
Finite-state Controllers of POMDPs via Parameter Synthesis*

Sebastian Junges¹, Nils Jansen², Ralf Wimmer³, Tim Quatmann¹,
Leonore Winterer³, Joost-Pieter Katoen¹, and Bernd Becker³

¹RWTH Aachen University, Aachen, Germany

²Radboud University, Nijmegen, The Netherlands

³Albert-Ludwigs-Universität Freiburg, Freiburg im Breisgau, Germany

Abstract

We study finite-state controllers (FSCs) for partially observable Markov decision processes (POMDPs) that are provably correct with respect to given specifications. The key insight is that computing (randomised) FSCs on POMDPs is equivalent to—and computationally as hard as—synthesis for parametric Markov chains (pMCs). This correspondence allows to use tools for synthesis in pMCs to compute correct-by-construction FSCs on POMDPs for a variety of specifications. Our experimental evaluation shows comparable performance to well-known POMDP solvers.

1 INTRODUCTION

Partially Observable MDPs. We intend to provide guarantees for planning scenarios given by dynamical systems with uncertainties. In particular, we want to synthesise a *strategy* for an agent that ensures certain desired behaviour (Howard, 1960). A popular formal model for planning subject to stochastic behaviour are Markov decision processes (MDPs) (Puterman, 1994). An MDP is a nondeterministic model in which the agent chooses to perform an action under full knowledge of the environment it is operating in. The outcome of the action is a probability distribution over the system states. Many applications, however, allow only *partial observability* of the current system state (Kaelbling et al., 1998; Thrun et al., 2005; Wongpiromsarn and Frazzoli, 2012; Russell and Norvig, 2010). For such applications, MDPs are extended to *partially observable Markov decision processes* (POMDPs). While the agent acts within the environment, it encounters certain *observations*, according to which it

can infer the likelihood of the system being in a certain state. This likelihood is called the *belief state*. Executing an action leads to an update of the belief state according to new observations. The belief state together with an update function form a (typically uncountably infinite) MDP, referred to as the *belief MDP* (Shani et al., 2013).

The POMDP Synthesis Problem. For (PO)MDPs, a *randomised strategy* is a function that resolves the non-determinism by providing a probability distribution over actions at each time step. In general, strategies depend on the full history of the current evolution of the (PO)MDP. If a strategy depends only on the current state of the system, it is called *memoryless*. For MDPs, memoryless strategies suffice to induce optimal values according to our measures of interest (Puterman, 1994). Contrarily, POMDPs require strategies taking the full observation history into account (Ross, 1983), e. g. in case of infinite-horizon objectives. Moreover, strategies inducing *optimal* values are computed by assessing the entire belief MDP (Madani et al., 1999; Braziunas, 2003; Szer and Charpillat, 2005; Norman et al., 2017), rendering the problem undecidable (Chatterjee et al., 2016c).

POMDP strategies can be represented by *infinite-state controllers*. For computational tractability, strategies are often restricted to finite memory; this amounts to using *randomised finite-state controllers* (FSCs) (Meuleau et al., 1999). We often refer to strategies as FSCs. Already the computation of a memoryless strategy adhering to a specification is NP-hard, SQRT-SUM-hard, and in PSPACE (Vlassis et al., 2012). While optimal values cannot be guaranteed, small memory in combination with *randomisation* may supersede large memory in many cases (Chatterjee et al., 2004; Amato et al., 2010).

Correct-by-Construction Strategy Computation. In this paper, we synthesise FSCs for POMDPs. We require these FSCs to be provably correct for specifications such as indefinite-horizon properties like expected reward or reach-avoid probabilities. State-of-the-art POMDP

*Supported by the DFG RTG 2236 “UnRAVeL”.

Table 1: Correspondence

POMDP under FSC	pMC
states \times memory	states
same observation	same parameter
strategy	parameter instantiation

solvers mainly consider expected discounted reward measures (Walraven and Spaan, 2017), which are a subclass of indefinite horizon properties (Kolobov et al., 2012).

Our key observation is that for a POMDP the *set of all FSCs* with a fixed memory bound can be succinctly represented by a *parametric Markov chain* (pMC) (Daws, 2004). Transitions of pMCs are given by functions over a finite set of parameters rather than constant probabilities. The *parameter synthesis* problem for pMCs is to determine parameter instantiations that satisfy (or refute) a given specification. We show that the pMC parameter synthesis problem and the POMDP strategy synthesis problem are equally hard. This correspondence not only yields complexity results (Hutschenreiter et al., 2017), but particularly enables using a plethora of methods for parameter synthesis implemented in sophisticated and optimised parameter synthesis tools like PARAM (Hahn et al., 2010), PRISM (Kwiatkowska et al., 2011), and PROPHESY (Dehnert et al., 2015). They turn out to be competitive alternatives to dedicated POMDP solvers. Moreover, as we are solving slightly different problems, our methods are orthogonal to, e. g., PRISM-POMDP (Norman et al., 2017) and solve-POMDP (Walraven and Spaan, 2017).

We detail our contributions and the structure of the paper, which starts with necessary formalisms in Sect. 2. A longer version of the paper (Junges et al., 2017) contains some additional material.

Section 3: We establish the correspondence of POMDPs and pMCs, see Tab. 1. The product of a POMDP and an FSC yields a POMDP with state-memory pairs, which we map to states in the pMC. If POMDP states share *observations*, the corresponding pMC states share *parameters* at emanating transitions. A *strategy* of the POMDP corresponds to a *parameter instantiation* in the pMC.

Section 4: We show the opposite direction, namely a transformation from pMCs to POMDPs. This result establishes that the synthesis problems for POMDPs and pMCs are equally hard. Technically, we identify the practically relevant class of *simple pMCs*, which coincides with POMDPs under memoryless strategies.

Section 5: Typical restrictions on parameter instantiations concern whether parameters may be assigned the probability zero. We discuss effects of such restrictions to the resulting POMDP strategies.

Section 6: We evaluate the computation of correct-by-

construction FSCs using pMC synthesis techniques. To that end, we explain how particular parameter synthesis approaches deliver optimal or near-optimal FSCs. Then, we evaluate the approach on a range of typical POMDP benchmarks. We observe that often a small amount of memory suffices. Our approach is competitive to state-of-the-art POMDP solvers and is able to synthesise small, almost-optimal FSCs.

Related Work. In addition to the cited works, (Meuleau et al., 1999) uses a branch-&-bound method to find optimal FSCs for POMDPs. A SAT-based approach computes FSCs for qualitative properties (Chatterjee et al., 2016a). For a survey of decidability results and algorithms for broader classes of properties refer to (Chatterjee et al., 2016c,b). Work on parameter synthesis (Hutschenreiter et al., 2017; Filieri et al., 2011) might contain additions to the methods considered here.

2 PRELIMINARIES

A *probability distribution* over a finite or countably infinite set X is a function $\mu: X \rightarrow [0, 1] \subseteq \mathbb{R}$ with $\sum_{x \in X} \mu(x) = \mu(X) = 1$. The set of all distributions on X is $Distr(X)$. The support of a distribution μ is $\text{supp}(\mu) = \{x \in X \mid \mu(x) > 0\}$. A distribution is *Dirac* if $|\text{supp}(\mu)| = 1$.

Let $V = \{p_1, \dots, p_n\}$ be a finite set of *parameters* over the domain \mathbb{R} and let $\mathbb{Q}[V]$ be the set of multivariate polynomials over V . An *instantiation* for V is a function $u: V \rightarrow \mathbb{R}$. Replacing each parameter p in a polynomial $f \in \mathbb{Q}[V]$ by $u(p)$ yields $f[u] \in \mathbb{R}$.

Decision problems can be considered as languages describing all positive instances. A language $L_1 \subseteq \{0, 1\}^*$ is *polynomial (many-one or Karp) reducible* to $L_2 \subseteq \{0, 1\}^*$, written $L_1 \leq_P L_2$, if there exists a polynomial-time computable function $f: \{0, 1\}^* \rightarrow \{0, 1\}^*$ such that for all $w \in \{0, 1\}^*$, $w \in L_1 \iff f(w) \in L_2$. Polynomial reductions are essential to define complexity classes, cf. (Papadimitriou, 1994).

2.1 PARAMETRIC MARKOV MODELS

Definition 1 (pMDP) A parametric Markov decision process (*pMDP*) M is a tuple $M = (S, s_1, Act, V, \mathcal{P})$ with a finite (or countably infinite) set S of states, initial state $s_1 \in S$, a finite set Act of actions, a finite set V of parameters, and a transition function $\mathcal{P}: S \times Act \times S \rightarrow \mathbb{Q}[V]$.

The *available actions* in $s \in S$ are $A(s) = \{a \in Act \mid \exists s' \in S : \mathcal{P}(s, a, s') \neq 0\}$. W.l.o.g. we assume $\forall s, s' \in$

$S. \forall a \in Act. P(s, a, s') \neq 0 \wedge P(s, a, s'') \neq 1 \Rightarrow \exists s'' \neq s'. P(s, a, s'') \neq 0$. We assume that pMDP M contains no deadlock states, i. e. $A(s) \neq \emptyset$ for all $s \in S$. A path of a pMDP M is an (in)finite sequence $\pi = s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \dots$, where $s_0 = s_I$, $s_i \in S$, $a_i \in A(s_i)$, and $\mathcal{P}(s_i, a_i, s_{i+1}) \neq 0$ for all $i \in \mathbb{N}$. For finite π , $\text{last}(\pi)$ denotes the last state of π . The set of (in)finite paths of M is Paths_{fin}^M (Paths^M).

Definition 2 (MDP) A Markov decision process (MDP) is a pMDP where $\mathcal{P}: S \times Act \times S \rightarrow [0, 1] \subseteq \mathbb{R}$ and for all $s \in S$ and $a \in A(s)$, $\sum_{s' \in S} \mathcal{P}(s, a, s') = 1$.

A (parametric) discrete-time Markov chain ((p)MC) is a (p)MDP with $|A(s)| = 1$ for all $s \in S$. For a (p)MC D , we may omit the actions and use the notation $D = (S, s_I, V, P)$ with a transition function P of the form $P: S \times S \rightarrow \mathbb{Q}[V]$.

Applying an instantiation $u: V \rightarrow \mathbb{R}$ to a pMDP or pMC M , denoted $M[u]$, replaces each polynomial f in M by $f[u]$. $M[u]$ is also called the instantiation of M at u . Instantiation u is well-defined for M if the replacement yields probability distributions, i. e. if $M[u]$ is an MDP or an MC, respectively.

Strategies. To resolve the nondeterministic action choices in MDPs, so-called strategies determine at each state a distribution over actions to take. This decision may be based on the history of the current path.

Definition 3 (Strategy) A strategy σ for (p)MDP M is a function $\sigma: \text{Paths}_{fin}^M \rightarrow \text{Distr}(Act)$ s. t. $\text{supp}(\sigma(\pi)) \subseteq Act(\text{last}(\pi))$ for all $\pi \in \text{Paths}_{fin}^M$. The set of all strategies of M is Σ^M .

A strategy σ is memoryless if $\text{last}(\pi) = \text{last}(\pi')$ implies $\sigma(\pi) = \sigma(\pi')$ for all $\pi, \pi' \in \text{Paths}_{fin}^M$. It is deterministic if $\sigma(\pi)$ is a Dirac distribution for all $\pi \in \text{Paths}_{fin}^M$. A strategy that is not deterministic is randomised.

A strategy σ for an MDP M resolves all nondeterministic choices, yielding an induced Markov chain M^σ , for which a probability measure over infinite paths is defined by the cylinder set construction (Baier and Katoen, 2008).

Definition 4 (Induced Markov Chain) For an MDP $M = (S, s_I, Act, \mathcal{P})$ and a strategy $\sigma \in \Sigma^M$, the MC induced by M and σ is given by $M^\sigma = (\text{Paths}_{fin}^M, s_I, P^\sigma)$ where:

$$P^\sigma(\pi, \pi') = \begin{cases} \mathcal{P}(\text{last}(\pi), a, s') \cdot \sigma(\pi)(a) & \text{if } \pi' = \pi a s' \\ 0 & \text{otherwise.} \end{cases}$$

2.2 PARTIAL OBSERVABILITY

Definition 5 (POMDP) A partially observable MDP (POMDP) is a tuple $\mathcal{M} = (M, Z, O)$, with $M = (S, s_I, Act, \mathcal{P})$ the underlying MDP of \mathcal{M} , Z a finite set of observations and $O: S \rightarrow Z$ the observation function.

We require that states with the same observations have the same set of enabled actions, i. e. $O(s) = O(s')$ implies $A(s) = A(s')$ for all $s, s' \in S$. We define $A(z) = A(s)$ if $O(s) = z$. More general observation functions (Roy et al., 2005; Shani et al., 2013) take the last action into account and provide a distribution over Z . There is a transformation of the general case to the POMDP definition used here that blows up the state space polynomially (Chatterjee et al., 2016b). In Fig. 1(a), a fragment of the underlying MDP of a POMDP has two different observations, indicated by the state colouring.

We lift the observation function to paths: For $\pi = s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \dots s_n \in \text{Paths}_{fin}^M$, the associated observation sequence is $O(\pi) = O(s_0) \xrightarrow{a_0} O(s_1) \xrightarrow{a_1} \dots O(s_n)$. Several paths in the underlying MDP may yield the same observation sequence. Strategies have to take this restricted observability into account.

Definition 6 An observation-based strategy σ for a POMDP \mathcal{M} is a strategy for the underlying MDP M such that $\sigma(\pi) = \sigma(\pi')$ for all $\pi, \pi' \in \text{Paths}_{fin}^M$ with $O(\pi) = O(\pi')$. $\Sigma^{\mathcal{M}}$ is the set of observation-based strategies for \mathcal{M} .

An observation-based strategy selects actions based on observations along a path and the past actions. Applying the strategy to a POMDP yields an induced MC as in Def. 4, resolving all nondeterminism and partial observability. To represent observation-based strategies with finite memory, we define finite-state controllers (FSCs). A randomised observation-based strategy for a POMDP \mathcal{M} with (finite) k memory is represented by an FSC \mathcal{A} with k memory nodes. If $k = 1$, the FSC describes a memoryless strategy. We often refer to observation-based strategies as FSCs.

Definition 7 (FSC) A finite-state controller (FSC) for a POMDP \mathcal{M} is a tuple $\mathcal{A} = (N, n_I, \gamma, \delta)$, where N is a finite set of memory nodes, $n_I \in N$ is the initial memory node, γ is the action mapping $\gamma: N \times Z \rightarrow \text{Distr}(Act)$, and δ is the memory update $\delta: N \times Z \times Act \rightarrow \text{Distr}(N)$. The set $\text{FSC}_k^{\mathcal{M}}$ denotes the set of FSCs with k memory nodes, called k -FSCs. Let $\sigma_{\mathcal{A}} \in \Sigma^{\mathcal{M}}$ denote the observation-based strategy represented by \mathcal{A} .

From a node n and the observation z in the current state of the POMDP, the next action a is chosen from $A(z)$

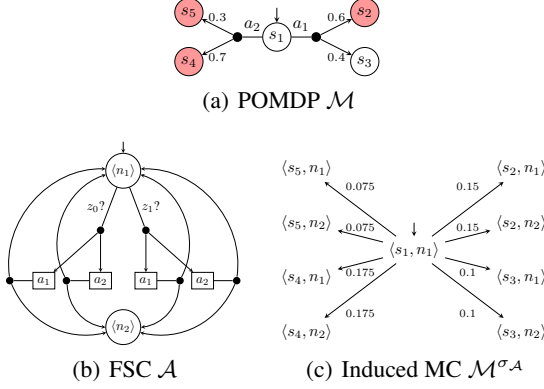


Figure 1: (a) The POMDP \mathcal{M} with observations $O(s_1) = O(s_3) = z_0$ (white) and $O(s_2) = O(s_4) = O(s_5) = z_1$ (red). (b) The associated (partial) FSC \mathcal{A} has two memory nodes. (c) A part of MC $\mathcal{M}^{\sigma\mathcal{A}}$ induced by \mathcal{M} and \mathcal{A} .


randomly as given by $\gamma(n, z)$. Then, the successor node of the FSC is determined randomly via $\delta(n, z, a)$.

Example 1 Fig. 1(b) shows an excerpt of an FSC \mathcal{A} with two memory nodes. From node n_1 , the action mapping distinguishes observations z_0 and z_1 . The solid dots indicate a probability distribution from $\text{Distr}(\text{Act})$. For readability, all distributions are uniform and we omit the action mapping for node n_2 .

Now recall the POMDP \mathcal{M} from Fig. 1(a). The induced MC $\mathcal{M}^{\sigma\mathcal{A}}$ is shown in Fig. 1(c). Assume \mathcal{M} is in state s_1 and \mathcal{A} in node n_1 . Based on the observation $z_0 := O(s_1)$, $\sigma_{\mathcal{A}}$ chooses action a_1 with probability $\delta(n_1, z_0)(a_1) = 0.5$ leading to the probabilistic branching in the POMDP. With probability 0.6, \mathcal{M} evolves to state s_2 . Next, the FSC \mathcal{A} updates its memory node; with probability $\delta(n_1, z_0, a_1)(n_1) = 0.5$, \mathcal{A} stays in n_1 . The corresponding transition from $\langle s_1, n_1 \rangle$ to $\langle s_2, n_1 \rangle$ in $\mathcal{M}^{\sigma\mathcal{A}}$ has probability $0.5 \cdot 0.6 \cdot 0.5 = 0.15$.

2.3 SPECIFICATIONS

For a POMDP \mathcal{M} , a set $G \subseteq S$ of goal states, a set $B \subseteq S$ of bad states, and a threshold $\lambda \in [0, 1)$, we consider quantitative reach-avoid specifications $\varphi = \mathbb{P}_{>\lambda}(\neg B \cup G)$. The specification φ is satisfied for a strategy $\sigma \in \Sigma^{\mathcal{M}}$ if the probability $\Pr^{\mathcal{M}^{\sigma}}(\neg B \cup G)$ of reaching a goal state in \mathcal{M}^{σ} without entering a bad state in between exceeds λ , denoted by $\mathcal{M}^{\sigma} \models \varphi$. The task is to compute such a strategy provided that one exists. For an MDP M , there is a memoryless deterministic strategy inducing the maximal probability $\Pr_{\max}^M(\neg B \cup G)$ (Condon, 1992). For a POMDP \mathcal{M} , however, observation-based strategies with infinite memory as in Def. 6 are



Act	\mathcal{P}	Node	Result
$a_1 : p$	0.6	$n_1 : q_1$	$0.6 \cdot p \cdot q_1$
		$n_2 : 1 - q_1$	$0.6 \cdot p \cdot (1 - q_1)$
	0.4	$n_1 : q_1$	$0.4 \cdot p \cdot q_1$
		$n_2 : 1 - q_1$	$0.4 \cdot p \cdot (1 - q_1)$
$a_2 : 1 - p$	0.7	$n_1 : q_2$	$0.7 \cdot (1 - p) \cdot q_2$
		$n_2 : 1 - q_2$	$0.7 \cdot (1 - p) \cdot (1 - q_2)$
	0.3	$n_1 : q_2$	$0.3 \cdot (1 - p) \cdot q_2$
		$n_2 : 1 - q_2$	$0.3 \cdot (1 - p) \cdot (1 - q_2)$

(b) Parameterised transition probabilities

Figure 2: Induced parametric Markov chain for FSCs.

necessary (Ross, 1983) to attain $\Pr_{\max}^{\mathcal{M}}(\neg B \cup G)$. The problem of proving the satisfaction of φ is therefore undecidable (Chatterjee et al., 2016c). In our experiments, we also use *undiscounted expected reachability reward properties* (Baier and Katoen, 2008).

3 FROM POMDPs TO PMCS

Our goal is to make pMC synthesis methods available for POMDPs. In this section we provide a transformation from a POMDP \mathcal{M} to a pMC D . We consider the following decision problems.

Problem 1 ($\exists k$ -FSC) Given a POMDP \mathcal{M} , a specification φ , and a (unary encoded) memory bound $k > 0$, is there a k -FSC \mathcal{A} with $\mathcal{M}^{\sigma\mathcal{A}} \models \varphi$?

Problem 2 ($\exists \text{INST}$) For a pMC D and a specification φ , does a well-defined instantiation u exist s.t. $D[u] \models \varphi$?

Theorem 1 $\exists k\text{-FSC} \leq_P \exists \text{INST}$.

The remainder of the section outlines the proof, the converse direction is addressed in Sect. 4. Consider a POMDP \mathcal{M} , a specification φ , and a memory bound $k > 0$ for which $\exists k$ -FSC is to be solved. The degrees of freedom to select a k -FSC are given by the possible choices for γ and δ . For each γ and δ , we get a different induced MC, but these MCs are *structurally similar* and can be represented by a single pMC.

Example 2 Recall Fig. 1 and Ex. 1. The action mapping γ and the memory update δ have arbitrary but fixed

probability distributions. For a_1 , we represent the probability $\gamma(n_1, z_0)(a_1) =: p$ by $p \in [0, 1]$. The memory update yields $\delta(n_1, z_0, a_1)(n_1) =: q_1 \in [0, 1]$ and $\delta(n_1, z_0, a_1)(n_2) =: 1 - q_1$, respectively. Fig. 2(a) shows the induced pMC for action choice a_1 . For instance, the transition from $\langle s_1, n_1 \rangle$ to $\langle s_2, n_1 \rangle$ is labelled with polynomial $p \cdot 0.6 \cdot q_1$.

We collect all polynomials for observation z_0 in Fig. 2(b). The result column describes a parameterised distribution over tuples of states and memory nodes. Thus, instantiations for these polynomials need to sum up to one.

As the next step, we define the pMC that results from combining a k -FSC with a POMDP. The idea is to assign parameters as arbitrary probabilities to action choices. Each observation has one *remaining action* given by a mapping $Remain: Z \rightarrow Act$. $Remain(z) \in A(z)$ is the action to which, after choosing probabilities for all other actions in $A(z)$, the remaining probability is assigned. A similar principle holds for the remaining memory node.

Definition 8 (Induced pMC for a k -FSC on POMDPs)
Let $\mathcal{M} = (M, Z, O)$ be a POMDP with $M = (S, s_I, Act, \mathcal{P})$ and let $k > 0$ be a memory bound. The induced pMC $D_{\mathcal{M},k} = (S_{\mathcal{M},k}, s_{I,\mathcal{M},k}, V_{\mathcal{M},k}, P_{\mathcal{M},k})$ is defined by:

- $S_{\mathcal{M},k} = S \times \{0, \dots, k-1\}$, $s_{I,\mathcal{M},k} = \langle s_I, 0 \rangle$,
- $V_{\mathcal{M},k} = \{p_a^{z,n} \mid z \in Z, n \in \{0, \dots, k-1\}, a \in A(z), a \neq Remain(z)\} \cup \{q_{a,n'}^{z,n} \mid z \in Z, n, n' \in \{0, \dots, k-1\}, n' \neq k-1, a \in A(z)\}$,
- $P_{\mathcal{M},k}(s, s') = \sum_{a \in A(s)} H(s, s', a)$ for all $s, s' \in S'$,

where $H: S_{\mathcal{M},k} \times S_{\mathcal{M},k} \times Act \rightarrow \mathbb{R}$ is for $z = O(s)$ defined by $H(\langle s, n \rangle, \langle s', n' \rangle, a) =$

$$\mathcal{P}(s, a, s') \cdot \left\{ \begin{array}{ll} p_a^{z,n}, & \text{if } a \neq Remain(z) \\ 1 - \sum_{b \neq a} p_b^{z,n}, & \text{if } a = Remain(z) \end{array} \right\} \cdot \left\{ \begin{array}{ll} q_{a,n'}^{z,n}, & \text{if } n' \neq k-1 \\ 1 - \sum_{\bar{n} \neq n'} q_{a,\bar{n}}^{z,n}, & \text{if } n' = k-1 \end{array} \right\}.$$

Intuitively, $H(s, s', a)$ describes the probability mass from s to s' in the induced pMC that is contributed by action a . The three terms correspond to the terms as seen in the first three columns of Tab. 2(b).

Example 3 Consider the POMDP in Fig. 3(a) and let $k = 1$. The induced pMC is given in Fig. 3(b). The three actions from s_0 have probability p_1 , p_2 , and $1 - p_1 - p_2$ for the remaining action a_3 . From the indistinguishable states s_1, s_3 , actions have probability q and $1 - q$, respectively.




Figure 3: From POMDPs to pMCs ($k = 1$)

By construction, the induced pMC describes the set of all induced MCs:

Theorem 2 (Correspondence Theorem) For POMDP \mathcal{M} , memory bound k , and the induced pMC $D_{\mathcal{M},k}$:

$$\{D_{\mathcal{M},k}[u] \mid u \text{ well-defined}\} = \{\mathcal{M}^{\sigma_A} \mid \mathcal{A} \in FSC_k^{\mathcal{M}}\}.$$

In particular, every well-defined instantiation u describes an FSC $\mathcal{A}_u \in FSC_k^{\mathcal{M}}$.

By the correspondence, we can thus evaluate an instantiation of the induced pMC to assess whether the corresponding k -FSC satisfies a given specification.

Corollary 1 Given an induced pMC $D_{\mathcal{M},k}$ and a specification φ : For every well-defined instantiation u of $D_{\mathcal{M},k}$ and the corresponding k -FSC \mathcal{A}_u we have:

$$\mathcal{M}^{\sigma_{\mathcal{A}_u}} \models \varphi \iff D_{\mathcal{M},k}[u] \models \varphi.$$

The number of parameters in the induced pMC $D_{\mathcal{M},k}$ is given by $\mathcal{O}(|Z| \cdot k^2 \cdot \max_{z \in Z} |A(z)|)$.

4 FROM PMCS TO POMDPS (AND BACK AGAIN)

In the previous section we have shown that $\exists k$ -FSC is at least as hard as $\exists INST$. We now discuss whether both problems are equally hard: The open question is whether we can reduce $\exists INST$ to $\exists k$ -FSC.

A straightforward reduction maintains the states of the pMC in the POMDP, or even yields a POMDP with the same graph structure (the topology) as the pMC. The next example shows that this naive reduction is impossible.

Example 4 In the pMC in Fig. 4 the parameter p occurs in two different distributions (at s_0 and s_2). For defining a reduction where the resulting POMDP has the same set of states, there are two options for the observation




Figure 4: Non-simple pMC

function at the states s_0 and s_2 : Either $O(s_0) = O(s_2)$ or $O(s_0) \neq O(s_2)$. The intuition is that every (parametric) transition in the pMC corresponds to an action choice in a POMDP. Then $O(s_0) = O(s_2)$ is impossible as s_0 and s_2 have a different number of outgoing transitions (outdegree). Adding a self-loop to s_2 does not alleviate the problem. Moreover, $O(s_0) \neq O(s_2)$ is impossible, as a strategy could distinguish s_0 and s_2 and assign different probabilities to p .

The pMC in the example is problematic as the parameters occur at the outgoing transitions of states in different combinations. We restrict ourselves to an important subclass¹ of pMCs which we call *simple pMCs*. A pMC is simple if for all states $s, s', P(s, s') \in \mathbb{Q} \cup \{p, 1-p \mid p \in V\}$. Consequently, we can map states to parameters, and use this map to define the observations. Then, the transformation from a POMDP to a pMC is the reverse of the transformation from Def. 8. In the remainder, we detail this correspondence. The correspondence also establishes a construction to compute k -FSCs via parameter synthesis on *simple* pMCs. Current tool-support (cf. Sect. 6) for simple pMCs is more mature than for the more general pMCs obtained via Def. 8.

Let *simple- \exists INST* be the restriction of \exists INST to simple pMCs. Similarly, let *simple- \exists 1-FSC* be a variant of \exists 1-FSC that only considers *simple* POMDPs.

Definition 9 (Binary/Simple POMDP) A POMDP is binary, if $|A(s)| \leq 2$ for all $s \in S$. A binary POMDP is simple, if for all $s \in S$


$$|A(s)| = 2 \implies \forall a \in A(s) \exists s' \in S : P(s, a, s') = 1.$$

We establish the following relation between the POMDP and pMC synthesis problems, which asserts that the problems are equivalently hard.


Theorem 3 For any $L_1, L_2 \in \{ \exists k\text{-FSC}, \exists 1\text{-FSC}, \text{simple-}\exists 1\text{-FSC}, \text{simple-}\exists \text{INST} \}$, $L_1 \leq_P L_2$.

The proof is a direct consequence of the Lemmas 1-4 below, as well as the facts that every 1-FSC is a k -FSC, and every simple POMDP is a POMDP.

¹All pMC benchmarks from the PARAM webpage (PARAM Website, 2015) are simple pMCs.



(a) Four actions in a POMDP (b) Four actions in a binary POMDP



(c) Binary POMDP (d) Simple POMDP (e) Simple pMC

Figure 5: POMDP \leftrightarrow simple pMC

The induced pMC $D_{\mathcal{M},1}$ of a simple POMDP \mathcal{M} is also simple. Consequently, Sect. 3 yields:

Lemma 1 *simple- \exists 1-FSC* \leq_P *simple- \exists INST*.

4.1 FROM SIMPLE PMCS TO POMDPS

Theorem 4 Every simple pMC D with n states and m parameters is isomorphic to $D_{\mathcal{M},1}$ for some simple POMDP \mathcal{M} with n states and m observations.

We refrain from a formal proof: The construction is the reverse of Def. 8, with observations $\{z_p \mid p \in V_D\}$. In a simple pMC, the outgoing transitions are either all parameter free, or of the form $p, 1-p$. The parameter-free case is transformed into a POMDP state with a single action (and any observation). The parametric case is transformed into a state with two actions with Dirac-distributions attached. As observation we use z_p .

Lemma 2 *simple- \exists INST* \leq_P *simple- \exists 1-FSC*.

4.2 MAKING POMDPS SIMPLE

We present a reduction from $\exists 1\text{-FSC}$ to *simple- \exists 1-FSC* by translating a (possibly not simple) POMDP into a *binary* POMDP and subsequently into a *simple* POMDP. Examples are given in Fig. 5(a-e). We emphasise that our construction only preserves the expressiveness of 1-FSCs. Details are given in (Junges et al., 2017).

Lemma 3 $\exists 1\text{-FSC} \leq_P \text{simple-}\exists 1\text{-FSC}$.




Figure 6: Unfolding a POMDP for two memory states

4.3 FROM K -FSCS TO 1-FSCS

For a POMDP \mathcal{M} and memory bound $k > 1$, we construct a POMDP \mathcal{M}_k such that \mathcal{M} satisfies a specification φ under some k -FSC iff \mathcal{M}_k satisfies φ under some 1-FSC.

Definition 10 (k -Unfolding) Let $\mathcal{M} = (M, Z, O)$ be a POMDP with $M = (S, s_I, Act, \mathcal{P})$, and $k > 1$. The k -unfolding of \mathcal{M} is the POMDP $\mathcal{M}_k = (M_k, Z_k, O_k)$ with $M_k = (S_k, s_{I,k}, Act_k, \mathcal{P}_k)$ defined by:

- $S_k = S \times \{0, \dots, k-1\}$, $s_{I,k} = \langle s_I, 0 \rangle$,
- $Act_k = Act \times \{0, \dots, k-1\}$
- $\mathcal{P}_k(\langle s, n \rangle, \langle a, \bar{n} \rangle, \langle s', n' \rangle) = \begin{cases} \mathcal{P}(s, a, s') & n' = \bar{n}, \\ 0 & \text{else.} \end{cases}$

and $Z_k = Z \times \{0, \dots, k-1\}$, $O_k(\langle s, n \rangle) = \langle O(s), n \rangle$.

Intuitively, \mathcal{M}_k stores the current memory node into its state space. At state $\langle s, n \rangle$ of \mathcal{M}_k , a 1-FSC can not only choose between the available actions $A(s)$ in \mathcal{M} but also between different successor memory nodes.

Fig. 6 shows this process for $k = 2$. All states of the POMDP are copied once. Different observations allow to determine in which copy of a state – and therefore, which memory cell – we currently are. Additionally, all actions are duplicated to model the option for a strategy to switch the memory cell.

The induced pMC $D_{\mathcal{M}_k,1}$ of the k -unfolding of \mathcal{M} has the same topology as the induced pMC $D_{\mathcal{M},k}$ of \mathcal{M} with memory bound k . In fact, both pMCs have the same instantiations.

Proposition 1 For POMDP \mathcal{M} and memory bound k :

$$\{D_{\mathcal{M}_k,1}[u] \mid u \text{ well-def.}\} = \{D_{\mathcal{M},k}[u] \mid u \text{ well-def.}\}.$$

The intuition is that in both pMCs the parameter instantiations reflect arbitrary probability distributions over the same set of successor states. In the transition probability function of the induced pMC $D_{\mathcal{M},k}$ of \mathcal{M} we can also substitute the multiplications of parameters $p_a^{z,n}$ and $q_a^{z,n}$,




Figure 7: MDP \mathcal{M}

by single parameters. This transformation yields a *substituted induced pMC* which is isomorphic to the induced pMC $D_{\mathcal{M}_k,1}$ of the k -unfolding of \mathcal{M} . Details are given in (Junges et al., 2017).

Proposition 1 and Thm. 2 imply that induced MCs of \mathcal{M} under k -FSCs coincide with induced MCs of \mathcal{M}_k under 1-FSCs: $\{\mathcal{M}^{\sigma_A} \mid \mathcal{A} \in FSC_k^{\mathcal{M}}\} = \{\mathcal{M}_k^{\sigma_A} \mid \mathcal{A} \in FSC_1^{\mathcal{M}_k}\}$.

Lemma 4 $\exists k\text{-FSC} \leq_P \exists 1\text{-FSC}$.

5 STRATEGY RESTRICTIONS

Two simplifying restrictions on the parameters are usually made in parameter synthesis for pMCs:

- Each transition is assigned a strictly positive probability (*graph-preserving*).
- Each transition is assigned at least probability $\varepsilon > 0$ (ε -*preserving*).

For simple pMCs, the restrictions correspond to parameters instantiations over $(0, 1)$ or $[\varepsilon, 1 - \varepsilon]$, respectively.

Accordingly, we define restrictions to POMDP strategies that correspond to such restricted parameter instantiations.

Definition 11 (Non-zero Strategies) A strategy σ is non-zero if $\sigma(\pi)(a) > 0$ for all $\pi \in \text{Paths}_{\min}^M$, $a \in A(\text{last}(\pi))$, and min- ε if additionally $\sigma(\pi)(a) \geq \varepsilon > 0$.

Non-zero strategies enforce $\text{supp}(\sigma(s)) = A(s)$. Example 5 shows the impact on reachability probabilities.

Example 5 The MDP M in Fig. 7 has a choice between actions a_1 and a_2 at state s_0 . If action a_1 is chosen with probability zero, the probability to reach s_1 from s_0 becomes zero, and the corresponding parameter instantiation is not graph-preserving. Contrarily, if a_1 is chosen with any positive probability, as would be enforced by a non-zero strategy, the probability to reach s_1 is one.

Proposition 2 Let \mathcal{M} be a POMDP. An instantiation u on $D_{\mathcal{M},1}$ is graph-preserving (ε -preserving), iff $\sigma_{\mathcal{A}_u}$ is non-zero (min- ε).

Still, for the considered specifications, we can, w. l. o. g., restrict ourselves to FSCs that induce non-zero strategies.

Theorem 5 Let \mathcal{M} be a POMDP, k a memory bound and $\varphi = \mathbb{P}_{>\lambda}(\neg B \cup G)$. Either $\forall \mathcal{A} \in k\text{-FSC} : \mathcal{M}^{\sigma_{\mathcal{A}}} \not\models \varphi$ or $\exists \mathcal{A}' \in k\text{-FSC} : \mathcal{M}^{\sigma_{\mathcal{A}'}} \models \varphi$ with $\sigma_{\mathcal{A}'}$ non-zero.

The statement is shown in (Junges et al., 2017) by considering the corresponding pMC.

6 EMPIRICAL EVALUATION

We established the correspondence between the synthesis problems for POMDPs and pMCs. Now, we discuss the available methods for pMC parameter synthesis, and how they may be exploited or adapted to synthesise FSCs. We distinguish three key problems:

1. *Find a correct-by-construction strategy for a POMDP and a specification.* To construct such a strategy, one needs to find a parameter valuation for the pMC that provably satisfies the specification. Most solution techniques focused on pMCs with a few parameters, rendering the problem at hand infeasible. Recently, efficient approaches emerged that are either based on particle swarm optimisation (PSO) (Chen et al., 2013) or on convex optimisation (Amato et al., 2010; Cubuktepe et al., 2017), in particular using quadratically-constrained quadratic programming (QCQP) (Cubuktepe et al., 2018). We employ PSO and QCQP for our evaluation.

2. *Prove that no FSC exists for a POMDP and a specification.* Proving the absence of an FSC with the given memory bound allows us to show ε -optimality of a previously synthesised strategy. Two approaches exist: An approximative technique called *parameter lifting* (Quatmann et al., 2016) and a method based on SAT-modulo-theories (SMT) solving (de Moura and Bjørner, 2008).

3. *Provide a closed-form solution for the underlying measure of a specification in form of a function over the induced parameters of an FSC.* The function may be used for further analysis, e. g. of the sensitivity of decisions or parameter values, respectively. To compute this function, all of the parameter synthesis tools PARAM (Hahn et al., 2010), PRISM (Kwiatkowska et al., 2011), Storm (Dehnert et al., 2017), and PROPhESY (Dehnert et al., 2015) employ a technique called *state elimination* (Daws, 2004).

Implementation and Setup. We extended the tool Storm (Dehnert et al., 2017) to parse and store POMDPs, and implemented several transformation options to pMCs. Most notably, Storm supports k -unfolding, the product with several restricted FSCs such as counters that can be incremented at will, and several types of transformation to (simple) pMCs.

We evaluate on a HP BL685C G7 with 48 2GHz cores, a 16GB memory limit, and 1800 seconds time

Table 2: Benchmarks

Id	Name	Tp.	POMDP \mathcal{M}			PRISM-POMDP		SolvePOMDP		MDP
			States	Bran.	Obs.	Result	Time	Result	Time	Res
1	NRP (8)	\mathbb{P}	125	161	41	[.125, .24]	20		TO	1.0
2	Grid (4)	\mathbb{E}	17	62	3	[3.97, 4.13]	1038	4.13	0.4	3.2
3	Netw (3,4,8)	\mathbb{E}	2729	4937	361				TO	0.83
4	Crypt (5)	\mathbb{P}	4885	11733	890				MO	1.0
5	Maze (2)	\mathbb{E}	16	58	8	[5.11, 5.23]	3.9	5.23	16	4.0
6	Load (8)	\mathbb{E}	16	28	5	[10.5, 10.5]	1356	10.5	7.6	10.5
7	Slippery (4)	\mathbb{P}	17	59	4			0.93	95	1.0

limit. The compared methods are single-threaded. We took the POMDPs from PRISM-POMDP (Norman et al., 2017), and additional maze, load/unload examples from (Meuleau et al., 1999), and a slippery gridworld with traps inspired by (Russell and Norvig, 2010). Table 2 gives details. The specifications (Tp.) are either the minimisation of expected costs from an initial state until reaching a specified target set (\mathbb{E}), or the maximisation the probability of reaching from an initial state a target set without hitting a bad state before (\mathbb{P}). We list the number of states, branches ($\sum |A(s)|$), and observations in each POMDP. As a baseline, we provide the results and run time of the model-checking tool PRISM-POMDP, and the point-based solver SolvePOMDP (Walraven and Spaan, 2017), obtained with default settings. Both tools compute optimal memory-unbounded strategies and are prototypes. The last column contains the result on the underlying, fully observable MDP. The experiments contain minimal expected rewards, which are analysed by a straightforward extension of maximal reachability probabilities. All pMCs computed are simple pMCs, as PROPhESY typically benefits from the simpler structure. PROPhESY has been invoked with the default set-up.

6.1 FINDING STRATEGIES

We evaluate how quickly a strategy that satisfies the specification is synthesised. We vary the threshold used in the specification, as well as the structure of the FSC.

Results. We summarise the obtained results in Tab. 3. For each instance (Id), we define three thresholds (Ts), ordered from challenging (i. e. close to the optimum) to less challenging. For different types of FSCs (FSC, F=full, C=counter) and memory bounds (k), we obtain pMCs with the given number of states, transitions and parameters. Full-FSCs are fully connected, in counter-FSCs memory state m is succeeded by either m or $m + 1$. For each threshold (T1, T2, T3), we report the run time of the two methods PSO and QCQP, respectively. T1 is chosen to be nearly optimal for all benchmarks. A dash indicates a combination of memory and threshold that is not realisable according to the results in Sect. 6.2. TO/MO denote violations of the time/memory limit, respectively.

Table 3: Synthesizing strategies

Id	Ts	FSC/k	States	Trans	Pars	T1		T2		T3	
						pso	qcqp	pso	qcqp	pso	qcqp
1	.124/.11/.09	F/1	75	118	8	<1	<1	<1	<1	<1	<1
		F/2	205	420	47	2	<1	2	<1	2	<1
		F/4	921	1864	215	9	2	9	2	10	2
		F/8	3889	7824	911	43	15	42	14	42	14
2	4.15/4.5/5.5	F/1	47	106	3	-	-	-	-	Err	<1
		F/2	183	390	15	7.4	11	4	9	2	<1
		F/4	719	1486	63	TO	64	39	91	14	8
		F/8	2845	5788	255	TO	700	TO	946	254	69
3	9/10/15	F/1	3268	13094	276	TO	TO	TO	43	22	4
		F/2	16004	46153	1783	TO	TO	TO	877	152	28
		C/2	11270	36171	1168	TO	TO	TO	358	100	62
		C/4	27183	82145	2940	TO	MO	TO	MO	476	MO
4	.249/.2/.15	F/1	3366	6534	364	18	25	18	15	18	12
		F/2	25713	51608	3907	330	MO	350	MO	326	MO
5	5.2/15/25	F/1	30	64	8	-	-	TO	TO	<1	TO
		F/2	137	294	49	TO	TO	14	TO	2	TO
		F/4	587	1214	219	93	TO	TO	TO	26	TO
		F/8	2421	4924	919	TO	TO	1034	TO	115	TO
6	10.6/10.9/82.5	C/2	99	212	33	TO	TO	3.7	TO	<1	TO
		C/4	231	476	81	7	TO	6	TO	3	TO
		F/1	16	33	1	-	-	-	-	<1	TO
		F/2	77	160	11	9	TO	6	TO	<1	TO
7	.929/.928/.927	F/4	354	721	63	20	TO	21	63	3	TO
		F/1	87	184	3	TO	TO	<1	1	<1	<1
		F/2	285	592	15	4	TO	4	20	3	22
		F/4	1017	2080	63	76	767	71	205	67	187
F/8	3825	7744	255	TO	TO	TO	TO	TO	TO		

Evaluation. Strategies for thresholds which are suboptimal (T3) are synthesised faster. If the memory bound is increased, the number of parameters quickly grows and the performance of the methods degrades. Additional experiments showed that the number of states has only a minor effect on the performance. The simpler FSC topology for a counter alleviates the blow-up of the pMC and is successfully utilised to find strategies for larger instances.

Trivially, a k -FSC is also a valid $(k+i)$ -FSC for some $i \in \mathbb{N}$. Yet, the larger number of parameters make searching for $(k+i)$ -FSCs significantly more difficult. We furthermore observe that the performance of PSO and QCQP is incomparable, and both methods have their merits.

Summarising, many of the POMDPs in the benchmarks allow good performance via FSCs with small memory. **We find nearly-optimal, and small, FSCs for POMDP benchmarks with thousands of states within seconds.**

6.2 PROVING ϵ -OPTIMALITY

We now focus on evaluating how fast pMC techniques prove the absence of a strategy satisfying the specification. In particular, we consider proving that for a specific threshold, no strategy induces a better value. Such a proof allows us to draw conclusions about the (ϵ -)optimality of a strategy synthesised in Sect. 6.1.

Results. Table 4(a) shows the run times to prove that for the POMDP in column *Id*, there exists no strategy of type FSC with k memory that performs better than threshold T . The row indicated by * was obtained with SMT. All

Table 4: ϵ -optimality and closed-form computation

(a) Proving absence				(b) Closed-form sol.		
Id	FSC/k	T	time	Id	FSC/k	time
2	F/1	5	<1	1	F/1	<1
3	F/1	5	8	1	F/2	97
3	F/4	5	183	2	F/1	155
4	F/1	0.25	2*	3	F/1	464
5	F/1	10	3	4	F/1	<1
5	F/2	5	TO	5	F/1	116
6	F/1	82	<1	6	F/1	<1
6	F/8	10.5	1	7	F/1	TO
7	F/1	0.94	5			

other results were obtained with parameter lifting.

Evaluation. The methods generally prove tight bounds for $k=1$. For $k>1$, the high number of parameters yields a mixed impression, the performance depends on the benchmark. **We find proofs for non-trivial bounds up to $k=8$, even if the pMC has hundreds of parameters.**

6.3 CLOSED-FORM SOLUTIONS

Results. Table 4(b) indicates running times to compute a closed-form solution, i. e. a rational function that maps k -FSCs to the induced probability.

Evaluation. Closed form computation is limited to small memory bounds. The rational functions obtained vary wildly in their structure. For (4), the result is a constant function which is trivial to analyse, while for (3), we obtained rational functions with roughly one million terms, rendering further evaluation expensive.

7 CONCLUSION

This paper connects two active research areas, namely verification and synthesis for POMDPs and parameter synthesis for Markov models. We see benefits for both areas. On the one hand, the rich application area for POMDPs in, e. g. robotics, yields new challenging benchmarks for parameter synthesis and can drive the development of more efficient methods. On the other hand, parameter synthesis tools and techniques extend the state-of-the-art approaches for POMDP analysis. Future work will also concern a thorough investigation of *permissive schedulers*, that correspond to regions of parameter instantiations, in concrete motion planning scenarios.

References

- Amato, C., Bernstein, D. S., and Zilberstein, S. (2010). Optimizing fixed-size stochastic controllers for POMDPs and decentralized POMDPs. *Autonomous Agents and Multi-Agent Systems*, 21(3):293–320.
- Baier, C. and Katoen, J.-P. (2008). *Principles of Model Checking*. MIT Press.
- Braziunas, D. (2003). POMDP solution methods. *University of Toronto*.
- Chatterjee, K., Chmelik, M., and Davies, J. (2016a). A symbolic SAT-based algorithm for almost-sure reachability with small strategies in POMDPs. In *AAAI*, pages 3225–3232. AAAI Press.
- Chatterjee, K., Chmelik, M., Gupta, R., and Kanodia, A. (2016b). Optimal cost almost-sure reachability in POMDPs. *Artif. Intell.*, 234:26–48.
- Chatterjee, K., Chmelik, M., and Tracol, M. (2016c). What is decidable about partially observable Markov decision processes with ω -regular objectives. *Journal of Computer and System Sciences*, 82(5):878–911.
- Chatterjee, K., De Alfaro, L., and Henzinger, T. A. (2004). Trading memory for randomness. In *QEST*. IEEE.
- Chen, T., Hahn, E. M., Han, T., Kwiatkowska, M. Z., Qu, H., and Zhang, L. (2013). Model repair for Markov decision processes. In *TASE*, pages 85–92. IEEE CS.
- Condon, A. (1992). The complexity of stochastic games. *Inf. Comput.*, 96(2):203–224.
- Cubuktepe, M., Jansen, N., Junges, S., Katoen, J., Pappas, I., Poonawala, H. A., and Topcu, U. (2017). Sequential convex programming for the efficient verification of parametric MDPs. In *TACAS (2)*, volume 10206 of *LNCS*, pages 133–150.
- Cubuktepe, M., Jansen, N., Junges, S., Katoen, J.-P., and Topcu, U. (2018). Synthesis in pMDPs: A tale of 1001 parameters. *CoRR*, abs/1803.02884.
- Daws, C. (2004). Symbolic and parametric model checking of discrete-time Markov chains. In *ICTAC*, volume 3407 of *LNCS*, pages 280–294. Springer.
- de Moura, L. M. and Björner, N. (2008). Z3: An efficient SMT solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340. Springer.
- Dehnert, C., Junges, S., Jansen, N., Corzilius, F., Volk, M., Bruintjes, H., Katoen, J.-P., and Ábrahám, E. (2015). PROPhESY: A probabilistic parameter synthesis tool. In *CAV*, volume 9206 of *LNCS*, pages 214–231. Springer.
- Dehnert, C., Junges, S., Katoen, J., and Volk, M. (2017). A Storm is coming: A modern probabilistic model checker. In *CAV (2)*, volume 10427 of *LNCS*, pages 592–600. Springer.
- Filieri, A., Ghezzi, C., and Tamburrelli, G. (2011). Runtime efficient probabilistic model checking. In *ICSE*, pages 341–350. ACM.
- Hahn, E. M., Hermanns, H., and Zhang, L. (2010). Probabilistic reachability for parametric Markov models. *Software Tools for Technology Transfer*, 13(1):3–19.
- Howard, R. A. (1960). *Dynamic Programming and Markov Processes*. The MIT Press.
- Hutschenreiter, L., Baier, C., and Klein, J. (2017). Parametric Markov chains: PCTL complexity and fraction-free Gaussian elimination. In *GandALF*, volume 256 of *EPTCS*, pages 16–30.
- Junges, S., Jansen, N., Wimmer, R., Quatmann, T., Winterer, L., Katoen, J., and Becker, B. (2017). Permissive finite-state controllers of POMDPs using parameter synthesis. *CoRR*, abs/1710.10294.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artif. Intell.*, 101(1):99–134.
- Kolobov, A., Mausam, and Weld, D. S. (2012). A theory of goal-oriented MDPs with dead ends. In *UAI*, pages 438–447. AUAI Press.
- Kwiatkowska, M., Norman, G., and Parker, D. (2011). PRISM 4.0: Verification of probabilistic real-time systems. In *CAV*, volume 6806 of *LNCS*, pages 585–591. Springer.
- Madani, O., Hanks, S., and Condon, A. (1999). On the undecidability of probabilistic planning and infinite-horizon partially observable Markov decision problems. In *AAAI*, pages 541–548. AAAI Press.
- Meuleau, N., Kim, K.-E., Kaelbling, L. P., and Cassandra, A. R. (1999). Solving POMDPs by searching the space of finite policies. In *UAI*, pages 417–426. Morgan Kaufmann Publishers Inc.
- Norman, G., Parker, D., and Zou, X. (2017). Verification and control of partially observable probabilistic systems. *Real-Time Systems*, 53(3):354–402.
- Papadimitriou, C. H. (1994). *Computational complexity*. Addison-Wesley.
- PARAM Website (2015). <http://depend.cs.uni-sb.de/tools/param/>.
- Puterman, M. L. (1994). *Markov Decision Processes*. John Wiley and Sons.
- Quatmann, T., Dehnert, C., Jansen, N., Junges, S., and Katoen, J. (2016). Parameter synthesis for Markov models: Faster than ever. In *ATVA*, volume 9938 of *LNCS*, pages 50–67. Springer.
- Ross, S. M. (1983). *Introduction to Stochastic Dynamic Programming*. Academic Press, Inc.
- Roy, N., Gordon, G. J., and Thrun, S. (2005). Finding approximate POMDP solutions through belief compression. *J. Artif. Intell. Res.*, 23:1–40.
- Russell, S. J. and Norvig, P. (2010). *Artificial Intelligence – A Modern Approach (3. ed.)*. Pearson Education.
- Shani, G., Pineau, J., and Kaplow, R. (2013). A survey of point-based POMDP solvers. *Autonomous Agents and Multi-Agent Systems*, 27(1):1–51.
- Szer, D. and Charpillet, F. (2005). An optimal best-

- first search algorithm for solving infinite horizon DEC-POMDPs. In *ECML*, pages 389–399. Springer.
- Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics*. The MIT Press.
- Vlassis, N., Littman, M. L., and Barber, D. (2012). On the computational complexity of stochastic controller optimization in POMDPs. *ACM Trans. on Computation Theory*, 4(4):12:1–12:8.
- Walraven, E. and Spaan, M. T. J. (2017). Accelerated vector pruning for optimal POMDP solvers. In *AAAI*, pages 3672–3678. AAAI Press.
- Wongpiromsarn, T. and Frazzoli, E. (2012). Control of probabilistic systems under dynamic, partially known environments with temporal logic specifications. In *CDC*, pages 7644–7651. IEEE.