiction $1 < f(\pi, t)$. For $l \notin G$, we can stepwise infer

$$\begin{aligned}
-\dot{f}_{\max}(l, t) &= \max \left\{ \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\max}(l', t) - f_{\max}(l, t)) \quad \Big| \right. \\
&\quad \left. a \text{ is enabled in } l \right\} \\
&\geq \sum_{\text{distribution}} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\max}(l', t) - f_{\max}(l, t)) \\
&\geq \sum_{\text{distribution}} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (f_{\max}(l', t) - f(\pi, t)) && \text{(with } \\
&\quad f_{\max}(l, t) \leq f(\pi, t)) \\
&\geq \sum_{\text{distribution}} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot \left(f(\pi \xrightarrow{a, t} l', t) - f(\pi, t) \right) && \text{(with I.H.)} \\
&= -\dot{f}(\pi, t).
\end{aligned}$$

Taking into account that $f_{\max}(l, t_{\max})$ and $f(\pi, t_{\max})$ are both initialised to 0 for $l \notin G$, we can, using the just demonstrated $f_{\max}(l, t) \leq f(\pi, t) \Rightarrow \dot{f}_{\max}(l, t) \leq \dot{f}(\pi, t)$, infer $f_{\max}(l, t) \geq f(\pi, t)$ for all $t \in [0, t_{\max}]$: This inequation holds on the right fringe of the interval (initialisation), and when we follow the curves of $f(l, t)$ and $f_{\max}(l, t)$ to the left along $[0, t_{\max}]$, then every time f would catch up with f_{\max} , f cannot fall steeper than f_{\max} (where ‘fall’ takes the usual left-to-right view, in the right-to-left direction we consider one should maybe say ‘cannot have a steeper ascend’) at such a position, and hence cannot not get above f_{\max} .

In particular, $f(l_0, t_0) \leq f_{\max}(l_0, t_0)$, which contradicts the initial assumption.

The min case can be proven accordingly, which provides a full proof of Lemma 1.

(Note that the extension to CTMDPs with both discrete and continuous locations is provided in Section 5.)

C Reproof of Theorem 5

To lift Theorem 6 to the full class of CTMGs, we reprove Theorem 5 in the style of the proof of Theorem 3. Recall that Theorem 3 establishes the existence of an optimal cylindrical scheduler *using* the existence of an optimal measurable scheduler, and the form of the (differential) equations defining the time-bounded reachability probability for it. The proof given in this appendix can therefore not be used to supersede the proof in the paper.

First we observe from the proof of Theorem 5 that, for discrete locations $l \in L_d$, the equations

$$f_{\text{opt}}(l, t) = \text{opt}_{a \in \text{Act}(l)} \sum_{l' \in L} \mathbf{P}(l, a_l, l') \cdot f_{\text{opt}}(l', t) \text{ for } t \in [0, t_{\max}],$$

holds for $\text{opt} \in \{\min, \max\}$, and that they together with the differential equations for the continuous locations (the differential equations remain unchanged), define f_{opt} .

The difference in the proof of Theorem 8 compared to the proof of Theorem 3 are marked in [blue](#).

Theorem 8. *For a single player continuous-time Markov game with only a reachability player, there is an optimal deterministic scheduler with finitely many switching points.*

Proof. We have seen that the true optimal reachability probability is defined by a system of [equations and](#) differential equations. In this proof we consider the effect of starting with the ‘correct’ values for a time $t \in [0, t_{\max}]$, but *locally fix a positional strategy* for a small left or right ε -environment of t . That is, we consider only schedulers that keep their decision constant for a (sufficiently) small time ε before or after t .

Given a CTMG \mathcal{M} , we consider the [equations and](#) differential equations that describe the development of the reachability probability for each location l under a positional deterministic strategy D :

$$-\dot{Pr}_l^D(\tau) = \sum_{l' \in L} \mathbf{R}(l, a_l, l') \cdot (Pr_{l'}^D(\tau) - Pr_l^D(\tau)) \quad \text{for } l \in L_c,$$

$$Pr_l^D(\tau) = \sum_{l' \in L} \mathbf{P}(l, a_l, l') \cdot Pr_{l'}^D(\tau) \quad \text{for } l \in L_d,$$

where a_l is the action chosen at l by D , starting at the support point $f_{\max}(l, t)$.

Different to the development of the true probability, the development of these linear differential equations provides us with smooth functions. This provides us with more powerful techniques when comparing two locally positional strategies: Each deterministic scheduler defines a system $\dot{y} = Ay$ of ordinary homogeneous linear differential equations with constant coefficients.

As a result, the solutions $Pr_l^D(\tau)$ of these differential equations—and hence their differences $Pr_l^{D'}(\tau) - Pr_l^D(\tau)$ —can be written as finite sums $\sum_{i=1}^n P_i(\tau)e^{\lambda_i \tau}$, where P_i is a polynomial and the λ_i may be complex. Consequently, these functions are holomorphic.

Using the identity theorem for holomorphic functions, t can only be a limit point of the set of 0 points of $Pr_l^{D'}(\tau) - Pr_l^D(\tau)$ if $Pr_l^{D'}(\tau)$ and $Pr_l^D(\tau)$ are identical on an ε -environment of t . The same applies to their derivations: $\dot{Pr}_l^{D'}(\tau) - \dot{Pr}_l^D(\tau)$ either has no limit point in t , or $\dot{Pr}_l^{D'}(\tau)$ and $\dot{Pr}_l^D(\tau)$ are identical on an ε -environment of t .

For the remainder of the proof, we fix, for a given time t , a sufficiently small $\varepsilon > 0$ such that, for each pair of schedulers D and D' [the following holds](#): for every location $l \in L_c$, $\dot{Pr}_l^{D'}(\tau) - \dot{Pr}_l^D(\tau)$ is either < 0 , $= 0$, or > 0 on the complete interval $L_\varepsilon^t = (t - \varepsilon, t) \cap [0, t_{\max}] \ni \tau$, and, possibly with different sign, for the complete interval $R_\varepsilon^t = (t, t + \varepsilon) \cap [0, t_{\max}] \ni \tau$; [and for every location](#) $l \in L_d$, $Pr_l^D(\tau) - Pr_l^{D'}(\tau)$ is either < 0 , $= 0$, or > 0 on the complete interval

$L_\varepsilon^t = (t - \varepsilon, t) \cap [0, t_{\max}] \ni \tau$, and, possibly with different sign, for the complete interval $R_\varepsilon^t = (t, t + \varepsilon) \cap [0, t_{\max}] \ni \tau$.

We argue the case for the left ε -environment L_ε^t . In the ‘>’ case for a location l , we say that D is *l-better* than D' . We call D *preferable* over D' if D' is not *l-better* than D for any location l , and *better* than D' if D is preferable over D' and *l-better* for some $l \in L$.

If D' is *l-better* than D in exactly a non-empty set $L_b \subset L$ of locations, then we can obviously use D' to construct a strategy D'' that is better than D by switching to the strategies of D' in exactly the locations L_b .

Since we choose our strategies from a finite domain—the deterministic positional schedulers—this can happen only finitely many times. Hence we can stepwise *strictly* improve a strategy, until we have constructed a strategy D_{\max} that is preferable over all others.

By the definition of being preferable over all other strategies, D_{\max} satisfies

$$-\dot{Pr}_l^{D_{\max}}(\tau) = \max_{a \in Act(l)} \sum_{l' \in L} \mathbf{R}(l, a, l') \cdot (Pr_{l'}^{D_{\max}}(\tau) - Pr_l^{D_{\max}}(\tau)) \quad \text{for all } \tau \in L_\varepsilon^t, l \in L_c,$$

$$Pr_l^{D_{\max}}(\tau) = \max_{a \in Act(l)} \sum_{l' \in L} \mathbf{P}(l, a, l') \cdot Pr_{l'}^{D_{\max}}(\tau) \quad \text{for all } \tau \in L_\varepsilon^t, l \in L_d.$$

We can use the same method for the right ε -environment R_ε^t , and pick the decision for t arbitrarily; we use the decision from the respective left ε environment.

Now we have fixed, for an ε -environment of an arbitrary $t \in [0, t_{\max}]$, an optimal scheduler with at most one switching point. As this is possible for all points in $[0, t_{\max}]$, the sets $I_\varepsilon^t = L_\varepsilon^t \cup R_\varepsilon^t$ define an open cover of $[0, t_{\max}]$. Using the compactness of $[0, t_{\max}]$, we infer a finite sub-cover, which establishes the existence of a strategy with a finite number of switching points. \square

Again, the proof for single player safety games runs accordingly.

Theorem 9. *For a single player continuous-time Markov game with only a safety player, there is an optimal deterministic scheduler with finitely many switching points.*

D From Late to Early Scheduling

Our main motivation for introducing discrete transitions is not the slightly improved generality of the model (nice though it is as a side result), but the introduction of a framework that covers *both* early schedulers (which have to fix an action when *entering* a location), and the late schedulers used in the paper.

Late schedulers are naturally subsumed in our model, as the schedulers we assume are the more powerful late schedulers. To embed early schedulers as well, it suffices to use a simple translation: we ‘split’ every continuous location l_c into a fresh discrete location l_c^d , and one fresh continuous location l_c^a for each action $a \in Act(l_c)$ enabled in l_c .




Fig. 5. An informal example, depicting the idea of the encoding of an early scheduling CTMG (left) in a late scheduling CTMG (right).

Every incoming transition to l_c is re-routed to l_c^d , l_c^d has an outgoing transition a that surely leads to l_c^a ($\mathbf{P}(l_c^d, a, l_c^a) = 1$) for each action $a \in \text{Act}(l_c)$ enabled in l_c , and no other outgoing transition. In l_c^a , we have $\text{Act}(l_c^a) = \{a\}$, and the entries in $\mathbf{R}(l_c^a, a, l)$ are the entries taken from $\mathbf{R}(l_c, a, l)$ for discrete locations l , and re-routed to the respective l^d for continuous locations. Probability mass assigned to l_c is moved to l_c^d by the translation, and if l_c is a goal state, so are l_c^d and the l_c^a 's.

Intuitively, every occurrence of $\xrightarrow{*,t} l_c \xrightarrow{a,t'} \rightarrow$ is replaced by $\xrightarrow{*,t} l_c^d \xrightarrow{a,t} l_c^a \xrightarrow{a,t'} \rightarrow$; l_c is the beginning of the path, $l_c \xrightarrow{a,t'} \rightarrow$ is replaced by $l_c^d \xrightarrow{a,t} l_c^a \xrightarrow{a,t'} \rightarrow$.

Obviously, there is a trivial bijection between early schedulers for a thus translated CTMDP (or, indeed, CTMG), and the late schedulers in the mapping: the actions chosen in a discrete location are doubled, and there is no alternative to doubling it.

As a consequence, the existence of finite deterministic optimal control extends to early scheduling.

From Early to Late Scheduling. As a side remark, we would like to point out that a similar translation can be used to reduce finding optimal control for late schedulers to finding optimal control for early schedulers. Following up on the remark that—assuming late scheduling—we can work with the uniformisation of a CTMG when seeking co-optimal control, it suffices to establish such a translation for uniform CTMGs, that is, for CTMGs where the transition rate $\mathbf{R}(l_c, a, L)$ is constant for all continuous locations $l_c \in L$, and all their enabled actions $a \in \text{Act}(l_c)$. And for such uniform CTMGs we can move the decision into a fresh discrete location *after* the continuous location, just as we moved it to a fresh discrete location *before* the continuous location in our reduction from late to early scheduling.