A Characterization of Meaningful Schedulers for Continuous-Time Markov Decision Processes

Nicolás Wolovick⋆[1] and Sven Johr⋆⋆[2]

[1] Fa.M.A.F., Universidad Nacional de Córdoba, Ciudad Universitaria, 5000 Córdoba, Argentina

[2] Universität des Saarlandes, FR 6.2 Informatik, D-66123 Saarbrücken, Germany

Abstract. Continuous-time Markov decision process are an important variant of labelled transition systems having nondeterminism through labels and stochasticity through exponential fire-time distributions. Nondeterministic choices are resolved using the notion of a *scheduler*. In this paper we characterize the class of *measurable* schedulers, which is the most general one, and show how a measurable scheduler induces a unique probability measure on the sigma-algebra of infinite paths. We then give evidence that for particular reachability properties it is sufficient to consider a subset of measurable schedulers. Having analyzed schedulers and their induced probability measures we finally show that each probability measure on the sigma-algebra of infinite paths is indeed induced by a measurable scheduler which proves that this class is complete.

1 Introduction

Continuous-time Markov decision processes (CTMDP) [1,2,3,4] are an important class of finite labelled transition systems (LTS). They have external nondeterminism through the interaction with edge labels and internal stochasticity given by the rates of negative exponential distributions. Besides their applications in, e. g., stochastic control theory [2], stochastic scheduling [5,6] and dynamic power management [7], these systems are interesting in their own, because they introduce a continuous quantity, namely the fire-time of transitions.

The external nondeterminism is usually reduced to probabilism using the notion of *schedulers*, also called *adversaries* or *policies*. Given that the system is in a particular state, a randomized scheduler eliminates the nondeterminism by making a decision for a certain probability distribution over edge labels. This decision making can be based on the present state alone, or together with all previous states, edge labels and time points, leading to memoryless or timed history-dependent schedulers, respectively. Besides, it is also possible to have deterministic schedulers as a particularization of randomized schedulers. However, as shown in [8], timed history-dependent schedulers are strictly more powerful than schedulers abstracting from time. Therefore time has an explicit role in the global behavior of the model of

⋆ Supported by the Programme Alßan, the European Union Programme of High Level Scholarships for Latin America, scholarship No.(E04D030410AR). Partially supported by PIP 6391: Métodos y Fundamentos para el Análisis de Corrección y Desempeño de Programas Concurrentes y Aleatorios, funded by CONICET, and PICT 26135: Verification of Probabilistic Distributed Systems, funded by Agencia.

⋆⋆ This work is supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center "Automatic Verification and Analysis of Complex Systems" (SFB/TR 14 AVACS). See `www.avacs.org` for more information.

CTMDPs and cannot, in principle, be avoided. We take the most general class of schedulers, namely randomized timed history-dependent (THR) schedulers as our starting point.

Time is continuous in its nature and thus particular issues quickly arise. Many of them refer to practical aspects as discretization, abstraction, approximation and computation, but others like measurability are more fundamental [9]. It is well known, that there are sets, namely *Vitali sets*, in the real interval $[0, 1]$ that cannot be measured. A measure is a function that generalizes the usual notions of cardinality, length, area and probability. With these unreasonable sets, all the questions about our system asking for a quantitative answer start to be endangered. The motivation of this paper is to tackle this particular problem.

We give a *soundness* result, showing that measurable schedulers generate a well-defined probability measure for timed paths. For this we first construct a combined transition that merges scheduler and CTMDP probabilism in a single transition probability. We establish that measurability of combined transitions is inherited from schedulers and vice versa, improving a result from [9]. Based on transition probabilities we construct a probability measure on the set of infinite timed paths with the aid of canonical product measure theorems.

A proper subclass of measurable schedulers inducing positive probability measures on infinite paths is identified. Besides that and the construction of scheduler dependent probability measures, we also answer the question if there is a proper subclass of measurable schedulers which is easily definable and generates arbitrary measurable schedulers (called *probabilistically complete*) in the affirmative.

Finally a *completeness* result is presented, showing that every timed path probability is generated by some measurable scheduler. To do this we motivate the need of this probability to be related or *compatible* with the underlying CTMDP. This compatibility cannot be stated using the probability measure. Instead, (continuous) conditional probabilities are obtained by the so-called Radon-Nikodym derivatives. However, those conditional probabilities turn out to be insufficient for the intended compatibility result. Taking advantage of our particular setting (timed paths as denumerable product of discrete spaces and positive reals), we use *disintegrability* to obtain transition probabilities instead of conditional ones. Deconstructing the probability into transition probability lead us to state compatibility and the completeness theorem.

Outline. The rest of the paper is organized as follows. In Section 2 we give an overview of related work. Section 3 establishes mathematical background and notation, including the CTMDP model with schedulers, as well as the measurability problem. Section 4 develops a timed path probability measure w.r.t. measurable schedulers. In Section 5 we propose the simple scheduler subclass and discuss its limiting properties. Section 6 shows that measurable schedulers can capture every probability measure generated by a CTMDP. Finally, Section 7 concludes the paper.

2 Related work

In systems with continuous state spaces featuring nondeterminism and probabilism, measurability issues of schedulers have been studied from the point of view of timed systems [10], extensions of discrete probabilistic automata [9] and Markov decision process (MDP) with Borel state space [11].

In [10] continuous probabilistic timed automata are defined and its concrete semantics is given in terms of dense Markov process once the nondeterminism is resolved by a timed history-dependent scheduler. Although the model is not comparable with CTMDPs because of the denseness of the state space, scheduler measurability issues also appear. The authors disregard this particular problem by considering it pathological.

Stochastic transition systems [9] are a generalization of CTMDPs and the measurability of schedulers is an important issue here. Cattani et al. show how to construct a measure on so called *executions*. For this they have to, and indeed do discard nonmeasurable schedulers. In particular, although we discuss our results in a more restrictive modelling formalism, this work is an improvement over their construction in two aspects. First, we use standard measure-theoretic results in conjunction with product space measures, leading to results that are more compact and provide deeper insight than the results presented in [9]. Secondly, the need for measurable functions in the construction of the probability measure is in this paper directly inherited from the definition of scheduler measurability. Instead, in [9] measurability properties are kind of directly given for the according functions [9, Definition 5] without giving a connection to schedulers. We added to this work a complete characterization of measurable schedulers, something that has not been considered in [9].

A result similar to our deconstruction discussed in Section 6 is given in the context of epistemic game theory [12]. The proof-tools used there are disintegration theorems for the product of two spaces [13] and the so-called *Ionescu-Tulcea* theorem. However, we prove a similar theorem by resorting to the well-known Radon-Nikodym theorem together with canonical product measures for denumerable product spaces.

To the best of our knowledge, the decomposition of particular probability measures on the infinite behavior of probabilistic systems with nondeterminism leading to measurable schedulers has not been considered in the literature so far.

3 Background and problem statement

3.1 Mathematical notation and background

By $\mathbb{R}$ and $\mathbb{N}$ the set of *real numbers* and *natural numbers* are denoted, $\mathbb{R}^+$ is the set of positive real numbers including 0. *Disjoint union* is denoted $\Omega_1 \uplus \Omega_2$, while *partial function application* is denoted $f(\cdot, \omega_2)$.

Given a set Ω and a collection $\mathcal{F}$ of subsets of Ω, we call $\mathcal{F}$ a *σ-algebra* iff $\Omega \in \mathcal{F}$ and $\mathcal{F}$ is closed under complement and denumerable disjoint union. The *σ-algebra generated by* the family $\mathcal{H} \in 2^\Omega$ is the minimal σ-algebra containing $\mathcal{H}$. We call the pair $(\Omega, \mathcal{F})$ a *measurable space*. A *measurable set* is denoted as $A \in \mathcal{F}$. When dealing with a discrete set Ω, we take the powerset σ-algebra 2^Ω as the default, however we still denote a *measurable set* as $A \in \mathcal{F}$ instead of the simpler $A \subseteq \Omega$. Let $(\Omega_i, \mathcal{F}_i), i = 1, 2, \ldots, n$ be arbitrary measurable spaces and $\Omega = \Omega_1 \times \Omega_2 \times \cdots \times \Omega_n$. A *measurable rectangle* in Ω is a set $A = A_1 \times A_2 \times \cdots \times A_n$, where $A_i \in \mathcal{F}_i$. The *product space σ-algebra* denoted $\mathcal{F}_{\Omega_1 \times \Omega_2 \times \cdots \times \Omega_n}$ is the σ-algebra generated by measurable rectangles. Given a measurable space $(\Omega, \mathcal{F})$, a σ-additive function $\mu : \mathcal{F} \to \mathbb{R}^+$ is called *measure*, and if $\mu(\Omega) = 1$ it is called *probability measure*. The triple $(\Omega, \mathcal{F}, \mu)$ is a *measure space* or *probability space* depending on μ. For a measurable set $(\Omega, \mathcal{F})$ we denote by $\mathrm{Distr}(\Omega)$ the *set of all probability distributions* over Ω, and given $\mu \in \mathrm{Distr}(\Omega)$ a *support* of μ is a measurable set A such that $\mu(A) = 1$.

A function $f : (\Omega_1, \mathcal{F}_1) \to (\Omega_2, \mathcal{F}_2)$ is *measurable* if $\forall A_2 \in \mathcal{F}_2,\ f^{-1}(A_2) \in \mathcal{F}_1$, i.e., the inverse function maps measurables to measurables. A measurable predicate $P : \Omega \to \mathit{Bool}$ in a measure space $(\Omega, \mathcal{F}, \mu)$ is *μ-almost everywhere valid*, P a.e. $[\mu]$, iff $\mu(P^{-1}(\mathit{false})) = 0$, that is the set of non-valid points is negligible. We call a function $f : \Omega_1 \times \mathcal{F}_2 \to [0,1]$ *transition probability* or *Markov kernel* iff for all $\omega_1 \in \Omega_1$, $f(\omega_1, \cdot)$ is a probability measure on $(\Omega_2, \mathcal{F}_2)$ and for all $A_2 \in \mathcal{F}_2$, $f(\cdot, A_2)$ is a measurable function.

3.2 Continuous-time Markov Decision Processes

In the following we present the basic definitions of continuous-time Markov decision processes. We use essentially the same notation as in [3].

Definition 1 (CTMDP). *A* continuous-time Markov decision process *(CTMDP) is a tuple $(S, L, \mathbf{R}, \Pr_{init})$ where S is a finite non-empty set of* states*, L is a finite non-empty set of* transition labels *also called* actions*, $\mathbf{R} : S \times L \times S \to \mathbb{R}^+$ is the three-dimensional* rate matrix*, $\Pr_{init} \in \mathit{Distr}(S)$ is the* initial distribution *over states.*

Given a CTMDP tuple $\mathcal{C} = (S, L, \mathbf{R}, \Pr_{\text{init}})$, we define the projection function $\mathcal{C}_{\Pr_{\text{init}}} \doteq \Pr_{\text{init}}$, and so on for the other coordinates. For set $Q \subseteq S$ we denote by $\mathbf{R}(s, a, Q) \doteq \sum_{s' \in Q} \mathbf{R}(s, a, s')$ the cumulative rate to leave state s under label a. In $L_s \doteq \{a \in L \mid \mathbf{R}(s, a, S) > 0\}$ we collect all labels that belong to transitions emanating from s.

Behavior. The behavior of a CTMDP is as follows. $\mathbf{R}(s, a, s') > 0$ denotes that there exists a transition from s to s' under label a where $\mathbf{R}(s, a, s')$ corresponds to the rate of a negative exponential distribution. When s has more than one successor state under label a, one of them is selected according to the *race condition.* The discrete branching probability $\mathbf{P}_s(a, s')$ from s to s' under label a is given by $\mathbf{P}_s(a, s') \doteq \frac{\mathbf{R}(s,a,s')}{\mathbf{R}(s,a,S)}$, where $\mathbf{R}(s, a, S)$ is the overall exit rate of s under label a. The probability that one of the successors of s is reached within time t is given by $1 - e^{-\mathbf{R}(s,a,S)\cdot t}$. In CTMCs this time is also referred to as the *sojourn time* in state s. Given that $a \in L$ is chosen, the sojourn time in $s \in S$ is determined by a negative exponential distribution with rate $\mathbf{R}(s, a, S)$.

Figure 1 shows a simple CTMDP. Arrows indicate transitions between states. They are labelled by an action and a rate, e. g., under label a there exists a transition from s_1 to s_3 whose fire-time is exponentially distributed with rate 1.

3.3 Timed paths

A *timed path* σ in CTMDP $\mathcal{C} = (S, L, \mathbf{R}, \Pr_{\text{init}})$ is a possibly infinite sequence of states, labels and time points, i. e., $\sigma \in \left(S \times (L \times \mathbb{R}^+ \times S)^*\right) \uplus \left(S \times (L \times \mathbb{R}^+ \times S)^\omega\right)$. For a given path $\sigma = s_0 a_1 t_1 s_1 \ldots$ we denote by $\mathit{first}(\sigma) = s_0$ its first state, $\sigma[k]$ denotes its $(k+1)$-st state, e. g., $\sigma[0] = \mathit{first}(\sigma)$. A finite path $\sigma' = s_0 a_1 t_1 s_1 \ldots a_k t_k s_k$ has *length k*, denoted as $|\sigma'| = k$ and its last state equals $\mathit{last}(\sigma) = s_k$. For arbitrary paths σ we denote by σ^i the *prefix of length i* of σ, i. e., $\sigma^i = s_0 a_1 t_1 s_1 \ldots a_i t_i s_i$. For finite $\sigma' = s_0 a_1 t_1 s_1 \ldots a_k t_k s_k$, the prefix of length $i > k$ equals σ'. Path^n, Path^*, Path^ω and Path denote the sets of paths of length n, *finite* paths, *infinite* paths and the union thereof, respectively, where $\mathit{Path} = \mathit{Path}^* \uplus \mathit{Path}^\omega$ and

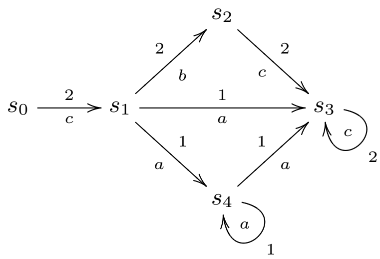

Fig. 1. Continuous-time Markov decision process

$Path^* = \biguplus_{n \in \mathbb{N}} Path^n$. A *time abstract* path is a timed path where time points are omitted. It will be clear from context which type of path we use in particular sections. Thus, the above given definitions carry over to time abstract paths.

Paths are defined as *arbitrary* sequences of states, labels and time points w.r.t. $\mathcal{C}$. We do not require that $\mathbf{R}(s_i, a_i, s_{i+1}) > 0$ but we will overcome this obstacle by giving all power to our schedulers. Given CTMDP $\mathcal{C}$ we distinguish between *observable* and *unobservable* paths w.r.t. $\mathcal{C}$. A path $\sigma = s_0 a_1 t_1 s_1 \ldots$ is called observable iff $\mathbf{R}(s_i, a_{i+1}, s_{i+1}) > 0$ for $i = 0, 1, \ldots$. Otherwise we call σ *unobservable*. As will be evident in Section 4, the probability measure on paths induced by a given scheduler will evaluate to 0 for unobservable paths. But before we introduce the measure, we have to define the measurable space over paths.

σ-Algebra over timed paths. We endow states and labels with the powerset σ-algebra, and $\mathbb{R}^+$ with the standard Borel σ-algebra. The σ-algebra $\mathcal{F}_{Path^n}$ is the standard product space σ-algebra, $\mathcal{F}_{S \times (L \times \mathbb{R}^+ \times S)^n}$, generated by the rectangles $Q_0 \times M_1 \times \cdots \times M_n$, where $Q_0 \in \mathcal{F}_S$ and $M_i \in \mathcal{F}_{L \times \mathbb{R}^+ \times S}$. The σ-algebra $\mathcal{F}_{Path^\omega}$ is defined using the concept of *cylinders* [14, Definition 2.7.1]. Given the set of timed paths Λ with $|\Lambda| = k$, $C_\Lambda \doteq \{\sigma \in Path^\omega \mid \sigma^k \in \Lambda\}$ defines a cylinder with *base* Λ. A *measurable cylinder* has measurable base, and a *measurable rectangle* (in $Path^\omega$) is a cylinder whose base is a measurable rectangle. The σ-algebra $\mathcal{F}_{Path^\omega}$ is the minimal σ-algebra generated by either measurable cylinders or measurable rectangles. Finally $\mathcal{F}_{Path}$ is standard σ-algebra for the disjoint union of $\{\mathcal{F}_{Path^n}\}_{n \in \mathbb{N}}$ and $\mathcal{F}_{Path^\omega}$.

3.4 Schedulers

A scheduler resolves the nondeterminism inherent in a CTMDP. Most generally, schedulers can be considered as functions from finite paths to probability distributions over labels. Such a general definition allows for all special cases as, e. g., given in [3].

Definition 2 (Scheduler over CTMDP). *Let $\mathcal{C}$ be a CTMDP with label set L. A* scheduler *D* over *$\mathcal{C}$ is a function $D : Path^* \to Distr(L)$ such that the support of $D(\sigma)$ is equal to $L_{last(\sigma)}$.*

The support condition says that the scheduler distributes its whole mass among the outgoing labels. We can also denote D as a two argument function $D : Path^* \times \mathcal{F}_L \to [0, 1]$, being a measure in its second argument $D(\sigma, \cdot)$ for all finite path σ. We now turn our attention to the *measurability problem*. This is motivated by the fact that there exist "bad-behaved" schedulers. Reconsider the CTMDP depicted in Figure 1 and the randomized timed history-dependent (THR) scheduler D with

$$D(s_0cts_1) = \begin{cases} \delta_a & \text{if } t \in \mathcal{V}, \\ \delta_b & \text{otherwise,} \end{cases}$$

where $\delta_i, i \in \{a, b\}$ denotes that action i is chosen with probability 1 and $\mathcal{V}$ is the *nonmeasurable Vitali set* in $[0, 1]$, i. e., for each time point in a nonmeasurable set, D will choose deterministically a and for all other time points, D will choose deterministically b. So the question, "*What is the probability to reach s_3 in two steps?*", cannot be answered, because it necessarily involves the evaluation of a measure in the set $\{t|D(s_0cts_1) = \delta_a\} = \mathcal{V}$ that is nonmeasurable.

In the subsequent definition we restrict to the class of schedulers that respects measurability issues.

Definition 3 (Measurable scheduler). *We call a scheduler D over CTMDP C* measurable scheduler *iff for all $A \in \mathcal{F}_L$, $D(\cdot, A) : Path^* \to [0, 1]$ is a measurable function.*

4 Constructing timed path probability

Probabilistic schedulers allow to quantify uncertainty. Therefore composing CTMDP negative exponential distributions with a particular scheduler, we can construct probability measures usually called *combined transitions* [9].

Definition 4 (Combined transition). *For a given CTMDP $\mathcal{C}$, scheduler D and finite path σ, the* combined transition $\mu_D : Path^* \times \mathcal{F}_{L\times\mathbb{R}^+\times S} \to [0, 1]$ *is defined in the measurable rectangles by*

$$\mu_D(\sigma, A\times\mathcal{I}\times Q) \doteq \sum_{a\in A} D(\sigma, \{a\})\cdot\mathbf{P}_{last(\sigma)}(a, Q)\cdot\int_{\mathcal{I}} \mathbf{R}(last(\sigma), a, S)e^{-\mathbf{R}(last(\sigma),a,S)\cdot t}dt \ . \quad (1)$$

The second factor in the terms of the above summation is the discrete branching probability from $last(\sigma)$ to Q by the label a, while the third one is the probability that the transition triggers within measurable time set $\mathcal{I}$.

We have, thanks to support restriction in the scheduler, $\mu_D(\sigma, L \times \mathbb{R}^+ \times S) = 1$, however μ_D is just defined for the $\mathcal{F}_{L\times\mathbb{R}^+\times S}$ generators. The extension to the whole σ-algebra is standard and follows by applying Carathéodory Extension Theorem [14, Theorem 1.3.10] to the field of finite disjoint unions of measurable rectangles [14, Problem 2.6.1].

Lemma 1 (Combined transition is a probability measure). *Given a CTMDP $\mathcal{C}$ and a finite path σ, the combined transition $\mu_D(\sigma, \cdot)$ as defined in (1) extends uniquely to a probability measure on $\mathcal{F}_{L\times\mathbb{R}^+\times S}$.*

Additionally we deduce the following lemma, where the *only if* implication is given by Lemma 1 and the *if* implication can be established by considering Definitions 2 and 4.

Lemma 2. *Given a CTMDP $\mathcal{C}$ and scheduler D, $\forall \sigma \in \mathcal{F}_{Path^*}$, it is the case that $D(\sigma, \cdot) : \mathcal{F}_L \to [0,1]$ is a measure iff $\forall \sigma \in \mathcal{F}_{Path^*}$, we have $\mu_D(\sigma, \cdot) : \mathcal{F}_{L \times \mathbb{R}^+ \times S} \to [0,1]$ is a measure.*

Next we prove that combined transition measurability is inherited from scheduler measurability and vice versa, and this will be crucial to integrate and disintegrate a probability measure on timed paths.

Theorem 1 (Combined transition measurability). *Given a CTMDP $\mathcal{C}$ and scheduler D, $\forall A \in \mathcal{F}_L$, it is the case that $D(\cdot, A) : Path^* \to [0,1]$ is measurable iff $\forall M \in \mathcal{F}_{L \times \mathbb{R}^+ \times S}$, we have $\mu_D(\cdot, M) : Path^* \to [0,1]$ is measurable.*

Proof. We prove each direction separately.

Only if: Having $D(\cdot, A)$ measurable, function $\mu_D(\cdot, A \times \mathcal{I} \times Q)$ is measurable in its first argument because projection functions like *last*() are measurable, as well as functions coming from powerset σ-algebras like $\mathcal{F}_S$ and $\mathcal{F}_L$. Closure properties of measurable functions (sum, product, composition) make the entire expression measurable.

We have to extend this result to the whole σ-algebra $\mathcal{F}_{L \times \mathbb{R}^+ \times S}$. We resort to the *good sets principle* [14]. Let $\mathcal{G} = \{M \in \mathcal{F}_{L \times \mathbb{R}^+ \times S} \mid \mu_D(\cdot, M) \textit{ is a measurable function}\}$ and show that the set $\mathcal{G}$ forms a σ-algebra. Let $\{M_i\}_{i \in \mathbb{N}}$ be a disjoint collection such that $M_i \in \mathcal{G}$, then by σ-additivity $\mu_D(\cdot, \biguplus_{i \in \mathbb{N}} M_i) = \sum_{i \in \mathbb{N}} \mu_D(\cdot, M_i)$, and this is measurable by closure properties of measurable functions, therefore $\mathcal{G}$ is closed under denumerable disjoint union. It is also closed under complement as $\mu_D(\cdot, M^c) = \mu_D(\cdot, L \times \mathbb{R}^+ \times S) - \mu_D(\cdot, M)$. Hence all sets in $\sigma(\mathcal{G})$ are good, and since rectangles are included in $\mathcal{G}$ we have $\sigma(\mathcal{G}) = \mathcal{F}_{L \times \mathbb{R}^+ \times S}$, concluding $\mu_D(\cdot, M)$ is measurable for all $M \in \mathcal{F}_{L \times \mathbb{R}^+ \times S}$.

If: By hypothesis the function $\mu_D(\cdot, A \times \mathbb{R}^+ \times S)$ is measurable for all $A \in \mathcal{F}_L$, but is easy to check that $\mu_D(\cdot, A \times \mathbb{R}^+ \times S) = D(\cdot, A)$, finishing our result. □

Given the previous results, we can deduce the following corollary.

Corollary 1. *Let D be a measurable scheduler on CTMDP $\mathcal{C}$. D is a transition probability iff μ_D is so.*

From now on we use canonical product measure theory [14, Sections 2.6, 2.7] to define finite and infinite timed path probability measure, where the combined transition plays a central role.

Definition 5 (Probability measure for measurable rectangle finite paths). *Let $\mathcal{C}$ be a CTMDP, D a measurable scheduler, and μ_D their combined transition. The* probability measure for finite paths consisting of measurable rectangles Pr_D *is given inductively as follows:*

$$Pr_D^0(S_0) = \mathcal{C}_{\mathrm{Pr}_{init}}(S_0) \ , \tag{2}$$

$$Pr_D^{n+1}(\Lambda \times M_{n+1}) = \int_{\sigma \in \Lambda} \mu_D(\sigma, M_{n+1}) dPr_D^n(\sigma) \ . \tag{3}$$

where Pr_{init} is the initial distribution, Λ is a measurable rectangle of $\mathcal{F}_{Path^n}$, and $M_{n+1} \in \mathcal{F}_{L \times \mathbb{R}^+ \times S}$.

Combined transitions are measurable, thus the above Lebesgue integral is well-defined. However, the definition only captures the measurable rectangles and we have to extend the measure to the complete σ-algebra. The unique extension of Pr_D^n is given by [14, Theorem 2.6.7].

Lemma 3 (Probability measure for finite paths). *For each n there is a unique probability measure Pr_D^n over the σ-algebra $(Path^n, \mathcal{F}_{Path^n})$ that extends Pr_D^n defined in (2),(3).*

Note that Pr_D^n is not a probability measure for finite paths but a denumerable set of probability measures one for each path length. They can be put together in a single probability measure space, namely the infinite timed path measure space $(Path^\omega, \mathcal{F}_{Path^\omega}, Pr_D^\omega)$ [14, Theorem 2.7.2] [15, Ionescu-Tulcea Theorem].

Lemma 4 (Probability measure for infinite paths). *Given the measurable space $(Path^\omega, \mathcal{F}_{Path^\omega})$, if we define a probability measure on measurable rectangle bases as in Definition 5, then there is a unique probability measure Pr_D^ω on $\mathcal{F}_{Path^\omega}$ such that for all n, $Pr_D^\omega(C_\Lambda) = Pr_D^n(\Lambda)$ with $\Lambda \in \mathcal{F}_{Path^n}$.*

The previous results can be summarized in the following theorem.

Theorem 2 (Soundness). *Given a measurable scheduler D over CTMDP $\mathcal{C}$, then a probability measure Pr_D^ω over $\mathcal{F}_{Path^\omega}$ can be constructed.*

Lemma 3 and Lemma 4 define probability measures for each of the disjoint measurable spaces $(\biguplus_{n\in\mathbb{N}}(Path^n, \mathcal{F}_{Path^n})) \uplus (Path^\omega, \mathcal{F}_{Path^\omega})$. A unifying measure (that is not a probability) on the disjoint union space can be defined in a standard way and is denoted by Pr_D. It is important to remark that under this way of constructing Pr, there is no CTMDP that can "hide" a nonmeasurable scheduler, namely having a well defined timed path probability, implies that the combined transition is measurable, and by Theorem 1 this implies a measurable scheduler.

5 Meaningful schedulers

In the previous section we have defined a measure on paths induced by measurable schedulers by using measure theoretic results and in particular abstract Lebesgue integration [14]. Now we examine the integral in more detail and characterize various classes of measurable schedulers that respect more than just the initial probability distribution over states. This boils down to investigate basic properties of abstract Lebesgue integrals over combined transitions. These classes of schedulers are of special interest because they comprise all schedulers that contribute to, e. g., particular reachability properties like: "what is the maximum probability to reach a set B of states within t time units" [8].

The following results are taken from [14] and give the theoretical background for our observations. Each nonnegative Borel measurable function f is the limit of a sequence of increasing simple functions that are nonnegative and finite-valued [14, Theorem 1.5.5]. And so, given a measurable scheduler D the combined transition $\mu_D(\cdot, M)$ is measurable for each $M \in \mathcal{F}_{L\times\mathbb{R}^+\times S}$ and is the limit of an increasing sequence of simple functions, say, $\mu_D^i(\cdot, M)$. Formally, we define $\mu_D^i(\cdot, M)$, for fixed $M \in \mathcal{F}_{L\times\mathbb{R}^+\times S}$ as

$$\mu_D^i(\sigma, M) \doteq \frac{k-1}{2^i} \text{ if } \frac{k-1}{2^i} \leq \mu_D(\sigma, M) < \frac{k}{2^i}, k = 1, 2, \ldots, 2^i \ . \tag{4}$$

Due to convergence properties of the abstract Lebesgue integral [14, Theorem 1.6.2] it holds that $Pr_D^n(\Lambda \times M) = \lim_{i \mapsto \infty} \int_{\sigma \in \Lambda} \mu_D^i(\sigma, M) dPr_D^{n-1}(\sigma)$, i. e., the probability distribution over infinite paths induced by D coincides with the limiting sequence of integrals over μ_D^i. Recall, that the probability measure over infinite paths is defined w.r.t. finite path-prefixes. Thus we deduce for given Λ with $|\Lambda| = n-1$.

$$\begin{aligned} Pr_D^n(C_{\Lambda a \mathcal{I} s}) &= \int_{\sigma \in \Lambda} \mu_D(\sigma, (a, \mathcal{I}, s)) dPr_D^{n-1}(C_\sigma) \\ &= \lim_{i \mapsto \infty} \int_{\sigma \in \Lambda} \mu_D^i(\sigma, (a, \mathcal{I}, s)) dPr_D^{n-1}(C_\sigma) \\ &= \lim_{i \mapsto \infty} \left[\sum_{j=1}^{2^i} \frac{j-1}{2^i} \cdot Pr_D^{n-1}(A_j^i) \right] \ , \end{aligned} \tag{5}$$

where $A_j^i \doteq \left\{ \sigma \in \Lambda \mid \frac{j-1}{2^i} \leq \mu_D(\sigma, (a, \mathcal{I}, s)) < \frac{j}{2^i} \right\}$. The summation in Equation 5 is a direct consequence of the definition of Lebesgue integration over simple functions. As a result we see, that the induced probability measure evaluates to zero iff the limiting sequence evaluates to zero.

Almost Always Measure Zero Schedulers. First of all we characterize the class of schedulers that gives measure zero to almost all sets of paths of a given CTMDP $\mathcal{C}$. We refer to the schedulers of this class as *almost always measure zero schedulers.* In this class we summarize all schedulers for which it is not possible to find cylinders on which they give positive probability. We restrict to particular cylinders as, e. g., cylinders $C_s, s \in S$ give for each scheduler D that $Pr_D^\omega(C_s)$ equals $\text{Pr}_{\text{init}}(s)$. Dependent on scheduler D we can always give at least one cylinder base Λ of length two such that $Pr_D(C_\Lambda)$ is positive, e. g., for $\Lambda = s_0 a_1 I_1 s_1$, where I_1 is a right-open interval with $|I_1| > 0$ and $D(s_0, \{a_1\}) > 0$. This motivates the following definition.

Definition 6 (Almost always measure zero scheduler). *We call scheduler D* almost always measure zero scheduler *(AAMZ) iff $Pr_D(C_\Lambda) = 0$ for all Λ with $|\Lambda| \geq 2$.*

The probability measure Pr_D^ω on the set of infinite paths induced by scheduler D evaluates to 0 iff μ_D equals 0 a.e. $[Pr_D]$. Due to the inductive definition of Pr_D^ω it can be observed, that $\mu_D(\sigma', (a', I', s')) = 0$ at some stage of the computation for all $(\sigma', (a', I', s'))$. We now describe for which type of schedulers this is generally true. For this we use the notion of time abstract paths and show that AAMZ schedulers do not give positive probability to all paths σ with corresponding time abstract path σ' that is observable in the system. We put this restriction in our analysis since the probability measure of unobservable paths is always zero, independently of the scheduler. Thus, the probability measure only depends on the set of time points $\mathcal{I}$ that has to be considered.

Based on $\mathcal{I}$ we compute the sojourn time distribution inside of the combined transition. As sojourn times are distributed according to negative exponential distributions, this reduces to the standard computation of a Riemann integral. Thus it is sufficient to investigate when this integral evaluates to zero which is the case for all point-intervals I. Thus, if $\mathcal{I}$ comprises only point-intervals the probability measure will evaluate to zero. This means, if we can not group paths together in a set Λ' for which D behaves *friendly*, the probability measure will not be positive.

Definition 7 (Friendly scheduler). *Let $\sigma = s_0 a_1 t_1 s_1 \dots a_k t_k s_k$ be an arbitrary finite path. We say that D behaves* friendly on σ *iff*

$$\forall t_i : \exists I_i : t_i \in I_i \ \wedge \ |I_i| > 0 \ \wedge \ \forall t \in I_i : D(s_0 a_1 t_1 s_1 \dots a_i t s_i, \{a_{i+1}\}) > 0, \qquad (6)$$

where I_i is an interval on $\mathbb{R}^+$.

A scheduler that does not behave friendly on path σ is called *unfriendly.* We can establish the following lemma.

Lemma 5 (Almost always measure zero). *If D behaves unfriendly on all finite paths σ then D is an AAMZ.*

Suppose D is a scheduler behaving unfriendly on all finite paths. Further assume that we have as cylinder base $\Lambda = s_0 a_1 \mathcal{I}_1 s_1 \dots a_k \mathcal{I}_k s_k$. It holds for D (by Definition 7) that $\mathcal{I}_i$ cannot be partitioned into intervals I_i^j with $|I_i^j| > 0$ such that $D(s_0 a_1 t_1 s_1 \dots a_i t_i s_i, \{a_{i+1}\}) > 0$ for all $t_i \in I_i^j$. As a consequence, we can only partition Λ into singular paths σ but $Pr_D^\omega(C_\sigma) = 0$.

Restricted Class of Schedulers. Now we discuss a restricted class of schedulers that gives positive probability to particular sets of infinite paths. This restriction depends on the property we want to check for a given CTMDP, i. e., when checking a timed reachability property like "what is the maximum probability to reach set B within time t" we are only interested in schedulers that *contribute* to that probability in a sense that they give positive probability to all sets of paths hitting B before time t. It is thus sufficient to consider schedulers that behave friendly on at least one time abstract path that hits B.

The following example shows how simple the computation of the probability measure is, when the given scheduler behaves friendly on the given set of paths. Assume that $\Lambda = s_0 a_1 I_1 s_1 \dots a_k I_k s_k$ is a given set of finite paths (which all share the same time-abstract path), where I_i is a nonsingular interval of length greater than 0. Now suppose scheduler D is such that each decision is *consistent* in each of the I_i, i. e., $D(s_0 a_0 t_0 s_1 \dots a_i t_i^m s_{i+1}) = D(s_0 a_0 t_0 s_1 \dots a_i t_i^n s_{i+1})$ for all $t_i^m, t_i^n \in I_i$. The probability of cylinder C_Λ induced by D is given by

$$\begin{aligned}
Pr_D^\omega(C_\Lambda) &= Pr_D^\omega(C_{\Lambda^{k-1} a_k I_k s_k}) \\
&= \int_{\sigma \in \Lambda'} \mu_D(\sigma, (a_k, I_k, s_k)) dPr_D^\omega(C_\sigma) \\
&= \int_{\sigma \in \Lambda'} \mathbf{1}_{\Lambda'}(\sigma) \mu_D(\sigma, (a_k, I_k, s_k)) dPr_D^\omega(C_\sigma) \\
&= \mu_D(\Lambda', (a_k, I_k, s_k)) \cdot Pr_D^\omega(C_{\Lambda'}) \\
&\ \vdots \\
&= \Pr_{\text{init}}(s_0) \cdot \mu_D(s_0, (a_1, I_1, s_1)) \cdot \mu_D(s_0 a_1 I_1 s_1, (a_2, I_2, s_2)) \cdot \dots \\
&\quad \cdot \mu_D(s_0 a_1 I_1 s_1 \dots a_{k-1} I_{k-1} s_{k-1}, (a_k, I_k, s_k)) \ ,
\end{aligned}$$

where we use, e. g., $D(\Lambda')$ to denote $D(\sigma)$ for arbitrary $\sigma \in \Lambda'$. In this case non zero probability is given as long as $D(s_0 a_1 I_1 s_1 \dots a_i I_i s_i, \{a_{i+1}\}) > 0$ which is a weak form of a friendly behaving scheduler.

Simple Scheduler. In this section we give the definition of an important class of schedulers, namely *simple schedulers*. Simple functions are functions that take on only finitely many different values. Along this line we define simple schedulers to have a finite range only, too. In particular, simple scheduler D gives a partition P of the set of finite paths in finitely many blocks B^j, for $j, n \in \mathbb{N}$ with $n > 0$ and $j \leq n$. Given arbitrary block B^j of P it holds that $D(\sigma) = D(\sigma')$ for all $\sigma, \sigma' \in B_i^j$. Thus, a simple scheduler D is a simple function on the set of finite paths. As a direct consequence, it can be observed that when D is simple, μ_D is also a simple function in its first argument, i. e., $\mu_D(\cdot, M)$ is a simple function for given $M \in \mathcal{F}_{L\times\mathbb{R}^+\times S}$. As we have discussed earlier in this section every measurable combined transition is the limit of simple combined transitions. This, and the fact that each simple combined transition can only be generated by a simple scheduler leads to the result that the set of simple schedulers can be used to generate arbitrary measurable schedulers and thus, this class is *probabilistically complete*.

Simple schedulers can be partitioned in two sets, namely *friendly simple* and *AAMZ simple* schedulers. An easy example of a friendly simple scheduler is as follows. Suppose label a is in L_s for all $s \in S$. A scheduler that always chooses a with probability 1 is friendly and simple. As an example of an AAMZ let D be a scheduler and $\mu_1, \mu_2 \in \text{Distr}(L)$. When D schedules for all σ where $t_{|\sigma|}$ is an irrational number μ_1, and μ_2 otherwise, then D is simple and an AAMZ scheduler. Simplicity follows from the fact that D is only supposed to decide between μ_1 and μ_2. There exists no path σ such that Equation 6 can be fulfilled for D, thus D is an unfriendly scheduler and by definition an AAMZ scheduler.

6 Deconstructing timed path probability

In Section 4 a measure Pr_D for finite and infinite paths was constructed from a measurable scheduler D. The goal now is to recover a measurable scheduler from an arbitrary probability measure Pr^ω (note there is no subscript D indicating the scheduler generating it) on the infinite timed path σ-algebra. This shows that measurable schedulers are sufficient to generate all quantitative behaviors of a CTMDP.

Given a CTMDP not every $Pr^\omega \in \text{Distr}(Path^\omega)$ is related or *compatible* with it. For example in the CTMDP $\mathcal{C}$ below, a probability measure Pr^ω such that $Pr^\omega(s_0 a[0,1] s_1 a[0,1] s_1) > 0$ is not compatible with $\mathcal{C}$. We postpone the definition of compatibility until we settle down some continuous space measure theoretic results.

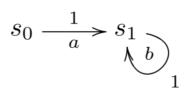

For a given timed path σ finishing at s, we are aiming at identifying two independent random sources: the scheduler selecting a label after history σ, and the sojourn time probability induced by the rate matrix.

Given Pr^ω, the tool to compute this product probability is a *conditional probability*, because we want to know the probability of the event $C_{\sigma a \mathcal{I} s}$ given that event C_σ happened.

$$Pr^\omega(C_{\sigma a \mathcal{I} s'} \mid C_\sigma) = \frac{Pr^\omega(C_{\sigma a \mathcal{I} s'} \cap C_\sigma)}{Pr^\omega(C_\sigma)} = D(\sigma, \{a\}) \cdot \mathbf{R}(s, a, s') \int_{\mathcal{I}} e^{-\mathbf{R}(s,a,S)\cdot t} dt \; . \quad (7)$$

For continuous systems like CTMDPs the numerator will usually be zero, so we resort to the continuous version of conditional probability the *Radon-Nikodym derivative* [16]. For this, some definitions are needed.

Definition 8. *The* marginal *of* Pr^ω *for the first* n *steps or coordinates is* $Pr^n \in Distr(Path^n)$ *where* $Pr^n(\Lambda) \doteq Pr^\omega(\Lambda \times (L \times \mathbb{R}^+ \times S)^\omega) = Pr^\omega(C_\Lambda)$ *and* $|\Lambda| = n$.

Definition 9. *Given two probability measures* $\mu, \nu \in Distr(X)$*, we say that* μ is absolutely continuous with respect to ν*, notation* $\mu \ll \nu$*, if* $\forall A \in \mathcal{F}_X,\ \nu(A) = 0 \Rightarrow \mu(A) = 0$.

It is straightforward to see $Pr^{n+1}(\cdot \times M) \ll Pr^n$, $\forall M \in \mathcal{F}_{L\times\mathbb{R}^+\times S}$, and this is the condition to apply Radon-Nikodym [14, Theorem 2.2.1] to obtain *(continuous) conditional probability* $f^n : Path^n \times \mathcal{F}_{L\times\mathbb{R}^+\times S} \to [0,1]$, such that $\forall \Lambda \in \mathcal{F}_{Path^n}, M \in \mathcal{F}_{L\times\mathbb{R}^+\times S}$:

1. $f^n(\cdot, M) : Path^n \to [0,1]$ is measurable,
2. $Pr^{n+1}(\Lambda \times M) = \int_\Lambda f^n(\sigma, M) dPr^n(\sigma)$.

Note that conditional probability is defined up-to sets of Pr^n-measure 0, therefore there are (potentially) infinite *versions* of it. This object shares with transition probability the measurability in its first argument. However if the measurable space is not restricted conveniently, it may be the case that from all versions of the conditional probability, none of them is a probability measure in its second argument [17, Section 6.4] [18, Problem 33.13]. Note this would be inconvenient for our purposes since equation (7) extends straightforwardly to a probability measure for all timed paths σ.

In our particular measure space (product of discrete spaces and positive reals), transition probabilities exists among conditional probabilities, and according to [13,19] we use *disintegration.* Namely Pr^ω can be disintegrated into transition probabilities, and this is where the compatibility with CTMDP can be defined.

Lemma 6 (Disintegration). *Given a probability measure* $Pr^\omega \in Distr(Path^\omega)$*, then there is a transition probability* $\mu : Path^* \times \mathcal{F}_{L\times\mathbb{R}^+\times S} \to [0,1]$ *unique* Pr^ω*-almost everywhere, that generates this probability measure.*

Proof. Using the same arguments of conditional probability we have for each $M \in \mathcal{F}_{L\times\mathbb{R}^+\times S}$ there is a measurable function $\mu^n(\cdot, M) : Path^n \to [0,1]$ such that $\forall \Lambda \in \mathcal{F}_{Path^n}$:

$$Pr^{n+1}(\Lambda \times M) = \int_\Lambda \mu^n(\sigma, M) dPr^n(\sigma) \tag{8}$$

From all the versions μ^n has, there is transition probability if the underlying space is *analytic* [20, Theorem 2.2] (the existence argument is over *regular conditional probabilities* but this problem is equivalent to transition probabilities [19, Theorem 3.1]). Examples of analytic spaces includes the discrete ones and the reals [9], and adding the fact that analytic spaces are closed under finite and denumerable product, we conclude $\mathcal{F}_{Path^n}$ is a space where we can choose a particular version of μ^n being a transition probability. The union $\cup_{0\leq n}\mu^n$ is denoted as $\mu : Path^* \times \mathcal{F}_{L\times\mathbb{R}^+\times S} \to [0,1]$ and gives the transition dynamics of the CTMDP.

Finally using $Pr^0 \in$ Distr(S) and μ, by [14, Theorem 2.6.7] expression (8) extends uniquely to the marginals Pr^n. As $Pr^\omega(C_\Lambda) = Pr^n(\Lambda)$, by [14, Theorem 2.7.2]

Pr^ω is the unique extension of this marginals to the infinite timed path probability measure, concluding that Pr^0 and μ generate Pr^ω. □

Having the transition probabilities underlying Pr^ω, we can state precisely what *compatible* means.

Definition 10 (Compatibility). *Given a CTMDP $\mathcal{C}$, we say probability measure Pr^ω on $\mathcal{F}_{Path^\omega}$ is* compatible *with $\mathcal{C}$ if: i) $Pr^0 = \mathcal{C}_{\mathrm{Pr}_{init}}$, ii) there is some scheduler D and transition probability μ from Lemma 6 that satisfies:*

$$\mu(\sigma, \{a\} \times \mathcal{I} \times \{s'\}) = D(\sigma, \{a\}) \cdot \mathbf{R}(last(\sigma), a, s') \int_{\mathcal{I}} e^{-\mathbf{R}(last(\sigma),a,S)\cdot t} dt \ . \quad (9)$$

It is clear that Definition 4 satisfies (9), and as both are additive in its second argument's first and third component, it is easy to show that $\mu(\sigma, A \times \mathcal{I} \times Q) = \mu_D(\sigma, A \times \mathcal{I} \times Q)$. Using Lemma 1, the main theorem follows.

Theorem 3 (Completeness). *Given a probability measure Pr^ω over $\mathcal{F}_{Path^\omega}$ compatible with CTMDP $\mathcal{C}$, then the underlying transition probability is generated by some measurable scheduler.*

Following the discussion of Section 5 it can be observed that given a positive probability measure Pr^ω for, e. g., C_Λ with $|\Lambda| \geq 2$, yields a friendly scheduler that generates the underlying transition probability. In contrast, when Pr^ω is almost always zero in the sense of Definition 6 the obtained scheduler falls in the class of AAMZ schedulers.

7 Conclusion

This work studies CTMDPs where external nondeterminism is resolved through timed history-dependent randomized schedulers. Taking time into account is the key difference that makes continuity and its related problems unavoidable, and this is where measurability issues arise.

We have obtained a complete characterization of those measurable schedulers, first using them to construct (integrate) a probability measure on timed paths, and later to deconstruct (disintegrate) an arbitrary probability measure on timed paths compatible with the CTMDP into a measurable scheduler. We showed that friendly schedulers generate positive probability measures and, vice versa, deconstructing positive measures yields friendly schedulers. We also showed an easily definable and probabilistically complete subclass of measurable schedulers.

As mentioned in Section 2, most of the results presented in this paper, however, do also carry over to less restrictive systems, e. g., stochastic transition systems [9]. We leave the consideration of a similar characterization of schedulers and the precise adaption of definitions and theorems to future work.

Acknowledgements. We would like to thank Holger Hermanns and Pedro D'Argenio for their inspiring ideas and helpful comments. We are also indebted to the anonymous referees for their valuable comments.

References

1. Bertsekas, D.P.: Dynamic Programming and Optimal Control. Volume II. Athena Scientific (1995)
2. Feinberg, E.A.: Continuous Time Discounted Jump Markov Decision Processes: A Discrete-Event Approach. Mathematics of Operations Research **29** (2004) 492–524
3. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley (1994)
4. Sennot, L.: Stochastic Dynamic Programming and the Control of Queueing Systems. John Wiley & Sons (1999)
5. Abdeddaïm, Y., Asarin, E., Maler, O.: On Optimal Scheduling under Uncertainty. TACAS (2003) 240–253
6. Bruno, J., Downey, P., Frederickson, G.N.: Sequencing Tasks with Exponential Service Times to Minimize the Expected Flow Time or Makespan. J. ACM **28**(1) (1981) 100–113
7. Qiu, Q., Pedram., M.: Dynamic power managment based on continuous-time Markov decision processes. In: Proceedings DAC. (1999) 555 – 561
8. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Efficient computation of time-bounded reachability probabilities in uniform continuous-time markov decision processes. In: TACAS: International Workshop on Tools and Algorithms for the Construction and Analysis of Systems, LNCS. (2004)
9. Cattani, S., Segala, R., Kwiatkowska, M.Z., Norman, G.: Stochastic Transition Systems for Continuous State Spaces and Non-determinism. In: FoSSaCS. (2005) 125–139
10. Kwiatkowska, M., Norman, G., Segala, R., Sproston, J.: Verifying quantitative properties of continuous probabilistic timed automata. In Palamidessi, C., ed.: Proc. CONCUR 2000 - Concurrency Theory. Volume 1877 of Lecture Notes in Computer Science., Springer (2000) 123–137
11. O.Hernndez-Lerma, Lassere, J.B.: Discrete-time Markov control processes: Basic optimality criteria. Volume 30 of Appl. Math. Springer-Verlag, Berlin and New York (1996)
12. K.Grabiszewski: Type space with disintegrability. Draft (2005)
13. Valadier, M.: Désintégration d'une mesure sur un produit. C.R. Acad. Sc. Paris **276**(1) (1973) 33–35 Serie A.
14. Ash, R.B., Doléans-Dade, C.A.: Probability & Measure Theory. second edn. Academic Press (2000)
15. Shiryaev, A.N.: Probability. second edn. Springer (1995)
16. Panangaden, P.: Measure and probability for concurrency theorists. TCS: Theoretical Computer Science **253** (2001)
17. Panangaden, P.: Stochastic Techniques in Concurrency. Unpublished Lecture Notes from a course given at BRICS (1997)
18. Billingsley, P.: Probability and Measure. Wiley, New York, NY (1986) Second Edition.
19. Jr., D.L., Fragoso, M., Ruffino, P.: Regular conditional probability, disintegration of probability and radon spaces. Proyecciones **23**(1) (2004) 15–29
20. Edalat, A.: Semi-pullbacks and bisimulation in categories of markov processes. Mathematical Structures in Computer Science **9**(5) (1999) 523–543