Scalable control synthesis for stochastic systems via structural IMDP abstractions

Frederik Baymler Mathiesen[1]
Sofie Haesaert[2]
Luca Laurenti[1]

[1] Delft University of Technology
[2] Eindhoven University of Technology

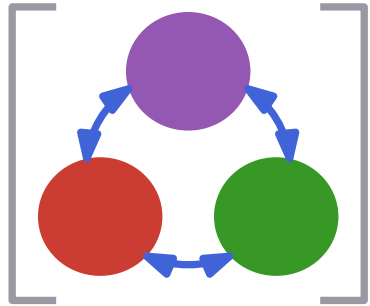

TU Delft

01

Verification of stochastic systems

TU Delft

System definition

$$x_{k+1} = f(x_k, u_k, w_k)$$

The transition kernel of the system is defined by

$$T(X \mid x, u) := \sum_{r=1}^{K} \alpha_r(x, u) \int_X \mathcal{N}(x' \mid \mu^r(x, u), \Sigma^r(x, u)) dx', \tag{1}$$

Assumption: $\Sigma^r(x, u)$ is diagonal for every (x, u).

Example: Additive noise systems

For a system $x_{k+1} = f(x_k, u_k) + w_k$ with $w_k \sim \mathcal{N}(0, \Sigma)$, the transition kernel is:

$$T(X \mid x, u) = \mathcal{N}(X \mid f(x, u), \Sigma)$$

Problem statement

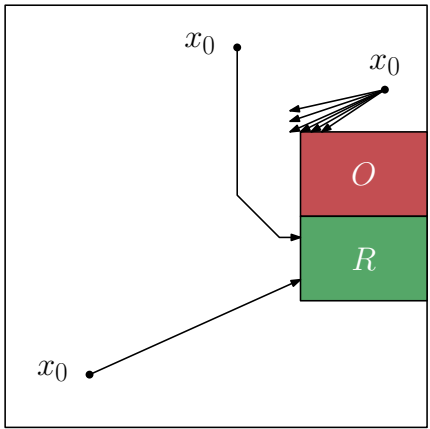

Problem

$$\max_{\pi_x} P_{\text{ra}}(R, O, x_0, \pi_x, H)$$

Abstraction

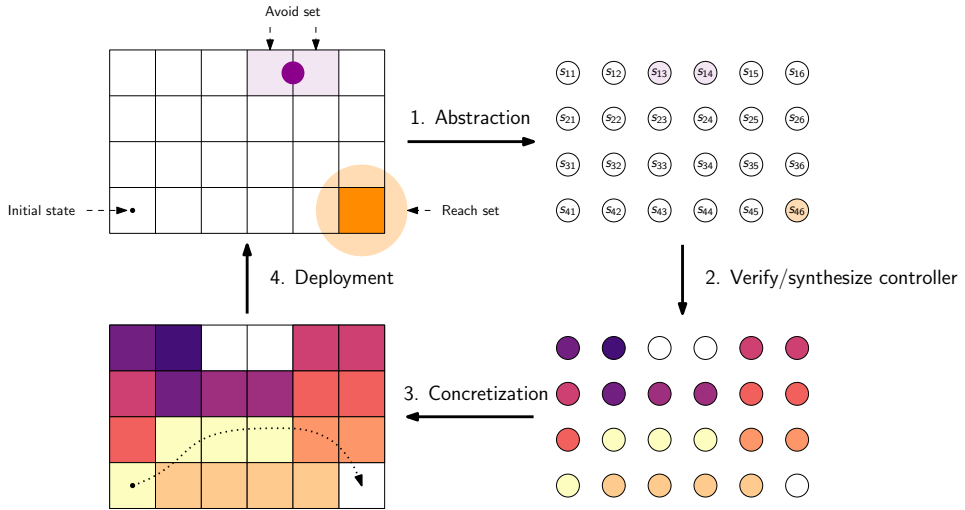

Interval Markov Decision Processes

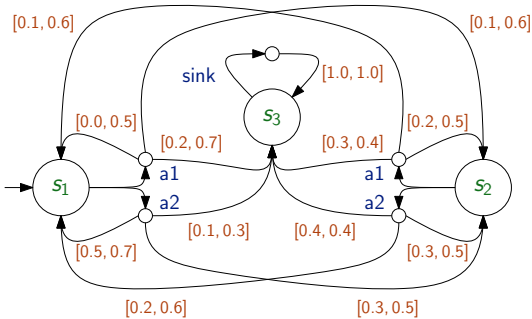

$$M = (S, A, \underline{P}, \overline{P})$$

where $\underline{P}, \overline{P} : S \times A \times S \to [0, 1]$ with $\sum_{s' \in S} \underline{P}(s, a, s') \leq 1 \leq \sum_{s' \in S} \overline{P}(s, a, s')$

Interval Markov Decision Processes

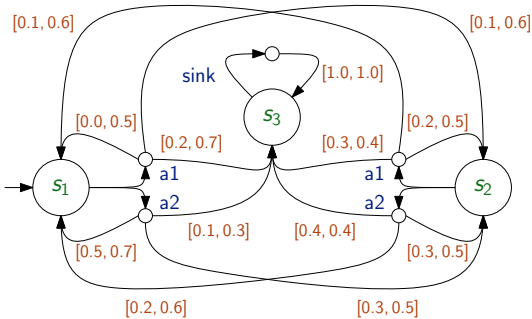

$$M = (S, A, \Gamma)$$

where $\Gamma = \{\Gamma_{s,a}\}_{s \in S, a \in A}$ with $\Gamma_{s,a} \in \operatorname{int}\operatorname{amb}(S)$.

Abstraction to IMDPs

$$x(k+1) = Ax(k) + \omega, \quad \omega \sim \mathcal{N}(0, \Sigma)$$

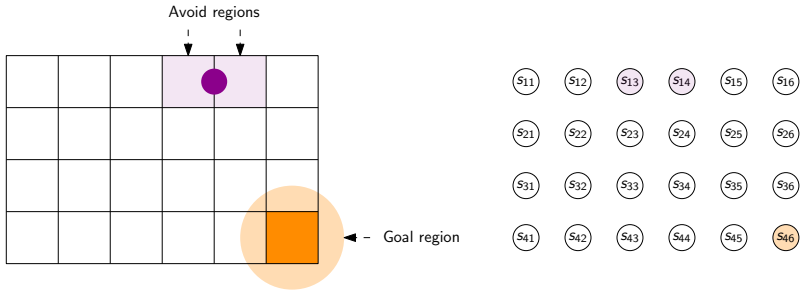

Abstraction to IMDPs

$$x(k+1) = Ax(k) + \omega, \quad \omega \sim \mathcal{N}(0, \Sigma)$$

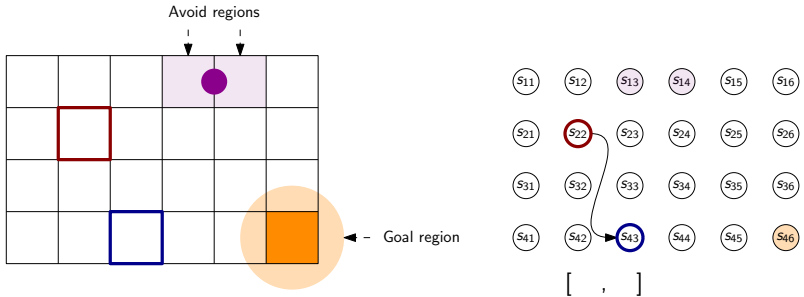

$$\Gamma_s = \{\gamma_s \in \mathcal{D}(S) : \forall t \in S, \min_{x \in s} T(t \mid x) \leq \gamma_s(t) \leq \max_{x \in s} T(t \mid x)\}.$$

Abstraction to IMDPs

$$x(k+1) = Ax(k) + \omega, \quad \omega \sim \mathcal{N}(0, \Sigma)$$

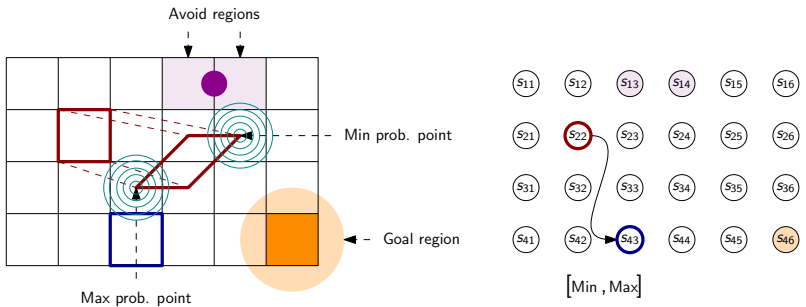

$$\Gamma_s = \{\gamma_s \in \mathcal{D}(S) : \forall t \in S, \min_{x \in s} T(t \mid x) \leq \gamma_s(t) \leq \max_{x \in s} T(t \mid x)\}.$$

Abstraction to IMDPs

$$x(k+1) = Ax(k) + \omega, \quad \omega \sim \mathcal{N}(0, \Sigma)$$

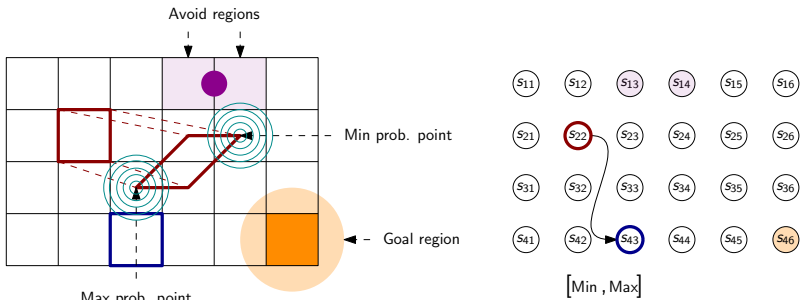

$$\min_{x \in s} T(t \mid x) = \min_{x \in s} \int_{t} \mathcal{N}(y \mid Ax, \Sigma)\, dy \geq \prod_{i=1}^{n} \min_{x \in s} \int_{\underline{t}_i}^{\bar{t}_i} \mathcal{N}(y_i \mid (Ax)_i, \Sigma_{ii})\, dy_i.$$

What is the problem? - Scalability!

Paper	Max dimension	# states
Cauchi et al. [1]	11	8190
Wooding et al. [2]	14	16384
Badings et al. [3]	10	7000
Adams et al. [4]	5	15000
Mathiesen et. al. [5]	6	25921

[1] N. Cauchi, L. Laurenti, M. Lahijanian, A. Abate, M. Kwiatkowska, and L. Cardelli (2019). "Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems". In: *Proceedings of the 22nd ACM international conference on hybrid systems: computation and control*, pp. 240–251

[2] B. Wooding and A. Lavaei (2024). "IMPaCT: Interval MDP Parallel Construction for Controller Synthesis of Large-Scale STochastic Systems". In: *Quantitative Evaluation of Systems and Formal Modeling and Analysis of Timed Systems*

[3] T. Badings, L. Romao, A. Abate, and N. Jansen (2023). "Probabilities are not enough: Formal controller synthesis for stochastic dynamical models with epistemic uncertainty". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. 12, pp. 14701–14710

[4] S. Adams, M. Lahijanian, and L. Laurenti (2022). "Formal control synthesis for stochastic neural network dynamic models". In: *IEEE Control Systems Letters* 6, pp. 2858–2863

[5] F. B. Mathiesen, M. Lahijanian, and L. Laurenti (2024). "IntervalMDP.jl: Accelerated Value Iteration for Interval Markov Decision Processes". In: *IFAC-PapersOnLine* 58.11. 8th IFAC Conference on Analysis and Design of Hybrid Systems ADHS 2024, pp. 1–6

02

Orthogonally decoupled IMDPs

Preserving marginal probability bounds

$$x(k+1) = Ax(k) + \omega, \quad \omega \sim \mathcal{N}(0, \Sigma)$$

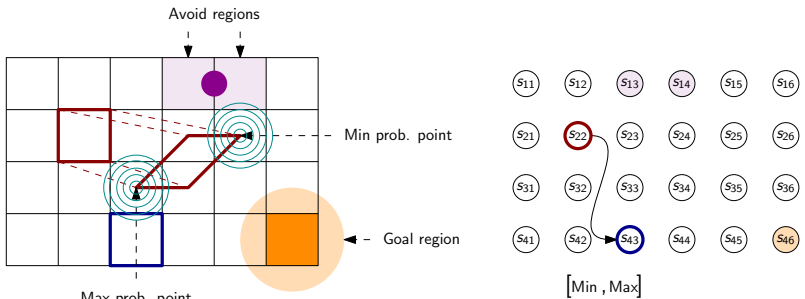

$$\min_{x \in s} T(t \mid x) = \min_{x \in s} \int_{t} \mathcal{N}(y \mid Ax, \Sigma)\, dy \geq \prod_{i=1}^{n} \min_{x \in s} \int_{\underline{t}_i}^{\bar{t}_i} \mathcal{N}(y_i \mid (Ax)_i, \Sigma_{ii})\, dy_i.$$

Orthogonal decoupling

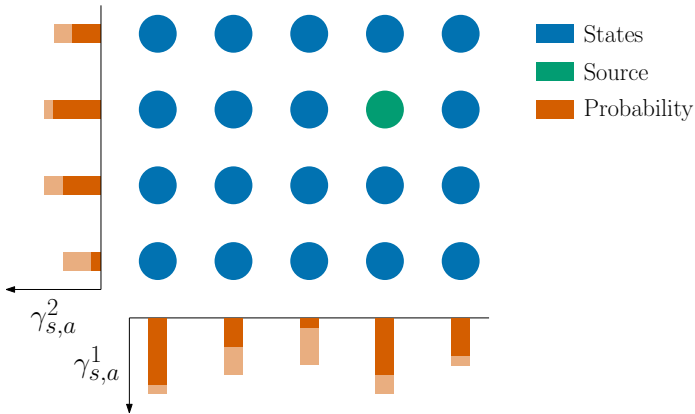

Definition: odIMDPs

An orthogonally decoupled IMDP (odIMDP) with n marginals is a tuple $M = (S, A, \Gamma)$ where

- $S = S_1 \times \cdots \times S_n$ is a finite set of joint states with S_i being the set of states in marginal i,
- A is a finite set of actions, and
- $\Gamma = \{\Gamma_{s,a}\}_{s \in S, a \in A}$ are sets of feasible transition probability distributions where $\Gamma_{s,a} = \bigotimes_{i=1}^{n} \Gamma^{i}_{s,a}$ with $\Gamma^{i}_{s,a} \in \text{int amb}(S_i)$.

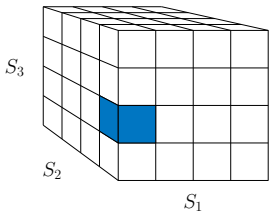

03

Abstraction

Univariate Gaussian case

$$\min_{x \in s} \int_{\underline{t}}^{\bar{t}} \mathcal{N}(x' \mid \mu(x,a), \sigma(x,a)) dx' =$$

$$\min_{x \in s} F_{x'|\mu(x,a),\sigma(x,a)}(\bar{t}) - F_{x'|\mu(x,a),\sigma(x,a)}(\underline{t})$$

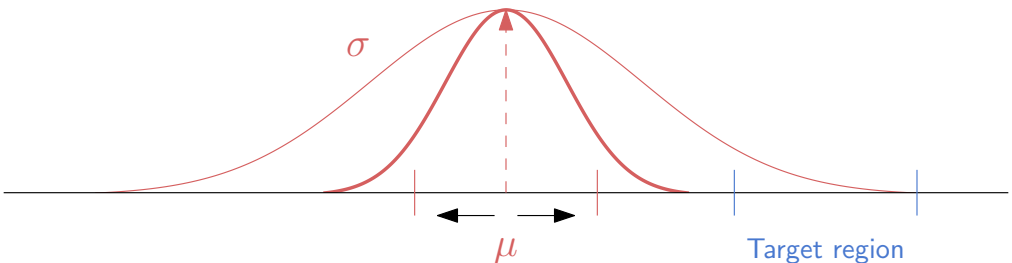

Multivariate Gaussian case

$$\prod_{i=1}^{n} \min_{x \in s} F_{x_i' | \mu_i(x,a), \Sigma_{ii}(x,a)}(\bar{t}) - F_{x_i' | \mu_i(x,a), \Sigma_{ii}(x,a)}(\underline{t})$$

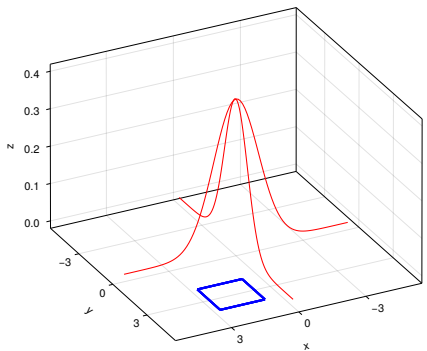

Multivariate Gaussian case

$$\prod_{i=1}^{n} \min_{x \in s} F_{x_i' | \mu_i(x,a), \Sigma_{ii}(x,a)}(\bar{t}) - F_{x_i' | \mu_i(x,a), \Sigma_{ii}(x,a)}(\underline{t})$$

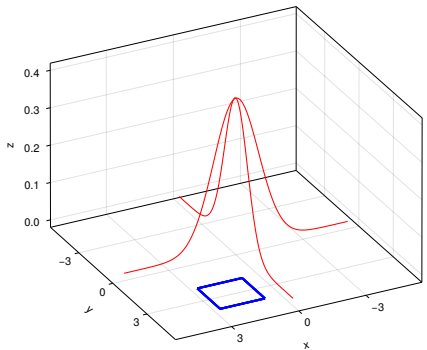

Proof of containment!

04

Experiments

Benchmarks [6] [7] [8]

Name	Dimension	Dynamics type	Property	# inputs	# states
Car parking	2	Additive linear	Reach-avoid	9	1600
Robot reachability	2	Additive linear	Reachability	121	400
Robot reach-avoid	2	Additive linear	Reach-avoid	441	1600
Building automation system	4	Additive affine	Safety	4	1225
Van der Pol	2	Additive polynomial	Reachability	11	2500
NNDM Cartpole	4	Additive NNDM	Safety	2	192000
6D linear model	6	Additive linear	Safety	0	262144
7D linear model	7	Additive linear	Safety	0	2097152
Dubin's car GP	3	GP with DKL	Reach-avoid	7	25600
Stochastic switched linear	2	Gaussian mixture	Reach-avoid	0	1600

[6] B. Van Huijgevoort, O. Schön, S. Soudjani, and S. Haesaert (2023). "SySCoRe: Synthesis via stochastic coupling relations". In: *Proceedings of the 26th ACM international conference on hybrid systems: Computation and control*

[7] B. Wooding and A. Lavaei (2024). "IMPaCT: Interval MDP Parallel Construction for Controller Synthesis of Large-Scale STochastic Systems". In: *Quantitative Evaluation of Systems and Formal Modeling and Analysis of Timed Systems*

[8] R. Reed, L. Laurenti, and M. Lahijanian (2023). "Promises of deep kernel learning for control synthesis". In: *IEEE Control Systems Letters*

IMDP-based approaches: computation performance

Figure: Peak memory

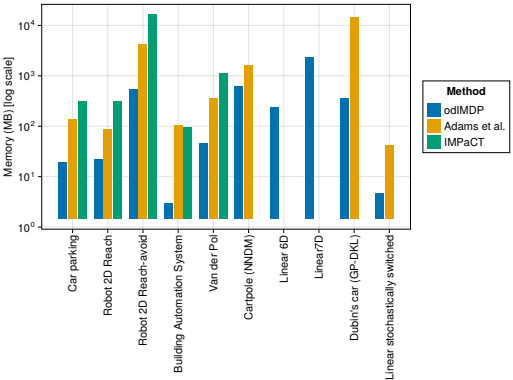

Figure: Abstraction and certification time

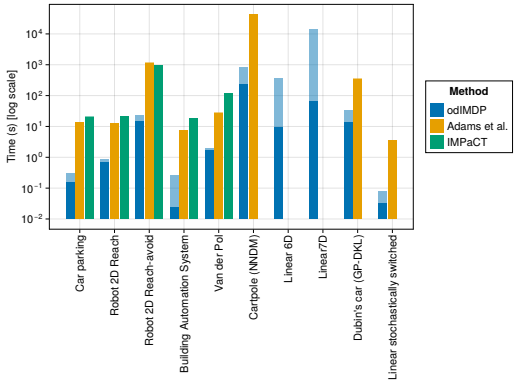

IMDP-based approaches: satisfaction probability

Figure: Mean satisfaction probability

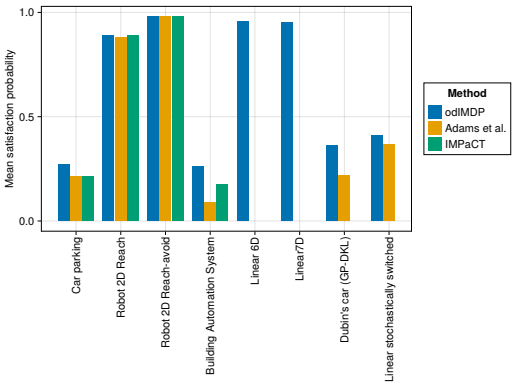

Figure: Mean error

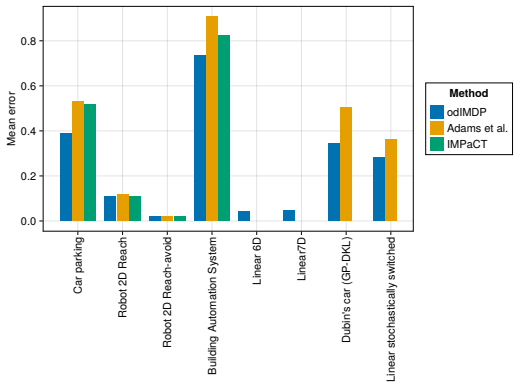

Comparison against SySCoRe (MDP)

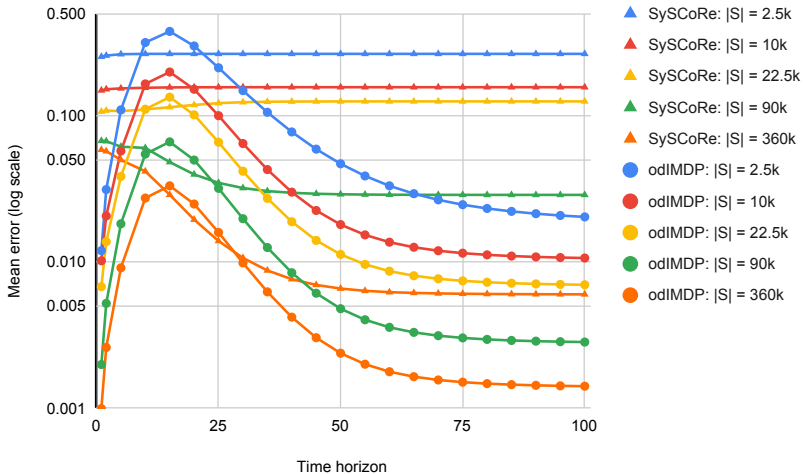

05

Discussion

Conclusion & Discussion

- Structural information in the abstraction:
  - reduces memory usage,
  - speeds up the computation, and
  - reduces conservativity
- Limitations:
  - Computing (robust) value iteration is very hard
  - Structure is necessary
- How far can we push structural abstractions?

IntervalMDP.jl

Refinement / adaptive gridding

Refinement procedures can be very beneficial. However, due to the reliance on the grid partitioning for odIMDPs, we have to rethink refinement procedures.

Bibliography I

Adams, S., M. Lahijanian, and L. Laurenti (2022). “Formal control synthesis for stochastic neural network dynamic models”. In: *IEEE Control Systems Letters* 6, pp. 2858–2863.

Badings, T., L. Romao, A. Abate, and N. Jansen (2023). “Probabilities are not enough: Formal controller synthesis for stochastic dynamical models with epistemic uncertainty”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. 12, pp. 14701–14710.

Cauchi, N., L. Laurenti, M. Lahijanian, A. Abate, M. Kwiatkowska, and L. Cardelli (2019). “Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems”. In: *Proceedings of the 22nd ACM international conference on hybrid systems: computation and control*, pp. 240–251.

Mathiesen, F. B., M. Lahijanian, and L. Laurenti (2024). “IntervalMDP.jl: Accelerated Value Iteration for Interval Markov Decision Processes”. In: *IFAC-PapersOnLine* 58.11. 8th IFAC Conference on Analysis and Design of Hybrid Systems ADHS 2024, pp. 1–6.

Reed, R., L. Laurenti, and M. Lahijanian (2023). “Promises of deep kernel learning for control synthesis”. In: *IEEE Control Systems Letters*.

Bibliography II

Van Huijgevoort, B., O. Schön, S. Soudjani, and S. Haesaert (2023). “SySCoRe: Synthesis via stochastic coupling relations”. In: *Proceedings of the 26th ACM international conference on hybrid systems: Computation and control*.

Wooding, B. and A. Lavaei (2024). “IMPaCT: Interval MDP Parallel Construction for Controller Synthesis of Large-Scale STochastic Systems”. In: *Quantitative Evaluation of Systems and Formal Modeling and Analysis of Timed Systems*.

07

Backup

Analysis of ambiguity sets

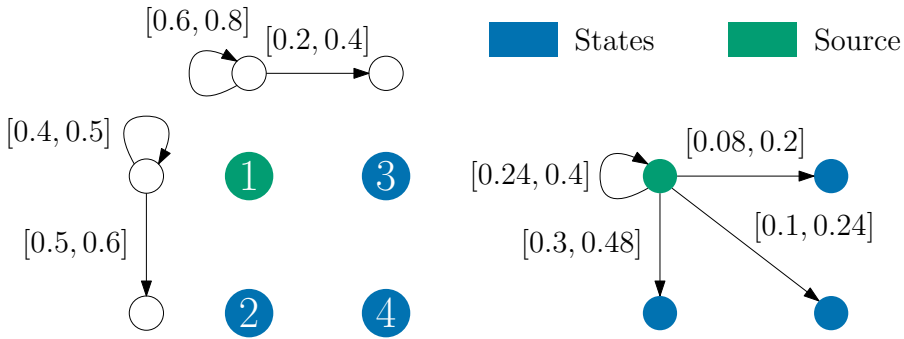

Interval bounds multiplication: value iteration

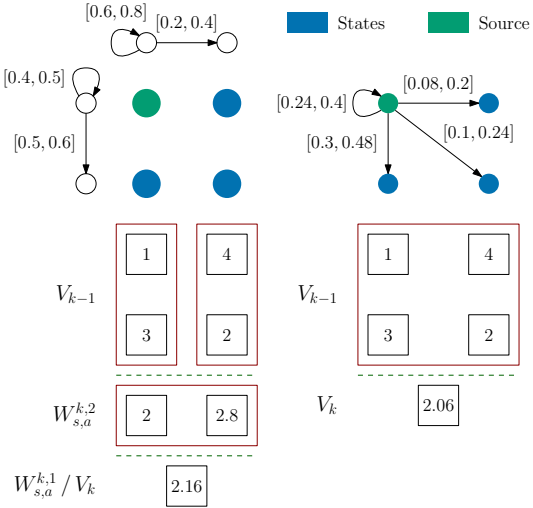

Eigentransformation of Gaussian systems

Facts about matrices

- An eigendecomposition of a matrix A is a factorization $A = Q\Lambda Q^{-1}$ where Λ has the eigenvalues of A on the diagonal and Q is a matrix where each column is an eigenvector of A for the corresponding eigenvector in Λ.
- Any covariance matrix is symmetric and thus normal.
- An eigendecomposition always exists for normal matrices.

Gaussian system:

$$T(\cdot|x, a) = \mathcal{N}(\cdot|f(x, a), \Sigma)$$

a. Find an eigendecomposition $Q\Lambda Q^{-1}$ of the (fixed) covariance matrix Σ.
b. Transform the the dynamics $\tilde{T}(\cdot|z, a) = \mathcal{N}(\cdot|Qf(Q^{-1}z, a), \Lambda)$.
c. Follow standard abstraction procedure.

Gaussian case

For each (s, a), we define intervals $[\underline{\mu}^i_{s,a}, \overline{\mu}^i_{s,a}]$ and $[\underline{\Sigma}^i_{s,a}, \overline{\Sigma}^i_{s,a}]$ such that for all $x \in s$:

$$\underline{\mu}^i_{s,a} \leq \mu(x, a)_i \leq \overline{\mu}^i_{s,a} \quad \text{and} \quad \underline{\Sigma}^i_{s,a} \leq \Sigma(x, a)_{ii} \leq \overline{\Sigma}^i_{s,a}. \tag{2}$$

Then the interval bounds can be computed as:

$$\underline{p}^i_{s,a}(t^i) = \min_{\mu^i \in [\underline{\mu}^i_{s,a}, \overline{\mu}^i_{s,a}],\ \Sigma^i \in \{\underline{\Sigma}^i_{s,a}, \overline{\Sigma}^i_{s,a}\}} \int_{\underline{t}_i}^{\bar{t}_i} \mathcal{N}(y_i \mid \mu^i, \Sigma^i) dy_i, \tag{3}$$

$$\overline{p}^i_{s,a}(t^i) = \max_{\mu^i \in [\underline{\mu}^i_{s,a}, \overline{\mu}^i_{s,a}],\ \Sigma^i \in \{\underline{\Sigma}^i_{s,a}, \overline{\Sigma}^i_{s,a}\}} \int_{\underline{t}_i}^{\bar{t}_i} \mathcal{N}(y_i \mid \mu^i, \Sigma^i) dy_i. \tag{4}$$

Non-trivial mixtures

Non-diagonal covariance!

A mixture of Gaussians, each with diagonal covariance, does not necessarily have diagonal covariance! For a distribution $\sum_{r=1}^{K} \alpha_r \mathcal{N}(\mu^r, \Sigma^r)$, the joint covariance is:

$$\Sigma = \sum_{r=1}^{K} \alpha_r \Sigma^r + \sum_{r=1}^{K} \alpha_r (\mu^r - \bar{\mu})(\mu^r - \bar{\mu})^\top \tag{5}$$

Non-trivial mixtures

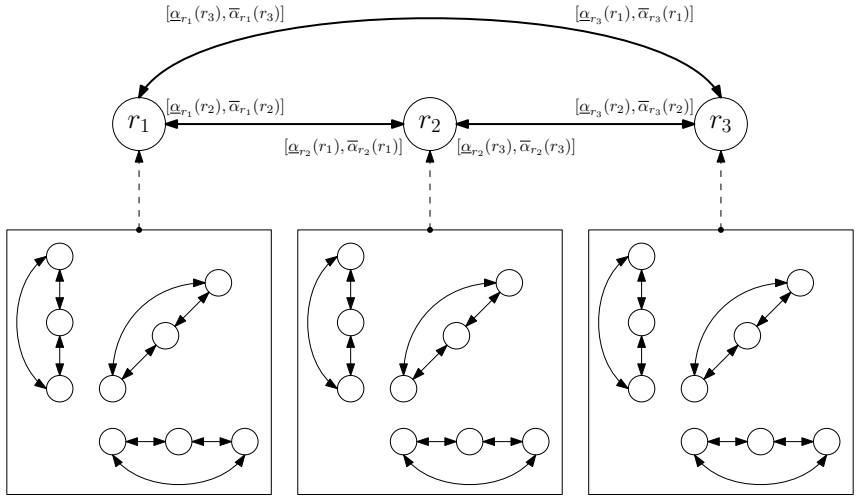

Non-trivial mixtures

Mixtures of odIMDPs

A mixture of K odIMDPs with n marginals is a tuple $M = (S, A, \Gamma, \Gamma^{\alpha})$ where

- $S = S_1 \times \cdots \times S_n$ is a finite set of joint states with S_i being the set of states in marginal i,
- A is a finite set of actions,
- $\Gamma = \{\Gamma_{r,s,a}\}_{r \in K, s \in S, a \in A}$ are a set of K product ambiguity sets, where $\Gamma_{r,s,a} = \bigotimes_{i=1}^{n} \Gamma^{i}_{r,s,a}$ with $\Gamma^{i}_{r,s,a} \in \operatorname{int amb}(S_i)$, and
- $\Gamma^{\alpha} = \{\Gamma^{\alpha}_{s,a}\}_{s \in S, a \in A}$ are sets of feasible weightings distributions, where $\Gamma^{\alpha}_{s,a} \in \operatorname{int amb}(K)$. A feasible weighting distribution for a source-action pair (s, a) is denoted by $\alpha_{s,a} \in \Gamma^{\alpha}_{s,a}$.

Value iteration

$$
\begin{aligned}
V_0(s) &= \mathbf{1}_R(s) \\
V_k(s) &= \operatorname{opt}^{\pi}_{a \in A} g\left(s, W^k_{s,a}\right),
\end{aligned}
\tag{6}
$$

where

$$
\begin{aligned}
W^{k,n+1}_{s,a}(t^1, \ldots, t^n) &= V_{k-1}(t) \\
W^{k,i}_{s,a}(t^1, \ldots, t^{i-1}) &= \operatorname{opt}^{\gamma}_{\gamma^i_{s,a} \in \Gamma^i_{s,a}} \sum_{t^i \in S_i} W^{k,i+1}_{s,a}(t^1, \ldots, t^i) \gamma^i_{s,a}(t^i), \\
&\qquad \text{for } i = 2, \ldots, n \\
W^k_{s,a} := W^{k,1}_{s,a} &= \operatorname{opt}^{\gamma}_{\gamma^1_{s,a} \in \Gamma^1_{s,a}} \sum_{t^1 \in S_1} W^{k,2}_{s,a}(t^1) \gamma^1_{s,a}(t^1).
\end{aligned}
\tag{7}
$$

IMDP-based approaches - computation performance

Table: Time is measured in seconds, and memory is measured in MB. OOM denotes out of memory.

Benchmark	Our method			Adams et al.			IMPaCT		
	Abs. time	Cert. time	Mem.	Abs. time	Cert. time	Mem.	Abs. time	Cert. time	Mem.
Car parking	0.150	0.146	19.2	13.497	0.255	138.1	19.570	0.846	304.9
Robot reachability	0.665	0.168	21.6	12.659	0.129	88.0	20.611	0.765	306.7
Robot reach-avoid	14.259	7.641	547.2	1136.720	6.739	4143.8	918.856	37.526	16388.0
Building automation system	0.023	0.237	3.1	7.273	0.285	105.1	17.564	0.318	96.0
Van der Pol	1.658	0.257	45.5	27.255	0.827	353.6	113.235	3.233	1093.4
NNDM Cartpole	236.154	550.747	610.8	42326.400	3.751	1590.9	Incompatible dynamics		
6D linear model	8.959	359.887	237.3	OOM			OOM ($\approx 1.1TB$)		
7D linear model	66.231	13903.540	2310.8	Timeout (24h)			OOM ($\approx 70.4TB$)		
Dubin's car GP	13.816	19.562	352.2	336.940	31.333	14265.8	Incompatible dynamics		
Stochastic switched linear	0.033	0.045	4.5	3.391	0.038	41.0	NLopt failure		

IMDP-based approaches - satisfaction probability

Table: $\check{V}$ denotes the lower bound satisfaction probability and ϵ the mean error. For Adams et al. and IMPaCT, we also report the difference δ in $\check{V}$ to our method, where positive means that our method yields a higher satisfaction probability and vice versa.

Benchmark	Our method		Adams et al.				
	Mean $\check{V}$	ϵ	Mean $\check{V}$	ϵ	Min δ	Max δ	Mean δ
Car parking	0.269	0.3885	0.213	0.5315	0.0040	0.1428	0.0560
Robot reachability	0.889	0.1108	0.881	0.1186	0.0060	0.0119	0.0079
Robot reach-avoid	0.980	0.0199	0.979	0.0208	0.0005	0.0027	0.0008
Building automation system	0.263	0.7336	0.090	0.9076	0.0510	0.2297	0.1733
Van der Pol	0.069	0.3367	0.051	0.4178	0.0000	0.0529	0.0177
NNDM Cartpole	0.004	0.7634	0.000	0.7184	0.0000	0.4101	0.0037
6D linear model	0.958	0.0419			OOM		
7D linear model	0.952	0.0483			Timeout		
Dubin's car GP	0.362	0.3461	0.216	0.5046	0.0000	0.8383	0.1458
Stochastic switched linear	0.411	0.2828	0.366	0.3605	0.0000	0.0979	0.0456

IMDP-based approaches - satisfaction probability

Table: $\check{V}$ denotes the lower bound satisfaction probability and ϵ the mean error. For Adams et al. and IMPaCT, we also report the difference δ in $\check{V}$ to our method, where positive means that our method yields a higher satisfaction probability and vice versa.

Benchmark	Our method		IMPaCT				
	Mean $\check{V}$	ϵ	Mean $\check{V}$	ϵ	Min δ	Max δ	Mean δ
Car parking	0.269	0.3885	0.213	0.5183	0.0040	0.1422	0.0556
Robot reachability	0.889	0.1108	0.890	0.1098	-0.0058	0.0022	-0.0010
Robot reach-avoid	0.980	0.0199	0.980	0.0202	-0.0001	0.0018	0.0003
Building automation system	0.263	0.7336	0.174	0.8237	0.0304	0.1131	0.0897
Van der Pol	0.069	0.3367	Unreliable results				
NNDM Cartpole	0.004	0.7634	Incompatible dynamics				
6D linear model	0.958	0.0419	OOM				
7D linear model	0.952	0.0483	OOM				
Dubin's car GP	0.362	0.3461	Incompatible dynamics				
Stochastic switched linear	0.411	0.2828	NLopt failure				

IMDP-based approaches: (local) satisfaction probability

Region-by-region difference: $V_{\text{odIMDP}}(s) - V_{\text{other}}(s)$.

Figure: Minimum difference

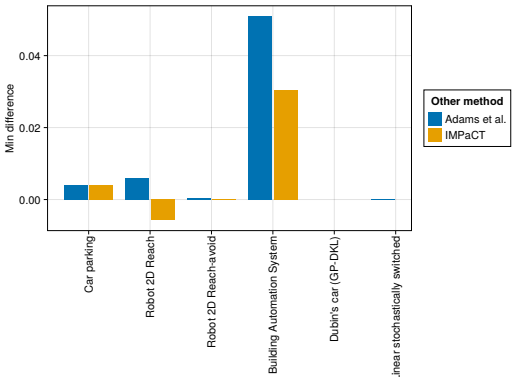

Figure: Mean difference

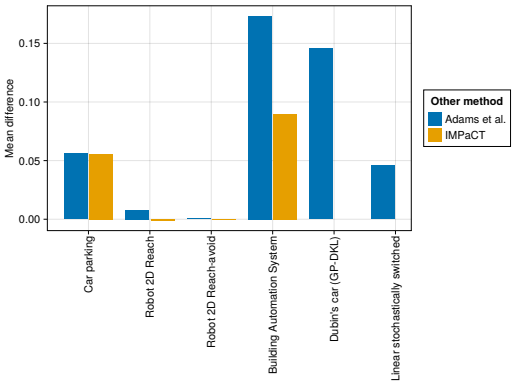

SySCoRe [9] computation performance

Table: The memory usage and the mean error ε at 100 time steps for SySCoRe and odIMDPs on the robot reach-avoid benchmark.

# regions	SySCoRe		odIMDPs	
	Mem. (MB)	ε	Mem. (MB)	ε
2500	0.80	0.2650	35.55	0.0204
10000	3.20	0.1568	279.51	0.0107
22500	7.20	0.1254	937.86	0.0070
90000	28.81	0.0288	7459.02	0.0028
360000	115.21	0.0060	59499.08	0.0014

[9] B. Van Huijgevoort, O. Schön, S. Soudjani, and S. Haesaert (2023). "SySCoRe: Synthesis via stochastic coupling relations". In: *Proceedings of the 26th ACM international conference on hybrid systems: Computation and control*