2015 IEEE 54th Annual Conference on Decision and Control (CDC)
December 15-18, 2015. Osaka, Japan

Correct-by-design output feedback of LTI systems

S. Haesaert, A. Abate and P.M.J. Van den Hof

Abstract— **Current state-of-the-art correct-by-design controllers are designed for full-state measurable systems. This work extends the applicability of correct-by-design controllers to partially observable linear, time-invariant (LTI) models. Towards the certification of the synthesised controllers, approximate simulation relations are leveraged to attain a quantification for the accuracy of introduced approximations. Additionally, the robustness of the approach allows the extension to models with the presence of probabilistic disturbances on state transitions and on output measurements. In a case study from smart buildings we evaluate the new output-based correct-by-design controller on a physical system with limited sensor information.**

Index Terms— **partially observable LTI systems, stochastic systems, correct-by-design control, approximate simulations**

I. INTRODUCTION

Reliable and autonomous operation of many complex engineering systems demands guaranteed behaviour over the full spectrum of their operating conditions. This is the case for applications in avionics, automotive, transportation systems, dependable electronics, semiconductors [1], and in general for systems where safety is critical and where mistakes lead to impactful economical losses.

Within the computer sciences, verification and synthesis of critical hardware and software has been translated to industrial practice by tools and techniques from the domain of formal methods [2]. Employing well-structured specifications, such as properties expressed over linear-time temporal logics (LTL), automated and computer-aided tools have been developed for the verification and synthesis of models of the industrial systems of interest. To meet new demands from domains dealing with complex new applications, these methods require to be extended to hold over models of physical systems. Recent research [3], [4], [5] pursues the verification of models of physical systems with continuous state spaces: of special interest is the safe-by-construction automatic synthesis of controllers. These correct-by-design controllers are however incompatible with systems for which models with exact knowledge of the dynamics and full-state measurements are not available.

Contributions: This work newly extends correct-by-design controllers for linear time invariant (LTI) models [3] to controllers that employ sensor outputs or partial state measurements. Towards the certification of the synthesised controllers, as in [3] our new control architectures come with quantitative certificates on approximations accuracy. Further, since the dynamics of physical systems are often disturbed (in a probabilistic sense) and the associated sensors are noisy, we require the new output-based controllers to show quantifiable robustness with respect to stochastic disturbances on state evolutions and output measurements.

Related work: Design methods for classical optimal control problems [6] of models with (noisy) output measurements can be distinguished in direct designs based on the input-output behaviour of the system, and in methods exploiting the separation of estimation and control. The former class includes frequency-domain and robust control methods. Alternatively, whenever applicable (as in the optimal linear quadratic Gaussian problem) the separation theorem [7] allows for the distinct design of an observer, estimating the state, and of a state feedback controller, finally yielding a combined output-feedback controller.

Within the formal methods literature, limited efforts have targeted the synthesis of controllers over models without full-state observations. Existing results exclusively target finite-state models. [8] studies the synthesis for partially observable models by searching the space of output-feedback controllers via counterexample-guided refinements. A heuristic algorithm in [9] finds controllers satisfying almost surely LTL properties over partially-observable Markov decision processes. The work of [10] extends PCTL* model checking to hidden Markov models.

For fully observable Markov processes with general state spaces, verification and controller synthesis problems are reviewed in [11], and are generally investigated over simplified (abstract) models that can be formally related to the given ones. Specifically, an abstract model is shown to be in an (approximate) relation with the original one, and the approximation is quantified either via metrics defined over the conditional kernels [12], or via metrics bounding the distance between the output trajectories [13]. In contrast, this work will use the definition of approximate bisimulation relations, similar to those in [14], to quantify the *expected* deviation of noisy trajectories affected by stochastic disturbances.

Structure of the article: After a review of state-based, correct-by-design controller architectures in Section II, in Section III we design a certified, output-based controller, introducing a state observer and the notion of output-based interface. Section IV discusses robustness to stochastic disturbances, both on state transitions and on sensor measurements. Section V evaluates the design methodology on a case study in the area of Smart Buildings. The proofs of the statements can be found in [15].

S. Haesaert and P.M.J. Van den Hof are with the Control Systems group in the Department of Electrical Engineering, Eindhoven University of Technology, The Netherlands. A. Abate is with the Department of Computer Science, University of Oxford, UK. This work is supported by the Netherlands Organisation for Scientific Research (NWO) and the Dutch Institute of Systems and Control (DISC), by the European Commission IAPP project AMBI 324432, and by the John Fell OUP Research Fund.

978-1-4799-7885-4/15/$31.00 ©2015 IEEE

II. PRELIMINARIES AND PROBLEM STATEMENT

We intend to synthesise a certifiable output-based controller for a physical system represented by the LTI model

$$\mathbf{M}:\begin{cases} x(t+1) &= Ax(t)+Bu(t)\\ y(t) &= Cx(t)\\ z(t) &= Hx(t),\end{cases} \tag{1}$$

where $x(t)\in\mathbb{R}^n$ is the state, initialised by $x(0)\in\mathbb{X}_0\subset\mathbb{R}^n$, the control input is $u(t)\in\mathbb{R}^m$, and $y(t)\in\mathbb{R}^p$ is the measured output available for control. A,B,C are real matrices of appropriate dimensions. The signals $z(t)\in\mathbb{R}^q$, mapped from the state space via the linear map Hx, are used to define performance and properties. This in unlike [10], which defines specifications over the signals $y(t)$. In contrast to the measured output $y(t)$, the structure of which is physically specified by the sensors attached to the system, the choice of H can be adapted to the design requirements, and include $H=C$ and $H=I$ as special cases.

A. Transition systems and simulation relations

Definition 1 (Transition system [4]): A transition system is a tuple $\Sigma=(\mathscr{X},\mathscr{X}_0,\mathscr{A},\rightarrow,\mathscr{Z},\mathscr{H})$, where

- $\mathscr{X}$ is a (possibly infinite) set of states;
- $\mathscr{X}_0$ is a (possibly infinite) set of initial states;
- $\mathscr{A}$ is a (possibly infinite) set of actions;
- $\rightarrow\subseteq\mathscr{X}\times\mathscr{A}\times\mathscr{X}$ is a transition relation;
- $\mathscr{Z}$ is a (possible infinite) set of observations;
- $\mathscr{H}:\mathscr{X}\rightarrow\mathscr{Z}$ is a map assigning to each $x\in\mathscr{X}$ an observation $\mathscr{H}(x)\in\mathscr{Z}$.

A metric transition system is a transition system endowed with a metric over the observation space $\mathscr{Z}$. □

This work considers non-blocking transition systems, where every state $x\in\mathscr{X}$ is associated to a non empty transition relation. The behaviour generated by Σ is denoted as $\mathscr{B}(\Sigma)$ and consists of all infinite sequences $z_0,z_1,z_2,...$ for which there exists an initialised path $(x_0,u_0),(x_1,u_1),(x_2,u_2),\ldots$, with $x_0\in\mathscr{X}_0$, $(x_i,u_i,x_{i+1})\in\rightarrow$, and $z_i=\mathscr{H}(x_i)$ for all $i\in\mathbb{N}$.

The LTI model $\mathbf{M}$ can be reinterpreted as a transition system characterised by a tuple $(\mathbb{R}^n,\mathbb{X}_0,\mathbb{R}^m,\rightarrow,\mathbb{R}^q,H)$, with a state space $x\in\mathbb{R}^n$, a set of initial states $x(0)\in\mathbb{X}_0$, and transitions $\rightarrow:=\{x,u,x'|x'=Ax+Bu\}$. Additionally, H assigns observation $z\in\mathbb{R}^q$ to $x\in\mathbb{R}^n$: $z=Hx$. Note that this transition system has uniquely defined transitions, since for every state-action pair there is a unique state transition. If, in addition, the initial state is defined deterministically $\mathscr{X}_0:=\{x_0\}$ then the transition system is called *deterministic*.

The verification of LTI models can be attained by abstracting them as finite-state ones and leveraging symbolic approaches [4]. Pairs of models can be related as follows.

Definition 2 (Simulation relation [4]): Let $\Sigma_a=(\mathscr{X}_a,\mathscr{X}_{a0},\mathscr{A}_a,\rightarrow_a,\mathscr{Z}_a,\mathscr{H}_a)$ and $\Sigma_b=(\mathscr{X}_b,\mathscr{X}_{b0},\mathscr{A}_b,\rightarrow_b,\mathscr{Z}_b,\mathscr{H}_b)$ be transition systems with the same output sets $\mathscr{Z}_a=\mathscr{Z}_b$. A binary relation $\mathscr{R}\subset\mathscr{X}_a\times\mathscr{X}_b$ is said to be a simulation relation from Σ_a to Σ_b if the following three conditions are satisfied: 1) for every $x_{a0}\in\mathscr{X}_a$, there exists $x_{b0}\in\mathscr{X}_b$ with $(x_{a0},x_{b0})\in\mathscr{R}$; 2) for every $(x_a,x_b)\in\mathscr{R}$ we have $\mathscr{H}_a(x_a)=\mathscr{H}_b(x_b)$; 3) for every $(x_a,x_b)\in\mathscr{R}$ we have that $x_a\xrightarrow{u_a}_a x_a{}'$ in Σ_a implies the existence of $x_b\xrightarrow{u_b}_b x_b{}'$ in Σ_b satisfying $(x_a{}',x_b{}')\in\mathscr{R}$. We say that Σ_a is simulated by Σ_b, denoted as $\Sigma_a\preceq_{\mathscr{S}}\Sigma_b$, if there exists a simulation relation from Σ_a to Σ_b. The models Σ_a and Σ_b are bisimilar, i.e., $\Sigma_a\sim_{\mathscr{B}}\Sigma_b$, if there exists relation $\mathscr{R}$ that is a simulation relation from Σ_a to Σ_b and for which $\mathscr{R}^{-1}$ is also a simulation relation from Σ_b to Σ_a. □

Note that this similarity relation over the set of transition system implies a relation over the behaviour of the transition systems [4], more precisely if $\Sigma_a\preceq_{\mathscr{S}}\Sigma_b$ then $\mathscr{B}(\Sigma_a)\subseteq\mathscr{B}(\Sigma_b)$, and if $\Sigma_a\sim_{\mathscr{B}}\Sigma_b$ then $\mathscr{B}(\Sigma_a)=\mathscr{B}(\Sigma_b)$.

B. Formal specifications and control design

Let us consider a specification of interest ψ for which the desired behaviour is represented by a transition system Σ_ψ [4]. Then a control synthesis problem for Σ can be formulated as the search of a controller $\mathbf{C}$ such that the controlled transition system, i.e., $\mathbf{C}\times\Sigma$ *satisfies* the specification, namely (a.) if $\mathbf{C}\times\Sigma\preceq_{\mathscr{S}}\Sigma_\psi$ or (b.) if $\mathbf{C}\times\Sigma\sim_{\mathscr{B}}\Sigma_\psi$. The notation $\mathbf{C}\times\Sigma$ refers to the composition of the controller $\mathbf{C}$ with model Σ: the actions of the obtained transition system are defined by the controller $\mathbf{C}$, whereas the internal state of $\mathbf{C}$ is updated based on information available from Σ.

If Σ_a and Σ_b are deterministic transition systems and $\Sigma_a\preceq_{\mathscr{S}}\Sigma_b$, then for every sequence of actions for Σ_a, there exists a corresponding sequence for Σ_b such that the observed behaviour is the same [16]. Definition 2 suggests the refinement of a controller for Σ_a to Σ_b via condition 3): for ever choice of u_a, picked by the controller for Σ_a, there exists a suitable input u_b. In practice this allows synthesis problems to be first solved on a simplified, and possibly finite, abstraction (Σ_a), before refinement over a concrete, complex model (Σ_b).

Towards robust notions of satisfaction, approximate versions of simulation relations [4] can be considered over metric transition systems. Consider two given metric transition systems with a shared output space $\mathscr{Z}$ and a metric $\mathbf{d}$. The relation $\mathscr{R}\subset\mathscr{X}_a\times\mathscr{X}_b$ is said to be an ε-approximate simulation relation from $\mathscr{X}_a$ to $\mathscr{X}_b$ iff conditions 1) and 3) of Definition 2 hold, and additionally if for every $(x_a,x_b)\in\mathscr{R}$ we have $\mathbf{d}(\mathscr{H}_a(x_a)-\mathscr{H}_b(x_b))\leq\varepsilon$. We say that Σ_a is *approximately simulated* by Σ_b, denoted by $\Sigma_a\preceq^{\varepsilon}_{\mathscr{S}}\Sigma_b$, if there exists an ε-approximate simulation relation from Σ_a to Σ_b. The models Σ_a and Σ_b are approximately *bisimilar*, i.e., $\Sigma_a\sim^{\varepsilon}_{\mathscr{B}}\Sigma_b$, iff there exists a relation $\mathscr{R}$ that is an ε-approximate simulation relation from Σ_a to Σ_b and for which $\mathscr{R}^{-1}$ is an ε-approximate simulation relation from Σ_b to Σ_a.

C. State-of-the-art correct-by-design controller synthesis

Suppose that an LTI model $x(t+1)=Ax(t)+Bu(t)$ is given, and that it has a finite-valued observation map that induces a partition over the observation space $\mathbb{R}^q$. Under assumptions on the controllability of the model, on the linear independence of the columns of its input matrix B, and on the observation map [3], [4], the LTI model can be bisimulated by a finite transition system. Alternatively, under less stringent conditions it is possible to synthesise a finite approximate bisimulation of the given model [4], [5]: further, for every controller synthesised on the finite-state abstraction there exists a refined controller for the original model, with the same closed-loop behaviour.

In the remainder of this work we assume that given a model $\mathbf{M}$ and a model Σ_ψ for the specification, both with the same output space, we have obtained a controlled model $\bar{\mathbf{M}}_\mathbf{C}$, which is such that $\bar{\mathbf{M}}_\mathbf{C} \sim_\mathscr{B} \Sigma_\psi$ [4]. $\bar{\mathbf{M}}_\mathbf{C} := \mathbf{C} \times \mathbf{M}$ denotes the composition of model $\mathbf{M}$ with the correct-by-design controller $\mathbf{C}$, where $\mathbf{C}$ takes as input the state of $\mathbf{M}$ and returns an action to $\mathbf{M}$. This controlled model has hybrid states $(\bar{x}, q)$ with $\bar{x} \in \mathbb{R}^n$ and $q \in Q$, where Q is a finite set. Its dynamics are defined as

$$\bar{\mathbf{M}}_\mathbf{C} : \begin{cases} \bar{x}(t+1) &= A\bar{x}(t) + B\bar{u}_q(\bar{x}(t)) \\ q(t+1) &= \delta(\bar{x}(t), q(t)), \end{cases} \tag{2}$$

and initialised by $(\bar{x}(0), q(0)) \in \bigcup_{q_0 \in Q_0} (\{q_0\} \times \mathbb{X}_0(q))$. Let us remark that the discrete states of this model follow from the states of a finite transition system, approximately bisimilar to the continuous-state model $\mathbf{M}$, and from the discrete states of the specification model Σ_ψ. Hence the discrete state q is initialised based on the specification model Σ_ψ and the initial state $\bar{x}(0)$. Note that $\bar{u}_q(\bar{x}(t))$ is a function that maps the current state to an action.

D. Problem statement

Suppose that there exists a state-based, correct-by-design controller for a fully-observed LTI model, with closed-loop dynamics denoted by $\bar{\mathbf{M}}_\mathbf{C}$ as in (2). The objective of this work is to design an output-based controller, a controller that only requires the measured signal $y(t)$ and that can therefore be deployed on the model in (1). Additionally, it is required that the new controller guarantees an upper-bound on the deviation from the state-based control in (2).

In the following we use the notion of *interface function*. Interface functions originate from the work in [16] on hierarchical control design based on (approximate) simulation relations: the construction of a controller over a simplified model is *refined* to a concrete model while maintaining the same guarantees over the controlled behaviour.

Definition 3 (Interface function): Let $\Sigma_a = (\mathscr{X}_a, \mathscr{X}_{a0}, \mathscr{A}_a, \rightarrow_a, \mathscr{Z}_a, \mathscr{H}_a)$ and $\Sigma_b = (\mathscr{X}_b, \mathscr{X}_{b0}, \mathscr{A}_b, \rightarrow_b, \mathscr{Z}_b, \mathscr{H}_b)$ be deterministic transition systems with the same output sets $\mathscr{Z}_a = \mathscr{Z}_b$. A relation $\mathscr{R} \subset \mathscr{X}_a \times \mathscr{X}_b$ is an ε-approximate simulation relation from $\mathscr{X}_a$ to $\mathscr{X}_b$, and $\mathscr{F} : \mathscr{A}_a \times \mathscr{X}_a \times \mathscr{X}_b \rightarrow \mathscr{A}_b$ is its related interface, if the following three conditions are satisfied: 1) for every $x_{a0} \in \mathscr{X}_{a0}$, there exists $x_{b0} \in \mathscr{X}_{b0}$ with $(x_{a0}, x_{b0}) \in \mathscr{R}$; 2) for every $(x_a, x_b) \in \mathscr{R}$, $\mathbf{d}(\mathscr{H}_a(x_a) - \mathscr{H}_b(x_b)) \leq \varepsilon$; 3) for every $(x_a, x_b) \in \mathscr{R}$ we have that $x_a \xrightarrow{u_a}_a x_a'$ in Σ_a implies $x_b \xrightarrow{u_b}_b x_b'$ in Σ_b with $u_b = \mathscr{F}(u_a, x_a, x_b)$, satisfying $(x_a', x_b') \in \mathscr{R}$. The feedback composition of Σ_a and Σ_b is denoted as $\Sigma_a \times_\mathscr{F} \Sigma_b$. □

Note that the existence of an (approximate) simulation relation implies the existence of an interface, i.e., for all ε-approximately simulated and deterministic transition systems there exists at least one interface function.

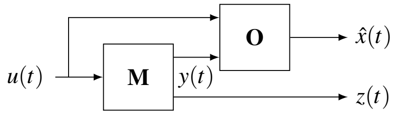

Fig. 1: Interconnection model/observer, $\mathbf{M}\|\mathbf{O}(\mathbf{M})$

In practice Definition 3 entails that the dynamics corresponding to the feedback-composed models $\Sigma_a \times_\mathscr{F} \Sigma_b$ do not differ more than ε. Hence, a controller composed on Σ_a can be refined to Σ_b via the interface $\mathscr{F}$, without affecting its closed-loop accuracy more than ε.

Let us define a specific class of interfaces denoted as *sensor-based interfaces*, which are defined exclusively based on sensor information from Σ_b, namely $\mathscr{F}_g : \mathscr{A}_a \times \mathscr{X}_a \times g(\mathscr{X}_b) \rightarrow \mathscr{A}_b$, where g is the sensor function. In the particular instance of (1), the sensor function is $g(x(t)) := Cx(t)$. These structures are of interest to us, as they define the set of interfaces that can be practically implemented for controller refinement on partially observable systems.

III. OBSERVER-BASED CORRECT-BY-DESIGN CONTROLLER SYNTHESIS

In this section we propose a new design methodology for output-based controller refinement. We first design an observer that extends the sensors output with state estimates, see Fig. 1. Then as in Fig. 2 we define a linear, sensor-based interface function between $\bar{\mathbf{M}}_\mathbf{C}$ (the state-based, correct-by-design controlled model) and the model/observer interconnection from Fig. 1.

A. Observer-based design

Consider a Luenberger observer denoted as $\mathbf{O}$:

$$\begin{aligned} \hat{x}(t+1) &= A\hat{x}(t) + Bu(t) + L\left(y(t) - \hat{y}(t)\right), \\ \hat{y}(t) &= C\hat{x}(t), \end{aligned} \tag{3}$$

with gain matrix L such that $A - LC$ is stable if (A, C) is detectable [6]. The observer is initialised as $\hat{x}(0)$, and uses the outputs from $\mathbf{M}$ to estimate its internal state. The composition of $\mathbf{M}$ with its observer $\mathbf{O}(\mathbf{M})$ is denoted as $\mathbf{M}\|\mathbf{O}(\mathbf{M})$ and portrayed in Fig. 1.

Denote the sensor-based interface as

$$\mathscr{F}_g(\bar{u}, \bar{x}, \hat{x}) = \bar{u} + K(\bar{x} - \hat{x}), \tag{4}$$

where $\bar{u}$ is the action selected by $\bar{\mathbf{M}}_\mathbf{C}$ (this role is played by $\bar{u}_q$ in (2)). For this linear interface we demand that matrix $A - BK$ is stable. Note that the interface is sensor-based (as defined in Section II), since the state estimate $\hat{x}$ of x can be obtained from the sensor function of $\mathbf{M}\|\mathbf{O}(\mathbf{M})$, thus $g(x, \hat{x}) = \hat{x}$.

The overall controlled model $\bar{\mathbf{M}}_\mathbf{C} \times_{\mathscr{F}_g} (\mathbf{M}\|\mathbf{O}(\mathbf{M}))$, denoted as $\mathbf{M}_\mathbf{C}$, is the result of interfacing the two structures discussed above, as depicted in Fig. 2. This has dynamics evolving over the continuous state space $\mathbb{R}^{3n}$ as:

$$\begin{aligned} \bar{x}(t+1) &= A\bar{x}(t) + B\bar{u}_q(\bar{x}(t)) \\ \hat{x}(t+1) &= (A-LC)\hat{x}(t) + Bu(t) + LCx(t) \\ x(t+1) &= Ax(t) + Bu(t) \\ u(t) &= \mathscr{F}_g(\bar{u}_q(\bar{x}(t)), \bar{x}(t), \hat{x}(t)) \end{aligned} \tag{5}$$

in combination with the discrete transitions $q(t+1) = \delta(\bar{x}(t), q(t))$ from (2).

Remark 4: As depicted in Fig. 2, we have designed an output-based controller by combining a given state-based controller with an observer. However, unlike classical results where a state-based controller is employed over estimated states from an observer, in this work we have interfaced the state-based controlled model $\bar{\mathbf{M}}_\mathbf{C}$ with the model/observer interconnection $\mathbf{M}\|\mathbf{O}(\mathbf{M})$, as in Fig. 1. This allows one to reason explicitly about the accuracy of the overall output-controlled system, based on the accuracy of the sensor-based interface function. In special cases the proposed architecture can reduce to the classical approach. □

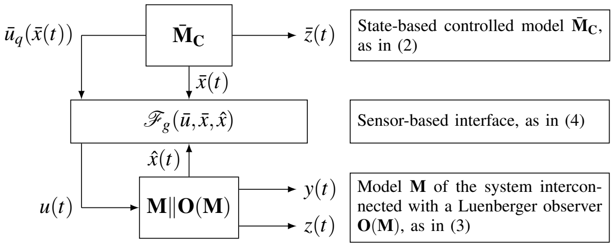
$\bar{u}_q(\bar{x}(t))$
$\bar{\mathbf{M}}_\mathbf{C}$
$\bar{z}(t)$
State-based controlled model $\bar{\mathbf{M}}_\mathbf{C}$, as in (2)
$\bar{x}(t)$
$\mathscr{F}_g(\bar{u}, \bar{x}, \hat{x})$
Sensor-based interface, as in (4)
$\hat{x}(t)$
$u(t)$
$\mathbf{M}\|\mathbf{O}(\mathbf{M})$
$y(t)$
$z(t)$
Model $\mathbf{M}$ of the system interconnected with a Luenberger observer $\mathbf{O}(\mathbf{M})$, as in (3)

Fig. 2: Observer-based correct-by-design controller synthesis. The overall interconnection is denoted as $\mathbf{M}_\mathbf{C}$.

B. Quantification of the overall accuracy

The controlled model $\bar{\mathbf{M}}_\mathbf{C}$, with traces $\bar{x}(t)$ as in (2)-(5), maps to the specification space as $\bar{z}(t) = H\bar{x}(t)$. Let a metric over this space $\mathbb{R}^q$ be defined as $\|\cdot\|_2$. Of interest is the distance between the system output $z(t)$ as in (1) and $\bar{z}(t)$, when the system is controlled via the interconnection of Fig. 2. From the definition of the sensor-based interface, the following result holds.

Theorem 5: The function in (4) is a sensor-based interface between $\bar{\mathbf{M}}_\mathbf{C}$ and $\mathbf{M}\|\mathbf{O}(\mathbf{M})$ with precision ε, where

$$\varepsilon := \sqrt{\operatorname{trace}\left([H\ H]Q[H\ H]^T\right)} \tag{6}$$

$$\text{with } \begin{bmatrix} \hat{x}(0)-\bar{x}(0) \\ x(0)-\hat{x}(0) \end{bmatrix} \begin{bmatrix} \hat{x}(0)-\bar{x}(0) \\ x(0)-\hat{x}(0) \end{bmatrix}^T - Q \preceq 0 \tag{7}$$

$$\begin{bmatrix} A-BK & LC \\ 0 & A-LC \end{bmatrix} Q \begin{bmatrix} A-BK & LC \\ 0 & A-LC \end{bmatrix}^T - Q \preceq 0. \tag{8}$$

□

Thus the distance between $\bar{z}(t)$ and $z(t)$ is bounded by ε if there exists a Q for which (7) and (8) are satisfied. A stability assumption on matrices $A-BK$ and $A-LC$ guarantees this [6]. Note that since both $\bar{x}(0)$ and $\hat{x}(0)$ are included in the design space, it would not make much sense to select $\bar{x}(0) \neq \hat{x}(0)$ for the initialisation. Hence, the accuracy depends on the initial states of the models only via $x(0) - \hat{x}(0)$. In case the initial state $x(0)$ is only known up to a set $\mathbb{X}_0$, the guarantee in Theorem 5 is required to hold over all $x(0) \in \mathbb{X}_0$.

IV. STOCHASTIC DISTURBANCES: ROBUSTNESS

We extend the previous results supposing that the physical system $\mathbf{M}$ is disturbed by stochastic noise. More precisely, state transitions are affected by additive noise $\mathbf{w}_1(t)$ with realisations $w_1(t) \sim \mathbf{w}_1(t)$ taking values in $\mathbb{R}^{d_1}$, whereas sensor measurements are disturbed by noise sources $\mathbf{w}_2(t)$, with realisations $w_2(t) \sim \mathbf{w}_2(t)$ in $\mathbb{R}^{d_2}$. (We denote random variables $\mathbf{x}$ as bold faced, in contrast to their realisations $x \sim \mathbf{x}$.) Each of the noise sources is supposed to be independent and identically distributed over time, with zero mean and unit variance. This assumption holds for a typical Gaussian process noise with distribution $w_1(t) \sim \mathscr{N}(0, I_{d_1 \times d_1})$. The resulting stochastic model is

$$\mathbf{M}: \begin{cases} x(t+1) &= Ax(t) + Bu(t) + Fw_1(t) \\ y(t) &= Cx(t) + Ew_2(t) \\ z(t) &= Hx(t), \end{cases} \tag{9}$$

where the matrices F, E, are again real-valued matrices of appropriate dimensions. The model is initialised as $\mathbf{x}(0) \sim \mathscr{N}(x_0, P_0)$.

With reference to the previous section, the control design strategy is as follows:

A. Let $\bar{\mathbf{M}}$ be a noiseless version of $\mathbf{M}$ in (9), and $\bar{\mathbf{M}}_\mathbf{C}$ be the composition of $\bar{\mathbf{M}}$ with its correct-by-design controller;
B. Design a state observer $\mathbf{O}(\mathbf{M})$ for $\mathbf{M}$;
C. Design a linear interface function $\mathscr{F}_g$ stabilising $A-BK$;
D. Implement the control structure in Fig. 2, and denote the resulting controlled stochastic model as $\mathbf{M}_\mathbf{C} := \bar{\mathbf{M}}_\mathbf{C} \times_{\mathscr{F}_g} (\mathbf{M}\|\mathbf{O}(\mathbf{M}))$.

The initial conditions for $\mathbf{M}_\mathbf{C}$, namely $\bar{x}(0), \hat{x}(0)$, are selected as part of the control design problem: as discussed earlier, we pick $\bar{x}(0) = \hat{x}(0)$. Further, let $q(0)$ be any discrete state such that $(\bar{x}(0), q(0)) \in \bigcup_{q_0 \in Q_0} (\{q_0\} \times \mathbb{X}_0(q))$.

In order to analyse the behaviour of the controlled stochastic model $\mathbf{M}_\mathbf{C}$ with respect to a metric of interest, let us embed $\mathbf{M}_\mathbf{C}$ into the formalism of deterministic transition systems (cf. Definition 1) as in [14]. The model can be represented as a symbolic transition system $\Sigma^*(\mathbf{M}_\mathbf{C})$, with states encompassing random variables $\mathbf{x}_\mathbf{C}(t)$ representing the distribution of $x_\mathbf{C}(t) \sim \mathbf{x}_\mathbf{C}(t)$, with $x_\mathbf{C}(t) \in \mathbb{R}^{3n}$ as in (5). Consider the metric output space $\mathscr{Z}$, to which the states are mapped as $\mathbf{z}_\mathbf{C}(t) = H\mathbf{x}_\mathbf{C}(t)$. Further consider the metric $\mathbf{d}^*(\mathbf{z}_1 - \mathbf{z}_2) = \mathbb{E}(\|\mathbf{z}_1 - \mathbf{z}_2\|_2)$, with $\|\cdot\|_2$ the Euclidean norm. Denote the class of all transition systems with the metric output space $\mathscr{Z}$ as $\mathscr{T}^*$.

Both the specification model Σ_ψ and the correct-by-design controlled model $\bar{\mathbf{M}}_\mathbf{C}$ can be trivially embedded in $\mathscr{T}^*$ via singleton distributions: we denote the corresponding symbolic transition systems as Σ^*_ψ and $\Sigma^*(\bar{\mathbf{M}}_\mathbf{C})$, respectively. We obtain:

Theorem 6: Transition system $\Sigma^*(\mathbf{M}_\mathbf{C})$ is approximately bisimulated by $\Sigma^*(\bar{\mathbf{M}}_\mathbf{C})$ with precision ε obtained as

$$\varepsilon := \sqrt{\operatorname{trace}([H\ H]Q[H\ H]^T)} \tag{10}$$

$$\text{with } \begin{bmatrix} 0 & 0 \\ 0 & (x_0-\hat{x}(0))(x_0-\hat{x}(0))^T \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & P_0 \end{bmatrix} - Q \preceq 0 \tag{11}$$

$$\begin{bmatrix} A-BK & LC \\ 0 & A-LC \end{bmatrix} Q \begin{bmatrix} A-BK & LC \\ 0 & A-LC \end{bmatrix}^T + \begin{bmatrix} LEE^T L^T & -LEE^T L^T \\ -LEE^T L^T & FF^T + LEE^T L^T \end{bmatrix} - Q \preceq 0. \ \square \tag{12}$$

As a consequence[1] of Theorem 6 it follows that if $\Sigma^*(\bar{\mathbf{M}}_\mathbf{C}) \preceq_{\mathscr{S}} \Sigma^*_\psi$, then $\Sigma^*(\mathbf{M}_\mathbf{C}) \preceq^{\varepsilon}_{\mathscr{S}} \Sigma^*_\psi$, and if $\Sigma^*(\bar{\mathbf{M}}_\mathbf{C}) \sim_{\mathscr{B}} \Sigma^*_\psi$, then $\Sigma^*(\mathbf{M}_\mathbf{C}) \sim^{\varepsilon}_{\mathscr{B}} \Sigma^*_\psi$. Finally note that (12) is known to admit positive matrices Q for which ε is finite if $A-BK$ and $A-LC$ are both stable matrices [6].

A. Selection of the matrix gains L and K

Thus far we have assumed that L and K are chosen so that they stabilise $A-LC$ and $A-BK$. It is known that, as long as the model is detectable and stabilisable, these gains exist [6]. A constructive approach to obtain L,K in a semi-optimal manner follows from Theorem 6. Omitting the initialisation, the computation of the precision level defined in (10) together with (12) for given L and K is equivalent to $\varepsilon = \lim_{t\to\infty} \sqrt{\mathbf{E}\|\Delta z(t)\|_2^2}$ for

$$\Delta x(t+1) = \begin{bmatrix} A-BK & LC \\ 0 & A-LC \end{bmatrix} \Delta x(t) + \begin{bmatrix} 0 & LE \\ F & -LE \end{bmatrix} \begin{bmatrix} w_1(t) \\ w_2(t) \end{bmatrix} \tag{13}$$

$$\Delta z(t) = Hx(t), \tag{14}$$

for given white noise sequences $w_1(t), w_2(t)$. As such the optimisation problem leading to L and K has been recast in the familiar LQG stochastic control problem [7] for which it is known that the optimal observer gain L and the optimal state-feedback gain K can be computed separately. The optimal observer gain with respect to the LQG problem is the Kalman filter gain, $L^* = (APC^T)(CPC^T + EE^T)^{-1}$ s.t. $P = APA^T - (APC^T)(CPC^T + EE^T)^{-1}(CPA^T) + FF^T$. On the other hand, the optimal state-feedback gain K solves a quadratic control problem, that is $K^* = (B^T SB)^{-1} B^T SA$ s.t. $S = A^T SA - A^T SB(B^T SB)^{-1} B^T SA + H^T H$. In the next case study this will be computed via the generalised eigenproblem algorithm [17] implemented in MATLAB. Note that since there is no trade-off between the state error and the magnitude of the control gain, the state feedback gain will push the control to deadbeat control [6]: this behaviour can be easily remedied by extending the observation space H with D_H, such that the extended performance signal becomes $z_e(t) = \begin{bmatrix} z^T(t) & z_u^T(t) \end{bmatrix}^T$, with $z_u(t) = D_H Kx(t)$, or equivalently with $z_u(t) = D_H(u(t) - \bar{u}(t))$.

V. CASE STUDY IN SMART BUILDINGS

We are interested in the advanced energy management of an office building. As a motivation for output-based controllers, consider a building that is divided in two connected zones, each with a radiator regulating the heat in each zone via the controlled boiler water temperature [18]. Due to a sensor fault in the second zone, only the temperature in the first zone and the ambient (outside) temperatures are measured. The temperature fluctuations in the two zones and the ambient temperature are modelled via $\mathbf{M}$ as [18]

$$x(t+1) = \Xi x(t) + \Gamma u(t) + Fw_1(t) \tag{15}$$

$$y(t) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} x(t) + Ew_2(t), \quad z(t) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} x(t), \tag{16}$$

with stable dynamics

$$\Xi = \begin{bmatrix} 0.8725 & 0.0625 & 0.0375 \\ 0.0625 & 0.8775 & 0.0250 \\ 0 & 0 & 0.9900 \end{bmatrix}, \quad \Gamma = \begin{bmatrix} 0.0650 & 0 \\ 0 & 0.0600 \\ 0 & 0 \end{bmatrix},$$

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. Note that since Ξ is stable, it follows that (Ξ, Γ) is stabilisable and $(\Xi, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix})$ is detectable. The state variables are initiated as $x(0) = [16 \ 14 \ -5]^T$. The constants in matrix Ξ are selected to represent the heat exchange rate between the individual zones and the heat loss rate of each zone to the ambient; those in Γ represent the rate of heat supplied by the radiators to the two zones, respectively. The disturbances are modelled as independent and identically distributed standard normal

[1] Note that we have trivially assumed that this (bi-)simulation relation between the transition system $\Sigma(\bar{\mathbf{M}}_\mathbf{C})$ and Σ_ψ is maintained when embedding them in $\mathscr{T}^*$ via Dirac distributions [14].

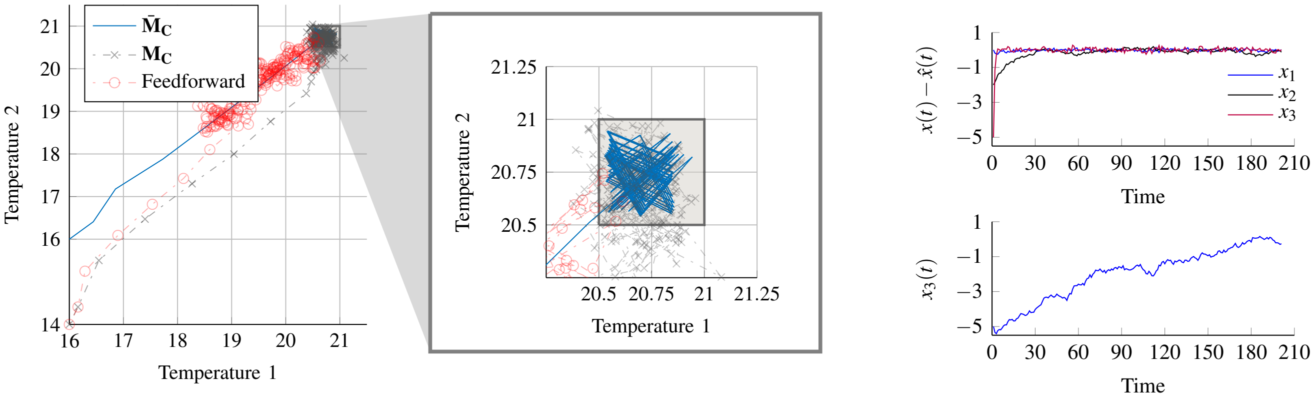

(a) Simulation outcomes for controlled models: $\bar{\mathbf{M}}_\mathbf{C}$ denotes state-based control of the noiseless model realisation ([5]); $\mathbf{M}_\mathbf{C}$ is the output-based control of the Gaussian process model (15); Feedforward denotes feedforward design using $\bar{\mathbf{M}}_\mathbf{C}$.

(b) (Upper plot) Error in state estimation for $\mathbf{M}_\mathbf{C}$; (Lower plot) Deviation from mean ambient temperature.

Fig. 3: Case study in smart buildings

TABLE I: *Error bounds – Accuracy of the controlled systems based on the interface. An initialisation is given by ε_{x_0}, for the perfect initialisation, or for $t \to \infty$ the system the accuracy is given as ε_∞. The estimates $\hat{\varepsilon}_{x_0,100}$ and ε_∞ are computed as $\sqrt{\hat{\mathbf{E}}_{1:100}\|z(t)-\bar{z}(t)\|_2^2}$ and $\sqrt{\hat{\mathbf{E}}_{10^2:4\times10^3}\|z(t)-\bar{z}(t)\|_2^2}$ respectively, with the empirical mean computed as $\mathbf{E}_{i:j}x=\frac{1}{j-i}\sum_{k=i}^{j}x(k)$.*

	ε_{x_0}	ε_∞	$\hat{\varepsilon}_{x_0,100}$	$\hat{\varepsilon}_\infty$
$\bar{\mathbf{M}}_\mathbf{C} \times_{\mathscr{F}_{ff}} \mathbf{M}$	3.9618	0.4890	1.9961	0.4845
$\mathbf{M}_\mathbf{C}$	2.1194	0.1284	0.5184	0.1240

distributions $w_{1,2}(t)$, rescaled by

$$F = \begin{bmatrix} .05 & -.02 & 0 \\ -.02 & .05 & 0 \\ 0 & 0 & 0.1 \end{bmatrix} \text{ and } E = \begin{bmatrix} .05 & 0 \\ 0 & .05 \end{bmatrix}.$$

The upper block in F represents random heat transfers, caused for example by people moving within and between zones, whereas the lower, right-diagonal element represents the stochastic nature of the fluctuation in the outside temperature. The values in E define the standard deviation of the additive disturbance on the temperature sensors in the first zone and in the ambient. $y(t)$ is the stochastic signal that can be measured, whereas the specification is defined over $z(t)$ (zone temperatures).

The objective is to design an output-based, correct-by-design controller, such that the temperature trajectories $z(t) = (x_1(t), x_2(t))$ eventually both take values in the interval $[20.5,\ 21]^2$, and remain within this interval thereafter. Formally this LTL property is expressed as "eventually always". The controller is initialised with $\hat{x}(0) = [16\ 16\ 0]^T$: this deviation from $x(0)$ is selected to model a realistic situation occurring after a sensor failure in zone 2 is discovered.

The dynamics of the noiseless model $\bar{\mathbf{M}}$ are solely governed over the first two states, where the correct-by-design controller for the given specification is designed. We synthesise $\bar{\mathbf{M}}_\mathbf{C}$ by PESSOA [5], where the discrete-time dynamics are further discretised over state and action spaces: we have selected a state quantisation of .05 over the range $[15, 25]^2$, and an input quantisation of .05 over $[10, 30]^2$. Fig. 3a displays (continuous blue line) the state trajectory of the obtained correct-by-design system $\bar{\mathbf{M}}_\mathbf{C}$: it can be observed that the controller regulates the model to eventually remain within the target region.

Next, we are interested in extending the designed controller to the concrete (noisy) model of the system based on noisy output measurements of the first zone and of the ambient. As a first attempt we implement the controller based on a feedforward architecture, where $\mathscr{F}_{ff} := \bar{u}(t)$. This is what we would obtain applying the results in [14]. It can be observed in Fig. 3a (circled red realisation) that a trajectory $(x_1(t), x_2(t))$ in $\bar{\mathbf{M}}_\mathbf{C} \times_{\mathscr{F}_{ff}} \mathbf{M}$ deviates substantially from the desired temperature range. In Table I the accuracy of this feedforward interface is given. As a second design, we implement the structure in Fig. 2, where the gains K, L, as detailed in Subsection IV-A, are selected as the optimal LQ and Kalman gains, respectively. The resulting design values are

$$L = \begin{bmatrix} 0.5201 & 0.0333 \\ -0.2239 & 0.0262 \\ 0.0022 & 0.8196 \end{bmatrix} \text{ and } K = \begin{bmatrix} 13.4231 & 0.9615 & 0.5769 \\ 1.0417 & 14.6250 & 0.4167 \end{bmatrix}.$$

A trajectory (crossed grey line in Fig. 3a) realised from $\mathbf{M}_\mathbf{C} = \bar{\mathbf{M}}_\mathbf{C} \times_{\mathscr{F}_g} (\mathbf{M} \| \mathbf{O}(\mathbf{M}))$ and based on the previous noise realisation ends up close to the desired temperature range. This substantial improvement with respect to the feedforward interface is also quantified in Table I . Fig. 3b displays the error of the state estimation $x(t) - \hat{x}(t)$ of $\mathbf{M}_\mathbf{C}$ (upper plot): it can be observed that the estimated state converges to the exact state. The lower plot in Fig. 3b provides a simulation of the deviation of the ambient temperature from its mean.

VI. Conclusions and Future Work

In this work we have shown that correct-by-design controllers can be extended to work on stochastic partially-observable LTI systems, as long as the LTI system is detectable and stabilisable. Future work will concern extensions to non-linear dynamics and the development of tailored notions of probabilistic approximations.

References

[1] M. Y. Vardi, "From philosophical to industrial logics," in *Logic and Its Applications*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 89–115.

[2] E. M. Clarke, "The birth of model checking," in *25 Years of Model Checking*. Springer, 2008, pp. 1–26.

[3] P. Tabuada and G. J. Pappas, "Linear time logic control of discrete-time linear systems," *Automatic Control, IEEE Transactions on*, vol. 51, no. 12, pp. 1862–1877, 2006.

[4] P. Tabuada, *Verification and control of hybrid systems*. Boston, MA: Springer US, 2009.

[5] M. Mazo Jr, A. Davitian, and P. Tabuada, "PESSOA: towards the automatic synthesis of correct-by-design control software," in *Work-in-progress HSCC*, 2010.

[6] G. F. Franklin, J. D. Powell, and M. L. Workman, *Digital control of dynamic systems*, 2nd ed. Addison-Wesley Menlo Park, 1990.

[7] H. Witsenhausen, "Separation of estimation and control for discrete time systems," *Proc. of IEEE*, vol. 59, no. 11, pp. 1557–1566, 1971.

[8] S. Giro and M. N. Rabe, "Verification of partial-information probabilistic systems using counterexample-guided refinements," in *Proc. on Automated Technology for Verification and Analysis*, ser. LNCS. Springer, 2012, pp. 333–348.

[9] K. Chatterjee, M. Chmelik, R. Gupta, and A. Kanodia, "Qualitative analysis of POMDPs with temporal logic specifications for robotics applications," *CoRR*, vol. abs/1409.3360, 2014.

[10] L. Zhang, H. Hermanns, and D. N. Jansen, "Logic and model checking for hidden Markov models," in *Proc. on Formal Techniques for Networked and Distributed Systems*. Springer, 2005, pp. 98–112.

[11] A. Abate, "Approximation metrics based on probabilistic bisimulations for general state-space markov processes: A survey," *Electronic Notes in Theoretical Computer Science*, vol. 297, pp. 3 – 25, 2013, proceedings of the first workshop on Hybrid Autonomous Systems.

[12] S. E. Z. Soudjani and A. Abate, "Adaptive and sequential gridding for abstraction and verification of stochastic processes," *SIAM Journal on Applied Dynamical Systems*, vol. 12, no. 2, pp. 921–956, 2012.

[13] A. Julius and G. Pappas, "Approximations of stochastic hybrid systems," *IEEE Trans. on Automatic Control*, vol. 54, no. 6, pp. 1193–1203, Jun. 2009.

[14] M. Zamani, P. Mohajerin Esfahani, R. Majumdar, A. Abate, and J. Lygeros, "Symbolic control of stochastic systems via approximately bisimilar finite abstractions," *IEEE Trans. on Automatic Control,*, vol. 59, no. 12, pp. 3135–3150, Dec 2014.

[15] S. Haesaert, P. M. J. Van den Hof, and A. Abate, "Observer-based correct-by-design controller synthesis," *ArXiv e-prints*, Sep. 2015, 1509.03427.

[16] A. Girard and G. J. Pappas, "Hierarchical control system design using approximate simulation," *Automatica*, vol. 45, pp. 566–571, 2009.

[17] W. F. Arnold III and A. J. Laub, "Generalized eigenproblem algorithms and software for algebraic Riccati equations," *Proceedings of the IEEE*, vol. 72, no. 12, pp. 1746–1754, 1984.

[18] O. Holub and K. Macek, "HVAC simulation model for advanced diagnostics," in *Symp. Intelligent Signal Processing*, 2013, pp. 93–96.