SIAM J. CONTROL OPTIM.
Vol. 55, No. 4, pp. 2333–2367

VERIFICATION OF GENERAL MARKOV DECISION PROCESSES BY APPROXIMATE SIMILARITY RELATIONS AND POLICY REFINEMENT*

SOFIE HAESAERT†, SADEGH ESMAEIL ZADEH SOUDJANI‡, AND ALESSANDRO ABATE‡

Abstract. In this work we introduce new approximate similarity relations that are shown to be key for policy (or control) synthesis over general Markov decision processes. The models of interest are discrete-time Markov decision processes, endowed with uncountably infinite state spaces and metric output (or observation) spaces. The new relations, underpinned by the use of metrics, allow, in particular, for a useful trade-off between deviations over probability distributions on states, and distances between model outputs. We show that the new probabilistic similarity relations, inspired by a notion of simulation developed for finite-state models, can be effectively employed over general Markov decision processes for verification purposes, and specifically for control refinement from abstract models.

Key words. policy refinement, approximate probabilistic simulation relations, correct-by-construction, verification

AMS subject classifications. 93E03, 93E99, 68Q60

DOI. 10.1137/16M1079397

Notation.

$\mathcal{R}$	Relation over $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$
$\bar{\mathcal{R}}$	Relation over $\bar{\mathcal{R}} \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2)$ obtained via lifting from $\mathcal{R}$, as per Definition 5
$\bar{\mathcal{R}}_\delta$	Relation over $\bar{\mathcal{R}} \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2)$ obtained via the approximate lifting with a deviation in probability bounded with δ obtained from $\mathcal{R}$, as per Definition 8
$\equiv_{\mathcal{R}_{eq}}$	Relation between two probability spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ based on the equivalence relation $\mathcal{R}_{eq} \subseteq (\mathbb{X}_1 \sqcup \mathbb{X}_2) \times (\mathbb{X}_1 \sqcup \mathbb{X}_2)$, à la [15], as reviewed in Appendix B
$\equiv^{\delta}_{\mathcal{R}_{eq}}$	Approximate relation between two probability spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ based on the equivalence relation $\mathcal{R}_{eq} \subseteq (\mathbb{X}_1 \sqcup \mathbb{X}_2) \times (\mathbb{X}_1 \sqcup \mathbb{X}_2)$, à la [1], as reviewed in Appendix B
$\preceq$	Probabilistic simulation relation; see Definition 6
$\approx$	Probabilistic bisimulation relation; see Definition 7
$\preceq^{\delta}_{\epsilon}$	ϵ, δ-approximate probabilistic simulation relation; see Definition 9

1. Introduction. The formal verification of computer systems allows for the quantification of their properties and for their correct functioning. While verifica-

*Received by the editors June 10, 2016; accepted for publication (in revised form) March 14, 2017; published electronically August 3, 2017. This work extends upon the preliminary conference paper in *Proceedings of the* 13*th International Conference on Quantitative Evaluation of Systems*, *QEST* 2016, Quebec City, Canada, Springer, 2016, pp. 227–243 [26] and includes the proofs of statements, an extensive literature review, and more elaborate examples.

http://www.siam.org/journals/sicon/55-4/M107939.html

†Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven 5600 MB, The Netherlands (s.haesaert@tue.nl).

‡Department of Computer Science, University of Oxford, Oxford, OX1 3QD, UK (sadegh.soudjani@cs.ox.ac.uk, aabate@cs.ox.ac.uk).

tion has classically focused on finite-state models, with the ever more ubiquitous embedding of digital components into physical systems, richer models are needed and correct functioning can only be expressed over the combined behavior of both the digital computer and the surrounding physical system. It is, in particular, of interest to synthesize the part of the computer software that controls or interacts with the physical system automatically, with low likelihood of malfunctioning. Furthermore, when computers interact with physical systems such as biological processes, power networks, and smart grids, stochastic models are key. Consider, as an example, a power network for which we would like to quantify the likelihood of blackouts and to synthesize strategies to minimize this.

Systems with uncertainty and nondeterminism can be naturally modeled as Markov decision processes (MDPs). In this work, we focus on general Markov decision processes (gMDPs) that have uncountable state spaces as well as metric output spaces. The characterization of properties over such processes cannot in general be attained analytically [3], so an alternative is to approximate these models by simpler processes that are prone to be mathematically analyzed or algorithmically verified [20, 21], such as finite-state MDPs [22]. Clearly, it is then key to provide formal guarantees on this approximation step, such that solutions of the verification or synthesis problem for a property on the simpler process can be extended to the original model. Our verification problems include the synthesis of a policy (or a control strategy) that maximizes the likelihood of the specification of interest.

In this work we develop a new notion of approximate similarity relation, aimed at attaining a computationally efficient controller synthesis over MDPs with metric output spaces. We show that it is possible to obtain a control strategy for a gMDP as a refinement of a strategy synthesized for an abstract model, at the expense of accuracy defined on a similarity relation between them, which quantifies bounded deviations in transition probabilities and output distances. In summary, we provide results allowing us to quantitatively relate the outcome of verification problems performed over the simpler (abstract) model to the original (concrete) model, and further to refine control strategies synthesized over the abstract model to strategies for the original model.

The use of similarity relations on *finite-state* probabilistic models has been broadly investigated, either via exact notions of probabilistic simulation and bisimulation relations [31, 36, 37, 29], or (more recently) via approximate notions [17, 18]. On the other hand, similar notions over *general, uncountable state spaces* have only recently been studied: available relations either hinge on stability requirements on model outputs [30, 44] (established via martingale theory or contractivity analysis) or, alternatively, enforce structural abstractions of a model [16] by exploiting continuity conditions on its probability laws [1, 2].

In this work, we want to quantify properties with a certified precision *both* in the deviation of the probability laws for finite-time events (as in the classical notion of probabilistic bisimulation) and of the output trajectories (as studied for dynamical models). Additionally, we impose no strict requirements on the dynamics of the given gMDP and its abstraction. To these ends, we first extend the exact probabilistic simulation and bisimulation relations based on lifting for finite-state probabilistic automata and stochastic games [29, 36, 37, 45] to gMDPs (section 3). We then generalize these notions to allow for errors on the probability laws *and* deviations over the output space (section 4). Two case studies in the area of smart buildings (section 5) are used to evaluate these new approximate probabilistic simulation relations. We start this paper with a comparison to existing simulation relations in literature.

Related literature. Unlike cognate work [1, 30] which recently appeared, we are interested in similarity relations that allow refining over the concrete model a control strategy synthesized on the abstract one. We zoom in on relations that, quite like the alternating notions in [5, 42] for nonprobabilistic models and in [45] for stochastic ones, quantitatively bound the difference in the controllable behavior of pairs of models (namely, a gMDP and its abstraction).

To attain this, we extend the simulation relations defined in [29, 36], which are connected to the preceding work in [31]. The latter has also inspired the notions of probabilistic (bi)simulation of labeled Markov processes (LMPs) in [14, 15] and their approximate versions [16, 17, 18]. In Appendix B we show how over a class of Markov processes (without controls), the proposed approximate similarity relation practically generalizes notions of probabilistic (bi)simulations of LMPs [14] based on zigzag morphisms, [15] based on equivalence relations, and their approximate versions [16, 17, 18] based on binary relations.

Since the seminal work in [31], extensions of exact (bi)simulation notions have been developed for specific model classes: these include [12, 40, 41, 35], which all provide exact similarity relations tailored to models classes different from LMPs.

2. Verification of gMDPs: Problem setup.

2.1. Preliminaries and notations. Given two sets A and B, the Cartesian product of A and B is given as $A \times B = \{(a, b) : a \in A \text{ and } b \in B\}$. The disjoint union of A and B is denoted as $A \sqcup B$ and consists of the combination of the members of A and B, where the original set membership is the distinguishing characteristic that forces the union to be disjoint, i.e., $A \sqcup B = (A \times \{0\}) \bigcup (B \times \{1\})$. As is usual for $C \subset A \sqcup B$ we denote $C \cap A = \{a \in A : (a, 0) \in C\}$. For the sets A and B a relation $\mathcal{R} \subset A \times B$ is a subset of their Cartesian product that relates elements $x \in A$ with elements $y \in B$, denoted as $x\mathcal{R}y$. We use the following notation for the mappings: $\mathcal{R}(\tilde{A}) := \{y : x\mathcal{R}y,\ x \in \tilde{A}\}$ and $\mathcal{R}^{-1}(\tilde{B}) := \{x : x\mathcal{R}y,\ y \in \tilde{B}\}$ for $\tilde{A} \subseteq A$ and $\tilde{B} \subseteq B$. A relation over a set defines a preorder if it is reflexive, for all $x \in A : x\mathcal{R}x$; and transitive, for all $x, y, z \in A$: if $x\mathcal{R}y$ and $y\mathcal{R}z$ then $x\mathcal{R}z$. A relation $\mathcal{R} \subseteq A \times A$ is an equivalence relation if it is reflexive, transitive, and symmetric, for all $x, y \in A$: if $x\mathcal{R}y$ then $y\mathcal{R}x$.

A measurable space is a pair $(\mathbb{X}, \mathcal{F})$ with sample space $\mathbb{X}$ and σ-algebra $\mathcal{F}$ defined over $\mathbb{X}$, which is equipped with a topology. As a specific instance of $\mathcal{F}$ consider the Borel measurable space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$. In this work, we restrict our attention to Polish spaces and generally consider the Borel σ-field [9]. Recall that a Polish space is a separable completely metrizable topological space. In other words, the space admits a topological isomorphism to a complete metric space which is dense with respect to a countable subset. A simple example of such a space is the real line.

A probability measure $\mathbb{P}(\cdot)$ for $(\mathbb{X}, \mathcal{F})$ is a nonnegative map, $\mathbb{P}(\cdot) : \mathcal{F} \to [0, 1]$ such that $\mathbb{P}(\mathbb{X}) = 1$ and such that for all countable collections $\{A_i\}_{i=1}^{\infty}$ of pairwise disjoint sets in $\mathcal{F}$, it holds that $\mathbb{P}(\bigcup_i A_i) = \sum_i \mathbb{P}(A_i)$. Together with the measurable space, such a probability measure $\mathbb{P}$ defines the probability space, which is denoted as $(\mathbb{X}, \mathcal{F}, \mathbb{P})$ and has realizations $x \sim \mathbb{P}$. Let us further denote the set of all probability measures for a given measurable pair $(\mathbb{X}, \mathcal{F})$ as $\mathcal{P}(\mathbb{X}, \mathcal{F})$. For a probability space[1] $(\mathbb{X}, \mathcal{F}_{\mathbb{X}}, \mathbb{P})$ and a measurable space $(\mathbb{Y}, \mathcal{F}_{\mathbb{Y}})$, a $(\mathbb{Y}, \mathcal{F}_{\mathbb{Y}})$-valued *random variable* is a function $y : \mathbb{X} \to \mathbb{Y}$ that is $(\mathcal{F}_{\mathbb{X}}, \mathcal{F}_{\mathbb{Y}})$-measurable, and which induces the probability

[1]The index $\mathbb{X}$ in $\mathcal{F}_{\mathbb{X}}$ distinguishes the given σ-algebra on $\mathbb{X}$ from that on $\mathbb{Y}$, which is denoted as $\mathcal{F}_{\mathbb{Y}}$. Whenever possible this index will be dropped.

measure $y_*\mathbb{P}$ in $\mathcal{P}(\mathbb{Y}, \mathcal{F}_\mathbb{Y})$. For a given set $\mathbb{Y}$ a metric or distance function $\mathbf{d}_\mathbb{Y}$ is a function $\mathbf{d}_\mathbb{Y} : \mathbb{Y} \times \mathbb{Y} \to \mathbb{R}_0^+$ satisfying the following conditions: for all $y_1, y_2, y_3 \in \mathbb{Y}$: $\mathbf{d}_\mathbb{Y}(y_1, y_2) = 0$ iff $y_1 = y_2$; $\mathbf{d}_\mathbb{Y}(y_1, y_2) = \mathbf{d}_\mathbb{Y}(y_2, y_1)$; and $\mathbf{d}_\mathbb{Y}(y_1, y_3) \leq \mathbf{d}_\mathbb{Y}(y_1, y_2) + \mathbf{d}_\mathbb{Y}(y_2, y_3)$.

2.2. gMDP models—syntax and semantics. gMDPs are related to control Markov processes [1] and MDPs [7, 34, 27], and formalized as follows.

DEFINITION 1 (MDP). *The tuple* $\mathbf{M} = (\mathbb{X}, \pi, \mathbb{T}, \mathbb{U})$ *defines a discrete-time MDP over an uncountable state space* $\mathbb{X}$*, and is characterized by* $\mathbb{T}$*, a conditional stochastic kernel that assigns to each point* $x \in \mathbb{X}$ *and control* $u \in \mathbb{U}$ *a probability measure* $\mathbb{T}(\cdot \mid x, u)$ *over* $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$*. For any set* $A \in \mathcal{B}(\mathbb{X})$*,* $\mathbb{P}_{x,u}(x(t+1) \in A) = \int_A \mathbb{T}(dy \mid x(t) = x, u)$*, where* $\mathbb{P}_{x,u}$ *denotes the conditional probability* $\mathbb{P}(\cdot \mid x, u)$*. The initial probability distribution is* $\pi : \mathcal{B}(\mathbb{X}) \to [0, 1]$*.*

At every state the state transition depends nondeterministically on the choice of $u \in \mathbb{U}$. When chosen according to a distribution $\mu_u : \mathcal{B}(\mathbb{U}) \to [0, 1]$, we refer to the stochastic control input as μ_u. Moreover the transition kernel is denoted as $\mathbb{T}(\cdot|x, \mu_u) = \int_\mathbb{U} \mathbb{T}(\cdot|x, u)\mu_u(du) \in \mathcal{P}(\mathbb{X}, \mathcal{B}(\mathbb{X}))$. Given a string of inputs $u(0), u(1), \ldots, u(N)$ over a finite time horizon $\{0, 1, \ldots, N\}$, and an initial condition x_0 (sampled from distribution π), the state at the $(t+1)$st time instant, $x(t+1)$, is obtained as a realization of the controlled Borel-measurable stochastic kernel $\mathbb{T}(\cdot \mid x(t), u(t))$—these semantics induce paths (or executions) of the MDP.

DEFINITION 2 (gMDP). $\mathbf{M} = (\mathbb{X}, \pi, \mathbb{T}, \mathbb{U}, h, \mathbb{Y})$ *is a discrete-time gMDP consisting of an MDP combined with output space* $\mathbb{Y}$ *and a measurable output mapping* $h : \mathbb{X} \to \mathbb{Y}$*. A metric* $\mathbf{d}_\mathbb{Y}$ *decorates the output space* $\mathbb{Y}$*.*

The gMDP semantics are directly inherited from those of the MDP. Further, output traces of gMDP are obtained as mappings of MDP paths, namely, $\{y(t)\}_{0:N} := y(0), y(1), \ldots, y(N)$, where $y(t) = h\big(x(t)\big)$. Denote the class of all gMDP with the metric output space $\mathbb{Y}$ as $\mathcal{M}_\mathbb{Y}$. Note that gMDP can be regarded as a superclass of the known LMPs [16] as elucidated in [2].

Example 1. Consider the stochastic process

$$\mathbf{M} : x(t+1) = f(x(t), u(t)) + e(t), \qquad y(t) = h(x(t)) \in \mathbb{Y}$$

with variables $x(t), u(t), e(t)$, taking values in $\mathbb{R}^n$, representing the state, control input,[2] and noise terms, respectively. The process is initialized as $x(0) \sim \pi$, and driven by $e(t)$, a white noise sequence with zero-mean normal distributions and covariance matrix Σ_e. This stochastic process, defined as a dynamical model, is a gMDP characterized by a tuple $(\mathbb{R}^n, \pi, \mathbb{T}, \mathbb{R}^n, h, \mathbb{Y})$, where the conditional transition kernel is defined as $\mathbb{T}(\cdot \mid x, u) = \mathcal{N}\left(f(x(t), u(t)), \Sigma_e\right)$, a normal probability distribution with mean $f(x(t), u(t))$ and covariance matrix Σ_e.

A policy is a selection of control inputs based on the past history of states and controls. We allow controls to be selected via universally measurable maps [7] from the state to the control space, so that time-bounded properties such as safety can be maximized [3]. When the selected controls are only dependent on the current states, and thus conditionally independent of history (or memoryless), the policy is referred to as Markov.

[2] In other domains one also refers to the control variables as actions (machine learning, stochastic games) or as external nondeterminism (computer science).

DEFINITION 3 (Markov policy). *For a gMDP* $\mathbf{M} = (\mathbb{X}, \pi, \mathbb{T}, \mathbb{U}, h, \mathbb{Y})$*, a Markov policy* μ *is a sequence* $\mu = (\mu_1, \mu_2, \mu_3, \ldots)$ *of universally measurable maps* $\mu_t = \mathbb{X} \to \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$ $t = 0, 1, 2, \ldots,$ *from the state space* $\mathbb{X}$ *to the set of controls.*

Recall that a function $f : \mathbb{Z}_1 \to \mathbb{Z}_2$ is universally measurable if the inverse image of every Borel set is measurable with respect to every complete probability measure on $\mathbb{Z}_1$ that measures all Borel subsets of $\mathbb{Z}_1$.

The execution $\{x(t), t \in [0, N]\}$ initialized by $x_0 \in \mathbb{X}$ and controlled with Markov policy μ is a stochastic process defined on the canonical sample space $\Omega := \mathbb{X}^{N+1}$ endowed with its product topology $\mathcal{B}(\Omega)$. This stochastic process has a probability measure $\mathbb{P}$ uniquely defined by the transition kernel $\mathbb{T}$, policy μ, and initial distribution π [7, Prop. 7.45].

Of interest are time-dependent properties such as those expressed as specifications in a temporal logic of choice. This leads to problems where one maximizes the probability that a sequence of labeled sets is reached within a time limit and in the right order. One can intuitively realize that, in general, the optimal policy leading to the maximal probability is not a Markov (memoryless) policy, as introduced in Definition 3. We introduce the notion of a control strategy, and define it as a broader, memory-dependent version of the Markov policy above. This strategy is formulated as a Markov process that takes as an input the state of the to-be-controlled gMDP.

DEFINITION 4 (control strategy). *A control strategy* $\mathbf{C} = (\mathbb{X}_\mathbf{C}, x_{\mathbf{C}0}, \mathbb{X}, \mathbb{T}^t_\mathbf{C}, h^t_\mathbf{C})$ *for a gMDP* $\mathbf{M}$ *with state space* $\mathbb{X}$ *and control space* $\mathbb{U}$ *over the time horizon* $t = 0, 1, 2, \ldots, N$ *is an* inhomogeneous Markov process *with state space* $\mathbb{X}_\mathbf{C}$*; an initial state* $x_{\mathbf{C}0}$*; inputs* $x \in \mathbb{X}$*; time-dependent, universally measurable kernels* $\mathbb{T}^t_\mathbf{C}$*,* $t = 0, 1, \ldots, N$*; and with universally measurable output maps* $h^t_\mathbf{C} : \mathbb{X}_\mathbf{C} \to \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$*,* $t = 1, \ldots, N$*, with elements* $\mu \in \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$*.*

Unlike a Markov policy, the control strategy is, in general, dependent on the history, as it has an internal state that can be used to remember relevant past events. As elucidated in Algorithm 1, note that the first control $u(0)$ is selected by drawing $x_\mathbf{C}(1)$ according to $\mathbb{T}^0_\mathbf{C}(\cdot\,|x_\mathbf{C}(0), x(0))$, where $x_\mathbf{C}(0) = x_{\mathbf{C}0}$, and selecting $u(0)$ from measure $\mu^0_\mathbf{C} = h^0_\mathbf{C}(x_\mathbf{C}(1))$.[3] The control strategy applied to $\mathbf{M}$ can be both stochastic (as a realization of $\mathbb{T}^0_\mathbf{C}(\cdot\,|\,x_\mathbf{C}(0), x(0))$), a function of the initial state $x(0)$, and of time.

The execution $\{(x(t), x_\mathbf{C}(t)), t \in [0, N]\}$ of a gMDP $\mathbf{M}$ controlled with strategy $\mathbf{C}$ is defined on the canonical sample space $\Omega := (\mathbb{X} \times \mathbb{X}_\mathbf{C})^{N+1}$ endowed with its product topology $\mathcal{B}(\Omega)$. This stochastic process is associated with a unique probability measure $\mathbb{P}_{\mathbf{C}\times\mathbf{M}}$, since the stochastic kernels $\mathbb{T}^t_\mathbf{C}$ for $t \in [0, N]$ and $\mathbb{T}$ are Borel measurable and composed via universally measurable policies [7, Prop. 7.45].

2.3. gMDP verification and strategy refinement: Problem statement. We qualitatively introduce the main problem that we want to solve in this work: how can one provide a general framework to synthesize control policies over a formal abstraction $\tilde{\mathbf{M}}$ of a concrete complex model $\mathbf{M}$ with the understanding that $\tilde{\mathbf{M}}$ is much simpler to be manipulated (analytically or computationally) than $\mathbf{M}$ is? We approach this problem by defining a simulation relation under which a control strategy $\tilde{\mathbf{C}}$ for the abstract Markov process $\tilde{\mathbf{M}}$ implies the existence of a control strategy $\mathbf{C}$ for $\mathbf{M}$, so that we can quantify differences in the stochastic transition kernels and in the output trajectories for the two controlled models. This allows us to derive bounds

[3] Note that the stochastic transitions for the control strategy and the gMDP are selected in an alternating fashion. The output map of the strategy is indexed based on the time instant at which the resulting policy will be applied to the gMDP.

Algorithm 1 Execution of the controlled model $\mathbf{C} \times \mathbf{M}$.

set $t := 0$ and $x_{\mathbf{C}}(0) := x_{\mathbf{C}0}$
draw $x(0) \sim \pi$ {from $\mathbf{M}$}
while $t < N$ **do**
  draw $x_{\mathbf{C}}(t+1) \sim \mathbb{T}^t_{\mathbf{C}}(\cdot \,|x_{\mathbf{C}}(t), x(t))$ {from $\mathbf{C}$}
  set $\mu_t := h^t_{\mathbf{C}}(x_{\mathbf{C}}(t+1))$, draw $u(t)$ from μ_t
  draw $x(t+1) \sim \mathbb{T}(\cdot \,|x(t), u(t)))$ {from $\mathbf{M}$}
  set $t := t + 1$
end while

on the probability of satisfaction of a specification for $\mathbf{M} \times \mathbf{C}$ from the satisfaction probability of modified specifications for $\tilde{\mathbf{M}} \times \tilde{\mathbf{C}}$. We will show that with this setup we can deal with finite-horizon temporal properties, including safety verification as a relevant instance.

The results in this paper are to be used in parallel with optimization, both for selecting the control refinement and for synthesizing a policy on the abstract model. It has been shown in [7] that stochastic optimal control even for a system on a "basic" space can lead to measurability issues: in order to avoid these issues we follow [7, 17] and the developed theory for Polish spaces and Borel (or universally) measurable notions. Throughout the paper we will give as clarifying examples Markov processes evolving, as in Example 1, over Euclidean spaces which are a special instances of Polish spaces. This allows us to elucidate the theory.

3. Exact (bi)simulation relations based on lifting.

3.1. Introduction. In this section, we define probabilistic simulation and bisimulation relations that are, respectively, a preorder and an equivalence relation on $\mathcal{M}_{\mathbb{Y}}$. Before introducing these relations, we first extend Segala's [36] and Segala and Lynch's [37] notion of *lifting* to uncountable state spaces, which allows us to equate the transition kernels of two given gMDPs. Thereafter, we leverage liftings to define (bi)simulation relations over $\mathcal{M}_{\mathbb{Y}}$, which characterize the similarity in the controllable behaviors of the two gMDPs. Subsequently we show that these similarity relations also imply controller refinement, i.e., within the similarity relation a control strategy for a given gMDP can be refined to a controller for another gMDP. In the next section, we show that this exact notion of similarity allows a more general notion of approximate probabilistic simulation. The new notions of similarity relations extend the known exact notions in [31], and the approximate notions of [17, 18]. Additionally, we will show that these results can be naturally extended to allow for both differences in probability and deviations in the outputs of the two gMDPs.

We work with pairs of gMDPs put in a relationship, denoting them with numerical indices $(\mathbf{M}_1, \mathbf{M}_2)$, with the intention to apply the developed notions to an abstraction $\tilde{\mathbf{M}}$ of a concrete model $\mathbf{M}$, respectively.

3.2. Lifting for gMDPs. Consider two gMDPs $\mathbf{M}_1, \mathbf{M}_2 \in \mathcal{M}_{\mathbb{Y}}$ mapping to a common output space $\mathbb{Y}$ with metric $\mathbf{d}_{\mathbb{Y}}$. For $\mathbf{M}_1 = (\mathbb{X}_1, \pi_1, \mathbb{T}_1, \mathbb{U}_1, h_1, \mathbb{Y})$ and $\mathbf{M}_2 = (\mathbb{X}_2, \pi_2, \mathbb{T}_2, \mathbb{U}_2, h_2, \mathbb{Y})$ at given state-action pairs $x_1 \in \mathbb{X}_1, u_1 \in \mathbb{U}_1$ and $x_2 \in \mathbb{X}_2, u_2 \in \mathbb{U}_2$, respectively, we want to relate the corresponding transition kernels, namely, the probability measures $\mathbb{T}_1(\cdot \mid x_1, u_1) \in \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $\mathbb{T}_2(\cdot \mid x_2, u_2) \in \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$.

Similarly to the coupling of measures in $\mathcal{P}(\mathbb{X}, \mathcal{F})$ [4, 32], consider the *coupling* of two arbitrary probability spaces $(\mathbb{X}_1, \mathcal{F}_1, \mathbb{P}_1)$ and $(\mathbb{X}_2, \mathcal{F}_2, \mathbb{P}_2)$ (cf. [38, 39]). A

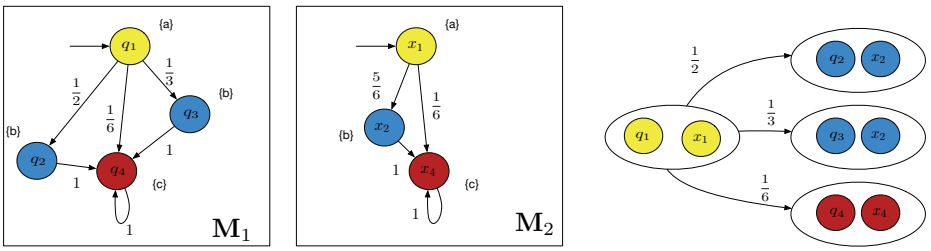

FIG. 1. *Finite-state Markov processes* $\mathbf{M}_1$ *and* $\mathbf{M}_2$ *(left and middle) with* $S = \{q_1, q_2, q_3, q_4\}$ *and* $T = \{x_1, x_2, x_4\}$ *the respective state spaces. The states are labeled with three different colors. Lifting probabilities of the transition kernels for* (q_1, x_1) *are given on the edges of the rightmost figure.*

probability measure $\mathbb{P}_c$ defined on $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{F})$ *couples* the two spaces if the projections p_1, p_2 with $x_1 = p_1(x_1, x_2)$ and $x_2 = p_2(x_1, x_2)$, define respectively, an $(\mathbb{X}_1, \mathcal{F}_1)$- and an $(\mathbb{X}_2, \mathcal{F}_2)$-valued random variable, such that $\mathbb{P}_1 = p_{1*}\mathbb{P}_c$ and $\mathbb{P}_2 = p_{2*}\mathbb{P}_c$. For *finite- or countable-state* stochastic processes a related concept has been introduced in [29] and has been referred to as lifting in [36, 37]: the transition probabilities are coupled using a weight function in a way that respects a given relation over the combined state spaces. In this work, rather than using weight functions over a countable or finite domain [36], we introduce lifting as a coupling of measures over Polish spaces and their corresponding Borel measurable σ-fields.

Since we assume that the state spaces are Polish and have a corresponding Borel σ-field for the given probability spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1), \mathbb{P}_1)$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2), \mathbb{P}_2)$ with $\mathbb{P}_1 := \mathbb{T}_1(\cdot \mid x_1, u_1)$ and $\mathbb{P}_2 := \mathbb{T}_2(\cdot \mid x_2, u_2)$, the natural choice for the σ-algebra becomes $\mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2) = \mathcal{B}(\mathbb{X}_1) \otimes \mathcal{B}(\mathbb{X}_2)$[4] and the question of finding a coupling can be reduced to finding a probability measure in $\mathcal{P}(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2))$.

DEFINITION 5 (lifting for general state spaces). *Let* $\mathbb{X}_1, \mathbb{X}_2$ *be two sets with associated measurable spaces* $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ *and* $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ *and let the Borel measurable set* $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ *be a relation. We denote by* $\bar{\mathcal{R}} \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ *the corresponding lifted relation, so that* $\Delta \bar{\mathcal{R}} \Theta$ *holds if there exists a probability space* $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ *(equivalently, a lifting* $\mathbb{W}$*) satisfying*

1. *for all* $X_1 \in \mathcal{B}(\mathbb{X}_1)$*:* $\mathbb{W}(X_1 \times \mathbb{X}_2) = \Delta(X_1)$*;*
2. *for all* $X_2 \in \mathcal{B}(\mathbb{X}_2)$*:* $\mathbb{W}(\mathbb{X}_1 \times X_2) = \Theta(X_2)$*;*
3. *for the probability space* $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ *it holds that* $x_1 \mathcal{R} x_2$ *with probability* 1 *or, equivalently, that* $\mathbb{W}(\mathcal{R}) = 1$*.*

With reference to the connection with the notion of coupling, an equivalent definition of lifting is obtained be replacing 1 and 2 by the condition that for $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$, the projections p_1, p_2 with $x_1 = p_1(x_1, x_2)$ and $x_2 = p_2(x_1, x_2)$, we can define $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$- and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$-valued random variables $\Delta = p_{1*}\mathbb{W}$ and $\Theta = p_{2*}\mathbb{W}$. An example is portrayed in Figure 1 containing two models $\mathbf{M}_1, \mathbf{M}_2$ and a relation (denoted by equally labeled/colored pairs of states), where the transition kernels for a pair of states is lifted with respect to the relation.

Remark 2. Notice that the extension of the notion of lifting to general spaces has required the use of *measures*, rather than weight functions over a countable or finite domain, as in [36, 29]. We have required that the σ-algebra $\mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$ contains

[4] $\mathcal{B}(\mathbb{X}_1) \otimes \mathcal{B}(\mathbb{X}_2)$ denotes the product σ-algebra of $\mathcal{B}(\mathbb{X}_1)$ and $\mathcal{B}(\mathbb{X}_2)$.

not only sets of the form $X_1 \times \mathbb{X}_2$ and $\mathbb{X}_1 \times X_2$, but also specifically the sets that characterize the relation $\mathcal{R}$. Since the spaces $\mathbb{X}_1$ and $\mathbb{X}_2$ have been assumed to be Polish, it holds that every open (closed) set in $\mathbb{X}_1 \times \mathbb{X}_2$ belongs to $\mathcal{B}(\mathbb{X}_1) \otimes \mathcal{B}(\mathbb{X}_2) = \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$ [9, Lemma 6.4.2]. As an instance consider the diagonal relation $\mathcal{R}_{diag} := \{(x, x) : x \in \mathbb{X}\}$ over $\mathbb{X} \times \mathbb{X}$, of importance for examples introduced later. This is a Borel measurable set [9, Theorem 6.5.7].

3.3. Exact probabilistic (bi)simulation relations via lifting. Similarly to the alternating notions for probabilistic game structures in [45], we provide a simulation that relates any input chosen for the first process with one for the second process. As such, we allow for more elaborate handling of the inputs than in the probabilistic simulation relations discussed in [17, 18], and further pave the way towards the inclusion of output maps. We extend the notions in [29, 36, 45] by allowing for more general Polish spaces. Further, we introduce the notion of *interface function* in order to connect the controllable behavior of two gMDPs:

$$\mathcal{U}_v : \mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2 \to \mathcal{P}(\mathbb{U}_2, \mathcal{B}(\mathbb{U}_2)),$$

where we require that $\mathcal{U}_v$ is a Borel measurable function. This means that $\mathcal{U}_v$ induces a Borel measurable stochastic kernel, again denoted by $\mathcal{U}_v$, over $\mathbb{U}_2$ given $(u_1, x_1, x_2) \in \mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2$. The notion of interface function is known in the context of correct-by-design controller synthesis and of hierarchical controller refinement [23, 42]. For the objective of hierarchical controller refinement, an interface function implements (or refines) any control action synthesized over the abstract model to an action for the concrete model. In order to establish an exact simulation relation between abstract and concrete models, we can attempt to refine the control actions from one model to the other by choosing an interface function that matches their stochastic behaviors. On the other hand in the next section, the interface function will be used to establish approximate simulation relations: for this goal, the optimal selection of the interface function is the one that optimizes the accuracy of the relation. This is a topic of ongoing research.

In this work we extend standard interface functions for deterministic systems by allowing randomized actions $\mu_2 \in \mathcal{P}(\mathbb{U}_2, \mathcal{B}(\mathbb{U}_2))$. The lifting of the transition kernels for the chosen interface generates a stochastic kernel $\mathbb{W}_{\mathbb{T}}$ conditional on the values of signals in $\mathbb{U}_1$ and in $\mathbb{X}_1 \times \mathbb{X}_2$. Let us trivially extend the interface function to $\mathcal{U}_v(\mu_1, x_1, x_2) := \int_{\mathbb{U}_1} \mathcal{U}_v(u_1, x_1, x_2)\mu_1(du_1)$.

Definition 6 (probabilistic simulation). *Consider two gMDPs $\mathbf{M}_i, i = 1, 2$, $\mathbf{M}_i = (\mathbb{X}_i, \pi_i, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y})$. The gMDP $\mathbf{M}_1$ is stochastically simulated by $\mathbf{M}_2$ if there exists an interface function $\mathcal{U}_v$ and a relation $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2 \in \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$ for which there exists a Borel measurable stochastic kernel $\mathbb{W}_{\mathbb{T}}(\,\cdot\,|u_1, x_1, x_2)$ on $\mathbb{X}_1 \times \mathbb{X}_2$ given $\mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2$, such that*

1. *for all $(x_1, x_2) \in \mathcal{R}$, $h_1(x_1) = h_2(x_2)$;*
2. *for all $(x_1, x_2) \in \mathcal{R}$, for all $u_1 \in \mathbb{U}_1$, $\mathbb{T}_1(\cdot|x_1, u_1)\ \bar{\mathcal{R}}\ \mathbb{T}_2(\cdot|x_2, \mathcal{U}_v(u_1, x_1, x_2))$ with lifted probability measure $\mathbb{W}_{\mathbb{T}}(\,\cdot\,|u_1, x_1, x_2)$;*
3. *$\pi_1 \overline{\mathcal{R}} \pi_2$.*

The relationship between the two models is denoted as $\mathbf{M}_1 \preceq \mathbf{M}_2$.

The Borel measurability for both $\mathcal{U}_v$ (see above) and $\mathbb{W}_{\mathbb{T}}$ (as in this definition), which is technically needed for the well-posedness of the controller refinement, can be relaxed to universal measurability, as will be discussed in the appendix.

DEFINITION 7 (probabilistic bisimulation). *Under the same conditions as above,* $\mathbf{M}_1$ *is a probabilistic bisimulation of* $\mathbf{M}_2$ *if there exists a relation* $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ *such that* $\mathbf{M}_1 \preceq \mathbf{M}_2$ *w.r.t.* $\mathcal{R}$ *and* $\mathbf{M}_2 \preceq \mathbf{M}_1$ *w.r.t. the inverse relation* $\mathcal{R}^{-1} \subseteq \mathbb{X}_2 \times \mathbb{X}_1$. $\mathbf{M}_1$ *and* $\mathbf{M}_2$ *are said to be probabilistically bisimilar, which is denoted* $\mathbf{M}_1 \approx \mathbf{M}_2$.

For every gMDP $\mathbf{M}$: $\mathbf{M} \preceq \mathbf{M}$ and $\mathbf{M} \approx \mathbf{M}$. This can be seen by considering the diagonal relation $\mathcal{R}_{diag} = \{(x_1, x_2) \in \mathbb{X} \times \mathbb{X} \mid x_1 = x_2\}$ and selecting equal inputs for the associated interfaces. The resulting equal transition kernels $\mathbb{T}(\cdot|x,u)\bar{\mathcal{R}}_{diag}\mathbb{T}(\cdot|x,u)$ are lifted by the measure $\mathbb{W}_{\mathbb{T}}(dx_1' \times dx_2'|u, x_1, x_2) = \delta_{x_1'}(dx_2')\mathbb{T}(dx_1'|x_1, u)$, where $\delta_{x_1'}$ denotes the Dirac distribution located at x_1'.

Example 3 (lifting for diagonal relations).

(a) Consider the gMDP ($\mathbf{M}_1$) introduced in Example 1 and a slight variation of it ($\mathbf{M}_2$), given as stochastic dynamic processes,

$$\begin{aligned} &\mathbf{M}_1 : x(t+1) = f(x(t), u(t)) + e(t), && y(t) = h(x(t)),\\ &\mathbf{M}_2 : x(t+1) = f(x(t), u(t)) + \tilde{e}(t) + \tilde{u}(t), && y(t) = h(x(t)) \end{aligned}$$

with variables $x(t), x(t+1), u(t), \tilde{u}(t), e(t), \tilde{e}(t)$ taking values in $\mathbb{R}^n$, and with dynamics initialized with the same probability distribution at $t = 0$ and driven by white noise sequences $e(t), \tilde{e}(t)$, both with zero-mean normal distributions and with variances $\Sigma_e, \Sigma_{\tilde{e}}$, respectively. Notice that if $\Sigma_e - \Sigma_{\tilde{e}}$ is positive definite then $\mathbf{M}_1 \preceq \mathbf{M}_2$. To see this, select the control input pair $(u_2, \tilde{u}_2) \in \mathbb{U}_2$ as $u_2 = u_1$, and $\tilde{u}_2$ according to the zero-mean normal distribution with variance $\Sigma_e - \Sigma_{\tilde{e}}$, then the associated *interface* is $\mathcal{U}_v(\,\cdot\,|u_1, x_1, x_2) = \delta_{u_1}(du_2)\mathcal{N}(d\tilde{u}_2|0, \Sigma_e - \Sigma_{\tilde{e}})$. For this interface the stochastic dynamics of the two processes are equal, and can be lifted with $\mathcal{R}_{diag}$.

(b) Similarly as above, consider two gMDPs modeled as Gaussian processes

$$\begin{aligned} &\mathbf{M}_1 : x(t+1) = (A + BK)x(t) + Bu(t) + e(t), && y(t) = h(x(t)),\\ &\mathbf{M}_2 : x(t+1) = Ax(t) + Bu(t) + e(t), && y(t) = h(x(t)) \end{aligned}$$

with variables $x(t), x(t+1), e(t)$ taking values in $\mathbb{R}^n$ and $u(t) \in \mathbb{R}^m$, matrices $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$, $K \in \mathbb{R}^{m\times n}$. Then $\mathbf{M}_1 \preceq \mathbf{M}_2$, since in $\mathcal{R}_{diag}$ for every action u_1 chosen for $\mathbf{M}_1$, the choice of interface $u_2 = u_1 + Kx_2$ for $\mathbf{M}_2$ results in the same transition kernel for the second model.

Remark 4. Over $\mathcal{M}_{\mathbb{Y}}$, the class of gMDPs with a shared output space, the relation $\preceq$ is a preorder, since it is reflexive (see Example 3) and transitive (see, later, Corollary 1). Moreover the relation $\approx$ is an equivalence relation as it is also symmetric (as argued below in Definition 7).

3.4. Controller refinement via probabilistic simulation relations. The ideas underlying the controller refinement are first discussed, after which it is shown that the refined controller induces a strategy as per Definition 4. Finally the equivalence of properties defined over the controlled gMDPs is shown.

Consider two gMDPs $\mathbf{M}_i = (\mathbb{X}_i, \pi_i, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y})$, $i = 1, 2$, with $\mathbf{M}_1 \preceq \mathbf{M}_2$. Given the entities $\mathcal{U}_v$ and $\mathbb{W}_{\mathbb{T}}$ associated with $\mathbf{M}_1 \preceq \mathbf{M}_2$, the distribution of the next state x_2' of $\mathbf{M}_2$ is given as $\mathbb{T}_2(\cdot \mid x_2, \mathcal{U}_v(u_1, x_1, x_2))$, and is equivalently defined via the lifted measure as the marginal of $\mathbb{W}_{\mathbb{T}}(\cdot|u_1, x_1, x_2)$ on $\mathbb{X}_2$. Therefore, the distribution of the combined next state (x_1', x_2'), defined as $\mathbb{W}_{\mathbb{T}}(\,\cdot\,|u_1, x_1, x_2)$, can be expressed as

$$\mathbb{W}_{\mathbb{T}}(dx_1' \times dx_2'|u_1, x_1, x_2) = \mathbb{W}_{\mathbb{T}}(dx_1'|x_2', u_1, x_1, x_2)\mathbb{T}_2(dx_2'|x_2, \mathcal{U}_v(u_1, x_1, x_2)),$$

where $\mathbb{W}_{\mathbb{T}}(dx_1'|x_2', u_1, x_1, x_2)$ is referred to as the conditional probability given x_2' (c.f. [10, Corollary 3.1.2]).[5] Similarly, the conditional measure for the initialization $\mathbb{W}_\pi$ is denoted as $\mathbb{W}_\pi(dx_1(0) \times dx_2(0)) = \mathbb{W}_\pi(dx_1(0)|x_2(0))\pi_2(dx_2(0))$.

Now suppose that we have a control strategy for $\mathbf{M}_1$, referred to as $\mathbf{C}_1$, and we want to construct the refined control strategy $\mathbf{C}_2$ for $\mathbf{M}_2$, which is such that events defined over the output space have equal probability. This refinement procedure follows directly from the interface and the conditional probability distributions, and is described in Algorithm 2. This execution algorithm is separated into the refined control strategy $\mathbf{C}_2$ and its gMDP $\mathbf{M}_2$. $\mathbf{C}_2$ is composed of $\mathbf{C}_1$, the stochastic kernel $\mathbb{W}_{\mathbb{T}}$, and the interface $\mathcal{U}_v$, and it remembers the previous state of $\mathbf{M}_2$ (cf. line 8 in Algorithm 2).

THEOREM 1 (refined control strategy). *Let gMDPs $\mathbf{M}_1$ and $\mathbf{M}_2$ be related as $\mathbf{M}_1 \preceq \mathbf{M}_2$, and consider the control strategy $\mathbf{C}_1 = (\mathbb{X}_{\mathbf{C}_1}, x_{\mathbf{C}_1 0}, \mathbb{X}_1, \mathbb{T}^t_{\mathbf{C}_1}, h^t_{\mathbf{C}_1})$ for $\mathbf{M}_1$ as given. Then there exists at least one refined control strategy*

$$\mathbf{C}_2 = (\mathbb{X}_{\mathbf{C}_2}, x_{\mathbf{C}_2 0}, \mathbb{X}_2, \mathbb{T}^t_{\mathbf{C}_2}, h^t_{\mathbf{C}_2}),$$

as defined in Definition 4, with

- *state space $\mathbb{X}_{\mathbf{C}_2} := \mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2$ with elements $x_{\mathbf{C}_2} = (x_{\mathbf{C}_1}, x_1, x_2)$;*
- *initial state $x_{\mathbf{C}_2 0} := (x_{\mathbf{C}_1 0}, 0, 0)$;*
- *input variable $x_2 \in \mathbb{X}_2$, namely, the state variable of $\mathbf{M}_2$;*
- *time-dependent stochastic kernels $\mathbb{T}^t_{\mathbf{C}_2}$, defined as*
  $\mathbb{T}^0_{\mathbf{C}_2}(dx_{\mathbf{C}_2}|x_{\mathbf{C}_2 0}, x_2(0)) := \mathbb{T}^0_{\mathbf{C}_1}(dx_{\mathbf{C}_1}|x_{\mathbf{C}_1 0}, x_1)\mathbb{W}_\pi(dx_1|x_2)\delta_{x_2(0)}(dx_2)$ *and*
  $\mathbb{T}^t_{\mathbf{C}_2}(dx'_{\mathbf{C}_2}|x_{\mathbf{C}_2}(t), x_2(t))$
  $:= \mathbb{T}^t_{\mathbf{C}_1}(dx'_{\mathbf{C}_1}|x_{\mathbf{C}_1}, x_1')\mathbb{W}_{\mathbb{T}}(dx_1'|x_2', h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_2, x_1)\delta_{x_2(t)}(dx_2')$ *for $t \in [1, N]$;*
- *measurable output maps $h^t_{\mathbf{C}_2}(x_{\mathbf{C}_1}, \tilde{x}_1, x_2) := \mathcal{U}_v(h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_1, x_2)$.*

Algorithm 2 Refinement of control strategy $\mathbf{C}_1$ as $\mathbf{C}_2$.

1: set $t := 0$
2: draw $x_2(0)$ from π_2,
3: draw $x_1(0)$ from $\mathbb{W}_\pi(\cdot \mid x_2(0))$.
4: **loop**
5: given $x_1(t)$, select $u_1(t)$ according $\mathbf{C}_1$,
6: set $\mu_{2t} := \mathcal{U}_v(u_1(t), x_1(t), x_2(t))$,
7: draw $x_2(t+1)$ from $\mathbb{T}_2(\cdot \mid x_2(t), \mu_{2t})$,
8: draw $x_1(t+1)$ from $\mathbb{W}_{\mathbb{T}}(\cdot | x_2(t+1), u_1(t), x_1(t), x_2(t))$,
9: set $t := t + 1$.
10: **end loop**

Both the time-dependent stochastic kernels $\mathbb{T}^t_{\mathbf{C}_2}$ and the output maps $h^t_{\mathbf{C}_2}$ for $t \in [0, N]$, are universally measurable, since Borel measurable maps are universally measurable and the latter are closed under composition [7, Chap. 7].

Since, by the above construction of $\mathbf{C}_2$, the output spaces of the controlled systems $\mathbf{C}_1 \times \mathbf{M}_1$ and $\mathbf{C}_2 \times \mathbf{M}_2$ have equal distribution, it follows that measurable events have equal probability, as stated next and proved in the appendix.

[5] Beyond Borel measurability, this also holds when the kernels are universally measurable, as corresponding universally measurable regular conditional probability measures are obtained [19].

THEOREM 2. *If* $\mathbf{M}_1 \preceq \mathbf{M}_2$ *then for all control strategies* $\mathbf{C}_1$ *there exists a control strategy* $\mathbf{C}_2$ *such that, for all measurable events* $A \in \mathcal{B}\left(\mathbb{Y}^{N+1}\right)$,

$$\mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}\left(\{y_1(t)\}_{0:N} \in A\right) = \mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}\left(\{y_2(t)\}_{0:N} \in A\right).$$

The above theorem shows that probabilities of events are preserved over the controlled systems $\mathbf{C}_1 \times \mathbf{M}_1$ and $\mathbf{C}_2 \times \mathbf{M}_2$. This result can be used to infer the preservation of stochastic properties, as has been done for the similarity relation given by [16, 17, 18] and to which we make a comparison in the appendix.

4. New ϵ, δ-approximate (bi)simulation relations via lifting.

4.1. Motivation and δ-lifting. The requirement on an exact simulation relation between two models is evidently restrictive. Consider the following example, where two Markov processes have a bounded output deviation.

Example 5 (models with a shared noise source). Consider an output space $\mathbb{Y} := \mathbb{R}^d$ with a metric $\mathbf{d}_{\mathbb{Y}}(x, y) := \|x - y\|$ (the Euclidean norm), and two gMDPs expressed as noisy dynamic processes:

$$\begin{aligned} \mathbf{M}_1 &: x_1(t+1) = f(x_1(t), u_1(t)) + e_1(t), \qquad & y_1(t) &= h(x_1(t)), \\ \mathbf{M}_2 &: x_2(t+1) = f(x_2(t), u_2(t)) + e_2(t), \qquad & y_2(t) &= h(x_2(t)), \end{aligned}$$

where f and h are both globally Lipschitz, satisfying $\|f(x_1, u) - f(x_2, u)\| \leq L\|x_1 - x_2\|$ for $0 < L < 1$, and, in addition, $\|h(x_1) - h(x_2)\| \leq H\|x_1 - x_2\|$ for a $0 < H$ valid for all $x_1, x_2 \in \mathbb{R}^n$ and for all u. Suppose that the probability distributions of the random variable e_1 and of e_2 depend on a shared noise source ω, with $\omega \in \Omega$ and distribution $\mathbb{P}_\omega$, and are such that $e_1(t) = g_1(\omega(t))$ and $e_2(t) = g_2(\omega(t))$. Assume now that there exists a value $c \in \mathbb{R}$, such that $\mathbb{P}_\omega\left[\|g_1(\omega) - g_2(\omega)\| < c\right] = 1$. Then for every pair of states $x_1(t)$ and $x_2(t)$ of $\mathbf{M}_1$ and $\mathbf{M}_2$, respectively, the difference between state transitions is bounded as $\|x_1(t+1) - x_2(t+1)\| \leq L\|x_1(t) - x_2(t)\| + c$ with probability 1. By induction it can be shown that if $\|x_1(0) - x_2(0)\| \leq \frac{c}{1-L}$ then for all $t \geq 0$, $\|x_1(t) - x_2(t)\| \leq \frac{c}{1-L}$ and $\|y_1(t) - y_2(t)\| \leq \frac{cH}{1-L}$.

Even though the difference in the output of the two models is bounded by the quantity $\frac{cH}{1-L}$ with probability 1, it is impossible to provide an approximation error using either the method in [30] (hinging on stochastic stability assumptions), nor using (approximate) relations as in [17, 18]: with the former approach, for the same input sequence $u(t)$ the output trajectories of $\mathbf{M}_1$ and $\mathbf{M}_2$ have bounded difference, but do not converge to each other; with the latter approach, the relation defined via a normed difference cannot satisfy the required notion of transitivity.

As mentioned before and highlighted in the previous Example 5, we are interested in introducing a new approximate version of the notion of probabilistic simulation relation, which allows for both δ-differences in the stochastic transition kernels, and ϵ-differences in the output trajectories. For the former prerequisite, we relax the requirements on the lifting in Definition 5; subsequently, we define the resulting approximate (bi)simulation relation according to the latter prerequisite on the outputs.

DEFINITION 8 (δ-lifting for general state spaces). *Let* $\mathbb{X}_1, \mathbb{X}_2$ *be two sets with associated measurable spaces* $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)), (\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$, *and let* $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ *be a relation for which* $\mathcal{R} \in \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$. *We denote by* $\bar{\mathcal{R}}_\delta \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ *the corresponding lifted relation (acting on* $\Delta\bar{\mathcal{R}}_\delta\Theta$*), if there exists a probability space* $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ *satisfying*

1. *for all* $X_1 \in \mathcal{B}(\mathbb{X}_1)$*:* $\mathbb{W}(X_1 \times \mathbb{X}_2) = \Delta(X_1)$*;*
2. *for all* $X_2 \in \mathcal{B}(\mathbb{X}_2)$*:* $\mathbb{W}(\mathbb{X}_1 \times X_2) = \Theta(X_2)$*;*
3. *for the probability space* $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ *it holds that* $x_1 \mathcal{R} x_2$ *with probability at least* $1-\delta$ *or, equivalently, that* $\mathbb{W}(\mathcal{R}) \geq 1-\delta$*.*

We leverage Definition 8 to introduce a new approximate similarity relation that encompasses both approximation requirements, obtaining the following ϵ, δ-approximate probabilistic simulation.

DEFINITION 9 (ϵ, δ-approximate probabilistic simulation). *Consider two gMDPs* $\mathbf{M}_i = (\mathbb{X}_i, \pi_i, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y}), i = 1, 2$*, over a shared metric output space* $(\mathbb{Y}, \mathbf{d}_{\mathbb{Y}})$*.* $\mathbf{M}_1$ *is* ϵ, δ*-stochastically simulated by* $\mathbf{M}_2$ *if there exists an interface function* $\mathcal{U}_v$ *and a relation* $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ *for which there exists a Borel measurable stochastic kernel* $\mathbb{W}_{\mathbb{T}}(\,\cdot\,|u_1, x_1, x_2)$ *on* $\mathbb{X}_1 \times \mathbb{X}_2$ *given* $\mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2$*, such that*

1. *for all* $(x_1, x_2) \in \mathcal{R}$*,* $\mathbf{d}_{\mathbb{Y}}\left(h_1(x_1), h_2(x_2)\right) \leq \epsilon$*;*
2. *for all* $(x_1, x_2) \in \mathcal{R}$*, for all* $u_1 \in \mathbb{U}_1$*:* $\mathbb{T}_1(\cdot|x_1, u_1)\ \bar{\mathcal{R}}_\delta\ \mathbb{T}_2(\cdot|x_2, \mathcal{U}_v(u_1, x_1, x_2))$*, with lifted probability measure* $\mathbb{W}_{\mathbb{T}}(\,\cdot\,|u_1, x_1, x_2)$*;*
3. $\pi_1 \bar{\mathcal{R}}_\delta \pi_2$*.*

The simulation relation is denoted as $\mathbf{M}_1 \preceq_\epsilon^\delta \mathbf{M}_2$*.*

DEFINITION 10 (ϵ, δ-approximate probabilistic bisimulation). *Under the same conditions as before* $\mathbf{M}_1$ *is an* ϵ, δ*-probabilistic bisimulation of* $\mathbf{M}_2$ *if there exists a relation* $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ *such that* $\mathbf{M}_1 \preceq_\epsilon^\delta \mathbf{M}_2$ *w.r.t.* $\mathcal{R}$ *and* $\mathbf{M}_2 \preceq_\epsilon^\delta \mathbf{M}_1$ *w.r.t.* $\mathcal{R}^{-1} \subset \mathbb{X}_2 \times \mathbb{X}_1$*.* $\mathbf{M}_1$ *and* $\mathbf{M}_2$ *are said to be* ϵ, δ*-probabilistically bisimilar, denoted as* $\mathbf{M}_1 \approx_\epsilon^\delta \mathbf{M}_2$*.*

In this section we have provided similarity relations quantifying the difference between two Markov processes. The end use of the introduced similarity relations is to quantify the probability of events of a gMDP via its abstraction and to refine controllers; this is achieved in the next section.

4.2. Controller refinement via approximate simulation relations. Consider two gMDPs $\mathbf{M}_1$ and $\mathbf{M}_2$, for which $\mathbf{M}_1$ is the abstraction of the concrete model $\mathbf{M}_2$. The following result is an approximate version of Theorem 2 and presents the main result of this paper, namely, the approximate equivalence of properties defined over the gMDPs $\mathbf{M}_1$ and $\mathbf{M}_2$.

THEOREM 3. *If* $\mathbf{M}_1 \preceq_\epsilon^\delta \mathbf{M}_2$ *then for all control strategies* $\mathbf{C}_1$ *there exists a control strategy* $\mathbf{C}_2$ *such that, for all measurable events* $A \subset \mathbb{Y}^{N+1}$*,*

$$\mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}\left(\{y_1(t)\}_{0:N} \in A_{-\epsilon}\right) - \gamma \leq \mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}(\{y_2(t)\}_{0:N} \in A) \leq \mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}(\{y_1(t)\}_{0:N} \in A_\epsilon) + \gamma$$

with constant $1-\gamma := (1-\delta)^{N+1}$ *and with the* ϵ*-expansion of* A *defined as*

$$A_\epsilon := \left\{\{y_\epsilon(t)\}_{0:N} | \exists \{y(t)\}_{0:N} \in A : \max_{t \in [0,N]} \mathbf{d}_{\mathbb{Y}}(y_\epsilon(t), y(t)) \leq \epsilon\right\}$$

and similarly the ϵ*-contraction defined as* $A_{-\epsilon} := \{\{y(t)\}_{0:N} | \{\{y(t)\}_{0:N}\}_\epsilon \subset A\}$*, where* $\{\{y(t)\}_{0:N}\}_\epsilon$ *is the pointwise* ϵ*-expansion of* $\{y(t)\}_{0:N}$*.*

The above theorem allows us to reason about probabilistic properties, such as bounded safety and reachability. The extension beyond this set of properties has been left to future work. Of special interest is the computation of properties via set containment, as introduced in [29] for finite-state stochastic systems.

While the details of the proof can be found in the appendix, its key aspect is the existence of a refined control strategy $\mathbf{C}_2$, which we detail next. Given a control

strategy $\mathbf{C}_1$ over the time horizon $t \in \{0, \dots, N\}$, there is a control strategy $\mathbf{C}_2$ that refines $\mathbf{C}_1$ over $\mathbf{M}_2$. The control strategy is conceptually given in Algorithm 3. While the state (x_1, x_2) of $\mathbf{C}_2$ is in $\mathcal{R}$, the control refinement from $\mathbf{C}_1$ follows in the same way (cf. Algorithm 3, lines 4–9) as for the exact case of section 3.4. Hence, similarly to the control refinement for exact probabilistic simulations, the *basic ingredients* of $\mathbf{C}_2$ are the states x_1 and x_2, whose stochastic transition to the pair (x_1', x_2') is governed, first, by a point distribution $\delta_{x_2(t)}(dx_2')$ based on the measured state $x_2(t)$ of $\mathbf{M}_2$; and, subsequently, by the lifted probability measure $\mathbb{W}_{\mathbb{T}}(dx_1' \mid x_2', u_1, x_2, x_1)$, *conditioned* on x_2'.

On the other hand, whenever the state (x_1, x_2) leaves $\mathcal{R}$ the control chosen by strategy $\mathbf{C}_1$ cannot be refined to $\mathbf{M}_2$; instead, an alternative control strategy $\mathbf{C}_{rec}$ has to be used to control the residual trajectory of $\mathbf{M}_2$. The choice is of no importance to the result in Theorem 3. This stage of the execution (cf. Algorithm 3, lines 11–15) referred to as *recovery* makes the choice of the overall control strategy $\mathbf{C}_2$ nonunique. In practice we will only synthesize the control strategy over a finite time.

By splitting the execution in Algorithm 3 into a control strategy and a gMDP $\mathbf{M}_2$, we can again obtain the refined control strategy.

THEOREM 4 (refined control strategy). *Let gMDPs* $\mathbf{M}_1$ *and* $\mathbf{M}_2$ *with* $\mathbf{M}_1 \preceq_\epsilon^\delta \mathbf{M}_2$ *and control strategy* $\mathbf{C}_1 = (\mathbb{X}_{\mathbf{C}_1}, x_{\mathbf{C}_1 0}, \mathbb{X}_1, \mathbb{T}^t_{\mathbf{C}_1}, h^t_{\mathbf{C}_1})$ *for* $\mathbf{M}_1$ *be given. Then for any given recovery control strategy* $\mathbf{C}_{rec}$*, a refined control strategy, denoted* $\mathbf{C}_2 = (\mathbb{X}_{\mathbf{C}_2}, x_{\mathbf{C}_2 0}, \mathbb{X}_2, \mathbb{T}^t_{\mathbf{C}_2}, h^t_{\mathbf{C}_2})$*, can be obtained as an* inhomogenous Markov process *with two discrete modes of operation,* {*refinement*} *and* {*recovery*}*, based on Algorithm* 3.

Algorithm 3 Refinement of $\mathbf{C}_1$ as $\mathbf{C}_2$.

1: set $t := 0$ {Start}
2: draw $x_2(0)$ from π_2
3: draw $x_1(0)$ from $\mathbb{W}_\pi(\cdot \mid x_2(0))$
4: **while** $(x_1(t), x_2(t)) \in \mathcal{R}$ **do** {Refine}
5: given $x_1(t)$, select $u_1(t)$ from $\mathbf{C}_1$,
6: set input $\mu_{2t} := \mathcal{U}_v(u_1(t), x_1(t), x_2(t))$,
7: draw $x_2(t+1)$ from $\mathbb{T}_2(\cdot \mid x_2(t), \mu_{2t})$,
8: draw $x_1(t+1)$ from $\mathbb{W}_{\mathbb{T}}(\cdot \mid x_2(t+1), u_1(t), x_1(t), x_2(t))$,
9: set $t := t + 1$
10: **end while**
11: **loop** {Recover}
12: given $x_2(t)$, select μ_t (from $\mathbf{C}_{rec}$),
13: draw $x_2(t+1)$ from $\mathbb{T}_2(\cdot \mid x_2(t), \mu_t)$,
14: set $t := t + 1$
15: **end loop**

The details of the tuple $(\mathbb{X}_{\mathbf{C}_2}, x_{\mathbf{C}_2 0}, \mathbb{X}_2, \mathbb{T}^t_{\mathbf{C}_2}, h^t_{\mathbf{C}_2})$ are given in the appendix, together with the proof of the theorem. They follow from Algorithm 3, in a similar way as Theorem 1 follows from Algorithm 2.

4.3. Examples and properties.

Example 6 (models with a shared noise source—continued from above). Based on the relation $\mathcal{R} := \{(x_1, x_2) : \|x_1 - x_2\| \leq \frac{c}{1-L}\}$ it can be shown that $\mathbf{M}_1 \approx_\epsilon^0 \mathbf{M}_2$ with $\epsilon = \frac{Hc}{1-L}$, since, first, it holds that $\mathbf{d}_{\mathbb{Y}}(h(x_1) - h(x_2)) \leq \epsilon$ for all $(x_1, x_2) \in \mathcal{R}$ with $\mathbf{d}_{\mathbb{Y}} = \|\cdot\|$. Additionally, for all $(x_1, x_2) \in \mathcal{R}$ and for any input u_1 the selection $u_2 = u_1$ is such that $\mathbb{T}_1(\cdot|x_1, u_1)\bar{\mathcal{R}}_0\mathbb{T}_2(\cdot|x_2, u_1)$; note that $\bar{\mathcal{R}}_0$ is equal to $\bar{\mathcal{R}}$ (the

lifted relation from $\mathcal{R}$). The lifted stochastic kernel is $\mathbb{W}_{\mathbb{T}}(dx_1' \times dx_2'|u_1, x_1, x_2) := \int_\omega \delta_{f(x_1,u_1)+g_1(\omega)}(dx_1')\delta_{f(x_2,u)+g_2(\omega)}(dx_2')\mathbb{P}_\omega(d\omega)$; this stochastic kernel is Borel measurable if $f(x_1, u_1) + g_1(\omega)$ and $f(x_2, u) + g_2(\omega)$ are assumed Borel measurable maps. Note that the employed identity interface is also Borel measurable.

Example 7 (relationship to model with truncated noise). Consider the stochastic dynamical process $\mathbf{M}_1 : x(t+1) = f(x(t), u(t)) + e(t)$ with output mapping $y(t) = h(x(t))$, operating over the Euclidean state space $\mathbb{R}^n$, and driven by a white noise sequence $e(t) \in \mathbb{R}^n$ with distribution $\mathbb{P}_e$. The output space $y \in \mathbb{Y} \subseteq \mathbb{R}^d$ is endowed with the Euclidean norm $\mathbf{d}_{\mathbb{Y}} = \|\cdot\|$. Select a domain $D \subset \mathbb{R}^n$ so that, at any given time instant t, $e(t) \in D$ with probability $1 - \delta$. Then define a truncated white noise sequence $\tilde{e}(t)$ with distribution $\mathbb{P}_e(\cdot \mid D)$. The resulting model $\mathbf{M}_2$ driven by $\tilde{e}(t)$ is $\mathbf{M}_2 : x(t+1) = f(x(t), u(t)) + \tilde{e}(t)$, with the same output mapping $y(t) = h(x(t))$. We show that $\mathbf{M}_2$ is a $0, \delta$-approximate probabilistic bisimulation of $\mathbf{M}_1$, i.e., $\mathbf{M}_1 \approx_0^\delta \mathbf{M}_2$. Select $\mathcal{R} := \{(x_1, x_2) \text{ for } x_1, x_2 \in \mathbb{R}^n | x_1 = x_2\}$, and choose as interface the identity one, i.e., $\mathcal{U}_v(u_1, x_1, x_2) = u_1$. A viable lifting measure is

$$\mathbb{W}_{\mathbb{T}}(dx_1' \times dx_2'|u_1, x_1, x_2) := \int_{e\in D} \delta_{x_1'}(dx_2')\delta_{t_1(e)}(dx_1')\mathbb{P}_e(de) + \int_{e\in\mathbb{R}^n\setminus D} \delta_{t_1(e)}(dx_1')\mathbb{P}_e(de) \int_{\tilde{e}} \delta_{t_2(\tilde{e})}(dx_2')\mathbb{P}_e(d\tilde{e}|D) \tag{1}$$

with $t_1(e) = f(x_1, u_1) + e$ and $t_2(\tilde{e}) = f(x_2, u_1) + \tilde{e}$.

Example 8 (relationship between noiseless and truncated-noise models). Consider the model with truncated noise $\mathbf{M}_2$ as defined in Example 7. In what sense is $\mathbf{M}_2$ approximated by its noiseless version $\mathbf{M}_3$, namely, $\mathbf{M}_3 : x(t+1) = f(x(t), u(t))$, $y(t) = h(x(t))$? Under requirements on the Lipschitz continuity $\|f(x_1, u) - f(x_2, y)\| \leq L\|x_1 - x_2\|$ $0 < L < 1$, $\|h(x_1) - h(x_2)\| \leq H\|x_1 - x_2\|$, and on the boundedness of D and of $c = \max_{d\in D} \|d\|$, Example 5 can be leveraged by concluding that $\mathbf{M}_2 \approx_\epsilon^0 \mathbf{M}_3$ with $\epsilon = \frac{Hc}{1-L}$.[6]

In Examples 7 and 8 we have that $\mathbf{M}_1$ is approximated by $\mathbf{M}_2$, which is subsequently approximated by $\mathbf{M}_3$. The following theorem and corollary attain a quantitative answer on the question whether $\mathbf{M}_1$ is approximated by $\mathbf{M}_3$.

THEOREM 5 (transitivity of $\preceq_\epsilon^\delta$). *Consider three gMDPs $\mathbf{M}_i$, $i = 1, 2, 3$, defined by tuples $(\mathbb{X}_i, \pi_i, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y})$. If*

- *$\mathbf{M}_1$ is ϵ_a, δ_a-stochastically simulated by $\mathbf{M}_2$, and*
- *$\mathbf{M}_2$ is ϵ_b, δ_b-stochastically simulated by $\mathbf{M}_3$,*

then $\mathbf{M}_1$ is $(\epsilon_a + \epsilon_b), (\delta_a + \delta_b)$-stochastically simulated by $\mathbf{M}_3$. Equivalently, if

$$\mathbf{M}_1 \preceq_{\epsilon_a}^{\delta_a} \mathbf{M}_2 \text{ and } \mathbf{M}_2 \preceq_{\epsilon_b}^{\delta_b} \mathbf{M}_3, \text{ then } \mathbf{M}_1 \preceq_{\epsilon_a+\epsilon_b}^{\delta_a+\delta_b} \mathbf{M}_3.$$

Next, as a corollary of this theorem, we derive properties of the notion of approximate bisimulation, and discuss the transitivity of the (exact) notions of simulation and of bisimulation relations. The latter implies that the simulation relation (cf. Definition 6) is a preorder, and that the bisimulation relation (cf. Definition 7) is an equivalence relation over the category of gMDP $\mathcal{M}_{\mathbb{Y}}$.

[6] Alternatively, if $\mathbf{M}_2$ with nondeterministic input $\tilde{e} \in D$ is an ϵ_a-alternating bisimulation [42] of $\mathbf{M}_3$ then $\mathbf{M}_2 \approx_{\epsilon_a}^0 \mathbf{M}_3$.

COROLLARY 1 (transitivity properties). *Following Theorem* 5,

- *if* $\mathbf{M}_1 \approx_{\epsilon_a}^{\delta_a} \mathbf{M}_2$ *and* $\mathbf{M}_2 \approx_{\epsilon_b}^{\delta_b} \mathbf{M}_3$ *then* $\mathbf{M}_1 \approx_{\epsilon_a+\epsilon_b}^{\delta_a+\delta_b} \mathbf{M}_3$, *and*
- *if* $\mathbf{M}_1 \preceq \mathbf{M}_2$ *and* $\mathbf{M}_2 \preceq \mathbf{M}_3$ *then* $\mathbf{M}_1 \preceq \mathbf{M}_3$, *and*
- *if* $\mathbf{M}_1 \approx \mathbf{M}_2$ *and* $\mathbf{M}_2 \approx \mathbf{M}_3$ *then* $\mathbf{M}_1 \approx \mathbf{M}_3$.

Here notice that for $\mathcal{R}_{13} := \{(x_1, x_3) | \exists x_2 \in \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{12}, (x_2, x_3) \in \mathcal{R}_{23}\}$ we show that if $\Delta_1 \bar{\mathcal{R}}_{12\delta_a} \Delta_2$ and $\Delta_2 \bar{\mathcal{R}}_{23\delta_b} \Delta_3$ then $\Delta_1 \bar{\mathcal{R}}_{13(\delta_a+\delta_b)} \Delta_3$, where the used lifting measure $\mathbb{W}_{\mathbb{T}}$ is a function of the respective liftings $\mathbb{W}_{\mathbb{T}12}$ and $\mathbb{W}_{\mathbb{T}23}$, i.e., for all $x_1, x_3 \in \mathcal{R}_{13}$ $\exists x_2 \in \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{12}, (x_2, x_3) \in \mathcal{R}_{23}$, $\mathbb{W}_{\mathbb{T}}$ is given as

$$\mathbb{W}_{\mathbb{T}}(dx_1' \times dx_3' | u_1, x_1, x_2) = \int_{\mathbb{X}_2} \mathbb{W}_{23}(dx_3' | x_2', \mathcal{U}_{v12}(u_1, x_1, x_2), x_2, x_3) \mathbb{W}_{12}(dx_1' \times dx_2' | u_1, x_1, x_2).$$

Furthermore, the interface $\mathcal{U}_{v13}$ is the composition of $\mathcal{U}_{v12}$ and $\mathcal{U}_{v23}$. The proof of Theorem 5 and Corollary 1 can be found in the appendix.

Example 9 (combination of Examples 7 and 8 via Corollary 1). For the models in Examples 7 and 8 we can conclude that $\mathbf{M}_1 \approx_{\epsilon}^{\delta} \mathbf{M}_3$. This means that a stochastic system as in $\mathbf{M}_1$ in Example 7 can be approximated via its deterministic counterpart, and that the approximation error can be expressed via the probability (i.e., amount of truncation; cf. Example 7) and the output error (i.e., Example 8). This allows for explicit trading off between output deviation and deviation in probability.

5. Case studies.

5.1. Introduction: Energy management in smart buildings. We are interested in developing advanced solutions for the energy management of smart buildings. In this work we first describe a simple example with a 3-dimensional model of the thermal dynamics in an office building: we consider a simple building that is divided into two connected zones, each with a radiator affecting the heat exchange in that zone by controlling the water temperature in a boiler. With this case study we aim at elucidating the theory of the previous sections. In the third subsection we work with a more realistic model of an office building: this 5-dimensional model shows how the given approximate similarity relations can be used for the design of controllers that verifiably satisfy properties expressed as quantitative specifications. In the final subsection, we discuss how to use the approximate simulation relations for gMDPs that cannot be described by linear Gaussian processes dynamics.

5.2. First case study. A model of the temperature dynamics in an office building with two zones to heat [25, 28] assumes that the temperature fluctuations in the two zones, as well as the ambient temperature dynamics, can be modeled as a Gaussian process

$$\mathbf{M} : x(t+1) = Ax(t) + Bu(t) + Fe(t), \qquad y(t) = \left[\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{smallmatrix}\right] x(t) \tag{2}$$

with stable dynamics characterized by matrices

$$A = \begin{bmatrix} 0.8725 & 0.0625 & 0.0375 \\ 0.0625 & 0.8775 & 0.0250 \\ 0 & 0 & 0.9900 \end{bmatrix}, \quad B = \begin{bmatrix} 0.0650 & 0 \\ 0 & 0.60 \\ 0 & 0 \end{bmatrix}, \quad F = \begin{bmatrix} 0.05 & -0.02 & 0 \\ -0.02 & 0.05 & 0 \\ 0 & 0 & 0.1 \end{bmatrix},$$

where $x_{1,2}(t)$ are the temperatures in zone 1 and 2, respectively; $x_3(t)$ is the deviation of the ambient temperature from its mean; and $u(t) \in \mathbb{R}^2$ is the control input. The disturbance $e(t)$ is a white noise sequence with standard Gaussian distributions for all $t \in \mathbb{R}^+$. The state variables are initiated as $x(0) = [16\ 14\ \text{-}5]^T$. This stochastic process

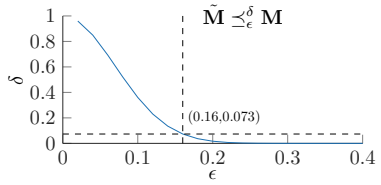

FIG. 2. *Trade-off between the output error ϵ and the probability error δ for the δ,ϵ-approximate probabilistic simulation $\tilde{\mathbf{M}} \preceq_\epsilon^\delta \mathbf{M}$. We have selected the pair $(\epsilon, \delta) = (0.16, 0.073)$ as an ideal trade-off.*

can be written as a gMDP, as detailed in Example 1. As the model abstraction, we select the controllable and deterministic dynamics of the mean of the state variables, and consequently omit the ambient temperature and the additive noise term:

$$\tilde{\mathbf{M}} : \begin{cases} \tilde{x}(t+1) & = \tilde{A}\tilde{x}(t) + \tilde{B}\tilde{u}(t) \in \mathbb{R}^2 \text{ with} \quad \tilde{A} := \left[\begin{smallmatrix} 0.8725 & 0.0625 \\ 0.0625 & 0.8775 \end{smallmatrix}\right], \\ \tilde{y}(t) & = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right]\tilde{x}(t), \qquad\qquad\qquad\quad \tilde{B} := \left[\begin{smallmatrix} 0.0650 & 0 \\ 0 & 0.60 \end{smallmatrix}\right]. \end{cases} \tag{3}$$

We then obtain that, as intuitive, $\tilde{\mathbf{M}} \preceq_\epsilon^\delta \mathbf{M}$. In order to compute specific values of ϵ and δ, we select the relation $\mathcal{R} := \{(\tilde{x}, x) \in \mathbb{R}^2 \times \mathbb{R}^3 \mid \sqrt{(\tilde{x}_1 - x_1)^2 + (\tilde{x}_2 - x_2)^2} \leq \epsilon\}$ and the interface function $\mathcal{U}_v(\tilde{u}, \tilde{x}, x) = \tilde{u} + \tilde{B}^{-1}(\tilde{A}\tilde{x} - \bar{A}x)$, with $\bar{A}V = \left[\begin{smallmatrix} 0.8725 & 0.0625 & 0.0375 \\ 0.0625 & 0.8775 & 0.0250 \end{smallmatrix}\right]$. The structure of the interface is arbitrary: in the specific instance the interface is selected to optimally correct the difference in room temperatures at the next time step.

A stochastic kernel $\mathbb{W}_\mathbb{T}$ for the lifting is $\mathbb{W}_\mathbb{T}(d\tilde{x}' \times dx' \mid \tilde{u}, \tilde{x}, x) = \int_e \delta_{\tilde{f}}(d\tilde{x}')\, \delta_{f(e)}(dx')\mathcal{N}(de \,|\, 0, I)$ with $\tilde{f} = \tilde{A}\tilde{x} + \tilde{B}\tilde{u}$ and $f(e) = Ax + B\mathcal{U}_v(\tilde{u}, \tilde{x}, x) + Fe$. The lower bound on $\mathbb{W}_\mathbb{T}(\mathcal{R} \mid \tilde{u}, \tilde{x}, x) \leq 1 - \delta$ has been computed and traded off against the output deviation, as in Figure 2.

We are interested in the goal, expressed for the model $\mathbf{M}$, of increasing the likelihood of trajectories reaching the target set $T = [20.5,\ 21]^2$ and staying there thereafter. For the abstract model we have developed a strategy, as in[25], satisfying by construction the property expressed in linear-time temporal logic-like notation with the formula $\varphi = \Diamond\Box T$ and shrunken to $\varphi_{-\epsilon}$ (as per Theorem 3). This strategy is synthesized as a correct-by-construction controller using PESSOA [33], where the discrete-time dynamics in (3) are further discretized over state and action spaces: we have selected a state quantization of 0.05 over the range $[15, 25]^2$ for the two state variables, and an input quantization of 0.05 over the set $[10, 30]^2$. It can be observed that the controller regulates the abstract model $\tilde{\mathbf{M}}$ to eventually remain within the target region, as shown in Figure 3. We now want to verify that, indeed, when refined to the concrete stochastic model, this strategy implies the reaching and staying in the safe set up to some probabilistic error. The refined strategy is obtained from this control strategy as discussed in section 4.2, and recovers from exits out of the relation $\mathcal{R}$ by resetting the abstract states in the relation.

In a simulation study reported in Figure 3, we have executed the refined control strategy over a time horizon of 200 steps. Observe that for the execution displayed in the top/left plot the behavior of the controlled concrete model $\mathbf{M}$ remains close to that of $\tilde{\mathbf{M}}$. Only at 4 incidences (circled) does the output error exceed the level $\epsilon = 0.16$. This reflects our expectations, since at any point in time the probability that the output error exceeds the level $\epsilon = 0.16$ over the following X time steps is provably less than $1 - (1 - \delta)^X \approx X\delta = 0.073X$, as per Theorem 3, which leads to an upper bound of 15 occurrences. Within this case study, whenever the state of the

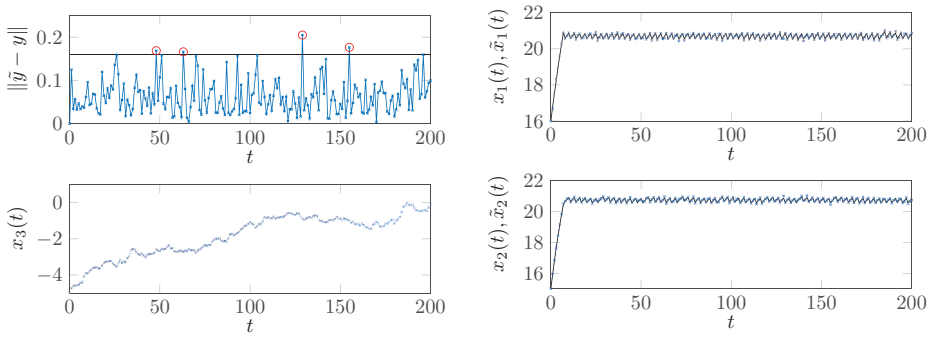

FIG. 3. *Refined control for deterministic model applied to* **M**. *The figure (top left) evaluates the accuracy of the approximation, and gives, with red circles, the instances in which the relation is left. The plot (bottom left) shows the ambient temperature. The plots on the right display the temperature inside the two rooms. The small blue crosses give the actual temperature in the rooms* (x_1, x_2) *whereas the deterministic simulation of* $(\tilde{x}_1, \tilde{x}_2)$ *is drawn in black and mostly covered by the crosses.*

abstract and concrete models leave the relation $\mathcal{R}$, then the recovery strategy consists of resetting the state of the abstract model and continuing with the refined control strategy. Thanks to the use of the ϵ-contraction $\varphi_{-\epsilon}$ of the concrete specification φ, model **M** will still abide by φ with a high confidence.

5.3. Second case study. We consider a realistic model for an office building, with the dynamics obtained from [6]. With a time sampling of 5 minutes, the following model describes stochastic temperature fluctuations around a known mean value:

$$\mathbf{M}_{\text{office}}: \quad \begin{cases} x_b(t+1) &= \Xi x_b(t) + \Gamma q(t) + B_p w_p(t) + B_s \Phi_s(t) + B_a T_a(t), \\ y(t) &= [\,0\ 1\ 0\ 0\,] x_b(t), \end{cases}$$

$$[\,\Xi|\Gamma|B_p|B_s|B_a\,] = \left[\begin{array}{cccc|c|c|c|c} 0.4487 & 0.216 & 0.2164 & 0.1186 & 2.65\text{e-}5 & 1.0939\text{e-}4 & 6.60\text{e-}4 & 2.96\text{e-}4 \\ 0.216 & 0.1778 & 0.3719 & 0.2334 & 7.45\text{e-}5 & 2.16\text{e-}4 & 1.31\text{e-}3 & 8.79\text{e-}4 \\ 0.09639 & 0.1657 & 0.6569 & 0.08082 & 2.06\text{e-}4 & 7.45\text{e-}5 & 4.49\text{e-}4 & 1.93\text{e-}4 \\ 0.005234 & 0.0103 & 0.008007 & 0.9708 & 0.07\text{e-}5 & 3.92\text{e-}6 & 2.36\text{e-}5 & 5.67\text{e-}3 \end{array}\right].$$

The output $y(t)$ models the temperature deviation of the internal air. The 4-dimensional state of the model, obtained from a frequency-based identification procedure, represents the fluctuation of internal temperatures in the building, including the building envelope and the interior [6, TiTeThTs model], where the influence of mean value dynamics have been eliminated from the model. The objective of this model is to capture the influence of stochastic effects acting upon the system and control them via the heater with input $q(t)$. The model represents the stochastic disturbances on the building temperature. We foresee three major sources of stochastic disturbance to the system, as explained next.

The first, $w_p(t)$, is the randomness of the heat generated by people in the building. An average person generates 100 Watt [W] under normal circumstances. We presume that the occupancy of the office adds a random element to this average number, which we capture as an independently and identically distributed random signal with Gaussian distribution and a standard deviation equal to 20% per person: when there are $n_p := 10$ people in the office this standard deviation becomes $\sqrt{n_p} \times 20$ [W].

The second source of stochastic disturbance is the ambient temperature, for which we model the stochastic deviation $T_a(t)$ from accurate weather forecasts. As this deviation is correlated over time, this is modeled as a first-order colored noise, with a time constant of 20 minutes. The choice of the time constant gives a measure of

correlation in time [43], so we use it to choose the time over which there is a significant correlation between successive values of $T_a(t)$. Additionally, we choose it such that the stationary variance is equal to 1, i.e., $\mathbf{E}\left[T_a(t)^2\right] = 1$. The resulting weather model is a first-order (1-dimensional) model $T_a(t+1) = 0.7788T_a + 0.6273w_w(t)$, which is driven by a white noise source with standard Gaussian distribution, namely, $w_w(t) \sim \mathcal{N}(0, I)$.

The third and final source of disturbance $\Phi_s(t)$ is the energy flow from solar radiation. Though measurable, this disturbance cannot be predicted exactly and has a high impact on the temperature inside the office. The impact depends on the effective window area of the building, which has been estimated as 6.03 [m^2] in [6]. Based on the measured solar radiation in [6], we model this disturbance as a white noise source with standard deviation of 0.1 [$\mathrm{kW/m}^2$].

Including the weather model for T_a, which requires encompassing the noise signal $w_w(t)$, leads to the following 5-dimensional model for the temperature fluctuations in the office building:

$$\mathbf{M} = (A, B, B_w, C): \quad \begin{cases} x(t+1) & = Ax(t) + B_w w(t) + Bu(t), \\ y(t) & = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \end{bmatrix} x(t), \end{cases}$$

$$\left[A \mid B \mid B_w\right] = \left[\begin{array}{ccccc|c|ccc} 0.4487 & 0.216 & 0.2164 & 0.1186 & 2.96\text{e-}4 & 0.1326 & 0.006918 & 0.06596 & 0 \\ 0.216 & 0.1778 & 0.3719 & 0.2334 & 8.789\text{e-}4 & 0.3725 & 0.01372 & 0.1308 & 0 \\ 0.09639 & 0.1657 & 0.6569 & 0.08082 & 1.928\text{e-}4 & 1.029 & 0.004712 & 0.04492 & 0 \\ 0.005234 & 0.0103 & 8.007\text{e-}3 & 0.9708 & 0.005667 & 4.309\text{e-}3 & 2.485\text{e-}4 & 0.002369 & 0 \\ 0 & 0 & 0 & 0 & 0.7788 & 0 & 0 & 0 & 0.6273 \end{array}\right].$$

In order to avoid numerical ill-conditioning issues, both the heat input $q(t)$ (expressed in kW) and the corresponding matrix Γ have been replaced by scaled versions, namely, the input signal $u(t)$ and the input matrix B. At full throttle the heating input $q(t) = 5$[kW] corresponds to the scaled input $u(t) = 1$. Similarly, the three noise sources discussed above have been normalized together with the respective system matrices, so that $w(t)$ is the new driving noise, as a white noise sequence with a standard Gaussian distribution, encompassing the unpredicted heat caused by people, solar radiation, and weather fluctuations.

We are interested in controlling the obtained stochastic system $\mathbf{M}$ to verify a quantitative property over its output signal, which is the inner air temperature. More precisely, we want to maximize the probability that the deviation of the inner air temperature stays within a 0.5 degrees difference from the nominal temperature, over an horizon of 30 minutes. This property can be encoded as a probabilistic computation tree logic (PCTL) specification for the discrete time model as follows: $\mathbb{P}_{\geq p}\left(\square^6[|y| < 0.5]\right)$, where p is a parameter to be optimized over.

In order to solve this type of probabilistic safety problems we would normally employ formal abstractions, as implemented in the software tool FAUST^2 [22]. However, a straightforward use of the tool on the nonautonomous 5-dimensional model does not yield tight guarantees. Hence, we first obtain several reduced-order models; then, over the input range of interest, we quantify the corresponding ϵ, δ-approximate probabilistic bisimulation relations; finally, we design a controller over the obtained formal abstractions with FAUST^2, and refine it to the original 5-dimensional model of the office building. In the refinement step we tune the trade-off between the conservativeness with respect to heating inputs and the accuracy of the approximation.

Model abstraction. We use model order reduction via balanced truncations, as implemented in MATLAB, to obtain lower-order approximations preserving the dynamics of interest. We seek to obtain either first- or second-order models, from two types of concrete dynamics: first, the native dynamics of model $\mathbf{M} = (A, B, B_w, C)$,

and second, the dynamics of model $\mathbf{M}' = (A + BF, B, B_w, C)$. In the latter case, the state-feedback gain F is chosen[7] so that it reduces the importance of the controllable modes of the system: $F = \begin{bmatrix} 0.48456 & 0.39865 & 0.85352 & 0.56387 & 0.0024252 \end{bmatrix}$.

As a result, we obtain four reduced-order models $\mathbf{M}_i = (A_i, B_i, B_{wi}, C_i)$ $(i = 1, 2, 3, 4)$ of $\mathbf{M}$ via balanced truncation:[8]

$$\mathbf{M}_i : \quad \begin{cases} x_s(t+1) & = A_i x_s(t) + B_{wi} w(t) + B_i u_s(t), \\ y_s(t) & = C_i x_s(t), \end{cases} \tag{4}$$

where the resulting matrices are given in the appendix.

Models $\mathbf{M}_1$ and $\mathbf{M}_3$ are obtained based on $\mathbf{M} = (A, B, B_w, C)$, whereas $\mathbf{M}_2$ and $\mathbf{M}_4$ are based on the dynamics of $\mathbf{M}' = (A + BF, B, B_w, C)$. As expected the quality of the reduced models depends on the choice of $\mathbf{M}'$ or $\mathbf{M}$: in the former case, the part of the dynamics that we cannot compensate for with a control is approximated best, whereas for $\mathbf{M}$ the most prominent dynamics are approximated best, notwithstanding how well they can be controlled.

Approximate probabilistic simulation relations. The reduced models $\mathbf{M}_1$, $\mathbf{M}_2$, $\mathbf{M}_3, \mathbf{M}_4$ are approximations of $\mathbf{M}$ and it is expected that, even when using an interface function, the error between these reduced models and $\mathbf{M}$ will increase with the input u_s. Therefore we quantify the performance of $\mathbf{M}_i$ for $i = 1, 2, 3, 4$ only over a bounded input set $\mathbb{U}_s := \{u_s \in \mathbb{R} \mid u_s^2 \leq c_1\}$. To choose a relevant c_1, suppose we would take constant c_1 of $0.25 = 0.5^2$, then this would be equal to an allowed deviation of 50 percent of the maximal input for the nominal heat input, which is 5[kW] for the original system. As we only want to correct the heating with respect to stochastic fluctuations we take the more realistic value for c_1 of $0.2^2 = 0.04$.

Let us now compute the parameters pair (ϵ, δ) establishing the relationship $\mathbf{M}_i \preceq_\epsilon^\delta \mathbf{M}$ between reduced-order and concrete models. Similarly to the work [23] on hierarchical control based on model reduction we consider a putative relation between the two state spaces as

$$\mathcal{R} := \left\{(x, x_s) \mid (x - Px_s)^T M (x - Px_s) \leq \epsilon^2\right\}$$

with properly-sized matrices M and P, satisfying the Sylvester equation $PA_i = AP + BQ$ for a choice of Q, and $C_i = CP$, and so that $M - C^T C$ is positive semidefinite, namely, $M - C^T C \succeq 0$. Introduce the interface $\mathcal{U}_v : \mathbb{U}_s \times \mathbb{X}_s \times \mathbb{X} \to \mathbb{U}$ as

$$u = Ru_s + Qx_s + K(x - Px_s),$$

and notice that $\mathcal{U}_v$ is a function of both P and Q above, alongside the additional design variables R and K (to be further discussed shortly). The interface function is chosen to reduce the differences in the observed stochastic behaviors of the two systems. It refines any choice of u_s to a control input u; as such, it implements any control strategy for $\mathbf{M}_i$ to the original model $\mathbf{M}$. In this case study we have considered a concrete model that is controllable, linear, time-invariant, and driven by an additive stochastic noise. The chosen interface $\mathcal{U}_v$, with design variables Q, K, and R, fully parameterizes the set of possible interfaces that refine controls synthesized over a reduced model that is deterministic, linear, and time-invariant, as suggested in [23].

[7] The gain term is obtained with the `dare`$(A, B, C^T C, 0.02)$ command in MATLAB.

[8] This results from the application of the `balred` function in MATLAB.

TABLE 1

ϵ, δ-simulation relation trade-off for the reduced-order models. The table gives for each model and δ the computed ϵ.

δ	1	$10^{-\frac{1}{3}}$	$10^{-\frac{2}{3}}$	10^{-1}	$10^{-\frac{4}{3}}$	$10^{-\frac{5}{3}}$	10^{-2}	$10^{-\frac{7}{3}}$	$10^{-\frac{8}{3}}$	10^{-3}
$\mathbf{M}_1$	0.1233	0.4803	0.6247	0.7347	0.827	0.9082	0.9816	1.049	1.112	1.171
$\mathbf{M}_2$	0.01445	0.1037	0.132	0.1534	0.1714	0.1871	0.2014	0.2145	0.2267	0.2381
$\mathbf{M}_3$	0.05206	0.7612	0.997	1.175	1.325	1.456	1.575	1.684	1.785	1.881
$\mathbf{M}_4$	0.1839	0.3029	0.3358	0.3604	0.3809	0.3988	0.415	0.4298	0.4435	0.4564

Let us next focus on the characterization of the relation $\mathbf{M}_i \preceq_\epsilon^\delta \mathbf{M}$. Condition 1 in Definition 9, namely, for all $(x, x_s) \in \mathcal{R} : d_{\mathbb{Y}}(y(t), y_s(t)) \leq \epsilon$, holds since $\|y - y_s\|^2 = \|Cx - CPx_s\|^2$ and $(x - Px_s)^T C^T C(x - Px_s) \leq (x - Px_s)^T M(x - Px_s)$, and the latter is bounded by ϵ^2 for $(x, x_s) \in \mathcal{R}$.

For condition 2, i.e., for all(x, x_s) and

$$\forall u_s \in \mathbb{U}_s : \mathbb{T}_s\left(\cdot \mid x_s, u_s\right) \bar{\mathcal{R}}_\delta \mathbb{T}\left(\cdot \mid x, \mathcal{U}_v(u_s, x_s, x)\right),$$

we construct a lifted probability measure $\mathbb{W}_{\mathbb{T}}(\cdot \mid u_s, x_s, x)$ based on the shared input noise $w(t)$. From this lifting measure, the original transition kernels can easily be recovered by marginalizing over $\mathbb{X}_s$ and over $\mathbb{X}$, respectively, as $\mathbb{T}\left(\cdot \mid x, u\right) = \mathcal{N}(\cdot | Ax + B\mathcal{U}_v(u_s, x_s, x), B_w B_w^T)$, and $\mathbb{T}_s\left(\cdot \mid x_s, u_s\right) = \mathcal{N}(\cdot | A_s x_s + B_s u_s, B_{wi} B_{wi}^T)$. The last condition requires that, with probability at least $1 - \delta$, the pair $(x', x_s') \in \mathcal{R}$ is distributed as $(x', x_s') \sim \mathbb{W}_{\mathbb{T}}\left(\cdot \mid u_s, x_s, x\right)$. This condition can be encoded as for all $w^T w \leq c_w$, for all $(x, x_s) \in \mathcal{R}$, for all $u_s \in \mathbb{U}_s$ it holds that $(x' - x_s') \in \mathcal{R}$. Note that the latter can be written as $(x' - Px_s')^T M(x' - Px_s') \leq \epsilon^2$, where

$$x' - Px_s' = (A + BK)(x - Px_s) + (B_w - PB_{wi})w + (BR - PB_s)u_s. \tag{5}$$

The conditions above can be expressed as a single matrix inequality via the S-procedure [11]. We know that $w \sim \mathcal{N}(0, I)$, $w^T w$ has a chi-square distribution with 2 degrees of freedom. Thus for a required level of $1 - \delta$, we select c_w as $c_w = \chi_2^{-1}(1 - \delta)$ and solve the resulting constraints with respect to ϵ for given values of K, P, Q, and R, for each of the reduced models $\mathbf{M}_i$ using the convex optimization toolbox CVX [24]. Note that χ_2^{-1} is the chi-square inverse cumulative distribution function with 2 degrees of freedom. The gains K and R are selected together with M by alternately optimizing their choice. The chosen P and Q follow from the Sylvester equation, for which additional freedom is used to minimize the influence of w and u_s in (5).

Table 1 provides a number of ϵ, δ values, derived from the approximate probabilistic simulation relation, for each of the models $\mathbf{M}_i$. Notice that for increasing values of δ, ϵ decreases to a positive lower bound: this lower bound is a function of the size of the set $\mathbb{U}_s$. Based on these outcomes, we have decided to proceed with $\mathbf{M}_2$.

Control synthesis over abstract model $\mathbf{M}_2$: use of FAUST^2. For a given choice of ϵ, δ we follow Theorem 3 and modify the given PCTL property $\psi := \mathbb{P}_{\geq p}\left(\square^6[|y| < 0.5]\right)$ to obtain $\psi_{\epsilon,\delta} := \mathbb{P}_{\geq p+\gamma}\left(\square^6[|y| < 0.5 - \epsilon]\right)$. Here γ gives the accumulation of the error in the probability over the time horizon of interest: for this case we have $1 - \gamma := (1 - \delta)^6$, which is $\gamma \approx 6\delta$. We then apply FAUST^2 to obtain a grid-based approximation of the safety probability over the six time steps of the formula (which adds up to 30 minutes in the model), with an accuracy of 0.1. More precisely, we first quantize the input space (this, on its own, generates an exact simulation), then we apply FAUST^2 [22] over the obtained continuous space, finite action

model. For this work we have optimized the algorithms in FAUST^2 to use less memory for models with Gaussian noise: by first decoupling the noise by means of a simple state transform, the storage of the discretized probability transitions can be done in a structured and more efficient manner. This leads us to perform the computations with 2.6×10^7 grid points to attain the desired accuracy of 0.1 (more precisely, 0.0983) with a 2.6 GHz Intel Core i5 with 16 GB memory within less then 20 minutes. We finally obtain that the modified safety property is satisfied with probability of at least $0.8412 - 0.0983 = 0.7429$ for the reduced-order model $\mathbf{M}_2$ initialized at zero.

Control refinement: simulation results. We refine the policy obtained from FAUST^2 for the reduced-order model $\mathbf{M}_2$ to the original model $\mathbf{M}$. Recall that we expect this refined policy to have a quantifiable safety, expressed via the property ψ, which is a requirement that the inner air temperature remains within the bound $y_s \in [-0.5, 0.5]$ of the nominal temperature during the next 30 minutes. The safety probability for the concrete model $\mathbf{M}$ initialized at the origin is lower bounded by the computed probability $p = (0.7429 - \gamma) = (0.7429 - 0.0585) = 0.6844$ (this is according to Theorem 3).

We empirically validate this result as follows. We first initialize the system and the state of the reduced-order model (in the controller) at the origin. Then we perform 10^5 Monte Carlo simulations and observe that executions of the reduced-order model remain in the modified safe set 85.81 percent of the time, whereas they exit it 14.19 percent of the time. For the same noise sequences, the controlled 5-dimensional model, where the control is refined based on the interface introduced before, stays in the *original* safe set 99.9 percent of the time, and exits it in 0.10 percent of the time. The concrete model is further seen to stay within the *modified* safe set 86.05 percent of the time, which is much closer to the computed probability for the reduced-order model. Notice that these empirical outcomes are expected to be higher than indicated in the error bounds, as these bounds are conservative especially when considering states starting in the middle of the relation.

Similarly, starting at the edge of the modified safe set $y_s \in [0.2986, -0.2986]$ of the reduced-order model, we have considered the initialization as follows: $x_s(0) = [-0.4229\ -0.2987]^T$ and $x(0) = Px_s(0)$, where P has been discussed above. For this initial state 0.7289 is the lower bound on the safety probability for the reduced-order model, and $p = 0.6704$ for the full-order model. With 10^5 empirical Monte Carlo runs, we obtain that the reduced-order model stays in the modified safe set 84.30 percent of the time, whereas the concrete model with the refined control policy stays in the safe set in 99.87 percent of the runs. Similar results were obtained upon initializing at other points on the edges of the (modified) safe set, or on the edge of the relation.

5.4. Discussion: Computing similarity relations beyond linear Gaussian dynamics. The use of interface functions and approximate similarity relations for the refinement of control strategies has been studied, amongst others, by [23] for deterministic models and by [42] for nonlinear models. In the case studies above, we have extended these results to gMDPs with Gaussian linear dynamics. Similarly, the study of approximate relations for gMDPs with more general dynamics can be tackled by direct extension of methods for approximate similarity relations on deterministic models [42]. It is likewise expected that tailored methods for the approximate stochastic simulation relations will yield less conservative and more computationally efficient results.

6. Conclusions. In this work we have discussed new and general approximate similarity relations for gMDPs, and shown that they can be effectively employed for

abstraction-based verification goals as well as for controller synthesis and refinement over quantitative specifications. The new relations, in particular, allow for a useful trade-off between the deviations in probability distribution on states and the deviations between model outputs. We have extended results on control refinement for deterministic LTI systems to construct interface functions effectively. For this and other model classes within the set of gMDPs the algorithmic construction of appropriate interface functions together with the optimal quantification of the ϵ, δ-approximate similarity relation is a topic of further research. Alongside practical applications of the developed notions, current efforts focus on further generalization of Theorem 3 to specific quantitative properties expressed via temporal logics.

Appendix A. Details on case study and use of FAUST^2. The model reduction procedure via balanced truncation[9] yields four reduced-order models $\mathbf{M}_i = (A_i, B_i, B_{wi}, C_i)$ $i = 1, 2, 3, 4$:

$$\mathbf{M}_i : \quad \begin{cases} x_s(t+1) & = A_i x_s(t) + B_{wi} w(t) + B_i u_s(t), \\ y_s(t) & = C_i x_s(t), \end{cases}$$

which are characterized by the following constant matrices

$$\begin{aligned}
&\mathbf{M}_1 : A_1 = \left[\begin{smallmatrix} 0 & -0.8572 \\ 1 & 1.857 \end{smallmatrix}\right], B_1 = \left[\begin{smallmatrix} -0.5343 \\ 0.5523 \end{smallmatrix}\right], B_{w1} = \left[\begin{smallmatrix} -5.916\text{e-}3 & -0.0564 & 8.62\text{e-}3 \\ 6.138\text{e-}3 & 0.05852 & -6.739\text{e-}3 \end{smallmatrix}\right], C_1 = \left[\begin{smallmatrix} 0 & 1 \end{smallmatrix}\right], \\
&\mathbf{M}_2 : A_2 = \left[\begin{smallmatrix} 0 & -0.05267 \\ 0.125 & -0.1081 \end{smallmatrix}\right], B_2 = \left[\begin{smallmatrix} 0.8917 \\ 0.3725 \end{smallmatrix}\right], B_{w2} = \left[\begin{smallmatrix} 0.01925 & 0.1835 & 0.002356 \\ 0.01372 & 0.1308 & 3.229\text{e-}5 \end{smallmatrix}\right], C_2 = \left[\begin{smallmatrix} 0 & 1 \end{smallmatrix}\right], \\
&\mathbf{M}_3 : A_3 = \left[\begin{smallmatrix} 0.9951 \end{smallmatrix}\right], B_3 = \left[\begin{smallmatrix} 0.1194 \end{smallmatrix}\right], B_{w3} = \left[\begin{smallmatrix} 0.001497 & 0.01427 & 0.01467 \end{smallmatrix}\right], C_3 = \left[\begin{smallmatrix} 1 \end{smallmatrix}\right], \\
&\mathbf{M}_4 : A_4 = \left[\begin{smallmatrix} 0.1203 \end{smallmatrix}\right], B_4 = \left[\begin{smallmatrix} 0.3829 \end{smallmatrix}\right], B_{w4} = \left[\begin{smallmatrix} 0.01257 & 0.1198 & 0.0002907 \end{smallmatrix}\right], C_4 = \left[\begin{smallmatrix} 1 \end{smallmatrix}\right].
\end{aligned}$$

Models $\mathbf{M}_1$ and $\mathbf{M}_3$ are obtained from $\mathbf{M} = (A, B, B_w, C)$, whereas $\mathbf{M}_2$ and $\mathbf{M}_4$ are based on the dynamics of $\mathbf{M}' = (A + BF, B, B_w, C)$. We have synthesized F to be $\left[\begin{smallmatrix} 0.4846 & 0.3986 & 0.8535 & 0.5639 & 0.002425 \end{smallmatrix}\right]$. As expected, the reduced models depend on the choice of $\mathbf{M}'$ or $\mathbf{M}$: in the former case, the part of the dynamics that we cannot compensate for with a control is approximated best, whereas for $\mathbf{M}$ the most prominent dynamics are approximated best.

Approximate probabilistic simulation relation. We quantify the performance of $\mathbf{M}_i$ for $i = 1, 2, 3, 4$ only over a bounded input set $\mathbb{U}_s := \{u_s \in \mathbb{R} \mid u_s^2 \leq c_1\}$.

Subsequently solving the Sylvester equations for Q, P, and R, tuning a stabilizing interface gain K, and then using the S-procedure as described in [11] to compute ϵ, δ (cf. Table 2), and M, we finally obtain the following matrices for the reduced-order models. For $\mathbf{M}_1$ we take $R := 1.403$, and we obtain

$$\begin{aligned}
&Q := \left[\begin{smallmatrix} -0.08954 & -0.07712 \end{smallmatrix}\right], \qquad K := \left[\begin{smallmatrix} -0.5717 & -0.4705 & -0.9859 & -0.6213 & -0.002364 \end{smallmatrix}\right], \\
&P := \begin{bmatrix} -1.061 & 0.09045 \\ 0 & 1 \\ -2.295 & -0.9696 \\ 9.064 & 8.775 \\ 0 & 0 \end{bmatrix}, \qquad M := \begin{bmatrix} 0.4797 & 0.1476 & 0.3298 & 0.1397 & -0.001306 \\ 0.1476 & 1.104 & 0.1592 & 0.06704 & -0.00359 \\ 0.3298 & 0.1592 & 0.2862 & 0.1207 & -0.001327 \\ 0.1397 & 0.06704 & 0.1207 & 0.1744 & 0.003174 \\ -0.001306 & -0.00359 & -0.001327 & 0.003174 & 0.003676 \end{bmatrix}.
\end{aligned}$$

Note that the latter is optimized for $\delta = 10^{-2}$.

For $\mathbf{M}_2$ we take $R := 1.004$, and obtain

$$\begin{aligned}
&Q := \left[\begin{smallmatrix} -1.857 & 1.406 \end{smallmatrix}\right], \qquad K := \left[\begin{smallmatrix} -0.3553 & -0.2931 & -0.65 & -0.4739 & -0.002547 \end{smallmatrix}\right], \\
&P := \begin{bmatrix} -0.6186 & 0.2348 \\ 0 & 1 \\ 2.562 & -2.314 \\ -0.009378 & 0.001329 \\ 0 & 0 \end{bmatrix}, \qquad M := \begin{bmatrix} 0.2416 & 0.06342 & 0.3159 & 0.1299 & 0.00106 \\ 0.06342 & 1.772 & 0.07267 & 0.02663 & 0.0007664 \\ 0.3159 & 0.07267 & 0.4191 & 0.1728 & 0.001395 \\ 0.1299 & 0.02663 & 0.1728 & 0.08168 & 0.000351 \\ 0.00106 & 0.0007664 & 0.001395 & 0.000351 & 0.0001456 \end{bmatrix}.
\end{aligned}$$

[9] This is obtained from the application of the `balred` function in MATLAB.

TABLE 2
Trade-off for parameters ϵ, δ in the simulation relation.

δ	1	$10^{-\frac{1}{3}}$	$10^{-\frac{2}{3}}$	10^{-1}	$10^{-\frac{4}{3}}$	$10^{-\frac{5}{3}}$	10^{-2}	$10^{-\frac{7}{3}}$	$10^{-\frac{8}{3}}$	10^{-3}
$\mathbf{M}_1$	0.1233	0.4803	0.6247	0.7347	0.827	0.9082	0.9816	1.049	1.112	1.171
$\mathbf{M}_2$	0.01445	0.1037	0.132	0.1534	0.1714	0.1871	0.2014	0.2145	0.2267	0.2381
$\mathbf{M}_3$	0.05206	0.7612	0.997	1.175	1.325	1.456	1.575	1.684	1.785	1.881
$\mathbf{M}_4$	0.1839	0.3029	0.3358	0.3604	0.3809	0.3988	0.415	0.4298	0.4435	0.4564

Again M is chosen based on the S procedure to optimize ϵ for $\delta = 10^{-2}$. For $\mathbf{M}_3$, take $R := 0.3074$ and obtain

$$Q := -0.0008755, \qquad K := \begin{bmatrix} -0.5796 & -0.477 & -0.9978 & -0.6265 & -0.00236 \end{bmatrix},$$

$$P := \begin{bmatrix} 1.004 \\ 1 \\ 1.006 \\ 0.9713 \\ 0 \end{bmatrix}, \qquad M := \begin{bmatrix} 8.584 & -4.974 & 4.929 & 2.078 & 0.1158 \\ -4.974 & 3.944 & -3.106 & -1.31 & -0.05919 \\ 4.929 & -3.106 & 3.917 & 1.653 & 0.06135 \\ 2.078 & -1.31 & 1.653 & 0.7024 & 0.02595 \\ 0.1158 & -0.05919 & 0.06135 & 0.02595 & 0.01179 \end{bmatrix}.$$

Note that M is chosen based on the S-procedure to optimize ϵ for $\delta = 10^{-2}$.

For $\mathbf{M}_4$, we take $R := 0.8996$ and

$$Q := -0.6961, \qquad K := \begin{bmatrix} -0.5307 & -0.4366 & -0.9241 & -0.5946 & -0.002391 \end{bmatrix},$$

$$P := \begin{bmatrix} -1.191 \\ 1 \\ 1.242 \\ -0.01296 \\ 0 \end{bmatrix}, \qquad M := \begin{bmatrix} 0.03949 & -0.01465 & 0.06076 & 0.02542 & 1.999e{-}05 \\ -0.01465 & 1.788 & 0.1162 & 0.05143 & -0.0005164 \\ 0.06076 & 0.1162 & 0.128 & 0.05469 & -2.765e{-}05 \\ 0.02542 & 0.05143 & 0.05469 & 0.04108 & -0.0004062 \\ 1.999e{-}05 & -0.0005164 & -2.765e{-}05 & -0.0004062 & 0.0003725 \end{bmatrix}.$$

A.1. FAUST^2 computations on a 2-dimensional model. For a given x, u pair the probability distribution of the next state is distributed with the following stochastic density kernel $t_x(\bar{x} \mid x, u) \sim \mathcal{N}(\cdot; A_i x + B_i u, \Sigma)$, where $\Sigma := B_{w_2} B_{w_2}^T$.

We resort to the algorithms implemented in [22] to maximize the probability of a stochastic event. We set up a stochastic dynamic programming scheme, leading to a final value function providing the probability of the property as

$$V_0(x) = \mathbb{P}\left[\Box^6(|y(t)| \leq 0.5 - \varepsilon)\right].$$

Define the safe set $\mathcal{A} := \mathbb{R} \times [-0.5 + \varepsilon, 0.5 - \varepsilon] \subset \mathbb{X} = \mathbb{R}^2$, then the property to be maximized can be written as $V_0(x) = \mathbb{P}\left[\Box^6 \mathcal{A}\right]$.

A.1.1. The error computation. Assume there are constants H_1, H_2, such that

$$\int_{\mathbb{R}^2} |t_x(\bar{x} \mid x, u) - t_x(\bar{x} \mid x', u)| d\bar{x} \leq H_1 |x_1' - x_1| + H_2 |x_2' - x_2|. \tag{6}$$

This gives a linearly increasing error $N(H_1\Delta_1 + H_2\Delta_2)$, where Δ_i is the grid size in the ith coordinate direction of the state space. Let us compute the two constants next. Starting from

$$t_x(\bar{x} \mid x, u) = \frac{1}{\sqrt{(2\pi)^2 \det(\sigma)}} \exp\left[-\frac{1}{2} (\bar{x} - A_i x - B_i u)^T \Sigma^{-1} (\bar{x} - A_i x - B_i u)\right],$$

define $m = \begin{bmatrix} m_1 \\ m_2 \end{bmatrix} = A_i x + B_i u$ and $\Sigma^{-1} = \begin{bmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{bmatrix} = L^T L$. Then

$$t_x(\bar{x} \mid x, u) = \frac{1}{\sqrt{(2\pi)^2 \det(\sigma)}} \exp\left[-\|L\bar{x} - Lm\|^2\right].$$

Define a change of variables with $v = L\bar{x} \to dv = |\det(L)|d\bar{x}$. Then the error computation follows from the maximal difference between the probability density distributions [22] as given in (6) and can be rewritten as follows:

$$\int_{\mathbb{R}^2} \left| \frac{1}{\sqrt{(2\pi)^2 \det(\Sigma)}} \left(\exp\left[-\frac{1}{2}\|v - Lm\|^2 \right] - \exp\left[-\frac{1}{2}\|v - Lm'\|^2 \right] \right) \right| \frac{dv}{\det(L)}.$$

Note that $\Sigma^{-1} = L^T L$, hence, $|\det(L)| = \frac{1}{\sqrt{\det(\Sigma)}}$ and, consequently,

$$= \int_{\mathbb{R}^2} \frac{1}{2\pi} \left| \left(\exp\left[-\frac{1}{2}\|v - Lm\|^2 \right] - \exp\left[-\frac{1}{2}\|v - Lm'\|^2 \right] \right) \right| dv.$$

Now we can transform a 2-dimensional integral into two 1-dimensional integrals:

$$\begin{aligned} &\leq \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} \left| \left(\exp\left[-\frac{1}{2}\|v_1 - L_1 m_1\|^2 \right] - \exp\left[-\frac{1}{2}\|v_1 - L_1 m_1'\|^2 \right] \right) \right| dv_1 \\ &\quad + \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} \left| \left(\exp\left[-\frac{1}{2}\|v_2 - L_2 m_2\|^2 \right] - \exp\left[-\frac{1}{2}\|v_2 - L_2 m_2'\|^2 \right] \right) \right| dv_2 \\ &\leq \frac{2|L_1 m - L_1 m'|}{\sqrt{2\pi}} + \frac{2|L_2 m - L_2 m'|}{\sqrt{2\pi}} \leq \frac{2}{\sqrt{2\pi}} \left(|L_1 A_i (x - x')| + |L_2 A_i (x - x')| \right). \end{aligned}$$

Define $\left[\begin{smallmatrix} \bar{a}_{11} & \bar{a}_{12} \\ \bar{a}_{21} & \bar{a}_{22} \end{smallmatrix} \right] = LA_i$. Then for (6) we have $H_1 = \frac{2}{\sqrt{2\pi}}(|\bar{a}_{11}| + |\bar{a}_{21}|)$, $H_2 = \frac{2}{\sqrt{2\pi}}(|\bar{a}_{12}| + |\bar{a}_{22}|)$.

Appendix B. Connections to literature and measurability issues. In this section we establish quantitative connections between the notion of approximate similarity that we have introduced for gMDPs and known and established concepts that have been discussed in the literature for processes that are special cases of gMDPs.

As measurability issues are key in this discussion we would like to first point out that the results in this paper can be extended to analytical spaces with universally measurable kernels. When we allow the gMDPs to have universally measurable kernels, we need to show the existence of a conditional probability measure $\mathbb{W}_{\mathbb{T}}(dx_1'|x_2', u_1, x_1, x_2)$: for this we refer to [19] which discusses the existence of universally measurable regular conditional probabilities.

B.1. Early results for Markov chains with finite-state spaces. From the perspective of testing, the concept of probabilistic bisimulation has been first introduced in [31], based on a relational notion, and later used to define equivalence between LMPs [14]. LMPs are different from gMDPs in that transitions are not governed by actions but by observable labels, and the acceptance of a label (and the consequent transition) defines the behavior of such a process. LMPs are defined over a finite-state space $\mathbb{S}$, a set of labels L, and stochastic transition kernels $\mathbb{T}_l : \mathbb{S} \times \mathbb{S} \to [0, 1]$ that are finitely indexed by $l \in L$. There is a strong relationship between LMPs and standard MDPs with labels [2], despite their different semantics.

Definition 11 (probabilistic bisimulation (relational notion)). *Let $T = (\mathbb{S}, \mathbb{P}_{l \in L}, L)$ be a labeled Markov chain, with L the finite set of labels. Then a probabilistic bisimulation $\equiv_p$ is an equivalence on $\mathbb{S}$ such that, whenever $s \equiv_p t$, the following holds:*

$$\forall l \in L : \forall A \in \mathbb{S}/\equiv_p, \sum_{s' \in A} \mathbb{T}_l(s|s') = \sum_{s' \in A} \mathbb{T}_l(t|s').$$

Two states s and t are said to be probabilistically bisimilar ($s \sim_{SL} t$) if the pair (s,t) is contained in a probabilistic bisimulation relation.

An extension of this definition is used to compare two separate processes by combining their state spaces (as a disjoint union) and defining the probabilistic bisimulation on the obtained extended state space [14]. (More details on this operation is given in the following subsection for continuous state-space models.)

For countable-state probabilistic processes combining probability and nondeterminism, [36, 37] has discussed probabilistic simulations based on a lifting notion—this has inspired the extension (over more general models) that is elaborated in this work. Over finite- or countable-state sets, [36, Lemma 8.2.2] has shown that lifting coincides with $\mathcal{R}_{eq}$-equivalence of the corresponding probability distributions.

B.2. Exact bisimulation relations for models with continuous state spaces. The early notion of bisimulation between labeled Markov chains [31] has been extended to processes (again denoted as LMPs) defined over analytical state spaces in [14], by employing zigzag morphisms. This work combines and extends earlier results on zigzag-based bisimulations [8, 13, 19], provides the fundamental measure theoretical results to support bisimulations over continuous spaces, and shows their logical characterization and their transitivity property. Alternatively, but equivalent to the zigzag definition, the follow-up work in [15] discusses an extension of the relational notion in [31], based on the concept of measurable $\mathcal{R}_{eq}$-closed sets.

Suppose that we have an LMP $\mathbf{S} = (\mathbb{X}, \mathcal{B}(\mathbb{X}), \mathbb{T}_l, L)$ with a finite label set $l \in L$ and with $\mathbb{X}$ being a Polish space. Note that, unlike the discrete-space case, this process is defined together with a Borel σ-algebra $\mathcal{B}(\mathbb{X})$. Then based on [15], an equivalence relation, denoted $\mathcal{R}_{eq}$, defines a bisimulation if for any $x_1 \mathcal{R}_{eq} x_2$ and for any measurable $\mathcal{R}_{eq}$-closed set B (or equivalently for every measurable set $B \subset \mathbb{X}/\mathcal{R}_{eq}$) it holds that

$$\mathbb{T}_l(B|x_1) = \mathbb{T}_l(B|x_2) \ \forall l \in L.$$

As an extension, a bisimulation between two different LMPs $\mathbf{S}_i = (\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \mathbb{T}_{l,i}, L)$, $i = 1, 2$, can be constructed by working on the disjoint union of their state spaces. More precisely, an equivalence relation $\mathcal{R}_{eq}$ over $\mathbb{X}_1 \sqcup \mathbb{X}_2$ defines a bisimulation if for every $x_1 \mathcal{R}_{eq} x_2$ (where $x_1 \in \mathbb{X}_1$ and $x_2 \in \mathbb{X}_2$) and for every $\mathcal{R}_{eq}$-closed set B, it holds that

$$\mathbb{T}_{l,1}(B \cap \mathbb{X}_1|x_1) = \mathbb{T}_{l,2}(B \cap \mathbb{X}_2|x_2) \ \forall l \in L.$$

An example of an equivalence relation over the disjoint union between two heterogeneous spaces, along with the induced quotient space, is given in Figure 4(a). The discussed notion of equivalence between LMPs crucially depends on the equivalence of the *probability spaces* $(\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \mathbb{P}_i)$ with probability measures $\mathbb{P}_i := \mathbb{T}_{l,i}(\cdot \mid x_i)$, given for a fixed l and state x_i. For an equivalence relation $\mathcal{R}_{eq}$ over $\mathbb{X}_1 \sqcup \mathbb{X}_2$, the probability spaces are equivalent if for every measurable $\mathcal{R}_{eq}$-closed set B it holds that

$$\mathbb{P}_1(B \cap \mathbb{X}_1) = \mathbb{P}_2(B \cap \mathbb{X}_2),$$

which is denoted as $\mathbb{P}_1 \equiv_{\mathcal{R}_{eq}} \mathbb{P}_2$.

This type of equivalence between probability spaces has also been used for bisimulation relations between control Markov processes [1], a simpler instance of the gMDP framework discussed in this work. As such, it is a natural extension of the notion in [14, 15] from LMPs to control Markov processes.

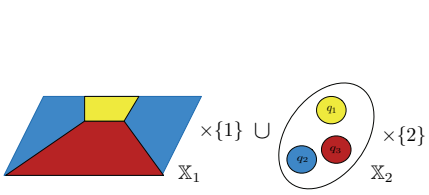

(a) An equivalence relation $\mathcal{R}_{eq}$ over the disjoint union $\mathbb{X}_1 \sqcup \mathbb{X}_2$, where two elements from each set are in the relation if they share the same color.

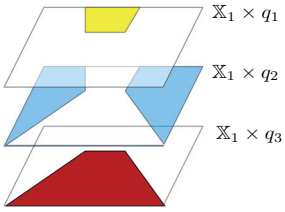

(b) Relation $\mathcal{R}$ over the Cartesian product of $\mathbb{X}_1 \subset \mathbb{R}^2$ and $\mathbb{X}_2 = \{q_1, q_2, q_3\}$, induced by the equivalence relation $\mathcal{R}_{eq}$. Elements of the relation are colored.

FIG. 4. *Example of an equivalence relation over the disjoint union of two heterogeneous spaces, and the corresponding relation over their Cartesian product.*

An equivalence relation defined over the disjoint union of $\mathbb{X}_1$ and $\mathbb{X}_2$, i.e., $\mathcal{R}_{eq} \subset (\mathbb{X}_1 \sqcup \mathbb{X}_2) \times (\mathbb{X}_1 \sqcup \mathbb{X}_2)$, can also be expressed as a relation over their Cartesian product, namely, $\mathcal{R} := \{(x_1, x_2) \in \mathbb{X}_1 \times \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{eq}\}$. As an example, we provide in Figure 4(b) the relation over the Cartesian product of two spaces, corresponding to the equivalence relation defined in Figure 4(a) over their disjoint union. This connection raises the question of whether probability spaces related via $\mathcal{R}_{eq}$ are also in a lifted relation. When working with finite or countable sets, we know that this connection holds [36]. On the other hand, for continuous or uncountable spaces this depends on the absence of measure-theoretical issues, and will be studied in depth to answer when the following claim holds.

Claim B.1. Consider two measure spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ and an equivalence relation $\mathcal{R}_{eq}$ that induces a relation over $\mathbb{X}_1 \times \mathbb{X}_2$ as $\mathcal{R} := \{(x_1, x_2) \in \mathbb{X}_1 \times \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{eq}\}$. Then,

- for any two probability measures $\Delta \in \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $\Theta \in \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$, we have
$$\Delta \bar{\mathcal{R}} \Theta \text{ if and only if } \Delta \equiv_{\mathcal{R}_{eq}} \Theta;$$
- for any two universally measurable transition kernels $\mathbb{T}_1$ and $\mathbb{T}_2$, there exists a universally measurable kernel $\mathbb{W}_{\mathbb{T}}$ that lifts the transition kernels for $\mathcal{R}$ as required in Definition 6.

In order to prove this claim and to construct the lifted measure based on an equivalence relation, we exploit the notion of zigzag morphism [14, 19] and its properties. More precisely, consider a tuple $(\mathbb{X}, \mathcal{B}(\mathbb{X}), \mathbb{T})$, with $\mathbb{X}$ a Polish space and $\mathbb{T} : \mathbb{X} \times \mathcal{B}(\mathbb{X}) \to [0, 1]$ a transition probability function.

DEFINITION 12 (morphism). *A function $f : (\mathbb{X}, \mathcal{B}(\mathbb{X}), \mathbb{T}) \to (\mathbb{X}', \mathcal{B}(\mathbb{X}'), \mathbb{T}')$ is a* morphism *if it is a continuous surjective map $f : \mathbb{X} \to \mathbb{X}'$, such that for all $s \in \mathbb{X}$ and for all $B \in \mathcal{B}(\mathbb{X})$,*
$$\mathbb{T}(f^{-1}(B)|s) = \mathbb{T}'(B|f(s)),$$
i.e., it is preserving transition probabilities.

Consider two LMPs $\mathbf{S}_i = (\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \{k_{l,i} | l \in L\})$ with a shared finite set of labels L, then a morphism f is a *zigzag morphism* if it preserves the two transition probability functions for all $l \in L$. Two LMPs $\mathbf{S}_1$ and $\mathbf{S}_2$ are *probabilistically bisimilar* if there is a generalized span of zigzag morphisms between them [14]; namely, if there exists

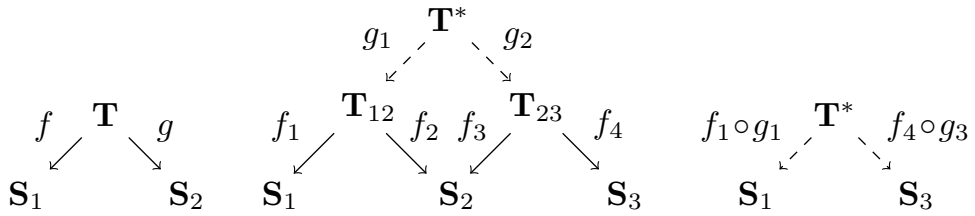

(a) Generalized span of zigzag morphisms

(b) Construct $\mathbf{T}^*$ as a semi-pullback of co-span $\mathbf{T}_{12} \to \mathbf{S}_2 \leftarrow \mathbf{T}_{22}$.b

(c) Transitive bisimulation based on semi-pullback

FIG. 5. *Probabilistic bisimulation between* $\mathbf{S}_1$ *and* $\mathbf{S}_2$ *established by zigzag morphism. Transitivity of probabilistic bisimulations* $\mathbf{S}_1$ *and* $\mathbf{S}_2$ *and* $\mathbf{S}_2$ *and* $\mathbf{S}_3$ *follows as a semi-pullback.*

an LMP $\mathbf{T}$ (with universally measurable transition kernels) and zigzag morphisms f and g from $\mathbf{T}$ to $\mathbf{S}_1$ and $\mathbf{S}_2$, respectively (see Figure 5(a)). In order to prove that this notion of probabilistic bisimulation is transitive, [19] has shown that

- the category of Markov processes with universally measurable transition probability functions $\mathbb{T}$ on Polish spaces and with surjective and continuous transition probability preserving maps has *semipullbacks* [19, Corollary 5.3];
- the category of probability measures $\mathbb{P}$ on Polish spaces and measure-preserving surjective maps has semipullbacks [19, Corollary 5.4].

By adding a labeling to the transition probability function $\mathbb{T}$, one can trivially show the existence of semipullbacks on an LMP. Moreover, the transitivity of probabilistic bisimulations follows based on semipullbacks: if $\mathbf{S}_1$ is probabilistically bisimilar to $\mathbf{S}_2$, which is also bisimilar to $\mathbf{S}_3$, then $\mathbf{S}_1$ and $\mathbf{S}_3$ are bisimilar, as in Figure 5(b).

Let us go back to Claim B.1. First, recall that, as depicted in Figure 4(a), an equivalence relation $\mathcal{R}_{eq}$ over $\mathbb{X}_1 \sqcup \mathbb{X}_2$ induces a quotient space, denoted by $\mathcal{Q} := (\mathbb{X}_1 \sqcup \mathbb{X}_2)/\mathcal{R}_{eq}$, and partitions the unionized state space by disjoint sets, namely, $\bigcup_{q\in\mathcal{Q}} q = \mathbb{X}_1 \sqcup \mathbb{X}_2$ and $q_1 \cap q_2 = \emptyset$ for $q_1 \neq q_2$, $q_1, q_2 \in \mathcal{Q}$. Thus starting from the Markov processes $\mathbf{S}_1 = (\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1), \mathbb{T}_1)$ and $\mathbf{S}_2 = (\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2), \mathbb{T}_2)$, we show that the claim holds under either of the following two conditions.

CONDITION 1 (Polish quotient space). *The equivalence relation of interest* $\mathcal{R}_{eq}$ *induces a quotient space* $(\mathcal{Q}, \mathcal{F})$ *that is Polish and the maps from* $\mathbb{X}_1$ *and* $\mathbb{X}_2$ *to the quotient space* $f_1 : \mathbb{X}_1 \to \mathcal{Q}$ *and* $f_2 : \mathbb{X}_2 \to \mathcal{Q}$ *are measurable and surjective.*

CONDITION 2 (analytic Borel quotient space). *The equivalence relation of interest* $\mathcal{R}_{eq}$ *induces a quotient space that is analytical as in* [14, 19] *and the maps from* $\mathbb{X}_1$ *and* $\mathbb{X}_2$ *to the quotient space* $f_1 : \mathbb{X}_1 \to \mathcal{Q}$ *and* $f_2 : \mathbb{X}_2 \to \mathcal{Q}$ *are measurable and surjective.*

Notice that Condition 1 implies Condition 2, and further note that f_1 and f_2 are constructed based on the injection ι_1 and ι_2, i.e., $\iota_i : \mathbb{X}_i \to \mathbb{X}_1 \sqcup \mathbb{X}_2$ for $i = 1, 2$, composed with $q : \mathbb{X}_1 \sqcup \mathbb{X}_2 \to Q$.

Then we can construct the quotient Markov process as the tuple $\mathbf{S} := (\mathcal{Q}, \mathcal{F}, \mathbb{T})$ such that $(\mathcal{Q}, \mathcal{F})$ is a Borel measurable space with $\mathcal{Q} = (\mathbb{S}_1 \sqcup \mathbb{S}_2)/\mathcal{R}_{eq}$, and $\mathcal{F}$ is defined as $\mathcal{F} := \{E \subset \mathcal{Q} : q^{-1}(E) \in \mathcal{B}(\mathbb{S}_1 \sqcup \mathbb{S}_2)\}$. The stochastic transition kernel $\mathbb{T}$ is constructed as in [14, Proof of Proposition 9.4]. For any $B \in \mathcal{F}$ it holds that

$$\mathbb{T}(B|t) = \mathbb{T}_1(f_1^{-1}(B)|s) \quad \text{with } s \in f_1^{-1}(t) \tag{7}$$

and $\mathbb{T}(B|\cdot)$ is Borel measurable.

Then f_1 and f_2 are zigzag morphisms from, respectively, $\mathbf{S}_1$ and $\mathbf{S}_2$ to $\mathbf{S}$, and they form a cospan. Based on [19] we now know that there exists a Markov process $\mathbf{W} := ((\mathbb{X}_1 \times \mathbb{X}_2), \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$, which is a semipullback, and where $\mathbb{W}$ lifts the relation over $\mathbb{X}_1 \times \mathbb{X}_2$ and defines a universally measurable stochastic kernel. If $\mathbf{S_1}$, $\mathbf{S_2}$, and $\mathbf{S}$ have analytical Borel spaces (this includes Polish spaces) and universally measurable transition kernels then $\mathbb{W} : \mathcal{R} \times \mathcal{B}(\times)$ is defined as

$$\mathbb{W}\left(dx_1' \times dx_2' \mid (x_1, x_2)\right) = \int_{q' \in Q} \mathbb{T}_1(dx_1' \mid x_1, q')\mathbb{T}_2(dx_2' \mid x_2, q')\mathbb{T}(dq' \mid f_1(x_1)), \tag{8}$$

where $\mathbb{T}_i(dx_i' \mid x_i, q')$ for $i = 1, 2$ are universally measurable regular conditional probability distributions, such that for measurable subsets $X_i \subset \mathbb{X}_i$ and $Q \subset \mathcal{Q}$ it holds that

$$\mathbb{T}_i(X_i \cap f_1^{-1}(Q) \mid x_i) = \int_Q \mathbb{T}_i(dx_i' \mid x_i, q')\mathbb{T}(dq' \mid f_1(x_1)).$$

The details of this reasoning follow from [19] together with the existence proof for the regular conditional probability distributions.

Remark 10 (measurability assumptions). The measurability assumption above is a nontrivial but natural assumption, since, as proven for LMPs, any equivalence relation on $\mathbb{X}_1 \sqcup \mathbb{X}_2$ based on logics induces a quotient LMP that has an analytical Borel space and measurable canonical maps [14, Proposition 9.4].

B.3. Approximate probabilistic bisimulation relations. A relaxation of exact equivalence relations in a probabilistic context has been introduced first for (finite-state) labeled Markov chains in [16], and later employed in [18].

DEFINITION 13. *A relation $\mathcal{R} \subseteq S \times S$ is a (probabilistic) ϵ-simulation if whenever $s\mathcal{R}t$ then for all labels $l \in L$, and sets in the event space $X \in \Sigma$, it holds that*

$$\mathbb{T}_l(\mathcal{R}(X)|t) \geq \mathbb{T}_l(X|s) - \epsilon.$$

Note that the relation is not required to be an equivalence relation, hence, it does not induce a partitioning of the state space. For continuous-space systems, [1] has discussed an approximate (bi)simulation notion derived from the finite-state definition. This definition relates to an approximate equivalence of the probability spaces $(\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \mathbb{P}_i)$, $i = 1, 2$, as follows. For an equivalence relation $\mathcal{R}_{eq}$ over $\mathbb{X}_1 \sqcup \mathbb{X}_2$ the probability spaces are approximately equivalent if for every measurable $\mathcal{R}_{eq}$-closed set B it holds that

$$|\mathbb{P}_1(B \cap \mathbb{X}_1) - \mathbb{P}_2(B \cap \mathbb{X}_2)| \leq \delta,$$

which is denoted as $\mathbb{P}_1 \equiv^{\delta}_{\mathcal{R}_{eq}} \mathbb{P}_2$.

THEOREM 6. *Consider two measure spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ and an equivalence relation $\mathcal{R}_{eq}$ satisfying Condition* 1. *Then for any two probability measures $\Delta \in \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $\Theta \in \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ we have that*

$$\Delta \equiv^{\delta}_{\mathcal{R}_{eq}} \Theta \text{ if and only if } \Delta \bar{\mathcal{R}}_{\delta} \Theta$$

with, as standard, $\mathcal{R} := \{(x_1, x_2) \in \mathbb{X}_1 \times \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{eq}\}$.

Proof.

1. $\Delta \bar{\mathcal{R}}_\delta \Theta \implies \Delta \equiv^{\delta}_{\mathcal{R}_{eq}} \Theta$.

If $\Delta \bar{\mathcal{R}}_\delta \Theta$ then for each $C \subset (\mathbb{X}_1 \sqcup \mathbb{X}_2)/\mathcal{R}_{eq}$ with subsets $\tilde{S} = \mathbb{X}_1 \cap C \in \mathcal{B}(\mathbb{X}_1)$ and $\tilde{T} = \mathbb{X}_2 \cap C \in \mathcal{B}(\mathbb{X}_2)$, $|\Delta(\tilde{S}) - \Theta(\tilde{T})| \leq \delta$ because $\mathbb{W}(\tilde{S} \times (\mathbb{X}_2 \setminus \tilde{T})) \leq \delta$ and $\mathbb{W}((\mathbb{X}_1 \setminus \tilde{S}) \times \tilde{T}) \leq \delta$. This can be shown as follows:

$$\Delta(\tilde{S}) \leq \Delta(\tilde{S}) + \mathbb{W}((\mathbb{X}_1 \setminus \tilde{S}) \times \tilde{T}) = \Theta(\tilde{T}) + \mathbb{W}(\tilde{S} \times (\mathbb{X}_2 \setminus \tilde{T})) \leq \Theta(\tilde{T}) + \delta$$

and, repeating the reasoning starting from $\Theta(\tilde{T})$, we get $\Theta(\tilde{T}) \leq \Delta(\tilde{S}) + \delta$ and $|\Delta(\tilde{S}) - \Theta(\tilde{T})| \leq \delta$.

2. $\Delta \equiv^{\delta}_{\mathcal{R}_{eq}} \Theta \implies \Delta \bar{\mathcal{R}}_\delta \Theta$.

Under Condition 1 we have that the quotient space has the Borel measure space $(\mathcal{Q}, \mathcal{F})$, where $\mathcal{Q}$ is Polish. Additionally we have measurable mappings $f_i : \mathbb{X}_1 \to \mathcal{Q}$. We denote the induced probability measures $f_{1*}\Delta \in \mathcal{P}(\mathcal{Q}, \mathcal{F})$ and $f_{2*}\Theta \in \mathcal{P}(\mathcal{Q}, \mathcal{F})$. Denote a measure that lifts these over the diagonal relation as $\mathbb{W}_{\mathcal{Q}} \in \mathcal{P}(\mathcal{Q}^2, \mathcal{F}^2)$. This is equivalent to maximal coupling of $f_{1*}\Delta$ and $f_{2*}\Theta$. Specifically for Polish spaces we take the γ-coupling given as $\mathbb{W}_{\mathcal{Q}} := \gamma(f_{1*}\Delta, f_{2*}\Theta) \in \mathcal{P}(\mathcal{Q}^2, \mathcal{F}^2)$ [4] based on [32, section 1.5] and given as follows.

DEFINITION B.2. *Let Z be a Borel space and let $\nu, \tilde{\nu} \in (Z)$ be two probability measures on it. The γ-coupling of $(\nu, \tilde{\nu})$ is a measure $\gamma \in (Z^2)$ given by*

$$\gamma(\nu, \tilde{\nu}) := \Psi_Z(\nu \wedge \tilde{\nu}) + \mathbf{1}_{[0,1)}(\|\nu \wedge \tilde{\nu}\|).\frac{(\nu - \tilde{\nu})^+ \otimes (\nu - \tilde{\nu})^-}{1 - \|\nu - \tilde{\nu}\|},$$

where $\Psi_Z : Z \to Z^2$ is the diagonal map on Z given by $\Psi_Z : z \mapsto (z, z)$.

The lifted measure over $\mathbb{W} \in \mathcal{P}(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2))$ is given as

$$\mathbb{W} := \int_{Q \times Q} \Delta(dx_1 \mid q_1)\Theta(dx_2 \mid q_2)\mathbb{W}_{\mathcal{Q}}(dq_1 \times dq_2).$$

∎

Appendix C. Proofs of theorems and corollaries.

C.1. Control refinement proofs, Theorems 1–4. Let us consider the controller refinement for exact simulation relations first. The execution

$$\{(x_2(t), x_{\mathbf{C}_2}(t)) | t \in [0, N]\}$$

is defined on the canonical space $\Omega = (\mathbb{X}_2 \times \mathbb{X}_{\mathbf{C}_2})^{N+1}$, and has a unique probability measure $\mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}$. Therefore in Algorithm 1, in order to write the execution of the refined control $\mathbf{C}_2$ and of the gMDP $\mathbf{M}_2$, we have included the state of $\mathbf{M}_2$ for one transition in the state of the refined control strategy. Therefore, while the execution of Algorithm 1 ranges over $\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2$, the execution of the controlled system with $\mathbf{C}_2$ ranges over $\mathbb{X}_{\mathbf{C}_2} \times \mathbb{X}_2 = (\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2) \times \mathbb{X}_2$. The marginal of $\mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}$ on $\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2$ defines the measure for the execution in Algorithm 1.

Since, by the above construction of $\mathbf{C}_2$, the output spaces of the closed loop systems $\mathbf{C}_1 \times \mathbf{M}_1$ and $\mathbf{C}_2 \times \mathbf{M}_2$ have equal distributions, it follows that measurable events have equal probability, as stated next.

Proof of Theorem 2. If $\{h_1(x_1(t)) | t \in [0, N]\} \in A$ and $(x_1(t), x_2(t)) \in \mathcal{R}$ for all $t \in [0, N]$ then $\{h_2(x_2(t)) | t \in [0, N]\} \in A$.

Let us rewrite the stochastic kernel of the combined transition of $\mathbf{C}_2$ and $\mathbf{M}_2$ for $t = 0$ as[10]

$$\mathbb{T}^0_{\mathbf{C}_2\times\mathbf{M}_2}(dx_{\mathbf{C}_2} \times dx_2) = \mathbb{T}^0_{\mathbf{C}_1}(dx_{\mathbf{C}_1}|x_{\mathbf{C}_1 0}, x_1)\mathbb{W}_\pi(dx_1|x_2)\delta_{x_2(0)}(dx_2)\pi(dx_2(0)).$$

Marginalized on $\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2$, this becomes (by definition of $\mathbb{W}_\pi$)

$$\begin{aligned}\mathbb{T}^0_{\mathbf{C}_2\times\mathbf{M}_2}(dx_{\mathbf{C}_1} \times dx_1 \times dx_2) &= \mathbb{T}^0_{\mathbf{C}_1}(dx_{\mathbf{C}_1}|x_{\mathbf{C}_1 0}, x_1)\mathbb{W}_\pi(dx_1|x_2)\pi(dx_2)\\ &= \mathbb{T}^0_{\mathbf{C}_1}(dx_{\mathbf{C}_1}|x_{\mathbf{C}_1 0}, x_1)\mathbb{W}_\pi(dx_2|x_1)\pi(dx_1).\end{aligned}$$

Further marginalized on $\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1$, this becomes

$$\mathbb{T}^0_{\mathbf{C}_2\times\mathbf{M}_2}(dx_{\mathbf{C}_1} \times dx_1) = \mathbb{T}^0_{\mathbf{C}_1}(dx_{\mathbf{C}_1}|x_{\mathbf{C}_1 0}, x_1)\pi(dx_1) = \mathbb{T}^0_{\mathbf{C}_1\times\mathbf{M}_1}(dx_{\mathbf{C}_1} \times dx_1).$$

For $t \in [1, N]$, the stochastic kernel marginalized on $\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2$ is

$$\begin{aligned}&\mathbb{T}^t_{\mathbf{C}_2\times\mathbf{M}_2}(dx'_{\mathbf{C}_1} \times dx'_1 \times dx'_2)\\ &\quad= \mathbb{T}^t_{\mathbf{C}_2}(dx'_{\mathbf{C}_1}|x_{\mathbf{C}_1}, x'_1)\mathbb{W}_\mathbb{T}(dx'_1|x'_2, h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_2, x_1)\mathbb{T}_2(dx'_2|x_2, h^t_{\mathbf{C}_2}(x_{\mathbf{C}_2}))\\ &\quad= \mathbb{T}^t_{\mathbf{C}_1}(dx'_{\mathbf{C}_1}|x_{\mathbf{C}_1}, x'_1)\mathbb{W}_\mathbb{T}(dx'_1 \times dx'_2|h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_2, x_1)\end{aligned}$$

and can be further marginalized on $\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1$ to obtain $\mathbb{T}^t_{\mathbf{C}_1\times\mathbf{M}_1}$. Note that since $\mathbb{W}_\mathbb{T}(\mathcal{R}|h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_2, x_1) = 1$ for $(x_1, x_2) \in \mathcal{R}$ it holds with probability 1 that $(x_1(t), x_2(t)) \in \mathcal{R}$ for $t \in [0, N]$. Therefore we can deduce that

$$\mathbb{P}_{\mathbf{C}_1\times\mathbf{M}_1}\left(\{y_1(t)\}_{0:N} \in A\right) = \mathbb{P}_{\mathbf{C}_2\times\mathbf{M}_2}\left(\{y_2(t)\}_{0:N} \in A\right). \qquad \square$$

To prove Theorems 4 and 3 we leverage their exact versions (Theorems 1 and 2). We first show the existence of a refined control strategy in the case of an approximate simulation relation; c.f. Theorem 4. Then we leverage these results to prove Theorem 3.

Theorem 4 states the following. Let gMDPs $\mathbf{M}_1$ and $\mathbf{M}_2$ with $\mathbf{M}_1 \preceq^\delta_\epsilon \mathbf{M}_2$ and control strategy $\mathbf{C}_1 = (\mathbb{X}_{\mathbf{C}_1}, x_{\mathbf{C}_1 0}, \mathbb{X}_1, \mathbb{T}^t_{\mathbf{C}_1}, h^t_{\mathbf{C}_1})$ for $\mathbf{M}_1$ be given. Then for every given recovery control strategy $\mathbf{C}_{rec}$, a refined control strategy $\mathbf{C}_2 = (\mathbb{X}_{\mathbf{C}_2}, x_{\mathbf{C}_2 0}, \mathbb{X}_2, \mathbb{T}^t_{\mathbf{C}_2}, h^t_{\mathbf{C}_2})$ can be obtained as an *inhomogenous Markov process* with two discrete modes of operation, {refinement} and {recovery}, based on Algorithm 2. More specifically, a possible choice of a refined control strategy is built up as follows:

- state space $\mathbb{X}_{\mathbf{C}_2} := \{\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2 \times \{\text{refine}\}\} \cup \mathbb{X}_{\mathbf{C}_{rec}} \times \{\text{recover}\}$ with elements $x_{\mathbf{C}_2} = (x_{\mathbf{C}_1}, x_1, x_2, \text{refine})$ and $x_{\mathbf{C}_2} = (x_{\mathbf{C}_{rec}}, \text{recover})$;
- initial state $x_{\mathbf{C}_2 0} := (x_{\mathbf{C}_1 0}, 0, 0, \text{refinement})$;
- control inputs $x_2 \in \mathbb{X}_2$ accepted;
- time dependent stochastic kernel $\mathbb{T}^t_{\mathbf{C}_2}$, defined for $t = 0$ as

$$\begin{aligned}\mathbb{T}^0_{\mathbf{C}_2}(dx^{\text{refine}}_{\mathbf{C}_2}|x_{\mathbf{C}_2 0}, x_2(0)) :=& \mathbb{T}^0_{\mathbf{C}_1}(dx_{\mathbf{C}_1}|x_{\mathbf{C}_1 0}, x_1)\mathbf{1}_\mathcal{R}(x_1, x_2)\\ &\times \mathbb{W}_\pi(dx_1|x_2)\delta_{x_2(0)}(dx_2),\\ \mathbb{T}^0_{\mathbf{C}_2}(dx^{\text{recover}}_{\mathbf{C}_2}|x_{\mathbf{C}_2 0}, x_2(0)) :=& \mathbb{T}^0_{init,rec}(dx_{\mathbf{C}_{rec}}|x_2)\mathbf{1}_{(\mathbb{X}_1\times\mathbb{X}_2)\setminus\mathcal{R}}(x_1, x_2)\\ &\times \mathbb{W}_\pi(dx_1|x_2)\delta_{x_2(0)}(dx_2),\end{aligned}$$

[10] For brevity, a part of the argument of the stochastic kernel has been omitted.

and for $t \in [1, N]$ over the {refine} operating mode

$$\mathbb{T}^t_{\mathbf{C}_2}(dx^{\text{refine}'}_{\mathbf{C}_2} | x^{\text{refine}}_{\mathbf{C}_2}(t), x_2(t)) := \mathbb{T}^t_{\mathbf{C}_1}(dx'_{\mathbf{C}_1} | x_{\mathbf{C}_1}, x'_1)\mathbf{1}_{\mathcal{R}}(x'_1, x'_2) \times \mathbb{W}_{\mathbb{T}}(dx'_1 | x'_2, h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_2, x_1)\delta_{x_2(t)}(dx'_2),$$

$$\mathbb{T}^t_{\mathbf{C}_2}(dx^{\text{recover}'}_{\mathbf{C}_2} | x^{\text{refine}}_{\mathbf{C}_2}(t), x_2(t)) := \mathbb{T}^t_{init,rec}(dx'_{\mathbf{C}_{rec}} | x'_2)\mathbf{1}_{(\mathbb{X}_1 \times \mathbb{X}_2)\setminus\mathcal{R}}(x'_1, x'_2) \times \mathbb{W}_{\mathbb{T}}(dx'_1 | x'_2, h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_2, x_1)\delta_{x_2(t)}(dx'_2),$$

defined based on a stochastic kernel $\mathbb{T}^t_{init,rec}$ $t \in [0, N]$ initiates the recovery strategy on the fly and is contained in the choice of recovery strategy; and for $t \in [1, N]$ for the recover operating mode

$$\mathbb{T}^t_{\mathbf{C}_2}(dx^{\text{recover}'}_{\mathbf{C}_2} | x^{\text{recover}}_{\mathbf{C}_2}(t), x_2(t)) := \mathbb{T}^t_{\mathbf{C}_{rec}}(dx'_{\mathbf{C}_{rec}} | x_{\mathbf{C}_{rec}}(t), x_2(t));$$

- universally measurable output map

$$h^t_{\mathbf{C}_2}(x_{\mathbf{C}_2}) := \begin{cases} \mathcal{U}_v(h^t_{\mathbf{C}_1}(x_{\mathbf{C}_1}), x_1, x_2) & \text{for refine}, \\ h^t_{\mathbf{C}_{rec}}(x_{\mathbf{C}_{rec}}) & \text{for recover}. \end{cases}$$

The refined control strategy is composed of the control strategy $\mathbf{C}_1$, the recovery strategy $\mathbf{C}_{rec}$, the stochastic kernel $\mathbb{W}_{\mathbb{T}}$, and the interface $\mathcal{U}_v$. Both the time-dependent stochastic kernels $\mathbb{T}^t_{\mathbf{C}_2}$ and the output maps $h^t_{\mathbf{C}_2}$ for $t \in [0, N]$, can be shown to be universally measurable, since Borel measurable maps (and kernels) are universally measurable and the latter are closed under composition [7, Chap. 7].

Now we need to use this control strategy to prove Theorem 3.

Proof of Theorem 3. Given $\mathbf{C}_{rec}$ consider an auxiliary recover strategy $\mathbf{C}^*_{rec}$ such that it has stochastic kernels over $\mathbb{X}_{\mathbf{C}_{rec}} \times \mathbb{X}_1 \times \mathbb{X}_{\mathbf{C}_1}$:

$$\mathbb{T}^t_{\mathbf{C}^*_{rec}}(dx'_{\mathbf{C}^*_{rec}} | x_{\mathbf{C}^*_{rec}}(t), x_2(t)) = \mathbb{T}^t_{\mathbf{C}_{rec}}(dx'_{\mathbf{C}_{rec}} | x_{\mathbf{C}_{rec}}(t), x_2(t))\mathbb{T}^t_{\mathbf{C}_1 \times \mathbf{M}_1}(dx'_{\mathbf{C}_1 \times \mathbf{M}_1} | x_{\mathbf{C}_1 \times \mathbf{M}_1}(t))$$

where $\mathbb{T}^t_{\mathbf{C}_1 \times \mathbf{M}_1}(dx'_{\mathbf{C}_1 \times \mathbf{M}_1} | x_{\mathbf{C}_1 \times \mathbf{M}_1}(t))$ is the stochastic kernel over $\mathbb{X}_{\mathbf{C}_1 \times \mathbf{M}_1} := \mathbb{X}_1 \times \mathbb{X}_{\mathbf{C}_1}$. Due to the independence of this kernel the probability distribution $\mathbb{P}_{\mathbf{C}^*_2 \times \mathbf{M}_2}$ of $\mathbf{M}_2$ controlled by $\mathbf{C}^*_2$ is, when marginalized on the canonical sample space $(\mathbb{X}_{\mathbf{C}_2} \times \mathbb{X}_{\mathbf{M}_2})^{N+1}$, equal to $\mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}$.

Now using the same arguments as in the proof of Theorem 2 we know that for all measurable sets $L \subset \mathbb{Y}^{N+1}$

$$\mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}(\{h_1(x_1(t))\}_{0:N} \in L) = \mathbb{P}_{\mathbf{C}^*_2 \times \mathbf{M}_2}(\{h_1(x_1(t))\}_{0:N} \in L).$$

The probability

$$\mathbb{P}_{\mathbf{C}^*_2 \times \mathbf{M}_2}((x_1(t), x_2(t)) \in \mathcal{R} \text{ for } t \in [0, N]) \geq (1 - \delta)^{N+1}.$$

This can be shown by induction starting from $t = 0$, and by showing that at every time step and for every pair of states the probability of staying in $\mathcal{R}$ is at least $1 - \delta$. Now note that if $\{h_1(x_1(t))\} \in A_{-\epsilon}$ and $(x_1(t), x_2(t)) \in \mathcal{R}$ for $t \in [0, N]$ then $\{y(t)\}_{0:N} \in A$. As a consequence

$$\mathbb{P}_{\mathbf{C}^*_2 \times \mathbf{M}_2}(\{h_1(x_1(t))\}_{0:N} \in A_{-\epsilon} \wedge (x_1(t), x(t)) \in \mathcal{R} \text{ for } t \in [0, N]) \leq \mathbb{P}_{\mathbf{C}^*_2 \times \mathbf{M}_2}(\{h_2(x_2(t))\}_{0:N} \in A) = \mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}(\{h_2(x_2(t))\}_{0:N} \in A).$$

Now using the union bounding argument we also have that

$$\begin{aligned}
&\mathbb{P}_{\mathbf{C}_2^*\times\mathbf{M}_2}(\{h_1(x_1(t))\}_{0:N}\in A_{-\epsilon})-(1-\delta)^{N+1}\\
&\quad\le \mathbb{P}_{\mathbf{C}_2^*\times\mathbf{M}_2}(\{h_1(x_1(t))\}_{0:N}\in A_{-\epsilon}\wedge (x_1(t),x(t))\in\mathcal{R}\text{ for } t\in[0,N]),\\
&1-\mathbb{P}_{\mathbf{C}_2^*\times\mathbf{M}_2}(\{h_1(x_1(t))\}_{0:N}\in A_{-\epsilon}\wedge (x_1(t),x(t))\in\mathcal{R}\text{ for } t\in[0,N])\\
&\le (1-\mathbb{P}_{\mathbf{C}_2^*\times\mathbf{M}_2}(\{h_1(x_1(t))\}_{0:N}\in A_{-\epsilon}))\\
&\qquad\qquad +(1-\mathbb{P}_{\mathbf{C}_2^*\times\mathbf{M}_2}(((x_1(t),x(t))\in\mathcal{R}\text{ for } t\in[0,N]))\\
&\le (1-\mathbb{P}_{\mathbf{C}_2^*\times\mathbf{M}_2}(\{h_1(x_1(t))\}_{0:N}\in A_{-\epsilon}))+(1-(1-\delta)^{N+1}).
\end{aligned}$$

We have deduced that

$$\mathbb{P}_{\mathbf{C}_1\times\mathbf{M}_1}(\{h_1(x_1(t))\}_{0:N}\in A_{-\epsilon})-(1-(1-\delta)^{N+1})\le\mathbb{P}_{\mathbf{C}_2\times\mathbf{M}_2}(\{h_2(x_2(t))\}_{0:N}\in A).$$

If $\{h_2(x_2(t))\}_{0:N}\in A$ and $(\tilde{x}(t),x(t))\in\mathcal{R}$ then $\{h_1(x_1(t))\}_{0:N}\in A_\epsilon$. Thus via similar arguments it can be deduced that

$$\mathbb{P}_{\mathbf{C}_2\times\mathbf{M}_2}(\{h_2(x_2(t))\}_{0:N}\in A)\le\mathbb{P}_{\mathbf{C}_1\times\mathbf{M}_1}(\{h_1(x_1(t))\}_{0:N}\in A_\epsilon)+(1-(1-\delta)^{N+1}).$$

∎

C.2. Proof of transitivity statements.

Proof of Theorem 5 *and Corollary* 1. Since $\mathbf{M}_1\preceq^{\delta_a}_{\epsilon_a}\mathbf{M}_2$ and $\mathbf{M}_2\preceq^{\delta_b}_{\epsilon_b}\mathbf{M}_3$ there exist

- relations $\mathcal{R}_{12}\subset\mathbb{X}_1\times\mathbb{X}_2$ and $\mathcal{R}_{23}\subset\mathbb{X}_2\times\mathbb{X}_3$ that satisfy the required conditions in Definition 9;
- interface $\mathcal{U}_{v12}:\mathbb{U}_1\times\mathbb{X}_1\times\mathbb{X}_2\to\mathcal{P}(\mathbb{U}_2,\mathcal{B}(\mathbb{U}_2))$ and $\mathcal{U}_{v23}:\mathbb{U}_2\times\mathbb{X}_2\times\mathbb{X}_3\to\mathcal{P}(\mathbb{U}_3,\mathcal{B}(\mathbb{U}_3))$;
- corresponding stochastic kernels $\mathbb{W}_{\mathbb{T}12}$ and $\mathbb{W}_{\mathbb{T}23}$.

Define the relation $\mathcal{R}_{13}\subset\mathbb{X}_1\times\mathbb{X}_3$ as

$$\mathcal{R}_{13}:=\{(x_1,x_3)\in\mathbb{X}_1\times\mathbb{X}_3\mid\exists x_2\in\mathbb{X}_2:(x_1,x_2)\in\mathcal{R}_{12},(x_2,x_3)\in\mathcal{R}_{23}\}.$$

Then for all $(x_1,x_3)\in\mathcal{R}_{13}$ there exists an $x_2\in\mathbb{X}_2:(x_1,x_2)\in\mathcal{R}_{12},(x_2,x_3)\in\mathcal{R}_{23}$. More specifically, define a Borel-measurable function $F:\mathbb{X}_1\times\mathbb{X}_3\to\mathbb{X}_2$ such that for all$(x_1,x_3)\in\mathcal{R}_{13}$ for the mapping $x_2=F(x_1,x_3)$ it holds that $(x_1,x_2)\in\mathcal{R}_{12},(x_2,x_3)\in\mathcal{R}_{23}$.

We have for all $(x_1,x_3)\in\mathcal{R}_{13}$ and $x_2=F(x_1,x_3)$

1. $\mathbf{d}\left(h_1(x_1(t)),h_3(x_3)\right)\le\mathbf{d}\left(h_1(x_1(t)),h_2(x_2(t))\right)+\mathbf{d}\left(h_2(x_2(t)),h_3(x_3)\right)\le\epsilon_a+\epsilon_b$;
2. for all $u_1\in\mathbb{U}_1:\mathbb{T}_1(\cdot|x_1,u_1)\ \bar{\mathcal{R}}_{12,\delta_a}\ \mathbb{T}_2(\cdot|x_2,\mathcal{U}_{v12}(u_1,x_1,x_2))$ and for all $u_2\in\mathbb{U}_2:\mathbb{T}_2(\cdot|x_2,u_2)\ \bar{\mathcal{R}}_{23,\delta_b}\ \mathbb{T}_3(\cdot|x_3,\mathcal{U}_{v23}(u_2,x_2,x_3))$ and $\mathbb{W}_{\mathbb{T}23}\in\mathcal{P}(\mathbb{X}_2\times\mathbb{X}_3,\mathcal{B}(\mathbb{X}_2\times\mathbb{X}_3))$ lifted with $\mathbb{W}_{\mathbb{T}12}(\cdot|u_1,x_1,x_2)$ and $\mathbb{W}_{\mathbb{T}23}(\cdot|u_2,x_2,x_3)$.

Let us derive the stochastic kernel $\mathbb{W}_{\mathbb{T}13}$ by combining $\mathbb{W}_{\mathbb{T}12}$ and $\mathbb{W}_{\mathbb{T}23}$ and marginalizing over $\mathbb{X}_2$,

$$\begin{aligned}
\mathbb{W}_{\mathbb{T}13}(dx_1'\times dx_3'|u_1,x_1,x_2,x_3)=\int_{\mathbb{X}_2}&\mathbb{W}_{\mathbb{T}23}(dx_3'\mid x_2',\mathcal{U}_v(u_1,x_1,x_2),x_2,x_3)\\
&\times\mathbb{W}_{\mathbb{T}12}(dx_1'\times dx_2'|u_1,x_1,x_2).
\end{aligned}$$

Composed with the mapping F we get a Borel-measurable stochastic kernel $\mathbb{W}_{\mathbb{T}13}(dx_1'\times dx_3'|u_1,x_1,x_3):=\mathbb{W}_{\mathbb{T}13}(dx_1'\times dx_3'|x_1,F(x_1,x_3),x_3)$. In the following, we

drop the argument of the stochastic kernel. Note that $\mathbb{T}_2(dx_2|x_2,\mu_{u,2}) = \mathbb{W}_{\mathbb{T}12}(\mathbb{X}_1 \times dx_2) = \mathbb{W}_{\mathbb{T}23}(dx_2 \times \mathbb{X}_3)$. For lifting we have to proof that $\mathbb{W}_{\mathbb{T}13}(\mathcal{R}_{13}) \geq 1-\delta_a-\delta_b$ or, equivalently, that $\mathbb{W}_{\mathbb{T}13}(\mathbb{X}_1 \times \mathbb{X}_3 \setminus \mathcal{R}_{13}) \leq \delta_a+\delta_b$, namely,

$$\begin{aligned}
&\mathbb{W}_{\mathbb{T}13}(\mathbb{X}_1 \times \mathbb{X}_3 \setminus \mathcal{R}_{13}) = \int_{\mathbb{X}_1}\int_{\mathbb{X}_2}\int_{\mathbb{X}_3\setminus\mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 \mid x_2)\mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)\\
&= \int_{\mathcal{R}_{12}}\int_{\mathbb{X}_3\setminus\mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 \mid x_2)\mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)\\
&\quad + \int_{\mathbb{X}_1}\int_{\mathbb{X}_2\setminus\mathcal{R}_{12}(x_1)}\int_{\mathbb{X}_3\setminus\mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 \mid x_2)\mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)
\end{aligned}$$

for all $(x_1,x_2) \in \mathcal{R}_{12} : \mathcal{R}_{23}(x_2) \subseteq \mathcal{R}_{13}(x_1)$

$$\begin{aligned}
&\leq \int_{\mathbb{X}_2}\int_{\mathbb{X}_3\setminus\mathcal{R}_{23}(x_2)}\int_{\mathcal{R}_{12}^{-1}(x_2)} \mathbb{W}_{\mathbb{T}12}(dx_1 \mid x_2)\mathbb{W}_{\mathbb{T}23}(dx_2 \times dx_3)\\
&\quad + \int_{\mathbb{X}_1}\int_{\mathbb{X}_2\setminus\mathcal{R}_{12}(x_1)}\int_{\mathbb{X}_3\setminus\mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 \mid x_2)\mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)\\
&\leq \int_{\mathbb{X}_2}\int_{\mathbb{X}_3\setminus\mathcal{R}_{23}(x_2)}\int_{\mathbb{X}_1} \mathbb{W}_{\mathbb{T}12}(dx_1 \mid x_2)\mathbb{W}_{\mathbb{T}23}(dx_2 \times dx_3)\\
&\quad + \int_{\mathbb{X}_1}\int_{\mathbb{X}_2\setminus\mathcal{R}_{12}(x_1)}\int_{\mathbb{X}_3} \mathbb{W}_{\mathbb{T}23}(dx_3 \mid x_2)\mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)\\
&= \int_{\mathbb{X}_2}\int_{\mathbb{X}_3\setminus\mathcal{R}_{23}(x_2)} \mathbb{W}_{\mathbb{T}23}(dx_2 \times dx_3) + \int_{\mathbb{X}_1}\int_{\mathbb{X}_2\setminus\mathcal{R}_{12}(x_1)} \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)\\
&\leq \delta_a+\delta_b.
\end{aligned}$$

In addition it has to hold that $\mathbb{W}_{\mathbb{T}13}(X_1 \times \mathbb{X}_3) = \mathbb{T}_1(\cdot|x_1,\mu_{u,1})$, namely,

$$\begin{aligned}
\mathbb{W}_{\mathbb{T}13}(X_1 \times \mathbb{X}_3) &= \int_{X_1}\int_{\mathbb{X}_3}\int_{\mathbb{X}_2} \mathbb{W}_{\mathbb{T}23}(dx_3 \mid x_2)\mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)\\
&= \int_{X_1}\int_{\mathbb{X}_2}\int_{\mathbb{X}_3} \mathbb{W}_{\mathbb{T}23}(dx_3 \mid x_2)\mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2)\\
&= \mathbb{W}_{\mathbb{T}12}(X_1 \times \mathbb{X}_2) = \mathbb{T}_1(\cdot|x_1,\mu_{u,1}).
\end{aligned}$$

The condition $\mathbb{W}_{\mathbb{T}13}(\mathbb{X}_1 \times X_3) = \mathbb{T}_3(\cdot|x_3,\mu_{u,3})$ can be proven via similar arguments. In conclusion $\mathbb{T}_1(\cdot|x_1,\mu_{u,1})\bar{\mathcal{R}}_{13,\delta_a+\delta_b}\mathbb{T}_3(\cdot|x_3,\mu_{u,3})$. To complete the proof we can show, using the same arguments as before, that if $\pi_1\bar{\mathcal{R}}_{12,\delta_a}\pi_2$ and if $\pi_2\bar{\mathcal{R}}_{23,\delta_b}\pi_3$ then $\pi_1\bar{\mathcal{R}}_{13,\delta_a+\delta_b}\pi_3$. ∎

REFERENCES

[1] A. Abate, *Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: a survey*, Electron. Notes Theor. Comput. Sci., 297 (2013), pp. 3–25.

[2] A. Abate, M. Kwiatkowska, G. Norman, and D. Parker, *Probabilistic model checking of labelled Markov processes via finite approximate bisimulations*, in Horizons of the Mind – P. Panangaden Festschrift, Springer, Cham, Switzerland,, 2014, pp. 40–58.

[3] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, *Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems*, Automatica J. IFAC, 44 (2008), pp. 2724–2734.

[4] A. Abate, F. Redig, and I. Tkachev, *On the effect of perturbation of conditional probabilities in total variation*, Statist. Probab. Lett., 88 (2014), pp. 1–80.

[5] R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi, *Alternating refinement relations*, in CONCUR '98, Springer, London, 1998, pp. 163–178.

[6] P. Bacher and H. Madsen, *Identifying suitable models for the heat dynamics of buildings*, Energy Build., 43 (2011), pp. 1511–1522.
[7] D. Bertsekas and S. E. Shreve, *Stochastic Optimal Control: The Discrete Time Case*, Athena Scientific, Belmont, MA, 1996.
[8] R. Blute, J. Desharnais, A. Edalat, and P. Panangaden, *Bisimulation for labelled Markov processes*, in Proceeding of the 12th Annual IEEE Symposium on Logic in Computer Science, 1997 IEEE, IEEE Computer Society, Los Alamitos, CA, 1997, pp. 149–158.
[9] V. I. Bogachev, *Measure Theory*, Springer, Berlin, 2007.
[10] V. S. Borkar, *Probability Theory: An Advanced Course*, Springer, New York, 1995.
[11] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, 2004.
[12] M. L. Bujorianu, J. Lygeros, and M. C. Bujorianu, *Bisimulation for general stochastic hybrid systems*, in International Workshop on Hybrid Systems: Computation and Control, Springer, Berlin, 2005, pp. 198–214.
[13] J. Desharnais, A. Edalat, and P. Panangaden, *A logical characterization of bisimulation for labeled Markov processes*, in Proceedings of the Thirteenth Annual IEEE Symposium on Logic in Computer Science, 1998, IEEE, IEEE Computer Society, Los Alamitos, CA, 1998, pp. 478–487.
[14] J. Desharnais, A. Edalat, and P. Panangaden, *Bisimulation for labelled Markov processes*, Inform. and Comput., 179 (2002), pp. 163–193.
[15] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden, *Approximating labelled Markov processes*, Inform. and Comput., 184 (2003), pp. 160–200.
[16] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden, *Metrics for labelled Markov processes*, Theoret. Comput. Sci., 318 (2004), pp. 323–354.
[17] J. Desharnais, F. Laviolette, and M. Tracol, *Approximate analysis of probabilistic processes: Logic, simulation and games*, Conference on Quantitative Evaluation of Systems, IEEE Computer Society, Los Alamitos, CA, 2008, pp. 264–273.
[18] A. D'Innocenzo, A. Abate, and J.-P. Katoen, *Robust PCTL model checking*, in Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control, ACM, New York, 2012, pp. 275–285.
[19] A. Edalat, *Semi-pullbacks and bisimulation in categories of Markov processes*, Math. Structures Comput. Sci., 9 (1999), pp. 523–543.
[20] S. Esmaeil Zadeh Soudjani and A. Abate, *Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes*, SIAM J. Appl. Dyn. Syst., 12 (2013), pp. 921–956.
[21] S. Esmaeil Zadeh Soudjani and A. Abate, *Quantitative approximation of the probability distribution of a Markov process by formal abstractions*, Log. Methods Comput. Sci., 11 (2015), 8.
[22] S. Esmaeil Zadeh Soudjani, C. Gevaerts, and A. Abate, *FAUST ^{2}: Formal Abstractions of Uncountable-STate STochastic Processes*, in Tools and Algorithms for the Construction and Analysis of Systems (TACAS), Lecture Notes in Comput. Sci. 9035, Springer, Berlin, 2015, pp. 272–286.
[23] A. Girard and G. J. Pappas, *Hierarchical control system design using approximate simulation*, Automatica J. IFAC, 45 (2009), pp. 566–571.
[24] M. Grant and S. Boyd, *CVX: Matlab Software for Disciplined Convex Programming, version* 2.1, http://cvxr.com/cvx (2014).
[25] S. Haesaert, A. Abate, and P. M. J. Van den Hof, *Correct-by-design output feedback of LTI systems*, in Proceedings of Conference on Decision and Control, IEEE, Piscataway, NJ, 2015, pp. 6159–6164.
[26] S. Haesaert, A. Abate, and P. M. J. Van den Hof, *Verification of general Markov decision processes by approximate similarity relations and policy refinement*, in 13th International Conference on Quantitative Evaluation of Systems, QEST 2016, Quebec City, Canada, Springer, 2016, pp. 227–243.
[27] O. Hernández-Lerma and J. B. Lasserre, *Discrete-Time Markov Control Processes*, Appl. Math. 30, Springer, New York, 1996.
[28] O. Holub and K. Macek, *HVAC simulation model for advanced diagnostics*, in Symposium on Intelligent Signal Processing, IEEE, Piscataway, NJ, 2013, pp. 93–96.
[29] B. Jonsson and K. G. Larsen, *Specification and refinement of probabilistic processes*, in Proceedings of the Sixth Annual IEEE Symposium on Logic in Computer Science, IEEE Computer Society, Los Alamitos, CA, 1991, pp. 266–277.
[30] A. A. Julius and G. J. Pappas, *Approximations of stochastic hybrid systems*, IEEE Trans. Automat. Control, 54 (2009), pp 1193–1203.

[31] K. G. Larsen and A. Skou, *Bisimulation through probabilistic testing*, Inform. and Comput., 94 (1991), pp. 1–28.
[32] T. Lindvall, *Lectures on the coupling method*, Dover, Mineola, NY, 2002.
[33] M. Mazo Jr, A. Davitian, and P. Tabuada, *PESSOA: Towards the automatic synthesis of correct-by-design control software*, presented at Work-in-Progress HSCC, 2010.
[34] S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*, Comm. Control Engrg., Springer, London, 1993.
[35] G. Pola, C. Manes, A. J. van der Schaft, and M. D. Di Benedetto, *On equivalence notions for discrete-time stochastic control systems*, in 2015 54th IEEE Conference on Decision and Control, IEEE, Piscataway, NJ, 2015, pp. 1180–1185.
[36] R. Segala, *Modeling and Verification of Randomized Distributed Real-Time Systems*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, 1995.
[37] R. Segala and N. Lynch, *Probabilistic simulations for probabilistic processes*, Nordic J. Comput., 2 (1995), pp. 250–273.
[38] H. J. Skala, *The existence of probability measures with given marginals*, Ann. Probab., 21 (1993), pp. 136–142.
[39] V. Strassen, *The existence of probability measures with given marginals*, Ann. Math. Statist., 36 (1965), pp. 423–439.
[40] S. Strubbe and A. Van Der Schaft, *Bisimulation for communicating piecewise deterministic Markov processes (CPDPS)*, in International Workshop on Hybrid Systems: Computation and Control, Springer, Berlin, 2005, pp. 623–639.
[41] S. Strubbe and A. van der Schaft, *Communicating piecewise deterministic Markov processes*, in Stochastic Hybrid Systems, Springer, Berlin, 2006, pp. 65–104.
[42] P. Tabuada, *Verification and Control of Hybrid Systems: A Symbolic Approach*, Springer, New York, 2009.
[43] C. W. Therrien, *Discrete Random Signals and Statistical Signal Processing*, Prentice Hall, Englewood Cliffs, NJ, 1992.
[44] M. Zamani, P. M. Esfahani, R. Majumdar, A. Abate, and J. Lygeros, *Symbolic control of stochastic systems via approximately bisimilar finite abstractions*, IEEE Trans. Automat. Control, 59 (2014), pp. 3135–3150.
[45] C. Zhang and J. Pang, *On probabilistic alternating simulations*, in Theoretical Computer Science, C. Calude and V. Sassone, eds., IFIP Adv. Inf. Commun. Technol. 323, Springer, Berlin, 2010, pp. 71–85.